



HAL
open science

Dynamical mean-field theory for stochastic gradient descent in Gaussian mixture classification

Francesca Mignacco, Florent Krzakala, Pierfrancesco Urbani, Lenka Zdeborová

► **To cite this version:**

Francesca Mignacco, Florent Krzakala, Pierfrancesco Urbani, Lenka Zdeborová. Dynamical mean-field theory for stochastic gradient descent in Gaussian mixture classification. *Advances in Neural Information Processing Systems*, Dec 2020, Online, Canada. hal-02983444

HAL Id: hal-02983444

<https://hal.science/hal-02983444v1>

Submitted on 22 Nov 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Dynamical mean-field theory for stochastic gradient descent in Gaussian mixture classification

Francesca Mignacco¹

Florent Krzakala^{2,3}

Pierfrancesco Urbani¹

Lenka Zdeborová^{1,4}

¹ Institut de physique théorique, Université Paris-Saclay, CNRS, CEA, Gif-sur-Yvette, France

² Laboratoire de Physique, CNRS, École Normale Supérieure, PSL University, Paris, France

³ IdePHICS Laboratory, EPFL, Switzerland

⁴ SPOC Laboratory, EPFL, Switzerland

Correspondence to: francesca.mignacco@ipht.fr

Abstract

We analyze in a closed form the learning dynamics of stochastic gradient descent (SGD) for a single layer neural network classifying a high-dimensional Gaussian mixture where each cluster is assigned one of two labels. This problem provides a prototype of a non-convex loss landscape with interpolating regimes and a large generalization gap. We define a particular stochastic process for which SGD can be extended to a continuous-time limit that we call stochastic gradient flow. In the full-batch limit, we recover the standard gradient flow. We apply dynamical mean field theory from statistical physics to track the dynamics of the algorithm in the high-dimensional limit via a self-consistent stochastic process. We explore the performance of the algorithm as a function of control parameters shedding light on how it navigates the loss landscape.

1 Introduction

Understanding how stochastic gradient descent (SGD) manages to train artificial neural networks with good generalization capabilities by exploring the high-dimensional non-convex loss landscape is one of the central problems in theory of machine learning. A popular attempt to explain this behavior is by showing that the loss landscape itself is simple, with no spurious (i.e. leading to bad test error) local minima. Some empirical evidence instead leads to the conclusion that the loss landscape of state-of-the-art deep neural networks actually has spurious local (or even global) minima and stochastic gradient descent is able to find them [1, 2]. Still, the stochastic gradient descent algorithm, initialized at random, leads to good generalization properties in practice. It became clear that a theory that would explain this success needs to account for the whole trajectory of the algorithm. Yet this remains a challenging task, certainly for the state-of-the-art deep networks trained on real datasets.

Related work — A detailed description of the whole trajectory taken by the (stochastic) gradient descent was so far obtained only in several special cases. First such case are deep linear networks where the dynamics of gradient descent has been analyzed [3, 4]. While this line of works has led to very interesting insights about the dynamics, linear networks lack the expressivity of the non-linear ones and the large time behavior of the algorithm can be obtained with a simple spectral algorithm. Moreover, the analysis of dynamics in deep linear networks was not extended to the case of stochastic gradient descent. Second case where the trajectory of the algorithm was understood in detail is the *one-pass* (online) stochastic gradient descent for two-layer neural networks with a small hidden layer in the teacher-student setting [5-9]. However, the one-pass assumption made in those

analyses is far from what is done in practice and is unable to access the subtle difference between the training and test error that leads to many of the empirical mysteries observed in deep learning. A third very interesting line of research that recently provided insight about the behavior of stochastic gradient descent concerns two layer networks with divergingly wide hidden layer. This mean-field limit [10-12] maps the dynamics into the space of functions where its description is simpler and the dynamics can be written in terms of a closed set of differential equations. It is not clear yet whether this analysis can be extended in a sufficiently explicit way to deeper or finite width neural networks. The term *mean-field* has been used in several contexts in machine learning [13-18]. Note that the term in the aforementioned works refers to a variety of approximations and concepts. In this work we use it with the same meaning as in [19-21]. Most importantly, the term mean-field in our case has nothing to do with the width of an eventual hidden layer. We refer to [22] for a broader methodological review of mean-field methods and their applications to neural networks.

Our present work inscribes in the above line of research offering the dynamical mean-field theory (DMFT) formalism [19-21] leading to a closed set of integro-differential equations to track the full trajectory of the gradient descent (stochastic or not) from random initial condition in the high-dimensional limit for in-general non-convex losses. While in general the DMFT is a heuristic statistical physics method, it has been amenable to rigorous proof in some cases [23]. This is hence an important future direction for the case considered in the present paper. The DMFT has been applied recently to a high-dimensional inference problem in [24, 25] studying the spiked matrix-tensor model. However, this problem does not allow a natural way to study the stochastic gradient descent or to explore the difference between training and test errors. In particular, the spiked matrix-tensor model does not allow for the study of the so-called interpolating regime, where the loss function is optimized to zero while the test error remains positive. As such, its landscape is intrinsically different from supervised learning problems since in the former the spurious minima proliferate at high values of the loss while the good ones lie at the bottom of the landscape. Instead, deep networks have both spurious and good minima at 100% training accuracy and their landscape resembles much more the one of continuous constraint satisfaction problems [26, 27].

Main contributions — We study a natural problem of supervised classification where the input data come from a high-dimensional Gaussian mixture of several clusters, and all samples in one cluster are assigned to one of two possible output labels. We then consider a single-layer neural network classifier with a general non-convex loss function. We analyze a stochastic gradient descent algorithm in which, at each iteration, the batch used to compute the gradient of the loss is extracted at random, and we define a particular stochastic process for which SGD can be extended to a continuous time limit that we call stochastic gradient flow (SGF). In the full-batch limit we recover the standard Gradient Flow (GF). We describe the high-dimensional limit of the randomly initialized SGF with the DMFT that leads to a description of the dynamics in terms of a self-consistent stochastic process that we compare with numerical simulations. In particular, we show that the finite batch size can have a beneficial effect in the test error and acts as an effective regularization that prevents overfitting.

2 Setting and definitions

In all what follows, we will consider the high-dimensional setting where the dimension of each point in the dataset is $d \rightarrow \infty$ and the size of the training set $n = \alpha d$, being α a control parameter that we keep of order one. We consider a training set made of n points

$$\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_n)^\top \in \mathbb{R}^{n \times d} \quad \text{with labels} \quad \mathbf{y} = (y_1, \dots, y_n)^\top \in \{+1, -1\}^n \quad (1)$$

The patterns \mathbf{x}_μ are given by

$$\mathbf{x}_\mu = c_\mu \frac{\mathbf{v}^*}{\sqrt{d}} + \sqrt{\Delta} \mathbf{z}_\mu \quad \mathbf{z}_\mu \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_d) \quad \mu = 1, \dots, n. \quad (2)$$

Without loss of generality, we choose a basis where $\mathbf{v}^* = (1, 1, \dots, 1) \in \mathbb{R}^d$.

Two-cluster dataset: We will illustrate our results on a two-cluster example where the coefficients c_μ are taken at random $c_\mu = \pm 1$ with equal probability. Therefore one has two symmetric clouds of Gaussian points centered around two vectors \mathbf{v}^* and $-\mathbf{v}^*$. The labels of the data points are fixed by $y_\mu = c_\mu$. If the noise level Δ of the number of samples is small enough, the two Gaussian clouds are

linearly separable by an hyperplane, as specified in detail in [28], and therefore a single layer neural network is enough to perform the classification task in this case. We hence consider learning with the simplest neural network that classifies the data according to $\hat{y}_\mu(\mathbf{w}) = \text{sgn}[\mathbf{w}^\top \mathbf{x}_\mu / \sqrt{d}]$.

Three-cluster dataset: We consider also an example of three clusters where a good generalization error cannot be obtained by separating the points linearly. In this case we define $c_\mu = 0$ with probability $1/2$, and $c_\mu = \pm 1$ with probability $1/2$. The labels are then assigned as

$$y_\mu = -1 \text{ if } c_\mu = 0, \text{ and } y_\mu = 1 \text{ if } c_\mu = \pm 1. \quad (3)$$

One has hence three clouds of Gaussian points, two external and one centered in zero. In order to fit the data we consider a single layer-neural network with the door activation function, defined as

$$\hat{y}_\mu(\mathbf{w}) = \text{sgn} \left[\left(\frac{\mathbf{w}^\top \mathbf{x}_\mu}{\sqrt{d}} \right)^2 - L^2 \right]. \quad (4)$$

The onset parameter L could be learned, but we will instead fix it to a constant.

Loss function: We study the dynamics of learning by the empirical risk minimization of the loss

$$\mathcal{H}(\mathbf{w}) = \sum_{\mu=1}^n \ell \left[y_\mu \phi \left(\frac{\mathbf{w}^\top \mathbf{x}_\mu}{\sqrt{d}} \right) \right] + \frac{\lambda}{2} \|\mathbf{w}\|_2^2, \quad (5)$$

where we added a Ridge regularization term. The activation function ϕ is given by

$$\phi(x) = \begin{cases} x & \text{linear for the two-cluster dataset} \\ x^2 - L^2 & \text{door for the three-cluster dataset.} \end{cases} \quad (6)$$

The DMFT analysis is valid for a generic loss function ℓ . However, for concreteness in the result section we will focus on the logistic loss $\ell(v) = \ln(1 + e^{-v})$. Note that in this setting the two-cluster dataset leads to convex optimization, with a unique minimum for finite λ , and implicit regularization for $\lambda = 0$ [29], and was analyzed in detail in [30, 28]. Still the performance of stochastic gradient descent with finite batch size cannot be obtained in *static* ways. The three-cluster dataset, instead, leads to a generically non-convex optimization problem which can present many spurious minima with different generalization abilities when the control parameters such as Δ and α are changed. We note that our analysis can be extended to neural networks with a small hidden layer [31], this would allow to study the role of overparametrization, but it is left for future work.

