


# Correspondence between neuroevolution and gradient descent

Stephen Whitelam <sup>1</sup>✉, Viktor Selin<sup>2</sup>, Sang-Won Park<sup>1</sup> & Isaac Tamblyn<sup>2,3,4</sup>✉

We show analytically that training a neural network by conditioned stochastic mutation or neuroevolution of its weights is equivalent, in the limit of small mutations, to gradient descent on the loss function in the presence of Gaussian white noise. Averaged over independent realizations of the learning process, neuroevolution is equivalent to gradient descent on the loss function. We use numerical simulation to show that this correspondence can be observed for finite mutations, for shallow and deep neural networks. Our results provide a connection between two families of neural-network training methods that are usually considered to be fundamentally different.

<sup>1</sup>Molecular Foundry, Lawrence Berkeley National Laboratory, 1 Cyclotron Road, Berkeley, CA 94720, USA. <sup>2</sup>Department of Physics, University of Ottawa, Ottawa, ON K1N 6N5, Canada. <sup>3</sup>National Research Council of Canada, Ottawa, ON K1N 5A2, Canada. <sup>4</sup>Vector Institute for Artificial Intelligence, Toronto, ON M5G 1M1, Canada. ✉email: [swhitelam@lbl.gov](mailto:swhitelam@lbl.gov); [isaac.tamblyn@uottawa.ca](mailto:isaac.tamblyn@uottawa.ca)

In broad terms there are two types of method used to train neural networks, divided according to whether or not they explicitly evaluate gradients of the loss function. Gradient-based methods include the backpropagation algorithm<sup>1–8</sup>. The non-gradient-based methods (sometimes called “black box” methods) include stochastic processes in which changes to a neural network are proposed and accepted with certain probabilities, and encompass Monte Carlo<sup>9,10</sup> and genetic algorithms<sup>11–13</sup>. Both gradient-based and non-gradient-based methods have been used to train neural networks for a variety of applications, and, where comparison exists, perform similarly well<sup>14–17</sup>. For instance, recent numerical work shows that stochastic evolutionary strategies applied to neural networks are competitive with explicit gradient-based methods on hard reinforcement-learning problems<sup>16</sup>.

Gradient-based- and non-gradient-based strategies are different in implementation and are sometimes thought of as entirely different approaches<sup>18</sup>. Here, we show that the two sets of methods have a fundamental connection. We demonstrate analytically an equivalence between the dynamics of neural-network training under conditioned stochastic mutations, and under gradient descent. This connection follows from one identified in the 1990s between the overdamped Langevin dynamics and Metropolis Monte Carlo dynamics of a particle in an external potential<sup>19,20</sup>. In the limit of small Monte Carlo trial moves, those things are equivalent. Similarly, we show here that a single copy of a neural network (a single individual) exposed to parameter mutations that are accepted probabilistically is equivalent, in the limit of small mutation size, to gradient descent on the loss function in the presence of Gaussian white noise. The details of the resulting dynamics depend on the details of the acceptance criterion, and encompass both standard- and clipped gradient descent. Such a mutation scheme corresponds to the simple limit of the set of processes called “neuroevolution”<sup>13,16,21–23</sup>. This connection demonstrates explicitly that optimization without access to gradients can, nonetheless, enact noisy gradient descent on the loss function.

In simple gradient descent, equivalent to noise-free overdamped Langevin dynamics, the parameters (weights and biases)  $\mathbf{x}$  of a neural network evolve with training time according to the prescription  $\dot{\mathbf{x}} = -\alpha \nabla U(\mathbf{x})$ . Here,  $\alpha$  is a learning rate, and  $\nabla U(\mathbf{x})$  is the gradient of a loss function  $U(\mathbf{x})$  with respect to the network parameters. Now consider a simple neuroevolution scheme in which we propose a mutation  $\mathbf{x} \rightarrow \mathbf{x} + \epsilon$  of all neural-network parameters, where  $\epsilon$  is a set of independent Gaussian random numbers of zero mean and variance  $\sigma^2$ . Let us accept the proposal with the Metropolis probability  $\min(1, e^{-\beta \Delta U})$ . Here,  $\beta$  is a reciprocal temperature, and  $\Delta U$  is the change of the loss function under the proposal. This is a Metropolis Monte Carlo algorithm, a Markovian dynamics that constitutes a form of importance sampling, and a common choice in the physics literature<sup>9,10,24</sup>. In physical systems,  $\beta$  is inversely proportional to the physical temperature, and we consider finite values of  $\beta$  in order to make contact with that literature. However, in the context of training a neural network it is interesting to consider the zero-temperature limit  $\beta = \infty$ , where mutations are accepted only if the loss does not increase. That regime is not normally considered in particle-based simulations.

Our main results can be summarized as follows. When  $\beta \neq \infty$  the weights of the network evolve, to leading order in  $\sigma$ , as  $\dot{\mathbf{x}} = -(\beta \sigma^2 / 2) \nabla U(\mathbf{x})$  plus Gaussian white noise. Averaged over independent realizations of the learning process, this form of neuroevolution is therefore equivalent to simple gradient descent, with learning rate  $\beta \sigma^2 / 2$ . In the limit  $\beta = \infty$ , where mutations are accepted only if the loss function does not increase, weights under

neuroevolution evolve instead as  $\dot{\mathbf{x}} = -(\sigma / \sqrt{2\pi}) |\nabla U(\mathbf{x})|^{-1} \nabla U(\mathbf{x})$  plus Gaussian white noise, which corresponds to clipped gradient descent on  $U(\mathbf{x})$ <sup>25</sup>. Note that conditioning the acceptance of neural-network parameter mutations on the change of the loss function for a single copy of that network is sufficient to enact gradient descent: a population of individuals is not required.

In this paper, we use the term “neuroevolution” to refer to a sequence of mutation steps applied to the parameters of a single copy of a neural network and accepted probabilistically. In general, neuroevolutionary algorithms encompass a broader variety of processes, including mutations of populations of communicating neural networks<sup>16</sup> and mutations of network topologies<sup>21,26,27</sup>. Similarly, the set of procedures for particles that can be described as “Monte Carlo algorithms” is large, and ranges from local moves of single particles—roughly equivalent to the procedure used here—to nonlocal moves and moves of collections of particles<sup>24,28–31</sup>. The dynamics of those collective-move Monte Carlo algorithms and of the more complicated neuroevolutionary schemes<sup>21,26,27</sup> do not correspond to simple gradient descent. Here, we demonstrate a correspondence between one member of this set of algorithms and gradient descent, the implication being that, given any potentially complicated set of neuroevolutionary methods, it is enough to add a simple mutation-acceptance protocol in order to ensure that gradient descent is also approximated. The neuroevolution-gradient descent correspondence is similar to the proofs that neural networks with enough hidden nodes can represent any smooth function<sup>32</sup>: it does not necessarily suggest how to solve a given problem, but provides understanding of the limits and capacity of the tool and its relation to other methods of learning.

Our work provides a rigorous connection between gradient descent and what is arguably the simplest form of neuroevolution. It complements studies that demonstrate a numerical similarity between gradient-based methods and population-based evolutionary methods<sup>16,17</sup>, and studies that show analytically that the gradients approximated by those methods are, under certain conditions, equivalent to the finite-difference gradient<sup>33–35</sup>.

