

MEAN FIELD ANALYSIS OF NEURAL NETWORKS: A LAW OF LARGE NUMBERS*

JUSTIN SIRIGNANO[†] AND KONSTANTINOS SPILIOPOULOS[‡]

Abstract. Machine learning, and in particular neural network models, have revolutionized fields such as image, text, and speech recognition. Today, many important real-world applications in these areas are driven by neural networks. There are also growing applications in engineering, robotics, medicine, and finance. Despite their immense success in practice, there is limited mathematical understanding of neural networks. This paper illustrates how neural networks can be studied via stochastic analysis and develops approaches for addressing some of the technical challenges which arise. We analyze one-layer neural networks in the asymptotic regime of simultaneously (a) large network sizes and (b) large numbers of stochastic gradient descent training iterations. We rigorously prove that the empirical distribution of the neural network parameters converges to the solution of a nonlinear partial differential equation. This result can be considered a law of large numbers for neural networks. In addition, a consequence of our analysis is that the trained parameters of the neural network asymptotically become independent, a property which is commonly called “propagation of chaos.”

Key words. stochastic analysis, weak convergence, machine learning

AMS subject classifications. 60-08, 60F99, 60H10, 62M45

DOI. 10.1137/18M1192184

1. Introduction. Neural networks have achieved immense practical success over the past decade. Neural networks are nonlinear statistical models whose parameters are estimated from data using stochastic gradient descent. They have been employed as critical components of many important technologies in a variety of industries. This practical success has sparked significant interest in their mathematical analysis. Currently, there is limited mathematical understanding of neural networks. This paper analyzes the asymptotic behavior of neural networks, rigorously proving that the empirical distribution of their parameters converges to the solution of a nonlinear partial differential equation (PDE).

Neural network models have revolutionized fields such as image, text, and speech recognition. They are actively used in a variety of applications. In image recognition, neural networks are able to accurately identify and recognize objects in images using only the raw pixels. Neural networks are used for image recognition in applications such as self-driving cars, image searches on search engines such as Google, and facial recognition for security systems (see [31], [21], [8], and [53]). In speech recognition, neural networks are used to develop computer systems that automatically understand human speech (see [31], [3], [33], and [56]). Applications include voice control of certain systems in vehicles, transcription (automatically converting human speech to written text), interactive voice response for customer service, and spoken commands for smartphones. In text recognition, neural networks are used to automatically trans-

*Received by the editors June 4, 2018; accepted for publication (in revised form) November 8, 2019; published electronically March 10, 2020.

<https://doi.org/10.1137/18M1192184>

Funding: The work of the second author was partially supported by the National Science Foundation through grant DMS 1550918.

[†]Department of Industrial and Systems Engineering, University of Illinois at Urbana–Champaign, Champaign, IL 61820 (jasirign@illinois.edu).

[‡]Department of Mathematics and Statistics, Boston University, Boston, MA 02215 (kspiliop@math.bu.edu).

late text from one language (e.g., English) to another language (e.g., Italian); see [55] and [50]. They have also been used for automatically generating summaries of long documents; see [39] and [9].

In addition, there is growing interest in applying neural networks to engineering, robotics, medicine, and finance. Neural networks are being used in reduced-form models of the Navier–Stokes equation in turbulent conditions (see [34] and [35]). [51], [22], and [40] describe applications in robotics. Neural networks have been used to identify cancer [17] and to model protein folding [1]. In finance, neural networks have been used to model loan default and prepayment risk [43] and to model high frequency financial data [44]. Neural networks have also been used to solve high-dimensional PDEs in financial applications [45].

Due to the impact that neural networks have had on practical applications, there is a significant interest in better understanding their mathematical properties. However, the existing literature is relatively limited, with only a few recent papers such as [5], [36], and [52]. There also exist classical results regarding the approximation power of neural networks [4], [24], [25].

Our result characterizes neural networks with a single hidden layer in the asymptotic regime of large network sizes and large numbers of stochastic gradient descent iterations. We rigorously prove that the empirical distribution of the neural network parameters will weakly converge to a distribution. This distribution satisfies a nonlinear PDE. The proof relies upon weak convergence analysis for interacting particle systems. The result can be considered a “law of large numbers” for neural networks when both the network size and the number of stochastic gradient descent steps grow to infinity.

Recently, [54] rigorously established a weak convergence result for a class of machine learning algorithms. Weak convergence analysis has been widely used in other fields (for example, see [18], [19], [12], [13], [14], [6], and [23] for a nonexhaustive list). In fact, mean field analysis has been actively used for many years to study biological neural networks and physical systems of interacting particles; see, for example, [15], [27], [37], [48], and the references therein.

Upon completion of this work, we became aware of the very recent work of [38], where a related PDE limit result for neural networks is derived; see also the recent work [41]. Our convergence analysis, setup, and assumptions are different. In [38], it is assumed that the gradient of the neural network is a priori globally Lipschitz and bounded and under this assumption a similar PDE result as well as certain rates of convergence are established. In our work, we do not assume that the gradient of the neural network is a priori globally Lipschitz or bounded. Often, neural network models (and their gradients) are not globally Lipschitz and not bounded. In this paper, we assume that the data come from a distribution that has its first moments bounded and that the initialization is done according to distributions with certain moment bounds. Based on this assumption, we rigorously prove relative compactness of the prelimit measure-valued process, identification of its limit, and uniqueness of the limit point in the appropriate space. Our method of proof leverages on weak convergence analysis in an appropriate Skorokhod space for measure-valued processes (similar to the approaches in [54] and [18]). In particular, the relative compactness and uniqueness proof addresses the challenge of neural networks not being a priori globally Lipschitz nor globally bounded using the structure of the stochastic gradient descent algorithm.

Consider the one-layer neural network

$$(1.1) \quad g_{\theta}^N(x) = \frac{1}{N} \sum_{i=1}^N c^i \sigma(w^i \cdot x),$$

where for every $i \in \{1, \dots, N\}$, $c^i \in \mathbb{R}$ and $x, w^i \in \mathbb{R}^d$. For notational convenience we shall interpret $w^i \cdot x = \sum_{j=1}^d w^{i,j} x^j$ as the standard scalar inner product. The neural network model has parameters $\theta = (c^1, \dots, c^N, w^1, \dots, w^N) \in \mathbb{R}^{(1+d)N}$, which must be estimated from data.

The neural network (1.1) takes a linear function of the original data, applies an elementwise nonlinearity using the function $\sigma : \mathbb{R} \rightarrow \mathbb{R}$, and then takes another linear function to produce the output. The activation function $\sigma(\cdot)$ is a nonlinear function such as a sigmoid or tanh function. The quantity $\sigma(w^i \cdot x)$ is referred to as the i th “hidden unit,” and the vector $(\sigma(w^1 \cdot x), \dots, \sigma(w^N \cdot x))$ is called the “hidden layer.” The number of units in the hidden layer is N .

The objective function is

$$(1.2) \quad L^N(\theta) = \frac{1}{2} \mathbb{E}_{Y,X}[(Y - g_{\theta}^N(X))^2],$$

where the data (Y, X) is assumed to have a joint distribution $\pi(dx, dy)$. We shall write \mathcal{X}, \mathcal{Y} for the state spaces of X and Y , respectively. The parameters $\theta = (c^1, \dots, c^N, w^1, \dots, w^N)$ are estimated using stochastic gradient descent:

$$(1.3) \quad \begin{aligned} c_{k+1}^i &= c_k^i + \frac{\alpha}{N} (y_k - g_{\theta_k}^N(x_k)) \sigma(w_k^i \cdot x_k), \\ w_{k+1}^{i,j} &= w_k^{i,j} + \frac{\alpha}{N} (y_k - g_{\theta_k}^N(x_k)) c_k^i \sigma'(w_k^i \cdot x_k) x_k^j, \quad j = 1, \dots, d, \end{aligned}$$

where α is the learning rate and $(x_k, y_k) \sim \pi(dx, dy)$. Stochastic gradient descent minimizes (1.2) using a sequence of noisy (but unbiased) gradient descent steps $\nabla_{\theta}[(y_k - g_{\theta_k}^N(x_k))^2]$. Note that typically $\nabla_{\theta}[(y - g_{\theta}^N(x))^2]$ is not a priori globally Lipschitz nor globally bounded as a function of θ . Stochastic gradient descent typically converges more rapidly than gradient descent for large datasets. For this reason, stochastic gradient descent is widely used in machine learning.

Define the empirical measure

$$(1.4) \quad \nu_k^N(dc, dw) = \frac{1}{N} \sum_{i=1}^N \delta_{c_k^i, w_k^i}(dc, dw).$$

The neural network’s output can be rewritten in terms of the empirical measure:

$$(1.5) \quad g_{\theta_k}^N(x) = \langle c \sigma(w \cdot x), \nu_k^N \rangle.$$

$\langle f, h \rangle$ denotes the

The scaled empirical measure is

$$(1.6) \quad \mu_t^N = \nu_{\lfloor Nt \rfloor}^N.$$

At any time t , the scaled empirical measure μ_t^N is a random element of $D_E([0, T]) = D([0, T]; E)^1$ with $E = \mathcal{M}(\mathbb{R}^{1+d})$. We study the convergence in distribution of μ_t^N in the Skorokhod space $D_E([0, T])$.

¹ $D_S([0, T])$ is the set of maps from $[0, T]$ into S which are right-continuous and which have left-hand limits.

Our main results are stated below. Theorem 1.2 (and the associated Remark 1.3) is a law of large numbers describing the distribution of the trained parameters when N is large. Theorem 1.6 describes the behavior of individual parameters when N is large. Theorem 1.6 is a “propagation of chaos” result. Section 1.1 presents several insights provided by these asymptotic results.

We shall work on a filtered probability space $(\Omega, \mathcal{F}, \mathbb{P})$ on which all the random variables are defined. The probability space is equipped with a filtration \mathfrak{F}_t that is right continuous and \mathfrak{F}_0 contains all \mathbb{P} -negligible sets.

At this point, let us recall the definition of chaoticity. Let q be a probability measure on a Polish space \mathcal{Z} and, for $N \in \mathbb{N}$, let \mathbb{Q}^N be a symmetric probability measure on the product space \mathcal{Z}^N . Then $(\mathbb{Q}^N)_{N \in \mathbb{N}}$ is called q -chaotic if, for every $k \in \mathbb{N}$, the joint distribution law of the first k marginals of \mathbb{Q}^N converges weakly to the product measure $\otimes^k q$.

We impose the following assumption.

Assumption 1.1.

- The activation function $\sigma \in C_b^2(\mathbb{R})$, i.e., σ is twice continuously differentiable and bounded.
- The sequence of data samples (x_k, y_k) is independent and identically distributed (i.i.d.) from a probability distributed $\pi(dx, dy)$ such that $\mathbb{E} \|x_k\|^4 + \mathbb{E}|y_k|^4$ is bounded.
- The randomly initialized parameters (c_0^i, w_0^i) are i.i.d. with a distribution $\bar{\mu}_0$ such that $\mathbb{E}[\exp(q|c_0^i|)] < C$ for some $0 < q < \infty$ and $\mathbb{E}[\|w_0^i\|^4] < C$.

Notice that initial distributions on c^i with compact support or exponential tails satisfy the condition on the moment generating function of $|c_0^i|$ in the assumption above. Under Assumption 1.1, the initial empirical measure satisfies $\mu_0^N \xrightarrow{d} \bar{\mu}_0$ as $N \rightarrow \infty$. In addition, due to our assumption on the distribution of the (x_k, y_k) data and of the initialization $(c_0^i, w_0^i)_{i=1}^N$, the joint distribution of $(c_k^i, w_k^i)_{i=1}^N \in (\mathbb{R}^{1+d})^{\otimes N}$ is exchangeable and, consequently, ν_k^N is a Markov chain in the space of probability measures on E .

THEOREM 1.2. *Assume Assumption 1.1. The scaled empirical measure μ_t^N converges in distribution to a limit measure $\bar{\mu}_t$ with values in $D_E([0, T])$ as $N \rightarrow \infty$. For every $f \in C_b^2(\mathbb{R}^{1+d})$, $\bar{\mu}$ is the unique deterministic solution of the measure evolution equation*

$$\begin{aligned} & \langle f, \bar{\mu}_t \rangle \\ &= \langle f, \bar{\mu}_0 \rangle + \int_0^t \left(\int_{\mathcal{X} \times \mathcal{Y}} \alpha(y - \langle c' \sigma(w' \cdot x), \bar{\mu}_s \rangle) \langle \sigma(w \cdot x) \partial_c f, \bar{\mu}_s \rangle \pi(dx, dy) \right) ds \\ & \quad + \int_0^t \left(\int_{\mathcal{X} \times \mathcal{Y}} \alpha(y - \langle c' \sigma(w' \cdot x), \bar{\mu}_s \rangle) \langle c \sigma'(w \cdot x) x \cdot \nabla_w f, \bar{\mu}_s \rangle \pi(dx, dy) \right) ds \\ &= \langle f, \bar{\mu}_0 \rangle + \int_0^t \left(\int_{\mathcal{X} \times \mathcal{Y}} \alpha(y - \langle c' \sigma(w' \cdot x), \bar{\mu}_s \rangle) \langle \nabla(c \sigma(w \cdot x)) \cdot \nabla f, \bar{\mu}_s \rangle \pi(dx, dy) \right) ds, \end{aligned}$$

where $\nabla f = (\partial_c f, \nabla_w f)$.

