

TRAINABILITY AND ACCURACY OF NEURAL NETWORKS: AN INTERACTING PARTICLE SYSTEM APPROACH

GRANT M. ROTSKOFF AND ERIC VANDEN-EIJNDEN

ABSTRACT. Neural networks, a central tool in machine learning, have demonstrated remarkable, high fidelity performance on image recognition and classification tasks. These successes evince an ability to accurately represent high dimensional functions, but rigorous results about the approximation error of neural networks after training are few. Here we establish conditions for global convergence of the standard optimization algorithm used in machine learning applications, stochastic gradient descent (SGD), and quantify the scaling of its error with the size of the network. This is done by reinterpreting SGD as the evolution of a particle system with interactions governed by a potential related to the objective or “loss” function used to train the network. We show that, when the number n of units is large, the empirical distribution of the particles descends on a convex landscape towards the global minimum at a rate independent of n , with a resulting approximation error that universally scales as $O(n^{-1})$. These properties are established in the form of a Law of Large Numbers and a Central Limit Theorem for the empirical distribution. Our analysis also quantifies the scale and nature of the noise introduced by SGD and provides guidelines for the step size and batch size to use when training a neural network. We illustrate our findings on examples in which we train neural networks to learn the energy function of the continuous 3-spin model on the sphere. The approximation error scales as our analysis predicts in as high a dimension as $d = 25$.

CONTENTS

1. Introduction	2
1.1. Problem set-up	3
1.2. Main results and organization	5
2. Functional formulation of the learning problem	6
2.1. Universal Approximation Theorem	6
2.2. Convexification at distributional level	8
3. Training by gradient descent on the exact loss	8
3.1. Empirical distribution and nonlinear Liouville equation	9
3.2. Law of Large Numbers (LLN)—mean field limit	10
3.3. Long time behavior—global convergence	11
3.4. Central Limit Theorem (CLT)	13
3.5. Scaling of the fluctuations at long and very long times	15
4. Training by online stochastic gradient descent	18
4.1. Limiting stochastic differential equation	18
4.2. Dean’s equation for particles with correlated noise	19
4.3. LLN for SGD	20
4.4. CLT for SGD	20
4.5. Fluctuations in SGD at long and very long times	22
5. Illustrative example: 3-spin model on the high-dimensional sphere	22
5.1. Learning with Gaussian kernels	23
5.2. Learning with single layer networks with sigmoid nonlinearity	24

We thank Joan Bruna and Weinan E for discussions about the approximation error of neural networks, and Sylvia Serfaty for her insights about interacting particle systems. We are also grateful to Andrea Montanari, Matthieu Wyart, and the anonymous referee for their helpful comments.

6. Concluding remarks	26
Appendix A. Training at finite (but small) temperature	27
A.1. Dean’s equation	28
A.2. Multiple-scale expansion	29
A.3. Law of Large Numbers at finite temperature	30
A.4. Global convergence on $O(n)$ timescales	30
A.5. Central Limit Theorem at finite temperature	32
References	34

1. INTRODUCTION

While both speech recognition and image classification remain active areas of research, extraordinary progress has been made on both problems—ones that appeared intractable only a decade ago [LBH15]. By harvesting the power of neural networks while simultaneously benefiting from advances in computational hardware, complex tasks such as automatic language translation are now routinely performed by computers with a high degree of reliability. The underlying explanation for these significant advances seems to be related to the expressive power of neural networks, and their ability to accurately represent high dimensional functions.

These successes open exciting possibilities in applied and computational mathematics that are only beginning to be explored [BP07, SDT⁺17, KLY18, BN17, EHJ17, BEJ17, ZHW⁺18]. Any numerical calculation that uses a given function begins with a finite-dimensional approximation of that function. Because standard approximations, e.g., Galerkin truncations or finite element decompositions, suffer from the curse of dimensionality, it is nearly impossible to scale such methods to large dimensions d . Fundamentally, these representations are linear combinations of basis functions. The issue arises because the dimensionality of the representation is equal to that of the truncation. Neural networks, on the other hand, are highly nonlinear in their adjusting parameters. As a result, the effective dimensionality of a neural network is much higher than its total number of parameters, which may explain the impressive function approximation capabilities observed in practice, even when d is large. Characterizing this observation with analysis is non-trivial though, precisely because the representation of a function by a neural network is nonlinear in its parameters. This renders many of the standard tools of numerical analysis useless, since they are in large part based on linear algebra.

The significant achievements of machine learning have inspired many efforts to provide theoretical justification to a vast and growing body of empirical knowledge. At the core of our understanding of the approximation properties of neural networks are the well-known “Universal Approximation Theorems” that specify the conditions under which a neural network can represent a target function with arbitrary accuracy [Cyb89, Bar93, PS91]. Despite the power of these results, they do not indicate how the network parameters should be optimized to achieve maximal accuracy in practice [BL04]. In particular, these theorems do not provide general guidance on how the error of the network scales with its size at the end of training. Several recent papers have focused on the analysis of the shape and properties of the objective or “loss” function landscape [SGBAL14, CHM⁺14, BSG⁺18]. These studies have mainly focused on the fine features of this landscape, trying to understand how non-convex it is and making analogies with glassy landscapes. Additionally, some analysis has been performed in cases where the number of parameters vastly exceeds the amount of training data, a setting that guarantees convexity and dramatically simplifies the landscape. Further studies have examined the dynamics of the parameters on the loss landscape to understand the properties of optimization procedures based on stochastic gradient descent.

In this paper, we adopt a different perspective which enables powerful tools for analysis. Similar to what was recently proposed in [MMN18, SS18, CB18], we view the parameters in the network as particles and the loss function as a potential that dictates the interaction between them. Correspondingly, training the network can be interpreted as the evolution of the particles in this interaction potential. Using the interchangeability of the n interacting particles / parameters in the neural representation, we focus on their empirical distribution and analyze its properties when n is large using standard limit theorems [KL13, Ser15, LS17, Ser17]. This viewpoint allows us to bypass many of the difficulties that arise with approaches that attempt to study the dynamics of the individual particles. In particular:

- (1) We derive an evolution equation for the empirical distribution of the particles, and show that it evolves by gradient descent in the 2-Wasserstein metric on a convex energy landscape. This observation allows us to assert that convergence towards equilibrium of the empirical distribution occurs on a time scale that is independent of n to leading order—similar results were obtained in [MMN18, SS18, CB18]. The results are obtained in the form of Law of Large Numbers (LLN) for the empirical distribution of the parameters. As a consequence, we rederive the Universal Approximation Theorem and establish that it can be realized dynamically.
- (2) We quantify the fluctuations of the empirical distribution at finite n above its limit. We show that these fluctuations are of order $O(n^{-1/2})$ and controlled at all $t < \infty$. In addition, we establish conditions under which these fluctuations heal and become $O(n^{-1})$ as $t \rightarrow \infty$. These results rely on a Central Limit Theorem (CLT) and indicate that the neural network approximation error is universal and scales as $O(n^{-1})$ as $n \rightarrow \infty$ in any d .

These results are established first in situations where gradient descent (GD) on the loss function is used to optimize or “train” the parameters in the network, and then shown to also apply in the context of stochastic gradient descent (SGD). In the latter case, our analysis sheds light on the nature of the noise introduced in SGD, and indicates how the time step and the batch size should be scaled to achieve the optimal error. We briefly elaborate on these statements below, first precisely formulating the problem.

1.1. Problem set-up. Given a function $f : \Omega \rightarrow \mathbb{R}$ defined on the closed manifold $\Omega \subseteq \mathbb{R}^d$, consider its approximation by a neural network of the form

$$(1) \quad f^{(n)}(\mathbf{x}) = \frac{1}{n} \sum_{i=1}^n c_i \hat{\varphi}(\mathbf{x}, \mathbf{z}_i)$$

where $n \in \mathbb{N}$, $(c_i, \mathbf{z}_i) \in D \equiv \mathbb{R} \times \hat{D}$ are parameters to be learned for $i = 1, \dots, n$, and $\varphi : \Omega \times D \rightarrow \mathbb{R}$ is some function—we assume throughout this paper that \hat{D} is a closed manifold in \mathbb{R}^N . The function $\hat{\varphi}$ is usually referred to as the ‘nonlinearity’ or ‘unit’ and n as the width of the network. To simplify notations, we use $\boldsymbol{\theta} = (c_i, \mathbf{z}) \in D$ and $\varphi(\mathbf{x}, \boldsymbol{\theta}) = c \hat{\varphi}(\mathbf{x}, \mathbf{z})$, in terms of which (1) reads

$$(2) \quad f^{(n)}(\mathbf{x}) = \frac{1}{n} \sum_{i=1}^n \varphi(\mathbf{x}, \boldsymbol{\theta}_i)$$

Many models used in machine learning can be cast in the form (1)-(2):

- **Radial basis function networks.** In this case $\hat{D} \equiv \Omega$ and $\hat{\varphi}(\mathbf{x}, \mathbf{z}) \equiv \phi(\mathbf{x} - \mathbf{z})$ where ϕ is some kernel, for example that of a radial function such as

$$\phi(\mathbf{x}) = \exp\left(-\frac{1}{2}\kappa|\mathbf{x}|^2\right)$$

where $\kappa > 0$ is a fixed constant.

- **Single hidden layer neural networks.** In this case, $\hat{D} \subset \mathbb{S}^d$ and $\hat{\varphi}(\mathbf{x}, \mathbf{z}) = \hat{\varphi}(\mathbf{x}, \mathbf{a}, b)$ with e.g. $\mathbf{a} \in \mathbb{S}^{d-1}$, $b \in [-1, 1]$, and

$$\varphi(\mathbf{x}, \mathbf{a}, b) = h(\mathbf{a} \cdot \mathbf{x} + b)$$

where $h : \mathbb{R} \rightarrow \mathbb{R}$ is e.g. a sigmoid function $h(z) = 1/(1 + e^{-z})$ or a rectified linear unit (ReLU) $h(z) = \max(z, 0)$.

- **Iterated neural networks.** These are structurally similar to single hidden layer neural networks. For example, to construct a two-layer network we take h as above and for $m \in \mathbb{N}$, $m \leq d$ define $\mathbf{h}^{(1)} : \mathbb{R}^m \rightarrow \mathbb{R}^m$ such that

$$h_j^{(1)}(\mathbf{v}) = h(v_j), \quad \mathbf{v} = (v_1, \dots, v_m) \in \mathbb{R}^m, \quad j = 1, \dots, m$$

then set

$$f^{(n)}(\mathbf{x}) = \frac{1}{n} \sum_{i=1}^n c_i h(\mathbf{a}_i^{(0)} \cdot \mathbf{h}^{(1)}(A_i^{(1)} \mathbf{x} + \mathbf{b}_i^{(1)}) + b_i^{(0)})$$

where $\mathbf{a}_i^{(0)} \in \mathbb{R}^m$, $b_i^{(0)} \in \mathbb{R}$, $A_i^{(1)} \in \mathbb{R}^{m \times d}$, $\mathbf{b}_i^{(1)} \in \mathbb{R}^m$, $i = 1, \dots, n$. Therefore here we have $\mathbf{z} = (\mathbf{a}^{(0)}, b^{(0)}, A^{(1)}, \mathbf{b}^{(1)}) \in \hat{D} \subset \mathbb{R}^{m+1+m \times d+m}$ (where with a slight abuse of notation we view the matrix $A^{(1)}$ has a vector in $\mathbb{R}^{m \times d}$). Three-layer networks, etc. can be constructed similarly. Note that our results apply to deep neural networks when their final layer grows large, with fixed depth.

To measure the discrepancy between the target function f and its neural network approximation $f^{(n)}$, we need to introduce a distance, or loss function, between f and $f^{(n)}$. A natural candidate often used in practice is

$$(3) \quad \mathcal{L}[f^{(n)}] = \frac{1}{2} \int_{\Omega} |f(\mathbf{x}) - f^{(n)}(\mathbf{x})|^2 \nu(d\mathbf{x}) = \frac{1}{2} \mathbb{E}_{\nu} |f - f^{(n)}|^2$$

where ν , the data distribution, is some positive measure on Ω such that $\nu(\Omega) < \infty$ (for example the Hausdorff measure on Ω , which we will denote by $d\mathbf{x}$). We can view $\mathcal{L}[f^{(n)}]$ as an objective function for $\{\boldsymbol{\theta}\}_{i=1}^n$:

$$(4) \quad \mathcal{L}[f^{(n)}] = C_f - \frac{1}{n} \sum_{i=1}^n F(\boldsymbol{\theta}_i) + \frac{1}{2n^2} \sum_{i,j=1}^n K(\boldsymbol{\theta}_i, \boldsymbol{\theta}_j)$$

where $C_f = \frac{1}{2} \mathbb{E}_{\nu} |f|^2$ and we defined

$$(5) \quad F(\boldsymbol{\theta}) = c \hat{F}(\mathbf{z}), \quad K(\boldsymbol{\theta}, \boldsymbol{\theta}') = cc' \hat{K}(\mathbf{z}, \mathbf{z}')$$

with

$$(6) \quad \hat{F}(\mathbf{z}) = \mathbb{E}_{\nu} [f \hat{\varphi}(\cdot, \mathbf{z})], \\ \hat{K}(\mathbf{z}, \mathbf{z}') = \mathbb{E}_{\nu} [\hat{\varphi}(\cdot, \mathbf{z}) \hat{\varphi}(\cdot, \mathbf{z}')] \equiv \hat{K}(\mathbf{z}', \mathbf{z}).$$

Trying to minimize (4) over $\{\boldsymbol{\theta}_i\}_{i=1}^n$ leads to difficulties, however, since this is potentially (and presumably) a non-convex optimization problem, which typically has local minimizers. In particular, if we perform training by making $\{\boldsymbol{\theta}_i\}_{i=1}^n$ evolve via gradient descent (GD) over the loss, i.e. if we use

$$(7) \quad \dot{\boldsymbol{\theta}}_i = \nabla F(\boldsymbol{\theta}_i) - \frac{1}{n} \sum_{j=1}^n \nabla K(\boldsymbol{\theta}_i, \boldsymbol{\theta}_j),$$

there is no guarantee *a priori* that these parameters will reach the global minimum of the loss or even a local minimum with a value for the loss that is close to that of the global minimum. As a result determining the value of (3) after training (and its scaling with n , say) is nontrivial. It is therefore natural to ask:

How accurate is the approximation (1)-(2) if we optimize $\{\boldsymbol{\theta}_i\}_{i=1}^n$ by applying the algorithms commonly used in machine learning?

This is the main question we investigate in the present paper.

1.2. **Main results and organization.** We will consider the evolution in time of the representation

$$(8) \quad f_t^{(n)} = \frac{1}{n} \sum_{i=1}^n \varphi(\cdot, \boldsymbol{\theta}_i(t))$$

and study the behavior of this function for large n and large t . To this end we use tools from interacting particle systems, and also build on known results about the nature of the loss function (4) in the limit as $n \rightarrow \infty$ —these results are recalled in Sec. 2.

In Sec. 3 we consider the situation where $\boldsymbol{\theta}_i(t)$ are the solution of the GD flow (7)—as explained below, this is somewhat of an idealized situation since we typically must work with the empirical loss rather than the exact one, but it is more easily amenable to analysis. By looking at the evolution of the empirical distribution of $\{\boldsymbol{\theta}_i(t)\}_{i=1}^n$ rather than that of the parameters themselves, we establish a Law of Large Numbers (LLN) for $f_t^{(n)}$ namely that $\lim_{n \rightarrow \infty} f_t^{(n)} = f_t$, where f_t evolves as

$$(9) \quad \partial_t f_t(\mathbf{x}) = - \int_{\Omega} M_t(\mathbf{x}, \mathbf{x}') (f_t(\mathbf{x}') - f(\mathbf{x}')) \nu(d\mathbf{x}')$$

where f is the target function and $M_t(\mathbf{x}, \mathbf{x}')$ a positive semi-definite kernel whose form is explicit—see Proposition 3.3. The evolution equation (9) can be interpreted as GD for f_t over the loss in some metric inherited from the 2-Wasserstein metric, and in Proposition 3.5 we show the flow converges to the target function, i.e.,

$$(10) \quad \lim_{t \rightarrow \infty} f_t = \lim_{t \rightarrow \infty} \lim_{n \rightarrow \infty} f_t^{(n)} = f.$$

We also establish that the limit in n and t commute. Regarding the scaling of the fluctuations above f_t when n is finite, in Proposition 3.7 we establish a Central Limit Theorem (CLT) that asserts that these fluctuations are of size $O(n^{-1/2})$, i.e. $n^{1/2}(f_t^{(n)} - f_t)$ has a limit in law as $n \rightarrow \infty$. In addition, in Proposition 3.8 we show that these fluctuations are controlled at all times, and in Proposition 3.9 that under certain conditions they heal as $t \rightarrow \infty$, in the sense that

$$(11) \quad f_{a_n}^{(n)} = f + O(n^{-1}) \quad \text{as } n \rightarrow \infty \quad \text{with } a_n / \log n \rightarrow \infty.$$

In Sec. 4 we analyze the typical situation in which it is not possible to calculate (3) or (6) exactly. Rather, we must approximate these expectations using a “training set”, i.e. a set of points $\{\mathbf{x}_p\}_{p=1}^P$ distributed according to ν on which f is known, so that instead of $\mathcal{L}[f^{(n)}]$ we must use

$$(12) \quad \mathcal{L}_P[f^{(n)}] = \frac{1}{P} \sum_{p=1}^P |f(\mathbf{x}_p) - f^{(n)}(\mathbf{x}_p)|^2$$

and instead of \hat{F} and \hat{K}

$$(13) \quad \hat{F}_P(\mathbf{z}) = \frac{1}{P} \sum_{p=1}^P f(\mathbf{x}_p) \hat{\varphi}(\mathbf{x}_p, \mathbf{z}), \quad \hat{K}_P(\mathbf{z}, \mathbf{z}') = \frac{1}{P} \sum_{p=1}^P \hat{\varphi}(\mathbf{x}_p, \mathbf{z}) \hat{\varphi}(\mathbf{x}_p, \mathbf{z}').$$

If in (7) we replace $F(\boldsymbol{\theta})$ and $K(\boldsymbol{\theta}, \boldsymbol{\theta}')$ by their empirical estimates over a subset of the training points $F_P(\boldsymbol{\theta}) = c\hat{F}_P(\mathbf{z})$ and $K_P(\boldsymbol{\theta}, \boldsymbol{\theta}') = cc'\hat{K}_P(\mathbf{z}, \mathbf{z}')$, we arrive at what is referred to as stochastic gradient descent (SGD)—the method of choice to train neural networks. We focus on situations in which we can redraw the training set as often as we need, namely, at every step during the learning process, an algorithm called online learning. In this case, in the limit as the optimization time step Δt used in SGD tends to zero, SGD becomes asymptotically equivalent to an SDE whose drift terms coincide with those of GD but with multiplicative noise terms added. In this context, we establish that (9) and (10) also hold if we choose the size P of the batch used in (13) at every SGD step such that as $P = O(n^{2\alpha})$ with $\alpha > 0$. Regarding the scaling of the fluctuations, if we set $\alpha \in (0, 1)$, we lose accuracy and (11) is replaced by

$$(14) \quad f_{a_n}^{(n)} = f + O(n^{-\alpha}) \quad \text{as } n \rightarrow \infty \quad \text{with } a_n / \log n \rightarrow \infty.$$

However if $\alpha \geq 1$, we get (11) back (meaning also that there is no advantage to take α bigger than 1). These results are stated in Propositions 4.3 and 4.4.

In Sec. 5 we illustrate these results, using a spherical p -spin model with $p = 3$ as test function to represent with a neural network. We show that the network accurately approximates this function in up to $d = 25$ dimensions, with a scaling of the error consistent with the results established in Secs. 3 and 4. These results are obtained using both a radial basis function network, and a single hidden layer network using sigmoid functions.

