



HAL
open science

Stochastic Gradient Descent outperforms Gradient Descent in recovering a high-dimensional signal in a glassy energy landscape

Persia Jana Kamali, Pierfrancesco Urbani

► **To cite this version:**

Persia Jana Kamali, Pierfrancesco Urbani. Stochastic Gradient Descent outperforms Gradient Descent in recovering a high-dimensional signal in a glassy energy landscape. 2024. hal-04778393

HAL Id: hal-04778393

<https://hal.science/hal-04778393v1>

Preprint submitted on 12 Nov 2024

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.

Stochastic Gradient Descent outperforms Gradient Descent in recovering a high-dimensional signal in a glassy energy landscape

Persia Jana Kamali¹ and Pierfrancesco Urbani¹

¹*Université Paris-Saclay, CNRS, CEA, Institut de physique théorique, 91191, Gif-sur-Yvette, France*

Stochastic Gradient Descent (SGD) is an out-of-equilibrium algorithm used extensively to train artificial neural networks. However very little is known on to what extent SGD is crucial for to the success of this technology and, in particular, how much it is effective in optimizing high-dimensional non-convex cost functions as compared to other optimization algorithms such as Gradient Descent (GD). In this work we leverage dynamical mean field theory to benchmark its performances in the high-dimensional limit. To do that, we consider the problem of recovering a hidden high-dimensional non-linearly encrypted signal, a prototype high-dimensional non-convex hard optimization problem. We compare the performances of SGD to GD and we show that SGD largely outperforms GD for sufficiently small batch sizes. In particular, a power law fit of the relaxation time of these algorithms shows that the recovery threshold for SGD with small batch size is smaller than the corresponding one of GD.

INTRODUCTION

Stochastic Gradient Descent (SGD) is an optimization algorithm that is at the basis of many deep learning (DL) models trained on large datasets. It can be viewed as an approximated Gradient Descent (GD). In GD the gradient of the empirical loss function is used to update the parameters (or weights) of artificial neural networks. It can be usually computed as a sum of terms, each coming from a datapoint in the training dataset. SGD consists in approximating the gradient by summing only datapoints in a particular fraction of the dataset, the minibatch, and then changing the minibatch during time in order to explore the entire dataset (a full exploration of the dataset is called an epoch in training dynamics). While the main reason for using SGD is clearly the reduction of the computational complexity of the optimization algorithm - computing the full gradient of the loss is costly as soon as the dataset is huge - the effect of SGD on the performances of deep neural networks is largely unknown. Given its ubiquitous use in DL, understanding this point has become a central problem in this research field.

In generic situations, training DL models involves the optimization of a high-dimensional non-convex cost function and a common belief is that SGD introduces a noise that is beneficial for optimization as much as in simulated annealing thermal noise triggers activated process that may lead to barrier jumping and better optimization performances [1]. However the extent to which SGD is effective in annealing the empirical loss is largely unknown due to the fact that it is an intrinsically out-of-equilibrium driven dynamics and therefore its stationary measure is unknown. In fact it has been even shown that SGD can easily find bad local minima which perform rather badly [2].

In the last years many research directions have been explored to clarify the effectiveness of SGD. One of them [3–5] has suggested that both GD and SGD dynamics

perform an implicit regularization of the neural network parameters during training and this is beneficial for the performances of the networks. However, especially for the SGD case, these results are often limited to particular optimization settings and therefore this leaves open the possibility that the implicit regularization picture may not hold generically. Additionally, the version of SGD considered is typically the one of online (or one-pass) SGD [6–12]. In this case, the gradient of the loss is computed via just one datapoint which is never used again. It follows that in this setting there is no notion of training error or epochs and therefore online SGD is slightly far from practical situations. In another research line, it has been suggested that overparametrized models trained via SGD are found in minimizers of the loss which are flat and that flatness positively correlates with generalization performances [13, 14]. However this perspective seems to be challenged by more recent studies [15]. An exact treatment of SGD in these cases is still lacking.

A parallel research line has focused on understanding the performances of gradient based algorithms in high-dimensional inference problems [16–22]. While the loss landscape of these problems is different from the one of overparametrized neural networks where a manifold of minimizers at zero loss is found [23], they represent prototypical high-dimensional non-convex hard optimization problems and, for this reason, they are ideal to benchmark optimization algorithms. Indeed in these cases, the ground states of the loss function are sharp minima having perfect generalization properties but are hard to find given that they are very rare (typically one or two configurations) surrounded by exponentially many spurious minima with bad generalization properties. For this class of problems a lot of work has been devoted to study the performances of GD and sampling algorithms such as the Langevin algorithm [16, 17, 19] and the performances of online SGD [21, 24, 25]. Recently in [20] one of these prototypical problems, phase retrieval [26, 27], has been proposed to assess the effectiveness of the SGD algorithm.

However, while numerical simulations show that SGD (and in particular a variant of SGD called persistent-SGD) seems to perform better than GD, it is unclear how these results extend in the high dimensional limit. The main difficulty is that the theoretical analysis of SGD is performed using Dynamical Mean Field Theory (DMFT) [28] which however has severe limitations and cannot be used to extrapolate the behavior of the algorithms at long times.