Remark 1.3. Since weak convergence to a constant implies convergence in probability, Theorem 1.2 leads to the stronger result of convergence in probability

$$\lim_{N \rightarrow \infty} \mathbb{P} \{d_E(\mu^N, \bar{\mu}) \geq \delta\} = 0$$

for every $\delta > 0$ and where d_E is the metric for $D_E([0, T])$.

COROLLARY 1.4. Assume Assumption 1.1. Suppose that $\bar{\mu}_0$ admits a density $p_0(c, w)$ and there exists a unique solution to the nonlinear PDE

$$\begin{aligned} & \frac{\partial p(t, c, w)}{\partial t} \\ &= -\alpha \int_{\mathcal{X} \times \mathcal{Y}} \left((y - \langle c' \sigma(w' \cdot x), p(t, c', w') \rangle) \frac{\partial}{\partial c} [\sigma(w \cdot x) p(t, c, w)] \right) \pi(dx, dy) \\ & \quad - \alpha \int_{\mathcal{X} \times \mathcal{Y}} \left((y - \langle c' \sigma(w' \cdot x), p(t, c', w') \rangle) x \cdot \nabla_w [c \sigma'(w \cdot x) p(t, c, w)] \right) \pi(dx, dy), \\ & p(0, c, w) = p_0(c, w), \end{aligned}$$

such that $p(t, c, w)$ vanishes as $|c|, \|w\| \rightarrow \infty$. Then, we have that the solution to the measure evolution equation (1.7) is such that

$$\bar{\mu}_t(dc, dw) = p(t, c, w) dcdw.$$

Remark 1.5. Notice that by setting $\theta = (c, w)$, the PDE for $p(t, \theta)$ in Corollary 1.4 can be written as

$$\begin{aligned} & \frac{\partial p(t, \theta)}{\partial t} = -\alpha \operatorname{div}_{\theta} (p(t, \theta) \nabla_{\theta} v(\theta, p(t, \cdot))), \\ (1.7) \quad & p(0, \theta) = p_0(\theta), \end{aligned}$$

where $\operatorname{div}_{\theta}$ is the divergence operator with respect to the variable θ and $v(\theta, p(t, \cdot))$ is defined as

$$v(\theta, p(t, \cdot)) = \int_{\mathcal{X} \times \mathcal{Y}} \left(\left(y - \int_{\mathbb{R}^{1+d}} c' \sigma(w' \cdot x) p(t, c', w') dc' dw' \right) c \sigma(w \cdot x) \right) \pi(dx, dy).$$

In addition, Theorem 1.2 and Corollary 1.4 imply that the objective function $L^N(\theta)$ from (1.2) satisfies

$$\begin{aligned} (1.8) \quad & \lim_{N \rightarrow \infty} L^N(\theta_{[Nt]}) = \bar{L}(p(t, \cdot)) = \frac{1}{2} \int_{\mathcal{X} \times \mathcal{Y}} (y - \bar{g}(x, p(t, \cdot)))^2 \pi(dx, dy), \quad \text{where} \\ & \bar{g}(x, p(t, \cdot)) = \int_{\mathbb{R}^{1+d}} c \sigma(w \cdot x) p(t, c, w) dcdw. \end{aligned}$$

As it is well known in the literature (see, for example, [2], [10], [28]), the PDE (1.7) is a gradient flow for the limiting objective function (1.8) in the space of probability measures on \mathbb{R}^{1+d} endowed with the Wasserstein metric. As also extensively discussed in [2], but also pointed out in the recent works by [11], [38], [26], this means that the trajectory's $t \mapsto p(t, \cdot)$ goal is to minimize the limit objective function $\bar{L}(p)$ as defined by (1.8). For more details on the related optimal transportation theory we refer the interested reader to Lemma 10.3.16 in [2] for general results and to Proposition 4.2 in [11] and Corollary 3.3 in [26] for a more specialized discussion in the special case of limiting problems of the type (1.7)–(1.8).

In Theorem 1.6 we prove that the neural network has the “propagation of chaos” property.

THEOREM 1.6. Assume Assumption 1.1. Consider $T < \infty$ and let $t \in (0, T]$. Define the probability measure $\rho_t^N \in \mathcal{M}(\mathbb{R}^{(1+d)N})$ where

$$\rho_t^N(dx^1, \dots, dx^N) = \mathbb{P}[(c_{[Nt]}^1, w_{[Nt]}^1) \in dx^1, \dots, (c_{[Nt]}^N, w_{[Nt]}^N) \in dx^N].$$

Then, the sequence of probability measures ρ^N is $\bar{\mu}$ -chaotic. That is, for $k \in \mathbb{N}$

$$(1.9) \quad \begin{aligned} & \lim_{N \rightarrow \infty} \langle f_1(x^1) \times \cdots \times f_k(x^k), \rho^N(dx^1, \dots, dx^N) \rangle \\ &= \prod_{i=1}^k \langle f_i, \bar{\mu} \rangle \quad \forall f_1, \dots, f_k \in C_b^2(\mathbb{R}^{1+d}). \end{aligned}$$

1.1. Insights from law of large numbers and numerical studies. The law of large numbers (1.7) suggests several interesting characteristics of trained neural networks (at least in the setting studied in this paper).

- As $N \rightarrow \infty$, the neural network converges (in probability) to a deterministic model. This is despite the fact that the neural network is randomly initialized and it is trained on a random sequence of data samples via stochastic gradient descent.
- The learning rate α was assumed to be constant and to not decay with time. However, notice that the hidden layer has been normalized by $1/N$ and it is this normalization by $1/N$ in the hidden layer that replaces the role of the learning rate decay, enabling convergence.
- As also discussed in Remark 1.5, the PDE (1.7) is a gradient flow for the limiting objective function (1.8) in the space of probability measures on \mathbb{R}^{1+d} endowed with the Wasserstein metric. Hence, the limiting law of large numbers goal is to minimize the limit objective function $\bar{L}(p)$ as defined by (1.8).
- The propagation of chaos result (1.9) indicates that, as $N \rightarrow \infty$, the dynamics of the weights (c_k^i, w_k^i) will become independent of the dynamics of the weights (c_k^j, w_k^j) for any $i \neq j$. Note that the dynamics (c_k^i, w_k^i) are still random due to the random initialization. However, the dynamics of the i th set of weights will be uncorrelated with the dynamics of the j th set of weights in the limit as $N \rightarrow \infty$.

In order to illustrate some aspects of the theoretical results of this paper, we performed the following numerical study.

Figure 1 displays the convergence of the distribution of the parameters in a trained neural network as the number of hidden units $N \rightarrow \infty$. The neural network has a single hidden layer followed by a softmax function. Figure 1 reports the distribution of the parameters connecting the hidden layer to the softmax function. The distributions are presented as histograms. The neural network is trained on the MNIST dataset, which is a standard image dataset in machine learning [32]. The dataset includes 60,000 images of handwritten numbers $\{0, 1, 2, \dots, 9\}$. The neural network is trained to identify the handwritten numbers using only the image pixels as an input (i.e., it learns to recognize images as a human would). In the MNIST dataset, each image has 784 pixels. A pixel takes values in $\{0, 1, \dots, 255\}$.² Neural networks can achieve 98% to 99% out-of-sample accuracy on the MNIST dataset.

Figure 1 shows that the distribution of parameters converges to a fixed distribution as $N \rightarrow \infty$. This can be seen by the fact that the distributions for $N = 10,000$, $N = 100,000$, and $N = 250,000$ are nearly identical. A priori it is unclear if the distribution of neural network parameters should converge as $N \rightarrow \infty$. Our theory and numerical results confirm that this is indeed the case. Indeed, as N gets large, we see that the empirical distribution of the parameters connecting the hidden layer to the softmax function converges to a specific deterministic distribution.

²The pixel values are normalized to $[0, 1]$ for the purposes of training the neural network.




FIG. 1. *Distribution of parameters for a neural network trained on MNIST dataset. Clockwise: $N = 1,000$, $N = 10,000$, $N = 100,000$, and $N = 250,000$ hidden units.*

1.2. Overview of the proof. The rest of the paper is organized as follows. Section 2 proves relative compactness of the family $\{\mu^N\}_{N \in \mathbb{N}}$. Section 3 identifies the limit point of any convergent subsequence. The limit point must satisfy the measure evolution equation (1.7). Section 4 proves uniqueness of the evolution equation (1.7) via a fixed point argument. Then, by Prokhorov's theorem, these results prove that the sequence of probability measures π^N of the processes μ^N weakly converge to π , the probability measure of the process $\bar{\mu}$ satisfying (1.7). These results are collected together in section 5 to prove Theorem 1.2, Corollary 1.4, and Theorem 1.6. We conclude with a discussion of our results in section 6.

2. Relative compactness. We now prove relative compactness of the family $\{\mu^N\}_{N \in \mathbb{N}}$ in $D_E([0, T])$, where $E = \mathcal{M}(\mathbb{R}^{1+d})$. It is sufficient to show compact containment and regularity of the μ^N 's (see, for example, Chapter 3 of [16]). We start by proving in the following lemma a crucial a priori bound for the stochastic gradient descent iterates.

LEMMA 2.1. *Consider the system (1.3). Then, for $k \leq TN$ and uniformly in $i \in \mathbb{N}$, there exists a constant $C < \infty$ such that*

$$\mathbb{E} [|c_k^i| + \|w_k^i\|] \leq C.$$

In particular, we have

$$\sup_{N \in \mathbb{N}, k/N \leq T} \frac{1}{N} \sum_{i=1}^N \mathbb{E} [|c_k^i| + \|w_k^i\|] \leq C.$$

For the purposes of presentation, the proof of Lemma 2.1 will be given at the end of this section. First, we prove compact containment for the measure-valued process $\{\mu_t^N, t \in [0, T]\}_{N \in \mathbb{N}}$.

LEMMA 2.2. For each $\eta > 0$, there is a compact subset \mathcal{K} of E such that

$$\sup_{N \in \mathbb{N}, 0 \leq t \leq T} \mathbb{P}[\mu_t^N \notin \mathcal{K}] < \eta.$$

Proof. For each $L > 0$, define $K_L = [-L, L]^{1+d}$. Then, we have that K_L is a compact subset of \mathbb{R}^{1+d} , and for each $t \geq 0$ and $N \in \mathbb{N}$,

$$\mathbb{E}[\mu_t^N(\mathbb{R}^{1+d} \setminus K_L)] = \frac{1}{N} \sum_{i=1}^N \mathbb{P}[|c_{[Nt]}^i| + \|w_{[Nt]}^i\| \geq L] \leq \frac{C}{L},$$

where the constant $C < \infty$ is from Lemma 2.1. The rest of the proof is standard now (see, for example, Lemma 6.1 of [18]). We define the compact subsets of $E = \mathcal{M}(\mathbb{R}^{1+d})$

$$\hat{K}_L = \left\{ \nu : \nu(\mathbb{R}^{1+d} \setminus K_{(L+j)^2}) < \frac{1}{\sqrt{L+j}} \quad \forall j \in \mathbb{N} \right\}$$

and we observe that

$$\begin{aligned} \mathbb{P}[\mu_t^N \notin \hat{K}_L] &\leq \sum_{j=1}^{\infty} \mathbb{P}\left[\mu_t^N(\mathbb{R}^{1+d} \setminus K_{(L+j)^2}) > \frac{1}{\sqrt{L+j}}\right] \leq \sum_{j=1}^{\infty} \frac{\mathbb{E}[\mu_t^N(\mathbb{R}^{1+d} \setminus K_{(L+j)^2})]}{1/\sqrt{L+j}} \\ &\leq \sum_{j=1}^{\infty} \frac{C}{(L+j)^2/\sqrt{L+j}} \leq \sum_{j=1}^{\infty} \frac{C}{(L+j)^{3/2}}. \end{aligned}$$

Given now that $\lim_{L \rightarrow \infty} \sum_{j=1}^{\infty} \frac{C}{(L+j)^{3/2}} = 0$, the proof of the lemma is concluded. \square

We now establish regularity of the μ^N 's. Define the function $q(z_1, z_2) = \min\{|z_1 - z_2|, 1\}$, where $z_1, z_2 \in \mathbb{R}$.

LEMMA 2.3. Let $f \in C_b^2(\mathbb{R}^{1+d})$. For any $p \in (0, 1)$, there is a constant $C < \infty$ such that for $0 \leq u \leq \delta$, $0 \leq v \leq \delta \wedge t$, $t \in [0, T]$,

$$\mathbb{E}[q(\langle f, \mu_{t+u}^N \rangle, \langle f, \mu_t^N \rangle) q(\langle f, \mu_t^N \rangle, \langle f, \mu_{t-v}^N \rangle) | \mathcal{F}_t^N] \leq C\delta^p + \frac{C}{N}.$$

Proof. We start by noticing that a Taylor expansion gives for $0 \leq s \leq t \leq T$

$$\begin{aligned} |\langle f, \mu_t^N \rangle - \langle f, \mu_s^N \rangle| &= |\langle f, \mu_t^N \rangle - \langle f, \mu_s^N \rangle| = \left| \left\langle f, \nu_{[Nt]}^N \right\rangle - \left\langle f, \nu_{[Ns]}^N \right\rangle \right| \\ &\leq \frac{1}{N} \sum_{i=1}^N |f(c_{[Nt]}^i, w_{[Nt]}^i) - f(c_{[Ns]}^i, w_{[Ns]}^i)| \\ &\leq \frac{1}{N} \sum_{i=1}^N |\partial_c f(\bar{c}_{[Nt]}^i, \bar{w}_{[Nt]}^i)| |c_{[Nt]}^i - c_{[Ns]}^i| \\ &\quad + \frac{1}{N} \sum_{i=1}^N \|\nabla_w f(\bar{c}_{[Nt]}^i, \bar{w}_{[Nt]}^i)\| \|w_{[Nt]}^i - w_{[Ns]}^i\| \end{aligned} \tag{2.1}$$

for points \bar{c}^i, \bar{w}^i in the segments connecting $c_{[Ns]}^i$ with $c_{[Nt]}^i$ and $w_{[Ns]}^i$ with $w_{[Nt]}^i$, respectively.