3 Stochastic gradient-descent training dynamics

Discrete SGD dynamics — We consider the discrete gradient-descent dynamics for which the weight update is given by

$$\mathbf{w}_j(t + \eta) = \mathbf{w}_j(t) - \eta \left[\lambda \mathbf{w}_j(t) + \sum_{\mu=1}^n s_\mu(t) \Lambda' \left(y_\mu, \frac{\mathbf{w}(t)^\top \mathbf{x}_\mu}{\sqrt{d}} \right) \frac{\mathbf{x}_{\mu,j}}{\sqrt{d}} \right] \quad (7)$$

where we have introduced the function $\Lambda(y, h) = \ell(y\phi(h))$ and we have indicated with a prime the derivative with respect to h , $\Lambda'(y, h) = y\ell'(y\phi(h))\phi'(h)$. We consider the following initialization of the weight vector $\mathbf{w}(0) \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_d R)$, where $R > 0$ is a parameter that tunes the average length of the weight vector at the beginning of the dynamics[†]. The variables $s_\mu(t)$ are *i.i.d.* binary random variables. Their discrete-time dynamics can be chosen in two ways:

- In classical **SGD** at iteration t one extracts the samples with the following probability distribution

$$s_\mu(t) = \begin{cases} 1 & \text{with probability } b \\ 0 & \text{with probability } 1 - b \end{cases} \quad (8)$$

and $b \in (0, 1]$. In this way for each time iteration one extracts on average $B = bn$ patterns at random on which the gradient is computed and therefore the batch size is given by B . Note that if $b = 1$ one gets full-batch gradient descent.

[†]The DMFT equations we derive can be easily generalized to the case in which the initial distribution over \mathbf{w} is different. We only need it to be separable and independent of the dataset.

- **Persistent-SGD** is defined by a stochastic process for $s_\mu(t)$ given by the following probability rules

$$\begin{aligned}\text{Prob}(s_\mu(t + \eta) = 1 | s_\mu(t) = 0) &= \frac{1}{\tau} \eta \\ \text{Prob}(s_\mu(t + \eta) = 0 | s_\mu(t) = 1) &= \frac{(1 - b)}{b\tau} \eta,\end{aligned}\tag{9}$$

where $s_\mu(0)$ is drawn from the probability distribution (8). In this case, for each time slice one has on average $B = bn$ patterns that are *active* and enter in the computation of the gradient. The main difference with respect to the usual SGD is that one keeps the same patterns and the same minibatch for a characteristic time τ . Again, setting $b = 1$ one gets full-batch gradient descent.

Stochastic gradient flow — To write the DMFT we consider a continuous-time dynamics defined by the $\eta \rightarrow 0$ limit. This limit is not well defined for the usual SGD dynamics described by the rule (8) and we consider instead its *persistent* version described by Eq. (9). In this case the stochastic process for $s_\mu(t)$ is well defined for $\eta \rightarrow 0$ and one can write a continuous time equation as

$$\dot{w}_j(t) = -\lambda w_j(t) - \sum_{\mu=1}^n s_\mu(t) \Lambda' \left(y_\mu, \frac{\mathbf{w}(t)^\top \mathbf{x}_\mu}{\sqrt{d}} \right) \frac{x_{\mu,j}}{\sqrt{d}},\tag{10}$$

Again, for $b = 1$ one recovers the gradient flow. We call Eq. (10) stochastic gradient flow (SGF).

4 Dynamical mean-field theory for SGF

We will now analyze the SGF in the infinite size limit $n \rightarrow \infty$, $d \rightarrow \infty$ with $\alpha = n/d$ and b and τ fixed and of order one. In order to do that, we use dynamical mean-field theory (DMFT). The derivation of the DMFT equations is given in the supplementary material, here we will just present the main steps. The derivation extends the one reported in [32] for the non-convex perceptron model [26] (motivated there as a model of glassy phases of hard spheres). The main differences of the present work with respect to [32] are that here we consider a finite-batch gradient descent and that our dataset is structured while in [32] the derivation was done for full-batch gradient descent and random i.i.d. inputs and i.i.d. labels, i.e. a case where one cannot investigate generalization error and its properties. The starting point of the DMFT is the dynamical partition function

$$Z_{\text{dyn}} = \int_{\mathbf{w}(0)=\mathbf{w}^{(0)}} \mathcal{D}\mathbf{w}(t) \prod_{j=1}^d \delta \left[-\dot{w}_j(t) - \lambda w_j(t) - \sum_{\mu=1}^n s_\mu(t) \Lambda' \left(y_\mu, \frac{\mathbf{w}(t)^\top \mathbf{x}_\mu}{\sqrt{d}} \right) \frac{x_{\mu,j}}{\sqrt{d}} \right],\tag{11}$$

where $\mathcal{D}\mathbf{w}(t)$ stands for the measure over the dynamical trajectories starting from $\mathbf{w}(0)$. Since $Z_{\text{dyn}} = 1$ (it is just an integral of a Dirac delta function) [33] one can average directly Z_{dyn} over the training set, the initial condition and the stochastic processes of $s_\mu(t)$. We indicate this average with the brackets $\langle \cdot \rangle$. Hence we can write

$$Z_{\text{dyn}} = \left\langle \int \mathcal{D}\mathbf{w}(t) \mathcal{D}\hat{\mathbf{w}}(t) e^{S_{\text{dyn}}} \right\rangle,\tag{12}$$

where we have defined

$$S_{\text{dyn}} = \sum_{j=1}^d \int_0^{+\infty} dt i \hat{w}_j(t) \left(-\dot{w}_j(t) - \lambda w_j(t) - \sum_{\mu=1}^n s_\mu(t) \Lambda' \left(y_\mu, \frac{\mathbf{w}(t)^\top \mathbf{x}_\mu}{\sqrt{d}} \right) \frac{x_{\mu,j}}{\sqrt{d}} \right).\tag{13}$$

and we have introduced a set of fields $\hat{\mathbf{w}}(t)$ to produce the integral representation of the Dirac delta function. The average over the training set can be then performed explicitly, and the dynamical partition function Z_{dyn} is expressed as an integral of an exponential with extensive exponent in d :

$$Z_{\text{dyn}} = \int \mathcal{D}\mathbf{Q} \mathcal{D}\mathbf{m} e^{dS(\mathbf{Q}, \mathbf{m})},\tag{14}$$

where \mathbf{Q} and \mathbf{m} are two dynamical order parameters defined in the supplementary material. Therefore, the dynamics in the $d \rightarrow \infty$ limit satisfies a large deviation principle and we can approximate Z_{dyn}

with its value at the saddle point of the action S . In particular, one can show that the saddle point equations for the parameters \mathbf{Q} and \mathbf{m} can be recast into a self consistent stochastic process for a variable $h(t)$ related to the typical behavior of $\mathbf{w}(t)^\top \mathbf{z}_\mu / \sqrt{d}$, which evolves according to the stochastic equation:

$$\partial_t h(t) = -(\lambda + \hat{\lambda}(t))h(t) - \sqrt{\Delta} s(t) \Lambda'(y(c), r(t) - Y(t)) + \int_0^t dt' M_R(t, t') h(t') + \xi(t), \quad (15)$$

where we have denoted by $r(t) = \sqrt{\Delta} h(t) + m(t)(c + \sqrt{\Delta} h_0)$ and $m(t)$ is the *magnetization*, namely $m(t) = \mathbf{w}(t)^\top \mathbf{v}^* / d$. The details of the computation are provided in the supplementary material. There are several sources of stochasticity in Eq. (15). First, one has a dynamical noise $\xi(t)$ that is Gaussian distributed and characterized by the correlations

$$\langle \xi(t) \rangle = 0, \quad \langle \xi(t) \xi(t') \rangle = M_C(t, t'). \quad (16)$$

Furthermore, the starting point $h(0)$ of the stochastic process is random and distributed according to

$$P(h(0)) = e^{-h(0)^2 / (2R)} / \sqrt{2\pi R}. \quad (17)$$

Moreover, one has to introduce a quenched Gaussian random variable h_0 with mean zero and average one. We recall that the random variable $c = \pm 1$ with equal probability in the two-cluster model, while $c = 0, \pm 1$ in the three-cluster one. The variable $y(c)$ is therefore $y(c) = c$ in the two-cluster case, and is given by Eq. (3) in the three-cluster one. Finally, one has a dynamical stochastic process $s(t)$ whose statistical properties are specified in Eq. (9). The magnetization $m(t)$ is obtained from the following deterministic differential equation

$$\partial_t m(t) = -\lambda m(t) - \mu(t), \quad m(0) = 0^+. \quad (18)$$

The stochastic process for $h(t)$, the evolution of $m(t)$, as well as the statistical properties of the dynamical noise $\xi(t)$ depend on a series of kernels that must be computed self consistently and are given by

$$\begin{aligned} \hat{\lambda}(t) &= \alpha \Delta \langle s(t) \Lambda''(y(c), r(t)) \rangle, \\ \mu(t) &= \alpha \left\langle s(t) \left(c + \sqrt{\Delta} h_0 \right) \Lambda'(y(c), r(t)) \right\rangle, \\ M_C(t, t') &= \alpha \Delta \langle s(t) s(t') \Lambda'(y(c), r(t)) \Lambda'(y(c), r(t')) \rangle, \\ M_R(t, t') &= \alpha \Delta \frac{\delta}{\delta Y(t')} \langle s(t) \Lambda'(y(c), r(t)) \rangle \Big|_{Y=0}. \end{aligned} \quad (19)$$

In Eq. (19) the brackets denote the average over all the sources of stochasticity in the self-consistent stochastic process. Therefore one needs to solve the stochastic process in a self-consistent way. Note that $Y(t)$ in Eq. (15) is set to zero and we need it only to define the kernel $M_R(t, t')$. The set of Eqs. (15), (18) and (19) can be solved by a simple straightforward iterative algorithm. One starts with a guess for the kernels and then runs the stochastic process for $h(t)$ several times to update the kernels. The iteration is stopped when a desired precision on the kernels is reached [34].