Our work joins this research line by considering a different model which has the same landscape structure of previously considered models yet it allows an exact theoretical description of the SGD dynamics in the high-dimensional limit. Indeed the DMFT equations are rather different from [20] and this allows to track SGD dynamics at long times and small batch size and to establish quantitatively its superiority with respect to GD.

THE MODEL AND THE ALGORITHMS

We consider an N -dimensional signal $\underline{w}^* = \{w_1^*, \dots, w_N^*\} \in \mathbb{R}^N$ uniformly distributed on the sphere $|\underline{w}^*|^2 = N$ and a set of $M = \alpha N$ non-linear measurements

$$y_\mu = \frac{1}{N} \sum_{i < j} J_{ij}^\mu w_i^* w_j^* \quad \mu = 1, \dots, M \quad (1)$$

where the entries of the symmetric matrices J^μ are independent and distributed according to a Gaussian law with zero mean and variance given by

$$\overline{J_{ij}^\mu J_{lm}^\nu} = \delta_{\mu\nu} \delta_{il} \delta_{jm} \quad i < j, l < m \quad (2)$$

The control parameter α is called the sample complexity.

A natural inference task is to reconstruct \underline{w}^* knowing the dataset of measurements and measurement matrices, $\{y_\mu, J^\mu\}_{\mu=1, \dots, M}$. We expect that this problem becomes easier when α is large and it is more difficult at small sample complexity. A way to solve it is to define a loss function given by

$$H[\underline{x}] = \frac{1}{2} \sum_{\mu=1}^M \left(y_\mu - \frac{1}{N} \sum_{i < j} J_{ij}^\mu w_i w_j \right)^2 \equiv \sum_{\mu=1}^M v_\mu(\underline{w}). \quad (3)$$

The ground state of H has been studied extensively in [29] when restricted to the case in which $|\underline{w}^*|^2$ is fixed. For $\alpha > 1$, H has two trivial ground states given by \underline{w}^* and $-\underline{w}^*$ [29]. However we will show that the model is glassy and has many spurious minima, poorly correlated with the signal, which do trap the dynamics for sufficiently small α and which coexist with the simple ground state structure. We are interested in reconstructing \underline{w}^* by minimizing H via a class of gradient based algorithms

defined by the following dynamical rule

$$\underline{w}(t+1) = \underline{w}(t) - \frac{\eta}{b} \sum_{\mu=1}^M \sigma_\mu(t) \frac{\partial v_\mu}{\partial \underline{w}}. \quad (4)$$

In this case, time is discretized into small time steps of length η also called the learning rate. The variables $\sigma_\mu(t) = \{0, 1\}$ are selection variables that control which datapoint enters in the computation of the gradient. We take $\sigma_\mu(t)$ to be random: $\sigma_\mu(t) = 1$ with probability $b \in (0, 1]$ and $\sigma_\mu(t) = 0$ with probability $1 - b$ [30]. The parameter b controls the batch size, namely the average fraction of datapoints entering the computation of the gradient. We assume that the variables $\sigma_\mu(t)$ are extracted independently at each time step (even if a time-dependent correlation can be also added in principle [30]) and are totally uncorrelated in the datapoint index μ . The particular case $b = 1$ corresponds to GD. The initial condition for the dynamics in Eq. (4) is assumed to be random with $|\underline{w}(0)|^2 = N$. However (4) does not constrain the norm of \underline{w} .

We are interested in measuring the performances of this class of algorithms as a function of the batch size. This can be done by defining the mean square displacement (MSD) between the signal and the current state of the system $\underline{w}(t)$

$$\Delta(t) = \frac{1}{N} \sum_{i=1}^N |w_i(t) - w_i^*|^2. \quad (5)$$

Therefore we would like to study the behavior of $\Delta(t)$ as a function of the batch size b . Since the SGD noise for $b < 1$ is an out-of-equilibrium state-dependent noise, the stationary probability distribution of (4) is unknown and therefore the only way to extract the performances of the SGD algorithm is to track it directly in the large N limit. This can be done by using dynamical mean field theory (DMFT) [28, 31].

DYNAMICAL MEAN FIELD THEORY

The MSD can be rewritten as

$$\Delta(t) = 1 - 2m(t) + C(t, t) \quad (6)$$

where the magnetization $m(t)$ and the correlation function $C(t, t')$ are defined as

$$m(t) = \underline{w}(t) \cdot \underline{w}^* / N \quad C(t, t') = \underline{w}(t) \cdot \underline{w}(t') / N. \quad (7)$$

Therefore a computation of these dynamical order parameters in the large N limit allows to track the corresponding dynamics of the MSD. These quantities concentrate at large N on their typical values. The equations describing their exact dynamical behavior can be

obtained via a DMFT analysis. For the model described in this work we follow the derivation of the DMFT equations outlined in a related model [32–34] and we generalize it to include (i) the fact that the norm of the vector $\underline{w}(t)$ is unbounded, (ii) the presence of a hidden signal

in the loss function and (iii) the effect of the selection variables $\sigma_\mu(t)$. Here we present the resulting DMFT equations whose derivation is reported in the Appendix of this manuscript.

