



OPEN

Stochastic Gradient Descent-like relaxation is equivalent to Metropolis dynamics in discrete optimization and inference problems

Maria Chiara Angelini^{1,2,✉}, Angelo Giorgio Cavaliere³, Raffaele Marino⁴ & Federico Ricci-Tersenghi^{1,2,5}

Is Stochastic Gradient Descent (SGD) substantially different from Metropolis Monte Carlo dynamics? This is a fundamental question at the time of understanding the most used training algorithm in the field of Machine Learning, but it received no answer until now. Here we show that in discrete optimization and inference problems, the dynamics of an SGD-like algorithm resemble very closely that of Metropolis Monte Carlo with a properly chosen temperature, which depends on the mini-batch size. This quantitative matching holds both at equilibrium and in the out-of-equilibrium regime, despite the two algorithms having fundamental differences (e.g. SGD does not satisfy detailed balance). Such equivalence allows us to use results about performances and limits of Monte Carlo algorithms to optimize the mini-batch size in the SGD-like algorithm and make it efficient at recovering the signal in hard inference problems.

Algorithms have become ubiquitous in modern life¹, with numerous applications in fields ranging from finance and business to healthcare and transportation. They are essential tools for making sense of vast amounts of data and making predictions about complex systems². Despite their prevalence, however, sometimes we do not fully understand how they work or the underlying mathematical principles that govern their behavior.

In this manuscript, we focus on two algorithms for solving discrete optimization and inference problems: (1) a Monte Carlo algorithm (MC) with the Metropolis updating rule³ and (2) a Stochastic Gradient Descent (SGD) like algorithm. The former has a theory beyond it, while the latter is introduced in this paper for the first time, as a version for discrete problems of the well-known and widely used SGD algorithm⁴. We aim to show the existence of a strong equivalence between the dynamics of the two algorithms and highlight its significance for both theory and practice. Our goal is to provide a comprehensive introduction to this equivalence, offering insights into its potential applications.

The MC algorithms are the state of the art for sampling complex functions. However, they are quite demanding in terms of computing resources, and so they are often replaced by faster, but less controlled algorithms. The example of SGD for the minimization of the loss function is very significant: while SGD works very efficiently thanks to several tricks⁵ its performances are in general not well known⁶. Making a strong connection between the two kinds of algorithms would help a lot in building a solid theory for SGD-like algorithms.

The effectiveness of MC methods is witnessed by their success in hard discrete combinatorial problems^{7–12}. The theory beyond MC algorithms is very strong thanks to the theory of Markov chains and many concepts borrowed from the principles of statistical mechanics (e.g. ergodicity, relaxation to equilibrium, phase transitions)¹³. Such a strong connection makes it natural to use in MC simulations the concept of temperature even if one is just interested in computing the optimal configurations (the solutions to a discrete combinatorial problem or the minimizers of a loss function). The temperature parameter T in MC algorithms controls the degree of

¹Dipartimento di Fisica, Sapienza Università di Roma, P.le Aldo Moro 5, 00185 Rome, Italy. ²Istituto Nazionale di Fisica Nucleare, Sezione di Roma I, P.le A. Moro 5, 00185 Rome, Italy. ³Cybermedia Center, Osaka University, Toyonaka, Osaka 560-0043, Japan. ⁴Dipartimento di Fisica e Astronomia, Università degli studi di Firenze, Via Giovanni Sansone 1, 50019 Sesto Fiorentino, FI, Italy. ⁵Institute of Nanotechnology (NANOTEC)-CNR, Rome unit, P.le A. Moro 5, 00185 Rome, Italy. ✉email: maria.chiara.angelini@roma1.infn.it

randomness in the exploration of the energy/cost/loss function. If one wants to use the MC method for computing the global minimum, one can either run the algorithm at $T = 0$ or change slowly the temperature from an initial value to $T = 0$: this is the so-called Simulated Annealing algorithm¹⁴. A key property that allows for a solid theory of MC is the so-called detailed balance condition that ensures the algorithm admits a limiting distribution at large times^{15,16}.

SGD^{17–19} is a popular optimization algorithm used in the development of state-of-the-art machine learning²⁰ and deep learning models^{21,22}, which have shown tremendous success in numerous fields, becoming indispensable tools for many advanced applications^{23–31}. It is an extension of the gradient descent algorithm³² that uses a subset of the training data to compute the gradient of the objective function at each iteration. The use of random subsets makes the algorithm stochastic, and it allows for faster convergence on very large datasets. The algorithm works by iteratively updating the model parameters using a mini-batch of training data at each iteration. In SGD there is no temperature and the parameter that controls the degree of randomness in the exploration of the energy landscape is given by the size of the mini-batch used^{16,33,34}. The introduction of the mini-batch was forced by practical reasons, because the evaluation of the current state of neural weights on the full training set was often computationally too hard, and in many cases practically impossible, due to huge training sets. Quite surprisingly, the introduction of the mini-batch led to better optimization and generalization but this success is lacking a full physical interpretation. No analyses have been performed to understand if the SGD satisfies the detail balance condition and to identify the asymptotic sampling distribution.

Some attempts had been made to model SGD (using the central limit theorem) as an approximated GD plus the introduction of Gaussian Langevin noise, with variance depending on the parameters (size of the mini-batch, learning rate, ...) ^{35–40}. However, some other works underlined how the noise could not always be Gaussian, the dynamics undergoing possibly Levy-flights^{41,42}.

Recently, using dynamical mean-field theory⁴³, a technique borrowed from statistical mechanics, the whole dynamical trajectory of SGD has been tracked analytically^{44–47}. In particular, in ref.⁴⁶, an effective temperature for SGD is obtained from the fluctuation-dissipation relation⁴⁸ (ideas in this direction had been given also in^{39,49}). This effective temperature has then been related to the parameters of the model (mini-batch size and learning rate). However, it is known that, for complex systems, the effective temperature could be different from the temperature of a thermal bath in contact with the system: the effective temperature could be affected also by the topology of the energy landscape and is time-scale-dependent⁵⁰, while in ref.⁴⁶ the effective temperature is just shown at large enough waiting times.

To directly identify a connection between the size of the chosen mini-batch in an SGD and the temperature of an equivalent thermal bath, we conduct a detailed analysis of the behavior of two algorithms in the domain of discrete optimization and inference: a Metropolis MC algorithm at temperature T , and an SGD-like algorithm on the q -coloring problem, both in its random and planted versions^{51–53}. The q -coloring problem is a well-known hard problem in the field of discrete optimization. Since we are dealing with a discrete optimization problem, the SGD-like algorithm is seen as a Metropolis dynamics at zero temperature that uses a subset of the training data to inject randomness into the exploration of the energy landscape. Our goal is to establish a relationship between the thermal fluctuations that govern the dynamics of MC and the fluctuations resulting from the size of the mini-batch, which enables the SGD-like algorithm to explore the energy landscape. To accomplish this, we conduct both numerical and analytical analyses. We find that the dynamics of the SGD-like algorithm with a given mini-batch size do not satisfy the detail balance, however, it follows closely the MC dynamics at a given temperature.

The paper is organized as follows. In section 2 we describe the q -coloring model and the MC and SGD-like algorithms. In section 3 we present the numerical analysis of both algorithms. In section 4 we report our analytical computations showing that SGD-like does not satisfy detailed balance, but seems to satisfy a sort of averaged version of this condition. In section 5 we discuss and summarize our results.

