

The merged-staircase property: a necessary and nearly sufficient condition for SGD learning of sparse functions on two-layer neural networks

Emmanuel Abbe

École polytechnique Fédérale de Lausanne

Enric Boix-Adserà

Department of Electrical Engineering, MIT

Theodor Misiakiewicz

Department of Statistics, Stanford University

Editors: Po-Ling Loh and Maxim Raginsky

Abstract

It is currently known how to characterize functions that neural networks can learn with SGD for two extremal parametrizations: neural networks in the linear regime, and neural networks with no structural constraints. However, for the main parametrization of interest —non-linear but regular networks— no tight characterization has yet been achieved, despite significant developments.

We take a step in this direction by considering depth-2 neural networks trained by SGD in the mean-field regime. We consider functions on binary inputs that depend on a latent low-dimensional subspace (i.e., small number of coordinates). This regime is of interest since it is poorly understood how neural networks routinely tackle high-dimensional datasets and adapt to latent low-dimensional structure without suffering from the curse of dimensionality. Accordingly, we study SGD-learnability with $O(d)$ sample complexity in a large ambient dimension d .

Our main results characterize a hierarchical property —the merged-staircase property— that is both *necessary and nearly sufficient* for learning in this setting. We further show that non-linear training is necessary: for this class of functions, linear methods on any feature map (e.g., the NTK) are not capable of learning efficiently. The key tools are a new “dimension-free” dynamics approximation result that applies to functions defined on a latent space of low-dimension, a proof of global convergence based on polynomial identity testing, and an improvement of lower bounds against linear methods for non-almost orthogonal functions.

Keywords: two-layer neural networks, feature learning, necessary and sufficient conditions, merged-staircase functions, mean-field dynamics.

1. Introduction

Major research activity has recently been devoted to understanding what function classes can be learned by SGD on neural networks. Two extremal cases are well understood. On one extreme, neural networks can be parametrized to collapse under SGD to linear models, for which a clear picture has been drawn (Jacot et al., 2018; Li and Liang, 2018; Du et al., 2018, 2019; Allen-Zhu et al., 2019b,a; Arora et al., 2019; Zou et al., 2020; Oymak and Soltanolkotabi, 2020). On the other extreme, neural networks with zero parametrization constraint (besides polynomial size) have been shown to be able to emulate essentially any efficient learning algorithm (Abbe and Sandon, 2020;

Abbe et al., 2021b) albeit with non-regular¹ architectures. So both of these extremes admit a fairly complete characterization. However, none of these seem to capture the right behavior behind deep learning, or more specifically, behind *non-linear but regular networks*. Such networks are known to go beyond linear learning (Bach, 2017; Ghorbani et al., 2021b; Daniely and Malach, 2020; Refinetti et al., 2021; Allen-Zhu and Li, 2019; Ghorbani et al., 2019; Yehudai and Shamir, 2019; Allen-Zhu and Li, 2020; Li et al., 2020) (even though the NTK can be competitive on several instances (Geiger et al., 2020)), and seem to exploit structural properties of the target functions in order to efficiently build their features.

Can we thus characterize learning in the non-linear regime for regular networks? Various important results have been developed in this direction, we focus here on the most relevant to us. Chizat and Bach (2018); Mei et al. (2018); Rotskoff and Vanden-Eijnden (2018); Sirignano and Spiliopoulos (2020) show that for a certain scaling at initialization, the SGD dynamics on large-width neural networks concentrates on a fully non-linear dynamics, the mean-field dynamics, described by a Wasserstein gradient flow, contrasting with the linear dynamics of the NTK regime (Jacot et al., 2018). In Allen-Zhu and Li (2019, 2020), the power of deep networks is demonstrated by showing how SGD and quadratic activations can efficiently learn a non-trivial teacher class hierarchically, with the notion of backward feature correction (Allen-Zhu and Li, 2020).

However, no tight necessary and sufficient characterization of what functions are learnable emerges from these works. The difficulty being that tight necessity results are difficult to obtain in such a setting since SQ-like arguments (Blum et al., 1994; Feldman et al., 2017b; Kearns, 1998; Blum et al., 2003; Feldman, 2016; Yang, 2005; Feldman et al., 2017a; Steinhardt et al., 2015; Abbe and Sandon, 2020; Abbe et al., 2021b; Goel et al., 2020) are not expected to be tight (besides for the extreme case of unconstrained networks (Abbe and Sandon, 2020; Abbe et al., 2021b)), and sufficiency results are significantly more difficult to obtain due to the more complex (non-linear) dynamics of SGD training.

Is there hope to characterize tight necessary and sufficient conditions for function classes to be learnable by standard SGD on standard neural networks?

As a first attempt in that direction, we focus in this paper on a natural setting: *learning sparse functions* on the d -dimensional hypercube, i.e., functions that depend on a small latent (unknown) subset of coordinates of the input. We further restrict the optimization regime considered to *two-layer neural networks* trained by one-pass batch-SGD in the *mean-field regime*. This allows us to study a regime of optimization that goes beyond the linear regime while averaging out some of the complexity of studying non-linear SGD.

The motivation for the setting of learning sparse function is three-fold: (1) Linear (fixed features) methods do not adapt to latent sparsity, and suffer from the *curse of dimensionality* (Bach, 2017). (2) On the contrary, Bach (2017); Schmidt-Hieber (2020) shows that neural networks can overcome this curse and learn sparse functions sample-efficiently. However, these works do not provide tractable algorithms and the question of when SGD-trained neural networks can adapt to sparsity remains largely open. (3) Some sparse functions, such as monomials, are known to be much harder to learn than others from SQ-like lower bounds (Kearns, 1998; Blum et al., 1994; Abbe et al., 2021b), and we expect SGD to inherit some of this complex behavior. Therefore, the problem of

1. Here we refer to ‘regular’ for architectures used in tangent kernel results or more generally architectures used in neural network applications.

learning sparse functions presents a clear-cut separation between fixed-feature and feature learning methods, and can help understand the limits of SGD-training on neural networks.

To gain insights on the interaction between SGD and the function structure that allows adaptivity to sparsity, we will ask the following question: *Can one characterize necessary and sufficient conditions for a low-dimensional latent function to be learnable by standard SGD on standard neural networks in an arbitrarily large ambient dimension?* More precisely, we will consider a P -dimensional latent function $h_* : \{+1, -1\}^P \rightarrow \mathbb{R}$ and consider learning sparse functions $f_* : \{+1, -1\}^d \rightarrow \mathbb{R}$ with $f_*(\mathbf{x}) = h_*(\mathbf{z})$ for arbitrary ambient dimension d and latent subset of coordinates $\mathbf{z} := \mathbf{x}_{\mathcal{I}} = (x_{i_1}, \dots, x_{i_P})$. As motivating examples, consider the two functions:

$$h_1(\mathbf{z}) = z_1 z_2 z_3, \quad h_2(\mathbf{z}) = z_1 + z_1 z_2 + z_1 z_2 z_3.$$

Both of these functions depend on only 3 coordinates (i.e., they are 3-sparse), and because of the presence of the degree-3 monomial both require $\Omega(d^3)$ samples to be learned by a linear method. However, are these functions equivalent for SGD-trained neural networks? If not, can we obtain a fine-grained analysis that separates them?

In this paper, we introduce the following notion: we say that a latent function h_* is **strongly SGD-learnable in $O(d)$ -scaling**, if $O(d)$ samples are enough to learn $f_*(\mathbf{x}) = h_*(\mathbf{z})$ for arbitrary latent subspace \mathbf{z} and dimension d , using batch-SGD on a two-layer neural network in the mean-field regime. The main contribution of this paper is then to characterize with a **necessary and nearly sufficient condition** the class of functions that are strongly SGD-learnable in $O(d)$ -scaling. This is achieved with the **merged-staircase property** (MSP), stating that the non-zero Fourier coefficients of g can be ordered as subsets $\{S_1, \dots, S_r\}$ such that for any $i \in [r]$,

$$|S_i \setminus \cup_{j=1}^{i-1} S_j| \leq 1.$$

For instance, $h_2(\mathbf{z}) = z_1 + z_1 z_2 + z_1 z_2 z_3$ has Fourier coefficients (monomials) that can be ordered as $\{\{1\}, \{1, 2\}, \{1, 2, 3\}\}$, and each new set is incremented only by one element at each time. So h_2 satisfies the MSP (or is an MSP function with a slight abuse of terminology) and so is the function $z_1 + z_1 z_2 + z_2 z_3 + z_1 z_2 z_3$. However, the function $h_1(\mathbf{z}) = z_1 z_2 z_3$ directly makes a leap to a degree-3 Fourier coefficient and does not therefore satisfy the MSP. Our main results thus imply that h_2 can be learned with $O(d)$ samples in this regime, but not h_1 . The *near* sufficiency part in our result stands for the fact that the sufficiency result is proved for “generic” merged-staircase functions, i.e., excluding a measure zero subclass. This ‘genericity’ is in fact needed, as we provide degenerate examples in Section 4 for which the strong SGD-learnability in $O(d)$ -scaling is indeed not achievable.

The terminology MSP comes from the fact that this condition generalizes the basic staircase property introduced in [Abbe et al. \(2021a\)](#), which only encompasses nested chains of coefficients with $|S_i \setminus S_{i-1}| = 1$, such as the vanilla staircase function (e.g., h_2) and slight generalizations with multiple chains. In [Abbe et al. \(2021a\)](#) it is shown that staircase functions are learnable by neural nets that are deep but sparse, and with an unconventional gradient-based training algorithm (see Section 1.2 for further discussion). Further [Abbe et al. \(2021a\)](#) does not provide necessary conditions for learning, nor fine-grained complexity guarantees (beyond ‘polynomial’).

Finally, while strong SGD-learnability is defined for a fixed latent function and fixed P , the number of samples required to fit MSP functions remains polynomial in d for P growing sufficiently slowly in d . This is of interest because in this regime, we can show that the considered functions are

not learnable by any linear methods with any sample complexity (or feature space dimension) that is polynomial (using contribution (4) below). Thus the merged-staircase functions of such degree are efficiently learnable by SGD on networks of depth two but not by linear methods.

1.1. Summary of main results

Recall that any function $h_* : \{+1, -1\}^P \rightarrow \mathbb{R}$ can be decomposed in the Fourier-Walsh basis as $h_*(\mathbf{z}) = \sum_{S \subseteq [P]} \hat{h}_*(S) \chi_S(\mathbf{z})$, where $\hat{h}_*(S) := \langle h_*, \chi_S \rangle$, $\chi_S(\mathbf{z}) := \prod_{i \in S} z_i$, where we denoted the inner-product between two functions $\langle f, g \rangle := \mathbb{E}_{\mathbf{z}}[f(\mathbf{z})g(\mathbf{z})]$ with $\mathbf{z} \sim \text{Unif}(\{-1, +1\}^P)$. This corresponds to expressing the function $h_*(\mathbf{z})$ as a weighted sum of orthogonal monomials $\chi_S(\mathbf{z})$, with weights $\hat{h}_*(S)$ called the Fourier coefficients of h_* .

We now formally define the Merged-Staircase Property. Let us call any $\mathcal{S} \subseteq 2^{[P]}$ a *set structure*.

Definition 1 We say that $\mathcal{S} = \{S_1, \dots, S_m\} \subseteq 2^{[P]}$ is a Merged-Staircase Property (MSP) set structure if the sets can be ordered so that for each $i \in [m]$, $|S_i \setminus \cup_{i' < i} S_{i'}| \leq 1$.

Definition 2 (Merged-Staircase Property) Let $\mathcal{S} \subset 2^{[P]}$ be the non-zero Fourier coefficients of h_* , i.e., $\hat{h}_*(S) \neq 0$ iff $S \in \mathcal{S}$. We say that h_* satisfies the merged-staircase property (MSP) if \mathcal{S} is a MSP set structure.

In words, h_* satisfies the MSP if the monomials in its Fourier decomposition can be ordered sequentially such that the supports of the monomials grow by at most one at a time. Examples of MSP functions include vanilla staircases (i.e., $z_1 + z_1 z_2 + \dots + \prod_{i=1}^P z_i$), $z_1 + z_1 z_2 + z_2 z_3 + z_3 z_4$, or $z_1 + z_2 + z_3 + z_4 + z_1 z_2 z_3 z_4$, but not $z_1 + z_1 z_2 z_3 + z_1 z_2 z_3 z_4$, $z_1 + z_1 z_2 + z_3 z_4$, or $z_1 z_2 z_3$. We briefly summarize our results here:

- (1) **Dimension-free dynamics and equivalent characterization of strong SGD-learnability.** We introduce a *dimension-free dynamics* (independent of d), which correspond to the gradient flow associated to learning h_* with a certain two-layer neural network in the space of distributions on \mathbb{R}^{P+2} . We show h_* is strongly $O(d)$ -SGD learnable if and only if this dimension-free dynamics can reach 0 risk when initialized with first-layer weights at 0.
- (2) **MSP necessity.** We show that for non-MSP h_* , the associated dimension-free dynamics stays bounded away from 0. From the previous equivalence, we deduce that MSP is necessary for a function to be strongly $O(d)$ -SGD-learnable.
- (3) **MSP near-sufficiency.** We first show that vanilla staircases are strongly $O(d)$ -SGD-learnable for smooth activation functions as long as $\sigma^{(r)}(0) \neq 0$ for $r = 0, \dots, P$.

For general MSP functions, however, some symmetric MSP functions have degenerate dynamics and are not strongly $O(d)$ -SGD-learnable (see Section 4). We show instead that MSP h_* are almost surely strongly $O(d)$ -SGD-learnable. I.e., the degenerate examples are a measure-zero set. This is proved for generic degree- L polynomial activations, and we explain how one can extend this result to generic smooth activations in the appendix.

- (4) **Superpolynomial separation with linear methods.** One can take MSP functions (e.g., vanilla staircases) with P slowly growing with d so that the overall sample complexity of the above neural network results stay as $d^{O(1)}$, while we show that any linear method requires a sample complexity of $d^{\omega_d(1)}$.

These main results are further achieved with several side results of independent interest: (i) The approximation of the standard mean-field dynamics by the *dimension-free dynamics*, valid for P -sparse target functions and $d \gg P$. We provide a new version of the non-asymptotic bounds from [Mei et al. \(2018, 2019\)](#), which now compares SGD with this dimension-independent dynamics; (ii) A new proof technique to study layer-wise SGD dynamics which reduces the proof of global convergence to a polynomial identity testing problem, i.e., whether a certain polynomial is non-identically zero; (iii) An improvement of prior dimension lower-bounds for linear (kernel) methods ([Hsu et al., 2021](#); [Hsu](#); [Kamath et al., 2020](#)) that is tighter for function classes that are non-almost orthogonal (such as staircase functions, allowing for contribution (4) above).

The rest of the paper is organized as follows. The next section overviews related work. Section 2 provides a formal definition of strong SGD-learnability in $O(d)$ -scaling. In Section 3, we introduce the dimension-free dynamics and the equivalence with strong $O(d)$ -SGD-learnability. The MSP necessary condition is then derived as a direct consequence of this equivalence. In Section 4, we provide our sufficient conditions for strong $O(d)$ -SGD-learnability. In Section 5, we discuss how this implies a separation with linear methods.

1.2. Further related literature

[Abbe et al. \(2021a\)](#) introduces a class of staircase functions, which our merged-staircase function class generalizes. They show that staircase functions are learnable by some neural nets with a gradient-based training algorithm. However, the approach remains non-standard: (i) the network’s layers are sparse in order to guide the construction of the features; (ii) a coordinate descent variant of SGD is used that differs from the classical SGD algorithm. Further, the analysis is carried in the ‘polynomial scaling lens’ rather than a finer sample complexity, and no necessity results are derived. In contrast, we provide here both a necessary and nearly sufficient characterization for SGD-learning on a two-layer neural networks in the fine-grained $O(d)$ -scaling.

Multiple works have used mean-field (also called distributional) dynamics to approximate the SGD trajectory. Relevant to us is [Chizat and Bach \(2020\)](#) which showed that neural networks trained in the mean-field regime converge to a max-margin classifier that is independent of the dimension for latent low-dimensional target functions. However, these works do not provide quantitative results in terms of sample-complexity. A notable exception is [Mei et al. \(2018\)](#) which studies classifying anisotropic gaussians: they show that the mean-field dynamics concentrates on a simplified low-dimensional dynamics as $d \rightarrow \infty$. However, this simplification is due to rotational invariance of the problem and not the sparsity of the target function.

In approximation theory, it has been understood for a long time that sparse functions are naturally well approximated by neural networks ([Barron, 1993](#)). Recent work ([Bach, 2017](#); [Schmidt-Hieber, 2020](#); [Ghorbani et al., 2021b](#); [Celentano et al., 2021](#)) have shown that neural networks can learn sparse functions more sample-efficiently than linear methods. However, these works do not provide tractable algorithms.

Finally, a string of works ([Yehudai and Shamir, 2019](#); [Allen-Zhu and Li, 2019](#); [Allen-Zhu and Li, 2020](#); [Li et al., 2020](#); [Daniely and Malach, 2020](#); [Refinetti et al., 2021](#); [Ghorbani et al., 2021a,b](#); [Malach et al., 2021](#); [Karp et al., 2021](#); [Suzuki and Akiyama, 2020](#)) have shown separation results between gradient-trained neural networks and fixed-features models. We refer to Appendix B of [Malach et al. \(2021\)](#) for a detailed survey. In particular, [Daniely and Malach \(2020\)](#) considers the learning of parity functions, with a modified input distribution that gives correlation to the response

and allows for domain extraction; it also uses the population dynamics (infinite samples). In [Malach and Shalev-Shwartz \(2020\)](#), the learning of Boolean circuits of logarithmic depth is considered via neural networks with layer-wise gradient descent, but with an architecture that is required to match the Boolean circuit being learned, i.e., not with a ‘regular’ or ‘blackbox’ architecture. Lastly, [Basri et al. \(2019\)](#); [Cao et al. \(2021\)](#) show that during training, SGD on 2-layer ReLU networks learns faster the lower frequency components of a target function, in similar spirit to low degree monomials, but the approach relies on the linear regime rather than the non-linear regime of interest here, and suffers from an exponential dependency on the degree.

2. Strong SGD-learnability in $O(d)$ -scaling

Consider n iid data points $(\mathbf{x}_i, y_i)_{i \in [n]}$ with covariates $\mathbf{x}_i \sim \text{Unif}(\{+1, -1\}^d)$ and responses $y_i = f_*(\mathbf{x}) + \varepsilon_i$ with bounded independent noise $\mathbb{E}[\varepsilon_i] = 0$. We assume that $f_* : \{+1, -1\}^d \rightarrow \mathbb{R}$ is a sparse function with latent P -dimensional function h_* , i.e., there exists an (unknown) subset of coordinates $\mathbf{z} = \mathbf{x}_{\mathcal{I}} = (x_{i_1}, \dots, x_{i_P})$ (the signal part of the input) such that $f(\mathbf{x}) = h_*(\mathbf{z})$. We consider fitting this data using a two-layer fully-connected neural network with N hidden units and weights $\Theta := (\theta_j)_{j \in [N]} = (a_j, \mathbf{w}_j)_{j \in [N]} \in \mathbb{R}^{N(d+1)}$:

$$\hat{f}_{\text{NN}}(\mathbf{x}; \Theta) = \frac{1}{N} \sum_{j \in [N]} a_j \sigma(\langle \mathbf{w}_j, \mathbf{x} \rangle). \quad (2\text{-NN})$$

We train the parameters Θ using batch-SGD with square loss and batch size b . We allow for time-varying step sizes $\{(\eta_k^a, \eta_k^w)\}_{k \geq 0}$, and layer-wise ℓ_2 -regularization with parameters $\lambda^a, \lambda^w \geq 0$. Given samples $\{(\mathbf{x}_{ki}, y_{ki})_{i \in [b]}\}_{k \geq 0}$ and initialization $(\theta_j^0)_{j \in [N]} \sim \text{iid } \rho_0$, the weights are updated at each step:

$$\theta_j^{k+1} = \theta_j^k + \frac{1}{b} \sum_{i \in [b]} \{y_{ki} - \hat{f}_{\text{NN}}(\mathbf{x}_{ki}; \Theta^k)\} \cdot \mathbf{H}_k \nabla_{\theta} [a_j^k \sigma(\langle \mathbf{w}_j^k, \mathbf{x}_{ki} \rangle)] - \mathbf{H}_k \Lambda \theta_j^k, \quad (\text{bSGD})$$

where we introduced $\mathbf{H}_k = \text{diag}(\eta_k^a, \eta_k^w \cdot \mathbf{I}_d)$ and $\Lambda = \text{diag}(\lambda^a, \lambda^w \cdot \mathbf{I}_d)$. We will be interested in the prediction error (test error) $R(f_*, \hat{f}) = \mathbb{E}_{\mathbf{x}} [\{f_*(\mathbf{x}) - \hat{f}(\mathbf{x})\}^2]$.

We first consider a general definition for a class of sparse functions to be learnable. We take $\{P(d)\}_{d \geq 1}$ a sequence of integers (here, we allow the sparsity parameter P to grow with d) and consider a general class of functions defined as $\mathcal{H} = \{\mathcal{H}_{P(d)}\}_{d \geq 1}$ with $\mathcal{H}_{P(d)} \subseteq L^2(\{+1, -1\}^{P(d)})$.

Definition 3 (SGD-learnability in $O(d^\alpha)$ -scaling) *We say that a function class \mathcal{H} is SGD-learnable in $O(d^\alpha)$ -scaling if the following hold for some $C(\cdot, \mathcal{H}) : \mathbb{R}_{>0} \rightarrow \mathbb{R}_{>0}$. For any $h \in \mathcal{H}_{P(d)}$, $\varepsilon > 0$ there exist hyperparameters $(N, b, \sigma, \lambda^a, \lambda^w, \{\eta_k^a, \eta_k^w\}_{k \in [0, k_0]})$ and initialization ρ_0 , such that: (1) for a sample size² $n \leq C(\varepsilon, \mathcal{H})d^\alpha$; and (2) for any $\mathcal{I} \subseteq [d]$, $|\mathcal{I}| = P(d)$, and target function $f_*(\mathbf{x}) = h(\mathbf{x}_{\mathcal{I}})$, k_0 steps of batch stochastic gradient descent (bSGD) achieves prediction error ε with prob. at least $9/10$.*

This definition covers many scenarios that occur in practice where the practitioner is allowed to tune the hyperparameters of the dynamics. While this choice leaves the question of tractability open,

2. Note that $n \leq bk_0$ with equality if we assume fresh samples at each iteration, as in the next definition.

we note that the requirement that learnability must hold uniformly over all possible latent subspaces excludes many irregular scenarios. Furthermore, the next definition will require strong regularity on the hyperparameters, and our sufficiency results will hold for simple choices of hyperparameters.

In order to introduce strong SGD-learnability, we will restrict the previous definition in three major ways: (1) we consider a fixed dimension P and a fixed function $\mathcal{H} = \{h_*\}$, which is still nontrivial to learn since we do not know the set $\mathcal{I} \subseteq [d]$ such that $f_*(\mathbf{x}) = h_*(\mathbf{x}_{\mathcal{I}})$; (2) we consider the scaling³ of $\alpha = 1$; (3) we restrain the hyperparameters to be in either of two regimes (i) small batch size $b = o(d)$ and step size $\eta = o(1)$ trained for $\Theta(1/\eta)$ steps (“continuous”); and (ii) large batch size $b = \Theta(d)$ and step size $\eta = \Theta(1)$ trained for a total number of $\Theta(1)$ steps (“discrete”). For the sake of presentation, we will only present the continuous regime in the main text and defer the presentation of the discrete regime to Appendix C. We will assume that the hyperparameters obey the following for some constant K (independent of d):

- A0. (*Activation*) $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ is three times differentiable with $\|\sigma^{(k)}\|_{\infty} \leq K$ for $k = 0, \dots, 3$.
- A1. (*One-pass*) We have fresh samples at each steps, meaning $\{(\mathbf{x}_{k_i}, y_{k_i})\}_{k \geq 0, i \in [b]}$ are iid. Furthermore, the response variable is bounded $|y| \leq K$.
- A2. (*Initialization*) The initialization verifies $(a_i^0, \sqrt{d} \cdot \mathbf{w}_i^0) \sim \mu_a \otimes \mu_w^{\otimes d}$ where the distributions $\mu_a, \mu_w \in \mathcal{P}(\mathbb{R})$ are independent of d with $|a| \leq K$ on the support of μ_a and μ_w is symmetric and K^2 -sub-Gaussian. We will denote $m_2^w := \mathbb{E}_{W \sim \mu_w} [W^2]^{1/2}$.
- A3. (*Boundedness and lipschitzness of hyperparameters*) There exists a constant $\eta > 0$ such that $\eta_k^a, \eta_k^w \leq \eta K$, $|\eta_{k+1}^a - \eta_k^a| \leq \eta^2 K$ and $|\eta_{k+1}^w - \eta_k^w| \leq \eta^2 K$. Furthermore, $\lambda^a, \lambda^w \leq K$.

Definition 4 (Strong SGD-learnability in $O(d)$ -scaling) *We say that a function $h_* : \{-1, +1\}^P \rightarrow \mathbb{R}$ is strongly $O(d)$ -SGD-learnable if the following hold for some $C(\cdot, h_*)$, $T(\cdot, h_*) : \mathbb{R}_{>0} \rightarrow \mathbb{R}_{>0}$. For any $\varepsilon > 0$, $d \geq C(\varepsilon, h_*)$, $n \geq C(\varepsilon, h_*)d$ and $e^d \geq N \geq C(\varepsilon, h_*)$, there exists hyperparameters $(\sigma, b, \lambda^a, \lambda^w, \{\eta_k^a, \eta_k^w\}_{k \in [0, k_0]})$ and initialization ρ_0 satisfying A0-A3, $b \leq d$ and $k_0 = n/b \leq T(\varepsilon, h_*)/\eta$ s.t. for any $\mathcal{I} \subseteq [d]$, $|\mathcal{I}| = P$ and target function $f_*(\mathbf{x}) = h_*(\mathbf{x}_{\mathcal{I}})$, k_0 steps of batch stochastic gradient descent (bSGD) achieves test error ε with prob. at least $9/10$.*

Conditions A0-A3 guarantee that as long as d, n, N are taken sufficiently large, there exists a continuous mean-field dynamics that well-approximates batch-SGD up to (continuous) time T depending on ε, h_* . An analogous statement is true for strong-SGD-learnability in the “discrete regime”, except convergence is to a family of limiting discrete-time dynamics (deferred to Appendix C). This allows us to get a necessary condition for strong-learnability by studying the limiting dynamics (see next section).

Finally, we note that for any degree- k sparse function h_* , any linear method (e.g., arbitrary kernel or random feature methods) will require $\Omega(d^k)$ samples to fit functions $f_*(\mathbf{x}) = h_*(\mathbf{z})$ uniformly well over all latent subspaces $\mathbf{z} = \mathbf{x}_{\mathcal{I}}$ (see Section 5 for a formal statement). As emphasized in the introduction, this bound is not adaptive to the sparsity parameter P . In particular, any non-linear h_* that is strongly $O(d)$ -SGD-learnable provides a separation result between SGD-trained neural networks and linear methods.

3. Extending our results to $\alpha > 1$, and establishing how this relates to the ‘leap’ in the staircase definition (i.e., how can one jump monomial degrees) is a natural future direction to this work.

3. Continuous dimension-free dynamics and necessary condition

For simplicity, the results in this section are stated in the ‘continuous regime’ of strong SGD-learnability. Discrete versions can be found, with little modification, in Appendix C.

Mean-field approximation: A recent line of work (Chizat and Bach, 2018; Mei et al., 2018; Rotkoff and Vanden-Eijnden, 2018; Sirignano and Spiliopoulos, 2020; Mei et al., 2019) showed that one-pass batch-SGD (bSGD) can be well approximated in some regime by a continuous dynamics in the space of probability distributions on \mathbb{R}^{d+1} , which we will refer to as the *mean-field dynamics*.

Before describing this limiting dynamics, we first introduce a few definitions. To any distribution $\rho \in \mathcal{P}(\mathbb{R}^{d+1})$, we associate the infinite-width neural network

$$\hat{f}_{\text{NN}}(\mathbf{x}; \rho) = \int a\sigma(\langle \mathbf{w}, \mathbf{x} \rangle)\rho(d\mathbf{a} d\mathbf{w}). \quad (1)$$

In particular, (2-NN) corresponds to taking the empirical distribution $\hat{\rho}^{(N)} = N^{-1} \sum_{j \in [N]} \delta_{\theta_j}$. We assume further that there exist functions $\xi^a, \xi^w : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ and a parameter $\eta > 0$ (the time discretization) such that $\eta_k^a = \eta \xi^a(k\eta)$ and $\eta_k^w = \eta \xi^w(k\eta)$. We replace Assumption A3 by :

A3'. ξ^a, ξ^w are bounded Lipschitz $\|\xi^a\|_\infty, \|\xi^w\|_\infty, \|\xi^a\|_{\text{Lip}}, \|\xi^w\|_{\text{Lip}} \leq K$ and $\lambda^a, \lambda^w \leq K$.

Note that for any $(\eta_k^a, \eta_k^w)_{k \geq 0}$ obeying A3, there exists functions ξ^a, ξ^w such that A3' holds with same constant K . Conversely, any η discretization of ξ^a, ξ^w obeys A3 with constants η, K .

Consider the empirical distribution of the weights $\hat{\rho}_k^{(N)}$ after k batch-SGD steps, i.e., $\hat{\rho}_k^{(N)} = N^{-1} \sum_{j \in [N]} \delta_{\theta_j^k}$. For large N and small step size η , setting $k = t/\eta$, $\hat{\rho}_k^{(N)}$ is well approximated by a distribution $\rho_t \in \mathcal{P}(\mathbb{R}^{d+1})$ that evolves according to the following PDE:

$$\begin{aligned} \partial_t \rho_t &= \nabla_{\boldsymbol{\theta}} \cdot (\rho_t \mathbf{H}(t) \nabla_{\boldsymbol{\theta}} \psi(\boldsymbol{\theta}; \rho_t)), \\ \psi(\boldsymbol{\theta}; \rho_t) &= a \mathbb{E}_{\mathbf{x}} \left[\{ \hat{f}_{\text{NN}}(\mathbf{x}; \rho_t) - f_*(\mathbf{x}) \} \sigma(\langle \mathbf{w}, \mathbf{x} \rangle) \right] + \frac{1}{2} \boldsymbol{\theta}^\top \boldsymbol{\Lambda} \boldsymbol{\theta}, \end{aligned} \quad (\text{MF-PDE})$$

with initial distribution ρ_0 , and where we introduced $\mathbf{H}(t) = \text{diag}(\xi^a(t), \xi^w(t) \mathbf{I}_d)$. This PDE corresponds to a Wasserstein gradient flow on the square-loss test error $R(\rho) := \mathbb{E}[\{f_*(\mathbf{x}) - \hat{f}_{\text{NN}}(\mathbf{x}; \rho)\}^2]$ with regularization $\int \boldsymbol{\theta}^\top \boldsymbol{\Lambda} \boldsymbol{\theta} \rho(d\boldsymbol{\theta})$ and learning schedule $\mathbf{H}(t)$.

Dimension-free dynamics: For a sparse function $f_*(\mathbf{x}) := h_*(\mathbf{z})$, the (MF-PDE) concentrates to a dimension-free dynamics when $d \rightarrow \infty$. Decompose the input $\mathbf{x} = (\mathbf{z}, \mathbf{r})$ and the weights $\mathbf{w}_i = (\mathbf{u}_i, \mathbf{v}_i)$ with $\mathbf{u}_i \in \mathbb{R}^P$ aligned with \mathbf{z} and $\mathbf{v}_i \in \mathbb{R}^{d-P}$ with \mathbf{r} . By Assumption A2, \mathbf{w}_i^0 has iid symmetric coordinates, which implies that $\hat{f}_{\text{NN}}(\mathbf{x}; \rho_0)$ does not depend on \mathbf{r} . In fact, by symmetry of (MF-PDE), the mean-field solution $\hat{f}_{\text{NN}}(\mathbf{x}; \rho_t)$ stays independent of \mathbf{r} throughout the dynamics:

$$\forall t \geq 0, \quad \hat{f}_{\text{NN}}(\mathbf{x}; \rho_t) = \int a^t \sigma(\langle \mathbf{x}, \mathbf{w}^t \rangle) \rho_t(d\boldsymbol{\theta}^t) = \int a^t \mathbb{E}_{\mathbf{r}} [\sigma(\langle \mathbf{u}^t, \mathbf{z} \rangle + \langle \mathbf{v}^t, \mathbf{r} \rangle)] \rho_t(d\boldsymbol{\theta}^t), \quad (2)$$

and we denote with a slight abuse of notation, $\hat{f}_{\text{NN}}(\mathbf{z}; \rho_t) := \hat{f}_{\text{NN}}(\mathbf{x}; \rho_t)$.

With $\mathbf{r} \sim \text{Unif}(\{-1, +1\}^{d-P})$, one can show that $\langle \mathbf{v}^t, \mathbf{r} \rangle$ can be well approximated by $\|\mathbf{v}^t\|_2 G$ with $G \sim \text{N}(0, 1)$ when $d \gg P$ and $t = O_d(1)$. We introduce effective parameters $\bar{\boldsymbol{\theta}}^t = (\bar{a}^t, \bar{\mathbf{u}}^t, \bar{s}^t)$

with distribution $\bar{\rho}_t \in \mathcal{P}(\mathbb{R}^{P+2})$, and we replace the neural network (6) by an effective neural network (with a slight abuse of notation)

$$\hat{f}_{\text{NN}}(\mathbf{z}; \bar{\rho}_t) := \int \bar{a}^t \mathbb{E}_G [\sigma(\langle \bar{\mathbf{u}}^t, \mathbf{z} \rangle + \bar{s}^t G)] \bar{\rho}_t(d\bar{\boldsymbol{\theta}}^t). \quad (3)$$

We see that $\hat{f}_{\text{NN}}(\cdot; \bar{\rho}_t)$ can be seen as a two layer neural network in dimension P , with adaptive Gaussian smoothing. Taking $d \rightarrow \infty$ with P fixed, $(a^0, \mathbf{u}^0, \|\mathbf{v}^0\|_2)$ (with distribution ρ_0 satisfying A2) converges in distribution to $(\bar{a}^0, \bar{\mathbf{u}}^0, \bar{s}^0) \sim \bar{\rho}_0$ with $\bar{a}^0 \sim \mu_a$, $\bar{\mathbf{u}}^0 = \mathbf{0}$ and $\bar{s}^0 = m_2^w$, and the dynamics (MF-PDE) simplifies into the following dimension-free dynamics

$$\begin{aligned} \partial_t \bar{\rho}_t &= \nabla_{\bar{\boldsymbol{\theta}}} \cdot (\bar{\rho}_t \bar{\mathbf{H}}(t) \nabla_{\bar{\boldsymbol{\theta}}} \psi(\bar{\boldsymbol{\theta}}; \bar{\rho}_t)), \\ \psi(\bar{\boldsymbol{\theta}}; \bar{\rho}_t) &= \frac{1}{2} \mathbb{E}_{\mathbf{z}, G} \left[\{ \hat{f}_{\text{NN}}(\mathbf{z}; \bar{\rho}_t) - f_*(\mathbf{z}) \} \bar{a} \sigma(\langle \bar{\mathbf{u}}, \mathbf{z} \rangle + \bar{s} G) \right] + \frac{1}{2} \bar{\boldsymbol{\theta}}^\top \bar{\boldsymbol{\Lambda}} \bar{\boldsymbol{\theta}}, \end{aligned} \quad (\text{DF-PDE})$$

where $\bar{\mathbf{H}}(t) = \text{diag}(\xi^a(t), \xi^w(t) \mathbf{I}_{P+1})$ and $\bar{\boldsymbol{\Lambda}} = \text{diag}(\lambda^a, \lambda^w \mathbf{I}_{P+1})$. Equivalently, (DF-PDE) can be seen as a Wasserstein gradient flow over the test error $R(\bar{\rho}) = \mathbb{E}_{\mathbf{z}} [\{h_*(\mathbf{z}) - \hat{f}_{\text{NN}}(\mathbf{z}; \bar{\rho})\}^2]$ in the space $\mathbf{z} \sim \text{Unif}(\{+1, -1\}^P)$ with initialization $\bar{\rho}_0$ and regularization $\int \bar{\boldsymbol{\theta}}^\top \bar{\boldsymbol{\Lambda}} \bar{\boldsymbol{\theta}} \bar{\rho}(d\bar{\boldsymbol{\theta}})$. We put further intuition for this result in Appendix B.1.

The following theorem provides a non-asymptotic bound between the (bSGD) solution $\hat{f}_{\text{NN}}(\cdot; \hat{\rho}_t^{(N)})$ and the (DF-PDE) solution $\hat{f}_{\text{NN}}(\cdot; \bar{\rho}_t)$:

Theorem 5 *Assume conditions A0-A2, A3' hold, and let $T \geq 1$. There exist constants K_0 and K_1 depending only on the constants in A0-A2, A3' (in particular, independent of d, P, T), such that for any $b \leq d$, $N \leq e^d$, $\eta \leq e^{-K_0 T^3} b / (d + \log(N))$, we have*

$$\left\| \hat{f}_{\text{NN}}(\cdot; \boldsymbol{\Theta}^k) - \hat{f}_{\text{NN}}(\cdot; \bar{\rho}_{k\eta}) \right\|_{L^2} \leq e^{K_1 T^7} \left\{ \sqrt{\frac{P + \log(d)}{d}} + \sqrt{\frac{\log N}{N}} + \sqrt{\frac{d + \log N}{b}} \sqrt{\eta} \right\},$$

for all $k \in [T/\eta] \cap \mathbb{N}$, with probability at least $1 - 1/N$.

The proof of Thm. 5 can be found in App. B.2.1. We first extend the results in Mei et al. (2019) to bound the difference between (bSGD) and (MF-PDE) dynamics, and then we use a propagation-of-chaos argument to bound the distance between the (MF-PDE) and (DF-PDE) solutions.

Equivalence with SGD-learnability: From Theorem 5, (DF-PDE) is a good approximation of (bSGD) as long as $d, N, 1/\eta$ are taken sufficiently large while keeping $T = \eta m/b$ bounded. This leads to the equivalence described in the introduction (the proof can be found in Appendix B.2.1):

Theorem 6 *A function $h_* : \{+1, -1\}^P \rightarrow \mathbb{R}$ is strongly $O(d)$ -SGD-learnable if and only if for any $\varepsilon > 0$, there exists $\lambda^a, \lambda^w \geq 0$ and Lipschitz $\xi^a, \xi^w : \mathbb{R}_{>0} \rightarrow \mathbb{R}_{>0}$, such that $\inf_{t \geq 0} R(\bar{\rho}_t) < \varepsilon$.*

For generic activation, we have $\inf_{\bar{\rho}} R(\bar{\rho}) = 0$. Hence, Theorem 6 states that h_* is strongly $O(d)$ -SGD-learnable if and only if the global minimizer is dynamically reachable by a gradient flow initialized at $\bar{\rho}_0 = \mu_a \otimes \delta_{\bar{\mathbf{u}}^0 = \mathbf{0}} \otimes \delta_{\bar{s}^0 = c}$. See Appendix A for additional discussions and numerical illustrations. In Figure 1, we plotted a comparison between (bSGD) and (DF-PDE) for $h_*(\mathbf{z}) = z_1 + z_1 z_2 + z_1 z_2 z_3 + z_1 z_2 z_3 z_4$ and shifted sigmoid activation $\sigma(x) = (1 + e^{-x+0.5})^{-1}$. We fix




Figure 1: Comparison between (bSGD) and (DF-PDE) dynamics for $h_*(\mathbf{z}) = z_1 + z_1z_2 + z_1z_2z_3 + z_1z_2z_3z_4$. Left: Test error. Right: Fourier coefficients of $\hat{f}_{\text{NN}}(\mathbf{x}; \Theta^{t/\eta})$ and $\hat{f}_{\text{NN}}(\mathbf{z}; \bar{\rho}_t)$. The dashed-dotted black lines correspond to (DF-PDE) and the continuous colored line to (bSGD). The test errors and Fourier coefficients are evaluated with $m = 300$ test samples and for (bSGD), we report the average and 95% confidence interval over 10 experiments.

$d = N = 100$, $b = 150$, $\lambda^a = \lambda^w = 0$, $\eta_k^a = \eta_k^w = 1/2$, $\mu_a = \text{Unif}([+1, -1])$ and $\mu_w = \text{N}(0, 1)$. Let us emphasize a few prominent features of this plot: 1) The (DF-PDE) approximation tracks well (bSGD) until convergence even for moderate $d, N, b/\eta$, despite a convergence with nontrivial structure. 2) The monomials are picked up sequentially to a nonnegligible amount with increasing degree, which agrees with the intuition that lower-degree monomials guide SGD to learn higher degree monomials. 3) (DF-PDE) reaches a global minimum, which by Theorem 6 implies that h_* is strongly SGD-learnable in $O(d)$ -scaling.

MSP is necessary: We can show that the (DF-PDE) dynamics with h_* without MSP cannot reach arbitrarily small test error when initialized with $\bar{\rho}_0$. By Theorem 6, this implies that MSP is necessary for strong SGD-learnability in $O(d)$ -scaling.

Theorem 7 *Let $h_* : \{+1, -1\}^P \rightarrow \mathbb{R}$ be a function without MSP. Then there exists $c > 0$ such that for any $\xi^a, \xi^w : \mathbb{R} \rightarrow \mathbb{R}$ and regularizations $\lambda^a, \lambda^w \geq 0$, we have $\inf_{t \geq 0} R(\bar{\rho}_t) \geq c$.*

This result is based on the following simple observation: for h_* without MSP and with $\bar{\mathbf{u}}^0 = \mathbf{0}$ initialization, some coordinates stay equal to 0 throughout the dynamic, i.e., $\bar{u}_i^t = 0$. In that case, any Fourier coefficient that contains $i \in S$ is not learned: $\mathbb{E}_{\mathbf{z}}[\hat{f}_{\text{NN}}(\mathbf{z}; \bar{\rho}_t)\chi_S(\mathbf{z})] = 0$. We report the proof to Appendix B.2.3 and simply detail one example $h_*(\mathbf{z}) = z_1 + z_1z_2z_3$. Consider the first-layer weight evolution $\dot{\bar{u}}_i^t = \bar{a}^t \mathbb{E}[(z_1 + z_1z_2z_3)\sigma'(\langle \bar{\mathbf{u}}^t, \mathbf{z} \rangle)z_i]$ (for the sake of intuition, we take $\hat{f}_{\text{NN}}(\mathbf{z}; \bar{\rho}_t) = 0$ and $\bar{\mathbf{s}}^t = 0$). Notice that the evolution equations are symmetric under exchange $\bar{u}_2^t \leftrightarrow \bar{u}_3^t$ with $\bar{u}_2^0 = \bar{u}_3^0 = 0$ and therefore $\bar{u}_2^t = \bar{u}_3^t =: u_{23}^t$. Denoting $z_{23} = z_2 + z_3$ and integrating out z_{23} ,

$$\dot{u}_{23}^t = \bar{a}^t \mathbb{E}_{\mathbf{z}}[z_1 z_{23} \sigma'(z_1 \bar{u}_1^t + z_{23} u_{23}^t)] = \bar{a}^t u_{23}^t \mathbb{E}_{z_1}[z_1 \sigma''(z_1 \bar{u}_1^t + r(u_{23}^t))],$$

for some $r(u_{23}^t) \in [-2u_{23}^t, 2u_{23}^t]$ using the mean value theorem. Recalling that $u_{23}^0 = 0$, we deduce that $u_{23}^t = \bar{u}_2^t = \bar{u}_3^t = 0$.

4. Sufficient conditions for strong SGD-learnability

In the previous section, we saw that having MSP is necessary for strong $O(d)$ -SGD-learnability. Is the converse true? Is any MSP function strongly SGD-learnable in the $O(d)$ -scaling?

Degenerate cases: It turns out that one first has to exclude some special cases. Some MSP functions present degenerate dynamics due to their symmetries and are not strongly $O(d)$ -SGD-learnable. For example, take $h_*(\mathbf{z}) = z_1 + z_2 + z_1z_3 + z_2z_4$, which is invariant by permutation $(1, 2, 3, 4) \leftrightarrow (2, 1, 4, 3)$ of its input coordinates. During the (DF-PDE) dynamics, $\bar{u}_1^t = \bar{u}_2^t$ and $\bar{u}_3^t = \bar{u}_4^t$, which implies that a solution with $\mathbb{E}_{\mathbf{z}}[\hat{f}_{\text{NN}}(\mathbf{z}; \bar{\rho}_t)z_1z_3] = \mathbb{E}_{\mathbf{z}}[\hat{f}_{\text{NN}}(\mathbf{z}; \bar{\rho}_t)z_2z_3]$ is found and therefore the risk of the (DF-PDE) dynamics is always bounded away from zero. See Section A for numerical simulations and further discussion on degenerate MSPs.

Generic MSP functions are learnable: To bypass this difficulty, we prove a learnability result that holds for “generic” MSP functions – i.e., that holds almost surely over a random choice of non-zero Fourier coefficients. Formally, for any set structure $\mathcal{S} = \{S_1, \dots, S_m\} \subseteq 2^{[P]}$, let us define a measure over functions that have those Fourier coefficients.

Definition 8 For any set structure $\mathcal{S} \subseteq 2^{[P]}$ define the measure $\mu_{\mathcal{S}}$ over functions $h_* : \{+1, -1\}^P \rightarrow \mathbb{R}$ induced by taking $h_*(\mathbf{z}) = \sum_{S \subseteq [P]} \alpha_S \chi_S(\mathbf{x})$, where the Fourier coefficients satisfy $\alpha_S = 0$ if $S \notin \mathcal{S}$, and $(\alpha_S)_{S \in \mathcal{S}}$ have Lebesgue measure on $\mathbb{R}^{|\mathcal{S}|}$.

Our main sufficiency result shows that the degenerate cases are a measure-zero set. In this sense, there are very few bad examples, and so MSP structure is “nearly” sufficient for strong $O(d)$ -SGD-learnability.

Theorem 9 For any MSP set structure $\mathcal{S} \subseteq 2^{[P]}$, h_* is strongly $O(d)$ -SGD-learnable almost surely with respect to $\mu_{\mathcal{S}}$, using activation function $\sigma(x) = (1+x)^L$ where $L = 2^{8P}$.⁴

The converse to this result is implied by the necessity result of the previous section, which states that for any h_* with non-zero Fourier coefficients (set structure) \mathcal{S} that is *not* MSP, h_* is *not* strongly $O(d)$ -SGD-learnable. While we prove Theorem 9 for a particular activation, we note that the proof implies that the same is true for any degree- L polynomial activation almost surely over its $(L+1)$ -coefficients (see Theorem 38 in Appendix E). In Appendix F we show how this result extends to generic smooth (non-polynomial) activations as long as a certain polynomial is not identically 0 for a given set structure (which we show with a small technical caveat).

Vanilla staircase, learnable without genericity: In the special case of functions with “vanilla staircase” structure we do not need a genericity assumption, and we require weaker assumptions on the activation function.

Theorem 10 Let h_* be of the form $h_*(\mathbf{z}) = \alpha_1z_1 + \alpha_2z_1z_2 + \dots + \alpha_Pz_1z_2 \cdots z_P$ where $\alpha_i \neq 0$ for $i \in [P]$. Then h_* is strongly $O(d)$ -SGD-learnable using any activation function $\sigma \in \mathcal{C}^{2^{P-1}+1}(\mathbb{R})$ with nonzero derivatives $\sigma^{(r)}(0) \neq 0$ for $r = 0, \dots, P$.

4. Technically speaking, for the strong SGD-learnability definition we cannot take $\sigma(x) = (1+x)^L$ as it is not bounded. However, we take an activation function that equals $(1+x)^L$ on the interval $(-1, 1)$ and is bounded elsewhere.

Proof ideas The proofs for Theorems 9 and 10 follow a similar approach. From the equivalence stated in Theorem 6, it is sufficient to display, for each $\varepsilon > 0$, hyperparameters such that the (DF-PDE) dynamics reaches ε -risk. We choose $\lambda^a = \lambda^w = 0$ (no regularization) and initialization $\mu_a = \text{Unif}([-1, +1])$ and $\mu_w = \delta_0$ (this choice simplifies the analysis as $\bar{s}^t = \bar{s}^0 = 0$). We split the learning in two phases: in *Phase 1*, we train the first layer weights $\bar{\mathbf{u}}^t$ for time $t \in [0, T_1]$ while keeping $\bar{a}^t = \bar{a}^0$ fixed, and in *Phase 2*, we train the second layer weights \bar{a}^t for time $t \in [T_1, T_2]$ while keeping $\bar{\mathbf{u}}^t = \bar{\mathbf{u}}^{T_1}$ fixed.

At the end of Phase 1, denote $(\bar{a}^0, \bar{\mathbf{u}}^{T_1}(\bar{a}^0))$ the weights obtained from the evolution (DF-PDE) from initialization $(\bar{a}^0, \bar{\mathbf{u}}^0 = \mathbf{0})$ (note that $\bar{\mathbf{u}}^{T_1}(\bar{a}^0)$ is a deterministic function of \bar{a}^0). Phase 2 corresponds to a linear training phase with kernel $K^{T_1}(\mathbf{z}, \mathbf{z}') = \mathbb{E}_{\bar{a}^0 \sim \mu_a} [\sigma(\langle \bar{\mathbf{u}}^{T_1}(\bar{a}^0), \mathbf{z} \rangle) \sigma(\langle \bar{\mathbf{u}}^{T_1}(\bar{a}^0), \mathbf{z}' \rangle)]$. In particular, the risk decreases as $\exp(-\lambda_{\min}(\mathbf{K}^{T_1})t)$ during this phase, where we denote by $\mathbf{K}^{T_1} = (K^{T_1}(\mathbf{z}, \mathbf{z}'))_{\mathbf{z}, \mathbf{z}' \in \{+1, -1\}^P}$ the kernel matrix. Showing global convergence reduces to showing that $\lambda_{\min}(\mathbf{K}^{T_1}) > 0$ for some T_1 and taking $T_2 = T_1 + \log(1/\varepsilon)/\lambda_{\min}(\mathbf{K}^{T_1})$.

The goal of the analysis in Phase 1 is therefore to prove this lower bound on the eigenvalues of the kernel matrix. Phase 1 corresponds to a nonlinear dynamics, and is a priori unclear how to analyze. In the case of vanilla staircases, we show that it is enough to track the leading order in t for each coordinates $(\bar{u}_i^t)_{i \in [P]}$ and take T_1 small enough. For example, when learning $h_*(\mathbf{z}) = z_1 + z_1 z_2 + z_1 z_2 z_3$, for small time $t \leq T_1$ we can roughly show that $\bar{u}_1^t(a) \propto at$, that $\bar{u}_2^t(a) \propto at^2$, and that $\bar{u}_3^t(a) \propto at^4$. In other words, the weight corresponding to z_1 increases in magnitude the fastest, followed by the weight corresponding to z_2 , and then weight corresponding to z_3 . We can then use this explicit calculation to lower bound the eigenvalues of \mathbf{K}^{T_1} , crucially using that the second-layer weights $\bar{a}^0 \sim \mu_a$ are chosen at random, which ensures that the neurons are diverse enough. See Appendix D for the detailed proof.

For general MSP set structure, it is not enough to only track the weights $\bar{\mathbf{u}}^t(\bar{a}^0)$ to leading order in t . We show instead that it suffices to lower bound a kernel matrix $\hat{\mathbf{K}}^{T_1}$ obtained from a simplified dynamics $\hat{\mathbf{u}}^t(\bar{a}^0)$. The weights $\hat{u}_i^t(\bar{a}^0)$ can be written in terms of polynomials in the second-layer weights \bar{a}^0 , the Fourier coefficients $(\alpha_S)_{S \in \mathcal{S}}$ and the derivatives of the activation $(\sigma^{(r)}(0))_{r=0, \dots, L}$, with coefficients defined explicitly by a recurrence relation and only depending on the set structure \mathcal{S} . Using algebraic facts about the linear independence of large powers of polynomials we show that $\det(\hat{\mathbf{K}}^{T_1})$ is a nonzero polynomial in the second-layer weights and Fourier coefficients. Therefore, plugging in random second-layer weights $\bar{a}^0 \sim \text{Unif}([-1, 1])$, and random Fourier coefficients we show that $\det(\hat{\mathbf{K}}^{T_1}) \neq 0$ almost surely, by anti-concentration of polynomials. This implies in particular that $\sum_{S \in \mathcal{S}} \alpha_S$ is almost surely strongly $O(d)$ -SGD-learnable. See Appendix E.

5. Separation with linear methods

It is known that linear methods with $\text{poly}(d)$ many features or samples cannot learn the class of degree- P monomials if P grows with the input dimension d (Hsu et al., 2021; Hsu; Kamath et al., 2020). One way of proving this is by using SQ lower bounds (Blum et al., 1994), which imply lower bounds on linear methods (Kamath et al., 2020). However, this proof strategy fails for staircase functions of growing degree, since the hierarchical structure makes these efficiently SQ learnable by sequentially querying the monomials of increasing degree⁵. We thus need a lower-bound on linear methods that goes beyond SQ lower-bounds, which we obtain by using subspace projections.

5. Making at most d queries per degree for vanilla staircases, e.g., at most $d \log d$ queries when $P = \log d$.

Consider a general linear method which is defined by a Hilbert space $(\mathcal{H}, \langle \cdot, \cdot \rangle_{\mathcal{H}})$, a feature map $\psi : \{+1, -1\}^d \rightarrow \mathcal{H}$, an empirical loss function $L : \mathbb{R}^{2n} \rightarrow \mathbb{R} \cup \infty$ and a regularization parameter $\lambda > 0$. Given data points $(y_i, \mathbf{x}_i)_{i \in [n]}$, the linear method constructs a prediction model $\hat{f}(\mathbf{x}) := \langle \hat{\mathbf{a}}, \psi(\mathbf{x}) \rangle_{\mathcal{H}}$ where $\hat{\mathbf{a}} \in \mathcal{H}$ is obtained by minimizing the regularized empirical risk functional

$$\hat{\mathbf{a}} = \arg \min_{\mathbf{a} \in \mathcal{H}} \left\{ L((y_i, \langle \mathbf{a}, \psi(\mathbf{x}_i) \rangle))_{i \in [n]} + \lambda \|\mathbf{a}\|_{\mathcal{H}}^2 \right\}. \quad (4)$$

We will further denote $q = \dim(\mathcal{H})$. Popular examples include random feature models (q is equal to the number of random features) and kernel methods ($q = \infty$ typically). While the optimization problem (4) is over a (potentially) infinite dimensional space \mathcal{H} , it is an easy exercise to verify that $\hat{\mathbf{a}} \in \text{span}\{\psi(\mathbf{x}_i) : i \in [n]\}$ which has dimension bounded by $\min(n, q)$.

We consider learning a class of functions $\mathcal{F}_M = \{f_1, \dots, f_M\} \subseteq L^2(\mathcal{X})$ from n evaluations at points $(\mathbf{x}_i)_{i \in [n]}$. For any linear method which, for $j \in [M]$, outputs the model \hat{f}_j obtained by (4) on $(f_j(\mathbf{x}_i) + \varepsilon_{ij}, \mathbf{x}_i)_{i \in [n]}$, we define the average prediction error on \mathcal{F} as $\bar{R}_n(\mathcal{F}_M) = \frac{1}{M} \sum_{j \in [M]} \mathbb{E}_{\mathcal{X}}[(f_j(\mathbf{x}) - \hat{f}_j(\mathbf{x}))^2]$. Several lower bounds on the sample complexity have appeared in the literature such as Hsu et al. (2021); Hsu; Kamath et al. (2020), here we present an improvement on Hsu et al. (2021); Hsu that is tighter for target functions that are not almost orthogonal, and an improvement of Kamath et al. (2020) that is tighter for functions like vanilla staircases of growing degree (which have polynomial SQ-dimension).