The paper is structured as follows. We summarize the neuroevolution-gradient descent correspondence in section “Results”, and derive it in section “Methods”. Our derivation uses ideas developed in ref. 20 to treat physical particles, and applies them to neural networks: we consider a different set-up (in effect, we work with a single particle in a high-dimensional space, rather than with many particles in three-dimensional space) and proposal rates, and we consider the limit  $\beta = \infty$  that is rarely considered in the physics literature but is natural in the context of a neural network. We can associate the state  $\mathbf{x}$  of the neural network with the position of a particle in a high-dimensional space, and the loss function  $U(\mathbf{x})$  with an external potential. The result is a rewriting of the correspondence between Langevin dynamics and Monte Carlo dynamics as a correspondence between the simplest forms of gradient descent and neuroevolution. Just as the Langevin-Monte Carlo correspondence provides a basis for understanding why Monte Carlo simulations of particles can approximate real dynamics<sup>31,36–41</sup>, so the neuroevolution-gradient descent correspondence shows how we can effectively perform gradient descent on the loss function without explicit calculation of gradients. The correspondence holds exactly only in the limit of vanishing mutation scale, but we use numerics to show in section “Numerical illustration of the neuroevolution-gradient descent correspondence” that it can be observed for neuroevolution done with finite mutations and gradient descent enacted with a finite timestep. We conclude in section “Conclusions”.

**Results**

In this section, we summarize the main analytic results of this paper. These results are derived in section “Methods”.

Consider a neural network with  $N$  parameters (weights and biases)  $\mathbf{x} = \{x_1, \dots, x_N\}$ , and a loss  $U(\mathbf{x})$  that is a deterministic function of the network parameters. The form of the network does not enter the proof, and so the result applies to neural networks of any architecture (we shall illustrate this point numerically by considering both deep and shallow nets). The loss function may also depend upon other parameters, such as a set of training data, as in supervised learning, or a set of actions and states, as in reinforcement learning; the correspondence we shall describe applies regardless.

**Gradient descent.** Under the simplest form of gradient descent, the parameters  $x_i$  of the network evolve according to numerical integration of

$$\frac{dx_i}{dt} = -\alpha \frac{\partial U(\mathbf{x})}{\partial x_i}. \tag{1}$$

Here, time  $t$  measures the progress of training, and  $\alpha$  is the learning rate<sup>3–7</sup>.

**Neuroevolution.** Now consider training the network by neuroevolution, defined by the following Monte Carlo protocol.

1. Initialize the neural-network parameters  $\mathbf{x}$  and calculate the loss function  $U(\mathbf{x})$ . Set time  $t = 0$ .
2. Propose a change (or “mutation”) of each neural-network parameter by an independent Gaussian random number of zero mean and variance  $\sigma^2$ , so that

$$\mathbf{x} \rightarrow \mathbf{x} + \boldsymbol{\epsilon}, \tag{2}$$

where  $\boldsymbol{\epsilon} = \{\epsilon_1, \dots, \epsilon_N\}$  and  $\epsilon_i \sim \mathcal{N}(0, \sigma^2)$ .

3. Accept the mutation with the Metropolis probability  $\min(1, e^{-\beta[U(\mathbf{x}+\boldsymbol{\epsilon})-U(\mathbf{x})]})$ , and otherwise reject it. In the latter case we return to the original neural network. The parameter  $\beta$  can be considered to be a reciprocal evolutionary temperature.
4. Increment time  $t \rightarrow t + 1$ , and return to step 2.

For finite  $\beta$ , and in the limit of small mutation scale  $\sigma$ , the parameters of the neural network evolve under this procedure according to the Langevin equation

$$\frac{dx_i}{dt} = -\frac{\beta\sigma^2}{2} \frac{\partial U(\mathbf{x})}{\partial x_i} + \xi_i(t), \tag{3}$$

where  $\xi$  is a Gaussian white noise with zero mean and variance  $\sigma^2$ :

$$\langle \xi_i(t) \rangle = 0, \quad \langle \xi_i(t)\xi_j(t') \rangle = \sigma^2 \delta_{ij} \delta(t - t'). \tag{4}$$

Eq. (3) describes an evolution of the neural-network parameters  $x_i$  that is equivalent to gradient descent with learning rate  $\alpha = \beta\sigma^2/2$  in the presence of Gaussian white noise. Averaging over independent stochastic trajectories of the learning process (starting from identical initial conditions) gives

$$\frac{d\langle x_i \rangle}{dt} = -\frac{\beta\sigma^2}{2} \frac{\partial U(\mathbf{x})}{\partial x_i}, \tag{5}$$

which has the same form as the gradient descent equation (1). Thus, when averaged over many independent realizations of the learning process, the neuroevolution procedure 1–4, with finite  $\beta$ , is equivalent in the limit of small mutation scale to gradient descent on the loss function.

In the case  $\beta = \infty$ , where mutations are only accepted if the loss function does not increase, the parameters of the network evolve

according to the Langevin equation

$$\frac{dx_i}{dt} = -\frac{\sigma}{\sqrt{2\pi}} \frac{1}{|\nabla U(\mathbf{x})|} \frac{\partial U(\mathbf{x})}{\partial x_i} + \eta_i(t), \tag{6}$$

where  $\eta$  is a Gaussian white noise with zero mean and variance  $\sigma^2/2$ :

$$\langle \eta_i(t) \rangle = 0, \quad \langle \eta_i(t)\eta_j(t') \rangle = \frac{\sigma^2}{2} \delta_{ij} \delta(t - t'). \tag{7}$$

The form (6) is different to (3), because the gradient in the first term is normalized by the factor  $|\nabla U(\mathbf{x})| = \sqrt{\sum_{i=1}^N (\partial U(\mathbf{x})/\partial x_i)^2}$ , which serves as an effective coordinate-dependent rescaling (or vector normalization) of the timestep, but (6) nonetheless describes a form of gradient descent on the loss function  $U(\mathbf{x})$ . Note that the drift term in (6) is of lower order in  $\sigma$  than the diffusion term (which is not the case for finite  $\beta$ ). In the limit of small  $\sigma$ , (6) describes an effective process in which uphill moves in loss cannot be made, consistent with the stochastic process from which it is derived.

Averaged over independent realizations of the learning process, (6) reads

$$\frac{d\langle x_i \rangle}{dt} = -\frac{\sigma}{\sqrt{2\pi}} \frac{1}{|\nabla U(\mathbf{x})|} \frac{\partial U(\mathbf{x})}{\partial x_i}. \tag{8}$$

The results (3) and (6) show that training a network by making random mutations of its parameters is, in the limit of small mutations, equivalent to noisy gradient descent on the loss function.

Writing  $\dot{U}(\mathbf{x}) = \dot{\mathbf{x}} \cdot \nabla U(\mathbf{x})$ , using (3) and (6), and averaging over noise shows the evolution of the mean loss function under neuroevolution to obey, in the limit of small  $\sigma$ ,

$$\langle \dot{U}(\mathbf{x}) \rangle = \begin{cases} -\frac{\beta\sigma^2}{2} (\nabla U(\mathbf{x}))^2 & \text{if } \beta \neq \infty \\ -\frac{\sigma}{\sqrt{2\pi}} |\nabla U(\mathbf{x})| & \text{if } \beta = \infty \end{cases}, \tag{9}$$

equivalent to evolution under the noise-free forms of gradient descent (5) and (8).

In section “Numerical illustration of the neuroevolution-gradient descent correspondence” we illustrate numerically the correspondence described here, and show that it can be observed numerically for non-vanishing mutations and finite integration steps. In section “Methods” we detail the derivation of the correspondence.

**Discussion**

**Numerical illustration of the neuroevolution-gradient descent correspondence.** In this section, we demonstrate the neuroevolution-gradient descent correspondence numerically. We consider single-layer neural networks for the cases of infinite and finite  $\beta$ , and a deep network for the case of infinite  $\beta$ .

*Shallow net,  $\beta = \infty$ .* In order to observe correspondence numerically, the neuroevolution mutation scale  $\sigma$  must be small enough that correction terms neglected in the expansion leading to (25) and (55) are small. The required range of  $\sigma$  is difficult to know in advance, but straightforward to determine empirically: below the relevant value of  $\sigma$ , the results of neuroevolution will be statistically similar when scaled in the manner described below.