Model and algorithms

In the following, we focus on the coloring problem, both in its random and planted versions: the random coloring problem is an example of an optimization problem, while the planted coloring problem is a simple example of a hard inference problem. In the random q -coloring problem, given a random graph of N vertices and mean degree c , we have to assign one among the q available colors to each vertex in a way to avoid monochromatic edges, i.e. adjacent vertices with the same color. In the planted version of the q -coloring problem, one first creates the planted solution by dividing the N nodes into q groups and assigning to all the nodes of a given group the same color (two groups cannot share the same color): we will call this the *planted solution* $\{s^*\}$ ⁵³. Then a random graph of mean degree c compatible with the planted solution is created by adding $M = cN/2$ edges, which are randomly chosen among all those not connecting vertices of the same color. Constructing the graph in this way, the probability distribution of the degree of a node (that is the number of edges that are attached to a given node), is a Poissonian distribution with mean c . We will simply call c as the “connectivity” of the graph in the following. Finally one asks to recover the planted solution just from the knowledge of the planted graph. This setting is particularly suitable because we can compare different algorithms, already knowing which should be the best output. For small connectivities c , there exists a large number of coloring compatible with the graph, and it is impossible to identify the planted one, while for very large c the planted configuration is the only configuration compatible with the graph and can potentially be identified in a time that grows polynomially with the size of the problem. The best-known algorithms for the recovery of the planted coloring solution are message-passing algorithms that succeed when the connectivity is higher than $c_l = (q - 1)^{253}$, finding the planted assignment in a time that is almost linear in the size of the graph. We will call s_i the color of the i -th node of the graph, and s_j

can take values from 1 to q . Given a configuration \mathbf{s} of colors for all the nodes, we can associate to it an energy, or a cost-function, that simply counts the number of monochromatic edges

$$H(\mathbf{s}) = \sum_{(i,j) \in E} \delta_{s_i, s_j}, \quad (1)$$

where E is the edge set of the graph, and $|E| = M$. In the planted setting, we know there exists at least one configuration with zero energy (\mathbf{s}^*), and thus we can compare different algorithms whose aim is to minimize the energy, knowing the optimal solution.

The simplest algorithm is probably the one that tries to minimize directly the cost function defined in Eq. (1): a sort of gradient-descent (GD) algorithm, but for discrete variables. It works as follows: We start with a random initial configuration in which we assign randomly to each node one among the q possible colors. Then at each step, we choose a node uniform at random (u.a.r.) and we propose a new color for it chosen u.a.r. among the remaining $q - 1$ colors. We accept the new color only if the energy in Eq. (1) decreases or stays constant. The attentive reader could say that the MC algorithm at zero temperature is not exactly a discretized version of GD but is much more similar to the so-called Coordinate Gradient Descent⁵⁴. However coordinate descent algorithms are demonstrated to have competitive or under some hypothesis even better performances than full-gradient descent algorithms⁵⁵ and thus, for simplicity of notation, we will just name our algorithm as GD-like. The GD-like algorithm finds a solution to the planted coloring problem only for connectivities larger than $c_{GD}(N)$ and the threshold seems to scale logarithmically with N : $c_{GD}(N) \simeq A \log(N)$, with $A = O(1)$, as shown in the Supplementary material. The threshold for the GD-like algorithm thus diverges in the large N limit. GD-like algorithm is thus highly inefficient (remind that the currently best algorithms find solutions as soon as $c > c_l = (q - 1)^2$ ^{53,56}, and the threshold does not scale with N). As shown in different contexts, stochasticity could help to find better solutions to the problem. In the following, we present two ways to introduce stochasticity: the first one through the use of a finite “temperature” and the second one through the use of a “mini-batch”.

In a statistical mechanics approach, we introduce a temperature $T \equiv 1/\beta$ and a Gibbs-Boltzmann-like probability measure on the configurations

$$P_{GB}(\mathbf{s}) = \frac{e^{-\beta H(\mathbf{s})}}{\sum_{\mathbf{s}'} e^{-\beta H(\mathbf{s}')}} = \frac{e^{-\beta \sum_{(i,j) \in E} \delta_{s_i, s_j}}}{\sum_{\mathbf{s}'} e^{-\beta \sum_{(i,j) \in E} \delta_{s'_i, s'_j}}}. \quad (2)$$

In the $T = 0$ limit, $P_{GB}(\mathbf{s})$ becomes the uniform distribution over the configurations with zero energy if they exist. They always exist in the planted coloring problem, while they disappear for $c > c_{COL/UNCOL}$ in the random coloring problem. $c_{COL/UNCOL} = 13.669(2)$ when $q = 5$ ⁵². The distribution $P_{GB}(\mathbf{s})$ with $T > 0$ relaxes the hard constraints into soft ones: configurations with monochromatic edges are admitted, but with a probability that is more suppressed the lower the temperature. We can then construct a standard Metropolis MC algorithm at temperature T . We start from a random configuration of colors and a single MC Sweep corresponds to the attempt to update the colors of N variables following the Metropolis rule: a node i is picked u.a.r. between the N nodes, we propose to change its color s_i into s'_i , choosing it u.a.r. between the remaining $q - 1$ colors. We accept the proposed move with probability 1 if the number of monochromatic edges does not increase, or with probability $e^{-\beta \Delta E}$ if the number of monochromatic edges increases by ΔE after the update. When $T = 0$ one only accepts moves that do not increase the number of monochromatic edges and the MC algorithm reduces to the GD-like algorithm. The performances of MC algorithms in inference problems like the planted coloring one have been derived in Ref. ⁵⁶. Using $T > 0$ allows us to obtain much better performances than working at $T = 0$. At a fixed connectivity c , for temperatures small enough $T < T_{\text{glassy}}(c)$, the MC is trapped by spurious glassy states, almost orthogonal to the planted one. T_{glassy} is often called the *dynamical transition temperature*, indicated as T_d in the statistical mechanics’ language of disordered systems. On the opposite, for too high temperatures $T > T_{\text{PM}}(c)$, the MC is attracted by a paramagnetic state that implies a flat probability measure over the configurations, without any information at all about the planted state. For intermediate temperatures $T_{\text{glassy}}(c) < T < T_{\text{PM}}(c)$, the paramagnetic state is locally unstable, the glassy states are still not formed, and thus the MC is attracted by the planted state: this is the only temperature range where recovery is possible. We know how to compute the values of $T_{\text{glassy}}(c)$ and $T_{\text{PM}}(c)$ in the thermodynamic limit⁵⁶. For $q = 5$ -coloring, when $c < 18$ we have $T_{\text{glassy}}(c) > T_{\text{PM}}(c)$, and recovery is not possible by the MC algorithm. Thus $c_{\text{MC}} = 18$ is the recovery threshold for MC algorithms in this case.

