
Phase diagram of Stochastic Gradient Descent in high-dimensional two-layer neural networks

Rodrigo Veiga

IdePHICS, EPFL, Lausanne
IFUSP, USP, São Paulo

Ludovic Stephan

IdePHICS, EPFL, Lausanne

Bruno Loureiro

IdePHICS, EPFL, Lausanne

Florent Krzakala

IdePHICS, EPFL, Lausanne

Lenka Zdeborová

SPOC, EPFL, Lausanne

Abstract

Despite the non-convex optimization landscape, over-parametrized shallow networks are able to achieve global convergence under gradient descent. The picture can be radically different for narrow networks, which tend to get stuck in badly-generalizing local minima. Here we investigate the cross-over between these two regimes in the high-dimensional setting, and in particular investigate the connection between the so-called mean-field/hydrodynamic regime and the seminal approach of Saad & Solla. Focusing on the case of Gaussian data, we study the interplay between the learning rate, the time scale, and the number of hidden units in the high-dimensional dynamics of stochastic gradient descent (SGD). Our work builds on a deterministic description of SGD in high-dimensions from statistical physics, which we extend and for which we provide rigorous convergence rates.

Descent-based algorithms such as stochastic gradient descent (SGD) and its variants are the workhorse of modern machine learning. They are simple to implement, efficient to run and most importantly: they work well in practice. A detailed understanding of the performance of SGD is a major topic in machine learning. Quite recently, significant progress was achieved in the context of learning in shallow neural networks. In a series of works, it was shown that the optimisation of wide two-layer neural networks can be mapped to a convex problem in the space of probability distributions over the weights [1–4]. This remarkable result implies global convergence of two-layer networks towards perfect learning provided that the number of hidden neurons is large, the learning rate is sufficiently small and enough data is at disposition. This line of work is commonly referred to as the mean-field or the hydrodynamic limit of neural networks. Mathematically, these works showed that one could describe the entire dynamics using a partial differential equation (PDE) in d dimensions.

In a different, and older, line of work, one-pass SGD for two-layer neural networks with a *finite number* p of hidden units, synthetic Gaussian input data and teacher-generated labels has been widely studied starting with the seminal work of [5]. These works consider the limit of high-dimensional data and show, in particular, that the stochastic process driven by gradient updates converge to a set of p^2 deterministic ordinary differential equations (ODEs) as the input dimension $d \rightarrow \infty$ and the learning rate is proportional to $1/d$. The validity of these ODEs in this limit was proven by [6]. However, the picture drawn from the analysis of these ODEs is slightly different from the mean-field/hydrodynamic picture: in this case SGD can get stuck for long time in minima associated to no specialization of the hidden units to the teacher hidden units, and even when it converges to specializing minima, it fails to perfectly learn (i.e. to achieve zero population risk). In fact, in this analysis, the interplay between the limit of the learning rate going to zero and $d \rightarrow \infty$ appeared to be fundamental.

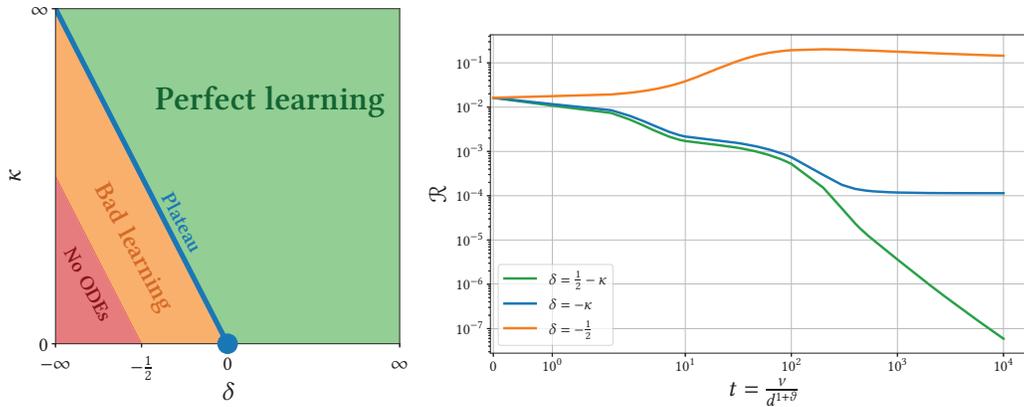


Figure 1: A summary of our contribution on SGD learning regimes for two-layer neural networks in the high-dimensional input layer limit $d \rightarrow \infty$. Left: phase diagram according to Eqs. (1): Perfect learning region ($\kappa + \delta > 0$), Plateau line ($\kappa + \delta = 0$), Bad learning region ($-1/2 < \kappa + \delta < 0$) and No ODEs region ($\kappa + \delta < -1/2$). Right: A solution of the ODEs in all regions, with matching colors (using $\kappa = 0.301$, $p = 8$, $k = 4$, $\rho_{rs} = \delta_{rs}$, $\Delta = 10^{-3}$, $\sigma(x) = \text{erf}(x/\sqrt{2})$). The time scaling is not uniform through the phase diagram: $\vartheta = \kappa + \delta$ on green and blue regimes and $\vartheta = 2(\kappa + \delta)$ on the orange region. The green curve decays as a power law to zero excess error.

One should naturally wonder about the link between these two sets of works with, on the one hand a d -dimensional PDE (with large p), and on the other a p^2 -dimensional ODE (with large d). In this work we aim to build a bridge between these two approaches for studying one-pass SGD.

Our starting point is the framework from [5], which we build upon and expand to a much broader range of choices of learning rate, time scales, and hidden layer width. This allows us to provide a sharp characterisation of the performance of SGD for two-layer neural networks in high-dimensions. We show it depends on the precise way in which the limit is taken, and in particular on how the quantity of data, the hidden layer width, & the learning rate scales as $d \rightarrow \infty$. For different choices of scaling, we can observe scenarios such as perfect learning, imperfect learning with an unavoidable error, or even no learning at all. As a consequence of our analysis, we provide a phase diagram (see Fig.1)¹ describing the possible scenarios arising in the high-dimensional setting. Our **main contributions** are as follow:

C1 We rigorously show that the dynamics of SGD can be captured by a set of deterministic ODEs, considerably extending the proof of [6] to accommodate for general time scalings defined by an arbitrary learning rate, and a general range of hidden layer width. We provide much finer non-asymptotic guarantees which are crucial for our subsequent analysis.

C2 From the analysis of the ODEs, we derive a phase diagram of SGD for two-layer neural networks in the high-dimensional input layer limit $d \rightarrow \infty$. In particular, scaling both the learning rate γ and hidden layer width p with the input dimension d as

$$\gamma \propto d^{-\delta}, \quad p \propto d^\kappa, \quad (1)$$

we identify 4 different learning regimes which are summarized in Fig. 1(left):

- Perfect learning (green region, $\kappa > -\delta$): we show that perfect learning (zero population risk) can be asymptotically achieved with $n \sim d^{1+\kappa+\delta}$ samples even for tasks with additive noise.
- Plateau (blue line $\kappa = -\delta$): learning reaches a plateau related to the noise strength. The point $\kappa = \delta = 0$ goes back to the classical work of [5].
- Bad learning (orange region $-1/2 < \kappa + \delta < 0$): here the noise dominates the learning process.
- No ODEs (red region $\kappa + \delta < -1/2$): the stochastic process associated to SGD is not guaranteed to converge to a set of deterministic ODEs. This region is thus outside the scope of our analysis.