Concluding remarks are made in Sec. 6 and in an Appendix we establish a finite-temperature variant (Langevin dynamics) of (11) which applies when additive noise terms are added in the GD equations (7). This result reads

$$(15) \quad \lim_{T \rightarrow -\infty} f_t^{(n)} = f + n^{-1} \tilde{f}_t + o(n^{-1}) \quad \text{with} \quad \tilde{f}_t = \beta^{-1} \epsilon^* + \beta^{-1/2} \tilde{\epsilon}_t$$

where T is the time at which we initiate the training. Here $\beta > 0$ is a parameter playing the role of inverse temperature, $\epsilon^* : \Omega \rightarrow \mathbb{R}$ is some given (non-random) function and $\tilde{\epsilon}_t : \Omega \rightarrow \mathbb{R}$ is a Gaussian process with mean zero and covariance $\mathbb{E}[\tilde{\epsilon}_t(\mathbf{x})\tilde{\epsilon}_t(\mathbf{x}')] \propto \delta(\mathbf{x} - \mathbf{x}')$. Note that (15) gives (11) back after quenching (i.e. by sending $\beta \rightarrow \infty$). The result in (15) is stated in Proposition A.5

As we have emphasized, our approach has strong ties with the statistical mechanics of systems of large numbers of interacting particles. Our main aim here is to introduce a framework showing how results and concepts developed in this context are useful to address questions in machine learning. Conversely, we seek to illustrate that ML provides new mathematical questions about an interesting class of particle systems. With this in mind, we adopt a presentation style that relies on formal asymptotic arguments to derive our results, though we are confident that providing rigorous proofs to our propositions is achievable. To a certain extent, this program was already started in [MMN18, CB18, SS18].

2. FUNCTIONAL FORMULATION OF THE LEARNING PROBLEM

As discussed in Bach [Bac17], it is useful to give conditions under which (1) has a limit as $n \rightarrow \infty$, for two main reasons: First it shows which functions can be represented as in (1) if we allow the number of units n to grow to infinity, and second, while the loss function (3) may be non-convex for $\{\theta_i\}_{i=1}^n$, the limiting functional for the parameter distribution is convex.

2.1. Universal Approximation Theorem. Consider the space \mathcal{F}_1 of all functions that can be represented as

$$(16) \quad f = \int_{\hat{D}} \hat{\phi}(\cdot, \mathbf{z}) \gamma(d\mathbf{z})$$

where γ is some (signed) Radon measure on D with finite total variation (L^1 -norm), $|\gamma|_{\text{TV}} = \int_{\hat{D}} |\gamma(d\mathbf{z})| < \infty$: we will denote the space of these Radon measures by $\mathcal{M}(\hat{D})$ and that of probability measures by $\mathcal{M}_+(\hat{D})$. The space \mathcal{F}_1 is important in our context, since any $f \in \mathcal{F}_1$ can be realized as

$$(17) \quad f = \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n c_i \hat{\phi}(\cdot, \mathbf{z}_i)$$

by drawing $\{c_i, \mathbf{z}_i\}_{i \in \mathbb{N}}$ as follows. Start from the Jordan decomposition for γ [Bil99],

$$(18) \quad \gamma = \gamma_+ - \gamma_-,$$

where γ_+ and γ_- are positive measures with $\text{supp } \gamma_+ \cup \text{supp } \gamma_- = \text{supp } \gamma$ and $\text{supp } \gamma_+ \cap \text{supp } \gamma_- = \emptyset$. Using this decomposition, we can draw \mathbf{z}_i 's independently from $(\gamma_+ + \gamma_-)/|\gamma|_{\text{TV}} \in \mathcal{M}_+(\hat{D})$,

where $|\gamma|_{TV} = \int_{\hat{D}} (\gamma_+(d\mathbf{z}) + \gamma_-(d\mathbf{z})) < \infty$, and set $c_i = +|\gamma|_{TV}$ if $\mathbf{z}_i \in \text{supp } \gamma_+$ and $c_i = -|\gamma|_{TV}$ if $\mathbf{z}_i \in \text{supp } \gamma_-$. By the Law of Large Numbers we then have

$$(19) \quad \gamma_n = \frac{1}{n} \sum_{i=1}^n c_i \delta_{\mathbf{z}_i} \rightarrow \gamma \quad \text{as } n \rightarrow \infty$$

which implies (17). We can also use the Central Limit Theorem to get an approximation error at finite n , a calculation we carry out in Sec. 3.4.

Since the space \mathcal{F}_1 depends on the choice of unit $\hat{\varphi}$, to characterize it we make:

Assumption 2.1. *Both the input space Ω and the feature space \hat{D} are closed (i.e. compact with no boundaries) smooth Riemannian manifolds. The unit is continuously differentiable in \mathbf{z} , i.e. $\forall \mathbf{x} \in \Omega, \hat{\varphi}(\mathbf{x}, \cdot) \in C^1(\hat{D})$.*

Assumption 2.2 (Discriminating unit). *The unit satisfies*

$$(20) \quad \int_{\Omega} g(\mathbf{x}) \hat{\varphi}(\mathbf{x}, \cdot) \nu(d\mathbf{x}) = 0 \quad \text{a.e. in } \hat{D} \Rightarrow g = 0 \quad \text{a.e. in } \Omega$$

The differentiability of $\hat{\varphi}$ in \mathbf{z} is required to guarantee uniqueness of the GD flow.

Theorem 2.3 (Universal Approximation Theorem [Cyb89, Bar93, PS91]). *Under Assumptions 2.1 and 2.2, \mathcal{F}_1 is a dense subspace of $L^2(\Omega, \nu)$, i.e. given any $f \in L^2(\Omega, \nu)$ and $\epsilon > 0$, there exists $\gamma^* \in \mathcal{M}(\hat{D})$ such that $|\gamma^*|_{TV} < \infty$ and*

$$(21) \quad f^* = \int_{\hat{D}} \hat{\varphi}(\cdot, \mathbf{z}) \gamma^*(d\mathbf{z}) \in \mathcal{F}_1$$

satisfies

$$(22) \quad \|f - f^*\|_{L^2(\Omega, \nu)} \leq \epsilon.$$

A similar theorem was originally stated in [Cyb89]. Since its proof is elementary let us reproduce it here:

Proof. The space \mathcal{F}_1 is a linear subspace of $L^2(\Omega, \nu)$ since, if $f = \int_{\hat{D}} \hat{\varphi}(\cdot, \mathbf{z}) \gamma(d\mathbf{z}) \in \mathcal{F}_1$,

$$\begin{aligned} \|f\|_{L^2(\Omega, \nu)}^2 &= \int_{\Omega} \left(\int_{\hat{D}} \hat{\varphi}(\mathbf{x}, \mathbf{z}) \gamma(d\mathbf{z}) \right)^2 \nu(d\mathbf{x}) \\ &= \int_{\hat{D} \times \hat{D}} \hat{K}(\mathbf{z}, \mathbf{z}') \gamma(d\mathbf{z}) \gamma(d\mathbf{z}') \\ &\leq \|\hat{K}\|_{\infty} |\gamma|_{TV}^2 < \infty \end{aligned}$$

where we used $\|\hat{K}\|_{\infty} = \sup_{(\mathbf{z}, \mathbf{z}') \in \hat{D} \times \hat{D}} |\hat{K}(\mathbf{z}, \mathbf{z}')| < \infty$ which follows from Assumption 2.1. To show that \mathcal{F}_1 is dense in $L^2(\Omega, \nu)$, we proceed by contradiction. Assuming that \mathcal{F}_1 is not dense, by the Hahn-Banach theorem, there exists a nonzero linear functional $L : L^2(\Omega, \nu) \rightarrow \mathbb{R}$ such that $Lf = 0$ for all $f \in \mathcal{F}_1$. By the Riesz representation theorem, the action of L on f can be represented as the inner product in $L^2(\Omega, \nu)$ between f and some $g \in L^2(\Omega, \nu)$, i.e. there must exist $g \neq 0$ such that for all $f = \int_{\hat{D}} \hat{\varphi}(\cdot, \mathbf{z}) \gamma(d\mathbf{z}) \in \mathcal{F}_1$ (i.e. all $\gamma \in \mathcal{M}(\hat{D})$ with finite variation)

$$\begin{aligned} 0 &= \int_{\Omega} g(\mathbf{x}) \left(\int_{\hat{D}} \hat{\varphi}(\mathbf{x}, \mathbf{z}) \gamma(d\mathbf{z}) \right) \nu(d\mathbf{x}) \\ &= \int_{\hat{D}} \left(\int_{\Omega} g(\mathbf{x}) \hat{\varphi}(\mathbf{x}, \mathbf{z}) \nu(d\mathbf{x}) \right) \gamma(d\mathbf{z}), \end{aligned}$$

This requires that

$$0 = \int_{\Omega} g(\mathbf{x}) \hat{\varphi}(\mathbf{x}, \cdot) \nu(d\mathbf{x}) \quad \text{a.e. in } \hat{D}.$$

which, by Assumption 2.2, implies that $g = 0$ a.e. in Ω , a contradiction. \square

From now on, we will make

Assumption 2.4. *The target function is representable by the network, i.e., $f \in \mathcal{F}_1$.*

This means that we can take $f = f^*$ in Theorem 2.3.

2.2. Convexification at distributional level. Another advantage of taken the $n \rightarrow \infty$ limit of (1) is that it turns (4) into a quadratic objective function for γ :

$$(23) \quad \mathcal{L}[\int_D \hat{\varphi}(\cdot, \mathbf{z}) \gamma(d\mathbf{z})] = C_f - \int_{\hat{D}} \hat{F}(\mathbf{z}) \gamma(d\mathbf{z}) + \frac{1}{2} \int_{\hat{D} \times \hat{D}} \hat{K}(\mathbf{z}, \mathbf{z}') \gamma(d\mathbf{z}) \gamma(d\mathbf{z}')$$

This means that minimizing (23) over γ rather than (4) over $\{\boldsymbol{\theta}_i\}_{i=1}^n$ is conceptually simpler. In particular, any minimizer γ^* of (23) solves the linear Euler-Lagrange equation:

$$(24) \quad \forall \mathbf{z} \in \hat{D} : \quad \hat{F}(\mathbf{z}) = \int_{\hat{D}} \hat{K}(\mathbf{z}, \mathbf{z}') \gamma^*(d\mathbf{z}')$$

and the loss evaluated on any γ^* has value zero. Indeed, using the definitions of \hat{F} and \hat{K} in (6), (24) can be written as

$$(25) \quad \int_{\Omega} \hat{\varphi}(\mathbf{x}, \mathbf{z}) \left(f(\mathbf{x}) - \int_{\hat{D}} \hat{\varphi}(\mathbf{x}, \mathbf{z}') \gamma^*(d\mathbf{z}') \right) \nu(d\mathbf{x}) = 0$$

which, by Assumptions 2.2 and 2.4, has a solution such that $f = \int_{\hat{D}} \hat{\varphi}(\cdot, \mathbf{z}) \gamma^*(d\mathbf{z})$ and, as a result, the loss evaluated on $\int_{\hat{D}} \hat{\varphi}(\cdot, \mathbf{z}) \gamma^*(d\mathbf{z})$ is zero.

Of course, the results above are not necessarily an assurance of convergence in practice. Indeed, we do not know how to pick $\gamma \in \mathcal{M}(\hat{D})$ to represent an $f \in \mathcal{F}_1$ nor can we manipulate these Radon measures explicitly: rather we will have to learn finite n approximations of the form $\gamma^{(n)} = n^{-1} \sum_{i=1}^n c_i \delta_{\mathbf{z}_i}$ by adjusting the parameters $\{\boldsymbol{\theta}_i\}_{i=1}^n = \{c_i, \mathbf{z}_i\}_{i=1}^n$ dynamically. Furthermore, even though the energy can be expressed in term of $\gamma^{(n)}$, as we will see the dynamics can only be closed at the level of the empirical distribution

$$(26) \quad \mu^{(n)}(dc, d\mathbf{z}) = \frac{1}{n} \sum_{i=1}^n \delta_{c_i}(dc) \delta_{\mathbf{z}_i}(d\mathbf{z}) \equiv \frac{1}{n} \sum_{i=1}^n \delta_{\boldsymbol{\theta}_i}(d\boldsymbol{\theta}) = \mu^{(n)}(d\boldsymbol{\theta})$$

with $\gamma^{(n)}$ given by

$$(27) \quad \gamma^{(n)} = \int_{\mathbb{R}} c \mu^{(n)}(dc, \cdot).$$

Viewed as a functional of $\mu \in \mathcal{M}_+(D)$ such that $\int_{\mathbb{R}} c \mu(dc, \cdot) = \gamma \in \mathcal{M}(\hat{D})$, the loss function (23) becomes

$$(28) \quad \begin{aligned} \mathcal{E}[\mu] &= C_f - \int_D F(\boldsymbol{\theta}) \mu(d\boldsymbol{\theta}) + \frac{1}{2} \int_{D \times D} K(\boldsymbol{\theta}, \boldsymbol{\theta}') \mu(d\boldsymbol{\theta}) \mu(d\boldsymbol{\theta}') \\ &= \frac{1}{2} \mathbb{E}_{\nu} \left(f - \int_D \varphi(\cdot, \boldsymbol{\theta}) \mu(d\boldsymbol{\theta}) \right)^2 \geq 0 \end{aligned}$$

3. TRAINING BY GRADIENT DESCENT ON THE EXACT LOSS

Here we assume that we train the network by evolving dynamically the parameters $\{\boldsymbol{\theta}_i(t)\}_{i=1}^n$ according to the GD flow (7), which we recall is given by the coupled ODEs,

$$(29) \quad \dot{\boldsymbol{\theta}}_i = \nabla F(\boldsymbol{\theta}_i) - \frac{1}{n} \sum_{j=1}^n \nabla K(\boldsymbol{\theta}_i, \boldsymbol{\theta}_j),$$

for $i = 1, \dots, n$. As we show in Sec. 4, (29) shares many properties with the stochastic gradient descent (SGD) used in applications, though in SGD a multiplicative noise term persists in the equations. The ODEs in (29) are the GD flow on the energy:

$$(30) \quad E(\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_n) = nC_f - \sum_{i=1}^n F(\boldsymbol{\theta}_i) + \frac{1}{2n} \sum_{i,j=1}^n K(\boldsymbol{\theta}_i, \boldsymbol{\theta}_j)$$

This energy is simply the loss function in (4) rescaled by n .

We consider (29) with initial conditions such that every $\theta_i(0)$ for $i = 1, \dots, n$ is drawn independently from some probability distribution μ_{in} satisfying

Assumption 3.1. *The distribution μ_{in} is such that: (i) its support contains a smooth manifold that separates the regions in $D = \mathbb{R} \times \hat{D}$ where $c > c_0$ and $c < -c_0$ for some large enough $c_0 > 0$; (ii) $\gamma_{\text{in}} = \int_{\mathbb{R}} c \mu_{\text{in}}(dc, \cdot)$ has finite total variation, $|\gamma_{\text{in}}|_{TV} < \infty$; and (iii) $\forall b \in \mathbb{R} : \int_{\mathbb{R} \times \hat{D}} e^{bc} \mu_{\text{in}}(dc, d\mathbf{z}) < \infty$.*

Note that property (i) guarantees that $\hat{\mu}_{\text{in}} = \int_{\mathbb{R}} \mu_{\text{in}}(dc, \cdot)$ has full support in \hat{D} , $\text{supp } \hat{\mu}_{\text{in}} = \hat{D}$ —we show below that this property is preserved by the dynamics. Distributions μ_{in} that satisfy Assumption 3.1 include e.g.

$$\delta_0(dc) \hat{\mu}_{\text{in}}(d\mathbf{z}) \quad \text{and} \quad (2\pi)^{-1/2} e^{-\frac{1}{2}c^2} dc \hat{\mu}_{\text{in}}(d\mathbf{z}),$$

if $\text{supp } \hat{\mu}_{\text{in}} = \hat{D}$ in both. We denote the measure for the infinite set $\{\theta_i(0)\}_{i \in \mathbb{N}}$ constructed this way by \mathbb{P}_{in} . Initial conditions of this type are used in practice.

3.1. Empirical distribution and nonlinear Liouville equation. To proceed, we consider the empirical distribution

$$(31) \quad \mu_t^{(n)} = \frac{1}{n} \sum_{i=1}^n \delta_{\theta_i(t)}$$

in terms of which we can express (8) as

$$(32) \quad f_t^{(n)} = \frac{1}{n} \sum_{i=1}^n \varphi(\cdot, \theta_i(t)) = \int_{D \times R} \varphi(\cdot, \boldsymbol{\theta}) \mu_t^{(n)}(d\boldsymbol{\theta}).$$

The empirical distribution (31) is useful to work with because it satisfies a nonlinear Liouville type equation

$$(33) \quad \partial_t \mu_t^{(n)} = \nabla \cdot \left(\nabla V(\boldsymbol{\theta}, [\mu_t^{(n)}]) \mu_t^{(n)} \right)$$

where we defined

$$(34) \quad V(\boldsymbol{\theta}, [\mu]) = -F(\boldsymbol{\theta}) + \int_D K(\boldsymbol{\theta}, \boldsymbol{\theta}') \mu(d\boldsymbol{\theta}')$$

Throughout, we will interpret (33) in the standard weak sense, as in (36) below. When there is a Laplacian term in (33) this equation is called the McKean-Vlasov equation [McK66, DG87, Gär88, Szn91]; with an additional noise term added it is often referred to as Dean's equation [Dea99]. To prove asymptotic trainability results, we analyze the properties of the solution to this equation as $n \rightarrow \infty$ and $t \rightarrow \infty$.

Derivation of (33). Let $\chi : D \rightarrow \mathbb{R}$ be a test function, and consider

$$(35) \quad \int_D \chi(\boldsymbol{\theta}) \mu_t^{(n)}(d\boldsymbol{\theta}) = \frac{1}{n} \sum_{i=1}^n \chi(\boldsymbol{\theta}_i(t))$$

Taking the time derivative of this equation and using (29) we deduce

$$(36) \quad \begin{aligned} & \int_D \chi(\boldsymbol{\theta}) \partial_t \mu_t^{(n)}(d\boldsymbol{\theta}) \\ &= \frac{1}{n} \sum_{i=1}^n \nabla \chi(\boldsymbol{\theta}_i(t)) \cdot \dot{\boldsymbol{\theta}}_i(t) \\ &= \frac{1}{n} \sum_{i=1}^n \nabla \chi(\boldsymbol{\theta}_i(t)) \cdot \left(\nabla F(\boldsymbol{\theta}_i(t)) - \frac{1}{n} \sum_{j=1}^n \nabla K(\boldsymbol{\theta}_i(t), \boldsymbol{\theta}_j(t)) \right) \\ &= \int_D \nabla \chi(\boldsymbol{\theta}) \cdot \left(\nabla F(\boldsymbol{\theta}) - \int_D \nabla K(\boldsymbol{\theta}, \boldsymbol{\theta}') \mu_t^{(n)}(d\boldsymbol{\theta}') \right) \mu_t^{(n)}(d\boldsymbol{\theta}) \end{aligned}$$

This is the weak form of (33).

3.2. Law of Large Numbers (LLN)—mean field limit. Since we know that, \mathbb{P}_{in} -almost surely as $n \rightarrow \infty$, $\mu_0^{(n)} \rightarrow \mu_{\text{in}}$ by the Law of Large Numbers, we can take the limit as $n \rightarrow \infty$ of (33) to formally deduce:

Proposition 3.2. *Let $\mu_t^{(n)}$ be given by (31) with $\{\theta_i(t)\}_{i=1}^n$ the solution of (29) with initial condition drawn from \mathbb{P}_{in} . Then, as $n \rightarrow \infty$, $\mu_t^{(n)} \rightarrow \mu_t$ a.s. where μ_t satisfies*

$$(37) \quad \partial_t \mu_t = \nabla \cdot (\nabla V(\boldsymbol{\theta}, [\mu_t]) \mu_t) \quad \mu_0 = \mu_{\text{in}},$$

interpreted in the weak sense.

Note that (37) is the same as (33) but with a different initial condition. Note also that (37) is the GD flow in the Wasserstein metric [Vil09, AGS05]. Indeed this equation can be written as the $\tau \rightarrow 0$ limit of the Jordan-Kinderlehrer-Otto (JKO) proximal scheme [JKO98]

$$(38) \quad \mu_{t+\tau} \in \operatorname{argmin} \left(\mathcal{E}[\mu] + \frac{1}{2} \tau^{-1} W_2^2(\mu, \mu_t) \right), \quad \mu_0 = \mu_{\text{in}}$$

where $W_2(\mu, \mu_t)$ is the 2-Wasserstein distance between μ and μ_t and $\mathcal{E}[\mu]$ is defined in (28). Finally, note that the weak solutions of (37) satisfy: for any test function $\chi: D \rightarrow \mathbb{R}$,

$$(39) \quad \int_D \chi(\boldsymbol{\theta}) \mu_t(d\boldsymbol{\theta}) = \int_D \chi(\boldsymbol{\Theta}_t(\boldsymbol{\theta})) \mu_{\text{in}}(d\boldsymbol{\theta})$$

where $\boldsymbol{\Theta}_t(\boldsymbol{\theta})$ solves is given in terms of characteristics

$$(40) \quad \dot{\boldsymbol{\Theta}}_t(\boldsymbol{\theta}) = -\nabla V(\boldsymbol{\Theta}_t(\boldsymbol{\theta}), [\mu_t]), \quad \boldsymbol{\Theta}_0(\boldsymbol{\theta}) = \boldsymbol{\theta}.$$

Of course, (39) is not explicit since (40) depends on μ_t , but this representation formula is useful in the sequel. In particular, notice that it implies that: (i) $\mu_t \in \mathcal{M}_+(D)$ for all $t < \infty$ since the velocity field in (40) is globally Lipschitz on $\mathbb{R} \times \hat{D}$ by Assumption 2.1 and hence the solutions to this equation exist for all $t < \infty$; and (ii) $\operatorname{supp} \hat{\mu}_t = \hat{D}$ with $\hat{\mu}_t = \int_{\mathbb{R}} \mu_t(d\mathbf{x}, \cdot)$ by Assumption 3.1, and $\operatorname{supp} \mu_t = D$ if $\operatorname{supp} \mu_{\text{in}} = D$.