The second way to introduce stochasticity in the GD-like algorithm is the use of a sort of “mini-batch” and thus we call the resulting algorithm the *Stochastic-GD-like* (SGD-like) algorithm. This algorithm works as follows: We start with a random initial configuration. Then at each step, we choose a node i u.a.r. and we propose a new color for it u.a.r. among the remaining $q - 1$ colors. Then we extract a fraction $B \cdot M$ of edges u.a.r. among the M total ones. We accept the new color for node i only if the energy, computed as in Eq. (1) but restricting the sum only over the “mini-batch” of $B \cdot M$ chosen edges, does not increase. The $B \cdot M$ edges over which the energy is evaluated are extracted u.a.r. at each move. One could also introduce a persistency parameter that models the time that a given edge remains in the extracted mini-batch, following Ref. ⁴⁴, but we have numerically seen that the performances of the algorithm deteriorate. A single SGD-like Sweep corresponds to the attempt to update the color of N variables. B is a parameter in the range $(0, 1]$: if $B = 1$ the SGD-like algorithm reduces to the GD-like algorithm, that is the MC algorithm with $T = 0$. The introduction of a finite number of observations (edges in this case) to compute the energy cost is reminiscent of the practice, widely used in the learning process of neural networks, of computing the gradient of the loss function only on a mini-batch of data. However, the SGD-like algorithm is introduced here for the first time to perform optimization and inference tasks. It is in principle quite different from the MC algorithm: we show in the subsequent sections that the SGD-like algorithm does not

satisfy detailed balance, while the MC algorithm does. However, in the next section we will show that, despite the differences, the two algorithms share common behaviors.

Numerical results

In this section, we analyze the performances of the SGD-like algorithm in solving the planted $q = 5$ -coloring problem and compare it to the MC algorithm. We consider graphs with average connectivity $c = 19$, but the behavior is qualitatively the same for each value of $c > c_{MC} = 18$. For $c < 18$ the two algorithms cannot recover the planted solution but try to optimize the cost function reaching random configurations uncorrelated with the planted ones: we show in the Supplementary Material that also in this region the two algorithms behave similarly. For $c = 19$, the MC algorithm finds a paramagnetic solution for $T > T_{PM} \simeq 0.491$, gets trapped in spurious glassy states for $T < T_{glassy} \simeq 0.454$ and finds a low energy state highly correlated with the planted solution for $0.454 < T < 0.491$ ⁵⁶. In Fig. 1 we show that the SGD-like algorithm has a qualitatively similar behavior changing the fraction B of considered edges: for B too small ($B = 0.86$) or too large ($B = 0.95$), the SGD-like algorithm gets stuck in some high energy configurations, while it manages to find a low energy state only for an intermediate B value ($B = 0.9$). This is also confirmed by the data in the inset, showing the overlap Q between the configuration s reached by the SGD-like algorithm and the planted solution s^* defined as

$$Q = \frac{\max_{\pi \in S_q} \sum_i \delta_{s_i^*, \pi(s_i)} / N - 1/q}{1 - 1/q}, \quad (3)$$

where S_q is the group of permutations of q elements. The overlap takes value $Q = 0$ for a random guess, while $Q = 1$ if there is a perfect recovery of the planted solution. The inset of Fig. 1 shows that the overlap is high only for the intermediate value $B = 0.9$.

Having shown that the qualitative behaviors of SGD-like and MC algorithms are the same, in the following we make a quantitative comparison in the three regions.

Comparison between MC and SGD-like algorithms in the recovery region

Let us start from the recovery region at intermediate values of B (or T). In Fig. 1 we have shown the behavior for just one realization of the planted graph: for $B = 0.9$ the energy first relaxes towards a plateau of high energy and then abruptly jumps towards a low-energy configuration strongly correlated with the planted one. We call *nucleation time* t_{nucl} the time corresponding to the sudden decrease of the energy: operatively, we can define it as the time at which the overlap reaches values $Q > 0.8$ ($t_{nucl} \simeq 8 \cdot 10^3$ in Fig. 1). For graphs of finite sizes, there are sample-to-sample fluctuations in t_{nucl} . Thus, to perform a quantitative comparison between MC and SGD-like algorithms, we consider the average over many different samples.

At the time of comparing MC and SGD-like algorithms, we need a criterion to match the temperature T and the mini-batch size B . We use the value of the energy at the plateau, declaring a matching pair (T, B) when the MC and SGD-like algorithms relax to a plateau with the same energy. Once we found a matching pair (T, B) , we noticed that both the relaxation towards the plateau and the average nucleation time surprisingly coincide in the two algorithms. In Fig. 2 we plot the average energy, averaged over many samples for three different sizes $N = 10^3, 10^4, 10^5$, as a function of the running time t for both MC and SGD-like algorithms. When temperature T and mini-batch size B are matched as explained above the energy relaxation and the mean nucleation time are very similar.

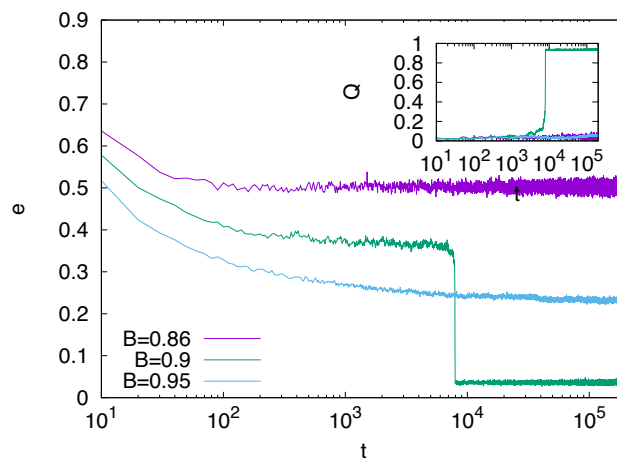


Figure 1. Intensive energy reached by the SGD-like algorithm with three different values of B as a function of time, for a single system of size $N = 10^4$ and mean connectivity $c = 19$. For $B = 0.86$ the algorithm ends in a paramagnetic state, for $B = 0.95$ it ends in some spurious glassy states, while for $B = 0.9$ the algorithm manages to find the low-energy planted state. Inset: overlap as defined in Eq. (3) between the planted state and the configurations visited at time t by the SGD-like algorithm with the same values of B as in the main Figure.

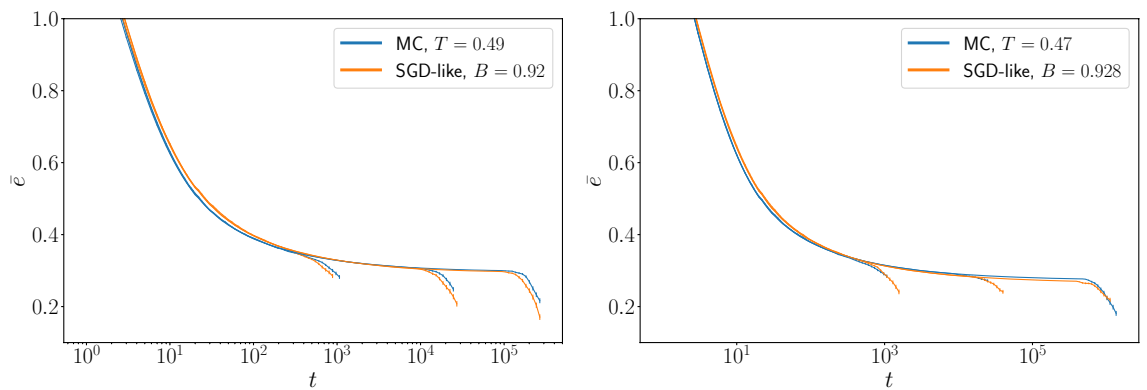


Figure 2. Average energy \bar{e} mediated over 280 samples for three different sizes $N = 10^3, 10^4, 10^5$ (from left to right), for MC and SGD-like algorithms. **Left:** In MC at a temperature $T = 0.49$ and SGD-like algorithm with a mini-batch parameter $B = 0.92$ the energy relaxes in a very similar way and the average nucleation times coincide. **Right:** as in the left panel, but with MC run at $T = 0.47$ and an SGD-like algorithm with $B = 0.928$.