Note that, in order to solve Eqs. (15), (18) and (19), one needs to discretize time. In the result section 5 in order to compare with numerical simulations, we will take the time-discretization of DMFT equal to the learning rate in the simulations. In the time-discretized DMFT, this allows us to extract the variables $s(t)$ either from (8) (SGD) or (9) (Persistent-SGD). In the former case this gives us a SGD-inspired discretization of the DMFT equations.

Finally, once the self-consistent stochastic process is solved, one has access also to the dynamical correlation functions $C(t, t') = \mathbf{w}(t) \cdot \mathbf{w}(t') / d$, encoded in the dynamical order parameter \mathbf{Q} that appears in the large deviation principle of Eq. (14). $C(t, t')$ concentrates for $d \rightarrow \infty$ and therefore is

controlled by the equations

$$\begin{aligned}
\partial_t C(t', t) &= -\tilde{\lambda}(t)C(t, t') + \int_0^t ds M_R(t, s)C(t', s) + \int_0^{t'} ds M_C(t, s)R(t', s) \\
&\quad - m(t') \left(\int_0^t ds M_R(t, s)m(s) + \mu(t) - \hat{\lambda}(t)m(t) \right) \quad \text{if } t \neq t', \\
\frac{1}{2} \partial_t C(t, t) &= -\tilde{\lambda}(t)C(t, t) + \int_0^t ds M_R(t, s)C(t, s) + \int_0^t ds M_C(t, s)R(t, s) \\
&\quad - m(t) \left(\int_0^t ds M_R(t, s)m(s) + \mu(t) - \hat{\lambda}(t)m(t) \right), \\
\partial_t R(t, t') &= -\tilde{\lambda}(t)R(t, t') + \delta(t - t') + \int_{t'}^t ds M_R(t, s)R(s, t'),
\end{aligned} \tag{20}$$

where we used the shorthand notation $\tilde{\lambda}(t) = \lambda + \hat{\lambda}(t)$. We consider the linear response regime, and $R(t, t') = \sum_i \delta w_i(t) / \delta H_i(t')$ is a response function that controls the variations of the weights when their dynamical evolution is affected by an infinitesimal local field $H_i(t)$. Coupling a local field $H_i(t)$ to each variable $w_i(t)$ changes the loss function as follows: $\mathcal{H}(\mathbf{w}(t)) \rightarrow \mathcal{H}(\mathbf{w}(t)) - \sum_{i=1}^d H_i(t)w_i(t)$, resulting in an extra term $H_i(t)$ to the right hand side of Eq. (10). We then consider the limit $H_i(t) \rightarrow 0$. It is interesting to note that the second of Eqs. (20) controls the evolution of the norm of the weight vector $C(t, t)$ and even if we set $\lambda = 0$ we get that it contains an effective regularization $\hat{\lambda}(t)$ that is dynamically self-generated [35].

Dynamics of the loss and the generalization error — Once the solution for the self-consistent stochastic process is found, one can get several interesting quantities. First, one can look at the training loss, which can be obtained as

$$e(t) = \alpha \langle \Lambda(y, r(t)) \rangle, \tag{21}$$

where again the brackets denote the average over the realization of the stochastic process in Eq. (15). The training accuracy is given by

$$a(t) = 1 - \langle \theta(-y\phi(r(t))) \rangle \tag{22}$$

and, by definition, it is equal to one as soon as all vectors in the training set are correctly classified. Finally, one can compute the generalization error. At any time step, it is defined as the fraction of mislabeled instances:

$$\varepsilon_{\text{gen}}(t) = \frac{1}{4} \mathbb{E}_{\mathbf{X}, \mathbf{y}, \mathbf{x}_{\text{new}}, y_{\text{new}}} \left[(y_{\text{new}} - \hat{y}_{\text{new}}(\mathbf{w}(t)))^2 \right], \tag{23}$$

where $\{\mathbf{X}, \mathbf{y}\}$ is the training set, \mathbf{x}_{new} is an unseen data point and \hat{y}_{new} is the estimator for the new label y_{new} . The dependence on the training set here is hidden in the weight vector $\mathbf{w}(t) = \mathbf{w}(t, \mathbf{X}, \mathbf{y})$. In the two-cluster case one can easily show that

$$\varepsilon_{\text{gen}}(t) = \frac{1}{2} \text{erfc} \left(\frac{m(t)}{\sqrt{2\Delta C(t, t)}} \right). \tag{24}$$

Conversely, for the door activation trained on the three-cluster dataset we get

$$\varepsilon_{\text{gen}}(t) = \frac{1}{2} \text{erfc} \left(\frac{L}{\sqrt{2\Delta C(t, t)}} \right) + \frac{1}{4} \left(\text{erf} \left(\frac{L - m(t)}{\sqrt{2\Delta C(t, t)}} \right) + \text{erf} \left(\frac{L + m(t)}{\sqrt{2\Delta C(t, t)}} \right) \right). \tag{25}$$

5 Results

In this section, we compare the theoretical curves resulting from the solution of the DMFT equations derived in Sec. 4 to numerical simulations. This analysis allows to gain insight into the learning dynamics of stochastic gradient descent and its dependence on the various control parameters in the two models under consideration.

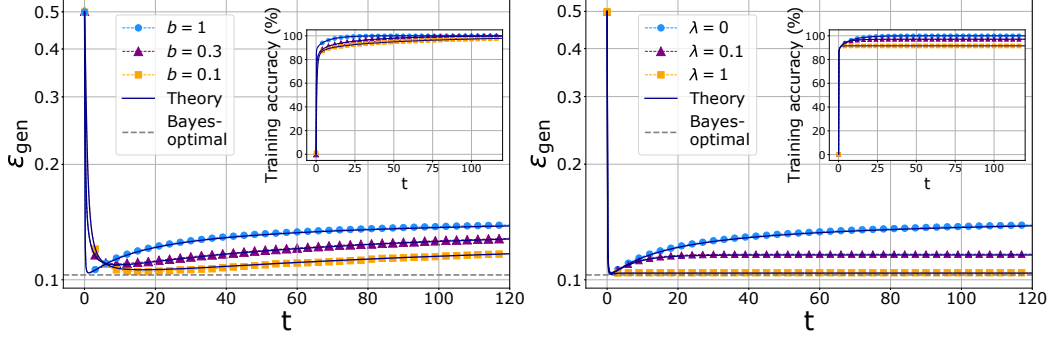


Figure 1: **Left:** Generalization error as a function of the training time for Persistent-SGD in the two-cluster model, with $\alpha = 2$, $\Delta = 0.5$, $\lambda = 0$, $1/\tau = 0.6$ and different batch sizes $b = 1, 0.3, 0.1$. The continuous lines mark the numerical solution of DMFT equations, while the symbols are the results of simulations at $d = 500$, $\eta = 0.2$, and $R = 0.01$. The dashed grey line marks the Bayes-optimal error from [28]. **Right:** Generalization error as a function of the training time for full-batch gradient descent in the two-cluster model with different regularization $\lambda = 0, 0.1, 1$ and the same parameters as in the left panel. In each panel, the inset shows the training accuracy as a function of the training time.

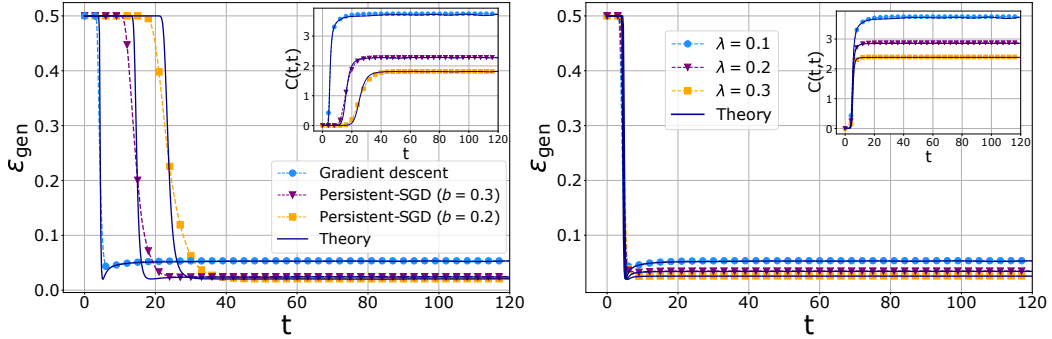


Figure 2: **Left:** Generalization error as a function of the training time in the three-cluster model, at fixed $\alpha = 3$, $\Delta = 0.05$, $L = 0.7$, $\lambda = 0.1$, for full-batch gradient descent and Persistent-SGD with different batch size $b = 0.2, 0.3$ and activation rate $1/\tau = b$. The continuous lines mark the numerical solution of DMFT equations, the symbols represent simulations at $\eta = 0.2$, $R = 0.01$, and $d = 5000$. **Right:** Generalization error as a function of training time for full-batch gradient descent in the three-cluster model, at fixed $\alpha = 3$, $\Delta = 0.05$, $L = 0.7$, $\eta = 0.2$, $R = 0.01$, and different regularization $\lambda = 0.1, 0.2, 0.3$. The simulations are done at $d = 5000$. In each panel, the inset shows the norm of the weights as a function of the training time.