The dynamics of $f_t = \lim_{n \rightarrow \infty} f_t^{(n)}$. We now discuss the implications of the limiting PDE for the evolution of

$$(41) \quad \lim_{n \rightarrow \infty} f_t^{(n)} = \lim_{n \rightarrow \infty} \int_D \varphi(\cdot, \boldsymbol{\theta}) \mu_t^{(n)}(d\boldsymbol{\theta}) = \int_D \varphi(\cdot, \boldsymbol{\theta}) \mu_t(d\boldsymbol{\theta}) \equiv f_t$$

To begin, note that from (34) we can express $V(\boldsymbol{\theta}, [\mu_t])$ as

$$(42) \quad V(\boldsymbol{\theta}, [\mu_t]) = \int_{\Omega} (f_t(\mathbf{x}) - f(\mathbf{x})) \varphi(\mathbf{x}, \boldsymbol{\theta}) v(d\mathbf{x})$$

As a result (37) can be written as

$$(43) \quad \partial_t \mu_t = \nabla \cdot \left(\int_{\Omega} \nabla_{\boldsymbol{\theta}} \varphi(\mathbf{x}, \boldsymbol{\theta}) (f_t(\mathbf{x}) - f(\mathbf{x})) v(d\mathbf{x}) \mu_t \right)$$

and we deduce, using (41),

$$(44) \quad \begin{aligned} \partial_t f_t &= \int_D \varphi(\cdot, \boldsymbol{\theta}) \partial_t \mu_t(d\boldsymbol{\theta}) \\ &= - \int_D \nabla_{\boldsymbol{\theta}} \varphi(\cdot, \boldsymbol{\theta}) \cdot \left(\int_{\Omega} \nabla_{\boldsymbol{\theta}} \varphi(\mathbf{x}', \boldsymbol{\theta}) (f_t(\mathbf{x}') - f(\mathbf{x}')) v(d\mathbf{x}') \mu_t(d\boldsymbol{\theta}) \right) \end{aligned}$$

Interchanging the order of integration gives:

Proposition 3.3 (LLN). *Let $f_t^{(n)}$ be given by (32) with $\{\theta_i(t)\}_{i=1}^n$ solution of (29) with initial condition drawn from \mathbb{P}_{in} . Then, as $n \rightarrow \infty$, $f_t^{(n)} \rightarrow f_t$ a.s. pointwise, where f_t satisfies*

$$(45) \quad \partial_t f_t(\mathbf{x}) = - \int_{\Omega} M([\mu_t], \mathbf{x}, \mathbf{x}') (f_t(\mathbf{x}') - f(\mathbf{x}')) v(d\mathbf{x}')$$

where we defined the kernel

$$\begin{aligned}
 (46) \quad M([\mu], \mathbf{x}, \mathbf{x}') &= \int_D \nabla_{\boldsymbol{\theta}} \varphi(\mathbf{x}, \boldsymbol{\theta}) \cdot \nabla_{\boldsymbol{\theta}} \varphi(\mathbf{x}', \boldsymbol{\theta}) \mu(d\boldsymbol{\theta}) \\
 &= \int_{\mathbb{R} \times \hat{D}} (c^2 \nabla_{\mathbf{z}} \hat{\varphi}(\mathbf{x}, \mathbf{z}) \cdot \nabla_{\mathbf{z}} \hat{\varphi}(\mathbf{x}', \mathbf{z}) + \hat{\varphi}(\mathbf{x}, \mathbf{z}) \hat{\varphi}(\mathbf{x}', \mathbf{z})) \mu(dc, d\mathbf{z}).
 \end{aligned}$$

The kernel (46) is symmetric in \mathbf{x} and \mathbf{x}' for any $\mu \in \mathcal{M}(D)$ and positive semidefinite if $\mu \in \mathcal{M}_+(D)$ since, given any $r \in L^2(\Omega, \nu)$, we then have

$$\begin{aligned}
 (47) \quad &\int_{\Omega^2} r(\mathbf{x}) r(\mathbf{x}') M([\mu], \mathbf{x}, \mathbf{x}') \nu(d\mathbf{x}) \nu(d\mathbf{x}') \\
 &= \int_{\mathbb{R} \times \hat{D}} (c^2 |\nabla_{\mathbf{z}} R(\mathbf{z})|^2 + |R(\mathbf{z})|^2) \mu(dc, d\mathbf{z}) \geq 0
 \end{aligned}$$

where

$$(48) \quad R(\mathbf{z}) = \int_{\Omega} r(\mathbf{x}) \hat{\varphi}(\mathbf{x}, \mathbf{z}) \nu(d\mathbf{x}).$$

Equation (45) also confirms that f_t evolves on a quadratic landscape, namely the loss function (3) itself: Indeed this equation can be written as

$$(49) \quad \partial_t f_t(\mathbf{x}) = - \int_{\Omega} M([\mu_t], \mathbf{x}, \mathbf{x}') D_{f_t(\mathbf{x}')} \mathcal{L}[f_t] \nu(d\mathbf{x}')$$

where $D_{f(\mathbf{x})}$ denotes the gradient with respect to $f(\mathbf{x})$ in the $L^2(\Omega, \nu)$ -norm, i.e. given a functional $\mathcal{F}[f]$,

$$(50) \quad \forall h : \Omega \rightarrow \mathbb{R} : \lim_{z \rightarrow 0} \frac{d}{dz} \mathcal{F}[f + zh] = \langle h, D_f \mathcal{F}[f] \rangle_{L^2(\Omega, \nu)} = \int_{\Omega} h(\mathbf{x}) D_{f(\mathbf{x})} \mathcal{F}[f] \nu(d\mathbf{x}).$$

That is, $D_{f(\mathbf{x})}$ reduces to $\delta/\delta f(\mathbf{x})$ if $\nu(d\mathbf{x}) = d\mathbf{x}$.

3.3. Long time behavior—global convergence. Let us now analyze the long-time solutions of (37) for the weak limit μ_t of $\mu_t^{(n)}$ and (45) for the limit f_t of $f_t^{(n)}$. As is well-known, (37) has more stationary points than $\mathcal{E}[\mu]$ has minimizers. Since (37) is the Wasserstein GD flow on $\mathcal{E}[\mu]$, a direct calculation shows that $E_t = \mathcal{E}[\mu_t]$ satisfies

$$(51) \quad \frac{dE_t}{dt} = - \int_D |\nabla V(\boldsymbol{\theta}, [\mu_t])|^2 \mu_t(d\boldsymbol{\theta})$$

This equation implies that the stationary points μ^s of (37) are the solutions of

$$(52) \quad \nabla V(\boldsymbol{\theta}, [\mu^s]) = 0 \quad \text{for } \boldsymbol{\theta} \in \text{supp } \mu^s.$$

This should be contrasted with the minimizers of $\mathcal{E}[\mu]$, which satisfy:

$$(53) \quad \begin{cases} V(\boldsymbol{\theta}, [\mu^*]) \geq \bar{V}[\mu^*] & \text{for } \boldsymbol{\theta} \in D \\ V(\boldsymbol{\theta}, [\mu^*]) = \bar{V}[\mu^*] & \text{for } \boldsymbol{\theta} \in \text{supp } \mu^* \end{cases}$$

where $\bar{V}[\mu] = \int_D V(\boldsymbol{\theta}, [\mu]) \mu(d\boldsymbol{\theta})$. In general, we cannot guarantee that the solutions to (52) also solve (53). However, due to the specific form of the unit, $\varphi(\mathbf{x}, \boldsymbol{\theta}) = c\hat{\varphi}(\mathbf{x}, \mathbf{z})$, the rate of decay of the energy (51) reads

$$\begin{aligned}
 (54) \quad \frac{dE}{dt} &= - \int_{\mathbb{R} \times \hat{D}} (c^2 |\nabla \hat{V}(\mathbf{z}, [\mu_t])|^2 + |\hat{V}(\mathbf{z}, [\mu_t])|^2) \mu_t(dc, d\mathbf{z}) \\
 &= - \int_{\mathbb{R} \times \hat{D}} c^2 |\nabla \hat{V}(\mathbf{z}, [\mu_t])|^2 \mu_t(dc, d\mathbf{z}) - \int_{\hat{D}} |\hat{V}(\mathbf{z}, [\mu_t])|^2 \hat{\mu}_t(d\mathbf{z})
 \end{aligned}$$

where $\hat{\mu}_t = \int_{\mathbb{R}} \mu_t(dc, \cdot)$ and

$$(55) \quad \hat{V}(\mathbf{z}, [\mu]) = -\hat{F}(\mathbf{z}) + \int_{\mathbb{R} \times \hat{D}} c' \hat{K}(\mathbf{z}, \mathbf{z}') \mu(dc', d\mathbf{z}')$$

(54) implies that the stationary points μ^s of (37) satisfy

$$(56) \quad \hat{V}(\mathbf{z}, [\mu^s]) = 0 \quad \text{for } \mathbf{z} \in \text{supp } \hat{\mu}^s = \int_{\mathbb{R}} \mu^s(dc, \cdot)$$

As a result $V(\boldsymbol{\theta}, [\mu^s]) = c\hat{V}(\mathbf{z}, [\mu^s]) = 0$ for $\boldsymbol{\theta} = (c, \mathbf{z}) \in \text{supp } \mu^s$, and this shows that the second equation in (53) is automatically satisfied, noting that $\bar{V} = 0$ for a global minimizer.

To show that the first equation of (56) also holds, we establish that $\hat{V}(\mathbf{z}, [\mu^s]) = 0$ everywhere in \hat{D} . We proceed by contradiction: Suppose that μ_t converges to some μ^s such that $\hat{V}(\mathbf{z}, [\mu^s]) \neq 0$ for $\mathbf{z} \in \hat{D}_s^c$ where \hat{D}_s^c is the complement in \hat{D} of $\hat{D}_s = \text{supp } \hat{\mu}^s$ —the relevant case is when \hat{D}_s^c has nonzero Hausdorff measure in \hat{D} . Looking at the characteristic equations (40) written in terms of $\boldsymbol{\Theta}_t = (C_t, \mathbf{Z}_t)$ as

$$(57) \quad \begin{cases} \dot{C}_t(c, \mathbf{z}) = -\hat{V}(\mathbf{Z}_t(c, \mathbf{z}), [\mu_t]), & C_0(c, \mathbf{z}) = c \\ \dot{\mathbf{Z}}_t(c, \mathbf{z}) = -C_t(c, \mathbf{z})\nabla\hat{V}(\mathbf{Z}_t(c, \mathbf{z}), [\mu_t]), & \mathbf{Z}_0(c, \mathbf{z}) = \mathbf{z} \end{cases}$$

Since we know that $\text{supp } \hat{\mu}_t = \hat{D}$ at all positive time $t < \infty$, in order that $\mu_t \rightarrow \mu^s$ as $t \rightarrow \infty$, all the mass must be expelled from \hat{D}_s^c . That is, all the solutions to (57) must leave this domain, or at least accumulate at its boundary, while at the same time we must have $\lim_{t \rightarrow \infty} \hat{V}(\mathbf{z}, [\mu_t]) \neq 0$. To show that this scenario is impossible, note that (using the fact that \hat{D} is compact)

$$(58) \quad \forall \delta > 0 \quad \exists t_c > 0 : \sup_{\mathbf{z}} |\hat{V}(\mathbf{z}, [\mu_t]) - \hat{V}(\mathbf{z}, [\mu_s])| \leq \delta \quad \text{if } t \geq t_c.$$

This means that, for $t \geq t_c$, to leading order in δ (57) reads

$$(59) \quad \begin{cases} \dot{C}_t(c, \mathbf{z}) = -\hat{V}(\mathbf{Z}_t(c, \mathbf{z}), [\mu^s]), & C_0(c, \mathbf{z}) = c \\ \dot{\mathbf{Z}}_t(c, \mathbf{z}) = -C_t(c, \mathbf{z})\nabla\hat{V}(\mathbf{Z}_t(c, \mathbf{z}), [\mu^s]), & \mathbf{Z}_0(c, \mathbf{z}) = \mathbf{z} \end{cases}$$

which is the GD flow on

$$(60) \quad c\hat{V}(\mathbf{z}, [\mu^s])$$

Suppose that $\hat{V}(\mathbf{z}, [\mu^s]) > 0$ somewhere in \hat{D}_s^c —the case when $\hat{V}(\mathbf{z}, [\mu^s]) < 0$ somewhere in \hat{D}_s^c can be treated similarly. Since \hat{D}_s^c is compact, $\hat{V}(\mathbf{z}, [\mu^s])$ must then have a maximum in \hat{D}_s^c , i.e. (using the differentiability of the unit) there exists $\mathbf{z}_1 \in \hat{D}_s^c$ with $\mathbf{z}_1 \notin \partial\hat{D}_s^c$ and such that $\nabla\hat{V}(\mathbf{z}_1, [\mu^s]) = 0$, $\hat{V}(\mathbf{z}_1, [\mu^s]) = \hat{V}_1 > 0$, and $\hat{V}(\mathbf{z}_1, [\mu^s]) > \hat{V}(\mathbf{z}, [\mu^s])$ for $\mathbf{z} \in \hat{D}_s^c$. Consider the solutions to (59) for initial (c, \mathbf{z}) such that $\mathbf{Z}_t(c, \mathbf{z})$ is very close to \mathbf{z}_1 at $t = t_c$ —these solutions must exist since $\text{supp } \mu_t = \hat{D}$ for all $t < \infty$. If among these solutions there are some such that $C_t(c, \mathbf{z})$ is negative at time $t = t_c$ (which is always the case if $\text{supp } \mu_{\text{in}} = D$ since $\text{supp } \mu_t = D$ for all $t < \infty$ in that case), then by (59) $C_t(c, \mathbf{z})$ becomes more negative and $\mathbf{Z}_t(c, \mathbf{z})$ gets closer to \mathbf{z}_1 for $t > t_c$. If all $C_{t+\tau}(c, \mathbf{z})$ are positive at time $t = t_c(\delta)$, then the corresponding $\mathbf{Z}_t(c, \mathbf{z})$ go away from \mathbf{z}_1 for as long as their $C_t(c, \mathbf{z})$ remains positive; however, eventually some $C_t(c, \mathbf{z})$ become negative (since $C_{t+\tau}(c, \mathbf{z}_1) = C_t(c, \mathbf{z}_1) - \tau\hat{V}(\mathbf{z}_1, [\mu^s]) = C_t(c, \mathbf{z}_1) - \tau\hat{V}_1$ under (59)), at which point we go back to the first case and $\mathbf{Z}_t(c, \mathbf{z})$ gets closer to \mathbf{z}_1 . Either way, we can always find solutions with $\mathbf{Z}_t(c, \mathbf{z})$ sufficiently close to \mathbf{z}_1 at time $t_c(\delta)$ that will eventually converge to \mathbf{z}_1 rather than exiting \hat{D}_s^c , a contradiction with our assumption that all solutions must either exit this domain or accumulate at its boundary. This argument is based on (59) rather than the original (57), but by setting δ small enough (and t_c large enough) we can make the terms left over in (59) arbitrarily small so that they do not affect the result.

This concludes the justification that the stationary points μ^s of (37) are such that $\hat{V}(\mathbf{z}, [\mu^s]) = 0$ everywhere in \hat{D} , i.e. they are minimizers of $\mathcal{E}[\mu]$, which from (56) implies

$$(61) \quad \forall \mathbf{z} \in \hat{D} : \quad 0 = \int_{\Omega} \hat{\varphi}(\mathbf{x}, \mathbf{z}) \left(f(\mathbf{x}) - \int_{\hat{D}} \hat{\varphi}(\mathbf{x}, \mathbf{z}') \gamma^s(d\mathbf{z}') \right) v(d\mathbf{x})$$

where $\gamma^s = \int_{\mathbb{R}} c\mu^s(dc, \cdot)$. As a result, by Assumptions 2.2 and 2.4,

$$(62) \quad f = \int_{\hat{D}} \hat{\varphi}(\cdot, \mathbf{z}) \gamma^s(d\mathbf{z}),$$

In other other words, we have established:

Proposition 3.4 (Global convergence). *Let μ_t be the solution to (37) for the initial condition $\mu_0 = \mu_{in}$ that satisfy Assumption 3.1. If $\mu_t \rightarrow \mu^* \in \mathcal{M}_+(D)$ as $t \rightarrow \infty$, then under Assumptions 2.2 and 2.4 μ^* is a minimizer of $\mathcal{E}[\mu]$ and we have*

$$(63) \quad \lim_{t \rightarrow \infty} \int_D \varphi(\cdot, \boldsymbol{\theta}) \mu_t(d\boldsymbol{\theta}) = \int_D \varphi(\cdot, \boldsymbol{\theta}) \mu^*(d\boldsymbol{\theta}) = f.$$

Note that we assume that μ_t converges to some probability measure to establish this proposition. This is because we cannot exclude *a priori* that the dynamics eventually loses mass at infinity, e.g. if some of the solutions of the characteristic equation (40) eventually diverge as $t \rightarrow \infty$. We do not expect this scenario to occur for most initial conditions and one way to preclude it altogether is to add regularizing terms in the loss function.

The argument that leads to (63) would be simple if it were the case that $\hat{\mu}^* = \int_{\mathbb{R}} \mu^*(dc, \cdot)$ has full support in \hat{D} . Indeed this would imply that the kernel (46) evaluated on μ_t is positive definite not only for all $t \geq 0$ but also in the limit as $t \rightarrow \infty$ and hence the only fixed point of (45) is f . It is reasonable to assume that $\text{supp } \hat{\mu}^* = \hat{D}$ because: (i) $\text{supp } \hat{\mu}_t = \hat{D}$ for all $t < \infty$ as mentioned before and (ii) there is no energetic incentive to shrink the support, even when $t \rightarrow \infty$. This see why, note that if μ^* is an energy minimizer such that $\text{supp } \hat{\mu}^* \neq \hat{D}$, then a direct calculation shows that for any $\alpha \in (0, 1)$ and any $\hat{\mu} \in \mathcal{M}_+(\hat{D})$ with $\text{supp } \hat{\mu} = \hat{D}$,

$$(64) \quad \mu^{**}(dc, dz) = (1 - \alpha)^2 \mu^*((1 - \alpha)dc, dz) + \alpha \delta_0(dc) \hat{\mu}(dz)$$

is also a energy minimizer in $\mathcal{M}_+(D)$ such that $\hat{\mu}^{**} = \int_{\mathbb{R}} \mu^{**}(dc, \cdot)$ has support \hat{D} .

In Appendix, we analyze the behavior of μ_t on a longer timescale and show that, with noise and certain regularizing terms added in (29), μ_t reaches a unique fixed point $\mu^* \in \mathcal{M}_+(D)$ such that $\int_D \log(d\mu^*/d\mu^0) d\mu^* < \infty$, where μ^0 is some prior measure used for regularization.