We then concentrate our attention on the distribution of the nucleation times for the two algorithms varying the problem size N . In Fig. 3 we show the averaged nucleation times \bar{t}_{nucl} and the corresponding standard deviation $\sigma_{t_{\text{nucl}}}$ for the same matching pairs (T, B) used in Fig. 2. We find that the values of both \bar{t}_{nucl} and $\sigma_{t_{\text{nucl}}}$ are very similar in the two algorithms for any value of N . So the two algorithms not only are quantitatively similar in their energy relaxation (which becomes N -independent in the large N limit, see Fig. 2), but they also have the same finite size effects. Thus making the (T, B) matching even more impressive.

It is worth stressing that the growth of both \bar{t}_{nucl} and $\sigma_{t_{\text{nucl}}}$ is roughly linear in N (the dashed lines in Fig. 3 are linear functions of N). The negative curvature of the data shown in Fig. 3 (in a double logarithmic scale) suggests the growth laws may be sub-linear in the large N limit, but it is not the scope of the present work to estimate them precisely. We remind the reader that the time is measured in sweeps and each sweep is made of N single variable updates. The actual time is thus almost quadratic in the problem size, in agreement with the results in Ref. ⁵⁶. In the Supplementary Material, we show the total running time in seconds and also that the entire probability distribution of t_{nucl} turns out to be the same for the two algorithms.

Comparison between MC and SGD-like algorithms in the paramagnetic and the glassy regions

Having shown the equivalence of MC and SGD-like algorithms in the recovery region, we consider now the region where both converge to a paramagnetic state. For the MC algorithm, this region is defined by the condition $T > T_{\text{PM}}$, where the exact value for T_{PM} can be computed in the large N limit by standard techniques from statistical mechanics⁵³ and reads $T_{\text{PM}}^{-1} = -\log \left[\frac{c-(q-1)^2}{q-1+c} \right]$ ($T_{\text{PM}} \simeq 0.491$ for $q = 5$ and $c = 19$). We find that, analogously, the SGD-like algorithm reaches a paramagnetic state for B smaller than a certain threshold. In the left panel of Fig. 4 we show the behavior of the energy obtained by an MC algorithm for $T > T_{\text{PM}}$. As already explained, relaxation gets stuck in a plateau of quite high energy. We then identify the corresponding value of B for the SGD-like algorithm matching the plateau energy. Similarly to the recovery region, once we match the values of T and B from the plateau energy condition, the entire energy relaxation coincides in the two algorithms. At this point, to be sure that the phase reached by the SGD-like algorithm at small B is paramagnetic as for the MC algorithm, we look at two-times correlation functions and in particular we compute $c(t, t_w) = \sum_{i=1}^N s_i(t_w) s_i(t + t_w) / N$, which is a standard observable in the statistical mechanics analysis of

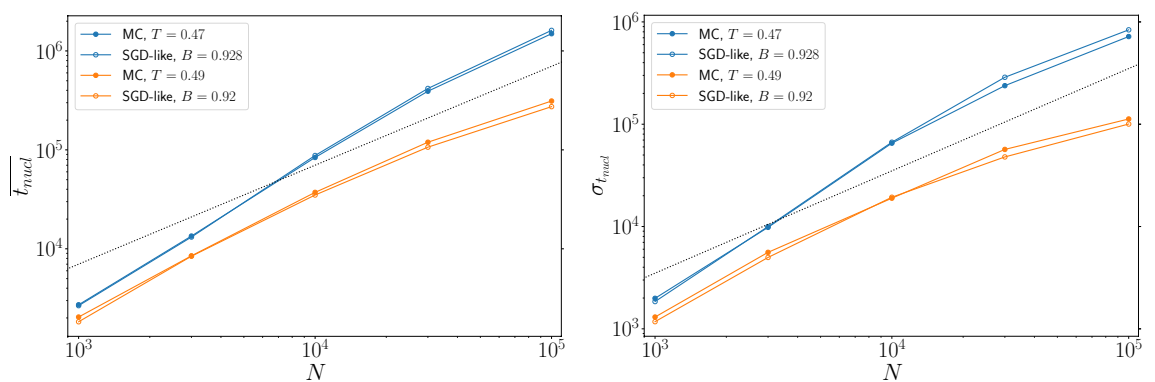


Figure 3. Averaged nucleation time \bar{t}_{nucl} (left) and the corresponding standard deviation $\sigma_{t_{\text{nucl}}}$ (right) for the same matching (T, B) pairs used in Fig. 2. Both \bar{t}_{nucl} and $\sigma_{t_{\text{nucl}}}$ are very similar for the two algorithms and grow almost linearly with the problem size.

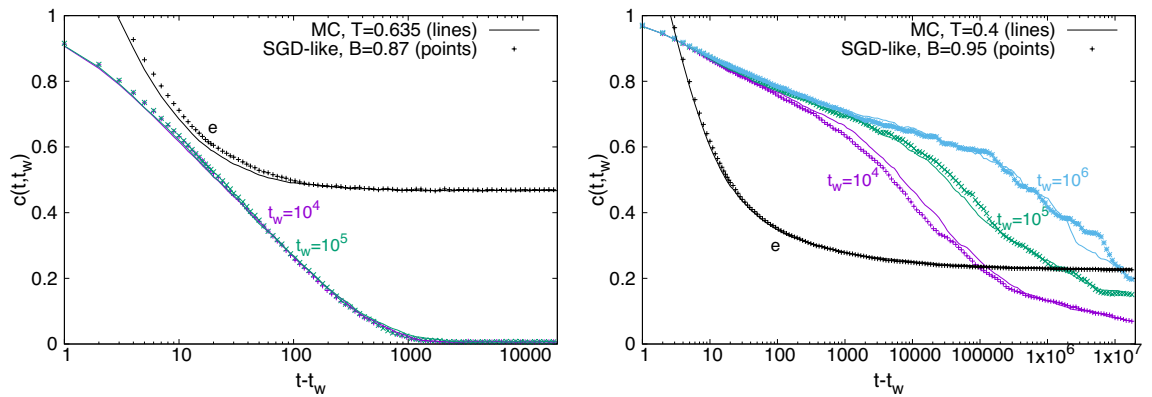


Figure 4. Data for a single system of size $N = 10^5$ and $c = 19$. **Left:** Intensive energy e and two-time correlations $c(t, t_w)$ at different values of t_w for MC algorithm at $T = 0.635$ and a SGD-like algorithm with $B = 0.87$. Values of T and B have been matched equating the plateau energy, and in turn, make the entire relaxation process of the two algorithms very similar. **Right:** As in the left panel, for the MC algorithm at $T = 0.4$ (lines) and the SGD-like algorithm with $B = 0.95$ (points). Both algorithms end in a low-temperature, aging regime, signaled by $c(t, t_w)$ depending on t_w .

disordered systems. In the left panel of Fig. 4 we show the behavior of $c(t, t_w)$ for MC and SGD-like algorithms, for different values of t_w : again the data match perfectly for the two algorithms. The observation that $c(t, t_w)$ rapidly decays to zero and is t_w -independent provides a clear indication that a paramagnetic state has been reached.