The left panel of Fig. 1 shows the learning dynamics of the Persistent-SGD in the two-cluster model without regularization $\lambda = 0$. We clearly see a good match between the numerical simulations and the theoretical curves obtained from DMFT, notably also for small values of batchsize b and dimension $d = 500$. The figure shows that there exist regions in control parameter space where Persistent-SGD is able to reach 100% training accuracy, while the generalization error is bounded away from zero. Remarkably, we observe that the additional noise introduced by decreasing the batch size b results in a shift of the early-stopping minimum of the generalization error at larger times and that, on the time window we show, a batch size smaller than one has a beneficial effect on the generalization error at long times. The right panel illustrates the role of regularization in the same model trained with full-batch gradient descent, presenting that regularization has a similar influence on the learning curve as small batch-size but without the slow-down incurred by Persistent-SGD.

The influence of the batch size b and the regularization λ for the three-cluster model is shown in Fig. 2. We see an analogous effect as for the two-clusters in Fig. 1. In the inset of Fig. 2 we show the norm of the weights as a function of the training time. Both with the smaller mini-batch size and larger regularization the norm is small, testifying further that the two play a similar role in this case.

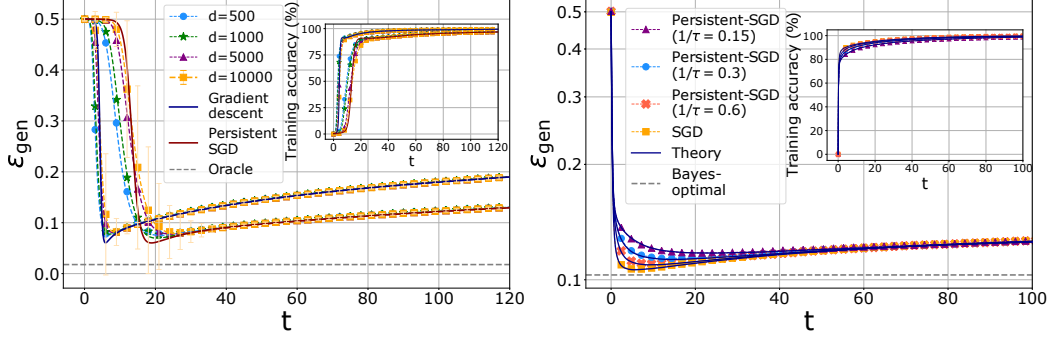


Figure 3: **Left:** Generalization error as a function of the training time for full-batch gradient descent and Persistent-SGD with $1/\tau = b = 0.3$ in the three-cluster model, at fixed $\alpha = 2$, $\Delta = 0.05$, $L = 0.7$ and $\lambda = 0$. The continuous lines mark the numerical solution of DMFT equations, the symbols represent simulations at $\eta = 0.2$, $R = 1$, and increasing dimension $d = 500, 1000, 5000, 10000$. Error bars are plotted for $d = 10000$. The dashed lines mark the oracle error (see supplementary material). **Right:** Generalization error as a function of the training time for Persistent-SGD with different activation rates $1/\tau = 0.15, 0.3, 0.6$ and classical SGD in the two-cluster model, both with $b = 0.3$, $\alpha = 2$, $\Delta = 0.5$, $\lambda = 0$, $\eta = 0.2$, $R = 0.01$. The continuous lines mark the numerical solution of DMFT equations (in case of SGD we use the SGD-inspired discretization), while the symbols represent simulations at $d = 500$. The dashed lines mark the Bayes-optimal error from [28]. In each panel, the inset displays the training accuracy as a function of time.

One difference between the two-cluster and the three-cluster models we observe concerns the behavior of the generalization error at small times. Actually, for the three-cluster model, good generalization is reached because of finite-size effects. Indeed, the corresponding loss function displays a \mathbb{Z}_2 symmetry according to which for each local minimum \mathbf{w} there is another one $-\mathbf{w}$ with exactly the same properties. Note that this symmetry is inherited from the activation function ϕ (6), which is even. This implies that if $d \rightarrow \infty$, the generalization error would not move away from 0.5 in finite time. However, when d is large but finite, at time $t = 0$ the weight vector has a finite projection on \mathbf{v}^* which is responsible for the dynamical symmetry breaking and eventually for a low generalization error at long times. In order to obtain an agreement between the theory and simulations, we initialize $m(t)$ in the DMFT equations with its corresponding finite- d average value at $t = 0$. In the left panel of Fig. 3 we show that while this produces a small discrepancy at intermediate times that diminishes with growing size, at long times the DMFT tracks perfectly the evolution of the algorithm.

The right panel of Fig. 3 summarizes the effect of the characteristic time τ in the Persistent-SGD, i.e. the typical persistence time of each pattern in the training mini-batch. When τ decreases, the Persistent-SGD algorithm is observed to be getting a better early-stopping generalization error and the dynamics gets closer to the usual SGD dynamics. As expected, the $\tau \rightarrow 0$ limit of the Persistent-SGD converges to the SGD. It is remarkable that the SGD-inspired discretization of the DMFT equations, that is in principle an ad-hoc construction as the corresponding flow-limit in which the derivation holds does not exist, shows a perfect agreement with the numerics.

Fig. 4 presents the influence of the weight norm at initialization R on the dynamics, for the two-cluster (left) and three-cluster (right) model. For the two-cluster case, the gradient descent algorithm with all-zeros initialization “jumps” on the Bayes-optimal error at the first iteration as derived in [28], and in this particular setting the generalization error is monotonically increasing in time. As R increases the early stopping error gets worse. At large times all the initializations converge to the same value of the error, as they must, since this is a full-batch gradient descent without regularization that at large times converges to the max-margin estimator according to [29]. For the three-cluster model we observe a qualitatively similar behavior.

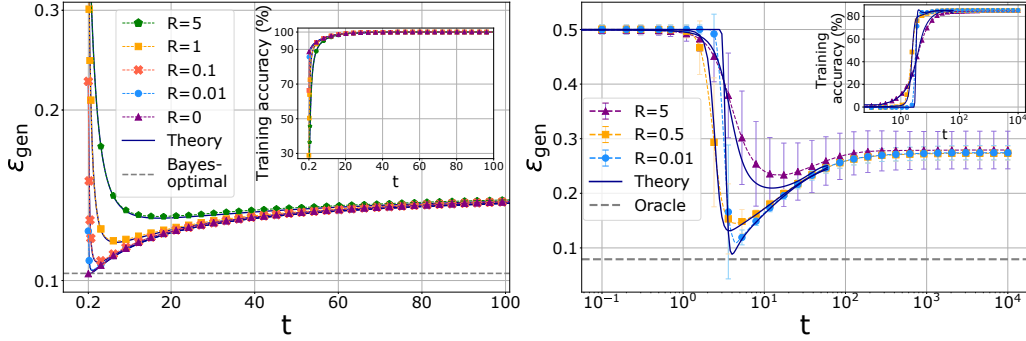


Figure 4: **Left:** Generalization error as a function of training time for full-batch gradient descent in the two-cluster model, at fixed $\alpha = 2$, $\Delta = 0.5$, $\lambda = 0$, $\eta = 0.2$, and different initialization variances $R = 0, 0.01, 0.1, 1, 5$. The continuous lines mark the numerical solution of DMFT equations, while the symbols represent simulations at $d = 500$. The dashed lines mark the Bayes-optimal error from [28]. The y -axis is cut for better visibility. **Right:** Generalization error as a function of training time for full-batch gradient descent in the three-cluster model, at fixed $\alpha = 3$, $\Delta = 0.1$, $\lambda = 0$, $\eta = 0.1$ and different initialization variances $R = 0.01, 0.5, 5$. The continuous lines mark the numerical solution of DMFT equations, while the symbols represent simulations at $d = 1000$. The dashed grey line marks the oracle error (see supplementary material). In each panel, the inset shows the training accuracy as a function of time.

Broader Impact

Our work is theoretical in nature, and as such the potential societal consequences are difficult to foresee. We anticipate that deeper theoretical understanding of the functioning of machine learning systems will lead to their improvement in the long term.

Acknowledgments and Disclosure of Funding

This work was supported by "Investissements d'Avenir" LabExpALM (ANR-10-LABX-0039-PALM), the ERC under the European Union's Horizon 2020 Research and Innovation Program 714608-SMiLe, as well as by the French Agence Nationale de la Recherche under grant ANR-17-CE23-0023-01 PAIL and ANR-19-P3IA-0001 PRAIRIE.

References

- [1] Itay Safran and Ohad Shamir. Spurious local minima are common in two-layer relu neural networks. *arXiv preprint arXiv:1712.08968*, 2017.
- [2] Shengchao Liu, Dimitris Papailiopoulos, and Dimitris Achlioptas. Bad global minima exist and sgd can reach them. *arXiv preprint arXiv:1906.02613*, 2019.
- [3] Siegfried Bös and Manfred Opper. Dynamics of training. In *Advances in Neural Information Processing Systems*, pages 141–147, 1997.
- [4] Andrew M Saxe, James L McClelland, and Surya Ganguli. Exact solutions to the nonlinear dynamics of learning in deep linear neural networks. *arXiv preprint arXiv:1312.6120*, 2013.
- [5] David Saad and Sara A Solla. Exact solution for on-line learning in multilayer neural networks. *Physical Review Letters*, 74(21):4337, 1995.
- [6] David Saad and Sara A Solla. On-line learning in soft committee machines. *Physical Review E*, 52(4):4225, 1995.
- [7] David Saad. *On-line learning in neural networks*, volume 17. Cambridge University Press, 2009.