Relation to previous work – Deterministic dynamical descriptions of one-pass stochastic gradient descent in high-dimensions have a long tradition in the statistical physics community, starting with

¹Code: https://github.com/rodsveiga/phdiag_sgd

single- and two-layer neural networks with few hidden units [7–11]. The seminal work by [5] overcame previous limitations by constructing a set of deterministic ODEs for two-layer networks with any finite number of hidden units, paving the way for a series of important contributions [12–14, 6]. This line of work corresponds to the $\kappa = \delta = 0$ case of Fig. 1 (left). One of our goals is to generalize this picture beyond fixed hidden layer size and learning rate.

A more recent line of work investigating the dynamics of SGD is the so-called *mean-field limit* [1, 15, 2–4], which connects the SGD dynamics of large-width two-layer neural networks to a diffusion equation in the hidden layer weight density. In particular, [15] provide non-asymptotic convergence bounds for sufficiently small learning rates, corresponding to the green region of Figure 1 (with $p \rightarrow \infty$). The mean-field approach computes the empirical distribution (in \mathbb{R}^d) of the hidden layer weights, while we focus on the macroscopic overlaps between the teacher and student weights.

1 Setting

Consider a supervised learning regression task. The data set is composed of n pairs $(\mathbf{x}^\nu, y^\nu)_{\nu \in [n]} \in \mathbb{R}^{d+1}$ identically and independently sampled from $\mathbb{P}(\mathbf{x}, y)$. The probability $\mathbb{P}(\mathbf{x})$ is assumed to be known and $\mathbb{P}(y|\mathbf{x})$ is modelled by a two layer neural network called the *teacher*. Given a feature vector $\mathbf{x}^\nu \in \mathbb{R}^d$, the respective label $y^\nu \in \mathbb{R}$ is defined as the output of a network with k hidden units, fixed weights $\mathbf{W}^* \in \mathbb{R}^{k \times d}$ and an activation function $\sigma : \mathbb{R} \rightarrow \mathbb{R}$:

$$y^\nu = f(\mathbf{x}^\nu, \mathbf{W}^*) + \sqrt{\Delta} \zeta^\nu ; \quad f(\mathbf{x}^\nu, \mathbf{W}^*) = \frac{1}{k} \sum_{r=1}^k \sigma \left(\frac{\mathbf{w}_r^{*\top} \mathbf{x}^\nu}{\sqrt{d}} \right) = \frac{1}{k} \sum_{r=1}^k \sigma(\lambda_r^{*\nu}), \quad (2)$$

with $\mathbf{w}_r^* \equiv [\mathbf{W}^*]_r \in \mathbb{R}^d$ as the r -th row of the matrix \mathbf{W}^* and $\lambda_r^{*\nu} \equiv \mathbf{w}_r^{*\top} \mathbf{x}^\nu / \sqrt{d} \in \mathbb{R}$ as the r -th component of the teacher *local field* vector $\boldsymbol{\lambda}^{*\nu} \in \mathbb{R}^k$. The parameter $\Delta \geq 0$ controls the strength of additive label noise: $\zeta^\nu \sim \mathbb{P}(\zeta^\nu)$ such that $\mathbb{E}_{\zeta^\nu \sim \mathbb{P}(\zeta^\nu)}[\zeta] = 0$ and $\mathbb{E}_{\zeta^\nu \sim \mathbb{P}(\zeta^\nu)}[\zeta^2] = 1$.

Given a new sample $\mathbf{x} \sim \mathbb{P}(\mathbf{x})$ outside the training data, the goal is to obtain an estimation $\hat{f}(\mathbf{x})$ for the respective label y . The error is quantified by a loss function $\mathcal{L}(y, \hat{f}(\mathbf{x}, \boldsymbol{\Theta}))$, where $\boldsymbol{\Theta}$ is an arbitrary set of parameters to be learned from data.

In this manuscript we are interested in the problem of estimating \mathbf{W}^* with another two-layer neural network with the same activation function, which we will refer to as the *student*. The student network has p hidden units and a matrix of weights $\mathbf{W} \in \mathbb{R}^{p \times d}$ to be *learned* from the data. Given a feature vector $\mathbf{x} \sim \mathbb{P}(\mathbf{x})$ the student prediction for the respective label is given as

$$\hat{f}(\mathbf{x}, \mathbf{W}) = \frac{1}{p} \sum_{j=1}^p \sigma \left(\frac{\mathbf{w}_j^\top \mathbf{x}}{\sqrt{d}} \right) = \frac{1}{p} \sum_{j=1}^p \sigma(\lambda_j^\nu), \quad (3)$$

where $\mathbf{w}_j \equiv [\mathbf{W}]_j \in \mathbb{R}^d$ is the j -th row of the matrix \mathbf{W} and $\lambda_j \equiv \mathbf{w}_j^\top \mathbf{x} / \sqrt{d} \in \mathbb{R}$ is defined as the j -th component of the student *local field* vector $\boldsymbol{\lambda} \in \mathbb{R}^p$.

One-pass gradient descent – Typically, one minimizes the *empirical risk* over the full data set. Instead, learning with *one-pass* gradient descent minimizes directly the *population risk*:

$$\mathcal{R}(\mathbf{W}, \mathbf{W}^*) \equiv \mathbb{E}_{\mathbf{x}, y \sim \mathbb{P}(\mathbf{x}, y)} \left[\mathcal{L} \left(f(\mathbf{x}, \mathbf{W}^*), \hat{f}(\mathbf{x}, \mathbf{W}) \right) \right]. \quad (4)$$

Given a *single* sample (\mathbf{x}^ν, y^ν) the weights are updated sequentially by the gradient descent rule:

$$\mathbf{w}_j^{\nu+1} = \mathbf{w}_j^\nu - \gamma \nabla_{\mathbf{w}_j} \mathcal{L} \left(y^\nu, \hat{f}(\mathbf{x}^\nu, \mathbf{W}) \right), \quad (5)$$

with $\nu \in [n]$ and $j \in [p]$. The parameter $\gamma > 0$ is the learning rate. Despite being a simplification with respect to batch learning, one-pass gradient descent is an amenable surrogate for the theoretical analysis of non-convex optimization, since at each step the gradient is computed with a fresh data sample, which is equivalent to performing SGD directly on the population risk. In particular, in this manuscript we assume realizability $p \geq k$, and focus our analysis on the square loss $\mathcal{L}(y, \hat{y}) = \frac{1}{2}(y - \hat{y})^2$, leading to

$$\mathbf{w}_j^{\nu+1} = \mathbf{w}_j^\nu + \frac{\gamma}{p\sqrt{d}} \sigma'(\lambda_j^\nu) \mathcal{E}^\nu \mathbf{x}^\nu, \quad (6)$$

where

$$\mathcal{E}^\nu \equiv \frac{1}{k} \sum_{r=1}^k \sigma(\lambda_r^{*\nu}) - \frac{1}{p} \sum_{l=1}^p \sigma(\lambda_l^\nu) + \sqrt{\Delta} \zeta^\nu, \quad (7)$$