The last region left to analyze is the one at low temperatures for the MC algorithm that corresponds to high values of B for the SGD-like algorithm. We already know that the two algorithms perfectly coincide when $T = 0$ for MC and $B = 1$ for SGD-like, since they reduce to the GD-like algorithm. From Ref. ⁵⁶ we also know that, for $q = 5$ and $c = 19$, the MC algorithm reaches glassy states almost uncorrelated with the planted signal when $T < T_{\text{glassy}} \simeq 0.454$ ⁵². Analogously, we find that the SGD-like algorithm reaches a glassy state for B higher than a certain threshold. In the right panel of Fig. 4 we show the behavior of the energy obtained by the MC algorithm at $T < T_{\text{glassy}}$. Since the energy gets stuck in a plateau we use the value of the energy at the plateau to match T and B values (as we did above for the other regions). It is clear that, once the right (T, B) matching is made, the dynamics of the two algorithms are very similar. To be sure that the state reached by the SGD-like algorithm at high B is a glassy one as for the MC algorithm, also in this case we look at two-times correlation function $c(t, t_w)$. In the right panel of Fig. 4 we show the behavior of $c(t, t_w)$ for the MC and SGD-like algorithms for different values of t_w : again the data match for the two algorithms. Moreover, now $c(t, t_w)$ does not rapidly decay to zero and strongly depends on t_w , at variance with what happens in the paramagnetic region: this is the so-called *aging* behavior, typical of glassy states⁵⁷. Please note that the aging behavior of $c(t, t_w)$ is a clear indication that the two algorithms are in the off-equilibrium regime, and their very similar behavior in this regime is highly non-trivial and unexpected.

Does the SGD-like algorithm satisfy the detailed balance condition?

Given the strong resemblance in the dynamical behavior of the two algorithms, it is natural to ask if this similarity comes from some fundamental property that they (may) have in common. The Metropolis MC algorithm is built to satisfy detailed balance, which is a sufficient (and strong) condition to guarantee convergence of the Markov process to the target equilibrium distribution when ergodicity is not spontaneously broken. For a Gibbs-Boltzmann measure at inverse temperature $\beta = 1/T$ the detailed balance condition reads

$$\frac{p(A \rightarrow B)}{p(B \rightarrow A)} = e^{-\beta(E_B - E_A)}, \tag{4}$$

where $p(A \rightarrow B)$ coincides with the acceptance probability for the move $A \rightarrow B$ in the case of a symmetric proposal function (as in the single-spin-flip dynamics we adopt).

The question that we address here is whether Eq. (4) is still valid for the SGD-like algorithm. Since the SGD-like algorithm does not involve temperature, being based on greedy dynamics, we look for a relation of the kind

$$\frac{p(A \rightarrow B)}{p(B \rightarrow A)} = e^{-f(B)(E_B - E_A)} \tag{5}$$

for some function $f(B)$ of the size of the mini-batch. If Eq. (5) is satisfied, then we expect the SGD-like algorithm to converge, when it is possible, to the Gibbs-Boltzmann equilibrium distribution at inverse temperature $f(B)$. The relation $\beta = f(B)$ would then provide an analytic mapping between the two algorithms. As we are going to show in the following, we have found that detailed balance is in this case not exactly satisfied, since the ratio of the transition probabilities between two states is not a function of their energy difference only. We are nonetheless able to quantify numerically the deviations from detailed balance. Interestingly, an arithmetic average over the

approximate $T - B$ relations one can obtain for the different choices of the energy levels is sufficient to obtain a good approximation to the experimental $T(B)$ curve we obtain from numerical simulations, suggesting that the SGD-like algorithm is, in practice, performing close to detailed balance.

To write the transition probabilities between configurations A and B in the mini-batch case, it is convenient to introduce the following nomenclature: we call s the number of interactions that are satisfied in configuration A but become violated in configuration B . Conversely, u is the number of interactions that are violated in configuration A but become satisfied in configuration B . The total number of interactions is $M = cN/2 \gg 1$. We are restricting ourselves to single-spin-flip dynamics. This implies that, since each spin is involved in roughly c interactions (being c the average connectivity), the total number of interactions that change their nature (satisfied/unsatisfied) when going from A to B is bounded by $s + u \leq O(c) = O(1)$ for $M, N \rightarrow \infty$. Also, from their definition, we have that $s - u = E_B - E_A = O(1)$. We now study the transition probability for the direct process $p(A \rightarrow B)$. First, we extract a mini-batch containing $M \cdot B$ interactions among the M total ones. There are 3 kinds of interactions in the system: s interactions of “type- s ”, u interactions of “type- u ”, and finally $(M - s - u)$ interactions that do not change their satisfied/unsatisfied nature upon going from A to B , and thus do not contribute to the energy difference between the two configurations. The probability that a uniformly extracted mini-batch of size $M \cdot B$ selects N_s interactions of type- s and N_u interactions of type- u follows a multivariate hypergeometric distribution

$$P_M(N_s, N_u, B, s, u) = \frac{\binom{s}{N_s} \binom{u}{N_u} \binom{M - s - u}{M \cdot B - N_s - N_u}}{\binom{M}{M \cdot B}}. \tag{6}$$

Since $N_s \leq s$ and $N_u \leq u$ by construction, we have that these quantities are $O(1)$ in the large M limit. We can then evaluate the $\lim_{M \rightarrow \infty} P_M$ by expanding $\binom{M - s - u}{M \cdot B - N_s - N_u}$ and $\binom{M}{M \cdot B}$. Neglecting $1/M$ correction in the exponent, we get

$$\lim_{M \rightarrow \infty} P_M(N_s, N_u, B, s, u) = P(N_s, N_u, B, s, u) = \binom{s}{N_s} \binom{u}{N_u} (1 - B)^{s+u-N_s-N_u} B^{N_s+N_u}. \tag{7}$$

Since the SGD-like dynamics is greedy (in the energy function \tilde{E} calculated over the mini-batch), the Metropolis probability for accepting a move is simply 1 if $E_B - E_A \leq 0$, and 0 otherwise. In the same way we wrote $s - u = E_B - E_A$ before, it also holds $N_s - N_u = \tilde{E}_B - \tilde{E}_A$. This means that in all the cases for which $N_s \leq N_u$ the move is accepted and we have

$$p(A \rightarrow B) = \sum_{N_u=0}^u \sum_{N_s=0}^{\min[N_u, s]} P(N_s, N_u, B, s, u). \tag{8}$$

For the reverse process $p(B \rightarrow A)$, we consider the same configurations A and B . We also keep the same definitions of s and u . The difference with the previous case is that now we want to go from B to A , and this only happens when $\tilde{E}_B - \tilde{E}_A \geq 0$, that is when $N_s - N_u \geq 0$. Putting together the results

$$\frac{p(A \rightarrow B)}{p(B \rightarrow A)} = \frac{\sum_{N_u=0}^u \sum_{N_s=0}^{\min[N_u, s]} P(N_s, N_u, B, s, u)}{\sum_{N_s=0}^s \sum_{N_u=0}^{\min[N_s, u]} P(N_s, N_u, B, s, u)} \equiv g(B, s, u). \tag{9}$$