We consider a simple supervised-learning problem in which we train a neural network to express the function  $f_0(\theta) = \sin(2\pi\theta)$  on the interval  $\theta \in [0, 1)$ . We calculated the loss using  $K = 1000$  points on the interval,

$$U(\mathbf{x}) = \frac{1}{K} \sum_{j=0}^{K-1} [f_{\mathbf{x}}(j/K) - f_0(j/K)]^2, \tag{10}$$

where

$$f_{\mathbf{x}}(\theta) = \sum_{i=0}^{M-1} x_{3i+1} \tanh(x_{3i+2}\theta + x_{3i+3}) \quad (11)$$

is the output of a single-layer neural network with one input node, one output node,  $M=30$  hidden nodes, and  $N=3M$  parameters  $x_i$ . These parameters are initially chosen to be Gaussian random numbers with zero mean and variance  $\sigma_0^2 = 10^{-4}$ . The correspondence is insensitive to the choice of initial conditions, and we shall show that it holds for different choices of initial network.

We performed gradient descent with learning rate  $\alpha = 10^{-5}$ . We chose the learning rate arbitrarily, and verified that the results of gradient-descent simulations were unchanged upon a changing learning rate by a factor of 10 and 1/10. We used Euler integration of the noise-free version of Eq. (6), updating all weights  $x_i$  at each timestep  $t_{\text{gd}} = 1, 2, \dots$  according to the prescription

$$x_i(t_{\text{gd}} + 1) = x_i(t_{\text{gd}}) - \frac{\alpha}{|\nabla U(\mathbf{x})|} \frac{\partial U(\mathbf{x})}{\partial x_i}, \quad (12)$$

where

$$\frac{\partial U(\mathbf{x})}{\partial x_i} = \frac{2}{K} \sum_{j=0}^{K-1} [f_{\mathbf{x}}(j/K) - f_0(j/K)] \frac{\partial f_{\mathbf{x}}(j/K)}{\partial x_i}, \quad (13)$$

and

$$\frac{\partial f_{\mathbf{x}}(\theta)}{\partial x_i} = \begin{cases} \tanh(\theta x_{i+1} + x_{i+2}) & \text{if } i, \text{ mod } 3 = 1 \\ \theta x_{i-1} \text{sech}^2(\theta x_i + x_{i+1}) & \text{if } i \text{ mod } 3 = 2 \\ x_{i-2} \text{sech}^2(\theta x_{i-1} + x_i) & \text{if } i \text{ mod } 3 = 0. \end{cases}$$

We did neuroevolution following the Monte Carlo procedure described in section ‘‘Neuroevolution’’, in the limit  $\beta = \infty$ , i.e., we accepted only moves that did not increase the loss function. We chose the mutation scale

$$\sigma = \lambda \alpha \sqrt{2\pi}, \quad (14)$$

where  $\lambda$  is a parameter. According to (6) and (12), this prescription sets the neuroevolution timescale  $t_{\text{evol}}$  to be a factor  $\lambda$  times that of the gradient-descent timescale. Thus, one neuroevolution step corresponds to  $\lambda$  integration steps of the gradient descent procedure. In figures, we compare gradient descent with neuroevolution as a function of common (scaled) time  $t = \alpha t_{\text{gd}} = \alpha \lambda t_{\text{evol}}$ .

In Fig. 1(a) we show the evolution of four individual weights under neuroevolution (using mutation scale  $\lambda = 1/10$ ) and gradient descent (weights are distinguishable because they always have the same initial values). The correspondence predicted analytically can be seen numerically: individual neuroevolution trajectories (gray) fluctuate around the gradient descent result (black), and when averaged over individual trajectories the results of neuroevolution (green) approximate those of gradient descent. In Fig. 1(b) we show the individual and averaged values of the weights of neuro-evolved networks at time  $t = 10$  compared to those of gradient descent. In general, the weights generated by averaging over neuroevolution trajectories approximate those of gradient descent, with some discrepancy seen in the values of the largest weights. In Fig. 1(c) we show the loss under neuroevolution and gradient descent. As predicted by (9), averaged neuroevolution and gradient descent are equivalent.

In Supplementary Fig. 1 we show similar quantities using a different choice of initial neural network; the correspondence between neuroevolution and gradient descent is again apparent.

In Fig. 2(a) we show the time evolution of a single weight of the network under gradient descent and neuroevolution, the latter for three sizes of mutation step  $\sigma$ . As  $\sigma$  increases, the size of

fluctuations of individual trajectories about the mean increase, as predicted by (6). As a result, more trajectories are required to estimate the average, and for fixed number of trajectories (as used here) the estimated average becomes less precise. In addition, as  $\sigma$  increases, the assumptions underlying the correspondence derivation eventually break down, in which case the neuroevolution average will not converge to the gradient descent result even as more trajectories are sampled.

In Fig. 2(b) we show the mean-squared difference of the parameter vector of the model under gradient descent and neuroevolution,

$$\Delta(t) \equiv \frac{1}{N} \sum_{i=1}^N \left( x_i^{\text{gradient}}(t) - \langle x_i^{\text{evolution}}(t) \rangle \right)^2. \quad (15)$$

Here,  $N$  is the number of network parameters;  $x_i^{\text{gradient}}(t)$  is the time evolution of neural-network parameter  $i$  under the gradient descent equation Eq. (12); and  $\langle x_i^{\text{evolution}}(t) \rangle$  is the mean value of neural-network parameter  $i$  over the ensemble of neuroevolution trajectories. The smaller the neuroevolution step size, the smaller is  $\Delta(t)$ , and the closer the neuroevolution-gradient descent correspondence.

In Supplementary Fig. 2 we show the evolution with time of the loss for different mutation scales (the left-hand plot is a reproduction of Fig. 1(c)). The trend shown is similar to that of the weights in Fig. 2.

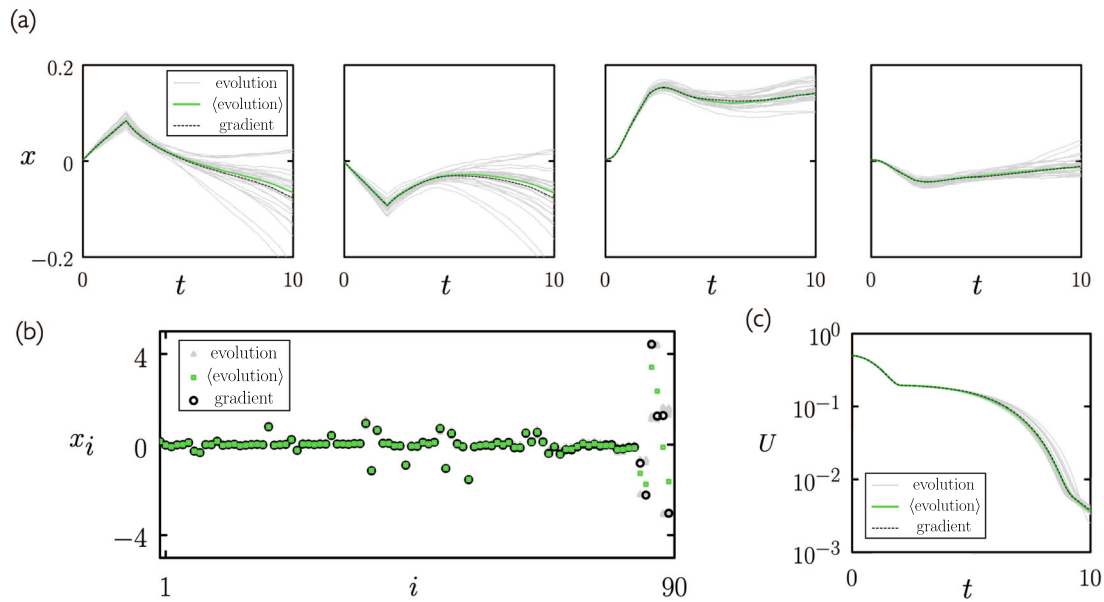
*Deep net*,  $\beta = \infty$ . One feature of the correspondence derivation is that the architecture of the neural network does not appear. As long as the loss  $U(\mathbf{x})$  is a deterministic function of the neural-network parameters  $\mathbf{x}$ , correspondence between gradient descent and neuroevolution will be observed if the mutation scale is small enough (it is likely that what constitutes ‘‘small enough’’ does depend on neural-network architecture, as well as the problem under study. The required mutation scale can be determined empirically, even without access to gradient-descent results: when correspondence holds, the results of neuroevolution simulations will be statistically similar, when scaled as we have described).