Therefore, from the above expression we can see that to monitor the population risk along the learning dynamics it is sufficient to track the joint distribution of the local fields $(\boldsymbol{\lambda}, \boldsymbol{\lambda}^*)$. For Gaussian data $\mathbb{P}(\boldsymbol{x}) = \mathcal{N}(\boldsymbol{x}|\mathbf{0}, \mathbb{1})$, one can replace the expectation $\mathbb{E}_{\boldsymbol{x}, y \sim \mathbb{P}(\boldsymbol{x}, y)}[\cdot]$ by $\mathbb{E}_{\boldsymbol{\lambda}, \boldsymbol{\lambda}^* \sim \mathcal{N}(\boldsymbol{\lambda}, \boldsymbol{\lambda}^*|\mathbf{0}, \boldsymbol{\Omega})}[\cdot]$ and fully describe the dynamics through the following sufficient statistics, known in the statistical physics literature as *macroscopic variables*:

$$\begin{aligned} \mathbf{Q}^\nu &\equiv \mathbb{E}_{\boldsymbol{x}, y \sim \mathbb{P}(\boldsymbol{x}, y)}[\boldsymbol{\lambda}^\nu \boldsymbol{\lambda}^{\nu\top}] = \frac{1}{d} \mathbf{W}^\nu \mathbf{W}^{\nu\top} \\ \mathbf{M}^\nu &\equiv \mathbb{E}_{\boldsymbol{x}, y \sim \mathbb{P}(\boldsymbol{x}, y)}[\boldsymbol{\lambda}^\nu \boldsymbol{\lambda}^{*\nu\top}] = \frac{1}{d} \mathbf{W}^\nu \mathbf{W}^{*\top} \\ \mathbf{P} &\equiv \mathbb{E}_{\boldsymbol{x}, y \sim \mathbb{P}(\boldsymbol{x}, y)}[\boldsymbol{\lambda}^* \boldsymbol{\lambda}^{*\top}] = \frac{1}{d} \mathbf{W}^* \mathbf{W}^{*\top}, \end{aligned} \quad (8)$$

with matrix elements, called *order parameters* in the statistical physics literature, denoted by $q_{jl}^\nu \equiv [\mathbf{Q}^\nu]_{jl}$, $m_{jr}^\nu \equiv [\mathbf{M}^\nu]_{jr}$ and $\rho_{rs} \equiv [\mathbf{P}]_{rs}$. The *macroscopic state* of the system at the learning step ν is given by the *overlap matrix* $\boldsymbol{\Omega}^\nu \in \mathbb{R}^{(p+k) \times (p+k)}$:

$$\boldsymbol{\Omega}^\nu = \begin{bmatrix} \mathbf{Q}^\nu & \mathbf{M}^\nu \\ \mathbf{M}^{\nu\top} & \mathbf{P} \end{bmatrix}, \quad (9)$$

and the population risk is completely determined by the macroscopic state:

$$\mathcal{R}(\boldsymbol{\Omega}) = \frac{1}{2} \mathbb{E}_{\boldsymbol{\lambda}, \boldsymbol{\lambda}^* \sim \mathcal{N}(\boldsymbol{\lambda}, \boldsymbol{\lambda}^*|\mathbf{0}, \boldsymbol{\Omega})} \mathbb{E}_{\zeta \sim \mathbb{P}(\zeta)} \left[\left(\hat{f}(\boldsymbol{\lambda}) - f(\boldsymbol{\lambda}^*) \right)^2 \right]. \quad (10)$$

The training dynamics (5) defines a discrete-time stochastic process for the evolution of the overlap matrix $\left\{ \boldsymbol{\Omega}^\nu \in \mathbb{R}^{(p+k) \times (p+k)}, \nu \in [n] \right\}$, with \mathbf{P} fixed and \mathbf{Q}^ν and \mathbf{M}^ν updated as:

$$\begin{aligned} q_{jl}^{\nu+1} - q_{jl}^\nu &= \frac{\gamma}{pd} \underbrace{(\mathcal{E}_j^\nu \lambda_l^\nu + \mathcal{E}_l^\nu \lambda_j^\nu)}_{\text{learning}} + \frac{\gamma^2 \|\boldsymbol{x}\|^2}{p^2 d^2} \underbrace{\mathcal{E}_j^\nu \mathcal{E}_l^\nu}_{\text{variance}}, \\ m_{jr}^{\nu+1} - m_{jr}^\nu &= \frac{\gamma}{pd} \underbrace{\mathcal{E}_j^\nu \lambda_r^{*\nu}}_{\text{learning}} \end{aligned} \quad (11)$$

with $\nu \in [n]$, $j, l \in [p]$, $r \in [k]$ and $\mathcal{E}_j^\nu \equiv \sigma'(\lambda_j^\nu) \mathcal{E}^\nu$. In what follows, we will make the concentration assumption $\|\boldsymbol{x}\|^2 = d$; this will be justified in the proof of Theorem 2.1.

We emphasize in (11) the specific role played by each term in the right hand-side. The "learning" terms are the fundamental ones, that actually drive the learning of the teacher by the student. We show in Appendix C.3 that these "learning" terms are identical to those obtained in the gradient flow approximation of SGD, whose performance is the topic of many works [1–4]. Those are *precisely* the terms that draw the population risk towards zero. However, in our setting there is an additional variance term (so that this flow approximation is incomplete) that corresponds to the fluctuations of $\mathcal{L}(\boldsymbol{x}, \mathbf{W}, \mathbf{W}^*)$ around its expected value $\mathcal{R}(\mathbf{W}, \mathbf{W}^*)$. In particular, this is where the effects of the noise ζ can be felt. These terms were sometimes denoted as (I_2) and (I_4) in [16]. We shall see that the additional "variance" term is the one responsible for the plateau in the critical (blue) region of Figure 1, while its contribution vanishes in the perfect learning (green) region.

Additionally, albeit our work particularizes to Gaussian input data, we believe our conclusion, and the phase diagram discussed in Figure 1, to hold beyond this restricted case. Indeed, while the Gaussian assumption is crucial to reach a particular set of ODEs and their analytic expression, the approach can be applied to more complex data distribution, as long as one can track the sufficient statistics required to have a closed set of equations. For instance, [17] obtained very similar equations for an arbitrary mixture of Gaussians—that would obey the same scaling analysis as ours—while [18–20] proved that many complex distributions behave as Gaussians in high-dimensional setting, including, e.g. realistic GAN-generated data. We thus expect our conclusions to be robust in this respect.