We now ask if the quantity $G(B, s, u) \equiv -\frac{\log g(B, s, u)}{(s-u)}$ is independent of the choice of the starting and ending configurations, that is from s and u . To this end, in the left panel of Fig. 5 we plot $G(B, s, u)$ as a function of B , for all the possible values of s and u . We consider for definiteness the case $c = 19$. Even though the expression for G is independent of the value of c , the connectivity affects the admissible values of s and u . The left panel of Fig. 5 shows a clear dependence of G on s and u , which implies $G(B, s, u) \neq f(B)$, that is the detailed balance is not satisfied. In the right panel of Fig. 5 we show the inverse of the arithmetic average over all the possible choices of s and u , $(\overline{G(B, s, u)})^{-1}$, as a function of B ; we believe this is a good proxy for the $T(B)$ relation. We also show in the same panel the maximum and minimum values of $(G(B, s, u))^{-1}$. Despite the large variability of this effective temperature, the averaged value $(\overline{G(B, s, u)})^{-1}$ stays very close to the points in the (B, T) plane that were obtained from our numerical experiments through the matching condition. This indicates that the SGD-like algorithm, although not respecting exactly the detailed balance, is in some sense effectively very close to satisfying it. We want to stress again that the detailed balance is a sufficient but not necessary condition for the existence of an equilibrium measure. It could be possible that the SGD-like algorithm only satisfies the more general balance equation, and thus admits the same equilibrium measure as the MC algorithm. However checking for full balance condition is quite involving both numerically and analytically. We also highlight that, even if we could manage to prove that both algorithms share the same equilibrium measure, this will not fully explain the impressive similarity of the two algorithms numerically found in the short-time, out-of-equilibrium regime, or in the long-time aging regime inside the glassy phase. The identification of the deep reason for the matching in the out-of-equilibrium regions certainly will need further research in the future.

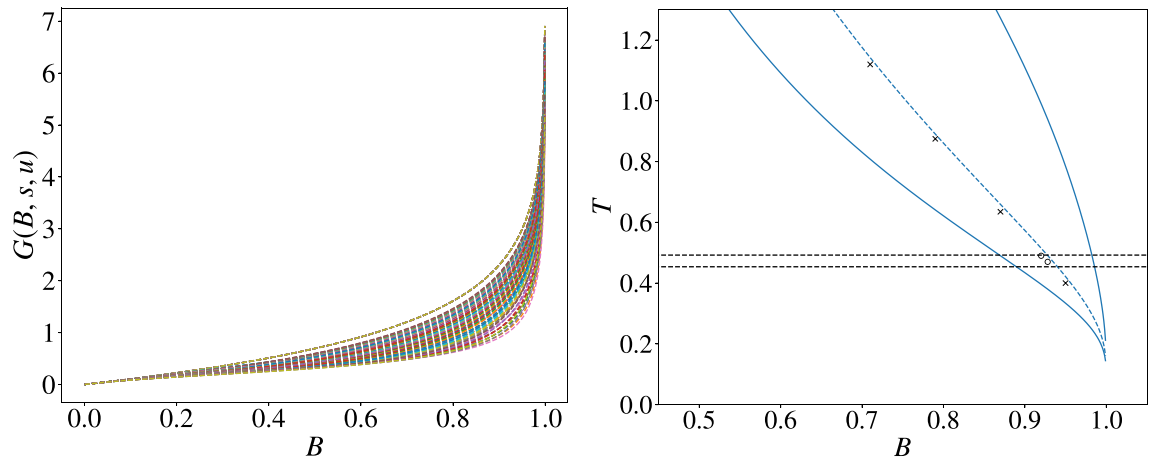


Figure 5. **Left:** $G(B, s, u)$ as a function of B , for all possible values of s and u for $c = 19$: the fact that it depends on s and u implies the failure of detailed balance (see text for more details) **Right:** The blue dashed line corresponds to $(G(B, s, u))^{-1}$, the blue full lines correspond to the maximum and minimum values of $(G(B, s, u))^{-1}$ over the possible choices of s and u . Crosses and circles correspond to (B, T) pairs obtained in previous sections through the matching condition on the plateau energy, which in turn make the entire dynamics of the two algorithms very similar. Horizontal dashed lines correspond to T_{PM} and T_{glassy} , the only recovery region for the MC algorithm being $T_{\text{glassy}} < T < T_{\text{PM}}$.

Discussion and conclusions

In this work, we have performed for the first time a quantitative comparison between an algorithm very similar in spirit to the well-known Stochastic Gradient Descent and a Monte Carlo algorithm based on the Metropolis updating rule. We have considered discrete optimization and inference problems. The model we have studied is always the same—the planted 5-coloring problem on random graphs of mean degree $c = 19$ —but depending on the parameters of the algorithms the dynamics can be very different: in the retrieval region, the planted signal is recovered with high probability and thus the algorithms perform an inference task, while outside this region the signal is undetectable⁵⁶ and thus the algorithms are performing an optimization task over random problems.

To perform a quantitative analysis between the two algorithms we have proposed a condition to map the temperature T of the MC algorithm to the mini-batch size B in the SGD-like algorithm. This matching condition is very simple and just requires to have the same energy in the plateau reached by both algorithms. In some sense, both the temperature and the mini-batch size are conjugated parameters to the energy (this is obvious for the MC algorithm, while it is an interesting observation for the SGD-like algorithm).

Once the (T, B) mapping has been established, we have performed extensive numerical simulations for the two algorithms, both in the retrieval phase and outside it, at equilibrium and in the out-of-equilibrium regime. Surprisingly, we find in all the above regimes a striking similarity in any observable we have measured: the energy relaxation, the nucleation time, and the auto-correlation function. The conclusion is highly unexpected: the evolution of the SGD-like algorithm is very close to that of the MC algorithm.

To justify theoretically the above findings, we have checked whether the SGD-like algorithm would satisfy the detailed balance condition which is at the heart of the Metropolis MC update. We have found that the SGD-like algorithm does not satisfy exactly the detailed balance, but it seems to satisfy it on average. This looks like a promising explanation (completely new to the best of our knowledge) that could support the unexpected similarity between the two algorithms.

The similarity between the MC and SGD-like algorithms found in this work opens a lot of possible future research paths. While the analysis of SGD-like algorithms cannot be made exactly and often relies on several hypotheses and approximations, the Metropolis MC dynamics has been studied in great detail and many predictions about it are available (because of the detailed balance condition). In particular, we have used the knowledge of the equilibrium measure at which the MC should tend at infinite times to make predictions on the type of states—glassy, paramagnetic or planted—that would be reached by MC as well as by SGD. We believe that the surprising findings presented in this work can stimulate the application of techniques borrowed from the study of the Metropolis MC dynamics to SGD-like algorithms (eventually adapted to take into account that detailed balance seems to be satisfied in an average sense).

We have analyzed the coloring model, which is a sparse discrete problem for optimization or inference. We have chosen this particular model because we prefer studying a case in which exact predictions for the results of MC algorithms are known: in the case of coloring, MC performances were largely studied in Ref.⁵⁶. Being the coloring a discrete problem, we had to design a generic version of an SGD-like algorithm that could be applied to discrete variables. However, in standard practical applications, SGD is an algorithm working on continuous problems. We should then check the equivalence between MC and SGD in a case with continuous variables. Unfortunately, the behavior of MC algorithms has not been studied systematically in those cases. We plan to extend our work in this sense in the following, both characterizing the behavior of MC and SGD. In this direction, Refs.^{58–60} have shown that a MC dynamics is equivalent, in the limit of small updates of the parameters, to

gradient descent in the presence of Gaussian white noise. However this result is valid only for small updates of the parameters, it relates MC to GD with Gaussian noise that is in principle different from SGD and the equivalence breaks down when the underlying measure is not ergodic, that is when extensive barriers need to be crossed to sample the whole measure, as in the case for complex energy landscapes. We leave for future work a better check (i.e. without the above limitations) for a quantitative correspondence between MC and SGD in problems with continuous variables and complex energy landscapes.