To demonstrate invariance to architecture we repeat the previous comparison, now using a deep neural network (we train the net to reproduce the target function  $f_0(\theta) = \sin(\pi\theta)$  on the interval  $\theta \in [-1, 1]$ ). The network has 8 fully-connected hidden layers, each 32 nodes wide, and 7489 total parameters. As before we use tanh activations on the hidden nodes, and have one input node and one output.

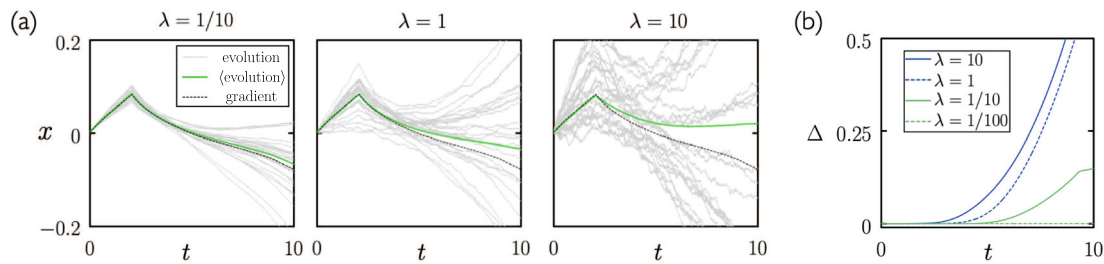
Results are shown in Fig. 3. In panel (a) we show the evolution of two parameters of the network, under gradient descent and for neuroevolution with step-size parameter  $\lambda = 1/10$  [see (17)]. As for the shallow net, the correspondence is apparent. Neuroevolution averages (green lines) are taken over 100 trajectories. In panel (b) we show the loss, for gradient descent and two different neuroevolution step-size parameters. As expected, the correspondence is more precise for smaller  $\lambda$ . As before, for large enough  $\lambda$  the correspondence breaks down: see Supplementary Fig. 3.

In Fig. 4 we show all parameters of the deep net at training time  $t = 5$ , under the two dynamics. We show the results of gradient descent in black, and independent neuroevolution trajectories in gray. As predicted analytically, the neuroevolution results fall either side of the gradient-descent result, and the network constructed by averaging over independent neuroevolution trajectories (green) is essentially identical to the network produced by gradient descent.


In Fig. 5, we illustrate the dynamics of learning and the scale of the loss function by showing a comparison between the target function  $f_0(\theta) = \sin(\pi\theta)$  and the net function  $f_{\mathbf{x}}(\theta)$ . We show the



**Fig. 1 Numerical illustration of the equivalence of gradient descent and neuroevolution in the case  $\beta = \infty$ .** **a** Evolution with time of 4 of the 90 weights of the neural network (11) under the gradient descent equation Eq. (12) (black: “gradient” stands for “gradient descent”), and under neuroevolution with mutation scale  $\lambda = 1/10$  [see (14)]. Here and in subsequent panels we show 25 independent neuroevolution trajectories (gray, marked “evolution”) and the average of 1000 independent trajectories (green, marked “(evolution)”). **b** All weights at time  $t = 10$  under the two methods. **c** Loss as a function of time.



**Fig. 2 The smaller the neuroevolution mutation scale, the closer the neuroevolution-gradient descent correspondence.** **a** Time evolution of a single neural-network weight, for three different neuroevolution mutation scales [see (14)]. **b** The quantity  $\Delta(t)$ , Eq. (15), is a measure of the difference between networks evolved under gradient descent and neuroevolution.



**Fig. 3 As Fig. 1, now for a deep neural network with eight hidden layers.** Panel **a** shows 2 of the 7489 parameters of the network. Panel **b** shows the loss, for two different neuroevolution step-size parameters [see (17)].

latter at three different training times, for gradient descent and neuroevolution trajectories.


*Shallow net, finite  $\beta$ .* In this section, we illustrate the gradient descent-neuroevolution correspondence for finite  $\beta$ . We consider the same supervised-learning problem as before, and set the network width to 256 nodes. We did neuroevolution with the Metropolis acceptance rate with reciprocal temperature parameter  $\beta = 10^3$ . This choice is arbitrary, but is representative of a wide range of finite values of  $\beta$ . Finite-temperature simulations

are common in particle-based systems<sup>24</sup>. Here, temperature has no particular physical significance, but comparing simulations done at finite and infinite  $\beta$  makes the point that different choices of neuroevolution acceptance rate result in a dynamics equivalent to different gradient-descent protocols.

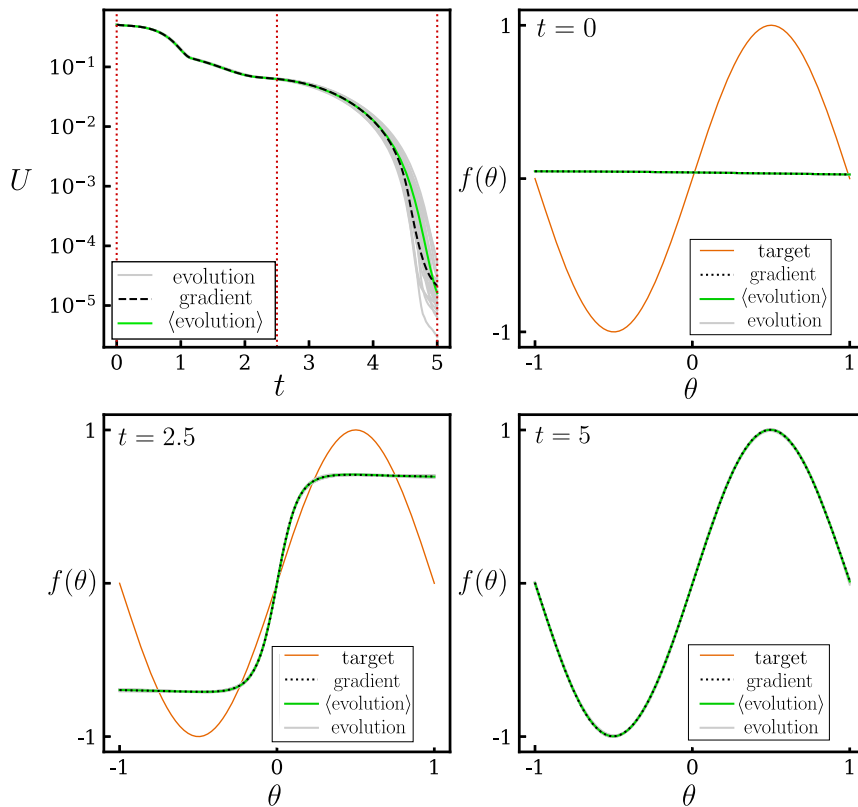
We did gradient descent using the integration scheme

$$x_i(t_{gd} + 1) = x_i(t_{gd}) - \alpha \frac{\partial U(\mathbf{x})}{\partial x_i}, \tag{16}$$

where  $\alpha = 10^{-4}$  is the learning rate. Comparing (3) and (16), we



**Fig. 4 Numerical illustration of correspondence for a deep net.** We show all 7489 parameters  $x_i$  of the deep net of Fig. 3 at training time  $t = 5$ . As predicted analytically, the neural network created by averaging (green symbols) over a noninteracting population of neuroevolutionary processes (gray symbols) is essentially the same network as that produced by gradient descent (black symbols). For clarity, weights are ordered by their final values under gradient descent.