We can summarize the results of Secs. 3.2 and 3.3 into:

Proposition 3.5 (LLN & global convergence). *Let $f_t^{(n)}$ be given by (32) with $\{\boldsymbol{\theta}_i(t)\}_{i=1}^n$ solution of (29) with initial condition drawn from \mathbb{P}_{in} . Then under the conditions of Proposition 3.4 we have*

$$(65) \quad \lim_{n \rightarrow \infty} f_t^{(n)} = f_t \quad \text{pointwise, } \mathbb{P}_{in}\text{-almost surely}$$

where f_t solves (45) and satisfies

$$(66) \quad \lim_{t \rightarrow \infty} f_t = f \quad \text{a.e. in } \Omega.$$

The convergence in (66) is equivalent to the statement in Proposition 3.4. Notice that, since the evolution of f_t occurs via (45), which is independent of n , for any $\delta > 0$ we can find t_c independent of n such that for $t > t_c$, $\mathbb{E}_v |f_t - f|^2 < \delta$. Since for any $\delta > 0$ and $t > 0$ we can also find n_c such that for $n > n_c$, $\mathbb{E}_v |f_t^{(n)} - f_t|^2 < \delta$, we can interchange the limits in n and t in Theorem 3.5, i.e. we also have

$$(67) \quad \lim_{n \rightarrow \infty} \lim_{t \rightarrow \infty} f_t^{(n)} = f.$$

3.4. Central Limit Theorem (CLT). Let us now consider the fluctuations of $\mu_t^{(n)}$ around its limit μ_t . To this end, we define $\omega_t^{(n)}$ via:

$$(68) \quad \omega_t^{(n)} = n^{1/2} (\mu_t^{(n)} - \mu_t),$$

Explicitly, (68) means:

$$(69) \quad \omega_t^{(n)} = n^{-1/2} \sum_{i=1}^n (\delta_{\boldsymbol{\theta}_i(t)} - \mu_t)$$

The scaling factor $n^{1/2}$ is set by the fluctuations in the initial conditions: if we pick a test function $\chi: D \rightarrow \mathbb{R}$ the CLT tells us that under \mathbb{P}_{in}

$$(70) \quad \int_D \chi(\boldsymbol{\theta}) \omega_0^{(n)}(d\boldsymbol{\theta}) = n^{-1/2} \sum_{i=1}^n \tilde{\chi}(\boldsymbol{\theta}_i(0)) \rightarrow N(0, C_\chi) \quad \text{in law as } n \rightarrow \infty$$

where $\tilde{\chi}(\boldsymbol{\theta}) = \chi(\boldsymbol{\theta}) - \int_D \chi(\boldsymbol{\theta}) \mu_{\text{in}}(d\boldsymbol{\theta})$ and $N(0, C_\chi)$ denotes the Gaussian random variable with mean zero and variance

$$(71) \quad C_\chi = \int_D |\tilde{\chi}(\boldsymbol{\theta})|^2 \mu_{\text{in}}(d\boldsymbol{\theta}),$$

To see what happens at later times, we derive an equation for $\omega_t^{(n)}$ by subtracting (37) from (33) and using (68)

$$(72) \quad \partial_t \omega_t^{(n)} = \nabla \cdot \left(\omega_t^{(n)} \nabla V(\boldsymbol{\theta}, [\mu_t]) + \left(\mu_t + n^{-1/2} \omega_t^{(n)} \right) \nabla F(\boldsymbol{\theta}, [\omega_t^{(n)}]) \right)$$

where we defined

$$(73) \quad F(\boldsymbol{\theta}, [\mu]) = \int_D K(\boldsymbol{\theta}, \boldsymbol{\theta}') \mu(d\boldsymbol{\theta}')$$

If we take the limit as $n \rightarrow \infty$, the term involving $n^{-1/2} \omega_t^{(n)}$ at the right hand side of (72) disappears (we quantify its rate of convergence to zero in more detail in Sec. 3.5) and we formally deduce that

Proposition 3.6. *Let $\omega_t^{(n)}$ be given by (69) with $\{\boldsymbol{\theta}_i(t)\}_{i=1}^n$ solution of (29) with initial conditions draw from \mathbb{P}_{in} and μ_t solution to (43). Then*

$$(74) \quad \omega_t^{(n)} \rightharpoonup \omega_t \quad \text{in law as } n \rightarrow \infty$$

where ω_t satisfies

$$(75) \quad \partial_t \omega_t = \nabla \cdot \left(\omega_t \nabla V(\boldsymbol{\theta}, [\mu_t]) + \mu_t \nabla F(\boldsymbol{\theta}, [\omega_t]) \right)$$

to be solved in the weak sense with the Gaussian initial conditions read from (70):

$$(76) \quad \int_D \chi(\boldsymbol{\theta}) \omega_0(d\boldsymbol{\theta}) = N(0, C_\chi)$$

Note that since the mean of ω_0 is zero initially and (75) is linear, this mean remains zero for all times, and we can focus on the evolution of its covariance:

$$(77) \quad \Sigma_t(d\boldsymbol{\theta}, d\boldsymbol{\theta}') = \mathbb{E}_{\text{in}}[\omega_t(d\boldsymbol{\theta}) \omega_t(d\boldsymbol{\theta}')]$$

From (75) it satisfies

$$(78) \quad \begin{aligned} \partial_t \Sigma_t &= \nabla \cdot \left(\Sigma_t \nabla V(\boldsymbol{\theta}, [\mu_t]) + \mu_t(d\boldsymbol{\theta}) \nabla G(\boldsymbol{\theta}, d\boldsymbol{\theta}', [\Sigma_t]) \right) \\ &+ \nabla' \cdot \left(\Sigma_t \nabla V(\boldsymbol{\theta}', [\mu_t]) + \mu_t(d\boldsymbol{\theta}') \nabla G(\boldsymbol{\theta}', d\boldsymbol{\theta}, [\Sigma_t]) \right) \end{aligned}$$

where we defined

$$(79) \quad G(\boldsymbol{\theta}, \cdot, [\Sigma]) = \int_D K(\boldsymbol{\theta}, \boldsymbol{\theta}'') \Sigma(d\boldsymbol{\theta}'', \cdot)$$

Equation (78) should be interpreted in the weak sense and solved for the initial condition

$$(80) \quad \Sigma_0(d\boldsymbol{\theta}, d\boldsymbol{\theta}') = \mu_{\text{in}}(d\boldsymbol{\theta}) \delta_{\boldsymbol{\theta}}(d\boldsymbol{\theta}')$$

The dynamics of $g_t = \lim_{n \rightarrow \infty} n^{1/2} (f_t^{(n)} - f_t)$. We can also test these equations against the unit, to deduce that, as $n \rightarrow \infty$,

$$\begin{aligned} g_t^{(n)} &= \int_D \varphi(\cdot, \boldsymbol{\theta}) \omega_t^{(n)}(d\boldsymbol{\theta}) = n^{1/2} (f_t^{(n)} - f_t) \\ &= n^{-1/2} \sum_{i=1}^n (\varphi(\cdot, \boldsymbol{\theta}_i(t)) - f_t) \end{aligned} \quad (81)$$

converges in law, $g_t^{(n)} \rightarrow g_t$, where g_t is a Gaussian process satisfying

$$\begin{aligned} \partial_t g_t &= - \int_{\Omega} M(\mathbf{x}, \mathbf{x}', [\omega_t]) (f_t(\mathbf{x}') - f(\mathbf{x}')) \nu(d\mathbf{x}') \\ &\quad - \int_{\Omega} M(\mathbf{x}, \mathbf{x}', [\mu_t]) g_t(\mathbf{x}') \nu(d\mathbf{x}') \end{aligned} \quad (82)$$

This equation should be solved with Gaussian initial conditions with mean zero and covariance

$$\begin{aligned} C_0(\mathbf{x}, \mathbf{x}') &= \mathbb{E}_{\text{in}}[g_0(\mathbf{x}) g_0(\mathbf{x}')] = \int_D \varphi(\mathbf{x}, \boldsymbol{\theta}) \varphi(\mathbf{x}', \boldsymbol{\theta}) \mu_{\text{in}}(d\boldsymbol{\theta}) \\ &\quad - \int_{D \times D} \varphi(\mathbf{x}, \boldsymbol{\theta}) \varphi(\mathbf{x}', \boldsymbol{\theta}') \mu_{\text{in}}(d\boldsymbol{\theta}) \mu_{\text{in}}(d\boldsymbol{\theta}') \end{aligned} \quad (83)$$

Since (82) is linear the mean of g_t remains zero at all times and we can again focus on the evolution of its covariance:

$$C_t(\mathbf{x}, \mathbf{x}') = \mathbb{E}_{\text{in}}[g_t(\mathbf{x}) g_t(\mathbf{x}')] \quad (84)$$

We obtain

$$\begin{aligned} \partial_t C_t &= - \int_{\Omega} N(\mathbf{x}, \mathbf{x}', \mathbf{x}'', [\Sigma_t]) (f_t(\mathbf{x}'') - f(\mathbf{x}'')) \nu(d\mathbf{x}'') \\ &\quad - \int_{\Omega} M(\mathbf{x}, \mathbf{x}'', [\mu_t]) C_t(\mathbf{x}'', \mathbf{x}) \nu(d\mathbf{x}'') \\ &\quad - \int_{\Omega} M(\mathbf{x}', \mathbf{x}'', [\mu_t]) C_t(\mathbf{x}'', \mathbf{x}') \nu(d\mathbf{x}'') \end{aligned} \quad (85)$$

where Σ_t solves (78) and

$$\begin{aligned} N(\mathbf{x}, \mathbf{x}', \mathbf{x}'', [\Sigma]) &= \int_{D \times D} \nabla_{\boldsymbol{\theta}} \varphi(\mathbf{x}, \boldsymbol{\theta}) \cdot \nabla_{\boldsymbol{\theta}} \varphi(\mathbf{x}'', \boldsymbol{\theta}) \varphi(\mathbf{x}', \boldsymbol{\theta}') \Sigma(d\boldsymbol{\theta}, d\boldsymbol{\theta}') \\ &\quad + \int_{D \times D} \nabla_{\boldsymbol{\theta}} \varphi(\mathbf{x}', \boldsymbol{\theta}) \cdot \nabla_{\boldsymbol{\theta}} \varphi(\mathbf{x}'', \boldsymbol{\theta}) \varphi(\mathbf{x}, \boldsymbol{\theta}') \Sigma(d\boldsymbol{\theta}, d\boldsymbol{\theta}') \end{aligned} \quad (86)$$

Summarizing, we have established:

Proposition 3.7 (CLT). *Let $g_t^{(n)}$ be given by (81) with $\{\boldsymbol{\theta}_i(t)\}_{i=1}^n$ solution of (29) with initial conditions draw from \mathbb{P}_{in} and μ_t solution to (43). Then, as $n \rightarrow \infty$, $g_t^{(n)} \rightarrow g_t$ in law, where g_t is the zero mean Gaussian process whose covariance solves to (85) for the initial condition (83).*

3.5. Scaling of the fluctuations at long and very long times. To analyze the behavior of the fluctuations as $t \rightarrow \infty$, we revisit the results from the last section from a different perspective. Suppose that, instead of (69) and (81), we would consider

$$\bar{\omega}_t^{(n)} = n^{-1/2} \sum_{i=1}^n (\delta_{\boldsymbol{\theta}_i(t)} - \mu_t) \quad (87)$$

and

$$\bar{g}_t^{(n)} = n^{-1/2} \sum_{i=1}^n \left(\varphi(\cdot, \boldsymbol{\theta}_i(t)) - \int_D \varphi(\cdot, \boldsymbol{\theta}) \mu_t(d\boldsymbol{\theta}) \right) \quad (88)$$

where $\Theta_i(t)$ are independent copies of the mean-field characteristic equation (40). Then, direct calculations show that $\bar{\omega}_t^{(n)} \rightarrow \bar{\omega}_t$ and $\bar{g}_t^{(n)} \rightarrow \bar{g}_t$ in law as $n \rightarrow \infty$, where $\bar{\omega}_t$ and \bar{g}_t are Gaussian processes with mean zero and covariance given explicitly by

$$(89) \quad \bar{\Sigma}_t(d\boldsymbol{\theta}, d\boldsymbol{\theta}') = \mathbb{E}_{\text{in}}[\bar{\omega}_t(d\boldsymbol{\theta})\bar{\omega}_t(d\boldsymbol{\theta}')] = \mu_t(d\boldsymbol{\theta})\delta_{\boldsymbol{\theta}}(d\boldsymbol{\theta}') - \mu_t(d\boldsymbol{\theta})\mu_t(d\boldsymbol{\theta}')$$

and

$$(90) \quad \bar{C}_t(\mathbf{x}, \mathbf{x}') = \mathbb{E}_{\text{in}}[\bar{g}_t(\mathbf{x})\bar{g}_t(\mathbf{x}')] = \int_D \varphi(\mathbf{x}, \boldsymbol{\theta})\varphi(\mathbf{x}', \boldsymbol{\theta})\mu_t(d\boldsymbol{\theta}) - f_t(\mathbf{x})f_t(\mathbf{x}')$$

We can also easily write down evolution equations for $\bar{\omega}_t$ and \bar{g}_t : they read

$$(91) \quad \partial_t \bar{\omega}_t = \nabla \cdot (\bar{\omega}_t \nabla V(\boldsymbol{\theta}, [\mu_t]))$$

and

$$(92) \quad \partial_t \bar{g}_t = - \int_{\Omega} M(\mathbf{x}, \mathbf{x}', [\bar{\omega}_t]) (f_t(\mathbf{x}') - f(\mathbf{x}')) \nu(d\mathbf{x}')$$

Let us focus on this last equation: it is similar to (82), but without the last term, $-\int_{\Omega} M(\mathbf{x}, \mathbf{x}', [\mu_t])g_t(\mathbf{x}')\nu(d\mathbf{x}')$. Since the kernel M is positive semi-definite, we know that the solutions to (82) are controlled by those of (92). In particular,

$$(93) \quad \mathbb{E}_{\text{in}} \int_{\Omega} |g_t(\mathbf{x})|^2 \nu(d\mathbf{x}) = \int_{\Omega} C_t(\mathbf{x}, \mathbf{x}) \nu(d\mathbf{x}) \leq \int_{\Omega} \bar{C}_t(\mathbf{x}, \mathbf{x}) \nu(d\mathbf{x})$$

If we assume that $\mu_t \rightarrow \mu^* \in \mathcal{M}_+(D)$ as $t \rightarrow \infty$, from (90) we have

$$(94) \quad \lim_{t \rightarrow \infty} \int_{\Omega} \bar{C}_t(\mathbf{x}, \mathbf{x}) \nu(d\mathbf{x}) = \int_D K(\boldsymbol{\theta}, \boldsymbol{\theta}) \mu^*(d\boldsymbol{\theta}) - \int_{\Omega} |f(\mathbf{x})|^2 \nu(d\mathbf{x})$$

and therefore

$$(95) \quad \lim_{t \rightarrow \infty} \int_{\Omega} C_t(\mathbf{x}, \mathbf{x}) \nu(d\mathbf{x}) \leq \int_D K(\boldsymbol{\theta}, \boldsymbol{\theta}) \mu^*(d\boldsymbol{\theta}) - \int_{\Omega} |f(\mathbf{x})|^2 \nu(d\mathbf{x}).$$

Because

$$(96) \quad \int_{\Omega} C_t(\mathbf{x}, \mathbf{x}) \nu(d\mathbf{x}) = \lim_{n \rightarrow \infty} n \mathbb{E}_{\text{in}} \int_{\Omega} |f_t^{(n)}(\mathbf{x}) - f_t(\mathbf{x})|^2 \nu(d\mathbf{x})$$

the previous result gives a Monte-Carlo type error bound on the loss. Note that this bound is only nontrivial if

$$(97) \quad \begin{aligned} \int_D K(\boldsymbol{\theta}, \boldsymbol{\theta}) \mu^*(d\boldsymbol{\theta}) &= \int_{\mathbb{R} \times \hat{D}} c^2 \hat{K}(\mathbf{z}, \mathbf{z}) \mu^*(dc, d\mathbf{z}) \\ &\leq \|\hat{K}\|_{\infty} \int_{\mathbb{R} \times \hat{D}} c^2 \mu^*(dc, d\mathbf{z}) < \infty \end{aligned}$$

During training, we have $\int_{\mathbb{R} \times \hat{D}} c^2 \mu_t(dc, d\mathbf{z}) < \infty$ for all $t < \infty$, and to guarantee that this moment does not blow up as $t \rightarrow \infty$, or more generally to control its value in that limit, we may need to add a regularizing term to the loss function. If (97) holds, there is a situation in which we can even deduce a better bound: if $\text{supp } \hat{\mu}^* = \hat{D}$ and, then $M(\mathbf{x}, \mathbf{x}, [\mu_t])$ is positive definite for all $t \geq 0$ and in the limit $t \rightarrow \infty$, indicating that the last term in (82) is always dissipative. In this case the argument above shows that

$$(98) \quad \lim_{t \rightarrow \infty} \int_{\Omega} C_t(\mathbf{x}, \mathbf{x}) \nu(d\mathbf{x}) = 0.$$

Summarizing:

Proposition 3.8 (Fluctuations at long times). *Let $f_t^{(n)}$ be given by with $\{\theta_i(t)\}_{i=1}^n$ solution of (32) with initial conditions draw from \mathbb{P}_{in} and f_t solution to (45). Then, under the conditions of Proposition 3.4 and assuming that (97) holds, we have*

$$(99) \quad \begin{aligned} & \lim_{t \rightarrow \infty} \lim_{n \rightarrow \infty} n \mathbb{E}_{in} \int_{\Omega} |f_t^{(n)}(\mathbf{x}) - f_t(\mathbf{x})|^2 v(d\mathbf{x}) \\ & \leq \int_D K(\boldsymbol{\theta}, \boldsymbol{\theta}) \mu^*(d\boldsymbol{\theta}) - \int_{\Omega} |f(\mathbf{x})|^2 v(d\mathbf{x}) \end{aligned}$$

In addition, if $\text{supp } \hat{\mu}^* = \hat{D}$, we have

$$(100) \quad \lim_{t \rightarrow \infty} \lim_{n \rightarrow \infty} n \mathbb{E}_{in} \int_{\Omega} |f_t^{(n)}(\mathbf{x}) - f_t(\mathbf{x})|^2 v(d\mathbf{x}) = 0.$$

In situations where $\text{supp } \hat{\mu}^* = \hat{D}$ and (100) holds, we see that the fluctuations, initially detectable on the scale $n^{-1/2}$, become higher order as time increases. To understand the scale at which the fluctuations eventually settle, consider

$$(101) \quad \tilde{\omega}_t^{(n)}(d\boldsymbol{\theta}) = n^{\xi(t)} \sum_{i=1}^n (\delta_{\theta_i(t)}(\boldsymbol{\theta}) - \mu_t(d\boldsymbol{\theta}))$$

where $\xi(t)$ is some time-dependent exponent to be specified. By proceeding as we did to derive (72), we have that $\tilde{\omega}_t^{(n)}$ satisfies

$$(102) \quad \begin{aligned} \partial_t \tilde{\omega}_t^{(n)} &= \nabla \cdot \left(\tilde{\omega}_t^{(n)} \nabla V(\boldsymbol{\theta}, [\mu_t]) + \mu_t \nabla F(\boldsymbol{\theta}, [\tilde{\omega}_t^{(n)}]) \right) \\ &+ n^{-\xi(t)} \nabla \cdot \left(\omega_t^{(n)} \nabla F(\boldsymbol{\theta}, [\tilde{\omega}_t^{(n)}]) \right) + \dot{\xi}(t) \log n \tilde{\omega}_t^{(n)}. \end{aligned}$$

In order to take the limit as $n \rightarrow \infty$ of this equation, we need to consider carefully the behavior of the factors in (102) that contain n explicitly, that is, $\omega_t^{(n)} \nabla F(\boldsymbol{\theta}, [\tilde{\omega}_t^{(n)}])$ and $\dot{\xi}(t) \log n \tilde{\omega}_t^{(n)}$. Regarding the former, for any $p \in \mathbb{N}$ and $\xi \in \mathbb{R}$,

$$(103) \quad \mathbb{E}_{in} \left(n^{-\xi} \int_{D \times D} \chi(\boldsymbol{\theta}) \chi(\boldsymbol{\theta}') \tilde{\omega}_0^{(n)}(d\boldsymbol{\theta}) \tilde{\omega}_0^{(n)}(d\boldsymbol{\theta}') \right)^p = O\left(n^{(\xi-1)p}\right),$$

which can be verified by a direct calculation. For example if $p = 1$, this expectation is $n^{(\xi-1)} C_{\chi}$ where C_{χ} is given in (71). Equation (103) implies that $n^{-\xi} \tilde{\omega}_0^{(n)}(d\boldsymbol{\theta}) \tilde{\omega}_0^{(n)}(d\boldsymbol{\theta}') \rightarrow 0$ weakly in L^{2p} at $t = 0$ for any $\xi < 1$. To see whether we can bring the fluctuations to that scale, notice that if we set

$$(104) \quad \dot{\xi}(t) \log n = o(1)$$

the last term at the right hand side of (102) is also higher order—(104) means that we can vary $\xi(t)$, but only slowly. (104) can be achieved by choosing e.g.

$$(105) \quad \xi(t) = \bar{\xi}(t/a_n)$$

with $\bar{\xi}(0) = \frac{1}{2}$, $\bar{\xi}'(u) > 0$, $\lim_{u \rightarrow \infty} \bar{\xi}(u) = < 1$, and a_n growing with n and such that $\lim_{n \rightarrow \infty} a_n / \log n = \infty$. With this choice, both the last two terms at the right hand of (102) are a small perturbation that vanishes as $n \rightarrow \infty$. Therefore, if we test $\tilde{\omega}_t^{(n)}$ against the unit, and define

$$(106) \quad \tilde{g}_t^{(n)} = n^{\xi(t)} (f_t^{(n)} - f_t) = \int_D \varphi(\cdot, \boldsymbol{\theta}) \tilde{\omega}_t^{(n)}(d\boldsymbol{\theta})$$

we know that, if $\xi(t)$ is as in (105) and $\text{supp } \hat{\mu}^* = \hat{D}$, this field will be controlled and go to zero eventually. Summarizing we have established:

Proposition 3.9 (Fluctuations at very long times). *Assume that the conditions of Proposition 3.8 hold and $\text{supp } \hat{\mu}^* = \hat{D}$. Then*

$$(107) \quad \forall \xi < 1 \quad : \quad \lim_{n \rightarrow \infty} n^{2\xi} \mathbb{E}_{in} \int_{\Omega} |f_{a_n}^{(n)}(\mathbf{x}) - f(\mathbf{x})|^2 v(d\mathbf{x}) = 0$$

if a_n grows with n and is such that $\lim_{n \rightarrow \infty} a_n / \log n = \infty$.