Proposition 11 *Let $\Omega \subseteq L^2(\mathcal{X})$ a linear subspace. Let $\mathcal{F}_M = \{f_1, \dots, f_M\} \subset L^2(\mathcal{X})$ such that $\|\mathbf{P}_{\Omega} f_j\|_{L^2}^2 = 1 - \kappa$ and $\|\mathbf{P}_{\Omega}^{\perp} f_j\|_{L^2}^2 = \kappa$ for all $f_j \in \mathcal{F}_M$. For any linear method, if $\bar{R}_n(\mathcal{F}_M) \leq 1 - \eta$, then we must have*

$$\min(n, q) \geq \frac{\eta - \kappa}{\max_{i \in [M]} \frac{1}{M} \sum_{j \in [M]} |\langle f_i, \mathbf{P}_{\Omega} f_j \rangle_{L^2}|}. \quad (5)$$

Define Ω_k to be the subspace spanned by all degree- k monomials $L^2(\{+1, -1\}^P)$. Consider a function $h_{\text{Poly-}k} \in L^2(\{+1, -1\}^P)$ such that $\|\mathbf{P}_{\Omega_k} h_{\text{Poly-}k}\|_{L^2}^2 = \|\mathbf{P}_{\Omega_k}^{\perp} h_{\text{Poly-}k}\|_{L^2}^2 = 1/2$ and $\mathbf{P}_{\Omega_k} h_{\text{Poly-}k}$ is supported over m monomials. Second, consider $h_{\text{Str-}P}(\mathbf{z}) = P^{-1/2}(z_1 + z_1 z_2 + \dots + z_1 \dots z_P)$ the degree- P staircase. We consider two sets of functions obtained by $h_{\text{Poly-}k}$ and $h_{\text{Str-}P}$ with all the permutations of their input signal: $\mathcal{F}_{\ell} = \left\{ f = h_{\ell} \circ \tau : \tau \in \Pi(d) \right\}, \ell \in \{\text{Poly-}k, \text{Str-}P\}$, where $\Pi(d)$ corresponds to the group of all permutation on $[d]$. Applying Proposition 11, we get the following sample-complexity lower bounds:

Proposition 12 *For any linear method, if $\bar{R}_n(\mathcal{F}_{\text{Poly-}k}) \leq 1/2 \cdot (1 - \eta)$ then we must have $\min(n, q) \geq \frac{\eta}{m} \binom{d}{k}$. Similarly, if $\bar{R}_n(\mathcal{F}_{\text{Str-}P}) \leq 1 - \eta$ then we must have $\min(n, q) \geq \frac{\eta}{2} \binom{d}{\lfloor \frac{d}{2} \rfloor}$.*

Note that kernel and random features methods achieve the lower bound for $\mathcal{F}_{\text{Poly-}k}$ (Ghorbani et al., 2021a; Mei et al., 2021). Comparing Proposition 11 with the result of Section 4, we get the following separation results between SGD-trained neural networks and linear methods:

- (1) SGD on two-layer neural networks outperforms linear methods almost surely on non-linear MSP functions ($n = O(d)$ versus $n = \Omega_d(d^k)$ for degree- k MSP).
- (2) We obtain from Proposition 12 that for any $P = \omega_d(1)$, linear methods must have $\min(n, q) = d^{\omega_d(1)}$ to learn the vanilla staircase of degree P , while Theorem 10 can still guarantee a sample complexity of $d^{O(1)}$ for P growing slowly enough with d .

6. Conclusion and future directions

In this paper, we considered learning sparse functions in arbitrarily large ambient dimension, using two-layer neural networks trained by batch-SGD in the mean-field regime. We proved that the merged-staircase property is a necessary and nearly-sufficient condition for such functions to be learnable on such models in $O(d)$ sample-complexity. The near sufficiency part, which excludes a measure-zero subset, is unavoidable as there exist symmetric MSP functions with degenerate dynamics that are not strongly SGD-learnable in $O(d)$ -scaling. This provides a regime where one can achieve a tight characterization of functions that are learnable by regular SGD on regular neural networks, while going beyond the linear regime.

One venue for future work is to characterize more precisely the set of degenerate MSP functions: current examples correspond to MSP functions with some group invariance (see Appendix A) which arise naturally in applications. More importantly, the current MSP condition hinges on the particular setting considered in this paper: P fixed (or sufficiently slowly growing), $O(d)$ sample-complexity and two-layer neural networks. In particular, this definition ignores composition-order, i.e., how many monomials are composed to create a new one, and how many fresh variables are involved in such a composition. We anticipate more complex categories to appear as we move away from this setting. For example, we conjecture that l -leap MSP (i.e., $|S_i \setminus \cup_{j=1}^{i-1} S_j| \leq l$) are SGD-learnable in $O(\varphi_l(d))$ -scaling with $\varphi_2(d) = d \log(d)$ and $\varphi_l(d) = \tilde{O}(d^{l/2})$ for $l > 2$ (this paper considers $l = 1$ and showed $\varphi_1(d) = d$ is tight in the mean-field regime). Furthermore, the compositionality-order, i.e., the number k of monomials that can be composed in order to produce new monomials, should also be factored in for a finer complexity analysis once P (and k) are no longer constant. The depth of the architecture is also expected to play a role when P is diverging: for instance, it is shown in Abbe et al. (2021a) that vanilla staircases (i.e., nested chains with $k = l = 1$) are learnable with P -layer neural networks (but unconventional gradient-based training) in $\text{poly}(d, P, 1/\varepsilon)$ samples, while our proof techniques (Theorem 9 and Theorem 10, see Appendix G for the statement of the explicit bounds) yield $O(e^{e^P} d/\varepsilon^C)$ sample-complexity⁶ with two-layer neural networks. More generally, the sparsity parameter P will not be necessarily the right complexity measure for deeper networks: for example, some functions with small leap l and large P will be easier to learn than some functions with smaller P but larger l .

Finally, it is natural to seek counterparts of the results in this work and counterparts of the staircase notions for other Hilbert spaces, such as L^2 functions with respect to the Gaussian measure.

Acknowledgements

We thank Guy Bresler, Dheeraj Nagaraj, and Nati Srebro for stimulating discussions. We thank the Simons Foundations and the NSF for supporting us through the Collaboration on the Theoretical Foundations of Deep Learning (deepfoundations.ai). This work was done (in part) while E.B. and T.M. were visiting the Simons Institute for the Theory of Computing and while E.B. was visiting the Bernoulli Center at EPFL.

6. One can show $e^{O(P)}$ -dependency is sufficient if all degree-1 monomials are included in the MSP set structure – see Appendix G.3. We conjecture that such an exponential scaling is needed, i.e., deeper neural networks are required to achieve a $\text{poly}(P)$ dependency.

References

- Emmanuel Abbe and Colin Sandon. On the universality of deep learning. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 20061–20072. Curran Associates, Inc., 2020. URL <https://proceedings.neurips.cc/paper/2020/file/e7e8f8e5982b3298c8addedf6811d500-Paper.pdf>.
- Emmanuel Abbe, Enric Boix-Adserà, Matthew S. Brennan, Guy Bresler, and Dheeraj Nagaraj. The staircase property: How hierarchical structure can guide deep learning. *CoRR*, abs/2108.10573, 2021a. URL <https://arxiv.org/abs/2108.10573>.
- Emmanuel Abbe, Prithish Kamath, Eran Malach, Colin Sandon, and Nathan Srebro. On the power of differentiable learning versus PAC and SQ learning. *CoRR*, abs/2108.04190, 2021b. URL <https://arxiv.org/abs/2108.04190>.
- Zeyuan Allen-Zhu and Yuanzhi Li. What can resnet learn efficiently, going beyond kernels? 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., 2019.
- Zeyuan Allen-Zhu and Yuanzhi Li. Backward feature correction: How deep learning performs deep learning. *arXiv preprint arXiv:2001.04413*, 2020.
- Zeyuan Allen-Zhu and Yuanzhi Li. Backward feature correction: How deep learning performs deep learning. *arXiv*, abs/2001.04413, 2020. URL <https://arxiv.org/abs/2001.04413>.
- Zeyuan Allen-Zhu, Yuanzhi Li, and Yingyu Liang. Learning and generalization in overparameterized neural networks, going beyond two layers. In *Proceedings of the 33rd International Conference on Neural Information Processing Systems*, pages 6158–6169, 2019a.
- Zeyuan Allen-Zhu, Yuanzhi Li, and Zhao Song. A convergence theory for deep learning via overparameterization. In *International Conference on Machine Learning*, pages 242–252. PMLR, 2019b.
- 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, 2019.
- Francis Bach. Breaking the curse of dimensionality with convex neural networks. *The Journal of Machine Learning Research*, 18(1):629–681, 2017.
- Francis Bach and Eric Moulines. Non-strongly-convex smooth stochastic approximation with convergence rate $o(1/n)$. *Advances in neural information processing systems*, 26, 2013.
- Andrew R Barron. Universal approximation bounds for superpositions of a sigmoidal function. *IEEE Transactions on Information theory*, 39(3):930–945, 1993.
- Ronen Basri, David Jacobs, Yoni Kasten, and Shira Kritchman. The convergence rate of neural networks for learned functions of different frequencies. *arXiv preprint arXiv:1906.00425*, 2019.

- Avrim Blum, Merrick L. Furst, Jeffrey C. Jackson, Michael J. Kearns, Yishay Mansour, and Steven Rudich. Weakly learning DNF and characterizing statistical query learning using fourier analysis. In *Proceedings of the Twenty-Sixth Annual ACM Symposium on Theory of Computing, 23-25 May 1994, Montréal, Québec, Canada*, pages 253–262, 1994. URL <https://doi.org/10.1145/195058.195147>.
- Avrim Blum, Adam Kalai, and Hal Wasserman. Noise-tolerant learning, the parity problem, and the statistical query model. *J. ACM*, 50(4):506–519, July 2003. ISSN 0004-5411. doi: 10.1145/792538.792543. URL <http://doi.acm.org/10.1145/792538.792543>.
- Yuan Cao, Zhiying Fang, Yue Wu, Ding-Xuan Zhou, and Quanquan Gu. Towards understanding the spectral bias of deep learning. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*, pages 2205–2211, 2021.
- Anthony Carbery and James Wright. Distributional and l_q norm inequalities for polynomials over convex bodies in \mathbb{R}^n . *Mathematical research letters*, 8(3):233–248, 2001.
- Michael Celentano, Theodor Misiakiewicz, and Andrea Montanari. Minimum complexity interpolation in random features models. *arXiv preprint arXiv:2103.15996*, 2021.
- Lénaïc Chizat and Francis Bach. On the global convergence of gradient descent for over-parameterized models using optimal transport. *Advances in Neural Information Processing Systems*, 31:3036–3046, 2018.
- Lénaïc Chizat and Francis Bach. Implicit bias of gradient descent for wide two-layer neural networks trained with the logistic loss. In *Conference on Learning Theory*, pages 1305–1338. PMLR, 2020.
- Amit Daniely and Eran Malach. Learning parities with neural networks. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.
- Simon Du, Jason Lee, Haochuan Li, Liwei Wang, and Xiyu Zhai. Gradient descent finds global minima of deep neural networks. In *International Conference on Machine Learning*, pages 1675–1685. PMLR, 2019.
- Simon S Du, Xiyu Zhai, Barnabas Poczos, and Aarti Singh. Gradient descent provably optimizes over-parameterized neural networks. In *International Conference on Learning Representations*, 2018.
- Vitaly Feldman. A general characterization of the statistical query complexity. *arXiv preprint arXiv:1608.02198*, 2016.
- Vitaly Feldman, Elena Grigorescu, Lev Reyzin, Santosh S. Vempala, and Ying Xiao. Statistical algorithms and a lower bound for detecting planted cliques. *J. ACM*, 64(2):8:1–8:37, April 2017a. ISSN 0004-5411. doi: 10.1145/3046674. URL <http://doi.acm.org/10.1145/3046674>.

- Vitaly Feldman, Cristóbal Guzmán, and Santosh Vempala. Statistical query algorithms for mean vector estimation and stochastic convex optimization. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '17, pages 1265–1277, Philadelphia, PA, USA, 2017b. Society for Industrial and Applied Mathematics. URL <http://dl.acm.org/citation.cfm?id=3039686.3039768>.
- 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, nov 2020. doi: 10.1088/1742-5468/abc4de. URL <https://doi.org/10.1088/1742-5468/abc4de>.
- Behrooz Ghorbani, Song Mei, Theodor Misiakiewicz, and Andrea Montanari. Limitations of lazy training of two-layers neural networks. In *Proceedings of the 33rd International Conference on Neural Information Processing Systems*, pages 9111–9121, 2019.
- Behrooz Ghorbani, Song Mei, Theodor Misiakiewicz, and Andrea Montanari. Linearized two-layers neural networks in high dimension. *The Annals of Statistics*, 49(2):1029–1054, 2021a.
- Behrooz Ghorbani, Song Mei, Theodor Misiakiewicz, and Andrea Montanari. When do neural networks outperform kernel methods? *Journal of Statistical Mechanics: Theory and Experiment*, 2021(12):124009, 2021b.
- Surbhi Goel, Aravind Gollakota, Zhihan Jin, Sushrut Karmalkar, and Adam Klivans. Superpolynomial lower bounds for learning one-layer neural networks using gradient descent. 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 3587–3596. PMLR, 13–18 Jul 2020. URL <https://proceedings.mlr.press/v119/goel20a.html>.
- Daniel Hsu. Dimension lower bounds for linear approaches to function approximation.
- Daniel Hsu, Clayton Sanford, Rocco A Servedio, and Emmanouil-Vasileios Vlatakis-Gkaragkounis. On the approximation power of two-layer networks of random relus. *arXiv preprint arXiv:2102.02336*, 2021.
- 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.
- Prateek Jain, Sham M Kakade, Rahul Kidambi, Praneeth Netrapalli, Venkata Krishna Pillutla, and Aaron Sidford. A markov chain theory approach to characterizing the minimax optimality of stochastic gradient descent (for least squares). *arXiv preprint arXiv:1710.09430*, 2017.
- Prateek Jain, Sham Kakade, Rahul Kidambi, Praneeth Netrapalli, and Aaron Sidford. Parallelizing stochastic gradient descent for least squares regression: mini-batching, averaging, and model misspecification. *Journal of Machine Learning Research*, 18, 2018.
- Prateek Jain, Dheeraj Nagaraj, and Praneeth Netrapalli. Making the last iterate of sgd information theoretically optimal. In *Conference on Learning Theory*, pages 1752–1755. PMLR, 2019.

- Pritish Kamath, Omar Montasser, and Nathan Srebro. Approximate is Good Enough: Probabilistic Variants of Dimensional and Margin Complexity. In Jacob Abernethy and Shivani Agarwal, editors, *Proceedings of Thirty Third Conference on Learning Theory*, volume 125 of *Proceedings of Machine Learning Research*, pages 2236–2262. PMLR, 09–12 Jul 2020. URL <https://proceedings.mlr.press/v125/kamath20b.html>.
- Stefani Karp, Ezra Winston, Yuanzhi Li, and Aarti Singh. Local signal adaptivity: Provable feature learning in neural networks beyond kernels. *Advances in Neural Information Processing Systems*, 34, 2021.
- Michael Kearns. Efficient noise-tolerant learning from statistical queries. *Journal of the ACM (JACM)*, 45(6):983–1006, 1998.
- Yuanzhi Li and Yingyu Liang. Learning overparameterized neural networks via stochastic gradient descent on structured data. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, pages 8168–8177, 2018.
- Yuanzhi Li, Tengyu Ma, and Hongyang R. Zhang. Learning over-parametrized two-layer neural networks beyond NTK. In *Conference on Learning Theory, COLT 2020, 9-12 July 2020, Virtual Event [Graz, Austria]*, volume 125 of *Proceedings of Machine Learning Research*, pages 2613–2682. PMLR, 2020. URL <http://proceedings.mlr.press/v125/li20a.html>.
- Eran Malach and Shai Shalev-Shwartz. The implications of local correlation on learning some deep functions. *Advances in Neural Information Processing Systems*, 33, 2020.
- Eran Malach, Pritish Kamath, Emmanuel Abbe, and Nathan Srebro. Quantifying the benefit of using differentiable learning over tangent kernels. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, pages 7379–7389. PMLR, 2021. URL <http://proceedings.mlr.press/v139/malach21a.html>.
- 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.
- Song Mei, Theodor Misiakiewicz, and Andrea Montanari. Mean-field theory of two-layers neural networks: dimension-free bounds and kernel limit. In *Conference on Learning Theory*, pages 2388–2464. PMLR, 2019.
- Song Mei, Theodor Misiakiewicz, and Andrea Montanari. Generalization error of random feature and kernel methods: Hypercontractivity and kernel matrix concentration. *Applied and Computational Harmonic Analysis*, 2021. ISSN 1063-5203. doi: <https://doi.org/10.1016/j.acha.2021.12.003>.
- DJ Newman and Morton Slater. Waring’s problem for the ring of polynomials. *Journal of Number Theory*, 11(4):477–487, 1979.
- Phan-Minh Nguyen and Huy Tuan Pham. A rigorous framework for the mean field limit of multi-layer neural networks. *arXiv preprint arXiv:2001.11443*, 2020.

- 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.
- Maria Refinetti, Sebastian Goldt, Florent Krzakala, and Lenka Zdeborova. Classifying high-dimensional gaussian mixtures: Where kernel methods fail and neural networks succeed. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 8936–8947. PMLR, 18–24 Jul 2021.
- Grant M Rotskoff and Eric Vanden-Eijnden. Neural networks as interacting particle systems: Asymptotic convexity of the loss landscape and universal scaling of the approximation error. *stat*, 1050:22, 2018.
- Johannes Schmidt-Hieber. Nonparametric regression using deep neural networks with relu activation function. *The Annals of Statistics*, 48(4):1875–1897, 2020.
- Ohad Shamir and Tong Zhang. Stochastic gradient descent for non-smooth optimization: Convergence results and optimal averaging schemes. In *International conference on machine learning*, pages 71–79. PMLR, 2013.
- Justin Sirignano and Konstantinos Spiliopoulos. Mean field analysis of neural networks: A central limit theorem. *Stochastic Processes and their Applications*, 130(3):1820–1852, 2020.
- Jacob Steinhardt, Gregory Valiant, and Stefan Wager. Memory, communication, and statistical queries. In *Electronic Colloquium on Computational Complexity*, 2015.
- Taiji Suzuki and Shunta Akiyama. Benefit of deep learning with non-convex noisy gradient descent: Provable excess risk bound and superiority to kernel methods. *arXiv preprint arXiv:2012.03224*, 2020.
- Roman Vershynin. Introduction to the non-asymptotic analysis of random matrices. *arXiv preprint arXiv:1011.3027*, 2010.
- 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.
- Ke Yang. New lower bounds for statistical query learning. *Journal of Computer and System Sciences*, 70(4):485–509, 2005.
- Gilad Yehudai and Ohad Shamir. On the power and limitations of random features for understanding neural networks. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, 8-14 December 2019, Vancouver, BC, Canada*, pages 6594–6604, 2019.
- Difan Zou, Yuan Cao, Dongruo Zhou, and Quanquan Gu. Gradient descent optimizes overparameterized deep relu networks. *Machine Learning*, 109(3):467–492, 2020.

Contents

1	Introduction	1
1.1	Summary of main results	4
1.2	Further related literature	5
2	Strong SGD-learnability in $O(d)$-scaling	6
3	Continuous dimension-free dynamics and necessary condition	8
4	Sufficient conditions for strong SGD-learnability	11
5	Separation with linear methods	12
6	Conclusion and future directions	14
A	Additional numerical simulations	22
B	Proofs for continuous mean-field and dimension-free dynamics	25
B.1	Justification for the dimension-free dynamics equations	25
B.2	Proof of the results in Section 3	27
B.2.1	Proof of Theorem 5	27
B.2.2	Proof of Theorem 6	27
B.2.3	Proof of Theorem 7	28
B.3	Proof of Proposition 15	29
B.4	Proof of Theorem 16: bound between (MF-PDE) and (DF-PDE) dynamics	30
B.4.1	Proof of Lemma 18	32
B.5	Auxiliary lemmas	37
C	Strong SGD-learnability in the discrete-time regime	40
C.1	Discrete time mean-field and dimension-free dynamics	40
C.2	Proof of Theorem 23	41
C.3	Auxiliary lemma	42
D	Vanilla staircase functions are strongly $O(d)$-SGD-learnable: Proof of Theorem 10	44
D.1	Outline of the proof	45
D.2	Approximating the u_i^t with polynomials	47
E	Generic MSP functions are strongly $O(d)$-SGD-learnable: Proof of Theorem 9 (discrete-time regime)	50
E.1	Outline of the proof	51
E.1.1	Phase 2 (linear training)	52
E.1.2	Phase 1 (nonlinear training)	52
E.2	Proof of Lemmas 39, 40, and 41	56
E.3	Proof of Lemma 42	57
E.4	Proof of Lemma 43	57
E.5	Proof of Lemma 45	58

E.5.1	Reducing to minimal MSP set structures	58
E.5.2	Computing the weights to leading order	58
E.5.3	Applying linear independence of powers of polynomials	60
F	Generic MSP functions are strongly $O(d)$-SGD-learnable with continuous-time dynamics and activation perturbation	62
F.1	Statement of the result	62
F.2	Discussion on the perturbation of the activation	63
F.3	Outline of the proof	63
F.4	Approximating the u_i^t with polynomials	65
F.5	Simplified dynamics without interaction term	68
F.6	Recurrence relation of the coefficients in the simplified dynamics	69
F.7	Reduction to analyzing the simplified dynamics	70
F.8	Proving learnability of generic MSP functions, Theorem 56	75
F.9	Proof of Lemma 73	76
G	Explicit sample-complexity bounds in all parameters	81
G.1	Vanilla staircase functions	81
G.1.1	Technical lemma	82
G.2	Merged-staircase functions in the smoothed complexity model	83
G.3	Merged-staircase functions when including all degree-1 monomials	84
G.3.1	The bump and gradient bump functions	84
G.3.2	Statement of Proposition 86 and Corollary 87	86
G.3.3	Proof of Proposition 86 and Corollary 87	86
H	Lower bounds on learning with linear methods	90
I	Technical results	94
I.1	Bound between batch-SGD and discrete mean-field dynamics	94
I.2	Discrete-time analysis of SGD on second layer	95
I.3	Bounding the loss reduces to lower bounding the eigenvalues of the kernel matrix	99
I.4	Anti-concentration of polynomials	104

Appendix A. Additional numerical simulations

In this Appendix, we provide further background and numerical illustrations on the strong $O(d)$ -SGD learning setting, merged-staircase functions and the dimension-free dynamics.

Global convergence of the dimension-free dynamics. As stated in Theorem 6, a function $h_* : \{+1, -1\}^P \rightarrow \mathbb{R}$ is strongly $O(d)$ -SGD-learnable if and only if a Wasserstein gradient flow on $R(\bar{\rho}) = \mathbb{E}_{\mathbf{z}}[(h_*(\mathbf{z}) - \hat{f}_{\text{NN}}(\mathbf{z}; \bar{\rho}))^2]$ (the (DF-PDE) dynamics) can converge to the global optimizer when initialized with $\bar{\mathbf{a}}^0 \sim \mu_a$, $\bar{\mathbf{u}}^0 = \mathbf{0}$ and $\bar{s}^0 = m_2^w$ (in this paragraph, we consider the continuous regime, as this is the regime that has attracted the most attention). Showing global convergence results for such dynamics is generically challenging: $R(\bar{\rho})$ presents many bad stationary points (e.g., measures $\bar{\rho}$ that are distributed on an insufficient number of atoms to represent h_*). While some progress has been made to show such results in the context of mean-field neural networks (Chizat and Bach, 2018; Nguyen and Pham, 2020; Wojtowytsch, 2020), existing global convergence results assume typically that (1) the dynamics converges to a limiting distribution as $t \rightarrow \infty$; and (2) we are given a good “spread-out” initialization. Condition (2) usually holds for initialization with bounded density on an open set around 0. Condition (1) is more challenging and presents counter-examples that are difficult to rule out. In Section 4, we avoid these difficulties by considering layer-wise training: global convergence reduces to showing that the final linear-training phase converges to 0-risk, which is implied by a certain kernel matrix being full rank.

While global convergence proofs are challenging, the (DF-PDE) dynamics is a low dimensional problem and can be efficiently solved numerically. In the rest of this section, we provide a few numerical simulations to illustrate phenomena alluded to in the main text. We will fix the activation to be a shifted sigmoid $\sigma(x) = (1 + e^{-x+1})^{-1}$, and choose learning schedules $\xi^a(t) = \xi^w(t) = 1$, zero regularization parameters $\lambda^a = \lambda^w = 0$, and initialization $\mu_a = \text{Unif}([+1, -1])$ and $\bar{s}^0 = 1$. In Figure 2, we consider four MSP functions and plot the evolution of their Fourier coefficients during the (DF-PDE) dynamics. In particular, the two top row examples, $h_1(\mathbf{z}) = z_1 + z_1 z_2 + z_3 + z_1 z_2 z_3 z_4$ and $h_2(\mathbf{z}) = z_1 + z_1 z_2 + z_2 z_3 + z_3 z_4 + z_1 z_2 z_3 z_4$, converge to the global minimum and are therefore strongly $O(d)$ -SGD-learnable. The bottom row examples, $h_3(\mathbf{z}) = z_1 + z_1 z_2 + z_3 + z_3 z_4$ and $h_4(\mathbf{z}) = z_1 + z_2 + z_3 + z_1 z_2 z_3$, do not converge and have risks bounded away from 0. Functions h_3 and h_4 are two examples of G -invariant MSP functions.

G -invariant MSP functions. We call h_* a G -invariant MSP function if h_* is invariant under a group of transformations, i.e., there exists $\tau : \{+1, -1\}^P \rightarrow \{+1, -1\}^P$ (invertible and $\tau \neq \text{id}$) such that $h_*(\tau(\mathbf{z})) = h_*(\mathbf{z})$. For these functions, the $\bar{\mathbf{u}}^t$ weight distribution remains invariant by this same group of transformations during the (DF-PDE) dynamics, regardless of the choice of parameters. For example, h_3 is invariant by permutation $(1, 2, 3, 4) \leftrightarrow (3, 4, 1, 2)$ of its input, and $\bar{u}_1^t = \bar{u}_3^t$, $\bar{u}_2^t = \bar{u}_4^t$ for all $t \geq 0$. h_4 is invariant by any permutation, and $\bar{u}_1^t = \bar{u}_2^t = \bar{u}_3^t$ for all $t \geq 0$. For G -invariant MSP functions, the weights $\bar{\mathbf{u}}^t$ remain constrained in a linear subspace of dimension $< P$. We can then prove that a function is not strongly $O(d)$ -SGD-learnable if no global minimizers lie on this subspace. For example, this is the case of h_3 : as argued in Section 4 of the main text, $\mathbb{E}_{\mathbf{z}}[\hat{f}_{\text{NN}}(\mathbf{z}; \bar{\rho}_t) z_1 z_2] = \mathbb{E}_{\mathbf{z}}[\hat{f}_{\text{NN}}(\mathbf{z}; \bar{\rho}_t) z_1 z_3]$, and $\hat{f}_{\text{NN}}(\cdot; \bar{\rho}_t)$ can never coincide with h_3 . For h_4 , however, there exists $\bar{\rho}$ with $\bar{\mathbf{u}} = u(1, 1, 1)$ that achieves 0 risk (indeed, define $X = z_1 + z_2 + z_3$, then $h_*(X) = (X^2 - 1)/2 + (-1)^{(X+3)/2}$ and can be fitted with a cubic polynomial). It is difficult to check whether for such functions, the (DF-PDE) dynamics will converge to 0-risk for some

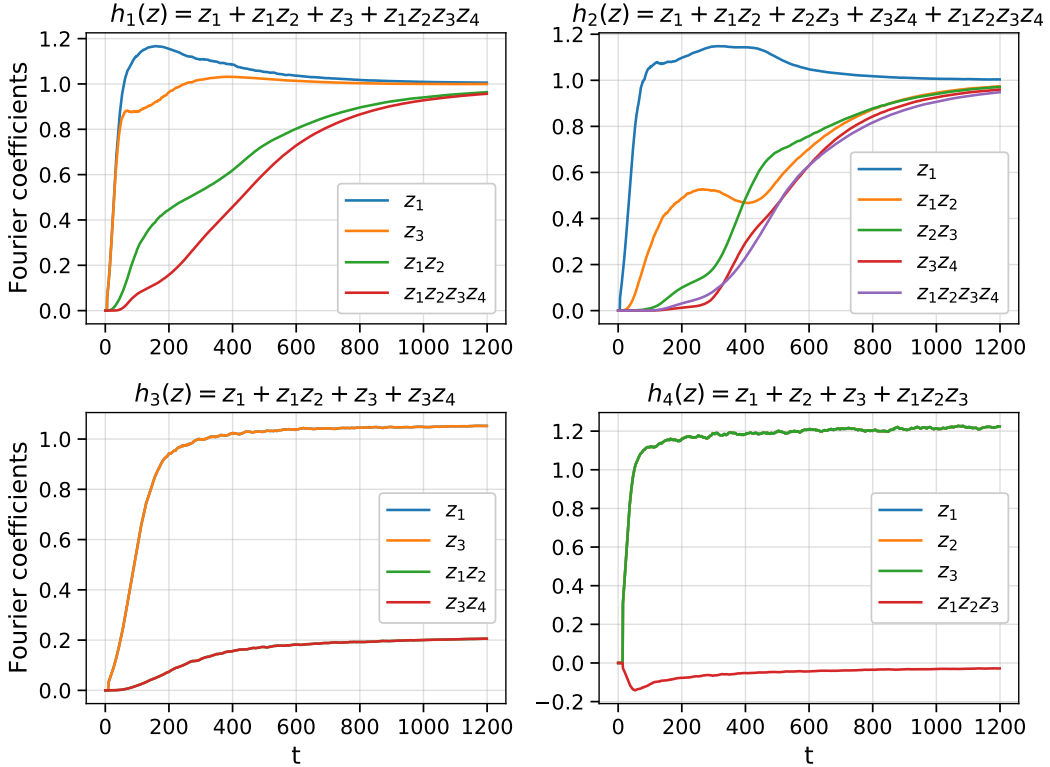


Figure 2: Evolution of the Fourier coefficients during the (DF-PDE) dynamics for 4 MSP functions.

initialization and activation function. (Let us just mention that some G -invariant MSP functions are indeed strongly $O(d)$ -SGD-learnable, such as $h_*(z) = z_1 + z_2$.)

Such G -invariant functions appear naturally in applications and we believe that understanding their dynamics is an important future direction. However, in this paper we consider instead to perturb the Fourier coefficients, which breaks the symmetries, and we show that any MSP function is strongly $O(d)$ -SGD-learnable almost surely over this perturbation. In Figure 3, we plot the evolution of the Fourier coefficients of the original $h_4(z) = z_1 + z_2 + z_3 + z_1z_2z_3$ and its perturbation $\tilde{h}_4(z) = z_1 + 0.99z_2 + 1.01z_3 + z_1z_2z_3$. We see that \tilde{h}_4 is no longer G -invariant and \tilde{h}_4 is strongly $O(d)$ -SGD-learnable.

Remark 13 *In this paper, we only prove that the set of MSP functions that are not strongly $O(d)$ -SGD-learnable is of Lebesgue measure 0. We do not characterize this set beyond this and do not prove that G -invariant MSP functions coincides with this set (in particular, we do not show that G -invariant MSP functions are the only functions that might not be strongly $O(d)$ -SGD-learnable).*

We conclude this section with a final comment about the necessity condition of MSP, which holds only when considering *arbitrarily large* d .

Escaping the saddle-space. The proof that non-MSP functions are not strongly $O(d)$ -SGD-learnable relies on the fact that, when d goes to infinity, the initialization $u_i^0 \rightarrow 0$ for $i \in [P]$. However, for d




Figure 3: Evolution of Fourier coefficients during the (DF-PDE) dynamics for the degenerate MSP $h_4(\mathbf{z}) = z_1 + z_2 + z_3 + z_1 z_2 z_3$ (left) and perturbed MSP function $\tilde{h}_4(\mathbf{z}) = z_1 + 0.99z_2 + 1.01z_3 + z_1 z_2 z_3$ (right).




Figure 4: Fourier coefficients of one-pass (bSGD) and (DF-PDE) solutions throughout the dynamics for $h_*(\mathbf{z}) = z_1 + z_1 z_2 z_3 + z_1 z_2 z_3 z_4$ (left) and $h_*(\mathbf{z}) = z_1 + z_1 z_2 z_3 z_4$ (right).

fixed, $u_i^0 \approx 1/\sqrt{d}$ and waiting sufficiently long, one-pass (bSGD) escapes the neighborhood of the subspace $u_i^0 = 0$. In this case, the time to escape the subspace has to grow with d , and we are not in the $O(d)$ -scaling anymore (indeed $n = Tb/\eta \approx Td$ for one pass (bSGD)). In Figure 4, we consider the same experimental setting as Figure 1 but with h_* missing one (left) or two (right) stairs. We see that (DF-PDE) remains trapped in the saddle-space, while one-pass (bSGD) escapes around $n \approx d^2$ and $n \approx d^{2.7}$ respectively. This agrees with the intuition that staircases with larger leaps are harder to learn with SGD.

Appendix B. Proofs for continuous mean-field and dimension-free dynamics

In this appendix, we provide proofs and discussions for the results presented in Section 3, which corresponds to the ‘continuous-time regime’ of strong $O(d)$ -SGD-learnability. A discrete version of these results and proofs are presented in Appendix C and require little modifications.

Throughout this section, we will denote by K a constant that depends only on the constants in Assumptions A0-A2, A3’ (in particular, K is independent of d, P, T). The value of this constant is allowed to change from line to line.

B.1. Justification for the dimension-free dynamics equations

Here, we provide more details and intuition on how to derive the equations of the dimension-free dynamics (DF-PDE). We report to Section B.2.1 a rigorous proof of Theorem 5, which shows a non-asymptotic bound between (bSGD) and (DF-PDE) dynamics.

First, by Assumption A2, the coordinates of \mathbf{w}^0 are iid and symmetric and therefore

$$\begin{aligned}\hat{f}_{\text{NN}}(\mathbf{x}; \rho_0) &= \int a^0 \sigma(\langle \mathbf{u}^0, \mathbf{z} \rangle + \langle \mathbf{v}^0, \mathbf{r} \rangle) \rho_0(d\boldsymbol{\theta}^0) \\ &= \int a^0 \mathbb{E}_{\mathbf{r}}[\sigma(\langle \mathbf{u}^0, \mathbf{z} \rangle + \langle \mathbf{v}^0, \mathbf{r} \rangle)] \rho_t(d\boldsymbol{\theta}^0) =: \hat{f}_{\text{NN}}(\mathbf{z}; \rho_0).\end{aligned}$$

By symmetry of (MF-PDE), the following lemma shows that the neural network stays independent of the uninformative part \mathbf{r} of the input during the whole trajectory.

Lemma 14 *The solution $(\rho_t)_{t \geq 0}$ of (MF-PDE) with initialization ρ_0 satisfying A2, obeys:*

$$\forall t \geq 0, \quad \hat{f}_{\text{NN}}(\mathbf{x}; \rho_t) = \int a^t \mathbb{E}_{\mathbf{r}}[\sigma(\langle \mathbf{u}^t, \mathbf{z} \rangle + \langle \mathbf{v}^t, \mathbf{r} \rangle)] \rho_t(d\boldsymbol{\theta}^t) =: \hat{f}_{\text{NN}}(\mathbf{z}; \rho_t). \quad (6)$$

Proof [Proof of Lemma 14] It is sufficient to show that for any $\mathbf{r} \in \{-1, +1\}^{d-P}$, the weights $(a^t, \mathbf{u}^t, \mathbf{v}^t \odot \mathbf{r})$ have the same distribution as $(a^t, \mathbf{u}^t, \mathbf{v}^t) \sim \rho_t$ where $\mathbf{v}^t \odot \mathbf{r} = (v_1^t r_1, \dots, v_{d-P}^t r_{d-P})$. Consider $\rho_t^{\#r} = \rho_t \circ \varphi_{\mathbf{r}}$ where $\varphi_{\mathbf{r}}(\boldsymbol{\theta}) = (a, \mathbf{u}, \mathbf{v} \odot \mathbf{r})$. First, notice that by assumption on ρ_0 , we have $\rho_0^{\#r} = \rho_0$. Second, it is easy to check that for any bounded function $h : \mathbb{R}^{d+1} \rightarrow \mathbb{R}$, we have

$$\begin{aligned}\frac{d}{dt} \int h(\boldsymbol{\theta}) \rho_t^{\#r}(d\boldsymbol{\theta}) &= \frac{d}{dt} \int h(\varphi_{\mathbf{r}}(\boldsymbol{\theta})) \rho_t(d\boldsymbol{\theta}) \\ &= - \int \langle \nabla_{\boldsymbol{\theta}} h(\varphi_{\mathbf{r}}(\boldsymbol{\theta})), \mathbf{H}(t) \nabla_{\boldsymbol{\theta}} \psi(\boldsymbol{\theta}; \rho_t) \rangle \rho_t(d\boldsymbol{\theta}) \\ &= - \int \langle (\mathbf{1}, \mathbf{r}) \odot \nabla_{\boldsymbol{\theta}} h(\boldsymbol{\theta}), \mathbf{H}(t) (\mathbf{1}, \mathbf{r}) \odot \nabla_{\boldsymbol{\theta}} \psi(\boldsymbol{\theta}; \rho_t^{\#r}) \rangle \rho_t^{\#r}(d\boldsymbol{\theta}) \\ &= - \int \langle \nabla_{\boldsymbol{\theta}} h(\boldsymbol{\theta}), \mathbf{H}(t) \nabla_{\boldsymbol{\theta}} \psi(\boldsymbol{\theta}; \rho_t^{\#r}) \rangle \rho_t^{\#r}(d\boldsymbol{\theta}),\end{aligned}$$

where we used in the third line that $\psi(\varphi_{\mathbf{r}}(\boldsymbol{\theta}); \rho_t) = \psi(\boldsymbol{\theta}; \rho_t^{\#r})$. Hence $\rho_t^{\#r}$ is the solution of the (MF-PDE) dynamics with initialization $\rho_0^{\#r} = \rho_0$. Hence by uniqueness of the solution, we deduce that $\rho_t^{\#r} = \rho_t$ for any $t \geq 0$. \blacksquare

As mentioned in the main text, one can show that conditional on \mathbf{v}^t , the noise part of the signal $\langle \mathbf{v}^t, \mathbf{r} \rangle$ for $\mathbf{r} \sim \text{Unif}(\{-1, +1\}^{d-P})$ is well approximated by $\|\mathbf{v}^t\|_2 G$ with $G \sim \mathcal{N}(0, 1)$, as long as $\max_i |v_i^t| / \|\mathbf{v}^t\|_2$ remains small. This is the case at $t = 0$ by Assumption A2 that the v_i^0 are iid and sub-Gaussian, and we show in Section B.4 that it remains true up to times $T = O_d(1)$. This motivates the introduction of effective parameters $\bar{\boldsymbol{\theta}}^t = (\bar{a}^t, \bar{\mathbf{u}}^t, \bar{s}^t)$ with effective distribution $\bar{\rho}_t \in \mathcal{P}(\mathbb{R}^{P+2})$. The new parameter \bar{s}^t plays the role of $\|\mathbf{v}^t\|_2$ and we replace the neural network (6) by an effective neural network

$$\hat{f}_{\text{NN}}(\mathbf{z}; \bar{\rho}_t) = \int \bar{a}^t \mathbb{E}_G [\sigma(\langle \bar{\mathbf{u}}^t, \mathbf{z} \rangle + \bar{s}^t G)] \bar{\rho}_t(d\bar{\boldsymbol{\theta}}^t). \quad (7)$$

The evolution equations of $(a^t, \mathbf{u}^t, \|\mathbf{v}^t\|_2)$ associated to the (MF-PDE) dynamics are given by

$$\begin{aligned} \frac{d}{dt} a^t &= \xi^a(t) \mathbb{E}_{\mathbf{z}, \mathbf{r}} \left[\{f_*(\mathbf{z}) - \hat{f}_{\text{NN}}(\mathbf{z}; \rho_t)\} \sigma(\langle \mathbf{u}^t, \mathbf{z} \rangle + \langle \mathbf{v}^t, \mathbf{r} \rangle) \right] - \xi^a(t) \lambda^a a^t, \\ \frac{d}{dt} \mathbf{u}^t &= \xi^w(t) a^t \mathbb{E}_{\mathbf{z}, \mathbf{r}} \left[\{f_*(\mathbf{z}) - \hat{f}_{\text{NN}}(\mathbf{z}; \rho_t)\} \sigma'(\langle \mathbf{u}^t, \mathbf{z} \rangle + \langle \mathbf{v}^t, \mathbf{r} \rangle) \mathbf{z} \right] - \xi^w(t) \lambda^w \mathbf{u}^t, \\ \frac{d}{dt} \|\mathbf{v}^t\|_2 &= \xi^w(t) a^t \mathbb{E}_{\mathbf{z}, \mathbf{r}} \left[\{f_*(\mathbf{z}) - \hat{f}_{\text{NN}}(\mathbf{z}; \rho_t)\} \sigma'(\langle \mathbf{u}^t, \mathbf{z} \rangle + \langle \mathbf{v}^t, \mathbf{r} \rangle) \langle \mathbf{v}^t / \|\mathbf{v}^t\|_2, \mathbf{r} \rangle \right] \\ &\quad - \xi^w(t) \lambda^w \|\mathbf{v}^t\|_2, \end{aligned} \quad (8)$$

where we used that $\frac{d}{dt} \|\mathbf{v}^t\|_2 = \frac{1}{\|\mathbf{v}^t\|_2} \langle \mathbf{v}^t, \frac{d}{dt} \mathbf{v}^t \rangle$ to write the last equation.

For P fixed and $d \rightarrow \infty$, we see that the distribution of $(a^0, \mathbf{u}^0, \|\mathbf{v}^0\|_2)$ converges in distribution to $(a^0, \mathbf{0}, m_2^w)$ which we denote $\bar{\rho}_0$ and we recall that $m_2^w = \mathbb{E}_{\mu_w}[W^2]^{1/2}$. As argued above, the mean-field neural network $\hat{f}_{\text{NN}}(\mathbf{z}; \rho_t)$ converges to $\hat{f}_{\text{NN}}(\mathbf{z}; \bar{\rho}_t)$ for any $t \geq 0$. The evolution equations of $\bar{\rho}_t$ can be obtained by taking $d \rightarrow \infty$ in Eq. (8), and replacing $\langle \mathbf{v}^t, \mathbf{r} \rangle$ by $\bar{s}^t G$:

$$\begin{aligned} \frac{d}{dt} \bar{a}^t &= \xi^a(t) \mathbb{E}_{\mathbf{z}, G} \left[\{f_*(\mathbf{z}) - \hat{f}_{\text{NN}}(\mathbf{z}; \bar{\rho}_t)\} \sigma(\langle \bar{\mathbf{u}}^t, \mathbf{z} \rangle + \bar{s}^t G) \right] - \xi^a(t) \lambda^a \bar{a}^t, \\ \frac{d}{dt} \bar{\mathbf{u}}^t &= \xi^w(t) \bar{a}^t \mathbb{E}_{\mathbf{z}, G} \left[\{f_*(\mathbf{z}) - \hat{f}_{\text{NN}}(\mathbf{z}; \bar{\rho}_t)\} \sigma'(\langle \bar{\mathbf{u}}^t, \mathbf{z} \rangle + \bar{s}^t G) \mathbf{z} \right] - \xi^w(t) \lambda^w \bar{\mathbf{u}}^t, \\ \frac{d}{dt} \bar{s}^t &= \xi^w(t) \bar{a}^t \mathbb{E}_{\mathbf{z}, G} \left[\{f_*(\mathbf{z}) - \hat{f}_{\text{NN}}(\mathbf{z}; \bar{\rho}_t)\} \sigma'(\langle \bar{\mathbf{u}}^t, \mathbf{z} \rangle + \bar{s}^t G) G \right] - \xi^w(t) \lambda^w \bar{s}^t. \end{aligned} \quad (9)$$

Denoting $\bar{\mathbf{H}}(t) = \text{diag}(\xi^a(t), \xi^w(t) \mathbf{I}_{P+1})$ and regularization parameters $\bar{\boldsymbol{\Lambda}} = \text{diag}(\lambda^a, \lambda^w \mathbf{I}_{P+1})$, the equations (9) are the parameter evolution equations associated to the following PDE in the space of probability distributions on \mathbb{R}^{P+2} :

$$\begin{aligned} \partial_t \bar{\rho}_t &= \nabla_{\bar{\boldsymbol{\theta}}} \cdot (\bar{\rho}_t \bar{\mathbf{H}}(t) \nabla_{\bar{\boldsymbol{\theta}}} \psi(\bar{\boldsymbol{\theta}}; \bar{\rho}_t)), \\ \psi(\bar{\boldsymbol{\theta}}; \bar{\rho}_t) &= \frac{1}{2} \mathbb{E}_{\mathbf{z}, G} \left[\{\hat{f}_{\text{NN}}(\mathbf{z}; \bar{\rho}_t) - f_*(\mathbf{z})\} \bar{a} \sigma(\langle \bar{\mathbf{u}}, \mathbf{z} \rangle + \bar{s} G) \right] + \frac{1}{2} \bar{\boldsymbol{\theta}}^\top \bar{\boldsymbol{\Lambda}} \bar{\boldsymbol{\theta}}, \end{aligned} \quad (10)$$

Equivalently, this PDE corresponds to the gradient flow in the Wasserstein space (with $\bar{\mathbf{H}}(t)$ rescaling) over the regularized risk functional:

$$E(\bar{\rho}) = \frac{1}{2} \mathbb{E}_{\mathbf{z}} [\{h_*(\mathbf{z}) - \hat{f}_{\text{NN}}(\mathbf{z}; \bar{\rho})\}^2] + \frac{1}{2} \int \bar{\boldsymbol{\theta}}^\top \bar{\boldsymbol{\Lambda}} \bar{\rho}(d\bar{\boldsymbol{\theta}}).$$

B.2. Proof of the results in Section 3

In this section, we gather the proofs for the results on the dimension free dynamics and the necessity condition. The longer and more technical arguments are deferred to Sections B.3 and B.4.

B.2.1. PROOF OF THEOREM 5

We use the mean-field dynamics (MF-PDE) as an intermediary dynamics for the bound. Theorem 5 is a direct consequence of the following two bounds:

Proposition 15 *Assume conditions A0-A2, A3', and let $T \geq 1$. There exists constants K_0 and K_1 depending only on the constants in A0-A2, A3' (in particular, independent of d, P, T), such that for any $\eta \leq e^{-K_0 T^3} [\{b/(d + \log(N))\} \wedge 1]$, we have*

$$\sup_{k \in [T/\eta] \cap \mathbb{N}} \|\hat{f}_{\text{NN}}(\cdot; \Theta^k) - \hat{f}_{\text{NN}}(\cdot; \rho_{k\eta})\|_{L^2} \leq K_1 e^{K_1 T^3} \left\{ \sqrt{\frac{\log N}{N}} + \left[\sqrt{\frac{d + \log N}{b}} \vee 1 \right] \sqrt{\eta} \right\},$$

with probability at least $1 - 1/N$.

This proposition follows from a straightforward extension of Mei et al. (2019) to batch-SGD and anisotropic step sizes, and can be found in Section B.3. In particular, Proposition 15 implies that, if we consider $T, K = O_d(1)$, then $N = \Omega_d(1)$ and $1/\eta = \Omega_d(d/b)$ are sufficient for the mean-field PDE to be an accurate approximation of batch-SGD up to time T (recall that $T = \eta n/b$ by one-pass assumption and therefore $n = O_d(d)$).

Theorem 16 *Assume conditions A0-A2, A3', and let $T \geq 1$. There exists a constant K_1 depending only on the constants in A0-A2, A3' (in particular, independent of d, P, T), such that*

$$\sup_{t \in [0, T]} \|\hat{f}_{\text{NN}}(\cdot; \rho_t) - \hat{f}_{\text{NN}}(\cdot; \bar{\rho}_t)\|_{L^2} \leq K_1 e^{K_1 T^7} \sqrt{\frac{P + \log(d)}{d}}.$$

The proof of Theorem 16 can be found in Section B.4.

B.2.2. PROOF OF THEOREM 6

Fix $\varepsilon > 0$. Consider Lipschitz $\xi^a, \xi^w : \mathbb{R}_{>0} \rightarrow \mathbb{R}_{>0}$ and $\lambda^a, \lambda^w \geq 0$ such that $\lim_{t \rightarrow \infty} R(\bar{\rho}_t) < \varepsilon/6$. Take T such that $R(\bar{\rho}_T) = \varepsilon/6$. Let K be sufficiently large such that A₀ – A₂ are satisfied, and A'₃ is satisfied on $[0, T]$. By Theorem 5, there exists constants K_0, K_1 that only depend on K such that the bound holds with probability at least $1 - 1/N$ for $\eta \leq e^{-K_0 T^3} [\{b/(d + \log(N))\} \wedge 1]$. Consider c_1 such that $c_1/\log(c_1) = 81K_1^2 e^{2K_1 T^3}/\varepsilon$ and take $C(h_*, \varepsilon) = 2(T \vee 1)(c_1 \vee e^{K_0 T^3})$ and $T(\varepsilon, h_*) = T$. Then for any $d \geq C(\varepsilon, h_*)$, $n \geq C(\varepsilon, h_*)d$ and $e^d \geq N \geq C(\varepsilon, h_*)$, taking $b = d + \log(N)$, we have $\sqrt{\eta} \leq \varepsilon e^{-K_1 T^6}/(9K_1)$, and with probability $1 - 1/N$, taking $k_0 = \lfloor T/\eta \rfloor$,

$$\|\hat{f}_{\text{NN}}(\cdot; \Theta^{k_0}) - \hat{f}_{\text{NN}}(\cdot; \bar{\rho}_{\eta k_0})\|_{L^2} \leq \frac{\sqrt{\varepsilon}}{3}.$$

Hence

$$\begin{aligned} R(h_*, \hat{f}_{\text{NN}}(\cdot; \Theta^{k_0})) &\leq 3R(h_*, \hat{f}_{\text{NN}}(\cdot; \bar{\rho}_T)) + 3\|\hat{f}_{\text{NN}}(\cdot; \Theta^{k_0}) - \hat{f}_{\text{NN}}(\cdot; \bar{\rho}_{\eta k_0})\|_{L^2}^2 \\ &\quad + 3\|\hat{f}_{\text{NN}}(\cdot; \bar{\rho}_T) - \hat{f}_{\text{NN}}(\cdot; \bar{\rho}_{\eta k_0})\|_{L^2}^2 \\ &\leq \frac{\varepsilon}{2} + \frac{\varepsilon}{3} + \frac{\varepsilon}{9} \leq \varepsilon. \end{aligned}$$

Conversely, assume that h_* is strongly SGD-learnable in $O(d)$ -scaling. Let $(b, \lambda^a, \lambda^w, \{\eta_k^a, \eta_k^w\}_{k \in [0, k_0]})$ be the hyperparameters that satisfy strong learnability for $\varepsilon/4$: in particular, $k_0 = n/b$ and $R(\Theta^{k_0}) \leq \varepsilon/4$ with probability at least $9/10$. Take $T = n\eta/b \leq T(h_*, \varepsilon/4)$ and let η^a, η^w be piecewise linear functions such that $\eta^a(\eta k) = \eta_k^a/\eta$ and $\eta^w(\eta k) = \eta_k^w/\eta$. Consider $(\bar{\rho}_t)_{t \geq 0}$ the solution of (DF-PDE) with $\eta^a, \eta^w, \lambda^a, \lambda^w$. From Theorem 5, there exists constants K_0 and K_1 that only depend on $\eta^a, \eta^w, \lambda^a, \lambda^w$ through the constants in assumption A3, such that

$$\begin{aligned} & \|\hat{f}_{\text{NN}}(\cdot; \Theta^{k_0}) - \hat{f}_{\text{NN}}(\cdot; \bar{\rho}_{\eta k_0})\|_{L^2} \\ & \leq K_1 e^{K_1 c(h_*, \varepsilon/4)^6} \left\{ \sqrt{\frac{P + \log(d)}{d}} + \sqrt{\frac{\log N}{N}} + \left[\sqrt{\frac{d + \log N}{b}} \vee 1 \right] \sqrt{\eta} \right\}, \end{aligned}$$

with probability at least $1 - 1/N$. We can therefore take d, N, n sufficiently big such that the right-hand side is less than $\sqrt{\varepsilon}/2$. On the intersection of this event and the event $R(\Theta^{k_0}) \leq \varepsilon/4$ (which happens with positive probability), we have

$$R(\bar{\rho}_{\eta k_0}) \leq 2R(\Theta^{k_0}) + 2\|\hat{f}_{\text{NN}}(\cdot; \Theta^{k_0}) - \hat{f}_{\text{NN}}(\cdot; \bar{\rho}_{\eta k_0})\|_{L^2}^2 \leq \frac{\varepsilon}{2} + \frac{\varepsilon}{2} = \varepsilon,$$

which finishes the proof.

B.2.3. PROOF OF THEOREM 7

Consider $h_* : \{+1, -1\}^P \rightarrow \mathbb{R}$ that is not MSP, and denote by $\{\hat{h}(S)\}_{S \in \mathcal{S}^*}$ its non-zero Fourier coefficients $h_*(z) = \sum_{S \in \mathcal{S}^*} \hat{h}(S) \chi_S(z)$. Denote by $\bar{\mathcal{S}}_* \subset \mathcal{S}_*$ the biggest subset of \mathcal{S}_* such that $\bar{\mathcal{S}}_* = \{S_1, \dots, S_r\}$ can be ordered with $|S_i \setminus \cup_{j=1}^{i-1} S_j| \leq 1$ for any $i \in [r]$. By assumption $\mathcal{S}_* \setminus \bar{\mathcal{S}}_*$ is not empty, and for any $S \in \mathcal{S}_* \setminus \bar{\mathcal{S}}_*$, there exists at least two coordinates $i_1, i_2 \in S$ such that $i_1, i_2 \notin \cup_{S \in \bar{\mathcal{S}}_*} S$. Denote $\Omega = [P] \setminus (\cup_{S \in \bar{\mathcal{S}}_*} S)$. We show that $\bar{u}_i^t = 0$ during the whole dynamics for every $i \in \Omega$. In particular, this implies that for any $S \in \mathcal{S}_* \setminus \bar{\mathcal{S}}_*$, $\mathbb{E}_z[\chi_S(z) \hat{f}_{\text{NN}}(z; \bar{\rho}_t)] = 0$, and

$$R(h_*; \hat{f}(\cdot; \bar{\rho}_t)) \geq \sum_{S \in \mathcal{S}_* \setminus \bar{\mathcal{S}}_*} \hat{h}(S)^2 > 0.$$

This lower bound does not depend on the details of the dynamics (parameters $\xi^a, \xi^w, \lambda^a, \lambda^w$, activation and initialization μ_a, m_2^w). Let $z_{i,+}$ and $z_{i,-}$ denote the vector $z \in \{+1, -1\}^P$, with $z_i = +1$ and $z_i = -1$ respectively, and note that by Lemma 19, $|\bar{u}_i^t| \leq K(1+t)$. Using Assumption A0, we have by integrating out z_i :

$$\begin{aligned} |\mathbb{E}_{z_i}[\hat{f}_{\text{NN}}(z; \bar{\rho}_t) \sigma'(\langle z, \bar{u}^t \rangle) z_i]| & \leq \frac{1}{2} |\hat{f}_{\text{NN}}(z_{i,+}; \bar{\rho}_t) \{\sigma'(\langle z_{i,+}, \bar{u}^t \rangle) - \sigma'(\langle z_{i,-}, \bar{u}^t \rangle)\}| \\ & \quad + \frac{1}{2} |\hat{f}_{\text{NN}}(z_{i,-}; \bar{\rho}_t) \{\sigma'(\langle z_{i,-}, \bar{u}^t \rangle) - \sigma'(\langle z_{i,+}, \bar{u}^t \rangle)\}| \\ & \leq K(1+t) \|\sigma\|_\infty \|\sigma''\|_\infty |\bar{u}_i^t| + K(1+t) \|\sigma'\|_\infty \|\sigma'\|_\infty |\bar{u}_i^t| \\ & \leq K(1+t) |\bar{u}_i^t|. \end{aligned}$$

Similarly, for $i \in \Omega$ and $S \in \bar{\mathcal{S}}_*$ (in particular, since $i \notin S$)

$$|\mathbb{E}_z[\chi_S(z) \sigma'(\langle z, \bar{u}^t \rangle) z_i]| \leq \|\sigma''\|_\infty |\bar{u}_i^t|,$$

while if $S \in \mathcal{S}_* \setminus \bar{\mathcal{S}}_*$, then there exists $j \in \Omega \cap S$ with $j \neq i$, hence

$$|\mathbb{E}_{\mathbf{z}}[\chi_S(\mathbf{z})\sigma'(\langle \mathbf{z}, \bar{\mathbf{u}}^t \rangle)z_i]| \leq \|\sigma''\|_\infty |\bar{u}_j^t|.$$

Denoting $m_\Omega^t = \max_{i \in \Omega} |\bar{u}_i^t|$ (recall $m_\Omega^0 = 0$), we conclude that for any $i \in \Omega$:

$$\left| \frac{d}{dt} \bar{u}_i^t \right| = |\bar{a}^t| |\mathbb{E}_{\mathbf{z}}[(h_*(\mathbf{z}) - \hat{f}_{\text{NN}}(\mathbf{z}; \bar{\rho}_t))\sigma'(\langle \bar{\mathbf{u}}^t, \mathbf{z} \rangle)z_i]| \leq K(1+t)^2 m_\Omega^t,$$

and therefore $m_\Omega^t = 0$ during the whole dynamics.

B.3. Proof of Proposition 15

The proof is an application of an extension of Theorem 1.(B) in [Mei et al. \(2019\)](#) to batch-SGD and anisotropic step sizes. This extension is straightforward and we simply list below the two main differences with the proof in Appendix C of [Mei et al. \(2019\)](#):

- Recall that we defined the regularized risk $E(\rho) := \frac{1}{2}R(\rho) + \frac{1}{2} \int \boldsymbol{\theta}^\top \boldsymbol{\Lambda} \boldsymbol{\theta} \rho(d\boldsymbol{\theta})$. We have

$$\frac{d}{dt} E(\rho_t) = - \int \|\nabla_{\boldsymbol{\theta}} \psi(\boldsymbol{\theta}^t; \rho_t)\|_{\mathbf{H}(t)\rho_t}^2 d\boldsymbol{\theta}^t \leq 0,$$

where $\psi(\boldsymbol{\theta}^t; \rho_t)$ is defined in Eq. (MF-PDE) and we denoted $\|\mathbf{v}\|_{\mathbf{A}} = \|\mathbf{A}^{1/2}\mathbf{v}\|_2$. We conclude that $E(\rho_t)$ is nonincreasing. The rest of the proof only uses that $\mathbf{H}(t)$ verifies $\|\mathbf{H}\|_\infty, \|\mathbf{H}\|_{\text{Lip}} \leq K$.