Moreover, the coloring problem that we have analyzed is sparse, because the connectivity of each variable stays finite in the large N limit. One should also check that the same equivalence between MC and SGD holds also for a dense problem, in which the connectivity of nodes could scale with N . However, we think that our work is a fundamental preliminary step that opens a new perspective on the problem and is then quite simple to generalize to more complex problems.

One more important message to take from the present work is the need to optimize the mini-batch size. This is a parameter that is often set to a given value without any good reason. For the planted coloring problem we have studied here, we have shown that only setting the mini-batch size to a value in the retrieval range allows the SGD-like algorithm to converge to the optimal solution. We believe this can be true in many other contexts and the optimization of the mini-batch size should become a standard process among machine learning practitioners. In our specific case, the planted coloring model near to the MC recovery transition, the optimal batch size is nearly 90%, that seems quite challenging for ML applications. However, this is just a specific case, in a regime of parameters very close to the critical connectivity c_{MC} . But in turn, in any realistic ML context, one is generically further away from critical thresholds, also as an effect of strong over-parametrization and we expect the margin for B also to be wider. Moreover, the case we have studied is a sparse case, in which the number of data (that corresponds to edges in our case) is linearly proportional to the size of the signal that we want to recover. If the mini-batch size is too small, the recovery problem quickly becomes impossible because the underlying graph becomes disconnected, and local information cannot propagate to large distances. This is not the common situation in ML problems, that could be instead classified as dense problems. In such cases, we expect an optimal batch size that could be much lower than the one in sparse problems.

We want also to add a last, more general comment. The findings of our paper fit in a more general debate about the nature of fluctuations that drive a system. In fact, the question about the equivalence of thermal fluctuations with respect to other forms of noise driving a system, such as the stochastic selection in our SGD-like algorithm, also arises in other contexts. For instance, the definition of an effective temperature is possible in a variety of non-equilibrium systems: some examples are dense tapped granular systems⁶¹; looking at an intruder immersed in a vibro-fluidized granular medium at small packing fraction (but not at high ones)⁶²; (self-propelled) active matter systems^{63,64}. Moreover two-time aging correlation functions equivalent to the ones of glassy thermal systems are found in frustrated models for granular materials⁶⁵. The problem we looked at is thus just another situation in which a different source of noise leads to consequences analogous to thermal noise.

Data availability

The code source to reproduce data shown in the paper is available at https://github.com/RaffaeleMarino/SGDlike_eq_Mdyn.

Received: 10 November 2023; Accepted: 20 May 2024

Published online: 21 May 2024

References

- Cormen, T. H., Leiserson, C. E., Rivest, R. L. & Stein, C. *Introduction to Algorithms* (MIT press, 2022).
- Cugliandolo, L. F. A scientific portrait of Giorgio Parisi: Complex systems and much more. *J. Phys.: Complex* **4**, 011001 (2023).
- Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H. & Teller, E. Equation of state calculations by fast computing machines. *J. Chem. Phys.* **21**, 1087–1092 (1953).
- Amari, S.-I. Backpropagation and stochastic gradient descent method. *Neurocomputing* **5**, 185–196 (1993).
- Bottou, L. *Stochastic Gradient Descent Tricks*. *Neural Networks: Tricks of the Trade* 2nd edn, 421–436 (Springer, 2012).
- Marino, R. & Ricci-Tersenghi, F. Phase transitions in the mini-batch size for sparse and dense two-layer neural networks. *Mach. Learn.: Sci. Technol.* **5**, 015015 (2024).
- Papadimitriou, C. H. & Steiglitz, K. *Combinatorial Optimization: Algorithms and Complexity* (Courier Corporation, 1998).
- Marino, R. & Kirkpatrick, S. Hard optimization problems have soft edges. *Sci. Rep.* **13**, 3671 (2023).
- Marino, R. & Kirkpatrick, S. Large independent sets on random d -regular graphs with fixed degree d . *Computation* **11**, 206 (2023).
- Angelini, M. C. Parallel tempering for the planted clique problem. *J. Stat. Mech: Theory Exp.* **2018**, 073404 (2018).
- Angelini, M. C. & Ricci-Tersenghi, F. Monte Carlo algorithms are very effective in finding the largest independent set in sparse random graphs. *Phys. Rev. E* **100**, 013302 (2019).
- Mohseni, M. *et al.* Nonequilibrium Monte Carlo for unfreezing variables in hard combinatorial optimization. [arXiv:2111.13628](https://arxiv.org/abs/2111.13628) (2021).
- Huang, K. *Statistical Mechanics* (Wiley, 2008).
- Kirkpatrick, S., Gelatt, C. D. & Vecchi, M. P. Optimization by simulated annealing. *Science* **220**, 671–680 (1983).
- Van Kampen, N. G. *Stochastic Processes in Physics and Chemistry* 1 (Elsevier, 1992).
- Kastner, M. Monte Carlo methods in statistical physics: Mathematical foundations and strategies. *Commun. Nonlinear Sci. Numer. Simul.* **15**, 1589–1602 (2010).
- Robbins, H. & Monro, S. A stochastic approximation method. *Ann. Math. Stat.* **1951**, 400–407 (1951).
- Bottou, L. Online algorithms and stochastic approximations. *Online Learn. Neural Netw.* **1998**, 896 (1998).
- LeCun, Y., Bottou, L., Orr, G. B. & Müller, K.-R. *Neural Networks: Tricks of the Trade* 9–50 (Springer, 2002).
- Bishop, C. M. & Nasrabadi, N. M. *Pattern Recognition and Machine Learning* 4 (Springer, 2006).
- Goodfellow, I., Bengio, Y. & Courville, A. *Deep Learning* (MIT Press, 2016).
- Grohs, P. & Kutyniok, G. *Mathematical Aspects of Deep Learning* (Cambridge University Press, 2022).
- Marino, R. & Macris, N. Solving non-linear Kolmogorov equations in large dimensions by using deep learning: A numerical comparison of discretization schemes. *J. Sci. Comput.* **94**, 8 (2023).