2 Main results

Although $t_0 = \nu/d$ seems to be the most natural time scaling in the high-dimensional limit $d \rightarrow \infty$, if γ and p are allowed to vary with d the right-hand side (RHS) of Eqs. (11) can diverge and render the ODE approximation obsolete. Instead, for a given time scaling δt , we can rewrite Eqs. (11) as

$$\begin{aligned} \frac{q_{jl}^{\nu+1} - q_{jl}^\nu}{\delta t} &= \frac{\gamma}{pd \delta t} (\mathcal{E}_j^\nu \lambda_l^\nu + \mathcal{E}_l^\nu \lambda_j^\nu) + \frac{\gamma^2}{p^2 d \delta t} \mathcal{E}_j^\nu \mathcal{E}_l^\nu, \\ \frac{m_{jr}^{\nu+1} - m_{jr}^\nu}{\delta t} &= \frac{\gamma}{pd \delta t} \mathcal{E}_j^\nu \lambda_r^{*\nu} \end{aligned} \quad (12)$$

In Theorem 2.1 we prove that as $d \rightarrow \infty$, Ω^ν converges to the solution of the ODE:

$$\frac{d}{dt} \bar{\Omega}(t) = \psi(\bar{\Omega}(t)), \quad (13)$$

where $\psi : \mathbb{R}^{(p+k) \times (p+k)} \rightarrow \mathbb{R}^{(p+k) \times (p+k)}$ is the expected value of the RHS of Eqs. (12), provided that this solution stays bounded. This enhances the result of [6] by providing convergence rates to the ODEs encompassing all scalings adopted hereafter:

Theorem 2.1 (Deterministic scaling limit of stochastic processes). *Let $\tau \in \mathbb{R}$ be the continuous time horizon and $\delta t = \delta t(d)$ be a time scaling factor such that the following assumptions hold:*

1. *the time scaling δt satisfies for some constant c ,*

$$\delta t \geq c \max\left(\frac{\gamma}{pd}, \frac{\gamma^2}{p^2 d}\right) \quad (14)$$

2. *the activation function σ is L -Lipschitz,*

3. *the function $\psi : \mathbb{R}^{(p+k) \times (p+k)} \rightarrow \mathbb{R}^{(p+k) \times (p+k)}$ is L' -Lipschitz.*

Then, there exists a constant $C > 0$ (depending on c, L, L') such that for any $0 \leq \nu \leq \lfloor \tau/\delta t \rfloor$, the following inequality holds:

$$\mathbb{E} \|\Omega^\nu - \bar{\Omega}(\nu \delta t)\|_\infty \leq e^{C\tau} \log(p) \sqrt{\delta t}. \quad (15)$$

Our proof is based on techniques introduced in [21] (namely, their Lemma 2) which studies a different problem with related proof techniques. The proof involves decomposing $\Omega^{\nu+1}$ as

$$\Omega^{\nu+1} = \Omega^\nu + \delta t \psi(\Omega^\nu) + (\Omega^{\nu+1} - \Omega^\nu - \delta t \psi(\Omega^\nu)),$$

where the two first terms can be considered as a deterministic discrete process, and the last term is a martingale increment. The main challenge lies in showing that the martingale contribution stays bounded throughout the considered time period.

Although the method is similar to [6], there are a number of differences between the two approaches. First, our proof fixes a number of holes in [6], in particular bounding q_{jj}^ν by a sufficiently slowly diverging function of ν . Additionally, the techniques used in this paper yield a dependency in p that is nearly negligible, while the previous methods imply bounds that are much too coarse for our needs.

The function ψ can be computed explicitly for various choices of σ , which allows to check Assumption 3 directly. We provide in Appendix C the necessary computations for $\sigma(x) = \text{erf}(x/\sqrt{2})$; those for the ReLU unit can be found in [22]. It can be checked that in the ReLU case, the function ψ is not Lipschitz around the matrices Ω satisfying

$$\Omega_{jl} = \sqrt{\Omega_{jj} \Omega_{ll}}$$

for some $j \neq l$. However, in every case we have a weaker square-root-Lipschitz property: there exists $C \in \mathbb{R}$ such that

$$\|\psi(\Omega) - \psi(\Omega')\| \leq C \left\| \sqrt{\Omega} - \sqrt{\Omega'} \right\|$$

for any Ω, Ω' . Since the square root function is Lipschitz whenever the eigenvalues of Ω are bounded away from zero (see e.g. [23]), Assumption 3 is implied by the condition

$$\Omega^\nu \succeq \epsilon I_{p+k};$$

however, this assumption is much stronger, and becomes unrealistic in the specialization phase (as well as when $p \gg d$).

Theorem 2.1 allows us to safely navigate through Fig. 1 by keeping track of convergence rates of the discrete process to a set ODEs. The interplay between learning rate and hidden layer width defines the time scaling δt and the trade-off between the linear contribution on \mathcal{E}_j and the quadratic one, playing a central role on whether the network achieves perfect learning or not. Specifically, consider the following learning rate and hidden layer width scaling with d : $\gamma = \gamma_0/d^\delta$, $p = p_0 d^\kappa$ where $\gamma_0 \in \mathbb{R}^+$ and $p_0 \in \mathbb{N}$ are constants. The exponent $\delta \in \mathbb{R}$ can be either greater or smaller than zero, while $\kappa \in \mathbb{R}^+$. Replacing these scalings on Eqs. (11), we find (where we have chosen $\gamma_0 = p_0$ without loss of generality):

$$q_{jl}^{\nu+1} - q_{jl}^\nu = \frac{1}{d^{1+\kappa+\delta}} \underbrace{(\mathcal{E}_j^\nu \lambda_l^\nu + \mathcal{E}_l^\nu \lambda_j^\nu)}_{\text{learning}} + \frac{1}{d^{1+2(\kappa+\delta)}} \underbrace{\mathcal{E}_j^\nu \mathcal{E}_l^\nu}_{\text{noise}}, \quad (16a)$$

$$m_{jr}^{\nu+1} - m_{jr}^\nu = \frac{1}{d^{1+\kappa+\delta}} \underbrace{\mathcal{E}_j^\nu \lambda_r^{*\nu}}_{\text{learning}}, \quad (16b)$$

Since the distribution of the label noise $\mathbb{P}(\zeta)$ is such that $\mathbb{E}_{\zeta \sim \mathbb{P}(\zeta)}[\zeta] = 0$, the linear contribution in \mathcal{E}_j is noiseless in the high-dimensional limit $d \rightarrow \infty$, and therefore we will refer to it as the *learning term*. The noise enters in the equations through the variance computed on the quadratic contribution $\mathcal{E}_j \mathcal{E}_l$, which we will refer to as the *noise term*; intuitively, it is a high-dimensional variance correction which hinders learning. In order to satisfy (14), we shall take

$$\delta t = \max \left(\frac{1}{d^{1+\kappa+\delta}}, \frac{1}{d^{1+2(\kappa+\delta)}} \right). \quad (17)$$

When $\kappa + \delta \neq 0$, this implies that either the learning term or the noise term scale like a negative power of d , and is negligible with respect to the other term. It is then easy to check that at a finite time horizon τ , the resulting ODE behaves as if the negligible term was not present. We refer to Theorem B.1 in the appendix for a quantitative proof of this phenomenon. Let us now describe the different regimes depicted in Fig. 1.

Blue line (plateau) – When γ and p are scaled such that $\kappa = -\delta$, Eqs. (16) converge to

$$\frac{dq_{jl}}{dt_0} = \mathbb{E}_{\lambda, \lambda^* \sim \mathcal{N}(\lambda, \lambda^* | \mathbf{0}, \Omega)} [\mathcal{E}_j \lambda_l + \mathcal{E}_l \lambda_j] + \mathbb{E}_{\lambda, \lambda^* \sim \mathcal{N}(\lambda, \lambda^* | \mathbf{0}, \Omega)} \mathbb{E}_{\zeta \sim \mathbb{P}(\zeta)} [\mathcal{E}_j \mathcal{E}_l], \quad (18a)$$

$$\frac{dm_{jr}}{dt_0} = \mathbb{E}_{\lambda, \lambda^* \sim \mathcal{N}(\lambda, \lambda^* | \mathbf{0}, \Omega)} [\mathcal{E}_j \lambda_r^*], \quad (18b)$$

with $\delta t_0 \equiv 1/d$. This regime is an extension of [5] for which $\kappa = \delta = 0$. The convergence rate to the ODEs scales with $d^{-1/2} \log(d)$, and the phenomenology we observe for $\kappa = \delta$ is consistent with previous works studying the setting $\kappa = \delta = 0$; namely the existence of an asymptotic plateau proportional to the noise level. For instance, the asymptotic population risk \mathcal{R}_∞ is known to be proportional to $\gamma \Delta$ [6] when $\kappa = \delta = 0$ and the dynamics is driven by a rescaled version of Eqs. (18). Since the noise term does not vanish under this scaling, perfect learning to zero population risk is not possible. There is always an asymptotic plateau related to the noise level Δ , and the learning rate γ .