- [8] 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 *Advances in Neural Information Processing Systems*, pages 6979–6989, 2019.
- [9] Sebastian Goldt, Marc Mézard, Florent Krzakala, and Lenka Zdeborová. Modelling the influence of data structure on learning in neural networks. *arXiv preprint arXiv:1909.11500*, 2019.
- [10] Grant M Rotskoff and Eric Vanden-Eijnden. Neural networks as interacting particle systems: Asymptotic convexity of the loss landscape and universal scaling of the approximation error. *arXiv preprint arXiv:1805.00915*, 2018.
- [11] 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.
- [12] Lenaic Chizat and Francis Bach. On the global convergence of gradient descent for over-parameterized models using optimal transport. In *Advances in neural information processing systems*, pages 3036–3046, 2018.
- [13] Ben Poole, Subhaneil Lahiri, Maithra Raghu, Jascha Sohl-Dickstein, and Surya Ganguli. Exponential expressivity in deep neural networks through transient chaos. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems 29*, pages 3360–3368. Curran Associates, Inc., 2016.
- [14] Samuel S. Schoenholz, Justin Gilmer, Surya Ganguli, and Jascha Sohl-Dickstein. Deep information propagation, 2017.
- [15] Greg Yang, Jeffrey Pennington, Vinay Rao, Jascha Sohl-Dickstein, and Samuel S. Schoenholz. A mean field theory of batch normalization. In *International Conference on Learning Representations*, 2019.
- [16] Song Mei, Theodor Misiakiewicz, and Andrea Montanari. Mean-field theory of two-layers neural networks: dimension-free bounds and kernel limit, 2019.
- [17] Dar Gilboa, Bo Chang, Minmin Chen, Greg Yang, Samuel S. Schoenholz, Ed H. Chi, and Jeffrey Pennington. Dynamical isometry and a mean field theory of lstms and grus, 2019.
- [18] Roman Novak, Lechao Xiao, Yasaman Bahri, Jaehoon Lee, Greg Yang, Daniel A. Abolafia, Jeffrey Pennington, and Jascha Sohl-dickstein. Bayesian deep convolutional networks with many channels are gaussian processes. In *International Conference on Learning Representations*, 2019.
- [19] Marc Mézard, Giorgio Parisi, and Miguel A. Virasoro. *Spin glass theory and beyond*. World Scientific, Singapore, 1987.
- [20] Antoine Georges, Gabriel Kotliar, Werner Krauth, and Marcelo J Rozenberg. Dynamical mean-field theory of strongly correlated fermion systems and the limit of infinite dimensions. *Reviews of Modern Physics*, 68(1):13, 1996.
- [21] Giorgio Parisi, Pierfrancesco Urbani, and Francesco Zamponi. *Theory of Simple Glasses: Exact Solutions in Infinite Dimensions*. Cambridge University Press, 2020.
- [22] Marylou Gabrié. Mean-field inference methods for neural networks. *Journal of Physics A: Mathematical and Theoretical*, 53(22):223002, may 2020.
- [23] G Ben Arous, Alice Guionnet, et al. Symmetric langevin spin glass dynamics. *The Annals of Probability*, 25(3):1367–1422, 1997.
- [24] Stefano Sarao Mannelli, Giulio Biroli, Chiara Cammarota, Florent Krzakala, Pierfrancesco Urbani, and Lenka Zdeborová. Marvels and pitfalls of the langevin algorithm in noisy high-dimensional inference. *Physical Review X*, 10(1):011057, 2020.
- [25] Stefano Sarao Mannelli, Florent Krzakala, Pierfrancesco Urbani, and Lenka Zdeborova. Passed & spurious: Descent algorithms and local minima in spiked matrix-tensor models. In *international conference on machine learning*, pages 4333–4342, 2019.

- [26] Silvio Franz, Giorgio Parisi, Maxim Sevelev, Pierfrancesco Urbani, and Francesco Zamponi. Universality of the sat-unsat (jamming) threshold in non-convex continuous constraint satisfaction problems. *SciPost Physics*, 2(3):019, 2017.
- [27] Silvio Franz, Sungmin Hwang, and Pierfrancesco Urbani. Jamming in multilayer supervised learning models. *Physical review letters*, 123(16):160602, 2019.
- [28] Francesca Mignacco, Florent Krzakala, Yue M Lu, and Lenka Zdeborová. The role of regularization in classification of high-dimensional noisy gaussian mixture. *arXiv preprint arXiv:2002.11544*, 2020.
- [29] Saharon Rosset, Ji Zhu, and Trevor J Hastie. Margin maximizing loss functions. In *Advances in neural information processing systems*, pages 1237–1244, 2004.
- [30] Zeyu Deng, Abla Kammoun, and Christos Thrampoulidis. A model of double descent for high-dimensional binary linear classification. *arXiv preprint arXiv:1911.05822*, 2019.
- [31] H Sebastian Seung, Manfred Opper, and Haim Sompolinsky. Query by committee. In *Proceedings of the fifth annual workshop on Computational learning theory*, pages 287–294, 1992.
- [32] Elisabeth Agoritsas, Giulio Biroli, Pierfrancesco Urbani, and Francesco Zamponi. Out-of-equilibrium dynamical mean-field equations for the perceptron model. *Journal of Physics A: Mathematical and Theoretical*, 51(8):085002, 2018.
- [33] Cirano de Dominicis. Technics of field renormalization and dynamics of critical phenomena. In *J. Phys.(Paris), Colloq*, pages C1–247, 1976.
- [34] H Eissfeller and M Opper. New method for studying the dynamics of disordered spin systems without finite-size effects. *Physical review letters*, 68(13):2094, 1992.
- [35] Daniel Soudry, Elad Hoffer, Mor Shpigel Nacson, Suriya Gunasekar, and Nathan Srebro. The implicit bias of gradient descent on separable data. *The Journal of Machine Learning Research*, 19(1):2822–2878, 2018.

Dynamical mean-field theory for stochastic gradient descent in Gaussian mixture classification – supplementary material

Francesca Mignacco¹

Florent Krzakala^{2,3}

Pierfrancesco Urbani¹

Lenka Zdeborová^{1,4}

¹ Institut de physique théorique, Université Paris-Saclay, CNRS, CEA, Gif-sur-Yvette, France

² Laboratoire de Physique, CNRS, École Normale Supérieure, PSL University, Paris, France

³ IdePHICS Laboratory, EPFL, Switzerland

⁴ SPOC Laboratory, EPFL, Switzerland

Correspondence to: francesca.mignacco@ipht.fr

A Derivation of the dynamical mean-field equations

The derivation of the self-consistent stochastic process discussed in the main text can be obtained using tools of statistical physics of disordered systems. In particular, it has been done very recently for a related model, the spherical perceptron with random labels, in [1]. Our derivation extends the known DMFT equations by including

- structure in the data;
- a stochastic version of gradient descent as discussed in the main text;
- the relaxation of the spherical constraint over the weights and the introduction of a Ridge regularization term.

There are at least two ways to write the DMFT equations. One is by using field theoretical techniques; otherwise one can employ a dynamical version of the so-called *cavity method* [2]. Here we opt for the first option that is generically very compact and immediate and it has a form that resembles very much a *static* treatment of the Gibbs measure of the problem [3]. We use a supersymmetric (SUSY) representation to derive the dynamical mean-field (DMFT) equations [1, 4]. We do not report all the details, that can be found in [1] along with an alternative derivation based on the cavity method, but we limit ourselves to provide the main points. We first consider the dynamical partition function, corresponding to Eq. (11) in the main text

$$\begin{aligned}
 Z_{\text{dyn}} = & \left\langle \int \left[\frac{d\mathbf{w}^{(0)}}{(2\pi)^{\frac{d}{2}}} e^{-\frac{1}{2}\|\mathbf{w}^{(0)}\|_2^2} \right] \int_{\mathbf{w}(0)=\mathbf{w}^{(0)}} \mathcal{D}\mathbf{w}(t) \right. \\
 & \left. \times \prod_{j=1}^d \delta \left[-\dot{w}_j(t) - \lambda w_j(t) - \sum_{\mu=1}^n s_{\mu}(t) \Lambda' \left(y_{\mu}, \frac{\mathbf{w}(t)^{\top} \mathbf{x}_{\mu}}{\sqrt{d}} \right) \frac{x_{\mu,j}}{\sqrt{d}} \right] \right\rangle, \tag{A.1}
 \end{aligned}$$

where the brackets $\langle \cdot \rangle$ stand for the average over $s_{\mu}(t)$, y_{μ} and the realization of the noise in the training set. The average over the initial condition is written explicitly. Note that we choose an initial condition that is Gaussian, but we could have chosen a different probability measure over the initial configuration of the weights. The equations can be generalized to other initial conditions as soon as they do not depend on quenched random variables that enter in the stochastic gradient descent (SGD) dynamics and their distribution is separable. As observed in the main text, we have

that $Z_{\text{dyn}} = \langle Z_{\text{dyn}} \rangle = 1$. We can write the integral representation of the Dirac delta function in Eq. [A.1](#) by introducing a set of fields $\hat{\mathbf{w}}(t)$

$$Z_{\text{dyn}} = \left\langle \int \mathcal{D}\mathbf{w}(t) \mathcal{D}\hat{\mathbf{w}}(t) e^{S_{\text{dyn}}} \right\rangle, \quad (\text{A.2})$$

where the dynamical action S_{dyn} is defined as in Eq. [\(I3\)](#) of the main text

$$S_{\text{dyn}} = \sum_{j=1}^d \int_0^{+\infty} dt i \hat{w}_j(t) \left(-\dot{w}_j(t) - \lambda w_j(t) - \sum_{\mu=1}^n s_{\mu}(t) \Lambda' \left(y_{\mu}, \frac{\mathbf{w}(t)^{\top} \mathbf{x}_{\mu}}{\sqrt{d}} \right) \frac{x_{\mu,j}}{\sqrt{d}} \right). \quad (\text{A.3})$$