**Fig. 5 Illustration of the learning dynamics.** Loss versus time and comparison of the target function  $f_0(\theta) = \sin(\pi\theta)$  and the net function  $f_{\mathbf{x}}(\theta)$  at three different training times, for the learning dynamics of Fig. 3 ( $\lambda = 1/10$  for the neuroevolution results).

set the neuroevolution mutation scale to be

$$\sigma = \lambda \sqrt{2\alpha/\beta}, \quad (17)$$

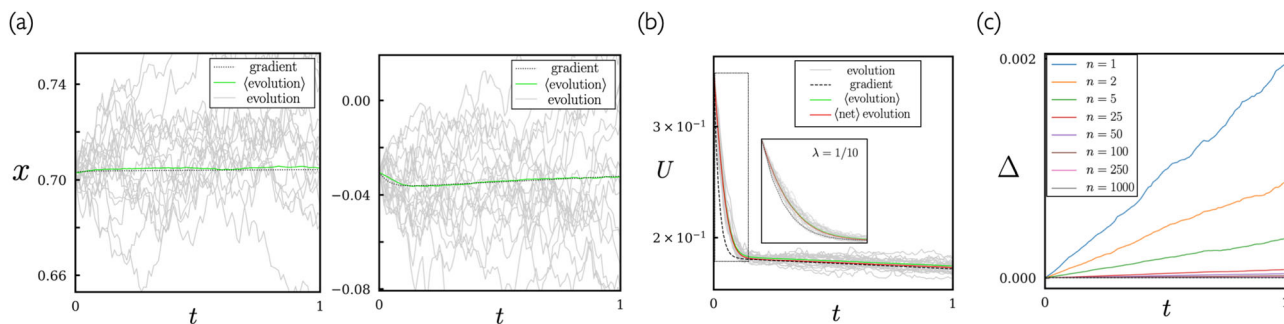
where  $\lambda$  is a parameter. Thus, one neuroevolution step corresponds to  $\lambda$  integration steps of the gradient descent procedure. In figures, we compare gradient descent with neuroevolution as a function of common (scaled) time  $t = \alpha t_{\text{gd}} = \alpha \lambda t_{\text{evol}}$ .

Results are shown in Fig. 6, for step-size parameter  $\lambda = 1$ . In panel (a) we show the evolution with time of two of the weights of the network. The noise associated with neuroevolution at this value of  $\beta$  is considerable: individual trajectories (gray lines) bear little resemblance to the gradient-descent result (black line). However, the population average (green line, average of  $n = 1000$  trajectories) shows the expected correspondence. The

correspondence is less precise than that shown previously, because we use a larger effective step size and each trajectory is much noisier than its infinite- $\beta$  counterpart.

In Fig. 6(b) we show the loss, with line colors corresponding to the quantities of panel (a). In addition, we show the loss of the average network produced by neuroevolution (red line),  $U(\langle \mathbf{x} \rangle)$ , which, if correspondence holds, should be equal to  $\langle U(\mathbf{x}) \rangle$  (green line). The initial fast relaxation of the loss (the boxed region) shows a difference between gradient descent and averaged neuroevolution results; doing neuroevolution for smaller step-size parameter  $\lambda = 1/10$  (inset) reduces this difference, as expected.

In panel (c) we show the parameter  $\Delta$ , Eq. (15), a measure of the difference between the average network  $\langle \mathbf{x} \rangle$  produced by neuroevolution and the network produced by gradient descent, as



**Fig. 6** Illustration of the neuroevolution-gradient descent correspondence for the case of finite  $\beta$ . Panel **a** shows two parameters of the network, panel **b** shows the loss, and panel **c** shows the parameter  $\Delta$ , Eq. (15), a measure of the difference between the average network produced by neuroevolution and the network produced by gradient descent.

a function of  $n$ , the number of trajectories included in the average. If correspondence holds, this quantity should vanish in the limit of large  $n$ ; the observed trend is consistent with this behavior.

In Supplementary Fig. 4, we compare a gradient-descent trajectory with a set of neuroevolution trajectories, periodically resetting the latter to the gradient-descent solution. The periodic resetting tests the correspondence for a range of initial conditions. The correspondence between gradient descent and the averaged neuroevolution trajectory is approximate (averages were taken over 152 trajectories, fewer than in Supplementary Fig. 4) but apparent.

**Conclusions.** We have shown analytically that training a neural network by neuroevolution of its weights is equivalent, in the limit of small mutation scale, to noisy gradient descent on the loss function. Conditioning neuroevolution on the Metropolis acceptance criterion at finite evolutionary temperature is equivalent to a noisy version of simple gradient descent, while at infinite reciprocal evolutionary temperature the procedure is equivalent to clipped gradient descent on the loss function. Averaged over noise, the evolutionary procedures correspond to forms of gradient descent on the loss function. This correspondence is described by Equations (3), (5), (6), and (8).

Correspondence in the sense described above means that each neural-network parameter evolves the same way as a function of time under the two dynamics. Correspondence implies that the convergence properties of the two methods are the same (see e.g., Fig. 3(b)) and that the neural networks produced by the same methods are the same (see e.g., Fig. 4). The generalization properties of those networks will then also be the same.

The correspondence is formally exact only in the limit of zero mutation scale, and holds approximately for small but finite mutations. It will fail when the assumptions underlying the derivation are violated, such as when the terms neglected in (22) and (23) are not small, or when the passage from (28) to (29) is not valid because the change  $\beta U(\mathbf{x} + \epsilon) - \beta U(\mathbf{x})$  is not small. It is straightforward to determine empirically where correspondence holds, even without access to gradient-descent results: the results of neuroevolution, with time scaled as described, will be statistically similar when the mutation size is small enough. The time duration for which the correspondence holds increases with decreasing mutation scale (see e.g., Fig. 2(b)). We have shown here that the correspondence can be observed for a range of mutation scales, and for different neural-net architectures.

More generally, several dynamical regimes are contained within the neuroevolution master equation (19), according to the scale  $\sigma$  of mutations: for vanishing  $\sigma$ , neuroevolution is formally noisy gradient descent on the loss function; for small but nonvanishing  $\sigma$  it approximates noisy gradient descent enacted by

explicit integration with a finite timestep; for larger  $\sigma$  it enacts a dynamics different to gradient descent, but one that can still learn; and for sufficiently large  $\sigma$  the steps taken are too large for learning to occur on accessible timescales. An indication of these various regimes can be seen in Fig. 2 and Supplementary Fig. 2.

Separate from the question of its precise temporal evolution, the master equation (19) has a well-defined stationary distribution  $\rho_0(\mathbf{x})$ . Requiring the brackets on the right-hand of (19) to vanish ensures that  $P(\mathbf{x}, t) \rightarrow \rho_0(\mathbf{x})$  becomes independent of time. Inserting (20) into (19) and requiring normalization of  $\rho_0(\mathbf{x})$  reveals the stationary distribution to be the Boltzmann one,  $\rho_0(\mathbf{x}) = e^{-\beta U(\mathbf{x})} / \int d\mathbf{x}' e^{-\beta U(\mathbf{x}')}$ . For finite  $\beta$  the neuroevolution procedure is ergodic, and this distribution will be sampled given sufficiently long simulation time. For  $\beta \rightarrow \infty$  we have  $\rho_0(\mathbf{x}) \rightarrow \delta(U(\mathbf{x}) - U_0)$ , where  $U_0$  is the global energy minimum; in this case the system is not ergodic (moves uphill in  $U(\mathbf{x})$  are not allowed) and there is no guarantee of reaching this minimum.