Green region (perfect learning) – If $\kappa > -\delta$ we can define the time scaling $\delta t_{\kappa+\delta} \equiv 1/d^{1+\kappa+\delta}$. By Theorem 2.1, Eqs. (16) converge to the following deterministic set of ODEs:

$$\frac{dq_{jl}}{dt_{\kappa+\delta}} = \mathbb{E}_{\lambda, \lambda^* \sim \mathcal{N}(\lambda, \lambda^* | \mathbf{0}, \Omega)} [\mathcal{E}_j \lambda_l + \mathcal{E}_l \lambda_j] + \mathcal{O} \left(\frac{\mathbb{E}_{\lambda, \lambda^* \sim \mathcal{N}(\lambda, \lambda^* | \mathbf{0}, \Omega)} \mathbb{E}_{\zeta \sim \mathbb{P}(\zeta)} [\mathcal{E}_j \mathcal{E}_l]}{d^{\kappa+\delta}} \right), \quad (19a)$$

$$\frac{dm_{jr}}{dt_{\kappa+\delta}} = \mathbb{E}_{\lambda, \lambda^* \sim \mathcal{N}(\lambda, \lambda^* | \mathbf{0}, \Omega)} [\mathcal{E}_j \lambda_r^*], \quad (19b)$$

at a rate proportional to $d^{-(1+\kappa+\delta)/2} \log(d)$, where we have highlighted that the noise term vanishes with $d^{-(\kappa+\delta)}$. Hence, as long as $\kappa > -\delta$ the noise does not play any role on the dynamics. This setting could be understood by taking an effect learning rate $\gamma_{\text{eff}} \propto d^{-\kappa-\delta}$ on $\mathcal{R}_\infty \propto \gamma \Delta$, which

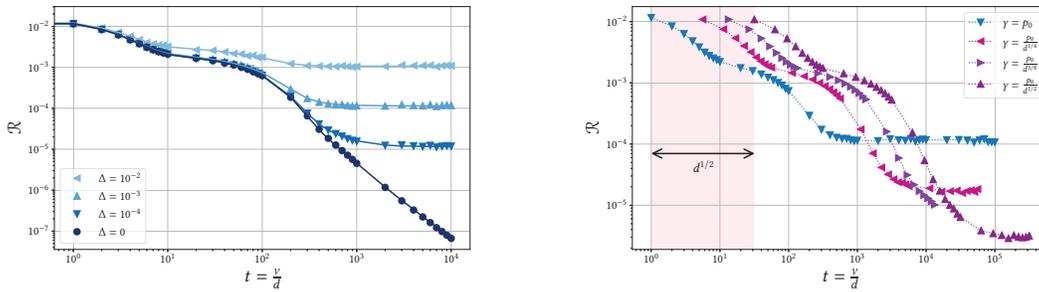


Figure 2: Left: Population risk dynamics for $\kappa = \delta = 0$ (Saad & Solla scaling) : $p_0 = 8$, $k = 4$, $\rho_{rs} = \delta_{rs}$. Activation function: $\sigma(x) = \text{erf}(x/\sqrt{2})$. Data distribution: $\mathbb{P}(\mathbf{x}) = \mathcal{N}(\mathbf{x}|\mathbf{0}, \mathbb{1})$. Dots represent simulations ($d = 1000$), while solid lines are obtained by integration of the ODEs given by Eqs. (18). Right: Simulations ($d = 1000$, $p_0 = 8$) for $\kappa = 0$ comparing different choices of the exponent δ . Noise level: $\Delta = 10^{-3}$. The final plateau is proportional to the learning rate: $\mathcal{R}_\infty \propto \gamma \Delta$.

leads to zero population risk, i.e. perfect learning, in the high dimensional limit $d \rightarrow \infty$. We validate this claim by a finite size analysis in the next section.

As discussed, the time scaling determines the number of data samples required to complete one learning step on the continuous scale. The bigger $\kappa + \delta$, the more attenuated the noise term, thus the closer to perfect learning. The trade-off is that the bigger $\kappa + \delta$, the larger the number of samples needed is, since $n = \tau d^{1+\kappa+\delta}$. Given a realizable learning task, one would thus rather choose the parameters to attain the perfect learning region, but being as close as possible to the plateau line for not increasing too much the needed number of samples. We remark that [15] provides an alternative deterministic approximation in this regime, with non-asymptotic bounds, whenever $p \gg 1$; this is the so-called mean-field approximation, with known convergence guarantees [2].

Orange region (bad learning) – We now step in the unusual situation where the learning rate grows faster with d than the hidden layer width: $\kappa < -\delta$. In this case, by (17) the noise term dominates over the dynamics. Defining the time scaling $\delta t_{2(\kappa+\delta)} \equiv 1/d^{1+2(\kappa+\delta)}$, we have

$$\frac{dq_{jl}}{dt_{2(\kappa+\delta)}} = \mathbb{E}_{\lambda, \lambda^* \sim \mathcal{N}(\lambda, \lambda^* | \mathbf{0}, \Omega)} \mathbb{E}_{\zeta \sim \mathbb{P}(\zeta)} [\mathcal{E}_j \mathcal{E}_l] + \mathcal{O}(d^{\kappa+\delta}), \quad (20a)$$

$$\frac{dm_{jr}}{dt_{2(\kappa+\delta)}} = \mathcal{O}(d^{\kappa+\delta}). \quad (20b)$$

According to Theorem 2.1 the convergence rate of Eqs. (16) to Eqs. (20) scales with $d^{-(1/2+\kappa+\delta)} \log(d)$. Therefore the existence of the noisy ODEs above is circumscribed to the region $-1/2 < \kappa + \delta < 0$ and presents a convergence trade-off absent in the other regimes: the faster of the contributions of Eqs. (16) goes to zero, the worse is the convergence rate. In the present case, the more the learning term is attenuated, i.e. the more negative is $\kappa + \delta$, the worse the dynamics is described by Eqs. (20). Although the weights are updated, the correlation between the teacher and the student weights parametrized by the overlap matrix M remains fixed on its initial value M^0 , which is a fixed point of the dynamics under this scaling. Unsurprisingly, this leads to poor generalization capacity.

Red region (no ODEs) – If $\kappa + \delta < -1/2$, the stochastic process driven by the weight dynamics does not converge to deterministic ODEs under the assumptions of Theorem 2.1. We are then not able to state any claim about this regime.