A.1 SUSY formulation

The dynamical action S_{dyn} [\(A.3\)](#) can be rewritten in a supersymmetric form, by extending the time coordinate to include two Grassman coordinates θ and $\bar{\theta}$, i.e. $t_a \rightarrow a = (t_a, \theta_a, \bar{\theta}_a)$. The dynamic variable $\mathbf{w}(t_a)$ and the auxiliary variable $i \hat{\mathbf{w}}(t_a)$ are encoded in a super-field

$$\mathbf{w}(a) = \mathbf{w}(t_a) + i \theta_a \bar{\theta}_a \hat{\mathbf{w}}(t_a). \quad (\text{A.4})$$

From the properties of Grassman variables [\[5\]](#)

$$\begin{aligned} \theta^2 = \bar{\theta}^2 = \theta \bar{\theta} + \bar{\theta} \theta &= 0, \\ \int d\theta &= \int d\bar{\theta} = 0, \quad \int d\theta \theta = \int d\bar{\theta} \bar{\theta} = 1, \\ \partial_{\theta} g(\theta) &= \int d\theta g(\theta) \quad \text{for a generic function } g, \end{aligned} \quad (\text{A.5})$$

it follows that

$$\int da f(\mathbf{w}(a)) = \int_0^{+\infty} dt_a i \hat{\mathbf{w}}(t_a) f'(\mathbf{w}(t_a)). \quad (\text{A.6})$$

We can use Eq. [\(A.6\)](#) to rewrite S_{dyn} . We obtain

$$S_{\text{dyn}} = -\frac{1}{2} \int da db \mathcal{K}(a, b) \mathbf{w}(a)^{\top} \mathbf{w}(b) - \sum_{\mu=1}^n \int da s_{\mu}(a) \Lambda(y_{\mu}, h_{\mu}(a)), \quad (\text{A.7})$$

where we have defined $h_{\mu}(a) \equiv \mathbf{w}(a)^{\top} \mathbf{x}_{\mu} / \sqrt{d}$ and we have implicitly defined the kernel $\mathcal{K}(a, b)$ such that

$$-\frac{1}{2} \int da db \mathcal{K}(a, b) \mathbf{w}(a)^{\top} \mathbf{w}(b) = \sum_{j=1}^d \int_0^{+\infty} dt i \hat{w}_j(t) (-\dot{w}_j(t) - \lambda w_j(t)). \quad (\text{A.8})$$

By inserting the definition of $h_{\mu}(a)$ in the partition function, we have

$$\begin{aligned} Z_{\text{dyn}} &= \left\langle \int \mathcal{D}\mathbf{w}(a) \mathcal{D}h_{\mu}(a) \mathcal{D}\hat{h}_{\mu}(a) \exp \left[-\frac{1}{2} \int da db \mathcal{K}(a, b) \mathbf{w}(a)^{\top} \mathbf{w}(b) \right. \right. \\ &\quad \left. \left. - \sum_{\mu=1}^n \int da s_{\mu}(a) \Lambda(y_{\mu}, h_{\mu}(a)) \right] \exp \left[\sum_{\mu=1}^n \int da i \hat{h}_{\mu}(a) \left(h_{\mu}(a) - \frac{\mathbf{w}(a)^{\top} \mathbf{x}_{\mu}}{\sqrt{d}} \right) \right] \right\rangle. \end{aligned} \quad (\text{A.9})$$

Let us consider the last factor in the integral in [\(A.9\)](#). We can perform the average over the random vectors $\mathbf{z}_{\mu} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_d)$, denoted by an overline, as

$$\begin{aligned} &\overline{\exp \left[\sum_{\mu=1}^n \int da i \hat{h}_{\mu}(a) \left(h_{\mu}(a) - \frac{\mathbf{w}(a)^{\top} \mathbf{x}_{\mu}}{\sqrt{d}} \right) \right]} \\ &= \exp \left[\sum_{\mu=1}^n \int da i \hat{h}_{\mu}(a) \left(h_{\mu}(a) - c_{\mu} m(a) - \sqrt{\frac{\Delta}{d}} \mathbf{w}(a)^{\top} \mathbf{z}_{\mu} \right) \right] \\ &= \exp \left[\sum_{\mu=1}^n \int da i \hat{h}_{\mu}(a) \left(h_{\mu}(a) - c_{\mu} m(a) \right) - \frac{\Delta}{2} \sum_{\mu=1}^n \int da db Q(a, b) \hat{h}_{\mu}(a) \hat{h}_{\mu}(b) \right], \end{aligned} \quad (\text{A.10})$$

where we have defined

$$\begin{aligned} m(a) &= \frac{1}{d} \mathbf{w}(a)^\top \mathbf{v}^*, \\ Q(a, b) &= \frac{1}{d} \mathbf{w}(a)^\top \mathbf{w}(b). \end{aligned} \quad (\text{A.11})$$

By inserting the definitions of $m(a)$ and $Q(a, b)$ in the partition function, we obtain

$$Z_{\text{dyn}} = \int \mathcal{D}\mathbf{Q} \mathcal{D}\mathbf{m} e^{dS(\mathbf{Q}, \mathbf{m})}, \quad (\text{A.12})$$

where $\mathbf{Q} = \{Q(a, b)\}_{a, b}$, $\mathbf{m} = \{m(a)\}_a$ and

$$\begin{aligned} S(\mathbf{Q}, \mathbf{m}) &= \frac{1}{2} \log \det (Q(a, b) - m(a)m(b)) - \frac{1}{2} \int da db \mathcal{K}(a, b) Q(a, b) + \alpha \log \mathcal{Z}, \\ \mathcal{Z} &= \left\langle \int \mathcal{D}h(a) \mathcal{D}\hat{h}(a) \exp \left[-\frac{\Delta}{2} \int da db Q(a, b) \hat{h}(a) \hat{h}(b) \right. \right. \\ &\quad \left. \left. + \int da i\hat{h}(a) (h(a) - cm(a)) - \int da s(a) \Lambda(y, h(a)) \right] \right\rangle. \end{aligned} \quad (\text{A.13})$$

We have used that the samples are i.i.d. and removed the index $\mu = 1, \dots, n$. The brackets denote the average over the random variable c , that has the same distribution as the c_μ , over y , distributed as y_μ , and over the random process of $s(t)$, defined by Eq. (9) in the main text. If we perform the change of variable $Q(a, b) \leftarrow Q(a, b) + m(a)m(b)$, we obtain

$$\begin{aligned} S(\mathbf{Q}, \mathbf{m}) &= \frac{1}{2} \log \det Q(a, b) - \frac{1}{2} \int da db \mathcal{K}(a, b) (Q(a, b) + m(a)m(b)) + \alpha \log \mathcal{Z}, \\ \mathcal{Z} &= \left\langle \int \mathcal{D}h(a) \mathcal{D}\hat{h}(a) e^{S_{\text{loc}}} \right\rangle, \end{aligned} \quad (\text{A.14})$$

where the effective local action S_{loc} is given by

$$\begin{aligned} S_{\text{loc}} &= -\frac{\Delta}{2} \int da db Q(a, b) \hat{h}(a) \hat{h}(b) - \frac{\Delta}{2} \left(\int da \hat{h}(a) m(a) \right)^2 \\ &\quad + \int da i\hat{h}(a) (h(a) - cm(a)) - \int da s(a) \Lambda(y, h(a)). \end{aligned} \quad (\text{A.15})$$

Performing a Hubbard-Stratonovich transformation on $\exp \left[-\frac{\Delta}{2} \left(\int da \hat{h}(a) m(a) \right)^2 \right]$ and a set of transformations on the fields $h(a)$, we obtain that we can rewrite \mathcal{Z} as

$$\begin{aligned} \mathcal{Z} &= \left\langle \int \frac{dh_0}{\sqrt{2\pi}} e^{-\frac{h_0^2}{2}} \int \mathcal{D}h(a) \mathcal{D}\hat{h}(a) \exp \left[-\frac{1}{2} \int da db Q(a, b) \hat{h}(a) \hat{h}(b) \right. \right. \\ &\quad \left. \left. + \int da i\hat{h}(a) h(a) - \int da s(a) \Lambda \left(y, \sqrt{\Delta} h(a) + m(a)(c + \sqrt{\Delta} h_0) \right) \right] \right\rangle. \end{aligned} \quad (\text{A.16})$$

A.2 Saddle-point equations

We are interested in the large d limit of Z_{dyn} , in which, according to Eq. (A.12), the partition function is dominated by the saddle-point value of $S(\mathbf{Q}, \mathbf{m})$:

$$\begin{cases} \left. \frac{\delta S(\mathbf{Q}, \mathbf{m})}{\delta Q(a, b)} \right|_{(\mathbf{Q}, \mathbf{m})=(\tilde{\mathbf{Q}}, \tilde{\mathbf{m}})} = 0 \\ \left. \frac{\delta S(\mathbf{Q}, \mathbf{m})}{\delta m(a)} \right|_{(\mathbf{Q}, \mathbf{m})=(\tilde{\mathbf{Q}}, \tilde{\mathbf{m}})} = 0 \end{cases}. \quad (\text{A.17})$$