- The concentration between the batch-SGD and gradient descent (Appendix C.5 in [Mei et al. \(2019\)](#)) uses that there is an extra $1/b$ factor in the sub-Gaussian constant.

The proof of Proposition 15 simply amounts to checking that our setting (with Assumptions A0-A2, A3') falls under the general framework of Theorem 1.(B) in [Mei et al. \(2019\)](#).

Proof [Proof of Proposition 15] First, from conditions A0 and A1, we have $\|\sigma\|_\infty \leq K$ and $|y_k| \leq K$. Furthermore, note that \mathbf{x} is a sub-Gaussian vector and σ' is bounded ($\|\sigma'\|_\infty \leq K$ by condition A0). Then, for any $\mathbf{w} \in \mathbb{R}^d$, the gradient $\nabla_{\mathbf{w}} \sigma(\langle \mathbf{x}, \mathbf{w} \rangle) = \mathbf{x} \sigma'(\langle \mathbf{x}, \mathbf{w} \rangle)$ is K -sub-Gaussian. Hence, assumption A₂ of [Mei et al. \(2019\)](#) is verified.

Denote $v(\mathbf{w}) = \mathbb{E}_{\mathbf{x}}[f_*(\mathbf{x})\sigma(\langle \mathbf{x}, \mathbf{w} \rangle)]$ and $u(\mathbf{w}_1, \mathbf{w}_2) = \mathbb{E}_{\mathbf{x}}[\sigma(\langle \mathbf{x}, \mathbf{w}_1 \rangle)\sigma(\langle \mathbf{x}, \mathbf{w}_2 \rangle)]$. Consider $\mathbf{n} \in \mathbb{R}^d$ with $\|\mathbf{n}\|_2 = 1$. Then, we have

$$\begin{aligned} \langle \nabla v(\mathbf{w}), \mathbf{n} \rangle &= \mathbb{E}_{\mathbf{x}}[f_*(\mathbf{x})\sigma'(\langle \mathbf{x}, \mathbf{w} \rangle)\langle \mathbf{n}, \mathbf{x} \rangle] \leq K \mathbb{E}[\langle \mathbf{n}, \mathbf{x} \rangle^2]^{1/2} = K, \\ \langle \nabla_{\mathbf{w}_1} u(\mathbf{w}_1, \mathbf{w}_2), \mathbf{n} \rangle &= \mathbb{E}_{\mathbf{x}}[\sigma'(\langle \mathbf{x}, \mathbf{w}_1 \rangle)\langle \mathbf{n}, \mathbf{x} \rangle \sigma(\langle \mathbf{x}, \mathbf{w}_2 \rangle)] \leq K^2 \mathbb{E}[\langle \mathbf{n}, \mathbf{x} \rangle^2]^{1/2} = K^2, \\ |\langle \nabla^2 v(\mathbf{w}), \mathbf{n}^{\otimes 2} \rangle| &= \mathbb{E}_{\mathbf{x}}[|f_*(\mathbf{x})\sigma''(\langle \mathbf{x}, \mathbf{w} \rangle)|\langle \mathbf{n}, \mathbf{x} \rangle^2] \leq K \mathbb{E}[\langle \mathbf{n}, \mathbf{x} \rangle^2] = PK. \end{aligned}$$

Finally, consider $\mathbf{n}_1, \mathbf{n}_2 \in \mathbb{R}^d$ with $\|\mathbf{n}_1\|_2^2 + \|\mathbf{n}_2\|_2^2 = 1$. Then,

$$\begin{aligned} |\langle \nabla_{(\mathbf{w}_1, \mathbf{w}_2)}^2 u(\mathbf{w}_1, \mathbf{w}_2), (\mathbf{n}_1, \mathbf{n}_2)^{\otimes 2} \rangle| &\leq \mathbb{E}_{\mathbf{x}}[|\sigma''(\langle \mathbf{x}, \mathbf{w}_1 \rangle)\sigma(\langle \mathbf{x}, \mathbf{w}_2 \rangle)|\langle \mathbf{n}_1, \mathbf{x} \rangle^2] \\ &\quad + 2\mathbb{E}_{\mathbf{x}}[|\sigma'(\langle \mathbf{x}, \mathbf{w}_1 \rangle)\sigma'(\langle \mathbf{x}, \mathbf{w}_2 \rangle)\langle \mathbf{n}_1, \mathbf{x} \rangle\langle \mathbf{n}_2, \mathbf{x} \rangle|] \\ &\quad + \mathbb{E}_{\mathbf{x}}[|\sigma(\langle \mathbf{x}, \mathbf{w}_1 \rangle)\sigma''(\langle \mathbf{x}, \mathbf{w}_2 \rangle)|\langle \mathbf{n}_2, \mathbf{x} \rangle^2] \\ &\leq 4K^2. \end{aligned}$$

We conclude that $\|\nabla v(\mathbf{w})\|_2, \|\nabla u(\mathbf{w}_1, \mathbf{w}_2)\|_2, \|\nabla^2 v(\mathbf{w})\|_{\text{op}}, \|\nabla^2 u(\mathbf{w}_1, \mathbf{w}_2)\|_{\text{op}} \leq K$, and assumption A3 in [Mei et al. \(2019\)](#) is verified. \blacksquare

B.4. Proof of Theorem 16: bound between (MF-PDE) and (DF-PDE) dynamics

We will assume throughout this section that the assumptions and the setting of Theorem 16 hold. In particular, we will use Assumptions A0-A2, A3' without mention when clear from context. For clarity, we will write the proof in the case $\xi^a(t) = \xi^w(t) = 1$ and $\lambda^a = \lambda^w = 0$. The general case follows easily, using $\|\xi^a\|_\infty, \|\xi^w\|_\infty, \lambda^a, \lambda^w \leq K$ by Assumption A3'.

We bound the distance between the mean-field and the dimension-free solutions by coupling the two dynamics through their initialization. Denote $\boldsymbol{\theta}^t = (a^t, \mathbf{u}^t, \mathbf{v}^t)$ and $\bar{\boldsymbol{\theta}}^t = (\bar{a}^t, \bar{\mathbf{u}}^t, \bar{s}^t)$ the parameters obtained by the evolution equations (8) and (9) from initial parameters $\boldsymbol{\theta}^0 = (a^0, \mathbf{u}^0, \mathbf{v}^0)$ and $\bar{\boldsymbol{\theta}}^0 = (\bar{a}^0, \bar{\mathbf{u}}^0, \bar{s}^0)$ respectively. Recall that we initialize independently $a^0 \sim \mu_a$ and $\sqrt{d} \cdot (\mathbf{u}^0, \mathbf{v}^0) \sim \mu_w^{\otimes d}$, and $\bar{a}^0 \sim \mu_a$, $\bar{\mathbf{u}}^0 = \mathbf{0}$ and $\bar{s}^0 = m_2^w := \mathbb{E}_{W \sim \mu_w}[W^2]^{1/2}$. We couple the two dynamics by taking $a^0 = \bar{a}^0$ (because of $(\mathbf{u}^t, \mathbf{v}^t)$ this coupling is not deterministic), and denote γ_t the obtained joint distribution on $(\boldsymbol{\theta}^t, \bar{\boldsymbol{\theta}}^t)$.

The goal is to bound

$$\begin{aligned} & \left\| \hat{f}_{\text{NN}}(\cdot; \rho_t) - \hat{f}_{\text{NN}}(\cdot; \bar{\rho}_t) \right\|_{L^2}^2 \\ &= \mathbb{E}_{\mathbf{z}} \left[\left(\int \left\{ a^t \mathbb{E}_{\mathbf{r}}[\sigma(\langle \mathbf{u}^t, \mathbf{z} \rangle + \langle \mathbf{v}^t, \mathbf{r} \rangle)] - \bar{a}^t \mathbb{E}_G[\sigma(\langle \bar{\mathbf{u}}^t, \mathbf{z} \rangle + \bar{s}^t G)] \right\} \gamma_t(d\boldsymbol{\theta}^t d\bar{\boldsymbol{\theta}}^t) \right)^2 \right], \end{aligned} \quad (11)$$

where we used Lemma 14 to remove the dependency in \mathbf{r} . It will be useful to introduce the residuals of the dynamics: $\hat{g}(\mathbf{z}; \rho_t) = h_*(\mathbf{z}) - \hat{f}_{\text{NN}}(\mathbf{z}; \rho_t)$ and $\hat{g}(\mathbf{z}; \bar{\rho}_t) = h_*(\mathbf{z}) - \hat{f}_{\text{NN}}(\mathbf{z}; \bar{\rho}_t)$. Recall that we denote by $R(\rho_t) = \mathbb{E}[(f_*(\mathbf{z}) - \hat{f}_{\text{NN}}(\mathbf{z}; \rho_t))^2] = \mathbb{E}_{\mathbf{z}}[\hat{g}(\mathbf{z}; \rho_t)^2]$ and $R(\bar{\rho}_t) = \mathbb{E}_{\mathbf{z}}[\hat{g}(\mathbf{z}; \bar{\rho}_t)^2]$ the prediction risks at time t .

The value of the integrand in Eq. (11) only depends on $\{(a^t, \mathbf{u}^t, \langle \mathbf{v}^t, \mathbf{r} \rangle)\}_{t \geq 0}$ and $\{(\bar{a}^t, \bar{\mathbf{u}}^t, \bar{s}^t G)\}_{t \geq 0}$ with $\mathbf{r} \sim \text{Unif}(\{-1, +1\}^{d-P})$ and $G \sim \text{N}(0, 1)$ independent of $\{(\boldsymbol{\theta}^t, \bar{\boldsymbol{\theta}}^t)\}_{t \geq 0}$. Conditioning on $(\boldsymbol{\theta}^t, \bar{\boldsymbol{\theta}}^t)$, we consider the 1-Wasserstein distance

$$W_1(\langle \mathbf{v}^t, \mathbf{r} \rangle, \bar{s}^t G) \leq \sqrt{2\pi} \|\mathbf{v}^t\|_2 - \bar{s}^t + W_1(\langle \mathbf{v}^t, \mathbf{r} \rangle, \|\mathbf{v}^t\|_2 G), \quad (12)$$

where we recall that W_1 is defined by

$$\begin{aligned} W_1(X, Y) &= \inf_{\gamma \in \Gamma(X, Y)} \mathbb{E}_{(X, Y) \sim \gamma} [|X - Y|] \\ &= \sup_{f: \mathbb{R} \rightarrow \mathbb{R}, \|f\|_{\text{Lip}} \leq 1} |\mathbb{E}[f(X)] - \mathbb{E}[f(Y)]|. \end{aligned}$$

Lemma 20 in Section B.5 shows that

$$W_1(\langle \mathbf{v}^t, \mathbf{r} \rangle, \|\mathbf{v}^t\|_2 G) \leq 3 \frac{\|\mathbf{v}^t\|_3^3}{\|\mathbf{v}^t\|_2^2} \leq 3 \max_{i \in [d-P]} |v_i^t|. \quad (13)$$

The following lemma bounds the right hand-side through the value of \mathbf{v}^t at initialization:

Lemma 17 *Consider the same setting and assumptions as Theorem 16. There exists a constant K independent of d, P and depending only on the Assumptions A0-A2, A3' such that for any $T \geq 0$,*

$$\sup_{t \in [0, T]} \max_{i \in [d-P]} |v_i^t| \leq K e^{KT^2} \max_{i \in [d-P]} |v_i^0|.$$

Proof [Proof of Lemma 17] We have

$$\begin{aligned} \left| \frac{d}{dt} v_i^t \right| &= \left| a^t \mathbb{E}_{\mathbf{z}, \mathbf{r}} [\hat{g}(\mathbf{z}; \rho_t) \sigma'(\langle \mathbf{u}^t, \mathbf{z} \rangle + \langle \mathbf{v}^t, \mathbf{r} \rangle) r_i] \right| \\ &= \left| a^t v_i^t \mathbb{E}_{\mathbf{z}, \mathbf{r}_{-i}} [\hat{g}(\mathbf{z}; \rho_t) \sigma''(\langle \mathbf{u}^t, \mathbf{z} \rangle + \langle \mathbf{v}_{-i}^t, \mathbf{r}_{-i} \rangle + \xi_{v_i^t})] \right| \\ &\leq |a^t| \cdot |v_i^t| \cdot \mathbb{E}_{\mathbf{z}} [\hat{g}(\mathbf{z}; \rho_t)^2]^{1/2} \|\sigma''\|_\infty \leq K(1+T)|v_i^t|, \end{aligned}$$

where we expanded the expectation on r_i in the second line and used the mean value theorem, and used Eq. (40) in Lemma 19 in the last line. We deduce that

$$|v_i^t| \leq e^{K(1+T)t} |v_i^0|,$$

which concludes the proof. \blacksquare

Using Lemma 17 in the bound (13) yields (conditional on θ^0):

$$\sup_{t \in [0, T]} W_1(\langle \mathbf{v}^t, \mathbf{r} \rangle, \|\mathbf{v}^t\|_2 G) \leq K e^{KT^2} \max_{i \in [d-P]} |v_i^0|. \quad (14)$$

By Lemma 21 in Section B.5, the following holds for any fixed $q \in \mathbb{N}$, $q \leq K$,

$$\begin{aligned} \int \left\{ \sup_{t \in [0, T]} W_1(\langle \mathbf{v}^t, \mathbf{r} \rangle, \|\mathbf{v}^t\|_2 G) \right\}^q \rho_t(d\theta) &\leq K e^{KT^2} \mathbb{E}_{\sqrt{d} \cdot \mathbf{v}^0 \sim \mu_w^{\otimes (d-P)}} \left[\max_{i \in [d-P]} |v_i^0|^{2q} \right] \\ &\leq K e^{KT^2} \left(\frac{\log d}{d} \right)^{q/2}. \end{aligned} \quad (15)$$

Using Eq. (15) and the coupling described above, we will bound (11). Introduce the random quantity

$$\delta(t) = |a^t - \bar{a}^t| \vee \|\mathbf{u}^t - \bar{\mathbf{u}}^t\|_2 \vee |\bar{s}^t - \|\mathbf{v}^t\|_2|, \quad (16)$$

and the square root of its second moment

$$\Delta(t) = \left(\int \delta(t)^2 \gamma_t(d\theta^t d\bar{\theta}^t) \right)^{1/2}. \quad (17)$$

We will show the following technical bounds:

Lemma 18 *Consider the same setting and assumptions as Theorem 16. There exists a constant K independent of d, P and depending only on the Assumptions A0-A2, A3' such that for any $T \geq 0$,*

$$\begin{aligned} \|\hat{f}_{\text{NN}}(\cdot; \rho_t) - \hat{f}_{\text{NN}}(\cdot; \bar{\rho}_t)\|_{L^2} &= \mathbb{E}_{\mathbf{z}} \left[\left\{ \hat{g}(\mathbf{z}; \rho_t) - \hat{g}(\mathbf{z}; \bar{\rho}_t) \right\}^2 \right]^{1/2} \\ &\leq K(1+T)\Delta(t) + K e^{KT^2} \sqrt{\frac{\log d}{d}}, \end{aligned} \quad (18)$$

where

$$\frac{d}{dt} \Delta(t) \leq K(1+T)^6 \Delta(t) + K e^{KT^2} \sqrt{\frac{\log d}{d}}. \quad (19)$$

From this lemma, we can now complete the proof of Theorem 16:

Proof [Proof of Theorem 16] From Gronwall's lemma applied to Eq. (52) in Lemma 18, we have

$$\Delta(t) \leq \left[\Delta(0) + Ke^{KT^2} \sqrt{\frac{\log d}{d}} \right] e^{K(1+T)^6 t},$$

where

$$\begin{aligned} \Delta(0)^2 &= \int \left\{ |a^0 - \bar{a}^0| \vee \|\mathbf{u}^0 - \bar{\mathbf{u}}^0\|_2 \vee |\bar{s}^t - \|\mathbf{v}^0\|_2| \right\}^2 \gamma_0(d\boldsymbol{\theta}^0 d\bar{\boldsymbol{\theta}}^0) \\ &\leq \mathbb{E}_{\sqrt{d} \cdot \mathbf{u}^0 \sim \mu_w^{\otimes P}} [\|\mathbf{u}^0\|_2^2] + \mathbb{E}_{\sqrt{d} \cdot \mathbf{v}^0 \sim \mu_w^{\otimes P}} [(\|\mathbf{v}^0\|_2 - m_2^w)^2] \\ &\leq K \frac{P}{d} + \frac{K}{d} + \frac{KP}{d} \leq K \frac{P}{d}. \end{aligned}$$

Injecting this bound in Eq. (18) concludes the proof. \blacksquare

B.4.1. PROOF OF LEMMA 18

Throughout the proof, we will use the following decomposition for any differentiable $\varphi : \mathbb{R} \rightarrow \mathbb{R}$:

$$\begin{aligned} & \left| \mathbb{E}_{\mathbf{r}, G} [\varphi(\langle \mathbf{u}^t, \mathbf{z} \rangle + \langle \mathbf{v}^t, \mathbf{r} \rangle) - \varphi(\langle \bar{\mathbf{u}}^t, \mathbf{z} \rangle + \bar{s}^t G)] \right| \\ & \leq \left| \mathbb{E}_{\mathbf{r}} [\varphi(\langle \mathbf{u}^t, \mathbf{z} \rangle + \langle \mathbf{v}^t, \mathbf{r} \rangle) - \varphi(\langle \bar{\mathbf{u}}^t, \mathbf{z} \rangle + \langle \mathbf{v}^t, \mathbf{r} \rangle)] \right| \\ & \quad + \left| \mathbb{E}_{\mathbf{r}, G} [\varphi(\langle \bar{\mathbf{u}}^t, \mathbf{z} \rangle + \langle \mathbf{v}^t, \mathbf{r} \rangle) - \varphi(\langle \bar{\mathbf{u}}^t, \mathbf{z} \rangle + \|\mathbf{v}^t\|_2 G)] \right| \\ & \quad + \left| \mathbb{E}_{\mathbf{r}, G} [\varphi(\langle \bar{\mathbf{u}}^t, \mathbf{z} \rangle + \|\mathbf{v}^t\|_2 G) - \varphi(\langle \bar{\mathbf{u}}^t, \mathbf{z} \rangle + \bar{s}^t G)] \right| \\ & \leq \|\varphi'\|_\infty |\langle \mathbf{u}^t - \bar{\mathbf{u}}^t, \mathbf{z} \rangle| + \|\varphi'\|_\infty W_1(\langle \mathbf{v}^t, \mathbf{r} \rangle, \|\mathbf{v}^t\|_2 G) + \|\varphi'\|_\infty |\|\mathbf{v}^t\|_2 - \bar{s}^t| \mathbb{E}[|G|]. \end{aligned} \tag{20}$$

The proof consists in carefully bounding the evolution of the distance between the parameters in the two dynamics.

Step 1. Bound on $\|\hat{f}_{\text{NN}}(\cdot; \rho_t) - \hat{f}_{\text{NN}}(\cdot; \bar{\rho}_t)\|_{L^2}$.

We can bound the difference between the two functions with

$$\begin{aligned} & |\hat{f}_{\text{NN}}(\mathbf{z}; \rho_t) - \hat{f}_{\text{NN}}(\mathbf{z}; \bar{\rho}_t)| \\ & = \left| \int a^t \mathbb{E}_{\mathbf{r}} [\sigma(\langle \mathbf{u}^t, \mathbf{z} \rangle + \langle \mathbf{v}^t, \mathbf{r} \rangle)] \rho_t(d\boldsymbol{\theta}^t) - \int \bar{a}^t \mathbb{E}_G [\sigma(\langle \bar{\mathbf{u}}^t, \mathbf{z} \rangle + \bar{s}^t G)] \bar{\rho}_t(d\bar{\boldsymbol{\theta}}^t) \right| \\ & \leq \text{(I)} + \text{(II)}, \end{aligned}$$

where

$$\begin{aligned} \text{(I)} &= \left| \int (a^t - \bar{a}^t) \mathbb{E}_{\mathbf{r}} [\sigma(\langle \mathbf{u}^t, \mathbf{z} \rangle + \langle \mathbf{v}^t, \mathbf{r} \rangle)] d\gamma_t \right|, \\ \text{(II)} &= \left| \int \bar{a}^t \mathbb{E}_{\mathbf{r}, G} [\{\sigma(\langle \mathbf{u}^t, \mathbf{z} \rangle + \langle \mathbf{v}^t, \mathbf{r} \rangle) - \sigma(\langle \bar{\mathbf{u}}^t, \mathbf{z} \rangle + \bar{s}^t G)\}] d\gamma_t \right|. \end{aligned}$$

The first term can simply be bounded by

$$\text{(I)} \leq \|\sigma\|_\infty \int |a^t - \bar{a}^t| d\gamma_t \leq K \Delta(t), \tag{21}$$

while we use Eq. (20) for the second term

$$\begin{aligned}
 (\text{II}) &\leq \|\bar{a}^t\|_\infty \|\sigma'\|_\infty \int \left\{ |\langle \mathbf{u}^t - \bar{\mathbf{u}}^t, \mathbf{z} \rangle| + \|\mathbf{v}^t\|_2 - \bar{s}^t + W_1(\langle \mathbf{v}^t, \mathbf{r} \rangle, \|\mathbf{v}^t\|_2 G) \right\} d\gamma_t \\
 &\leq K(1+T) \int |\langle \mathbf{u}^t - \bar{\mathbf{u}}^t, \mathbf{z} \rangle| d\gamma_t + K(1+T)\Delta(t) + Ke^{KT^2} \sqrt{\frac{\log d}{d}},
 \end{aligned} \tag{22}$$

where we used Eq. (40) in Lemma 19 and Eq. (15) with $q = 1$.

Combining bounds (21) and (22) and by Jensen's inequality,

$$\begin{aligned}
 &\|\hat{f}_{\text{NN}}(\cdot; \rho_t) - \hat{f}_{\text{NN}}(\cdot; \bar{\rho}_t)\|_{L^2} \\
 &\leq K(1+T) \left(\int \mathbb{E}_{\mathbf{z}} [|\langle \mathbf{u}^t - \bar{\mathbf{u}}^t, \mathbf{z} \rangle|^2] d\gamma_t \right)^{1/2} + K(1+T)\Delta(t) + Ke^{KT^2} \sqrt{\frac{\log d}{d}} \\
 &\leq K(1+T)\Delta(t) + Ke^{KT^2} \sqrt{\frac{\log d}{d}},
 \end{aligned}$$

which proves Eq. (18).

Step 2. Bound on $(a^t - \bar{a}^t)^2$.

Let us bound the derivative

$$\begin{aligned}
 \left| \frac{d}{dt}(a^t - \bar{a}^t) \right| &= \left| (a^t - \bar{a}^t) \mathbb{E}_{\mathbf{z}, \mathbf{r}, G} \left[\hat{g}(\mathbf{z}; \rho_t) \sigma(\langle \mathbf{u}^t, \mathbf{z} \rangle + \langle \mathbf{v}^t, \mathbf{r} \rangle) - \hat{g}(\mathbf{z}; \bar{\rho}_t) \sigma(\langle \bar{\mathbf{u}}^t, \mathbf{z} \rangle + \bar{s}^t G) \right] \right| \\
 &\leq (\text{I}) + (\text{II}),
 \end{aligned}$$

where

$$\begin{aligned}
 (\text{I}) &= \left| \mathbb{E}_{\mathbf{z}, \mathbf{r}, G} \left[\left\{ \hat{g}(\mathbf{z}; \rho_t) - \hat{g}(\mathbf{z}; \bar{\rho}_t) \right\} \sigma(\langle \mathbf{u}^t, \mathbf{z} \rangle + \langle \mathbf{v}^t, \mathbf{r} \rangle) \right] \right|, \\
 (\text{II}) &= \left| \mathbb{E}_{\mathbf{z}, \mathbf{r}, G} \left[\hat{g}(\mathbf{z}; \bar{\rho}_t) \left\{ \sigma(\langle \mathbf{u}^t, \mathbf{z} \rangle + \langle \mathbf{v}^t, \mathbf{r} \rangle) - \sigma(\langle \bar{\mathbf{u}}^t, \mathbf{z} \rangle + \bar{s}^t G) \right\} \right] \right|.
 \end{aligned}$$

Noting that $\hat{g}(\mathbf{z}; \rho_t) - \hat{g}(\mathbf{z}; \bar{\rho}_t) = \hat{f}_{\text{NN}}(\mathbf{z}; \bar{\rho}_t) - \hat{f}_{\text{NN}}(\mathbf{z}; \rho_t)$, the first term can be bounded as in step 1 by

$$(\text{I}) \leq \|\sigma\|_\infty \mathbb{E}_{\mathbf{z}} \left[\left\{ \hat{g}(\mathbf{z}; \rho_t) - \hat{g}(\mathbf{z}; \bar{\rho}_t) \right\}^2 \right]^{1/2} \leq K(1+T)\Delta(t) + Ke^{KT^2} \sqrt{\frac{\log d}{d}}. \tag{23}$$

For the second term, we use Eq. (41) in Lemma 19 and the decomposition (20):

$$\begin{aligned}
 (\text{II}) &\leq \|\hat{g}(\cdot; \bar{\rho}_t)\|_\infty \mathbb{E}_{\mathbf{z}, \mathbf{r}, G} \left[\left| \sigma(\langle \mathbf{u}^t, \mathbf{z} \rangle + \langle \mathbf{v}^t, \mathbf{r} \rangle) - \sigma(\langle \bar{\mathbf{u}}^t, \mathbf{z} \rangle + \bar{s}^t G) \right| \right] \\
 &\leq K(1+T) \left\{ \mathbb{E}_{\mathbf{z}} [|\langle \mathbf{u}^t - \bar{\mathbf{u}}^t, \mathbf{z} \rangle|] + \|\mathbf{v}^t\|_2 - \bar{s}^t + W_1(\langle \mathbf{v}^t, \mathbf{r} \rangle, \|\mathbf{v}^t\|_2 G) \right\} \\
 &\leq K(1+T)\delta(t) + K(1+T)W_1(\langle \mathbf{v}^t, \mathbf{r} \rangle, \|\mathbf{v}^t\|_2 G).
 \end{aligned} \tag{24}$$

Combining Eqs. (23) and (24) and applying Cauchy-Schwarz inequality yield

$$\begin{aligned}
 &\frac{1}{2} \left| \frac{d}{dt} \int (a^t - \bar{a}^t)^2 d\gamma_t \right| \\
 &\leq \Delta(t)^2 + \int \left| \frac{d}{dt}(a^t - \bar{a}^t) \right|^2 d\gamma_t \\
 &\leq K(1+T)^2 \Delta(t)^2 + Ke^{KT^2} \frac{\log d}{d} + K(1+T)^2 \int W_1(\langle \mathbf{v}^t, \mathbf{r} \rangle, \|\mathbf{v}^t\|_2 G)^2 \gamma_t(d\bar{\theta}^t) \\
 &\leq K(1+T)^2 \Delta(t)^2 + Ke^{KT^2} \frac{\log d}{d},
 \end{aligned} \tag{25}$$

where we used the bound (13) on W_1 and Eq. (15) in Lemma 21 in the last line. We deduce that for $t \in [0, T]$,

$$\int (a^t - \bar{a}^t)^2 d\gamma_t \leq K e^{KT^2} \frac{\log d}{d} + K(1+T)^2 \int_0^T \Delta(t)^2 dt, \quad (26)$$

where we used that $a^0 = \bar{a}^0$ at initialization.

Step 3. Bound on $\|\mathbf{u}^t - \bar{\mathbf{u}}^t\|_2^2$.

Again we first bound the derivative:

$$\left| \frac{1}{2} \frac{d}{dt} \|\mathbf{u}^t - \bar{\mathbf{u}}^t\|_2^2 \right| = \left| \langle \mathbf{u}^t - \bar{\mathbf{u}}^t, \frac{d}{dt} (\mathbf{u}^t - \bar{\mathbf{u}}^t) \rangle \right| \leq \text{(I)} + \text{(II)} + \text{(III)},$$

where

$$\begin{aligned} \text{(I)} &= \left| (a^t - \bar{a}^t) \mathbb{E}_{\mathbf{z}, r} [\hat{g}(\mathbf{z}; \rho_t) \sigma'(\langle \mathbf{u}^t, \mathbf{z} \rangle + \langle \mathbf{v}^t, \mathbf{r} \rangle) \langle \mathbf{z}, \mathbf{u}^t - \bar{\mathbf{u}}^t \rangle] \right|, \\ \text{(II)} &= \left| \bar{a}^t \mathbb{E}_{\mathbf{z}, r} [\{\hat{g}(\mathbf{z}; \rho_t) - \hat{g}(\mathbf{z}; \bar{\rho}_t)\} \sigma'(\langle \mathbf{u}^t, \mathbf{z} \rangle + \langle \mathbf{v}^t, \mathbf{r} \rangle) \langle \mathbf{z}, \mathbf{u}^t - \bar{\mathbf{u}}^t \rangle] \right|, \\ \text{(III)} &= \left| \bar{a}^t \mathbb{E}_{\mathbf{z}, r, G} [\hat{g}(\mathbf{z}; \bar{\rho}_t) \{\sigma'(\langle \mathbf{u}^t, \mathbf{z} \rangle + \langle \mathbf{v}^t, \mathbf{r} \rangle) - \sigma'(\langle \bar{\mathbf{u}}^t, \mathbf{z} \rangle + \bar{s}^t G)\} \langle \mathbf{z}, \mathbf{u}^t - \bar{\mathbf{u}}^t \rangle] \right|. \end{aligned}$$

These terms are bounded respectively by

$$\begin{aligned} \text{(I)} &\leq |a^t - \bar{a}^t| \|\hat{g}(\cdot; \rho_t)\|_\infty \|\sigma'\|_\infty \mathbb{E}[\langle \mathbf{z}, \mathbf{u}^t - \bar{\mathbf{u}}^t \rangle^2]^{1/2} \\ &\leq K(1+T) \delta(t)^2, \\ \text{(II)} &\leq \|\bar{a}^t\|_\infty \|\sigma'\|_\infty \mathbb{E}_{\mathbf{z}} [\{\hat{g}(\mathbf{z}; \rho_t) - \hat{g}(\mathbf{z}; \bar{\rho}_t)\}^2]^{1/2} \mathbb{E}_{\mathbf{z}} [\langle \mathbf{z}, \mathbf{u}^t - \bar{\mathbf{u}}^t \rangle^2]^{1/2} \\ &\leq K(1+T)^2 \delta(t) \Delta(t) + K e^{KT^2} \delta(t) \sqrt{\frac{\log d}{d}} \\ &\leq K(1+T)^2 [\Delta(t)^2 + \delta(t)^2] + K e^{KT^2} \frac{\log d}{d}, \end{aligned} \quad (27)$$

and

$$\begin{aligned} \text{(III)} &\leq \|\bar{a}^t\|_\infty \|\hat{g}(\cdot; \bar{\rho}_t)\|_\infty \|\mathbf{u}^t - \bar{\mathbf{u}}^t\|_2 \\ &\quad \times \mathbb{E}_{\mathbf{z}} \left[\left(\mathbb{E}_{r, G} [\sigma'(\langle \mathbf{u}^t, \mathbf{z} \rangle + \langle \mathbf{v}^t, \mathbf{r} \rangle) - \sigma'(\langle \bar{\mathbf{u}}^t, \mathbf{z} \rangle + \bar{s}^t G)] \right)^2 \right]^{1/2} \\ &\leq K(1+T)^2 \delta(t)^2 + K(1+T)^2 W_1(\langle \mathbf{v}^t, \mathbf{r} \rangle, \|\mathbf{v}^t\|_2 G)^2. \end{aligned} \quad (28)$$

Combining inequalities (27) and (28) yields

$$\frac{1}{2} \left| \frac{d}{dt} \int \|\mathbf{u}^t - \bar{\mathbf{u}}^t\|_2^2 d\gamma_t \right| \leq K(1+T)^2 \Delta(t)^2 + K e^{KT^2} \frac{\log d}{d}, \quad (29)$$

where we again used the bound (15). We deduce that for $t \in [0, T]$,

$$\int \|\mathbf{u}^t - \bar{\mathbf{u}}^t\|_2^2 d\gamma_t \leq K e^{KT^2} \frac{\log d}{d} + K \frac{P}{d} + K(1+T)^2 \int_0^T \Delta(t)^2 dt, \quad (30)$$

where we used that $\bar{\mathbf{u}}^0 = \mathbf{0}$ and $\sqrt{d} \cdot \mathbf{u}^0 \sim \mu_w^{\otimes P}$ at initialization, and $\mathbb{E}_{\mu_w^{\otimes P}} [\|\mathbf{u}\|_2^2] \leq KP/d$.

Step 4. Bound on $|\bar{s}^t - \|\mathbf{v}^t\|_2|^2$.

First, notice that we have the following simple upper bounds on the evolution of $\|\mathbf{v}^t\|_2$ and \bar{s}^t :

$$\begin{aligned} \left| \frac{d}{dt} \|\mathbf{v}^t\|_2 \right| &\leq \|a^t\|_\infty \|\sigma'\|_\infty \mathbb{E}_{\mathbf{z}} [\hat{g}(\mathbf{z}; \rho_t)^2]^{1/2} \mathbb{E}_{\mathbf{r}} [\langle \mathbf{v}^t, \mathbf{r} \rangle^2 / \|\mathbf{v}^t\|_2^2]^{1/2} \leq K(1+T), \\ \left| \frac{d}{dt} \bar{s}^t \right| &\leq \|\bar{a}^t\|_\infty \|\sigma'\|_\infty \mathbb{E}_{\mathbf{z}} [\hat{g}(\mathbf{z}; \bar{\rho}_t)^2]^{1/2} \mathbb{E}_G [G^2]^{1/2} \leq K(1+T), \end{aligned}$$

which yields

$$\sup_{t \in [0, T]} \|\mathbf{v}^t\|_2 \leq \|\mathbf{v}^0\|_2 + K(1+T)^2, \quad \sup_{t \in [0, T]} \bar{s}^t \leq K(1+T)^2. \quad (31)$$

Furthermore, we have by Gaussian integration by part

$$\frac{d}{dt} \bar{s}^t = \bar{s}^t \bar{a}^t \mathbb{E}_{\mathbf{z}, G} [\hat{g}(\mathbf{z}; \bar{\rho}_t) \sigma''(\langle \bar{\mathbf{u}}^t, \mathbf{z} \rangle + \bar{s}^t G)]. \quad (32)$$

Similarly, we have by expanding the expectation over the r_i 's and using the mean-value theorem:

$$\begin{aligned} \frac{d}{dt} \|\mathbf{v}^t\|_2 &= a^t \mathbb{E}_{\mathbf{z}, \mathbf{r}} [\hat{g}(\mathbf{z}; \rho_t) \sigma'(\langle \mathbf{u}^t, \mathbf{z} \rangle + \langle \mathbf{v}^t, \mathbf{r} \rangle) \langle \mathbf{v}^t / \|\mathbf{v}^t\|_2, \mathbf{r} \rangle] \\ &= \frac{a^t}{\|\mathbf{v}^t\|_2} \sum_{i \in [d-P]} \mathbb{E}_{\mathbf{z}, \mathbf{r}} [\hat{g}(\mathbf{z}; \rho_t) \sigma'(\langle \mathbf{u}^t, \mathbf{z} \rangle + \langle \mathbf{v}^t, \mathbf{r} \rangle) v_i^t r_i] \\ &= \frac{a^t}{\|\mathbf{v}^t\|_2} \sum_{i \in [d-P]} (v_i^t)^2 \mathbb{E}_{\mathbf{z}, \mathbf{r}} [\hat{g}(\mathbf{z}; \rho_t) \sigma''(\langle \mathbf{u}^t, \mathbf{z} \rangle + \langle \mathbf{v}_{-i}^t, \mathbf{r}_{-i} \rangle + \xi_i)] \\ &= a^t \|\mathbf{v}^t\|_2 \mathbb{E}_{\mathbf{z}, \mathbf{r}} [\hat{g}(\mathbf{z}; \rho_t) \sigma''(\langle \mathbf{u}^t, \mathbf{z} \rangle + \langle \mathbf{v}^t, \mathbf{r} \rangle)] + M_t, \end{aligned} \quad (33)$$

where

$$\begin{aligned} M_t &= \frac{a^t}{\|\mathbf{v}^t\|_2} \sum_{i \in [d-P]} (v_i^t)^2 \\ &\quad \times \mathbb{E}_{\mathbf{z}, \mathbf{r}} [\hat{g}(\mathbf{z}; \rho_t) \{ \sigma''(\langle \mathbf{u}^t, \mathbf{z} \rangle + \langle \mathbf{v}_{-i}^t, \mathbf{r}_{-i} \rangle + \xi_i) - \sigma''(\langle \mathbf{u}^t, \mathbf{z} \rangle + \langle \mathbf{v}^t, \mathbf{r} \rangle) \}] \end{aligned} \quad (34)$$

We can now bound the evolution in time of $(\|\mathbf{v}^t\|_2 - \bar{s}^t)$. Using the expressions in Eqs. (32) and (33), we decompose

$$\begin{aligned} &\left| \frac{d}{dt} \{ \|\mathbf{v}^t\|_2 - \bar{s}^t \} \right| \\ &= \left| a^t \|\mathbf{v}^t\|_2 \mathbb{E}_{\mathbf{z}, \mathbf{r}} [\hat{g}(\mathbf{z}; \rho_t) \sigma''(\langle \mathbf{u}^t, \mathbf{z} \rangle + \langle \mathbf{v}^t, \mathbf{r} \rangle)] + M_t - \bar{a}^t \bar{s}^t \mathbb{E}_{\mathbf{z}, G} [\hat{g}(\mathbf{z}; \bar{\rho}_t) \sigma''(\langle \bar{\mathbf{u}}^t, \mathbf{z} \rangle + \bar{s}^t G)] \right| \\ &\leq \text{(I)} + \text{(II)} + \text{(III)} + \text{(IV)} + |M_t|, \end{aligned}$$

where

$$\begin{aligned} \text{(I)} &= \left| (\|\mathbf{v}^t\|_2 - \bar{s}^t) a^t \mathbb{E}_{\mathbf{z}, \mathbf{r}} [\hat{g}(\mathbf{z}; \rho_t) \sigma''(\langle \mathbf{u}^t, \mathbf{z} \rangle + \langle \mathbf{v}^t, \mathbf{r} \rangle)] \right|, \\ \text{(II)} &= \left| (a^t - \bar{a}^t) \bar{s}^t \mathbb{E}_{\mathbf{z}, \mathbf{r}} [\hat{g}(\mathbf{z}; \rho_t) \sigma''(\langle \mathbf{u}^t, \mathbf{z} \rangle + \langle \mathbf{v}^t, \mathbf{r} \rangle)] \right|, \\ \text{(III)} &= \left| \bar{s}^t \bar{a}^t \mathbb{E}_{\mathbf{z}, \mathbf{r}} [\{ \hat{g}(\mathbf{z}; \rho_t) - \hat{g}(\mathbf{z}; \bar{\rho}_t) \} \sigma''(\langle \mathbf{u}^t, \mathbf{z} \rangle + \langle \mathbf{v}^t, \mathbf{r} \rangle)] \right|, \\ \text{(IV)} &= \left| \bar{a}^t \bar{s}^t \mathbb{E}_{\mathbf{z}, G} [\hat{g}(\mathbf{z}; \bar{\rho}_t) \{ \sigma''(\langle \mathbf{u}^t, \mathbf{z} \rangle + \langle \mathbf{v}^t, \mathbf{r} \rangle) - \sigma''(\langle \bar{\mathbf{u}}^t, \mathbf{z} \rangle + \bar{s}^t G) \}] \right|. \end{aligned}$$

These four quantities can be bounded as previously:

$$\begin{aligned}
 \text{(I)} &\leq \|\mathbf{v}^t\|_2 - \bar{s}^t \|a^t\|_\infty \|\sigma''\|_\infty \|\hat{g}(\cdot; \rho_t)\|_{L^2} \leq K(1+T)\delta(t), \\
 \text{(II)} &\leq |a^t - \bar{a}^t| \cdot |\bar{s}^t| \cdot \|\sigma''\|_\infty \|\hat{g}(\cdot; \rho_t)\|_{L^2} \leq K(1+T)\delta(t), \\
 \text{(III)} &\leq \|\bar{a}^t\|_\infty \cdot |\bar{s}^t| \cdot \|\sigma''\|_\infty \|\hat{g}(\cdot; \rho_t) - \hat{g}(\cdot; \bar{\rho}_t)\|_{L^2} \\
 &\leq K(1+T)^3 \Delta(t) + K e^{KT^2} \sqrt{\frac{\log d}{d}}, \\
 \text{(IV)} &\leq \|\bar{a}^t\|_\infty \cdot |\bar{s}^t| \cdot \|\sigma''\|_\infty \|\hat{g}(\cdot; \bar{\rho}_t)\|_{L^2} \\
 &\quad \times \mathbb{E}_{\mathbf{z}, \mathbf{r}, G} \left[\left\{ \sigma''(\langle \mathbf{u}^t, \mathbf{z} \rangle + \langle \mathbf{v}^t, \mathbf{r} \rangle) - \sigma''(\langle \bar{\mathbf{u}}^t, \mathbf{z} \rangle + \bar{s}^t G) \right\}^2 \right]^{1/2} \\
 &\leq K(1+T)^2 \delta(t) + K(1+T)^2 W_1(\langle \mathbf{v}^t, \mathbf{r} \rangle, \|\mathbf{v}^t\|_2 G).
 \end{aligned} \tag{35}$$

For the last term, we use Eq. (34) and that $|\xi_i| \leq |v_i^t|$:

$$\begin{aligned}
 &|M_t| \\
 &\leq \frac{\|a^t\|_\infty}{\|\mathbf{v}^t\|_2} \sum_{i \in [d-P]} (v_i^t)^2 \\
 &\quad \times \left| \mathbb{E}_{\mathbf{z}, \mathbf{r}} \left[\hat{g}(\mathbf{z}; \rho_t) \left\{ \sigma''(\langle \mathbf{u}^t, \mathbf{z} \rangle + \langle \mathbf{v}_{-i}^t, \mathbf{r}_{-i} \rangle + \xi_i) - \sigma''(\langle \mathbf{u}^t, \mathbf{z} \rangle + \langle \mathbf{v}^t, \mathbf{r} \rangle) \right\} \right] \right| \\
 &\leq 2 \frac{\|a^t\|_\infty}{\|\mathbf{v}^t\|_2} \|\hat{g}(\cdot; \rho_t)\|_{L^2} \|\sigma'''\|_\infty \sum_{i \in [d-P]} |v_i^t|^3 \\
 &\leq K \|a^t\|_\infty \|\mathbf{v}^t\|_2 \max_{i \in [d-P]} |v_i^t| \leq K(1+T) [\|\mathbf{v}^0\|_2 + K(1+T)] \max_{i \in [d-P]} |v_i^t|.
 \end{aligned} \tag{36}$$

Combining Eqs. (35) and (36) and applying Cauchy-Schwarz inequality yield

$$\begin{aligned}
 \frac{1}{2} \left| \frac{d}{dt} \int (\|\mathbf{v}^t\|_2 - \bar{s}^t)^2 d\gamma_t \right| &\leq \Delta(t)^2 + \int \left| \frac{d}{dt} \{ \|\mathbf{v}^t\|_2 - \bar{s}^t \} \right|^2 d\gamma_t \\
 &\leq K(1+T)^6 \Delta(t)^2 + K e^{KT^2} \frac{\log d}{d},
 \end{aligned} \tag{37}$$

where we used that

$$\int \|\mathbf{v}^0\|_2^2 \max_{i \in [d-P]} |v_i^t|^2 d\rho_t \leq \left(\int \|\mathbf{v}^0\|_2^4 d\rho_t \right)^{1/2} \left(\int \max_{i \in [d-P]} |v_i^t|^4 d\rho_t \right)^{1/2} \leq K e^{KT^2} \frac{\log d}{d}.$$

We deduce that for $t \in [0, T]$,

$$\int \|\|\mathbf{v}^t\|_2 - \bar{s}^t\|^2 d\gamma_t \leq K e^{KT^2} \frac{\log d}{d} + K \frac{P}{d} + K(1+T)^6 \int_0^T \Delta(t)^2 dt, \tag{38}$$

where we used that $\sqrt{d} \cdot \mathbf{v}^0 \sim \mu_w^{\otimes (d-P)}$ and $\bar{s}^0 = \mathbb{E}_{\mu_w} [W^2]^{1/2}$ at initialization, and

$$\begin{aligned}
 \int \|\|\mathbf{v}^0\|_2 - \bar{s}^0\|^2 d\gamma_0 &\leq \mathbb{E}_{\mathbf{v}^0} \left[\|\|\mathbf{v}^0\|_2^2 - (\bar{s}^0)^2 \right] \\
 &\leq \mathbb{E}_{\mathbf{v}^0} \left[\left\{ \|\mathbf{v}^0\|_2^2 - \mathbb{E}_{\mathbf{v}^0} [\|\mathbf{v}^0\|_2^2] \right\}^2 \right]^{1/2} + \left| \bar{s}^0 - \mathbb{E}_{\mathbf{v}^0} [\|\mathbf{v}^0\|_2^2] \right| \leq K \frac{P}{d}.
 \end{aligned}$$

Step 5. Concluding the proof.

We can now combine inequalities (26), (30) and (30) to get

$$\begin{aligned}\Delta(t)^2 &= \int \{|a^t - \bar{a}^t| \vee \|\mathbf{u}^t - \bar{\mathbf{u}}^t\|_2 \vee \|\mathbf{v}^t\|_2 - \bar{s}^t\}^2 d\gamma_t \\ &\leq Ke^{KT^2} \frac{\log d}{d} + K \frac{P}{d} + K(1+T)^6 \int_0^T \Delta(t)^2 dt,\end{aligned}$$

which concludes the proof.

B.5. Auxiliary lemmas

Lemma 19 *Denote the residuals of the dynamics $\hat{g}(\mathbf{z}; \rho_t) = h_*(\mathbf{z}) - \hat{f}_{\text{NN}}(\mathbf{z}; \rho_t)$ and $\hat{g}(\mathbf{z}; \bar{\rho}_t) = h_*(\mathbf{z}) - \hat{f}_{\text{NN}}(\mathbf{z}; \bar{\rho}_t)$. By the properties of gradient flows, the risks*

$$\|\hat{g}(\cdot; \rho_t)\|_{L^2} \leq \|\hat{g}(\cdot; \rho_0)\|_{L^2} \leq K, \quad \|\hat{g}(\cdot; \bar{\rho}_t)\|_{L^2} \leq \|\hat{g}(\cdot; \bar{\rho}_0)\|_{L^2} \leq K. \quad (39)$$

In particular, this implies

$$\sup_{t \in [0, T]} \|a^t\|_\infty \vee \|\bar{a}^t\|_\infty \leq K(1+T), \quad (40)$$

$$\sup_{t \in [0, T]} \|\hat{g}(\cdot; \rho_t)\|_\infty \vee \|\hat{g}(\cdot; \bar{\rho}_t)\|_\infty \leq K(1+T). \quad (41)$$

Proof [Proof of Lemma 19] By definition $(\rho_t)_{t \geq 0}$ and $(\bar{\rho}_t)_{t \geq 0}$ are the solutions of a gradient flow:

$$\frac{d}{dt} R(\rho_t) = - \int \|\psi(\boldsymbol{\theta}; \rho_t)\|_2^2 \rho_t(d\boldsymbol{\theta}) \leq 0,$$

and therefore

$$R(\rho_t) = \|\hat{g}(\cdot; \rho_t)\|_{L^2}^2 \leq \|\hat{g}(\cdot; \rho_0)\|_{L^2}^2 \leq 2\|f_*\|_\infty^2 + 2\|a^0\|_\infty^2 \|\sigma\|_\infty^2 \leq K,$$

and similarly for $\|\hat{g}(\cdot; \bar{\rho}_t)\|_{L^2}$.

Furthermore, by Jensen inequality,

$$\left| \frac{d}{dt} a^t \right| = \left| \mathbb{E}_{\mathbf{z}, \mathbf{r}} [\hat{g}(\mathbf{z}; \rho_t) \sigma(\langle \mathbf{u}^t, \mathbf{z} \rangle + \langle \mathbf{v}^t, \mathbf{r} \rangle)] \right| \leq \|\sigma\|_\infty \|\hat{g}(\cdot; \rho_t)\|_{L^2} \leq K.$$

We deduce that

$$|a^t| \leq |a^0| + \int_0^t \left| \frac{d}{ds} a^s \right| ds \leq K(1+T).$$

A similar result holds for \bar{a}^t . Finally,

$$\|\hat{g}(\cdot; \rho_t)\|_\infty \leq \|a^t\|_\infty \|\sigma\|_\infty + \|f_*\|_\infty \leq K(1+T),$$

which concludes the proof. ■

Lemma 20 (Berry-Esseen bound in Wasserstein metric) *Let $(X_i)_{i \geq 1}$ be independent random variables with mean zero. Denote $v_n = \sum_{i \in [n]} \mathbb{E}[X_i^2]$ and $S = v_n^{-1/2} \sum_{i \in [n]} X_i$. Then*

$$W_1(S, G) \leq \frac{3}{v_n^{3/2}} \sum_{i \in [n]} \mathbb{E}[|X_i|^3], \quad (42)$$

where we denoted $G \sim \mathbf{N}(0, 1)$.

Proof [Proof of Lemma 20] This is a simple application of Stein's method. Consider f twice differentiable such that $\|f\|_\infty \leq 1$, $\|f'\|_\infty \leq \sqrt{2/\pi}$ and $\|f''\|_\infty \leq 2$. Introduce $S_i = S - v_n^{-1/2} X_i = v_n^{-1/2} \sum_{j \neq i} X_j$. By expanding, we get

$$\begin{aligned} \mathbb{E}[Sf(S)] &= v_n^{-1/2} \sum_{i \in [n]} \mathbb{E}[X_i f(S)] \\ &= v_n^{-1/2} \sum_{i \in [n]} \mathbb{E}[X_i (S - S_i) f'(S_i)] + M \\ &= v_n^{-1} \sum_{i \in [n]} \mathbb{E}[X_i^2] \mathbb{E}[f'(S_i)] + M, \end{aligned} \quad (43)$$

where, by Taylor's theorem,

$$\begin{aligned} |M| &= \left| v_n^{-1/2} \sum_{i \in [n]} \mathbb{E}[X_i \{f(S) - f(S_i) - (S - S_i) f'(S_i)\}] \right| \\ &\leq \frac{1}{2} v_n^{-1/2} \sum_{i \in [n]} \mathbb{E}[|X_i (S - S_i)|^2 \|f''\|_\infty] \\ &\leq v_n^{-3/2} \sum_{i \in [n]} \mathbb{E}[|X_i|^3]. \end{aligned} \quad (44)$$

Finally, note that

$$\begin{aligned} \left| v_n^{-1} \sum_{i \in [n]} \mathbb{E}[X_i^2] \mathbb{E}[f'(S_i)] - \mathbb{E}[f'(S)] \right| &\leq v_n^{-1} \|f''\|_\infty \sum_{i \in [n]} \mathbb{E}[X_i^2] \mathbb{E}[|S - S_i|] \\ &\leq 2v_n^{-3/2} \sum_{i \in [n]} \mathbb{E}[X_i^2] \mathbb{E}[|X_i|] \\ &\leq 2v_n^{-3/2} \sum_{i \in [n]} \mathbb{E}[|X_i|^3], \end{aligned} \quad (45)$$

where we used Jensen's inequality in the last line. Combining bounds Eqs. (44) and (45) in the identity (43) yields

$$\left| \mathbb{E}[Sf(S)] - \mathbb{E}[f'(S)] \right| \leq 3v_n^{-3/2} \sum_{i \in [n]} \mathbb{E}[|X_i|^3].$$

The result follows by Stein's lemma. ■

Lemma 21 *Let $d \geq 2$ be an integer. Consider $\{X_i\}_{i \in [d]}$ iid (τ^2/d) -sub-Gaussian random variables with 0 mean. Then for any $q \in \mathbb{N}$, there exists a universal constant $C_q > 0$ such that*

$$\mathbb{E} \left[\max_{i \in [d]} |X_i|^q \right] \leq C_q \left(\tau^2 \frac{\log d}{d} \right)^{q/2}.$$

Proof [Proof of Lemma 21] By sub-Gaussianity, there exists a universal constant $c > 0$ such that

$$\begin{aligned} \mathbb{P} \left(\max_{i \in [d]} |X_i|^q > t \right) &= 1 - \{1 - \mathbb{P}(|X_i| > t^{1/q})\}^d \leq 1 - \left\{1 - 2e^{-cdt^{2/q}/\tau^2}\right\}^d \\ &\leq 1 \wedge \left(2de^{-cdt^{2/q}/\tau^2}\right). \end{aligned}$$

Consider $t_c = \left(\tau^2 \frac{\log(2d)}{\kappa_q cd}\right)^{q/2}$ with $\kappa_q = 2^{2(q-2)/q^2 \vee 0}$, such that $2de^{-cd\kappa_q t_c^{2/q}/\tau^2} = 1$. Then, we have the following upper bound:

$$\begin{aligned} \mathbb{E} \left[\max_{i \in [d]} |X_i|^q \right] &= \int_0^\infty \mathbb{P} \left(\max_{i \in [d]} |X_i|^q > t \right) dt \\ &\leq t_c + \int_{t_c}^\infty 2de^{-cdt^{2/q}/\tau^2} dt \\ &\leq t_c + 2de^{-cd\kappa_q t_c^{2/q}/\tau^2} \int_0^\infty e^{-cd\kappa_q t^{2/q}/\tau^2} dt = \left(\tau^2 \frac{\log(2d)}{\kappa_q cd}\right)^{q/2} + c_q \left(\frac{\tau^2}{\kappa_q cd}\right)^{q/2}, \end{aligned}$$

which concludes the proof. ■

Appendix C. Strong SGD-learnability in the discrete-time regime

In this appendix, we define strong SGD-learnability in the discrete-time regime, i.e., for large batch size b and large η . We keep the same assumptions A0-A2, and replace Assumption A3 by

D3. (*Boundedness of hyperparameters*) We have $\eta_k^a, \eta_k^w \leq K$ and $\lambda^a, \lambda^w \leq K$.

While the continuous-time regime requires step size η to be small enough compared to n/b , the discrete-time regime requires the batch size b to be big enough compared to cn for $c \ll 1$ (recall $b \leq n$ by one-pass assumption) in the discrete regime.

Definition 22 (Strong SGD-learnability in $O(d)$ -scaling (discrete time)) *We say that a function $h_* : \{-1, +1\}^P \rightarrow \mathbb{R}$ is strongly $O(d)$ -SGD-learnable if the following hold for some $C(\cdot, h_*)$, $T(\cdot, h_*) : \mathbb{R}_{>0} \rightarrow \mathbb{R}_{>0}$. For any $\varepsilon > 0$, $d \geq C(\varepsilon, h_*)$, $n \geq C(\varepsilon, h_*)d$ and $e^d \geq N \geq C(\varepsilon, h_*)$, there exists hyperparameters $(\sigma, b, \lambda^a, \lambda^w, \{\eta_k^a, \eta_k^w\}_{k \in [0, k_0]})$ and initialization ρ_0 satisfying A0-A2, D3 and $k_0 = n/b \leq T(\varepsilon, h_*)$ such that for any $\mathcal{I} \subseteq [d]$, $|\mathcal{I}| = P$ and target function $f_*(\mathbf{x}) = h_*(\mathbf{x}_{\mathcal{I}})$, k_0 steps of batch stochastic gradient descent (**BSGD**) achieves test error ε with probability at least $9/10$.*

Again, conditions A0-A2, D3 guarantee that as long as d, n, N are taken sufficiently large, there exist a discrete mean-field dynamics that well-approximates batch-SGD up to a constant number of steps that depends on ε, h_* .

C.1. Discrete time mean-field and dimension-free dynamics

We first give the discrete time mean-field dynamics to which batch-SGD converges. Recall that when $N \rightarrow \infty$ and $\eta \rightarrow 0$, the dynamics converge to the continuous (**MF-PDE**). Here instead, we fix the step sizes and consider $N, b \rightarrow \infty$, and get the following discrete mean-field dynamics $(\rho_k)_{k \geq 0}$ (with $\rho_k \in \mathcal{P}(\mathbb{R}^{d+1})$) described by the initialization $(a^0, \mathbf{w}^0) \sim \rho_0$ and the recurrence relation: $(a^{k+1}, \mathbf{w}^{k+1}) \sim \rho_{k+1}$ the distribution of the updated weights

$$\begin{aligned} a^{k+1} &= (1 - \eta_k^a \lambda^a) a^k + \eta_k^a \mathbb{E}_{\mathbf{x}} [\{f_*(\mathbf{x}) - \hat{f}_{\text{NN}}(\mathbf{x}; \rho_k)\} \sigma(\langle \mathbf{x}, \mathbf{w}^k \rangle)], \\ \mathbf{w}^{k+1} &= (1 - \eta_k^w \lambda^w) \mathbf{w}^k + \eta_k^w a^k \mathbb{E}_{\mathbf{x}} [\{f_*(\mathbf{x}) - \hat{f}_{\text{NN}}(\mathbf{x}; \rho_k)\} \sigma'(\langle \mathbf{x}, \mathbf{w}^k \rangle) \mathbf{x}], \end{aligned} \quad (\text{d-MF-PDE})$$

where $(a^k, \mathbf{w}^k) \sim \rho_k$.