Initialization and convergence – There are two additional features worth commenting on the high-dimensional dynamics and its connection to the mean-field/hydrodynamic approach, regarding initialization and the specialization transition. In the ODE approach we discuss here, we always observe a first plateau where the teacher-student overlaps are all the same. This means all the hidden layer neurons learned the same linear separator. At this point, the two-layer network is essentially linear. This is called a *unspecialized* network in [16, 5]. In fact, this is a perfectly normal phenomenon, as with few samples even the Bayes-optimal solution would be unspecialized [24]. Only by running

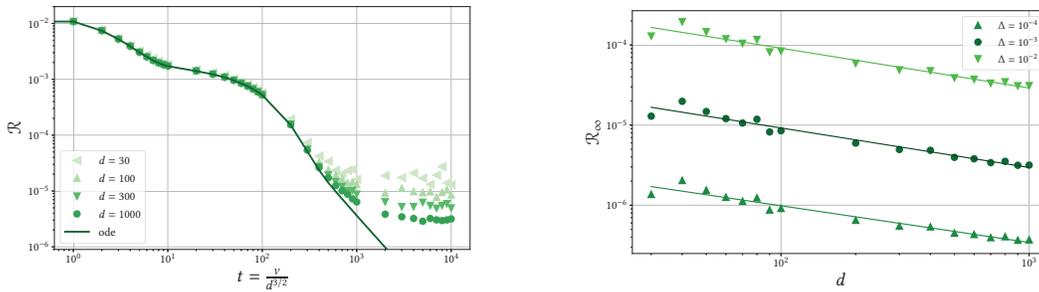


Figure 3: Networks parameters $p_0 = 8, k = 4, \rho_{rs} = \delta_{rs}$. Activation function: $\sigma(x) = \text{erf}(x/\sqrt{2})$. Data distribution: $\mathbb{P}(\mathbf{x}) = \mathcal{N}(\mathbf{x}|\mathbf{0}, \mathbf{1})$. Left: Population risk dynamics for $\kappa = 0$ and $\delta = 1/2$. Fixed noise $\Delta = 10^{-3}$ and varying d . Dots represent simulations, while the solid line is obtained by integration of the ODEs given by Eqs. (19). The data are compatible with the claim that as $d \rightarrow \infty$ the curve converges to zero population risk. Right: Asymptotic population risk \mathcal{R}_∞ from simulations (dots) as a function of d for different noise levels under the scaling $\kappa = 0$ and $\delta = 1/2$. The fitted straight lines have slopes $-0.458, -0.494, -0.497$, for $\Delta = 10^{-4}, 10^{-3}, 10^{-2}$, respectively.

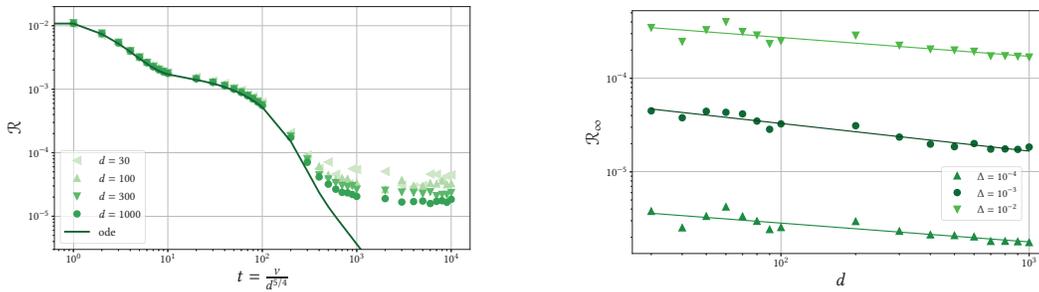


Figure 4: Networks parameters $p_0 = 8, k = 4, \rho_{rs} = \delta_{rs}$. Activation function: $\sigma(x) = \text{erf}(x/\sqrt{2})$. Data distribution: $\mathbb{P}(\mathbf{x}) = \mathcal{N}(\mathbf{x}|\mathbf{0}, \mathbf{1})$. Left: Population risk dynamics for $\kappa = 0$ and $\delta = 1/4$. Fixed noise $\Delta = 10^{-3}$ and varying d . Dots represent simulations, while the solid line is obtained by integration of the ODEs given by Eqs. (19). The data are compatible with the claim that as $d \rightarrow \infty$ the curve converges to zero population risk. Right: Asymptotic population risk \mathcal{R}_∞ from simulations (dots) as a function of d for different noise levels under the scaling $\kappa = 0$ and $\delta = 1/4$. The fitted straight lines have slopes $-0.201, -0.295, -0.201$, for $\Delta = 10^{-4}, 10^{-3}, 10^{-2}$, respectively.

the dynamics long enough the student hidden neurons start to *specialize*, each of them learning a different sub-function so that the two-layer network can learn the non-trivial teacher.

Let us make two comments on this phenomenon: (i) while the "linear" learning in the unspecialized regime may remind the reader of the linear learning in the lazy regime [25, 26] of neural nets, the two phenomena are *completely* different. In lazy training, the learning is linear because weights change very little, so that the effective network is a linear approximation of the initial one. Here, instead, the weights are changing *considerably*, but each hidden neuron learns essentially the same function. (ii) If the ODEs are initialized with weights uncorrelated with the teacher, then the unspecialized regime is a fixed point of the ODEs: the student thus never specializes, at any time. Strikingly, such condition arises as well in the analysis of mean-field equations (see e.g. Theorem 2 in [27] that discusses the need to have *spread* initial conditions with a non-zero overlap with the teacher) to guarantee global convergence. This raises the question about the precise dependence of the learning on the initialization condition in the high-dimensional regime, where a random start gets a vanishing ($1/\sqrt{d}$) overlap. This is a challenging problem that only recently has been studied (though in a simpler setting) in [28–30] who showed it yields an additional $\log(d)$ time-dependence. Generalizing these results for high-dimensional two-layer nets is an open question which we leave for future work.

3 Discussion, special cases, and simulations

To illustrate the phase diagram of Fig. 1, we present now several special cases for which we can perform simulations or numerically solve the set of ODEs. Henceforth, we take $\sigma(x) = \text{erf}(x/\sqrt{2})$,

for which the expectations of the ODEs and of the population risk, Eq. (10), can be calculated analytically [5]. The explicit expressions are presented in Appendix C. Teacher weights are such that $\rho_{rs} = \delta_{rs}$. The initial student weights are chosen such that the dimension d can be varied without changing the initial conditions \mathbf{Q}^0 , \mathbf{M}^0 , \mathbf{P} and consequently the initial population risk \mathcal{R}_0 . A detailed discussion can be found in Appendix D.

Saad & Solla scaling $\kappa = \delta = 0$ — We start by recalling the well-known setting characterized by the point $\kappa = \delta = 0$. The convergence of the stochastic process for fixed learning rate and hidden layer width to Eqs. (18) was first obtained heuristically by [5]. In Fig. 2 (left) we recall this classical result by plotting the population risk dynamics for different noise levels. Dots represent simulations, while solid lines are obtained by integration of the ODEs, Eqs. (18).

Learning is characterized by two phases after the initial decay. The first is the unspecialized plateau where all the teacher-student overlaps are approximately the same: $m_{jr} \approx m$. Waiting long enough, the dynamics reaches the *specialization* phase, where the student neurons start to *specialize*, i.e., their overlaps with one of the teacher neurons increase and consequently the population risk decreases. This specialization is discussed extensively in [5]. If $\Delta = 0$, the population risk goes asymptotically to zero. Instead, if $\Delta \neq 0$, the specialization phase presents a second plateau related to the noise Δ .