This proposition can be stated as (11). It shows a remarkable self-healing property of the dynamics: the fluctuations at scale $O(n^{-1/2})$ of $f_t^{(n)}$ around f_t that were present initially decrease in amplitude as time progresses, and become $O(n^{-1})$ or smaller as $t \rightarrow \infty$.

4. TRAINING BY ONLINE STOCHASTIC GRADIENT DESCENT

In most applications, it is not possible to evaluate the expectation over the data in (6) defining $\hat{F}(\mathbf{z})$ and $\hat{K}(\mathbf{z}, \mathbf{z}')$. This is especially true for $\hat{F}(\mathbf{z})$, since we typically have limited access to $f(\mathbf{x})$: often, we can only evaluate it pointwise or only know its value on a discrete set of points. In these cases, we typically need to approximate the expectation in (6) by a sum over a finite subset of \mathbf{x} 's obtained by sampling from the measure ν .

If we were to fix this training data set, $\{\mathbf{x}_p\}_{p=1}^P$, and denote by $\nu_P = P^{-1} \sum_{p=1}^P \delta_{\mathbf{x}_p}$ the corresponding empirical measure, then all the results in Sec. 3 apply at empirical level if we replace everywhere ν by ν_P . This, however, is not the question we are typically interested in, which is rather:

How does the test error (that is, the error obtained using the exact loss defined with the original ν) scale if we train the network on the empirical loss associated to ν_P ?

Here we will address this question in the specific setting of ‘‘online’’ learning algorithms, in which we can draw a training data set of batch size P at every step of the learning. This effectively assumes that we have access to infinite data, but cannot use it all at the same time, and the finite size of the batch introduces noise into the learning algorithm. The algorithm in which the gradient is estimated from a subset of training data at each step is known as stochastic gradient descent. It reads

$$(108) \quad \hat{\boldsymbol{\theta}}_i(t + \Delta t) = \hat{\boldsymbol{\theta}}_i(t) + \nabla F_P(t, \hat{\boldsymbol{\theta}}_i^P(t)) \Delta t - \frac{1}{n} \sum_{j=1}^n \nabla K_P(t, \hat{\boldsymbol{\theta}}_i(t), \hat{\boldsymbol{\theta}}_j(t)) \Delta t$$

where $i = 1, \dots, n$, $\Delta t > 0$ is some time-step, and we defined

$$(109) \quad \begin{aligned} F_P(t, \boldsymbol{\theta}) &= \frac{1}{P} \sum_{p=1}^P f(\mathbf{x}_p(t)) \varphi(\mathbf{x}_p(t), \boldsymbol{\theta}), \\ K_P(t, \boldsymbol{\theta}, \boldsymbol{\theta}') &= \frac{1}{P} \sum_{p=1}^P \varphi(\mathbf{x}_p(t), \boldsymbol{\theta}) \varphi(\mathbf{x}_p(t), \boldsymbol{\theta}') \end{aligned}$$

in which $\{\mathbf{x}_p(t)\}_{p=1}^P$ are P iid variables which are redrawn from ν independently at every time step t . Next we analyze how the result from Sec. 3 must be modified when we use (109) rather than (29) to perform the training.

4.1. Limiting stochastic differential equation. To analyze the properties of (108), we start by noticing that the term

$$(110) \quad \mathbf{R}_i(\vec{\boldsymbol{\theta}}) = \nabla F_P(t, \boldsymbol{\theta}_i) - \frac{1}{n} \sum_{j=1}^n \nabla K_P(t, \boldsymbol{\theta}_i, \boldsymbol{\theta}_j), \quad \vec{\boldsymbol{\theta}} = (\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_n)$$

is an unbiased estimator of the right hand side of the GD equation (29). Indeed, conditional on $\{\boldsymbol{\theta}_i\}_{i=1}^n$ fixed, we have

$$(111) \quad \mathbb{E}_\nu \mathbf{R}_i(\vec{\boldsymbol{\theta}}) = \nabla F(\boldsymbol{\theta}_i) - \frac{1}{n} \sum_{j=1}^n \nabla K(\boldsymbol{\theta}_i, \boldsymbol{\theta}_j)$$

This means that, if we split the right hand side of (108) into its expectation plus a zero-mean fluctuation, the expression resembles an Euler-Maruyama scheme for a stochastic differential equation (SDE), except that the scaling of the noise term involves Δt rather than $\sqrt{\Delta t}$. To write this SDE explicitly, we compute the covariance of $\mathbf{R}(\vec{\boldsymbol{\theta}})$ conditional on $\{\boldsymbol{\theta}_i\}_{i=1}^n$ fixed,

$$(112) \quad \text{cov}_\nu[\mathbf{R}_i(\vec{\boldsymbol{\theta}}), \mathbf{R}_j(\vec{\boldsymbol{\theta}}')] = A([f - f^{(n)}], \boldsymbol{\theta}_i, \boldsymbol{\theta}'_j)$$

where $f^{(n)} = n^{-1} \sum_{i=1}^n \varphi(\cdot, \boldsymbol{\theta}_i)$ and we defined

$$(113) \quad \begin{aligned} A([f], \boldsymbol{\theta}, \boldsymbol{\theta}') &= \mathbb{E}_v[|f|^2 \nabla_{\boldsymbol{\theta}} \varphi(\cdot, \boldsymbol{\theta}) \otimes \nabla_{\boldsymbol{\theta}'} \varphi(\cdot, \boldsymbol{\theta}')] \\ &\quad - \mathbb{E}_v[f \nabla_{\boldsymbol{\theta}} \varphi(\cdot, \boldsymbol{\theta})] \otimes \mathbb{E}_v[f \nabla_{\boldsymbol{\theta}'} \varphi(\cdot, \boldsymbol{\theta}')] \end{aligned}$$

The SDE capturing the behavior of the solution to (108) is

$$(114) \quad d\boldsymbol{\theta}_i = \nabla F(\boldsymbol{\theta}_i) dt - \frac{1}{n} \sum_{j=1}^n \nabla K(\boldsymbol{\theta}_i, \boldsymbol{\theta}_j) dt + \sqrt{\sigma} d\mathbf{B}_i,$$

where $\sigma = \Delta t / P$ and $\{d\mathbf{B}_i\}_{i=1}^n$ is a white-noise process with quadratic variation

$$(115) \quad \langle d\mathbf{B}_i, d\mathbf{B}_j \rangle = A([f - f^{(n)}], \boldsymbol{\theta}_i, \boldsymbol{\theta}_j) dt.$$

More precisely [LTE15, HLLL17],

Lemma 4.1. *Given any test functions $\chi : D \rightarrow \mathbb{R}$ and any $T > 0$, there is a constant $C > 0$ such that*

$$(116) \quad \sup_{0 \leq k\Delta t \leq T} \left| \frac{1}{n} \sum_{i=1}^n (\mathbb{E} \chi(\hat{\boldsymbol{\theta}}_i(k\Delta t)) - \mathbb{E} \chi(\boldsymbol{\theta}_i(k\Delta t))) \right| \leq C \Delta t.$$

where $\hat{\boldsymbol{\theta}}_i(t)$ and $\boldsymbol{\theta}_i(t)$ denote the solutions to (108) and (120), respectively.

This lemma is a direct consequence of the fact that (108) can be viewed as the Euler-Maruyama discretization scheme for (114), and this scheme has weak order of accuracy 1. Note that if we let $\Delta t \rightarrow 0$, (114) reduces to the ODEs in (29) since $\sigma = \Delta t / P \rightarrow 0$ in that limit. We should stress, however, that this limit is not reached in practice since the scheme (108) is used at small but finite Δt . We analyze next what happens when we adjust the size of σ by changing Δt and/or the batch size P .

4.2. Dean's equation for particles with correlated noise. Lemma 4.1 indicates that we can analyze the properties of (114) instead of those of (108). To this end, we derive an equation for the empirical distribution $\mu_t^{(n)}$ in (31) when $\{\boldsymbol{\theta}_i(t)\}_{i=1}^n$ satisfy the SDE (114); this calculation is operationally similar to the derivation of (33) but takes into account the extra drift term and the noise term in (114). By applying Itô's formula to (35) we deduce

$$(117) \quad \begin{aligned} d \int_D \chi(\boldsymbol{\theta}) \mu_t^{(n)}(d\boldsymbol{\theta}) &= \frac{1}{n} \sum_{i=1}^n \nabla \chi(\boldsymbol{\theta}_i(t)) \cdot d\boldsymbol{\theta}_i(t) \\ &\quad + \frac{\sigma}{2n} \sum_{i=1}^n \nabla \nabla \chi(\boldsymbol{\theta}_i(t)) : A([f - f_t^{(n)}], \boldsymbol{\theta}_i(t), \boldsymbol{\theta}_i(t)) dt \end{aligned}$$

where $f_t^{(n)} = n^{-1} \sum_{i=1}^n \varphi(\cdot, \boldsymbol{\theta}_i(t)) = \int_D \varphi(\cdot, \boldsymbol{\theta}) \mu_t^{(n)}(d\boldsymbol{\theta})$. Using (114) and the definition of $\mu_t^{(n)}$, this relation can be written as

$$(118) \quad \begin{aligned} d \int_D \chi(\boldsymbol{\theta}) \mu_t^{(n)}(d\boldsymbol{\theta}) &= \int_D \nabla \chi(\boldsymbol{\theta}) \cdot \nabla V(\boldsymbol{\theta}, [\mu_t^{(n)}]) \mu_t^{(n)}(d\boldsymbol{\theta}) dt \\ &\quad + \frac{\sigma}{2} \int_D \nabla \nabla \chi(\boldsymbol{\theta}) : A([f - f_t^{(n)}], \boldsymbol{\theta}, \boldsymbol{\theta}) \mu_t^{(n)}(d\boldsymbol{\theta}) dt \\ &\quad + \frac{\sqrt{\sigma}}{n} \sum_{i=1}^n \nabla \chi(\boldsymbol{\theta}_i(t)) \cdot d\mathbf{B}_i(t) \end{aligned}$$

The drift terms in this equation are expressed in term of $\mu_t^{(n)}$; for the noise term, notice that its quadratic variation is

$$(119) \quad \begin{aligned} &\left\langle \frac{\sqrt{\sigma}}{n} \sum_{i=1}^n \nabla \chi(\boldsymbol{\theta}_i(t)) \cdot d\mathbf{B}_i(t), \frac{\sqrt{\sigma}}{n} \sum_{i=1}^n \nabla \chi(\boldsymbol{\theta}_i(t)) \cdot d\mathbf{B}_i(t) \right\rangle \\ &= \sigma \int_{D \times D} \nabla \chi(\boldsymbol{\theta}) \nabla \chi(\boldsymbol{\theta}') : A([f - f_t^{(n)}], \boldsymbol{\theta}, \boldsymbol{\theta}') \mu_t^{(n)}(d\boldsymbol{\theta}) \mu_t^{(n)}(d\boldsymbol{\theta}') dt \end{aligned}$$

This means that, in law, (118) is equivalent to

$$(120) \quad \begin{aligned} d \int_D \chi(\boldsymbol{\theta}) \mu_t^{(n)}(d\boldsymbol{\theta}) &= \int_D \nabla \chi(\boldsymbol{\theta}) \cdot \nabla V(\boldsymbol{\theta}, [\mu_t^{(n)}]) \mu_t^{(n)}(d\boldsymbol{\theta}) dt \\ &+ \frac{\sigma}{2} \int_D \nabla \nabla \chi(\boldsymbol{\theta}) : A([f - f_t^{(n)}], \boldsymbol{\theta}, \boldsymbol{\theta}) \mu_t^{(n)}(d\boldsymbol{\theta}) dt \\ &+ \sqrt{\sigma} \int_D \nabla \chi(\boldsymbol{\theta}) \cdot d\boldsymbol{\eta}_t^{(n)}(d\boldsymbol{\theta}) \end{aligned}$$

where $d\boldsymbol{\eta}_t^{(n)}(d\boldsymbol{\theta})$ is vector-valued random measure, white in time, and with quadratic variation

$$(121) \quad \langle d\boldsymbol{\eta}_t^{(n)}(d\boldsymbol{\theta}), d\boldsymbol{\eta}_t^{(n)}(d\boldsymbol{\theta}') \rangle = A([f - f_t^{(n)}], \boldsymbol{\theta}, \boldsymbol{\theta}') \mu_t^{(n)}(d\boldsymbol{\theta}) \mu_t^{(n)}(d\boldsymbol{\theta}') dt$$

The first term at the right hand side of (120) is the same as in the weak form of (33). This is because these terms come from the drift terms in (114), which also coincide with those in (29). However, (120) also contains additional terms that were absent in (33)—note that these terms are different from those in the standard Dean's equation, because the noise term in (114) is correlated between the particles, instead of being independent.

4.3. LLN for SGD. If we want the result established in Proposition 3.5 to apply and also for the approximation error to vanish as $n \rightarrow \infty$, we need to make the additional terms in (120) compared to (33) higher order. This can be done by scaling σ with some inverse power of n . Specifically, we will set

$$(122) \quad \sigma = an^{-2\alpha} \quad \text{for some } a > 0 \quad \text{and } \alpha > 0$$

This scaling can be achieved by choosing, e.g., $P = O(n^{2\alpha})$, which amounts to increasing the batch size with n . The choice (122) implies that the last two terms in (120) disappear in the limit as $n \rightarrow \infty$. Therefore, we formally conclude that $\mu_t^{(n)} \rightarrow \mu_t$ as $n \rightarrow \infty$, where μ_t solves the same deterministic equation (37) as before. This implies that $\lim_{n \rightarrow \infty} f_t^{(n)} = f_t = \int_D \varphi(\cdot, \boldsymbol{\theta}) \mu_t(d\boldsymbol{\theta})$ satisfies (45) and is such that $f_t \rightarrow f$ as $t \rightarrow \infty$. In particular, both the LLN and the global convergence result in Proposition 3.5 still hold if the assumption in this proposition are met and we use the solution of (127) in (32). In turn, we can also conclude that this proposition holds up to discretization errors in Δt if we use the solution of (108) in (32). Importantly, the covariance associated with the estimator for the gradient, defined in (113), satisfies

$$(123) \quad \forall (\boldsymbol{\theta}, \boldsymbol{\theta}') \in D \times D : \quad \lim_{t \rightarrow \infty} A([f_t - f], \boldsymbol{\theta}, \boldsymbol{\theta}') = 0.$$

This property will be useful later.

4.4. CLT for SGD. Turning our attention to the fluctuations of $\mu_t^{(n)}$ around μ_t , notice that there are two sources of them: some are intrinsic to the discrete nature of the particles apparent in $\mu_t^{(n)}$, and scale as $O(n^{-1/2})$ for all $t < \infty$ and possibly as $O(n^{-\xi})$ for any $\xi < 1$ as $t \rightarrow \infty$, as discussed in Sec. 3.5. Other fluctuations come from the noise term in (120), and scale as $O(n^{-\alpha})$ when (122) holds. The Itô drift terms proportional to $\sigma = an^{-2\alpha}$ in (120) always make higher order contributions.

We first consider $t < \infty$ and subsequently examine the limit $t \rightarrow \infty$ in Sec. 4.5. In the present case, we first observe that if $\alpha \geq \frac{1}{2}$, then for all $t < \infty$ the fluctuations due to the noise in (120) are negligible compared to the intrinsic ones from discreteness, and we are back to the GD situation studied in Sec. 3. In contrast, if $\alpha \in (0, \frac{1}{2})$, for all $t < \infty$ the fluctuations due to the noise in (120) dominate the intrinsic ones from discreteness, so let us focus on this case from now on. To quantify these fluctuations, we can introduce $n^\alpha(\mu_t^{(n)} - \mu_t)$, write an equation for this scaled discrepancy, and take the limit as $n \rightarrow \infty$. The derivation proceeds akin to the derivation of (75)

and leads to the conclusion that, as $n \rightarrow \infty$, $n^\alpha(\mu_t^{(n)} - \mu_t) \rightarrow \omega_t^{(\alpha)}$ in law which satisfies

$$(124) \quad \begin{aligned} d \int_D \chi(\boldsymbol{\theta}) \omega_t^{(\alpha)}(d\boldsymbol{\theta}) &= \int_D \nabla \chi(\boldsymbol{\theta}) \cdot \nabla V(\boldsymbol{\theta}, [\mu_t]) \omega_t^{(\alpha)}(d\boldsymbol{\theta}) dt \\ &+ \int_D \nabla \chi(\boldsymbol{\theta}) \cdot \nabla F(\boldsymbol{\theta}, [\mu_t^{(\alpha)}]) \mu_t(d\boldsymbol{\theta}) dt \\ &+ \sqrt{a} \int_D \nabla \chi(\boldsymbol{\theta}) \cdot d\boldsymbol{\eta}_t(d\boldsymbol{\theta}) \end{aligned}$$

in which $d\boldsymbol{\eta}_t(d\boldsymbol{\theta})$ is vector valued random measure, white in time, and with quadratic variation (compare (121))

$$(125) \quad \langle d\boldsymbol{\eta}_t(d\boldsymbol{\theta}), d\boldsymbol{\eta}_t(d\boldsymbol{\theta}') \rangle = A([f - f_t], \boldsymbol{\theta}, \boldsymbol{\theta}') \mu_t(d\boldsymbol{\theta}) \mu_t(d\boldsymbol{\theta}') dt$$

Equation (124) should be solved with zero initial condition, since the $O(n^{-1/2})$ fluctuations arising from the initial condition are higher order compared to scaling $O(n^{-\alpha})$ we picked to obtain (124). Since (124) is linear in $\omega_t^{(\alpha)}$ with additive noise, it indicates that $\omega_t^{(\alpha)}$ a Gaussian process with mean zero and thereby fully characterized by its covariance (we omit the equation for brevity). This also implies that

$$(126) \quad n^\alpha(f_t^{(n)} - f) \rightarrow g_t^{(\alpha)} \quad \text{in law as } n \rightarrow \infty$$

where $g_t^{(\alpha)}$ is a Gaussian process whose evolution equation (cf. the derivation of (82)) gives,

$$(127) \quad \begin{aligned} dg_t^{(\alpha)} &= - \int_\Omega M([\omega_t^\alpha], \mathbf{x}, \mathbf{x}') (f_t(\mathbf{x}) - f(\mathbf{x}')) v(d\mathbf{x}') dt \\ &- \int_\Omega M([\mu_t], \mathbf{x}, \mathbf{x}') g_t^{(\alpha)}(\mathbf{x}') v(d\mathbf{x}') dt + \sqrt{a} d\zeta_t(\mathbf{x}) \end{aligned}$$

where $M([\mu], \mathbf{x}, \mathbf{x}')$ is given in (46), and the quadratic variation of $d\zeta_t$ is that of $\int_D \varphi(\cdot, \boldsymbol{\theta}) d\boldsymbol{\eta}_t(d\boldsymbol{\theta})$. Explicitly,

$$(128) \quad \begin{aligned} \langle d\zeta_t(\mathbf{x}), d\zeta_t(\mathbf{x}') \rangle &= \int_\Omega M([\mu_t], \mathbf{x}, \mathbf{x}'') M([\mu_t], \mathbf{x}', \mathbf{x}'') |f_t(\mathbf{x}'') - f(\mathbf{x}'')|^2 dv(\mathbf{x}'') dt \\ &- \int_\Omega M([\mu_t], \mathbf{x}, \mathbf{x}'') (f_t(\mathbf{x}'') - f(\mathbf{x}'')) dv(\mathbf{x}'') \\ &\times \int_\Omega M([\mu_t], \mathbf{x}', \mathbf{x}'') (f_t(\mathbf{x}'') - f(\mathbf{x}'')) dv(\mathbf{x}'') dt. \end{aligned}$$

The SDE (127) should be solved with zero initial condition, $g_0^{(\alpha)} = 0$. Since it is linear in $g_t^{(\alpha)}$ with additive noise, it defines a Gaussian process with mean zero and is specified by its covariance

$$(129) \quad C_t^{(\alpha)}(\mathbf{x}, \mathbf{x}') = \mathbb{E}[g_t^{(\alpha)}(\mathbf{x}) g_t^{(\alpha)}(\mathbf{x}')]$$

where \mathbb{E} denotes expectation over the noise $d\zeta_t$ (that is, over the data in the batches used in SGD). With this calculation, we have established

Proposition 4.2 (CLT for SGD). *Consider*

$$(130) \quad g_t^{(\alpha, n)} = n^{\alpha-1} \sum_{i=1}^n (\varphi(\cdot, \boldsymbol{\theta}_i(t)) - f_t) = n^\alpha (f_t^{(n)} - f_t)$$

with $\{\boldsymbol{\theta}_i(t)\}_{i=1}^n$ solution to the SDE (114) with $\sigma = an^{-2\alpha}$, $\alpha \in (0, \frac{1}{2})$, and f_t solution to (45). Then, as $n \rightarrow \infty$, $g_t^{(\alpha, n)}$ converges in law towards the Gaussian process $g_t^{(\alpha)}$ solution of (127) for $g_0^{(\alpha)} = 0$.

