

Learning time-scales in two-layers neural networks

Raphaël Berthier*, Andrea Montanari†, Kangjie Zhou‡

April 19, 2024

Abstract

Gradient-based learning in multi-layer neural networks displays a number of striking features. In particular, the decrease rate of empirical risk is non-monotone even after averaging over large batches. Long plateaus in which one observes barely any progress alternate with intervals of rapid decrease. These successive phases of learning often take place on very different time scales. Finally, models learnt in an early phase are typically ‘simpler’ or ‘easier to learn’ although in a way that is difficult to formalize.

Although theoretical explanations of these phenomena have been put forward, each of them captures at best certain specific regimes. In this paper, we study the gradient flow dynamics of a wide two-layer neural network in high-dimension, when data are distributed according to a single-index model (i.e., the target function depends on a one-dimensional projection of the covariates). Based on a mixture of new rigorous results, non-rigorous mathematical derivations, and numerical simulations, we propose a scenario for the learning dynamics in this setting. In particular, the proposed evolution exhibits separation of timescales and intermittency. These behaviors arise naturally because the population gradient flow can be recast as a singularly perturbed dynamical system.

Keywords: Deep learning, Neural network, Gradient flow, Dynamical system, Non-convex optimization, Incremental learning

Mathematics Subject Classification: 34E15, 37N40, 68T07

Communicated by Joan Bruna

*EPFL, email address: raphael.berthier1@gmail.com

†Department of Electrical Engineering and Department of Statistics, Stanford University, email address: montanar@stanford.edu

‡Department of Statistics, Stanford University, email address: kangjie@stanford.edu

Contents

1	Introduction	3
2	Setting and canonical learning order	4
3	Further related work	8
4	The large-network, high-dimensional limit	9
4.1	Reduction to d -independent flow	10
4.2	Elimination of the products $\langle u_i, u_j \rangle$	11
4.3	Connection with mean field theory	12
4.4	A general formulation	12
5	Numerical solution	14
6	Timescales hierarchy in the gradient flow dynamics	17
6.1	Matched asymptotic expansions	18
6.2	First time scale: constant component	19
6.3	Second time scale: linear component I	20
6.4	Third time scale: linear component II	23
6.5	Conjectured behavior for larger time scales	25
7	Stochastic gradient descent and finite sample size	27
8	Discussion	29
A	Proof of Proposition 1	36
B	Appendix to Section 4	37
B.1	Proof of Proposition 2	37
B.2	Proof of Corollary 1	38
B.3	Proof of Proposition 3	39
B.4	Derivation of the mean field dynamics (29)	43
B.5	Details of the alternative mean field approach	43
B.6	Proof of Proposition 4	45
C	Calculations for the analysis of mean-field gradient flow	46
C.1	Solution of Eq. (89)	46
C.2	Induced approximation of the risk	48
C.3	Proof of Theorem 1	49
D	Proofs of Theorem 2 and 3: learning with projected SGD	53
D.1	Difference between GF and GD	54
D.2	Difference between GD and SGD	58
D.3	Difference between SGD and projected SGD	59
D.4	Proof of Theorem 3	62
E	Counterexamples to the canonical learning order	63
E.1	Case 1: $\sigma_k = 0$ for some $k \in \mathbb{N}$	63
E.2	Case 2: $\varphi_0 = \dots = \varphi_k = 0$ for some $k \geq 1$	63
E.3	Case 3: $\varphi_k = 0$ for some $k \geq 1$	64

1 Introduction

It is a recurring empirical observation that the training dynamics of neural networks exhibits a whole range of surprising behaviors:

1. *Plateaus.* Plotting the training and test error as a function of SGD steps, using either small stepsize or large batches to average out stochasticity, reveals striking patterns. These error curves display long plateaus where barely anything seems to be happening, which are followed by rapid drops [41, 48, 39].
2. *Time-scales separation.* The time window for this rapid descent is much shorter than the time spent in the plateaus. Additionally, subsequent phases of learning take increasingly longer times [18, 8].
3. *Incremental learning.* Models learnt in the first phases of learning appear to be simpler than in later phases. Among others, [5] demonstrated that easier examples in a dataset are learned earlier; [28] showed that models learnt in the first phase of training correlate well with linear models; [22] showed that, in many simplified models, the dynamics of gradient descent explores the solution space in an incremental order of complexity; [39] demonstrated that, in certain settings, a function that approximates well the target is only learnt past the point of overfitting.

Understanding these phenomena is not a matter of intellectual curiosity. In particular, incremental learning plays a key role in our understanding of generalization in deep learning. Indeed, in this scenario, stopping the learning at a certain time t amounts to controlling the complexity of the model learnt. The notion of complexity corresponds to the order in which the space of models is explored.

While a number of groups have developed models to explain these phenomena, it is fair to say that a complete picture is still lacking. An exhaustive overview of these works is out of place here. We will outline three possible explanations that have been developed in the past, and provide more pointers in Section 3.

Theory #1: Dynamics near singular points. Several early works [41, 17, 44] pointed out that the parametrization of multi-layer neural networks presents symmetries and degeneracies. For instance, the function represented by a multi-layer perceptron is invariant under permutations of the neurons in the same layer. As a consequence, the population risk has multiple local minima connected through saddles or other singular sub-manifolds. Dynamics near these sub-manifolds naturally exhibits plateaus. Further, random or agnostic initializations typically place the network close to such submanifolds.

Theory #2: Linear networks. Following the pioneering work of [7], a number of authors, most notably [43, 30], studied the behavior of deep neural networks with linear activations. While such networks can only represent linear functions, the training dynamics is highly non-linear. As demonstrated in [43], learning happens through stages that correspond to the singular value decomposition of the input-output covariance. Time scales are determined by the singular values.

Theory #3: Kernel regime. Following an initial insight of [26], a number of groups proved that, for certain initializations, the training dynamics and model learnt by overparametrized neural networks is well approximated by certain linearly parametrized models. In the limit of very wide networks, the training dynamics of these models converges in turn to the training dynamics of kernel ridge(less) regression (KRR) with respect to a deterministic kernel (independent of the random initialization.) We refer to [9] for an overview and pointers to this literature. Recently

[21] show that, in high dimension, the learning dynamics of KRR also exhibits plateaus and waterfalls, and learns functions of increasing complexity over a diverging sequence of timescales.

While each of these theories offers useful insights, it is important to realize that they do not agree on the basic mechanism that explains plateaus, time-scales separation, and incremental learning. In theory #1, plateaus are associated to singular manifolds and high-dimensional saddles, while in theories #2 and #3 they are related to a hierarchy of singular values of a certain matrix. In #2, the relevant singular values are the ones of the input-output covariance, and the fact that these singular values are well separated is postulated to be a property of the data distribution. In contrast, in #3 the relevant singular values are the eigenvalues of the kernel operator, and hence completely independent of the output (the target function). In this case, eigenvalues which are very different are proved to exist under natural high-dimensional distributions.

Not only these theories propose different explanations, but they are also motivated by very different simplified models. Theory #1 has been developed only for networks with a small number of hidden units. Theory #2 only applies to networks with multiple output units, because otherwise the input-output covariance is a $d \times 1$ matrix and hence has only one non-trivial singular value. Finally, theory #3 applies under the conditions of the linear (a.k.a. lazy) regime, namely large overparametrization and suitable initialization (see, e.g., [9]).

In order to better understand the origin of plateaus, time-scales separation, and incremental learning, we attempt a detailed analysis of gradient flow for two-layer neural networks. We consider a simple data-generation model, and propose a precise scenario for the behavior of learning dynamics. We do not assume any of the simplifying features of the theories described above: activations are non-linear; the number of hidden neurons is large; we place ourselves outside the linear (lazy) regime.

Our analysis is based on methods from dynamical systems theory: singular perturbation theory and matched asymptotic expansions. Unfortunately, we fall short of providing a general rigorous proof of the proposed scenario, but we can nevertheless prove it in several special cases and provide a heuristic argument supporting its generality.

The rest of the paper is organized as follows. Section 2 describes our data distribution, learning model, and the proposed scenario for the learning dynamics. We review further related work in Section 3. Section 4 describes the reduction of the gradient flow to a ‘mean field’ dynamics that will be the starting point of our analysis. Section 5 presents numerical evidence of the proposed learning scenario. Finally, Sections 6 to 7 present our analysis of the learning dynamics.

Notations. In this paper, we use the classical asymptotic notations. The notations $f(\varepsilon) = o(g(\varepsilon))$ or $g(\varepsilon) = \omega(f(\varepsilon))$ as $\varepsilon \rightarrow 0$ both denote that $|f(\varepsilon)|/|g(\varepsilon)| \rightarrow 0$ in the limit $\varepsilon \rightarrow 0$. The notations $f(\varepsilon) = O(g(\varepsilon))$ or $g(\varepsilon) = \Omega(f(\varepsilon))$ both denote that the ratio $|f(\varepsilon)|/|g(\varepsilon)|$ remains upper bounded in the limit. The notation $f(\varepsilon) = \Theta(g(\varepsilon))$ or $f(\varepsilon) \asymp g(\varepsilon)$ denote that $f(\varepsilon) = O(g(\varepsilon))$ and $g(\varepsilon) = O(f(\varepsilon))$ both hold. Finally, $f(\varepsilon) \sim g(\varepsilon)$ denotes that $f(\varepsilon)/g(\varepsilon) \rightarrow 1$ in the limit.

2 Setting and canonical learning order

We are given pairs $\{(x_i, y_i)\}_{i \leq n}$, where $x_i \in \mathbb{R}^d$ is a feature vector and $y_i \in \mathbb{R}$ is a response variable. We are interested in cases in which the feature vector is high-dimensional but does not contain strong structure, but the response depends on a low-dimensional projection of the data. We assume the simplest model of this type, the so-called single-index model:

$$y_i = \varphi(\langle u_*, x_i \rangle), \quad x_i \sim \mathbf{N}(0, I_d), \quad u_* \in \mathbb{S}^{d-1}, \quad (1)$$

where $\varphi : \mathbb{R} \rightarrow \mathbb{R}$ is a link function, $\mathbf{N}(0, I_d)$ denotes the standard multivariate Gaussian distribution in dimension d , and $\mathbb{S}^{d-1} := \{v \in \mathbb{R}^d : \|v\|_2 = 1\}$. We study the ability to learn

model (1) using a two-layers neural network with m hidden neurons:

$$f(x; a, u) = \frac{1}{m} \sum_{i=1}^m a_i \sigma(\langle u_i, x \rangle), \quad a_1, \dots, a_m \in \mathbb{R}, \quad u_1, \dots, u_m \in \mathbb{S}^{d-1}, \quad (2)$$

where $(a, u) := (a_1, \dots, a_m, u_1, \dots, u_m)$ collectively denotes all the model's parameter and $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ is the activation function of the neural network. The factor $1/m$ in the definition is relevant for the initialization and learning rate. We anticipate that we will initialize the a_i 's to be of order one, which results in second layer coefficients $a_i/m = \Theta(1/m)$.

Remark 2.1. Standard initializations in deep learning frameworks yield second-layer coefficients $a_i/m = \Theta(1/\sqrt{m})$ [29, 23, 24]. However, it is increasingly clear that this initialization presents fundamental limitations for large m . Notably, two-layers networks with this initialization converges to kernel methods [37], and the latter cannot learn ridge functions from polynomially many samples [20, 47].

It is well understood that, in order to drive the learning process outside the kernel regime (for $m \rightarrow \infty$), it is necessary to set $a_i/m = \Theta(1/m)$. This is often referred to as the ‘mean-field initialization’ [32, 13, 19, 1]. We notice that suitable generalizations of the mean-field initialization are currently used in state-of-the-art implementations [45, 46].

The bulk of our work will be devoted to the analysis of projected gradient flow in $(a_i, u_i)_{1 \leq i \leq m}$ on the population risk

$$\mathcal{R}(a, u) = \frac{1}{2} \mathbb{E} \{ (y - f(x; a, u))^2 \} \quad (3)$$

$$= \frac{1}{2} \mathbb{E} \left\{ \left(\varphi(\langle u_*, x \rangle) - \frac{1}{m} \sum_{i=1}^m a_i \sigma(\langle u_i, x \rangle) \right)^2 \right\}. \quad (4)$$

In Section 7, we will bound the distance between stochastic gradient descent (SGD) and gradient flow in population risk. As a consequence, we will establish finite sample generalization guarantees for SGD learning.

Projected gradient flow with respect to the risk $\mathcal{R}(a, u)$ is defined by the following ordinary differential equations (ODEs):

$$\partial_t(\varepsilon a_i) = -m \partial_{a_i} \mathcal{R}(a, u), \quad (5)$$

$$\partial_t u_i = -m(I_d - u_i u_i^\top) \nabla_{u_i} \mathcal{R}(a, u). \quad (6)$$

Here, ε can be viewed as the relative step size, namely the ratio between the first and second-layer step sizes. It is useful to make a few remarks about the definition of gradient flow:

- The projection $I_d - u_i u_i^\top$ ensures that u_i remains on the unit sphere \mathbb{S}^{d-1} .
- The overall scaling of time is arbitrary, and the matching to SGD steps will be carried out in Section 7. The factors m on the right-hand side are introduced for mathematical convenience, since the partial derivatives are of order $1/m$.
- As aforementioned, the factor ε introduced in the gradient flow of the a_i 's plays the role of the relative step size. Throughout the paper, we will keep ε as a free parameter independent of m , and study the evolution of gradient flow for small ε . This corresponds to a setting in which the second-layer coefficients are learned much faster than the first-layer weights. We emphasize however that the small ε limit is taken after the large m, d limits. Thus, despite the second-layer weights are learnt faster, the evolution of first layer weights will be crucial, and lead to true feature learning.




Figure 1: Cartoon illustration of the evolution of the population risk within the canonical learning order of Definition 1.

We assume the initialization to be random with i.i.d. components $(a_{i,\text{init}}, u_{i,\text{init}})$:

$$(a_{i,\text{init}}, u_{i,\text{init}}) \sim P_A \otimes \text{Unif}(\mathbb{S}^{d-1}), \quad (7)$$

where P_A is a probability measure on \mathbb{R} . The unique solution of the gradient flow ODEs with this initialization will be denoted by $(a(t), u(t))$. We will be interested in the case of large networks ($m \rightarrow \infty$) in high dimension ($d \rightarrow \infty$). As shown below, the two limits commute (over fixed time horizons).

Our main finding is that, in a number of cases, φ is learnt incrementally. Namely, the function $f(x; a(t), u(t))$ evolves over time according to a sequence of polynomial approximations of $\varphi(\langle u_*, x \rangle)$. These polynomial approximations are given by the decomposition of φ in $L^2(\mathbb{R}, \phi(x)dx)$, where $\phi(x)$ is the standard normal density: $\phi(x) = \exp(-x^2/2)/\sqrt{2\pi}$. (For notational simplicity, we will use the shorthand L^2 instead of $L^2(\mathbb{R}, \phi(x)dx)$ in the sequel.)

In order to describe the polynomial approximations learnt during the training more explicitly, we decompose φ and σ into normalized Hermite polynomials:

$$\varphi(z) = \sum_{k=0}^{\infty} \varphi_k \text{He}_k(z), \quad \sigma(z) = \sum_{k=0}^{\infty} \sigma_k \text{He}_k(z). \quad (8)$$

Here, He_k denotes the k -th Hermite polynomial, normalized so that $\|\text{He}_k\|_{L^2(\mathbb{R}, \phi(x)dx)} = 1$.

As we will see, the incremental learning behavior arises for small ε . By the law of large numbers (see below), the following almost sure limit exists (provided P_A is square integrable)

$$\mathcal{R}_{\text{init}} := \lim_{m \rightarrow \infty} \lim_{d \rightarrow \infty} \mathcal{R}(a_{\text{init}}, u_{\text{init}}) = \frac{1}{2} \left(\varphi_0 - \sigma_0 \int a P_A(da) \right)^2 + \frac{1}{2} \sum_{k \geq 1} \varphi_k^2. \quad (9)$$

We are now in position to describe the scenario that we will study in the rest of the paper.

Definition 1. We say that the canonical learning order holds up to level L for a certain target function φ , activation σ , and distribution P_A , if the followings hold:

1. The limit below exists:

$$\mathcal{R}_\infty(t, \varepsilon) = \lim_{m \rightarrow \infty} \lim_{d \rightarrow \infty} \mathcal{R}(a(t), u(t)). \quad (10)$$

2. There exist constants $c_2, \dots, c_{L+1} > 0$ such that the following asymptotic holds as $\varepsilon \rightarrow 0$, $t \rightarrow 0$:

$$\mathcal{R}_\infty(t, \varepsilon) \xrightarrow{\varepsilon \rightarrow 0, t \rightarrow 0} \begin{cases} \mathcal{R}_{\text{init}} & \text{if } t = o(\varepsilon), \\ \frac{1}{2} \sum_{k \geq 1} \varphi_k^2 & \text{if } t = \omega(\varepsilon) \text{ and } t = \frac{1}{4|\sigma_1 \varphi_1|} \varepsilon^{1/2} \log \frac{1}{\varepsilon} - \omega(\varepsilon^{1/2}), \\ \frac{1}{2} \sum_{k \geq 2} \varphi_k^2 & \text{if } t = \frac{1}{4|\sigma_1 \varphi_1|} \varepsilon^{1/2} \log \frac{1}{\varepsilon} + \omega(\varepsilon^{1/2}) \text{ and } t = c_2 \varepsilon^{1/4} - \omega(\varepsilon^{1/3}), \\ \frac{1}{2} \sum_{k \geq l} \varphi_k^2 & \text{if } t = c_{l-1} \varepsilon^{1/2(l-1)} + \omega(\varepsilon^{1/l}) \text{ and } t = c_l \varepsilon^{1/2l} - \omega(\varepsilon^{1/(l+1)}), \\ & \text{for all } 3 \leq l \leq L+1. \end{cases}$$

Figure 1 provides a cartoon illustration of the canonical learning order.

At first sight, the setting of Eq. (2) is overly restrictive because we require $\|u_i\|_2 = 1$ and we do not have offsets in the activations. Therefore, it might seem that $s = 1$ and $\varphi = \sigma$ is required in order to approximate arbitrarily well the target function. In contrast, the next proposition shows that the network (2) enjoys universal approximation properties.

Proposition 1. Assume that σ is Lipschitz continuous and generic in the following sense: the decomposition of σ into Hermite polynomials does not have any coefficient equal to 0. For any Lipschitz function $\varphi : \mathbb{R} \rightarrow \mathbb{R}$, $\|u_*\|_2 = 1$, and $x \sim \mathcal{N}(0, I_d)$ such that $\mathbb{E}\{\varphi(\langle u_*, x \rangle)^2\} < \infty$, there exists a sequence $m \rightarrow \infty$ and $a^{(m)}, u^{(m)}$ with $\|u_i^{(m)}\|_2 = 1$ such that

$$\lim_{d \rightarrow \infty} \lim_{m \rightarrow \infty} \mathbb{E}\{(\varphi(\langle u_*, x \rangle) - f(x; a^{(m)}, u^{(m)}))^2\} = 0.$$

This result is not surprising in view of the arguments in the next sections, which suggest that indeed gradient flow constructs such an approximation for a broad class of functions of the form $f_*(x) = \varphi(\langle u_*, x \rangle)$. We nevertheless give an independent proof in Appendix A.

A specific realization of our general setup is determined by the triple (σ, φ, P_A) . In the rest of the paper, we will provide evidence showing that the canonical learning order holds in a number of cases. Nevertheless, we can also construct examples in which it does not hold:

- If one or more of the Hermite coefficients of the activation vanish, then the canonical learning order does not hold for general φ . Specifically, if $\sigma_k = 0$, then for any t the function $f(x; a(t), u(t))$ remains orthogonal to $\text{He}_k(\langle u_*, x \rangle)$. In particular, if $\varphi_k \neq 0$ then the risk remains bounded away from zero for every t . We refer to Appendix E.1 for a formal statement.
- If the first $k+1$ Hermite coefficients of φ vanish, $\varphi_0 = \dots = \varphi_k = 0$, $k \geq 1$, then the canonical learning order does not hold. (See Appendix E.2 for the proof.)
- In fact, we expect the canonical learning order might fail every time one or more of the coefficients φ_k vanish, for $k \geq 1$. Appendix E.3 provides some heuristic justification for this failure.

Remark 2.2. We can compare the canonical learning order described here to the ones in earlier literature and described as theory #1, #2, #3 in the introduction. There appears points of contact, but also important differences with both theory #1 and #3:

- As in theory #1, the plateaus and separation of time scales arise because the trajectory of gradient flow is approximated by a sequence of motions along submanifolds in the space of parameters (a, u) . Along the l -th such submanifold $f(x; a, u)$ is well-approximated by a degree- l polynomial. Escaping each submanifold takes an increasingly longer time.

This is reminiscent of the motion between saddles investigated in earlier work [41, 17, 44]. However, unlike in earlier work, we will see that this applies to networks with a large (possibly diverging) number of hidden neurons. Also, we identify the subsequent phases of learning with the polynomial decomposition of Eq. (8).

- As in theory #3, subsequent phases of learning correspond to increasingly accurate polynomial approximations of the target function $\varphi(\langle u_*, x \rangle)$. However, the underlying mechanism and time scales are completely different. In the linear regime, the different time scales emerge because of increasingly small eigenvalues of the neural tangent kernel. In that case, the time required to learn degree- l polynomials is of order d^l [21].

In contrast, in the canonical learning order, polynomials of degree l are learnt on a time scale of order one in d (and only depending on the learning rate ε). This of course has important implications when approximating gradient flow by SGD. Within the linear regime, the sample size required to learn a polynomial of order l scales like d^l [21], while in the canonical learning order, it is only of order d (see Section 7).

3 Further related work

As we mentioned in the introduction, plateaus and time scales in the learning dynamics of kernel models were analyzed by [21]. A sharp analysis for the related random features model was developed by [12].

Our analysis builds upon the mean-field description of learning in two-layer neural networks, which was developed in a sequence of works, see, e.g., [32, 40, 13, 33]. In particular, we leverage the fact that, for the data distribution (1), the population risk function is invariant under rotations around the axis u_* , and this allows for a dimensionality reduction in the mean field description. Similar symmetry argument were used by [32] and, more recently, by [1].

The single-index model can be learnt using simpler methods than large two-layer networks. Limiting ourselves to the case of gradient descent algorithms, [31] proved that gradient descent with respect to the non-convex empirical risk $\hat{R}_n(u) := n^{-1} \sum_{i=1}^n (y_i - \varphi(u^\top x_i))^2$ converges to a near global optimum, provided φ is strictly increasing. [4] considered online SGD under more challenging learning scenarios and characterized the time (sample size) for $|\langle u, u_* \rangle|$ to become significantly larger than for a random unit vector u .

Learning in overparametrized two-layer networks under model (1) (or its variations) has been studied recently by several groups. In particular, [6] considers a training procedure which runs a single step gradient descent followed by freezing the first layer and performing ridge regression with respect to the second layer. This scheme is amenable to a precise characterization of the generalization error. [11] consider a similar scheme in which a first phase of gradient descent is run to achieve positive correlation with the unknown direction u_* . [14] also consider a two-phases scheme, and prove consistency and excess risk bounds for a more general class of target functions whereby the first equation in (1) is replaced by

$$y_i = \varphi(U_*^\top x_i) + \varepsilon_i, \quad U_* \in \mathbb{R}^{d \times k}, \varphi : \mathbb{R}^k \rightarrow \mathbb{R}, \quad (11)$$

with $k \ll d$. In particular, near optimal error bounds are obtained under a non-degeneracy condition on $\nabla^2 \varphi$.

[1] consider a similar model whereby $x \sim \text{Unif}(\{+1, -1\}^d)$, and $y = \varphi(x_S)$ where $S \subseteq [d]$, and $x_S = (x_i)_{i \in S}$ (i.e., x_S contains the coordinates of x indexed by entries of S). Under a

structural assumption on φ (the ‘merged staircase property’), and for $|S|$ fixed, they prove the two stages algorithm learns the target function with sample complexity of order d . This paper is technically related to ours in that it uses mean-field theory to obtain a characterization of learning in terms of a PDE in a reduced $(k+2)$ -dimensional space.

A similar model was studied by [8] that bounds the sample complexity by $d^{O(k)}$ for learning parities on k bits using gradient descent with large batches (if $k = O(1)$, [8] require $O(1)$ steps with batch size $d^{O(k)}$).

Let us emphasize that our objective is quite different from these works. We implement a simple online SGD algorithm with additional projection steps, and try to derive a precise picture of the successive phases of learning (in particular, we do not consider two-stage schemes or layer-by-layer learning). On the other hand, we focus on a relatively simple model.

To clarify the difference, it is perhaps useful to rephrase our claims in terms of sample complexity. While previous works show that the target function can be learnt with $O(d)$ samples, we claim that it is learnt by online SGD with test error r from about $C(r, \varepsilon)d$ samples and characterize the dependence of $C(r, \varepsilon)$ on r for small ε . (Falling short of a proof in the general case.)

After posting an initial version of this paper, we became aware that [3] independently derived equations similar to (15)-(19), (26), (130). There are technical differences, and hence we cannot apply their results directly. However, Section 4.3 and Appendix B.4 are analogous to their work.

4 The large-network, high-dimensional limit

The first step of our analysis is a reduction of the system of ODEs (5), (6), with dimension $m(d+1)$ to a system of ODEs in $2m$ dimensions. We will achieve this reduction in two steps:

- (i) First we reduce to a system in $m(m+3)/2$ dimensions for the variables $a_i, \langle u_i, u_j \rangle, \langle u_i, u_* \rangle$. This reduction is exact and is quite standard. It is done in Section 4.1.
- (ii) We then show that the products $\langle u_i, u_j \rangle$ can be eliminated, with an error $O(1/m)$. This is done in Section 4.2. As further discussed below, the resulting dynamics could also be derived from the mean field theory of [32, 40, 13, 33] (with the required modifications for the constraints $\|u_i\| = 1$).

In order to define formally the reduced system, we define the functions $U, V : [-1, 1] \rightarrow \mathbb{R}$ via:

$$V(s) := \mathbb{E}\{\varphi(G) \sigma(G_s)\} = \sum_{k \geq 0} \varphi_k \sigma_k s^k, \quad (G, G_s) \sim \mathcal{N}\left(0, \begin{bmatrix} 1 & s \\ s & 1 \end{bmatrix}\right), \quad (12)$$

$$U(s) := \mathbb{E}\{\sigma(G) \sigma(G_s)\} = \sum_{k \geq 0} \sigma_k^2 s^k. \quad (13)$$

Note that the above identities follow from [36, Proposition 11.31]. Throughout this section, we will make the following assumptions.

A1. The distribution of weights at initialization, P_A is supported on $[-M_1, M_1]$.

A2. The activation function is bounded: $\|\sigma\|_\infty \leq M_2$. Additionally, the functions V and U are bounded and of class C^2 , with uniformly bounded first and second derivatives over $s \in [-1, 1]$. A sufficient condition for this is

$$\sup \{\|\sigma'\|_{L^2}, \|\sigma''\|_{L^2}\} \leq M_2, \quad \sup \{\|\varphi\|_{L^2}, \|\varphi'\|_{L^2}, \|\varphi''\|_{L^2}\} \leq M_2.$$

A3. Responses are bounded, i.e., $\|\varphi\|_\infty \leq M_3$.

Remark 4.1. We hereby briefly explain the sufficiency of L^2 -boundedness of derivatives of σ and φ as claimed in Assumption A2. Suppose for example that $\|\sigma'\|_{L^2} \leq M_2$ and $\|\varphi'\|_{L^2} \leq M_2$, then we have

$$\sup_{s \in [-1, 1]} |V'(s)| \stackrel{(a)}{=} \sup_{s \in [-1, 1]} |\mathbb{E}\{\varphi'(G) \sigma'(G_s)\}| \stackrel{(b)}{\leq} \|\varphi'\|_{L^2} \|\sigma'\|_{L^2} \leq M_2^2, \quad (14)$$

where (a) follows from Gaussian integration by parts and (b) follows from Cauchy-Schwarz inequality.

4.1 Reduction to d -independent flow

Our first statement establishes reduction (i) mentioned above. The proof of this fact is presented in Appendix B.1.

Proposition 2 (Reduction to d -independent flow). *Define $s_i = \langle u_i, u_* \rangle$, $r_{ij} = \langle u_i, u_j \rangle$ for $i, j = 1, \dots, m$. Then, letting $R = (r_{ij})_{i,j \leq m}$, we have*

$$\mathcal{R}(a, u) = \mathcal{R}_{\text{red}}(a, s, R) := \frac{1}{2} \|\varphi\|_{L^2}^2 - \frac{1}{m} \sum_{i=1}^m a_i V(s_i) + \frac{1}{2m^2} \sum_{i,j=1}^m a_i a_j U(r_{ij}). \quad (15)$$

If $(a(t), u(t))$ solve the gradient flow ODEs (5)-(6) then $(a(t), s(t), R(t))$ are the unique solution of the following set of ODEs (note that $r_{ii} = 1$ identically)

$$\varepsilon \partial_t a_i = V(s_i) - \frac{1}{m} \sum_{j=1}^m a_j U(r_{ij}), \quad (16)$$

$$\partial_t s_i = a_i \left(V'(s_i)(1 - s_i^2) - \frac{1}{m} \sum_{j=1}^m a_j U'(r_{ij})(s_j - r_{ij}s_i) \right), \quad (17)$$

$$\partial_t r_{ij} = a_i \left(V'(s_i)(s_j - s_i r_{ij}) - \frac{1}{m} \sum_{p=1}^m a_p U'(r_{ip})(r_{jp} - r_{ip} r_{ij}) \right), \quad (18)$$

$$+ a_j \left(V'(s_j)(s_i - s_j r_{ij}) - \frac{1}{m} \sum_{p=1}^m a_p U'(r_{jp})(r_{ip} - r_{jp} r_{ij}) \right). \quad (19)$$

The input dimension d does not appear in the reduced ODEs, Eqs. (16) to (19), and only plays a role in the initialization of the s_i 's and the r_{ij} 's. Namely, since $u_{i,\text{init}} \sim \text{Unif}(\mathbb{S}^{d-1})$, we can represent $u_{i,\text{init}} = g_i / \|g_i\|_2$ with $g_i \sim \mathcal{N}(0, I_d/d)$. By concentration of $\|g_i\|_2$, this implies that, for $1 \leq i < j \leq m$, s_i, r_{ij} are approximately $\mathcal{N}(0, 1/d)$.

This discussion immediately yields the following consequence.

Corollary 1. *Let $(a(t), u(t))$ be the solution of the gradient flow ODEs (5), (6) with initialization (7), and let $(a^0(t), s^0(t), R^0(t))$ be the unique solution of Eqs. (16) to (19), with initialization $a_i^0(0) = a_i(0)$, $s_i^0(0) = 0$, $r_{ij}^0(0) = 0$ for $i \neq j$. Then, for any fixed T (possibly dependent on m but not on d), the followings holds with probability at least $1 - \exp(-C'm)$ over the i.i.d. initialization $(a_i(0), u_i(0))_{i \in [m]}$:*

$$\sup_{t \in [0, T]} |\mathcal{R}(a(t), u(t)) - \mathcal{R}_{\text{red}}(a^0(t), s^0(t), R^0(t))| \leq \frac{CM}{\sqrt{d}} \exp\left(MT(1+T)^2/\varepsilon^2\right), \quad (20)$$

$$\max \left(\sup_{t \in [0, T]} \frac{1}{\sqrt{m}} \|a(t) - a^0(t)\|_2, \frac{1}{\sqrt{m}} \sup_{t \in [0, T]} \|s(t) - s^0(t)\|_2 \right) \leq \frac{1}{\sqrt{d}} \cdot C \exp\left(MT(1+T)^2/\varepsilon^2\right), \quad (21)$$

$$\sup_{t \in [0, T]} \frac{1}{m} \|R(t) - R^0(t)\|_F \leq \frac{1}{\sqrt{d}} \cdot C \exp\left(MT(1+T)^2/\varepsilon^2\right). \quad (22)$$

Here C, C' are absolute constants and M only depends on the M_i 's in Assumptions A1-A3.

The proof of Corollary 1 is deferred to Appendix B.2. From now on, we will assume the initialization $s_i^0(0) = 0$, $r_{ij}^0(0) = 0$ for $i \neq j$, but drop the superscript 0 for notational simplicity. We notice in passing that the right-hand sides of Eqs. (20) to (22) are independent of m : this approximation step holds uniformly over m . (Note that the left hand sides are normalized by m as to yield the root mean square error per entry.)

4.2 Elimination of the products $\langle u_i, u_j \rangle$

In order to state the reduction (ii) outlined above, we define the mean field risk as

$$\mathcal{R}_{\text{mf}}(a, s) := \mathcal{R}_{\text{red}}(a, s, R = ss^\top) = \frac{1}{2} \|\varphi\|_{L^2}^2 - \frac{1}{m} \sum_{i=1}^m a_i V(s_i) + \frac{1}{2m^2} \sum_{i,j=1}^m a_i a_j U(s_i s_j). \quad (23)$$

Further, we denote by $\{a_i^{\text{mf}}(t), s_i^{\text{mf}}(t)\}_{i=1}^m$ the solution to the following ODEs:

$$\begin{aligned} \varepsilon \partial_t a_i &= V(s_i) - \frac{1}{m} \sum_{j=1}^m a_j U(s_i s_j), \\ \partial_t s_i &= a_i \left(1 - s_i^2\right) \left(V'(s_i) - \frac{1}{m} \sum_{j=1}^m a_j U'(s_i s_j) s_j \right). \end{aligned} \quad (24)$$

Note that (24) would be identical to (16)-(17) if we had $r_{ij} = s_i s_j$. A priori, this is not the case. However, the two systems of equations are close to each other for large m as made precise by our next proposition, which formalizes reduction (ii).

The intuitive explanation for the approximation $r_{ij} \approx s_i s_j$ is quite interesting. For large m , due to ‘propagation of chaos’, the neuron weights $\{(u_i, a_i)\}_{i \leq m}$ are approximately independent. Further, because of the symmetry of the problem under rotations that keep u_* fixed, weights $(u_i)_{i \leq n}$ are approximately uniformly distributed conditional on $s_i = \langle u_i, u_* \rangle$. As a consequence, decomposing $u_i = s_i u_* + u_i^\perp$, we have $r_{ij} = s_i s_j + \langle u_i^\perp, u_j^\perp \rangle$, with u_i^\perp, u_j^\perp approximately uniform on $\text{span}\{u_*\}^\perp$ and independent. Therefore, in high dimensions we have $r_{ij} \approx s_i s_j$.