We have focused on the simple limit of the set of neuroevolution algorithms, namely a non-interacting population of neural networks that experience sequential probabilistic mutations of their parameters. We have illustrated the correspondence at the level of population averages, Equations (5) and (8). However, no communication between individuals is required, and each individually observes the correspondence defined by Equations (3) and (6).

Our results are also relevant to population-based genetic algorithms in which members of the population are periodically reset to the identities of the individuals with lowest loss values<sup>11–13</sup>. For instance, when correspondence holds, individuals in the neuroevolution populations considered in this paper have an averaged loss equal to that of the corresponding gradient descent algorithm. Therefore, some individuals must have loss less than that of the corresponding gradient descent algorithm (see e.g. Fig. 1(a), and Fig. 3(b) for the case  $\lambda = 1/10$ ). This observation indicates the potential for such methods to be competitive with gradient-descent algorithms.

The neuroevolution-gradient descent correspondence we have identified follows from that between the overdamped Langevin dynamics and Metropolis Monte Carlo dynamics of a particle in an external potential<sup>19,20</sup>. Our work therefore adds to the existing set of connections between machine learning and statistical mechanics<sup>42,43</sup>, and continues a trend in machine learning of making use of old results: the stochastic and deterministic algorithms considered here come from the 1950s<sup>9,10</sup> and 1970s<sup>1–6</sup>, and are connected by ideas developed in the 1990s<sup>19,20</sup>.

**Methods**

**Derivation of the neuroevolution-gradient descent correspondence.** We start by considering the quantity  $P(\mathbf{x}, t)$ , the probability that a neural network has the set of parameters  $\mathbf{x}$  at time  $t$  under a given stochastic protocol. The time evolution of

this quantity is governed by the master equation<sup>44,45</sup>, which in generic form reads

$$\partial_t P(\mathbf{x}, t) = \int d\mathbf{x}' [P(\mathbf{x}', t) W_{\mathbf{x}-\mathbf{x}'}(\mathbf{x}') - P(\mathbf{x}, t) W_{\mathbf{x}'-\mathbf{x}}(\mathbf{x})]. \tag{18}$$

The two terms in (19) describe, respectively, gain and loss of the probability  $P(\mathbf{x}, t)$  (note that the probability to have some set of parameters is conserved, i.e.,  $\int d\mathbf{x} P(\mathbf{x}, t) = 1$ ). The symbol  $W_{\mathbf{x}'-\mathbf{x}}(\mathbf{x})$  (sometimes written  $W(\mathbf{x} \rightarrow \mathbf{x}')$ ) quantifies the probability of moving from the set of parameters  $\mathbf{x} + (\mathbf{x}' - \mathbf{x}) = \mathbf{x}'$ , and encodes the details of the stochastic protocol. For the neuroevolution procedure defined in section “Neuroevolution” we write (18) as

$$\partial_t P(\mathbf{x}, t) = \int d\epsilon [P(\mathbf{x} - \epsilon, t) W_\epsilon(\mathbf{x} - \epsilon) - P(\mathbf{x}, t) W_\epsilon(\mathbf{x})]. \tag{19}$$

Here,  $\epsilon$  denotes the set of random numbers (the “mutation”) by which the neural-network parameters are changed; the integral  $\int d\epsilon = \int_{-\infty}^{\infty} d\epsilon_i \dots \int_{-\infty}^{\infty} d\epsilon_N$  runs over all possible choices of mutations; and

$$W_\epsilon(\mathbf{x}) = p(\epsilon) \min(1, e^{-\beta[U(\mathbf{x}+\epsilon)-U(\mathbf{x})]}) \tag{20}$$

is the rate for going from the set of parameters  $\mathbf{x}$  to the set of parameters  $\mathbf{x} + \epsilon$ . Eq. (20) contains two factors. The first,

$$p(\epsilon) = \prod_{i=1}^N p(\epsilon_i) \quad \text{with} \quad p(\epsilon_i) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{\epsilon_i^2}{2\sigma^2}}, \tag{21}$$

quantifies the probability of proposing a set of Gaussian random numbers  $\epsilon$ . The second factor, the Metropolis’ “min” function in (20), quantifies the probability of accepting the proposed move from  $\mathbf{x}$  to  $\mathbf{x} + \epsilon$ ; recall that  $U(\mathbf{x})$  is the loss function.

We can pass from the master equation (19) to a Fokker-Planck equation by assuming a small mutation scale  $\sigma$ , and expanding the terms in (19) to second order in  $\sigma$ <sup>44,45</sup>. Thus

$$P(\mathbf{x} - \epsilon, t) \approx \left(1 - \epsilon \cdot \nabla + \frac{1}{2} (\epsilon \cdot \nabla)^2\right) P(\mathbf{x}, t) \tag{22}$$

and

$$W_\epsilon(\mathbf{x} - \epsilon) \approx \left(1 - \epsilon \cdot \nabla + \frac{1}{2} (\epsilon \cdot \nabla)^2\right) W_\epsilon(\mathbf{x}), \tag{23}$$

where  $\epsilon \cdot \nabla = \sum_{i=1}^N \epsilon_i \partial_i$  (note that  $\partial_i \equiv \frac{\partial}{\partial x_i}$ ). Collecting terms resulting from the expansion gives

$$\partial_t P(\mathbf{x}, t) \approx - \int d\epsilon (\epsilon \cdot \nabla) P(\mathbf{x}, t) W_\epsilon(\mathbf{x}) + \frac{1}{2} \int d\epsilon (\epsilon \cdot \nabla)^2 P(\mathbf{x}, t) W_\epsilon(\mathbf{x}). \tag{24}$$

Taking the integrals inside the sums, (24) reads

$$\partial_t P(\mathbf{x}, t) \approx - \sum_{i=1}^N \frac{\partial}{\partial x_i} (A_i(\mathbf{x}) P(\mathbf{x}, t)) + \frac{1}{2} \sum_{i,j=1}^N \frac{\partial^2}{\partial x_i \partial x_j} (B_{ij}(\mathbf{x}) P(\mathbf{x}, t)), \tag{25}$$

where

$$A_i(\mathbf{x}) \equiv \int d\epsilon \epsilon_i W_\epsilon(\mathbf{x}), \tag{26}$$

and

$$B_{ij}(\mathbf{x}) \equiv \int d\epsilon \epsilon_i \epsilon_j W_\epsilon(\mathbf{x}). \tag{27}$$

What remains is to calculate (26) and (27), which we do in different ways depending on the value of the evolutionary reciprocal temperature  $\beta$ .

*The case of finite  $\beta$ .* First we consider finite  $\beta$ , in which case we can evaluate (26) and (27) using the results of Refs. 19,20 (making small changes in order to account for differences in proposal rates between those papers and ours).