Let's now establish a bound on $|c_{[Nt]}^i - c_{[Ns]}^i|$ for $s < t \leq T$. Let $0 < p < 1$.

$$\begin{aligned}
\mathbb{E}|c_{[Nt]}^i - c_{[Ns]}^i| &= \mathbb{E} \left| \sum_{k=[Ns]}^{[Nt]-1} (c_{k+1} - c_k) \right| \\
&\leq \mathbb{E} \sum_{k=[Ns]}^{[Nt]-1} \left| \alpha(y_k - g_{\theta_k}^N(x_k)) \frac{1}{N} \sigma(w_k^i \cdot x_k) \right| \\
&\leq \frac{1}{N} \sum_{k=[Ns]}^{[Nt]-1} C \leq C(t-s) + \frac{C}{N} \\
&\leq C(t-s)^p \mathbf{1}_{t-s < 1} + C(t-s)^p T^{1/p} \mathbf{1}_{t-s \geq 1} + \frac{C}{N} \\
&\leq C(t-s)^p + \frac{C}{N},
\end{aligned}$$

where Assumption 1.1 was used. Let's now establish a bound on $\|w_{[Nt]}^i - w_{[Ns]}^i\|$ for $s < t \leq T$. Making use of the uniform bounds established in Lemma 2.1, we obtain similarly to the previous bound

$$\begin{aligned}
\mathbb{E} \|w_{[Nt]}^i - w_{[Ns]}^i\| &= \mathbb{E} \left\| \sum_{k=[Ns]}^{[Nt]-1} (w_{k+1} - w_k) \right\| \\
&\leq \mathbb{E} \sum_{k=[Ns]}^{[Nt]-1} \left\| \alpha(y_k - g_{\theta_k}^N(x_k)) \frac{1}{N} c_k^i \sigma'(w_k^i \cdot x_k) x_k \right\| \\
&\leq \frac{1}{N} \sum_{k=[Ns]}^{[Nt]-1} C \\
&\leq C(t-s) + \frac{C}{N} \leq C(t-s)^p + \frac{C}{N}.
\end{aligned}$$

Now, we return to (2.1). By Lemma 2.1, the quantities $(\bar{c}_{[Nt]}^i, \bar{w}_{[Nt]}^i)$ are bounded in expectation for $0 < s < t \leq T$. Therefore, for $0 < s < t \leq T$,

$$\mathbb{E} [|\langle f, \mu_t^N \rangle - \langle f, \mu_s^N \rangle| | \mathcal{F}_s^N] \leq C(t-s)^p + \frac{C}{N}.$$

where $C < \infty$ is some unimportant constant. Then, the statement of the lemma follows. \square

We can now prove the required relative compactness of the sequence $\{\mu^N\}_{N \in \mathbb{N}}$. This implies that every subsequence μ^N 's has a convergent subsubsequence.

LEMMA 2.4. *The sequence of probability measures $\{\mu^N\}_{N \in \mathbb{N}}$ is relatively compact in $D_E([0, T])$.*

Proof. Given Lemmas 2.2 and 2.3, Theorem 8.6 of Chapter 3 of [16] gives the statement of the lemma. (See also Remark 8.7 B of Chapter 3 of [16] regarding replacing \sup_N with \lim_N in the regularity condition B of Theorem 8.6.) \square

We conclude this section with the proof of the a priori bound of Lemma 2.1.

Proof of Lemma 2.1. We start by establishing some useful a priori bounds on c_k^i and w_k^i . The unimportant finite constant $C < \infty$ may change from line to line. We first observe that

$$\begin{aligned} |c_{k+1}^i| &\leq |c_k^i| + \alpha |y_k - g_{\theta_k}^N(x_k)| \frac{1}{N} |\sigma(w_k^i \cdot x_k)| \\ &\leq |c_k^i| + \frac{\alpha C |y_k|}{N} + \frac{C}{N^2} \sum_{i=1}^N |c_k^i|, \end{aligned}$$

where to derive the last line we used the definition of $g_{\theta_k}^N(x)$ via (1.1) and the uniform boundedness assumption on σ . Then, we subsequently obtain that

$$\begin{aligned} |c_k^i| &= |c_0^i| + \sum_{j=1}^k [|c_j^i| - |c_{j-1}^i|] \\ &\leq |c_0^i| + \sum_{j=1}^k \frac{\alpha C |y_{j-1}|}{N} + \frac{C}{N^2} \sum_{j=1}^k \sum_{i=1}^N |c_{j-1}^i|. \end{aligned}$$

This implies that

$$\frac{1}{N} \sum_{i=1}^N |c_k^i| \leq \frac{1}{N} \sum_{i=1}^N |c_0^i| + \sum_{j=1}^k \frac{\alpha C |y_{j-1}|}{N} + \frac{C}{N^2} \sum_{j=1}^k \sum_{i=1}^N |c_{j-1}^i|,$$

Let us now define $m_k^N = \frac{1}{N} \sum_{i=1}^N |c_k^i|$ and $b_k^N = \frac{1}{N} \sum_{i=1}^N |c_0^i| + \sum_{j=1}^k \frac{\alpha C |y_{j-1}|}{N}$. Then we have

$$m_k^N \leq b_k^N + \frac{C}{N} \sum_{j=1}^k m_j^N,$$

which by the discrete Gronwall lemma gives the bound

$$m_k^N \leq b_k^N + \frac{C}{N} \sum_{j=1}^k b_j^N e^{C \frac{j-i}{N}} \leq b_k^N + \frac{C}{N} \sum_{j=1}^k b_j^N$$

for a possibly different constant that may depend on T , where the relation $k/N \leq T$ was used in the last step. Going back now to the bound for c_k^i we obtain

$$|c_k^i| \leq |c_0^i| + b_k^N + \frac{C}{N} \sum_{j=1}^k m_j^N.$$

Raising this to power $1 \leq p \leq 4$, we have for a constant C_p that may depend on p

$$|c_k^i|^p \leq C_p \left[|c_0^i|^p + |b_k^N|^p + \frac{1}{N^p} \left| \sum_{j=1}^k b_j^N \right|^p + \frac{1}{N^{2p}} \left| \sum_{j=1}^k \sum_{i=1}^j b_i^N \right|^p \right].$$

Let us bound now each of the terms on the right-hand side of the last display. We have for some constant $C_p < \infty$ that may change from line to line

$$\begin{aligned} |b_k^N|^p &\leq C_p \left[\frac{1}{N} \sum_{i=1}^N |c_0^i|^p + \frac{k^{p-1}}{N^p} \sum_{j=1}^k |y_{j-1}|^p \right], \\ \frac{1}{N^p} \left| \sum_{j=1}^k b_j^N \right|^p &\leq \frac{k^{p-1}}{N^p} \sum_{j=1}^k |b_j^N|^p \leq C_p \frac{k^{p-1}}{N^p} \sum_{j=1}^k \left[\frac{1}{N} \sum_{i=1}^N |c_0^i|^p + \frac{j^{p-1}}{N^p} \sum_{i=1}^j |y_{i-1}|^p \right], \\ \frac{1}{N^{2p}} \left| \sum_{j=1}^k \sum_{i=1}^j b_i^N \right|^p &\leq \frac{k^{p-1}}{N^p} \sum_{j=1}^k \frac{j^{p-1}}{N^p} \sum_{i=1}^j |b_i^N|^p \\ &\leq C_p \frac{k^{p-1}}{N^p} \sum_{j=1}^k \frac{j^{p-1}}{N^p} \sum_{i=1}^j \left[\frac{1}{N} \sum_{i=1}^N |c_0^i|^p + \frac{j^{p-1}}{N^p} \sum_{\lambda=1}^i |y_{\lambda-1}|^p \right]. \end{aligned}$$

Plugging those bounds now into the previous bound for $|c_k^i|^p$, using the assumption $E|y_i|^p \leq C < \infty$ for all i and all $p \in [1, 4]$ and that $k/N \leq T$, we obtain that for all $i \in \mathbb{N}$ and all k such that $k/N \leq T$, the bound

$$(2.2) \quad \mathbb{E}|c_k^i|^p \leq C < \infty$$

for some constant C that may depend on p, T , and the bound on the activation function σ . We have also used the fact that $\mathbb{E}[|c_0^i|^p] < K_1$ due to $E[\exp(q|c_0|)] < K_2$ for some $0 < q < \infty$ in Assumption 1.1. Now, we turn to the bound for $\|w_k^i\|$. We start with the bound (using Young's inequality)

$$\begin{aligned} \|w_{k+1}^i\| &\leq \|w_k^i\| + \frac{C}{N} \left(|y_k| + \frac{1}{N} \sum_{i=1}^N |c_k^i| \right) |c_k^i| |\sigma'(w_k^i \cdot x_k)| \|x_k\| \\ &\leq \|w_k^i\| + C \left(\frac{1}{N} |y_k|^2 + \frac{1}{N^2} \sum_{i=1}^N |c_k^i|^2 + \frac{1}{N} |c_k^i|^2 \|x_k\|^2 \right) \\ &\leq \|w_k^i\| + C \left(\frac{1}{N} |y_k|^2 + \frac{1}{N^2} \sum_{i=1}^N |c_k^i|^2 + \frac{1}{N} |c_k^i|^4 + \frac{1}{N} \|x_k\|^4 \right) \end{aligned}$$

for a constant $C < \infty$ that may change from line to line. Taking now expectation, using Assumption 1.1, the a priori bound (2.2), and the fact that $k/N \leq T$ we obtain

$$\mathbb{E} \|w_k^i\| \leq C < \infty$$

for all $i \in \mathbb{N}$ and all k such that $k/N \leq T$, concluding the proof of the lemma. \square

3. Identification of the limit. We consider the evolution of the empirical measure ν_k^N via test functions $f \in C_b^2(\mathbb{R}^{1+d})$. A Taylor expansion yields

$$\begin{aligned} & \langle f, \nu_{k+1}^N \rangle - \langle f, \nu_k^N \rangle \\ &= \frac{1}{N} \sum_{i=1}^N \left(f(c_{k+1}^i, w_{k+1}^i) - f(c_k^i, w_k^i) \right) \\ &= \frac{1}{N} \sum_{i=1}^N \partial_c f(c_k^i, w_k^i) (c_{k+1}^i - c_k^i) + \frac{1}{N} \sum_{i=1}^N \nabla_w f(c_k^i, w_k^i) (w_{k+1}^i - w_k^i) \\ &\quad + \frac{1}{N} \sum_{i=1}^N \partial_c^2 f(\bar{c}_k^i, \bar{w}_k^i) (c_{k+1}^i - c_k^i)^2 + \frac{1}{N} \sum_{i=1}^N (c_{k+1}^i - c_k^i) \nabla_{cw} f(\bar{c}_k^i, \bar{w}_k^i) (w_{k+1}^i - w_k^i) \\ &\quad + \frac{1}{N} \sum_{i=1}^N (w_{k+1}^i - w_k^i)^\top \nabla_w^2 f(\bar{c}_k^i, \bar{w}_k^i) (w_{k+1}^i - w_k^i) \end{aligned}$$

for points \bar{c}_k^i, \bar{w}_k^i in the segments connecting c_{k+1}^i with c_k^i and w_{k+1}^i with w_k^i , respectively. Notice now that the uniform bounds of Lemma 2.1 and the relation (1.3) imply that as N gets large

$$\begin{aligned} & \langle f, \nu_{k+1}^N \rangle - \langle f, \nu_k^N \rangle \\ &= \frac{1}{N^2} \sum_{i=1}^N \partial_c f(c_k^i, w_k^i) \alpha(y_k - g_{\theta_k}^N(x_k)) \sigma(w_k^i \cdot x_k) \\ &\quad + \frac{1}{N^2} \sum_{i=1}^N \alpha(y_k - g_{\theta_k}^N(x_k)) c_k^i \sigma'(w_k^i \cdot x_k) \nabla_w f(c_k^i, w_k^i) \cdot x_k + O_p(N^{-2}). \end{aligned}$$

The term $O_p(N^{-2})$ is a result of $f \in C_b^2$, the bounds from Lemma 2.1 as well as the moment bounds on (x_k, y_k) from Assumption 1.1.³ We next define the drift and martingale components:

$$\begin{aligned} D_k^{1,N} &= \frac{1}{N} \int_{\mathcal{X} \times \mathcal{Y}} \alpha(y - \langle c\sigma(w \cdot x), \nu_k^N \rangle) \langle \sigma(w \cdot x) \partial_c f, \nu_k^N \rangle \pi(dx, dy), \\ D_k^{2,N} &= \frac{1}{N} \int_{\mathcal{X} \times \mathcal{Y}} \alpha(y - \langle c\sigma(w \cdot x), \nu_k^N \rangle) \langle c\sigma'(w \cdot x) x \cdot \nabla_w f, \nu_k^N \rangle \pi(dx, dy), \\ M_k^{1,N} &= \frac{1}{N} \alpha(y_k - \langle c\sigma(w \cdot x_k), \nu_k^N \rangle) \langle \sigma(w \cdot x_k) \nabla_c f, \nu_k^N \rangle - D_k^{1,N}, \\ M_k^{2,N} &= \frac{1}{N} \alpha(y_k - \langle c\sigma(w \cdot x_k), \nu_k^N \rangle) \langle c\sigma'(w \cdot x_k) x \cdot \nabla_w f, \nu_k^N \rangle - D_k^{2,N}. \end{aligned}$$

Combining the different terms together, we then obtain

$$\langle f, \nu_{k+1}^N \rangle - \langle f, \nu_k^N \rangle = D_k^{1,N} + D_k^{2,N} + M_k^{1,N} + M_k^{2,N} + O_p(N^{-2}).$$

Next, we define the scaled versions of $D^{1,N}, D^{2,N}, M^{1,N}$, and $M^{2,N}$:

$$\begin{aligned} D^{1,N}(t) &= \sum_{k=0}^{\lfloor Nt \rfloor - 1} D_k^{1,N}, & D^{2,N}(t) &= \sum_{k=0}^{\lfloor Nt \rfloor - 1} D_k^{2,N}, \\ M^{1,N}(t) &= \sum_{k=0}^{\lfloor Nt \rfloor - 1} M_k^{1,N}, & M^{2,N}(t) &= \sum_{k=0}^{\lfloor Nt \rfloor - 1} M_k^{2,N}. \end{aligned}$$

³Recall that when we write $Z = O_p(b)$ we mean that Z/b is stochastically bounded.