Proposition 3 (Reduction to flow in \mathbb{R}^{2m}). *Let $(a_i(t), s_i(t), r_{ij}(t))_{1 \leq i < j \leq m}$ be the unique solution of the ODEs (16)-(19) with initialization $s_i(0) = 0$, $r_{ij}(0) = 0$ for all $1 \leq i \neq j \leq m$. Let $(a_i^{\text{mf}}(t), s_i^{\text{mf}}(t))_{i \leq m}$ be the unique solution of the ODEs (24) with initialization $s_i^{\text{mf}}(0) = 0$, $a_i^{\text{mf}}(0) = a_i(0)$ for all $i \leq m$.*

If assumptions A1-A3 hold, then for any $T < \infty$ there exists a constant

$$C(T) = M \exp(MT(1+T)^2/\varepsilon^2) \quad (25)$$

(with M depending on the constants $\{M_i\}_{1 \leq i \leq 3}$ appearing in Assumptions A1-A3 only) such that:

$$\sup_{t \in [0, T]} \frac{1}{m} \sum_{i=1}^m \|(a_i(t), s_i(t)) - (a_i^{\text{mf}}(t), s_i^{\text{mf}}(t))\|_2^2 \leq \frac{C(T)}{m}.$$

Consequently,

$$\sup_{t \in [0, T]} |\mathcal{R}_{\text{red}}(a(t), s(t), R(t)) - \mathcal{R}_{\text{mf}}(a^{\text{mf}}(t), s^{\text{mf}}(t))| \leq \frac{C(T)}{\sqrt{m}}.$$

The proof of this proposition is deferred to Appendix B.3. Now, combining the propositions and corollaries in this section, we deduce that with high probability over the i.i.d. initialization,

$$\sup_{t \in [0, T]} |\mathcal{R}(a(t), u(t)) - \mathcal{R}_{\text{mf}}(a^{\text{mf}}(t), s^{\text{mf}}(t))| \leq \left(\frac{1}{\sqrt{d}} + \frac{1}{\sqrt{m}} \right) CM \exp(MT(1+T)^2/\varepsilon^2). \quad (26)$$

4.3 Connection with mean field theory

Consider the empirical distributions of the neurons:

$$\hat{\rho}_t := \frac{1}{m} \sum_{i=1}^m \delta_{(a_i(t), s_i(t))}, \quad (27)$$

$$\rho_t := \frac{1}{m} \sum_{i=1}^m \delta_{(a_i^{\text{mf}}(t), s_i^{\text{mf}}(t))}, \quad (28)$$

with $(a_i(t), s_i(t))_{i \leq m}$, $(a_i^{\text{mf}}(t), s_i^{\text{mf}}(t))_{i \leq m}$ as in the statement of Proposition 3, i.e., solving (respectively) Eqs. (16)-(19) and Eq. (24) with initial conditions as given there.

Then, it is immediate to show that ρ_t solves (in weak sense) the following continuity partial differential equation (PDE) (we refer to [2, 42] for the definition of weak solutions and basic properties, and Appendix B.4 for a short derivation.)

$$\partial_t \rho_t(a, s) = -\nabla \cdot (\rho_t \Psi(a, s; \rho_t)) \quad (29)$$

$$:= -(\partial_a (\rho_t \Psi_a(a, s; \rho_t)) + \partial_s (\rho_t \Psi_s(a, s; \rho_t))), \quad (30)$$

where $\Psi = (\Psi_a, \Psi_s)$ is given by

$$\Psi_a(a, s; \rho) = \varepsilon^{-1} \cdot \left(V(s) - \int_{\mathbb{R}^2} a_1 U(ss_1) \rho(da_1, ds_1) \right), \quad (31)$$

$$\Psi_s(a, s; \rho) = a(1 - s^2) \cdot \left(V'(s) - \int_{\mathbb{R}^2} a_1 s_1 U'(ss_1) \rho(da_1, ds_1) \right). \quad (32)$$

This equation can be extended to a flow in the whole space $(\mathcal{P}(\mathbb{R}^2), W_2)$ (all probability measures on \mathbb{R}^2 equipped with the second Wasserstein distance), and interpreted as gradient flow with respect to this metric in the following risk:

$$\mathcal{R}_{\text{mf},*}(\rho) := \frac{1}{2} \|\varphi\|_{L^2}^2 - \int aV(s) \rho(da, ds) + \frac{1}{2} \int a_1 a_2 U(s_1 s_2) \rho(da_1, ds_1) \rho(da_2, ds_2), \quad (33)$$

which is the obvious extension of $\mathcal{R}_{\text{mf}}(a, s)$ of Eq. (23) to general probability distributions. Proposition 3 implies that for any $T < \infty$, and under the above initial conditions,

$$\sup_{t \in [0, T]} W_2(\rho_t, \hat{\rho}_t) \leq \sqrt{\frac{M \exp(MT(1+T)^2/\varepsilon^2)}{m}}. \quad (34)$$

If we further denote by ρ_t^d the empirical distribution of $(a_i(t), s_i(t))$, $i \leq m$, when $s_i(0) = \langle u_i(0), u_* \rangle$, $u_i(0) \sim \text{Unif}(\mathbb{S}^{d-1})$, a further application of Corollary 1 yields

$$\sup_{t \in [0, T]} W_2(\rho_t^d, \rho_t) \leq \sqrt{\frac{M \exp(MT(1+T)^2/\varepsilon^2)}{m \wedge d}}. \quad (35)$$

Starting with [32, 13, 40], several authors used continuity PDEs of the form (29) to study the learning dynamics of two-layer neural networks. Following the physics tradition, this is referred to as the ‘mean-field theory’ of two-layer neural networks. Appendix B.5 sketches an alternative approach to prove bounds of the form (26), (35) using the results of [32, 33]. The present derivation has the advantages of yielding a sharper bound and of being self-contained.

4.4 A general formulation

As mentioned above, the system of ODEs in Eq. (24) is a special case of the Wasserstein gradient flow of Eq. (29) whereby we set $\rho_0 = m^{-1} \sum_{i=1}^m \delta_{(a_i^{\text{mf}}(0), s_i^{\text{mf}}(0))}$. In order to study the solutions of

Eq. (29) (hence Eq. (24)) we adopt the following framework. Let (Ω, ρ) denote a probability space. Let $a = a(\omega, t)$ and $s = s(\omega, t)$ ($\omega \in \Omega$, $t \geq 0$) be two measurable functions satisfying (dropping dependencies in t below)

$$\begin{aligned}\varepsilon \partial_t a(\omega) &= V(s(\omega)) - \int a(\nu) U(s(\omega)s(\nu)) d\rho(\nu), \\ \partial_t s(\omega) &= a(\omega) \left(1 - s(\omega)^2\right) \left(V'(s(\omega)) - \int a(\nu) U'(s(\omega)s(\nu)) s(\nu) d\rho(\nu) \right).\end{aligned}\tag{36}$$

If $\omega = i \in \Omega = \{1, \dots, m\}$ endowed with the uniform measure, we obtain the equations (24). In general, the push-forward ρ_t of the measure ρ through the map $\omega \in \Omega \mapsto (a(\omega, t), s(\omega, t)) \in \mathbb{R}^2$ satisfies the mean-field equation (29). As a consequence, the dynamics (36) can be viewed as a gradient flow on the risk

$$\mathcal{R}_{\text{mf},*}(\rho) = \frac{1}{2} \|\varphi\|^2 - \int a(\omega) V(s(\omega)) d\rho(\omega) + \frac{1}{2} \int a(\omega_1) a(\omega_2) U(s(\omega_1)s(\omega_2)) d\rho(\omega_1) d\rho(\omega_2). \tag{37}$$

We next characterize the landscape of the risk function $\mathcal{R}_{\text{mf},*}(\rho)$. In particular, we establish that under certain conditions, the global infimum of $\mathcal{R}_{\text{mf},*}(\rho)$ is 0.

Proposition 4. *The risk function $\mathcal{R}_{\text{mf},*}(\rho)$ can be expressed as*

$$\mathcal{R}_{\text{mf},*}(\rho) = \frac{1}{2} \sum_{k=0}^{\infty} \left(\varphi_k - \sigma_k \int a(\omega) s(\omega)^k d\rho(\omega) \right)^2. \tag{38}$$

Assume that $\sigma_k \neq 0$ for all $k \geq 0$, and that

$$\sum_{k=0}^{\infty} \sigma_k^2 < \infty, \quad \sum_{k=0}^{\infty} \varphi_k^2 < \infty.$$

Then, for any $\delta > 0$, there exists a triple (a, s, ρ) such that $a \in L^2(\rho)$, $s \in [-1, 1]$, and $\mathcal{R}_{\text{mf},*}(\rho) \leq \delta^2$.

This proposition is proved in Appendix B.6.

Remark 4.2. Proposition 4 complements Proposition 1 which establishes approximability of the target function $f_*(x) = \varphi(\langle u_*, x \rangle)$ using the networks (2) (Proposition 4 can be seen as an $m = d = \infty$ version of the latter). We note that the proofs of these propositions also provides insight into the structure of approximators. Namely, we can take the weights u_i to be i.i.d. with distribution $\rho(u)du$ that is symmetric under rotations around u_* , and $a_i = \alpha(u_i)$, $a(u) = \alpha(u)\rho(u)$ is concentrated close to $\langle u_*, u \rangle = 0$ (on a scale that can rely on the desired approximation error).

Indeed, the analysis of gradient flow in Section 6 reveals that the solutions found by gradient flow are of this nature. Namely, neurons develop a small but strictly positive alignment with u_* . The distribution and size of the alignment evolves over time.

Remark 4.3. The results in this section can be generalized to multi-index models: $y = \varphi(U_*^\top x)$ where $U_* \in O(d, k)$, the space of $d \times k$ orthogonal matrices. Further, the corresponding limiting dynamics become

$$\begin{aligned}\varepsilon \partial_t a(\omega) &= V(s(\omega)) - \int a(\nu) U(s(\omega)^\top s(\nu)) d\rho(\nu), \\ \partial_t s(\omega) &= a(\omega) \left(I_k - s(\omega) s(\omega)^\top \right) \left(\nabla V(s(\omega)) - \int a(\nu) U'(s(\omega)^\top s(\nu)) s(\nu) d\rho(\nu) \right).\end{aligned}$$

Here, $s(\omega) \in \mathbb{R}^k$ represents $U_*^\top u(\omega)$, and for $s \in \mathbb{R}^k$, $\|s\|_2 \leq 1$:

$$V(s) = \mathbb{E}[\varphi(G)\sigma(G_s)], \quad (G, G_s) \sim \mathcal{N}\left(0, \begin{bmatrix} I_k & s \\ s^\top & 1 \end{bmatrix}\right).$$

The definition of U is the same as before.

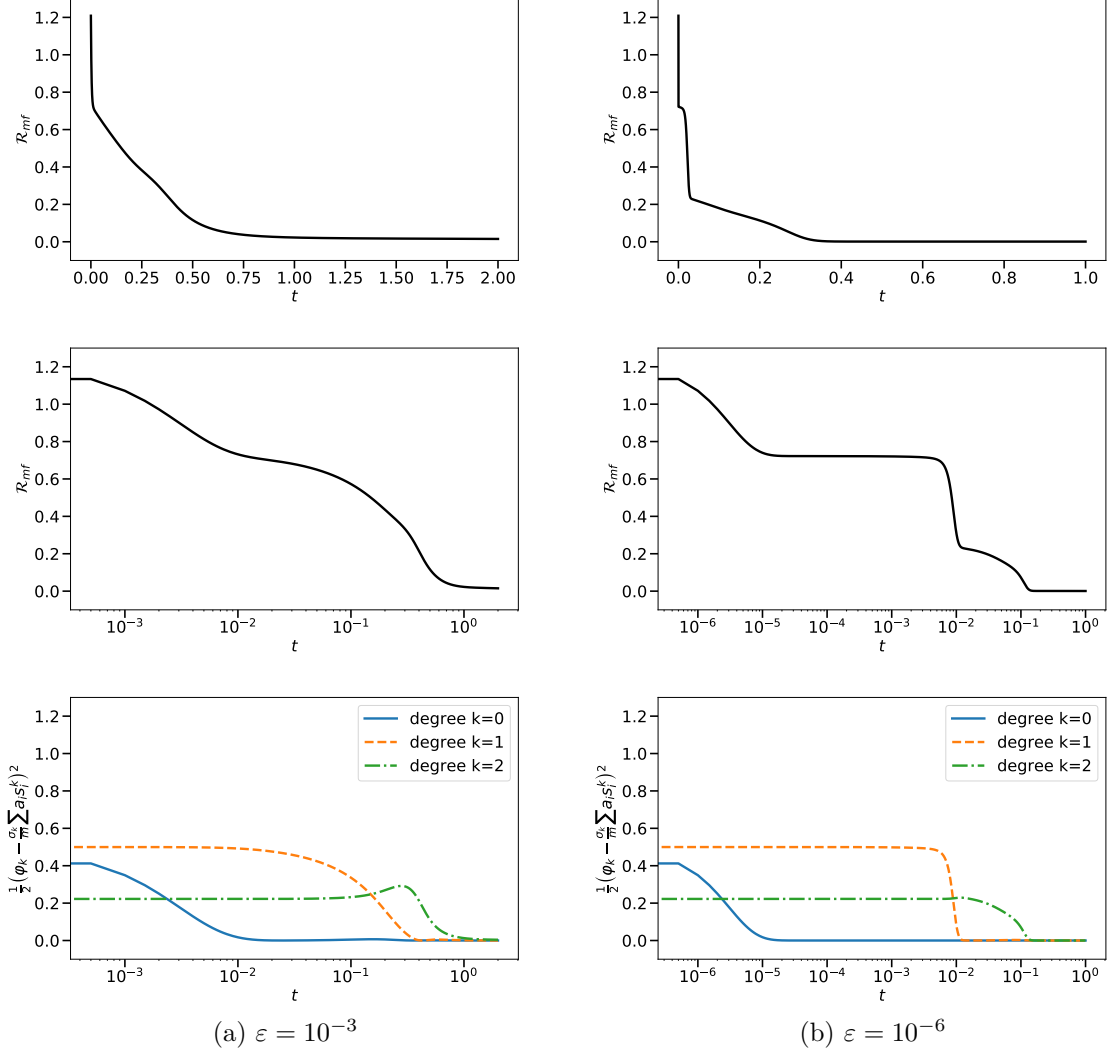


Figure 2: Simulation of the simplified neuron dynamics of Eqs. (24), with the target function of Eq. (39) and ReLU activations. We use learning rate ratios $\varepsilon = 10^{-3}$ (left) and $\varepsilon = 10^{-6}$ (right) and we use $m = 10$ neurons. First two rows: evolution of the risk \mathcal{R}_{mf} of Eq. (23), in linear and log-scales. Third row: evolution of the first three terms of the sum of (40).

5 Numerical solution

In Figure 2, we present the result of an Euler discretization of Eqs. (24) where φ is a degree-2 polynomial and σ is the ReLU activation: $\sigma(s) = \max(s, 0)$,

$$\begin{aligned} \varphi(s) &= \text{He}_0(s) - \text{He}_1(s) - \frac{2}{3}\text{He}_2(s) \\ &= \left(1 - \frac{2\sqrt{2}}{6}\right) - s - \frac{2\sqrt{2}}{6}s^2. \end{aligned} \quad (39)$$

These plots clearly display two of the features emphasized in the introduction: (i) plateaus separated by periods of rapid improvement of the risk; (ii) increasingly long timescales (notice the logarithmic time axis in the second and third row).

In order to examine the incremental learning structure, we rewrite the risk \mathcal{R}_{mf} of Eq. (23)




Figure 3: Same simulation as in Figure 2 (b). In these plots, we show the evolution of the a_i and the s_i for $i \in \{1, \dots, m\}$ following a discretization of Eqs. (24).

by decomposing φ and σ in the basis of Hermite polynomials

$$\mathcal{R}_{\text{mf}}(a, s) = \frac{1}{2} \sum_{k \geq 0} \left(\varphi_k - \frac{\sigma_k}{m} \sum_{i=1}^m a_i s_i^k \right)^2. \quad (40)$$

We observe that, for small ε , the Hermite coefficients of φ are learned sequentially, in the order of their degree. When ε is sufficiently small (right plots), this incremental learning happens in well separated phases. The plateaus and waterfalls in the plots of \mathcal{R}_{mf} correspond to the network learning increasingly higher degree polynomials.

In Figure 3 we plot the evolution of the values of the a_i and s_i , for $i \in \{1, \dots, m\}$. We observe that the overall order of magnitude of the a_i 's and the s_i 's increases when passing through the different phases of the incremental learning process. In the mean time, some of the a_i 's and s_i 's will undergo a sign change during the learning process, which is characterized by a sudden decrease and subsequent rapid increase in its magnitude.

Altogether, the results of Figures 2 and 3 are consistent with the canonical learning order up to level $L = 2$ as per Definition 1. While we conjecture that incremental learning also occurs for higher-order polynomials, we found this hard to observe in numerical simulations: we would need to take ε much smaller than in Figure 2, resulting in prohibitively large simulation costs.

First, as predicted in Definition 1, the times at which the components are learned are closer on a logarithmic scale as the degree increases. It is therefore increasingly difficult to observe time scales corresponding to higher degrees.

Second, we expect there to be a choice of the initialization $(a_{i,\text{init}}, u_{i,\text{init}})_{i \in [m]}$, activation and target function, for which not all the components of φ are actually learnt. We observed empirically that this happens easily for small m .

To conclude this section, in Figure 4 we compare the simplified neuron dynamics (MF) of Eq. (24) and the evolution of projected gradient descent for the original two-layer neural network (NN). From the plots we observe two remarkable phenomena: (1) the evolution of the risk for NN and MF are close to each other during the entire learning process for both large learning rate ratio ($\varepsilon = 1$) and small ($\varepsilon = 10^{-3}$), and their risk curves have the same qualitative behavior even if m is small ($m = 10$); (2) as we increase the value of m from 10 to 50, the alignment between the learning curves of NN and MF improves significantly. These observations justify our argument in Section 4.2 that the inter-neuron correlations r_{ij} are well approximated by $s_i s_j$ for wide networks.

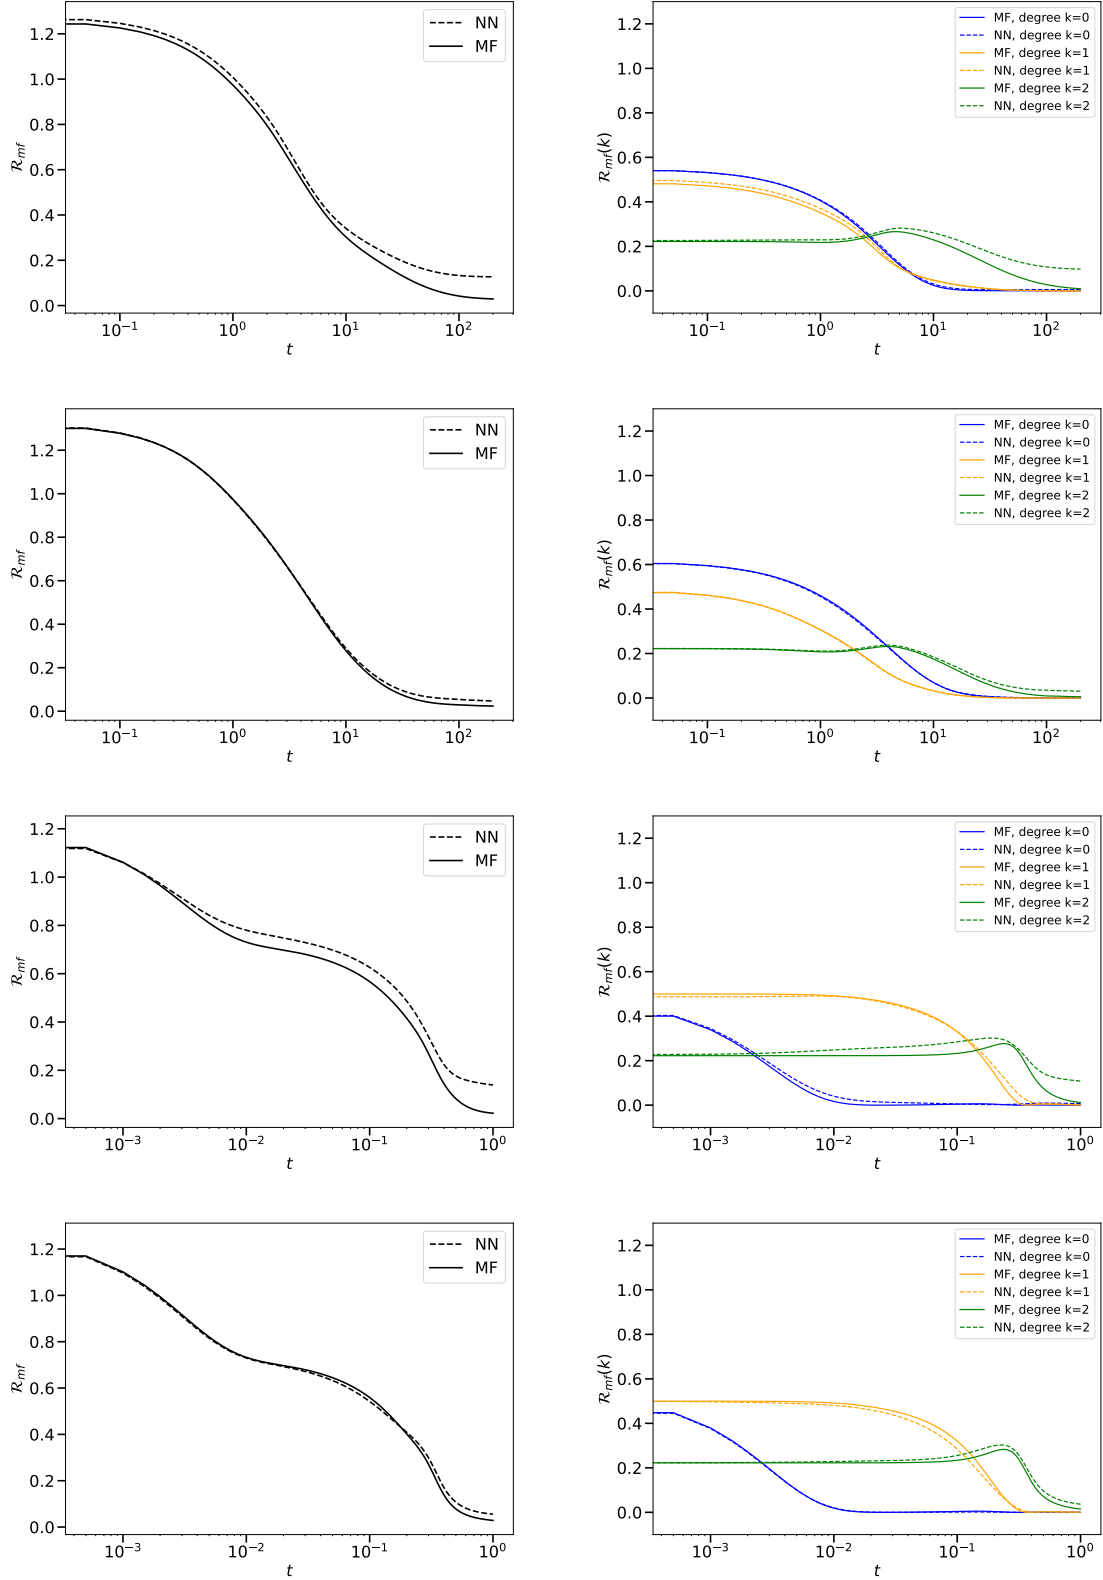


Figure 4: Comparison between the simplified neuron dynamics (24) (MF) and projected gradient descent for the two-layer neural network (2) (NN), with the same target function and activation as the simulations in Figure 2. We use four different combinations of (learning rate ratio, network width): $(\varepsilon, m) = (1, 10)$ (first row), $(\varepsilon, m) = (1, 50)$ (second row), $(\varepsilon, m) = (10^{-3}, 10)$ (third row), and $(\varepsilon, m) = (10^{-3}, 50)$ (fourth row). Left panel: evolution of the risk \mathcal{R}_{mf} for NN and MF on a logarithmic scale. Right panel: evolution of the first three components of \mathcal{R}_{mf} (constant, linear, and quadratic) for NN and MF.

6 Timescales hierarchy in the gradient flow dynamics

We are interested in the behavior of the solution of the ODEs (36), initialized from $s(\omega, 0) = 0$ for all ω (as per Proposition 3). The canonical learning order of Definition 1 concerns the behavior of solutions for $\varepsilon \rightarrow 0$. This type of questions can be addressed within the theory of dynamical systems using *singular perturbation theory* [25]. Here, ‘singular’ refers to the fact that ε multiplies one of the highest-order derivatives. In Eq. (36), ε multiplies the differential term $\partial_t a(\omega)$, so that the ODE system becomes singular in the limit $\varepsilon \rightarrow 0$. In particular, it degenerates to the following system of differential-algebraic equations:

$$\begin{aligned} V(s(\omega)) &= \int a(\nu)U(s(\omega)s(\nu))d\rho(\nu), \\ \partial_t s(\omega) &= a(\omega) \left(1 - s(\omega)^2\right) \left(V'(s(\omega)) - \int a(\nu)U'(s(\omega)s(\nu))s(\nu)d\rho(\nu) \right). \end{aligned} \quad (41)$$

Due to singularity, the qualitative behavior of the above system is dramatically different from that of Eq. (36) with ε small but non-zero. This is in stark contrast to regular perturbation problems, for which the limiting dynamics will still be a system of differential equations with the same order and similar qualitative behavior as the perturbed system.

As a side remark, we note that the system (36) can be seen as a slow-fast dynamical system, where the $a(\omega)$ ’s are the fast variables and the $s(\omega)$ ’s are the slow variables [10]. Formally, the time derivative of the $a(\omega)$ ’s is multiplied by a factor $(1/\varepsilon)$. From a dynamical systems perspective, the present case is made complicated because of a bifurcation when the $s(\omega)$ ’s become non-zero.

The canonical learning order provides a detailed description of this bifurcation. We will motivate this scenario using a classical, but non-rigorous, technique of singular perturbation theory, called the *matched asymptotic expansion* [25, Chapter 2]. This technique decomposes the approximation of the solution in several time scales on which a regular approximation holds. These time scales are traditionally called *layers* in the literature; however, we avoid this terminology due to the potential confusion with the layers of the neural network.

We will work mainly using the Hermite representation of the dynamical ODEs (36), which we write down for the reader’s convenience:

$$\begin{aligned} \varepsilon \partial_t a(\omega) &= \sum_{k=0}^{\infty} \sigma_k s(\omega)^k \left(\varphi_k - \sigma_k \int a(\nu) s(\nu)^k d\rho(\nu) \right), \\ \partial_t s(\omega) &= a(\omega) \left(1 - s(\omega)^2\right) \sum_{k=1}^{\infty} k \sigma_k s(\omega)^{k-1} \left(\varphi_k - \sigma_k \int a(\nu) s(\nu)^k d\rho(\nu) \right). \end{aligned} \quad (42)$$

The rest of this section is organized as follows. We first give a brief overview of the method of matched asymptotic expansions and a summary of our main results regarding the learning timescales in Section 6.1. Sections 6.2-6.4 respectively describe the first three time scales of the matched asymptotic expansion of (42). This gives, for each time scale, an approximation of the $a(\omega)$, $s(\omega)$. In Appendix C.2, we detail how these sections induce an evolution of the risk alternating plateaus and rapid decreases, and support the standing learning scenario of Definition 1. Finally, in Section 6.5, we conjecture the behavior on longer time scales.

Notations. We denote $\mathbb{1}$ the constant function $\mathbb{1} : \omega \in \Omega \mapsto 1 \in \mathbb{R}$. Denote $\langle \cdot, \cdot \rangle_{L^2(\rho)}$ the dot product on $L^2(\rho)$ and $\|\cdot\|_{L^2(\rho)}$ the associated norm. For $x \in L^2(\rho)$, we denote x_{\perp} the orthogonal projection of x on the hyperplane $\mathbb{1}^{\perp}$ of $L^2(\rho)$ of functions orthogonal to $\mathbb{1}$:

$$x_{\perp}(\omega) = x(\omega) - \int x(\nu) d\rho(\nu).$$

We denote $a_{\text{init}}(\omega) = a(\omega, 0)$ and thus $a_{\perp, \text{init}}$ is the orthogonal projection of a_{init} on $\mathbb{1}^{\perp}$.

6.1 Matched asymptotic expansions

The method of matched asymptotic expansions is a common approach to finding approximate solutions of perturbed differential equations. In the present paper, we are mainly interested in applying this technique to approximate the solution to the specific singularly perturbed ODE system¹ of Eq. (42). Denoting by t the independent variable and by ε the perturbation parameter, the method of matched asymptotic expansions consists of the following three steps: (1) Divide the domain of t (generally a subinterval of \mathbb{R}) to several subdomains, which may overlap each other and depend on the perturbation parameter ε ; (2) Within each subdomain, find an accurate approximation to the perturbed system. This is usually achieved by expanding the perturbed system in powers of ε , and keeping only terms that are relevant to the current domain; (3) The approximate solutions obtained in Step (2) might not be valid in the overlap of two adjacent subdomains. To resolve this issue, these approximate solutions are then combined together through a process called “matching” to produce an approximation that is valid on the entire domain.

In our setting, the singularly perturbed system (42) takes the form of

$$\begin{aligned}\varepsilon \partial_t a(\omega) &= f(a(\omega), s(\omega)), \\ \partial_t s(\omega) &= g(a(\omega), s(\omega)).\end{aligned}$$

We will carry out explicit calculations for the first three time scales in Sections 6.2-6.4, respectively. Here is a summary of our main findings:

- In Section 6.2 we explore the learning of the constant component of the target function, which happens at the timescale $t = \Theta(\varepsilon)$. At the end of this phase, the mean-field risk (see (37)) evolves to

$$\mathcal{R}_{\text{mf},*} = \frac{1}{2} \sum_{k \geq 1} \varphi_k^2 + O(\varepsilon). \quad (43)$$

In other words, during this phase, gradient flow learns the constant term in φ . At the end of this time scale we have $a(\omega) = \Theta(1)$ and $s(\omega) = \Theta(\varepsilon)$.

- Then, in Section 6.3 we investigate the second time scale $t = t_2 \varepsilon^{1/2}$, $t_2 \leq c \log(1/\varepsilon)$, during which the $a(\omega)$ ’s and $s(\omega)$ ’s increase to a different order in ε . The result of this time scale is mainly technical and needed to understand the transition to the time scale of Section 6.4. We also perform the matching procedure to combine the approximate solution within this time scale to the one obtained in Section 6.2. At the end of this time scale we have $a(\omega) = \Theta(e^{c't_2})$ and $s(\omega) = \Theta(e^{c't_2} \varepsilon^{1/2})$.
- To understand the evolution of the risk relevant to learning the linear component of φ , we introduce a new time scale $t = \frac{1}{4|\varphi_1 \sigma_1|} \varepsilon^{1/2} \log \frac{1}{\varepsilon} + \Theta(\varepsilon^{1/2})$ in Section 6.4, and show that the linear component can be learned within this time scale. To be more accurate, at the end of this time scale we have

$$\mathcal{R}_{\text{mf},*} = \frac{1}{2} \sum_{k \geq 2} \varphi_k^2 + O(\varepsilon^{1/4}), \quad (44)$$

and $a(\omega) = \Theta(\varepsilon^{-1/4})$ and $s(\omega) = \Theta(\varepsilon^{1/4})$.

Finally, in Section 6.5, we conjecture the behavior of the approximate solutions and induced risks for longer time scales.

¹Although we keep calling this an ODE system, it is important to keep in mind that it takes place in an infinite-dimensional space.

6.2 First time scale: constant component

We define a “fast” time variable $t_1 = t/\varepsilon$ and replace it in Eq. (42). We expand the solutions $a(\omega)$ and $s(\omega)$ in powers of ε :

$$a(\omega) = a^{(0)}(\omega) + \varepsilon a^{(1)}(\omega) + \varepsilon^2 a^{(2)}(\omega) + \dots, \quad (45)$$

$$s(\omega) = s^{(0)}(\omega) + \varepsilon s^{(1)}(\omega) + \varepsilon^2 s^{(2)}(\omega) + \dots, \quad (46)$$

where $a^{(0)}(\omega), a^{(1)}(\omega), a^{(2)}(\omega), \dots, s^{(0)}(\omega), s^{(1)}(\omega), s^{(2)}(\omega), \dots$ are implicitly functions of t_1 . They are initialized at

$$a^{(0)}(\omega, t_1 = 0) = a_{\text{init}}(\omega), \quad a^{(1)}(\omega, t_1 = 0) = 0, \quad a^{(2)}(\omega, t_1 = 0) = 0, \quad \dots \quad (47)$$

$$s^{(0)}(\omega, t_1 = 0) = 0, \quad s^{(1)}(\omega, t_1 = 0) = 0, \quad s^{(2)}(\omega, t_1 = 0) = 0, \quad \dots \quad (48)$$

to be consistent with the initial condition $a(\omega, t_1 = 0) = a(\omega, t = 0) = a_{\text{init}}(\omega)$ and $s(\omega, t_1 = 0) = s(\omega, t = 0) = 0$.