$\tilde{Q}(a, b)$ is obtained from the equation

$$-\mathcal{K}(a, b) + Q^{-1}(a, b) + \frac{2\alpha}{\mathcal{Z}} \frac{\delta \mathcal{Z}}{\delta Q(a, b)} = 0. \quad (\text{A.18})$$

The saddle-point equation for $\tilde{m}(a)$ is instead

$$-\int db \mathcal{K}(a, b)m(b) + \frac{\alpha}{\mathcal{Z}} \frac{\delta \mathcal{Z}}{\delta m(a)} = 0. \quad (\text{A.19})$$

It can be easily shown by exploiting the Grassmann structure of Eqs. (A.18)-(A.19) that they lead to a self consistent stochastic process described by

$$\dot{h}(t) = -\tilde{\lambda}(t)h(t) - \sqrt{\Delta}s(t)\Lambda'(y, r(t) - Y(t)) + \int_0^t dt' M_R(t, t')h(t') + \xi(t), \quad (\text{A.20})$$

where the initial condition is drawn from $P(h(0)) \sim e^{-h(0)^2/(2R)}/\sqrt{2\pi}$, and $r(t) = \sqrt{\Delta}h(t) + m(t)(c + \sqrt{\Delta}h_0)$, with $P_0(h_0) \sim e^{-h_0^2/2}/\sqrt{2\pi}$. We have defined the auxiliary functions

$$\begin{aligned} \mu(t) &= \alpha \left\langle s(t) \left(c + \sqrt{\Delta}h_0 \right) \Lambda'(y, r(t)) \right\rangle \\ \hat{\lambda}(t) &= \alpha \Delta \left\langle s(t) \Lambda''(y, r(t)) \right\rangle \\ \tilde{\lambda}(t) &= \lambda + \hat{\lambda}(t) \end{aligned} \quad (\text{A.21})$$

and kernels

$$\begin{aligned} M_C(t, t') &= \alpha \Delta \left\langle s(t)s(t') \Lambda'(y, r(t)) \Lambda'(y, r(t')) \right\rangle, \\ M_R(t, t') &= \alpha \Delta^{3/2} \left\langle s(t)s(t') \Lambda'(y, r(t)) \Lambda''(y, r(t')) i\hat{h}(t') \right\rangle \\ &\equiv \alpha \Delta \frac{\delta}{\delta Y(t')} \left\langle s(t) \Lambda'(y, r(t)) \right\rangle \Big|_{Y=0}. \end{aligned} \quad (\text{A.22})$$

In addition, from (A.19), one can derive an ordinary differential equation for the magnetization

$$\dot{m}(t) = -\lambda m(t) - \mu(t). \quad (\text{A.23})$$

The brackets in the previous equations denote, at the same time, the average over the label y , the process $s(t)$, as well as the average over the noise $\xi(t)$ and both h_0 and $h(0)$, whose probability distributions are given by $P(h(0))$ and $P_0(h_0)$ respectively. In other words, one has a set of kernels, such as $M_R(t, t')$ and $M_C(t, t')$, that can be obtained as average over the stochastic process for $h(t)$ and therefore must be computed self-consistently.

Finally, Eq. (A.18) gives rise to Eq. (20) of the main text while Eq. (A.19) gives rise to the equation for the evolution of the magnetization. Note that the norm of the weight vector $\mathbf{w}(t)$ can be also computed by sampling the stochastic process

$$\begin{aligned} \dot{\mathbf{w}}(t) &= -\tilde{\lambda}(t)\mathbf{w}(t) + \int_0^t dt' M_R(t, t')(\mathbf{w}(t') - m(t')h_0) + \xi(t) + h_0(\hat{\lambda}(t)m(t) - \mu(t)), \\ P(\mathbf{w}_0) &= \frac{1}{\sqrt{2\pi R}} e^{-\mathbf{w}_0^2/(2R)}, \end{aligned} \quad (\text{A.24})$$

from which one gets

$$C(t, t') = \langle \mathbf{w}(t)^2 \rangle. \quad (\text{A.25})$$

A.3 Numerical solution of DMFT equations

The algorithm to solve the DMFT equations that are summed up in Eq. (A.20) is the most natural one. It can be understood in this way. The outcome of the DMFT is the computation of the kernels and functions appearing in it, namely $m(t)$, $M_C(t, t')$ and so on. They are determined as averages over the stochastic process that is defined through them. Therefore, one needs to solve the system of equations in a self-consistent way. The straightforward way to do that is to proceed by iterations:

1. We start from a random guess of the kernels, that we use to sample the stochastic process (A.20) several times;
2. We compute the averages over these multiple realizations to obtain the updates of the auxiliary functions (A.21) and kernels (A.22), along with the magnetization (A.23);
3. We use these new guesses to sample again multiple realizations of the stochastic process;
4. We repeat steps 2. and 3. until the kernels reach a fixed point.

As in all iterative solutions of fixed point equations, it is natural to introduce some damping in the update of the kernels to avoid wild oscillations. Note that the DMFT fixed point equations are deterministic, hence at given initial condition the solution is unique. Indeed, the kernels computed by DMFT are causal and a simple integration scheme of the equations is just extending them progressively in time starting from their initial value, which is completely deterministic given the initial condition for the stochastic process. This procedure has been first implemented in [6, 7] and recently developed further in other applications [8, 9]. However, DMFT has a long tradition in condensed matter physics [10] where more involved algorithms have been developed.

B Generalization error

The generalization error at any time step is defined as the fraction of mislabeled instances:

$$\varepsilon_{\text{gen}}(t) \equiv \frac{1}{4} \mathbb{E}_{\mathbf{X}, \mathbf{y}, \mathbf{x}_{\text{new}}, y_{\text{new}}} \left[(y_{\text{new}} - \hat{y}_{\text{new}}(\mathbf{w}(t)))^2 \right], \quad (\text{B.1})$$

where $\{\mathbf{X}, \mathbf{y}\}$ is the training set, \mathbf{x}_{new} is an unseen data point and \hat{y}_{new} is the estimator for the new label y_{new} . The dependence on the training set here is hidden in the weight vector $\mathbf{w}(t) = \mathbf{w}(t, \mathbf{X}, \mathbf{y})$.

B.1 Perceptron with linear activation function

In this case, the estimator for a new label is $\hat{y}_{\text{new}}(\mathbf{w}(t)) = \text{sign}(\mathbf{w}(t)^\top \mathbf{x}_{\text{new}})$. The generalization error in the infinite dimensional limit $d \rightarrow \infty$ has been computed in [11] and reads

$$\varepsilon_{\text{gen}}(t) = \frac{1}{2} \text{erfc} \left(\frac{m(t)}{\sqrt{2\Delta C(t, t)}} \right). \quad (\text{B.2})$$

B.2 Perceptron with door activation function

In this case, the estimator for a new label is $\hat{y}_{\text{new}}(\mathbf{w}(t)) = \text{sign}(\frac{1}{d}(\mathbf{w}(t)^\top \mathbf{x}_{\text{new}})^2 - L^2)$. From Eq. (B.1), we have that

$$\varepsilon_{\text{gen}}(t) = \frac{1}{2} (1 - \mathbb{E}_{\mathbf{X}, \mathbf{y}, \mathbf{x}_{\text{new}}, y_{\text{new}}} [y_{\text{new}} \cdot \hat{y}_{\text{new}}(\mathbf{w}(t))]). \quad (\text{B.3})$$

We consider the second term of (B.3)

$$\mathbb{E}_{\mathbf{X}, \mathbf{y}, \mathbf{x}_{\text{new}}, y_{\text{new}}} [y_{\text{new}} \cdot \hat{y}_{\text{new}}(\mathbf{w}(t))] = \mathbb{E}_{\mathbf{X}, \mathbf{y}, \mathbf{x}_{\text{new}}} \left[\text{sign} \left(\frac{y_{\text{new}}}{d} (\mathbf{w}(t)^\top \mathbf{x}_{\text{new}})^2 - y_{\text{new}} L^2 \right) \right]. \quad (\text{B.4})$$

In the high dimensional limit, the overlap between weight vector and data point at each time step concentrates

$$\frac{\mathbf{w}(t)^\top \mathbf{x}_{\text{new}}}{\sqrt{d}} = \frac{\mathbf{w}(t)^\top}{\sqrt{d}} \left(c_{\text{new}} \frac{\mathbf{v}^*}{\sqrt{d}} + \sqrt{\Delta} \mathbf{z}_{\text{new}} \right) \xrightarrow{d \rightarrow \infty} c_{\text{new}} m(t) + \sqrt{\Delta C(t, t)} z, \quad (\text{B.5})$$

where $z \sim \mathcal{N}(0, 1)$. Therefore, we obtain

$$\begin{aligned} & \mathbb{E}_{\mathbf{X}, \mathbf{y}, \mathbf{x}_{\text{new}}, y_{\text{new}}} [y_{\text{new}} \cdot \hat{y}_{\text{new}}(\mathbf{w}(t))] \simeq \\ & \simeq \mathbb{E}_{c_{\text{new}}, z, y_{\text{new}}} \left[\text{sign} \left(y_{\text{new}} \left(c_{\text{new}} m(t) + \sqrt{\Delta C(t, t)} z \right)^2 - y_{\text{new}} L^2 \right) \right] \\ & = \mathbb{P} \left(y_{\text{new}} \left(c_{\text{new}} m(t) + \sqrt{\Delta C(t, t)} z \right)^2 \geq y_{\text{new}} L^2 \right) \\ & \quad - \mathbb{P} \left(y_{\text{new}} \left(c_{\text{new}} m(t) + \sqrt{\Delta C(t, t)} z \right)^2 < y_{\text{new}} L^2 \right) \end{aligned} \quad (\text{B.6})$$

and the generalization error in the infinite dimensional limit $d \rightarrow \infty$ is

$$\varepsilon_{\text{gen}}(t) = (1-\rho)\text{erfc}\left(\frac{L}{\sqrt{2\Delta C(t,t)}}\right) + \frac{\rho}{2}\left(\text{erf}\left(\frac{L-m(t)}{\sqrt{2\Delta C(t,t)}}\right) + \text{erf}\left(\frac{L+m(t)}{\sqrt{2\Delta C(t,t)}}\right)\right). \quad (\text{B.7})$$