These equations are written in terms of an additional order parameter, the response function, given by

$$R(t, t') = \lim_{\underline{h} \rightarrow 0} \frac{1}{N} \sum_{i=1}^N \frac{\delta w_i(t)}{\delta h_i(t')} \quad (8)$$

and the field $\underline{h}(t)$ is introduced in the dynamical equations (4) via an infinitesimal tilt of the loss function $H \rightarrow H - \underline{h} \cdot \underline{w}$. Causality of the dynamical equations implies that $R(t, t') = 0$ if $t \leq t'$. The resulting DMFT equations are

$$\begin{aligned} m(t+1) &= m(t) - \eta^2 \alpha \left(\sum_{s=0}^t (\Lambda_R(t, s) C(t, s) + \Lambda_C(t, s) R(t, s)) m(s) - m(t) \sum_{s=0}^t \Lambda_R(t, s) \right) \\ C(t+1, t') &= C(t, t') + \eta \Omega_1(t, t') \quad \forall t' \leq t \\ R(t+1, t') &= \delta_{t, t'} - \eta^2 \alpha \sum_{s=t'+1}^t (\Lambda_R(t, s) C(t, s) + \Lambda_C(t, s) R(t, s)) R(s, t') \\ C(t+1, t+1) &= C(t, t) + 2\eta \Omega_1(t, t) + \eta^2 \Omega_2(t) \\ \Omega_1(t, t') &= \alpha \eta \left[m(t) m(t') \sum_{s=0}^t \Lambda_R(t, s) - \sum_{s=0}^{t'} \Lambda_C(t, s) C(t, s) R(t', s) - \sum_{s=0}^t (\Lambda_R(t, s) C(t, s) + \Lambda_C(t, s) R(t, s)) C(t', s) \right] \\ \Omega_2(t) &= \alpha^2 \eta^2 \sum_{s, s'=0}^t (\Lambda_R(t, s) C(t, s) + \Lambda_C(t, s) R(t, s)) C(s, s') (\Lambda_R(t, s) C(t, s') + \Lambda_C(t, s) R(t, s')) \\ &\quad - 2\alpha^2 \eta^2 m(t) \left(\sum_{s=0}^t \Lambda_R(t, s) \right) \left(\sum_{s=0}^t (\Lambda_R(t, s) C(t, s) + \Lambda_C(t, s) R(t, s)) m(s) \right) \\ &\quad + 2\alpha^2 \eta^2 \sum_{s=0}^t \sum_{s'=0}^s (\Lambda_R(t, s) C(t, s) + \Lambda_C(t, s) R(t, s)) \Lambda_C(t, s') C(t, s') R(s, s') \\ &\quad - \alpha \Lambda_C(t, t) C(t, t) + \left(\alpha \eta \sum_{s=0}^t \Lambda_R(t, s) \right)^2 \end{aligned} \quad (9)$$

and must be integrated starting with an initial condition given by

$$m(0) = m_0 \quad C(0, 0) = C_0. \quad (10)$$

This means that we are extracting the initial configuration $\underline{w}(0)$ randomly with a uniform probability distribution such that it has a projection m_0 on \underline{w}^* and its squared norm is C_0 . The quantities $\Lambda_C(t, t')$ and $\Lambda_R(t, t')$ play the role of a self-energy in the Dyson equation which describes the evolution of m , C and R . Their expressions are given by

$$\begin{aligned} \Lambda_C(t, t') &= \frac{1}{2} \langle C_A^{(r)}(t, t') \sigma(t) + C_A^{(l)}(t', t) \sigma(t') \rangle_\sigma \\ \Lambda_R(t, t') &= \frac{1}{2} \langle R_A^{(r)}(t, t') \sigma(t) + R_A^{(l)}(t', t) \sigma(t') \rangle_\sigma \end{aligned} \quad (11)$$

where the random variable $\sigma(t)$ is a selection variable whose dynamical and statistical properties are the same as a typical selection variable $\sigma_\mu(t)$. Correspondingly, the brackets denote the average with respect to its probability distribution. Furthermore $C_A^{(l, r)}$ and $R_A^{(l, r)}$ are matrices which depend on $\sigma(t)$ and whose precise expressions are given in the Appendix.

RESULTS

The DMFT equations can be integrated numerically rather efficiently. At variance with other cases [20, 30] which are typically more difficult, here the only bottleneck is in the computation of the averages appearing in

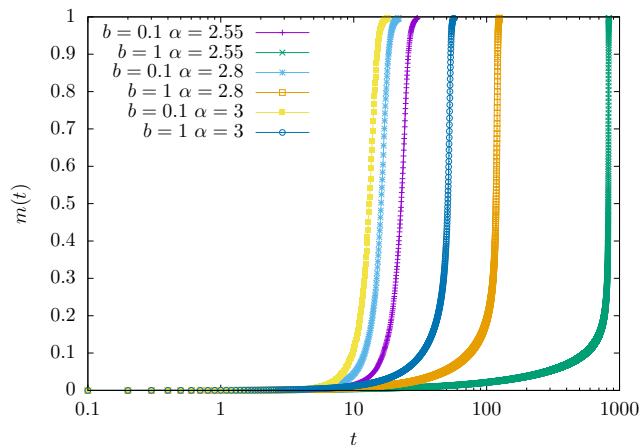


FIG. 1. The magnetization $m(t)$ as a function of the time t for different values of α and b . SGD is systematically faster than GD in reaching the signal. The curves are obtained by integrating the DMFT equations with $\mathcal{N} = 10^3$ samples.