We substitute the expansion in (42):

$$\partial_{t_1} a^{(0)}(\omega) + \varepsilon \partial_{t_1} a^{(1)}(\omega) + \dots \quad (49)$$

$$= \sum_{k=0}^{\infty} \sigma_k \left(s^{(0)}(\omega) + \varepsilon s^{(1)}(\omega) + \dots \right)^k \quad (50)$$

$$\times \left(\varphi_k - \sigma_k \int \left(a^{(0)}(\nu) + \varepsilon a^{(1)}(\nu) + \dots \right) \left(s^{(0)}(\nu) + \varepsilon s^{(1)}(\nu) + \dots \right)^k d\rho(\nu) \right), \quad (51)$$

$$\partial_{t_1} s^{(0)}(\omega) + \varepsilon \partial_{t_1} s^{(1)}(\omega) + \dots \quad (52)$$

$$= \varepsilon \left(a^{(0)}(\omega) + \varepsilon a^{(1)}(\omega) + \dots \right) \left(1 - \left(s^{(0)}(\omega) + \varepsilon s^{(1)}(\omega) + \dots \right)^2 \right) \quad (53)$$

$$\times \sum_{k=1}^{\infty} k \sigma_k \left(s^{(0)}(\omega) + \varepsilon s^{(1)}(\omega) + \dots \right)^{k-1} \quad (54)$$

$$\times \left(\varphi_k - \sigma_k \int \left(a^{(0)}(\nu) + \varepsilon a^{(1)}(\nu) + \dots \right) \left(s^{(0)}(\nu) + \varepsilon s^{(1)}(\nu) + \dots \right)^k d\rho(\nu) \right). \quad (55)$$

The basic assumption of matched asymptotic expansions is that terms of the same order in ε can be identified (with some limitations that we develop below). For now, let us identify terms of order $1 = \varepsilon^0$:

$$\partial_{t_1} a^{(0)}(\omega) = \sum_{k=0}^{\infty} \sigma_k \left(s^{(0)}(\omega) \right)^k \left(\varphi_k - \sigma_k \int a^{(0)}(\nu) \left(s^{(0)}(\nu) \right)^k d\rho(\nu) \right), \quad (56)$$

$$\partial_{t_1} s^{(0)}(\omega) = 0. \quad (57)$$

From (57) and (48), we have $s^{(0)}(\omega) = 0$: time $t_1 = O(1) \Leftrightarrow t = O(\varepsilon)$ is too short for the $s(\omega)$ to be of order 1.

Substituting $s^{(0)}(\omega) = 0$ in (56), we obtain

$$\partial_{t_1} a^{(0)}(\omega) = \sigma_0 \left(\varphi_0 - \sigma_0 \int a^{(0)}(\nu) d\rho(\nu) \right). \quad (58)$$

Recall that $\langle \cdot, \cdot \rangle_{L^2(\rho)}$ is the dot product on $L^2(\rho)$, $\mathbb{1}$ denotes the constant function $\mathbb{1} : \omega \in \Omega \mapsto 1 \in \mathbb{R}$ and a_{\perp} is the orthogonal projection of a on $\mathbb{1}^{\perp}$. Equation (58) can be rewritten as

$$\begin{aligned} \partial_{t_1} \langle a^{(0)}, \mathbb{1} \rangle_{L^2(\rho)} &= \sigma_0 \left(\varphi_0 - \sigma_0 \langle a^{(0)}, \mathbb{1} \rangle_{L^2(\rho)} \right), \\ \partial_{t_1} a_{\perp}^{(0)} &= 0, \end{aligned}$$

which gives after integration (using (47)):

$$\begin{aligned}\langle a^{(0)}, \mathbb{1} \rangle_{L^2(\rho)} &= e^{-\sigma_0^2 t_1} \langle a_{\text{init}}, \mathbb{1} \rangle_{L^2(\rho)} + (1 - e^{-\sigma_0^2 t_1}) \frac{\varphi_0}{\sigma_0}, \\ a_{\perp}^{(0)} &= a_{\perp, \text{init}}.\end{aligned}\tag{59}$$

At this point, we have determined $a^{(0)}(\omega)$ and $s^{(0)}(\omega)$, and thus $a(\omega) = a^{(0)}(\omega) + O(\varepsilon)$ and $s(\omega) = s^{(0)}(\omega) + O(\varepsilon)$ up to a $O(\varepsilon)$ precision, which is sufficient to obtain a $o(1)$ -approximation of the risk $\mathcal{R}_{\text{mf},*}$ (see Section C.2). However, note that we could obtain more precise estimates by identifying higher-order terms in (49)-(55). For instance, identifying the $O(\varepsilon)$ terms in (52)-(55), we obtain $\partial_{t_1} s^{(1)}(\omega) = a^{(0)}(\omega) \sigma_1 \varphi_1$. This shows that the $s(\omega)$ become non-zero, though only of order ε on the time scale $t_1 \asymp 1$; the inner-layer weights develop an infinitesimal correlation with the true direction u_* thanks to the linear component of σ and φ .

The approximation constructed above should be considered as valid on the time scale $t_1 \asymp 1 \Leftrightarrow t \asymp \varepsilon$. As $\varepsilon \rightarrow 0$, we obtain the following approximation of the risk (see Eq. (37) for definition, and Appendix C.2 for a detailed derivation):

$$\mathcal{R}_{\text{mf},*} = \frac{1}{2} e^{-2\sigma_0^2 t_1} \left(\varphi_0 - \sigma_0 \langle a_{\text{init}}, \mathbb{1} \rangle_{L^2(\rho)} \right)^2 + \frac{1}{2} \sum_{k \geq 1} \varphi_k^2 + O(\varepsilon).$$

This approximation breaks down when we reach a new time scale, at which the $s(\omega)$ are large enough for the $a(\omega)$ to be affected (at leading order) by the linear part of the functions. We detail the new time scale and its resolution in the next section.

6.3 Second time scale: linear component I

In this section, we seek a second, slower time scale, for which the behavior of the asymptotic expansion is different.

Identification of the scale. Consider $t_2 = \frac{t}{\varepsilon^\gamma}$, where $\gamma < 1$ is to be determined. We rewrite the system (42) using t_2 , and expand the solutions $a(\omega)$ and $s(\omega)$:

$$a(\omega) = a^{(0)}(\omega) + \varepsilon^\delta a^{(1)}(\omega) + \varepsilon^{2\delta} a^{(2)}(\omega) + \dots,\tag{60}$$

$$s(\omega) = \varepsilon^\delta s^{(1)}(\omega) + \varepsilon^{2\delta} s^{(2)}(\omega) + \dots,\tag{61}$$

where the exponent δ is also to be determined. (Since within the previous time scale we obtained $s(\omega) = O(\varepsilon)$, it is natural to assume $s^{(0)}(\omega) = 0$.)

Let us pause to comment on our method.

Similarly to what has been done in the previous time scale, we will substitute the expansions (60)-(61) in the equations (42) in order to compute the different terms in the expansion. However, this step also allows us to compute the exponents γ and δ , that give respectively the new time scale and the size of the $s(\omega)$'s.

Note that we should have proceeded similarly for the first time scale, by introducing a first time variable $t_1 = \frac{t}{\varepsilon^{\gamma'}}$, expanding $a(\omega), s(\omega)$ in powers $1, \varepsilon^{\delta'}, \varepsilon^{2\delta'}, \dots$, and determining γ' and δ' a posteriori. This would have led, indeed, to $\gamma' = 1$ and $\delta' = 1$. However, for simplicity, we preferred to fix these values that are natural a priori.

Finally, note that the expansions (45)-(46) and (60)-(61) are different, because they are valid on different time scales. In fact, the only coherence conditions that we require below is that the expansions match in a joint asymptotic where $t_1 = \frac{t}{\varepsilon} \rightarrow \infty$ and $t_2 = \frac{t}{\varepsilon^\gamma} \rightarrow 0$. We thus build different approximations for each one of the time scales, with some matching conditions; this justifies the name of *matched asymptotic expansion*.

We now return to our computations and substitute (60)-(61) in (42):

$$\begin{aligned}
\varepsilon^{1-\gamma} \partial_{t_2} a^{(0)}(\omega) + \dots &= \sum_{k=0}^{\infty} \sigma_k \left(\varepsilon^\delta s^{(1)}(\omega) + \dots \right)^k \\
&\quad \times \left(\varphi_k - \sigma_k \int \left(a^{(0)}(\nu) + \dots \right) \left(\varepsilon^\delta s^{(1)}(\nu) + \dots \right)^k d\rho(\nu) \right), \\
\varepsilon^\delta \partial_{t_2} s^{(1)}(\omega) + \dots &= \varepsilon^\gamma \left(a^{(0)}(\omega) + \dots \right) \left(1 - \left(\varepsilon^\delta s^{(1)}(\omega) + \dots \right)^2 \right) \sum_{k=1}^{\infty} k \sigma_k \left(\varepsilon^\delta s^{(1)}(\omega) + \dots \right)^{k-1} \\
&\quad \times \left(\varphi_k - \sigma_k \int \left(a^{(0)}(\nu) + \dots \right) \left(\varepsilon^\delta s^{(1)}(\nu) + \dots \right)^k d\rho(\nu) \right),
\end{aligned}$$

and thus

$$\varepsilon^{1-\gamma} \partial_{t_2} a^{(0)}(\omega) + O(\varepsilon^{1-\gamma+\delta}) = \sigma_0 \left(\varphi_0 - \sigma_0 \int a^{(0)}(\nu) d\rho(\nu) \right) \quad (62)$$

$$- \varepsilon^\delta \sigma_0^2 \int a^{(1)}(\nu) d\rho(\nu) + \varepsilon^\delta \sigma_1 \varphi_1 s^{(1)}(\omega) + O(\varepsilon^{2\delta}), \quad (63)$$

$$\varepsilon^\delta \partial_{t_2} s^{(1)}(\omega) + O(\varepsilon^{2\delta}) = \varepsilon^\gamma \sigma_1 \varphi_1 a^{(0)}(\omega) + O(\varepsilon^{\gamma+\delta}). \quad (64)$$

For the first time scale, we chose $\gamma = \delta = 1$, so that the terms of order ε^δ were negligible compared to $\varepsilon^{1-\gamma} \partial_{t_2} a^{(0)}(\omega)$ in (62). This means that the linear components σ_1, φ_1 of the functions had no effect on the $a(\omega)$ at leading order. We are now interested in a new time scale where $\varepsilon^{1-\gamma} \partial_{t_2} a^{(0)}(\omega)$ and $\varepsilon^\delta \sigma_1 \varphi_1 s^{(1)}(\omega)$ are of the same order, i.e., $1 - \gamma = \delta$; then the linear components play a role in the dynamics.

Further, for $s^{(1)}(\omega)$ to be non-zero, we need both sides of (64) to be of the same order, thus $\delta = \gamma$. Putting together, this gives $\gamma = \delta = 1/2$.

Derivation of the ODEs for this time scale. Let us summarize equations. For $t_2 = \frac{t}{\varepsilon^{1/2}}$ and

$$a(\omega) = a^{(0)}(\omega) + \varepsilon^{1/2} a^{(1)}(\omega) + \dots,$$

$$s(\omega) = \varepsilon^{1/2} s^{(1)}(\omega) + \dots,$$

we have from (62)-(64):

$$\varepsilon^{1/2} \partial_{t_2} a^{(0)}(\omega) = \sigma_0 \left(\varphi_0 - \sigma_0 \int a^{(0)}(\nu) d\rho(\nu) \right) \quad (65)$$

$$- \varepsilon^{1/2} \sigma_0^2 \int a^{(1)}(\nu) d\rho(\nu) + \varepsilon^{1/2} \sigma_1 \varphi_1 s^{(1)}(\omega) + O(\varepsilon), \quad (66)$$

$$\varepsilon^{1/2} \partial_{t_2} s^{(1)}(\omega) = \varepsilon^{1/2} \sigma_1 \varphi_1 a^{(0)}(\omega) + O(\varepsilon). \quad (67)$$

First, we identify the terms of order $1 = \varepsilon^0$:

$$0 = \sigma_0 \left(\varphi_0 - \sigma_0 \int a^{(0)}(\nu) d\rho(\nu) \right). \quad (68)$$

This means that the trajectory remains in the affine hyperplane defined by $\varphi_0 = \sigma_0 \int a^{(0)}(\nu) d\rho(\nu)$. Intuitively, the constant component of φ remains fitted by the neural network in this second time scale.

Second, we identify the terms of order $\varepsilon^{1/2}$ in (65)-(67):

$$\partial_{t_2} a^{(0)}(\omega) = -\sigma_0^2 \int a^{(1)}(\nu) d\rho(\nu) + \sigma_1 \varphi_1 s^{(1)}(\omega), \quad (69)$$

$$\partial_{t_2} s^{(1)}(\omega) = \sigma_1 \varphi_1 a^{(0)}(\omega). \quad (70)$$

Note that, in Eqs. (69)–(70), the time derivative of $a^{(1)}$ does not appear, and therefore the evolution of $a^{(1)}$ is not determined by these equations. In fact, $a^{(1)}$ is best interpreted as the Lagrange multiplier associated to the constraint (68). Namely, this is a free term that can be adjusted so that the solution of the system (69)–(70) satisfies the constraint (68). We can check unknown term in (69) leaves the right degree of freedom such that this is the case: we have

$$0 = \partial_{t_2} \left(\frac{\varphi_0}{\sigma_0} \right) \stackrel{(68)}{=} \partial_{t_2} \left(\int a^{(0)}(\omega) d\rho(\omega) \right) \stackrel{(69)}{=} -\sigma_0^2 \int a^{(1)}(\nu) d\rho(\nu) + \sigma_1 \varphi_1 \int s^{(1)}(\omega) d\rho(\omega).$$

In this last expression, the first unknown term can always compensate the second term so that the constraint is satisfied. The entire evolution of $a^{(1)}$ is determined by higher orders in the expansion.

To eliminate this Lagrange multiplier, we use again the compact notations:

$$\partial_{t_2} a^{(0)} = -\sigma_0^2 \langle a^{(1)}, \mathbb{1} \rangle_{L^2(\rho)} \mathbb{1} + \sigma_1 \varphi_1 s^{(1)}, \quad (71)$$

$$\partial_{t_2} s^{(1)} = \sigma_1 \varphi_1 a^{(0)}, \quad (72)$$

and thus

$$\partial_{t_2} a_{\perp}^{(0)} = \sigma_1 \varphi_1 s_{\perp}^{(1)}, \quad (73)$$

$$\partial_{t_2} s_{\perp}^{(1)} = \sigma_1 \varphi_1 a_{\perp}^{(0)}. \quad (74)$$

Matching. The initialization of the ODEs (71)–(72) for the second time scale is determined by a classical procedure that matches with the previous time scale. In this paragraph, we denote $\underline{a}, \underline{s}$ the approximation obtained in the first time scale (Section 6.2), and \bar{a}, \bar{s} the approximation in the second time scale, described above.

Consider an intermediate time scale $\tilde{t} = \frac{t}{\varepsilon^\alpha}$, $1/2 < \alpha < 1$, and assume $\tilde{t} \asymp 1$ so that

$$t_1 = \frac{t}{\varepsilon} = \frac{\tilde{t}}{\varepsilon^{1-\alpha}} \rightarrow \infty, \quad t_2 = \frac{t}{\varepsilon^{1/2}} = \varepsilon^{\alpha-1/2} \tilde{t} \rightarrow 0.$$

In this intermediate regime, we want the approximations provided on the first and the second time scales to match: $\underline{a}(\tilde{t})$ and $\bar{a}(\tilde{t})$ (resp. $\underline{s}(\tilde{t})$ and $\bar{s}(\tilde{t})$) should match to leading order.

From the first time scale approximation,

$$\underline{a} = \underline{a}^{(0)} + O(\varepsilon) \quad (75)$$

$$= \langle \underline{a}^{(0)}, \mathbb{1} \rangle_{L^2(\rho)} \mathbb{1} + \underline{a}_{\perp}^{(0)} + O(\varepsilon) \quad (76)$$

$$= \left[e^{-\sigma_0^2 t_1} \langle a_{\text{init}}, \mathbb{1} \rangle_{L^2(\rho)} + \left(1 - e^{-\sigma_0^2 t_1} \right) \frac{\varphi_0}{\sigma_0} \right] \mathbb{1} + a_{\perp, \text{init}} + O(\varepsilon) \quad (77)$$

$$= \left[e^{-\sigma_0^2 \tilde{t} / \varepsilon^{1-\alpha}} \langle a_{\text{init}}, \mathbb{1} \rangle_{L^2(\rho)} + \left(1 - e^{-\sigma_0^2 \tilde{t} / \varepsilon^{1-\alpha}} \right) \frac{\varphi_0}{\sigma_0} \right] \mathbb{1} + a_{\perp, \text{init}} + O(\varepsilon) \quad (78)$$

$$= \frac{\varphi_0}{\sigma_0} \mathbb{1} + a_{\perp, \text{init}} + o(1). \quad (79)$$

From the second time scale approximation,

$$\bar{a} = \bar{a}^{(0)}(t_2) + O(\varepsilon^{1/2}) = \bar{a}^{(0)}(\varepsilon^{\alpha-1/2} \tilde{t}) + O(\varepsilon^{1/2}) \quad (80)$$

$$= \bar{a}^{(0)}(0) + o(1). \quad (81)$$

By matching, Equations (79) and (81) should be coherent. Thus the ODE for the second time scale should be initialized from $\bar{a}^{(0)}(0) = \frac{\varphi_0}{\sigma_0} \mathbb{1} + a_{\perp, \text{init}}$.

Similarly, the matching procedure gives that the ODE for the second time scale should be initialized from $\bar{s}^{(1)} = 0$.

Solution. As we are done with the matching procedure, we now consider the solution in the second time scale only, that we denote again by a, s as in (71), (72). The matching procedure motivates us to consider the solution of (73)-(74) initialized at $a_{\perp}^{(0)}(0) = a_{\perp, \text{init}}, s_{\perp}^{(1)} = 0$. This gives

$$a_{\perp}^{(0)} = \cosh(\varphi_1 \sigma_1 t_2) a_{\perp, \text{init}}, \quad s_{\perp}^{(1)} = \sinh(\varphi_1 \sigma_1 t_2) a_{\perp, \text{init}}.$$

To conclude, we note that $\langle a^{(0)}, \mathbb{1} \rangle_{L^2(\rho)} = \frac{\varphi_0}{\sigma_0}$ is constrained by (68). Further, from (70),

$$\partial_{t_2} \langle s^{(1)}, \mathbb{1} \rangle_{L^2(\rho)} = \sigma_1 \varphi_1 \langle a^{(0)}, \mathbb{1} \rangle_{L^2(\rho)} = \sigma_1 \varphi_1 \frac{\varphi_0}{\sigma_0},$$

thus $\langle s^{(1)}, \mathbb{1} \rangle_{L^2(\rho)} = \sigma_1 \varphi_1 \frac{\varphi_0}{\sigma_0} t_2$.

Putting together, these equations give:

$$a^{(0)} = \frac{\varphi_0}{\sigma_0} \mathbb{1} + \cosh(\varphi_1 \sigma_1 t_2) a_{\perp, \text{init}}, \quad s^{(1)} = \sigma_1 \varphi_1 \frac{\varphi_0}{\sigma_0} t_2 \mathbb{1} + \sinh(\varphi_1 \sigma_1 t_2) a_{\perp, \text{init}}. \quad (82)$$

We observe that $a^{(0)}$ and $s^{(1)}$ diverge as $t_2 \rightarrow \infty$. This implies that our approximation on the second time scale must break down at a certain point. Indeed, we analyzed this time scale under the assumption that both $a^{(0)}$ and $s^{(1)}$ are of order 1. However, since $a^{(0)}$ and $s^{(1)}$ diverge exponentially as $t_2 \rightarrow \infty$, as per Eq. (82), this assumption breaks down when $t_2 \asymp \log(1/\varepsilon)$.

More precisely, in (65) (resp. (67)), the $O(\varepsilon)$ term includes a term of the form

$$-\varepsilon s^{(1)}(\omega) \sigma_1^2 \int a^{(0)}(\nu) s^{(1)}(\nu) d\rho(\nu) \quad \left(\text{resp. } -\varepsilon a^{(0)}(\omega) \sigma_1^2 \int a^{(0)}(\nu) s^{(1)}(\nu) d\rho(\nu) \right).$$

When $a^{(0)}$ and $s^{(1)}$ become of order $\varepsilon^{-1/4}$, this term becomes of order $\varepsilon^{1/4}$, which is then of the same order as the term $\varepsilon^{1/2} \sigma_1 \varphi_1 s^{(1)}(\omega)$ in (65) (resp. the term $\varepsilon^{1/2} \sigma_1 \varphi_1 a^{(0)}(\omega)$ in (67)). At this point, these terms can not be neglected anymore. From (82), we have

$$a^{(0)} \sim \frac{e^{|\varphi_1 \sigma_1| t_2}}{2} a_{\perp, \text{init}}, \quad s^{(1)} \sim \text{sign}(\varphi_1 \sigma_1) \frac{e^{|\varphi_1 \sigma_1| t_2}}{2} a_{\perp, \text{init}}, \quad t_2 \rightarrow \infty.$$

Therefore, $a^{(0)}$ and $s^{(1)}$ become of order $\varepsilon^{-1/4}$ at the time $t_2 \sim \frac{1}{4|\varphi_1 \sigma_1|} \log \frac{1}{\varepsilon}$, at which the approximation on the second time scale breaks down. We thus introduce a new time scale centered at this critical point.

6.4 Third time scale: linear component II

We now introduce the time $t_3 = t_2 - \frac{1}{4|\varphi_1 \sigma_1|} \log \frac{1}{\varepsilon}$. As t_3 is only a translation from t_2 , the ODEs in terms of t_3 are the same as the ones in term of t_2 . However, in this time scale, a and $\varepsilon^{1/2} s$ have diverged. In coherence with the discussion above, we seek expansions of the form

$$a = \varepsilon^{-1/4} a^{(-1)} + a^{(0)} + \varepsilon^{1/4} a^{(1)} + \dots, \quad (83)$$

$$s = \varepsilon^{1/4} s^{(1)} + \varepsilon^{1/2} s^{(2)} + \dots. \quad (84)$$

Similarly to the second time scale, we substitute (83)-(84) in (42) and obtain

$$\begin{aligned} \varepsilon^{1/4} \partial_{t_3} a^{(-1)}(\omega) &= -\varepsilon^{-1/4} \sigma_0^2 \int a^{(-1)}(\nu) d\rho(\nu) + \sigma_0 \left(\varphi_0 - \sigma_0 \int a^{(0)}(\nu) d\rho(\nu) \right) \\ &\quad - \varepsilon^{1/4} \sigma_0^2 \int a^{(1)}(\nu) d\rho(\nu) + \varepsilon^{1/4} \sigma_1 \left(\varphi_1 - \sigma_1 \int a^{(-1)}(\nu) s^{(1)}(\nu) d\rho(\nu) \right) s^{(1)}(\omega) + O(\varepsilon^{1/2}), \\ \varepsilon^{1/4} \partial_{t_3} s^{(1)}(\omega) &= \varepsilon^{1/4} \sigma_1 \left(\varphi_1 - \sigma_1 \int a^{(-1)}(\nu) s^{(1)}(\nu) d\rho(\nu) \right) a^{(-1)}(\omega) + O(\varepsilon^{1/2}). \end{aligned}$$

First, we identify the terms of order $\varepsilon^{-1/4}$:

$$0 = -\sigma_0^2 \int a^{(-1)}(\nu) d\rho(\nu) = -\sigma_0^2 \langle a^{(-1)}, \mathbf{1} \rangle_{L^2(\rho)}. \quad (85)$$

This means that a has no component diverging in ε in the direction of $\mathbf{1}$.

Second, we identify the terms of order $1 = \varepsilon^0$:

$$0 = \sigma_0 \left(\varphi_0 - \sigma_0 \int a^{(0)}(\nu) d\rho(\nu) \right) = \sigma_0 \left(\varphi_0 - \sigma_0 \langle a^{(0)}, \mathbf{1} \rangle_{L^2(\rho)} \right). \quad (86)$$

Put together with (85), this equation ensures that the constant component of φ remains learned on this third time scale.

Third, we identify the terms of order $\varepsilon^{1/4}$:

$$\begin{aligned} \partial_{t_3} a^{(-1)}(\omega) &= -\sigma_0^2 \int a^{(1)}(\nu) d\rho(\nu) + \sigma_1 \left(\varphi_1 - \sigma_1 \int a^{(-1)}(\nu) s^{(1)}(\nu) d\rho(\nu) \right) s^{(1)}(\omega), \\ \partial_{t_3} s^{(1)}(\omega) &= \sigma_1 \left(\varphi_1 - \sigma_1 \int a^{(-1)}(\nu) s^{(1)}(\nu) d\rho(\nu) \right) a^{(-1)}(\omega). \end{aligned} \quad (87)$$

Again, the term $-\sigma_0^2 \int a^{(1)}(\nu) d\rho(\nu)$ is best interpreted as the Lagrange multiplier associated to the constraints (85), (86). Using the compact notations,

$$\begin{aligned} \int a^{(-1)}(\nu) s^{(1)}(\nu) d\rho(\nu) &= \langle a^{(-1)}, s^{(1)} \rangle_{L^2(\rho)} = \langle a^{(-1)}, \mathbf{1} \rangle_{L^2(\rho)} \langle \mathbf{1}, s^{(1)} \rangle_{L^2(\rho)} + \langle a_{\perp}^{(-1)}, s_{\perp}^{(1)} \rangle_{L^2(\rho)} \\ &= \langle a_{\perp}^{(-1)}, s_{\perp}^{(1)} \rangle_{L^2(\rho)}, \end{aligned}$$

where in the last equality we use (85). Thus we can rewrite (87) as

$$\begin{aligned} \partial_{t_3} a^{(-1)} &= -\sigma_0^2 \langle a^{(1)}, \mathbf{1} \rangle_{L^2(\rho)} \mathbf{1} + \sigma_1 \left(\varphi_1 - \sigma_1 \langle a_{\perp}^{(-1)}, s_{\perp}^{(1)} \rangle_{L^2(\rho)} \right) s^{(1)}, \\ \partial_{t_3} s^{(1)} &= \sigma_1 \left(\varphi_1 - \sigma_1 \langle a_{\perp}^{(-1)}, s_{\perp}^{(1)} \rangle_{L^2(\rho)} \right) a^{(-1)}, \end{aligned} \quad (88)$$

and thus

$$\begin{aligned} \partial_{t_3} a_{\perp}^{(-1)} &= \sigma_1 \left(\varphi_1 - \sigma_1 \langle a_{\perp}^{(-1)}, s_{\perp}^{(1)} \rangle_{L^2(\rho)} \right) s_{\perp}^{(1)}, \\ \partial_{t_3} s_{\perp}^{(1)} &= \sigma_1 \left(\varphi_1 - \sigma_1 \langle a_{\perp}^{(-1)}, s_{\perp}^{(1)} \rangle_{L^2(\rho)} \right) a_{\perp}^{(-1)}. \end{aligned} \quad (89)$$

In Appendix C.1, we solve this system of ODEs and determine the initial condition by matching with the previous layer. The result is that

$$\begin{aligned} a^{(-1)} &= a_{\perp}^{(-1)} = \lambda a_{\perp, \text{init}}, \\ s^{(1)} &= s_{\perp}^{(1)} = \text{sign}(\sigma_1 \varphi_1) \lambda a_{\perp, \text{init}}, \end{aligned} \quad (90)$$

where $\lambda = \lambda(t_3)$ is the function

$$\lambda(t_3) = \frac{|\varphi_1|^{1/2}}{\left(|\sigma_1| \|a_{\perp, \text{init}}\|_{L^2(\rho)}^2 + 4|\varphi_1| e^{-2|\sigma_1 \varphi_1| t_3} \right)^{1/2}}. \quad (91)$$

This solution finishes to describe how the linear part of the function φ is learned. Plugging it into the equations for $a^{(-1)}$ and $s^{(1)}$, we get

$$\sigma_1 \int a^{(-1)}(\nu) s^{(1)}(\nu) d\rho(\nu) = \sigma_1 \langle a_{\perp}^{(-1)}, s_{\perp}^{(1)} \rangle_{L^2(\rho)} = \frac{\varphi_1 |\sigma_1| \|a_{\perp, \text{init}}\|_{L^2(\rho)}^2}{|\sigma_1| \|a_{\perp, \text{init}}\|_{L^2(\rho)}^2 + 4|\varphi_1| e^{-2|\sigma_1 \varphi_1| t_3}},$$

which converges to φ_1 as $t_3 \rightarrow \infty$. Consequently, we obtain the following approximation for $\mathcal{R}_{\text{mf},*}$ within this time scale (again, see Appendix C.2 for details):

$$\mathcal{R}_{\text{mf},*} = \frac{1}{2} \varphi_1^2 \left(1 - \frac{1}{1 + \frac{4|\varphi_1|}{\|\sigma_1\| \|a_{\perp, \text{init}}\|_{L^2(\rho)}^2} e^{-2|\sigma_1 \varphi_1| t_3}} \right)^2 + \frac{1}{2} \sum_{k \geq 2} \varphi_k^2 + O(\varepsilon^{1/4}), \quad \varepsilon \rightarrow 0.$$

6.5 Conjectured behavior for larger time scales

The analysis of the previous sections naturally suggests the existence of a sequence of cutoffs. At each time scale, a new polynomial component of φ is learned within a window that is much shorter than the time elapsed before that phase started. Along this sequence, we expect s and a to grow to increasingly larger scales in ε (but s remains $o(1)$ while a diverges).

More precisely, we assume that during the l -th phase, the network learns the degree- l component φ_l , and various quantities satisfy the following scaling behavior:

$$a = O(\varepsilon^{-\omega_l}), \quad s = O(\varepsilon^{\beta_l}), \quad t = O(\varepsilon^{\mu_l}), \quad (92)$$

where $\omega_l > 0$ is an increasing sequence and $\beta_l, \mu_l > 0$ are decreasing sequences. Further, while learning of this component takes place when $t = O(\varepsilon^{\mu_l})$, the actual evolution of the risk (and of the neural network) take place on much shorter scales, namely:

$$\Delta t = O(\varepsilon^{\nu_l}), \quad (93)$$

where ν_l is also decreasing, with $\nu_l > \mu_l$. The goal of this section is to provide heuristic arguments to conjecture the values of ω_l, β_l, μ_l and ν_l . We will base this conjecture on a rigorous analysis of a simplified model.

The simplified model is motivated by the expectation (supported by the heuristics and simulations in the previous sections) that learning each component happens independently from the details of the evolution on previous time scales. In the simplified model, the activation function $\sigma(x)$ is proportional to the l -th Hermite polynomial, namely $\sigma(x) = \sigma_l \text{He}_l(x)$. This is the component of σ that we expect to be relevant on the l -th time scale. The gradient flow equations (42) then read:

$$\begin{aligned} \varepsilon \partial_t a(\omega) &= \sigma_l s(\omega)^l \left(\varphi_l - \sigma_l \int a(\nu) s(\nu)^l d\rho(\nu) \right), \\ \partial_t s(\omega) &= a(\omega) (1 - s(\omega)^2) l \sigma_l s(\omega)^{l-1} \left(\varphi_l - \sigma_l \int a(\nu) s(\nu)^l d\rho(\nu) \right). \end{aligned} \quad (94)$$

with corresponding risk component

$$\mathcal{R}_l = \frac{1}{2} \left(\varphi_l - \sigma_l \int a(\nu) s(\nu)^l d\rho(\nu) \right)^2.$$

We capture the effect of learning dynamics on the previous time scales by the overall magnitude of the $a(\omega)$'s and $s(\omega)$'s at initialization. Namely, we choose the scale of initialization of the simplified model to be given by the end of the $(l-1)$ -th time scale, i.e., $a(\omega) \asymp \varepsilon^{-\omega_{l-1}}$ and $s(\omega) \asymp \varepsilon^{\beta_{l-1}}$. Further, in order for the $(l-1)$ -th component to be learned, namely

$$\int a(\nu) s(\nu)^{l-1} d\rho(\nu) \approx \frac{\varphi_{l-1}}{\sigma_{l-1}}, \quad (95)$$

we require $\omega_{l-1} = (l-1)\beta_{l-1}$ so that $\int a(\nu) s(\nu)^{l-1} d\rho(\nu) = \Theta(1)$. Analogously, we assume $\omega_l = l\beta_l$.

Based on this consideration, we introduce the rescaled variables

$$\tilde{a}(\omega) = \varepsilon^{\omega_l} a(\omega), \quad \tilde{s}(\omega) = \varepsilon^{-\beta_l} s(\omega), \quad \text{where } \tilde{a}(\omega, 0) \asymp \varepsilon^{\omega_l - \omega_{l-1}}, \tilde{s}(\omega, 0) \asymp \varepsilon^{\beta_{l-1} - \beta_l}.$$

Rewriting Eq. (94) in terms of $\tilde{a}(\omega)$'s and $\tilde{s}(\omega)$'s, and using $\omega_l = l\beta_l$, we get that

$$\begin{aligned} \varepsilon^{1-2l\beta_l} \partial_t \tilde{a}(\omega) &= \sigma_l \tilde{s}(\omega)^l \left(\varphi_l - \sigma_l \int \tilde{a}(\nu) \tilde{s}(\nu)^l d\rho(\nu) \right) \\ \varepsilon^{2\beta_l} \partial_t \tilde{s}(\omega) &= l\sigma_l \tilde{a}(\omega) \tilde{s}(\omega)^{l-1} \left(1 - \varepsilon^{2\beta_l} \tilde{s}(\omega)^2 \right) \left(\varphi_l - \sigma_l \int \tilde{a}(\nu) \tilde{s}(\nu)^l d\rho(\nu) \right). \end{aligned} \quad (96)$$

In order for the $\tilde{a}(\omega)$'s and $\tilde{s}(\omega)$'s to be learned simultaneously, we need $1 - 2l\beta_l = 2\beta_l$, which implies $\beta_l = 1/2(l+1)$. Making a further change of the time variable $t = \varepsilon^{\nu_l} \tau$, where $\nu_l = 2\beta_l = 1/(l+1)$, it follows that

$$\begin{aligned} \partial_\tau \tilde{a}(\omega) &= \sigma_l \tilde{s}(\omega)^l \left(\varphi_l - \sigma_l \int \tilde{a}(\nu) \tilde{s}(\nu)^l d\rho(\nu) \right) \\ \partial_\tau \tilde{s}(\omega) &= l\sigma_l \tilde{a}(\omega) \tilde{s}(\omega)^{l-1} \left(1 - \varepsilon^{2\beta_l} \tilde{s}(\omega)^2 \right) \left(\varphi_l - \sigma_l \int \tilde{a}(\nu) \tilde{s}(\nu)^l d\rho(\nu) \right). \end{aligned} \quad (97)$$

Moreover, rewriting the risk in terms of the rescaled variables \tilde{a}, \tilde{s} , $\mathcal{R}_l(\tau) = \mathcal{R}_l(\tilde{a}(\tau), \tilde{s}(\tau))$ satisfies the ODE:

$$\partial_\tau \mathcal{R}_l = -2\sigma_l^2 \mathcal{R}_l \cdot \int \tilde{s}(\omega)^{2(l-1)} \left(l^2 \tilde{a}(\omega)^2 \left(1 - \varepsilon^{2\beta_l} \tilde{s}(\omega)^2 \right) + \tilde{s}(\omega)^2 \right) d\rho(\omega). \quad (98)$$

Note that with our choice of β_l and ω_l , we have $\omega_l - \omega_{l-1} = \beta_{l-1} - \beta_l = 1/2l(l+1)$. This means that the $\tilde{a}(\omega)$'s and $\tilde{s}(\omega)$'s are initialized at the same scale, namely

$$\tilde{a}(\omega, 0), \tilde{s}(\omega, 0) = \Theta(\varepsilon^{1/2l(l+1)}). \quad (99)$$

The theorem below describes quantitatively the dynamics of the simplified model for small ε , and determines the value of μ_l (recall that $\nu_l = 1/(l+1)$):

Theorem 1 (Evolution of the simplified gradient flow). *Assume $l \geq 2$ and let $(\tilde{a}(\omega, \tau), \tilde{s}(\omega, \tau))_{\tau \geq 0}$ be the unique solution of the ODE system (97), initialized as per Eq. (99) (note in particular that $\sigma_l \varphi_l \tilde{a}(\omega, 0) \tilde{s}(\omega, 0)^l \asymp \varepsilon^{1/2l}$). Then the followings hold:*

(a) *Let us denote*

$$A = \left\{ \omega : \sigma_l \varphi_l \liminf_{\varepsilon \rightarrow 0} \varepsilon^{-1/2l} \tilde{a}(\omega, 0) \tilde{s}(\omega, 0)^l > 0 \right\} \quad (100)$$

and assume $\rho(A) > 0$. For $\Delta \in (0, \varphi_l^2/2)$, define

$$\tau(\Delta) = \inf \{ \tau \geq 0 : \mathcal{R}_l(\tilde{a}(\tau), \tilde{s}(\tau)) \leq \Delta \}. \quad (101)$$

Then, for any fixed Δ we have $\tau(\Delta) = \Theta(\varepsilon^{-(l-1)/2l(l+1)})$ as $\varepsilon \rightarrow 0$. Further, if ρ is a discrete probability measure, then there exists $\tau_(\varepsilon) = \Theta(\varepsilon^{-(l-1)/2l(l+1)})$ and, for any $\Delta > 0$ a constant $c_*(\Delta) > 0$ independent of ε such that*

$$\tau \leq \tau_*(\varepsilon) - c_*(\Delta) \Rightarrow \liminf_{\varepsilon \rightarrow 0} \mathcal{R}_l(\tilde{a}(\tau), \tilde{s}(\tau)) \geq \frac{1}{2} \varphi_l^2 - \Delta, \quad (102)$$

$$\tau \geq \tau_*(\varepsilon) + c_*(\Delta) \Rightarrow \limsup_{\varepsilon \rightarrow 0} \mathcal{R}_l(\tilde{a}(\tau), \tilde{s}(\tau)) \leq \Delta, \quad (103)$$

namely the l -th component is learnt in an $O(1)$ time window around $\tau_(\varepsilon) = \Theta(\varepsilon^{-(l-1)/2l(l+1)})$.*

(b) Similarly, we denote

$$B = \left\{ \omega : \sigma_l \varphi_l \limsup_{\varepsilon \rightarrow 0} \varepsilon^{-1/2l} \tilde{a}(\omega, 0) \tilde{s}(\omega, 0)^l < 0, \text{ and } \liminf_{\varepsilon \rightarrow 0} (\tilde{s}(\omega, 0)^2 / \tilde{a}(\omega, 0)^2) > l \right\}. \quad (104)$$

If $\rho(B) > 0$, then the same claims as in (a) hold.