C Oracle error

We call *oracle error* the classification error made by an ideal oracle that has access to the vector \mathbf{v}^* that characterizes the centers of the clusters in the two models under consideration (see Sec. 2 in the main text). We define the oracle's estimator \hat{y}_{new}^O given a new data point \mathbf{x}_{new} as

$$\hat{y}_{\text{new}}^O = \arg \max_{\tilde{y}_{\text{new}}} \mathbb{P}(\tilde{y}_{\text{new}} | \mathbf{x}_{\text{new}}), \quad (\text{C.1})$$

where the prior over the label \tilde{y}_{new} and the coefficient \tilde{c}_{new} along with the channel distribution

$$\mathbb{P}(\mathbf{x}_{\text{new}} | \tilde{c}_{\text{new}}) \propto \exp\left[-\frac{1}{2\Delta} \|\mathbf{x}_{\text{new}} - \frac{\tilde{c}_{\text{new}}}{\sqrt{d}} \mathbf{v}^*\|_2^2\right] \quad (\text{C.2})$$

are known. We can rewrite the probability in Eq. (C.1) as

$$\begin{aligned} \mathbb{P}(\tilde{y}_{\text{new}} | \mathbf{x}_{\text{new}}) &\propto \sum_{\tilde{c}_{\text{new}}=0,\pm 1} \mathbb{P}(\tilde{y}_{\text{new}}, \tilde{c}_{\text{new}}) \mathbb{P}(\mathbf{x}_{\text{new}} | \tilde{c}_{\text{new}}) = (1-\rho)\delta(\tilde{y}_{\text{new}} + 1)e^{-\frac{1}{2\Delta} \|\mathbf{x}_{\text{new}}\|_2^2} \\ &\quad + \frac{\rho}{2}\delta(\tilde{y}_{\text{new}} - 1)\left(e^{-\frac{1}{2\Delta} \|\mathbf{x}_{\text{new}} - \frac{1}{\sqrt{d}} \mathbf{v}^*\|_2^2} + e^{-\frac{1}{2\Delta} \|\mathbf{x}_{\text{new}} + \frac{1}{\sqrt{d}} \mathbf{v}^*\|_2^2}\right) \\ &= e^{-\frac{1}{2\Delta} \|\mathbf{x}_{\text{new}}\|_2^2} \left[(1-\rho)\delta(\tilde{y}_{\text{new}} + 1) + \rho\delta(\tilde{y}_{\text{new}} - 1)e^{-\frac{1}{2\Delta} \cosh\left(\frac{1}{\Delta\sqrt{d}} \mathbf{x}_{\text{new}}^\top \mathbf{v}^*\right)} \right]. \end{aligned} \quad (\text{C.3})$$

The oracle error is then

$$\varepsilon_{\text{gen}}^O = \mathbb{P}(\hat{y}_{\text{new}}^O \neq y_{\text{new}}) = (1-\rho)\mathbb{P}(\hat{y}_{\text{new}}^O = 1 | y_{\text{new}} = -1) + \rho\mathbb{P}(\hat{y}_{\text{new}}^O = -1 | y_{\text{new}} = 1). \quad (\text{C.4})$$

We can compute the two terms in the above equation separately

$$\begin{aligned} \mathbb{P}(\hat{y}_{\text{new}}^O = 1 | y_{\text{new}} = -1) &= \mathbb{P}\left(\rho e^{-\frac{1}{2\Delta}} \cosh\left(\frac{1}{\sqrt{\Delta d}} \mathbf{z}_{\text{new}}^\top \mathbf{v}^*\right) > 1 - \rho\right) \\ &= \mathbb{P}\left(\rho e^{-\frac{1}{2\Delta}} \cosh\left(\frac{\zeta_{\text{new}}}{\sqrt{\Delta}}\right) > 1 - \rho\right) = \text{erfc}\left(\sqrt{\frac{\Delta}{2}} \left| \text{arccosh}\left(\frac{(1-\rho)}{\rho} e^{1/2\Delta}\right) \right|\right), \end{aligned} \quad (\text{C.5})$$

and

$$\begin{aligned} \mathbb{P}(\hat{y}_{\text{new}}^O = -1 | y_{\text{new}} = 1) &= \mathbb{P}\left(1 - \rho > \rho e^{-\frac{1}{2\Delta}} \cosh\left(\frac{c_{\text{new}}}{\Delta} + \frac{1}{\sqrt{\Delta d}} \mathbf{z}_{\text{new}}^\top \mathbf{v}^*\right)\right) \\ &= \mathbb{P}\left(1 - \rho > \rho e^{-\frac{1}{2\Delta}} \cosh\left(\frac{c_{\text{new}}}{\Delta} + \frac{\zeta_{\text{new}}}{\sqrt{\Delta}}\right)\right) \\ &= \frac{1}{2} \left[\text{erf}\left(\frac{\Delta \left| \text{arccosh}\left(\frac{(1-\rho)}{\rho} e^{1/2\Delta}\right) \right| + 1}{\sqrt{2\Delta}}\right) + \text{erf}\left(\frac{\Delta \left| \text{arccosh}\left(\frac{(1-\rho)}{\rho} e^{1/2\Delta}\right) \right| - 1}{\sqrt{2\Delta}}\right) \right], \end{aligned} \quad (\text{C.6})$$

where $\mathbf{z}_{\text{new}} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_d)$, $\zeta_{\text{new}} \sim \mathcal{N}(0, 1)$, and $c_{\text{new}} = \pm 1$ with probability $\frac{1}{2}$. Finally, we obtain that the oracle error is

$$\begin{aligned} \varepsilon_{\text{gen}}^{BO} &= (1-\rho)\text{erfc}\left(\sqrt{\frac{\Delta}{2}} \left| \text{arccosh}\left(\frac{(1-\rho)}{\rho} e^{1/2\Delta}\right) \right|\right) \\ &\quad + \frac{\rho}{2} \left[\text{erf}\left(\frac{\Delta \left| \text{arccosh}\left(\frac{(1-\rho)}{\rho} e^{1/2\Delta}\right) \right| + 1}{\sqrt{2\Delta}}\right) + \text{erf}\left(\frac{\Delta \left| \text{arccosh}\left(\frac{(1-\rho)}{\rho} e^{1/2\Delta}\right) \right| - 1}{\sqrt{2\Delta}}\right) \right]. \end{aligned} \quad (\text{C.7})$$

References

- [1] Elisabeth Agoritsas, Giulio Biroli, Pierfrancesco Urbani, and Francesco Zamponi. Out-of-equilibrium dynamical mean-field equations for the perceptron model. *Journal of Physics A: Mathematical and Theoretical*, 51(8):085002, 2018.
- [2] Marc Mézard, Giorgio Parisi, and Miguel A. Virasoro. *Spin glass theory and beyond*. World Scientific, Singapore, 1987.
- [3] Jorge Kurchan. Supersymmetry, replica and dynamic treatments of disordered systems: a parallel presentation. *arXiv preprint cond-mat/0209399*, 2002.
- [4] Jorge Kurchan. Supersymmetry in spin glass dynamics. *Journal de Physique I*, 2(7):1333–1352, 1992.
- [5] Jean Zinn-Justin. *Quantum field theory and critical phenomena*. Clarendon Press, 1996.
- [6] H Eissfeller and M Opper. New method for studying the dynamics of disordered spin systems without finite-size effects. *Physical review letters*, 68(13):2094, 1992.
- [7] H Eissfeller and M Opper. Mean-field monte carlo approach to the sherrington-kirkpatrick model with asymmetric couplings. *Physical Review E*, 50(2):709, 1994.
- [8] Felix Roy, Giulio Biroli, Guy Bunin, and Chiara Cammarota. Numerical implementation of dynamical mean field theory for disordered systems: application to the lotka–volterra model of ecosystems. *Journal of Physics A: Mathematical and Theoretical*, 52(48):484001, 2019.
- [9] Alessandro Manacorda, Grégory Schehr, and Francesco Zamponi. Numerical solution of the dynamical mean field theory of infinite-dimensional equilibrium liquids. *The Journal of Chemical Physics*, 152(16):164506, 2020.
- [10] Antoine Georges, Gabriel Kotliar, Werner Krauth, and Marcelo J Rozenberg. Dynamical mean-field theory of strongly correlated fermion systems and the limit of infinite dimensions. *Reviews of Modern Physics*, 68(1):13, 1996.
- [11] Francesca Mignacco, Florent Krzakala, Yue M Lu, and Lenka Zdeborová. The role of regularization in classification of high-dimensional noisy gaussian mixture. *arXiv preprint arXiv:2002.11544*, 2020.