4.5. Fluctuations in SGD at long and very long times. The noise $d\zeta_t$ in (127) has the remarkable property that it self-quenches as $t \rightarrow \infty$ if the conditions of Proposition 3.5 are met and $f_t \rightarrow f$ as $t \rightarrow \infty$ and therefore, from (128):

$$(131) \quad \forall \mathbf{x}, \mathbf{x}' \in \Omega : \quad \lim_{t \rightarrow \infty} \langle d\zeta_t(\mathbf{x}), d\zeta_t(\mathbf{x}') \rangle = 0.$$

Since the first drift term in (127) also goes to zero when $f_t \rightarrow f$ and the second drift term is a damping term because $M([\mu_t], \mathbf{x}, \mathbf{x}')$ is positive definite for all $t < \infty$, we know that $g_t^{(\alpha)}$ will be controlled as $t \rightarrow \infty$, i.e. $C_t^{(\alpha)}$ as a limit. In addition, if $\text{supp } \hat{\mu}^* = \hat{D}$ where $\hat{\mu}^* = \int_{\mathbb{R}} \mu^*(dc, \cdot) = \lim_{t \rightarrow \infty} \int_{\mathbb{R}} \mu^*(dc, \cdot)$, then $M([\mu_t], \mathbf{x}, \mathbf{x}')$ is positive definite for all $t < \infty$ and in the limit as $t \rightarrow \infty$, and the solution to (127) goes to zero. Using the definition the Gaussian process $g_t^{(n, \alpha)}$ defined in (130), we can summarize this result into:

Proposition 4.3 (Fluctuations in SGD at long time). *Under the conditions of Proposition 3.5, if $f_t^{(n)}$ is given by (32) with $\{\theta_i(t)\}_{i=1}^n$ solution of (114) with $\sigma = an^{-2\alpha}$, $\alpha \in (0, \frac{1}{2})$, and initial condition drawn from \mathbb{P}_{in} , and f_t solves (45), then*

$$(132) \quad \lim_{t \rightarrow \infty} \lim_{n \rightarrow \infty} n^{2\alpha} \mathbb{E} \int_{\Omega} |f_t^{(n)}(\mathbf{x}) - f_t(\mathbf{x})|^2 \nu(d\mathbf{x}) = \lim_{t \rightarrow \infty} C_t^{(\alpha)}(\mathbf{x}, \mathbf{x}') \quad \text{exists}$$

In addition, if $\text{supp } \hat{\mu}^ = \hat{D}$, then this limit is zero.*

If $\text{supp } \hat{\mu}^* = \hat{D}$ where $\hat{\mu}^* = \int_{\mathbb{R}} \mu^*(dc, \cdot) = \lim_{t \rightarrow \infty} \int_{\mathbb{R}} \mu^*(dc, \cdot)$, then $M([\mu_t], \mathbf{x}, \mathbf{x}')$ is positive definite for all $t < \infty$ and in the limit as $t \rightarrow \infty$. In that case, the only fixed point of (127) is zero. Since in this case we also know that the fluctuations from the initial conditions disappear on scale $O(n^\xi)$ for any $\xi < 0$, we can proceed as in Sec. 3.5 and adjust α all the way up to 1 instead of $\frac{1}{2}$. That is, we can generalize Proposition 3.9 into

Proposition 4.4 (Fluctuations in SGD at very long times). *Under the conditions of Proposition 4.3, if $\text{supp } \hat{\mu}^* = \hat{D}$, then for any $\alpha \in (0, 1)$,*

$$(133) \quad \lim_{n \rightarrow \infty} n^{2\alpha} \mathbb{E} \int_{\Omega} |f_{a_n}^{(n)}(\mathbf{x}) - f(\mathbf{x})|^2 \nu(d\mathbf{x}) = 0$$

if a_n grows with n and is such that $\lim_{n \rightarrow \infty} a_n / \log n = \infty$ —here \mathbb{E} denotes expectation over both the initial condition, \mathbb{P}_{in} , and the noise in (114).

5. ILLUSTRATIVE EXAMPLE: 3-SPIN MODEL ON THE HIGH-DIMENSIONAL SPHERE

To test our results, we use a function known for its complex features in high-dimensions: the spherical 3-spin model, $f : S^{d-1}(\sqrt{d}) \rightarrow \mathbb{R}$, given by

$$(134) \quad f(\mathbf{x}) = \frac{1}{d} \sum_{p, q, r=1}^d a_{p, q, r} x_p x_q x_r, \quad \mathbf{x} \in S^{d-1}(\sqrt{d}) \subset \mathbb{R}^d$$

where the coefficients $\{a_{p, q, r}\}_{p, q, r=1}^d$ are independent Gaussian random variables with mean zero and variance one. The function (134) is known to have a number of critical points that grows exponentially with the dimensionality d [ABA13, SGBAL14, ABAČ12]. We note that previous works have sought to draw a parallel between the glassy 3-spin function and generic loss functions [CHM⁺14], but we are not exploring such an analogy here. Rather, we simply use the function (134) as a difficult target for approximation by neural networks. That is, throughout this section, we train networks to learn f with a particular realization of $a_{p, q, r}$ and study the accuracy of that representation as a function of the number of particles n .

5.1. **Learning with Gaussian kernels.** We first consider the case when $D = S^{d-1}(\sqrt{d})$ and we use

$$(135) \quad \varphi(\mathbf{x}, \mathbf{z}) = e^{-\frac{1}{2}\alpha|\mathbf{x}-\mathbf{z}|^2}$$

for some fixed $\alpha > 0$. In this case, the parameters are elements of the domain of the function (here the d -dimensional sphere). Note that, since $|\mathbf{x}| = |\mathbf{z}| = \sqrt{d}$, up to an irrelevant constant that can be absorbed in the weights c , we can also write (135) as

$$(136) \quad \varphi(\mathbf{x}, \mathbf{z}) = e^{-\alpha\mathbf{x}\cdot\mathbf{z}}$$

This setting allow us to simplify the problem. Using

$$(137) \quad f^{(n)}(\mathbf{x}) = \frac{1}{n} \sum_{i=1}^n c_i \varphi(\mathbf{x}, \mathbf{z}_i) = \frac{1}{n} \sum_{i=1}^n c_i e^{-\alpha\mathbf{x}\cdot\mathbf{z}_i},$$

we can use as alternative loss

$$(138) \quad \mathcal{L}[f^{(n)}] = -\frac{1}{n} \sum_{i=1}^n c_i f(\mathbf{z}_i) + \frac{1}{2n^2} \sum_{i,j=1}^n c_i c_j \varphi(\mathbf{z}_i, \mathbf{z}_j)$$

i.e. eliminate the need for data beside the set $\{\mathbf{z}_i\}_{i=1}^n$. In terms of the empirical distribution, the loss can be represented as

$$(139) \quad \mathcal{L}[f^{(n)}] = -\int_{\hat{D}} f(\mathbf{z}) \gamma^{(n)}(d\mathbf{z}) + \frac{1}{2} \int_{\hat{D} \times \hat{D}} \varphi(\mathbf{z}, \mathbf{z}') \gamma^{(n)}(d\mathbf{z}) \gamma^{(n)}(d\mathbf{z}')$$

where $\gamma^{(n)} = \int_{\mathbb{R}} c \mu^{(n)}(dc, \cdot)$. Viewed as an integral kernel, φ is positive definite, as a result the loss is a convex functional of $\gamma^{(n)}$ (or $\mu^{(n)}$). Hence, the results established above apply to this special case, as well. The GD flow on the loss (138) can now be written explicitly as

$$(140) \quad \begin{cases} \dot{\mathbf{z}}_i = c_i \nabla f(\mathbf{z}_i) + \frac{\alpha}{n} \sum_{j=1}^n c_i c_j \mathbf{z}_j e^{-\alpha\mathbf{z}_i \cdot \mathbf{z}_j} - \lambda_i \mathbf{z}_i \\ \dot{c}_i = f(\mathbf{z}_i) - \frac{1}{n} \sum_{j=1}^n c_j e^{-\alpha\mathbf{z}_i \cdot \mathbf{z}_j} \end{cases}$$

where $-\lambda_i \mathbf{z}_i$ is a Lagrange multiplier term added to enforce $|\mathbf{z}_i| = \sqrt{d}$ for all $i = 1, \dots, n$, $f(\mathbf{x})$ is given by (134), and $\nabla f(\mathbf{z})$ is given componentwise by

$$(141) \quad \frac{\partial f}{\partial z_p} = \frac{1}{d} \sum_{q,r=1}^d (a_{p,q,r} + a_{r,p,q} + a_{q,r,p}) z_q z_r,$$

As is apparent from (140) the advantage of using radial basis function networks (or, in fact, any unit $\hat{\varphi}$ which is (i) such that $\hat{D} = \Omega$ and (ii) positive definite) is that we can use $f(\mathbf{x})$ and the unit $\varphi(\mathbf{x}, \mathbf{z})$ directly, and do not need to evaluate $\hat{F}(\mathbf{z})$ and $\hat{K}(\mathbf{z}, \mathbf{z}')$ (that is, we need no batch). In other words, the cost of running (140) scales like $(dn)^2$, instead of $P(Nn)^2$ in the case of a general network optimized by SGD with a batch of size P and $\mathbf{z} \in \hat{D} \subset \mathbb{R}^N$. If we make P scale with n , like $P = Cn^{2\alpha}$ for some $C > 0$, as we need to do to obtain the scalings discussed in Sec. 4, the cost of SGD becomes $N^2 n^{2+2\alpha}$, which is quickly becomes much worse than $(dn)^2$ as n grows.

We tested the representation (137) in $d = 5$ using $n = 16, 32, 64, 128$, and 256 and setting $\alpha = 5/d = 1$. The training was done by running a time-discretized version of (140) with time step $\Delta t = 10^3$ for 2×10^5 steps: during the first 10^5 we added thermal noise to (140), which we then remove during the second half of the run. The representation (137) proves to be accurate even at rather low value of n : for example, the right panel of Fig. 1 shows a contour plot of the original function f and its representation $f^{(n)}$ with $n = 128$ through a slice of the sphere defined as

$$(142) \quad \mathbf{x}(\theta) = \sqrt{d} (\sin(\theta) \cos(\phi), \sin(\theta) \sin(\phi), \cos(\theta), 0, 0),$$

with $\theta \in [0, \pi]$ and $\phi \in [0, 2\pi)$. The level sets of both functions are in good agreement. Also shown on this figure is the projection on the slice of the position of the 64 particles on the sphere. In this result, the parameters c_i take values that are initially uniformly distributed by about $-40d^2 =$

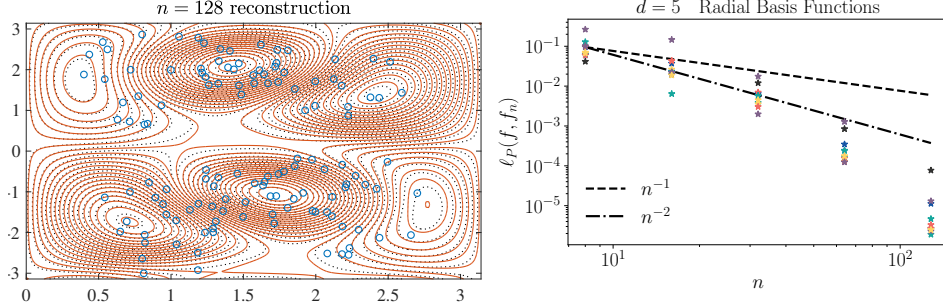


FIGURE 1. Left panel: Comparison between the level sets of the original function f in (134) (black dotted curves) and its approximation by the neural network in (137) with $n = 128$ and $d = 5$ in the slice defined by (142). Also shown are the projection in the slice of the particle position. Right panel: empirical loss in (143) vs n at the end of the calculation. The stars show the empirical loss for 10 independent realizations of the coefficients $a_{p,q,r}$ in (134).




FIGURE 2. The log of the empirical loss in (143) as a function of training time by SGD for the sigmoid neural network in $d = 10$ (left panel) and $d = 25$ (right panel). At time $t = 2 \times 10^6$, the batch size is increased to initiate a quench. The insets show the log of the empirical loss as a function of time during the final 10^5 time steps of training.

-10^3 and $40d^2 = 10^3$. To test the accuracy of the representation, we used the following Monte Carlo estimate of the loss function

$$(143) \quad \mathcal{L}_P[f_t^{(n)}] = \frac{1}{2P} \sum_{p=1}^P \left| f(\mathbf{x}_p) - f_t^{(n)}(\mathbf{x}_p) \right|^2.$$

This empirical loss function was computed with a batch of 10^6 points \mathbf{x}_p uniformly distributed on the sphere. The value (143) calculated at the end of the calculation is shown as a function of n in the right panel of Fig. 1: the empty circles show (143) for 4 individual realizations of the coefficient $a_{p,q,r}$ in (134), the full circle shows the average of (143) over these 4 realizations. The blue line scale as n^{-1} , the red one as n^{-2} : as can be seen, the empirical loss decays with n faster than n^{-1} , which is as expected.

5.2. Learning with single layer networks with sigmoid nonlinearity. To further test our predictions and also assess the learnability of high dimensional functions, we used 3-spin models in




FIGURE 3. Error scaling for single layer neural network with sigmoid nonlinearities. Upper row: $d = 10$; lower row: $d = 25$. The first column shows the empirical loss in (143), the second column shows (146), and the third column shows (146) with $\Theta(f)$ replaced by $\Theta(-f)$. The stars show the results for 10 different realizations of the coefficients $a_{p,q,r}$ in (134): the dashed lines decay as n^{-1} , consistent with the predictions in (133) and 3.7.

$d = 10$ and 25 dimensions, which we approximated with a single-layer neural network with sigmoid nonlinearity parameterized by $\mathbf{z} = (\mathbf{a}, b) \in D = \mathbb{R}^{d+1}$, with $\mathbf{a} \in \mathbb{R}^d$, $b \in \mathbb{R}$, and

$$(144) \quad \varphi(\mathbf{x}, \mathbf{z}) = h(\mathbf{a} \cdot \mathbf{x} + b).$$

This gives

$$(145) \quad f^{(n)}(\mathbf{x}) = \frac{1}{n} \sum_{i=1}^n c_i h(\mathbf{a}_i \cdot \mathbf{x} + b_i)$$

where $h(z) = 1/(1 + e^{-z})$. Simple networks like these, as opposed to deep neural with many parameters, provide greater assurance that we have trained sufficiently to test the scaling.

We trained the model in (145) using SGD with an initial batch size of $P = \lfloor n/5 \rfloor$ points uniformly sampled on the sphere for 2×10^6 time steps, resampling a new batch at every time step: this corresponds to choosing $\alpha = 1/2$ in the notation of Sec. 4. Towards the end of the trajectory, we initiated a partial quench by increasing the batch size to $P = \lfloor (n/5)^2 \rfloor$ (i.e $\alpha = 1$) which we run for an additional 2×10^5 time steps. Fig. 2 shows the empirical loss in (143) calculated over the batch as a function of training time during the optimization with $n = 256$ particles and $d = 10$ (left panel) and $d = 25$ (right panel). Note that the lack of intermediate plateaus in the loss during training is consistent with our conclusion that the dynamics effectively descends on a quadratic energy landscape (i.e. the loss function itself) at the level of the empirical distribution of the particles. After the quench the empirical loss shows substantially smaller fluctuations as a function of time which helps to reduce the fluctuating error. The inset shows the final 10^5 time steps in which there is negligible downward drift, indicating convergence towards stationarity at this batch size.

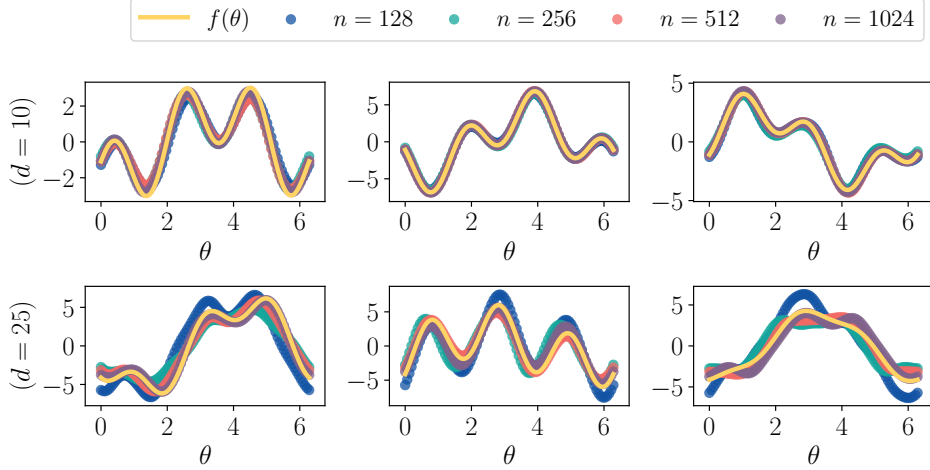


FIGURE 4. One dimensional slices through the $d = 10$ (upper row) and $d = 25$ (lower row) neural net representation $f^{(n)}$ are shown below a yellow curve with the target function f . In $d = 10$, the function representations clearly capture the main features of the target function, with only small scale deviations. In $d = 25$ there is remarkably good signal when $n = 1024$ while the smaller neural network is less able to faithfully represent the target function.

In these higher dimensional examples, we tested the scaling with three different observables. First, we considered the empirical loss function in (143) which we computed over a batch of size $\hat{P} = 10^5$ larger than P . As shown in the two right panels Fig. 3, $\mathcal{L}_{\hat{P}}[f_t^{(n)}]$ scales as n^{-1} , as expected. We also tested the estimate in (133) using

$$(146) \quad \frac{1}{\hat{P}} \sum_{p=1}^{\hat{P}} \Theta(f(\mathbf{x}_p)) \left(f(\mathbf{x}_p) - f_t^{(n)}(\mathbf{x}_p) \right),$$

and similarly with $\Theta(-f(\mathbf{x}_p))$: here Θ denotes the Heaviside function. The result is shown in the four right panels in Fig. 3: (146) scales as n^{-1} , consistent with (133) and our choice of $\alpha = 1$.