(c) If neither of the conditions at points (a), (b) holds, and

$$\sigma_l \varphi_l \limsup_{\varepsilon \rightarrow 0} \varepsilon^{-1/2l} \tilde{a}(\omega, 0) \tilde{s}(\omega, 0)^l < 0, \quad \limsup_{\varepsilon \rightarrow 0} (\tilde{s}(\omega, 0)^2 / \tilde{a}(\omega, 0)^2) < l \quad (105)$$

for almost every $\omega \in \Omega$. Then, for such $\omega \in \Omega$ and each $\Delta > 0$, there exists a constant $C_*(\omega, \Delta) > 0$ such that

$$\tau \geq C_*(\omega, \Delta) \varepsilon^{-(l-1)/2l(l+1)} \Rightarrow |\tilde{s}(\omega, \tau)| \leq \Delta \varepsilon^{1/2l(l+1)}, \quad (106)$$

meaning that $\tilde{s}(\omega, \tau)$ converges to 0 eventually.

We further note that $\tau = \Theta(\varepsilon^{-(l-1)/2l(l+1)}) \iff t = \Theta(\varepsilon^{\mu_l})$ with $\mu_l = 1/2l$, and $\tau = O(1) \iff t = O(\varepsilon^{\nu_l})$ with $\nu_l = 1/(l+1)$.

The proof of Theorem 1 is deferred to Appendix C.3.

Remark 6.1. Under the conditions of cases (a) and (b), we see that the degree- l component of the target function is learnt within an $O(\varepsilon^{1/(l+1)})$ time window around $t_*(l, \varepsilon) \asymp \varepsilon^{1/2l}$, which is consistent with the timescales conjectured in Definition 1.

Remark 6.2. Case (c) corresponds to $s(\omega)/s(\omega, 0)$ becoming close to 0 in time $t = O(\varepsilon^{\mu_l})$, and staying at 0. In other words, the neurons become orthogonal to the target direction and play no role in learning higher-degree components any longer.

Informally, case (c) couples the learning of different polynomial components. It can happen that the learning phase $l-1$ induces an effective initialization $(\tilde{a}(\omega, 0), \tilde{s}(\omega, 0))$ within the domain of case (c).

We expect this not to be the case for suitable choices of initialization (or equivalently P_A), φ , and σ . Establishing this would amount to establishing that the canonical learning order holds.

7 Stochastic gradient descent and finite sample size

So far we focused on analyzing the projected gradient flow (GF) dynamics with respect to the population risk, as defined in Eqs. (5)-(6). In this section, we extract the implications of our analysis of GF on online projected stochastic gradient descent, which is a projected version of the SGD dynamics (162).

For simplicity of notation, we denote by $z = (y, x) \in \mathbb{R} \times \mathbb{R}^d$ a datapoint and by $\theta_i = (a_i, u_i) \in \mathbb{R} \times \mathbb{S}^{d-1}$ the parameters of neuron i . For $z = (y, x)$ and $\rho^{(m)} = (1/m) \sum_{i=1}^m \delta_{\theta_i} = (1/m) \sum_{i=1}^m \delta_{(a_i, u_i)}$, we define

$$\begin{aligned} \hat{F}_i(\rho^{(m)}; z) &= \left(y - \frac{1}{m} \sum_{j=1}^m a_j \sigma(\langle u_j, x \rangle) \right) \sigma(\langle u_i, x \rangle), \\ \hat{G}_i(\rho^{(m)}; z) &= a_i \left(y - \frac{1}{m} \sum_{j=1}^m a_j \sigma(\langle u_j, x \rangle) \right) \sigma'(\langle u_i, x \rangle) x. \end{aligned}$$

The projected SGD dynamics is specified as follows:

$$\begin{aligned} \bar{a}_i(k+1) &= \bar{a}_i(k) + \varepsilon^{-1} \eta \hat{F}_i(\bar{\rho}^{(m)}(k); z_{k+1}) \\ \bar{u}_i(k+1) &= \text{Proj}_{\mathbb{S}^{d-1}} \left(\bar{u}_i(k) + \eta \hat{G}_i(\bar{\rho}^{(m)}(k); z_{k+1}) \right), \end{aligned} \quad (107)$$

where for $u \in \mathbb{R}^d$ and compact $S \subset \mathbb{R}^d$, $\text{Proj}_S(u) := \arg\min_{s \in S} \|s - u\|_2$, and $\bar{\rho}^{(m)} := (1/m) \sum_{i=1}^m \delta_{\bar{\theta}_i}$. Note that the (\bar{a}_i, \bar{u}_i) 's here are different from the (\bar{a}, \bar{s}) 's in Section 6.

We prove that, for small η , the projected SGD of Eq. (107) is close to the gradient flow of Eqs. (5)-(6). Throughout this section, we make the following assumptions similar to those assumed in Section 4:

A1. ρ_0 is supported on $[-M_1, M_1] \times \mathbb{S}^{d-1}$. Hence, $|a_i(0)| \leq M_1$ for all $i \in [m]$.

A2. The activation function is bounded: $\|\sigma\|_\infty \leq M_2$. Additionally, define for $u, u' \in \mathbb{R}^d$:

$$V(\langle u_*, u \rangle; \|u_*\|_2, \|u\|_2) = \mathbb{E} [\varphi(\langle u_*, x \rangle) \sigma(\langle u, x \rangle)], \quad (108)$$

$$U(\langle u, u' \rangle; \|u\|_2, \|u'\|_2) = \mathbb{E} [\sigma(\langle u, x \rangle) \sigma(\langle u', x \rangle)]. \quad (109)$$

We then require the functions V and U to be bounded and differentiable, with uniformly bounded and Lipschitz continuous gradients for all $\|u\|_2, \|u'\|_2 \leq 2$:

$$\|\nabla_u V\|_2 \leq M_2, \quad \|\nabla_u V - \nabla_{u'} V\|_2 \leq M_2 \|u - u'\|_2, \quad (110)$$

$$\|\nabla_{(u, u')} U\|_2 \leq M_2, \quad \|\nabla_{(u, u')} U - \nabla_{(u_1, u'_1)} U\|_2 \leq M_2 (\|u - u_1\|_2 + \|u' - u'_1\|_2). \quad (111)$$

Similar to Remark 4.1, we can show that a sufficient condition for Eq.s (110) and (111) is

$$\sup \{\|\sigma'\|_{L^2}, \|\sigma''\|_{L^2}\} \leq M'_2, \quad \sup \{\|\varphi\|_{L^2}, \|\varphi'\|_{L^2}, \|\varphi''\|_{L^2}\} \leq M'_2,$$

where the constant M'_2 depends uniquely on M_2 .

A3. Assume $(x, y) \sim \mathbb{P}$, then we require that $y \in [-M_3, M_3]$ almost surely. Moreover, we assume that for all $\|u\|_2 \leq 2$, both $\sigma(\langle u, x \rangle)$ and $\sigma'(\langle u, x \rangle)(x - \langle u, x \rangle u)$ are M_3 -sub-Gaussian.

The following theorem upper bounds the distance between gradient flow and projected stochastic gradient descent dynamics.

Theorem 2 (Difference between GF and Projected SGD). *Let $\theta_i(t) = (a_i(t), u_i(t))$ be the solution of the GF ordinary differential equations (5)-(6). There exists a constant M that only depends on the M_i 's from Assumptions A1-A3, such that for any $T, z \geq 0$ and*

$$\eta \leq \frac{1}{(d + \log m + z^2)M \exp((1 + 1/\varepsilon)MT(1 + T/\varepsilon)^2)},$$

the following holds with probability at least $1 - \exp(-z^2)$:

$$\sup_{k \in [0, T/\eta] \cap \mathbb{N}} \max_{i \in [m]} |\bar{a}_i(k)| \leq M(1 + T/\varepsilon), \quad (112)$$

$$\sup_{k \in [0, T/\eta] \cap \mathbb{N}} \max_{i \in [m]} \|\theta_i(k\eta) - \bar{\theta}_i(k)\|_2 \leq \left(\sqrt{d + \log m} + z\right) \quad (113)$$

$$\times M \exp \left(\left(1 + \frac{1}{\varepsilon}\right) MT \left(1 + \frac{T}{\varepsilon}\right)^2 \right) \sqrt{\eta}, \quad (114)$$

$$\sup_{k \in [0, T/\eta] \cap \mathbb{N}} |\mathcal{R}(\bar{a}(k), \bar{u}(k)) - \mathcal{R}(a(k\eta), u(k\eta))| \leq \left(\sqrt{d + \log m} + z\right) \quad (115)$$

$$\times M \exp \left(\left(1 + \frac{1}{\varepsilon}\right) MT \left(1 + \frac{T}{\varepsilon}\right)^2 \right) \sqrt{\eta}. \quad (116)$$

The proof is presented in Appendix D and follows the same scheme as in that of Theorem 1 part (B) in [33]. The main difference with respect to that theorem is here we are interested in projected SGD (and GF) instead of plain SGD (and GF), hence an additional step of approximation is required, and the a_i 's and u_i 's need to be treated separately. We next draw implications of the last result on learning by online SGD within the canonical learning order.

Theorem 3. Fix any $\delta > 0$. Assume φ, σ and the initialization P_A be such that the canonical learning order of Definition 1 holds up to level L for some $L \geq 2$, and that

$$\sum_{k \geq L+1} \varphi_k^2 \leq \frac{\delta}{2}. \quad (117)$$

Then, there exist constants $\varepsilon_* = \varepsilon_*(\delta)$, $T_0 = T_0(\delta)$, $T = T(\varepsilon, \delta) = T_0(\delta)\varepsilon^{1/(2L)}$ and $M = M(\varepsilon, \delta)$ that depend on ε, δ (together with φ, σ and P_A) such that the following happens. Assume $\varepsilon \leq \varepsilon_*(\delta)$ and m, d, z are such that $d \geq M$, $m \geq \max(M, z)$, and the step size η and number of samples (equivalently, number of steps) n satisfy

$$\eta = \frac{1}{M(d + \log m + z)}, \quad (118)$$

$$n = MT(d + \log m + z). \quad (119)$$

Then, with probability at least $1 - e^{-z}$, the projected gradient descent algorithm of Eq. (107) achieves population risk smaller than δ :

$$\mathbb{P}(\mathcal{R}(\bar{a}(n), \bar{u}(n)) \leq \delta) \geq 1 - e^{-z}. \quad (120)$$

The proof of Theorem 3 is deferred to Appendix D.4.

Remark 7.1. Within the lazy or neural tangent regime, learning the projection of the target function $\varphi(\langle u_*, x \rangle)$ onto polynomials of degree ℓ requires $n \gg d^\ell$ samples, and $m \gg d^{\ell-1}$ neurons [20, 34, 35].

In contrast, Theorem 3 shows that, within the canonical learning order, $O(d)$ samples and $O(1)$ neurons are sufficient. Further as per Theorem 2, the learning dynamics is accurately described by the GF analyzed in the previous sections.

8 Discussion

We conclude by discussing some of our findings as well as potential extensions of our work. As mentioned in the introduction, our initial motivation was to understand certain ubiquitous phenomena in the learning dynamics of multi-layer neural networks. A particularly striking phenomenon that we could reproduce in the present mathematical setting is the coexistence of plateaus in which the risk barely changes and sudden drops.

In the next paragraphs, we will briefly emphasize results or future directions that were not anticipated at the beginning of this work.

Implicit bias in function space. We provided evidence towards the canonical learning order of Definition 1. According to this scenario, the target function φ is learnt according to its decomposition into Hermite polynomials, with lower degree components learnt first. This theory applies to online SGD via Theorem 2 and Theorem 3. In this setting, the number of SGD steps correspond to the number of samples. Therefore, at a small sample size, SGD will fit a low degree polynomial approximation of the target function, with the degree increasing with samples.

A similar phenomenon is observed with (rotationally invariant) kernel methods [34], with one important difference. Here the number of samples always scale linearly in the degree, while for kernel methods, different polynomial degree correspond to different scalings with the dimension.

Implicit bias in parameter space. Our analysis tracks the evolution of the weights as well. As explained in Section 6, in order for the degree- k component of the target function to be well approximated (in the $d, m \rightarrow \infty$ limit), it is sufficient that $\sigma_k \int a(\nu) s(\nu)^k \rho(d\nu) = \varphi_k$. Here ν is an abstract neuron index, $a(\nu)$ is the second-layer weight and $s(\nu)$ is the projection of the first layer weight along the target direction u_* .

Naively, one would expect that, in order for learning to take place, first layer weights should be well aligned with u_* , i.e. $s(\nu)$ should concentrate close to one. However this is not the only way to satisfy the constraints $\sigma_k \int a(\nu) s(\nu)^k d\rho(\nu) = \varphi_k$. Indeed, our analysis in Section 6 indicates that gradient flow satisfies this constraint with $s = \Theta(\varepsilon^{\beta_k})$ and $a = \Theta(\varepsilon^{-\omega_k})$ with $\beta_k = 1/2(k+1)$, $\omega_k = k/2(k+1)$ (so that $\sigma_k \int a(\nu) s(\nu)^k \rho(d\nu)$ will be of order one) as $\varepsilon \rightarrow 0$. In other words, the alignment is small, and second layer weights are large. (In general, weights on multiple scales coexist.)

The role of the learning rate ε . The initialization of parameters and relative step-sizes play a key role in modern (non-convex) machine learning. The combination of the two scalings (initialization and relative stepsize) affects the learning dynamics. In order to clarify this point, we can consider a general parametrization (we keep $\|u_i\|_2 = 1$)

$$f(x; a, u) = \frac{1}{m^\gamma} \sum_{i=1}^m a_i \sigma(\langle u_i, x \rangle) =: \sum_{i=1}^m c_i \sigma(\langle u_i, x \rangle),$$

and gradient flow dynamics

$$\begin{aligned} \varepsilon \partial_t a_i &= -m \partial_{a_i} \mathcal{R}(a, u), \\ \partial_t u_i &= -m(I_d - u_i u_i^\top) \nabla_{u_i} \mathcal{R}(a, u). \end{aligned}$$

(Note that the learning rate in the second equation can be set to 1 without loss of generality, by rescaling the time axis.) Rewriting this in terms of the coefficients c_i , so that the function representation is kept fixed, we have

$$s \partial_t c_i = -m \partial_{c_i} \overline{\mathcal{R}}(c, u), \quad s = \varepsilon m^{2\gamma},$$

while the second equation remains unchanged. This parametrization allows us to compare various scalings in a uniform fashion.

- Mean field scaling [31, 13]: $s = \Theta(m^2)$, $|c_i(0)| = \Theta(m^{-1})$.
- In this paper: $s = \varepsilon m^2$, $|c_i(0)| = \Theta(m^{-1})$, $\varepsilon \rightarrow 0$ after $m \rightarrow \infty$.
- Classical scaling [29, 24]: $s = \Theta(1)$, $|c_i(0)| = \Theta(m^{-1/2})$.

As mentioned already, mean field scaling can exhibit better feature learning properties. In particular, the class of functions studied in the present paper can require much larger sample size to learn under the classical scaling [37, 20, 47]. The choice of initialization in this paper is the same as in the mean field literature, with the difference that the relative learning rate s is a factor ε smaller, hence making it—in a sense—slightly closer to the the classical scaling. It would be interesting to explore other scalings as well.

We also note that, while the limit of small ε is interesting, setting directly $\varepsilon = 0$ leads to a singular behavior². Formally, setting $\varepsilon = 0$ corresponds to keeping second layer weights equal to their optimal values: a correct analysis of this case requires to account for the role of stepsize and not just use the gradient flow approximation.

²No matter how we rescale time, in this case learning takes place instantly, up to a certain critical degree.

More complex network models. The choice of the neural network model in this paper was mainly dictated by the desire to avoid inessential technicalities. It would be important to move towards more realistic models.

First, we used projected gradient descent to constrain the weights' norms $\|u_i\| = 1$. While this is a common theoretical device in studying single-index models [4, 11], we believe that techniques developed here can be extended to the more general case. Analogously, we could add biases to the network architecture and hence replace Eq. (2) by

$$f(x; a, u, b) = \frac{1}{m} \sum_{i=1}^m a_i \sigma(\langle u_i, x \rangle + b_i), \quad a_1, b_1, \dots, a_m, b_m \in \mathbb{R}, \quad u_1, \dots, u_m \in \mathbb{R}^d, \quad (121)$$

With this change, the limiting mean-field dynamics will be an autonomous ODE system of $(a_i(t), b_i(t), s_i(t), r_i(t))_{i=1}^m$ where $r_i(t) = \|u_i(t)\|_2$. We expect that its evolution will be qualitatively similar to that of the simplified dynamics considered in the paper.

Second, the single-index model studied here is a simple example of target function which requires feature learning. An obvious generalization is to consider multi-index models, as already discussed in Remark 4.3.

Finally, it would be interesting to generalize our analysis to classification losses.

Acknowledgments

This work was supported by the NSF through award DMS-2031883, the Simons Foundation through Award 814639 for the Collaboration on the Theoretical Foundations of Deep Learning, the NSF grant CCF-2006489 and the ONR grant N00014-18-1-2729, and a grant from Eric and Wendy Schmidt at the Institute for Advanced Studies. Part of this work was carried out while Andrea Montanari was on partial leave from Stanford and a Chief Scientist at Ndata Inc dba Project N. The present research is unrelated to AM's activity while on leave.

References

- [1] Emmanuel Abbe, Enric Boix Adsera, and Theodor Misiakiewicz. The merged-staircase property: a necessary and nearly sufficient condition for sgd learning of sparse functions on two-layer neural networks. In *Conference on Learning Theory*, pages 4782–4887. PMLR, 2022.
- [2] Luigi Ambrosio, Nicola Gigli, and Giuseppe Savaré. *Gradient flows: in metric spaces and in the space of probability measures*. Springer Science & Business Media, 2005.
- [3] Luca Arnaboldi, Ludovic Stephan, Florent Krzakala, and Bruno Loureiro. From high-dimensional & mean-field dynamics to dimensionless odes: A unifying approach to sgd in two-layers networks. *arXiv preprint arXiv:2302.05882*, 2023.
- [4] Gerard Ben Arous, Reza Gheissari, and Aukosh Jagannath. Online stochastic gradient descent on non-convex losses from high-dimensional inference. *The Journal of Machine Learning Research*, 22(1):4788–4838, 2021.
- [5] Devansh Arpit, Stanisław Jastrzębski, Nicolas Ballas, David Krueger, Emmanuel Bengio, Maxinder S Kanwal, Tegan Maharaj, Asja Fischer, Aaron Courville, Yoshua Bengio, et al. A closer look at memorization in deep networks. In *International conference on machine learning*, pages 233–242. PMLR, 2017.
- [6] Jimmy Ba, Murat A Erdogdu, Taiji Suzuki, Zhichao Wang, Denny Wu, and Greg Yang. High-dimensional asymptotics of feature learning: How one gradient step improves the representation. In *Advances in Neural Information Processing Systems*, 2022.
- [7] Pierre Baldi and Kurt Hornik. Neural networks and principal component analysis: Learning from examples without local minima. *Neural networks*, 2(1):53–58, 1989.
- [8] Boaz Barak, Benjamin L Edelman, Surbhi Goel, Sham Kakade, Eran Malach, and Cyril Zhang. Hidden progress in deep learning: Sgd learns parities near the computational limit. *arXiv:2207.08799*, 2022.
- [9] Peter L Bartlett, Andrea Montanari, and Alexander Rakhlin. Deep learning: a statistical viewpoint. *Acta numerica*, 30:87–201, 2021.
- [10] Nils Berglund. Perturbation theory of dynamical systems. *arXiv preprint math/0111178*, 2001.
- [11] Alberto Bietti, Joan Bruna, Clayton Sanford, and Min Jae Song. Learning single-index models with shallow neural networks. *Advances in Neural Information Processing Systems*, 35:9768–9783, 2022.
- [12] Antoine Bodin and Nicolas Macris. Model, sample, and epoch-wise descents: exact solution of gradient flow in the random feature model. *Advances in Neural Information Processing Systems*, 34:21605–21617, 2021.
- [13] Lenaïc 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, 2018.
- [14] Alexandru Damian, Jason Lee, and Mahdi Soltanolkotabi. Neural networks can learn representations with gradient descent. In *Conference on Learning Theory*, pages 5413–5452. PMLR, 2022.

- [15] Encyclopedia of Mathematics. Bernoulli equation. http://encyclopediaofmath.org/index.php?title=Bernoulli_equation&oldid=40764.
- [16] Christopher Frye and Costas J Efthimiou. Spherical harmonics in p dimensions. *arXiv preprint arXiv:1205.3548*, 2012.
- [17] Kenji Fukumizu and Shun-ichi Amari. Local minima and plateaus in hierarchical structures of multilayer perceptrons. *Neural networks*, 13(3):317–327, 2000.
- [18] Behrooz Ghorbani, Song Mei, Theodor Misiakiewicz, and Andrea Montanari. Discussion of: “nonparametric regression using deep neural networks with relu activation function”. *The Annals of Statistics*, 48(4), 2020.
- [19] Behrooz Ghorbani, Song Mei, Theodor Misiakiewicz, and Andrea Montanari. When do neural networks outperform kernel methods? *Advances in Neural Information Processing Systems*, 33:14820–14830, 2020.
- [20] 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, 2021.
- [21] Nikhil Ghosh, Song Mei, and Bin Yu. The three stages of learning dynamics in high-dimensional kernel methods. In *International Conference on Learning Representations*, 2021.
- [22] Daniel Gissin, Shai Shalev-Shwartz, and Amit Daniely. The implicit bias of depth: How incremental learning drives generalization. *arXiv preprint arXiv:1909.12051*, 2019.
- [23] Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 249–256. JMLR Workshop and Conference Proceedings, 2010.
- [24] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE international conference on computer vision*, pages 1026–1034, 2015.
- [25] Mark Holmes. *Introduction to Perturbation Methods*. Springer Texts in Applied Mathematics, 2013.
- [26] Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: Convergence and generalization in neural networks. *Advances in neural information processing systems*, 31, 2018.
- [27] Chi Jin, Praneeth Netrapalli, Rong Ge, Sham M Kakade, and Michael I Jordan. A short note on concentration inequalities for random vectors with subgaussian norm. *arXiv preprint arXiv:1902.03736*, 2019.
- [28] Dimitris Kalimeris, Gal Kaplun, Preetum Nakkiran, Benjamin Edelman, Tristan Yang, Boaz Barak, and Haofeng Zhang. Sgd on neural networks learns functions of increasing complexity. *Advances in neural information processing systems*, 32, 2019.
- [29] Yann LeCun, Léon Bottou, Genevieve B Orr, and Klaus-Robert Müller. Efficient backprop. In *Neural networks: Tricks of the trade*, pages 9–50. Springer, 2002.
- [30] Zhiyuan Li, Yuping Luo, and Kaifeng Lyu. Towards resolving the implicit bias of gradient descent for matrix factorization: Greedy low-rank learning. *arXiv preprint arXiv:2012.09839*, 2020.

- [31] Song Mei, Yu Bai, and Andrea Montanari. The landscape of empirical risk for nonconvex losses. *The Annals of Statistics*, 46(6A):2747–2774, 2018.
- [32] 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.
- [33] 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.
- [34] 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*, 59:3–84, 2022.
- [35] Andrea Montanari and Yiqiao Zhong. The interpolation phase transition in neural networks: Memorization and generalization under lazy training. *The Annals of Statistics*, 50(5):2816–2847, 2022.
- [36] Ryan O’Donnell. *Analysis of boolean functions*. Cambridge University Press, 2014.
- [37] 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.
- [38] Allan Pinkus. Approximation theory of the mlp model in neural networks. *Acta numerica*, 8:143–195, 1999.
- [39] Alethea Power, Yuri Burda, Harri Edwards, Igor Babuschkin, and Vedant Misra. Grokking: Generalization beyond overfitting on small algorithmic datasets. *arXiv:2201.02177*, 2022.
- [40] Grant Rotskoff and Eric Vanden-Eijnden. Parameters as interacting particles: long time convergence and asymptotic error scaling of neural networks. *Advances in neural information processing systems*, 31, 2018.
- [41] David Saad and Sara A Solla. On-line learning in soft committee machines. *Physical Review E*, 52(4):4225, 1995.
- [42] Filippo Santambrogio. Optimal transport for applied mathematicians. *Birkhäuser, NY*, 55 (58-63):94, 2015.
- [43] Andrew M Saxe, James L McClelland, and Surya Ganguli. Exact solutions to the nonlinear dynamics of learning in deep linear neural networks. *arXiv preprint arXiv:1312.6120*, 2013.
- [44] Haikun Wei, Jun Zhang, Florent Cousseau, Tomoko Ozeki, and Shun-ichi Amari. Dynamics of learning near singularities in layered networks. *Neural computation*, 20(3):813–843, 2008.
- [45] Greg Yang and Edward J Hu. Feature learning in infinite-width neural networks. *arXiv:2011.14522*, 2020.
- [46] Greg Yang and Edward J Hu. Tensor programs iv: Feature learning in infinite-width neural networks. In *International Conference on Machine Learning*, pages 11727–11737. PMLR, 2021.
- [47] Gilad Yehudai and Ohad Shamir. On the power and limitations of random features for understanding neural networks. *Advances in Neural Information Processing Systems*, 32, 2019.

- [48] Yuki Yoshida and Masato Okada. Data-dependence of plateau phenomenon in learning with neural network—statistical mechanical analysis. *Advances in Neural Information Processing Systems*, 32, 2019.

A Proof of Proposition 1

By standard approximation theory arguments [38], it is sufficient to show that there exists an integrable function $a_d \in L^1(\mathbb{S}^{d-1}, \mu_0)$ such that

$$\lim_{d \rightarrow \infty} \mathbb{E}\left\{\left(\int a_d(u) \sigma(\langle u, x \rangle) \mu_0(du) - \varphi(\langle u_*, x \rangle)\right)^2\right\} = 0. \quad (122)$$

(We denote by μ_0 the uniform probability measure over \mathbb{S}^{d-1} .)

Denote by $P_{d,k}$ the Gegenbauer polynomial of order d and degree k (see, e.g., [34]). Namely, $(P_{d,k} : k \geq 0)$ form an orthogonal system with respect to the measure with density $\propto (1-t^2)^{(d-3)/2}$, $t \in [-1, 1]$. Recall that for fixed v, w of norm 1, the polynomials $P_{d,j}(\langle v, u \rangle), P_{d,k}(\langle w, u \rangle)$ are spherical harmonics satisfying

$$\int P_{d,j}(\langle v, u \rangle) P_{d,k}(\langle w, u \rangle) \mu_0(du) = \delta_{kj} P_{d,k}(\langle v, w \rangle). \quad (123)$$

Also, $P_{d,k}(1) = B_{d,k}$ is the dimension of the space of spherical harmonics of degree k , whence $(P_{d,k}(\cdot)/B_{d,k}^{1/2} : k \geq 0)$ form an orthonormal set. We will denote by $c_{d,k}(\sigma)$ the k -th coefficient of the expansion of $\sigma(\cdot \sqrt{d})$ in this basis, and similarly for $\varphi(\cdot \sqrt{d})$, with coefficients $c_{d,k}(\varphi)$, namely

$$\begin{aligned} \sigma(t\sqrt{d}) &= \sum_{k=0}^{\infty} \frac{c_{d,k}(\sigma)}{B_{d,k}^{1/2}} P_{d,k}(t), \\ \varphi(t\sqrt{d}) &= \sum_{k=0}^{\infty} \frac{c_{d,k}(\varphi)}{B_{d,k}^{1/2}} P_{d,k}(t). \end{aligned}$$

As shown for instance in [34], $\lim_{d \rightarrow \infty} c_{d,k}(\sigma) = c_k(\sigma)$ is the k -th Hermite coefficient of σ and similarly for $c_{d,k}(\varphi)$. In particular, $c_{d,k}(\sigma) \neq 0$ for all d large enough. For N a large integer let

$$a_d(u) = \sum_{k=0}^N \frac{c_{d,k}(\varphi)}{c_{d,k}(\sigma)} P_{d,k}(\langle u, u_* \rangle).$$

By Eq. (123), we have, for $\|z\| = \sqrt{d}$,

$$\begin{aligned} \int a_d(u) \sigma(\langle u, z \rangle) \mu_0(du) &= \sum_{k=0}^N \frac{c_{d,k}(\varphi)}{c_{d,k}(\sigma)} \frac{c_{d,k}(\sigma)}{B_{d,k}^{1/2}} P_{d,k}(\langle u_*, z \rangle / \sqrt{d}) \\ &= \sum_{k=0}^N \frac{c_{d,k}(\varphi)}{B_{d,k}^{1/2}} P_{d,k}(\langle u_*, z \rangle / \sqrt{d}). \end{aligned}$$

Denoting by z a uniform random vector on the sphere of radius \sqrt{d} , and $r = \|x\|_2 / \sqrt{d}$, we have

$$\begin{aligned} \mathbb{E}\left\{\left(\int a_d(u) \sigma(\langle u, x \rangle) \mu_0(du) - \varphi(\langle u_*, x \rangle)\right)^2\right\} &= \mathbb{E}\left\{\left(\int a_d(u) \sigma(r \langle u, z \rangle) \mu_0(du) - \varphi(r \langle u_*, z \rangle)\right)^2\right\} \\ &\stackrel{(*)}{\leq} \mathbb{E}\left\{\left(\int a_d(u) \sigma(\langle u, z \rangle) \mu_0(du) - \varphi(\langle u_*, z \rangle)\right)^2\right\} + \frac{C_N L^2}{d} \\ &\leq \mathbb{E}\{\varphi_{>N}(\langle u_*, z \rangle)^2\} + \frac{C_N L^2}{d}, \end{aligned}$$

where in $(*)$ we used concentration of χ -squared random variables, Lipschitzness of σ and φ , and that $\varphi_{>N}(t\sqrt{d})$ is the projection of $\varphi(t\sqrt{d})$ orthogonal to polynomials of degree at most N (with respect to the measure with density proportional to $(1-t^2)^{(d-3)/2}$ on $[-1, 1]$). Therefore

$$\limsup_{d \rightarrow \infty} \mathbb{E}\left\{\left(\int a_d(u) \sigma(\langle u, x \rangle) \mu_0(du) - \varphi(\langle u_*, x \rangle)\right)^2\right\} \leq \sum_{k=N+1}^{\infty} c_k(\varphi)^2.$$

The claim (122) follows by taking $N \rightarrow \infty$.