Similarly to the continuous regime, the discrete dynamics simplify when $d \rightarrow \infty$ with P fixed, to the following discrete dimension-free dynamics $(\bar{\rho}_k)_{k \geq 0}$ (with $\bar{\rho}_k \in \mathcal{P}(\mathbb{R}^{P+2})$) defined by the initialization $(\bar{a}^0, \bar{\mathbf{u}}^0, \bar{s}^0) \sim \bar{\rho}_0$ (with $\bar{a}^0 \sim \mu_a$, $\bar{\mathbf{u}}^0 = \mathbf{0}$ and $\bar{s}^0 = m_2^w$) and the recurrence relation

$$\begin{aligned} \bar{a}^{k+1} &= (1 - \eta_k^a \lambda^a) \bar{a}^k + \eta_k^a \mathbb{E}_{\mathbf{z}, G} [\{h_*(\mathbf{z}) - \hat{f}_{\text{NN}}(\mathbf{z}; \bar{\rho}_k)\} \sigma(\langle \mathbf{z}, \bar{\mathbf{u}}^k \rangle + \bar{s}^k G)], \\ \bar{\mathbf{u}}^{k+1} &= (1 - \eta_k^w \lambda^w) \bar{\mathbf{u}}^k + \eta_k^w \bar{a}^k \mathbb{E}_{\mathbf{z}, G} [\{h_*(\mathbf{z}) - \hat{f}_{\text{NN}}(\mathbf{z}; \bar{\rho}_k)\} \sigma'(\langle \mathbf{z}, \bar{\mathbf{u}}^k \rangle + \bar{s}^k G) \bar{\mathbf{u}}^k], \\ \bar{s}^{k+1} &= (1 - \eta_k^w \lambda^w) \bar{s}^k + \eta_k^w \bar{a}^k \mathbb{E}_{\mathbf{z}, G} [\{h_*(\mathbf{z}) - \hat{f}_{\text{NN}}(\mathbf{z}; \bar{\rho}_k)\} \sigma'(\langle \mathbf{z}, \bar{\mathbf{u}}^k \rangle + \bar{s}^k G) G]. \end{aligned} \quad (\text{d-DF-PDE})$$

We have the new non-asymptotic bound between the (**BSGD**) and (**d-DF-PDE**) dynamics, analogous to Theorem 5, but with a worse dependency on the number of iterations.

Theorem 23 *Assume conditions A0-A2,D3 hold, and let $k_0 \geq 0$. There exists a constant K depending only on the constants in A0-A2,D3 (in particular, independent of d, P, k_0), such that*

$$\begin{aligned} & \sup_{k=0, \dots, k_0} \left\| \hat{f}_{\text{NN}}(\cdot; \Theta^k) - \hat{f}_{\text{NN}}(\cdot; \bar{\rho}_k) \right\|_{L^2} \\ & \leq K e^{e^{K k_0^2}} \left\{ \sqrt{\frac{P + \log(d)}{d}} + \sqrt{\frac{\log N}{N}} + \sqrt{\frac{d + \log N}{b}} \right\}, \end{aligned} \quad (46)$$

with probability at least $1 - 1/N$.

From there, it is straightforward, following the same arguments as for Theorems 6 and 7, to get the equivalence of strong $O(d)$ -SGD-learnability in the discrete-time regime and global convergence of the discrete (**d-DF-PDE**) dynamics, and the MSP necessary condition:

Theorem 24 *A function $h_* : \{+1, -1\}^P \rightarrow \mathbb{R}$ is strongly $O(d)$ -SGD-learnable in the discrete-time regime if and only if for any $\varepsilon > 0$, there exists $\lambda^a, \lambda^w \geq 0$ and bounded step-sizes $\{\eta_k^a, \eta_k^w\}_{k \geq 0}$, such that $\inf_{k \in \mathbb{N}} R(\bar{\rho}_k) < \varepsilon$, where $\bar{\rho}_k$ is the solution of the discrete (**d-DF-PDE**) dynamics.*

Theorem 25 *Let $h_* : \{+1, -1\}^P \rightarrow \mathbb{R}$ be a function without MSP. Then there exists $c > 0$ such that for any regularizations $\lambda^a, \lambda^w \geq 0$ and step-sizes $\{\eta_k^a, \eta_k^w\}_{k \geq 0}$, we have $\inf_{k \in \mathbb{N}} R(\bar{\rho}_k) \geq c$.*

C.2. Proof of Theorem 23

The proof relies on first comparing the (**BSGD**) dynamics to the discrete mean-field dynamics (**d-MF-PDE**), using an extension of the results in Mei et al. (2019) to the discrete (**d-DF-PDE**) dynamics (see Appendix I.1).

Proposition 26 *Assume conditions A0-A2,D3, and let $k_0 \in \mathbb{N}$. There exists a constant K depending only on the constants in A0-A2,D3 (in particular, independent of d, P, T), such that*

$$\sup_{k=0, \dots, k_0} \left\| \hat{f}_{\text{NN}}(\cdot; \Theta^k) - \hat{f}_{\text{NN}}(\cdot; \rho_k) \right\|_{L^2} \leq K e^{e^{K k_0}} \left\{ \sqrt{\frac{\log N}{N}} + \sqrt{\frac{d + \log N}{b}} \right\},$$

with probability at least $1 - 1/N$.

The proof of this proposition follows from applying Proposition 92, with the assumptions already verified in Appendix B.3.

The proof of Theorem 23 then follows by combining the above result with the following bound between the discrete mean-field dynamics (**d-MF-PDE**) and the discrete dimension-free dynamics (**d-DF-PDE**):

Theorem 27 *Assume conditions A0-A2,D3, and let $k_0 \geq 0$. There exists a constant K depending only on the constants in A0-A2,D3 (in particular, independent of d, P, T), such that*

$$\sup_{k=0, \dots, k_0} \left\| \hat{f}_{\text{NN}}(\cdot; \rho_k) - \hat{f}_{\text{NN}}(\cdot; \bar{\rho}_k) \right\|_{L^2} \leq K e^{e^{K k_0^2}} \sqrt{\frac{P + \log(d)}{d}}.$$

Proof [Proof of Theorem 27] The proof follows similarly to the proof in the continuous case (see Section B.4) and we will simply highlight the differences. First, by the same argument as in the proof of Proposition 92, we replace the bounds from Lemma 19 by

$$\sup_{k=0,\dots,k_0} \|a^k\|_\infty \vee \|\bar{a}^k\|_\infty \leq K e^{Kk_0}, \quad (47)$$

$$\sup_{k=0,\dots,k_0} \|\hat{g}(\cdot; \rho_k)\|_\infty \vee \|\hat{g}(\cdot; \bar{\rho}_k)\|_\infty \leq K e^{Kk_0}. \quad (48)$$

and from Lemma 17 by

$$\sup_{k=0,\dots,k_0} \max_{i \in [d-P]} |v_i^k| \leq K e^{e^{Kk_0}} \max_{i \in [d-P]} |v_i^0|.$$

We define

$$\delta(k) = |a^k - \bar{a}^k| \vee \|\mathbf{u}^k - \bar{\mathbf{u}}^k\|_2 \vee |\bar{s}^k - \|\mathbf{v}^k\|_2|, \quad (49)$$

and the square root of its second moment

$$\Delta(k) = \left(\int \delta(k)^2 \gamma_t(d\boldsymbol{\theta}^k d\bar{\boldsymbol{\theta}}^k) \right)^{1/2}. \quad (50)$$

The proof follows by using discrete Grönwall lemma in Lemma 28 stated in the next section, which is the analogous of Lemma 18 in discrete time. \blacksquare

C.3. Auxiliary lemma

Lemma 28 *Consider the same setting and assumptions as Theorem 27. There exists a constant K independent of d, P and depending only on the Assumptions A0-A2, D3 such that for any $k_0 \in \mathbb{N}$,*

$$\begin{aligned} \|\hat{f}_{\text{NN}}(\cdot; \rho_k) - \hat{f}_{\text{NN}}(\cdot; \bar{\rho}_k)\|_{L^2} &= \mathbb{E}_{\mathbf{z}} \left[\left\{ \hat{g}(\mathbf{z}; \rho_k) - \hat{g}(\mathbf{z}; \bar{\rho}_k) \right\}^2 \right]^{1/2} \\ &\leq K e^{Kk_0} \Delta(k) + K e^{e^{Kk_0}} \sqrt{\frac{\log d}{d}}, \end{aligned} \quad (51)$$

where

$$\Delta(k) \leq K e^{e^{Kk_0}} \sqrt{\frac{P + \log d}{d}} + K e^{Kk_0^2} \sum_{j=0}^{k-1} \Delta(j). \quad (52)$$

Proof [Proof of Lemma 28] The proof proceeds similarly to the proof of Lemma 18 in Section B.4.1, where we use discrete Grönwall instead. Step 1 to Step 3 are very similar, using that

$$\begin{aligned} &\|\mathbf{u}^{k+1} - \bar{\mathbf{u}}^{k+1}\|_2^2 \\ &\leq \|\mathbf{u}^k - \bar{\mathbf{u}}^k\|_2^2 + \langle \mathbf{u}^{k+1} - \bar{\mathbf{u}}^{k+1}, (\mathbf{u}^{k+1} - \mathbf{u}^k) - (\bar{\mathbf{u}}^{k+1} - \bar{\mathbf{u}}^k) \rangle. \end{aligned}$$

For Step 4, notice that, denoting $\mathbf{P}^k = a^k \mathbb{E}_{\mathbf{z}, \mathbf{r}} [\hat{g}(\mathbf{z}; \rho_k) \sigma'(\langle \mathbf{u}^k, \mathbf{z} \rangle + \langle \mathbf{v}^k, \mathbf{r} \rangle) \mathbf{r}]$,

$$\|\mathbf{v}^{k+1}\|_2^2 = \|\mathbf{v}^k\|_2^2 + 2\eta_k \langle \mathbf{v}^k, \mathbf{P}^k \rangle + \eta_k^2 \|\mathbf{P}^k\|_2^2.$$

Note that integrating out r_i and with mean-value theorem, we get $P_i^k = v_i^k \mathbb{E}_{\mathbf{z}, \mathbf{r}} [\hat{g}(\mathbf{z}; \rho_k) \sigma''(\langle \mathbf{u}^k, \mathbf{z} \rangle + \langle \mathbf{v}_{-i}^k, \mathbf{r}_{-i} \rangle + \xi_i)]$. Denote $m_k = a^k \mathbb{E}_{\mathbf{z}, \mathbf{r}} [\hat{g}(\mathbf{z}; \rho_k) \sigma''(\langle \mathbf{u}^k, \mathbf{z} \rangle + \langle \mathbf{v}^k, \mathbf{r} \rangle)]$ and $\mathbf{T}^k = (T_i^k) = \mathbf{P}^k - m_k \mathbf{v}^k$ with

$$T_i^k = a^k v_i^k \mathbb{E}_{\mathbf{z}, \mathbf{r}} [\hat{g}(\mathbf{z}; \rho_k) (\sigma'(\langle \mathbf{u}^k, \mathbf{z} \rangle + \langle \mathbf{v}^k, \mathbf{r} \rangle + \xi_i) - \sigma''(\langle \mathbf{u}^k, \mathbf{z} \rangle + \langle \mathbf{v}^k, \mathbf{r} \rangle))].$$

Then, we can decompose

$$\|\mathbf{v}^{k+1}\|_2^2 = \|\mathbf{v}^k\|_2^2 (1 + \eta_k m_k)^2 + 2\eta_k (1 + \eta_k m_k) \langle \mathbf{v}^k, \mathbf{T}^k \rangle + \eta_k^2 \|\mathbf{T}^k\|_2^2.$$

Denote $\bar{m}_k = \bar{a}^k \mathbb{E}_{\mathbf{z}, G} [\hat{g}(\mathbf{z}; \bar{\rho}_k) \sigma''(\langle \bar{\mathbf{u}}^k, \mathbf{z} \rangle + \bar{s}^k G)]$ and note that $\bar{s}^{k+1} = \bar{s}^k (1 + \eta_k \bar{m}_k)$ (using Gaussian integration by parts).

We decompose:

$$\left| \|\mathbf{v}^{k+1}\|_2 - \bar{s}^{k+1} \right| \leq \text{(I)} + \text{(II)},$$

where

$$\begin{aligned} \text{(I)} &= \|\mathbf{v}^k\|_2 \left| \sqrt{(1 + \eta_k m_k)^2 + 2\eta_k (1 + \eta_k m_k) \langle \mathbf{v}^k / \|\mathbf{v}^k\|_2, \mathbf{T}^k \rangle + \eta_k^2 \|\mathbf{T}^k\|_2^2 / \|\mathbf{v}^k\|_2^2} - |1 + \eta_k m_k| \right|, \\ \text{(II)} &= \left| \|\mathbf{v}^k\|_2 |1 + \eta_k m_k| - \bar{s}^k |1 + \bar{m}_k| \right|. \end{aligned}$$

The first term is bounded by

$$\text{(I)} \leq 2\eta_k |\langle \mathbf{v}^k / \|\mathbf{v}^k\|_2, \mathbf{T}^k \rangle| + \eta_k \|\mathbf{T}^k\|_2.$$

Note that $\langle \mathbf{v}^k / \|\mathbf{v}^k\|_2, \mathbf{T}^k \rangle = M_k$ as defined in Eq. (34) and we can use the bound in Eq. (36):

$$|\langle \mathbf{v}^k / \|\mathbf{v}^k\|_2, \mathbf{T}^k \rangle| \leq K e^{Kk_0} [\|\mathbf{v}^0\|_2 + 1] \max_{i \in [d-P]} |v_i^k|.$$

Similarly, we have

$$\begin{aligned} \|\mathbf{T}^k\|_2^2 &= (a^k)^2 \sum_{i \in [d-P]} (v_i^k)^2 \\ &\quad \times \mathbb{E}_{\mathbf{z}, \mathbf{r}} \left[\hat{g}(\mathbf{z}; \rho_k) \left\{ \sigma''(\langle \mathbf{u}^k, \mathbf{z} \rangle + \langle \mathbf{v}_{-i}^k, \mathbf{r}_{-i} \rangle + \xi_i) - \sigma''(\langle \mathbf{u}^k, \mathbf{z} \rangle + \langle \mathbf{v}^k, \mathbf{r} \rangle) \right\} \right]^2 \\ &\leq 2 \|a^k\|_\infty^2 \|\hat{g}(\cdot; \rho_k)\|_{L^2}^2 \|\sigma'''\|_\infty^2 \sum_{i \in [d-P]} |v_i^k|^4 \\ &\leq K \|a^k\|_\infty^2 \|\mathbf{v}^k\|_2^2 \max_{i \in [d-P]} |v_i^k|^2 \leq K e^{Kk_0} [\|\mathbf{v}^0\|_2^2 + 1] \max_{i \in [d-P]} |v_i^k|^2. \end{aligned}$$

We deduce

$$\text{(I)} \leq K e^{Kk_0} [\|\mathbf{v}^0\|_2 + 1] \max_{i \in [d-P]} |v_i^k|. \quad (53)$$

We can further bound (II) using the same decomposition as in Eq. (35):

$$\text{(II)} \leq K e^{Kk_0} [\delta(k) + \Delta(k) + W_1(\langle \mathbf{v}^t, \mathbf{r} \rangle, \|\mathbf{v}^t\|_2 G)] + K e^{Kk_0} \sqrt{\frac{\log d}{d}}. \quad (54)$$

Combining Eqs. (53) and (54), we obtain

$$\Delta(k+1)^2 \leq K e^{Kk_0} \Delta(k)^2 + K e^{Kk_0} \frac{\log d}{d} + K e^{Kk_0} \frac{P}{d}.$$

This concludes the proof. \blacksquare

Appendix D. Vanilla staircase functions are strongly $O(d)$ -SGD-learnable: Proof of Theorem 10

We start by providing the proof that vanilla staircases are strongly $O(d)$ -SGD-learnable, as described in Theorem 10. This proof will outline the main ideas behind our global convergence results, without the technical complexity of dealing with general MSP set structure.

Assumption on activation function We will assume the following hold for the activation σ :

A0'. Let $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ be an activation function that satisfies Assumption A0. Furthermore, assume that for some $L \in \mathbb{N}$ and $\eta > 0$ such that σ is $L + 1$ differentiable on $(-\eta, \eta)$ with

$$\sup_{x \in (-\eta, \eta)} |\sigma^{(L+1)}(x)| \leq K.$$

We will denote $m_r = \sigma^{(r)}(0)$ and $\mathbf{m} = (m_0, \dots, m_L) \in \mathbb{R}^{L+1}$.

In particular, this assumption implies that we have the following polynomial approximations of σ and σ' around 0: for any $x \in (-\eta, \eta)$,

$$\begin{aligned} \left| \sigma(x) - \sum_{r=0}^L \frac{m_r}{r!} x^r \right| &\leq K|x|^{L+1}, \\ \left| \sigma'(x) - \sum_{r=0}^{L-1} \frac{m_{r+1}}{r!} x^r \right| &\leq K|x|^L. \end{aligned} \tag{55}$$

The Assumption A0' is simply to connect with the definition of strong $O(d)$ -SGD-learnability. However, in the proof below, we will choose hyperparameters such that $|\langle \mathbf{u}^t, \mathbf{z} \rangle| < \eta$ (the input of the activation) during the whole dynamics, so that A0 can be lifted. In particular, any activation that is $\mathcal{C}^{L+1}(\mathbb{R})$ will satisfy A0'.

Now recall that vanilla staircases $h_* : \{+1, -1\}^P \rightarrow \mathbb{R}$ are those functions of the form

$$h_*(\mathbf{z}) = \alpha_{\{1\}} z_1 + \alpha_{\{1,2\}} z_1 z_2 + \alpha_{\{1,2,3\}} z_1 z_2 z_3 + \dots + \alpha_{\{1,\dots,P\}} z_1 z_2 \dots z_P. \tag{56}$$

for some Fourier coefficients $\alpha_{\{1\}}, \dots, \alpha_{\{1,\dots,P\}} \in \mathbb{R} \setminus \{0\}$. We will prove that any such function h_* is strongly SGD-learnable in the $O(d)$ -scaling.

Choice of hyperparameters Recall from the equivalence with (DF-PDE) (Theorem 6) that it is sufficient to show for any $\varepsilon > 0$, there exist hyperparameters satisfying A0-A2, A3' such that (DF-PDE) dynamics reaches ε -risk. We consider the following hyperparameters:

- We do not regularize, i.e., $\lambda^a = \lambda^w = 0$.
- We initialize the first layer to deterministically 0 weights, and the second layer to uniform random weights. I.e., we take $\mu_a = \text{Unif}([+1, -1])$ and $\mu_w = \delta_0$. Although initializing the first layer to 0 may at first glance seem restrictive, there turns out to be enough randomness in the initialization of the second layer to ensure that the neural network learns. For the dimension-free dynamics, this corresponds to taking $(\bar{a}^0, \bar{\mathbf{u}}^0, \bar{s}^0) \sim \bar{\rho}_0$ with $\bar{a}^0 \sim \text{Unif}([+1, -1])$, $\bar{\mathbf{u}}^0 = \mathbf{0}$ and $\bar{s}^0 = 0$. In particular, $\bar{s}^t = 0$ during the whole dynamics, which allows for a simpler analysis.

- Our learning rate schedule has two phases:

Phase 1: We train the first layer weights $\bar{\mathbf{u}}^t$ while keeping the second layer weights fixed $\bar{\mathbf{a}}^t = \bar{\mathbf{a}}^0$. We set $\xi^a(t) = 0$ and $\xi^w(t) = 1$ for $t \in [0, T_1]$.

Phase 2: We train the second layer weights $\bar{\mathbf{a}}^t$ while keeping the first layer weights fixed at $\bar{\mathbf{u}}^t = \bar{\mathbf{u}}^{T_1}$. We set $\xi^a(t) = 1$ and $\xi^w(t) = 0$ for $t \in [T_1, T_2]$.

Remark 29 As written above, the learning rate schedules ξ^a, ξ^w are not Lipschitz at T_1 . Note that we can always do the following change of time variable on $[0, T_1]$: $t' = 2tT_1 - t^2$ such that $\xi^w(t') = 2(\sqrt{T_1} - t)_+$ is Lipschitz on $\mathbb{R}_{\geq 0}$ (and we have now $T_1' = T_1^2$). Similarly, we can do a change of time variable on $[T_1', T_2']$ such that $\xi_a(t) = 2 \min((t - \sqrt{T_1})_+, 1)$. We will proceed with the simpler learning schedule $\xi^a(t) = \mathbb{1}_{t \geq T_1}$ and $\xi^w(t) = \mathbb{1}_{t \leq T_1}$ with the understanding that we can do the above change of variables to obtain Lipschitz learning schedules and therefore fall under the assumptions of strong SGD learnability.

We restate the sufficient condition in the case of the vanilla staircase.

Theorem 30 (Theorem 10 restated) Suppose that $h_* : \{+1, -1\}^P \rightarrow \mathbb{R}$ has the vanilla staircase structure (56). Suppose also that the activation function σ has nonzero derivatives $m_r \neq 0$ for $r = 0, \dots, P$ and satisfies A0' for some $L > 2^{P-1}$. Then, for any $\varepsilon > 0$, there are $T_1, T_2 > 0$ such that training the dimension-free PDE with the above hyperparameters will learn h_* to accuracy ε . Therefore, h_* is strongly $O(d)$ -SGD-learnable.

D.1. Outline of the proof

Consider $(\bar{\rho}_t)_{t \geq 0}$ the solution of (DF-PDE) with the hyperparameters described above. Denote $\bar{\mathbf{u}}^t(\bar{\mathbf{a}}^0)$ the solution of the evolution equations (9) obtained from initialization $(\bar{\mathbf{a}}^0, \mathbf{0}, 0)$. For clarity, we will suppress some notations in the proof: we will denote \mathbf{u} instead of $\bar{\mathbf{u}}$, and a instead of \bar{a}^0 . We will further forget about $\bar{\mathbf{s}}^t = 0$ and simply consider $\bar{\rho}_t \in \mathcal{P}(\mathbb{R}^{P+1})$ the distribution of (a, \mathbf{u}) . This last simplification can be done since we initialize the first-layer weights to 0, so in particular $\bar{\mathbf{s}}^0 = 0$, and by the evolution equation of (9) we have $\bar{\mathbf{s}}^t = 0$ throughout training. Furthermore, we will denote K a generic constant that only depends on P and the constants in the assumptions. The value of K can change from line to line.

The proof analyzes Phase 1 and Phase 2 of training separately.

Phase 1 (nonlinear dynamics): In this phase, we train the first layer, which has nonlinear dynamics, and so it is a priori unclear how to analyze. Nevertheless, since h_* is specially structured, the structure in the weights during training is particularly simple and it is enough to track the smallest order terms in the weights.

Specifically, in Proposition 33 (see next section), we prove that there exist constants $c, C > 0$ such that for all $t \leq c$ and $k \in [P]$, we have $|u_k^t(a) - \hat{u}_k^t(a)| \leq Ct^{2^{k-1}+1}$, where

$$\hat{u}_k^t(a) = 2^{1-2^{k-1}} (at)^{2^{k-1}} \cdot \prod_{i \in [k]} (m_i \alpha_{\{1, \dots, i\}})^{2^{\max((k-1-i), 0)}}. \quad (57)$$

Denote $\nu_k(t) = 2^{1-2^{k-1}} t^{2^{k-1}} \prod_{i \in [k]} (m_i \alpha_{\{1, \dots, i\}})^{2^{\max((k-1-i), 0)}}$ such that $u_k^t(a) = \nu_k(t)(a^{2^{k-1}} + O(t))$.

Phase 2 (linear dynamics): In this phase, we train the second layer, and the training has linear dynamics. Denote $g_t(\mathbf{z}) := h_*(\mathbf{z}) - \hat{f}_{\text{NN}}(\mathbf{z}; \bar{\rho}_t)$ the residual function at time t . During this phase, we have the following evolution on the risk:

$$\frac{d}{dt} R(\bar{\rho}_t) = -\mathbb{E}_{\mathbf{z}, \mathbf{z}'} [g_t(\mathbf{z}) K^{T_1}(\mathbf{z}, \mathbf{z}') g_t(\mathbf{z}')], \quad (58)$$

where the kernel is given by

$$K^{T_1}(\mathbf{z}, \mathbf{z}') = \mathbb{E}_{a \sim \mu_a} [\sigma(\langle \mathbf{u}^{T_1}(a), \mathbf{z} \rangle) \sigma(\langle \mathbf{u}^{T_1}(a), \mathbf{z}' \rangle)]. \quad (59)$$

(This is indeed the kernel, since at the end of Phase 1, the distribution $\bar{\rho}_{T_1}$ of the parameters is given by $(a, \mathbf{u}^{T_1}(a))$ with $a \sim \mu_a$, and the first-layer weights are kept constant during Phase 2.)

Let us decompose these quantities in the Fourier basis: denote $g_t(S) = \mathbb{E}_{\mathbf{z}} [g_t(\mathbf{z}) \chi_S(\mathbf{z})]$ and $K^{T_1}(S, S') = \mathbb{E}_{\mathbf{z}, \mathbf{z}'} [K^{T_1}(\mathbf{z}, \mathbf{z}') \chi_S(\mathbf{z}) \chi_{S'}(\mathbf{z}')]$, and the vector $\mathbf{g}_t = (g_t(S))_{S \subseteq [P]}$ and matrix $\mathbf{K}^{T_1} = (K^{T_1}(S, S'))_{S, S' \subseteq [P]}$. Noting that $R(\bar{\rho}_t) = \|\mathbf{g}_t\|_2^2$, we have

$$\frac{d}{dt} \|\mathbf{g}_t\|_2^2 = -\mathbf{g}_t^\top \mathbf{K}^{T_1} \mathbf{g}_t \leq -\lambda_{\min}(\mathbf{K}^{T_1}) \|\mathbf{g}_t\|_2^2. \quad (60)$$

This implies that for $t \geq T_1$, we have $\|\mathbf{g}_t\|_2^2 \leq e^{-\lambda_{\min}(\mathbf{K}^{T_1})(t-T_1)} \|\mathbf{g}_{T_1}\|_2^2$. By assumption we have $\|\mathbf{g}_{T_1}\|_2^2 \leq K$ for a constant K .⁷ We deduce that, if we prove that $\lambda_{\min}(\mathbf{K}^{T_1}) > c$ for some constant $c > 0$, then it is sufficient to consider $T_2 = T_1 + \log(K/\varepsilon)/c$, to guarantee that $R(\bar{\rho}_{T_2}) \leq \varepsilon$. This would conclude the proof of strong $O(d)$ -SGD-learnability.

Lower bound on $\lambda_{\min}(\mathbf{K}^{T_1})$: It only remains to lower-bound $\lambda_{\min}(\mathbf{K}^{T_1})$. For this we use the structure on \mathbf{u}^{T_1} that we prove holds in Phase 1. For all $S \subseteq [P]$, denote

$$\nu_S(T_1) = \prod_{k \in S} \nu_k(T_1) \text{ and } \beta(S) = \sum_{k \in S} 2^{k-1}.$$

From Lemma 32 (see next section), there exists a constant $C > 0$ depending only on P (and independent of T_1) such that for any $S \subseteq [P]$,

$$\left| \mathbb{E}_{\mathbf{z}} [\chi_S(\mathbf{z}) \sigma(\langle \mathbf{u}^{T_1}(a), \mathbf{z} \rangle)] - m_{|S|} \nu_S(T_1) a^{\beta(S)} \right| \leq C m_{|S|} \nu_S(T_1) T_1.$$

Denote $D_S = m_{|S|} \nu_S(T_1)$ and $\mathbf{D} = \text{diag}((D_S)_{S \subseteq [P]})$. We have

$$|K^{T_1}(S, S') - D_S D_{S'} \mathbb{E}_a [a^{\beta(S) + \beta(S')}]| \leq C D_S D_{S'} T_1.$$

Introduce $\mathbf{M} = (\mathbb{E}_{a \sim \mu_a} [a^{\beta(S) + \beta(S')}]_{S, S' \subseteq [P]})$, then we have

$$\mathbf{K}^{T_1} = \mathbf{D}(\mathbf{M} + \mathbf{\Delta})\mathbf{D},$$

where $\|\mathbf{\Delta}\|_{\text{op}} \leq C T_1 P$.

Note that $\beta(S)$ takes value $\{0, \dots, 2^P - 1\}$, and \mathbf{M} is the Gram matrix of the monomials $[1, X, \dots, X^{2^P-1}]$ in $L^2([+1, -1], \text{Unif})$, which are linearly independent. We deduce that $\lambda_{\min}(\mathbf{M})$ is bounded away from 0 (independent of T_1). We can therefore take $T_1 \leq \lambda_{\min}(\mathbf{M})/(2P)$, so that $\lambda_{\min}(\mathbf{M} + \mathbf{\Delta}) \geq \lambda_{\min}(\mathbf{M})/2$, and $\lambda_{\min}(\mathbf{K}^{T_1}) \geq \{\min_S D_S^2\} \lambda_{\min}(\mathbf{M})/2 > 0$.

7. This is since $\|h_*\|_\infty \leq K$ and $\|\hat{f}_{\text{NN}}(\cdot; \bar{\rho}_t)\|_\infty \leq \mathbb{E}[|a| \|\sigma\|_\infty] \leq K$, because we choose initialization with $|a| \leq 1$. So $\|\mathbf{g}_{T_1}\|_2 \leq \sqrt{2^P} \|\mathbf{g}_{T_1}\|_\infty \leq \sqrt{2^P} (2K)$.

D.2. Approximating the u_i^t with polynomials

First, we have the following simple bound on $\|\mathbf{u}^t\|_1$:

Claim 31 *There exists a constant C depending on K, P such that $\|\mathbf{u}^t\|_1 \leq Ct$.*

Proof [Proof of Claim 31] By Assumptions A0 and A1, we have $\|\sigma'\|_\infty, \|h_*\|_\infty \leq K$ and also $\|\hat{f}_{\text{NN}}(\cdot; \bar{\rho}_t)\|_\infty \leq \int |a| \|\sigma\|_\infty d\mu_a \leq K$. Combining these bounds, we get for $t \leq T_1$:

$$\left| \frac{d}{dt} u_k^t \right| = |a \mathbb{E}_{\mathbf{z}} [g_t(\mathbf{z}) \sigma'(\langle \mathbf{u}^t, \mathbf{z} \rangle) z_k]| \leq K$$

and therefore $\frac{d}{dt} \|\mathbf{u}^t\|_1 \leq PK$. Recalling, $\mathbf{u}^0 = \mathbf{0}$, we conclude $\|\mathbf{u}^t\|_1 \leq KPt$. \blacksquare

The following lemma give the leading order in t approximation of the Fourier coefficients of $\sigma(\langle \mathbf{u}^t, \mathbf{z} \rangle)$:

Lemma 32 *There exists a constant $c > 0$ that depend on η, K, P such that for any $t \leq c$, $S \subseteq [P]$ and $i \in \{0, 1\}$,*

$$\mathbb{E}_{\mathbf{z}} [\chi_S(\mathbf{z}) \sigma^{(i)}(\langle \mathbf{u}^t, \mathbf{z} \rangle)] = m_{|S|+i} \left(\prod_{k \in S} u_k^t \right) \cdot (1 + O(t)) + O(t^L). \quad (61)$$

Proof [Proof of Lemma 32] From Claim 31, we can choose c sufficiently small such that $\|\mathbf{u}^t\|_1 < \eta$, and $|\langle \mathbf{z}, \mathbf{u}^t \rangle| < \eta$. We can therefore use the polynomial approximation Eq. (55) of $\sigma^{(i)}$:

$$\mathbb{E}_{\mathbf{z}} [\chi_S(\mathbf{z}) \sigma(\langle \mathbf{u}^t, \mathbf{z} \rangle)] = \sum_{r=0}^{L-1} \frac{m_r}{r!} \mathbb{E}_{\mathbf{z}} [\chi_S(\mathbf{z}) \langle \mathbf{u}^t, \mathbf{z} \rangle^r] + O(t^L). \quad (62)$$

Note that $\mathbb{E}_{\mathbf{z}} [\chi_S(\mathbf{z}) \langle \mathbf{u}^t, \mathbf{z} \rangle^r] = 0$ for $r < |S|$, $\mathbb{E}_{\mathbf{z}} [\chi_S(\mathbf{z}) \langle \mathbf{u}^t, \mathbf{z} \rangle^r] = r! \prod_{k \in S} u_k^t$ for $r = |S|$, and for $|S| = l < r$ (such that $r - l = 2s$)

$$\begin{aligned} |\mathbb{E}_{\mathbf{z}} [\chi_S(\mathbf{z}) \langle \mathbf{u}^t, \mathbf{z} \rangle^r]| &\leq \frac{r!}{(2s)!} \cdot \left(\prod_{k \in S} u_k^t \right) \frac{(2s)!}{s! 2^s} \sum_{i_1, \dots, i_s \in [P]} (u_{i_1}^t)^2 \cdots (u_{i_s}^t)^2 \\ &= \frac{r!}{s! 2^s} \cdot \left(\prod_{k \in S} u_k^t \right) \cdot \|\mathbf{u}^t\|_2^{2s} = \left(\prod_{k \in S} u_k^t \right) \cdot O(t). \end{aligned}$$

Injecting these bounds in Eq. (62) yields the result. \blacksquare

We can now prove the main structural result on the \mathbf{u}^t , on which the rest of the proof relies.

Proposition 33 *Assume $L > 2^{P-1}$ and denote*

$$\hat{u}_k^t(a) = 2^{1-2^{k-1}} (at)^{2^{k-1}} \cdot \left\{ \prod_{i \in [k]} (\alpha_{\{1, \dots, i\}} m_i)^{2^{(k-1-i) \vee 0}} \right\}. \quad (63)$$

There exists constants $c, C > 0$ depending on η, K, P , such that for all $t \leq c$ and $k \in [P]$, $|u_k^t(a) - \hat{u}_k^t(a)| \leq Ct^{2^{k-1}+1}$.

Proof [Proof of Proposition 33] Denote $\hat{\mathbf{u}}^t = (\hat{u}_k^t)_{k \in [P]}$. Notice that

$$\frac{d}{dt} \hat{u}_k^t(a) = a \alpha_{\{1, \dots, k\}} m_k \prod_{j < k} \hat{u}_j^t(a).$$

Denote $\Delta_k^t = \sup_{s \in [0, t]} \sup_{a \in [-1, +1]} |u_k^t(a) - \hat{u}_k^t(a)|$. By Grönwall's lemma, it is sufficient to show that $\frac{d}{dt} \Delta_k^t \leq K(t^{2k-1} + \Delta_k^t)$ for some constant $K > 0$. We will consider c sufficiently small to apply Lemma 32.

We recall the evolution equations:

$$\frac{d}{dt} u_k^t = a \mathbb{E}_{\mathbf{z}} [h_*(\mathbf{z}) z_k \sigma'(\langle \mathbf{u}^t, \mathbf{z} \rangle)] - a \mathbb{E}_{\mathbf{z}} [\hat{f}_{\text{NN}}(\mathbf{z}; \bar{\rho}_t) z_k \sigma'(\langle \mathbf{u}^t, \mathbf{z} \rangle)].$$

Let us first show that $|\mathbb{E}_{\mathbf{z}} [\hat{f}_{\text{NN}}(\mathbf{z}) z_k \sigma'(\langle \mathbf{u}^t, \mathbf{z} \rangle)]| \leq \Delta_k^t + O(t^L)$. Denote the Fourier coefficients $\hat{f}_{\text{NN}}(S; \bar{\rho}_t) = \mathbb{E}_{\mathbf{z}} [\chi_S(\mathbf{z}) \hat{f}_{\text{NN}}(\mathbf{z}; \bar{\rho}_t)]$. If $k \in S$, then

$$\begin{aligned} |\hat{f}_{\text{NN}}(S; \bar{\rho}_t)| &\leq \int |a \mathbb{E}_{\mathbf{z}} [\chi_S(\mathbf{z}) \sigma(\langle \mathbf{u}^t(a), \mathbf{z} \rangle)]| \mu_a(\mathrm{d}a) \\ &\leq \int |m_{|S|} \prod_{i \in S} u_i^t(a) \cdot (1 + O(t))| \mu_a(\mathrm{d}a) + O(t^L) \leq K \Delta_k^t + O(t^L), \end{aligned} \quad (64)$$

where we used Lemma 32. Furthermore, note that $|\hat{f}_{\text{NN}}(S; \bar{\rho}_t)| \leq K$ for any $S \subseteq [P]$. By expanding \hat{f}_{NN} in the Fourier basis, we get

$$\begin{aligned} |\mathbb{E}_{\mathbf{z}} [\hat{f}_{\text{NN}}(\mathbf{z}) z_k \sigma'(\langle \mathbf{u}^t, \mathbf{z} \rangle)]| &\leq \sum_{S \subseteq [P]} |\hat{f}_{\text{NN}}(S; \bar{\rho}_t) \mathbb{E}_{\mathbf{z}} [\chi_{S \oplus k}(\mathbf{z}) \sigma'(\langle \mathbf{u}^t, \mathbf{z} \rangle)]| \\ &\leq K \Delta_k^t + O(t^L) + \sum_{S \subseteq [P], k \notin S} K |\mathbb{E}_{\mathbf{z}} [\chi_{S \cup k}(\mathbf{z}) \sigma'(\langle \mathbf{u}^t, \mathbf{z} \rangle)]| \\ &\leq K \Delta_k^t + O(t^L), \end{aligned}$$

where we used Eq. (64) in the second line and Lemma 32 in the third line. We see therefore that

$$\left| \frac{d}{dt} (u_k^t - \hat{u}_k^t) \right| \leq \left| a \mathbb{E}_{\mathbf{z}} [h_*(\mathbf{z}) z_k \sigma'(\langle \mathbf{u}^t, \mathbf{z} \rangle)] - a \alpha_{\{1, \dots, k\}} m_k \prod_{j < k} \hat{u}_j^t(a) \right| + K \Delta_k^t + O(t^L).$$

We can separate the first term into three contributions:

$$\left| \mathbb{E}_{\mathbf{z}} [h_*(\mathbf{z}) z_k \sigma'(\langle \mathbf{u}^t, \mathbf{z} \rangle)] - \alpha_{\{1, \dots, k\}} m_k \prod_{j < k} \hat{u}_j^t(a) \right| \leq \text{(I)} + \text{(II)} + \text{(III)},$$

where

$$\begin{aligned}
 \text{(I)} &= \sum_{i < k} \left| \alpha_{\{1, \dots, i\}} \mathbb{E}_{\mathbf{z}} [\chi_{\{1, \dots, i\} \cup \{k\}}(\mathbf{z}) \sigma'(\langle \mathbf{u}^t, \mathbf{z} \rangle)] \right| \leq K \Delta_k^t + O(t^L), \\
 \text{(II)} &= \left| \alpha_{\{1, \dots, k\}} \mathbb{E}_{\mathbf{z}} [\chi_{\{1, \dots, k-1\}}(\mathbf{z}) \sigma'(\langle \mathbf{u}^t, \mathbf{z} \rangle)] - \alpha_{\{1, \dots, k\}} m_k \prod_{j < k} \hat{u}_j^t(a) \right| \\
 &\leq K \left| m_k \prod_{i \in [k-1]} u_i^t \cdot (1 + O(t)) - \prod_{i \in [k-1]} \hat{u}_i^t \right| + O(t^L) \\
 &\leq K \sum_{i \in [k-1]} \Delta_i^t \prod_{j \in [k-1], j \neq i} (\Delta_j^t + |\hat{u}_j^t|) + O(t^L), \\
 \text{(III)} &= \sum_{i > k} \left| \alpha_{\{1, \dots, i\}} \mathbb{E}_{\mathbf{z}} [\chi_{\{1, \dots, i\} \setminus \{k\}}(\mathbf{z}) \sigma'(\langle \mathbf{u}^t, \mathbf{z} \rangle)] \right| \\
 &\leq K \prod_{j \in [k-1]} (\Delta_j^t + |\hat{u}_j^t|) \sum_{i > k+1} |u_i^t| + O(t^L) \leq Kt \prod_{i \in [k-1]} (\Delta_i^t + |\hat{u}_i^t|) + O(t^L),
 \end{aligned}$$

where we used in the last line that $\|\mathbf{u}^t\|_1 \leq Ct$ from Claim 31. In particular, notice that for any $i < k$, $\frac{d}{dt} \Delta_k^t \leq \Delta_i^t$. We can therefore prove recursively that $|\Delta_k^t| \leq O(t^{2^{k-1}+1})$ by noting that 1) $\Delta_1^t \leq Kt^2$; 2) $|\hat{u}_k^t| = \Theta(t^{2^{k-1}})$ and $\prod_{j < k} |\hat{u}_j^t| \leq Kt^{2^{k-1}-1}$; and 3) $t^L = O(t^{2^{k-1}})$ for any $k \in [P]$, and do not contribute to the leading terms. \blacksquare

Appendix E. Generic MSP functions are strongly $O(d)$ -SGD-learnable: Proof of Theorem 9 (discrete-time regime)

In this appendix, we prove Theorem 9, which states that generic functions with MSP structure are strongly SGD-learnable in the $O(d)$ -scaling. While the proof for vanilla staircases in Appendix D is done in the continuous-time regime, we use here the discrete-time regime as defined in Appendix C, with $O(1)$ -steps of size $\eta = \Theta(1)$. Furthermore, we will consider the activation function to be a degree- L polynomial, with L sufficiently large. In Appendix F, we provide a more general proof of this result for smooth (non-polynomial) activations (see Theorem 56) and using the continuous-time regime, with one technical caveat: the activation function needs to be perturbed at some point during training (the result holds almost surely over this perturbation, see Appendix F.2 for a discussion on this technical caveat).

Recall the definition of an MSP set structure.

Definition 34 *We say that $\mathcal{S} = \{S_1, \dots, S_m\}$ is a Merged-Staircase Property (MSP) set structure on the variables z_1, \dots, z_P if the sets are (without loss of generality) ordered so that for each $i \in [m]$, $|S_i \setminus \cup_{i' < i} S_{i'}| \leq 1$.*

Given an MSP set structure $\mathcal{S} \subseteq 2^{[P]}$ and a function $h_* : \{+1, -1\}^P \rightarrow \mathbb{R}$, we say that h_* has MSP structure \mathcal{S} if h_* can be written as

$$h_*(\mathbf{z}) = \sum_{S \in \mathcal{S}} \alpha_S \chi_S(\mathbf{z}),$$

where $\alpha_S \in \mathbb{R} \setminus \{0\}$ for all $S \in \mathcal{S}$. In other words, h_* has MSP structure \mathcal{S} if its nonzero Fourier coefficients are \mathcal{S} .

Ideally, we would like prove that for any MSP set structure \mathcal{S} , then any function h_* with nonzero Fourier coefficients \mathcal{S} is strongly $O(d)$ -SGD-learnable. However, there are degenerate examples of functions such as $h_*(\mathbf{z}) = z_1 + z_2 + z_1 z_3 + z_2 z_4$ which satisfy MSP structure but are not strongly $O(d)$ -SGD-learnable (see Section A). Therefore, it is not possible to prove a result that holds for every MSP function. The existence of degenerate functions satisfying MSP also adds difficulty to the problem of showing that specific functions satisfying MSP are learnable.

Nevertheless, in this section we are able to show that for any MSP set structure \mathcal{S} there are very few degenerate functions h_* . In fact, almost all functions with MSP structure \mathcal{S} are non-degenerate and are strongly $O(d)$ -SGD-learnable.

More precisely, for any set structure $\mathcal{S} \subseteq 2^{[P]}$, define the following measure over functions:

Definition 35 (Definition 8 restated) *For any set structure $\mathcal{S} \subseteq 2^{[P]}$ define the measure $\mu_{\mathcal{S}}$ over functions $h_* : \{+1, -1\}^P \rightarrow \mathbb{R}$ induced by taking $h_*(\mathbf{z}) = \sum_{S \subseteq [P]} \alpha_S \chi_S(\mathbf{x})$, where the Fourier coefficients satisfy $\alpha_S = 0$ if $S \notin \mathcal{S}$, and $(\alpha_S)_{S \in \mathcal{S}}$ have Lebesgue measure on $\mathbb{R}^{|\mathcal{S}|}$.*

For any MSP structure \mathcal{S} , we prove that h_* is almost surely strongly $O(d)$ -SGD-learnable with respect to $\mu_{\mathcal{S}}$:

Theorem 36 (Theorem 9 restated) *For any MSP set structure $\mathcal{S} \subseteq 2^{[P]}$, h_* is strongly $O(d)$ -SGD-learnable almost surely with respect to $\mu_{\mathcal{S}}$, using activation function $\sigma(x) = (1 + x)^L$ where $L = 2^{8P}$.*

Remark 37 We note that although $\sigma(x) = (1+x)^L$ does not satisfy Assumption A0, we can instead use an activation function such that $\sigma(x) = (1+x)^L$ in the interval $(-1, 1)$, and $\sigma(x)$ is smoothly thresholded outside this interval. In the proof, we control the growth of the first-layer weights and the input of the activation remains $|x| \leq 1$, so such a thresholding does not impact training.

We also prove the following variation on the theorem, which shows that we can take activation function that is a polynomial of degree $L \geq 2^{8P}$ with random coefficients. This proves that almost surely any polynomial activation will work, so it does not hold just for activation $(1+x)^L$:

Theorem 38 For any MSP structure $\mathcal{S} \subseteq 2^{[P]}$, and any $L \geq 2^{8P}$, if we draw $\mathbf{m} \sim \text{Unif}[-1, 1]^{\otimes L+1}$, then h_* is strongly $O(d)$ -SGD-learnable almost surely with respect to $\mu_{\mathcal{S}}$, using activation function $\sigma(x) = \sum_{i=0}^L m_i x^i$.

E.1. Outline of the proof

Choice of hyperparameters We train in the discrete-time regime with $\Theta(1)$ steps of size $\Theta(1)$ and $\Theta(d)$ batch size. Recall from (d-DF-PDE) (Theorem 24) that it is sufficient to show for any $\varepsilon > 0$, there exist hyperparameters satisfying A0-A2, D3 such that (d-DF-PDE) reaches ε -risk. We consider the following hyperparameters.

- We do not regularize. I.e., $\lambda^w = 0$, and $\lambda^a = \lambda^w = 0$, same as Section D.
- We initialize the first layer to deterministically 0 weights, and the second layer to uniform random weights. I.e., we take $\mu_a = \text{Unif}([+1, -1])$ and $\mu_W = \delta_0$. This is the same as in the vanilla staircase proof of Section D. For the dimension-free dynamics, this corresponds to taking $(\bar{a}^0, \bar{\mathbf{u}}^0, \bar{s}^0) \sim \bar{\rho}_0$ with $\bar{a}^0 \sim \text{Unif}([+1, -1])$, $\bar{\mathbf{u}}^0 = \mathbf{0}$ and $\bar{s}^0 = 0$. In particular, $\bar{s}^k = 0$ during the whole dynamics, which lets us ignore it and allows for a simpler analysis.
- Our learning rate schedule has two phases, with learning rate given by parameter $\eta > 0$:

Phase 1: For k_1 steps we train the first layer weights $\bar{\mathbf{u}}^k$ while keeping the second layer weights fixed $\bar{a}^k = \bar{a}^0$. We set $\eta_k^a = 0$ and $\eta_k^w = \eta$ for $k \in \{0, \dots, k_1 - 1\}$.

Phase 2: For k_2 steps we train the second layer weights \bar{a}^k while keeping the first layer weights fixed at $\bar{\mathbf{u}}^k = \bar{\mathbf{u}}^{k_1}$. We set $\eta_k^a = \eta$ and $\eta_k^w = 0$ for $k \in \{k_1, \dots, k_2 - 1\}$.

We also take $\eta > 0$ to be a small enough constant, and $b = \Omega(d)$ for a large enough constant depending on P, ε, η . For the first phase, we will train for $k_1 = P$ time steps, since this turns out to be sufficient to prove learnability. For the second phase, we train for $k_2 = \Theta(1)$ time steps, where k_2 is a constant depending on η, ε , and P , to be determined later. We prove that (d-DF-PDE) with such hyperparameters will reach ε -risk, which, by the equivalence stated Theorem 24, implies the strong SGD-learnability in $O(d)$ -scaling.

Assumption on the activation We will assume that on the interval $(-1, 1)$ our activation is given by a polynomial of degree at most L . I.e., for all $x \in (-1, 1)$, we have $\sigma(x) = \sum_{i=0}^L \frac{m_i}{i!} x^i$ for $\mathbf{m} = [m_0, \dots, m_L] = [\sigma(0), \sigma^{(1)}(0), \dots, \sigma^{(L)}(0)]$.

E.1.1. PHASE 2 (LINEAR TRAINING)

Let us first present the analysis of Phase 2. We train the second layer and keep the first layer weights fixed. This is kernel gradient descent with kernel $K^{k_1} : \{+1, -1\}^P \times \{+1, -1\}^P \rightarrow \mathbb{R}$ given by

$$K^{k_1}(\mathbf{z}, \mathbf{z}') = \mathbb{E}_{a \sim \mu_a} [\sigma(\langle \mathbf{u}^{k_1}(a), \mathbf{z} \rangle) \sigma(\langle \mathbf{u}^{k_1}(a), \mathbf{z}' \rangle)].$$

So the residual $g_k(\mathbf{z}) = h_*(\mathbf{z}) - \hat{f}_{\text{NN}}(\mathbf{z}; \bar{\rho}_k)$, evolves, for any $k \in \{k_1, \dots, k_2 - 1\}$, as:

$$g_{k+1}(\mathbf{z}) = g_k(\mathbf{z}) - \eta \mathbb{E}_{\mathbf{z}'} [K^{k_1}(\mathbf{z}, \mathbf{z}') g_k(\mathbf{z}')].$$

The evolution of the risk is given by:

$$\begin{aligned} R(\bar{\rho}_{k+1}) &= \frac{1}{2} \mathbb{E}[g_{k+1}(\mathbf{z})^2] \\ &= R(\bar{\rho}_k) - \eta \mathbb{E}_{\mathbf{z}, \mathbf{z}'} [g_k(\mathbf{z}) K^{k_1}(\mathbf{z}, \mathbf{z}') g_k(\mathbf{z}')] \\ &\quad + \frac{\eta^2}{2} \mathbb{E}_{\mathbf{z}, \mathbf{z}', \mathbf{z}''} [K^{k_1}(\mathbf{z}, \mathbf{z}') K^{k_1}(\mathbf{z}, \mathbf{z}'') g_k(\mathbf{z}') g_k(\mathbf{z}'')] \\ &\leq R(\bar{\rho}_k) - \eta \left(1 - \frac{\eta \lambda_{\max}(\mathbf{K}^{k_1})}{2^{P+1}}\right) \mathbb{E}_{\mathbf{z}, \mathbf{z}'} [g_k(\mathbf{z}) K^{k_1}(\mathbf{z}, \mathbf{z}') g_k(\mathbf{z}')], \end{aligned}$$

where $\mathbf{K}^{k_1} = (K^{k_1}(\mathbf{z}, \mathbf{z}'))_{\mathbf{z}, \mathbf{z}'}$ is the $2^P \times 2^P$ kernel matrix. Note that $\lambda_{\max}(\mathbf{K}^{k_1}) \leq \|\mathbf{K}^{k_1}\|_F \leq 2^P \|\sigma\|_\infty^2 \leq 2^P K^2$. So if we take any learning rate $\eta \leq 1/K^2$, we have

$$\begin{aligned} R(\bar{\rho}_{k+1}) &\leq R(\bar{\rho}_k) - \frac{\eta}{2} \mathbb{E}_{\mathbf{z}, \mathbf{z}'} [g_k(\mathbf{z}) K^{k_1}(\mathbf{z}, \mathbf{z}') g_k(\mathbf{z}')] \\ &\leq \left(1 - \frac{\eta \lambda_{\min}(\mathbf{K}^{k_1})}{2^P}\right) R(\bar{\rho}_k). \end{aligned}$$

Finally, note that $R(\bar{\rho}_{k_1}) \leq \frac{1}{2} (\|\sigma\|_\infty + \|h_*\|_\infty)^2 \leq 2K^2$, so if we take any $k_2 \geq k_1 + \log(\varepsilon/2K^2) / \log(1 - \eta \lambda_{\min}(\mathbf{K}^{k_1})/2^P)$, we ensure that $R(\bar{\rho}_{k_2}) \leq \varepsilon$. It remains only to show that $\lambda_{\min}(\mathbf{K}^{k_1}) \geq c$, for a constant $c > 0$ depending only on η, h_*, P, k_1 , and m .

E.1.2. PHASE 1 (NONLINEAR TRAINING)

Now let us show how to analyze Phase 1, and in particular how to prove that $\lambda_{\min}(\mathbf{K}^{k_1})$ is bounded away from 0.

Writing the weight evolution with a polynomial First, we show that if we train for a constant number k_1 of steps, then we can write the weights obtained by the dimension-free dynamics as a constant-degree polynomial in the second-layer weights. This is because the activation is a polynomial in the interval $(-1, 1)$, and the weights of the first layer do not grow enough to leave this interval.

Lemma 39 (Training dynamics are given by a polynomial) *Let $\xi = (\xi_{S,k})_{S \subseteq [P], 0 \leq k \leq k_1 - 1} \in \mathbb{R}^{2^P k_1}$, $\zeta \in \mathbb{R}$, and $\rho \in \mathbb{R}^{L+1}$ be variables.*

For each $i \in [P]$ define $p_{0,i}(\zeta, \xi, \rho) \equiv 0$. For each $0 \leq k \leq k_1 - 1$, define $p_{k+1,i}(\zeta, \xi, \rho)$ with the recurrence relation:

$$p_{k+1,i}(\zeta, \xi, \rho) = p_{k,i}(\zeta, \xi, \rho) + \zeta \rho_1 \xi_{\{i\},k} \\ + \zeta \sum_{r=1}^{L-1} \frac{\rho_{r+1}}{r!} \sum_{(i_1, \dots, i_r) \in [P]^r} \xi_{\{i\} \oplus \{i_1\} \oplus \dots \oplus \{i_r\}, k} \prod_{l=1}^r p_{k,i_l}(\zeta, \xi, \rho)$$

There is a constant $c > 0$ depending only on k_1, P, K , such that for any $0 < \eta < c$,

$$\bar{u}_i^k(a) = p_{k,i}(\eta a, \beta, \mathbf{m})$$

where $\beta = (\beta_{S,k})_{S \subseteq [P], 0 \leq k \leq k_1 - 1}$ has values given by, for all $S \subseteq [P]$,

$$\beta_{S,k} = \mathbb{E}[(-\hat{f}_{\text{NN}}(\mathbf{z}; \bar{\rho}_k) + h_*(\mathbf{z}))\chi_S(\mathbf{z})].$$

Because of the term $\hat{f}_{\text{NN}}(\cdot; \bar{\rho}_k)$, which evolves nonlinearly, this is nontrivial to directly analyze. However, if the step size η is taken small enough, then the interaction term $\hat{f}_{\text{NN}}(\cdot; \bar{\rho}_k)$ is small, of order $O(\eta k)$, and we show that it can be ignored. Formally, we define the simplified dynamics $\hat{u}^k(a)$ for each $0 \leq k \leq k_1$ by letting $\hat{u}^0(a) = \mathbf{0}$ and inductively setting for each $k \in \{0, \dots, k_1 - 1\}$,

$$\hat{u}^{k+1}(a) = \hat{u}^k(a) - \eta a_j^0 \mathbb{E}_{\mathbf{x}}[-h_*(\mathbf{z})\sigma'(\langle \hat{u}^k(a), \mathbf{z} \rangle)\mathbf{z}].$$

This differs from the definition of the dynamics for \bar{u}^k in that we have dropped the $\hat{f}_{\text{NN}}(\mathbf{z}; \rho_k)$ term in the update equation. By a similar argument, we may show:

Lemma 40 (Simplified training dynamics are given by a polynomial) *There is a constant $c > 0$ depending only on k_1, P, K , such that for any $0 < \eta < c$, any $i \in [P]$ and any $0 \leq k \leq k_1$, we have*

$$\hat{u}_i^k(a) = p_{k,i}(\eta a, \alpha, \mathbf{m}),$$

where we abuse notation (since $\alpha = (\alpha_S)_{S \subseteq [P]}$ otherwise) and let $\alpha = (\alpha_{S,k})_{S \subseteq [P], 0 \leq k \leq k_1 - 1}$ be given by

$$\alpha_{S,k} = \alpha_S = \mathbb{E}[h_*(\mathbf{z})\chi_S(\mathbf{z})]$$

We now show that the simplified dynamics \hat{u}^k is a good enough approximation to \bar{u}^k , and it suffices to analyze \hat{u}^k .