The scaled empirical measure satisfies, as N grows,

$$\begin{aligned} & \langle f, \mu_t^N \rangle - \langle f, \mu_0^N \rangle \\ &= \int_0^t \left(\int_{\mathcal{X} \times \mathcal{Y}} \alpha(y - \langle c\sigma(w \cdot x), \mu_s^N \rangle) \langle \sigma(w \cdot x) \nabla_c f, \mu_s^N \rangle \pi(dx, dy) \right) ds \\ &+ \int_0^t \left(\int_{\mathcal{X} \times \mathcal{Y}} \alpha(y - \langle c\sigma(w \cdot x), \mu_s^N \rangle) \langle c\sigma'(w \cdot x)x \cdot \nabla_w f, \mu_s^N \rangle \pi(dx, dy) \right) ds \\ &+ M^{1,N}(t) + M^{2,N}(t) + O_p(N^{-1}). \end{aligned}$$

In fact as we show below $M^{1,N}(t)$ and $M^{2,N}(t)$ converge to 0 in L^2 as $N \rightarrow \infty$.

LEMMA 3.1. *We have that*

$$\begin{aligned} \lim_{N \rightarrow \infty} \mathbb{E} \left[\left(M^{1,N}(t) \right)^2 \right] &= 0, \\ \lim_{N \rightarrow \infty} \mathbb{E} \left[\left(M^{2,N}(t) \right)^2 \right] &= 0. \end{aligned}$$

Proof. First, notice that

$$\begin{aligned} & \mathbb{E} \left[\left(\sum_{k=0}^{\lfloor Nt \rfloor - 1} \frac{1}{N} \alpha(y_k - \langle c\sigma(w \cdot x_k), \nu_k^N \rangle) \langle \sigma(w \cdot x_k) \partial_c f, \nu_k^N \rangle - D_k^{1,N} \right)^2 \right] \\ &= \sum_{j,k=0}^{\lfloor Nt \rfloor - 1} \mathbb{E} \left[\left(\frac{1}{N} \alpha(y_k - \langle c\sigma(w \cdot x_k), \nu_k^N \rangle) \langle \sigma(w \cdot x_k) \partial_c f, \nu_k^N \rangle - D_k^{1,N} \right) \right. \\ (3.1) \quad & \left. \times \left(\frac{1}{N} \alpha(y_j - \langle c\sigma(w \cdot x_j), \nu_j^N \rangle) \langle \sigma(w \cdot x_j) \partial_c f, \nu_j^N \rangle - D_j^{1,N} \right) \right]. \end{aligned}$$

Let \mathcal{F}_k^N be the σ -algebra generated by $(c_0^i, w_0^i)_{i=1}^N$ and $(x_j, y_j)_{j=0}^{k-1}$. If $j > k$, then

$$\begin{aligned} & \mathbb{E} \left[\left(\frac{1}{N} \alpha(y_k - \langle c\sigma(w \cdot x_k), \nu_k^N \rangle) \langle \sigma(w \cdot x_k) \partial_c f, \nu_k^N \rangle - D_k^{1,N} \right) \right. \\ & \times \left. \left(\frac{1}{N} \alpha(y_j - \langle c\sigma(w \cdot x_j), \nu_j^N \rangle) \langle \sigma(w \cdot x_j) \partial_c f, \nu_j^N \rangle - D_j^{1,N} \right) \right] \\ &= \mathbb{E} \left[\left(\frac{1}{N} \alpha(y_k - \langle c\sigma(w \cdot x_k), \nu_k^N \rangle) \langle \sigma(w \cdot x_k) \partial_c f, \nu_k^N \rangle - D_k^{1,N} \right) \right. \\ & \times \mathbb{E} \left[\left(\frac{1}{N} \alpha(y_j - \langle c\sigma(w \cdot x_j), \nu_j^N \rangle) \langle \sigma(w \cdot x_j) \partial_c f, \nu_j^N \rangle - D_j^{1,N} \right) \middle| \mathcal{F}_{j-1}^N \right] \Bigg] \\ &= \mathbb{E} \left[\left(\frac{1}{N} \alpha(y_k - \langle c\sigma(w \cdot x_k), \nu_k^N \rangle) \langle \sigma(w \cdot x_k) \partial_c f, \nu_k^N \rangle - D_k^{1,N} \right) \times 0 \right] \\ &= 0. \end{aligned}$$

Therefore, (3.1) reduces to

$$(3.2) \quad \mathbb{E} \left[\left(\sum_{k=0}^{\lfloor Nt \rfloor - 1} \frac{1}{N} \alpha(y_k - \langle c\sigma(w \cdot x_k), \nu_k^N \rangle) \langle \sigma(w \cdot x_k) \partial_c f, \nu_k^N \rangle - D_k^{1,N} \right)^2 \right] \\ = \sum_{k=0}^{\lfloor Nt \rfloor - 1} \mathbb{E} \left[\left(\frac{1}{N} \alpha(y_k - \langle c\sigma(w \cdot x_k), \nu_k^N \rangle) \langle \sigma(w \cdot x_k) \partial_c f, \nu_k^N \rangle - D_k^{1,N} \right)^2 \right].$$

Using (3.2), we have that

$$\mathbb{E} \left[\left(M^{1,N}(t) \right)^2 \right] \\ = \mathbb{E} \left[\left(\sum_{k=0}^{\lfloor Nt \rfloor - 1} \frac{1}{N} \alpha(y_k - \langle c\sigma(w \cdot x_k), \nu_k^N \rangle) \langle \sigma(w \cdot x_k) \partial_c f, \nu_k^N \rangle - D_k^{1,N} \right)^2 \right] \\ = \sum_{k=0}^{\lfloor Nt \rfloor - 1} \mathbb{E} \left[\left(\frac{1}{N} \alpha(y_k - \langle c\sigma(w \cdot x_k), \nu_k^N \rangle) \langle \sigma(w \cdot x_k) \partial_c f, \nu_k^N \rangle - D_k^{1,N} \right)^2 \right] \\ \leq \frac{2}{N^2} \sum_{k=0}^{\lfloor Nt \rfloor - 1} \mathbb{E} \left[\left(\alpha(y_k - \langle c\sigma(w \cdot x_k), \nu_k^N \rangle) \langle \sigma(w \cdot x_k) \partial_c f, \nu_k^N \rangle \right)^2 \right] \\ + \frac{2}{N^2} \sum_{k=0}^{\lfloor Nt \rfloor - 1} \mathbb{E} \left[\left(\int_{\mathcal{X} \times \mathcal{Y}} \alpha(y - \langle c\sigma(w \cdot x), \nu_k^N \rangle) \langle \sigma(w \cdot x) \partial_c f, \nu_k^N \rangle \pi(dx, dy) \right)^2 \right] \\ \leq \frac{C}{N^2} \lfloor Nt \rfloor.$$

The final inequality comes from the bounds proven in section 2 and Assumption 1.1. A similar bound can also be established for $\mathbb{E}[(M^{2,N}(t))^2]$. The result directly follows. \square

Let π^N be the probability measure of a convergent subsequence of $(\mu^N)_{0 \leq t \leq T}$. Each π^N takes values in the set of probability measures $\mathcal{M}(D_E([0, T]))$. Relative compactness, proven in section 2, implies that there is a subsequence π^{N_k} which weakly converges. We must prove that any limit point π of a convergent subsequence π^{N_k} will satisfy the evolution equation (1.7).

LEMMA 3.2. *Let π^{N_k} be a convergent subsequence with a limit point π . Then π is a Dirac measure concentrated on $\bar{\mu} \in D_E([0, T])$ and $\bar{\mu}$ satisfies the measure evolution equation (1.7).*

Proof. We define a map $F(\mu) : D_E([0, T]) \rightarrow \mathbb{R}_+$ for each $t \in [0, T]$, $f \in C_b^2(\mathbb{R}^{1+d})$, $g_1, \dots, g_p \in C_b(\mathbb{R}^{1+d})$, and $0 \leq s_1 < \dots < s_p \leq t$.

$$F(\mu) = \left| \langle f, \mu_t \rangle - \langle f, \mu_0 \rangle \right. \\ - \int_0^t \left(\int_{\mathcal{X} \times \mathcal{Y}} \alpha(y - \langle c' \sigma(w' \cdot x), \mu_s \rangle) \langle \sigma(w \cdot x) \partial_c f, \mu_s \rangle \pi(dx, dy) \right) ds \\ + \int_0^t \left(\int_{\mathcal{X} \times \mathcal{Y}} \alpha(y - \langle c' \sigma(w' \cdot x), \mu_s \rangle) \langle c \sigma'(w \cdot x) x \cdot \nabla_w f, \mu_s \rangle \pi(dx, dy) \right) ds \\ \left. \times \langle g_1, \mu_{s_1} \rangle \times \dots \times \langle g_p, \mu_{s_p} \rangle \right|.$$

Then, by the proof of Lemma 3.1, we obtain for large N

$$\begin{aligned}\mathbb{E}_{\pi^N}[F(\mu)] &= \mathbb{E}[F(\mu^N)] \\ &= \mathbb{E}\left|(M^{1,N}(t) + M^{2,N}(t) + O(N^{-1})) \prod_{i=1}^p \langle g_i, \mu_{s_i}^N \rangle\right| \\ &\leq \mathbb{E}[|M^{1,N}(t)|] + \mathbb{E}[|M^{2,N}(t)|] + O(N^{-1}) \\ &\leq \mathbb{E}[(M^{1,N}(t))^2]^{1/2} + \mathbb{E}[(M^{2,N}(t))^2]^{1/2} + O(N^{-1}) \\ &\leq C \left(\frac{1}{\sqrt{N}} + \frac{1}{N} \right).\end{aligned}$$

Therefore,

$$\lim_{N \rightarrow \infty} \mathbb{E}_{\pi^N}[F(\mu)] = 0.$$

Since $F(\cdot)$ is continuous and $F(\mu^N)$ is uniformly bounded (due to the uniform boundedness results of section 2),

$$\mathbb{E}_{\pi}[F(\mu)] = 0.$$

Since this holds for each $t \in [0, T]$, $f \in C_b^2(\mathbb{R}^{1+d})$, and $g_1, \dots, g_p \in C_b(\mathbb{R}^{1+d})$, $\bar{\mu}$ satisfies the evolution equation (1.7). \square

It remains to be proven that the evolution equation (1.7) has a unique solution. This is the content of section 4.

4. Uniqueness. We prove uniqueness of a solution to the evolution equation (1.7). We will set up a Picard type of iteration and prove that it has a unique fixed point through a contraction mapping. We start by noticing that we can write

$$\begin{aligned}\langle f, \bar{\mu}_t \rangle &= \langle f, \bar{\mu}_0 \rangle + \int_0^t \left(\int_{\mathcal{X} \times \mathcal{Y}} \alpha(y - \langle c' \sigma(w' \cdot x), \bar{\mu}_s \rangle) \langle \sigma(w \cdot x) \partial_c f, \bar{\mu}_s \rangle \pi(dx, dy) \right) ds \\ &\quad + \int_0^t \left(\int_{\mathcal{X} \times \mathcal{Y}} \alpha(y - \langle c' \sigma(w' \cdot x), \bar{\mu}_s \rangle) \langle c \sigma'(w \cdot x) x \nabla_w f, \bar{\mu}_s \rangle \pi(dx, dy) \right) ds. \\ (4.1) \quad &= \langle f, \bar{\mu}_0 \rangle + \int_0^t \langle G(z, Q(\bar{\mu}_s, \cdot)) \cdot \nabla f, \bar{\mu}_s \rangle ds,\end{aligned}$$

where for $z = (c, w_1, \dots, w_d) \in \mathbb{R}^{1+d}$, $Q(\bar{\mu}, x) = \langle c \sigma(w \cdot x), \bar{\mu} \rangle$ we have

$$G(z, Q(\bar{\mu}, \cdot)) = (G_1(z, Q(\bar{\mu}, \cdot)), G_2(z, Q(\bar{\mu}, \cdot))) \in \mathbb{R}^{1+d}$$

with

$$\begin{aligned}G_1(z, Q(\bar{\mu}, \cdot)) &= \int_{\mathcal{X} \times \mathcal{Y}} \alpha(y - Q(\bar{\mu}, x)) \sigma(w \cdot x) \pi(dx, dy) \in \mathbb{R}, \\ G_2(z, Q(\bar{\mu}, \cdot)) &= \int_{\mathcal{X} \times \mathcal{Y}} \alpha(y - Q(\bar{\mu}, x)) c \sigma'(w \cdot x) x \pi(dx, dy) \in \mathbb{R}^d.\end{aligned}$$

We remark here that a solution to (4.1), $\bar{\mu}_\cdot$, is associated to the nonlinear random process Z_t (see, for example, [29]) satisfying the random ordinary differential equation

(ODE)

$$\begin{aligned}
 Z_t &= Z_0 + \int_0^t G(Z_s, Q(\bar{\mu}_s, \cdot)) ds, \\
 Z_0 &\sim \bar{\mu}(0, c, w), \\
 \bar{\mu}_t &= \text{Law}(Z_t).
 \end{aligned}
 \tag{4.2}$$

This ODE is random due to the random initial data.