B Appendix to Section 4

B.1 Proof of Proposition 2

When $x \sim \mathcal{N}(0, I_d)$ and $u, u' \in \mathbb{S}^{d-1}$, $\begin{pmatrix} \langle u, x \rangle \\ \langle u', x \rangle \end{pmatrix} \sim \mathcal{N}\left(0, \begin{pmatrix} 1 & \langle u, u' \rangle \\ \langle u, u' \rangle & 1 \end{pmatrix}\right)$. Thus

$$\begin{aligned} \mathcal{R}(a, u) &= \frac{1}{2} \mathbb{E} \left(\varphi(\langle u_*, x \rangle) - \frac{1}{m} \sum_{i=1}^m a_i \sigma(\langle u_i, x \rangle) \right)^2 \\ &= \frac{1}{2} \mathbb{E} \left[\varphi(\langle u_*, x \rangle)^2 \right] - \frac{1}{m} \sum_{i=1}^m a_i \mathbb{E} [\varphi(\langle u_*, x \rangle) \sigma(\langle u_i, x \rangle)] + \frac{1}{2} \frac{1}{m^2} \sum_{i,j=1}^m a_i a_j \mathbb{E} [\sigma(\langle u_i, x \rangle) \sigma(\langle u_j, x \rangle)] \\ &= \frac{1}{2} \|\varphi\|_{L^2}^2 - \frac{1}{m} \sum_{i=1}^m a_i V(\langle u_*, u_i \rangle) + \frac{1}{2} \frac{1}{m^2} \sum_{i,j=1}^m a_i a_j U(\langle u_i, u_j \rangle) \\ &= \frac{1}{2} \|\varphi\|_{L^2}^2 - \frac{1}{m} \sum_{i=1}^m a_i V(s_i) + \frac{1}{2} \frac{1}{m^2} \sum_{i,j=1}^m a_i a_j U(r_{ij}). \end{aligned} \quad (124)$$

This proves (15). Equation (16) follows directly:

$$\varepsilon \partial_t a_i = -m \partial_{a_i} \mathcal{R}(a, u) = V(s_i) - \frac{1}{m} \sum_{j=1}^m a_j U(r_{ij}).$$

To obtain equations (17)-(19), we now take gradients in (124):

$$\begin{aligned} \partial_t u_i &= -m(I_d - u_i u_i^\top) \nabla_{u_i} \mathcal{R}(a, u) \\ &= a_i (I_d - u_i u_i^\top) \left(V'(\langle u_*, u_i \rangle) u_* - \frac{1}{m} \sum_{j=1}^m a_j U'(\langle u_i, u_j \rangle) u_j \right) \\ &= a_i \left(V'(\langle u_*, u_i \rangle) (u_* - u_i u_i^\top u_*) - \frac{1}{m} \sum_{j=1}^m a_j U'(\langle u_i, u_j \rangle) (u_j - u_i u_i^\top u_j) \right) \\ &= a_i \left(V'(s_i) (u_* - s_i u_i) - \frac{1}{m} \sum_{j=1}^m a_j U'(r_{ij}) (u_j - r_{ij} u_i) \right). \end{aligned}$$

Thus

$$\begin{aligned} \partial_t s_i &= \langle u_*, \partial_t u_i \rangle \\ &= a_i \left(V'(s_i) (\langle u_*, u_* \rangle - s_i \langle u_*, u_i \rangle) - \frac{1}{m} \sum_{j=1}^m a_j U'(r_{ij}) (\langle u_*, u_j \rangle - r_{ij} \langle u_*, u_i \rangle) \right) \\ &= a_i \left(V'(s_i) (1 - s_i^2) - \frac{1}{m} \sum_{j=1}^m a_j U'(r_{ij}) (s_j - r_{ij} s_i) \right). \end{aligned}$$

This gives (17). Finally, we perform a similar computation to compute $\partial_t r_{ij} = \langle \partial_t u_i, u_j \rangle + \langle u_i, \partial_t u_j \rangle$. We compute only the first term, as the second term can be obtained by inverting i and j :

$$\begin{aligned} \langle \partial_t u_i, u_j \rangle &= a_i \left(V'(s_i) (\langle u_j, u_* \rangle - s_i \langle u_j, u_i \rangle) - \frac{1}{m} \sum_{p=1}^m a_p U'(r_{ip}) (\langle u_j, u_p \rangle - r_{ip} \langle u_j, u_i \rangle) \right) \\ &= a_i \left(V'(s_i) (s_j - s_i r_{ij}) - \frac{1}{m} \sum_{p=1}^m a_p U'(r_{ip}) (r_{jp} - r_{ip} r_{ij}) \right). \end{aligned}$$

Adding the symmetric term $\langle u_i, \partial_t u_j \rangle$, we obtain (18)-(19).

B.2 Proof of Corollary 1

First, note that in the proof of Lemma 1, we obtain the following a priori estimate on the magnitude of the a_i^0 's:

$$\sup_{1 \leq i \leq m} |a_i^0(t)| \leq M \left(1 + \frac{t}{\varepsilon}\right), \quad \forall t \geq 0, \quad (125)$$

where M only depends on the M_i 's in Assumptions A1-A3. Using a similar argument as that in the proof of Proposition 3, we obtain that for any $t \in [0, T]$ and $i \in [m]$,

$$\begin{aligned} |\partial_t(a_i - a_i^0)| &\leq \frac{M}{\varepsilon} \left(|s_i - s_i^0| + \frac{1}{m} \sum_{j=1}^m |a_j - a_j^0| \right) + \frac{M(1+t/\varepsilon)}{\varepsilon} \cdot \frac{1}{m} \sum_{j=1}^m |r_{ij} - r_{ij}^0|, \\ |\partial_t(s_i - s_i^0)| &\leq M(1+t/\varepsilon) \cdot \left(|a_i - a_i^0| + \frac{1}{m} \sum_{j=1}^m |a_j - a_j^0| \right) \\ &\quad + M(1+t/\varepsilon)^2 \cdot \left(|s_i - s_i^0| + \frac{1}{m} \sum_{j=1}^m (|s_j - s_j^0| + |r_{ij} - r_{ij}^0|) \right), \end{aligned}$$

and for $1 \leq i \neq j \leq m$,

$$\begin{aligned} |\partial_t(r_{ij} - r_{ij}^0)| &\leq M(1+t/\varepsilon) \left(|a_i - a_i^0| + |a_j - a_j^0| + |s_i - s_i^0| + |s_j - s_j^0| + \frac{1}{m} \sum_{p=1}^m |a_p - a_p^0| \right) \\ &\quad + M(1+t/\varepsilon)^2 \cdot \left(|r_{ij} - r_{ij}^0| + \frac{1}{m} \sum_{p=1}^m (|r_{ip} - r_{ip}^0| + |r_{jp} - r_{jp}^0|) \right). \end{aligned}$$

Therefore, we deduce that

$$\begin{aligned} \partial_t \sum_{i=1}^m (a_i - a_i^0)^2 &\leq \frac{M}{\varepsilon} \sum_{i=1}^m (s_i - s_i^0)^2 + \frac{M(1+t/\varepsilon)}{\varepsilon} \left(\sum_{i=1}^m (a_i - a_i^0)^2 + \frac{1}{m} \sum_{i,j=1}^m (r_{ij} - r_{ij}^0)^2 \right), \\ \partial_t \sum_{i=1}^m (s_i - s_i^0)^2 &\leq M(1+t/\varepsilon) \sum_{i=1}^m (a_i - a_i^0)^2 + M(1+t/\varepsilon)^2 \left(\sum_{i=1}^m (s_i - s_i^0)^2 + \frac{1}{m} \sum_{i,j=1}^m (r_{ij} - r_{ij}^0)^2 \right), \\ \partial_t \left(\frac{1}{m} \sum_{i,j=1}^m (r_{ij} - r_{ij}^0)^2 \right) &\leq M(1+t/\varepsilon) \left(\sum_{i=1}^m (a_i - a_i^0)^2 + \sum_{i=1}^m (s_i - s_i^0)^2 \right) + M(1+t/\varepsilon)^2 \cdot \frac{1}{m} \sum_{i,j=1}^m (r_{ij} - r_{ij}^0)^2. \end{aligned}$$

Defining

$$G(t) = \sum_{i=1}^m (a_i(t) - a_i^0(t))^2 + \sum_{i=1}^m (s_i(t) - s_i^0(t))^2 + \frac{1}{m} \sum_{i,j=1}^m (r_{ij}(t) - r_{ij}^0(t))^2,$$

then we know that $G'(t) \leq (M(1+t)^2/\varepsilon^2)G(t)$. Applying Grönwall's inequality yields

$$G(t) \leq G(0) \exp \left(\int_0^t (M(1+s)^2/\varepsilon^2) ds \right) \leq G(0) \exp \left(Mt(1+t)^2/\varepsilon^2 \right), \quad \forall t \in [0, T].$$

Since $\{\langle u_i(0), u_* \rangle\}_{i \in [m]} \sim \text{i.i.d. } \mathcal{N}(0, 1/d)$ and for any $i \in [m]$, $\{\langle u_i(0), u_j(0) \rangle\}_{j \neq i} \sim \text{i.i.d. } \mathcal{N}(0, 1/d)$. Using standard concentration inequalities, we know that

$$G(0) = \sum_{i=1}^m \langle u_i(0), u_* \rangle^2 + \frac{1}{m} \sum_{i \neq j} \langle u_i(0), u_j(0) \rangle^2 \leq C \frac{m}{d} \quad (126)$$

with probability at least $1 - \exp(-C'm)$, where C and C' are both absolute constants. Therefore,

$$\sup_{t \in [0, T]} \|a(t) - a^0(t)\|_2 \leq C \sqrt{\frac{m}{d}} \exp\left(MT(1+T)^2/\varepsilon^2\right), \quad (127)$$

$$\sup_{t \in [0, T]} \|s(t) - s^0(t)\|_2 \leq C \sqrt{\frac{m}{d}} \exp\left(MT(1+T)^2/\varepsilon^2\right), \quad (128)$$

$$\sup_{t \in [0, T]} \|R(t) - R^0(t)\|_F \leq C \frac{m}{\sqrt{d}} \exp\left(MT(1+T)^2/\varepsilon^2\right). \quad (129)$$

Next we upper bound the risk difference, by direct calculation,

$$\begin{aligned} & |\mathcal{R}(a(t), u(t)) - \mathcal{R}_{\text{red}}(a^0(t), s^0(t), R^0(t))| \\ &= |\mathcal{R}_{\text{red}}(a(t), s(t), R(t)) - \mathcal{R}_{\text{red}}(a^0(t), s^0(t), R^0(t))| \\ &\leq \frac{1}{m} |a^\top V(s) - (a^0)^\top V(s^0)| + \frac{1}{2m^2} |a^\top U(R)a - (a^0)^\top U(R^0)a^0| \\ &\leq \frac{1}{m} \left(\sqrt{m} \|V\|_\infty \|a - a^0\|_2 + \|a^0\|_2 \|V'\|_\infty \|s - s^0\|_2 \right) \\ &\quad + \frac{1}{2m^2} \left(\|U(R) - U(R^0)\|_{\text{op}} \|a^0\|_2^2 + 2 \|U(R)\|_{\text{op}} \|a^0\|_2 \|a - a^0\|_2 \right) \\ &\leq \frac{M(1+t/\varepsilon)}{\sqrt{m}} (\|a - a^0\|_2 + \|s - s^0\|_2) + \frac{M(1+t/\varepsilon)^2}{2m} \|R - R^0\|_F \\ &\leq \frac{M}{\sqrt{d}} \exp\left(Mt(1+t)^2/\varepsilon^2\right) \end{aligned}$$

with probability at least $1 - \exp(-C'm)$, where the constant M only depends on the M_i 's from Assumptions A1-A3. The conclusion now follows from taking the supremum over all $t \in [0, T]$. This completes the proof of Corollary 1.

B.3 Proof of Proposition 3

We consider $r_{ij}^\perp = r_{ij} - s_i s_j = \langle u_i, u_j \rangle - \langle u_i, u_* \rangle \langle u_j, u_* \rangle$, the dot product between u_i and u_j that is out of the relevant subspace spanned by u_* . We show that these variables satisfy the ODEs

$$\begin{aligned} \partial_t r_{ij}^\perp &= -a_i \left(V'(s_i) \cdot s_i r_{ij}^\perp + \frac{1}{m} \sum_{p=1}^m a_p U'(r_{ip})(r_{jp}^\perp - r_{ip} r_{ij}^\perp) \right) \\ &\quad - a_j \left(V'(s_j) \cdot s_j r_{ij}^\perp + \frac{1}{m} \sum_{p=1}^m a_p U'(r_{jp})(r_{ip}^\perp - r_{jp} r_{ij}^\perp) \right). \end{aligned} \quad (130)$$

By definition of r_{ij}^\perp , we readily see that

$$\partial_t r_{ij}^\perp = \partial_t r_{ij} - s_i \partial_t s_j - s_j \partial_t s_i.$$

Plugging in Eq.s (17) to (19) gives that

$$\begin{aligned} \partial_t r_{ij}^\perp &= a_i \left(V'(s_i)(s_j s_i^2 - s_i r_{ij}) - \frac{1}{m} \sum_{p=1}^m a_p U'(r_{ip})(r_{jp} - s_j s_p - r_{ip} r_{ij} + s_i s_j r_{ip}) \right) \\ &\quad + a_j \left(V'(s_j)(s_i s_j^2 - s_j r_{ij}) - \frac{1}{m} \sum_{p=1}^m a_p U'(r_{jp})(r_{ip} - s_i s_p - r_{jp} r_{ij} + s_i s_j r_{jp}) \right) \\ &= -a_i \left(V'(s_i) s_i r_{ij}^\perp + \frac{1}{m} \sum_{p=1}^m a_p U'(r_{ip})(r_{jp}^\perp - r_{ip} r_{ij}^\perp) \right) \end{aligned}$$

$$-a_j \left(V'(s_j) s_j r_{ij}^\perp + \frac{1}{m} \sum_{p=1}^m a_p U'(r_{jp}) (r_{ip}^\perp - r_{jp} r_{ij}^\perp) \right).$$

This proves Eq. (130).

Lemma 1. *If Assumptions A1-A3 hold, then we have for any fixed $T > 0$:*

$$\sup_{t \in [0, T]} \sum_{i,j=1}^m r_{ij}^\perp(t)^2 \leq m \exp \left(MT(1+T)^2 / \varepsilon^2 \right).$$

Proof. To begin with, using Eq. (130), we obtain that

$$\begin{aligned} \partial_t \left(\sum_{i,j=1}^m r_{ij}^\perp(t)^2 \right) &= 2 \sum_{i,j=1}^m r_{ij}^\perp \times \partial_t r_{ij}^\perp \\ &= -4 \sum_{i,j=1}^m a_i r_{ij}^\perp \left(V'(s_i) \cdot s_i r_{ij}^\perp + \frac{1}{m} \sum_{p=1}^m a_p U'(r_{ip}) (r_{jp}^\perp - r_{ip} r_{ij}^\perp) \right). \end{aligned}$$

Using the ODEs for the a_i 's, we obtain that

$$\begin{aligned} |\partial_t a_i| &= \frac{1}{\varepsilon} \left| V(s_i) - \frac{1}{m} \sum_{j=1}^m a_j U(r_{ij}) \right| = \frac{1}{\varepsilon} \left| \mathbb{E} [\varphi(\langle u_*, x \rangle) \sigma(\langle u_i, x \rangle)] - \frac{1}{m} \sum_{j=1}^m a_j \mathbb{E} [\sigma(\langle u_i, x \rangle) \sigma(\langle u_j, x \rangle)] \right| \\ &= \frac{1}{\varepsilon} |\mathbb{E} [\sigma(\langle u_i, x \rangle) (y - f(x; a, u))]| \leq \frac{1}{\varepsilon} \mathbb{E} [\sigma(\langle u_i, x \rangle)^2]^{1/2} \mathbb{E} [(y - f(x; a, u))^2]^{1/2} \\ &\stackrel{(i)}{\leq} \frac{1}{\varepsilon} M_2 \sqrt{2\mathcal{R}(a(0), u(0))} \leq \frac{M}{\varepsilon}, \end{aligned}$$

where (i) follows from our assumptions and the fact that $\mathcal{R}(a(t), u(t)) \leq \mathcal{R}(a(0), u(0))$, since $\partial_t \mathcal{R}(a, u) \leq 0$ by gradient flow equations. Moreover, the constant M only depends on the M_i 's. Since $|a_i(0)| \leq M_1$ for all $i \in [m]$, we know that $|a_i(t)| \leq M(1 + t/\varepsilon)$ for all $t \geq 0$, thus leading to the following estimate:

$$\begin{aligned} \partial_t \left(\sum_{i,j=1}^m r_{ij}^\perp(t)^2 \right) &\leq 4 \sum_{i,j=1}^m M \left(1 + \frac{t}{\varepsilon} \right) |r_{ij}^\perp| \left(\|V'\|_\infty |r_{ij}^\perp| + M \left(1 + \frac{t}{\varepsilon} \right) \|U'\|_\infty \left(\frac{1}{m} \sum_{p=1}^m |r_{jp}^\perp| + |r_{ij}^\perp| \right) \right) \\ &\leq M \left(1 + \frac{t}{\varepsilon} \right)^2 \left(\sum_{i,j=1}^m r_{ij}^\perp(t)^2 + \frac{1}{m} \sum_{i,j,p=1}^m |r_{ij}^\perp(t)| \cdot |r_{jp}^\perp(t)| \right) \\ &\leq M \left(1 + \frac{t}{\varepsilon} \right)^2 \left(\sum_{i,j=1}^m r_{ij}^\perp(t)^2 + \frac{1}{2m} \sum_{i,j,p=1}^m (r_{ij}^\perp(t)^2 + r_{jp}^\perp(t)^2) \right) \\ &\leq M \left(1 + \frac{t}{\varepsilon} \right)^2 \sum_{i,j=1}^m r_{ij}^\perp(t)^2, \end{aligned}$$

where the constant M only depends on the M_i 's in our assumptions. At initialization, we know that $\sum_{i,j=1}^m r_{ij}^\perp(0)^2 = m$. Applying Grönwall's inequality yields that

$$\sum_{i,j=1}^m r_{ij}^\perp(t)^2 \leq m \exp \left(\int_0^t M \left(1 + \frac{s}{\varepsilon} \right)^2 ds \right), \quad \forall t \in [0, T],$$

which further implies that

$$\sup_{t \in [0, T]} \sum_{i,j=1}^m r_{ij}^\perp(t)^2 \leq m \exp \left(\int_0^T M \left(1 + \frac{t}{\varepsilon} \right)^2 dt \right) \leq m \exp \left(MT(1+T)^2 / \varepsilon^2 \right).$$

This completes the proof. \square

We show that

$$\sup_{t \in [0, T]} \sum_{i=1}^m \left((a_i(t) - a_i^{\text{mf}}(t))^2 + (s_i(t) - s_i^{\text{mf}}(t))^2 \right) \leq C(T). \quad (131)$$

To this end, we define $S(t) = \sum_{i=1}^m \left((a_i(t) - a_i^{\text{mf}}(t))^2 + (s_i(t) - s_i^{\text{mf}}(t))^2 \right)$. By our assumption, $S(0) = 0$. Moreover, using the same technique as in the proof of Lemma 1, we know that $|a_i^{\text{mf}}(t)| \leq M(1 + t/\varepsilon)$ for all $i \in [m]$. According to Eq.s (16)-(19) and Eq. (24), we deduce that

$$\begin{aligned} |\partial_t(a_i - a_i^{\text{mf}})| &= \frac{1}{\varepsilon} \cdot \left| V(s_i) - V(s_i^{\text{mf}}) - \frac{1}{m} \sum_{j=1}^m \left(a_j U(r_{ij}) - a_j^{\text{mf}} U(s_i^{\text{mf}} s_j^{\text{mf}}) \right) \right| \\ &\leq \frac{1}{\varepsilon} \cdot \left(|V(s_i) - V(s_i^{\text{mf}})| + \frac{1}{m} \sum_{j=1}^m |a_j U(r_{ij}) - a_j U(s_i s_j)| + \frac{1}{m} \sum_{j=1}^m |a_j U(s_i s_j) - a_j^{\text{mf}} U(s_i^{\text{mf}} s_j^{\text{mf}})| \right) \\ &\leq \frac{1}{\varepsilon} \cdot \left(\|V'\|_{\infty} |s_i - s_i^{\text{mf}}| + \frac{M(1 + t/\varepsilon)}{m} \|U'\|_{\infty} \sum_{j=1}^m |r_{ij}^{\perp}| \right) \\ &\quad + \frac{1}{m\varepsilon} \sum_{j=1}^m \left(\|U\|_{\infty} |a_j - a_j^{\text{mf}}| + M \left(1 + \frac{t}{\varepsilon} \right) \|U'\|_{\infty} (|s_i - s_i^{\text{mf}}| + |s_j - s_j^{\text{mf}}|) \right) \\ &\leq \frac{M}{\varepsilon} \left(1 + \frac{t}{\varepsilon} \right) \cdot \left(|s_i - s_i^{\text{mf}}| + \frac{1}{m} \sum_{j=1}^m (|s_j - s_j^{\text{mf}}| + |a_j - a_j^{\text{mf}}| + |r_{ij}^{\perp}|) \right), \end{aligned}$$

thus leading to the following estimate:

$$\begin{aligned} \sum_{i=1}^m (a_i - a_i^{\text{mf}}) \cdot \partial_t(a_i - a_i^{\text{mf}}) &\leq \frac{M}{\varepsilon} \left(1 + \frac{t}{\varepsilon} \right) \cdot \frac{1}{m} \sum_{i=1}^m |a_i - a_i^{\text{mf}}| \cdot \sum_{j=1}^m (|a_j - a_j^{\text{mf}}| + |s_j - s_j^{\text{mf}}|) \\ &\quad + \frac{M}{\varepsilon} \left(1 + \frac{t}{\varepsilon} \right) \cdot \left(\sum_{i=1}^m |a_i - a_i^{\text{mf}}| |s_i - s_i^{\text{mf}}| + \frac{1}{m} \sum_{i,j=1}^m |a_i - a_i^{\text{mf}}| |r_{ij}^{\perp}| \right) \\ &\stackrel{(i)}{\leq} \frac{M}{\varepsilon} \left(1 + \frac{t}{\varepsilon} \right) \cdot \left(\sum_{i=1}^m (a_i - a_i^{\text{mf}})^2 + \sum_{i=1}^m (s_i - s_i^{\text{mf}})^2 + \frac{1}{m} \sum_{i,j=1}^m (r_{ij}^{\perp})^2 \right) \\ &\stackrel{(ii)}{\leq} \frac{M}{\varepsilon} \left(1 + \frac{t}{\varepsilon} \right) \cdot \left(\sum_{i=1}^m (a_i - a_i^{\text{mf}})^2 + \sum_{i=1}^m (s_i - s_i^{\text{mf}})^2 + \exp(Mt(1 + t)^2/\varepsilon^2) \right), \end{aligned}$$

where in (i) we use the Cauchy-Schwarz inequality and the inequality of arithmetic and geometric means, and (ii) follows from the conclusion of Lemma 1. Similarly, we obtain that

$$\begin{aligned} |\partial_t(s_i - s_i^{\text{mf}})| &\leq \|V'\|_{\infty} |a_i - a_i^{\text{mf}}| + M \left(1 + \frac{t}{\varepsilon} \right) (\|V''\|_{\infty} + 2\|V'\|_{\infty}) |s_i - s_i^{\text{mf}}| \\ &\quad + \frac{1}{m} \left| a_i \sum_{j=1}^m a_j U'(r_{ij})(s_j - r_{ij}s_i) - a_i \sum_{j=1}^m a_j U'(s_i s_j)(1 - s_i^2)s_j \right| \\ &\quad + \frac{1}{m} \left| a_i \sum_{j=1}^m a_j U'(s_i s_j)(1 - s_i^2)s_j - a_i^{\text{mf}} \sum_{j=1}^m a_j^{\text{mf}} U'(s_i^{\text{mf}} s_j^{\text{mf}})(1 - (s_i^{\text{mf}})^2)s_j^{\text{mf}} \right| \\ &\leq M \left(1 + \frac{t}{\varepsilon} \right) \cdot (|a_i - a_i^{\text{mf}}| + |s_i - s_i^{\text{mf}}|) + \frac{2}{m} \cdot M \left(1 + \frac{t}{\varepsilon} \right) (\|U'\|_{\infty} + \|U''\|_{\infty}) \sum_{j=1}^m |a_j| |r_{ij}^{\perp}| \\ &\quad + M \left(1 + \frac{t}{\varepsilon} \right)^2 \cdot \frac{1}{m} \sum_{j=1}^m (|a_j - a_j^{\text{mf}}| + |s_j - s_j^{\text{mf}}|) \end{aligned}$$

$$\leq M \left(1 + \frac{t}{\varepsilon}\right)^2 \cdot \left(|a_i - a_i^{\text{mf}}| + |s_i - s_i^{\text{mf}}| + \frac{1}{m} \sum_{j=1}^m (|s_j - s_j^{\text{mf}}| + |a_j - a_j^{\text{mf}}| + |r_{ij}^\perp|) \right),$$

which further implies that

$$\sum_{i=1}^m (s_i - s_i^{\text{mf}}) \cdot \partial_t (s_i - s_i^{\text{mf}}) \leq M \left(1 + \frac{t}{\varepsilon}\right)^2 \cdot \left(\sum_{i=1}^m (a_i - a_i^{\text{mf}})^2 + \sum_{i=1}^m (s_i - s_i^{\text{mf}})^2 + \exp(Mt(1+t)^2/\varepsilon^2) \right).$$

Combining the above estimates, we finally deduce that

$$\begin{aligned} S'(t) &= 2 \sum_{i=1}^m ((a_i - a_i^{\text{mf}}) \cdot \partial_t (a_i - a_i^{\text{mf}}) + (s_i - s_i^{\text{mf}}) \cdot \partial_t (s_i - s_i^{\text{mf}})) \\ &\leq \frac{M(1+t)(1+t/\varepsilon)}{\varepsilon} \cdot \left(\sum_{i=1}^m (a_i - a_i^{\text{mf}})^2 + \sum_{i=1}^m (s_i - s_i^{\text{mf}})^2 + \exp(Mt(1+t)^2/\varepsilon^2) \right) \\ &= \frac{M(1+t)(1+t/\varepsilon)}{\varepsilon} \cdot (S(t) + \exp(Mt(1+t)^2/\varepsilon^2)). \end{aligned}$$

Applying Grönwall's inequality immediately implies

$$S(t) \leq \exp \left(Mt(1+t)^2/\varepsilon^2 + \int_0^t \frac{M(1+s)(1+s/\varepsilon)}{\varepsilon} ds \right) \leq \exp \left(Mt(1+t)^2/\varepsilon^2 \right), \quad (132)$$

which further leads to Eq. (131) and concludes the proof of Proposition 3. The “consequently” part can be shown via direct calculation, but we include it here for the sake of completeness. By definition, for any $t \in [0, T]$ we have

$$\begin{aligned} &|\mathcal{R}_{\text{red}}(a(t), s(t), R(t)) - \mathcal{R}_{\text{mf}}(a^{\text{mf}}(t), s^{\text{mf}}(t))| \\ &\leq \frac{1}{m} \left| \sum_{i=1}^m a_i V(s_i) - \sum_{i=1}^m a_i^{\text{mf}} V(s_i^{\text{mf}}) \right| + \frac{1}{2m^2} \left| \sum_{i,j=1}^m a_i a_j U(r_{ij}) - \sum_{i,j=1}^m a_i^{\text{mf}} a_j^{\text{mf}} U(s_i^{\text{mf}} s_j^{\text{mf}}) \right| \\ &\leq \frac{1}{m} \sum_{i=1}^m (\|V\|_\infty |a_i - a_i^{\text{mf}}| + M(1+t/\varepsilon) \|V'\|_\infty |s_i - s_i^{\text{mf}}|) + \frac{1}{2m^2} M(1+t/\varepsilon)^2 \|U'\|_\infty \sum_{i,j=1}^m |r_{ij}^\perp| \\ &\quad + \frac{1}{2m^2} \sum_{i,j=1}^m (M(1+t/\varepsilon) \|U\|_\infty (|a_i - a_i^{\text{mf}}| + |a_j - a_j^{\text{mf}}|) + M(1+t/\varepsilon)^2 \|U'\|_\infty (|s_i - s_i^{\text{mf}}| + |s_j - s_j^{\text{mf}}|)) \\ &\leq M(1+t/\varepsilon) \frac{1}{m} \sum_{i=1}^m |a_i - a_i^{\text{mf}}| + M(1+t/\varepsilon)^2 \frac{1}{m} \sum_{i=1}^m |s_i - s_i^{\text{mf}}| + M(1+t/\varepsilon)^2 \frac{1}{m^2} \sum_{i,j=1}^m |r_{ij}^\perp| \\ &\leq M(1+t/\varepsilon)^2 \cdot \left(\sqrt{\frac{1}{m} \sum_{i=1}^m ((a_i(t) - a_i^{\text{mf}}(t))^2 + (s_i(t) - s_i^{\text{mf}}(t))^2)} + \sqrt{\frac{1}{m^2} \sum_{i,j=1}^m |r_{ij}^\perp(t)|^2} \right) \\ &\leq \frac{1}{\sqrt{m}} M(1+t/\varepsilon)^2 \exp(Mt(1+t)^2/\varepsilon^2). \end{aligned}$$

Therefore,

$$\sup_{t \in [0, T]} |\mathcal{R}_{\text{red}}(a(t), s(t), R(t)) - \mathcal{R}_{\text{mf}}(a^{\text{mf}}(t), s^{\text{mf}}(t))| \leq \frac{M \exp(MT(1+T)^2/\varepsilon^2)}{\sqrt{m}}, \quad (133)$$

as desired.

B.4 Derivation of the mean field dynamics (29)

For any bounded continuous $f \in C_b(\mathbb{R}^2)$, we have

$$\begin{aligned}
\int_{\mathbb{R}^2} f(a, s) \partial_t \rho_t(da, ds) &= \partial_t \left(\int_{\mathbb{R}^2} f(a, s) \rho_t(da, ds) \right) = \partial_t \left(\frac{1}{m} \sum_{i=1}^m f(a_i^{\text{mf}}(t), s_i^{\text{mf}}(t)) \right) \\
&= \frac{1}{m} \sum_{i=1}^m (\partial_a f(a_i^{\text{mf}}(t), s_i^{\text{mf}}(t)) \cdot \partial_t a_i^{\text{mf}}(t) + \partial_s f(a_i^{\text{mf}}(t), s_i^{\text{mf}}(t)) \cdot \partial_t s_i^{\text{mf}}(t)) \\
&\stackrel{(i)}{=} \frac{1}{m} \sum_{i=1}^m (\partial_a f(a_i^{\text{mf}}(t), s_i^{\text{mf}}(t)) \cdot \Psi_a(a_i^{\text{mf}}(t), s_i^{\text{mf}}(t); \rho_t) + \partial_s f(a_i^{\text{mf}}(t), s_i^{\text{mf}}(t)) \cdot \Psi_s(a_i^{\text{mf}}(t), s_i^{\text{mf}}(t); \rho_t)) \\
&= \int_{\mathbb{R}^2} (\partial_a f(a, s) \cdot \Psi_a(a, s; \rho_t) + \partial_s f(a, s) \cdot \Psi_s(a, s; \rho_t)) \rho_t(da, ds) \\
&\stackrel{(ii)}{=} - \int_{\mathbb{R}^2} f(a, s) \cdot (\partial_a (\rho_t \Psi_a(a, s; \rho_t)) + \partial_s (\rho_t \Psi_s(a, s; \rho_t))) (da, ds),
\end{aligned}$$

where (i) follows from the ODE satisfied by the $(a_i^{\text{mf}}(t), s_i^{\text{mf}}(t))$'s, and in (ii) we use integration by parts. We thus obtain that

$$\partial_t \rho_t = -(\partial_a (\rho_t \Psi_a(a, s; \rho_t)) + \partial_s (\rho_t \Psi_s(a, s; \rho_t))) = -\nabla \cdot (\rho_t \Psi(a, s; \rho_t)),$$

which recovers Eq. (29).