Reducing to analyzing simplified dynamics We lower-bound $\lambda_{\min}(\mathbf{K}^{k_1})$ in terms of the determinant of a certain random matrix. Let $\zeta = [\zeta_1, \dots, \zeta_{2^P}]$ be a vector of 2^P variables. Define $M = M(\zeta, \xi, \rho) \in \mathbb{R}^{2^P \times 2^P}$ to be the matrix indexed by $\mathbf{z} \in \{+1, -1\}^P$ and $j \in [2^P]$ with entries

$$M_{\mathbf{z},j}(\zeta, \xi, \rho) = \sum_{r=0}^L \frac{\rho_r}{r!} \left(\sum_{i=1}^P z_i p_{k,i}(\zeta_j, \xi, \rho) \right)^r. \quad (65)$$

This matrix is motivated by the following fact:

Lemma 41 *There is a constant $c > 0$ depending only on k_1, P, K , such that for any $0 < \eta < c$, and any $\mathbf{a} = [a_1, \dots, a_{2P}] \in [-1, 1]^{2P}$, we have*

$$\begin{aligned} M_{\mathbf{z},j}(\eta\mathbf{a}, \boldsymbol{\beta}, \mathbf{m}) &= \sigma(\langle \hat{\mathbf{u}}^{k_1}(a_j), \mathbf{z} \rangle) \\ M_{\mathbf{z},j}(\eta\mathbf{a}, \boldsymbol{\alpha}, \mathbf{m}) &= \sigma(\langle \hat{\mathbf{u}}^{k_1}(a_j), \mathbf{z} \rangle) \end{aligned}$$

Using this we can show:

Lemma 42 *There is a constant $c > 0$ depending on K, P such that for any $0 < \eta < c$,*

$$\lambda_{\min}(\mathbf{K}^{k_1}) \geq c \mathbb{E}_{\mathbf{a} \sim \mu_{\mathbf{a}}^{\otimes 2P}} [\det(\mathbf{M}(\eta\mathbf{a}, \boldsymbol{\beta}, \mathbf{m}))^2].$$

On the other hand, we can prove a lower-bound on $\mathbb{E}[\det(\mathbf{M}(\eta\mathbf{a}, \boldsymbol{\beta}, \mathbf{m}))^2]$ simply by lower-bounding the sum of magnitudes of coefficients of $\det(\mathbf{M}(\boldsymbol{\zeta}, \boldsymbol{\alpha}, \mathbf{m}))$ when viewed as a polynomial in $\boldsymbol{\zeta}$. This is because of (a) the fact that $\det(\mathbf{M}(\boldsymbol{\zeta}, \boldsymbol{\alpha}, \mathbf{m}))$ and $\det(\mathbf{M}(\boldsymbol{\zeta}, \boldsymbol{\beta}, \mathbf{m}))$ have coefficients in $\boldsymbol{\zeta}$ that are $O(\eta)$ -close for η small, and (b) the fact that polynomials anti-concentrate over random inputs:

Lemma 43 *There is $D > 0$ depending only on P, k_1, L , and there are $C, c > 0$ depending only on P, k_1, K, L such that if we write*

$$\det(\mathbf{M}(\boldsymbol{\zeta}, \boldsymbol{\alpha}, \mathbf{m})) = \sum_{\boldsymbol{\gamma} \in \{0, \dots, D\}^{2P}} h_{\boldsymbol{\gamma}} \boldsymbol{\zeta}^{\boldsymbol{\gamma}},$$

then

$$\mathbb{E}_{\mathbf{a} \sim \mu_{\mathbf{a}}^{\otimes 2P}} [\det(\mathbf{M}(\eta\mathbf{a}, \boldsymbol{\beta}, \mathbf{m}))^2] \geq c \sum_{\boldsymbol{\gamma} \in \{0, \dots, D\}^{2P}} \eta^{2\|\boldsymbol{\gamma}\|_1} \max(0, |h_{\boldsymbol{\gamma}}| - C\eta)^2.$$

Combining the above lemmas, it holds that if $\det(\mathbf{M}(\boldsymbol{\zeta}, \boldsymbol{\alpha}, \mathbf{m}))$ is a nonzero polynomial in $\boldsymbol{\zeta}$, then h_* is strongly- $O(d)$ learnable:

Lemma 44 *Suppose that $\det(\mathbf{M}(\boldsymbol{\zeta}, \boldsymbol{\alpha}, \mathbf{m})) \not\equiv 0$ as a polynomial in $\boldsymbol{\zeta}$. Then the function $h_*(\mathbf{z}) = \sum_{S \subseteq [P]} \alpha_S \chi_S(\mathbf{z})$ is strongly $O(d)$ -SGD-learnable with any activation function σ that is equal to $\sigma(x) = \sum_{i=0}^L \frac{m_i}{i!} x^i$ on the interval $x \in (-1, 1)$.*

Proof Let k_1 be a constant depending on P , and let $C, c > 0$ be constants depending on k_1, P, K, L such that Lemmas 42 and 43 hold. Then taking any learning rate

$$0 < \eta < \min\left(c, \max_{\boldsymbol{\gamma} \in \{0, \dots, D\}^{2P}} |h_{\boldsymbol{\gamma}}| / (2C)\right),$$

we have

$$\lambda_{\min}(\mathbf{K}^{k_1}) \geq c^2 \eta^{2^P D} |h_{\boldsymbol{\gamma}}|^2 / 4 > 0, \tag{66}$$

which is a nonnegative constant that does not depend on d . So by the analysis of Phase 2 in Section E.1.1, we can set k_2 to be a large enough constant that $R(\bar{\rho}_{k_2}) \leq \varepsilon$. By Theorem 24 (which gives the equivalence between (d-DF-PDE) and strong $O(d)$ -SGD-learnability in the discrete-time setting), this implies strong $O(d)$ -SGD-learnability. \blacksquare

Analyzing simplified dynamics By the above arguments, the problem has been reduced to proving that $\det(\mathbf{M}(\zeta, \alpha, \mathbf{m})) \neq 0$ as a polynomial in ζ . In other words, by Lemma 41, this means that it suffices to analyze the simplified dynamics $\hat{\mathbf{u}}^k$.

We wish to prove that $\det(\mathbf{M}(\zeta, \alpha, \mathbf{m})) \neq 0$ almost surely over the choice of α . Since we take $h_*(\mathbf{z}) = \sum_{S \in \mathcal{S}} \alpha_S \chi_S(\mathbf{z})$ to be a generic function satisfying MSP, we could hope that it would be sufficient to prove that $\det(\mathbf{M}(\zeta, \xi, \mathbf{m})) \neq 0$ as a polynomial over ζ and ξ . However, there is an important technical subtlety. Although $\alpha = (\alpha_S)_{S \in \mathcal{S}}$ can be chosen to be generic, the vector $(\alpha_{S,k})_{S \subseteq [P], k \in \{0, \dots, k_1-1\}}$ has the constraints that $\alpha_{S,k} = \alpha_S$ for all S, k , and that $\alpha_S = 0$ for all $S \notin \mathcal{S}$. To take this into account, let $\phi = (\phi_S)_{S \in \mathcal{S}}$ be a vector of variables and define the following matrix $\mathbf{N}(\zeta, \phi, \rho) \in \mathbb{R}^{2^P \times 2^P}$, indexed by $\mathbf{z} \in \{+1, -1\}^P$ and $j \in [2^P]$:

$$N_{\mathbf{z},j}(\zeta, \phi, \rho) = M_{\mathbf{z},j}(\zeta, \xi, \rho) \Big|_{\xi_{S,k}=0 \text{ for all } S \notin \mathcal{S}, \text{ and } \xi_{S,k}=\phi_S \text{ for all } S \in \mathcal{S}}. \quad (67)$$

The matrix \mathbf{N} differs from \mathbf{M} only in that we have changed the variables from $(\xi_{S,k})_{S,k}$ to variables $(\phi_S)_{S \in \mathcal{S}}$, effectively incorporating the constraints on α . This is helpful, because suppose that we can prove that

$$\det(\mathbf{N}(\zeta, \phi, \mathbf{m})) \neq 0 \text{ as a polynomial in } \zeta \text{ and } \phi. \quad (68)$$

Then almost surely over the Lebesgue measure on $(\alpha_S)_{S \in \mathcal{S}}$, we have that $\det(\mathbf{N}(\zeta, \alpha, \mathbf{m})) \neq 0$ as a polynomial over ζ . And indeed, $\det(\mathbf{N}(\zeta, \alpha, \mathbf{m})) \equiv \det(\mathbf{M}(\zeta, \alpha, \mathbf{m})) \neq 0$, which is what we wanted to show. So it suffices to prove (68).

We prove (68) by analyzing the recurrence relations for $p_{k,i}$ to show that to first-order the polynomials $p_{k_1,i}$ are distinct for all $i \in [P]$, and then leveraging the algebraic result of Newman and Slater (1979) that large powers of distinct polynomials are linearly independent. We show:

Lemma 45 *Suppose that $L \geq 2^{8P}$ and let $m_i = i! \binom{L}{i}$ for all $0 \leq i \leq L$, corresponding to activation function $\sigma(x) = (1+x)^L$. Also let $k_1 = P$. Then $\det(\mathbf{N}(\zeta, \phi, \mathbf{m})) \neq 0$ (i.e., (68) holds).*

This also yields the immediate corollary:

Corollary 46 *Suppose that $L \geq 2^{8P}$, and let $\mathbf{m} \sim \text{Unif}[-1, 1]^{\otimes L+1}$, corresponding to a random polynomial activation function. Then $\det(\mathbf{N}(\zeta, \phi, \mathbf{m})) \neq 0$ (i.e., (68) holds) almost surely over \mathbf{m} .*

Proof Lemma 45 implies that $\det(\mathbf{N}(\zeta, \phi, \rho))$ is a nonzero polynomial in ζ, ϕ, ρ . Since we choose $\mathbf{m} \sim \text{Unif}[-1, 1]^{\otimes L+1}$, this means that $\det(\mathbf{N}(\zeta, \phi, \mathbf{m})) \neq 0$ almost surely over the choice of \mathbf{m} . ■

This allows us to prove Theorems 9 and 38.

Proof [Proof of Theorem 9] Taking $m_i = i! \binom{L}{i}$ corresponds to activation function $\sigma(x) = (1+x)^L$. By Lemma 45, we have $\det(\mathbf{N}(\zeta, \alpha, \mathbf{m})) \neq 0$ almost surely over α with respect to the Lebesgue measure. So by Lemma 44, $h_*(\mathbf{z}) = \sum_{S \in \mathcal{S}} \alpha_S \chi_S(\mathbf{z})$ is strongly $O(d)$ -SGD-learnable with activation $\sigma(x) = (1+x)^L$, almost surely over h_* with respect to μ_S . ■

Proof [Proof of Theorem 38] The argument is the same, except using Corollary 46. ■

E.2. Proof of Lemmas 39, 40, and 41

We show that if the learning rate η is small then for $0 \leq k \leq k_1$ the weights of $\bar{\mathbf{u}}^k$ and $\hat{\mathbf{u}}^k$ remain small enough that the activation σ only ever has inputs in the range $(-1, 1)$, meaning that we can treat the activation σ as exactly given by the polynomial $\sum_{i=0}^L \frac{m_i}{i!} x^i$.

Claim 47 *For any time step $0 \leq k \leq k_1$ any $j \in [N]$, and any learning rate $\eta < 1/(4K^2Pk)$, and any $a \in [-1, 1]$ we have*

$$\|\bar{\mathbf{u}}^k(a)\|_1, \|\hat{\mathbf{u}}^k(a)\|_1 \leq 2\eta K^2 P k \leq 1/2.$$

Proof The proof is by induction on k . The base case is clear since $\bar{\mathbf{u}}^0 = \hat{\mathbf{u}}^0 = \mathbf{0}$. For the inductive step, $\hat{f}_{\text{NN}}(\mathbf{z}; \bar{\rho}_k) \leq \mathbb{E}_{a \sim \mu_a} [a |\sigma(\langle \bar{\mathbf{u}}^k, \mathbf{z} \rangle)|] \leq \|\sigma\|_\infty \leq K$, since $a \sim \text{Unif}[-1, 1]$. Therefore

$$\begin{aligned} \|\bar{\mathbf{u}}^{k+1}(a)\|_1 &\leq \|\bar{\mathbf{u}}^k(a)\|_1 + \eta \|\mathbb{E}_{\mathbf{x}}[(\hat{f}_{\text{NN}}(\mathbf{z}; \bar{\rho}_k) + h_*(\mathbf{x})) a \sigma'(\langle \bar{\mathbf{u}}^k, \mathbf{z} \rangle) \mathbf{z}]\|_1 \\ &\leq \|\bar{\mathbf{u}}^k(a)\|_1 + 2\eta K^2 P \leq 2\eta K^2 P k. \end{aligned}$$

The bound for $\|\hat{\mathbf{u}}^k(a)\|_1$ is similar. ■

This allows us to prove Lemmas 39 and 40.

Proof [Proof of Lemmas 39 and 40] Let $\zeta \in \mathbb{R}$, $\boldsymbol{\xi} = (\xi_{S,k})_{S \subseteq [P], k \in \{0, \dots, k_1-1\}}$, and $\boldsymbol{\rho} \in \mathbb{R}^{L+1}$ be variables. Define $s_0, \dots, s_{k_1-1} : \{+1, -1\}^P \rightarrow \mathbb{R}$ to be $s_k(\mathbf{z}) = \sum_{S \subseteq [P]} \xi_{S,k} \chi_S(\mathbf{z})$. Consider the recurrence relation $\boldsymbol{\nu}^k \in \mathbb{R}^P$, where we initialize $\boldsymbol{\nu}^0 = \mathbf{0}$ and, for $0 \leq k \leq k_1 - 1$,

$$\boldsymbol{\nu}^{k+1} = \boldsymbol{\nu}^k + \zeta \mathbb{E}_{\mathbf{z}} \left[s_k(\mathbf{z}) \sum_{r=0}^{L-1} \frac{\rho_{r+1}}{r!} \langle \boldsymbol{\nu}^k, \mathbf{z} \rangle^r \mathbf{z} \right]. \quad (69)$$

Substituting in $\zeta = \eta a$ and $\boldsymbol{\rho} = \mathbf{m}$, this recurrence relation is satisfied by $\bar{\mathbf{u}}^k(a)$ with $s_k(\mathbf{z}) = -\hat{f}_{\text{NN}}(\mathbf{z}; \bar{\rho}_k) + h_*(\mathbf{z}) = \sum_S \beta_{S,k} \chi_S(\mathbf{z})$ and by $\hat{\mathbf{u}}^k(a)$ with $s_k(\mathbf{z}) = h_*(\mathbf{z}) = \sum_S \alpha_S \chi_S(\mathbf{z})$. This is because $|\langle \bar{\mathbf{u}}^k, \mathbf{z} \rangle|, |\langle \hat{\mathbf{u}}^k, \mathbf{z} \rangle| \leq 1/2 < 1$ by Claim 47 and in the interval $(-1, 1)$ $\sigma(x) = \sum_{r=0}^L \frac{m_r}{r!} x^r$.

It remains to show that

$$\nu_i^k = p_{k,i}(\zeta, \boldsymbol{\xi}, \boldsymbol{\rho}).$$

The proof is by induction on k . For $k = 0$, it is true that $p_{0,i}(\zeta, \boldsymbol{\xi}) = 0 = \nu_i^0$. For the inductive step, notice that for any $r \geq 1$ and $i \in [d]$, we can write

$$\begin{aligned} \mathbb{E}_{\mathbf{z}} [s_k(\mathbf{z}) \langle \boldsymbol{\nu}^k, \mathbf{z} \rangle^r z_i] &= \mathbb{E}_{\mathbf{z}} \left[s_k(\mathbf{z}) z_i \sum_{(i_1, \dots, i_r) \in [P]^r} \prod_{l=1}^r \nu_{i_l}^k z_{i_l} \right] \\ &= \sum_{(i_1, \dots, i_r) \in [P]^r} \mathbb{E}_{\mathbf{z}} \left[s_k(\mathbf{z}) \chi_i(\mathbf{z}) \prod_{l=1}^r \chi_{i_l}(\mathbf{z}) \right] \prod_{l=1}^r p_{k, i_l}(\zeta, \boldsymbol{\xi}, \boldsymbol{\rho}) \\ &= \sum_{(i_1, \dots, i_r) \in [P]^r} \xi_{\{i\} \oplus \{i_1\} \oplus \dots \oplus \{i_r\}, k} \prod_{l=1}^r p_{k, i_l}(\zeta, \boldsymbol{\xi}, \boldsymbol{\rho}), \end{aligned}$$

and $\mathbb{E}_{\mathbf{z}} [s_k(\mathbf{z}) \langle \boldsymbol{\nu}^k, \mathbf{z} \rangle^0 z_i] = \mathbb{E}_{\mathbf{z}} [s_k(\mathbf{z}) z_i] = \xi_{\{i\}, k}$. The inductive step follows by linearity of expectation. ■

Finally, we prove Lemma 41:

Proof [Proof of Lemma 41] This is immediate from Lemmas 39 and 40, using the fact from Claim 47 that $\|\bar{\mathbf{u}}^{k_1}(a)\|_1, \|\hat{\mathbf{u}}^{k_1}(a)\|_1 \leq 1/2$, so $\langle \bar{\mathbf{u}}^{k_1}(a), \mathbf{z} \rangle, \langle \hat{\mathbf{u}}^{k_1}(a), \mathbf{z} \rangle \in (-1/2, 1/2)$, and in this interval $\sigma(x) = \sum_{i=0}^L \frac{m_i}{i!} x^i$. ■

E.3. Proof of Lemma 42

Proof [Proof of Lemma 42] For short-hand write $\mathbf{B}(\mathbf{a}) = \mathcal{M}(\eta\mathbf{a}, \beta, \mathbf{m})$. By Lemma 41, $B_{z,j}(\mathbf{a}) = \sigma(\langle \bar{\mathbf{u}}^{k_1}(a_j), \mathbf{z} \rangle)$, so

$$\begin{aligned} \mathbf{K}_{z,z'}^{k_1} &= \mathbb{E}_{\mathbf{a} \sim \mu_{\mathbf{a}}} [\sigma(\langle \bar{\mathbf{u}}^{k_1}(a), \mathbf{z} \rangle) \sigma(\langle \bar{\mathbf{u}}^{k_1}(a), \mathbf{z}' \rangle)] \\ &= \mathbb{E}_{\mathbf{a} \sim \mu_{\mathbf{a}}^{\otimes 2P}} \left[\frac{1}{2^P} \sum_{j=1}^{2^P} \sigma(\langle \bar{\mathbf{u}}^{k_1}(a_j), \mathbf{z} \rangle) \sigma(\langle \bar{\mathbf{u}}^{k_1}(a_j), \mathbf{z}' \rangle) \right] \\ &= \frac{1}{2^P} \mathbb{E}_{\mathbf{a} \sim \mu_{\mathbf{a}}^{\otimes 2P}} [\mathbf{B}(\mathbf{a}) \mathbf{B}(\mathbf{a})^\top]. \end{aligned}$$

So $\lambda_{\min}(\mathbf{K}^{k_1}) \geq \frac{1}{2^P} \mathbb{E}_{\mathbf{a} \sim \mu_{\mathbf{a}}^{\otimes 2P}} [\lambda_{\min}(\mathbf{B}(\mathbf{a}) \mathbf{B}(\mathbf{a})^\top)]$, and

$$\lambda_{\min}(\mathbf{B}(\mathbf{a}) \mathbf{B}(\mathbf{a})^\top) \geq \det(\mathbf{B}(\mathbf{a}))^2 / (\lambda_{\max}(\mathbf{B}(\mathbf{a})))^{2^{P+1}-2} \geq c \det(\mathbf{B}(\mathbf{a}))^2,$$

for $c = 1/(2^P K)^{2^{P+1}-2} > 0$. ■

E.4. Proof of Lemma 43

Let us first show that $\beta_{S,k} = \mathbb{E}[(-\hat{f}_{\text{NN}}(\mathbf{z}; \bar{\rho}_k) + h_*(\mathbf{z})) \chi_S(\mathbf{z})]$ is close to $\alpha_{S,k} = \alpha_S = \mathbb{E}[h_*(\mathbf{z}) \chi_S(\mathbf{z})]$.

Claim 48 *There are constants $C, c > 0$ depending on k_1, P, K such that for any $0 < \eta < c$, any $S \subseteq [P]$, and any $k \in \{0, \dots, k_1 - 1\}$,*

$$|\beta_{S,k} - \alpha_{S,k}| \leq C\eta$$

Proof It suffices to show that $\|\hat{f}_{\text{NN}}(\cdot; \bar{\rho}_k)\|_\infty \leq C\eta$. This is true since Claim 47 implies $\|\hat{\mathbf{u}}^k(a)\|_1 \leq C\eta$, so $|\hat{f}_{\text{NN}}(\mathbf{z}; \bar{\rho}_k)| \leq \mathbb{E}_{\mathbf{a} \sim \mu_{\mathbf{a}}} [a\sigma(0) + |a| \|\sigma'\|_\infty |\langle \bar{\mathbf{u}}^k(a), \mathbf{z} \rangle|] \leq K \|\bar{\mathbf{u}}^k(a)\|_1 \leq KC\eta \leq C\eta$. ■

We now show the lemma.

Proof [Proof of Lemma 43]

Write $\det(\mathcal{M}(\zeta, \alpha, \mathbf{m})) = \sum_{\gamma \in \{0, \dots, D\}^{2^P}} \hat{h}_\gamma \zeta^\gamma$. Let us prove that there is a constant C depending on k_1, P, K, L such that $|h_\gamma - \bar{h}_\gamma| \leq C\eta$ for all γ . To see this, notice that $\det(\mathcal{M}(\zeta, \xi, \rho))$ is a polynomial in ζ, ξ, ρ , whose degree and coefficients depend only on k_1, P, L (this is because each entry of $\mathcal{M}(\zeta, \xi, \rho)$ is a polynomial in ζ, ξ, ρ with coefficients depending on k_1, P, L , and it is a $2^P \times 2^P$ matrix). Since $\|\mathbf{m}\|_\infty \leq K$ and $\|\alpha\|_\infty, \|\beta\|_\infty \leq 2K$, and $\|\alpha - \beta\|_\infty \leq C\eta$ by Claim 48, we conclude that there is a constant C depending on k_1, P, K, L such that $|h_\gamma - \bar{h}_\gamma| \leq C\eta$ for all γ .

By anti-concentration of polynomials (i.e., Lemma 104), we have that there exists a constant $c > 0$ depending on k_1, L, P such that

$$\begin{aligned} \mathbb{E}_{\mathbf{a} \sim \mu_{\mathbf{a}}^{\otimes 2P}} [\det(\mathbf{M}(\eta \mathbf{a}, \boldsymbol{\alpha}, \mathbf{m}))^2] &\geq c \sum_{\gamma \in \{0, \dots, D\}^{2P}} \eta^{2\|\gamma\|_1} |\bar{h}_\gamma|^2 \\ &\geq c \sum_{\gamma \in \{0, \dots, D\}^{2P}} \eta^{2\|\gamma\|_1} \max(0, |h_\gamma| - C\eta)^2, \end{aligned}$$

concluding the lemma. ■

E.5. Proof of Lemma 45

For this section, fix $\mathbf{m} \in \mathbb{R}^{L+1}$ to be $m_i = i! \binom{L}{i}$ for all $i \in \{0, \dots, L\}$. This corresponds to the activation function $\sigma(x) = (1+x)^L$.

E.5.1. REDUCING TO MINIMAL MSP SET STRUCTURES

To show that $\det(\mathbf{N}(\zeta, \phi, \mathbf{m})) \neq 0$, we first show that it suffices to consider “minimal” MSP set structures.

Claim 49 *Let $S' \subseteq S$ be such that S' is an MSP set structure. Then if*

$$\det(\mathbf{N}(\zeta, \phi, \mathbf{m})) \big|_{\phi_S=0} \text{ for all } S \in \mathcal{S} \setminus S' \neq 0,$$

we have

$$\det(\mathbf{N}(\zeta, \phi, \mathbf{m})) \neq 0.$$

Proof Substituting 0 for ϕ_S for all $S \in \mathcal{S} \setminus S'$. ■

Therefore it suffices to prove the lemma for minimal MSP structures. Without loss of generality (up to permutation of the variables), we assume that we can write

$$S' = \{S_1, \dots, S_P\},$$

where, for all $i \in [P]$,

$$i \in S_i \text{ and } S_i \subseteq [i].$$

Otherwise, we could remove a set from \mathcal{S} and still have a MSP set structure.

E.5.2. COMPUTING THE WEIGHTS TO LEADING ORDER

Let us define the polynomials $q_{k,i}$ in variables ζ, ϕ, ρ . For all $k \in \{0, \dots, k_1 - 1\}$ and $i \in [P]$,

$$q_{k,i}(\zeta, \phi, \rho) = p_{k,i}(\zeta, \boldsymbol{\xi}, \rho) \big|_{\xi_{S,k}=0} \text{ for all } S \notin S' \text{ and } \xi_{S,k} = \phi_S \text{ for all } S \in S'.$$

Therefore $\mathbf{N}(\zeta, \phi, \rho)$ has entries $N_{z,j}(\zeta, \phi, \rho) = \sum_{r=0}^L \frac{\rho_r}{r!} \left(\sum_{i=1}^P q_{k,i}(\zeta_j, \phi, \rho) \right)^r$. Let us explicitly compute the nonzero term of $q_{k,i}$ that is of lowest-degree in ζ . First, we show that many terms are zero.

Claim 50 *Recursively define $o_i = 1 + \sum_{i' \in S_i \setminus \{i\}} o_{i'}$ for all $i \in [P]$.⁸ Then $q_{k,i}(\zeta, \phi, \mathbf{m})$ has no nonzero terms of degree less than o_i in ζ .*

Proof The proof is by induction on k . In the base case of $k = 0$ it is true since $q_{0,i} \equiv 0$. In the inductive step, we assume it is true for all $k' \in \{0, \dots, k\}$ and we prove the claim for $k + 1$. By the recurrence dynamics,

$$\begin{aligned} & q_{k+1,i}(\zeta, \phi, \mathbf{m}) \\ &= q_{k,i}(\zeta, \phi, \mathbf{m}) + \zeta m_1 \phi_{\{i\}} \mathbf{1}(\{i\} \in \mathcal{S}) \\ &+ \zeta \sum_{r=1}^{L-1} \frac{m_{r+1}}{r!} \sum_{(i_1, \dots, i_r) \in [P]^r} \phi_{\{i\} \oplus \{i_1\} \oplus \dots \oplus \{i_r\}} \mathbf{1}(\{i\} \oplus \{i_1\} \oplus \dots \oplus \{i_r\} \in \mathcal{S}) \prod_{l=1}^r q_{k,i_l}(\zeta, \phi). \end{aligned}$$

The first term, $q_{k,i}(\zeta, \phi, \mathbf{m})$, is handled by the inductive hypothesis. The second term is nonzero only in the case that $\{i\} \in \mathcal{S}$, in which case $S_i = \{i\} \notin \mathcal{S}'$ and $o_i = 1$, so we do not have a contradiction. The last terms can be handled by the inductive hypothesis: for any (i_1, \dots, i_r) , each q_{k,i_l} has no terms of degree less than o_{i_l} in ζ . So $\zeta \prod_l q_{k,i_l}(\zeta, \alpha)$ has no terms of degree less than $1 + \sum_{l=1}^r o_{i_l}$ in ζ . We break into cases. *Case a.* If $\{i\} \oplus \{i_1\} \oplus \dots \oplus \{i_r\} = S_i$, then $S_i \setminus \{i\} \subset \{i_1, \dots, i_r\}$, so $1 + \sum_{l=1}^r o_{i_l} \geq o_i$, and so no new terms of degree less than o_i are added. *Case b.* If $\{i\} \oplus \{i_1\} \oplus \dots \oplus \{i_r\} = S_{i'}$ for some $i' \neq i$, then either $i \in \{i_1, \dots, i_r\}$, in which case $1 + \sum_{l=1}^r o_{i_l} > o_i$. Otherwise, we must have $i' > i$. But in this case $o_{i'} > o_i$ since $i \in S_{i'}$, so we also have $\sum_{l=1}^r o_{i_l} > o_i$ and again no new terms of degree less than o_i are added. In fact, only terms of degree strictly more than o_i are added. \blacksquare

Finally, we give a recurrence for the degree- o_i term in ζ of $q_{k,i}(\zeta, \phi, \mathbf{m})$. Because of the previous claim, when this term is nonzero, it is the smallest-degree nonzero term. Denote this term by $\tilde{q}_{k,i}(\phi, \mathbf{m}) = [\zeta^{o_i}] q_{k,i}(\zeta, \phi, \mathbf{m})$.

Claim 51 *If $k_1 \geq P$, then $\tilde{q}_{k_1,i}(\phi, \mathbf{m})$ is a nonzero monomial in the variables ϕ . Furthermore, for any $i \neq i' \in [P]$, the monomials $\tilde{q}_{k_1,i}(\phi, \mathbf{m})$ and $\tilde{q}_{k_1,i'}(\phi, \mathbf{m})$ are not constant multiples of each other.*

Proof Following the analysis of the previous claim used to prove that $[\zeta^l] q_{k,i}(\zeta, \phi, \mathbf{m}) = 0$ for all $l < o_i$, only certain terms contribute in the recurrence. So we can simplify it to:

$$\tilde{q}_{k+1,i}(\phi, \mathbf{m}) = \tilde{q}_{k,i}(\phi, \mathbf{m}) + m_{|S_i|} \phi_{S_i} \prod_{i' \in S_i \setminus \{i\}} \tilde{q}_{k,i'}(\phi, \mathbf{m}).$$

Define $s_i = 1$ for all i such that $|S_i| = 1$. And recursively define $s_i = 1 + \max\{s_{i'} : i' \in S_i \setminus \{i\}\}$ for all other $i \in [P]$. Inductively on k , for all $k < s_i$ we have $\tilde{q}_{k,i} \equiv 0$. This is clear from the base case $\tilde{q}_{0,i} \equiv 0$ and the recurrence.

Next, for all $k \geq s_i$ we prove that

$$\tilde{q}_{k,i}(\phi, \mathbf{m}) = \gamma_{k,i} \phi_{S_i} \prod_{i' \in S_i \setminus \{i\}} \tilde{q}_{s_{i'},i'}(\phi, \mathbf{m}) \neq 0 \quad (70)$$

8. The sum over an empty set is 0 by convention.

for some nonzero constant $\gamma_{k,i} > 0$ that depends on \mathbf{m} . This is proved inductively on k . For $k = s_i > 1$, we have

$$\tilde{q}_{k,i}(\phi, \mathbf{m}) = m_{|S_i|} \phi_{S_i} \prod_{i' \in S_i \setminus \{i\}} \tilde{q}_{k-1,i'}(\phi, \mathbf{m}) = m_{|S_i|} \phi_{S_i} \prod_{i' \in S_i \setminus \{i\}} \gamma_{k-1,i'} \tilde{q}_{s_{i'},i'}(\phi, \mathbf{m}),$$

so it is true since $m_{|S_i|} > 0$. For the inductive step, if $k > s_i$,

$$\tilde{q}_{k,i}(\phi, \mathbf{m}) = \left(\gamma_{k-1,i} + m_{|S_i|} \prod_{i' \in S_i \setminus \{i\}} \gamma_{k-1,i'} \right) \phi_{S_i} \prod_{i' \in S_i \setminus \{i\}} \tilde{q}_{s_{i'},i'}.$$

So $\gamma_{k,i} = \gamma_{k-1,i} + m_{|S_i|} \prod_{i' \in S_i \setminus \{i\}} \gamma_{k-1,i'} > 0$ since $m_{|S_i|} > 0$ by nonnegativity. This concludes the induction for (70).

Using this recurrence relation (70) for $\tilde{q}_{k,i}$, by induction on k we conclude that for any $k \geq P > s_i$ we have that $\tilde{q}_{k,i}(\phi, \mathbf{m})$ is a nonzero monomial. Also, $\tilde{q}_{k,i}(\phi, \mathbf{m})$ and $\tilde{q}_{k,i'}(\phi, \mathbf{m})$ are distinct for all $i \neq i'$, since if $s_i \geq s_{i'}$ then ϕ_{S_i} divides $\tilde{q}_{k,i}$, but it does not divide $\tilde{q}_{k,i'}$. ■

Recall that the interpretation of $q_{k,i}$ with respect to the simplified dynamics: for any second-layer weight $a \in [-1, 1]$, the first-layer weights after training the simplified dynamics are $\hat{\mathbf{u}}_i^{k_1}(a) = q_{k,i}(\eta a, \alpha, \mathbf{m})$. What we have shown in the previous two claims is that for any $i \neq i'$ to leading order $\hat{\mathbf{u}}_i^{k_1}(a)$ and $\hat{\mathbf{u}}_{i'}^{k_1}$ have different dependence on the Fourier coefficients α of the target function h_* . Now we use this to essentially show that $\langle \hat{\mathbf{u}}_i^{k_1}(a), \mathbf{z} \rangle$ and $\langle \hat{\mathbf{u}}_{i'}^{k_1}(a), \mathbf{z}' \rangle$ are distinct for all $\mathbf{z} \neq \mathbf{z}'$.

Claim 52 *Define*

$$r_{\mathbf{z}}(\zeta, \phi, \mathbf{m}) = \sum_i z_i q_{k_1,i}(\zeta, \phi, \mathbf{m}).$$

Then, for each distinct pair $\mathbf{z}, \mathbf{z}' \in \{+1, -1\}^P$, we have $r_{\mathbf{z}}(\zeta, \phi, \mathbf{m}) - r_{\mathbf{z}'}(\zeta, \phi, \mathbf{m}) \neq 0$ as a polynomial in ζ and ϕ .

Proof Recall the definition of o_i from Claim 50. Let $i \in [P]$ be such that $z_i \neq z'_i$ and o_i is minimized. By Claim 50,

$$\begin{aligned} [\zeta^{o_i}](r_{\mathbf{z}}(\zeta, \phi, \mathbf{m}) - r_{\mathbf{z}'}(\zeta, \phi, \mathbf{m})) &= \sum_{i' \text{ s.t. } o_{i'} = o_i, z_{i'} \neq z'_{i'}} (z_{i'} - z'_{i'}) [\zeta^{o_i}] q_{k_1,i'}(\zeta, \phi, \mathbf{m}) \\ &= \sum_{i' \text{ s.t. } o_{i'} = o_i, z_{i'} \neq z'_{i'}} (z_{i'} - z'_{i'}) \tilde{q}_{k_1,i'}(\phi, \mathbf{m}), \end{aligned}$$

but $\tilde{q}_{k_1,i'}$ are distinct nonzero monomials in ϕ by Claim 51. So $r_{\mathbf{z}}(\zeta, \phi, \mathbf{m}) - r_{\mathbf{z}'}(\zeta, \phi, \mathbf{m}) \neq 0$. ■

E.5.3. APPLYING LINEAR INDEPENDENCE OF POWERS OF POLYNOMIALS

We conclude the proof of the lemma by using the following result of Newman and Slater (1979) showing that large powers of distinct polynomials are linearly independent.

Proposition 53 (Remark 5.2 in Newman and Slater (1979)) *Let $R_1, \dots, R_m \in \mathbb{C}[\zeta]$ be non-constant polynomials such that for all $i \neq i' \in [m]$ we have $R_i(\zeta)$ is not a constant multiple of $R_{i'}(\zeta)$. Then for $L \geq 8m^2$ we have that $(R_1)^L, \dots, (R_m)^L \in \mathbb{C}[\zeta]$ are \mathbb{C} -linearly independent.*

We are ready to prove that $\det(\mathbf{N}(\zeta, \phi, \mathbf{m})) \neq 0$.

Proof [Proof of Lemma 45]

Let us fix $\alpha = (\alpha_S)_{S \in \mathcal{S}}$ such that for all $z \neq z'$ we have $r_z(\zeta, \alpha, \mathbf{m}) - r_{z'}(\zeta, \alpha, \mathbf{m}) \neq 0$ as polynomials in ζ . This can be ensured by drawing $\alpha_S \sim \text{Unif}[-1, +1]$ for all $S \in \mathcal{S}$, since for all $z \neq z'$ we have $r_z(\zeta, \phi, \mathbf{m}) - r_{z'}(\zeta, \phi, \mathbf{m}) \neq 0$ as polynomials in ζ, ϕ by Claim 52. Let us write $\tilde{r}_z(\zeta) = r_z(\zeta, \alpha, \mathbf{m})$ to emphasize that we have fixed the variables $\phi = \alpha$ and $\rho = \mathbf{m}$, and that we are looking at a polynomial over ζ .

Since we have chosen $m_i = i! \binom{L}{i}$ for all $i \in \{0, \dots, L\}$, we have

$$N_{z,j}(\zeta, \alpha, \mathbf{m}) = (1 + \tilde{r}_z(\zeta))^L.$$

From the recurrence relations ζ divides $q_{k_1,i}(\zeta, \phi, \mathbf{m})$ for each $i \in [P]$, so ζ divides $\tilde{r}_z(\zeta) = \sum_{i=1}^P z_i q_{k_1,i}(\zeta, \alpha, \mathbf{m})$. Therefore, no two polynomials $(1 + \tilde{r}_z(\zeta)), (1 + \tilde{r}_{z'}(\zeta))$ are constant multiples of each other for each distinct z, z' . Otherwise, if $(1 + \tilde{r}_z(\zeta)) \equiv c(1 + \tilde{r}_{z'}(\zeta))$, then we would have $1 = (1 + \tilde{r}_z(0)) = c(1 + \tilde{r}_{z'}(0)) = c$, which would imply $c = 1$, but $(1 + \tilde{r}_z(\zeta)) \neq (1 + \tilde{r}_{z'}(\zeta))$ since $\tilde{r}_z(\zeta)$ and $\tilde{r}_{z'}(\zeta)$ are distinct.

Construct the Wronskian matrix over the L th power polynomials $\{(1 + \tilde{r}_z(\zeta))^L\}_{z \in \{+1, -1\}^P}$. This is a $2^P \times 2^P$ matrix $\mathbf{H}(\zeta)$ whose entries are indexed by z and $j \in [2^P]$ and defined by:

$$H_{z,j}(\zeta) = \frac{\partial^{j-1}}{\partial \zeta^{j-1}} (1 + r_z(\zeta))^L.$$

By Proposition 53, the polynomials $\{(1 + \tilde{r}_z(\zeta))^L\}_{z \in \{+1, -1\}^P}$ are linearly-independent, so the Wronskian determinant is nonzero as a polynomial in ζ :

$$\det(\mathbf{H}(\zeta)) \neq 0.$$

Finally notice that we can write $\det(\mathbf{H}(\zeta)) = \frac{\partial}{\partial \zeta_2} \frac{\partial}{\partial \zeta_3^2} \dots \frac{\partial^{2^P-1}}{\partial \zeta_{2^P}} \det(\mathbf{N}(\zeta, \alpha, \mathbf{m})) \big|_{\zeta = \zeta_1 = \dots = \zeta_{2^P}}$.

Therefore $\det(\mathbf{N}(\zeta, \alpha, \mathbf{m})) \neq 0$ as a polynomial in ζ . So $\det(\mathbf{N}(\zeta, \phi, \mathbf{m})) \neq 0$ as a polynomial in ζ and ϕ . ■

Appendix F. Generic MSP functions are strongly $O(d)$ -SGD-learnable with continuous-time dynamics and activation perturbation

In this appendix, we provide a more general approach to proving strong $O(d)$ -SGD-learnability for generic MSP functions that goes beyond polynomial activation functions. The reason to include this second approach is two-fold:

1. We consider the continuous-time regime (as opposed to the discrete-time regime as in Appendix E), which is closer to practice, with small batch and step sizes. (Note that the extension to non-polynomial activations would also hold in discrete time.)
2. For continuous time and non-polynomial activations, the first layer weights $\bar{\mathbf{u}}^t$ are not polynomials in \bar{a}^0 anymore. However, we show that they can still be approximated by polynomials and that global convergence reduces to showing that certain (universal) polynomials are not identically 0.

Using this approach, we show in Theorem 56 that generic MSP functions are strongly $O(d)$ -SGD-learnable for smooth activation functions (as long as $\sigma^{(r)}(0) \neq 0$ for $r = 0, \dots, P$), with one technical caveat: we need to introduce a random perturbation to the activation function at one point during the training dynamics. While unnatural, this modification allows us to prove that the polynomials are non-zero for general MSP structure, using a ‘‘Vandermonde trick’’. See Section F.2 for a discussion on this technicality.

F.1. Statement of the result

Recall the definition of the measure over functions with MSP set structure \mathcal{S} :

Definition 54 (Definition 8 restated) For any set structure $\mathcal{S} \subseteq 2^{[P]}$ define the measure $\mu_{\mathcal{S}}$ over functions $h_* : \{+1, -1\}^P \rightarrow \mathbb{R}$ induced by taking $h_*(\mathbf{z}) = \sum_{S \subseteq [P]} \alpha_S \chi_S(\mathbf{x})$, where the Fourier coefficients satisfy $\alpha_S = 0$ if $S \notin \mathcal{S}$, and $(\alpha_S)_{S \in \mathcal{S}}$ have Lebesgue measure on $\mathbb{R}^{|\mathcal{S}|}$.

Choice of hyperparameters: Recall from the equivalence with (DF-PDE) (Theorem 6) that it is sufficient to show for any $\varepsilon > 0$, there exists hyperparameters satisfying A3' such that (DF-PDE) reaches ε -risk. We consider the following hyperparameters, which are the same as in the proof for the vanilla staircase in Section D:

- We do not regularize, i.e., $\lambda^a = \lambda^w = 0$, same as Section D.
- We initialize the first layer to deterministically $\bar{\mathbf{u}}^0 = \mathbf{0}$, and the second layer to uniform random weights on $[-1, -1]$, i.e., $\mu_a = \text{Unif}([+1, -1])$ and $\mu_W = \delta_0$.
- Our learning rate schedule is the same as in Section D,

Phase 1: We train the first layer weights $\bar{\mathbf{u}}^t$ while keeping the second layer weights fixed $\bar{a}^t = \bar{a}^0$. We set $\xi^a(t) = 0$ and $\xi^w(t) = 1$ for $t \in [0, T_1]$.

Phase 2: We train the second layer weights \bar{a}^t while keeping the first layer weights fixed at $\bar{\mathbf{u}}^t = \bar{\mathbf{u}}^{T_1}$. We set $\xi^a(t) = 1$ and $\xi^w(t) = 0$ for $t \in [T_1, T_2]$.

Remark 55 As in Section D, the learning rate schedules can be made Lipschitz at T_1 with a change of variables, falling under the assumptions of strong SGD learnability.

Perturbing the activation: We consider an activation function σ that verifies $A0'$, i.e., that is sufficiently smooth in a neighborhood of 0. However, unlike the proof for the vanilla staircase, we add the following technical caveat. At time T_1 , we randomly perturb the activation σ to get an activation σ_{pert} . We use activation σ in the training of Phase 1 when training (DF-PDE) during time $[0, T_1]$ but we use the perturbed activation function σ_{pert} when training (DF-PDE) during time $[T_1, T_2]$ in Phase 2. By perturbing the activation, we mean the following: let $0 < \tau_{pert} < 1$ be a parameter that controls the amount of perturbation. Draw $\rho_i \sim \text{Unif}([- \tau_{pert}, \tau_{pert}])$ for each $i \in \{0, \dots, 2^{8P}\}$. The perturbed activation is defined as $\sigma_{pert}(x) = \sigma(x) + \sum_{r=0}^{2^{8P}} \frac{\rho_r}{r!} x^r$.

The dynamics of (DF-PDE) in time $[0, T_1]$ with activation σ stitched together with the dynamics in time $[T_1, T_2]$ with activation σ_{pert} corresponds to an algorithm that falls under the definition of strong $O(d)$ -SGD-learnability, when extended to allow such a perturbation (in particular, the equivalent characterization and necessary condition in Theorems 6 and 7 would still hold). See Section F.2 for more discussion.

We restate the sufficient condition, proving that for any MSP set structure \mathcal{S} , generic functions h_* with that set structure \mathcal{S} are strongly $O(d)$ -SGD-learnable:

Theorem 56 *Consider $\mathcal{S} \subseteq 2^{[P]}$ a MSP set structure, and $0 < \tau_{pert} < 1$ a perturbation parameter. Assume that the activation function σ satisfies $A0'$ and has nonzero derivatives $\sigma^{(r)}(0) \neq 0$ for $r = 0, \dots, P$. Then, almost surely for h_* with respect to $\mu_{\mathcal{S}}$ and almost surely for perturbation $\rho \sim \text{Unif}([- \tau_{pert}, \tau_{pert}]^{2^{8P}})$, the following hold: for any $\varepsilon > 0$, there exist $T_1, T_2 > 0$ such that training with the above hyperparameters and activation perturbation will learn h_* to accuracy ε .*

This implies that almost surely over $\mu_{\mathcal{S}}$, h_ is strongly $O(d)$ -SGD-learnable (under the expanded definition of $O(d)$ -SGD-learnability where the SGD algorithm is allowed to perturb the activation function once).*

F.2. Discussion on the perturbation of the activation

The perturbation is convenient to show that a polynomial is not identically zero for arbitrary MSP set structure. Note that given a set structure \mathcal{S} , these polynomials are fully explicit (given by recurrence relations) and one can verify by hand that they have a non zero coefficient. It is an interesting direction to show this result directly without relying on perturbing the activation function. In the setting of discrete-time regime and polynomial activations (cf. Theorem 9), such a perturbation is not needed: the weights $\bar{\mathbf{u}}^{T_1}$ are exact polynomials of \bar{a}^0 and one can use algebraic tricks involving linear independence of powers of polynomials (see Proposition 53).

Note that we can extend the definition of strong SGD-learnability in $O(d)$ -scaling to allow such a perturbation. In that case, the dimension-free dynamics (DF-PDE) corresponds to gluing two dynamics with activations σ between $[0, T_1]$ and σ_{pert} between $[T_1, T_2]$. The equivalent characterization (Theorem 6) and necessary condition (Theorem 7) still hold using this extended definition.

F.3. Outline of the proof

Similarly to the proof for the vanilla staircase in Section D, the proof of Theorem 56 follows by analyzing the solution $\bar{\mathbf{u}}^t(\bar{a}^0)$ to the evolution equations (9) obtained from initialization $(\bar{a}^0, \mathbf{0}, 0)$. Again, for clarity, we will suppress some notations: we denote \mathbf{u} instead of $\bar{\mathbf{u}}$, and a instead of \bar{a}^0 . We also forget about $\bar{\mathbf{s}}^t = 0$ and simply consider $\bar{\rho}_t \in \mathcal{P}(\mathbb{R}^{P+1})$ the distribution of (a, \mathbf{u}) . This last

simplification can be done since we initialize the first-layer weights to 0, so in particular $\bar{s}^0 = 0$, and by the evolution equation of (9) we have $\bar{s}^t = 0$ throughout training. Furthermore, we will denote K a generic constant that only depends on P and the constants in the assumptions. The value of K can change from line to line.

For MSP functions beyond the vanilla staircase, the approach used to prove theorem 10 no longer works, and a finer-grained analysis is needed.⁹ The argument is more involved because we need to track higher-order corrections to \mathbf{u}^t . We present here the finer-grained analysis.

The proof analyzes Phase 1 and Phase 2 of training separately.

Phase 1 (nonlinear dynamics) We break our analysis of the nonlinear training in Phase 1 into several parts. The goal is to understand the evolution under the dimension-free PDE of each neuron's weights $(a, \mathbf{u}^t(a))$. Because we initialize the first layer to 0, it suffices to study the dynamics of \mathbf{u}^t , ignoring the dynamics of \bar{s}^t since it stays at $\bar{s}^t = \bar{s}^0 = 0$ throughout. The dynamics of \mathbf{u}^t are given by

$$\frac{d}{dt}\mathbf{u}^t = a\mathbb{E}_{\mathbf{z}}[g_t(\mathbf{z})\sigma'(\langle \mathbf{u}^t, \mathbf{z} \rangle)\mathbf{z}], \quad (71)$$

where $g_t(\mathbf{z}) = h_*(\mathbf{z}) - \hat{f}_{\text{NN}}(\mathbf{z}; \bar{\rho}_t)$ is the residual at time t .

Reducing to analyzing with polynomial approximation. Our first step is to analyze a polynomial approximation of \mathbf{u}^t instead of analyzing \mathbf{u}^t directly. Let $L > 0$ be an integer governing the degree of approximation. We will choose L to be a large enough constant depending on P . We first prove in Section F.4 that for small times t we can approximate the dynamics of $\mathbf{u}^t(a)$ by an approximate dynamics $\tilde{\mathbf{u}}^t$ defined as

$$\tilde{\mathbf{u}}^t(a) = \mathbf{Q}^t[a, a^2, a^3, \dots, a^L]^\top,$$

where $[a, a^2, a^3, \dots, a^L]$ denotes the vector with the powers of a , and $\mathbf{Q}^t \in \mathbb{R}^{P \times L}$ is a time-dependent matrix with $\mathbf{Q}^0 = \mathbf{0}$ and which is updated according to a certain non-linear dynamics defined as follows (this corresponds essentially to truncating the dynamics of \mathbf{u}^t by only keeping the order- L approximation). Let $g_t(\mathbf{z}) = h_*(\mathbf{z}) - \hat{f}_{\text{NN}}(\mathbf{z}; \bar{\rho}_t)$ denote the residual at time t . For $l = 1$,

$$\frac{d}{dt}Q_{i1}^t = \mathbb{E}_{\mathbf{z}}[z_i g_t(\mathbf{z}) m_1], \quad (72)$$

and for $2 \leq l \leq L$,

$$\frac{d}{dt}Q_{il}^t = \mathbb{E}_{\mathbf{z}} \left[z_i g_t(\mathbf{z}) \sum_{1 \leq r \leq L-1} \frac{m_{r+1}}{r!} \sum_{i_1, \dots, i_r \in [P]} \sum_{\substack{l_1, \dots, l_r \in [L] \\ \sum_{r'=1}^r l_{r'} = l-1}} \prod_{r'=1}^r z_{i_{r'}} Q_{i_{r'} l_{r'}}^t \right]. \quad (73)$$

We prove in Claim 59 that we have $\|\mathbf{u}^t(a) - \tilde{\mathbf{u}}^t(a)\| \leq O(t^L)$ for small enough times t , so it suffices to study $\tilde{\mathbf{u}}^t$ instead of \mathbf{u}^t . Of course, the dynamics of $\tilde{\mathbf{u}}^t$ still present a challenge to analyze.

9. Indeed, for MSP functions that are not vanilla staircases, $\mathbf{M} = (\mathbb{E}_a[a^{\beta(S)+\beta(S')}])_{S, S' \subseteq [P]}$ (introduced in Section D.1) can have some sets $S \neq S'$ such that $\beta(S) = \beta(S')$, and \mathbf{M} is not a positive matrix anymore.

Reducing to analyzing the simplified dynamics. One significant challenge is that the residual g_t is time-dependent, as it depends on $\hat{f}_{\text{NN}}(\cdot; \bar{\rho}_t)$. This interaction term complicates the picture significantly. However, if we train for small time t , then $|\hat{f}_{\text{NN}}(\mathbf{z}; \bar{\rho}_t)| \leq O(t)$, and we can expect the contribution of this term to be negligible. To make this intuition precise, in Section F.5 we introduce a time-dependent matrix $\hat{\mathbf{Q}}^t \in \mathbb{R}^{P \times L}$ which is initialized at $\hat{\mathbf{Q}}^0 = \mathbf{0}$ and which has the same evolution equations (72) and (73) as \mathbf{Q}^t , except with g_t replaced by h_* . We obtain a “simplified dynamics” by letting $\hat{\mathbf{u}}^t(a) = \hat{\mathbf{Q}}^t[a, a^2, a^3, \dots, a^L]^\top$. This is easier to analyze since it neglects the interaction term.

However, unlike the comparison of \mathbf{u}^t to its polynomial approximation $\tilde{\mathbf{u}}^t$, where we could prove that $\|\mathbf{u}^t - \tilde{\mathbf{u}}^t\| \leq O(t^L)$, it is not the case that the simplified dynamics $\hat{\mathbf{u}}$ give such a good accuracy approximation to \mathbf{u} in L_2 norm. Indeed, we may unfortunately have $\|\mathbf{u}^t - \hat{\mathbf{u}}^t\| \geq \Omega(t^2)$, which is a bound that would be far too loose for our analysis of higher-order terms in the dynamics. To overcome this issue, we prove that $|Q_{il}^t - \hat{Q}_{il}^t| \leq O(t^{l+1})$ for each $i \in [P], l \in [L]$. We then use the fact that $\tilde{\mathbf{u}}^t$ and $\hat{\mathbf{u}}^t$ are both polynomials in a with coefficients \mathbf{Q} and $\hat{\mathbf{Q}}$, respectively, to reduce to analyzing the $\hat{\mathbf{u}}^t$ dynamics (see Section F.7 for details).

Analyzing the simplified dynamics with a recurrence relation. We analyze the $\hat{\mathbf{u}}^t$ dynamics by deriving recurrence relations for the coefficients \hat{Q}_{il}^t . In particular, we may express each coefficient \hat{Q}_{il}^t as a polynomial in a, t , and the nonzero Fourier coefficients $\{\alpha_S\}_{S \in \mathcal{S}}$ of h_* (see Section F.6). This allows us to prove that almost surely over the choice of h_* each coordinate \hat{u}_i^t has distinct dynamics: namely, $\hat{u}_i^t - \hat{u}_{i'}^t \neq 0$ for all $i \neq i' \in [P]$. This is where we must use the fact that the MSP function h_* is “generic”, i.e., the coefficients $\{\alpha_S\}_{S \in \mathcal{S}}$ are chosen randomly. (In fact, we prove and use the stronger result that for any $\mathbf{z} \neq \mathbf{z}' \in \{+1, -1\}^P$, we have $\langle \hat{\mathbf{u}}^t, \mathbf{z} - \mathbf{z}' \rangle \neq 0$, and this difference has nonzero low-degree terms.)

Phase 2 (linear dynamics) The linear dynamics are analyzed by showing a lower-bound on $\lambda_{\min}(\mathbf{K}^{T_1})$, as was the case for the proof of the vanilla staircase in Section D. We show in Section F.7 (and similarly to the discrete case) that it is sufficient to show that a polynomial depending on the simplified dynamics is non-zero. In Sections F.8 and F.9, we show how this can be achieved using the perturbation on the activation function: one of the coefficient of the polynomial can be rewritten as the determinant of a Vandermonde matrix with entries $\{\langle \hat{\mathbf{u}}^t, \mathbf{z} \rangle\}_{\mathbf{z} \in \{+1, -1\}^P}$. Using that $\langle \hat{\mathbf{u}}^t, \mathbf{z} - \mathbf{z}' \rangle \neq 0$ for $\mathbf{z} \neq \mathbf{z}'$, this determinant is non zero and we conclude the proof.

F.4. Approximating the u_i^t with polynomials

As outlined above, we study the dynamics of the dimension-free PDE. Let us first analyze Phase 1, when we train for time T_1 using activation function σ , and keep the second layer fixed. In particular, we analyze the dynamics of $\mathbf{u}^t(a)$ given by eq. (71) and the initialization $\mathbf{u}^t = \mathbf{0}$. In the proof below, we sometimes omit the dependence on a and time t , e.g., writing \mathbf{u} instead of $\mathbf{u}^t(a)$, when the dependence on t and a is clear.

The first step of the proof is to approximate $\mathbf{u}^t(a)$ with a polynomial in a . Let $L > 0$ be an integer which corresponds to the degree of approximation. We prove in this section that we have the approximation $\mathbf{u}^t \approx \tilde{\mathbf{u}}^t$, where we define $\tilde{\mathbf{u}}^t$ as:

$$\tilde{\mathbf{u}}^t = \mathbf{Q}^t[a, a^2, a^3, \dots, a^L]^\top.$$

Here, recall that $\mathbf{Q}^t \in \mathbb{R}^{P \times L}$ is given by initializing $\mathbf{Q}^0 = \mathbf{0}$ and training with eq. (72) and eq. (73).

We first prove for each $l \in [L], i \in [L]$, that each coefficient Q_{il}^t of a^l scales as $O(t^l)$.

Claim 57 *There is a constant C depending on K, L, P such that for any $i \in [P], l \in [L]$, and $0 \leq t \leq T_1$, $|Q_{il}^t| \leq Ct^l$.*

Proof We prove this by induction on l . For the base case of $l = 1$, we know that

$$\left| \frac{\partial Q_{i1}^t}{\partial t} \right| = |\mathbb{E}_{\mathbf{z}}[z_i g_t(\mathbf{z}) m_1]| \leq K,$$

since $\|g_t\|_\infty \leq K$ throughout the dynamics, and $|m_1| \leq K$. So $|Q_{i1}^t| \leq Kt \leq C_1 t$ for a constant C_1 . For the inductive step, let $2 \leq l \leq L$ and suppose $|Q_{i'l'}^t| \leq C_{l'} t^{l'}$ for all $1 \leq l' < l$. Then

$$\begin{aligned} \left| \frac{\partial Q_{il}^t}{\partial t} \right| &\leq K \|g_t\|_\infty \mathbb{E}_{\mathbf{z}} \left[\sum_{1 \leq r \leq L-1} \sum_{i_1, \dots, i_r \in [P]} \sum_{\substack{l_1, \dots, l_r \in [L] \\ \sum_{r'=1}^r l_{r'} = l-1}} \prod_{r'=1}^r |Q_{i_r l_{r'}}^t| \right] \\ &\leq K \|g_t\|_\infty \mathbb{E}_{\mathbf{z}} \left[\sum_{1 \leq r \leq L-1} \sum_{i_1, \dots, i_r \in [P]} \sum_{\substack{l_1, \dots, l_r \in [L] \\ \sum_{r'=1}^r l_{r'} = l-1}} t^{l-1} \prod_{r'=1}^r |C_{l_{r'}}| \right] \\ &\leq KL(2P)^L \max(|C_1|, \dots, |C_{L-1}|) L t^{l-1} \leq C_l t^{l-1}. \end{aligned}$$

So $|Q_{il}^t| \leq C_l t^l$, defining C_l appropriately. \blacksquare

Let us prove that \mathbf{u}^t and $\tilde{\mathbf{u}}^t$ have norm $O(t)$.