The asymptotic population risk \mathcal{R}_∞ related to the 2nd plateau is proportional to $\gamma\Delta$ [6] in the high-dimensional limit $d \rightarrow \infty$ with p finite. As mentioned in the previous section, the expectation over $\mathcal{E}_j \mathcal{E}_l$ in Eq. (18a) prevents one from obtaining zero population risk for a noisy teacher.

Perfect learning for $\kappa = 0$ — In this section we study the line $\kappa = 0$ with $\delta > 0$ of Fig. 1, for which Eqs. (19) with $\kappa = 0$ hold. We show that perfect learning can be asymptotically achieved in the realizable setting for any finite hidden layer width $p = p_0$. Keeping δ and Δ fixed, we have done simulations increasing the input layer dimension d . In Fig. 3 (left) we set $\delta = 1/2$, $\Delta = 10^{-3}$ and vary the input layer dimension. The bigger d is, the closer we are to the ODE-derived noiseless result.

Gathering the asymptotic population risk from simulations for varying d and Δ we perform a finite-size analysis to study the dependence of \mathcal{R}_∞ with d . This shows that the noise term goes to zero under this setting. In Fig. 3 (right) we plot \mathcal{R}_∞ versus d from simulations (dots) for different noise levels. We fit lines under the log-log scale showing that $\mathcal{R}_\infty \propto d^{-\delta}$, as expected. Fig. 4 draws the same conclusion for $\delta = 1/4$.

As already stated, the interplay between the exponents directly affects the time scale. We end this subsection by graphically illustrating this fact through simulations. Setting the noise to $\Delta = 10^{-3}$ we compare the cases $\delta = 0, 1/4, 3/8, 1/2$ in Fig. 2 (right). All simulations are rendered on the scale $\delta t_0 = 1/d$ to illustrate the trade-off between asymptotic performance and training time.

Bad learning for $\kappa = 0$ — We now quickly discuss the uncommon case of γ growing with d within the orange region. In Fig. 5 (left) we compare simulations varying d with the solution of the ODEs given by Eqs. (20). Both lead to poor results compared to the green & blue regions. Moreover, this regime present strong finite-size effects, making it harder to observe the asymptotic ODEs at small sizes. However, the trend as d increases is very clear from the simulations. As discussed in Section 2, the more the learning term is attenuated on the ODEs, the worse they describe the dynamics.

Large hidden layer: $\kappa > 0$ — Finishing our voyage through Fig. 1 with examples, we briefly discuss the case where both input and hidden layer widths are large. Although Theorem 2.1 provides non-asymptotic guarantees for $\kappa > 0$, the number of coupled ODEs grows quadratically with p , making the task of solving them rather challenging. Thus, we present simulations that illustrate the regions of Fig. 1. Fixing $d = 100$ we show in Fig. 5 (right) learning curves for different values of κ and δ . The colors are chosen to match their respective regions in the phase diagram.

Due to the relatively small sizes used in Fig. 5 (right), the green dots seem to decrease towards perfect learning, even when $\delta < 0$, provided that κ is large enough, as is predicted by the phase diagram in Fig. 1. Moreover, since d is not large enough, when the parameters are within the orange region the finite-size effects actually dominates, similarly to Fig. 5. The learning contribution still plays a role and the asymptotic population risk is similar to the case $\kappa = \delta = 0$. Within the red region, which is out of scope of our theory, the simulation gets stuck on a plateau with larger population risk.

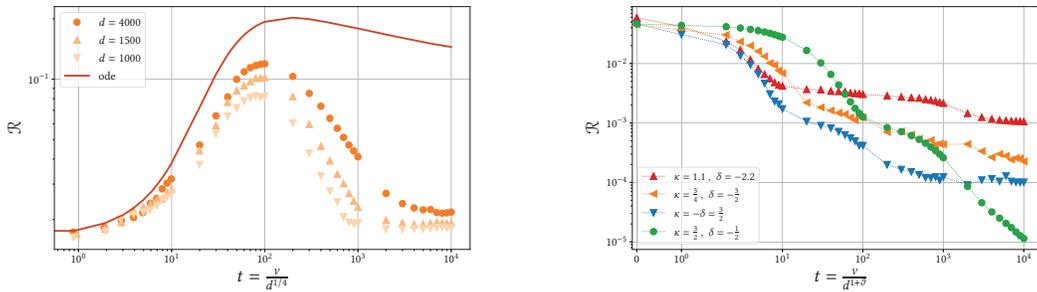


Figure 5: Left: Population risk dynamics for $\kappa = 0$ and $\delta = -3/8$. Fixed noise $\Delta = 10^{-3}$ and varying d on simulations. Dots represent simulations, while the solid line is obtained by integration of the ODEs given by Eqs. (20). Right: Simulations across different regions of Fig. 1. Networks parameters $d = 100$, $p = d^\kappa$, $\gamma = d^{-\delta}$, $k = 4$, $\rho_{rs} = \delta_{rs}$. Noise: $\Delta = 10^{-3}$. Activation function: $\sigma(x) = \text{erf}(x/\sqrt{2})$. Data distribution: $\mathbb{P}(\mathbf{x}) = \mathcal{N}(\mathbf{x}|\mathbf{0}, \mathbf{1})$. Time scaling: $\vartheta = \kappa + \delta$ for green and blue and $\vartheta = 2(\kappa + \delta)$ for orange. The colors match Fig. 1.

4 Conclusion

Building up on classical statistical physics approaches and extending them to a broad range of learning rate, time scales, and hidden layer width, we rendered a sharp characterisation of the performance of SGD for two-layer neural networks in high-dimensions. Our phase diagram describes the possible learning scenarios, characterizing learning regimes which had not been addressed by previous classical works using ODEs. Crucially, our key conclusions do not rely on an explicit solution, as our theory allows the characterization of the learning dynamics *without* solving the system of ODEs. The introduction of scaling factors is non-trivial and has deep implications. Our generalized description enlightens the trade-off between learning rate and hidden layer width, which has also been crucial in the mean-field theories.

Acknowledgements

We thank Gérard Ben Arous, Lenaïc Chizat, Maria Refinetti and Sebastian Goldt for discussions. We acknowledge funding from the ERC under the European Union’s Horizon 2020 Research and Innovation Program Grant Agreement 714608- SMiLe, and by the Swiss National Science Foundation grant SNFS OperaGOST, 200021_200390. RV was partially financed by the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior - Brasil (CAPES) - Finance Code 001. RV is grateful to EPFL and IdePHICS lab for their generous hospitality during the realization of this project.