B.5 Details of the alternative mean field approach

Let

$$\bar{\rho}_t = \frac{1}{m} \sum_{i=1}^m \delta_{(a_i(t), u_i(t))}, \quad (134)$$

where $(a_i(t), u_i(t))_{1 \leq i \leq m}$ is the solution of (5)–(6). $\bar{\rho}_t$ is a measure on $\mathbb{R} \times \mathbb{S}^{d-1}$ solving the continuity PDE

$$\begin{aligned}
\partial_t \bar{\rho}_t(a, u) &= -\nabla \cdot (\bar{\rho}_t \bar{\Psi}(a, u; \bar{\rho}_t)) \\
&= -\left(\partial_a (\bar{\rho}_t \bar{\Psi}_a(a, u; \bar{\rho}_t)) + \partial_u (\bar{\rho}_t \bar{\Psi}_u(a, u; \bar{\rho}_t)) \right),
\end{aligned} \quad (135)$$

where $\bar{\Psi} = (\bar{\Psi}_a, \bar{\Psi}_u)$ is given by

$$\begin{aligned}
\bar{\Psi}_a(a, u; \bar{\rho}) &= \varepsilon^{-1} \left(V(\langle u, u_* \rangle) - \int_{\mathbb{R} \times \mathbb{S}^{d-1}} a_1 U(\langle u, u_1 \rangle) \bar{\rho}(da_1, du_1) \right), \\
\bar{\Psi}_u(a, u; \bar{\rho}) &= a \left(I_d - uu^\top \right) \left(V'(\langle u, u_* \rangle) u_* - \int_{\mathbb{R} \times \mathbb{S}^{d-1}} a_1 U'(\langle u, u_1 \rangle) u_1 \bar{\rho}(da_1, du_1) \right).
\end{aligned}$$

A remarkable property of the equation (135) is that it preserves invariance to rotations orthogonal to u_* . Indeed, assume that $\bar{\rho}$ is invariant to rotations orthogonal to u_* . In this case, we show that $\bar{\Psi}_a(a, u; \bar{\rho})$ and $\langle u_*, \bar{\Psi}_u(a, u; \bar{\rho}) \rangle$ depend only on $s := \langle u, u_* \rangle$ and $s_1 := \langle u_1, u_* \rangle$. Let u^\perp (resp. u_1^\perp) denote the component of u (resp. u_1) orthogonal to u_* . Let R denote a random uniform rotation orthogonal to u_* . By the rotation invariance of $\bar{\rho}$,

$$\begin{aligned}
\bar{\Psi}_a(a, u; \bar{\rho}) &= \varepsilon^{-1} \left(V(\langle u, u_* \rangle) - \int_{\mathbb{R} \times \mathbb{S}^{d-1}} a_1 \mathbb{E}_R [U(\langle u, Ru_1 \rangle)] \bar{\rho}(da_1, du_1) \right) \\
&= \varepsilon^{-1} \left(V(s) - \int_{\mathbb{R} \times \mathbb{S}^{d-1}} a_1 \mathbb{E}_R [U(ss_1 + \langle u^\perp, Ru_1^\perp \rangle)] \bar{\rho}(da_1, du_1) \right).
\end{aligned}$$

The random variable $B^{(d)} = \left\langle \frac{u^\perp}{\|u^\perp\|}, R \frac{u_1^\perp}{\|u_1^\perp\|} \right\rangle$ is a one dimensional projection of a random variable uniform on the unit sphere of the hyperplane orthogonal to u_* ; thus it has the density $p_{B^{(d)}}(b) \propto (1 - b^2)^{d/2-2}$ (see, e.g., [16, Lemma 4.17]). Denote

$$U^{(d)}(s, s_1) = \mathbb{E}_{B^{(d)}} \left[U \left(ss_1 + (1 - s^2)^{1/2} (1 - s_1^2)^{1/2} B^{(d)} \right) \right],$$

then we have

$$\bar{\Psi}_a(a, u; \bar{\rho}) = \varepsilon^{-1} \left(V(s) - \int_{\mathbb{R} \times \mathbb{S}^{d-1}} a_1 U^{(d)}(s, s_1) \bar{\rho}(da_1, du_1) \right). \quad (136)$$

Further, we compute

$$\begin{aligned} \langle u_*, \bar{\Psi}_u(a, u; \bar{\rho}) \rangle &= \left\langle u_*, a \left(I_d - uu^\top \right) \right. \\ &\quad \left(V'(s)u_* - \int_{\mathbb{R} \times \mathbb{S}^{d-1}} a_1 \mathbb{E}_R \left[U'(ss_1 + \langle u^\perp, Ru_1^\perp \rangle) (s_1 u_* + Ru_1^\perp) \right] \bar{\rho}(da_1, du_1) \right) \left. \right\rangle \\ &= a \left[(1 - s^2) V'(s) - \int_{\mathbb{R} \times \mathbb{S}^{d-1}} a_1 \mathbb{E}_R \left[U'(ss_1 + \langle u^\perp, Ru_1^\perp \rangle) \langle u_*, (I_d - uu^\top)(s_1 u_* + Ru_1^\perp) \rangle \right] \bar{\rho}(da_1, du_1) \right]. \end{aligned}$$

In the equation above, we have $\langle u_*, (I_d - uu^\top)s_1 u_* \rangle = s_1(1 - s^2)$ and as $\langle u_*, Ru_1^\perp \rangle = 0$ a.s., we have

$$\langle u_*, (I_d - uu^\top)Ru_1^\perp \rangle = -\langle u, u_* \rangle \langle u, Ru_1^\perp \rangle = -s \langle u^\perp, Ru_1^\perp \rangle.$$

Thus we obtain

$$\begin{aligned} \langle u_*, \bar{\Psi}_u(a, u; \bar{\rho}) \rangle &= a(1 - s^2) \left[V'(s) - \int_{\mathbb{R} \times \mathbb{S}^{d-1}} a_1 \mathbb{E}_R \left[U'(ss_1 + \langle u^\perp, Ru_1^\perp \rangle) \left(s_1 - \frac{s}{1 - s^2} \langle u^\perp, Ru_1^\perp \rangle \right) \right] \bar{\rho}(da_1, du_1) \right]. \end{aligned}$$

Note that

$$\partial_s U^{(d)}(s, s_1) = \mathbb{E}_{B^{(d)}} \left[U' \left(ss_1 + (1 - s^2)^{1/2} (1 - s_1^2)^{1/2} B^{(d)} \right) \left(s_1 - \frac{s}{(1 - s^2)^{1/2}} (1 - s_1^2)^{1/2} B^{(d)} \right) \right]$$

and thus we have

$$\langle u_*, \bar{\Psi}_u(a, u; \bar{\rho}) \rangle = a(1 - s^2) \left[V'(s) - \int_{\mathbb{R} \times \mathbb{S}^{d-1}} a_1 \left(\partial_s U^{(d)} \right) (s, s_1) \bar{\rho}(da_1, du_1) \right]. \quad (137)$$

Of course, a discrete measure of the form (134) can not be invariant to rotations orthogonal to u_* . However, if the u_i are initialized uniformly on the unit sphere, then the measure $\bar{\rho}_0$ converges to a measure with the rotation invariance as $m \rightarrow \infty$. One can then apply the results of [33] to control the deviations from this limit. Let us thus assume that $\bar{\rho}_0$ satisfies the rotation invariance. Define the map $\varphi(a, u) = (a, \langle u, u_* \rangle)$. Then, from (136), (137), the push-forward ρ_t of $\bar{\rho}_t$ through the map φ satisfies the continuity equation

$$\begin{aligned} \partial_t \rho_t(a, s) &= -\nabla \cdot \left(\rho_t \Psi^{(d)}(a, s; \rho_t) \right) \\ &= - \left(\partial_a \left(\rho_t \Psi_a^{(d)}(a, s; \rho_t) \right) + \partial_s \left(\rho_t \Psi_s^{(d)}(a, s; \rho_t) \right) \right), \end{aligned}$$

where $\Psi^{(d)} = (\Psi_a^{(d)}, \Psi_s^{(d)})$ is given by

$$\Psi_a^{(d)}(a, s; \rho) = \varepsilon^{-1} \cdot \left(V(s) - \int_{\mathbb{R}^2} a_1 U^{(d)}(s, s_1) \rho(da_1, ds_1) \right),$$

$$\Psi_s^{(d)}(a, s; \rho) = a(1 - s^2) \cdot \left(V'(s) - \int_{\mathbb{R}^2} a_1 \partial_s U^{(d)}(s, s_1) \rho(da_1, ds_1) \right).$$

When $d \rightarrow \infty$, $p_{B^{(d)}}(b)db \propto (1 - b^2)^{d/2-2}db$ converges weakly to the Dirac mass $\delta_0(db)$. As a consequence,

$$U^{(d)}(s, s_1) \xrightarrow{d \rightarrow \infty} U(ss_1), \quad \partial_s U^{(d)}(s, s_1) \xrightarrow{d \rightarrow \infty} U'(ss_1)s_1.$$

As a consequence, in the limit $d \rightarrow \infty$, we recover the equations (29)–(32). Moreover, if $\bar{\rho}_0 = P_A \otimes \text{Unif}(\mathbb{S}^{d-1})$, then ρ_0 converges weakly to $P_A \otimes \delta_0(ds)$ as $d \rightarrow \infty$.

B.6 Proof of Proposition 4

First, note that the potential functions U and V admit the following expansion:

$$U(s) = \sum_{k=0}^{\infty} \sigma_k^2 s^k, \quad V(s) = \sum_{k=0}^{\infty} \varphi_k \sigma_k s^k.$$

As a consequence, we deduce that

$$\begin{aligned} \mathcal{R}_{\text{mf},*}(\rho) &= \frac{1}{2} \sum_{k=0}^{\infty} \varphi_k^2 - \sum_{k=0}^{\infty} \varphi_k \sigma_k \int a(\omega) s(\omega)^k d\rho(\omega) + \frac{1}{2} \sum_{k=0}^{\infty} \sigma_k^2 \int a(\omega_1) a(\omega_2) s(\omega_1)^k s(\omega_2)^k d\rho(\omega_1) d\rho(\omega_2) \\ &= \frac{1}{2} \sum_{k=0}^{\infty} \left(\varphi_k^2 - 2\varphi_k \sigma_k \int a(\omega) s(\omega)^k d\rho(\omega) + \sigma_k^2 \left(\int a(\omega) s(\omega)^k d\rho(\omega) \right)^2 \right) \\ &= \frac{1}{2} \sum_{k=0}^{\infty} \left(\varphi_k - \sigma_k \int a(\omega) s(\omega)^k d\rho(\omega) \right)^2. \end{aligned}$$

Now we show that the above risk can be arbitrarily small. We will choose ρ to be the Lebesgue measure on $[-1, 1]$ and $a \in L^2[-1, 1]$ so that $\int a(\omega) s(\omega)^k d\rho(\omega) = \int_{-1}^1 a(s) s^k ds$. Now, we define the following set of sequences

$$W = \left\{ \left(\sigma_k \int_{-1}^1 a(s) s^k ds \right)_{k \geq 0} \mid a \in L^2[-1, 1] \right\}.$$

Since $a \in L^2[-1, 1]$ and $(\sigma_k)_{k \geq 0} \in \ell^2$, we know that $W \subset \ell^2$, i.e., W is a linear subspace of ℓ^2 . Now it suffices to show that W is dense in ℓ^2 , which is equivalent to $W^\perp = \{0\}$, namely

$$v \in \ell^2, \quad v \perp W \implies v = 0.$$

Fix any such v and take $\mu \in W$ such that for all k , $\mu_k = \sigma_k \int_{-1}^1 a(s) s^k ds$ for some $a \in L^2[-1, 1]$. We then have

$$0 = \langle v, \mu \rangle = \sum_{k=0}^{\infty} v_k \mu_k = \sum_{k=0}^{\infty} v_k \sigma_k \int_{-1}^1 a(s) s^k ds = \int_{-1}^1 a(s) \cdot \left(\sum_{k=0}^{\infty} v_k \sigma_k s^k \right) ds,$$

where the last step follows from dominated convergence theorem. Indeed, by Hölder's inequality,

$$\sum_{k=0}^{\infty} |v_k \sigma_k| \leq \left(\sum_{k=0}^{\infty} v_k^2 \right)^{1/2} \left(\sum_{k=0}^{\infty} \sigma_k^2 \right)^{1/2} < \infty.$$

As a consequence, the function series $\sum_{k=0}^n \sigma_k v_k s^k$ uniformly absolutely converges to the continuous function $f(s) = \sum_{k=0}^{\infty} \sigma_k v_k s^k$ on $[-1, 1]$. The above argument then implies that for any $a \in L^2[-1, 1]$, $\int_{-1}^1 a(s) f(s) ds = 0$, which further implies that $f(s) \equiv 0$. Therefore, $\sigma_k v_k = 0$ for all $k \geq 0$. Since $\sigma_k \neq 0$ for all $k \geq 0$, we must have $v_k = 0$ for all $k \geq 0$, i.e., $v = 0$. This completes the proof of the density of W in ℓ^2 , and thus the proof of the Proposition.

C Calculations for the analysis of mean-field gradient flow

C.1 Solution of Eq. (89)

In order to solve the system (89), we start from an associated one-dimensional ODE.

Lemma 2. *The solution $\lambda = \lambda(t_3)$ of the ODE*

$$\partial_{t_3} \lambda = |\sigma_1| \left(|\varphi_1| - |\sigma_1| \|a_{\perp, \text{init}}\|_{L^2(\rho)}^2 \lambda^2 \right) \lambda \quad (138)$$

with initial condition $\lambda(0)$ is

$$\lambda = \frac{|\varphi_1|^{1/2}}{\left(|\sigma_1| \|a_{\perp, \text{init}}\|_{L^2(\rho)}^2 + \left(|\varphi_1| \lambda(0)^{-2} - |\sigma_1| \|a_{\perp, \text{init}}\|_{L^2(\rho)}^2 \right) e^{-2|\sigma_1| \varphi_1 |t_3|} \right)^{1/2}}. \quad (139)$$

Proof. For simplicity, denote $\alpha = |\sigma_1|$, $\beta = |\varphi_1|$ and $\gamma = |\sigma_1| \|a_{\perp, \text{init}}\|_{L^2(\rho)}^2$. Then

$$\partial_{t_3} \lambda = \alpha \left(\beta - \gamma \lambda^2 \right) \lambda.$$

This is Bernoulli differential equation (see, e.g., [15]). In this situation, the classical trick is to reduce the problem to a linear inhomogeneous first-order equation by considering

$$\begin{aligned} \partial_{t_3} \left(\lambda^{-2} \right) &= -2 \left(\partial_{t_3} \lambda \right) \lambda^{-3} = -2\alpha \left(\beta - \gamma \lambda^2 \right) \lambda^{-2} \\ &= 2\alpha(\gamma - \beta \lambda^{-2}). \end{aligned}$$

Solving this linear inhomogeneous first-order equation gives

$$\lambda^{-2} = \frac{\gamma}{\beta} + \left(\lambda(0)^{-2} - \frac{\gamma}{\beta} \right) e^{-2\alpha\beta t_3},$$

and thus

$$\lambda = \frac{\beta^{1/2}}{(\gamma + (\beta \lambda(0)^{-2} - \gamma) e^{-2\alpha\beta t_3})^{1/2}},$$

which is the claimed result. \square

Let $\lambda = \lambda(t_3)$ be a solution of (138) and consider

$$\begin{aligned} a^{(-1)} &= a_{\perp}^{(-1)} = \lambda a_{\perp, \text{init}}, \\ s^{(1)} &= s_{\perp}^{(1)} = \text{sign}(\sigma_1 \varphi_1) \lambda a_{\perp, \text{init}}. \end{aligned} \quad (140)$$

Then $a^{(-1)}, s^{(1)}$ are solutions of the constrained ODE system (85), (88). Indeed,

$$\left\langle a^{(-1)}, \mathbf{1} \right\rangle_{L^2(\rho)} = \lambda \left\langle a_{\perp, \text{init}}, \mathbf{1} \right\rangle_{L^2(\rho)} = 0,$$

thus the constraint (85) is satisfied. Further

$$\begin{aligned} \partial_{t_3} a_{\perp}^{(-1)} &= (\partial_{t_3} \lambda) a_{\perp, \text{init}} \\ &= |\sigma_1| \left(|\varphi_1| - |\sigma_1| \|a_{\perp, \text{init}}\|_{L^2(\rho)}^2 \lambda^2 \right) \lambda a_{\perp, \text{init}} \\ &= \text{sign}(\sigma_1 \varphi_1) \sigma_1 \left(\varphi_1 - \text{sign}(\sigma_1 \varphi_1) \sigma_1 \|a_{\perp, \text{init}}\|_{L^2(\rho)}^2 \lambda^2 \right) \lambda a_{\perp, \text{init}} \\ &= \sigma_1 \left(\varphi_1 - \sigma_1 \left\langle a_{\perp}^{(-1)}, s_{\perp}^{(1)} \right\rangle_{L^2(\rho)} \right) s_{\perp}^{(1)}. \end{aligned}$$

A similar computation shows that the differential equation for $s^{(1)}$ is also satisfied. This concludes that (140) is a valid candidate to solve the third time scale.

Matching. To determine the value of the initialization $\lambda(0)$ we perform a matching procedure with the previous time scale. In this paragraph, we denote $\underline{a}, \underline{s}$ the approximation obtained in the second time scale (Section 6.3), and \bar{a}, \bar{s} the approximation in the third time scale (Section 6.4 and above).

Consider an intermediate time scale $\tilde{t} = t_2 - c \log \frac{1}{\varepsilon}$ with $0 < c < \frac{1}{4|\sigma_1\varphi_1|}$. Assume $\tilde{t} \asymp 1$. Then

$$\begin{aligned} t_2 &= \tilde{t} + c \log \frac{1}{\varepsilon} \xrightarrow{\varepsilon \rightarrow 0} +\infty, \\ t_3 &= \tilde{t} - \left(\frac{1}{4|\sigma_1\varphi_1|} - c \right) \log \frac{1}{\varepsilon} \xrightarrow{\varepsilon \rightarrow 0} -\infty. \end{aligned}$$

From the approximation (82) on the second time scale,

$$\begin{aligned} \underline{a} &= \underline{a}^{(0)} + O(\varepsilon^{1/2}) \\ &= \frac{\varphi_0}{\sigma_0} \mathbf{1} + \cosh(\varphi_1 \sigma_1 t_2) a_{\perp, \text{init}} + O(\varepsilon^{1/2}) \\ &= \cosh\left(\varphi_1 \sigma_1 \left(\tilde{t} + c \log \frac{1}{\varepsilon}\right)\right) a_{\perp, \text{init}} + O(1) \\ &= \frac{1}{2} e^{|\varphi_1 \sigma_1| \tilde{t} - c|\varphi_1 \sigma_1|} a_{\perp, \text{init}} + O(1). \end{aligned} \tag{141}$$

From the approximation on the third time scale,

$$\begin{aligned} \bar{a} &= \varepsilon^{-1/4} \bar{a}^{(-1)} + O(1) \\ &= \varepsilon^{-1/4} \lambda a_{\perp, \text{init}} + O(1). \end{aligned}$$

Note that as $t_3 \rightarrow -\infty$, from (139),

$$\lambda \sim \frac{|\varphi_1|^{1/2}}{\left(|\varphi_1| \lambda(0)^{-2} - |\sigma_1| \|a_{\perp, \text{init}}\|_{L^2(\rho)}^2\right)^{1/2}} e^{|\sigma_1 \varphi_1| t_3}.$$

Thus

$$\begin{aligned} \bar{a} &\sim \varepsilon^{-1/4} \frac{|\varphi_1|^{1/2}}{\left(|\varphi_1| \lambda(0)^{-2} - |\sigma_1| \|a_{\perp, \text{init}}\|_{L^2(\rho)}^2\right)^{1/2}} e^{|\sigma_1 \varphi_1| \left(\tilde{t} - \left(\frac{1}{4|\sigma_1 \varphi_1|} - c\right) \log \frac{1}{\varepsilon}\right)} a_{\perp, \text{init}} \\ &\sim \frac{|\varphi_1|^{1/2}}{\left(|\varphi_1| \lambda(0)^{-2} - |\sigma_1| \|a_{\perp, \text{init}}\|_{L^2(\rho)}^2\right)^{1/2}} e^{|\sigma_1 \varphi_1| \tilde{t} - c|\sigma_1 \varphi_1|} a_{\perp, \text{init}}. \end{aligned} \tag{142}$$

By matching, Equations (141) and (142) should be coherent. This gives

$$\frac{1}{2} = \frac{|\varphi_1|^{1/2}}{\left(|\varphi_1| \lambda(0)^{-2} - |\sigma_1| \|a_{\perp, \text{init}}\|_{L^2(\rho)}^2\right)^{1/2}},$$

and thus

$$\lambda(t_3) = \frac{|\varphi_1|^{1/2}}{\left(|\sigma_1| \|a_{\perp, \text{init}}\|_{L^2(\rho)}^2 + 4|\varphi_1| e^{-2|\sigma_1 \varphi_1| t_3}\right)^{1/2}}. \tag{143}$$

One could check similarly that $s^{(1)}$ also satisfies the matching conditions under the same constraint, and thus that (140) are indeed the solutions of the third time scale.

C.2 Induced approximation of the risk

In this section, we show that the behavior of a and s derived in Sections 6.2–6.4 leads to an evolution of the risk alternating plateaus and rapid decreases, in agreement with the canonical learning order of Definition 1. For the convenience of the reader, we recall the expression (37) of the risk

$$\begin{aligned}\mathcal{R}_{\text{mf},*} &= \frac{1}{2}\|\varphi\|^2 - \int a(\omega)V(s(\omega))d\rho(\omega) + \frac{1}{2} \int a(\omega_1)a(\omega_2)U(s(\omega_1)s(\omega_2))d\rho(\omega_1)d\rho(\omega_2) \\ &= \frac{1}{2} \sum_{k \geq 0} \left(\varphi_k - \sigma_k \int a(\omega)s(\omega)^k d\rho(\omega) \right)^2.\end{aligned}$$

First time scale $t_1 = \frac{t}{\varepsilon}$ (Section 6.2). On this time scale, we have $a = O(1)$ and $s = O(\varepsilon)$. Thus for all $k \geq 1$, $\int a(\omega)s(\omega)^k d\rho(\omega) = O(\varepsilon)$ whence $\left(\varphi_k - \sigma_k \int a(\omega)s(\omega)^k d\rho(\omega) \right)^2 = \varphi_k^2 + O(\varepsilon)$.

Further, using (59),

$$\begin{aligned}\left(\varphi_0 - \sigma_0 \int a(\omega)d\rho(\omega) \right)^2 &= \left(\varphi_0 - \sigma_0 \langle a, \mathbb{1} \rangle_{L^2(\rho)} \right)^2 = \left(\varphi_0 - \sigma_0 \langle a^{(0)}, \mathbb{1} \rangle_{L^2(\rho)} + O(\varepsilon) \right)^2 \\ &= e^{-2\sigma_0^2 t_1} \left(\varphi_0 - \sigma_0 \langle a_{\text{init}}, \mathbb{1} \rangle_{L^2(\rho)} \right)^2 + O(\varepsilon).\end{aligned}$$

Thus as $\varepsilon \rightarrow 0$,

$$\mathcal{R}_{\text{mf},*} = \frac{1}{2} e^{-2\sigma_0^2 t_1} \left(\varphi_0 - \sigma_0 \langle a_{\text{init}}, \mathbb{1} \rangle_{L^2(\rho)} \right)^2 + \frac{1}{2} \sum_{k \geq 1} \varphi_k^2 + O(\varepsilon).$$

This describes, in a more detailed form, the first transition in Definition 1.

Second time scale $t_2 = \frac{t}{\varepsilon^{1/2}}$ (Section 6.3). On this time scale, we have $a = O(1)$ and $s = O(\varepsilon^{1/2})$. Thus for all $k \geq 1$, $\left(\varphi_k - \sigma_k \int a(\omega)s(\omega)^k d\rho(\omega) \right)^2 = \varphi_k^2 + O(\varepsilon^{1/2})$.

Further, using (68),

$$\left(\varphi_0 - \sigma_0 \int a(\omega)d\rho(\omega) \right)^2 = \left(\varphi_0 - \sigma_0 \int a^{(0)}(\omega)d\rho(\omega) + O(\varepsilon^{1/2}) \right)^2 = O(\varepsilon^{1/2}).$$

Thus as $\varepsilon \rightarrow 0$,

$$\mathcal{R}_{\text{mf},*} = \frac{1}{2} \sum_{k \geq 1} \varphi_k^2 + O(\varepsilon^{1/2}).$$

This second time scale does not induce any transition of the risk $\mathcal{R}_{\text{mf},*}$ (but was necessary to understand the divergence of a and $\varepsilon^{-1/2}s$).

Third time scale $t_3 = \frac{t}{\varepsilon^{1/2}} - \frac{1}{4|\sigma_1\varphi_1|} \log \frac{1}{\varepsilon}$ (Section 6.4). On this time scale, we have $a = O(\varepsilon^{-1/4})$ and $s = O(\varepsilon^{1/4})$. Thus for all $k \geq 2$, $\left(\varphi_k - \sigma_k \int a(\omega)s(\omega)^k d\rho(\omega) \right)^2 = \varphi_k^2 + O(\varepsilon^{1/4})$.

Further, using (85), (86),

$$\begin{aligned}\left(\varphi_0 - \sigma_0 \int a(\omega)d\rho(\omega) \right)^2 &= \left(\varphi_0 - \sigma_0 \varepsilon^{-1/4} \int a^{(-1)}(\omega)d\rho(\omega) - \sigma_0 \int a^{(0)}(\omega)d\rho(\omega) + O(\varepsilon^{1/4}) \right)^2 \\ &= O(\varepsilon^{1/4}).\end{aligned}$$

Finally, using (90), (91),

$$\left(\varphi_1 - \sigma_1 \int a(\omega)s(\omega)d\rho(\omega) \right)^2 = \left(\varphi_1 - \sigma_1 \langle a, s \rangle_{L^2(\rho)} \right)^2 = \left(\varphi_1 - \sigma_1 \langle a^{(-1)}, s^{(1)} \rangle_{L^2(\rho)} + O(\varepsilon^{1/4}) \right)^2$$

$$\begin{aligned}
&\stackrel{(a)}{=} \left(\varphi_1 - \sigma_1 \operatorname{sign}(\sigma_1 \varphi_1) \lambda^2 \|a_{\perp, \text{init}}\|_{L^2(\rho)}^2 \right)^2 + O(\varepsilon^{1/4}) \\
&\stackrel{(b)}{=} \left(\varphi_1 - \frac{\sigma_1 \operatorname{sign}(\sigma_1 \varphi_1) \|\varphi_1\| \|a_{\perp, \text{init}}\|_{L^2(\rho)}^2}{(|\sigma_1| \|a_{\perp, \text{init}}\|_{L^2(\rho)}^2 + 4|\varphi_1| e^{-2|\sigma_1 \varphi_1| t_3})} \right)^2 + O(\varepsilon^{1/4}) \\
&= \varphi_1^2 \left(1 - \frac{1}{1 + \frac{4|\varphi_1|}{|\sigma_1| \|a_{\perp, \text{init}}\|_{L^2(\rho)}^2} e^{-2|\sigma_1 \varphi_1| t_3}} \right)^2 + O(\varepsilon^{1/4}),
\end{aligned}$$

where in (a) we used (90) and in (b) (91). Thus as $\varepsilon \rightarrow 0$,

$$\mathcal{R}_{\text{mf},*} = \frac{1}{2} \varphi_1^2 \left(1 - \frac{1}{1 + \frac{4|\varphi_1|}{|\sigma_1| \|a_{\perp, \text{init}}\|_{L^2(\rho)}^2} e^{-2|\sigma_1 \varphi_1| t_3}} \right)^2 + \frac{1}{2} \sum_{k \geq 2} \varphi_k^2 + O(\varepsilon^{1/4}).$$

This describes, in a more detailed form, the second transition in Definition 1.

C.3 Proof of Theorem 1

Throughout the proof, we will use the shorthand $\mathcal{R}_l(\tau)$ to represent $\mathcal{R}_l(\tilde{a}(\tau), \tilde{s}(\tau))$. First, note that according to the ODE satisfied by \mathcal{R}_l (Eq. (98)), we know that \mathcal{R}_l must be non-increasing, thus for small enough $\varepsilon > 0$,

$$\begin{aligned}
\left| \varphi_l - \sigma_l \int \tilde{a}(\nu, \tau) \tilde{s}(\nu, \tau)^l d\rho(\nu) \right| &\leq \left| \varphi_l - \sigma_l \int \tilde{a}(\nu, 0) \tilde{s}(\nu, 0)^l d\rho(\nu) \right| \\
&\leq |\varphi_l| + O(\varepsilon^{1/2l}) \leq 2|\varphi_l|, \quad \forall \tau \geq 0.
\end{aligned}$$

Hence, we obtain the estimates:

$$\partial_\tau |\tilde{a}(\omega)| \leq |\partial_\tau \tilde{a}(\omega)| \leq 2|\sigma_l| |\varphi_l| |\tilde{s}(\omega)|^l, \quad \partial_\tau |\tilde{s}(\omega)| \leq |\partial_\tau \tilde{s}(\omega)| \leq 2l|\sigma_l| |\varphi_l| |\tilde{a}(\omega)| |\tilde{s}(\omega)|^{l-1}.$$

According to the comparison theorem for system of ODEs, we know that $|\tilde{a}(\omega, \tau)| \leq \hat{a}(\omega, \tau)$, $|\tilde{s}(\omega, \tau)| \leq \hat{s}(\omega, \tau)$ for all $\tau \geq 0$ where

$$\hat{a}(\omega, 0) = \max \{ |\tilde{a}(\omega, 0)|, |\tilde{s}(\omega, 0)| \}, \quad \hat{s}(\omega, 0) = l^{1/2} \hat{a}(\omega, 0) = l^{1/2} \max \{ |\tilde{a}(\omega, 0)|, |\tilde{s}(\omega, 0)| \},$$

and

$$\partial_\tau \hat{a}(\omega) = 2|\sigma_l| |\varphi_l| \hat{s}(\omega)^l, \quad \partial_\tau \hat{s}(\omega) = 2l|\sigma_l| |\varphi_l| \hat{a}(\omega) \hat{s}(\omega)^{l-1}. \quad (144)$$

The above system of ODEs can be solved analytically via integration. First, we note that

$$\partial_\tau (\hat{s}(\omega)^2) = 4l|\sigma_l| |\varphi_l| \hat{a}(\omega) \hat{s}(\omega)^l = l \partial_\tau (\hat{a}(\omega)^2),$$

which implies that (further note $\hat{s}(\omega, 0)^2 = l \hat{a}(\omega, 0)^2$)

$$\hat{s}(\omega, \tau) = l^{1/2} \hat{a}(\omega, \tau), \quad \forall \tau \geq 0. \quad (145)$$

The ODE system then reduces to $\partial_\tau \hat{a}(\omega) = 2l^{l/2} |\sigma_l| |\varphi_l| \hat{a}(\omega)^l$, which admits the solution

$$\hat{a}(\omega, \tau) = \left(\hat{a}(\omega, 0)^{-l+1} - 2l^{l/2} (l-1) |\sigma_l| |\varphi_l| \tau \right)^{-1/(l-1)}. \quad (146)$$

Since $\hat{a}(\omega, 0) = \Theta(\varepsilon^{1/2l(l+1)})$, we know that $\hat{a}(\omega, \tau), \hat{s}(\omega, \tau) = o(1)$ until $\tau = \Theta(\varepsilon^{-(l-1)/2l(l+1)}) - O(1) = \Theta(\varepsilon^{-(l-1)/2l(l+1)})$, which means that $\tilde{a}(\omega, \tau), \tilde{s}(\omega, \tau) = o(1)$ until $\tau = \Omega(\varepsilon^{-(l-1)/2l(l+1)})$. As a consequence,

$$\mathcal{R}_l(\tau) = \frac{1}{2} \left(\varphi_l - \sigma_l \int \tilde{a}(\nu, \tau) \tilde{s}(\nu, \tau)^l d\rho(\nu) \right)^2 = \frac{1}{2} \varphi_l^2 - o(1)$$

until $\tau = \Omega(\varepsilon^{-(l-1)/2l(l+1)})$. This means that the learning of the l -th component will not begin until $\tau = \Omega(\varepsilon^{-(l-1)/2l(l+1)})$, namely $\tau(\Delta) = \Omega(\varepsilon^{-(l-1)/2l(l+1)})$ for any fixed $\Delta > 0$. Note that the above argument applies to all of the settings in the theorem statement.

Next, we show that for any fixed $\Delta > 0$, $\tau(\Delta) = O(\varepsilon^{-(l-1)/2l(l+1)})$, which means that the l -th component can be learnt in $O(\varepsilon^{-(l-1)/2l(l+1)})$ time. To prove our claim by contradiction, assume that there exists $\Delta > 0$ and a sequence $\varepsilon_k \downarrow 0$, such that

$$\lim_{k \rightarrow \infty} \frac{\tau(\Delta)}{\varepsilon_k^{-(l-1)/2l(l+1)}} = +\infty. \quad (147)$$

By definition of $\tau(\Delta)$, we know that $\forall \tau \leq \tau(\Delta)$,

$$\left| \varphi_l - \sigma_l \int \tilde{a}(\nu, \tau) \tilde{s}(\nu, \tau)^l d\rho(\nu) \right| \geq \sqrt{2\Delta}.$$

Now, assume the condition of setting (a) holds and denote

$$A_{\varepsilon_0, \eta} = \left\{ \omega : \forall \varepsilon < \varepsilon_0, \min(|\tilde{a}(\omega, 0)|, |\tilde{s}(\omega, 0)|) > \eta \varepsilon^{1/2l(l+1)}, \sigma_l \varphi_l \tilde{a}(\omega, 0) \tilde{s}(\omega, 0)^l > 0 \right\}. \quad (148)$$

Then by definition and our assumption that $\tilde{a}(\omega, 0)$ is of the same order as $\tilde{s}(\omega, 0)$, we know that $A = \cup_{\varepsilon_0 > 0, \eta > 0} A_{\varepsilon_0, \eta}$. Since $\rho(A) > 0$, there exists $\varepsilon_0, \eta > 0$ such that $\rho(A_{\varepsilon_0, \eta}) > 0$. Note that here we can choose ε_0 and η to be arbitrarily small since the set $A_{\varepsilon_0, \eta}$ is non-increasing in ε_0 and η . For $\omega \in A_{\varepsilon_0, \eta}$ and $\tau \leq \tau(\Delta)$, we have

$$\begin{aligned} \partial_\tau (\tilde{a}(\omega)^2) &= 2\sigma_l \tilde{a}(\omega) \tilde{s}(\omega)^l \left(\varphi_l - \sigma_l \int \tilde{a}(\nu) \tilde{s}(\nu)^l d\rho(\nu) \right) \geq 2\sqrt{2\Delta} \left| \sigma_l \tilde{a}(\omega) \tilde{s}(\omega)^l \right| \\ \partial_\tau (\tilde{s}(\omega)^2) &= 2l\sigma_l \tilde{a}(\omega) \tilde{s}(\omega)^l \left(1 - \varepsilon^{2\beta_l} \tilde{s}(\omega)^2 \right) \left(\varphi_l - \sigma_l \int \tilde{a}(\nu) \tilde{s}(\nu)^l d\rho(\nu) \right) \\ &\geq 2l\sqrt{2\Delta} \left| \sigma_l \tilde{a}(\omega) \tilde{s}(\omega)^l \right| \left(1 - \varepsilon^{2\beta_l} \tilde{s}(\omega)^2 \right). \end{aligned}$$

Moreover, we know that at initialization, $|\tilde{a}(\omega, 0)|, |\tilde{s}(\omega, 0)| > \eta \varepsilon^{1/2l(l+1)}$. Using the ODE comparison theorem and a similar argument as that in proving $\tau(\Delta) = \Omega(\varepsilon^{-(l-1)/2l(l+1)})$, we deduce that for sufficiently large k such that $\varepsilon = \varepsilon_k < \varepsilon_0$, there exist constants $C, C' > 0$ that does not depend on ε satisfying the following: For all $\omega \in A_{\varepsilon_0, \eta}$ and $\tau \geq C\varepsilon^{-(l-1)/2l(l+1)}$,

$$\min \{ |\tilde{a}(\omega, \tau)|, |\tilde{s}(\omega, \tau)| \} \geq C', \sigma_l \varphi_l \tilde{a}(\omega, \tau) \tilde{s}(\omega, \tau)^l > 0.$$

This further implies that at time τ ,

$$\int \tilde{s}(\omega)^{2(l-1)} \left(l^2 \tilde{a}(\omega)^2 \left(1 - \varepsilon^{2\beta_l} \tilde{s}(\omega)^2 \right) + \tilde{s}(\omega)^2 \right) d\rho(\omega) \geq C'^{2l} \rho(A_{\varepsilon_0, \eta}) > 0. \quad (149)$$

According to Eq. (98), we know that \mathcal{R}_l will decrease to 0 exponentially fast in an $O(1)$ time window after $\tau = C\varepsilon^{-(l-1)/2l(l+1)}$, which contradicts our assumption (147). This proves that $\tau(\Delta) = O(\varepsilon^{-(l-1)/2l(l+1)})$ under setting (a). Next, we show that setting (b) can be reduced to setting (a). Under setting (b), let us denote

$$B_{\varepsilon_0, \eta} = \left\{ \omega : \forall \varepsilon < \varepsilon_0, \sigma_l \varphi_l \tilde{a}(\omega, 0) \tilde{s}(\omega, 0)^l < 0, \text{ and } \tilde{s}(\omega, 0)^2 > (l + \eta) \tilde{a}(\omega, 0)^2 \right\}.$$

Then similar to the previous argument, there exists $\varepsilon_0, \eta > 0$ such that $\rho(B_{\varepsilon_0, \eta}) > 0$, and further we can choose ε_0 and η to be arbitrarily small. For $\omega \in B_{\varepsilon_0, \eta}$, we have

$$\partial_\tau (\tilde{a}(\omega)^2) < 0, \partial_\tau (\tilde{s}(\omega)^2) < 0 \text{ at } \tau = 0.$$

Hence, both $\tilde{a}(\omega)^2$ and $\tilde{s}(\omega)^2$ will decrease at initialization. Moreover, Eq. (97) implies that

$$\partial_\tau(\tilde{a}(\omega)^2) = \frac{\partial_\tau(\tilde{s}(\omega)^2)}{l(1 - \varepsilon^{2\beta_l}\tilde{s}(\omega)^2)} = -\frac{\varepsilon^{-2\beta_l}}{l}\partial_\tau \log(1 - \varepsilon^{2\beta_l}\tilde{s}(\omega)^2).$$

Integrating both sides of the above equation, we obtain that

$$\tilde{a}(\omega, 0)^2 - \tilde{a}(\omega, \tau)^2 = -\frac{\varepsilon^{-2\beta_l}}{l} \left(\log(1 - \varepsilon^{2\beta_l}\tilde{s}(\omega, 0)^2) - \log(1 - \varepsilon^{2\beta_l}\tilde{s}(\omega, \tau)^2) \right), \quad (150)$$

which is close to $(\tilde{s}(\omega, 0)^2 - \tilde{s}(\omega, \tau)^2)/l$ as long as $\tilde{s}(\omega, \tau) = O(1)$. To be accurate, let us define

$$\tau_{a,\omega} = \inf\{\tau \geq 0 : \tilde{a}(\omega, \tau) = 0\},$$

then we know that $\tilde{s}(\omega, \tau_{a,\omega}) = \Omega(\varepsilon^{1/2l(l+1)})$ and $\tau_{a,\omega} = O(\varepsilon^{-(l-1)/2l(l+1)})$ under the assumption (147), where the latter claim can be proved through making the change of variable $\tilde{a}'(\omega) = \varepsilon^{-1/2l(l+1)}\tilde{a}(\omega)$ and $\tilde{s}'(\omega) = \varepsilon^{-1/2l(l+1)}\tilde{s}(\omega)$. Note that after the time point $\tau_{a,\omega}$, the sign of $\tilde{a}(\omega)$ changes. Hence, $\varphi_l \sigma_l \tilde{a}(\omega) \tilde{s}(\omega)^l > 0$, and $\tilde{a}(\omega, \tau)^2$ and $\tilde{s}(\omega, \tau)^2$ will begin to increase for $\tau \geq \tau_{a,\omega}$. Similarly, we can show that in $O(\varepsilon^{-(l-1)/2l(l+1)})$ time after $\tau_{a,\omega}$, both $\tilde{a}(\omega)$ and $\tilde{s}(\omega)$ become of order $\varepsilon^{1/2l(l+1)}$, and we still have $\varphi_l \sigma_l \tilde{a}(\omega) \tilde{s}(\omega)^l > 0$. This reduces our case (b) to case (a).