Claim 58 *There is a constant C depending only on K, P and a constant C' depending only on K, L, P such that for any $0 \leq t \leq T_1$, $\|\mathbf{u}^t\| \leq Ct$, $\|\tilde{\mathbf{u}}^t\| \leq C't$.*

Proof Note $\mathbf{u}^0 = \mathbf{0}$ and $\left\| \frac{\partial \mathbf{u}^t}{\partial t} \right\| \leq |a| \|g_t(\mathbf{z})\|_\infty \|\sigma'\|_\infty \|\mathbf{z}\| \leq (2K)K\sqrt{P} \leq C$. So $\|\mathbf{u}^t\| \leq Ct$. Similarly, $\|\tilde{\mathbf{u}}^t\| \leq \sum_{i \in [P], \sum_{l \in [L]} |Q_{il}^t| \leq PLCt \leq C't$ by Claim 57. \blacksquare

Let us prove that $\mathbf{u}^t \approx \tilde{\mathbf{u}}^t$ throughout the dynamics.

Claim 59 *There are constants $c, C > 0$ depending on K, L, P such that if $T_1 \leq c$ then for any $0 \leq t \leq T_1$, $\|\mathbf{u}^t - \tilde{\mathbf{u}}^t\| \leq Ct^L$.*

Proof The proof will use Gronwall's inequality. First, by triangle inequality

$$\frac{\partial}{\partial t} \|\mathbf{u}^t - \tilde{\mathbf{u}}^t\| \leq \left\| \frac{\partial \tilde{\mathbf{u}}^t}{\partial t} - \frac{\partial \mathbf{u}^t}{\partial t} \right\|.$$

Notice that

$$\begin{aligned} &\left\| \frac{\partial \tilde{\mathbf{u}}^t}{\partial t} - \frac{\partial \mathbf{u}^t}{\partial t} \right\| \\ &= \left\| \mathbb{E}[z g_t(\mathbf{z}) m_1] + \sum_{l=1}^L a^l \mathbb{E}_{\mathbf{z}} \left[z g_t(\mathbf{z}) \sum_{1 \leq r \leq L-1} \frac{m_{r+1}}{r!} \sum_{i_1, \dots, i_r \in [P]} \sum_{\substack{l_1, \dots, l_r \in [L] \\ \sum_{r'=1}^r l_{r'} = l-1}} \prod_{r'=1}^r z_{i_{r'}} Q_{i_{r'} l_{r'}}^t \right] - \frac{\partial \mathbf{u}^t}{\partial t} \right\| \\ &\leq \left\| \mathbb{E}[z g_t(\mathbf{z}) m_1] + \mathbb{E}_{\mathbf{z}} \left[z g_t(\mathbf{z}) \sum_{1 \leq r \leq L-1} \frac{m_{r+1}}{r!} h_{L,r}^t(\mathbf{z}) \right] - \frac{\partial \mathbf{u}^t}{\partial t} \right\|, \end{aligned}$$

where for any $1 \leq r \leq L - 1$,

$$h_{L,r}^t(\mathbf{z}) = \sum_{l=1}^L a^l \sum_{i_1, \dots, i_r \in [P]} \sum_{\substack{l_1, \dots, l_r \in [L] \\ \sum_{r'=1}^r l_{r'} = l-1}} \prod_{r'=1}^r z_{i_{r'}} Q_{i_{r'}, l_{r'}}^t,$$

which can be thought of as a degree- $(L - 1)$ approximation to $\langle \tilde{\mathbf{u}}^t, \mathbf{z} \rangle^r$, in the sense that

$$\begin{aligned} & |\langle \tilde{\mathbf{u}}^t, \mathbf{z} \rangle^r - h_{L,r}^t(\mathbf{z})| \\ &= \left| \langle \tilde{\mathbf{u}}^t, \mathbf{z} \rangle^r - \sum_{i_1, \dots, i_r \in [P]} \sum_{\substack{l_1, \dots, l_r \in [L] \\ 0 \leq \sum_{r'=1}^r l_{r'} \leq L-1}} \prod_{r'=1}^r z_{i_{r'}} Q_{i_{r'}, l_{r'}}^t a^{l_{r'}} \right| \\ &= \left| \sum_{i_1, \dots, i_r \in [P]} \left\{ \sum_{l_1, \dots, l_r \in [L]} \prod_{r'=1}^r z_{i_{r'}} Q_{i_{r'}, l_{r'}}^t a^{l_{r'}} - \sum_{\substack{l_1, \dots, l_r \in [L] \\ 0 \leq \sum_{r'=1}^r l_{r'} \leq L-1}} \prod_{r'=1}^r z_{i_{r'}} Q_{i_{r'}, l_{r'}}^t a^{l_{r'}} \right\} \right| \\ &= \left| \sum_{i_1, \dots, i_r \in [P]} \sum_{\substack{l_1, \dots, l_r \in [L] \\ L \leq \sum_{r'=1}^r l_{r'} \leq rL}} \prod_{r'=1}^r z_{i_{r'}} Q_{i_{r'}, l_{r'}}^t a^{l_{r'}} \right| \\ &\leq \sum_{i_1, \dots, i_r \in [P]} \sum_{\substack{l_1, \dots, l_r \in [L] \\ L \leq \sum_{r'=1}^r l_{r'} \leq rL}} \prod_{r'=1}^r C t^{l_{r'}} |a|^{l_{r'}} \\ &\leq P^r 2^{rL} C^r |a|^{L^2} \\ &\leq P^L 2^{L^2} C^L t^L \\ &\leq C t^L, \end{aligned}$$

for a constant C depending on K, L, P , where used Claim 57 to bound $Q_{i_{r'}, l_{r'}}^t$ and that $|a| \leq 1$ and $t < 1$ in the final bound.

We conclude that

$$\begin{aligned} \left\| \frac{\partial \tilde{\mathbf{u}}^t}{\partial t} - \frac{\partial \mathbf{u}^t}{\partial t} \right\| &\leq \left\| \mathbb{E}_{\mathbf{z}} \left[\mathbf{z} g_t(\mathbf{z}) \sum_{1 \leq r \leq L-1} \frac{m_{r+1}}{r!} (h_{L,r}^t(\mathbf{z}) - \langle \mathbf{u}^t, \mathbf{z} \rangle^r) \right] \right\| \\ &\leq 2K^2 PL \max_{1 \leq r \leq L-1, \mathbf{z} \in \{+1, -1\}^P} |h_{L,r}^t(\mathbf{z}) - \langle \mathbf{u}^t, \mathbf{z} \rangle^r| \\ &\leq 2K^2 PL \max_{1 \leq r \leq L-1, \mathbf{z} \in \{+1, -1\}^P} |h_{L,r}^t(\mathbf{z}) - \langle \tilde{\mathbf{u}}^t, \mathbf{z} \rangle^r| + |\langle \tilde{\mathbf{u}}^t, \mathbf{z} \rangle^r - \langle \mathbf{u}^t, \mathbf{z} \rangle^r| \\ &\leq 2K^2 PL (C t^L + r P^r \|\tilde{\mathbf{u}}^t - \mathbf{u}^t\|) \\ &\leq C t^L + C \|\tilde{\mathbf{u}}^t - \mathbf{u}^t\|, \end{aligned}$$

where for the second-to-last line we have used $\|\tilde{\mathbf{u}}^t\|, \|\mathbf{u}^t\| \leq C t \leq 1/(r\sqrt{P})$ if we take small enough time $T_1 \leq c$ for a constant $c > 0$ depending on L, P, K . The claim follows by Gronwall's inequality, since $\tilde{\mathbf{u}}^0 = \mathbf{u}^0 = \mathbf{0}$ and we train for time $T_1 \leq c < 1$. \blacksquare

F.5. Simplified dynamics without interaction term

We have introduced the dynamics $\tilde{\mathbf{u}}^t$ and proved that they give a $O(t^L)$ -approximation of the true dynamics \mathbf{u}^t . We now reduce further, to analyzing the dynamics of $\hat{\mathbf{u}}^t$, where we have dropped the \hat{f}_{NN} term, replacing $g_t(\mathbf{z})$ with $h_*(\mathbf{z})$ in the definition of the dynamics (72) and (73):

$$\hat{\mathbf{u}}(t) = \hat{\mathbf{Q}}^t [a, a^2, a^3, \dots, a^L]^\top,$$

where for $l = 1$,

$$\frac{\partial \hat{Q}_{i1}^t}{\partial t} = \mathbb{E}_{\mathbf{z}} [z_i h_*(\mathbf{z}) m_1],$$

and for $2 \leq l \leq L$,

$$\frac{\partial \hat{Q}_{il}^t}{\partial t} = \mathbb{E}_{\mathbf{z}} \left[z_i h_*(\mathbf{z}) \sum_{1 \leq r \leq L-1} \frac{m_{r+1}}{r!} \sum_{i_1, \dots, i_r \in [P]} \sum_{\substack{l_1, \dots, l_r \in [L] \\ \sum_{r'=1}^r l_{r'} = l-1}} \prod_{r'=1}^r z_{i_{r'}} \hat{Q}_{i_r' l_{r'}}^t \right].$$

To show that the new dynamics is close to the old dynamics, we first show that $\|\hat{f}_{\text{NN}}(\mathbf{z}; \rho_t)\|_\infty = O(t)$, is small when t is small:

Claim 60 *There is a constant C depending on K, L, P such that for all $0 \leq t \leq T_1$, $\|\hat{f}_{\text{NN}}(\mathbf{z}; \rho_t)\|_\infty \leq Ct$.*

Proof For any \mathbf{z} , $|\hat{f}_{\text{NN}}(\mathbf{z}; \rho_t) - \hat{f}_{\text{NN}}(\mathbf{z}; \rho_0)| \leq \mathbb{E}_a [|a\sigma(\langle \mathbf{u}^t(a), \mathbf{z} \rangle) - a\sigma(\langle \mathbf{u}^0(a), \mathbf{z} \rangle)|] \leq \mathbb{E}_a [K |\langle \mathbf{u}^t(a) - \mathbf{u}^0(a), \mathbf{z} \rangle|] \leq K \|\mathbf{u}^t(a) - \mathbf{0}\|_1 \leq KPCt \leq Ct$ by Claim 58 and K -Lipschitzness. And $\hat{f}_{\text{NN}}(\mathbf{z}; \rho_0) = \mathbb{E}_a [a\sigma(0)] = 0$, since $\mathbb{E}_a [a] = 0$ and $\mathbf{u}^0(a) = \mathbf{0}$. \blacksquare

We also prove the analogue of Claim 57 for $\hat{\mathbf{Q}}$:

Claim 61 *There is a constant C depending on K, L, P such that for all $i \in [P]$, $l \in [L]$, and $0 \leq t \leq T_1$, $|\hat{Q}_{il}^t| \leq Ct^l$. Also, $\|\hat{\mathbf{u}}^t\| \leq Ct$.*

Proof The bound on $|\hat{Q}_{il}^t|$ is the same as Claim 57, but using the bound $\|h_*\|_\infty \leq K$ instead of the bound $\|g_t\|_\infty \leq 2K$. The bound on $\|\hat{\mathbf{u}}^t\|$ is the same as Claim 58, using the bound on $|\hat{Q}_{il}^t|$. \blacksquare

We show that $|Q_{il}^t - \hat{Q}_{il}^t| \leq O(t^{l+1})$ for each $l \in [L]$:

Lemma 62 *There is a constant C depending on K, L, P such that for any $i \in [P]$, $l \in [L]$, $|Q_{il}^t - \hat{Q}_{il}^t| \leq Ct^{l+1}$.*

Proof We prove this by induction on l . For $l = 1$,

$$\left| \frac{\partial \hat{Q}_{i1}^t}{\partial t} - \frac{\partial Q_{i1}^t}{\partial t} \right| \leq \|h_*(\mathbf{z}) - g_t(\mathbf{z})\|_\infty |m_1| \leq K \|\hat{f}_{\text{NN}}(\mathbf{z}; \rho_t)\|_\infty \leq CKt \leq C_1 t,$$

by Claim 60, for some large enough constant C_1 . Therefore $|\hat{Q}_{i1}^t - Q_{i1}^t| \leq C_1 t^2$. For the inductive step, let $2 \leq l \leq L$, and assume that $|\hat{Q}_{il'}^t - Q_{il'}^t| \leq C_1 t^{l'+1}$ for all $1 \leq l' \leq l-1$. Then

$$\begin{aligned}
 & \left| \frac{\partial \hat{Q}_{il}^t}{\partial t} - \frac{\partial Q_{il}^t}{\partial t} \right| \\
 = & \left| \mathbb{E}_z \left[z_i \sum_{1 \leq r \leq L-1} \frac{m_{r+1}}{r!} \sum_{i_1, \dots, i_r \in [P]} \sum_{\substack{l_1, \dots, l_r \in [L] \\ \sum_{r'=1}^r l_{r'} = l-1}} \left\{ h_*(z) \prod_{r'=1}^r z_{i_{r'}} \hat{Q}_{i_{r'} l_{r'}}^t - g_t(z) \prod_{r'=1}^r z_{i_{r'}} Q_{i_{r'} l_{r'}}^t \right\} \right] \right| \\
 \leq & KLP^L \sum_{\substack{l_1, \dots, l_r \in [L] \\ \sum_{r'=1}^r l_{r'} = l-1}} \max_z \left| h_*(z) \prod_{r'=1}^r \hat{Q}_{i_{r'} l_{r'}}^t - g_t(z) \prod_{r'=1}^r Q_{i_{r'} l_{r'}}^t \right| \\
 \leq & C \sum_{\substack{l_1, \dots, l_r \in [L] \\ \sum_{r'=1}^r l_{r'} = l-1}} \left(\max_z |h_*(z) - g_t(z)| \left| \prod_{r'=1}^r \hat{Q}_{i_{r'} l_{r'}}^t \right| + |g_t(z)| \left| \prod_{r'=1}^r \hat{Q}_{i_{r'} l_{r'}}^t - \prod_{r'=1}^r Q_{i_{r'} l_{r'}}^t \right| \right) \\
 \leq & \sum_{\substack{l_1, \dots, l_r \in [L] \\ \sum_{r'=1}^r l_{r'} = l-1}} \left(CtC^r t^{l-1} + (2K) \left| \prod_{r'=1}^r \hat{Q}_{i_{r'} l_{r'}}^t - \prod_{r'=1}^r Q_{i_{r'} l_{r'}}^t \right| \right) \\
 \leq & Ct^l + C \sum_{\substack{l_1, \dots, l_r \in [L] \\ \sum_{r'=1}^r l_{r'} = l-1}} \sum_{r''=1}^r \left| \prod_{r'=1}^{r''} \hat{Q}_{i_{r'} l_{r'}}^t \prod_{r'=r''+1}^r Q_{i_{r'} l_{r'}}^t - \prod_{r'=1}^{r''-1} \hat{Q}_{i_{r'} l_{r'}}^t \prod_{r'=r''}^r Q_{i_{r'} l_{r'}}^t \right| \\
 \leq & Ct^l + C \sum_{\substack{l_1, \dots, l_r \in [L] \\ \sum_{r'=1}^r l_{r'} = l-1}} \sum_{r''=1}^r \left| \hat{Q}_{i_{r''} l_{r''}}^t - Q_{i_{r''} l_{r''}}^t \right| C^{r-1} t^{l-1-l_{r''}} \\
 \leq & Ct^l + C \sum_{\substack{l_1, \dots, l_r \in [L] \\ \sum_{r'=1}^r l_{r'} = l-1}} \sum_{r''=1}^r \left(C_{l_{r''}} t^{l_{r''}+1} \right) C^{r-1} t^{l-1-l_{r''}} \\
 \leq & C_l t^l,
 \end{aligned}$$

where the second-to-last-line was by the inductive hypothesis. Since $Q_{il}^0 = \hat{Q}_{il}^0 = 0$, we conclude $|\hat{Q}_{il}^t - Q_{il}^t| \leq C_l t^{l+1}$. \blacksquare

The above lemma will be used in Section F.7 to show that it suffices to analyze the dynamics of \hat{Q}^t instead of the dynamics of Q^t , and in turn instead of the dynamics of u^t .

F.6. Recurrence relation of the coefficients in the simplified dynamics

We prove that each entry of the matrix \hat{Q}^t is a polynomial in t and the Fourier coefficients of h_* , and we give a recurrence relation for the coefficients. Define $\alpha = \{\alpha_S\}_{S \subseteq [P]} \in \mathbb{R}^{2^P}$ where $\alpha_S = \mathbb{E}_z[\chi_S(z) h_*(z)]$, and $\mathbf{m} = (m_0, \dots, m_L) \in \mathbb{R}^{L+1}$ where $m_i = \sigma^{(i)}(0)$ for all $i \in \{0, \dots, L\}$.

Lemma 63 *For each $i \in [P]$, $l \in [L]$, we have $\hat{Q}_{il}^t = t^l p_{il}(\alpha, \mathbf{m})$, where p is a polynomial in the Fourier coefficients α of h_* and in the first L derivatives \mathbf{m} of σ . Furthermore, $\{p_{il}\}_{i \in [P], l \in [L]}$*

satisfies the recurrence relations $p_{i1} = \alpha_{\{i\}} m_1$ and

$$p_{il}(\boldsymbol{\alpha}, \mathbf{m}) = \frac{1}{l} \sum_{S \subseteq [P]} \sum_{1 \leq r \leq L-1} \sum_{i_1, \dots, i_r \in [P]} \sum_{\substack{l_1, \dots, l_r \in [L] \\ \sum_{r'=1}^r l_{r'} = l-1}} 1(\{i\} \oplus S \oplus \{i_1\} \cdots \oplus \{i_r\} = \emptyset) \frac{m_{r+1}}{r!} \alpha_S \prod_{r'=1}^r p_{i_r, l_{r'}}(\boldsymbol{\alpha}, \mathbf{m}).$$

Proof The proof is by induction on l . In the base case, for any $i \in [P]$,

$$\frac{\partial \hat{Q}_{i1}^t}{\partial t} = \mathbb{E}_{\mathbf{z}}[z_i h_*(\mathbf{z}) m_1] = \alpha_{\{i\}} m_1,$$

so $\hat{Q}_{i1}^t = t \alpha_{\{i\}} m_1$. For the inductive step, suppose that the lemma is true for all $i \in [P]$ and $l' \in \{1, \dots, l-1\}$. Then

$$\begin{aligned} & \frac{\partial \hat{Q}_{il}^t}{\partial t} \\ &= \mathbb{E}_{\mathbf{z}} \left[z_i \left\{ \sum_{S \subseteq [P]} \alpha_S \chi_S(\mathbf{z}) \right\} \sum_{1 \leq r \leq L-1} \frac{m_{r+1}}{r!} \sum_{i_1, \dots, i_r \in [P]} \sum_{\substack{l_1, \dots, l_r \in [L] \\ \sum_{r'=1}^r l_{r'} = l-1}} \prod_{r'=1}^r z_{i_{r'}} t^{l_{r'}} p_{i_r, l_{r'}}(\boldsymbol{\alpha}, \mathbf{m}) \right] \\ &= t^{l-1} \sum_{S \subseteq [P]} \sum_{1 \leq r \leq L-1} \sum_{i_1, \dots, i_r \in [P]} \sum_{\substack{l_1, \dots, l_r \in [L] \\ \sum_{r'=1}^r l_{r'} = l-1}} \left\{ \mathbb{E}_{\mathbf{z}} \left[z_i \chi_S(\mathbf{z}) \prod_{r'=1}^r z_{i_{r'}} \right] \right\} \frac{m_{r+1}}{r!} \alpha_S \prod_{r'=1}^r p_{i_r, l_{r'}}(\boldsymbol{\alpha}, \mathbf{m}) \\ &= t^{l-1} \sum_{S \subseteq [P]} \sum_{1 \leq r \leq L-1} \sum_{i_1, \dots, i_r \in [P]} \sum_{\substack{l_1, \dots, l_r \in [L] \\ \sum_{r'=1}^r l_{r'} = l-1}} 1(\{i\} \oplus S \oplus \{i_1\} \cdots \oplus \{i_r\} = \emptyset) \frac{m_{r+1}}{r!} \alpha_S \prod_{r'=1}^r p_{i_r, l_{r'}}(\boldsymbol{\alpha}, \mathbf{m}). \end{aligned}$$

The recurrence relation follows by integrating with respect to t . ■

We will subsequently prove that it suffices to study $\hat{\mathbf{u}}$, for which the recurrence relation in Lemma 63 becomes useful.

E.7. Reduction to analyzing the simplified dynamics

Let us study the training in Phase 2, where we train the second layer from time T_1 to time T_2 , while keeping the first layer fixed. Furthermore, we train with the perturbed activation function σ_{pert} . In order to prove that the training of the second layer converges, it is sufficient to prove that the kernel obtained as the linearization of the second layer weights, after the training in Phase 1 has condition number bounded by a constant C depending only on K, L, P . Define the kernel $\mathbf{K}^{T_1} : \{+1, -1\}^P \times \{+1, -1\}^P \rightarrow \mathbb{R}$ for times $t \geq T_1$ as

$$\mathbf{K}^{T_1}(\mathbf{z}, \mathbf{z}') = \mathbb{E}_{(a, \mathbf{u}^{T_1}) \sim \tilde{\rho}^{T_1}} [\sigma_{\text{pert}}(\langle \mathbf{u}^{T_1}, \mathbf{z} \rangle) \sigma_{\text{pert}}(\langle \mathbf{u}^{T_1}, \mathbf{z}' \rangle)].$$

In order to bound the learning in Phase 2, it is sufficient to bound the minimum eigenvalue of \mathbf{K}^{T_1} . To this aim, define the kernel $\tilde{\mathbf{K}}^{T_1} : \{+1, -1\}^P \times \{+1, -1\}^P \rightarrow \mathbb{R}$ corresponding to the $\tilde{\mathbf{u}}^t$ dynamics as:

$$\tilde{\mathbf{K}}^{T_1}(\mathbf{z}, \mathbf{z}') = \mathbb{E}_{a \sim \mu_a} [\hat{\sigma}_{\text{pert}}(\langle \tilde{\mathbf{u}}^{T_1}(a), \mathbf{z} \rangle) \hat{\sigma}_{\text{pert}}(\langle \tilde{\mathbf{u}}^{T_1}(a), \mathbf{z}' \rangle)],$$

where

$$\hat{\sigma}_{pert}(s) = \sum_{r=0}^L \frac{m_r + \rho_r}{r!} s^r$$

is the degree- L approximation of the perturbed activation function σ_{pert} . Recall that the perturbation ρ is chosen so that $\rho_r \sim \text{Unif}[-\tau_{pert}, \tau_{pert}]$ for all $r \in \{0, \dots, 2^{8P}\}$ and 0 otherwise.

We bound the minimum eigenvalue of \mathbf{K}^{T_1} by the minimum eigenvalue of $\tilde{\mathbf{K}}^{T_1}$ by showing that the kernel $\tilde{\mathbf{K}}^{T_1}$ is $O(t^L)$ close in spectral norm to the kernel \mathbf{K}^{T_1} .

Claim 64 *There are constants $c, C > 0$ depending on K, L, P such that, for all $T_1 \leq c$,*

$$\lambda_{\min}(\mathbf{K}^{T_1}) \geq \lambda_{\min}(\tilde{\mathbf{K}}^{T_1}) - C(T_1)^L$$

Proof This follows by proving that \mathbf{K}^{T_1} and $\tilde{\mathbf{K}}^{T_1}$ are close in spectral norm: i.e., $\|\mathbf{K}^{T_1} - \tilde{\mathbf{K}}^{T_1}\| \leq C(T_1)^L$. For any \mathbf{z}, \mathbf{z}' ,

$$\begin{aligned} & |\mathbf{K}^{T_1}(\mathbf{z}, \mathbf{z}') - \tilde{\mathbf{K}}^{T_1}(\mathbf{z}, \mathbf{z}')| \\ & \leq \mathbb{E}_{a \sim \mu_a} [|\sigma_{pert}(\langle \mathbf{u}^{T_1}(a), \mathbf{z} \rangle) \sigma_{pert}(\langle \mathbf{u}^{T_1}(a), \mathbf{z}' \rangle) - \hat{\sigma}_{pert}(\langle \tilde{\mathbf{u}}^{T_1}(a), \mathbf{z} \rangle) \hat{\sigma}_{pert}(\langle \tilde{\mathbf{u}}^{T_1}(a), \mathbf{z}' \rangle)|] \\ & \leq \mathbb{E}_{a \sim \mu_a} [|\sigma_{pert}(\langle \mathbf{u}^{T_1}(a), \mathbf{z} \rangle) \sigma_{pert}(\langle \mathbf{u}^{T_1}(a), \mathbf{z}' \rangle) - \sigma_{pert}(\langle \tilde{\mathbf{u}}^{T_1}(a), \mathbf{z} \rangle) \sigma_{pert}(\langle \tilde{\mathbf{u}}^{T_1}(a), \mathbf{z}' \rangle)| \\ & \quad + |\sigma_{pert}(\langle \tilde{\mathbf{u}}^{T_1}(a), \mathbf{z} \rangle) \sigma_{pert}(\langle \tilde{\mathbf{u}}^{T_1}(a), \mathbf{z}' \rangle) - \hat{\sigma}_{pert}(\langle \tilde{\mathbf{u}}^{T_1}(a), \mathbf{z} \rangle) \hat{\sigma}_{pert}(\langle \tilde{\mathbf{u}}^{T_1}(a), \mathbf{z}' \rangle)|] \\ & \leq \mathbb{E}_{a \sim \mu_a} [2K^2 \sqrt{P} \|\mathbf{u}^{T_1}(a) - \tilde{\mathbf{u}}^{T_1}(a)\| + C \|\tilde{\mathbf{u}}^{T_1}(a)\|^{L+1}] \\ & \leq C(T_1)^L, \end{aligned}$$

where we use that $\|\sigma\|_\infty, \|\sigma'\|_\infty \leq K$, $\|\mathbf{z}\|, \|\mathbf{z}'\| \leq \sqrt{P}$, and also the Taylor series error bound and the fact that $\|\tilde{\mathbf{u}}^t\| \leq Ct$ (by Claim 58), and $\|\mathbf{u} - \tilde{\mathbf{u}}\| \leq Ct^L$ by Claim 59 for $0 \leq t \leq T_1$. So $\|\mathbf{K}^{T_1} - \tilde{\mathbf{K}}^{T_1}\| \leq \|\mathbf{K}^{T_1} - \tilde{\mathbf{K}}^{T_1}\|_F \leq 2^P C(T_1)^L \leq C(T_1)^L$. \blacksquare

So if we can prove that $\lambda_{\min}(\tilde{\mathbf{K}}^{T_1}) \geq \Omega((T_1)^l)$ for any $l < L$, then for sufficiently small T_1 this implies that $\lambda_{\min}(\mathbf{K}^{T_1}) > c(T_1)^l$ for some constant $c > 0$ depending on K, L, P . This would prove that the condition number of \mathbf{K}^{T_1} is bounded by a constant independent of d .

We now show a strategy to prove that $\lambda_{\min}(\tilde{\mathbf{K}}^{T_1}) \geq \Omega((T_1)^l)$, by analyzing the $\hat{\mathbf{u}}^t$ dynamics instead of the $\tilde{\mathbf{u}}^t$ dynamics. We must use a much more delicate argument than the bound used to compare \mathbf{K}^{T_1} and $\tilde{\mathbf{K}}^{T_1}$. The reason is that we used $\|\mathbf{u}^{T_1} - \tilde{\mathbf{u}}^{T_1}\| \leq O((T_1)^L)$, but it is not necessarily true that $\tilde{\mathbf{u}}^{T_1}$ and $\hat{\mathbf{u}}^{T_1}$ are $O((T_1)^L)$ -close in L_2 norm. In fact, we typically have $\|\tilde{\mathbf{u}}^{T_1} - \hat{\mathbf{u}}^{T_1}\| \geq \Omega((T_1)^2)$. So we instead use the fact that $\tilde{\mathbf{u}}^{T_1}$ and $\hat{\mathbf{u}}^{T_1}$ are polynomials in a , and their coefficients are close as polynomials in a (previously proved in Lemma 62).

Let us first prove a lower-bound on $\lambda_{\min}(\tilde{\mathbf{K}}^{T_1})$ in terms of the determinant of a certain ‘‘feature matrix’’ $\tilde{M} : \{+1, -1\}^P \times [2^P] \rightarrow \mathbb{R}$ indexed by $\mathbf{z} \in \{+1, -1\}^P$ and $j \in [2^P]$ as

$$\tilde{M}(\mathbf{z}, j) = \hat{\sigma}_{pert}(\langle \tilde{\mathbf{u}}^{T_1}(a_j), \mathbf{z} \rangle),$$

where $\mathbf{a} = [a_1, \dots, a_{2^P}] \in \mathbb{R}^{2^P}$ is a vector of indeterminate variables. We bound $\lambda_{\min}(\tilde{\mathbf{K}}^{T_1})$ in terms of the expected magnitude of the determinant of \tilde{M} , for random $\mathbf{a} \sim \mu_a^{\otimes 2^P} = \text{Unif}([-1, 1]^{\otimes 2^P})$.

Claim 65 *There is a constant $c > 0$ depending on K, L, P such that, for all $T_1 \leq c$,*

$$\lambda_{\min}(\tilde{\mathbf{K}}^{T_1}) \geq c \mathbb{E}_{\mathbf{a} \sim \text{Unif}([-1,1]^{\otimes 2P})} [\det(\tilde{\mathbf{M}})^2].$$

Proof Since

$$[\tilde{\mathbf{M}}\tilde{\mathbf{M}}^\top](\mathbf{z}, \mathbf{z}') = \sum_{j \in [2^P]} \hat{\sigma}_{\text{pert}}(\langle \tilde{\mathbf{u}}^{T_1}(a_j), \mathbf{z} \rangle) \hat{\sigma}_{\text{pert}}(\langle \tilde{\mathbf{u}}^{T_1}(a_j), \mathbf{z}' \rangle),$$

we can write $\tilde{\mathbf{K}}^{T_1}$ in terms of this matrix product

$$\tilde{\mathbf{K}}^{T_1} = \frac{1}{2^P} \mathbb{E}_{\mathbf{a} \sim \text{Unif}([-1,1]^{\otimes 2P})} [\tilde{\mathbf{M}}\tilde{\mathbf{M}}^\top].$$

So, since $\tilde{\mathbf{M}}\tilde{\mathbf{M}}^\top$ is p.s.d.,

$$\lambda_{\min}(\tilde{\mathbf{K}}^{T_1}) \geq \frac{1}{2^P} \mathbb{E}_{\mathbf{a} \sim \text{Unif}([-1,1]^{\otimes 2P})} [\lambda_{\min}(\tilde{\mathbf{M}}\tilde{\mathbf{M}}^\top)]. \quad (74)$$

For any $a_j \in [-1, 1]$ and assuming $T_1 \leq c$ is small enough, we have $|\langle \tilde{\mathbf{u}}^{T_1}(a_j), \mathbf{z} \rangle| \leq c$ for some small enough constant $c > 0$ so that $|\hat{\sigma}_{\text{pert}}(\langle \tilde{\mathbf{u}}^{T_1}(a_j), \mathbf{z} \rangle)| \leq C$ for some large enough constant C depending on K, L, P . This means $\lambda_{\max}(\tilde{\mathbf{M}}\tilde{\mathbf{M}}^\top) \leq 2^{2P} C^2 \leq C$ almost surely. So $\lambda_{\min}(\tilde{\mathbf{M}}\tilde{\mathbf{M}}^\top) \geq (\lambda_{\max}(\tilde{\mathbf{M}}\tilde{\mathbf{M}}^\top))^{-2^P+1} \prod_{i=1}^{2^P} \lambda_i(\tilde{\mathbf{M}}\tilde{\mathbf{M}}^\top) \geq \det(\tilde{\mathbf{M}})^2 / C^{2^P-1} \geq c \det(\tilde{\mathbf{M}})^2$. This proves the claim when combined with the lower bound (74). \blacksquare

It remains to lower-bound the magnitude of the determinant of $\tilde{\mathbf{M}}$, for $\mathbf{a} \sim \text{Unif}([-1, 1]^{\otimes 2P})$. First, we note that the determinant is a polynomial in \mathbf{a} .

Claim 66 *For each $\gamma \in \{0, \dots, L^2\}^{2^P}$, there is a coefficient \tilde{h}_γ depending only on t, h_*, \mathbf{m} , and ρ such that*

$$\det(\tilde{\mathbf{M}}) = \sum_{\gamma \in \{0, \dots, L^2\}^{2^P}} \tilde{h}_\gamma \mathbf{a}^\gamma.$$

In other words, the determinant is a polynomial in \mathbf{a} of individual degree at most L^2 .

Proof For each $i \in [P]$, and $j \in [2^P]$, recall that $\tilde{\mathbf{u}}_i^t(a_j) = \sum_{l \in [L]} (a_j)^l Q_{il}^t$. Here Q_{il}^t depends only on $i, l, t, \alpha, \mathbf{m}$ and does not depend on \mathbf{a} . So for each $\mathbf{z} \in \{+1, -1\}^P$, $\tilde{\mathbf{M}}(\mathbf{z}, j) = \hat{\sigma}_{\text{pert}}(\langle \tilde{\mathbf{u}}^t(a_j), \mathbf{z} \rangle) = \sum_{r=0}^L \frac{m_r + \rho_r}{r!} \langle \tilde{\mathbf{u}}^t(a_j), \mathbf{z} \rangle^r = \sum_{r=0}^L \frac{m_r + \rho_r}{r!} \left(\sum_{i \in [P]} z_i \sum_{l \in [L]} (a_j)^l Q_{il}^t \right)^r$. So since each entry of $\tilde{\mathbf{M}}$ is a polynomial in \mathbf{a} , the determinant is also a polynomial in \mathbf{a} . \blacksquare

We can prove that in expectation over $\mathbf{a} \sim \text{Unif}([-1, 1]^{\otimes 2P})$ this determinant is nonzero if it has nonzero coefficients of low degree:

Claim 67 *There is a constant $c > 0$ depending on L and P such that for all $0 \leq t \leq T_1$,*

$$\mathbb{E}[\det(\tilde{\mathbf{M}})^2] \geq c \sum_{\gamma \in \{0, \dots, L^2\}^{2^P}} |\tilde{h}_\gamma|^2$$

Proof The proof is by writing $\det(\tilde{\mathbf{M}})$ in the Legendre basis, lower-bounding its coefficients in this basis, and using the orthogonality of the Legendre polynomials. This is Lemma 103. \blacksquare

This leaves the question of how to prove that $\det(\tilde{\mathbf{M}})$ is a nonzero polynomial with some nonzero term of degree t^l where $l < 2L$. Here we show that this problem can be reduced to analyzing the $\hat{\mathbf{u}}$ dynamics, which are simpler to analyze since they do not have the dependence on \hat{f}_{NN} and admit the recurrence relations of Lemma 63. Similarly to the definition of $\tilde{\mathbf{M}}$, we can define $\hat{\mathbf{M}} : \{+1, -1\}^P \times [2^P] \rightarrow \mathbb{R}$ by

$$\hat{\mathbf{M}}(\mathbf{z}, j) = \hat{\sigma}(\langle \tilde{\mathbf{u}}^{T_1}(a_j), \mathbf{z} \rangle).$$

Similarly to $\tilde{\mathbf{M}}$, we can prove that each entry of $\hat{\mathbf{M}}$ is a polynomial in \mathbf{a} .

Claim 68 *For each $\gamma \in \{0, \dots, L^2\}^{2^P}$, there is a coefficient \hat{h}_γ depending only on t, h_*, \mathbf{m} and ρ such that*

$$\det(\hat{\mathbf{M}}) = \sum_{\gamma \in \{0, \dots, L^2\}^{2^P}} \hat{h}_\gamma \mathbf{a}^\gamma.$$

In other words, the determinant is a polynomial in \mathbf{a} of individual degree at most L^2 .

Proof Same as the proof of Claim 66. \blacksquare

This is useful, since we can show that the coefficients of $\det(\hat{\mathbf{M}})$ are close to those of $\det(\tilde{\mathbf{M}})$.

Claim 69 *There is a constant $C > 0$ depending on K, L, P , such that for any $\gamma \in \{0, \dots, L^2\}^{2^P}$,*

$$|\tilde{h}_\gamma - \hat{h}_\gamma| \leq C(T_1)^{\|\gamma\|_1 + 1}.$$

Proof We write $\det(\tilde{\mathbf{M}})$ as a sum over permutations τ ,

$$\det(\tilde{\mathbf{M}}) = \sum_{\tau \in S_{2^P}} \text{sgn}(\tau) \prod_{j=1}^{2^P} \sum_{r=0}^L \frac{m_r + \rho_r}{r!} \left(\sum_{i \in [P]} \tau(j)_i \sum_{l \in [L]} (a_j)^l Q_{il}^{T_1} \right)^r.$$

Therefore,

$$\tilde{h}_\gamma = \sum_{\tau \in S_{2^P}} \text{sgn}(\tau) \sum_{r_1, \dots, r_{2^P} \in \{0, \dots, L\}} \prod_{j=1}^{2^P} \frac{m_{r_j} + \rho_{r_j}}{r_j!} \times \left[\delta_{r_j, 0} + (1 - \delta_{r_j, 0}) \left\{ \sum_{i_1, \dots, i_{r_j} \in [P]} \sum_{l_1, \dots, l_{r_j} \in [L]} \prod_{r'=1}^{r_j} \tau(j)_{i_{r'}} Q_{i_{r'} l_{r'}}^{T_1} \right\} \right]_{\sum_{r'=1}^{r_j} l_{r'} = \gamma_j},$$

and the same expression holds for \hat{h}_γ , with \hat{Q} replacing Q . Since $Q_{il}^t, \hat{Q}_{il}^t \leq Ct^l$ by Claim 57 and 61 and $Q_{il}^t - \hat{Q}_{il}^t \leq Ct^{l+1}$ by Lemma 62, we conclude by a triangle inequality and telescoping that $|\tilde{h}_\gamma - \hat{h}_\gamma| \leq C(T_1)^{l+1}$ for a constant C depending only on K, L, P . \blacksquare

Furthermore, in fact $\det(\hat{\mathbf{M}})$ has the special structure that each coefficient \hat{h}_γ is of size proportional to $(T_1)^{\|\gamma\|_1}$ if it is nonzero:

Claim 70 For any $\gamma \in \{0, \dots, L^2\}^{2^P}$, there is a polynomial $q_\gamma(\alpha, \mathbf{m}, \rho)$ such that

$$\hat{h}_\gamma = (T_1)^{\|\gamma\|_1} q_\gamma(\alpha, \mathbf{m}, \rho).$$

Proof By direct calculation,

$$\begin{aligned} \hat{h}_\gamma &= \sum_{\tau \in S_{2^P}} \text{sgn}(\tau) \sum_{r_1, \dots, r_{2^P} \in \{0, \dots, L\}} \\ &\quad \prod_{j=1}^{2^P} \frac{m_{r_j} + \rho_{r_j}}{r_j!} \left[\delta_{r_j, 0} + (1 - \delta_{r_j, 0}) \left\{ \sum_{i_1, \dots, i_{r_j} \in [P]} \sum_{\substack{l_1, \dots, l_{r_j} \in [L] \\ \sum_{r'=1}^{r_j} l_{r'} = \gamma_j}} \prod_{r'=1}^{r_j} \tau(j)_{i_{r'}} \hat{Q}_{i_{r'} l_{r'}}^{T_1} \right\} \right]. \end{aligned}$$

Since $\hat{Q}_{i_{r'} l_{r'}} = (T_1)^{l_{r'}} p_{i_{r'} l_{r'}}(\alpha, \mathbf{m})$ by Lemma 63, we have

$$\begin{aligned} \hat{h}_\gamma &= \sum_{\tau \in S_{2^P}} \text{sgn}(\tau) \sum_{r_1, \dots, r_{2^P} \in \{0, \dots, L\}} \\ &\quad \prod_{j=1}^{2^P} \frac{m_{r_j} + \rho_{r_j}}{r_j!} (T_1)^{\gamma_j} \left[\delta_{r_j, 0} + (1 - \delta_{r_j, 0}) \left\{ \sum_{i_1, \dots, i_{r_j} \in [P]} \sum_{\substack{l_1, \dots, l_{r_j} \in [L] \\ \sum_{r'=1}^{r_j} l_{r'} = \gamma_j}} \prod_{r'=1}^{r_j} \tau(j)_{i_{r'}} p_{i_{r'} l_{r'}}(\alpha, \mathbf{m}) \right\} \right]. \end{aligned}$$

We deduce that

$$\hat{h}_\gamma = (T_1)^{\|\gamma\|_1} q_\gamma(\alpha, \mathbf{m}, \rho),$$

where $q_\gamma(\alpha, \mathbf{m}, \rho)$ is the polynomial defined by

$$\begin{aligned} & q_\gamma(\alpha, \mathbf{m}, \rho) \\ &= \sum_{\tau \in S_{2^P}} \text{sgn}(\tau) \sum_{r_1, \dots, r_{2^P} \in \{0, \dots, L\}} \\ &\quad \prod_{j=1}^{2^P} \frac{m_{r_j} + \rho_{r_j}}{r_j!} \left[\delta_{r_j, 0} + (1 - \delta_{r_j, 0}) \left\{ \sum_{i_1, \dots, i_{r_j} \in [P]} \sum_{\substack{l_1, \dots, l_{r_j} \in [L] \\ \sum_{r'=1}^{r_j} l_{r'} = \gamma_j}} \prod_{r'=1}^{r_j} \tau(j)_{i_{r'}} p_{i_{r'} l_{r'}}(\alpha, \mathbf{m}) \right\} \right], \end{aligned}$$

which concludes the proof of the claim. \blacksquare

Combining the above claims we obtain a bound on the determinant of $\tilde{\mathbf{M}}$ in terms of the \hat{u} dynamics.

Claim 71 Suppose that for some $\gamma \in \{0, \dots, L^2\}^{2^P}$, we have $q_\gamma(\alpha, \mathbf{m}, \rho) \neq 0$. Then there is a small enough constant $c > 0$ depending on $K, L, P, \alpha, \mathbf{m}, \gamma$, such that for all $T_1 \leq c$,

$$\mathbb{E}[\det(\tilde{\mathbf{M}})^2] \geq c(T_1)^{2\|\gamma\|_1}.$$

Proof By combining Claims 67, 69, and 70, we know that there is a large enough constant $C > 0$ and small enough constant $c > 0$ such that

$$\mathbb{E}[\det(\tilde{\mathbf{M}})^2] \geq c \sum_{\gamma \in \{0, \dots, L^2\}^{2P}} |\min(0, (T_1)^{\|\gamma\|_1} q_\gamma(\boldsymbol{\alpha}, \mathbf{m}, \boldsymbol{\rho}) - C(T_1)^{\|\gamma\|_1+1})|^2.$$

Choosing $c > 0$ smaller than $|q_\gamma(\boldsymbol{\alpha}, \mathbf{m}, \boldsymbol{\rho})|/(2C)$ concludes the claim. \blacksquare

We conclude by combining all of the above claims to get the result of this subsection:

Lemma 72 *Suppose that for some $\gamma \in \{0, \dots, L^2\}^{2P}$ such that $\|\gamma\|_1 < L/2$ we have $q_\gamma(\boldsymbol{\alpha}, \mathbf{m}, \boldsymbol{\rho}) \neq 0$. Then there is a small enough constant $c > 0$ depending on $K, L, P, \boldsymbol{\alpha}, \mathbf{m}, \gamma, \boldsymbol{\rho}$ such that for all $T_1 \leq c$ we have*

$$\lambda_{\min}(\mathbf{K}^{T_1}) \geq c(T_1)^{2\|\gamma\|_1}.$$

Proof This is immediate by combining Claims 64, 65 and 71. \blacksquare

F.8. Proving learnability of generic MSP functions, Theorem 56

Here we give the final technical step to proving that generic MSP functions are learnable. The proof idea is to use Lemma 72 to lower-bound the minimum eigenvalue of the kernel matrix \mathbf{K}^{T_1} . By Lemma 72, it suffices to prove that for any minimal MSP structure $\mathcal{S} \subseteq 2^{[P]}$, if we plug in $\alpha_S = 0$ for all $S \notin \mathcal{S}$ the determinant $\det(\hat{\mathbf{M}})$ almost surely is a non-zero polynomial in t with nonzero low-order terms. In other words, the main technical lemma that remains to be proved is the following.

Lemma 73 *Let $\mathcal{S} \subseteq 2^{[P]}$ be any MSP set structure on P variables. Then there are constants $l_{\mathcal{S}}$ and $L_{\mathcal{S}}$ depending only on \mathcal{S} such that if we take the truncation to the dynamics to be $L \geq L_{\mathcal{S}}$ then $\det(\hat{\mathbf{M}}) |_{(\alpha_S)_{S \notin \mathcal{S}} = 0}$ is a polynomial in $t, \mathbf{a}, \{\boldsymbol{\alpha}\}_{S \in \mathcal{S}}, \mathbf{m}, \boldsymbol{\rho}$ that has a nonzero term with degree $l_{\mathcal{S}}$ in t .*

Before we show this lemma, let us see how it implies the main theorem.

Proof [Proof of Theorem 56] Let $L_{\mathcal{S}}$ and $l_{\mathcal{S}}$ be as in Lemma 73. Choose the approximation parameter $L = \max(L_{\mathcal{S}}, 2l_{\mathcal{S}} + 1)$ for defining the dynamics $\hat{\mathbf{u}}$. We know that $\det(\hat{\mathbf{M}}) |_{(\alpha_S)_{S \notin \mathcal{S}} = 0}$ is a polynomial in $t, \mathbf{a}, \{\boldsymbol{\alpha}\}_{S \in \mathcal{S}}, \mathbf{m}$, and $\boldsymbol{\rho}$ that has a nonzero term with degree $l_{\mathcal{S}}$ in t . Therefore, almost surely over plugging in the activation perturbation $\boldsymbol{\rho} = [\rho_0, \dots, \rho_{2^{8P}}] \sim \text{Unif}[-\tau_{\text{pert}}, \tau_{\text{pert}}]^{\otimes 2^{8P}}$, the generic Fourier coefficients on the MSP set structure $(\alpha_S)_{S \in \mathcal{S}} \sim \text{Unif}[-1, 1]^{\otimes |\mathcal{S}|}$, and the zero Fourier coefficients outside the MSP set structure $(\alpha_S)_{S \notin \mathcal{S}} = 0$, we must have that $\det(\hat{\mathbf{M}})$ is a polynomial in t with a nonzero term of degree $l_{\mathcal{S}}$.

Since $\mathcal{S} \subseteq 2^{[P]}$, L and $l_{\mathcal{S}}$ are upper-bounded by a constant C that depends only on P . So by Lemma 72, we conclude that almost surely over $\boldsymbol{\rho}$ and $\boldsymbol{\alpha}$ there is a constant $c > 0$ depending only on $K, P, \boldsymbol{\rho}, \boldsymbol{\alpha}$ such that we have $\lambda_{\min}(\mathbf{K}^{T_1}) \geq c(T_1)^{2l_{\mathcal{S}}}$ as long as $T_1 \leq c$. In particular, choosing $T_1 = c$, then $\lambda_{\min}(\mathbf{K}^{T_1}) \geq c^2$.

For $t \geq T_1$, let $\mathbf{g}_t = (g_t(\mathbf{z}))_{\mathbf{z} \in \{+1, -1\}^P}$ denote the residual vector where $g_t(\mathbf{z}) = h_*(\mathbf{z}) - \hat{f}_{\text{NN}}^{\text{pert}}(\mathbf{z}; \bar{\rho}_t)$. Here $\hat{f}_{\text{NN}}^{\text{pert}}$ is \hat{f}_{NN} but with the activation σ replaced by the perturbed activation σ_{pert}

that is used in Phase 2. Recall that during Phase 2 the dynamics are linear since we are training the second layer, and are governed by kernel \mathbf{K}^{-T_1} . We have following bound on the norm of the residuals for $t \geq T_1$:

$$\|\mathbf{g}_t\|_2^2 \leq e^{-\lambda_{\min}(\mathbf{K}^{T_1})(t-T_1)} \|\mathbf{g}_{T_1}\|_2^2.$$

Choose $T_1 = c$, and $T_2 = T_1 + \log(\|\mathbf{g}_{T_1}\|_2^2/\varepsilon)/c^2$ to achieve error $\varepsilon > 0$. Since $\|\mathbf{g}_{T_1}\|_2^2 \leq 2^P(\|h_*\|^2 + \|\hat{f}_{\text{NN}}^{\text{pert}}(\cdot; \bar{\rho}_{T_1})\|^2) \leq K$, we have that T_1 and T_2 are constants depending on K, P, α, ρ . This proves strong $O(d)$ -SGD learnability (with the variation that the activation function is perturbed at time T_1) almost surely over the Fourier coefficients α and the perturbation ρ . ■

F.9. Proof of Lemma 73

It only remains to show Lemma 73. To show this lemma, we will use the fact from Claim 70 that $\det(\hat{\mathbf{M}})$ is a polynomial in all relevant parameters: $t, \mathbf{a}, \alpha, \mathbf{m}, \rho$.

Claim 74 *There is a large enough integer D depending on L, P such that $\det(\hat{\mathbf{M}})$ is a polynomial of degree at most D in $\mathbf{a}, \mathbf{m}, \alpha, \rho$, and t .*

Proof This is by writing $\det(\hat{\mathbf{M}}) = \sum_{\gamma} t^{\|\gamma\|_1} \mathbf{a}^{\gamma} q_{\gamma}(\alpha, \mathbf{m}, \rho)$ where each q_{γ} is a polynomial, as proved in Claim 70. ■

To study this polynomial, we first reduce to studying “minimal” MSP set structures, defined as follows.

Definition 75 *We say that $\mathcal{S} = \{S_1, \dots, S_P\}$ is a minimal MSP set structure if the sets can be ordered such that for each $i \in [P]$ we have $S_i \subset [i]$ and $i \in S_i$.*

The following claim shows that it is sufficient to restrict our attention to minimal MSP set structures.

Claim 76 *Suppose that for every P there are constants $l_P, L_{P,0}$ depending only on P such that for any $L > L_{P,0}$, and every minimal MSP set structure $\mathcal{S} \subseteq 2^{[P]}$, the polynomial $\det(\hat{\mathbf{M}})|_{(\alpha_S)_{S \notin \mathcal{S}}=0}$ has a nonzero term with degree at most l_P in t .*

Then, for any $L > L_{P,0}$ and MSP set structure $\mathcal{S} \subseteq 2^{[P]}$, the polynomial $\det(\hat{\mathbf{M}})|_{(\alpha_S)_{S \notin \mathcal{S}'}=0}$ has a nonzero term with degree at most l_P in t .

Proof For any MSP set structure $\mathcal{S}' \subseteq 2^{[P]}$, up to a permutation of the variables there is a minimal MSP set structure $\mathcal{S} \subseteq 2^{[P]}$ such that $\mathcal{S} \subseteq \mathcal{S}'$. Since $\det(\hat{\mathbf{M}})|_{(\alpha_S)_{S \notin \mathcal{S}}=0}$ has a nonzero term with degree at most l_P , so does $\det(\hat{\mathbf{M}})|_{(\alpha_S)_{S \notin \mathcal{S}'}=0}$, because the former polynomial can be constructed from the latter by additionally setting $(\alpha_S)_{S \in \mathcal{S}' \setminus \mathcal{S}} = 0$, which could only zero out monomials. ■

Because of the above claim, for the remainder of this section, we fix a minimal MSP set structure $\mathcal{S} = \{S_1, \dots, S_P\}$. Let us analyze the behavior of the dynamics of $\hat{\mathbf{u}}$ on a function $h_*(\mathbf{z}) = \sum_{S \subseteq [P]} \alpha_S \chi_S(\mathbf{z})$ with this structure, i.e., with $\alpha_S = 0$ for all $S \notin \mathcal{S}$. Let us explicitly compute the leading order terms of the weights $\hat{\mathbf{u}}_i^t$ using the recurrence relations for the simplified dynamics. Recall that $\hat{\mathbf{u}}_i^t(\mathbf{a}) = \sum_{l=1}^L a^{t^l} p_{il}(\alpha, \mathbf{m})$.

Claim 77 Suppose that $L > 2^P$. For each $i \in [P]$, define

$$o_i = 1 + \sum_{i' \in S_i \setminus \{i\}} o_{i'}.$$

We have $p_{il}(\boldsymbol{\alpha}, \mathbf{m}) \big|_{(\alpha_S)_{S \notin \mathcal{S}} = \mathbf{0}} \equiv 0$ for all $l < o_i$, and for $l = o_i$ we have

$$p_{io_i}(\boldsymbol{\alpha}, \mathbf{m}) \big|_{(\alpha_S)_{S \notin \mathcal{S}} = \mathbf{0}} = \frac{\alpha_{S_i} m_{|S_i|}}{o_i} \prod_{i' \in S_i \setminus \{i\}} p_{i'o_{i'}},$$

with the convention that a product over an empty set is 1 and a sum over an empty set is 0.

Proof We prove this by induction on l using the recurrence relations for p_{il} derived in Lemma 63. For simplicity, we write $p_{il} = p_{il}((\alpha_S)_{S \in \mathcal{S}}, \mathbf{m}) = p_{il}(\boldsymbol{\alpha}, \mathbf{m}) \big|_{(\alpha_S)_{S \notin \mathcal{S}} = \mathbf{0}}$. First consider the base case of $l = 1$. For any i such that $o_i = 1$, we have $S_i = \{i\}$. Therefore, from the base case of the recurrence relations, we have $p_{io_i} = p_{i1} = t_{\alpha_{\{i\}}} m_1$. On the other hand, if $o_i > 1$, then $S_i \neq \{i\}$. By the minimality of the MSP structure we have $\{i\} \notin \mathcal{S}$ so $\alpha_{\{i\}} = 0$. Therefore $p_{i1} = t_{\alpha_{\{i\}}} m_1 = 0$.

For the inductive step, suppose $l \geq 2$ and that the result is true for $l' \in \{0, \dots, l-1\}$. Now consider any $S \in \mathcal{S}$, any $1 \leq r \leq L$ and any $(i_1, \dots, i_r) \in [P]^r$ such that $\{i\} \oplus S \oplus \{i_1\} \oplus \dots \oplus \{i_r\} = \emptyset$. Consider also any $l_1, \dots, l_r \in [L-1]^r$ such that $\sum_{r'} l_{r'} = l-1$. Each of these corresponds to a possible contribution to p_{il} in the recurrence relation of Lemma 63. Suppose that $l \leq o_i$.

Case 1: Suppose there is $i' \in \{i_1, \dots, i_r\}$ such that $o_{i'} \geq o_i$. Without loss of generality take $i' = i_1$. But since $l_1, \dots, l_r \leq l-1 < o_i \leq o_{i'}$, we have $p_{i_1 l_1} = p_{i' o_{i'}} = 0$ by the inductive hypothesis, so the terms in case 1 do not contribute.

Case 2: Suppose for all $r' \in [r]$ we have $o_{i_{r'}} < o_i$. Then $i \in S$ since otherwise $i \in \{i\} \oplus S \oplus \{i_1\} \oplus \dots \oplus \{i_r\}$ and of course $o_i \geq o_i$. If $S = S_{i'}$ for some $i' > i$, then we have $i' \in S$. And, as a consequence $i' \in \{i_1, \dots, i_r\}$, because otherwise $i' \in \{i\} \oplus S \oplus \{i_1\} \oplus \dots \oplus \{i_r\}$. However, $o_{i'} > o_i$ since $i \in S_{i'}$, so this is a contradiction. We conclude that $S = S_i$, and so $S_i \setminus \{i\} = \{i_1, \dots, i_r\}$. Since $\sum_{i' \in S_i \setminus \{i\}} o_{i'} = o_i - 1$ and $\sum_{r'=1}^r l_{r'} = l-1 \leq o_i$, we conclude that either *Case a:* there is some r' such that $l_{r'} < o_{i_{r'}}$, or *Case b:* $l_{r'} = o_{i_{r'}}$ for all $r' \in [r]$. In Case a, we have $p_{i_{r'} l_{r'}} = 0$ by the inductive hypothesis, so the term does not contribute to p_{il} . Case b occurs if and only if $l = o_i$ and i_1, \dots, i_r are a permutation of $S_i \setminus \{i\}$. There are exactly $(|S_i| - 1)!$ such terms, so the recurrence relation for p_{il} holds. \blacksquare

For any $\mathbf{z} \in \{+1, -1\}^P$, define the multivariable polynomial

$$q_{\mathbf{z}}(a, T_1, \boldsymbol{\alpha}, \mathbf{m}) = \langle \hat{\mathbf{u}}^{T_1}(a), \mathbf{z} \rangle = \sum_{l=1}^L \sum_{i=1}^P (aT_1)^l p_{il}(\boldsymbol{\alpha}, \mathbf{m}).$$

Claim 78 There is a constant L_0 depending on P such that for large enough truncation $L > L_0$, for any $\mathbf{z} \neq \mathbf{z}' \in \{+1, -1\}^P$, $\frac{\partial}{\partial a}(q_{\mathbf{z}} - q_{\mathbf{z}'})$ has a nonzero term of degree at most 2^{P-1} in T_1 .