Let us now define the following mappings. Let

$$F : D([0, T]; \mathbb{R}) \mapsto D([0, T]; M(\mathbb{R}^{1+d}))$$

be such that for a path $(R_t)_{t \in [0, T]} \in D([0, T]; \mathbb{R})$, we have that $F(R_\cdot) = \text{Law}(Y_\cdot)$, where Y_\cdot is given by

$$\begin{aligned}
 Y_t &= Y_0 + \int_0^t G(Y_s, R_s) ds, \\
 Y_0 &\sim \bar{\mu}(0, c, w).
 \end{aligned}$$

Now, let us also define the map $L : D([0, T]; M(\mathbb{R}^{1+d})) \mapsto D([0, T]; \mathbb{R})$ taking a measure-valued process μ_t and mapping it to $Q(\mu_t, x) = L(\mu)$, where

$$Q(\mu_t, x) = \langle c\sigma(w \cdot x), \mu_t \rangle.$$

Then, we consider the mapping $H : D([0, T]; M(\mathbb{R}^{1+d})) \mapsto D([0, T]; M(\mathbb{R}^{1+d}))$ defined via the composition of the mappings F and L , and we set $H = F \circ L$. Sometimes, in order to emphasize the dependence on T , we may write H_T for H .

It is clear that if $(\mu_t)_{t \in [0, T]}$ is a fixed point of H , then $\text{Law}(Z_t) = H_t(\mu_\cdot)$ is a solution to (4.1). Conversely, if $(Z_t)_{t \in [0, T]}$ is a solution to 4.2, then its law will be a fixed point of H , implying that $\text{Law}(Z_t) = H_t(\mu)$. In addition, if μ is a weak measure valued solution to (4.1), then it must be a fixed point of H and thus satisfy (4.2), proving our result.

Now, we need to show that H is a contraction mapping for $t \in [0, T]$. The first step is to show that in studying the fixed point of H , we can in fact consider $H : C([0, T]; M(\mathbb{R}^{1+d})) \mapsto C([0, T]; M(\mathbb{R}^{1+d}))$. This will allow us to work in $C([0, T]; M(\mathbb{R}^{1+d}))$ instead of working in the larger space $D([0, T]; M(\mathbb{R}^{1+d}))$ streamlining some elements of the proof.

For this reason we first derive some a priori bounds and study regularity for Z_t satisfying the random ODE given by (4.2), where $\bar{\mu}_t$ is the probability measure of the parameters at time t . Denoting by \mathbb{E} the expectation operator taken with respect to this measure (notice that here (x, y) are considered to be integration variables) we essentially consider the following system of random ODEs:

$$\begin{aligned}
 c_t &= c_0 + \int_0^t \alpha \int_{\mathcal{X} \times \mathcal{Y}} (y - \mathbb{E}[c_s \sigma(w_s \cdot x)]) \sigma(w_s \cdot x) \pi(dx, dy) ds, \\
 w_t &= w_0 + \int_0^t \alpha \int_{\mathcal{X} \times \mathcal{Y}} (y - \mathbb{E}[c_s \sigma(w_s \cdot x)]) c_s \sigma'(w_s \cdot x) x \pi(dx, dy) ds, \\
 (c_0, w_0) &\sim \bar{\mu}(0, c, w).
 \end{aligned}
 \tag{4.3}$$

Lemma 4.1 shows that there is regularity in time and it also provides us with some useful a priori uniform bounds.

LEMMA 4.1. *Let $1 \leq p \leq 4$ and $T < \infty$ be given. Then, there are constants $C_1, C_2, C_3 < \infty$, depending on p , such that*

$$\sup_{t \in [0, T]} |c_t|^p \leq C_1 (|c_0|^p + 1 + T^p C_2), \quad \mathbb{E} \sup_{t \in [0, T]} \|w_t\|^p \leq C_1 (\mathbb{E}|c_0|^p + \mathbb{E}\|w_0\|^p + T^p C_2)$$

and for every $0 \leq s \leq t \leq T$ we have that

$$|c_t - c_s|^p \leq C|t - s|^p, \quad \mathbb{E}\|w_t - w_s\|^p \leq C(\mathbb{E}|c_0|^p + 1)|t - s|^p.$$

Proof. Let's examine c_t first and establish a bound on its growth. The constant C may change from line to line and it may also depend upon the final time T and on p .

$$\begin{aligned} c_t &= c_0 + \int_0^t \alpha \int_{\mathcal{X} \times \mathcal{Y}} (y - \mathbb{E}[c_s \sigma(w_s \cdot x)]) \sigma(w_s \cdot x) \pi(dx, dy) ds, \\ c_t \sigma(w_t \cdot x) &= \sigma(w_t \cdot x) c_0 + \sigma(w_t \cdot x) \int_0^t \alpha \int_{\mathcal{X} \times \mathcal{Y}} (y - \mathbb{E}[c_s \sigma(w_s \cdot x)]) \sigma(w_s \cdot x) \pi(dx, dy) ds, \\ |c_t \sigma(w_t \cdot x)| &\leq C|c_0| + C \int_0^t \int_{\mathcal{X} \times \mathcal{Y}} |(y - \mathbb{E}[c_s \sigma(w_s \cdot x)])| \pi(dx, dy) ds. \end{aligned} \tag{4.4}$$

We have used the fact that $\sigma(\cdot)$ is bounded. Now, we will use the facts that $\mathbb{E}[|c_0|^p] < C$ and x, y have finite moments when integrated against $\pi(dx, dy)$ via Assumption 1.1. For a potentially different constant $C < \infty$

$$\begin{aligned} \mathbb{E}[|c_t \sigma(w_t \cdot x)|^p] &\leq C + C \int_0^t \int_{\mathcal{X} \times \mathcal{Y}} \mathbb{E}[|c_s \sigma(w_s \cdot x)|^p] \pi(dx, dy) ds, \\ \int_{\mathcal{X} \times \mathcal{Y}} \mathbb{E}[|c_t \sigma(w_t \cdot x)|^p] \pi(dx, dy) &\leq C + C \int_0^t \int_{\mathcal{X} \times \mathcal{Y}} \mathbb{E}[|c_s \sigma(w_s \cdot x)|^p] \pi(dx, dy) ds. \end{aligned}$$

Therefore, by Gronwall's inequality, there exists a constant $C_2 < \infty$ such that

$$\int_{\mathcal{X} \times \mathcal{Y}} \mathbb{E}[|c_s \sigma(w_s \cdot x)|^p] \pi(dx, dy) \leq C_2$$

for $0 \leq s \leq T$. Therefore, returning to (4.4) and recalling Assumption 1.1 we get that uniformly in $t \in [0, T]$, there exist constants $C_1, C_2 < \infty$ such that

$$\sup_{t \in [0, T]} |c_t|^p \leq C_1 (|c_0|^p + 1 + T^p C_2).$$

Let us now obtain the claimed bound on $\mathbb{E} \sup_{t \in [0, T]} \|w_t\|^p$. We obtain from (4.3) and Assumption 1.1 that

$$\begin{aligned} \|w_t\|^p &\leq \|w_0\|^p + t^{p-1} \int_0^t \alpha^p \left(\int_{\mathcal{X} \times \mathcal{Y}} (y - \mathbb{E}[c_s \sigma(w_s \cdot x)]) c_s \sigma'(w_s \cdot x) x \pi(dx, dy) \right)^p ds \\ &\leq \|w_0\|^p + Ct^{p-1} \int_0^t \left(\int_{\mathcal{X} \times \mathcal{Y}} |y - \mathbb{E}[c_s \sigma(w_s \cdot x)]|^2 \pi(dx, dy) \right)^{p/2} \\ &\quad \times \left(\int_{\mathcal{X} \times \mathcal{Y}} |c_s|^2 \|x\|^2 \pi(dx, dy) \right)^{p/2} ds \\ &\leq \|w_0\|^p + Ct^{p-1} \int_0^t |c_s|^p \left(\int_{\mathcal{X} \times \mathcal{Y}} (|y|^2 + \mathbb{E}[|c_s \sigma(w_s \cdot x)|^2] \pi(dx, dy)) \right)^{p/2} \\ &\quad \times \left(\int_{\mathcal{X} \times \mathcal{Y}} \|x\|^2 \pi(dx, dy) \right)^{p/2} ds \end{aligned}$$

and the claimed bound follows by taking supremum over all $t \in [0, T]$, expectation, and using the previously derived uniform bound for c_t . Let us now prove the second statement of the lemma. Similarly to the calculations above and using the uniform moment bounds on c_t and w_t together with Assumption 1.1, we have

$$\begin{aligned} |c_t - c_s|^p &\leq C|t - s|^{p-1} \int_s^t \int_{\mathcal{X} \times \mathcal{Y}} |y - \mathbb{E}[c_u \sigma(w_u \cdot x)]|^p \pi(dx, dy) du \\ &\leq C_3 |t - s|^p \end{aligned}$$

for some unimportant constant $C_3 < \infty$. Using an analysis almost identical to the one that led to the bound for $\mathbb{E} \sup_{t \in [0, T]} \|w_t\|^p$ gives us the desired bound for $\mathbb{E} \|w_t - w_s\|^p$. This concludes the proof of the lemma. \square

As a consequence of the regularity result in Lemma 4.1, (4.3) is a continuous process. Therefore, we can prove a contraction in $C([0, T]; M(\mathbb{R}^{1+d}))$ (instead of studying the process in the larger space $D([0, T]; M(\mathbb{R}^{1+d}))$).

Let us define for notational convenience $C_T = C([0, T], \mathbb{R}^{1+d})$ and let M_T be the set of probability measures on C_T . Consider an element $m \in M_T$. Motivated by the discussion before Lemma 4.1 let us set $\text{Law}(Y) = H(m, \cdot)$, where, slightly abusing notation, $Y = (c, w)$ with

$$\begin{aligned} c_t &= c_0 + \int_0^t \int_{\mathcal{X} \times \mathcal{Y}} \alpha \left(y - \langle G_{s,x}, m \rangle \right) \sigma(w_s \cdot x) \pi(dx, dy) ds, \\ w_t &= w_0 + \int_0^t \int_{\mathcal{X} \times \mathcal{Y}} \alpha \left(y - \langle G_{s,x}, m \rangle \right) c_s \sigma'(w_s \cdot x) x \pi(dx, dy) ds, \\ G_{s,x} &= c'_s \sigma(w'_s \cdot x), \\ (4.5) \quad (c_0, w_0) &\sim \bar{\mu}(0, c, w). \end{aligned}$$

For $m, m' \in M_T$ and $p \geq 1$ define the metric

$$D_{T,p}(m, m') = \inf \left\{ \left(\int_{C_T \times C_T} \sup_{s \leq T} \|x_s - y_s\|_p^p \wedge 1 d\nu(x, y) \right)^{1/p}, \nu \in P(m, m') \right\},$$

where $P(m, m')$ is the set of probability measures on $C_T \times C_T$ such that the marginal distributions are m and m' , respectively.

Now we show existence and uniqueness of a fixed point $\text{Law}(c_t, w_t)$ for the mapping H , as defined via (4.5). If a solution to (4.2) exists, then it must be a fixed point of H (defined via (4.5)). This is an immediate consequence of Lemma 4.1. Therefore, if H has a unique solution, there can be at most one solution to (4.2). If (4.2) has at most one solution, (4.1) has at most one solution. Therefore, if H has a unique fixed point, this proves uniqueness for (4.1).

Due to Lemma 4.1 we need only to consider the space of measures M_T that have bounded moments up to order $p = 4$. By known results (see, for example, [7]), the space of measures with $p = 4$ finite moments endowed with the $D_{T,4}$ metric is complete and separable. Due to closedness, the space of measures with bounded moments of order four is a complete and separable metric space when endowed with the Wasserstein metric $D_{T,4}$. Therefore, in the arguments below we work with the space of measures that have bounded the first four moments and we consider the metric $D_{T,4}$. We will show that there is a unique fixed point by proving a contraction.

Lemma 4.2 shows that for a large enough bound, $H(m)$ maps from a subspace of bounded moments to the same subspace of bounded moments.