We have proven that under settings (a) and (b), $\tau(\Delta) = \Theta(\varepsilon^{-(l-1)/2l(l+1)})$ for any fixed $\Delta \in (0, \varphi_l^2/2)$. This means that some of the neurons $(\tilde{a}(\omega), \tilde{s}(\omega))$ become of order $\Omega(1)$ and the l -th component of the target function is learnt at a timescale of order $\varepsilon^{-(l-1)/2l(l+1)}$. Next, we show that if the probability measure ρ is discrete, then the evolution of \mathcal{R}_l actually happens in an $O(1)$ time window. It suffices to prove that, for any $\Delta > 0$ a small constant ($\Delta < \varphi_l^2/4$),

$$\tau(\Delta) - \tau\left(\frac{\varphi_l^2}{2} - \Delta\right) = O(1) \quad (151)$$

as $\varepsilon \rightarrow 0$. Note that by continuity and monotonicity of \mathcal{R}_l , we have

$$\mathcal{R}_l(\tau(\Delta)) = \Delta, \quad \mathcal{R}_l\left(\tau\left(\frac{\varphi_l^2}{2} - \Delta\right)\right) = \frac{\varphi_l^2}{2} - \Delta, \quad \text{and} \quad \tau(\Delta) \geq \tau\left(\frac{\varphi_l^2}{2} - \Delta\right).$$

By definition of \mathcal{R}_l , we know that $\forall \tau \geq \tau(\varphi_l^2/2 - \Delta)$,

$$\left| \int \tilde{a}(\nu, \tau) \tilde{s}(\nu, \tau)^l d\rho(\nu) \right| \geq \frac{1}{|\sigma_l|} \left(|\varphi_l| - \sqrt{\varphi_l^2 - 2\Delta} \right) := r_l(\Delta) > 0.$$

Denote by $\{(\tilde{a}_i, \tilde{s}_i)\}_{i \in [m]}$ the realizations of $\{(\tilde{a}(\omega), \tilde{s}(\omega))\}_{\omega \in \Omega}$ under the discrete measure ρ , and by $\{p_i\}_{i \in [m]}$ the point masses of ρ . Then, we know that

$$r_l(\Delta) \leq \left| \int \tilde{a}(\nu, \tau) \tilde{s}(\nu, \tau)^l d\rho(\nu) \right| = \left| \sum_{j=1}^m p_j \tilde{a}_j(\tau) \tilde{s}_j(\tau)^l \right| \leq \sum_{j=1}^m p_j |\tilde{a}_j(\tau) \tilde{s}_j(\tau)^l|, \quad (152)$$

which implies that $\exists j \in [m]$, s.t. $|\tilde{a}_j(\tau) \tilde{s}_j(\tau)^l| \geq r_l(\Delta)$. Applying Lemma 3 yields

$$\int \tilde{s}(\omega)^{2(l-1)} \left(l^2 \tilde{a}(\omega)^2 (1 - \varepsilon^{2\beta_l} \tilde{s}(\omega)^2) + \tilde{s}(\omega)^2 \right) d\rho(\omega) \quad (153)$$

$$\geq p_j \tilde{s}_j^{2(l-1)} \left(l^2 \tilde{a}_j^2 (1 - \varepsilon^{2\beta_l} \tilde{s}_j^2) + \tilde{s}_j^2 \right) \geq \min_{j \in [m]} p_j \cdot c(l, r_l(\Delta)) > 0. \quad (154)$$

It then follows from Eq. (98) that \mathcal{R}_l will decrease to 0 exponentially fast, and Eq. (151) holds consequently. This completes the proof for settings (a) and (b).

We then focus on the case (c). By our assumption, for almost every ω there exists $\eta > 0$ (may depend on ω) such that

$$\sigma_l \varphi_l \tilde{a}(\omega, 0) \tilde{s}(\omega, 0)^l < -\eta \varepsilon^{1/2l}, \quad \tilde{s}(\omega, 0)^2 < (l - \eta) \tilde{a}(\omega, 0)^2$$

for sufficiently small ε . Therefore, $\tilde{s}(\omega, \tau)^2$ and $\tilde{a}(\omega, \tau)^2$ will keep decreasing until one of them reaches 0, which means that

$$\int \tilde{a}(\nu) \tilde{s}(\nu)^l d\rho(\nu) = o(1) \implies \left| \varphi_l - \sigma_l \int \tilde{a}(\nu) \tilde{s}(\nu)^l d\rho(\nu) \right| \geq \frac{|\varphi_l|}{2}. \quad (155)$$

According to Eq. (150) and the inequality $\tilde{s}(\omega, 0)^2 < (l - \eta) \tilde{a}(\omega, 0)^2$, $\tilde{a}(\omega, \tau)^2$ will not reach 0 until $\tilde{s}(\omega, \tau)^2$ reaches 0. Furthermore, for any $\tau \geq 0$,

$$\begin{aligned} \tilde{a}(\omega, \tau)^2 &= \tilde{a}(\omega, 0)^2 + \frac{\varepsilon^{-2\beta_l}}{l} \left(\log \left(1 - \varepsilon^{2\beta_l} \tilde{s}(\omega, 0)^2 \right) - \log \left(1 - \varepsilon^{2\beta_l} \tilde{s}(\omega, \tau)^2 \right) \right) \\ &\geq \tilde{a}(\omega, 0)^2 + \frac{\varepsilon^{-2\beta_l}}{l} \log \left(1 - \varepsilon^{2\beta_l} \tilde{s}(\omega, 0)^2 \right) \\ &\geq \frac{1}{l - \eta} \tilde{s}(\omega, 0)^2 - \frac{1}{l - \eta/2} \tilde{s}(\omega, 0)^2 \geq c(\eta, l, \omega) \varepsilon^{1/(l+1)}, \end{aligned}$$

thus leading to

$$\partial_\tau (\tilde{s}(\omega)^2) = 2l \sigma_l \tilde{a}(\omega) \tilde{s}(\omega)^l \left(1 - \varepsilon^{2\beta_l} \tilde{s}(\omega)^2 \right) \left(\varphi_l - \sigma_l \int \tilde{a}(\nu) \tilde{s}(\nu)^l d\rho(\nu) \right) \quad (156)$$

$$\leq -c(\eta, l, \omega) \varepsilon^{1/2l(l+1)} (\tilde{s}(\omega)^2)^{l/2}. \quad (157)$$

Using again the comparison theorem for ODE, we get that

$$|\tilde{s}(\omega, \tau)| \leq \left(|\tilde{s}(\omega, 0)|^{-l+2} + c(\eta, l, \omega) \varepsilon^{1/2l(l+1)} \tau \right)^{-1/(l-2)}. \quad (158)$$

Since $\tilde{s}(\omega, 0) \asymp \varepsilon^{1/2l(l+1)}$, it follows immediately that for any $\Delta > 0$, there exists a constant $C_*(\omega, \Delta) > 0$ such that

$$\tau \geq C_*(\omega, \Delta) \varepsilon^{-(l-1)/2l(l+1)} \implies |\tilde{s}(\omega, \tau)| \leq \Delta \varepsilon^{1/2l(l+1)}. \quad (159)$$

This completes the discussion for case (c), thus concluding the proof of Theorem 1.

Lemma 3. *Let $r > 0$ be a constant that does not depend on ε . Then there exists a constant $c = c(l, r) > 0$ that only depends on l and r such that the following holds: For any $a > 0$, $s > 0$ satisfying $as^l \geq r$ and $\varepsilon^{2\beta_l} s^2 \leq 1$, we have*

$$s^{2(l-1)} \left(l^2 a^2 \left(1 - \varepsilon^{2\beta_l} s^2 \right) + s^2 \right) \geq c. \quad (160)$$

Proof. If $s \geq 1$, then we immediately get

$$s^{2(l-1)} \left(l^2 a^2 \left(1 - \varepsilon^{2\beta_l} s^2 \right) + s^2 \right) \geq s^{2l} \geq 1.$$

Otherwise, $1 - \varepsilon^{2\beta_l} s^2 \geq 1/2$, and consequently

$$\begin{aligned} s^{2(l-1)} \left(l^2 a^2 \left(1 - \varepsilon^{2\beta_l} s^2 \right) + s^2 \right) &\geq s^{2(l-1)} \left(\frac{l^2 a^2}{2} + s^2 \right) \geq s^{2(l-1)} \left(\frac{l^2 r^2}{2s^{2l}} + s^2 \right) \\ &= \frac{l^2 r^2}{2s^2} + s^{2l} \geq c(l, r), \end{aligned}$$

where the last line follows from the AM-GM inequality. This completes the proof. \square

D Proofs of Theorem 2 and 3: learning with projected SGD

We will prove Theorem 2 which bounds the distance between GF and projected SGD in sub-Sections D.1 through D.3, with sub-Section D.4 devoted to the proof of Theorem 3. Throughout this section, we use M to refer to any constant that only depends on the M_i 's from Assumptions A1-A3, whereas the value of M can change from line to line. We start with an elementary lemma that establishes the Lipschitz continuity of the gradient flow trajectory:

Lemma 4 (A priori estimate). *There exists a constant M that only depends on the M_i 's, such that for all $t \geq 0$, ρ_t is supported on $[-M(1+t/\varepsilon), M(1+t/\varepsilon)] \times \mathbb{S}^{d-1}$, namely $|a_i(t)| \leq M(1+t/\varepsilon)$ for all $i \in [m]$. Moreover, for any $0 \leq s \leq t$, we have*

$$\begin{aligned} \sup_{j \in [m]} |a_j(s) - a_j(t)| &\leq \varepsilon^{-1} M(t-s), \\ \sup_{j \in [m]} \|u_j(s) - u_j(t)\|_2 &\leq M(1+t/\varepsilon)^2(t-s). \end{aligned}$$

Proof. First, notice that along the trajectory of gradient flow, the risk must be non-increasing. In fact, we have

$$\partial_t \mathcal{R} = -m \sum_{i=1}^m \left(\varepsilon^{-2} (\partial_{a_i} \mathcal{R})^2 + (\nabla_{u_i} \mathcal{R})^\top (I_d - u_i u_i^\top) (\nabla_{u_i} \mathcal{R}) \right) \leq 0.$$

Therefore, we obtain that

$$\begin{aligned} |\partial_t a_i| &= \varepsilon^{-1} \left| \mathbb{E}[y \sigma(\langle u_i, x \rangle)] - \frac{1}{m} \sum_{j=1}^m a_j \mathbb{E}[\sigma(\langle u_i, x \rangle) \sigma(\langle u_j, x \rangle)] \right| = \varepsilon^{-1} |\mathbb{E}[(y - \hat{y}) \sigma(\langle u_i, x \rangle)]| \\ &\leq \varepsilon^{-1} \mathbb{E}[(y - \hat{y})^2]^{1/2} \mathbb{E}[\sigma(\langle u_i, x \rangle)^2]^{1/2} \leq \varepsilon^{-1} \sqrt{2\mathcal{R}(0)} \|\sigma\|_{L^2} \leq \varepsilon^{-1} M, \end{aligned}$$

where the last line follows from our assumption. Since $|a_i(0)| \leq M$, we know that $|a_i(t)| \leq M(1+t/\varepsilon)$, and $|a_i(t) - a_i(s)| \leq \varepsilon^{-1} M(t-s)$. Moreover, according to Eq. (6), we have

$$\|\partial_t u_i\|_2 \leq |a_i| \left(\|V'\|_\infty + \|U'\|_\infty \sup_{j \in [m]} |a_j| \right) \leq M(1+t/\varepsilon)^2,$$

thus leading to

$$\|u_i(s) - u_i(t)\|_2 \leq M(1+t/\varepsilon)^2(t-s), \quad \forall i \in [m].$$

This completes the proof. \square

In what follows we define two discretized versions of Eqs (5) and (6), namely the gradient descent (GD) and stochastic gradient descent (SGD) dynamics. They will serve as important intermediate objects for our proof.

- Gradient descent: Let $\eta > 0$ be the step size, and let the initialization be the same as

gradient flow: $(\tilde{a}_i(0), \tilde{u}_i(0)) = (a_i(0), u_i(0))$ for all $i \in [m]$. We have for $k \in \mathbb{N}$,

$$\begin{aligned}
\tilde{a}_i(k+1) - \tilde{a}_i(k) &= -m\varepsilon^{-1}\eta\partial_{\tilde{a}_i(k)}\mathcal{R} \\
&= \varepsilon^{-1}\eta\left(V(\langle u_*, \tilde{u}_i(k) \rangle; \|u_*\|_2, \|\tilde{u}_i(k)\|_2) \right. \\
&\quad \left. - \frac{1}{m}\sum_{j=1}^m \tilde{a}_j(k)U(\langle \tilde{u}_i(k), \tilde{u}_j(k) \rangle; \|\tilde{u}_i(k)\|_2, \|\tilde{u}_j(k)\|_2) \right) \\
\tilde{u}_i(k+1) - \tilde{u}_i(k) &= -m\eta\left(I_d - \tilde{u}_i(k)\tilde{u}_i(k)^\top\right)\nabla_{\tilde{u}_i(k)}\mathcal{R} \\
&= \eta\tilde{a}_i(k)\left(I_d - \tilde{u}_i(k)\tilde{u}_i(k)^\top\right)\left(\nabla_{\tilde{u}_i(k)}V(\langle u_*, \tilde{u}_i(k) \rangle; \|u_*\|_2, \|\tilde{u}_i(k)\|_2) \right. \\
&\quad \left. - \frac{1}{m}\sum_{j=1}^m \tilde{a}_j(k)\nabla_{\tilde{u}_i(k)}U(\langle \tilde{u}_i(k), \tilde{u}_j(k) \rangle; \|\tilde{u}_i(k)\|_2, \|\tilde{u}_j(k)\|_2) \right),
\end{aligned} \tag{161}$$

where we recall from Eq.s (108) and (109):

$$\begin{aligned}
V(\langle u_*, \tilde{u}_i(k) \rangle; \|u_*\|_2, \|\tilde{u}_i(k)\|_2) &= \mathbb{E}[\varphi(\langle u_*, x \rangle)\sigma(\langle \tilde{u}_i(k), x \rangle)] = \mathbb{E}[y\sigma(\langle \tilde{u}_i(k), x \rangle)], \\
\nabla_{\tilde{u}_i(k)}V(\langle u_*, \tilde{u}_i(k) \rangle; \|u_*\|_2, \|\tilde{u}_i(k)\|_2) &= \mathbb{E}[\varphi(\langle u_*, x \rangle)\sigma'(\langle \tilde{u}_i(k), x \rangle)x] = \mathbb{E}[y\sigma'(\langle \tilde{u}_i(k), x \rangle)x], \\
U(\langle \tilde{u}_i(k), \tilde{u}_j(k) \rangle; \|\tilde{u}_i(k)\|_2, \|\tilde{u}_j(k)\|_2) &= \mathbb{E}[\sigma(\langle \tilde{u}_i(k), x \rangle)\sigma(\langle \tilde{u}_j(k), x \rangle)], \\
\nabla_{\tilde{u}_i(k)}U(\langle \tilde{u}_i(k), \tilde{u}_j(k) \rangle; \|\tilde{u}_i(k)\|_2, \|\tilde{u}_j(k)\|_2) &= \mathbb{E}[x\sigma'(\langle \tilde{u}_i(k), x \rangle)\sigma(\langle \tilde{u}_j(k), x \rangle)].
\end{aligned}$$

By convention, we have $V(s) = V(s; 1, 1)$ and $U(s) = U(s; 1, 1)$ for $s \in [-1, 1]$.

- One-pass stochastic gradient descent: Under the same choice of the step size and initialization, and let $\{(x_k, y_k)\}_{k \in \mathbb{N}^*}$ be i.i.d. samples from $P \in \mathcal{P}(\mathbb{R}^d \times \mathbb{R})$, where

$$P = \text{Law}(x, y), \quad x \sim \mathcal{N}(0, I_d), \quad y = \varphi(\langle u_*, x \rangle).$$

The iteration equations for one-pass SGD read:

$$\begin{aligned}
\underline{a}_i(k+1) - \underline{a}_i(k) &= \varepsilon^{-1}\eta\left(y_{k+1} - \frac{1}{m}\sum_{j=1}^m \underline{a}_j(k)\sigma(\langle \underline{u}_j(k), x_{k+1} \rangle)\right)\sigma(\langle \underline{u}_i(k), x_{k+1} \rangle) \\
\underline{u}_i(k+1) - \underline{u}_i(k) &= \eta\underline{a}_i(k)\left(I_d - \underline{u}_i(k)\underline{u}_i(k)^\top\right)\left(y_{k+1} - \frac{1}{m}\sum_{j=1}^m \underline{a}_j(k)\sigma(\langle \underline{u}_j(k), x_{k+1} \rangle)\right) \\
&\quad \times \sigma'(\langle \underline{u}_i(k), x_{k+1} \rangle)x_{k+1}.
\end{aligned} \tag{162}$$

Note that Eq. (162) can also be written as:

$$\begin{aligned}
\underline{a}_i(k+1) &= \underline{a}_i(k) + \varepsilon^{-1}\eta\hat{F}_i(\underline{\rho}^{(m)}(k); z_{k+1}) \\
\underline{u}_i(k+1) &= \underline{u}_i(k) + \eta\left(I_d - \underline{u}_i(k)\underline{u}_i(k)^\top\right)\hat{G}_i(\underline{\rho}^{(m)}(k); z_{k+1}).
\end{aligned}$$

D.1 Difference between GF and GD

For notational simplicity, we denote $\theta_i(t) = (a_i(t), u_i(t))$ for $i \in [m]$ and $t \geq 0$, and

$$\rho^{(m)}(t) = \frac{1}{m}\sum_{i=1}^m \delta_{\theta_i(t)} = \frac{1}{m}\sum_{i=1}^m \delta_{(a_i(t), u_i(t))}.$$

Similarly, $\tilde{\theta}_i(k) = (\tilde{a}_i(k), \tilde{u}_i(k))$, and

$$\tilde{\rho}^{(m)}(k) = \frac{1}{m} \sum_{i=1}^m \delta_{\tilde{\theta}_i(k)} = \frac{1}{m} \sum_{i=1}^m \delta_{(\tilde{a}_i(k), \tilde{u}_i(k))}.$$

Moreover, for $\theta = (a, u)$ and $\rho \in \mathcal{P}(\mathbb{R} \times \mathbb{R}^d)$, we define the following two functionals:

$$F(\theta, \rho) = V(\langle u_*, u \rangle; \|u_*\|_2, \|u\|_2) - \int_{\mathbb{R} \times \mathbb{R}^d} a' U(\langle u, u' \rangle; \|u\|_2, \|u'\|_2) \rho(da', du'),$$

$$G(\theta, \rho) = a \left(I_d - uu^\top \right) \left(\nabla_u V(\langle u_*, u \rangle; \|u_*\|_2, \|u\|_2) - \int_{\mathbb{R} \times \mathbb{R}^d} a' \nabla_u U(\langle u, u' \rangle; \|u\|_2, \|u'\|_2) \rho(da', du') \right),$$

and $H_\varepsilon(\theta, \rho) = (\varepsilon^{-1} F(\theta, \rho), G(\theta, \rho))$. Then, Eq.s (5) and (6) and Eq. (161) can be rewritten as

$$\frac{d}{dt} \theta_i(t) = H_\varepsilon(\theta_i(t), \rho^{(m)}(t)), \quad \tilde{\theta}_i(k+1) - \tilde{\theta}_i(k) = \eta H_\varepsilon(\tilde{\theta}_i(k), \tilde{\rho}^{(m)}(k)),$$

respectively. The lemma below will be used several times in the proof.

Lemma 5. Denoting $\rho^{(m)} = (1/m) \sum_{i=1}^m \delta_{\theta_i}$ and $\rho'^{(m)} = (1/m) \sum_{i=1}^m \delta_{\theta'_i}$. If $\|u_i\|_2 \leq C$ and $\|u'_i\|_2 \leq C$ for all $i \in [m]$ (C is any fixed absolute constant, for example, here we can take $C = 2$), then we have

$$\left| F(\theta_i, \rho^{(m)}) - F(\theta'_i, \rho'^{(m)}) \right| \leq M \left(\left(1 + \sup_{j \in [m]} |a_j| \right) \cdot \sup_{j \in [m]} \|u_j - u'_j\|_2 + \sup_{j \in [m]} |a_j - a'_j| \right), \quad (163)$$

$$\left\| G(\theta_i, \rho^{(m)}) - G(\theta'_i, \rho'^{(m)}) \right\|_2 \leq M \cdot \left(1 + \sup_{j \in [m]} |a_j| \right)^2 \cdot \sup_{j \in [m]} \|u_j - u'_j\|_2 \quad (164)$$

$$+ M \cdot \sup_{j \in [m]} |a_j - a'_j| \cdot \left(1 + \sup_{j \in [m]} |a_j| + \sup_{j \in [m]} |a_j - a'_j| \right), \quad (165)$$

where the constant M only depends on the M_i 's. As a consequence, we obtain that

$$\begin{aligned} & \left\| H_\varepsilon(\theta_i, \rho^{(m)}) - H_\varepsilon(\theta'_i, \rho'^{(m)}) \right\|_2 \leq \varepsilon^{-1} \left| F(\theta_i, \rho^{(m)}) - F(\theta'_i, \rho'^{(m)}) \right| + \left\| G(\theta_i, \rho^{(m)}) - G(\theta'_i, \rho'^{(m)}) \right\|_2 \\ & \leq (\varepsilon^{-1} + 1) M \cdot \left(\left(1 + \sup_{j \in [m]} |a_j| \right)^2 \cdot \sup_{j \in [m]} \|u_j - u'_j\|_2 + \sup_{j \in [m]} |a_j - a'_j| \cdot \left(1 + \sup_{j \in [m]} |a_j| + \sup_{j \in [m]} |a_j - a'_j| \right) \right) \\ & \leq (\varepsilon^{-1} + 1) M \cdot \left(\left(1 + \sup_{j \in [m]} |a_j| \right)^2 + \sup_{j \in [m]} \|\theta_j - \theta'_j\|_2 \right) \cdot \sup_{j \in [m]} \|\theta_j - \theta'_j\|_2. \end{aligned}$$

Proof. First, by triangle inequality, we have

$$\begin{aligned} & \left| F(\theta_i, \rho^{(m)}) - F(\theta'_i, \rho'^{(m)}) \right| \leq \left| V(\langle u_*, u_i \rangle; \|u_*\|_2, \|u_i\|_2) - V(\langle u_*, u'_i \rangle; \|u_*\|_2, \|u'_i\|_2) \right| \\ & + \frac{1}{m} \sum_{j=1}^m \left| a_j U(\langle u_i, u_j \rangle; \|u_i\|_2, \|u_j\|_2) - a'_j U(\langle u'_i, u'_j \rangle; \|u'_i\|_2, \|u'_j\|_2) \right| \\ & \leq \|\nabla V\|_\infty \|u_i - u'_i\|_2 + \frac{\|U\|_\infty}{m} \sum_{j=1}^m |a_j - a'_j| + \frac{\|\nabla U\|_\infty}{m} \sum_{j=1}^m |a_j| \cdot (\|u_i - u'_i\|_2 + \|u_j - u'_j\|_2) \\ & \leq M \left(\left(1 + \sup_{j \in [m]} |a_j| \right) \|u_i - u'_i\|_2 + \sup_{j \in [m]} |a_j - a'_j| + \sup_{j \in [m]} |a_j| \cdot \sup_{j \in [m]} \|u_j - u'_j\|_2 \right) \\ & \leq M \left(\left(1 + \sup_{j \in [m]} |a_j| \right) \cdot \sup_{j \in [m]} \|u_j - u'_j\|_2 + \sup_{j \in [m]} |a_j - a'_j| \right). \end{aligned}$$

Second, using again triangle inequality, we deduce that

$$\begin{aligned}
& \left\| G(\theta_i, \rho^{(m)}) - G(\theta'_i, \rho'^{(m)}) \right\|_2 \\
& \stackrel{(i)}{\leq} 2C |a_i| \|u_i - u'_i\|_2 \left(\|\nabla V\|_\infty + \frac{\|\nabla U\|_\infty}{m} \sum_{j=1}^m |a_j| \right) + C |a_i| \cdot \frac{\|\nabla U\|_\infty}{m} \sum_{j=1}^m |a_j| \cdot \|u_j - u'_j\|_2 \\
& \quad + C |a_i| \cdot \left(\|\nabla^2 V\|_\infty \|u_i - u'_i\|_2 + \frac{\|\nabla^2 U\|_\infty}{m} \sum_{j=1}^m |a_j| (\|u_i - u'_i\|_2 + \|u_j - u'_j\|_2) \right) \\
& \quad + C |a_i - a'_i| \left(\|\nabla V\|_\infty + \frac{\|\nabla U\|_\infty}{m} \sum_{j=1}^m |a_j| \right) + C (|a_i| + |a'_i - a_i|) \cdot \frac{\|\nabla U\|_\infty}{m} \sum_{j=1}^m |a_j - a'_j| \\
& \leq 5M \cdot \left(1 + \sup_{j \in [m]} |a_j| \right)^2 \cdot \sup_{j \in [m]} \|u_j - u'_j\|_2 + M \cdot \sup_{j \in [m]} |a_j - a'_j| \cdot \left(1 + 2 \sup_{j \in [m]} |a_j| + \sup_{j \in [m]} |a_j - a'_j| \right) \\
& \leq M \cdot \left(1 + \sup_{j \in [m]} |a_j| \right)^2 \cdot \sup_{j \in [m]} \|u_j - u'_j\|_2 + M \cdot \sup_{j \in [m]} |a_j - a'_j| \cdot \left(1 + \sup_{j \in [m]} |a_j| + \sup_{j \in [m]} |a_j - a'_j| \right),
\end{aligned}$$

where (i) follows from the inequality $\|u_i u_i^\top - (u'_i)(u'_i)^\top\|_{\text{op}} \leq 2C \|u_i - u'_i\|_2$, which is a result of the following direct calculation:

$$\|u_i u_i^\top - (u'_i)(u'_i)^\top\|_{\text{op}} = \sup_{\|x\|_2=1} |\langle x, u_i \rangle^2 - \langle x, u'_i \rangle^2| \leq 2C \sup_{\|x\|_2=1} |\langle x, u_i \rangle - \langle x, u'_i \rangle| = 2C \|u_i - u'_i\|_2.$$

This completes the proof of Lemma 5, since the “as a consequence” part follows naturally from the upper bounds obtained earlier. \square

Lemma 6. *Following the notation and assumption of Lemma 5, we have*

$$\left| \mathcal{R}(\rho^{(m)}) - \mathcal{R}(\rho'^{(m)}) \right| \leq M \cdot \left(\left(1 + \sup_{j \in [m]} |a_j| \right)^2 + \sup_{j \in [m]} \|\theta_j - \theta'_j\|_2 \right) \cdot \sup_{j \in [m]} \|\theta_j - \theta'_j\|_2.$$

Proof. By definition of the risk function and triangle inequality, we deduce that

$$\begin{aligned}
\left| \mathcal{R}(\rho^{(m)}) - \mathcal{R}(\rho'^{(m)}) \right| & \leq \frac{1}{m} \sum_{i=1}^m |a_i V(\langle u_*, u_i \rangle; \|u_*\|_2, \|u_i\|_2) - a'_i V(\langle u_*, u'_i \rangle; \|u_*\|_2, \|u'_i\|_2)| \\
& \quad + \frac{1}{m^2} \sum_{i,j=1}^m |a_i a_j U(\langle u_i, u_j \rangle; \|u_i\|_2, \|u_j\|_2) - a'_i a'_j U(\langle u'_i, u'_j \rangle; \|u'_i\|_2, \|u'_j\|_2)| \\
& \leq \frac{\|\nabla V\|_\infty}{m} \sum_{i=1}^m |a_i| \|u_i - u'_i\|_2 + \frac{\|V\|_\infty}{m} \sum_{i=1}^m |a_i - a'_i| \\
& \quad + \frac{\|U\|_\infty}{m^2} \sum_{i,j=1}^m (|a_i - a'_i| |a'_j| + |a_i| |a_j - a'_j|) \\
& \quad + \frac{\|\nabla U\|_\infty}{m^2} \sum_{i,j=1}^m |a_i| |a_j| \cdot (\|u_i - u'_i\|_2 + \|u_j - u'_j\|_2) \\
& \leq M \cdot \left(1 + \sup_{j \in [m]} |a_j| \right)^2 \cdot \sup_{j \in [m]} \|u_j - u'_j\|_2 \\
& \quad + M \cdot \sup_{j \in [m]} |a_j - a'_j| \cdot \left(1 + \sup_{j \in [m]} |a_j| + \sup_{j \in [m]} |a_j - a'_j| \right)
\end{aligned}$$

$$\leq M \cdot \left(\left(1 + \sup_{j \in [m]} |a_j| \right)^2 + \sup_{j \in [m]} \|\theta_j - \theta'_j\|_2 \right) \cdot \sup_{j \in [m]} \|\theta_j - \theta'_j\|_2.$$

This concludes the proof. \square

First, let us define the error function

$$\Delta(t) = \sup_{k \in [0, t/\eta] \cap \mathbb{N}} \max_{i \in [m]} \|\tilde{\theta}_i(k) - \theta_i(k\eta)\|_2,$$

and the stopping time $T_\Delta = \inf\{t \geq 0 : \Delta(t) \geq 1\}$. For $k \in \mathbb{N}$ and $t = k\eta \leq T_\Delta$, we have the following estimate:

$$\begin{aligned} \|\theta_i(t) - \tilde{\theta}_i(k)\|_2 &\leq \int_0^t \|H_\varepsilon(\theta_i(s), \rho^{(m)}(s)) - H_\varepsilon(\tilde{\theta}_i(\lfloor s/\eta \rfloor), \tilde{\rho}^{(m)}(\lfloor s/\eta \rfloor))\|_2 ds \\ &\leq \int_0^t \|H_\varepsilon(\theta_i(s), \rho^{(m)}(s)) - H_\varepsilon(\theta_i(\eta \lfloor s/\eta \rfloor), \rho^{(m)}(\eta \lfloor s/\eta \rfloor))\|_2 ds \\ &\quad + \int_0^t \|H_\varepsilon(\theta_i(\eta \lfloor s/\eta \rfloor), \rho^{(m)}(\eta \lfloor s/\eta \rfloor)) - H_\varepsilon(\tilde{\theta}_i(\lfloor s/\eta \rfloor), \tilde{\rho}^{(m)}(\lfloor s/\eta \rfloor))\|_2 ds. \end{aligned}$$

For any $s \in [0, t]$, by Lemma 4 and 5 we have (denote $[s] = \eta \lfloor s/\eta \rfloor$, and notice that we can take $C = 2$ since $t \leq T_\Delta$)

$$\begin{aligned} &\|H_\varepsilon(\theta_i(s), \rho^{(m)}(s)) - H_\varepsilon(\theta_i([s]), \rho^{(m)}([s]))\|_2 \\ &\leq (\varepsilon^{-1} + 1)M(1 + t/\varepsilon)^4(s - [s]) \leq (\varepsilon^{-1} + 1)M(1 + t/\varepsilon)^4\eta. \end{aligned}$$