Eq. (11). This can be done by extracting \mathcal{N} trajectories of the variable $\sigma(t)$ and approximating these averages by empirical averages. In the following we present the results of such integration. We fix the learning rate $\eta = 0.1$ and we take $C_0 = 1$. The loss function is invariant under spin-flip and therefore recovering the signal requires a spontaneous breaking of this symmetry. In a finite dimensional system, the small symmetry breaking field is given by the initial condition of the dynamics that always has a finite albeit of order $N^{-1/2}$ projection on the signal. In order to avoid the inconvenient of analyzing the equation in the small m_0 limit we fix its value to be small but finite. This implies that we are considering a hot start for the dynamics. Therefore from now on we fix[35] $m_0 = 10^{-4}$.

In order to describe the result of the numerical integration of the DMFT equations we first start by a qualitative analysis. In Fig.1 we report the behavior of $m(t)$ for GD ($b = 1$) and SGD ($b = 0.1, 0.2$) and decreasing values of the sample complexity α . We observe that (i) decreasing α the time it takes to achieve the signal increases for all algorithms and (ii) the SGD algorithm is qualitatively faster to get to it. In order to quantify this speedup we define the relaxation time to the signal $\tau(b)$ as the time it takes for $\Delta(t)$ to go below a threshold value which we fix to be 0.15 (and we have checked that shortly after this point all algorithms achieve the signal). In Fig.2 we plot the behavior of $\tau(b)$ for $b = 0.1, 0.2, 1$ from which we deduce that the SGD noise provides a clear speedup of the dynamics. To describe these findings even more quantitatively we assume that the relaxation time diverges as a power law at a recovery transition point $\alpha^*(b)$ and we fit the data with

$$\tau(b) \simeq \tau_0(b) |\alpha - \alpha^*(b)|^{-z(b)}. \quad (12)$$

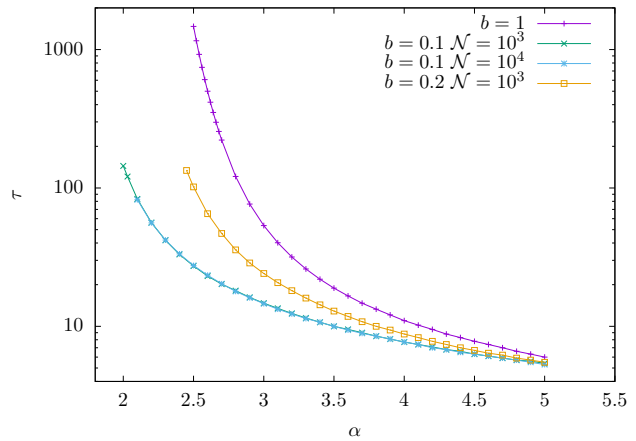


FIG. 2. The relaxation time to the signal as a function of α for GD and SGD ($b = 0.1, 0.2$). For $b = 0.1$ we show the result of the integration of the DMFT equation for $\sigma = 10^3$ and $\sigma = 10^4$. This shows that for what concerns the relaxation time, with 10^3 samples we are close to the asymptotic limit.

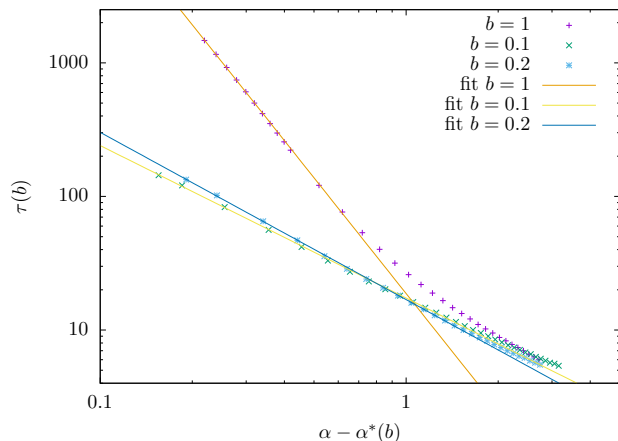


FIG. 3. Power law fit of the relaxation time with the law in Eq. (12). We find that $\alpha^*(b) \simeq \{1.84, 2.26, 2.28\}$ for $b = 0.1, 0.2, 1$.

In Fig.3 we plot the result of this fit. The fitting value of the recovery threshold $\alpha^*(b)$ is smaller when $b < 1$. In particular for we can compare the cases $b = 0.1$ and $b = 1$. The recovery threshold for GD is found at $\alpha^*(1) = 2.28$ and for this value of α SGD with $b = 0.1$ has a finite relaxation time to the signal[36]. Therefore the recovery threshold of SGD for small batch size is strictly smaller than the one of GD. This is the main result of our work.