To provide further confidence in the quality of the representations, we also made a visual comparison by plotting f and $f^{(n)}$ along great circles of the sphere. We do so by picking $i \neq j$ in $\{1, \dots, d\}$ and setting $\mathbf{x} = \mathbf{x}(\theta) = (x_1(\theta), \dots, x_d(\theta))$ with

$$(147) \quad x_i(\theta) = \sqrt{d} \cos(\theta), \quad x_j(\theta) = \sqrt{d} \sin(\theta), \quad x_k(\theta) = 0 \quad \forall k \neq i, j.$$

In Fig. 4 we plot $f(\mathbf{x}(\theta))$ and $f^{(n)}(\mathbf{x}(\theta))$ along three great circles for $d = 10$ and $d = 25$. As can be seen, the agreement is quite good and confirms the quality of the final fit. A strong signal is present in $d = 25$ with $n = 1024$, a remarkable fact when considering that if we had only two grid points per dimension, the total number of points in the grid would be $2^{25} = 33,554,432$.

6. CONCLUDING REMARKS

Viewing parameters as particles with the loss function as interaction potential enables us to leverage a powerful theoretical apparatus developed to analyze problems from statistical physics. Using these ideas, we can analyze the approximation quality and the trainability of neural network representations of high-dimensional functions. Several insights emerge from our analysis based on this viewpoint: First, these tools show the dynamical realizability of the Universal Approximation Theorems, a direct consequence of the Law of Large Numbers for the empirical distribution of the parameters. Specifically, we conclude that the empirical distribution effectively

descends on the quadratic loss function landscape when the number n of parameters in the network is large. This confirms the empirical observation that wide neural networks are trainable despite the non-convexity of the loss function viewed from the individual particles perspective (as opposed to that of their empirical distribution). Secondly, we have derived a Central Limit Theorem for the empirical distribution of the parameters, specifying the approximation error of the neural network representation and showing that it is universal.

We derived these results first in the context of gradient descent dynamics; however, our conclusions also apply to stochastic gradient descent. The analysis indicates how the parameters in SGD should be chosen, in particular how the batch size should be scaled with n given the time step used in the scheme, which can be done towards the end of training.

These results were derived for a quadratic loss, $\mathcal{L}[f^{(n)}] = \frac{1}{2}\mathbb{E}_v|f - f^{(n)}|^2$. However, they do generalize to other losses as long as they are convex in $f^{(n)}$.

We also worked in the limit of an infinite amount of training data, an idealized setting that does not address the error incurred from a finite data set. For a neural network trained on a dataset of P points, $\{\mathbf{x}_p\}_{p=1}^P$, we can decompose the ‘‘generalization’’ error into components that involve the approximation error and the error from the finiteness of the data,

$$(148) \quad \mathbb{E}_v|f - f_p^{(n)}|^2 \leq \mathbb{E}_v|f - f_p|^2 + \mathbb{E}_v|f_p - f_p^{(n)}|^2$$

where f_p and $f_p^{(n)}$ are the approximations of f we can get if we train the network on the empirical loss build on $\{\mathbf{x}_p\}_{p=1}^P$ with finitely ($n < \infty$) or infinitely ($n \rightarrow \infty$) many units, respectively. Our results give direct insight on the second term at the right hand side of (148). We leave assessments of the first term for future work.

Our numerical results not only confirm our predictions, they emphasize the capability of neural networks to represent high-dimensional function accurately with a relatively modest number of adjustable parameters. Needless to say, the computational achievements of neural networks open the door to developments in scientific computing that we are only beginning to grasp. Such applications may benefit from better understanding how the specific architecture of the neural networks affects the approximation error and trainability, not in the general terms of their scaling with n that we analyzed here, but in the details of the constant involved.

APPENDIX A. TRAINING AT FINITE (BUT SMALL) TEMPERATURE

For completeness, let us consider here the case when noise-terms are added in (29) and the ODEs become stochastic differential equations (SDEs). Additive noise addresses the non-uniqueness issues encountered in Sec. 3. To formulate the resulting SDEs, we need a distribution $\mu_0 \in \mathcal{M}_+(D)$, used to regularize the dynamics. We specify its properties via:

Assumption A.1. *The distribution μ_0 (i) has a density ρ_0 that is continuously differentiable, $\rho_0 \in C^1(D)$; (ii) is such that $\text{supp}(\mu_0) = D$; and (iii) satisfies*

$$(149) \quad \forall b \in \mathbb{R} : \int_{\mathbb{R}} e^{bc} \mu_0(dc, \cdot) < \infty \quad \text{and} \quad \int_{\mathbb{R}} c \mu_0(dc, \cdot) = 0.$$

We then replace (29) with the SDEs

$$(150) \quad \begin{aligned} d\boldsymbol{\theta}_i &= \nabla F(\boldsymbol{\theta}_i) dt - \frac{1}{n} \sum_{j=1}^n \nabla K(\boldsymbol{\theta}_i, \boldsymbol{\theta}_j) dt \\ &+ (\beta n)^{-1} \nabla \log \rho_0(\boldsymbol{\theta}_i) dt + \sqrt{2}(\beta n)^{-1/2} d\mathbf{W}_i, \end{aligned}$$

for $i = 1, \dots, n$. Here \mathbf{W}_i are n independent Wiener processes, taking values in D , and $\beta > 0$ is a parameter playing the role of inverse temperature and controlling the amplitude of a noise added to the dynamics. Note the specific scale on which the regularizing and the noise terms act in (150):

they are higher order perturbations. We comment on the choice of this scaling in Remark A.2 below. The SDEs (150) are overdamped Langevin equations associated with the energy:

$$(151) \quad \begin{aligned} E_\beta(\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_n) &= nC_f - \sum_{i=1}^n F(\boldsymbol{\theta}_i) + \frac{1}{2n} \sum_{i,j=1}^n K(\boldsymbol{\theta}_i, \boldsymbol{\theta}_j) \\ &\quad - (\beta n)^{-1} \sum_{i=1}^n \log \rho(\boldsymbol{\theta}_i), \end{aligned}$$

This energy is (30) plus a regularizing term (the one involving $-\log \rho_0$). Under Assumption A.1 this term guarantees that, for any $\beta > 0$, the following integral is finite

$$(152) \quad Z_n = \int_{D^n} e^{-n\beta E_\beta(\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_n)} d\boldsymbol{\theta}_1 \cdots d\boldsymbol{\theta}_n < \infty$$

which in turns implies that

$$(153) \quad Z_n^{-1} \exp(-n\beta E_\beta(\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_n))$$

is a normalized probability density on D^n . As a result, the solutions to (29) are ergodic with respect to the equilibrium distribution with density (153) for any $\beta > 0$.

A.1. Dean's equation. Let $\chi : D \rightarrow \mathbb{R}$ be a test function. Applying Itô's formula to $n^{-1} \sum_{i=1}^n \chi(\boldsymbol{\theta}_i(t)) = \int_D \chi(\boldsymbol{\theta}) \mu_t^{(n)}(d\boldsymbol{\theta})$ and using (150) gives

$$(154) \quad \begin{aligned} & d \int_D \chi(\boldsymbol{\theta}) \mu_t^{(n)}(d\boldsymbol{\theta}) \\ &= \frac{1}{n} \sum_{i=1}^n \nabla \chi(\boldsymbol{\theta}_i(t)) \cdot d\boldsymbol{\theta}_i(t) + \beta^{-1} \sum_{i=1}^n \Delta \chi(\boldsymbol{\theta}_i(t)) dt \\ &= \int_D \nabla \chi(\boldsymbol{\theta}) \cdot \left(\nabla F(\boldsymbol{\theta}) \mu_t^{(n)}(\boldsymbol{\theta}) - \int_D \nabla K(\boldsymbol{\theta}, \boldsymbol{\theta}') \mu_t^{(n)}(d\boldsymbol{\theta}') \mu_t^{(n)}(d\boldsymbol{\theta}) \right) dt \\ &\quad + (\beta n)^{-1} \int_D \nabla \chi(\boldsymbol{\theta}) \cdot \left(\nabla \log \rho_0(\boldsymbol{\theta}) \mu_t^{(n)}(d\boldsymbol{\theta}) \right) dt \\ &\quad + (\beta n)^{-1} \int_D \Delta \chi(\boldsymbol{\theta}) \mu_t^{(n)}(d\boldsymbol{\theta}) dt \\ &\quad + \sqrt{2} (\beta n^3)^{-1/2} \sum_{i=1}^n \nabla \chi(\boldsymbol{\theta}_i(t)) \cdot d\mathbf{W}_i(t) \end{aligned}$$

The drift terms in this equation are closed in terms of $\mu_t^{(n)}$; the noise term has a quadratic variation given by:

$$(155) \quad \begin{aligned} & \left\langle (\beta n^3)^{-1/2} \sum_{i=1}^n \nabla \chi(\boldsymbol{\theta}_i(t)) \cdot d\mathbf{W}_i(t), (\beta n^3)^{-1/2} \sum_{i=1}^n \nabla \chi(\boldsymbol{\theta}_i(t)) \cdot d\mathbf{W}_i(t) \right\rangle \\ &= \beta^{-1} n^{-3} \sum_{i=1}^n |\nabla \chi(\boldsymbol{\theta}_i(t))|^2 dt \\ &= \beta^{-1} n^{-2} \int_D |\nabla \chi(\boldsymbol{\theta})|^2 \mu_t^{(n)}(d\boldsymbol{\theta}) dt \end{aligned}$$

As a result, (154) is sometimes written formally as the stochastic partial differential equation (SPDE)

$$(156) \quad \begin{aligned} \partial_t \mu_t^{(n)} &= \nabla \cdot \left(\nabla F \mu_t^{(n)} + \int_D \nabla K(\boldsymbol{\theta}, \boldsymbol{\theta}') \mu_t^{(n)}(d\boldsymbol{\theta}') \mu_t^{(n)} \right) \\ &\quad - (\beta n)^{-1} \nabla \cdot \left(\nabla \log \rho_0 \mu_t^{(n)} \right) \\ &\quad + (\beta n)^{-1} \Delta \mu_t^{(n)} + \sqrt{2} \beta^{-1/2} n^{-1} \nabla \cdot \left(\sqrt{\mu_t^{(n)}} \cdot \dot{\boldsymbol{\eta}}_t \right) \end{aligned}$$

where $\dot{\boldsymbol{\eta}}_t = \dot{\boldsymbol{\eta}}_t(\theta)$ is a spatio-temporal white-noise so that the quadratic variation of the noise term in (156) is formally given by (155). This equation is referred to as Dean's equation. It is difficult to give (33) a precise meaning because it is not clear how to interpret the noise term. It remains useful to analyze the properties of $\mu_t^{(n)}$ as $n \rightarrow \infty$, however, which is what we will do next.

Remark A.2. We could also consider situations where in (150) n^{-1} is replaced by $n^{-\alpha}$ with $\alpha \in [0, 1)$. The case $\alpha = 0$ is treated in [MMN18]: with this scaling, the diffusive and regularizing terms in (156) are replaced by

$$\beta^{-1} \Delta \mu_t^{(n)} - \beta^{-1} \nabla \cdot (\nabla \log \rho_0 \mu_t^{(n)}),$$

and the noise terms by

$$\sqrt{2}(\beta n)^{-1/2} \nabla \cdot (\sqrt{\mu_t^{(n)}} \dot{\boldsymbol{\eta}}_t).$$

This means that these diffusive and regularizing terms affect the mean field limit equation for μ_t , whereas the noise terms remain higher order. In particular, in that case one can prove that $\mu_t^{(n)} \rightarrow \mu_t$ with μ_t that converges to a unique fixed point μ_β such that $\mu_\beta > 0$ a.e. in D but for which $\int_D \varphi(\cdot, \boldsymbol{\theta}) \mu_\beta(d\boldsymbol{\theta}) \neq f$ (there is a correction proportional to β^{-1}). When $\alpha \in (0, 1)$, the diffusive and regularizing terms in (156) are replaced by

$$\beta^{-1} n^{-\alpha} \Delta \mu_t^{(n)} - \beta^{-1} n^{-\alpha} \nabla \cdot (\nabla \log \rho_0 \mu_t^{(n)}),$$

and the noise terms by

$$\sqrt{2}(\beta n^{1+\alpha})^{-1/2} \nabla \cdot (\sqrt{\mu_t^{(n)}} \dot{\boldsymbol{\eta}}_t).$$

This means that none of these terms affect the mean field limit equation, but at next order, $O(n^{-\alpha})$, the diffusive and regularizing terms dominate whereas the noise terms remain higher order. In the case when $\alpha = 1$, on which we focus here, the diffusive, regularizing, and noise terms are perturbations on the $O(n^{-1})$ same scale, the same scale as the errors introduced by discretization effects (finite n) also present in GD.

A.2. Multiple-scale expansion. The advantage of adding noise terms in (150) is that it guarantees ergodicity of the solution to these SDEs with respect to the equilibrium distribution with density (153). Correspondingly, we focus on analyzing the long-time ergodicity properties of the empirical distribution satisfying (156). On long timescales, the memory of the initial conditions is lost, and we can directly pick the right scaling to analyze the fluctuations of $\mu_t^{(n)}$ around its limit μ_t : as discussed in Remark A.2 and confirmed below, this scale is $O(n^{-1})$, consistent with what we reach at long times with GD as discussed in Sec. 3.

We analyze (156) by formal asymptotic, using a two-timescale expansion. Consistent with the expected $O(n^{-1})$ scaling of the fluctuations, we look for a solution of this equation of the form

$$(157) \quad \mu_t^{(n)} = \mu_{t,\tau} + n^{-1} \omega_{t,\tau} + o(n^{-1}), \quad \tau = t/n.$$

We use the rescaled time $\tau = t/n$ to look at the solution to (156) on $O(n)$ timescales. Not only does this fix the behavior of $\mu_{t,\tau}$ on long timescales but also guarantees solvability of the equation for $\omega_{t,\tau}$. Treating t and τ as independent variables, (157) implies that

$$(158) \quad \partial_t \mu_t^{(n)} = \partial_t \mu_{t,\tau} + n^{-1} (\partial_\tau \mu_{t,\tau} + \partial_t \omega_{t,\tau})$$

Inserting (157) and (158) in (156) and collecting terms of the same order in n^{-1} , we arrive at the following two equations at order $O(1)$ and $O(n^{-1})$, respectively

$$(159) \quad \partial_t \mu_{t,\tau} = \nabla \cdot (\nabla V(\boldsymbol{\theta}, [\mu_{t,\tau}]) \mu_{t,\tau})$$

and

$$(160) \quad \begin{aligned} \partial_\tau \mu_{t,\tau} + \partial_t \omega_{t,\tau} &= \nabla \cdot (\nabla V(\boldsymbol{\theta}, [\mu_{t,\tau}]) \omega_{t,\tau} + \nabla F(\boldsymbol{\theta}, [\omega_{t,\tau}]) \mu_{t,\tau}) \\ &+ \beta^{-1} \Delta \mu_{t,\tau} - \beta^{-1} \nabla \cdot (\nabla \log \rho_0 \mu_{t,\tau}) \\ &+ \sqrt{2} \beta^{-1/2} \nabla \cdot (\sqrt{\mu_{t,\tau}} \dot{\boldsymbol{\eta}}_t) \end{aligned}$$

A.3. Law of Large Numbers at finite temperature. Since (159) is identical to (37), the results we established in Sec. 3.3 still hold at finite temperature. In particular, Proposition 3.5 applies. As we see below, we can obtain more information about μ_t by looking at the evolution of this function on longer timescales, and we will be able to deduce that $\text{supp } \mu_{t,\tau} = D$. This guarantees that (25) holds, so it can be removed from the assumptions needed in Proposition 3.5.

A.4. Global convergence on $O(n)$ timescales. An equation governing the evolution of $\mu_{t,\tau}$ on the rescaled time $\tau = t/n$ can be derived by time averaging (160) over t . This equation guarantees the solvability of (160). Since $\mu_{t,\tau} \rightarrow \mu_\tau$ as $t \rightarrow \infty$, where μ_τ is a stationary point of (159), we have

$$(161) \quad \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T \mu_{t,\tau} \omega_{t,\tau} dt = \mu_\tau \bar{\omega}_\tau$$

where

$$(162) \quad \bar{\omega}_\tau =: \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T \omega_{t,\tau} dt$$

in which we assume that the time-average of $\omega_{t,\tau}$ exists (which we check *a posteriori*). Using (161) and the fact that the white-noise terms time-average to zero almost surely, we deduce that the time-average of (160) is

$$(163) \quad \begin{aligned} \partial_\tau \mu_\tau &= \nabla \cdot (\nabla V(\boldsymbol{\theta}, [\mu_\tau]) \bar{\omega}_\tau + \nabla F(\boldsymbol{\theta}, [\bar{\omega}_\tau]) \mu_\tau) \\ &\quad + \beta^{-1} \Delta \mu_\tau - \beta^{-1} \nabla \cdot (\nabla \log \rho_0 \mu_\tau) \end{aligned}$$

Because of the presence of the diffusive term $\beta^{-1} \Delta \mu_\tau$ in (163), we can therefore conclude that on the timescales where this equation holds we must have $\mu_\tau > 0$ a.e. on D . This means that (25) holds and so $V(\boldsymbol{\theta}, [\mu_\tau]) = 0$ since $\int_D \varphi(\cdot, \boldsymbol{\theta}) \mu_\tau(d\theta) = f$ by Proposition 3.5. As a result (163) reduces to

$$(164) \quad \partial_\tau \mu_\tau = \nabla \cdot (\nabla F(\boldsymbol{\theta}, [\bar{\omega}_\tau]) \mu_\tau) + \beta^{-1} \Delta \mu_\tau - \beta^{-1} \nabla \cdot (\nabla \log \rho_0 \mu_\tau).$$

Since $V(\boldsymbol{\theta}, [\mu_\tau]) = 0$ needs to be satisfied, in (164) we can treat the term involving the factor $F(\boldsymbol{\theta}, [\bar{\omega}_\tau])$ as a Lagrange multiplier used to enforce this constraint. It is also easy to see that (164) is the Wasserstein GD flow on

$$(165) \quad \int_D (\beta^{-1} \log(d\mu/d\mu_0) + F(\boldsymbol{\theta}, [\bar{\omega}_\tau])) \mu(d\boldsymbol{\theta})$$

Since this energy is strictly convex, a direct consequence of these observations is that the stable fixed points of (164) are the minimizers of the energy (165) subject to the constraints that $V(\boldsymbol{\theta}, [\mu_\tau]) = 0$ and $\mu_\tau \in \mathcal{M}_+(D)$. These fixed points are reached on a timescale that is large compared the $O(n)$ timescale $\tau = t/n$.

Recalling that $\mu_{t,\tau}$ is the weak limit of $\mu_t^{(n)}$ as $n \rightarrow \infty$ and $V(\boldsymbol{\theta}, [\mu]) = -F(\boldsymbol{\theta}) + \int_D K(\boldsymbol{\theta}, \boldsymbol{\theta}') \mu(d\boldsymbol{\theta}')$, we can summarize these considerations into:

Proposition A.3. *If $\mu_t^{(n)}$ be the empirical distribution defined in (31) with $\{\boldsymbol{\theta}_i(t)\}_{i=1}^n$ the solution to (156). Then given any $b_n > 0$ such that $b_n/n \rightarrow \infty$ as $n \rightarrow \infty$, we have*

$$(166) \quad \mu_{b_n}^{(n)} \rightarrow \mu^* \quad \text{as } n \rightarrow \infty$$

where μ^* is the minimizer in $\mathcal{M}_+(D)$ of

$$(167) \quad \beta^{-1} \int_D \log(d\mu/d\mu_0) \mu(d\boldsymbol{\theta})$$

subject to

$$(168) \quad F(\boldsymbol{\theta}) = \int_D K(\boldsymbol{\theta}, \boldsymbol{\theta}') \mu^*(d\boldsymbol{\theta}') \quad \text{a.e. in } D$$

It is easy to see that the solution to the minimization problem in Proposition A.3 is such that

$$(169) \quad \int_D \log(d\mu^*/d\mu_0)\mu^*(d\boldsymbol{\theta}) < \infty \quad \text{and} \quad \text{supp } \mu^* = D.$$

The first condition says that the minimizer exists, which is clear since we can find test distributions $\mu \in \mathcal{M}_+(D)$ such that (i) $\int_{\mathbb{R}} c\mu(dc, \cdot) = \gamma^*$ where γ^* solves (21) (i.e. such that μ satisfies the constraint in (168)), and (ii) μ has finite entropy with respect to μ_0 . One such μ is

$$(170) \quad \mu(dc, d\mathbf{z}) = |\gamma^*|_{\text{TV}}^{-1} (\delta_{|\gamma^*|_{\text{TV}}}(dc)\gamma_+^*(d\mathbf{z}) + \delta_{-|\gamma^*|_{\text{TV}}}(dc)\gamma_-^*(d\mathbf{z}))$$

To prove that $\text{supp } \mu^* = D$, suppose by contradiction that the minimizer is such that $\mu^* = 0$ if $\boldsymbol{\theta} \in B$ with $\int_B \mu_0(d\boldsymbol{\theta}) > 0$. For $s \in [0, 1]$, consider $\mu^s = (1-s)\mu^* + s\mu_0$. A direct calculation shows that

$$(171) \quad \begin{aligned} \int_D \log(d\mu^s/d\mu_0)\mu^s(d\boldsymbol{\theta}) &= \int_B \log(d\mu^*/d\mu_0)\mu^*(d\boldsymbol{\theta}) \\ &+ s \log s \int_{B^c} \mu_0(d\boldsymbol{\theta}) + O(s) \end{aligned}$$

Since $s \log s \int_{B^c} \mu_0(d\boldsymbol{\theta}) < 0$ for $s \in (0, 1)$, (171) implies that for $s > 0$ small enough

$$(172) \quad \int_D \log(d\mu^s/d\mu_0)\mu^s(d\boldsymbol{\theta}) < \int_B \log(d\mu^*/d\mu_0)\mu^*(d\boldsymbol{\theta}),$$

a contradiction with our assumption that μ^* is the minimizer.