Using again Lemma 4 and 5, we obtain that

$$\begin{aligned} &\|H_\varepsilon(\theta_i(\eta \lfloor s/\eta \rfloor), \rho^{(m)}(\eta \lfloor s/\eta \rfloor)) - H_\varepsilon(\tilde{\theta}_i(\lfloor s/\eta \rfloor), \tilde{\rho}^{(m)}(\lfloor s/\eta \rfloor))\|_2 \\ &\leq (\varepsilon^{-1} + 1)M(1 + \varepsilon^{-1}s)^2\Delta(s) + (\varepsilon^{-1} + 1)M\Delta(s)^2, \end{aligned}$$

thus leading to

$$\begin{aligned} \Delta(t) &\leq (\varepsilon^{-1} + 1)Mt(1 + t/\varepsilon)^4\eta + (\varepsilon^{-1} + 1) \int_0^t (M(1 + \varepsilon^{-1}s)^2\Delta(s) + M\Delta(s)^2) ds \\ &\leq (\varepsilon^{-1} + 1)Mt(1 + t/\varepsilon)^4\eta + (\varepsilon^{-1} + 1) \int_0^t M(1 + \varepsilon^{-1}s)^2 \cdot \max(\Delta(s), \Delta(s)^2) ds. \end{aligned}$$

For $s \leq t \leq T_\Delta$, we have $\Delta(s)^2 \leq \Delta(s)$. Hence,

$$\Delta(t) \leq (\varepsilon^{-1} + 1)Mt(1 + t/\varepsilon)^4\eta + (\varepsilon^{-1} + 1) \int_0^t M(1 + \varepsilon^{-1}s)^2\Delta(s) ds.$$

Applying Grönwall's inequality yields

$$\Delta(t) \leq (\varepsilon^{-1} + 1)Mt(1 + t/\varepsilon)^4\eta \cdot \exp\left((\varepsilon^{-1} + 1)Mt(1 + t/\varepsilon)^2\right) \leq M \exp((\varepsilon^{-1} + 1)Mt(1 + t/\varepsilon)^2)\eta.$$

Therefore, for all $T \geq 0$ and $\eta \leq 1/(M \exp((\varepsilon^{-1} + 1)MT(1 + \varepsilon^{-1}T)^2))$, we have

$$\sup_{k \in [0, t/\eta] \cap \mathbb{N}} \max_{i \in [m]} \|\tilde{\theta}_i(k) - \theta_i(k\eta)\|_2 \leq M \exp((\varepsilon^{-1} + 1)Mt(1 + t/\varepsilon)^2)\eta \leq 1, \quad \forall t \leq \min(T, T_\Delta).$$

This proves $T \leq T_\Delta$, and consequently

$$\sup_{k \in [0, t/\eta] \cap \mathbb{N}} \max_{i \in [m]} \|\tilde{\theta}_i(k) - \theta_i(k\eta)\|_2 \leq M \exp((\varepsilon^{-1} + 1)Mt(1 + t/\varepsilon)^2)\eta, \quad \forall t \in [0, T],$$

which immediately implies that

$$\sup_{k \in [0, t/\eta] \cap \mathbb{N}} \max_{i \in [m]} |\tilde{a}_i(k)| \leq M(1 + t/\varepsilon) + 1 \leq M(1 + t/\varepsilon).$$

Finally, with the aid of Lemma 6, we get the following upper bound on the difference between the risk of gradient flow and gradient descent:

$$\begin{aligned} \sup_{k \in [0, t/\eta] \cap \mathbb{N}} \left| \mathcal{R}(\rho^{(m)}(k\eta)) - \mathcal{R}(\tilde{\rho}^{(m)}(k)) \right| &\leq M(M^2(1 + t/\varepsilon)^2 + 1)M \exp((\varepsilon^{-1} + 1)Mt(1 + t/\varepsilon)^2)\eta \\ &\leq M \exp((\varepsilon^{-1} + 1)Mt(1 + t/\varepsilon)^2)\eta. \end{aligned}$$

To summarize, we have the following:

Theorem 4 (Difference between GF and GD). *There exists a constant M that only depends on the M_i 's, such that for any $T \geq 0$ and*

$$\eta \leq \frac{1}{M \exp((\varepsilon^{-1} + 1)MT(1 + \varepsilon^{-1}T)^2)},$$

the following holds for all $t \in [0, T]$:

$$\sup_{k \in [0, t/\eta] \cap \mathbb{N}} \max_{i \in [m]} |\tilde{a}_i(k)| \leq M(1 + t/\varepsilon), \quad (166)$$

$$\sup_{k \in [0, t/\eta] \cap \mathbb{N}} \max_{i \in [m]} \left\| \tilde{\theta}_i(k) - \theta_i(k\eta) \right\|_2 \leq M \exp((\varepsilon^{-1} + 1)Mt(1 + \varepsilon^{-1}T)^2)\eta, \quad (167)$$

$$\sup_{k \in [0, t/\eta] \cap \mathbb{N}} \left| \mathcal{R}(\rho^{(m)}(k\eta)) - \mathcal{R}(\tilde{\rho}^{(m)}(k)) \right| \leq M \exp((\varepsilon^{-1} + 1)Mt(1 + \varepsilon^{-1}T)^2)\eta. \quad (168)$$

D.2 Difference between GD and SGD

The proof for this section is almost identical to Appendix C.5 in [33]. The only difference is that, here we need to verify that $(I_d - uu^\top)\sigma'(\langle u, x \rangle)x$ is an M_3 -sub-Gaussian random vector. This follows from the identity $(I_d - uu^\top)x = x - \langle u, x \rangle u$ and Assumption A3. We thus obtain the following interpolation bound between GD and SGD:

Theorem 5 (Difference between GD and SGD). *There exists a constant M that only depends on the M_i 's, such that for any $T, z \geq 0$ and*

$$\eta \leq \frac{1}{(d + \log m + z^2)M \exp((\varepsilon^{-1} + 1)MT(1 + \varepsilon^{-1}T)^2)},$$

the following happens with probability at least $1 - \exp(-z^2)$: For all $t \in [0, T]$, we have

$$\sup_{k \in [0, t/\eta] \cap \mathbb{N}} \max_{i \in [m]} |\underline{a}_i(k)| \leq M(1 + t/\varepsilon), \quad (169)$$

$$\sup_{k \in [0, t/\eta] \cap \mathbb{N}} \max_{i \in [m]} \left\| \tilde{\theta}_i(k) - \underline{\theta}_i(k) \right\|_2 \leq M \exp((\varepsilon^{-1} + 1)Mt(1 + \varepsilon^{-1}T)^2)\sqrt{\eta} \left(\sqrt{d + \log m} + z \right), \quad (170)$$

$$\sup_{k \in [0, t/\eta] \cap \mathbb{N}} \left| \mathcal{R}(\underline{\rho}^{(m)}(k)) - \mathcal{R}(\tilde{\rho}^{(m)}(k)) \right| \leq M \exp((\varepsilon^{-1} + 1)Mt(1 + \varepsilon^{-1}T)^2)\sqrt{\eta} \left(\sqrt{d + \log m} + z \right). \quad (171)$$

D.3 Difference between SGD and projected SGD

The aim of this section is to prove a coupling bound between the trajectory of SGD and that of projected SGD, thus finally leading to an upper bound on the difference between the risk of projected gradient flow and projected SGD. To begin with, let us fix $T, z \geq 0$ and choose

$$\eta \leq \frac{1}{(d + \log m + z^2)M \exp((\varepsilon^{-1} + 1)MT(1 + \varepsilon^{-1}T)^2)}$$

as in Theorem 2, where M is a large enough constant (to be determined later). Define

$$T_\theta = \inf \left\{ t \geq 0 : \max_{k \in [0, t/\eta] \cap \mathbb{N}} \max_{i \in [m]} |\bar{a}_i(k)| \geq 2M(1 + t/\varepsilon), \text{ or } \max_{k \in [0, t/\eta] \cap \mathbb{N}} \max_{i \in [m]} \|\bar{u}_i(k)\|_2 \geq 2 \right\},$$

then for $k \leq \min(T, T_\theta)/\eta$ and $i \in [m]$, we have (note that here $t = k\eta$)

$$\begin{aligned} \left\| \hat{G}_i(\bar{\rho}^{(m)}(k); z_{k+1}) \right\|_2 &\leq M |\bar{a}_i(k)| \left(1 + \max_{i \in [m]} |\bar{a}_i(k)| \right) \left\| \sigma'(\langle \bar{u}_i(k), x_{k+1} \rangle) x_{k+1} \right\|_2 \\ &\leq M(1 + t/\varepsilon)^2 \left\| \sigma'(\langle \bar{u}_i(k), x_{k+1} \rangle) x_{k+1} \right\|_2. \end{aligned}$$

Denoting $\mathcal{F}_k = \sigma(\bar{\theta}(0), z_1, \dots, z_k)$, we know from Assumption A3 that, conditioning on \mathcal{F}_k , $\sigma'(\langle \bar{u}_i(k), x_{k+1} \rangle) x_{k+1}$ is an M_3 -sub-Gaussian random vector. By well-known results on Euclidean norm of sub-Gaussian random vectors (see, e.g., [27]), we know that there exists a constant M satisfying

$$\mathbb{P} \left(\left\| \sigma'(\langle \bar{u}_i(k), x_{k+1} \rangle) x_{k+1} \right\|_2 \geq M \left(\sqrt{d} + \sqrt{\log(1/\delta)} \right) \right) \leq \delta.$$

Choosing $\delta = \eta \exp(-z^2)/(mT)$ and applying a union bound gives

$$\mathbb{P} \left(\max_{k \in [0, \min(T, T_\theta)/\eta] \cap \mathbb{N}} \max_{i \in [m]} \left\| \sigma'(\langle \bar{u}_i(k), x_{k+1} \rangle) x_{k+1} \right\|_2 \leq M \left(\sqrt{d + \log m} + z + T^2 \right) \right) \geq 1 - \exp(-z^2).$$

Therefore, with probability at least $1 - \exp(-z^2)$, for all $k \leq \min(T, T_\theta)/\eta$ and $i \in [m]$, we have

$$\left\| \hat{G}_i(\bar{\rho}^{(m)}(k); z_{k+1}) \right\|_2 \leq M(1 + t/\varepsilon)^2 \left(\sqrt{d + \log m} + z + T^2 \right).$$

The above bound also holds for the trajectory of SGD, namely after replacing $\bar{\rho}^{(m)}(k)$ with $\underline{\rho}^{(m)}(k)$. Now, let us define the approximation error $\Delta_i(k) = \underline{u}_i(k) - \bar{u}_i(k)$ for $i \in [m]$ and $k \in \mathbb{N}$, then we get the following decomposition:

$$\Delta_i(l) = \sum_{k=0}^{l-1} (\Delta_i(k+1) - \Delta_i(k)) = \sum_{k=0}^{l-1} (\mathbb{E}[\Delta_i(k+1) - \Delta_i(k) | \mathcal{F}_k] + Z_i(k+1)),$$

where $Z_i(k+1) = \Delta_i(k+1) - \Delta_i(k) - \mathbb{E}[\Delta_i(k+1) - \Delta_i(k) | \mathcal{F}_k]$ has zero mean. With our choice of η , one can verify that as long as $\max(d, m, z) \rightarrow \infty$, Lemma 7 is applicable to

$$u_1 = \underline{u}_i(k), \quad g_1 = \hat{G}_i(\underline{\rho}^{(m)}(k); z_{k+1}), \quad u_2 = \bar{u}_i(k), \quad g_2 = \hat{G}_i(\bar{\rho}^{(m)}(k); z_{k+1}).$$

Hence, we deduce from the definition of $\Delta_i(k)$ that

$$\begin{aligned} \Delta_i(k+1) - \Delta_i(k) &= (\underline{u}_i(k+1) - \underline{u}_i(k)) - (\bar{u}_i(k+1) - \bar{u}_i(k)) = (v_1 - u_1) - (v_2 - u_2) \\ &= \eta \left((I_d - u_1 u_1^\top) g_1 - (I_d - u_2 u_2^\top) g_2 \right) + O\left(\eta^2 \|g_2\|_2^2\right), \end{aligned}$$

thus leading to the following estimate:

$$\left\| \mathbb{E}[\Delta_i(k+1) - \Delta_i(k) | \mathcal{F}_k] \right\|_2 \leq \eta \left\| \mathbb{E} \left[(I_d - u_2 u_2^\top) (g_1 - g_2) | \mathcal{F}_k \right] \right\|_2 + \eta \left\| \mathbb{E} \left[(u_2 u_2^\top - u_1 u_1^\top) g_1 | \mathcal{F}_k \right] \right\|_2$$

$$\begin{aligned}
& + C\eta^2 \mathbb{E} \left[\|g_2\|_2^2 \middle| \mathcal{F}_k \right] \\
& \stackrel{(i)}{\leq} \eta \mathbb{E} \left[\|(g_1 - g_2)|\mathcal{F}_k\|_2 + C\eta \|u_1 - u_2\|_2 \|\mathbb{E}[g_1|\mathcal{F}_k]\|_2 + C\eta^2 \mathbb{E} \left[\|g_2\|_2^2 \middle| \mathcal{F}_k \right] \right] \\
& = \eta \mathbb{E} \left[\left\| \left(\widehat{G}_i(\underline{\rho}^{(m)}(k); z_{k+1}) - \widehat{G}_i(\bar{\rho}^{(m)}(k); z_{k+1}) \right) \middle| \mathcal{F}_k \right\|_2 \right. \\
& \quad + C\eta \|\bar{u}_i(k) - \underline{u}_i(k)\|_2 \cdot \left\| \mathbb{E} \left[\widehat{G}_i(\underline{\rho}^{(m)}(k); z_{k+1}) \middle| \mathcal{F}_k \right] \right\|_2 \\
& \quad \left. + C\eta^2 \mathbb{E} \left[\left\| \widehat{G}_i(\bar{\rho}^{(m)}(k); z_{k+1}) \right\|_2^2 \middle| \mathcal{F}_k \right] \right],
\end{aligned}$$

where (i) is due to the fact that $u_1, u_2 \in \sigma(\mathcal{F}_k)$, and $\|u_1 u_1^\top - u_2 u_2^\top\|_{\text{op}} \leq C \|u_1 - u_2\|_2$. According to the definition of \widehat{G}_i , we obtain that

$$\begin{aligned}
& \mathbb{E} \left[\left(\widehat{G}_i(\underline{\rho}^{(m)}(k); z_{k+1}) - \widehat{G}_i(\bar{\rho}^{(m)}(k); z_{k+1}) \right) \middle| \mathcal{F}_k \right] \\
& = \underline{a}_i(k) \left(\nabla_{\underline{u}_i(k)} V(\langle u_*, \underline{u}_i(k) \rangle; \|u_*\|_2, \|\underline{u}_i(k)\|_2) \right. \\
& \quad - \frac{1}{m} \sum_{j=1}^m \underline{a}_j(k) \nabla_{\underline{u}_i(k)} U(\langle \underline{u}_i(k), \underline{u}_j(k) \rangle; \|\underline{u}_i(k)\|_2, \|\underline{u}_j(k)\|_2) \Big) \\
& \quad - \bar{a}_i(k) \left(\nabla_{\bar{u}_i(k)} V(\langle u_*, \bar{u}_i(k) \rangle; \|u_*\|_2, \|\bar{u}_i(k)\|_2) \right. \\
& \quad \left. - \frac{1}{m} \sum_{j=1}^m \bar{a}_j(k) \nabla_{\bar{u}_i(k)} U(\langle \bar{u}_i(k), \bar{u}_j(k) \rangle; \|\bar{u}_i(k)\|_2, \|\bar{u}_j(k)\|_2) \right),
\end{aligned}$$

thus leading to (using the same argument as in the proof of Lemma 5)

$$\begin{aligned}
& \left\| \mathbb{E} \left[\left(\widehat{G}_i(\underline{\rho}^{(m)}(k); z_{k+1}) - \widehat{G}_i(\bar{\rho}^{(m)}(k); z_{k+1}) \right) \middle| \mathcal{F}_k \right] \right\|_2 \leq M \left(1 + \varepsilon^{-1} T \right)^2 \cdot \sup_{j \in [m]} \left\| \bar{u}_j(k) - \underline{u}_j(k) \right\|_2 \\
& + M \left(1 + \varepsilon^{-1} T + \sup_{j \in [m]} \left| \bar{a}_j(k) - \underline{a}_j(k) \right| \right) \cdot \sup_{j \in [m]} \left| \bar{a}_j(k) - \underline{a}_j(k) \right|,
\end{aligned}$$

and

$$\left\| \mathbb{E} \left[\widehat{G}_i(\underline{\rho}^{(m)}(k); z_{k+1}) \middle| \mathcal{F}_k \right] \right\|_2 \leq M(1 + \varepsilon^{-1} T)^2.$$

Moreover, by (conditional) sub-Gaussianity of the \widehat{G}_i 's, we know that

$$\mathbb{E} \left[\left\| \widehat{G}_i(\bar{\rho}^{(m)}(k); z_{k+1}) \right\|_2^2 \middle| \mathcal{F}_k \right] \leq M^2(1 + \varepsilon^{-1} T)^4 \mathbb{E} \left[\left\| \sigma'(\langle \bar{u}_i(k), x_{k+1} \rangle) x_{k+1} \right\|_2^2 \middle| \mathcal{F}_k \right] \leq M(1 + \varepsilon^{-1} T)^4 d.$$

Combining the above estimates, it then follows that

$$\begin{aligned}
\left\| \mathbb{E} [\Delta_i(k+1) - \Delta_i(k) \middle| \mathcal{F}_k] \right\|_2 & \leq \eta M \left(1 + \varepsilon^{-1} T \right)^2 \cdot \sup_{j \in [m]} \left\| \bar{u}_j(k) - \underline{u}_j(k) \right\|_2 + \eta^2 M(1 + \varepsilon^{-1} T)^4 d \\
& + \eta M \left(1 + \varepsilon^{-1} T + \sup_{j \in [m]} \left| \bar{a}_j(k) - \underline{a}_j(k) \right| \right) \cdot \sup_{j \in [m]} \left| \bar{a}_j(k) - \underline{a}_j(k) \right|.
\end{aligned}$$

Using the same proof technique as in Appendix C.5 of [33], we conclude that

$$\mathbb{P} \left(\max_{i \in [m]} \max_{l \in [0, \min(T, T_\theta)/\eta] \cap \mathbb{N}} \left\| \sum_{k=0}^{l-1} Z_i(k+1) \right\|_2 \geq M(1 + \varepsilon^{-1} T)^2 \left(\sqrt{d + \log m} + z + T^2 \right) \sqrt{T\eta} \right) \leq \exp(-z^2).$$

Similarly as in the proof of Theorem 4, we define

$$\Delta(t) = \max_{l \in [0, t/\eta] \cap \mathbb{N}} \max_{i \in [m]} \|\bar{\theta}_i(l) - \underline{\theta}_i(l)\|_2, \quad T_\Delta = \inf\{t \geq 0 : \Delta(t) \geq 1\}.$$

Then, for $l \leq \min(T, T_\theta, T_\Delta)/\eta$, we have

$$\begin{aligned} \sup_{i \in [m]} \|\bar{u}_i(l) - \underline{u}_i(l)\|_2 &= \sup_{i \in [m]} \|\Delta_i(l)\|_2 \leq \sup_{i \in [m]} \left\{ \sum_{k=0}^{l-1} \|\mathbb{E}[\Delta_i(k+1) - \Delta_i(k) | \mathcal{F}_k]\|_2 + \left\| \sum_{k=0}^{l-1} Z_i(k+1) \right\|_2 \right\} \\ &\leq \eta M(1 + \varepsilon^{-1}T)^2 \sum_{k=0}^{l-1} \Delta(k\eta) + l\eta^2 M(1 + \varepsilon^{-1}T)^4 d \\ &\quad + M(1 + \varepsilon^{-1}T)^2 \left(\sqrt{d + \log m} + z + T^2 \right) \sqrt{T\eta}. \end{aligned}$$

Proceeding with the same argument, it follows that

$$\sup_{i \in [m]} |\bar{a}_i(l) - \underline{a}_i(l)| \leq \varepsilon^{-1} \eta M(1 + \varepsilon^{-1}T)^2 \sum_{k=0}^{l-1} \Delta(k\eta) + \varepsilon^{-1} M(1 + \varepsilon^{-1}T)^2 \left(\sqrt{d + \log m} + z + T^2 \right) \sqrt{T\eta}.$$

Therefore, we finally conclude that

$$\begin{aligned} \Delta(l\eta) &\leq (\varepsilon^{-1} + 1) \eta M(1 + \varepsilon^{-1}T)^2 \sum_{k=0}^{l-1} \Delta(k\eta) + l\eta^2 M(1 + \varepsilon^{-1}T)^4 d \\ &\quad + (\varepsilon^{-1} + 1) M(1 + \varepsilon^{-1}T)^2 \left(\sqrt{d + \log m} + z + T^2 \right) \sqrt{T\eta} \\ &\leq (\varepsilon^{-1} + 1) \eta M(1 + \varepsilon^{-1}T)^2 \sum_{k=0}^{l-1} \Delta(k\eta) + (\varepsilon^{-1} + 1) M(1 + \varepsilon^{-1}T)^4 \left(\sqrt{d + \log m} + z + T^2 \right) \sqrt{T\eta}. \end{aligned}$$

Applying Grönwall's inequality (discrete version) yields that

$$\begin{aligned} \Delta(l\eta) &\leq (\varepsilon^{-1} + 1) M(1 + \varepsilon^{-1}T)^4 \left(\sqrt{d + \log m} + z + T^2 \right) \sqrt{T\eta} \\ &\quad \times \left(1 + (\varepsilon^{-1} + 1) l\eta M(1 + \varepsilon^{-1}T)^2 \exp\left((\varepsilon^{-1} + 1) l\eta M(1 + \varepsilon^{-1}T)^2\right) \right) \\ &\leq M \exp\left((\varepsilon^{-1} + 1) MT(1 + \varepsilon^{-1}T)^2\right) \left(\sqrt{d + \log m} + z + T^2 \right) \sqrt{\eta} \\ &\leq M \exp\left((\varepsilon^{-1} + 1) MT(1 + \varepsilon^{-1}T)^2\right) \left(\sqrt{d + \log m} + z \right) \sqrt{\eta}, \end{aligned}$$

as long as $\max(d, m, z) \rightarrow \infty$ with $T = O(1)$. Note that the above inequality holds for all $l \in [0, \min(T, T_\theta, T_\Delta)/\eta] \cap \mathbb{N}$ with probability at least $1 - \exp(-z^2)$, which further implies that $T_\theta, T_\Delta \geq T$, and consequently

$$\sup_{k \in [0, T/\eta] \cap \mathbb{N}} \max_{i \in [m]} |\bar{a}_i(k)| \leq 2M(1 + \varepsilon^{-1}T).$$

Applying again Lemma 6, we deduce that

$$\sup_{k \in [0, T/\eta] \cap \mathbb{N}} \left| \mathcal{R}(\underline{\rho}^{(m)}(k)) - \mathcal{R}(\bar{\rho}^{(m)}(k)) \right| \leq \left(\sqrt{d + \log m} + z \right) M \exp((\varepsilon^{-1} + 1) MT(1 + \varepsilon^{-1}T)^2) \sqrt{\eta}.$$

Combining the above estimates gives the following:

Theorem 6 (Difference between SGD and projected SGD). *There exists a constant M that only depends on the M_i 's, such that for any $T, z \geq 0$ and*

$$\eta \leq \frac{1}{(d + \log m + z^2) M \exp((\varepsilon^{-1} + 1) MT(1 + \varepsilon^{-1}T)^2)},$$

the following happens with probability at least $1 - \exp(-z^2)$: For all $t \in [0, T]$, we have

$$\sup_{k \in [0, t/\eta] \cap \mathbb{N}} \max_{i \in [m]} |\bar{a}_i(k)| \leq M(1 + t/\varepsilon), \quad (172)$$

$$\sup_{k \in [0, t/\eta] \cap \mathbb{N}} \max_{i \in [m]} \|\bar{\theta}_i(k) - \theta_i(k)\|_2 \leq M \exp((\varepsilon^{-1} + 1)Mt(1 + \varepsilon^{-1}T)^2) \sqrt{\eta} \left(\sqrt{d + \log m} + z \right), \quad (173)$$

$$\sup_{k \in [0, t/\eta] \cap \mathbb{N}} \left| \mathcal{R}(\underline{\rho}^{(m)}(k)) - \mathcal{R}(\bar{\rho}^{(m)}(k)) \right| \leq M \exp((\varepsilon^{-1} + 1)Mt(1 + \varepsilon^{-1}T)^2) \sqrt{\eta} \left(\sqrt{d + \log m} + z \right). \quad (174)$$

Theorem 2 then follows as a result of combining Theorem 4, Theorem 5, and Theorem 6.

Lemma 7. Let $v_1 = u_1 + \eta(I_d - u_1 u_1^\top)g_1$, $v_2 = \text{Proj}_{\mathbb{S}^{d-1}}(u_2 + \eta g_2)$, where $\|u_2\|_2 = 1$ and $\eta \|g_2\|_2 \leq 1/2$. Then we have

$$(v_1 - u_1) - (v_2 - u_2) = \eta \left((I_d - u_1 u_1^\top)g_1 - (I_d - u_2 u_2^\top)g_2 \right) + O\left(\eta^2 \|g_2\|_2^2\right).$$

Proof. Using Taylor expansion, we know that

$$\begin{aligned} v_2 &= \text{Proj}_{\mathbb{S}^{d-1}}(u_2 + \eta g_2) = (u_2 + \eta g_2) \left(1 + 2\eta \langle u_2, g_2 \rangle + \eta^2 \|g_2\|_2^2 \right)^{-1/2} \\ &= (u_2 + \eta g_2) \left(1 - \eta \langle u_2, g_2 \rangle + O(\eta^2 \|g_2\|_2^2) \right) \\ &= (1 - \eta \langle u_2, g_2 \rangle) u_2 + \eta g_2 + O(\eta^2 \|g_2\|_2^2) \\ &= u_2 + \eta(I_d - u_2 u_2^\top)g_2 + O(\eta^2 \|g_2\|_2^2), \end{aligned}$$

which implies

$$v_2 - u_2 = \eta(I_d - u_2 u_2^\top)g_2 + O(\eta^2 \|g_2\|_2^2).$$

The proof is completed by noting that

$$v_1 - u_1 = \eta(I_d - u_1 u_1^\top)g_1.$$

□

D.4 Proof of Theorem 3

By our assumption, we know that the canonical learning order holds up to level L , and that

$$\frac{1}{2} \sum_{k \geq L+1} \varphi_k^2 \leq \frac{\delta}{4}.$$

Then, according to Definition 1, there exists $\varepsilon_* = \varepsilon_*(\delta)$, $T_0 = T_0(\delta)$ such that for all $\varepsilon \leq \varepsilon_*$ and $T = T(\varepsilon, \delta) = T_0(\delta)\varepsilon^{1/2L}$, one has

$$\mathcal{R}_\infty(T, \varepsilon) = \lim_{m \rightarrow \infty} \lim_{d \rightarrow \infty} \mathcal{R}(a(T), u(T)) \leq \frac{\delta}{3}.$$

Moreover, from Section 4 we know that with probability at least $1 - e^{-C'm}$ over the i.i.d. initialization,

$$\sup_{t \in [0, T]} |\mathcal{R}(a(t), u(t)) - \mathcal{R}_\infty(t, \varepsilon)| \leq \left(\frac{1}{\sqrt{d}} + \frac{1}{\sqrt{m}} \right) CM' \exp(M'T(1 + T)^2/\varepsilon^2), \quad (175)$$

where M' only depends on (σ, φ, P_A) . Now we choose $\varepsilon \leq \varepsilon_*$ and $T = T(\varepsilon, \delta) = T_0(\delta)\varepsilon^{1/2L}$. It then follows that

$$\mathcal{R}(a(T), u(T)) \leq \frac{\delta}{3} + \left(\frac{1}{\sqrt{d}} + \frac{1}{\sqrt{m}} \right) CM' \exp(M'T^3/\varepsilon^2). \quad (176)$$

According to Theorem 2, we know that with probability at least $1 - \exp(-z)$,

$$|\mathcal{R}(\bar{a}(n), \bar{u}(n)) - \mathcal{R}(a(T), u(T))| \leq \sqrt{\eta(d + \log m + z)} M' \exp(M'T^3/\varepsilon^3) \quad (177)$$

with $n = T/\eta = T(\varepsilon, \delta)/\eta$. We now take

$$M = M(\varepsilon, \delta) = \max \left\{ \frac{9M'^2 \exp(2M'T(\varepsilon, \delta)^3/\varepsilon^3)}{\delta^2}, \frac{36C^2 M'^2 \exp(2M'T(\varepsilon, \delta)^3/\varepsilon^2)}{\delta^2} \right\}.$$

Then, by our choice of m and d , we know that $\mathcal{R}(a(T), u(T)) \leq 2\delta/3$. Further, taking

$$\eta = \frac{1}{M(d + \log m + z)}, \quad n = MT(d + \log m + z), \quad (178)$$

we obtain that

$$\mathcal{R}(\bar{a}(n), \bar{u}(n)) \leq \mathcal{R}(a(T), u(T)) + \frac{\delta}{3} \leq \delta. \quad (179)$$

The above happens with probability $1 - \exp(-C'm) - \exp(-z)$. Hence, our conclusion follows naturally from the assumption $m \geq z$.

E Counterexamples to the canonical learning order

E.1 Case 1: $\sigma_k = 0$ for some $k \in \mathbb{N}$

For any fixed $(a, u) = (a_i, u_i)_{1 \leq i \leq m}$, we have

$$\begin{aligned} \mathbb{E}[f(x; a, u) \text{He}_k(\langle u_*, x \rangle)] &= \frac{1}{m} \sum_{i=1}^m a_i \mathbb{E}[\sigma(\langle u_i, x \rangle) \text{He}_k(\langle u_*, x \rangle)] \\ &= \frac{1}{m} \sum_{i=1}^m a_i \sigma_k \langle u_i, u_* \rangle^k = 0. \end{aligned}$$

Moreover, the risk is always lower bounded by

$$\begin{aligned} \mathcal{R}(a, u) &= \frac{1}{2} \mathbb{E}[(\varphi(\langle u_*, x \rangle) - f(x; a, u))^2] \\ &= \frac{1}{2} \mathbb{E}[(\varphi_k \text{He}_k(\langle u_*, x \rangle) + (\varphi(\langle u_*, x \rangle) - \varphi_k \text{He}_k(\langle u_*, x \rangle)) - f(x; a, u))^2] \\ &\stackrel{(i)}{=} \frac{1}{2} \varphi_k^2 + \frac{1}{2} \mathbb{E}[(\varphi(\langle u_*, x \rangle) - \varphi_k \text{He}_k(\langle u_*, x \rangle) - f(x; a, u))^2] \geq \frac{1}{2} \varphi_k^2, \end{aligned}$$

where (i) follows from orthogonality between $\text{He}_k(\langle u_*, x \rangle)$ and $f(x; a, u)$.

E.2 Case 2: $\varphi_0 = \dots = \varphi_k = 0$ for some $k \geq 1$

We consider the reduced mean-field equations (24):

$$\begin{aligned} \varepsilon \partial_t a_i &= V(s_i) - \frac{1}{m} \sum_{j=1}^m a_j U(s_i s_j), \\ \partial_t s_i &= a_i (1 - s_i^2) \left(V'(s_i) - \frac{1}{m} \sum_{j=1}^m a_j U'(s_i s_j) s_j \right). \end{aligned}$$

Note that if $\varphi_0 = \varphi_1 = 0$, then $V'(s) = s \cdot v(s)$ for some continuous function v . Denoting $a = (a_1, \dots, a_m)$ and $s = (s_1, \dots, s_m)^\top$, the above equation regarding the evolution of the s_i 's can be written as

$$s'(t) = A(a(t), s(t))s(t),$$

where $A(a, s)$ is a matrix-valued function satisfying

$$A_{ij}(a, s) = a_i(1 - s_i^2) \left(v(s_i) \mathbf{1}_{i=j} - \frac{a_j}{m} U'(s_i s_j) \right), \quad \forall i, j \in [m].$$

Using the similar a priori estimate as in the proof of Lemma 1, we can show that

$$\sup_{t \in [0, T]} \|A(a(t), s(t))\|_{\text{op}} \leq C(T) < \infty$$

for any finite time T , which immediately implies that $s(t) \equiv 0$ for $t \in [0, T]$. Therefore, we won't be able to learn any component of φ with degree ≥ 1 .

E.3 Case 3: $\varphi_k = 0$ for some $k \geq 1$

We may assume $\sigma_k \neq 0$, and analyze the simplified ODE system (97), which reduces to

$$\begin{aligned} \partial_\tau \tilde{a}(\omega) &= -\sigma_k^2 \tilde{s}(\omega)^k \int \tilde{a}(\nu) \tilde{s}(\nu)^k d\rho(\nu) \\ \partial_\tau \tilde{s}(\omega) &= -k\sigma_k^2 \tilde{a}(\omega) \tilde{s}(\omega)^{k-1} \left(1 - \varepsilon^{2\beta_k} \tilde{s}(\omega)^2\right) \int \tilde{a}(\nu) \tilde{s}(\nu)^k d\rho(\nu). \end{aligned} \quad (180)$$

We thus obtain the following equations:

$$\begin{aligned} \partial_\tau \int \tilde{a}(\omega)^2 d\rho(\omega) &= -2\sigma_k^2 \left(\int \tilde{a}(\omega) \tilde{s}(\omega)^k d\rho(\omega) \right)^2 \leq 0, \\ \partial_\tau \int \tilde{s}(\omega)^2 d\rho(\omega) &= -2k\sigma_k^2 \int \tilde{a}(\omega) \tilde{s}(\omega)^k \left(1 - \varepsilon^{2\beta_k} \tilde{s}(\omega)^2\right) d\rho(\omega) \cdot \int \tilde{a}(\omega) \tilde{s}(\omega)^k d\rho(\omega) \\ &\leq 2k\sigma_k^2 \varepsilon^{2\beta_k} \int \tilde{a}(\omega) \tilde{s}(\omega)^{k+2} d\rho(\omega) \cdot \int \tilde{a}(\omega) \tilde{s}(\omega)^k d\rho(\omega), \end{aligned} \quad (181)$$

which means that for any $\tau \geq 0$,

$$\int \tilde{a}(\omega, \tau)^2 d\rho(\omega) \leq \int \tilde{a}(\omega, 0)^2 d\rho(\omega) = O(\varepsilon^{1/k(k+1)}) = o(1). \quad (182)$$

Therefore, most of the neurons cannot evolve to the magnitude of $\Omega(1)$ in the process of learning the k -th component, and therefore fails to provide an effective initialization for learning the next component φ_{k+1} .