LEMMA 4.2. *Consider (c_t, w_t) solving (4.5). There is a K_0 such that for any $K > K_0$, we can find $T_1 = T_1(K) < T$ such that $\langle \sup_{0 \leq t \leq T_1} (|c_t|^4 + \|w_t\|^4), m \rangle < K$ implies $\mathbb{E}[\sup_{0 \leq t \leq T_1} (|c_t|^4 + \|w_t\|^4)] < K$.*

Proof. Assume that the measure m is such that $\langle \sup_{0 \leq t \leq T} (|c_t|^4 + \|w_t\|^4), m \rangle < K$ (where $K < \infty$ will be chosen below). Using the same steps as in Lemma 4.1, we can show that for some $T_1 < T$ (to be chosen later),

$$\begin{aligned} \mathbb{E}\left[\sup_{0 \leq t \leq T_1} |c_t|^4\right] &\leq C_1 \left(1 + T_1^4 K\right), \\ \mathbb{E}\left[\sup_{0 \leq t \leq T_1} \|w_t\|^4\right] &\leq C_1 \left(1 + T_1^4 K \mathbb{E}\left[\sup_{0 \leq t \leq T_1} |c_t|^4\right]\right). \end{aligned}$$

Let $K_0 > 2C_1$ and let $K > K_0$. Then, if $T_1 \leq (\frac{K-2C_1}{2C_1K})^{1/4}$, $\mathbb{E}[\sup_{0 \leq t \leq T_1} |c_t|^4] \leq \frac{K}{2}$. If, in addition, we have $T_1 \leq (\frac{K-2C_1}{2C_1K^2})^{1/4}$, then we get $\mathbb{E}[\sup_{0 \leq t \leq T_1} \|w_t\|^4] \leq \frac{K_1}{2}$. Therefore, if

$$T_1 \leq \min \left\{ \left(\frac{K-2C_1}{2C_1K} \right)^{1/4}, \left(\frac{K-2C_1}{2C_1K^2} \right)^{1/4} \right\},$$

we have that $\mathbb{E}[\sup_{0 \leq t \leq T_1} (|c_t|^4 + \|w_t\|^4)] < K$, concluding the proof of the lemma. \square

We can now prove a contraction and then apply the Banach fixed-point theorem to prove that there is a unique fixed point.

For two elements $m^1, m^2 \in M_T$, let us set $\text{Law}(Y^i) = \text{Law}((c^i, w^i)) = H(m^i)$ for $t \in [0, T]$ with $i = 1, 2$. So, let (c_t^1, w_t^1) satisfying (4.5) with $m = m^1$ and (c_t^2, w_t^2) satisfying (4.5) with $m = m^2$. The processes (c_t^1, w_t^1) and (c_t^2, w_t^2) have the same initial conditions. That is,

$$\begin{aligned} (c_0^1, w_0^1) &= (c_0^2, w_0^2) = (c_0, w_0), \\ (c_0, w_0) &\sim \bar{\mu}(0, dc, dw). \end{aligned}$$

We now prove a contraction for the mapping H for some $0 < T_0 < T$. By definition, (c_t^1, w_t^1) and (c_t^2, w_t^2) have marginal distributions $H(m^1)$ and $H(m^2)$, respectively, on the time interval $[0, T_0]$. Once this is proven, we can extend this to the

entire interval $[0, T]$ since T_0 is not affected by the input measures m^1, m^2 or by which subinterval of $[0, T]$ we are considering. We have the following lemma.

LEMMA 4.3. *Let $m^1, m^2 \in M_T$ and $T < \infty$. Then, there exist constants $C_1, C_2 < \infty$ that may depend on T such that*

$$D_{t,4}^4(H(m^1), H(m^2)) \leq C_2 t^4 D_{t,4}^4(m^1, m^2) \left(\mathbb{E} e^{C_1 t |c_0|} \right)^{1/2}$$

for any $0 < t < T$. In addition, if $\bar{\mu}(0, dc, dw)$ has compact support, there exists a constant $C < \infty$ that may depend on T such that

$$D_{t,4}^4(H(m^1), H(m^2)) \leq C t^4 D_{t,4}^4(m^1, m^2).$$

Proof. Using the formula (4.5) we obtain

$$\begin{aligned} c_t^1 - c_t^2 &= \int_0^t \int_{\mathcal{X} \times \mathcal{Y}} \alpha \left(y - \langle G_{s,x}, m^1 \rangle \right) \sigma(w_s^1 \cdot x) \pi(dx, dy) ds \\ &\quad - \int_0^t \int_{\mathcal{X} \times \mathcal{Y}} \alpha \left(y - \langle G_{s,x}, m^2 \rangle \right) \sigma(w_s^2 \cdot x) \pi(dx, dy) ds, \\ &= \int_0^t \int_{\mathcal{X} \times \mathcal{Y}} \alpha y \left(\sigma(w_s^1 \cdot x) - \sigma(w_s^2 \cdot x) \right) \pi(dx, dy) ds \\ &\quad + \int_0^t \int_{\mathcal{X} \times \mathcal{Y}} \alpha \langle G_{s,x}, m^2 \rangle \left(\sigma(w_s^2 \cdot x) - \sigma(w_s^1 \cdot x) \right) \pi(dx, dy) ds \\ &\quad + \int_0^t \int_{\mathcal{X} \times \mathcal{Y}} \alpha \langle G_{s,x}, m^2 - m^1 \rangle \sigma(w_s^1 \cdot x) \pi(dx, dy) ds. \end{aligned}$$

First, let's address the mean field term. Recall that $\sigma'(\cdot)$ is bounded and that $\pi(dx, dy)$ has bounded marginal moments via Assumption 1.1. Therefore, for $0 \leq t \leq T$,

$$\begin{aligned} &\left| \int_0^t \int_{\mathcal{X} \times \mathcal{Y}} \langle c'_s \sigma(w'_s x), m^2 \rangle \left(\sigma(w_s^2 \cdot x) - \sigma(w_s^1 \cdot x) \right) \pi(dx, dy) ds \right| \\ &\leq C \int_0^t \int_{\mathcal{X} \times \mathcal{Y}} |\langle c'_s \sigma(w'_s x), m^2 \rangle| \|x\| \pi(dx, dy) \|w_s^2 - w_s^1\| ds \\ &\leq C \int_0^t \langle |c'_s|, m^2 \rangle \|w_s^2 - w_s^1\| \left(\int_{\mathcal{X} \times \mathcal{Y}} \|x\| \pi(dx, dy) \right) ds \leq C \int_0^t \|w_s^2 - w_s^1\| ds. \end{aligned}$$

We next bound the term $|\int_0^t \int_{\mathcal{X} \times \mathcal{Y}} \langle c'_s \sigma(w'_s x), m^2 - m^1 \rangle \sigma(w_s^1 \cdot x) \pi(dx, dy) ds|$. We have that

$$\begin{aligned} |c^2 \sigma(w^2 \cdot x) - c^1 \sigma(w^1 \cdot x)| &= |(c^2 - c^1) \sigma(w^2 \cdot x) + c^1 (\sigma(w^2 \cdot x) - \sigma(w^1 \cdot x))| \\ &\leq C_1 |c^2 - c^1| + C_2 \|x\| |c^1| \|w^2 - w^1\|. \end{aligned}$$

Let the random variables $(c_s^{2'}, w_s^{2'})$ have marginal distribution m^2 and $(c_s^{1'}, w_s^{1'})$

have marginal distribution m^1 . Then, for $0 \leq s \leq T$,

$$\begin{aligned}
 & \left| \int_0^t \int_{\mathcal{X} \times \mathcal{Y}} \langle c'_s \sigma(w'_s x), m^2 - m^1 \rangle \sigma(w_s^1 \cdot x) \pi(dx, dy) ds \right| \\
 & \leq C \int_0^t \int_{\mathcal{X} \times \mathcal{Y}} \mathbb{E} \left[\left| c_s^{2'} \sigma(w_s^{2'} \cdot x) - c_s^{1'} \sigma(w_s^{1'} \cdot x) \right| \right] \pi(dx, dy) ds \\
 & \leq C \int_0^t \int_{\mathcal{X} \times \mathcal{Y}} \mathbb{E} \left[|c_s^{2'} - c_s^{1'}| + \|x\| |c_s^{1'}| \|w_s^{2'} - w_s^{1'}\| \right] \pi(dx, dy) ds \\
 & \leq C \int_0^t \mathbb{E} \left[|c_s^{2'} - c_s^{1'}| + |c_s^{1'}| \|w_s^{2'} - w_s^{1'}\| \right] ds \\
 (4.6) \quad & \leq C \int_0^t \mathbb{E} \left[(1 + |c_s^{1'}|) (|c_s^{2'} - c_s^{1'}| + \|w_s^{2'} - w_s^{1'}\|) \right] ds,
 \end{aligned}$$

where again Assumption 1.1 was used for the moment bounds of $\pi(dx, dy)$. The inequality holds for any joint distribution $\gamma(m^1, m^2)$ of m^1 and m^2 . Then, the auxiliary calculations provided in Appendix A show that (4.6) can be bounded in terms of $D_{s,4}(m^1, m^2)$:

$$(4.7) \quad \left| \int_0^t \int_{\mathcal{X} \times \mathcal{Y}} \langle c'_s \sigma(w'_s x), m^2 - m^1 \rangle \sigma(w_s^1 \cdot x) \pi(dx, dy) ds \right| \leq C \int_0^t D_{s,4}(m^1, m^2) ds.$$

Thus, we overall get that there is a constant $C < \infty$ such that

$$|c_t^2 - c_t^1| \leq C \int_0^t \left[\|w_s^2 - w_s^1\| + D_{s,4}(m^1, m^2) \right] ds.$$

Similar calculations also give the necessary bound for $\|w_t^1 - w_t^2\|$. For completeness, the details are provided in Appendix B.

$$(4.8) \quad \|w_t^2 - w_t^1\| \leq C \int_0^t \left[|c_s^2 - c_s^1| + |c_s^1| \|w_s^2 - w_s^1\| + D_{s,4}(m^1, m^2) |c_s^2| \right] ds.$$

Hence, for $0 \leq s \leq T$, we have the bound

$$\begin{aligned}
 |c_s^2 - c_s^1| + \|w_s^2 - w_s^1\| & \leq C_1 \int_0^s (1 + |c_u^1|) \left[|c_u^2 - c_u^1| + \|w_u^2 - w_u^1\| \right] du \\
 & + C_2 \int_0^s (1 + |c_u^2|) D_{u,2}(m^1, m^2) du.
 \end{aligned}$$

Setting for notational convenience $Z_s = |c_s^2 - c_s^1| + \|w_s^2 - w_s^1\|$, the latter relation gives

$$\sup_{s \leq t} Z_s \leq C_1 \int_0^t (1 + |c_u^1|) \sup_{r \leq u} Z_r du + C_2 \int_0^t (1 + |c_u^2|) D_{u,4}(m^1, m^2) du,$$

which by the Gronwall lemma gives

$$\begin{aligned}
 \sup_{s \leq t} Z_s & \leq C_2 \left(\int_0^t (1 + |c_u^2|) D_{u,4}(m^1, m^2) du \right) \exp \left(C_1 \int_0^t (1 + |c_u^1|) du \right) \\
 & \leq C_2 D_{t,4}(m^1, m^2) \left(\int_0^t (1 + |c_u^2|) du \right) \exp \left(C_1 \int_0^t (1 + |c_u^1|) du \right),
 \end{aligned}$$

where we used the fact that $u \mapsto D_{u,2}(m^1, m^2)$ is monotonically increasing.