DISCUSSION

To summarize, we have considered a prototypical high-dimensional non-convex hard optimization problem to benchmark the performances of SGD. Thanks to an exact

treatment of this algorithm via DMFT, we have shown that SGD at small batch sizes outperforms GD in the sense that the recovery threshold of the former is strictly smaller than the corresponding one of the latter.

Close to the recovery threshold, the dynamics is a two step process: first the system lands on a stationary state not correlated (or poorly correlated) with the signal and then, on longer timescales, this stationary state becomes unstable and the system reaches the ground state of the loss function. This is similar to what happens in the GD case and for the thermal Langevin algorithm [16–18]. However while in these cases the intermediate stationary regime is described by aging [37], in the present case the SGD noise provides an out-of-equilibrium driven dynamics which leads the system away from aging and into a stationary, time-translational invariant (TTI) regime [38]. Therefore constructing a theory of the recovery threshold for SGD would require a detailed understanding of this intermediate dynamical regime which is left for future work.

Acknowledgments – P. J. Kamali acknowledges the Institute of Theoretical Physics of the University Paris-Saclay where this work has been done. P. Urbani wants to thank Antonio Sclocchi for discussions.

-
- [1] S. Kirkpatrick, C. D. Gelatt Jr, and M. P. Vecchi, Optimization by simulated annealing, *science* **220**, 671 (1983).
 - [2] S. Liu, D. Papailiopoulos, and D. Achlioptas, Bad global minima exist and sgd can reach them, *Advances in Neural Information Processing Systems* **33**, 8543 (2020).
 - [3] B. Neyshabur, R. Tomioka, and N. Srebro, In search of the real inductive bias: On the role of implicit regularization in deep learning, *arXiv preprint arXiv:1412.6614* (2014).
 - [4] D. Soudry, E. Hoffer, M. S. Nacson, S. Gunasekar, and N. Srebro, The implicit bias of gradient descent on separable data, *The Journal of Machine Learning Research* **19**, 2822 (2018).
 - [5] S. Gunasekar, J. Lee, D. Soudry, and N. Srebro, Characterizing implicit bias in terms of optimization geometry, in *International Conference on Machine Learning* (PMLR, 2018) pp. 1832–1841.
 - [6] D. Saad and S. A. Solla, On-line learning in soft committee machines, *Physical Review E* **52**, 4225 (1995).
 - [7] D. Saad and S. A. Solla, Exact solution for on-line learning in multilayer neural networks, *Physical Review Letters* **74**, 4337 (1995).
 - [8] D. Saad and M. Rattray, Globally optimal parameters for on-line learning in multilayer neural networks, *Physical review letters* **79**, 2578 (1997).
 - [9] D. Saad, *On-line learning in neural networks* (Cambridge University Press, 1999).
 - [10] A. C. Coolen and D. Saad, Dynamics of learning with restricted training sets, *Physical Review E* **62**, 5444 (2000).
 - [11] A. C. Coolen, D. Saad, and Y.-S. Xiong, On-line learning from restricted training sets in multilayer neural networks, *Europphysics Letters* **51**, 691 (2000).
 - [12] R. Veiga, L. Stephan, B. Loureiro, F. Krzakala, and L. Zdeborová, Phase diagram of stochastic gradient descent in high-dimensional two-layer neural networks, *Advances in Neural Information Processing Systems* **35**, 23244 (2022).
 - [13] Y. Yang, L. Hodgkinson, R. Theisen, J. Zou, J. E. Gonzalez, K. Ramchandran, and M. W. Mahoney, Taxonomizing local versus global structure in neural network loss landscapes, *Advances in Neural Information Processing Systems* **34**, 18722 (2021).
 - [14] N. Yang, C. Tang, and Y. Tu, Stochastic gradient descent introduces an effective landscape-dependent regularization favoring flat solutions, *Physical Review Letters* **130**, 237101 (2023).
 - [15] M. Andriushchenko, F. Croce, M. Müller, M. Hein, and N. Flammarion, A modern look at the relationship between sharpness and generalization, *arXiv preprint arXiv:2302.07011* (2023).
 - [16] S. S. Mannelli, F. Krzakala, P. Urbani, and L. Zdeborová, Passed & spurious: Descent algorithms and local minima in spiked matrix-tensor models, in *international conference on machine learning* (PMLR, 2019) pp. 4333–4342.
 - [17] S. S. Mannelli, G. Biroli, C. Cammarota, F. Krzakala, P. Urbani, and L. Zdeborová, Marvels and pitfalls of the langevin algorithm in noisy high-dimensional inference, *Physical Review X* **10**, 011057 (2020).
 - [18] S. Sarao Mannelli, G. Biroli, C. Cammarota, F. Krzakala, P. Urbani, and L. Zdeborová, Complex dynamics in simple neural networks: Understanding gradient flow in phase retrieval, *Advances in Neural Information Processing Systems* **33**, 3265 (2020).
 - [19] S. Sarao Mannelli and P. Urbani, Analytical study of momentum-based acceleration methods in paradigmatic high-dimensional non-convex problems, *Advances in Neural Information Processing Systems* **34**, 187 (2021).
 - [20] F. Mignacco, P. Urbani, and L. Zdeborová, Stochasticity helps to navigate rough landscapes: comparing gradient-descent-based algorithms in the phase retrieval problem, *Machine Learning: Science and Technology* **2**, 035029 (2021).
 - [21] G. Ben Arous, R. Gheissari, and A. Jagannath, High-dimensional limit theorems for sgd: Effective dynamics and critical scaling, *Advances in Neural Information Processing Systems* **35**, 25349 (2022).
 - [22] M. C. Angelini and F. Ricci-Tersenghi, Limits and performances of algorithms based on simulated annealing in solving sparse hard inference problems, *Physical Review X* **13**, 021011 (2023).
 - [23] F. Draxler, K. Veschgini, M. Salmhofer, and F. Hamprecht, Essentially no barriers in neural network energy landscape, in *International conference on machine learning* (PMLR, 2018) pp. 1309–1318.
 - [24] G. Ben Arous, R. Gheissari, and A. Jagannath, Algorithmic thresholds for tensor pca, (2020).
 - [25] G. B. Arous, R. Gheissari, and A. Jagannath, On-line stochastic gradient descent on non-convex losses from high-dimensional inference, *The Journal of Machine Learning Research* **22**, 4788 (2021).
 - [26] Y. Chi, Y. M. Lu, and Y. Chen, Nonconvex optimization meets low-rank matrix factorization: An overview, *IEEE Transactions on Signal Processing* **67**, 5239 (2019).
 - [27] J. Dong, L. Valzania, A. Maillard, T.-a. Pham, S. Gigan, and M. Unser, Phase retrieval: From computational imaging to machine learning: A tutorial, *IEEE Signal*