Proof Let us take a constant $L_0 = 2^P$. Then the low-order solutions to the recursion from Claim 77 are valid. There must be an index $i \in [P]$ such that $z_i \neq z'_i$. Choose $i \in \{i' : z_{i'} \neq z'_{i'}\}$ such that o_i

is minimized, breaking ties in favor larger i . Consider the terms of $\frac{\partial}{\partial a}(q_{\mathbf{z}} - q_{\mathbf{z}'})$ which are of degree o_i in T_1 . The degree o_i part is equal to

$$[T_1^{o_i}] \frac{\partial}{\partial a}(q_{\mathbf{z}} - q_{\mathbf{z}'}) = \sum_{i'=1}^P (z_{i'} - z'_{i'}) o_i a^{o_i-1} p_{i'o_i}(\boldsymbol{\alpha}, \mathbf{m}) = \sum_{\substack{i'=1 \\ z_{i'} \neq z'_{i'}}}^P (z_{i'} - z'_{i'}) o_i a^{o_i-1} p_{i'o_i}(\boldsymbol{\alpha}, \mathbf{m}).$$

Notice that if $z_{i'} \neq z'_{i'}$, then have $o_{i'} \geq o_i$ by the choice of i . And if $o_{i'} > o_i$ then $p_{i'o_i} \equiv 0$ by Claim 77. So

$$[T_1^{o_i}] \frac{\partial}{\partial a}(q_{\mathbf{z}} - q_{\mathbf{z}'}) = \sum_{\substack{i'=1 \\ o_{i'}=o_i}}^P (z_{i'} - z'_{i'}) o_i a^{o_i-1} p_{i'o_{i'}}(\boldsymbol{\alpha}, \mathbf{m}).$$

By the recurrence relations for $p_{i'o_{i'}}$ in Claim 77, one can see that $p_{i'o_i}$ is a monomial with degree 1 in α_{S_i} . On the other hand, for all $i' < i$, the polynomial $p_{i'o_{i'}}$ does not depend on α_{S_i} . Therefore $[T_1^{o_i}] \frac{\partial}{\partial a}(q_{\mathbf{z}} - q_{\mathbf{z}'})$ is a nonzero polynomial. So $\frac{\partial}{\partial a}(q_{\mathbf{z}} - q_{\mathbf{z}'})$ has a nonzero degree o_i term in T_1 . One can prove using the recurrence relation of Claim 77 inductively on i that $o_i \leq 2^{i-1}$. \blacksquare

Now consider the following matrix $\mathbf{N} \in \mathbb{R}^{2^P \times 2^P}$ indexed by $\mathbf{z} \in \{+1, -1\}^P$ and $j \in [2^P]$, and depending on some indeterminate scaling factor $\nu \in \mathbb{R}$,

$$N_{\mathbf{z},j} = \frac{\partial^{j-1}}{\partial a^{j-1}} \exp(\nu q_{\mathbf{z}})$$

We prove that $\det(\mathbf{N})$ has a low-order non-zero term in the analytic expansion of T_1 at 0. This is an auxiliary result that will allow us to prove the corresponding result for $\det(\hat{\mathbf{M}})$.

Claim 79 *There is a constant L_0 depending on $S \subseteq [2^P]$ such that for large enough $L > L_0$, there exists $l \leq 2^{3P}$ where*

$$\frac{\partial^l}{(\partial T_1)^l} \det(\mathbf{N}) |_{T_1=0}$$

equals a nonzero polynomial in $\nu, a, \boldsymbol{\alpha}, \mathbf{m}$.

Proof By the chain rule we may write $N_{\mathbf{z},j} = \exp(\nu q_{\mathbf{z}}) R_{\mathbf{z},j}$, for a function $R_{\mathbf{z},j}(a, T_1, \boldsymbol{\alpha}, \mathbf{m})$ defined inductively on j as $R_{\mathbf{z},1}(a, T_1, \boldsymbol{\alpha}, \mathbf{m}) \equiv 1$, and

$$R_{\mathbf{z},j+1} = R_{\mathbf{z},j} \frac{\partial}{\partial a} \nu q_{\mathbf{z}} + \frac{\partial}{\partial a} R_{\mathbf{z},j}.$$

So $\det(\mathbf{N}) = \left(\prod_{\mathbf{z} \in \{+1, -1\}^P} \exp(\nu q_{\mathbf{z}}) \right) \det(\mathbf{R})$, where \mathbf{R} is the matrix with entries $R_{\mathbf{z},j}$. Since each $R_{\mathbf{z},j}$ is a polynomial of degree $j-1$ in ν , $\det(\mathbf{R})$ is a polynomial of degree at most $\sum_{j=1}^{2^P} j - 1 = (2^P - 1)(2^P)/2 = \binom{2^P}{2}$ in ν . Let us consider the part of $\det(\mathbf{R})$ that has degree $\binom{2^P}{2}$ in ν . This must come from the degree $j-1$ part of each $R_{\mathbf{z},j}$, which can inductively be shown to be $[\nu^{j-1}] R_{\mathbf{z},j} = \left(\frac{\partial}{\partial a} q_{\mathbf{z}} \right)^{j-1}$. So $[\nu^{\binom{2^P}{2}}] \det(\mathbf{R}) = \det(\hat{\mathbf{R}})$, where $\hat{\mathbf{R}}$ is the matrix with entries

$$\hat{R}_{\mathbf{z},j} = \left(\frac{\partial}{\partial a} q_{\mathbf{z}} \right)^{j-1}.$$

This matrix is Vandermonde, so its determinant is (up to a factor of $+1$ or -1):

$$\det(\hat{\mathbf{R}}) = \prod_{z \neq z' \in \{+1, -1\}^P} \left(\frac{\partial}{\partial a} (q_z - q_{z'}) \right).$$

From Claim 78, we know that for each distinct z, z' , we have that $(\frac{\partial}{\partial a} (q_z - q_{z'}))$ has a nonzero term of degree at most 2^{P-1} in T_1 . Therefore $\det(\hat{\mathbf{R}})$ has a nonzero term of degree at most $\binom{2^P}{2} 2^{P-1} \leq 2^{3P}$ in T_1 . In particular, we have proved that $\det(\mathbf{R})$ is a polynomial in $\nu, a, T_1, \boldsymbol{\alpha}, \mathbf{m}$ that has a nonzero term of degree at most 2^{3P} in T_1 . Let $0 \leq l \leq 2^{3P}$ be the smallest l such that $[T_1^l] \det(\mathbf{R}) \neq 0$. Then we have

$$\begin{aligned} \frac{\partial^l}{(\partial T_1)^l} \det(\mathbf{N}) |_{T_1=0} &= \frac{\partial^l}{(\partial T_1)^l} \det(\mathbf{R}) \prod_z \exp(\nu q_z) |_{T_1=0} \\ &= \left(\frac{\partial^l}{(\partial T_1)^l} \det(\mathbf{R}) \right) \prod_z \exp(\nu q_z) |_{T_1=0} \\ &= \left(\frac{\partial^l}{(\partial T_1)^l} \det(\mathbf{R}) \right) |_{T_1=0}, \end{aligned}$$

since $q_z |_{T_1=0} \equiv 1$, since T_1 divides the polynomial q_z by its definition. ■

Now consider the following matrix $\hat{\mathbf{N}} \in \mathbb{R}^{2^P \times 2^P}$. We will eventually compare the determinant of $\hat{\mathbf{N}}$ to that of \mathbf{N} . Each entry of $\hat{\mathbf{N}}$ is a polynomial in $a, \boldsymbol{\alpha}, \mathbf{m}, \boldsymbol{\rho}$

$$\hat{N}_{z,j} = \frac{\partial^{j-1}}{\partial a^{j-1}} \hat{\sigma}_{pert}(q_z).$$

Let us prove that $\det(\hat{\mathbf{N}})$ has a low-order nonzero term in T_1 by comparing it to $\det(\mathbf{N})$.

Claim 80 *For any \mathcal{S} there is large enough truncation parameter L_0 , such that for $L > L_0$ there exists $l \leq 2^{3P}$ with $\frac{\partial^l}{(\partial T_1)^l} \det(\hat{\mathbf{N}}) |_{T_1=0} \neq 0$.*

Proof Suppose that we were to make the substitution $\rho_r = -m_r + \nu^r$ for each $r \in \{0, \dots, 2^{8P}\}$. Then we would get $\hat{N}_{z,j} = \sum_{r=0}^{2^{8P}} \frac{\nu^r}{r!} (q_z(a, T_1, \boldsymbol{\alpha}, \mathbf{m}))^r + \sum_{r=2^{8P}+1}^L \frac{m_r}{r!} (q_z(a, T_1, \boldsymbol{\alpha}, \mathbf{m}))^r$. Then since T_1 divides q_z and $\sum_{r=0}^{2^{8P}} \frac{\nu^r}{r!} s^r$ is the first few order expansion of $\exp(\nu s)$, for any $l \leq 2^{3P}$, we have

$$\begin{aligned} \frac{\partial^l}{(\partial T_1)^l} \det(\hat{\mathbf{N}}) |_{T_1=0} &= \frac{\partial^l}{(\partial T_1)^l} \sum_{\tau \in S_{2^P}} \text{sgn}(\tau) \prod_{j=1}^{2^P} \hat{N}_{\tau(j),j} |_{T_1=0} \\ &= \frac{\partial^l}{(\partial T_1)^l} \prod_{j=1}^{2^P} \left(\sum_{r=0}^{2^{8P}} \frac{\nu^r}{r!} (q_{\tau(j)}(a, T_1, \boldsymbol{\alpha}, \mathbf{m}))^r \right) |_{T_1=0} \\ &= \frac{\partial^l}{(\partial T_1)^l} \prod_{j=1}^{2^P} (\exp(\nu q_{\tau(j)}(a, T_1, \boldsymbol{\alpha}, \mathbf{m}))) |_{T_1=0} \\ &= \frac{\partial^l}{(\partial T_1)^l} \det(\mathbf{N}) |_{T_1=0}. \end{aligned}$$

Recall that by Claim 79, there is a $l \leq 2^{3P}$ such that $\frac{\partial^l}{(\partial T_1)^l} \det(\mathbf{N})|_{T_1=0}$ is a nonzero polynomial. Since we have derived the above by substituting $\rho_r = -m_r + \nu^r$, we must have that without substituting we have $\frac{\partial^l}{(\partial T_1)^l} \det(\hat{\mathbf{N}})|_{T_1=0}$ is a nonzero polynomial in $a, \alpha, \mathbf{m}, \rho$. ■

Furthermore, $\det(\hat{\mathbf{N}})$ is related to $\det(\hat{\mathbf{M}})$.

Claim 81 $\det(\hat{\mathbf{N}}) = \frac{\partial}{\partial a_2} \frac{\partial^2}{(\partial a_3)^2} \cdots \frac{\partial^{2^P-1}}{(\partial a_{2^P})^{2^P-1}} \det(\hat{\mathbf{M}})|_{a_1=a_2=\dots a_{2^P}=a}$.

Proof By linearity of the derivative,

$$\begin{aligned} \frac{\partial}{\partial a_2} \frac{\partial^2}{(\partial a_3)^2} \cdots \frac{\partial^{2^P-1}}{(\partial a_{2^P})^{2^P-1}} \det(\hat{\mathbf{M}}) &= \sum_{\tau \in S_{2^P}} \text{sgn}(\tau) \frac{\partial}{\partial a_2} \frac{\partial^2}{(\partial a_3)^2} \cdots \frac{\partial^{2^P-1}}{(\partial a_{2^P})^{2^P-1}} \prod_{j=1}^{2^P} \hat{\sigma}_{pert}(q_{\tau(j)}(a_j)) \\ &= \sum_{\tau \in S_{2^P}} \text{sgn}(\tau) \prod_{j=1}^{2^P} \frac{\partial^{j-1}}{(\partial a_j)^{j-1}} \hat{\sigma}_{pert}(q_{\tau(j)}(a_j)) \\ &= \det(\hat{\mathbf{N}}). \end{aligned}$$

■

Combining the above two claims allows us to conclude that there is a nonzero term in $\det(\hat{\mathbf{M}})$ that has low degree in T_1 . This concludes the proof of the lemma, which implies the theorem.

Proof [Proof of Lemma 73] By the above two claims, there is $l \leq 2^{3P}$ such that

$$\frac{\partial^l}{(\partial T_1)^l} \frac{\partial}{\partial a_2} \frac{\partial^2}{(\partial a_3)^2} \cdots \frac{\partial^{2^P-1}}{(\partial a_{2^P})^{2^P-1}} \det(\hat{\mathbf{M}})|_{a_1=a_2=\dots a_{2^P}=a, T_1=0} \neq 0.$$

This implies that $\det(\hat{\mathbf{M}})$ has a nonzero term of degree $l \leq 2^{3P}$ in T_1 . ■

Appendix G. Explicit sample-complexity bounds in all parameters

In this paper, we focused on the dependence of the sample complexity on the ambient dimension d . In particular, our main result shows that MSP is a necessary and nearly sufficient condition for a sparse function h_* to be learnable in $n = C(\varepsilon, \delta, h_*)d$ samples in the mean-field scaling (i.e., to achieve test error ε with probability $1 - \delta$). While this was not our goal, we note for the interested reader that our proof techniques provide explicit dependencies of the sample size in all parameters ε, δ, h_* . In this appendix, we gather these fully explicit sample-complexity bounds and leave for future work the task of improving them.

G.1. Vanilla staircase functions

As a first bound, let us naively use the propagation-of-chaos bound comparison between the dimension-free dynamics and batch-SGD (bSGD) presented in Theorem 5.

Proposition 82 *Consider learning a vanilla staircase:*

$$h_*(\mathbf{z}) = \alpha_{\{1\}}z_1 + \alpha_{\{1,2\}}z_1z_2 + \alpha_{\{1,2,3\}}z_1z_2z_3 + \dots + \alpha_{\{1,\dots,P\}}z_1z_2 \cdots z_P,$$

such that $\sum_{j \in [P]} |\alpha_{\{1,\dots,j\}}| \leq 1$ and denote $\alpha_* := \min_{j \in [P]} |\alpha_{\{1,\dots,j\}}|$. Then there exist a constant $K > 0$ that only depends on the activation σ and a numerical constant $C > 0$, such that the dynamics described in Section D reaches $\varepsilon > 0$ test error with probability at least $1 - 2^{-2^{CP}}$, with

$$n = Ce^{\log(1/\varepsilon)^C (K/\alpha_*)^{2^{CP}}} d, \quad N = Ce^{\log(1/\varepsilon)^C (K/\alpha_*)^{2^{CP}}}.$$

Proof [Proof of Proposition 82] The limiting mean-field dynamics verify at T_2 :

$$\|f_* - \hat{f}_{\text{NN}}(\cdot; \rho_{T_2})\|_{L^2}^2 \leq \varepsilon/4.$$

With $\mathbf{w}^0 = \mathbf{0}$ initialization, the mean-field and dimension free dynamics are the same and we can use the comparison bound between batch-SGD and the mean-field dynamics in Proposition 15, which yields that there exists a numerical constant $C > 0$ such that if

$$N = CK e^{KT_2^3}/\varepsilon, \quad \eta/b = e^{-KT_2^3}\varepsilon/(CK), \quad (75)$$

then $\|\hat{f}_{\text{NN}}(\cdot; \Theta^{k_2}) - \hat{f}_{\text{NN}}(\cdot; \rho_{T_2})\|_{L^2}^2 \leq \varepsilon/4$ with probability at least $1 - 1/N$. Note that we choose $T_2 = \log(K/\varepsilon)/\lambda_{\min}(\mathbf{K}^{T_1})$, and $T_1 = \lambda_{\min}(\mathbf{M})/2$. By Lemma 84, we have $\lambda_{\min}(\mathbf{M}) \geq 2^{-2^{-CP}}$. Furthermore, in Section D.1, we showed that

$$\lambda_{\min}(\mathbf{K}^{T_1}) \geq \{\min_S D_S^2\} \lambda_{\min}(\mathbf{M})/2,$$

where

$$\min_S D_S \geq (1/K) \cdot \prod_{s \in [P]} \nu_k(T_1) \geq (\alpha_*/K)^{2^{CP}}.$$

Injecting this lower bound on $\lambda_{\min}(\mathbf{K}^{T_1})$ in Eq. (75) yields the bounds in the proposition. \blacksquare

This first bound uses a worst case bound that depends exponentially on the training time, which scales as $e^{1/\lambda_{\min}(\mathbf{K}^{T_1})} \leq e^{e^P}$ because of Phase 2 of linear training. A more careful analysis of Phase 2 yields an error that scales as $1/\lambda_{\min}(\mathbf{K}^{T_1}) \leq e^{e^P}$ (see Section). This results in the following improved bound:

Proposition 83 *Follow the same setup as in Proposition 82. Then there exist a constant $K > 0$ that only depends on the activation σ and a numerical constant $C > 0$, such that h_* is strongly $O(d)$ -learnable with the following dependency on n and N :*

$$n = \left(\frac{K}{\alpha_*}\right)^{2^{CP}} \frac{\log(1/\varepsilon)}{\varepsilon^5} \log(1/\delta) d, \quad N = \left(\frac{K}{\alpha_*}\right)^{2^{CP}} \frac{\log(1/\delta)^2}{\varepsilon^3}.$$

Proof [Proof of Proposition 83] This follows from applying Lemma 101 presented in Section I.3, using Lemma 100 and the lower bound on the kernel matrix provided in the proof of Proposition 83. ■

G.1.1. TECHNICAL LEMMA

Lemma 84 *Consider $\mathbf{M} = (\mathbb{E}_a[a^{i+j-2}])_{i,j \in [2^P]} \in \mathbb{R}^{2^P \times 2^P}$ with $a \sim \text{Unif}([+1, -1])$, i.e., $M_{ij} = \frac{1}{i+j-1} \delta_{i+j \equiv 0[2]}$. Then there exists $C > 0$ independent of P such that*

$$\lambda_{\min}(\mathbf{M}) \geq 2^{-2^{CP}}.$$

Proof [Proof of Lemma 84] We follow a similar argument as in the proof of Lemma 103. First, note that

$$\lambda_{\min}(\mathbf{M}) = \inf_{\|\mathbf{u}\|_2=1} \mathbf{u}^\top \mathbf{M} \mathbf{u} = \inf_{\|\mathbf{u}\|_2=1} \mathbb{E}_a[h(a; \mathbf{u})^2],$$

where $h(a; \mathbf{u}) = \sum_{j=0}^{2^P-1} u_j a^j$. Consider P_l the degree- l Legendre polynomial on $[-1, 1]$ and denote

$$P_l(z) = \sum_{j=0}^l p_{l,j} z^j.$$

The polynomial $h(a; \mathbf{u})$ can be decomposed in this basis as

$$h(a; \mathbf{u}) = \sum_{l=0}^{2^P-1} g_l P_l(z).$$

In particular, we have $u_j = \sum_{l \geq j} g_l p_{l,j}$. Therefore,

$$\begin{aligned} 1 = \|\mathbf{u}\|_2^2 &= \sum_{j=0}^{D-1} \left(\sum_{l \geq j} g_l p_{l,j} \right)^2 \\ &\leq 2^{2^P} \cdot \left\{ \max_{j,l=0,\dots,2^P-1} p_{l,j}^2 \right\} \cdot \sum_{j,l=0}^{2^P-1} g_l^2 \leq 2^{2^{CP}} \cdot \mathbb{E}_a[h(a; \mathbf{u})^2], \end{aligned}$$

where we used that $|p_{l,l-2k}| = 2^{-l} \binom{l}{k} \binom{2l-2k}{l} \leq 2^{Cl}$ for some $C > 0$ and $|p_{l,l-2k-1}| = 0$ otherwise. This concludes the proof. ■

G.2. Merged-staircase functions in the smoothed complexity model

Similarly, the tighter analysis of second-layer training in Lemmas 100 and 101 can be used to provide a complexity bound for learning MSP functions under a smoothed model of complexity. This corresponds to making the dependencies of Theorem 9 on parameters other than d tighter.

Proposition 85 *Let $h_* : \{+1, -1\}^P \rightarrow \mathbb{R}$ be any function normalized so that $\max_{\mathbf{z}} |h_*(\mathbf{z})| \leq 1/P$ and $\mathcal{S} = \{S : \hat{h}_*(S) \neq 0\}$ is an MSP set structure. Then for any $0 < \mu < 1$, there is a function $\tilde{h} : \{+1, -1\}^P \rightarrow \mathbb{R}$ such that*

$$\tilde{h}(\mathbf{z}) = h_*(\mathbf{z}) + \sum_{S \in \mathcal{S}} c_S \chi_S(\mathbf{z})$$

for some c_S such that $|c_S| \leq \mu$, and such that $f_*(\mathbf{x}) = \tilde{h}(\mathbf{z})$ can be learned by SGD to ϵ error in $n = d \cdot \text{poly}(2^{2^{O(P)}}, (1/\mu)^{2^{2^{O(P)}}}, 1/\epsilon)$ samples with a neural network of width $N = \text{poly}(2^{2^{2^{O(P)}}}, (1/\mu)^{2^{2^{O(P)}}}, 1/\epsilon)$ neurons.

Proof The proof is the same as the proof of Theorem 9. The main difference is that we lower-bound $\lambda_{\min}(\mathbf{K}^{k_1})$ explicitly and apply Lemmas 100 and 101 as in Proposition 83. Let $\alpha = (\alpha_S)_{S \in 2^{[P]}}$, where $\alpha_S = \mathbb{E}_{\mathbf{z}}[\tilde{h}(\mathbf{z})\chi_S(\mathbf{z})]$. Also consider activation function $\sigma(u) = \sum_{i=0}^L \frac{m_i}{i!} u^i$ for $L = 2^{8P}$. By Lemmas 42 and 43 we know that we can bound the minimum eigenvalue of the expected kernel $\lambda_{\min}(\mathbf{K}^{k_1})$ in terms of the coefficients of $\det(\mathbf{M}(\zeta, \alpha, \mathbf{m}))$, viewed as a polynomial in ζ . Here, $\mathbf{M}(\zeta, \xi, \rho)$ is the $2^P \times 2^P$ matrix given by the recurrence relations in Lemma 39 and the definition in (65). Notice that by construction, for any $\mathbf{z} \in \{+1, -1\}^P$ and $j \in [2^P]$ we have that

$$(L!)^{L+1} \cdot M_{\mathbf{z},j}(\zeta, \xi, \rho)$$

is a polynomial in ζ, ξ, ρ with integral coefficients and has degree at most $(L!)^{O(L+1)} = 2^{2^{O(P)}}$. Therefore the polynomial $N_{\mathbf{z},j}(\zeta, \phi, \rho)$ constructed in (67) is such that

$$(L!)^{(L+1)2^P} \det(\mathbf{N}(\zeta, \phi, \rho))$$

has integral coefficients in ζ, ϕ, ρ and degree at most $(L!)^{O((L+1)2^P)} = 2^{2^{O(P)}}$. Let $\kappa = \min(\mu, 1/L^2)$, and choose $\mathbf{m} \sim \text{Unif}[-\kappa, \kappa]^{\otimes L}$ and $c_S \sim \text{Unif}[-\kappa, \kappa]$ for each $S \in \mathcal{S}$. By the polynomial anti-concentration of Lemma 104 and a Markov bound, with probability at least 9/10 we must have that $\det(\mathbf{M}(\zeta, \alpha, \mathbf{m}))^2$ has some monomial of ζ with coefficient at least $\kappa^{2^{2^{O(P)}}}$. This implies by (66) that

$$\lambda_{\min}(\mathbf{K}^{k_1}) \geq (1/2^{2^{O(P)}}) \eta^{2^{2^{O(P)}}} \kappa^{2^{2^{O(P)}}}.$$

Taking learning rate $\eta = 2^{-\Theta(P)}$ which is small enough, we obtain the result that

$$\lambda_{\min}(\mathbf{K}^{k_1}) \geq \min(\mu^{2^{2^{O(P)}}}, 1/2^{2^{2^{O(P)}}}).$$

The result when we combine with Lemmas 100 and 101. ■

G.3. Merged-staircase functions when including all degree-1 monomials

We conjecture that the optimal dependence on P for learning vanilla staircase functions with unregularized SGD by two-layer neural networks should be on the order of $\exp(O(P))$, but the results of Propositions 83 and 85 have P dependence on the order of at least $\exp(\exp(\Omega(P)))$ and $\exp(\exp \exp(\Omega(P)))$, respectively. Therefore, we focus here on improving our understanding of the P dependence. We prove in Proposition 86 that SGD can succeed with $\exp(O(P))$ sample complexity dependence, but our result has two qualifications:

- **Stronger non-degeneracy assumption.** We assume that all of the *degree-1* terms of h_* – i.e., $\hat{h}_*({1}), \dots, \hat{h}_*({P})$ – are nonzero and are sufficiently distinct. This is in contrast to Propositions 83 and 85, which did not make this assumption.
- **Ad hoc activation function.** We use an ad hoc activation function. This is in contrast with Propositions 83 and 85, which apply to “most” activation functions.

Because of the stronger non-degeneracy assumption, the neural network can learn even when we only train the first layer for $k_1 = 1$ step and then train the second layer for a sufficiently large number of steps, k_2 .

G.3.1. THE BUMP AND GRADIENT BUMP FUNCTIONS

We define the σ_{bump} and σ_{gradbump} functions that are used to construct our ad hoc activation function. For any $\alpha < \beta$ and $\gamma < (\beta - \alpha)/2$, define the “bump” function:

$$\sigma_{\text{bump}}(u; \alpha, \beta, \gamma) = \begin{cases} 1, & u \in [\alpha + \gamma, \beta - \gamma] \\ 0, & u \notin [\alpha, \beta] \\ (6(u - \alpha)^2/\gamma^2 - 15(u - \alpha)/\gamma + 10)(u - \alpha)^3/\gamma^3, & u \in [\alpha, \alpha + \gamma] \\ (6(\beta - u)^2/\gamma^2 - 15(\beta - u)/\gamma + 10)(\beta - u)^3/\gamma^3, & u \in [\beta - \gamma, \beta]. \end{cases}$$

This function is twice-differentiable since $\sigma_{\text{bump}}(\alpha) = \sigma'_{\text{bump}}(\alpha) = \sigma''_{\text{bump}}(\alpha) = \sigma_{\text{bump}}(\beta) = \sigma'_{\text{bump}}(\beta) = \sigma''_{\text{bump}}(\beta) = 0$ and it also satisfies

$$\|\sigma_{\text{bump}}\|_{\infty} \leq 1, \quad \|\sigma'_{\text{bump}}\|_{\infty} \leq \frac{2}{\gamma}, \quad \text{and} \quad \|\sigma''_{\text{bump}}\|_{\infty} \leq \frac{6}{\gamma^2}.$$

If we take γ small relative to $\beta - \alpha$, then this function is effectively an indicator on the set $[\alpha, \beta]$. See Figure 5 for an example.

Also define the “gradient bump” function:

$$\sigma_{\text{gradbump}}(u; \kappa) = \begin{cases} u, & u \in [-1, 1] \\ -9u^5 - 68u^4 - 198u^3 - 276u^2 - 184u - 48, & u \in [-2, -1] \\ -9u^5 + 68u^4 - 198u^3 + 276u^2 - 184u + 48, & u \in [1, 2] \\ 0, & u \notin [-2, 2]. \end{cases}$$

This function is twice-differentiable since the pieces agree up to second derivative. Furthermore, it satisfies

$$\sigma_{\text{gradbump}}(0) = 0, \quad \sigma'_{\text{gradbump}}(0) = 1,$$

and it and its first and second derivatives are bounded by $O(1)$. We will use this function to ensure that our activation function has nonzero gradient at zero. See Figure 6.

THE MERGED-STAIRCASE PROPERTY




Figure 5: The bump function $\sigma_{\text{bump}}(u; \alpha, \beta, \gamma)$ with $\alpha = 2, \beta = 8, \gamma = 1$.




Figure 6: The gradient bump function $\sigma_{\text{gradbump}}(u)$.

G.3.2. STATEMENT OF PROPOSITION 86 AND COROLLARY 87

Proposition 86 *Let $h_* : \{+1, -1\}^P \rightarrow \mathbb{R}$, normalized for convenience so that $\max_{\mathbf{z}} |h_*(\mathbf{z})| \leq 1/P$. Define $\mathbf{c} = [\hat{h}(\{1\}), \dots, \hat{h}(\{P\})]$ be the vector of Fourier coefficients of degree 1, and suppose that \mathbf{c} satisfies the following conditions for some $0 < \mu < 1$.*

- *Bounded away from zero: for all $\mathbf{z} \in \{+1, -1\}^P$,*

$$|\langle \mathbf{c}, \mathbf{z} \rangle| \geq \mu. \quad (76)$$

- *Diverse: for any $\mathbf{z} \neq \mathbf{z}' \in \{+1, -1\}^P$,*

$$\frac{\langle \mathbf{c}, \mathbf{z} \rangle}{\langle \mathbf{c}, \mathbf{z}' \rangle} \notin (1 - \mu, 1 + \mu). \quad (77)$$

Then $f_(\mathbf{x}) = h_*(\mathbf{z})$ is SGD-learnable to any error $0 < \varepsilon < 1$ with probability at least $1 - \delta$ in $d \cdot \text{poly}(2^P \log(1/\delta)/(\mu\varepsilon))$ samples on a network with $N = \text{poly}(2^P \log(1/\delta)/(\mu\varepsilon))$ neurons. Furthermore, this SGD learnability is with initialization $\rho_0 = \text{Unif}[-1, 1] \otimes \delta_{\mathbf{0}}$ and activation function*

$$\sigma(t) = \sigma_{\text{gradbump}}(u) + \gamma^2 \sigma_{\text{bump}}(u; \alpha, \beta, \gamma),$$

where $\gamma = \mu/2^{8P}$, $\alpha = \mu/2^{4P}$, $\beta = \alpha + \mu/2^{4P}$.

The following result shows that the condition of Proposition 86 is true under a smoothed complexity model, where we perturb any function h_* slightly on its degree-1 Fourier coefficients.

Corollary 87 *For any $h : \{+1, -1\}^P \rightarrow \mathbb{R}$ with the normalization $\max_{\mathbf{z}} |h_*(\mathbf{z})| \leq 1/(2P)$ and any $0 < \mu' < 1$, there is a $\tilde{h} : \{+1, -1\}^P \rightarrow \mathbb{R}$ satisfying the conditions of Proposition 86 with $\mu = \mu'/2^{O(P)}$ and $\tilde{h}(\mathbf{z}) = h(\mathbf{z}) + \langle \boldsymbol{\delta}, \mathbf{z} \rangle$ and $\|\boldsymbol{\delta}\|_{\infty} \leq \mu'$.*

G.3.3. PROOF OF PROPOSITION 86 AND COROLLARY 87

Proof [Proof of Proposition 86] Initialize with $\rho_0 = \text{Unif}[-1, 1] \otimes \delta_{\mathbf{0}}$. Let the regularization parameter on the second layer be $\lambda = \varepsilon \cdot \text{poly}(\mu/2^P)$ for some large enough polynomial. Train the first layer with one step of (bSGD), i.e., take $k_1 = 1$ with learning rate $\eta = 1$. Train the second layer for $k_2 \geq \text{poly}(2^P \log(1/\delta)/(\varepsilon\mu))$ steps for a large enough polynomial. Let the number of neurons be $N \geq \text{poly}(2^P \log(1/\delta)/(\varepsilon\mu))$ and let the batch size be any $b \geq \text{poly}(2^P \log(1/\delta)/(\varepsilon\mu)) \log(N)d$ for large enough polynomials.

Computing 1 step of dynamics Let us analyze the (d-DF-PDE) dynamics for $k_1 = 1$ step. Since $\bar{\mathbf{u}}^0(a) = \mathbf{0}$, and $\sigma(0) = 0$, $\sigma'(0) = 1$, there is a simple formula for $\bar{\mathbf{u}}^1(a)$:

$$\begin{aligned} \bar{\mathbf{u}}^1(a) &= \bar{\mathbf{u}}^0(a) - a \mathbb{E}_{\mathbf{z}}[(\hat{f}_{\text{NN}}(\mathbf{z}; \bar{\rho}_0) - h_*(\mathbf{z}))\sigma'(\langle \bar{\mathbf{u}}^0(a), \mathbf{z} \rangle)\mathbf{z}] \\ &= a \mathbb{E}_{\mathbf{z}}[h_*(\mathbf{z})\sigma'(0)\mathbf{z}] \\ &= a \mathbb{E}_{\mathbf{z}}[h_*(\mathbf{z})\mathbf{z}] \\ &= a\mathbf{c}, \end{aligned}$$

where $\mathbf{c} = [\hat{h}_*(\{1\}), \dots, \hat{h}_*(\{P\})]$. By the choice of activation, and since $|a| \leq 1$ and $|\langle \mathbf{c}, \mathbf{z} \rangle| \leq 1$, we have

$$\sigma(\langle \bar{\mathbf{u}}^1(a), \mathbf{z} \rangle) = a\langle \mathbf{c}, \mathbf{z} \rangle + \gamma^2 \sigma_{\text{bump}}(a\langle \mathbf{c}, \mathbf{z} \rangle; \alpha, \beta, \gamma), \quad (78)$$

Definition of events $E_{z,j}$ and F_j Write $\mathbf{u}_j^k = \mathbf{u}(a_j^0)$ to denote the first-layer weight corresponding to neuron j . For any $\mathbf{z} \in \{+1, -1\}^P$ and $j \in [N]$, define the event

$$E_{z,j} = \{a_j^0 \langle \mathbf{c}, \mathbf{z} \rangle \in [\alpha + \gamma, \beta - \gamma]\}.$$

Notice that if event $E_{z,j}$ holds, then by (77), for any $\mathbf{z}' \neq \mathbf{z}$ we have

$$a_j^0 \langle \mathbf{c}, \mathbf{z}' \rangle \notin [\beta(1 - \mu), \alpha(1 + \mu)] \supset [\alpha, \beta].$$

So by (78) and the definition of σ_{bump} , under event $E_{z,j}$, we have

$$\sigma(\langle \bar{\mathbf{u}}_j^1, \mathbf{z} \rangle) = a_j^0 \langle \mathbf{c}, \mathbf{z} \rangle + \gamma^2, \quad \text{and} \quad \sigma(\langle \bar{\mathbf{u}}_j^1, \mathbf{z}' \rangle) = a_j^0 \langle \mathbf{c}, \mathbf{z}' \rangle \text{ for all } \mathbf{z} \neq \mathbf{z}'. \quad (79)$$

Also, for each $j \in [N]$ define the event F_j , which is

$$F_j = \{a_j^0 \geq 1/2\}$$

Since $\min |\langle \mathbf{c}, \mathbf{z} \rangle| > 2\alpha$ for all \mathbf{c} , under event F_j we have $a_j^0 \langle \mathbf{c}, \mathbf{z} \rangle \notin [\alpha, \beta]$. So by (78), under event F_j

$$\sigma(\langle \bar{\mathbf{u}}_j^1, \mathbf{z} \rangle) = a_j^0 \langle \mathbf{c}, \mathbf{z} \rangle \text{ for all } \mathbf{z}. \quad (80)$$

Lower-bounding the event probabilities Recall that $a_j^0 \sim \text{Unif}[-1, 1]$. By (76), we know that $\min_{\mathbf{z}} |\langle \mathbf{c}, \mathbf{z} \rangle| \geq \mu > \beta - \gamma$, so for any \mathbf{z}, j

$$\mathbb{P}[E_{z,j}] = \frac{\beta - \alpha - 2\gamma}{2|\langle \mathbf{c}, \mathbf{z} \rangle|} \geq \mu/2^{4P+1}. \quad (81)$$

And clearly since $a_j^0 \sim \text{Unif}[-1, 1]$,

$$\mathbb{P}[F_j] = 1/4. \quad (82)$$

Let E_{good} be the event that

$$E_{\text{good}} = \left\{ \frac{\min_{\mathbf{z}} |\{j : E_{z,j} \text{ holds}\}|}{N} \geq \mu/2^{4P+2} \right\} \cap \left\{ \frac{|\{j : F_j \text{ holds}\}|}{N} \geq 1/5 \right\}.$$

By a Hoeffding bound and a union bound

$$\mathbb{P}[E_{\text{good}}] \geq 1 - \delta,$$

as long as

$$N \geq 2^{\Omega(P)} \log(1/\delta).$$

Lower-bounding the empirical kernel eigenvalues Note that σ satisfies $\|\sigma\|_\infty, \|\sigma'\|_\infty, \|\sigma''\|_\infty \leq O(1)$. Therefore, by Lemma 101, the final loss is bounded in terms of $\lambda_{\min}(\mathbf{M}\mathbf{M}^\top)$, where \mathbf{M} is the feature matrix matrix at iteration $k_1 = 1$ with entries

$$\mathbf{M}_{z,j} = \frac{1}{\sqrt{N}} \sum_{j=1}^N \sigma(\langle \bar{\mathbf{u}}_j^{k_1}, \mathbf{z} \rangle).$$

Under event E_{good} , let us prove that the minimum eigenvalue of the empirical kernel $\mathbf{M}\mathbf{M}^\top$ is lower-bounded by $\lambda_{\min}(\mathbf{M}\mathbf{M}^\top) \geq \mu^5 2^{-100P}$. Consider any test function $\phi : \{+1, -1\}^P \rightarrow \mathbb{R}$. Suppose by contradiction that

$$\mathbb{E}_{z,z'}[\phi(\mathbf{z})\{\mathbf{M}\mathbf{M}^\top\}_{z,z'}\phi(\mathbf{z}')] \leq \mu^5 2^{-100P} \mathbb{E}_z[\phi(\mathbf{z})^2]. \quad (83)$$

Under event E_{good} , we have

$$\begin{aligned} \mathbb{E}_{z,z'}[\phi(\mathbf{z})\{\mathbf{M}\mathbf{M}^\top\}_{z,z'}\phi(\mathbf{z}')] &\geq \frac{1}{N} \sum_{j:F_j \text{ holds}} \mathbb{E}_{z,z'}[\phi(\mathbf{z})\sigma(\langle \bar{\mathbf{u}}_j^1, \mathbf{z} \rangle)\sigma(\langle \bar{\mathbf{u}}_j^1, \mathbf{z}' \rangle)\phi(\mathbf{z}')] \\ &= \frac{1}{N} \sum_{j:F_j \text{ holds}} \mathbb{E}_{z,z'}[\phi(\mathbf{z})(a_j^0)^2 \langle \mathbf{c}, \mathbf{z} \rangle \langle \mathbf{c}, \mathbf{z}' \rangle \phi(\mathbf{z}')] \\ &\geq \frac{1}{20} \mathbb{E}_z[\phi(\mathbf{z}) \langle \mathbf{c}, \mathbf{z} \rangle]^2. \end{aligned}$$

In order to avoid contradiction with (83) it follows that

$$\mathbb{E}_z[\phi(\mathbf{z}) \langle \mathbf{c}, \mathbf{z} \rangle]^2 \leq 20\mu^5 \cdot 2^{-100P} \mathbb{E}_z[\phi(\mathbf{z})^2] \quad (84)$$

Similarly, for any $\mathbf{z}'' \in \{+1, -1\}^P$

$$\begin{aligned} \mathbb{E}_{z,z'}[\phi(\mathbf{z})\{\mathbf{M}\mathbf{M}^\top\}_{z,z'}\phi(\mathbf{z}')] &\geq \frac{1}{N} \sum_{j:E_{z'',j} \text{ holds}} \mathbb{E}_{z,z'}[\phi(\mathbf{z})\sigma(\langle \bar{\mathbf{u}}_j^1, \mathbf{z} \rangle)\sigma(\langle \bar{\mathbf{u}}_j^1, \mathbf{z}' \rangle)\phi(\mathbf{z}')] \\ &\geq \frac{1}{N} \sum_{j:E_{z'',j} \text{ holds}} \mathbb{E}_{z,z'}[\phi(\mathbf{z})(a_j^0 \langle \mathbf{c}, \mathbf{z} \rangle + \gamma^2 1(z = \mathbf{z}''))(a_j^0 \langle \mathbf{c}, \mathbf{z}' \rangle + \gamma^2 1(\mathbf{z}' = \mathbf{z}''))\phi(\mathbf{z}')] \\ &\geq \frac{1}{N} \sum_{j:E_{z'',j} \text{ holds}} \gamma^4 \phi(\mathbf{z}'')\phi(\mathbf{z}'') - 2\gamma^2 |\mathbb{E}_z[\phi(\mathbf{z}) \langle \mathbf{c}, \mathbf{z} \rangle \phi(\mathbf{z}'')]| - \mathbb{E}_z[\phi(\mathbf{z}) \langle \mathbf{c}, \mathbf{z} \rangle]^2 \\ &\geq \mu 2^{-4P-2} (\gamma^4 \phi(\mathbf{z}'')\phi(\mathbf{z}'') - 2\gamma^2 |\mathbb{E}_z[\phi(\mathbf{z}) \langle \mathbf{c}, \mathbf{z} \rangle \phi(\mathbf{z}'')]| - \mathbb{E}_z[\phi(\mathbf{z}) \langle \mathbf{c}, \mathbf{z} \rangle]^2) \\ &\geq \mu 2^{-4P-2} (\gamma^4 \phi(\mathbf{z}'')^2 - 2\gamma^2 |\phi(\mathbf{z}'')| \sqrt{20 \cdot \mu^5 2^{-100P} \mathbb{E}_z[\phi(\mathbf{z})^2]} - 20 \cdot \mu^5 2^{-100P} \mathbb{E}_z[\phi(\mathbf{z})^2]), \end{aligned}$$

where in the last line we use (84). So in order to avoid contradiction with (83) we conclude that

$$\phi(\mathbf{z}'')^2 \leq \mu^5 \frac{2^{-40P}}{\mu \gamma^4} \mathbb{E}_z[\phi(\mathbf{z})^2].$$

However, since $\mu^5 2^{-40P} / (\mu \gamma^4) = 2^{-8P} < 1$, we get a contradiction by taking $\mathbf{z}'' = \arg \max_{\mathbf{z}} \phi(\mathbf{z})^2$. Therefore, we conclude that $\lambda_{\min}(\{\mathbf{M}\mathbf{M}^\top\}_{z,z'}) \geq \mu^5 2^{-100P}$. Plugging this into the guarantees for the linear regression of the last layer (Lemma 101) concludes the proof.

■

Proof [Proof of Corollary 87] We prove that such a δ exists by the probabilistic method. Let $\mathbf{c} = \mathbb{E}_z[h(\mathbf{z})\mathbf{z}]$ and let $\tilde{\mathbf{c}} = \mathbb{E}_z[\tilde{h}(\mathbf{z})\mathbf{z}] = \mathbf{c} + \delta$. If we take random $\delta \sim \text{Unif}([0, \mu'])^{\otimes P}$ then for any distinct pair \mathbf{z}, \mathbf{z}' such that without loss of generality $z_1 = 1 \neq -1 = z'_1$, we have

$$\begin{aligned} \mathbb{P}_\delta[|\langle \tilde{\mathbf{c}}, \mathbf{z} - \mathbf{z}' \rangle| \leq \mu' 2^{-3P}] &= \mathbb{P}_\delta[|\langle \mathbf{c}, \mathbf{z} - \mathbf{z}' \rangle + \langle \delta, \mathbf{z} - \mathbf{z}' \rangle| \leq \mu' 2^{-3P}] \\ &= \mathbb{P}_\delta[\delta(z_1 - z'_1) \in \langle \mathbf{c}, \mathbf{z}' - \mathbf{z} \rangle + \langle \delta, \mathbf{z}' - \mathbf{z} \rangle + [-\mu' 2^{-3P}, \mu' 2^{-3P}]] \\ &= \mathbb{P}_\delta[2\delta_1 \in \langle \mathbf{c}, \mathbf{z}' - \mathbf{z} \rangle + \langle \delta_{-1}, \mathbf{z}'_{-1} - \mathbf{z}_{-1} \rangle + [-\mu' 2^{-3P}, \mu' 2^{-3P}]] \\ &\leq 2^{-3P+1}. \end{aligned}$$

Similarly, for any \mathbf{z} ,

$$\mathbb{P}_\delta[|\langle \tilde{\mathbf{c}}, \mathbf{z} \rangle| \leq \mu' 2^{-3P}] \leq 2^{-3P+1}.$$

Therefore, taking a union bound over the $\binom{2^P}{2} < 2^{2P-1}$ distinct pairs \mathbf{z}, \mathbf{z}' , we have

$$\mathbb{P}_\delta[|\langle \tilde{\mathbf{c}}, \mathbf{z} \rangle| \geq \mu' 2^{-3P} \text{ and } |\langle \tilde{\mathbf{c}}, \mathbf{z} - \mathbf{z}' \rangle| \geq \mu' 2^{-3P} \text{ for all } \mathbf{z} \neq \mathbf{z}'] > 1 - (2^{2P-1} + 2^P)(2^{-3P+1}) \geq 0,$$

so by the probabilistic method a deterministic choice of δ satisfying $|\langle \tilde{\mathbf{c}}, \mathbf{z} \rangle| \geq \mu' 2^{-3P}$ and $|\langle \tilde{\mathbf{c}}, \mathbf{z} \rangle - \langle \tilde{\mathbf{c}}, \mathbf{z}' \rangle| \geq \mu' 2^{-3P}$ for all $\mathbf{z} \neq \mathbf{z}'$ exists. These conditions are sufficient to satisfy the conditions of Proposition 86 with $\mu = \mu' / 2^{O(P)}$. ■

Appendix H. Lower bounds on learning with linear methods

We recall the general definition for linear methods from Section 5 in the main text. Given a Hilbert space $(\mathcal{H}, \langle \cdot, \cdot \rangle_{\mathcal{H}})$, a feature map $\psi : \{+1, -1\}^d \rightarrow \mathcal{H}$, an empirical loss function $L : \mathbb{R}^{2n} \rightarrow \mathbb{R} \cup \infty$ and a regularization parameter $\lambda > 0$, a linear method construct from data points $(y_i, \mathbf{x}_i)_{i \in [n]}$ a prediction model $\hat{f}(\mathbf{x}) := \langle \hat{\mathbf{a}}, \psi(\mathbf{x}) \rangle_{\mathcal{H}}$ where $\hat{\mathbf{a}} \in \mathcal{H}$ is obtained by minimizing a regularized empirical risk functional

$$\hat{\mathbf{a}} = \arg \min_{\mathbf{a} \in \mathcal{H}} \left[L((y_i, \langle \mathbf{a}, \psi(\mathbf{x}_i) \rangle)_{i \in [n]}) + \lambda \|\mathbf{a}\|_{\mathcal{H}}^2 \right]. \quad (85)$$

Recall that we denote $q = \dim(\mathcal{H})$.

Example 1 Popular examples of linear methods include

- (a) *Random Feature models:* take $(\mathbf{w}_i)_{i \in [N]} \sim_{iid} \tau$, $\mathbf{w}_i \in \mathcal{V}$, and an activation $\phi : \mathcal{X} \times \mathcal{V} \rightarrow \mathbb{R}$, then the Hilbert space and the feature map are defined by $\mathcal{H} = \text{span}\{\phi(\cdot; \mathbf{w}_i) : i \in [N]\}$ and $\psi(\mathbf{x}) = (\phi(\mathbf{x}; \mathbf{w}_1), \dots, \phi(\mathbf{x}; \mathbf{w}_N))$. For generic examples, we have $q = \dim(\mathcal{H}) = N$ almost surely.
- (b) *Kernel methods:* take \mathcal{H} a reproducing kernel Hilbert space (RKHS) with reproducing kernel $K : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$. There exists a Hilbert space $(\mathcal{F}, \langle \cdot, \cdot \rangle_{\mathcal{F}})$ (the feature space) and a feature map $\psi : \mathcal{X} \rightarrow \mathcal{F}$ such that $K(\mathbf{x}_1, \mathbf{x}_2) = \langle \psi(\mathbf{x}_1), \psi(\mathbf{x}_2) \rangle_{\mathcal{F}}$ and $\mathcal{H} = \{\langle \mathbf{a}, \psi(\cdot) \rangle_{\mathcal{F}} : \mathbf{a} \in \mathcal{F}\}$. We have typically $q = \dim(\mathcal{H}) = \infty$.

Ridge regression corresponds to taking the functional: $L((y_i, \hat{f}_i)_{i \in [n]}) = \frac{1}{n} \sum_{i \in [n]} (y_i - \hat{f}_i)^2$.

We will be interested in providing lower bounds on the number of samples necessary to learn some classes of functions for any linear methods. We first present the following general dimension-based (see discussion below) approximation lower bound that is a slight variation of [Hsu et al. \(2021\)](#); [Hsu; Kamath et al. \(2020\)](#); it improves on [Hsu et al. \(2021\)](#); [Hsu](#) for target functions that are not (almost) orthogonal, and it uses the operator norm of the gram matrix rather than its min-eigenvalue as in [Kamath et al. \(2020\)](#).

Proposition 88 (Dimension lower bound) *Let \mathcal{R} be a Hilbert space with inner product denoted by $\langle \cdot, \cdot \rangle_{\mathcal{R}}$. Fix $\mathcal{F} = \{f_1, \dots, f_M\} \subset \mathcal{R}$ a set of target functions with $\|f_i\|_{\mathcal{R}}^2 = \langle f_i, f_i \rangle_{\mathcal{R}} = 1$ for all $i \in [M]$. Let \mathcal{T} be a (potentially random) finite-dimensional subspace of \mathcal{R} , with $r = \mathbb{E}_{\mathcal{T}}[\dim(\mathcal{T})] < \infty$.*

Define the average¹⁰ approximation error of the target functions \mathcal{F} by the subspace \mathcal{T}

$$\varepsilon := \frac{1}{M} \sum_{i \in [M]} \mathbb{E}_{\mathcal{T}} \left[\inf_{g \in \mathcal{T}} \|g - f_i\|_{\mathcal{R}}^2 \right],$$

and $\mathbf{G} = (\langle f_i, f_j \rangle_{\mathcal{R}})_{i,j \in [M]}$ the Gram matrix associated to the f_i 's. Then

$$r \geq \frac{M}{\|\mathbf{G}\|_{\text{op}}} (1 - \varepsilon). \quad (86)$$

10. This is a lower-bound on the worst-case approximation error considered in [Kamath et al. \(2020\)](#).

Note that the results in [Hsu et al. \(2021\)](#); [Hsu](#) are simply obtained by using

$$\|\mathbf{G}\|_{\text{op}} \leq \|\mathbf{I}\|_{\text{op}} + \|\mathbf{G} - \mathbf{I}\|_{\text{op}} \leq 1 + \|\mathbf{G} - \mathbf{I}\|_F = 1 + \sqrt{\sum_{i \neq j} \langle f_i, f_j \rangle_{\mathcal{R}}^2}.$$

Proof [Proof of Proposition 88] In the proofs in [Hsu et al. \(2021\)](#); [Hsu](#), we simply replace the Boas-Bellman inequality by (for any $g \in \mathcal{R}$)

$$\begin{aligned} \sum_{i \in [M]} \langle g, f_i \rangle_{\mathcal{R}}^2 &= \left\langle g, \sum_{i \in [M]} \langle g, f_i \rangle_{\mathcal{R}} f_i \right\rangle_{\mathcal{R}} \\ &\leq \|g\|_{\mathcal{R}} \left(\sum_{ij \in [M]} \langle g, f_i \rangle_{\mathcal{R}} \langle g, f_j \rangle_{\mathcal{R}} \langle f_i, f_j \rangle_{\mathcal{R}} \right)^{1/2} \\ &= \|g\|_{\mathcal{R}} (\mathbf{b}^T \mathbf{G} \mathbf{b})^{1/2} \leq \|g\|_{\mathcal{R}} \|\mathbf{G}\|_{\text{op}}^{1/2} \|\mathbf{b}\|_2, \end{aligned} \quad (87)$$

where we denoted $\mathbf{b} = (\langle g, f_1 \rangle_{\mathcal{R}}, \dots, \langle g, f_M \rangle_{\mathcal{R}})$. Noticing that $\|\mathbf{b}\|_2^2$ is equal to the left-hand side of Eq. (87), we get

$$\sum_{i \in [M]} \langle g, f_i \rangle_{\mathcal{R}}^2 \leq \|\mathbf{G}\|_{\text{op}} \|g\|_{\mathcal{R}}^2,$$

which together with $\|g\|_{\mathcal{R}} \leq \frac{1}{M} \sum_{i=1}^M \|f_i\|_{\mathcal{R}} \leq 1$ yields the improved bound (86). \blacksquare

Let us explain how to derive lower-bounds on the performance of linear methods using Proposition 88. Consider $\mathcal{R} = L^2(\mathcal{X})$ and \mathcal{T} the space of functions $f = \langle \mathbf{a}, \psi(\cdot) \rangle_{\mathcal{F}}$ with $\mathbf{a} \in \text{span}\{\psi(\mathbf{x}_i) : i \in [n]\}$. We can consider \mathcal{T} random or fixed conditional on ψ (e.g., random feature map) and the \mathbf{x}_i 's. We always have $r \leq \min(p, n)$. Consider learning a set of M functions $\mathcal{F} = \{f_1, \dots, f_M\}$ with the linear estimator obtained by (85). From the above discussion, we must have that the estimator $\hat{f} \in \mathcal{T}$ and the generalization error is lower bounded by the approximation error $\|f_i - \hat{f}\|_{L^2}^2 \geq \inf_{g \in \mathcal{T}} \|f_i - g\|_{L^2}^2$. Therefore ε lower bound the average generalization error over learning \mathcal{F} . Therefore, Proposition 88 implies the following: if the average generalization error over \mathcal{F} is less than ε , then we must have

$$\min(n, q) \geq \frac{M}{\|\mathbf{G}\|_{\text{op}}} (1 - \varepsilon). \quad (88)$$

This bound is a dimension lower bound in the sense that it does not assume anything about the statistical model (e.g., the \mathbf{x}_i can be arbitrary and do not have to be independent), only that the estimator lies in a $\min(n, q)$ -dimensional subspace \mathcal{T} : this subspace can be a good approximation of M orthogonal functions only if $\min(n, q) \geq \Omega(M)$.

To get Proposition 11 in the main text, we make the following two modifications of the bound in Proposition 88. In Eq. (88), we upper bound $\|\mathbf{G}\|_{\text{op}} \leq \|\mathbf{G}\|_{1, \infty} = \max_{i \in [M]} \sum_{j \in [M]} |\langle f_i, f_j \rangle_{\mathcal{R}}|$. Second, some linear subspaces $\Omega \subseteq \mathcal{R}$ are harder to fit for linear methods (see for example [Ghorbani et al. \(2021a\)](#); [Mei et al. \(2021\)](#)). For instance, vanilla staircase functions of large degree contain monomials of large degree that have a large dimension lower-bound, but the overall staircase functions do not have a large dimension lower-bound per se. We next present a corollary that applies to any decomposition $\mathcal{R} = \Omega \oplus \Omega^\perp$, and distinguishes the error incurred on each of the two orthogonal subspaces. Denote \mathbf{P}_Ω and $\mathbf{P}_{\Omega^\perp} = \mathbf{I} - \mathbf{P}_\Omega$ the orthogonal projections onto Ω and Ω^\perp respectively.

Corollary 89 *Let \mathcal{R} be a Hilbert space with inner product denoted by $\langle \cdot, \cdot \rangle_{\mathcal{R}}$ and Ω a linear subspace of \mathcal{R} . Fix $\mathcal{F} = \{f_1, \dots, f_M\} \subset \mathcal{R}$ a set of target functions with $\|\mathbf{P}_{\Omega} f\|_{\mathcal{R}}^2 = 1$ for all $f \in \mathcal{F}$. Define $\mathbb{E}_{\mathcal{F}}$ the expectation over $f \sim \text{Unif}(\mathcal{F})$. Let \mathcal{T} be a (potentially random) finite-dimensional subspace of \mathcal{R} , with $r = \mathbb{E}_{\mathcal{T}}[\dim(\mathcal{T} \cap \Omega)] < \infty$.*

Define the average approximation error on Ω of the target functions \mathcal{F} by the subspace \mathcal{T}

$$\varepsilon := \frac{1}{M} \sum_{i \in [M]} \mathbb{E}_{\mathcal{T}} \left[\inf_{g \in \mathcal{T}} \|\mathbf{P}_{\Omega}(g - f_i)\|_{\mathcal{R}}^2 \right],$$

Then

$$r \geq \frac{1 - \varepsilon}{\max_{i \in [M]} \frac{1}{M} \sum_{j \in [M]} |\langle f_i, \mathbf{P}_{\Omega} f_j \rangle_{\mathcal{R}}|}. \quad (89)$$

This is a direct consequence of Proposition 88 with \mathcal{R} and \mathcal{T} replaced by Ω and $\mathcal{T} \cap \Omega$, and the target functions by $\mathcal{F}' = \{\mathbf{P}_{\Omega} f_1, \dots, \mathbf{P}_{\Omega} f_M\}$. Proposition 11 in the main text is simply Corollary 89 rewritten in the context of linear methods.