Let us also analyze in some more detail the constrained optimization problem in Proposition A.3 since this will be useful in the next section. If we denote by μ^* the minimizer of (167) subject to (168) and by λ^* the Lagrange multiplier used to satisfy the first constraint in (168), this Lagrange multiplier is given by

$$(173) \quad \lambda^*(\boldsymbol{\theta}) = \beta^{-1} \frac{\delta}{\delta F(\boldsymbol{\theta})} \int_D \log(d\mu^*/d\mu_0)\mu^*(d\boldsymbol{\theta}')$$

It is easy to see that μ^* is independent of β : indeed, we can drop the factor β^{-1} in front of (167) without affecting the minimization problem. This also means that the dependency of λ^* in β is explicit: Indeed from (173)

$$(174) \quad \lambda^*(\boldsymbol{\theta}) = \beta^{-1} \delta^*(\boldsymbol{\theta})$$

where $\delta^*(\boldsymbol{\theta})$ is given by

$$(175) \quad \delta^*(\boldsymbol{\theta}) = \frac{\delta}{\delta F(\boldsymbol{\theta})} \int_D \log(d\mu^*/d\mu_0)\mu^*(d\boldsymbol{\theta}')$$

This factor is independent of β since μ^* is. It will be useful later to work with the function $\epsilon^*(\mathbf{x})$ defined via the equation

$$(176) \quad \int_D \varphi(\mathbf{x}, \boldsymbol{\theta}) \delta^*(\boldsymbol{\theta}) d\boldsymbol{\theta} = \int_D \int_{\Omega} \varphi(\mathbf{x}, \boldsymbol{\theta}) \varphi(\mathbf{x}', \boldsymbol{\theta}) \epsilon^*(\mathbf{x}') \nu(d\mathbf{x}') d\boldsymbol{\theta}$$

This is the Euler-Lagrange equation for the minimizer of

$$(177) \quad \frac{1}{2} \int_D \left| \delta^*(\boldsymbol{\theta}) - \int_{\Omega} \epsilon(\mathbf{x}) \varphi(\mathbf{x}, \boldsymbol{\theta}) \nu(d\mathbf{x}) \right|^2 d\boldsymbol{\theta}$$

over ϵ . Therefore, (176) is also the equation for the least square solution of

$$(178) \quad \delta^*(\boldsymbol{\theta}) = \int_{\Omega} \epsilon^*(\mathbf{x}) \varphi(\mathbf{x}, \boldsymbol{\theta}) \nu(d\mathbf{x})$$

and such a least square solution exists for a modification of $\delta^*(\boldsymbol{\theta})$ which is arbitrarily close to it in $L^2(D)$: any such solution for a modification of $\delta^*(\boldsymbol{\theta})$ that is $O(n^{-1})$ away from it is good enough

for our purpose since the discrepancy can be absorbed in higher order terms in our expansion in n^{-1} . This solution is also unique by Assumption 2.2 and it can be expressed as

$$(179) \quad \epsilon^*(\mathbf{x}) = D_{f(\mathbf{x})} \int_D \log(d\mu^*/d\mu_0) \mu^*(d\boldsymbol{\theta})$$

where μ^* is viewed as a functional of $f(\mathbf{x})$ by using $F(\boldsymbol{\theta}) = \int_\Omega f(\mathbf{x}) \varphi(\mathbf{x}, \boldsymbol{\theta}) \nu(d\mathbf{x})$, and $D_{f(\mathbf{x})}$ denotes the gradient with respect to $f(\mathbf{x})$ in the $L^2(\Omega, \nu)$ -norm defined in (50). The equality (179) follows from (175) and the fact that $D_{f(\mathbf{x})} F(\boldsymbol{\theta}) = \varphi(\mathbf{x}, \boldsymbol{\theta})$.

Remark A.4. *Compared to the case treated in [MMN18] where the noise and regularizing terms in (150) are scaled as β^{-1} (high temperature) rather than $(\beta n)^{-1}$ (low temperature), we see that we can also conclude that μ_t converges as $t \rightarrow \infty$ to a distribution μ^* with $\text{supp } \mu^* = D$; however, the fixed point μ^* we obtain satisfies $\int_D \varphi(\cdot, \boldsymbol{\theta}) \mu^*(d\boldsymbol{\theta}) = f$, whereas the one obtained at high temperature introduces a correction proportional to β^{-1} in this relation. The price we pay by working at low temperature is that convergence in time may be slower if the initial condition $\mu_0 = \mu_{in}$ is such that (25) is not satisfied by the GD flow without noise: specifically, this convergence should occur on timescales that are intermediate between $O(1)$ and $O(n)$.*

A.5. Central Limit Theorem at finite temperature. Now that we have determined the behavior of $\lim_{n \rightarrow \infty} \mu_t^{(n)} = \mu_t$ at all times, we can stop distinguishing τ from t , and focus on ω_t . We already know that (168) constrain the average value of ω_t on long timescales, but we would also like to quantify this average value beyond what (168) implies, and also analyze the fluctuations around this average. To this end, let us use (164) in (160) and look at the resulting equation on timescales where μ_t has converged to μ^* , the minimizer specified in Proposition A.3, so that $V(\boldsymbol{\theta}, [\mu^*]) = 0$ and λ has converged to $\lambda^* = \beta^{-1} \delta^*$. This can be achieved by considering (160) with initial condition at $t = T$ and pushing back $T \rightarrow -\infty$. The resulting equation is

$$(180) \quad \begin{aligned} \partial_t \omega_t = & \nabla \cdot \left(-\beta^{-1} \nabla \delta^* \mu^* + \int_D \nabla K(\boldsymbol{\theta}, \boldsymbol{\theta}') \omega_t(d\boldsymbol{\theta}') \mu^* \right) \\ & + \sqrt{2} \beta^{-1/2} \nabla \cdot \left(\sqrt{\mu^*} \dot{\boldsymbol{\eta}}_t \right) \end{aligned}$$

Even though we derived it formally, the SPDE (186) can be given a precise meaning: since its drift is linear in ω_t and its noise is additive (recall that μ^* is a given, non-random, distribution), (186) defines ω_t as a Gaussian process. This also means that

$$(181) \quad g_t = \int_D \varphi(\cdot, d\boldsymbol{\theta}) \omega_t(d\boldsymbol{\theta})$$

is a Gaussian process. This is an important quantity since gives the error on f made in $f_t^{(n)}$ at order $O(n^{-1})$:

$$(182) \quad f_t^{(n)} = f + n^{-1} g_t + o(n^{-1})$$

Let us derive a closed equation for g_t from (180). To this end, notice first that we can use

$$(183) \quad \int_D K(\boldsymbol{\theta}, \boldsymbol{\theta}') \omega_t(d\boldsymbol{\theta}') = \int_\Omega \varphi(\mathbf{x}, \boldsymbol{\theta}) g_t(\mathbf{x}) \nu(d\mathbf{x})$$

to express the integral terms in (180) in terms of g_t . By taking the time derivative of (181) and using (180) together with (183) and (176) we derive:

$$\begin{aligned}
 \partial_t g_t &= \int_D \varphi(\cdot, \boldsymbol{\theta}) \partial_t \omega_t(d\boldsymbol{\theta}) \\
 &= \int_D \nabla_{\boldsymbol{\theta}} \varphi(\cdot, \boldsymbol{\theta}) \cdot \int_{\Omega} \nabla_{\boldsymbol{\theta}} \varphi(\mathbf{x}', \boldsymbol{\theta}) (g_t(\mathbf{x}') - \beta^{-1} \epsilon^*(\mathbf{x}')) \nu(d\mathbf{x}') \mu^*(d\boldsymbol{\theta}) \\
 &\quad - \sqrt{2} \beta^{-1/2} \int_D \nabla_{\boldsymbol{\theta}} \varphi(\cdot, \boldsymbol{\theta}) \cdot \sqrt{\mu^*} \dot{\boldsymbol{\eta}}_t \\
 &= \int_{\Omega} M([\mu^*], \mathbf{x}, \mathbf{x}') (g_t(\mathbf{x}') - \beta^{-1} \epsilon^*(\mathbf{x}')) \nu(d\mathbf{x}') \\
 &\quad - \sqrt{2} \beta^{-1/2} \int_D \nabla_{\boldsymbol{\theta}} \varphi(\cdot, \boldsymbol{\theta}) \cdot \sqrt{\mu^*} \dot{\boldsymbol{\eta}}_t
 \end{aligned} \tag{184}$$

where $M([\mu^*], \mathbf{x}, \mathbf{x}')$ is the kernel defined in (46). Since the quadratic variation of the noise term in this equation is

$$2\beta^{-1} M([\mu^*], \mathbf{x}, \mathbf{x}') dt \tag{185}$$

in law it is equivalent to

$$\begin{aligned}
 \partial_t g_t &= - \int_{\Omega} M([\mu^*], \mathbf{x}, \mathbf{x}') (g_t(\mathbf{x}') - \beta^{-1} \epsilon^*(\mathbf{x})) \nu(d\mathbf{x}') \\
 &\quad + \sqrt{2} \beta^{-1/2} \int_{\Omega} \sigma([\mu^*], \mathbf{x}, \mathbf{x}') \dot{\eta}_t(\mathbf{x}') d\mathbf{x}'
 \end{aligned} \tag{186}$$

where $\dot{\eta}_t(\mathbf{x})$ is a spatio-temporal white-noise, and $\sigma([\mu^*], \mathbf{x}, \mathbf{x}')$ is such that

$$\int_{\Omega} \sigma([\mu^*], \mathbf{x}, \mathbf{x}'') \sigma([\mu^*], \mathbf{x}', \mathbf{x}'') d\mathbf{x}'' = M([\mu^*], \mathbf{x}, \mathbf{x}') \tag{187}$$

Note that this decomposition exists since $\mu^* \in \mathcal{M}_+(D)$ with $\text{supp } \mu^* = D$ and hence $M([\mu^*], \mathbf{x}, \mathbf{x}')$ is positive-definite. The asymptotic mean and variance of g_t can be readily deduced from (186) by noting that this Ornstein-Uhlenbeck equation is in detailed-balance with respect to the Gibbs distribution associated with the energy

$$\frac{1}{2} \int_{\Omega} |g_t(\mathbf{x}') - \beta^{-1} \epsilon^*(\mathbf{x}')|^2 \nu(d\mathbf{x}'). \tag{188}$$

We can state this as

Proposition A.5 (CLT at finite temperature). *Let $f_t^{(n)}$ be given by (32) with $\{\boldsymbol{\theta}_i(t)\}_{i=1}^n$ solution to (156) with initial conditions specified at $t = T$. Then*

$$\lim_{T \rightarrow -\infty} \lim_{n \rightarrow \infty} n(f_t^{(n)} - f) = g_t \quad \text{in law} \tag{189}$$

where g_t is the stationary Gaussian process specified by (186) and whose mean and covariance satisfy: for any test function $\chi: \Omega \rightarrow \mathbb{R}$

$$\begin{aligned}
 \mathbb{E} \int_{\Omega} \chi(\mathbf{x}) g_t(\mathbf{x}) \nu(d\mathbf{x}) &= \beta^{-1} \int_{\Omega} \chi(\mathbf{x}) \epsilon^*(\mathbf{x}) \nu(d\mathbf{x}) \\
 \mathbb{E} \left(\int_{\Omega} \chi(\mathbf{x}) (g_t(\mathbf{x}) - \beta^{-1} \epsilon^*(\mathbf{x})) \nu(d\mathbf{x}) \right)^2 &= \beta^{-1} \int_{\Omega} |\chi(\mathbf{x})|^2 \nu(d\mathbf{x})
 \end{aligned} \tag{190}$$

where ϵ^* is given by (179)

Notice that if we quench the result in (190) (i.e. send $\beta \rightarrow \infty$), we arrive at the conclusion that $g_t \rightarrow 0$ as $t \rightarrow \infty$ in that case. This is consistent with what happens at zero-temperature, in the limit as $\xi \rightarrow 1$, see Proposition 3.9.

REFERENCES

- [ABA13] Antonio Auffinger and Gérard Ben Arous. Complexity of random smooth functions on the high-dimensional sphere. *The Annals of Probability*, 41(6):4214–4247, November 2013.
- [ABAČ12] Antonio Auffinger, Gérard Ben Arous, and Jiří Černý. Random Matrices and Complexity of Spin Glasses. *Communications on Pure and Applied Mathematics*, 66(2):165–201, 2012.
- [AGS05] Luigi Ambrosio, Nicola Gigli, and Giuseppe Savaré. *Gradient Flows in Metric Spaces and in the Space of Probability Measures*. Lectures in Mathematics ETH Zürich. Birkhäuser Verlag AG, Basel, second edition, 2005.
- [Bac17] Francis Bach. Breaking the curse of dimensionality with convex neural networks. *Journal of Machine Learning Research*, 18(19):1–53, 2017.
- [Bar93] A R Barron. Universal approximation bounds for superpositions of a sigmoidal function. *IEEE Transactions on Information Theory*, 39(3):930–945, May 1993.
- [BEJ17] C. Beck, W. E, and A. Jentzen. Machine learning approximation algorithms for high-dimensional fully nonlinear partial differential equations and second-order backward stochastic differential equations. [arXiv:1709.05963](https://arxiv.org/abs/1709.05963), September 2017.
- [Bil99] Patrick Billingsley. *Convergence of probability measures*. Wiley Series in Probability and Statistics: Probability and Statistics. John Wiley & Sons Inc., second edition, 1999.
- [BL04] Léon Bottou and Yann LeCun. Large scale online learning. In S. Thrun, L. K. Saul, and B. Schölkopf, editors, *Advances in Neural Information Processing Systems 16*, pages 217–224. MIT Press, 2004.
- [BN17] Jens Berg and Kaj Nyström. A unified deep artificial neural network approach to partial differential equations in complex geometries. [arXiv:1711.06464](https://arxiv.org/abs/1711.06464), November 2017.
- [BP07] Jörg Behler and Michele Parrinello. Generalized Neural-Network Representation of High-Dimensional Potential-Energy Surfaces. *Physical Review Letters*, 98(14):583, April 2007.
- [BSG⁺18] M. Baity-Jesi, L. Sagun, M. Geiger, S. Spigler, G. Ben Arous, C. Cammarota, Y. LeCun, M. Wyart, and G. Biroli. Comparing Dynamics: Deep Neural Networks versus Glassy Systems. [arXiv:1803.06969](https://arxiv.org/abs/1803.06969), March 2018.
- [CB18] Lénaïc Chizat and Francis Bach. On the global convergence of gradient descent for over-parameterized models using optimal transport. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems 31*, pages 3036–3046. Curran Associates, Inc., 2018.
- [CHM⁺14] Anna Choromanska, Mikael Henaff, Michael Mathieu, Gérard Ben Arous, and Yann LeCun. The Loss Surfaces of Multilayer Networks. [arXiv:1412.0233](https://arxiv.org/abs/1412.0233), November 2014.
- [Cyb89] G Cybenko. Approximation by superpositions of a sigmoidal function. *Mathematics of Control, Signals, and Systems*, 2(4):303–314, December 1989.
- [Dea99] David S Dean. Langevin equation for the density of a system of interacting Langevin processes. *Journal of Physics A: Mathematical and General*, 29(24):L613–L617, January 1999.
- [DG87] Donald Dawson and Jürgen Gärtner. Large deviations from the McKean-Vlasov limit for weakly interacting diffusions. *Stochastics*, 20(4):247–308, 1987.
- [EHJ17] W. E, J. Han, and A. Jentzen. Deep learning-based numerical methods for high-dimensional parabolic partial differential equations and backward stochastic differential equations. [arXiv:1706.04702](https://arxiv.org/abs/1706.04702), June 2017.
- [Gär88] Jürgen Gärtner. On the McKean-Vlasov Limit for Interacting Diffusions. *Mathematische Nachrichten*, 137(1):197–248, 1988.
- [HLLL17] Wenqing Hu, Chris Junchi Li, Lei Li, and Jian-Guo Liu. On the diffusion approximation of nonconvex stochastic gradient descent. [arXiv:1705.07562](https://arxiv.org/abs/1705.07562), May 2017.
- [JKO98] Richard Jordan, David Kinderlehrer, and Felix Otto. The Variational Formulation of the Fokker–Planck Equation. *SIAM Journal on Mathematical Analysis*, 29(1):1–17, January 1998.
- [KL13] Claude Kipnis and Claudio Landim. *Scaling limits of interacting particle systems*, volume 320. Springer Science & Business Media, 2013.
- [KLY18] Yuehaw Khoo, Jianfeng Lu, and Lexing Ying. Solving for high dimensional committor functions using artificial neural networks. [arXiv:1802.10275](https://arxiv.org/abs/1802.10275), February 2018.
- [LBH15] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521(7553):436–444, May 2015.
- [LS17] Thomas Leblé and Sylvia Serfaty. Large deviation principle for empirical fields of log and riesz gases. *Inventiones mathematicae*, 210(3):645–757, 2017.
- [LTE15] Qianxiao Li, Cheng Tai, and Weinan E. Stochastic modified equations and adaptive stochastic gradient algorithms. [arXiv:1511.06251](https://arxiv.org/abs/1511.06251), November 2015.
- [McK66] H P McKean. A class of markov processes associated with nonlinear parabolic equations. *Proceedings of the National Academy of Sciences*, 56(6):1907–1911, December 1966.
- [MMN18] 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.
- [PS91] J Park and I W Sandberg. Universal Approximation Using Radial-Basis-Function Networks. *Neural Computation*, 3(2):246–257, June 1991.

- [SDT⁺17] Elia Schneider, Luke Dai, Robert Q Topper, Christof Drechsel-Grau, and Mark E Tuckerman. Stochastic Neural Network Approach for Learning High-Dimensional Free Energy Surfaces. Physical Review Letters, 119(15):150601, October 2017.
- [Ser15] Sylvia Serfaty. Coulomb gases and Ginzburg-Landau vortices. Zurich Lectures in Advanced Mathematics. European Mathematical Society Publishing House, 2015.
- [Ser17] Sylvia Serfaty. Systems of Points with Coulomb Interactions. arXiv:1712.04095, December 2017.
- [SGBAL14] Levent Sagun, V. Ugur Guney, Gérard Ben Arous, and Yann LeCun. Explorations on high dimensional landscapes. arXiv:1412.6615, December 2014.
- [SS18] Justin Sirignano and Konstantinos Spiliopoulos. Mean Field Analysis of Neural Networks. arXiv:1805.01053, May 2018.
- [Szn91] Alain-Sol Sznitman. Topics in propagation of chaos. In Paul-Louis Hennequin, editor, Ecole d'Été de Probabilités de Saint-Flour XIX—1989, pages 165–251, Berlin, Heidelberg, 1991. Springer Berlin Heidelberg.
- [Vil09] Cédric Villani. Optimal Transport: Old and New. Number 338 in Grundlehren Der Mathematischen Wissenschaften. Springer, Berlin, 2009. OCLC: ocn244421231.
- [ZHW⁺18] L. Zhang, J. Han, H. Wang, R. Car, and W. E. DeePCG: constructing coarse-grained models via deep neural networks. arXiv:1802.08549, February 2018.

COURANT INSTITUTE OF MATHEMATICAL SCIENCES, NEW YORK UNIVERSITY, 251 MERCER STREET, NEW YORK, NY 10012

E-mail address: `rotskoff@cims.nyu.edu`, `eve2@cims.nyu.edu`