- Processing Magazine **40**, 45 (2023).
- [28] L. F. Cugliandolo, Recent applications of dynamical mean-field methods, arXiv preprint arXiv:2305.01229 (2023).
- [29] Y. V. Fyodorov, A spin glass model for reconstructing nonlinearly encrypted signals corrupted by noise, Journal of Statistical Physics **175**, 789 (2019).
- [30] F. Mignacco, F. Krzakala, P. Urbani, and L. Zdeborová, Dynamical mean-field theory for stochastic gradient descent in gaussian mixture classification, Advances in Neural Information Processing Systems **33**, 9540 (2020).
- [31] M. Mezard, G. Parisi, and M. A. Virasoro, *Spin glass theory and beyond* (World Scientific, Singapore, 1987).
- [32] P. J. Kamali and P. Urbani, Dynamical mean field theory for models of confluent tissues and beyond, arXiv preprint arXiv:2306.06420 (2023).
- [33] P. Urbani, A continuous constraint satisfaction problem for the rigidity transition in confluent tissues, Journal of Physics A: Mathematical and Theoretical **56**, 115003 (2023).
- [34] A. Montanari and E. Subag, Solving overparametrized systems of random equations: I. model and algorithms for approximate solutions, arXiv preprint arXiv:2306.13326 (2023).
- [35] Note that the central limit theorem implies that $m_0 = 10^{-4}$ is equivalent to having a system size of order $N \simeq 10^8$ which is impossible to simulate.
- [36] This is the best indication that SGD for $b = 0.1$ has a smaller recovery threshold than GD. Indeed fitting the relaxation time in the SGD case is hard given that we have data only on just a decade in time.
- [37] L. F. Cugliandolo, Dynamics of glassy systems, arXiv preprint cond-mat/0210312 (2002).
- [38] F. Mignacco and P. Urbani, The effective noise of stochastic gradient descent, Journal of Statistical Mechanics: Theory and Experiment **2022**, 083405 (2022).
- [39] J. Zinn-Justin, *Quantum field theory and critical phenomena*, Vol. 171 (Oxford university press, 2021).
- [40] T. Castellani and A. Cavagna, Spin-glass theory for pedestrians, Journal of Statistical Mechanics: Theory and Experiment **2005**, P05012 (2005).

Derivation of the DMFT equations

In this section we derive the dynamical mean field theory equations (DMFT) that describe the behavior of SGD (and GD when $b = 1$) in the infinite dimensional limit $N \rightarrow \infty$. The derivation closely follows the one of [32] and we will highlight the main differences.