Consider a set of target functions \mathcal{F} such that $\|\mathbf{P}_{\Omega} f\|_{L^2}^2 = 1 - \kappa$ and $\|(\mathbf{I} - \mathbf{P}_{\Omega})f\|_{L^2}^2 = \kappa$ for any $f \in \mathcal{F}$. If the averaged generalization error is less than $1 - u$, we can take $\varepsilon = (1 - u)/(1 - \kappa)$ and get

$$\min(n, q) \geq \frac{u - \kappa}{\max_{i \in [M]} \frac{1}{M} \sum_{j \in [M]} |\langle f_i, \mathbf{P}_{\Omega} f_j \rangle_{\mathcal{R}}|}. \quad (90)$$

Let us apply this bound to the examples described in the main text. We take $\mathcal{X} = \{+1, -1\}^d$. First consider Ω the span of all degree k monomials and a target function f_* such that $\|\mathbf{P}_{\Omega} f_*\|_{L^2}^2 = \frac{1}{2}$, and $\mathbf{P}_{\Omega} f_*$ is supported on m monomials $\{S_1, \dots, S_m\}$, with $S_i \subseteq [d]$, $|S_i| = k$:

$$\mathbf{P}_{\Omega} f_*(\mathbf{x}) = \sum_{i \in [m]} \alpha_{S_i} \chi_{S_i}(\mathbf{x}).$$

We consider the class of functions

$$\mathcal{F}_* = \left\{ f_*(\tau(\mathbf{x})) : \tau \in \Pi(d) \right\}.$$

\mathcal{F}_* is the smallest class of functions containing f_* that is invariant under a permutation of the input coordinates. The generalization error $\mathbb{E}_{\mathcal{F}_*} [\|f - \hat{f}\|_{L^2}^2]$ corresponds to the test error with uniform prior distribution over all permutation of the input space. Note that any method that is equivariant with respect to permutations (e.g., kernel methods with inner-product kernel) will have the same generalization error $\mathbb{E}_{\mathcal{T}} [\|f \circ \tau - \hat{f}\|_{L^2}^2]$ for any $\tau \in \Pi(d)$.

Applying Eq. (90), we obtain the following lower bound:

Proposition 90 *For any linear method, in order to get an average generalization error over \mathcal{F}_* that is smaller than $1/2 \cdot (1 - \eta)$, we must have*

$$\min(n, q) \geq \frac{\eta}{m} \binom{d}{k}.$$

Proof [Proof of Proposition 90] Fix $\tau_1 \in \Pi(d)$. We have

$$\begin{aligned} \mathbb{E}_{\tau_2} [|\langle f_* \circ \tau_1, \mathbf{P}_\Omega f_* \circ \tau_2 \rangle_{L^2}|] &= \sum_{ij \in [m]} \alpha_{S_i} \alpha_{S_j} \mathbb{E}_{\tau_2} [\langle \chi_{S_i} \circ \tau_1, \chi_{S_j} \circ \tau_2 \rangle_{L^2}] \\ &= \sum_{ij \in [m]} \alpha_{S_i} \alpha_{S_j} \cdot \frac{k!(d-k)!}{d!} \\ &\leq m \frac{k!(d-k)!}{d!} \sum_{i \in [m]} \alpha_{S_i}^2 = \frac{m k!(d-k)!}{2 d!}. \end{aligned}$$

We can then apply Eq. (90) with $u := 1/2 \cdot (1 + \eta)$ and $\kappa = 1/2$. \blacksquare

Proposition 90 shows that for k fixed, $n = \Omega_d(d^k/m)$ samples are necessary to learn \mathcal{F}_* . As a second example, consider the vanilla staircase function of degree P :

$$f_P(\mathbf{x}) = \frac{1}{\sqrt{P}} \sum_{i=1}^P x_1 \cdots x_i,$$

and the function class of all staircase functions of degree P :

$$\mathcal{F}_P = \left\{ f_P(\tau(\mathbf{x})) : \tau \in \Pi(d) \right\}.$$

Proposition 91 *Let $P \leq d/2$ and $\eta \in (0, 1)$. For any linear method, in order to get an average generalization error over \mathcal{F}_P that is smaller than $1 - \eta$, we must have*

$$\min(n, q) \geq \frac{\eta}{2} \binom{d}{\lfloor \frac{\eta P}{2} \rfloor}.$$

Proof [Proof of Proposition 91] Denote now $S_i = \{1, 2, \dots, i\}$ and \mathbf{P}_ℓ the projection on every monomials of degree at least ℓ . Notice that $\|\mathbf{P}_\ell f_P\|_{L^2}^2 = 1 - \frac{\ell-1}{P}$ and $\|(\mathbf{I} - \mathbf{P}_\ell) f_P\|_{L^2}^2 = \frac{\ell-1}{P}$.

Fix $\tau_1 \in \Pi(d)$. We have

$$\begin{aligned} \mathbb{E}_{\tau_2} [|\langle f_P \circ \tau_1, \mathbf{P}_\ell f_P \circ \tau_2 \rangle_{L^2}|] &= \frac{1}{P} \sum_{i=\ell}^P \mathbb{E}_{\tau_2} [\langle \chi_{S_i} \circ \tau_1, \chi_{S_i} \circ \tau_2 \rangle_{L^2}] \\ &= \frac{1}{P} \sum_{i=\ell}^P \frac{i!(d-i)!}{d!} \leq \frac{1}{\binom{d}{\ell}}. \end{aligned}$$

The proposition follows by taking $\ell = \lfloor \frac{\eta P}{2} \rfloor$ and applying Eq. (90) with \mathbf{P}_ℓ , $u = \eta$ and $\kappa = \frac{\ell-1}{P}$. \blacksquare

In our case, we are interested in $P = \omega_d(1)$. Letting η decay at moderate rate, such as $\eta = 1/\sqrt{P}$ in Proposition 91, we get the following superpolynomial lower bound on the number of samples $n \geq d^{\omega_d(1)}$.

Appendix I. Technical results

In this appendix, we gather a few technical results needed to prove the main results in this paper.

I.1. Bound between batch-SGD and discrete mean-field dynamics

While the results in [Mei et al. \(2019\)](#) are written for continuous-time dynamics, we note that their proof can be easily adapted to the discrete-time regime, as described in [Appendix C](#). More precisely, following the notations in [Mei et al. \(2019\)](#), we compare the solution $(\Theta^k)_{k \in \mathbb{N}}$ of batch-SGD:

$$\theta_j^{k+1} = \theta_j^k + \frac{1}{b} \sum_{i \in [b]} \{y_{ki} - \hat{f}_{\text{NN}}(\mathbf{x}_{ki}; \Theta^k)\} \cdot \mathbf{H}_k \nabla_{\theta} \sigma_*(\mathbf{x}_{ki}; \theta) - \mathbf{H}_k \Lambda \theta_j^k, \quad (91)$$

to the solution $(\rho_k)_{k \in \mathbb{N}}$ of the discrete mean-field dynamics:

$$\theta^{k+1} = \theta^k + E_{(y, \mathbf{x})} [\{y - \hat{f}_{\text{NN}}(\mathbf{x}; \rho_k)\} \mathbf{H}_k \nabla_{\theta} \sigma_*(\mathbf{x}; \theta)] - \mathbf{H}_k \Lambda \theta^k. \quad (92)$$

We consider the same assumptions as ([Mei et al., 2019](#), Theorem 1), with the difference that A1 is replaced by A1' : $\|\Lambda\|_{\text{op}}, \|\mathbf{H}_k\|_{\text{op}} \leq K$.

Proposition 92 (Discrete propagation-of-chaos) *Assume that conditions A1', A2-A4 in [Mei et al. \(2019\)](#) hold and let $k_0 \in \mathbb{N}$. There exists a constant K depending only on the constants in the assumptions (in particular independent of d, k_0) such that:*

(A) Fixed second-layer:

$$\sup_{k=0, \dots, k_0} \|\hat{f}_{\text{NN}}(\cdot; \Theta^k) - \hat{f}_{\text{NN}}(\cdot; \rho_k)\|_{L^2} \leq K e^{K k_0} \left\{ \frac{\sqrt{\log N} + z}{\sqrt{N}} + \frac{\sqrt{d + \log N} + z}{\sqrt{b}} \right\},$$

with probability at least $1 - e^{-z^2}$.

(B) Training both layers:

$$\sup_{k=0, \dots, k_0} \|\hat{f}_{\text{NN}}(\cdot; \Theta^k) - \hat{f}_{\text{NN}}(\cdot; \rho_k)\|_{L^2} \leq K e^{K k_0} \left\{ \frac{\sqrt{\log N} + z}{\sqrt{N}} + \frac{\sqrt{d + \log N} + z}{\sqrt{b}} \right\},$$

with probability at least $1 - e^{-z^2}$.

Proof [Proof of [Proposition 92](#)] The proof of this proposition follows by adapting the proof of ([Mei et al., 2019](#), Theorem 1) to the discrete setting described above (see also [Appendix B.3](#)). In particular, part (A) (fixed second layer coefficients) follows from the [Appendix B](#) in [Mei et al. \(2019\)](#): the comparison between discrete and continuous gradient is not needed anymore, and the only difference is in the first part of [Proposition 16](#) in [Mei et al. \(2019\)](#), which can simply be rewritten by noting that

$$\begin{aligned} & \|\bar{\theta}_i^{k+1} - \underline{\theta}_i^{k+1}\|_2^2 \\ & \leq |\langle \bar{\theta}_i^{k+1} - \underline{\theta}_i^{k+1}, \bar{\theta}_i^k - \underline{\theta}_i^k \rangle| + |\langle \bar{\theta}_i^{k+1} - \underline{\theta}_i^{k+1}, (\bar{\theta}_i^{k+1} - \bar{\theta}_i^k) - (\underline{\theta}_i^{k+1} - \underline{\theta}_i^k) \rangle| \\ & \leq \|\bar{\theta}_i^{k+1} - \underline{\theta}_i^{k+1}\|_2 \|\bar{\theta}_i^k - \underline{\theta}_i^k\|_2 + |\langle \bar{\theta}_i^{k+1} - \underline{\theta}_i^{k+1}, (\bar{\theta}_i^{k+1} - \bar{\theta}_i^k) - (\underline{\theta}_i^{k+1} - \underline{\theta}_i^k) \rangle|, \end{aligned}$$

and the rest of the proof follows similarly.

For part (B), the main difference comes from bounding a^k : we have

$$|a^{k+1}| \leq |a^k| + \mathbb{E}_{\mathbf{x}} [|\hat{g}_k(\mathbf{x})| |\sigma(\mathbf{x}; \mathbf{w}^k)|] \leq |a^k| + K + K \|a^k\|_{\infty},$$

where we denoted $\|a^k\|_{\infty} = \sup_{a \in \text{supp}(\rho_k)} |a|$ and $\hat{g}_k(\mathbf{x}) = f_*(\mathbf{x}) - \hat{f}_{\text{NN}}(\mathbf{z}; \rho_k)$ and used that by assumption $\|f_*\|_{\infty}, \|\sigma\|_{\infty} \leq K$ and $\|\hat{f}_{\text{NN}}\|_{\infty} \leq \int |a| \|\sigma\|_{\infty} \rho_k(d\boldsymbol{\theta}) \leq K$. We can then use the discrete Grönwall inequality to get $\|a^t\|_{\infty} \leq K e^{K k_0}$. This explains the worse dependency (double exponential) in k_0 in the bound, than for continuous time, where one can use properties of continuous gradient flows to get a bound on a^t linear in time. With this modification, the rest of the proof follow by adapting Appendix C in [Mei et al. \(2019\)](#), where we can assume that the activation function $\sigma_*(\mathbf{x}; \boldsymbol{\theta})$ is bounded by $K e^{K k_0}$. ■

I.2. Discrete-time analysis of SGD on second layer

In this section, we provide a tighter analysis for the training of the second-layer weights by SGD using the bias-variance decomposition, instead of a propagation-of-chaos argument. This technical tool is used in Section G to provide tighter sample-complexity bounds. Indeed, the second-layer weights are trained during a time $T = \log(1/\varepsilon)/\lambda_{\min}$, where λ_{\min} is the minimum eigenvalue of the kernel matrix associated to the neural network after training the first layer weights (see for example Section D.1). If we naively use the propagation-of-chaos comparison between the dimension-free dynamics and batch-SGD (bSGD) in Theorem 5, we need

$$e^{K_1 T^7} \left\{ \sqrt{\frac{P + \log(d)}{d}} + \sqrt{\frac{\log N}{N}} + \sqrt{\frac{d + \log N}{b}} \sqrt{\eta} \right\} \leq \varepsilon,$$

and therefore we need the sample size to scale as $n = d e^{\log(1/\varepsilon)^7 / \lambda_{\min}^7}$. In this section, we consider instead a direct analysis of (bSGD) on the second layer weights and we obtain the tighter requirement $n = d / (\varepsilon \lambda_{\min})^C$ for some constant $C > 0$.

Training setup: We assume that the weights of the first layer are fixed at some $\mathbf{W}^0 \in \mathbb{R}^{N \times d}$. We train the weights \mathbf{a} of the second layer and obtain a discrete-time dynamics \mathbf{a}^k on these weights. Namely, given initial weights $\mathbf{W}^0 \in \mathbb{R}^{N \times d}$ and $\mathbf{a}^0 \in \mathbb{R}^N$, we train \mathbf{a}^k with the batch-SGD dynamics (bSGD), with the step size $\eta > 0$ and regularization parameter $\lambda^a = \lambda > 0$ on \mathbf{a} .

$$\begin{aligned} \mathbf{w}_j^{k+1} &= \mathbf{w}_j^0, \\ a_j^{k+1} &= a_j^k + \eta \left[\frac{1}{b} \sum_{i \in [b]} \{y_{ki} - \hat{f}_{\text{NN}}(\mathbf{x}_{ki}; \boldsymbol{\Theta}^k)\} \sigma(\langle \mathbf{w}_j^k, \mathbf{x}_{ki} \rangle) \mathbf{x}_{ki} \right] - \eta \lambda a_j^k, \end{aligned}$$

where at each time step k we take fresh i.i.d. data samples $\{(y_{ki}, \mathbf{x}_{ki})\}_{i \in [b]}$. In the analysis, it is convenient to define the feature map which depends on \mathbf{W}^0 as

$$\phi(\mathbf{x}) = \frac{1}{\sqrt{N}} [\sigma(\langle \mathbf{w}_1^0, \mathbf{x} \rangle) \dots \sigma(\langle \mathbf{w}_N^0, \mathbf{x} \rangle)] \in \mathbb{R}^N.$$

With this notation, the neural network function while only the second layer is being trained can be written as:

$$\hat{f}_{\text{NN}}(\mathbf{x}; \mathbf{a}, \mathbf{W}^0) = \frac{1}{\sqrt{N}} \langle \mathbf{a}, \phi(\mathbf{x}) \rangle,$$

This is simply a linear function in \mathbf{a} with a fixed feature map ϕ depending on \mathbf{W}^0 . Therefore, the training of \mathbf{a} can be studied with classical ideas for analyzing linear methods. Namely, define the regularized loss¹¹:

$$\ell(\mathbf{a}) = \frac{1}{2} \mathbb{E}_{\mathbf{x}} \left[(h_*(z) - \frac{1}{\sqrt{N}} \langle \mathbf{a}, \phi(\mathbf{x}) \rangle)^2 \right] + \frac{\lambda}{2N} \|\mathbf{a}\|^2.$$

The main result of this subsection is that if there exists a low-norm “certificate” \mathbf{a}^{cert} that achieves small loss, then batch-SGD will achieve a loss that is approximately upper-bounded by the loss at \mathbf{a}^{cert} , after a short number of iterations. This is a key ingredient in our more quantitative bounds on the sample complexity in Section G.

Lemma 93 (Suffices to prove existence of certificate) *Let $B_1, B_2 > 0$ be such that $|y| \leq \sqrt{B_1}$ and $\|\phi(\mathbf{x})\|^2 \leq B_2$ almost surely. For any step size $0 < \eta \leq 1/(B_2 + \lambda)$, any $\delta > 0$, and any $\mathbf{a}^{\text{cert}} \in \mathbb{R}^N$, and any time step k , with probability at least $1 - \delta$*

$$\begin{aligned} \ell(\mathbf{a}^k) \leq & \ell(\mathbf{a}^{\text{cert}}) + 2(B_2 + \lambda) \left((1 - \eta\lambda)^{2k} \left(\frac{\|\mathbf{a}^0\|^2}{N} + \frac{2\ell(\mathbf{a}^{\text{cert}})}{\lambda} \right) \right. \\ & \left. + \left(\frac{9B_1B_2}{\lambda^2 b N} + \frac{18(B_2 + \lambda)^2 \ell(\mathbf{a}^{\text{cert}})}{\lambda^3 b} \right) \log(1/\delta) \right). \end{aligned}$$

Remark 94 (Relation of Lemma 93 to prior work) *The proof of this lemma relies on the well-known bias-variance decomposition idea in for analyzing least-squares regression with SGD (e.g., Bach and Moulines (2013); Jain et al. (2017)). However, we wish to prove a statement about the final iterate of SGD, and most works analyze averaged iterates rather than the final iterate. Standard bounds for last-iterate SGD (e.g., Shamir and Zhang (2013); Jain et al. (2019)) do not apply because they assume Lipschitzness of the loss function, which cannot be assumed because we use the squared loss. Furthermore, the bound of Jain et al. (2018) for last-iterate batch-SGD holds in expectation, rather than with high probability. Nevertheless, our proof is a straightforward modification of Jain et al. (2018), making stronger assumptions and obtaining a suboptimal rate in order to get a simpler proof.*

We first prove the following lemma, where \mathbf{a}^* is the minimizer of ℓ , guaranteed to be unique by strong convexity:

Lemma 95 *Under the same conditions as Lemma 93, for any time step $k > 0$, with probability at least $1 - \delta$ the following bound holds:*

$$\|\mathbf{a}^k - \mathbf{a}^*\| \leq (1 - \eta\lambda)^k \|\mathbf{a}^0 - \mathbf{a}^*\| + \frac{3(\sqrt{B_1 B_2} + (B_2 + \lambda) \|\mathbf{a}^*\|) \sqrt{\log(1/\delta)}}{\lambda \sqrt{b}}$$

11. The $1/N$ scaling in the regularization is needed because each neuron is scaled as $1/N$ in the expression for \hat{f}_{NN} .

Proof

Define $\mathbf{H} = \frac{1}{N}\mathbb{E}[\phi(\mathbf{x}) \otimes \phi(\mathbf{x})] + \frac{\lambda}{N}\mathbf{I}$ and $\mathbf{v} = \frac{1}{\sqrt{N}}\mathbb{E}[\phi(\mathbf{x})y]$. The loss at \mathbf{a} can be rewritten as:

$$\ell(\mathbf{a}) = \frac{1}{2}\mathbb{E}[y^2] - \langle \mathbf{v}, \mathbf{a} \rangle + \frac{1}{2}\langle \mathbf{a} \otimes \mathbf{a}, \mathbf{H} \rangle.$$

Define the gap to optimality $\mathbf{w}^k = \mathbf{a}^k - \mathbf{a}^*$. We track the evolution of \mathbf{w}^k . Let $\mathbf{P}^k = \mathbf{I} - \eta(\frac{1}{b}\sum_{i=1}^b \phi(\mathbf{x}_{ki}) \otimes \phi(\mathbf{x}_{ki}) + \lambda\mathbf{I})$. Note that $0 \preceq \mathbf{P}^k \preceq (1 - \eta\lambda\mathbf{I})$, since $\|\phi(\mathbf{x}_{ki})\|^2 \leq B_2$, and $\eta \leq 1/(B_2 + \lambda)$. We have

$$\begin{aligned} \mathbf{w}^{k+1} &= \mathbf{P}^k \mathbf{a}^k + \frac{\eta}{b} \sum_{i=1}^b y_{ki} \phi(\mathbf{x}_{ki}) - \mathbf{a}^* \\ &= \mathbf{P}^k \mathbf{w}^k - \frac{\eta}{b} \sum_{i=1}^b y_{ki} \phi(\mathbf{x}_{ki}) + \eta \left(\frac{1}{b} \sum_{i=1}^b \phi(\mathbf{x}_{ki}) \otimes \phi(\mathbf{x}_{ki}) + \lambda\mathbf{I} \right) \mathbf{a}^* \\ &= \mathbf{P}^k \mathbf{w}^k + \eta \boldsymbol{\zeta}^k, \end{aligned}$$

where for any $k \geq 0$, we have the noise vector

$$\boldsymbol{\zeta}^k = \frac{1}{b} \boldsymbol{\xi}^{ki} \text{ and } \boldsymbol{\xi}^{ki} = y_{ki} \phi(\mathbf{x}_{ki}) - (\phi(\mathbf{x}_{ki}) \otimes \phi(\mathbf{x}_{ki}) + \lambda\mathbf{I}) \mathbf{a}^*$$

Recursively expanding this, we have the well-known ‘‘bias-variance’’ decomposition

$$\mathbf{w}^k = \left(\prod_{l=0}^{k-1} \mathbf{P}^l \right) \mathbf{w}^0 + \eta \sum_{j=0}^{k-1} \left(\prod_{l=1}^j \mathbf{P}^l \right) \boldsymbol{\zeta}^j.$$

The first term tracks how close \mathbf{a}^t would be to \mathbf{a}^0 if there were no noise, and the second term controls how much error the noise in the batch-SGD contributes if we had started at the optimal solution \mathbf{a}^* .

By the spectral norm bound on \mathbf{P}^l and the triangle inequality, we have

$$\|\mathbf{w}^k\| \leq \underbrace{(1 - \eta\lambda)^k}_{\text{(Term 1)}} \|\mathbf{w}^0\| + \eta \underbrace{\sum_{j=0}^{k-1} (1 - \eta\lambda)^{t-1-j}}_{\text{(Term 2)}} \|\boldsymbol{\zeta}^j\|. \quad (93)$$

The first term is already essentially in the form that we want. Let us bound the second term:

Claim 96 For any time step j , and $i \in [b]$, $\|\boldsymbol{\xi}^{ji}\| \leq B_3 := \sqrt{B_1 B_2} + (B_2 + \lambda)\|\mathbf{a}^*\|$ almost surely.

Proof By triangle inequality and the almost-sure bounds on y_{ji} and $\phi(\mathbf{x}_{ji})$. ■

Claim 97 For any time-step j , $\mathbb{E}[\|\boldsymbol{\zeta}^j\|] \leq B_3/\sqrt{b}$.

Proof By the first-order optimality conditions on \mathbf{a}^* , we have $\mathbf{H}\mathbf{a}^* = \mathbf{v}$, so $\mathbb{E}[\boldsymbol{\xi}^{j_i}] = \mathbf{v} - \mathbf{H}\mathbf{a}^* = \mathbf{0}$. Furthermore, $\boldsymbol{\xi}^{j_i}$ and $\boldsymbol{\xi}^{j_{i'}}$ are independent for all $i \neq i' \in [b]$. By Cauchy-Schwarz, and the bound from Claim 96,

$$\mathbb{E}[\|\boldsymbol{\zeta}^j\|] \leq \sqrt{\mathbb{E}[\|\boldsymbol{\zeta}^j\|^2]} = \sqrt{\mathbb{E}\left[\frac{1}{b^2} \sum_{i,i'=1}^b \boldsymbol{\xi}^{j_i} \cdot \boldsymbol{\xi}^{j_{i'}}\right]} = \sqrt{\frac{1}{b} \mathbb{E}[\|\boldsymbol{\xi}^{j_1}\|^2]} \leq B_3/\sqrt{b}.$$

■

Claim 98 *The second term in eq. (93) is bounded with probability $1 - \delta$:*

$$\mathbb{P}\left[(\text{Term } 2) \geq \frac{3B_3\sqrt{\log(1/\delta)}}{\lambda\sqrt{b}}\right] \leq \delta$$

Proof Group the samples $\{(\mathbf{y}_{j_i}, \mathbf{x}_{j_i})\}_{j \in \{0, \dots, k-1\}, i \in [b]}$ into groups where i is the same, letting

$$Z_i = ((\mathbf{y}_{0,i}, \mathbf{x}_{0,i}), (\mathbf{y}_{1,i}, \mathbf{x}_{1,i}), \dots, (\mathbf{y}_{k-1,i}, \mathbf{x}_{k-1,i}))$$

denote the collection of i th samples at all time-steps, for all $i \in [b]$. We can write the term that we want to bound as a function of these samples as

$$(\text{Term } 2) = h(Z_1, \dots, Z_b) = \sum_{j=0}^{k-1} (1 - \eta\lambda)^{k-1-j} \|\boldsymbol{\zeta}^j\|.$$

By the previous claim, the expectation is bounded by

$$\mathbb{E}[h(Z_1, \dots, Z_b)] \leq \frac{B_3}{\lambda\sqrt{b}}.$$

We bound deviation from the expectation using McDiarmid's inequality. Note that if we replace Z_1 with a new independent draw $\tilde{Z}_1 = ((\tilde{\mathbf{y}}_{0,i}, \tilde{\mathbf{x}}_{0,i}), (\tilde{\mathbf{y}}_{1,i}, \tilde{\mathbf{x}}_{1,i}), \dots, (\tilde{\mathbf{y}}_{k-1,i}, \tilde{\mathbf{x}}_{k-1,i}))$, then by triangle inequality almost surely

$$|h(Z_1, Z_2, \dots, Z_b) - h(\tilde{Z}_1, Z_2, \dots, Z_b)| \leq \eta \sum_{j=0}^{k-1} (1 - \eta\lambda)^{k-1-j} \frac{1}{b} \|\boldsymbol{\xi}^{j,1} - \tilde{\boldsymbol{\xi}}^{j,1}\| \leq \frac{2B_3}{b\lambda}.$$

By symmetry, the same is true of replacing any Z_i with an independent draw \tilde{Z}_i . Furthermore, Z_1, \dots, Z_b are independent, so by McDiarmid's inequality

$$\mathbb{P}\left[h(Z_1, \dots, Z_b) \geq \mathbb{E}[h(Z_1, \dots, Z_b)] + \frac{2B_3\sqrt{\log(1/\delta)}}{\lambda\sqrt{b}}\right] \leq \delta.$$

Combining the above bounds yields the claim. ■

The above claim combined with eq. (93) proves the lemma. ■

Lemma 99 *Under the same conditions as Lemma 95, for any $\delta > 0$, with probability $1 - \delta$ we have*

$$\ell(\mathbf{a}^k) \leq \ell(\mathbf{a}^*) + \frac{(B_2 + \lambda)}{2N} \left((1 - \eta\lambda)^k \|\mathbf{a}^0 - \mathbf{a}^*\| + \frac{3(\sqrt{B_1 B_2} + (B_2 + \lambda)\|\mathbf{a}^*\|)\sqrt{\log(1/\delta)}}{\lambda\sqrt{b}} \right)^2$$

Proof The excess loss at \mathbf{a} equals

$$\begin{aligned} \ell(\mathbf{a}) - \ell(\mathbf{a}^*) &= \frac{1}{2} \langle \mathbf{a} \otimes \mathbf{a}, \mathbf{H} \rangle - \frac{1}{2} \langle \mathbf{a}^* \otimes \mathbf{a}^*, \mathbf{H} \rangle - \langle \mathbf{v}, \mathbf{a} - \mathbf{a}^* \rangle \\ &= \frac{1}{2} \langle (\mathbf{a} - \mathbf{a}^*) \otimes (\mathbf{a} - \mathbf{a}^*), \mathbf{H} \rangle + \langle (\mathbf{a} - \mathbf{a}^*) \otimes \mathbf{a}^*, \mathbf{H} \rangle - \langle \mathbf{v}, \mathbf{a} - \mathbf{a}^* \rangle \\ &= \frac{1}{2} \langle (\mathbf{a} - \mathbf{a}^*) \otimes (\mathbf{a} - \mathbf{a}^*), \mathbf{H} \rangle + \langle \mathbf{H} \mathbf{a}^* - \mathbf{v}, \mathbf{a} - \mathbf{a}^* \rangle \\ &= \frac{1}{2} \langle (\mathbf{a} - \mathbf{a}^*) \otimes (\mathbf{a} - \mathbf{a}^*), \mathbf{H} \rangle, \end{aligned}$$

where the last step is by the first-order optimality condition $\mathbf{H} \mathbf{a}^* = \mathbf{v}$. We conclude by using the bound on $\|\mathbf{a}^k - \mathbf{a}^*\|$ from Lemma 95, and the fact that $\|\mathbf{H}\| \leq (B_2 + \lambda)/N$. \blacksquare

We may now prove the main result of this subsection: i.e., that it suffices to prove that there is a certificate \mathbf{a}^{cert} achieving low loss.

Proof [Proof of Lemma 93] Because of the quadratic regularization term and the optimality of \mathbf{a}^* ,

$$\frac{\lambda}{2N} \|\mathbf{a}^*\|^2 \leq \ell(\mathbf{a}^*) \leq \ell(\mathbf{a}^{cert}),$$

so

$$\|\mathbf{a}^*\|^2 \leq \frac{2N}{\lambda} \ell(\mathbf{a}^{cert}).$$

Plugging this into the bound from Lemma 99, with probability $1 - \delta$,

$$\begin{aligned} \ell(\mathbf{a}^k) &\leq \ell(\mathbf{a}^*) + \frac{(B_2 + \lambda)}{N} \left(2(1 - \eta\lambda)^{2k} (\|\mathbf{a}^0\|^2 + \|\mathbf{a}^*\|^2) \right. \\ &\quad \left. + \frac{18(B_1 B_2 + (B_2 + \lambda)^2 \|\mathbf{a}^*\|^2) \log(1/\delta)}{\lambda^2 b} \right) \\ &\leq \ell(\mathbf{a}^{cert}) + 2(B_2 + \lambda) \left((1 - \eta\lambda)^{2k} \left(\frac{\|\mathbf{a}^0\|^2}{N} + \frac{2\ell(\mathbf{a}^{cert})}{\lambda} \right) \right. \\ &\quad \left. + \left(\frac{9B_1 B_2}{\lambda^2 b N} + \frac{18(B_2 + \lambda)^2 \ell(\mathbf{a}^{cert})}{\lambda^3 b} \right) \log(1/\delta) \right). \end{aligned}$$

I.3. Bounding the loss reduces to lower bounding the eigenvalues of the kernel matrix

We give a lemma that bounds the final loss of batch-SGD on the second layer weights, after an initial phase of training the first layer weights. In Phase 1, we train \mathbf{W} for k_1 iterations while keeping \mathbf{a} fixed. In Phase 2, we train \mathbf{a} for $k_2 - k_1$ iterations, while keeping \mathbf{W} fixed. We reduce the problem of bounding the final loss $R(\mathbf{a}^{k_2}, \mathbf{W}^{k_1})$ to the problem of lower bounding the kernel matrix associated with the neural network constructed by the dimension-free dynamics during phase 1. We will use this result to provide explicit sample-complexity bounds in Section G.

Training setup Let $\Theta^k = (\theta_1^k, \dots, \theta_N^k)$ denote the parameters on step k of batch-SGD training. For $j \in [N]$ we independently draw

$$\theta_j^0 = (a_j^0, \mathbf{w}_j^0) \sim \text{Unif}([-1, 1]) \times \delta_{\mathbf{0}}.$$

We train Θ with (bSGD) with regularization $\lambda^w = 0$ and $\lambda^a = \lambda > 0$ which is a parameter to be set. Let $\eta_1, \eta_2 > 0$ be the step sizes during phase 1 and phase 2 respectively.

- In Phase 1, for time step $k \in \{0, \dots, k_1 - 1\}$, we update according to

$$\begin{aligned} \mathbf{w}_j^{k+1} &= \mathbf{w}_j^k + \eta_1 a_j^k \left[\frac{1}{b} \sum_{i \in [b]} \{y_{ki} - \hat{f}_{\text{NN}}(\mathbf{x}_{ki}; \Theta^k)\} \sigma'(\langle \mathbf{w}_j^k, \mathbf{x}_{ki} \rangle) \mathbf{x}_{ki} \right] \\ a_j^{k+1} &= a_j^k = a_j^0. \end{aligned}$$

- In Phase 2, for time step $k \in \{k_1, \dots, k_2 - 1\}$, we update according to

$$\begin{aligned} \mathbf{w}_j^{k+1} &= \mathbf{w}_j^{k_1} \\ a_j^{k+1} &= a_j^k + \eta_2 \left[\frac{1}{b} \sum_{i \in [b]} \{y_{ki} - \hat{f}_{\text{NN}}(\mathbf{x}_{ki}; \Theta^k)\} \sigma'(\langle \mathbf{w}_j^k, \mathbf{x}_{ki} \rangle) \mathbf{x}_{ki} \right] - \eta \lambda a_j^k. \end{aligned}$$

Reduction to analyzing the limiting mean-field dynamics For $k \in \{0, \dots, k_1\}$, consider either ρ_k or $\rho_{k_1 \eta_1}$ the limiting dynamics in the discrete-time or in the continuous time setting, and let $\bar{\Theta}^k = (\bar{\theta}_1^k, \dots, \bar{\theta}_N^k)$ or $\bar{\Theta}^t = (\bar{\theta}_1^t, \dots, \bar{\theta}_N^t)$ be the parameters if we had trained with the limiting discrete or continuous-time mean-field dynamics, initialized at $\bar{\Theta}^0 = \Theta^0$. If we denote $\bar{\mathbf{w}}_j^k = (\bar{\mathbf{u}}_j^k, \bar{\mathbf{v}}_j^k)$ and $\bar{\mathbf{w}}_j^t = (\bar{\mathbf{u}}_j^t, \bar{\mathbf{v}}_j^t)$, recall that the discrete mean-field dynamics is given by (note that with $\bar{\mathbf{w}}^0 = \mathbf{0}$ initialization, the mean-field PDE and dimension free dynamics are the same)

$$\begin{aligned} \bar{\mathbf{u}}_j^{k+1} &= \bar{\mathbf{u}}_j^k + \eta_1 \bar{a}_j^0 \mathbb{E}_{\mathbf{x}} [\{ \hat{f}_{\text{NN}}(\mathbf{z}; \bar{\rho}_k) - h_*(\mathbf{z}) \} \sigma'(\langle \bar{\mathbf{u}}_j^k, \mathbf{z} \rangle) \mathbf{x}] \\ \bar{\mathbf{v}}_j^{k+1} &= \bar{\mathbf{v}}_j^k = \mathbf{0}, \end{aligned}$$

while for the continuous time mean-field dynamics

$$\begin{aligned} \frac{d}{dt} \bar{\mathbf{u}}_j^t &= \bar{a}_j^0 \mathbb{E}_{\mathbf{x}} [\{ \hat{f}_{\text{NN}}(\mathbf{z}; \bar{\rho}_t) - h_*(\mathbf{z}) \} \sigma'(\langle \bar{\mathbf{u}}_j^t, \mathbf{z} \rangle) \mathbf{x}] \\ \bar{\mathbf{v}}_j^t &= \mathbf{0}. \end{aligned}$$

For convenience, denote $T_1 = k_1$ or $T_1 = \eta k_1$, and $\bar{\mathbf{u}}_j^{T_1}$ the weights at the end of phase 1. Denote $\text{err} = \max_{j \in [N]} \|\mathbf{w}_j^{T_1} - \bar{\mathbf{w}}_j^{T_1}\|_2$. The propagation-of-chaos argument in Propositions 15 and 92 yields the following bound with probability $1 - 1/N$ (when training only the first layer-weights):

$$\text{err} \leq K e^{KT_1} \left\{ \sqrt{\frac{\log N}{N}} + \sqrt{\frac{d + \log N}{b}} \right\}. \quad (94)$$

We define $\mathbf{K}^{T_1} \in \mathbb{R}^{2^P \times 2^P}$ the kernel at the end of phase 1, with entries

$$K^{T_1}(\mathbf{z}, \mathbf{z}') = \int \sigma(\langle \bar{\mathbf{u}}, \mathbf{z} \rangle) \sigma(\langle \bar{\mathbf{u}}, \mathbf{z}' \rangle) \bar{\rho}_{T_1}(d\bar{\mathbf{u}}).$$

We further define the following matrix \mathbf{M} associated with the features computed by the mean-field dynamics from initialization Θ_0 . $\mathbf{M} \in \mathbb{R}^{2^P \times N}$ is indexed by $\mathbf{z} \in \{+1, -1\}^P$ and $j \in [N]$, and has entries

$$M_{\mathbf{z},j} = \frac{1}{\sqrt{N}} \sigma(\langle \bar{\mathbf{u}}_j^{T_1}, \mathbf{z} \rangle).$$

We call this the mean-field feature matrix (MF feature matrix, for short). Because $\mathbf{a}_j^0 \in \mathbb{R}^N$ are iid with $\bar{\rho}_0$, the $\bar{\mathbf{u}}_j^{T_1}$ are iid distributed with respect to $\bar{\rho}_{T_1}$ and $\mathbf{M}\mathbf{M}$ is a random approximation of \mathbf{K}^{T_1} . Indeed, we have the following bound:

Lemma 100 *There exists $C > 0$ such that if $N \geq 2^{CP} \log(1/\delta)^2 / \lambda_{\min}(\mathbf{K}^{T_1})^2$, then with probability at least $1 - \delta$, we have*

$$\lambda_{\min}(\mathbf{M}\mathbf{M}^T) \geq \lambda_{\min}(\mathbf{K}^{T_1})/2.$$

Proof [Proof of Lemma 100] Note that $\mathbf{M} = [\mathbf{m}_1, \dots, \mathbf{m}_N]$ has iid columns with covariance $\mathbb{N}\mathbb{E}[\mathbf{m}_j \mathbf{m}_j^T] = \mathbf{K}^{T_1}$. Furthermore $\|\mathbf{m}_j\|_2^2 \leq K^2$. The result follows by applying (Vershynin, 2010, Theorem 4.44). \blacksquare

We prove a bound on the risk Θ^{k_2} found using batch-SGD in terms of the minimum singular value of \mathbf{M} .

Lemma 101 *[Sufficient to prove \mathbf{M} is well-conditioned] There is a universal constant $C > 0$ such that for any $0 < \delta < 1, 0 < \varepsilon < 1$, any k_1 , if we pick hyperparameters*

$$\begin{aligned} \lambda &= c_1 \varepsilon \quad \text{for} \quad c_1 = \frac{\min(1, \lambda_{\min}(\mathbf{M}\mathbf{M}^T)^2)}{2^{2P+2} K^6} \\ \eta &\leq 1/(2K^2) \\ N &\geq c_2/\varepsilon^3, \quad \text{for} \quad c_2 = (CK^2)^{3k_1} \log(1/\delta)^2 / c_1^3 \\ b &\geq c_3 \log(N)d/\varepsilon^4, \quad \text{for} \quad c_3 = (4K^2 + c_1)^2 (CK^2)^{2k_1} \log(1/\delta) / c_1^3 \\ k_2 &\geq k_1 + c_4 \log(1/\varepsilon)/(\eta\varepsilon), \quad \text{for} \quad c_4 = \log(100(4K^2 + c_1))/c_1, \end{aligned}$$

then, with probability at least $1 - \delta$, the final loss is bounded by

$$R(\mathbf{a}^{k_2}, \mathbf{W}^{k_1}) \leq \varepsilon.$$

Let us first prove that if $\mathbf{M}\mathbf{M}^T$ is well-conditioned then a low-error certificate exists. As in Section 1.2, let

$$\ell(\mathbf{a}) = R(\mathbf{a}, \mathbf{W}^{k_1}) + \frac{\lambda}{2N} \|\mathbf{a}\|^2.$$

Lemma 102 *There is a universal constant $C > 0$ such that for any $\delta > 0$ and mini-batch size $b \geq \max(d, C \log(1/\delta))$, with probability at least $1 - \delta$, there is $\mathbf{a}^{cert} \in \mathbb{R}^N$ such that*

$$\ell(\mathbf{a}^{cert}) \leq (\text{err} + \lambda) 2^{2P} K^6 / \lambda_{\min}(\mathbf{M}\mathbf{M}^T)^2,$$

where $\text{err} = \text{err}(k_1, b, N, \delta)$ is the error bound for convergence to the limiting discrete or continuous-time mean-field dynamics.

Proof Let $\beta \in \mathbb{R}^{2^P}$ be the vector of values of h_* , indexed by $z \in \{+1, -1\}^P$:

$$\beta_z = h_*(z).$$

We construct a certificate for the least-squares training in Phase 2. Let

$$\mathbf{a}^{cert} = \sqrt{N} \mathbf{M}^T (\mathbf{M} \mathbf{M}^T)^{-1} \beta.$$

We bound the norm of \mathbf{a}^{cert} in terms of the minimum singular value of \mathbf{M} :

$$\begin{aligned} \|\mathbf{a}^{cert}\| &\leq \sqrt{N} \|\mathbf{M}\| \|\beta\| / \lambda_{\min}(\mathbf{M} \mathbf{M}^T) \\ &\leq 2^P K^2 \sqrt{N} / \lambda_{\min}(\mathbf{M} \mathbf{M}^T), \end{aligned} \quad (95)$$

for a constant K since $\|\mathbf{M}\| \leq \sqrt{2^P} \|\sigma\|_{\infty} \max_j |a_j^{k_1}| \leq \sqrt{2^P} K$, and $\|\beta\| \leq \sqrt{2^P} \|h_*\|_{\infty} \leq K$. We bound the risk given by using \mathbf{a}^{cert} as a certificate. By the triangle inequality,

$$R(\mathbf{a}^{cert}, \mathbf{W}^{k_1}) = \frac{1}{2} \mathbb{E}_{\mathbf{x}} [(f_*(\mathbf{x}) - \frac{1}{N} \sum_{j=1}^N a_j^{cert} \sigma(\langle \mathbf{w}_j^{k_1}, \mathbf{x} \rangle))^2] \leq \Upsilon_1 + \Upsilon_2, \quad (96)$$

where

$$\begin{aligned} \Upsilon_1 &= \mathbb{E}_{\mathbf{x}} [(f_*(\mathbf{x}) - \frac{1}{N} \sum_{j=1}^N a_j^{cert} \sigma(\langle \bar{\mathbf{w}}_j^{k_1}, \mathbf{x} \rangle))^2] \\ \Upsilon_2 &= \mathbb{E}_{\mathbf{x}} [(\frac{1}{N} \sum_{j=1}^N a_j^{cert} (\sigma(\langle \mathbf{w}_j^{k_1}, \mathbf{x} \rangle) - \sigma(\langle \bar{\mathbf{w}}_j^{k_1}, \mathbf{x} \rangle)))^2]. \end{aligned}$$

We first bound Υ_1 , using the fact that that $f_*(\mathbf{x}) = h_*(z)$ and $\frac{1}{\sqrt{N}} \sigma(\langle \bar{\mathbf{w}}_j^{k_1}, \mathbf{x} \rangle) = \frac{1}{\sqrt{N}} \sigma(\langle \bar{\mathbf{u}}_j^{k_1}, z \rangle) = M_{z,j}$:

$$\Upsilon_1 = \mathbb{E}_{\mathbf{x}=(z,r)} [(h_*(z) - \frac{1}{\sqrt{N}} \sum_{j=1}^N a_j^{cert} M_{z,j})^2] = \frac{1}{2^P} \|\beta - \frac{1}{\sqrt{N}} \mathbf{M} \mathbf{a}^{cert}\|^2 = 0. \quad (97)$$

For Υ_2 , use (a) Jensen's inequality, (b) K -Lipschitzness of σ , and (c) linearity of expectation:

$$\begin{aligned} \Upsilon_2 &\stackrel{(a)}{\leq} \frac{1}{N} \sum_{j=1}^N \mathbb{E}_{\mathbf{x}} [(a_j^{cert} (\sigma(\langle \mathbf{w}_j^{k_1}, \mathbf{x} \rangle) - \sigma(\langle \bar{\mathbf{w}}_j^{k_1}, \mathbf{x} \rangle)))^2] \\ &\stackrel{(b)}{\leq} \frac{1}{N} \sum_{j=1}^N K^2 |a_j^{cert}|^2 \mathbb{E}_{\mathbf{x}} [\langle \mathbf{w}_j^{k_1} - \bar{\mathbf{w}}_j^{k_1}, \mathbf{x} \rangle^2] \\ &\stackrel{(c)}{=} \frac{1}{N} \sum_{j=1}^N K^2 |a_j^{cert}|^2 \|\mathbf{w}_j^{k_1} - \bar{\mathbf{w}}_j^{k_1}\|_2^2 \end{aligned}$$

Finally, by the propagation of chaos bound in Eq. (94), at time T_1 with probability at least $1 - \delta$,

$$\max_{j \in [N]} \|\mathbf{w}_j^{k_1} - \bar{\mathbf{w}}_j^{k_1}\|_2^2 \leq \text{err} := \text{err}(k_1, b, N, \delta) = e^{KT_1} (\sqrt{\log(N/\delta)/N} + \sqrt{d \log(N/\delta)/b}),$$

where $C > 0$ is some universal constant. So

$$\Upsilon_2 \leq K^2 \|\mathbf{a}^{cert}\|^2 (\text{err})/N. \quad (98)$$

Combining eqs. (95) to (98), with probability at least $1 - \delta$,

$$\begin{aligned} \ell(\mathbf{a}^{cert}) &= R(\mathbf{a}^{cert}, \mathbf{W}^{k_1}) + \frac{1}{2N} \lambda \|\mathbf{a}^{cert}\|^2 \\ &\leq K^2 \|\mathbf{a}^{cert}\|^2 (\text{err})/N + \frac{1}{2N} \lambda \|\mathbf{a}^{cert}\|^2 \\ &\leq (\text{err} + \lambda) 2^{2P} K^6 / \lambda_{\min}(\mathbf{M}\mathbf{M}^T)^2 \end{aligned}$$

■

Now we may prove the main result of this subsection.

Proof [Proof of Lemma 101] By the choice of hyperparameters N and b , we have the following error bound between the batch-SGD dynamics on the first layer and the infinite-width, population limit,

$$\text{err}(T_1, N, b, \delta/2) = e^{KT_1} (\sqrt{\log(N/\delta)/N} + \sqrt{d \log(N/\delta)/b}) \quad (99)$$

$$\leq \varepsilon / (2^{2P+1} K^6 \lambda_{\min}(\mathbf{M}\mathbf{M}^T)^2) \quad (100)$$

By Lemma 102, with probability $1 - \delta/2$ over the choice of the initialization Θ^0 and the samples $\{(\mathbf{x}_{ki}, y_i)\}_{k \in \{0, \dots, k_1-1\}, i \in [b]}$ there exists a certificate \mathbf{a}^{cert} with loss:

$$\ell(\mathbf{a}^{cert}) \leq (\text{err} + \lambda) 2^{2P} K^6 / \lambda_{\min}(\mathbf{M}\mathbf{M}^T)^2 \leq \varepsilon/2.$$

Let us conclude by applying Lemma 93, which shows that with probability $1 - \delta/2$ the final error satisfies the following bound in terms of $\ell(\mathbf{a}^{cert})$:

$$\begin{aligned} \ell(\mathbf{a}^{k_2}) &\leq \ell(\mathbf{a}^{cert}) + 2(B_2 + \lambda) \left((1 - \eta\lambda)^{2(k_2-k_1)} \left(\frac{\|\mathbf{a}^0\|^2}{N} + \frac{2\ell(\mathbf{a}^{cert})}{\lambda} \right) \right. \\ &\quad \left. + \left(\frac{9B_1B_2}{\lambda^2 b N} + \frac{18(B_2 + \lambda)^2 \ell(\mathbf{a}^{cert})}{\lambda^3 b} \right) \log(2/\delta) \right). \end{aligned}$$

Here we take $B_1 = 4K^2 \geq |y_{ki}|^2$ almost surely, and $B_2 = K^2$, since

$$\max_{\mathbf{x}} \|\phi(\mathbf{x})\|^2 = \max_{\mathbf{x}} \frac{1}{N} \sum_{i=1}^N |\sigma(\langle \mathbf{w}_i^{k_1}, \mathbf{x} \rangle)|^2 \leq K^2$$

almost surely. Also, $\|\mathbf{a}^0\|^2 \leq N$. Finally, $(1 - \eta\lambda)^{2(k_2-k_1)} \leq \varepsilon / (100(4K^2 + c_1\varepsilon))$. So with probability $1 - \delta/2$,

$$\begin{aligned} \ell(\mathbf{a}^{k_2}) &\leq \ell(\mathbf{a}^{cert}) + 2(4K^2 + c_1\varepsilon) \left((1 - \eta\lambda)^{2(k_2-k_1)} (1 + 2/c_1) \right. \\ &\quad \left. + \left(\frac{36K^4}{c_1^2 \varepsilon^2 b N} + \frac{18(4K^2 + c_1\varepsilon)^2 \varepsilon}{c_1^3 \varepsilon^3 b} \right) \log(2/\delta) \right) \\ &\leq \ell(\mathbf{a}^{cert}) + \varepsilon/2 \leq \varepsilon. \end{aligned}$$

We conclude by noting $R(\mathbf{a}^{k_2}, \mathbf{W}^{k_1}) \leq \ell(\mathbf{a}^{k_2})$. ■

I.4. Anti-concentration of polynomials

We prove the technical lemma that polynomials anti-concentrate when evaluated at random inputs. Concretely, we lower-bound the variance of the polynomial evaluated at a random input based on the sum of the magnitudes of its coefficients. Our bound is crude, but suffices for our purposes. We remark that anti-concentration bounds for polynomials in terms of their variance (and other moments) are a well-studied subject. For instance, the seminal paper [Carbery and Wright \(2001\)](#) bounds the probability that a polynomial of random variables lies in an interval in terms of the variance (or other moments) of the polynomial. In contrast, we bound the variance based on the sum of magnitudes of the polynomial's coefficients.

Lemma 103 (Polynomial anticoncentration) *For any integers $D, m > 0$, there exists a constant $c > 0$ such that the following hold. For any polynomial $h : \mathbb{R}^m \rightarrow \mathbb{R}$ of the form:*

$$h(\mathbf{z}) = \sum_{\alpha \in \{0, \dots, D\}^m} h_{\alpha} \prod_{l \in [m]} z_l^{\alpha_l},$$

we have

$$\mathbb{E}_{\mathbf{u} \sim \text{Unif}([-1, 1]^{\otimes m})} [h(\mathbf{u})^2] \geq c \left(\sum_{\alpha} |h_{\alpha}| \right)^2.$$

Proof Define $M = \sum_{\alpha} |h_{\alpha}|$ and let $\mathbf{u} \sim \text{Unif}([-1, 1]^{\otimes m})$. Let us lower bound $\mathbb{E}_{\mathbf{u}} [h(\mathbf{u})^2]$. For this purpose, we decompose $h(\mathbf{z})$ in the multivariate Legendre polynomial basis, which is defined as follows. Let $P_l : \mathbb{R} \rightarrow \mathbb{R}$ denote the degree- l Legendre polynomial in one-dimension. In particular, they satisfy the orthogonality relations: for any $k, l \in \mathbb{Z}_{\geq 0}$,

$$\mathbb{E}_{u \sim \text{Unif}([-1, 1])} [P_k(u)P_l(u)] = \delta_{kl}.$$

For any $\alpha \in \mathbb{Z}_{\geq 0}^m$, the multivariate Legendre polynomial is then given by

$$P_{\alpha}(\mathbf{z}) = \prod_{l \in [m]} P_{\alpha_l}(z_l).$$

These polynomials inherit the orthogonality relations over the multivariate uniform distribution, i.e., for any $\alpha, \beta \in \mathbb{Z}_{\geq 0}^m$,

$$\mathbb{E}_{\mathbf{u}} [P_{\alpha}(\mathbf{u})P_{\beta}(\mathbf{u})] = \delta_{\alpha\beta}.$$

The polynomials $\{P_{\alpha}\}_{\alpha \in \{0, \dots, D\}^m}$ therefore form an orthonormal basis over the multivariate polynomials whose degree in each variable is bounded by D . Writing $h(\mathbf{z})$ in this basis, we get

$$h(\mathbf{z}) = \sum_{\alpha \in \{0, \dots, D\}^m} g_{\alpha} P_{\alpha}(\mathbf{z}),$$

for some coefficients $g_{\alpha} \in \mathbb{R}$. For each multivariate Legendre polynomial we also write its expansion

$$P_{\alpha}(z) = \sum_{\beta \in \{0, \dots, D\}^m} p_{\alpha, \beta} \prod_{l \in [m]} z_l^{\beta_l},$$

for some coefficients $p_{\alpha,\beta} \in \mathbb{R}$. Therefore, for any β , we have

$$h_\beta = \sum_{\alpha \in \{0,\dots,D\}^m} g_\alpha p_{\alpha,\beta},$$

and so

$$\begin{aligned} M = \sum_{\beta} |h_\beta| &\leq \sum_{\beta, \alpha \in \{0,\dots,D\}^m} |g_\alpha p_{\alpha,\beta}| \\ &\leq (D+1)^m \max_{\alpha', \beta' \in \{0,\dots,D\}^m} |p_{\alpha',\beta'}| \cdot \sum_{\alpha} |g_\alpha| \leq C \sum_{\alpha} |g_\alpha|, \end{aligned} \quad (101)$$

for some constant $0 < C < \infty$ depending on m, D .

Therefore,

$$\begin{aligned} \mathbb{E}_{\mathbf{u}}[h(\mathbf{u})^2] &= \sum_{\alpha, \alpha' \in \{0,\dots,D\}^m} \mathbb{E}_{\mathbf{u}}[g_\alpha g_{\alpha'} P_\alpha(\mathbf{u}) P_{\alpha'}(\mathbf{u})] && \text{(linearity of expectation)} \\ &= \sum_{\alpha \in \{0,\dots,D\}^m} (g_\alpha)^2 && \text{(orthogonality relations)} \\ &\geq \frac{1}{(D+1)^m} \left(\sum_{\alpha \in \{0,\dots,D\}^m} |g_\alpha| \right)^2 && \text{(Cauchy-Schwarz)} \\ &\geq \frac{1}{(D+1)^m} (M/C)^2 && \text{(by eq. (101))} \end{aligned}$$

■

We will also use the following corollary:

Lemma 104 (Polynomial anticoncentration for shifted input distribution) *For any integers $m, D > 0$ and constant $C > 0$, there exists a constant $c > 0$ such that the following hold. For any $0 < \rho < 1$, any $\mathbf{w} \in \mathbb{R}^m$ with $\|\mathbf{w}\|_\infty \leq C$, and any polynomial $h : \mathbb{R}^m \rightarrow \mathbb{R}$ of the form*

$$h(\mathbf{z}) = \sum_{\alpha \in \{0,\dots,D\}^m} h_\alpha \prod_{l \in [m]} z_l^{\alpha_l},$$

we have, writing $\|\alpha\|_1 = \sum_{i \in [m]} \alpha_i$,

$$\mathbb{E}_{\mathbf{u} \sim \text{Unif}([- \rho, \rho]^{\otimes m})} [h(\mathbf{w} + \mathbf{u})^2] \geq c \left(\sum_{\alpha} |h_\alpha| \rho^{\|\alpha\|_1} \right)^2.$$

Proof Fix $0 < \rho < 1$ and $\mathbf{w} \in \mathbb{R}^m$ such that $\|\mathbf{w}\|_\infty \leq C$. Write

$$g(\mathbf{z}) = h(\mathbf{w} + \rho \mathbf{z}) = \sum_{\alpha \in \{0,\dots,D\}^m} g_\alpha \prod_{l \in [m]} (z_l)^{\alpha_l}.$$

Then

$$h(\mathbf{z}) = g((\mathbf{z} - \mathbf{w})/\rho) = \sum_{\alpha} \sum_{\beta} g_\alpha g_{\alpha,\beta} \prod_{l \in [m]} (z_l)^{\beta_l},$$

where we denoted

$$q_{\alpha}(z) = \prod_{l \in [m]} ((z_l - w_l)/\rho)^{\alpha_l} = \sum_{\beta} q_{\alpha, \beta} \prod_{l \in [m]} (z_l)^{\beta_l}.$$

We have the following easy bound

$$|q_{\alpha, \beta}| \leq ((C + 1)/\rho)^{\|\alpha\|_1},$$

so that we have the upper bound

$$\begin{aligned} \sum_{\alpha} |h_{\alpha}| \rho^{\|\alpha\|_1} &= \sum_{\beta} \left| \sum_{\alpha} g_{\alpha} q_{\alpha, \beta} \right| \rho^{\|\alpha\|_1} \\ &\leq \sum_{\beta} \sum_{\alpha} |g_{\alpha}| |q_{\alpha, \beta}| \rho^{\|\alpha\|_1} \leq (C + 1)^{mD} (D + 1)^m \sum_{\alpha} |g_{\alpha}|. \end{aligned}$$

We deduce that

$$\sum_{\alpha} |g_{\alpha}| \geq \sum_{\alpha} |h_{\alpha}| \rho^{\|\alpha\|_1} (C + 1)^{-mD} (D + 1)^{-m}.$$

The lemma follows by noting that $h(\mathbf{w} + \mathbf{u})$, where $\mathbf{u} \sim \text{Unif}([- \rho, \rho])$, is equal in distribution to $g(\mathbf{v})$, where $\mathbf{v} \sim \text{Unif}([-1, 1])$, and applying Lemma 103. ■