Now, note that $|c_s^1|$ can be bounded in terms of $|c_0^1|$, the initial condition. In particular, using the boundedness of $\sigma(\cdot)$, the bound on $|\langle G_{s,x}, m^1 \rangle|$, the moments bounds for the distribution $\pi(dx, dy)$, and the fact that $0 \leq t \leq T < \infty$, we get for some constant $C < \infty$ that changes from line to line

$$\begin{aligned} |c_t^1| &\leq |c_0^1| + C \int_0^t \int_{\mathcal{X} \times \mathcal{Y}} |y - \langle G_{s,x}, m^1 \rangle| |\sigma(w_s^1 \cdot x)| \pi(dx, dy) ds \\ &\leq |c_0^1| + C. \end{aligned}$$

This bound holds for any $0 \leq t \leq T$. Then,

$$\sup_{s \leq t} Z_s \leq C_2 D_{t,4}(m^1, m^2) t (1 + |c_0^2|) e^{tC_1(1+|c_0^1|)}.$$

Raising this to the power four gives for some constants $C_1, C_2 < \infty$ different than before,

$$\sup_{s \leq t} Z_s^4 \leq C_2 D_{t,4}^4(m^1, m^2) t^4 (1 + |c_0^2|^4) e^{tC_1(1+|c_0^1|)}.$$

Now notice that

$$\begin{aligned} \sup_{s \leq t} [|c_s^1 - c_s^2|^4 + \|w_s^1 - w_s^2\|_4^4] \wedge 1 &\leq \sup_{s \leq t} [|c_s^1 - c_s^2|^4 + \|w_s^1 - w_s^2\|_4^4] \\ &\leq \sup_{s \leq t} [|c_s^1 - c_s^2|^4 + \|w_s^1 - w_s^2\|_1^4] \\ &\leq \sup_{s \leq t} Z_s^4. \end{aligned}$$

Combining the last two displays yields

$$\sup_{s \leq t} [|c_s^1 - c_s^2|^4 + \|w_s^1 - w_s^2\|_4^4] \wedge 1 \leq C_2 D_{t,4}^4(m^1, m^2) t^4 (1 + |c_0^2|^4) e^{tC_1(1+|c_0^1|)}.$$

Next, we take expectation and apply the Cauchy–Schwarz inequality on the right-hand side. We obtain

$$\begin{aligned} \mathbb{E} \sup_{s \leq t} [|c_s^1 - c_s^2|^4 + \|w_s^1 - w_s^2\|_4^4] \wedge 1 &\leq C_2 D_{t,4}^4(m^1, m^2) t^4 \mathbb{E} \left[(1 + |c_0^2|^4) e^{tC_1(1+|c_0^1|)} \right] \\ &\leq C_2 D_{t,4}^4(m^1, m^2) t^4 \left(\mathbb{E}(1 + |c_0^2|^8) \right)^{1/2} \left(\mathbb{E} e^{t2C_1(1+|c_0^1|)} \right)^{1/2}. \end{aligned}$$

The latter concludes the proof due to Assumption 1.1. \square

By Assumption 1.1 we have that there exists a $0 < q < \infty$ such that the moment generating function exists, i.e.,

$$M_C(q) := \mathbb{E} \left[\exp \left(q |c_0| \right) \right] < C_M.$$

Lemma 4.3 immediately proves there is a contraction on the interval $[0, T_0]$. Indeed,

$$D_{t,4}^4(H(m^1), H(m^2)) \leq C_2 t^4 D_{t,4}^4(m^1, m^2) \left[\mathbb{E} \exp \left(C_1 t |c_0^1| \right) \right]^{1/2}.$$

Therefore,

$$D_{t,4}(H(m^1), H(m^2)) \leq C_2^{1/4} t D_{t,4}(m^1, m^2) \left[\mathbb{E} \exp \left(C_1 t |c_0^1| \right) \right]^{1/8}.$$

Then, choose T_0 such that $C_2^{1/4} C_M^{1/8} T_0 < 1$, $C_1 T_0 < q$, and $T_0 \leq T_1$, where T_1 is from Lemma 4.2. That is, we choose $T_0 < \min \{ \frac{1}{C_2^{1/4} C_M^{1/8}}, \frac{q}{C_1}, T_1 \}$.

If $T_0 < \min \{ \frac{1}{C_2^{1/4} C_M^{1/8}}, \frac{q}{C_1}, T_1 \}$, then $D_{T_0,4}(H(m^1), H(m^2)) \leq k D_{T_0,4}(m^1, m^2)$ for $k < 1$ and we have proven uniqueness on the subinterval $[0, T_0]$ (via the Banach fixed-point theorem).

In fact, this directly proves our next lemma, Lemma 4.4, regarding uniqueness of the limit point over the entire interval $[0, T]$.

LEMMA 4.4. *Let $T < \infty$. The mapping $H_T = (F \circ F)_T$ has a unique fixed point.*

Proof. By Lemmas 4.2 and 4.3 and the Banach fixed-point theorem we readily obtain that there is $0 < T_0 < \infty$ such that $H_{T_0}(m)$ will be a contraction map leading to (4.5) having a unique solution on $[0, T_0]$. We then extend this construction to the whole interval $[0, T]$ by dividing the interval $[0, T]$ into subintervals $[0, T_0], [T_0, 2T_0], \dots, [T - T_0, T]$. In each subinterval, it can be shown that the solution is unique by proving a contraction as was done in Lemma 4.3, which can be done as T_0 can always be taken to be of the same magnitude, i.e., it does not depend on which subinterval is being examined. This concludes the proof. \square

5. Proof of the main results. We now collect the results to prove Theorem 1.2, Corollary 1.4, and Theorem 1.6.

Proof of Theorem 1.2. Let π^N be the probability measure corresponding to μ^N . Each π^N takes values in the set of probability measures $\mathcal{M}(D_E([0, T]))$. Relative compactness, proven in section 2, implies that every subsequence π^{N_k} has a further subsequence $\pi^{N_{k_m}}$ which weakly converges. Section 3 proves that any limit point π of $\pi^{N_{k_m}}$ will satisfy the evolution equation (1.7). Section 4 proves that the solution of the evolution equation (1.7) is unique. Therefore, by Prokhorov's theorem, π^N weakly converges to π , where π is the distribution of $\bar{\mu}$, the unique solution of (1.7). That is, μ^N converges in distribution to $\bar{\mu}$. \square

Proof of Corollary 1.4. The result follows from applying integration by parts to (1.7) using the assumption that $p(t, c, w) \rightarrow 0$ as $|c|, \|w\| \rightarrow \infty$. We also note that if a solution exists to (1.7), then it is unique due to the uniqueness of (1.7). \square

Proof of Theorem 1.6. By Theorem 1.2 we have that the scaled empirical measure μ^N converges in distribution in $D_E([0, T])$ toward a deterministic limit $\bar{\mu}$ with the joint distribution of the weights $(c_k^i, w_k^i)_{i=1}^N \in (\mathbb{R}^{1+d})^{\otimes N}$ being exchangeable. Then, by the Tanaka-Sznitman theorem (see, for example, Theorem 3.2 in [20] or [49]) we get that ρ^N will be $\bar{\mu}$ -chaotic. \square

6. Conclusion. In this paper we develop a law of large numbers result for neural networks with a single hidden layer as the number of hidden units and stochastic

gradient descent iterations grow. The limiting distribution of the parameters is rigorously shown to satisfy an explicitly stated first-order nonlinear deterministic PDE, in the form of a measure evolution equation. The limiting PDE is a function of the inputs to the model, such as the learning rate, activation function, and distribution of the observed data. A numerical study on the well-known MNIST dataset illustrates the theoretical results of this paper. In related work which builds upon the results in this paper, a central limit theorem has been proven for single-layer neural networks in [46] and a law of large numbers has been proven for deep neural networks in [47].

Appendix A. Proof of (4.7). We will show that (4.6) can be bounded in terms of $D_{s,4}(m^1, m^2)$. For notational convenience, define $Z_s := |c_s^{2,'} - c_s^{1,'}| + \|w_s^{2,'} - w_s^{1,'}\|$.

$$\begin{aligned}
\mathbb{E}[(1 + |c_s^{1,'}|)Z_s] &= \mathbb{E}[(1 + |c_s^{1,'}|)Z_s(\mathbf{1}_{Z_s \leq 1} + \mathbf{1}_{Z_s > 1})] \\
&= \mathbb{E}\left[(1 + |c_s^{1,'}|)((Z_s \wedge 1)\mathbf{1}_{Z_s \leq 1} + Z_s(Z_s \wedge 1)\mathbf{1}_{Z_s > 1})\right] \\
&= \mathbb{E}\left[(1 + |c_s^{1,'}|)(\mathbf{1}_{Z_s \leq 1} + Z_s\mathbf{1}_{Z_s > 1})(Z_s \wedge 1)\right] \\
&\leq \mathbb{E}\left[(1 + |c_s^{1,'}|)(1 + Z_s)(Z_s \wedge 1)\right] \\
&\leq \mathbb{E}\left[(1 + |c_s^{1,'}|)(1 + |c_s^{2,'}| + |c_s^{1,'}| + \|w_s^{1,'}\| + \|w_s^{2,'}\|)(Z_s \wedge 1)\right] \\
&\leq C \left[\mathbb{E}(1 + |c_s^{1,'}|^4 + |c_s^{2,'}|^4 + \|w_s^{1,'}\|^4 + \|w_s^{2,'}\|^4)\right]^{1/2} \left[\mathbb{E}(Z_s \wedge 1)^2\right]^{1/2} \\
&\leq C \left[\mathbb{E}(Z_s^4 \wedge 1)\right]^{1/4} \\
&\leq C \left[\mathbb{E}\left((|c_s^{2,'} - c_s^{1,'}|^4 + \|w_s^{2,'} - w_s^{1,'}\|_4^4) \wedge 1\right)\right]^{1/4}.
\end{aligned}$$

The sixth line uses the Cauchy–Schwarz inequality and Young’s inequality. The seventh line uses the facts that m^1 and m^2 have bounded fourth-order moments. The eighth line uses Young’s inequality. We have also used the facts that $(z \wedge 1)^4 = z^4 \wedge 1$ and $(Kz) \wedge 1 \leq K(z \wedge 1)$ when $z \geq 0$ and $K \geq 1$.

Therefore,

$$\begin{aligned}
&\left| \int_0^t \int_{\mathcal{X} \times \mathcal{Y}} \langle c'_s \sigma(w'_s x), m^2 - m^1 \rangle \sigma(w_s^1 \cdot x) \pi(dx, dy) ds \right| \\
&\leq C \int_0^t \left[\mathbb{E}\left((|c_s^{2,'} - c_s^{1,'}|^4 + \|w_s^{2,'} - w_s^{1,'}\|_4^4) \wedge 1\right) \right]^{1/4} ds \\
&\leq C \int_0^t \left[\mathbb{E}\left(\sup_{u \leq s} (|c_u^{2,'} - c_u^{1,'}|^4 + \|w_u^{2,'} - w_u^{1,'}\|_4^4) \wedge 1\right) \right]^{1/4} ds.
\end{aligned}$$

Since this inequality holds for any joint distribution $\gamma(m^1, m^2)$, we have that (4.7) holds.

Appendix B. Proof of (4.8). By definition, we have

$$\begin{aligned}
 w_t^1 - w_t^2 &= \int_0^t \int_{\mathcal{X} \times \mathcal{Y}} \alpha \left(y - \langle G_{s,x}, m^1 \rangle \right) c_s^1 \sigma'(w_s^1 \cdot x) x \pi(dx, dy) ds \\
 &\quad - \int_0^t \int_{\mathcal{X} \times \mathcal{Y}} \alpha \left(y - \langle G_{s,x}, m^2 \rangle \right) c_s^2 \sigma'(w_s^2 \cdot x) x \pi(dx, dy) ds \\
 &= \alpha \int_0^t \int_{\mathcal{X} \times \mathcal{Y}} x \left[y \left(c_s^1 \sigma'(w_s^1 \cdot x) - c_s^2 \sigma'(w_s^2 \cdot x) \right) \right. \\
 &\quad \left. + \left(\langle G_{s,x}, m^2 \rangle - \langle G_{s,x}, m^1 \rangle \right) c_s^2 \sigma'(w_s^2 \cdot x) \right. \\
 &\quad \left. + \langle G_{s,x}, m^1 \rangle \left(c_s^2 \sigma'(w_s^2 \cdot x) - c_s^1 \sigma'(w_s^1 \cdot x) \right) \right] \pi(dx, dy) ds \\
 &= \alpha \int_0^t \int_{\mathcal{X} \times \mathcal{Y}} x \left[y \left(c_s^1 - c_s^2 \right) \sigma'(w_s^2 \cdot x) + y c_s^1 \left(\sigma'(w_s^1 \cdot x) - \sigma'(w_s^2 \cdot x) \right) \right. \\
 &\quad \left. + \left(\langle G_{s,x}, m^2 \rangle - \langle G_{s,x}, m^1 \rangle \right) c_s^2 \sigma'(w_s^2 \cdot x) \right. \\
 &\quad \left. + \langle G_{s,x}, m^1 \rangle \left(c_s^2 \sigma'(w_s^2 \cdot x) - c_s^1 \sigma'(w_s^1 \cdot x) \right) \right] \pi(dx, dy) ds \\
 &= \alpha \int_0^t \int_{\mathcal{X} \times \mathcal{Y}} x \left[y \left(c_s^1 - c_s^2 \right) \sigma'(w_s^2 \cdot x) + y c_s^1 \left(\sigma'(w_s^1 \cdot x) - \sigma'(w_s^2 \cdot x) \right) \right. \\
 &\quad \left. + \left(\langle G_{s,x}, m^2 \rangle - \langle G_{s,x}, m^1 \rangle \right) c_s^2 \sigma'(w_s^2 \cdot x) \right. \\
 &\quad \left. + \langle G_{s,x}, m^1 \rangle \left(c_s^2 - c_s^1 \right) \sigma'(w_s^2 \cdot x) \right. \\
 &\quad \left. + \langle G_{s,x}, m^1 \rangle c_s^1 \left(\sigma'(w_s^2 \cdot x) - \sigma'(w_s^1 \cdot x) \right) \right] \pi(dx, dy) ds.
 \end{aligned}$$

Therefore, we have the bound

$$\begin{aligned}
 \|w_t^2 - w_t^1\| &\leq C \int_0^t \int_{\mathcal{X} \times \mathcal{Y}} \|x\| \left[|y| |c_s^2 - c_s^1| + |y| |c_s^1| |\sigma'(w_s^2 \cdot x) - \sigma'(w_s^1 \cdot x)| \right. \\
 &\quad \left. + |\langle G_{s,x}, m^2 \rangle - \langle G_{s,x}, m^1 \rangle| |c_s^2| \right. \\
 &\quad \left. + |\langle G_{s,x}, m^1 \rangle| |c_s^2 - c_s^1| \right. \\
 &\quad \left. + |\langle G_{s,x}, m^1 \rangle| |c_s^1| |\sigma'(w_s^2 \cdot x) - \sigma'(w_s^1 \cdot x)| \right] \pi(dx, dy) ds.
 \end{aligned}$$

Due to m^1 having bounded moments, $|\langle G_{s,x}, m^1 \rangle| < C$ (see section 4 for similar calculations). Due to Assumption 1.1 on σ ,

$$\begin{aligned}
 |\sigma'(w_s^2 \cdot x) - \sigma'(w_s^1 \cdot x)| &\leq C |w_s^2 \cdot x - w_s^1 \cdot x| = C \left| \sum_{i=1}^d (w_s^{2,i} x_i - w_s^{1,i} x_i) \right| \\
 &\leq C \sum_{i=1}^d |x_i| |w_s^{2,i} - w_s^{1,i}| \leq C \sum_{i=1}^d \|x\| |w_s^{2,i} - w_s^{1,i}| \\
 &= C \|x\| \|w_s^2 - w_s^1\|.
 \end{aligned}$$

Using these inequalities and the bounded moments of $\pi(dx, dy)$, we can calculate the upper bound

$$\|w_t^2 - w_t^1\| \leq C \int_0^t \left[|c_s^2 - c_s^1| + |c_s^1| \|w_s^2 - w_s^1\| + |\langle G_{s,x}, m^2 \rangle - \langle G_{s,x}, m^1 \rangle| |c_s^2| \right] ds.$$

Using the same approach as in the bound for (4.6) (see Appendix A), we have the bound

$$|\langle G_{s,x}, m^2 \rangle - \langle G_{s,x}, m^1 \rangle| |c_s^2| \leq CD_{s,4}(m^1, m^2) |c_s^2|.$$

Therefore, we have obtained (4.8).