We start with a trivial identity which follows from the causality of the dynamics:

$$1 = \mathcal{Z}_{dyn} = \left\langle \int \mathcal{D}\underline{w}(t) \prod_{i=1}^N \delta \left(-\dot{w}_i(t) - \frac{\partial H}{\partial w_i(t)} \right) \right\rangle \quad (13)$$

and the brackets denote an average over the selection variables σ_μ and the realization of the random matrices J^μ . Introducing a Fourier representation for the Dirac delta we get

$$\mathcal{Z}_{dyn} = \left\langle \int \mathcal{D}\underline{w}(t) \mathcal{D}\hat{\underline{w}}(t) \exp \left(i \sum_t \hat{\underline{w}}(t) \cdot \left(-(\underline{w}(t+1) + \underline{w}(t)) - \frac{\eta}{b} \sum_\mu \sigma_\mu(t) \frac{\partial v_\mu}{\partial \underline{w}(t)} \right) \right) \right\rangle \quad (14)$$

We introduce now Grassmann variables [39] so that we can rewrite:

$$\mathcal{Z}_{dyn} = \left\langle \int \mathcal{D}\underline{w}(a) \exp \left(-\frac{1}{2} \int da db \underline{w}(a) \mathcal{K}(a, b) \underline{w}(b) + M \mathcal{A}_{loc} \right) \right\rangle \quad (15)$$

and the superfield $\underline{w}(a)$ is defined on coordinate $a = (t_a, \theta_a)$:

$$\underline{w}(a) = \underline{w}(t_a) + i\theta_a \hat{\underline{w}}(t_a). \quad (16)$$

The kernel $\mathcal{K}(a, b)$ is implicitly defined such that:

$$-\frac{1}{2} \int da db \underline{w}(a) \mathcal{K}(a, b) \underline{w}(b) = - \sum_t i \hat{\underline{w}}(t) \cdot (\underline{w}(t+1) - \underline{w}(t)) \quad (17)$$

Following the same steps as in [32] we can show that the term $M \mathcal{A}_{loc}$ can be obtained as a function of the following dynamical order parameters

$$\begin{aligned} m(a) &= \frac{1}{N} \underline{w}^* \cdot \underline{w}(a) \\ Q(a, b) &= \frac{1}{N} \underline{w}(a) \cdot \underline{w}(b) - m(a)m(b). \end{aligned} \quad (18)$$

In particular we have

$$\begin{aligned}\mathcal{Z}_{dyn} &= \int \mathcal{D}Q \mathcal{D}m \exp N \mathcal{A}_{dyn}[Q, m] \\ \mathcal{A}_{dyn} &= -\frac{1}{2} \int da db \mathcal{K}(a, b) [Q(a, b) + m(a)m(b)] + \frac{1}{2} \ln \det(Q) + \alpha \mathcal{A}_{loc}\end{aligned}\quad (19)$$

and \mathcal{A}_{loc} is given by

$$\begin{aligned}\mathcal{A}_{loc} &= \frac{1}{2} \ln \left\langle \int \mathcal{D}h(a) \mathcal{D}\hat{h}(a) e^{S_{loc}^R} \right\rangle + \frac{1}{2} \ln \left\langle \int \mathcal{D}h(a) \mathcal{D}\hat{h}(a) e^{S_{loc}^L} \right\rangle \\ S_{loc}^{(R,L)} &= i \int da h(a) \hat{h}(a) - \int da \mathbf{s}(a) \frac{\hat{h}(a)^2}{2} - \frac{1}{2} \int dadb \hat{h}(a) \hat{U}^{(R,L)}(a, b) \hat{h}(b)\end{aligned}\quad (20)$$

and we have defined $U^{(R,L)}$ as

$$\begin{aligned}\hat{U}^{(L)}(a, b) &= \mathbf{s}(a) \left[\frac{1}{2} - \frac{m(a)^2}{2} - \frac{m(b)^2}{2} + \frac{(Q(a, b) + m(a)m(b))^2}{2} \right] \\ \hat{U}^{(R)}(a, b) &= \left[\frac{1}{2} - \frac{m(a)^2}{2} - \frac{m(b)^2}{2} + \frac{(Q(a, b) + m(a)m(b))^2}{2} \right] \mathbf{s}(b)\end{aligned}\quad (21)$$

and $\mathbf{s}(a)$ is an effective selection variable which has only a scalar component:

$$\mathbf{s}(a) = \frac{\sigma(t)}{b} . \quad (22)$$

The brackets in Eq. (20) denote the average with respect to \mathbf{s} . Performing the Gaussian integral over \hat{h} and h we find

$$\mathcal{A}_{loc} \propto \frac{1}{2} \ln \langle \det(I + \hat{U}^{(R)})^{-\frac{1}{2}} \rangle + \frac{1}{2} \ln \langle \det(I + \hat{U}^{(L)})^{-\frac{1}{2}} \rangle$$

and we have neglected irrelevant constant factors. The integral defining \mathcal{Z}_{dyn} can be evaluated through a saddle point. Taking the functional derivatives with respect to $Q(a, b)$ and $m(a)$ and then performing the shift $Q(a, b) \rightarrow Q(a, b) - m(a)m(b)$ we get the following equations:

$$0 = - \int dc \mathcal{K}(a, c) Q(c, b) + \delta(a, b) - \alpha \int dc \Lambda(a, c) Q(a, c) Q(c, b) - m(a)m(b) \int dc \Lambda(a, c) \quad (23)$$

$$0 = - \int db \mathcal{K}(a, b) m(b) - \alpha \left[\int dc \Lambda(a, c) Q(a, c) m(c) - m(a) \int dc \Lambda(a, c) \right] \quad (24)$$

and $\Lambda(a, b)$ is given by

$$\Lambda(a, b) = \frac{1}{2} \left\langle \left(1 + U^{(L)}\right)^{-1}(a, b) \mathbf{s}(b) + \left(1 + U^{(R)}\right)^{-1}(a, b) \mathbf{s}(a) \right\rangle \quad (25)$$

where

$$\begin{aligned}U^{(L)}(a, b) &= \mathbf{s}(a) \left[\frac{1}{2} - \frac{m(a)^2}{2} - \frac{m(b)^2}{2} + \frac{Q(a, b)^2}{2} \right] \\ U^{(R)}(a, b) &= \left[\frac{1}{2} - \frac{m(a)^2}{2} - \frac{m(b)^2}{2} + \frac{Q(a, b)^2}{2} \right] \mathbf{s}(b) .\end{aligned}\quad (26)$$