Eq. (26) can be evaluated as follows (writing  $U(\mathbf{x}) = U$  for brevity):

$$A_i(\mathbf{x}) = \int d\epsilon p(\epsilon) \epsilon_i \min(1, e^{-\beta[U(\mathbf{x}+\epsilon)-U(\mathbf{x})]}) \tag{28}$$

$$\approx \int d\epsilon p(\epsilon) \epsilon_i \min(1, 1 - \beta \epsilon \cdot \nabla U) \tag{29}$$

$$= \int d\epsilon p(\epsilon) \epsilon_i \Theta(-\epsilon \cdot \nabla U) + \int d\epsilon p(\epsilon) \epsilon_i (1 - \beta \epsilon \cdot \nabla U) \Theta(\epsilon \cdot \nabla U) \tag{30}$$

$$= \int d\epsilon p(\epsilon) \epsilon_i - \beta \int d\epsilon p(\epsilon) \Theta(\epsilon \cdot \nabla U) \sum_j \epsilon_i \epsilon_j \partial_j U \tag{31}$$

$$= -\beta \int d\epsilon p(\epsilon) \Theta(\epsilon \cdot \nabla U) \sum_j \epsilon_i \epsilon_j \partial_j U \tag{32}$$

$$= -\frac{\beta}{2} \int d\epsilon p(\epsilon) \sum_j \epsilon_i \epsilon_j \partial_j U \tag{33}$$

$$= -\frac{\beta \sigma^2}{2} \partial_i U. \tag{34}$$

In these expressions  $\Theta(x) = 1$  if  $x \geq 0$  and is zero otherwise. In going from (28) to (29) we have assumed that  $\beta \epsilon \cdot \nabla U(\mathbf{x})$  is small. This condition cannot be met for  $\beta = \infty$ ; that case is treated later. In going from (30) to (31) we have used the result  $\Theta(x) + \Theta(-x) = 1$ ; the first integral in (31) then vanishes by symmetry. The second integral, shown in (32), can be turned into (33) using the symmetry arguments given in Ref. 20, which we motivate as follows. Upon a change of sign of the integration variables,  $\epsilon \rightarrow -\epsilon$ , the value of the integral in (32) is unchanged and it is brought to a form that is identical except for a change of sign of the argument of the  $\Theta$  function. Adding the two forms of the integral removes the  $\Theta$  functions, giving the form shown in (33), and dividing by 2 restores the value of the original integral. (33) can be evaluated using standard results of Gaussian integrals.

Eq. (27) can be evaluated in a similar way:

$$B_{ij}(\mathbf{x}) = \int d\epsilon p(\epsilon) \epsilon_i \epsilon_j \min(1, e^{-\beta[U(\mathbf{x}+\epsilon)-U(\mathbf{x})]}) \tag{35}$$

$$\approx \int d\epsilon p(\epsilon) \epsilon_i \epsilon_j \min(1, 1 - \beta \epsilon \cdot \nabla U) \tag{36}$$

$$= \int d\epsilon p(\epsilon) \epsilon_i \epsilon_j \Theta(-\epsilon \cdot \nabla U) + \int d\epsilon p(\epsilon) \epsilon_i \epsilon_j (1 - \beta \epsilon \cdot \nabla U) \Theta(\epsilon \cdot \nabla U) \tag{37}$$

$$\approx \int d\epsilon p(\epsilon) \epsilon_i \epsilon_j \tag{38}$$

$$= \sigma^2 \delta_{ij}. \tag{39}$$

The  $\approx$  sign in (38) indicates that we have omitted terms of order  $\sigma^3$ .

Inserting (34) and (39) into (25) gives us, to second order in  $\sigma$ , the Fokker-Planck equation

$$\frac{\partial P(\mathbf{x}, t)}{\partial t} \approx - \sum_{i=1}^N \frac{\partial}{\partial x_i} \left( -\frac{\beta \sigma^2}{2} \frac{\partial U(\mathbf{x})}{\partial x_i} P(\mathbf{x}, t) \right) + \frac{1}{2} \sum_{i=1}^N \frac{\partial^2}{\partial x_i^2} (\sigma^2 P(\mathbf{x}, t)). \tag{40}$$

This equation is equivalent (the diffusion term is independent of  $\mathbf{x}$ , and so the choice of stochastic calculus is unimportant) to the  $N$  Langevin equations<sup>44,45</sup>

$$\frac{dx_i}{dt} = -\frac{\beta \sigma^2}{2} \frac{\partial U(\mathbf{x})}{\partial x_i} + \xi_i(t) \quad \forall i, \tag{41}$$

where  $\xi$  is a Gaussian white noise with zero mean and variance  $\sigma^2$ :

$$\langle \xi_i(t) \rangle = 0, \quad \langle \xi_i(t) \xi_j(t') \rangle = \sigma^2 \delta_{ij} \delta(t - t'). \tag{42}$$

Eq. (41) describes an evolution of the neural-network parameters  $x_i$  that is equivalent to gradient descent with learning rate  $\alpha = \beta \sigma^2 / 2$ , plus Gaussian white noise. Averaging over independent stochastic trajectories of the learning process (starting from identical initial conditions) gives

$$\frac{d\langle x_i \rangle}{dt} = -\frac{\beta \sigma^2}{2} \frac{\partial U(\mathbf{x})}{\partial x_i}, \tag{43}$$

which is equivalent to simple gradient descent on the loss function.

*The case  $\beta = \infty$ .* When  $\beta = \infty$  we only accept mutations that do not increase the loss function. To treat this case we return to (26) and take the limit  $\beta \rightarrow \infty$ :

$$A_i(\mathbf{x}) = \int d\epsilon p(\epsilon) \epsilon_i \Theta(-\epsilon \cdot \nabla U(\mathbf{x})). \tag{44}$$

We can make progress by introducing the integral form of the  $\Theta$  function (see e.g.,



ref. 46),

$$\begin{aligned} \Theta(-\epsilon \cdot \nabla U(\mathbf{x})) &= \int_{-\infty}^0 dz \delta(z - \epsilon \cdot \nabla U(\mathbf{x})) \\ &= \int_{-\infty}^0 \frac{dz}{2\pi} \int_{-\infty}^{\infty} d\omega e^{i\omega(z - \epsilon \cdot \nabla U(\mathbf{x}))}. \end{aligned} \tag{45}$$

Then (44) reads

$$A_i(\mathbf{x}) = \int_{-\infty}^0 \frac{dz}{2\pi} \int_{-\infty}^{\infty} d\omega e^{i\omega z} G_i^{(1)} \prod_{j \neq i} G_j^{(0)}, \tag{46}$$

where the symbols

$$G_j^{(n)} \equiv \frac{1}{\sqrt{2\pi\sigma^2}} \int_{-\infty}^{\infty} d\epsilon_j \epsilon_j^n e^{-\epsilon_j^2/(2\sigma^2) - i\omega\epsilon_j \partial_j U} \tag{47}$$

are standard Gaussian integrals. Upon evaluating them as

$$G_j^{(0)} = e^{-\frac{1}{2}\omega^2 \sigma^2 (\partial_j U)^2} \tag{48}$$

and

$$G_i^{(1)} = -i\omega\sigma^2 (\partial_i U) e^{-\frac{1}{2}\omega^2 \sigma^2 (\partial_i U)^2}, \tag{49}$$

(46) reads

$$A_i(\mathbf{x}) = - \int_{-\infty}^0 \frac{dz}{2\pi} \int_{-\infty}^{\infty} d\omega i\sigma^2 \omega (\partial_i U) e^{-\frac{1}{2}\omega^2 \sigma^2 |\nabla U|^2 + i\omega z} \tag{50}$$

$$\begin{aligned} &= \int_{-\infty}^0 \frac{dz}{2\pi} \sqrt{2\pi} \frac{\partial_i U}{\sigma |\nabla U|^3} z e^{-z^2/(2\sigma^2 |\nabla U|^2)} \\ &= - \frac{\sigma}{\sqrt{2\pi}} \frac{\partial_i U(\mathbf{x})}{|\nabla U(\mathbf{x})|}. \end{aligned} \tag{51}$$

The form (51) is similar to (34) in that it involves the derivative of the loss function  $U(\mathbf{x})$  with respect to  $x_i$ , but contains an additional normalization term,  $|\nabla U|$ . This term is sometimes introduced as a form of regularization in gradient-based methods<sup>25</sup>. Here, the form emerges naturally from the acceptance criterion that sees any move accepted if the move does not increase the loss function: as a result, the length of the step taken does not depend strongly on the size of the gradient.