REFERENCES

- [1] B. ALIPANAH, A. DELONG, M. WEIRAUCH, AND B. FREY, *Predicting the sequence specificities of DNA-and RNA-binding proteins by deep learning*, Nature Biotechnology, 33 (2015), pp. 831–838.
- [2] L. AMBROSIO, N. GIGLI, AND G. SAVARÉ, *Gradient Flows: In Metric Spaces and in the Space of Probability Measures*, Springer, New York, 2008.
- [3] S. ARIK, M. CHRZANOWSKI, A. COATES, G. DIAMOS, A. GIBIANSKY, Y. KANG, X. LI, J. MILLER, A. NG, J. RAIMAN, AND S. SENGPUTA, *Deep Voice: Real-Time Neural Text-to-Speech*, <https://arxiv.org/abs/1702.07825>, 2017.
- [4] A. BARRON, *Approximation and estimation bounds for artificial neural networks*, Mach. Learn., 14 (1994), pp. 115–133.
- [5] P. BARTLETT, D. FOSTER, AND M. TELGARSKY, *Spectrally-normalized margin bounds for neural networks*, Adv. Neural Inf. Process. Syst., 30 (2017), pp. 6241–6250.
- [6] L. BO AND A. CAPPONI, *Systemic risk in interbanking networks*, SIAM J. Financial Math., 6 (2015), pp. 386–424.
- [7] F. BOLLEY, *Separability and completeness for the Wasserstein distance*, in Séminaire de Probabilités XLI, C. Donati-Martin, M. Émery, A. Rouault, and C. Stricker, eds., Lecture Notes in Math. 1934, Springer, Berlin, 2008.
- [8] M. BOJARSKI, D. DEL TEST, D. DWORAKOWSKI, B. FIRNIER, B. FLEPP, P. GOYAL, L. JACKEL, M. MONFORT, U. MULLER, J. ZHANG, AND X. ZHANG, *End to End Learning for Self-Driving Cars*, <https://arxiv.org/abs/1604.07316>, 2016.
- [9] Z. CAO, W. LI, S. LI, AND F. WEI, *Improving multi-document summarization via text classification*, in Proceedings of the AAAI Conference on Artificial Intelligence, 2017, pp. 3053–3059.
- [10] J. A. CARRILLO, R. J. MCCANN, AND C. VILLANI, *Kinetic equilibration rates for granular media and related equations: Entropy dissipation and mass transportation estimates*, Rev. Mat. Iberoam., 19 (2003), pp. 971–1018.
- [11] L. CHIZAT AND F. BACH, *On the global convergence of gradient descent for over-parameterized models using optimal transport*, Adv. Neural Inf. Process. Syst., 32 (2018), pp. 3040–3050.
- [12] P. DAI PRA, W. RUNGGALDIER, E. SARTORI, AND M. TOLOTTI, *Large portfolio losses: A dynamic contagion model*, Ann. Appl. Probab., 19 (2009), pp. 347–394.
- [13] P. DAI PRA AND F. HOLLANDER, *McKean-Vlasov limit for interacting random processes in random media*, J. Stat. Phys., 84 (1996), pp. 735–772.
- [14] P. DAI PRA AND M. TOLOTTI, *Heterogeneous credit portfolios and the dynamics of the aggregate losses*, Stochastic Process. Appl., 119 (2009), pp. 2913–2944.
- [15] F. DELARUE, J. INGLIS, S. RUBENTHALER, AND E. TANRE, *Particle systems with a singular mean-field self-excitation. Application to neuronal networks*, Stochastic Process. Appl., 125 (2015), pp. 2451–2492.
- [16] S. ETHIER AND T. KURTZ, *Markov Processes: Characterization and Convergence*, 1986, Wiley, New York, 1986.
- [17] A. ESTEVA, B. KUPREL, R. NOVOA, J. KO, S. SWETTER, H. BLAU, AND S. THRUN, *Dermatologist-level classification of skin cancer with deep neural networks*, Nature, 542 (2017), pp. 115–118.
- [18] K. GIESECKE, K. SPILIOPOULOS, AND R. SOWERS, *Default clustering in large portfolios: Typical events*, Ann. Appl. Probab., 23 (2013), pp. 348–385.
- [19] K. GIESECKE, K. SPILIOPOULOS, R. SOWERS, AND J. SIRIGNANO, *Large portfolio asymptotics for loss from default*, Math. Finance, 25 (2015), pp. 77–114.
- [20] A. D. GOTTLIEB, *Markov Transitions and the Propagation of Chaos*, Ph.D. thesis, University of California, Berkeley, 1998.
- [21] I. GOODFELLOW, Y. BENGIO, AND A. COURVILLE, *Deep Learning*, MIT Press, Cambridge, MA, 2016.
- [22] S. GU, E. HOLLY, T. LILLCRAP, AND S. LEVINE, *Deep reinforcement learning for robotic manipulation with asynchronous off-policy updates*, in Proceedings of the IEEE Conference on Robotics and Automation, 2017, pp. 3389–3396.

- [23] B. HAMBLBY AND S. LEDGER, *A stochastic McKean-Vlasov equation for absorbing diffusions on the half-line*, Ann. Appl. Probab., 27 (2017), pp. 2698–2752.
- [24] K. HORNIK, M. STINCHCOMBE, AND H. WHITE, *Multilayer feedforward networks are universal approximators*, Neural Networks, 2 (1989), pp. 359–366.
- [25] K. HORNIK, *Approximation capabilities of multilayer feedforward networks*, Neural Networks, 4 (1991), pp. 251–257.
- [26] K. HU, Z. REN, D. ŠIŠKA, AND L. SZPRUCH, *Mean-Field Langevin Dynamics and Energy Landscape of Neural Networks*, <https://arxiv.org/abs/1905.07769>, 2019.
- [27] J. INGLIS AND D. TALAY, *Mean-field limit of a stochastic particle system smoothly interacting through threshold hitting-times and applications to neural networks with dendritic component*, SIAM J. Math. Anal., 47 (2015), pp. 3884–3916.
- [28] R. JORDAN, D. KINDERLEHRER, AND F. OTTO, *The variational formulation of the Fokker–Planck equation*, SIAM J. Math. Anal., 29 (1998), pp. 1–17.
- [29] V. N. KOLOKOLTSOV, *Nonlinear Markov Processes and Kinetic Equations*, Cambridge Tracts in Math. 182, Cambridge University Press, Cambridge, UK, 2010.
- [30] C. KUAN AND K. HORNIK, *Convergence of learning algorithms with constant learning rates*, IEEE Trans. Neural Networks, 2 (1991), pp. 484–489.
- [31] Y. LECUN, Y. BENGIO, AND G. HINTON, *Deep learning*, Nature, 521 (2015), pp. 436–444.
- [32] Y. LECUN, L. BOTTOU, Y. BENGIO, AND P. HAFNER, *Gradient-based learning applied to document recognition*, Proc. IEEE, 86 (1998), pp. 2278–2324.
- [33] Y. LEVIATHAN AND Y. MATIAS, *Google Duplex: An AI System for Accomplishing Real-World Tasks over the Phone*, Google AI Blog, ai.googleblog.com/2018/05/duplex-ai-system-for-natural-conversation.html, 2018.
- [34] J. LING, A. KURZAWSKI, AND J. TEMPLETON, *Reynolds averaged turbulence modelling using deep neural networks with embedded invariance*, J. Fluid Mech., 807 (2016), pp. 155–166.
- [35] J. LING, R. JONES, AND J. TEMPLETON, *Machine learning strategies for systems with invariance properties*, J. Comput. Phys., 318 (2016), pp. 22–35.
- [36] S. MALLAT, *Understanding deep convolutional neural networks*, Philos. Trans. A, 374 (2016), 20150203.
- [37] O. MOYNOT AND M. SAMUELIDES, *Large deviations and mean-field theory for asymmetric random recurrent neural networks*, Probab. Theory Related Fields, 123 (2002), pp. 41–75.
- [38] S. MEI, A. MONTANARI, AND P. NGUYEN, *A mean field view of the landscape of two-layer neural networks*, Proc. Natl. Acad. Sci. USA, 115 (2018), pp. 7665–7671.
- [39] R. NALLAPATI, B. ZHOU, C. GULCEHRE, AND B. XIANG, *Abstractive Text Summarization Using Sequence-to-Sequence RNNs and Beyond*, <https://arxiv.org/abs/1602.06023>, 2016.
- [40] H. PIERSON AND M. GASHLER, *Deep learning in robotics: A review of recent research*, Adv. Robotics, 31 (2017), pp. 821–835.
- [41] G. M. ROTSKOFF AND E. VANDEN-EIJNDEN, *Neural Networks as Interacting Particle Systems: Asymptotic Convexity of the Loss Landscape and Universal Scaling of the Approximation Error*, <https://arxiv.org/abs/1805.00915>, 2018.
- [42] D. SILVER ET AL., *Mastering the game of Go with deep networks and tree search*, Nature, 529 (2016), pp. 484–489.
- [43] J. SIRIGNANO, A. SADHWANI, AND K. GIESECKE, *Deep Learning for Mortgage Risk*, <https://arxiv.org/abs/1607.02470>, 2016.
- [44] J. SIRIGNANO AND R. CONT, *Universal Features of Price Formation in Financial Markets: Perspectives from Deep Learning*, <https://arxiv.org/abs/1803.06917>, 2018.
- [45] J. SIRIGNANO AND K. SPILIOPOULOS, *DGM: A deep learning algorithm for solving partial differential equations*, J. Comput. Phys., 375 (2018), pp. 1339–1364.
- [46] J. SIRIGNANO AND K. SPILIOPOULOS, *Mean field analysis of neural networks: A central limit theorem*, Stochastic Process. Appl., to appear.
- [47] J. SIRIGNANO AND K. SPILIOPOULOS, *Mean Field Analysis of Deep Neural Networks*, <https://arxiv.org/abs/1903.04440>, 2019.
- [48] H. SOMPOLINSKY, A. CRISANTI, AND H. SOMMERS, *Chaos in random neural networks*, Phys. Rev. Lett., 61 (1988), 259.
- [49] A.-S. SZNITMAN, *Topics in propagation of chaos*, in Ecole d’Eté de Probabilités de Saint-Flour XIX 1989, P.-L. Hennequin, ed., Lecture Notes in Math. 1464, Springer, Berlin, 1991, pp. 165–251.
- [50] I. SUTSKEVER, O. VINYALS, AND Q. LE, *Sequence to sequence learning with neural networks*, Adv. Neural Inf. Process. Syst., 27 (2014), pp. 3104–3112.
- [51] N. SUNDERHAUF, O. BROCK, W. CHEIRER, R. HADSELL, D. FOX, J. LEITNER, B. UPCROFT, P. ABBEEL, W. BURGARD, M. MILFORD, AND P. CORKE, *The limits and potentials of deep learning for robotics*, Internat. J. Robotics Res., 37 (2018), pp. 405–420.

- [52] M. TELGARSKY, *Benefits of Depth in Neural Networks*, <https://arxiv.org/abs/1602.04485>, 2016.
- [53] Y. TAIGMAN, M. YANG, M. RANZATO, AND L. WOLF, *Deepface: Closing the gap to human-level performance in face verification*, in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2014, pp. 1701–1708.
- [54] C. WANG, J. MATTINGLY, AND Y. LU, *Scaling Limit: Exact and Tractable Analysis of Online Learning Algorithms with Applications to Regularized Regression and PCA*, <https://arxiv.org/abs/1712.04332>, 2017.
- [55] Y. WU, M. SCHUSTER, Z. CHEN, Q. LE, M. NOROUZI, W. MACHEREY, M. KRIKUN, Y. CAO, Q. GAO, K. MACHEREY, AND J. KLINGNER, *Google’s Neural Machine Translation System: Bridging the Gap Between Human and Machine Translation*, <https://arxiv.org/abs/1609.08144>, 2016.
- [56] Y. ZHANG, W. CHAN, AND N. JAITLY, *Very deep convolutional networks for end-to-end speech recognition*, in IEEE International Conference on Acoustics, Speech, and Signal Processing, 2017, pp. 4845–4849.