References

- [1] 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. ISSN 0027-8424. doi:[10.1073/pnas.1806579115](https://doi.org/10.1073/pnas.1806579115).
- [2] 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*, volume 31. Curran Associates, Inc., 2018.
- [3] Grant Rotskoff and Eric Vanden-Eijnden. Trainability and accuracy of artificial neural networks: An interacting particle system approach. *Communications on Pure and Applied Mathematics*, 75(9):1889–1935, 2022. doi:<https://doi.org/10.1002/cpa.22074>.
- [4] Justin Sirignano and Konstantinos Spiliopoulos. Mean field analysis of neural networks: A central limit theorem. *Stochastic Processes and their Applications*, 130(3):1820–1852, 2020. ISSN 0304-4149. doi:<https://doi.org/10.1016/j.spa.2019.06.003>.
- [5] David Saad and Sara A. Solla. On-line learning in soft committee machines. *Phys. Rev. E*, 52: 4225–4243, Oct 1995. doi:[10.1103/PhysRevE.52.4225](https://doi.org/10.1103/PhysRevE.52.4225).
- [6] Sebastian Goldt, Madhu Advani, Andrew M Saxe, Florent Krzakala, and Lenka Zdeborová. Dynamics of stochastic gradient descent for two-layer neural networks in the teacher-student setup. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- [7] W Kinzel and P Ruján. Improving a network generalization ability by selecting examples. *Europhysics Letters (EPL)*, 13(5):473–477, nov 1990. doi:[10.1209/0295-5075/13/5/016](https://doi.org/10.1209/0295-5075/13/5/016).
- [8] O Kinouchi and N Caticha. Optimal generalization in perceptions. *Journal of Physics A: Mathematical and General*, 25(23):6243–6250, dec 1992. doi:[10.1088/0305-4470/25/23/020](https://doi.org/10.1088/0305-4470/25/23/020).
- [9] M Copelli and N Caticha. On-line learning in the committee machine. *Journal of Physics A: Mathematical and General*, 28(6):1615–1625, mar 1995. doi:[10.1088/0305-4470/28/6/016](https://doi.org/10.1088/0305-4470/28/6/016).
- [10] M Biehl and H Schwarze. Learning by on-line gradient descent. *Journal of Physics A: Mathematical and General*, 28(3):643–656, feb 1995. doi:[10.1088/0305-4470/28/3/018](https://doi.org/10.1088/0305-4470/28/3/018).
- [11] P Riegler and M Biehl. On-line backpropagation in two-layered neural networks. *Journal of Physics A: Mathematical and General*, 28(20):L507–L513, oct 1995. doi:[10.1088/0305-4470/28/20/002](https://doi.org/10.1088/0305-4470/28/20/002).
- [12] David Saad and Sara Solla. Dynamics of on-line gradient descent learning for multilayer neural networks. In D. Touretzky, M. C. Mozer, and M. Hasselmo, editors, *Advances in Neural Information Processing Systems*, volume 8. MIT Press, 1996.
- [13] Renato Vicente, Osame Kinouchi, and Nestor Caticha. Statistical mechanics of online learning of drifting concepts: A variational approach. *Machine learning*, 32(2):179–201, 1998. ISSN 0885-6125. doi:[10.1023/A:1007428731714](https://doi.org/10.1023/A:1007428731714).
- [14] David Saad, editor. *On-Line Learning in Neural Networks*. Publications of the Newton Institute. Cambridge University Press, Cambridge, 1999. ISBN 9780521652636. doi:[10.1017/CBO9780511569920](https://doi.org/10.1017/CBO9780511569920).
- [15] Song Mei, Theodor Misiakiewicz, and Andrea Montanari. Mean-field theory of two-layers neural networks: dimension-free bounds and kernel limit. In Alina Beygelzimer and Daniel Hsu, editors, *Proceedings of the Thirty-Second Conference on Learning Theory*, volume 99 of *Proceedings of Machine Learning Research*, pages 2388–2464. PMLR, 25–28 Jun 2019.
- [16] David Saad and Sara A. Solla. Exact solution for on-line learning in multilayer neural networks. *Phys. Rev. Lett.*, 74:4337–4340, May 1995. doi:[10.1103/PhysRevLett.74.4337](https://doi.org/10.1103/PhysRevLett.74.4337).

- [17] Maria Refinetti, Sebastian Goldt, Florent Krzakala, and Lenka Zdeborova. Classifying high-dimensional gaussian mixtures: Where kernel methods fail and neural networks succeed. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 8936–8947. PMLR, 18–24 Jul 2021.
- [18] Sebastian Goldt, Bruno Loureiro, Galen Reeves, Florent Krzakala, Marc Mézard, and Lenka Zdeborová. The gaussian equivalence of generative models for learning with two-layer neural networks. In *Proceedings of Machine Learning Research*, volume 145, pages 1–46. 2nd Annual Conference on Mathematical and Scientific Machine Learning, 2021.
- [19] Hong Hu and Yue M. Lu. Universality laws for high-dimensional learning with random features. *IEEE Transactions on Information Theory*, pages 1–1, 2022. doi:[10.1109/TIT.2022.3217698](https://doi.org/10.1109/TIT.2022.3217698).
- [20] Andrea Montanari and Basil N. Saeed. Universality of empirical risk minimization. In Po-Ling Loh and Maxim Raginsky, editors, *Proceedings of Thirty Fifth Conference on Learning Theory*, volume 178 of *Proceedings of Machine Learning Research*, pages 4310–4312. PMLR, 02–05 Jul 2022.
- [21] Chuang Wang, Yonina C. Eldar, and Yue M. Lu. Subspace estimation from incomplete observations: A high-dimensional analysis. *IEEE Journal of Selected Topics in Signal Processing*, 12(6):1240–1252, 2018. doi:[10.1109/JSTSP.2018.2877405](https://doi.org/10.1109/JSTSP.2018.2877405).
- [22] Yuki Yoshida, Ryo Karakida, Masato Okada, and Shun-ichi Amari. Statistical Mechanical Analysis of Online Learning with Weight Normalization in Single Layer Perceptron. *Journal of the Physical Society of Japan*, 86(4):044002, April 2017. ISSN 0031-9015. doi:[10.7566/JPSJ.86.044002](https://doi.org/10.7566/JPSJ.86.044002).
- [23] P. Del Moral and A. Niclas. A taylor expansion of the square root matrix function. *Journal of Mathematical Analysis and Applications*, 465(1):259–266, 2018. ISSN 0022-247X. doi:<https://doi.org/10.1016/j.jmaa.2018.05.005>.
- [24] Benjamin Aubin, Antoine Maillard, Jean Barbier, Florent Krzakala, Nicolas Macris, and Lenka Zdeborová. The committee machine: computational to statistical gaps in learning a two-layers neural network. *Journal of Statistical Mechanics: Theory and Experiment*, 2019(12):124023, dec 2019. doi:[10.1088/1742-5468/ab43d2](https://doi.org/10.1088/1742-5468/ab43d2).
- [25] Arthur Jacot, Franck Gabriel, and Clement Hongler. Neural tangent kernel: Convergence and generalization in neural networks. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018.
- [26] Lénaïc Chizat, Edouard Oyallon, and Francis Bach. On lazy training in differentiable programming. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- [27] Francis Bach and Lénaïc Chizat. Gradient descent on infinitely wide neural networks: Global convergence and generalization. *arXiv preprint arXiv:2110.08084*, 2021.
- [28] Yan Shuo Tan and Roman Vershynin. Phase retrieval via randomized kaczmarz: theoretical guarantees. *Information and Inference: A Journal of the IMA*, 8(1):97–123, 2019.
- [29] Gerard Ben Arous, Reza Gheissari, and Aukosh Jagannath. Online stochastic gradient descent on non-convex losses from high-dimensional inference. *Journal of Machine Learning Research*, 22(106):1–51, 2021.
- [30] Gérard Ben Arous, Reza Gheissari, and Aukosh Jagannath. Algorithmic thresholds for tensor PCA. *The Annals of Probability*, 48(4):2052 – 2087, 2020. doi:[10.1214/19-AOP1415](https://doi.org/10.1214/19-AOP1415).
- [31] Chuang Wang, Hong Hu, and Yue Lu. A solvable high-dimensional model of gan. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.