At this point we can unfold the Grassmann structure of the equations. In order to do that one can show by inspection that the Grassmann structure of the fields at their saddle point value is given by

$$\begin{aligned}Q(a, b) &= C(t_a, t_b) + \theta_a R(t_b, t_a) + \theta_b R(t_a, t_b) \\ \Lambda(a, b) &= \Lambda_C(t_a, t_b) + \theta_a \Lambda_R(t_b, t_a) + \theta_b \Lambda_R(t_a, t_b) \\ m(a) &= m(t_a)\end{aligned}\quad (27)$$

Substituting these expressions inside Eqs.(24) we get Eqs. (9) of the main text. We note that the Eqs.(24) do not contain the equation to propagate the diagonal part of the matrix $C(t, t)$. This can be derived either by taking the corresponding derivatives of \mathcal{A}_{dyn} with respect to these quantities or by following an alternative route which we explain now given that it is easier. The main idea is to write a self consistent process for an effective variable w . Using the same strategy as in [40], we get to

$$w(t+1) = w(t) - \alpha\eta^2 \sum_{s=0}^t [\Lambda_R(t, s)C(t, s) + \Lambda_C(t, s)R(t, s)] w(s) + \alpha\eta^2 m(t) \sum_0^t \Lambda_R(t, s) + \xi(t) \quad (28)$$

where the noise ξ has zero average and two point correlation function given by

$$\langle \xi(t)\xi(s) \rangle = -\alpha\Lambda_C(t, s)C(t, s). \quad (29)$$

Therefore we get that

$$C(t+1, t+1) = \left\langle \left(w(t) - \alpha\eta^2 \sum_{s=0}^t [\Lambda_R(t, s)C(t, s) + \Lambda_C(t, s)R(t, s)] w(s) + \alpha\eta^2 m(t) \sum_0^t \Lambda_R(t, s) + \xi(t) \right)^2 \right\rangle_{\xi} \quad (30)$$

which gives the corresponding equation reported in Eq. (9). In order to conclude our derivation we need to give precise expressions for the kernels $\Lambda_C(t, t')$ and $\Lambda_R(t, t')$. These can be derived unfolding the Grassmann structure in their definition in Eq. (25). In particular we obtain Eqs. (11) of the main text where $C_A^{(l)}$, $C_A^{(r)}$, $R_A^{(l)}$ and $R_A^{(r)}$ solve the following equations

$$\begin{aligned} \sum_{s=0}^t \mathcal{M}^{(R)}(t', s) \begin{pmatrix} C_A^{(r)}(t, s) \\ R_A^{(r)}(t, s) \end{pmatrix} &= \begin{pmatrix} 0 \\ \delta_{t, t'} / \eta \end{pmatrix} \quad t' \leq t \\ \mathcal{M}^{(R)}(t', s) &= \begin{pmatrix} \delta_{s, t'} + \frac{\eta}{b} C(s, t') R(t', s) \sigma(t') & \frac{\eta}{2b} [1 - m(s)^2 - m(t')^2 + C(s, t')^2] \sigma(t') \\ 0 & \delta_{s, t'} + \frac{\eta}{b} C(s, t') R(s, t') \sigma(t') \end{pmatrix}. \end{aligned} \quad (31)$$

$$\begin{aligned} \sum_{s=0}^t \mathcal{M}^{(L)}(t', s) \begin{pmatrix} C_A^{(l)}(t, s) \\ R_A^{(l)}(t, s) \end{pmatrix} &= \begin{pmatrix} 0 \\ \delta_{t, t'} / \eta \end{pmatrix} \quad t' \leq t \\ \mathcal{M}^{(L)}(t', s) &= \begin{pmatrix} \delta_{s, t'} + \frac{\eta}{b} C(s, t') R(t', s) \sigma(s) & \frac{\eta}{2b} [1 - m(s)^2 - m(t')^2 + C(s, t')^2] \sigma(s) \\ 0 & \delta_{s, t'} + \frac{\eta}{b} C(s, t') R(s, t') \sigma(s) \end{pmatrix}. \end{aligned} \quad (32)$$

This concludes the derivation of the DMFT equations. We underline that they have a causal structure and therefore they can be easily integrated numerically. We underline that the numerical integration for the GD case, $b = 1$, is computationally simpler given that there is no need to compute the averages of Eqs. (11) of the main text. This is the reason why we systematically achieve longer times in the GD case.