24. Marino, R. Learning from survey propagation: A neural network for MAX-E-3-SAT. *Mach. Learn. Sci. Technol.* **2**, 035032 (2021).
25. Baldassi, C., Lauditi, C., Malatesta, E. M., Perugini, G. & Zecchina, R. Unveiling the structure of wide flat minima in neural networks. *Phys. Rev. Lett.* **127**, 278301 (2021).
26. Baldassi, C. *et al.* Learning through atypical phase transitions in overparameterized neural networks. *Phys. Rev. E* **106**, 014116 (2022).
27. Lucibello, C., Pittorino, F., Perugini, G. & Zecchina, R. Deep learning via message passing algorithms based on belief propagation. *Mach. Learn.: Sci. Technol.* **3**, 035005 (2022).
28. Giambagli, L., Buffoni, L., Carletti, T., Nocentini, W. & Fanelli, D. Machine learning in spectral domain. *Nat. Commun.* **12**, 1330 (2021).
29. Buffoni, L., Civitelli, E., Giambagli, L., Chicchi, L. & Fanelli, D. Spectral pruning of fully connected layers. *Sci. Rep.* **12**, 11201 (2022).
30. Chicchi, L. *et al.* Training of sparse and dense deep neural networks: Fewer parameters, same performance. *Phys. Rev. E* **104**, 054312 (2021).
31. Chicchi, L., Fanelli, D., Giambagli, L., Buffoni, L. & Carletti, T. Recurrent Spectral Network (RSN): Shaping a discrete map to reach automated classification. *Chaos Soliton. Fract.* **168**, 113128 (2023).
32. Ruder, S. An overview of gradient descent optimization algorithms. [arXiv:1609.04747](https://arxiv.org/abs/1609.04747) (2016).
33. Masters, D. & Luschi, C. Revisiting small batch training for deep neural networks. [arXiv:1804.07612](https://arxiv.org/abs/1804.07612) (2018).
34. Lin, T. Stich, S. U. Patel, K. K. Jaggi, M. Don't use large mini-batches, use local sgd. [arXiv:1808.07217](https://arxiv.org/abs/1808.07217) (2018).
35. Mehta, P. *et al.* A high-bias, low-variance introduction to machine learning for physicists. *Phys. Rep.* **810**, 1–124 (2019).
36. Cheng, X., Yin, D., Bartlett, P. & Jordan, M. Stochastic gradient and Langevin processes. *Int. Conf. Mach. Learn.* **2020**, 1810–1819 (2020).
37. Marino, R. & Aurell, E. Advective-diffusive motion on large scales from small-scale dynamics with an internal symmetry. *Phys. Rev. E* **93**, 062147 (2016).
38. Aurell, E., Bo, S., Dias, M., Eichhorn, R. & Marino, R. Diffusion of a Brownian ellipsoid in a force field. *Europhys. Lett.* **114**, 30005 (2016).
39. Han, M., Park, J., Lee, T. & Han, J. H. Fluctuation-dissipation-type theorem in stochastic linear learning. *Phys. Rev. E* **104**, 034126 (2021).
40. Jastrzębski, S. *et al.* Three factors influencing minima in sgd. [arXiv:1711.04623](https://arxiv.org/abs/1711.04623) (2017).
41. Li, Z., Malladi, S. & Arora, S. On the validity of modeling sgd with stochastic differential equations (sdes). *Adv. Neural. Inf. Process. Syst.* **34**, 12712–12725 (2021).
42. Simsekli, U., Sagun, L. & Gurbuzbalaban, M. A tail-index analysis of stochastic gradient noise in deep neural networks. *Int. Conf. Mach. Learn.* **2019**, 5827–5837 (2019).
43. Mézard, M., Parisi, G. & Virasoro, M. A. *Spin Glass Theory and Beyond: An Introduction to the Replica Method and Its Applications* (World Scientific Publishing Company, 1987).
44. Mignacco, F., Krzakala, F., Urbani, P. & Zdeborová, L. Dynamical mean-field theory for stochastic gradient descent in Gaussian mixture classification. *Adv. Neural. Inf. Process. Syst.* **33**, 9540–9550 (2020).
45. Mignacco, F., Urbani, P. & Zdeborová, L. Stochasticity helps to navigate rough landscapes: Comparing gradient-descent-based algorithms in the phase retrieval problem. *Mach. Learn.: Sci. Technol.* **2**, 035029 (2021).
46. Mignacco, F. & Urbani, P. The effective noise of stochastic gradient descent. *J. Stat. Mech.: Theory Exp.* **2022**, 083405 (2022).
47. Kamali, P. J. & Urbani, P. Stochastic Gradient Descent outperforms Gradient Descent in recovering a high-dimensional signal in a glassy energy landscape. [arXiv:2309.04788](https://arxiv.org/abs/2309.04788) (2023).
48. Kubo, R. The fluctuation-dissipation theorem. *Rep. Prog. Phys.* **29**, 255 (1966).
49. Yaida, S. Fluctuation-dissipation relations for stochastic gradient descent. [arXiv:1810.00004](https://arxiv.org/abs/1810.00004) (2018).
50. Cugliandolo, L. F., Kurchan, J. & Peliti, L. Energy flow, partial equilibration, and effective temperatures in systems with slow dynamics. *Phys. Rev. E* **55**, 3898 (1997).
51. Jensen, T. R. & Toft, B. *Graph Coloring Problems* (Wiley, 2011).
52. Zdeborova, L. & Krzakala, F. Phase transitions in the coloring of random graphs. *Phys. Rev. E* **76**, 031131 (2007).
53. Krzakala, F. & Zdeborová, L. Hiding quiet solutions in random constraint satisfaction problems. *Phys. Rev. Lett.* **102**, 238701 (2009).
54. Wright, S. J. Coordinate descent algorithms. *Math. Program.* **151**, 3–34 (2015).
55. Nesterov, Y. Efficiency of coordinate descent methods on huge-scale optimization problems. *SIAM J. Optim.* **22**, 341–362 (2012).
56. Angelini, M. C. & Ricci-Tersenghi, F. Limits and performances of algorithms based on simulated annealing in solving Sparse hard inference problems. *Phys. Rev. X* **13**, 021011 (2023).
57. Cugliandolo, L. F. & Kurchan, J. Analytical solution of the off-equilibrium dynamics of a long-range spin-glass model. *Phys. Rev. Lett.* **71**, 173 (1993).
58. Kikuchi, K., Yoshida, M., Maekawa, T. & Watanabe, H. Metropolis Monte Carlo method as a numerical technique to solve the Fokker-Planck equation. *Chem. Phys. Lett.* **185**, 335–338 (1991).
59. Kikuchi, K., Yoshida, M., Maekawa, T. & Watanabe, H. Metropolis Monte Carlo method for Brownian dynamics simulation generalized to include hydrodynamic interactions. *Chem. Phys. Lett.* **196**, 57–61 (1992).
60. Whitelam, S., Selin, V., Park, S.-W. & Tamblyn, I. Correspondence between neuroevolution and gradient descent. *Nat. Commun.* **12**, 6317 (2021).
61. Makse, H. A. & Kurchan, J. Testing the thermodynamic approach to granular matter with a numerical model of a decisive experiment. *Nature* **415**, 614–617 (2002).
62. Gnoli, A., Puglisi, A., Sarracino, A. & Vulpiani, A. Nonequilibrium Brownian motion beyond the effective temperature. *PLoS ONE* **9**, e93720 (2014).
63. Bechinger, C. . Di. *et al.* Active particles in complex and crowded environments. *Rev. Mod. Phys.* **88**, 045006 (2016).
64. Dal Cengio, S., Levis, D. & Pagonabarraga, I. Fluctuation-dissipation relations in the absence of detailed balance: Formalism and applications to active matter. *J. Stat. Mech.: Theory Exp.* **2021**, 043201 (2021).
65. Nicodemi, M. & Coniglio, A. Aging in out-of-equilibrium dynamics of models for granular media. *Phys. Rev. Lett.* **82**, 916 (1999).

Acknowledgements

This work is supported by ICSC - Centro Nazionale di Ricerca in High Performance Computing, Big Data and Quantum Computing, funded by European Union - NextGenerationEU by PNRR MUR project PE0000013-FAIR, and by PRIN 2022 PNRR, project P20229PBZR “When deep learning is not enough: creating hard benchmarks and physics-inspired algorithms for combinatorial optimization problems”. R.M. is supported by #NEXT-GENERATIONEU (NGEU) and funded by the Ministry of University and Research (MUR), National Recovery and Resilience Plan (NRRP), project MNESYS (PE0000006) - A Multiscale integrated approach to the study of the nervous system in health and disease (DR. 1553 11.10.2022).

Author contributions

The authors have contributed equally to the work.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1038/s41598-024-62625-8>.

Correspondence and requests for materials should be addressed to M.C.A.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2024