In the limit  $\beta \rightarrow \infty$ , (27) reads

$$B_{ij}(\mathbf{x}) = \int d\epsilon p(\epsilon) \epsilon_i \epsilon_j \Theta(-\epsilon \cdot \nabla U(\mathbf{x})) \tag{52}$$

$$= \frac{1}{2} \int d\epsilon p(\epsilon) \epsilon_i \epsilon_j \tag{53}$$

$$= \frac{1}{2} \sigma^2 \delta_{ij}, \tag{54}$$

upon applying the symmetry arguments used to evaluate (32). Eq. (54) is half the value of the corresponding term for the case  $\beta \neq \infty$ , Eq. (39), because one term corresponds to Brownian motion in unrestricted space, the other to Brownian motion on a half-space.

Inserting (51) and (54) into (25) gives a Fokker–Planck equation equivalent to the  $N$  Langevin equations

$$\frac{dx_i}{dt} = - \frac{\sigma}{\sqrt{2\pi}} \frac{1}{|\nabla U(\mathbf{x})|} \frac{\partial U(\mathbf{x})}{\partial x_i} + \eta_i(t) \quad \forall i, \tag{55}$$

where  $\eta$  is a Gaussian white noise with zero mean and variance  $\sigma^2/2$ :

$$\langle \eta_i(t) \rangle = 0, \quad \langle \eta_i(t) \eta_j(t') \rangle = \frac{1}{2} \sigma^2 \delta_{ij} \delta(t - t'). \tag{56}$$

As an aside, we briefly consider the case of non-isotropic mutations, for which the Gaussian random update applied to parameter  $i$  has its own variance  $\sigma_i^2$ , i.e.,  $\epsilon_i \sim \mathcal{N}(0, \sigma_i^2)$  in step 2 of the procedure described in section “Neuroevolution”. In this case the derivation above is modified to have  $\sigma$  replaced by  $\sigma_i$  in (21). In the case of finite  $\beta$  the equations (41) and (42) retain their form with the replacement  $\sigma \rightarrow \sigma_i$ . In the case of infinite  $\beta$ , (56) retains its form with the replacement  $\sigma \rightarrow \sigma_i$  and (55) reads

$$\frac{dx_i}{dt} = - \frac{\sigma_i}{\sqrt{2\pi}} \frac{1}{|\tilde{\nabla} U(\mathbf{x})|} \tilde{\partial}_i U(\mathbf{x}) + \eta_i(t), \tag{57}$$

with  $\tilde{\partial}_i \equiv \sigma_i \frac{\partial}{\partial x_i}$  and  $|\tilde{\nabla} U| \equiv \sqrt{\sum_{j=1}^N (\tilde{\partial}_j U)^2}$ . Non-isotropic mutations are used in covariance matrix adaptation strategies<sup>47</sup>. Those schemes also evolve the step size parameter dynamically, and to model this more general case one must make  $\sigma$  a dynamical variable of the master equation, with update rules appropriate to the algorithm of interest.

## Data availability

Data can be generated using the source code at <https://github.com/reproducible-science/MC-GD-correspondence>.

## Code availability

Source code is available at <https://github.com/reproducible-science/MC-GD-correspondence>.

Received: 26 April 2021; Accepted: 4 October 2021;

Published online: 02 November 2021

## References

- Linnainmaa, S. Taylor expansion of the accumulated rounding error. *BIT Numer. Math.* **16**, 146–160 (1976).
- Werbos, P. J. Applications of advances in nonlinear sensitivity analysis. In: *System Modeling and Optimization. Lecture Notes in Control and Information Sciences* (eds Drenick R. F. & Kozin F.), vol 38. 762–770 (Springer, Berlin, Heidelberg, 1982). <https://doi.org/10.1007/BFb0006203>.
- Rumelhart, D. E., Durbin, R., Golden, R. & Chauvin, Y. In *Backpropagation: Theory, Architectures and Applications*, (eds Chauvin Y. & Rumelhart D. E.) 1–34 (Hillsdale: NJ. Lawrence Erlbaum, 1995).
- Rumelhart, D. E., Hinton, G. E. & Williams, R. J. Learning representations by back-propagating errors. *Nature* **323**, 533–536 (1986).
- Hecht-Nielsen, R. In *Neural Networks for Perception* (eds Wechsler H.) 65–93 (Elsevier, 1992).
- LeCun, Y. et al. Backpropagation applied to handwritten zip code recognition. *Neural Comput.* **1**, 541–551 (1989).
- Chauvin, Y. & Rumelhart, D. E. *Backpropagation: Theory, Architectures, and Applications* (Psychology Press, 1995).
- Schmidhuber, J. Deep learning in neural networks: an overview. *Neural Networks* **61**, 85–117 (2015).
- 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).
- Hastings, W. K. Monte Carlo sampling methods using markov chains and their applications. *Biometrika* **57**, 97–109 (1970).
- Holland, J. H. Genetic algorithms. *Sci. Am.* **267**, 66–73 (1992).
- Fogel, D. B. & Stayton, L. C. On the effectiveness of crossover in simulated evolutionary optimization. *BioSystems* **32**, 171–182 (1994).
- Montana, D. J. & Davis, L. Training feedforward neural networks using genetic algorithms. In *IJCAI*, Vol. 89 762–767 (1989).
- Mnih, V. et al. Playing Atari with deep reinforcement learning, Preprint at <https://arxiv.org/abs/1312.5602> (2013).
- Morse, G. & Stanley, K. O., Simple evolutionary optimization can rival stochastic gradient descent in neural networks. In *Proceedings of the Genetic and Evolutionary Computation Conference 2016* 477–484 (2016).
- Salimans, T., Ho, J., Chen, X., Sidor, S. & Sutskever, I. Evolution strategies as a scalable alternative to reinforcement learning. Preprint at <https://arxiv.org/abs/1703.03864> (2017).
- Zhang, X., Clune, J. & Stanley, K. O. On the relationship between the OpenAI evolution strategy and stochastic gradient descent. Preprint at <https://arxiv.org/abs/1712.06564> (2017).
- Sutton, R.S. & Barto, A.G. *Reinforcement Learning: An Introduction* (MIT press, 2018).
- 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).
- 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).
- Floreano, D., Dürr, P. & Mattiussi, C. Neuroevolution: from architectures to learning. *Evolution. Intell.* **1**, 47–62 (2008).
- Such, F. P. et al. Deep neuroevolution: genetic algorithms are a competitive alternative for training deep neural networks for reinforcement learning. Preprint at <https://arxiv.org/abs/1712.06567> (2017).
- Whitelam, S. & Tamblyn, I. Learning to grow: control of material self-assembly using evolutionary reinforcement learning. *Phys. Rev. E* **101**, 052604 (2020).
- Frenkel, D. & Smit, B. *Understanding Molecular Simulation: from Algorithms to Applications*, Vol. 1 (Academic Press, 2001).
- Pascanu, R., Mikolov, T. & Bengio, Y. On the difficulty of training recurrent neural networks. In *International Conference on Machine Learning* 1310–1318 (PMLR, 2013).

26. Stanley, K. O. & Miikkulainen, R. Evolving neural networks through augmenting topologies. *Evolution. Comput.* **10**, 99–127 (2002).
27. Stanley, K. O., Clune, J., Lehman, J. & Miikkulainen, R. Designing neural networks through neuroevolution. *Nat. Machine Intell.* **1**, 24–35 (2019).
28. Swendsen, R. H. & Wang, J.-S. Nonuniversal critical dynamics in monte carlo simulations. *Phys. Rev. Lett.* **58**, 86 (1987).
29. Wolff, U. Collective monte carlo updating for spin systems. *Phys. Rev. Lett.* **62**, 361 (1989).
30. Liu, J. & Luijten, E. Rejection-free geometric cluster algorithm for complex fluids. *Phys. Rev. Lett.* **92**, 035504 (2004).
31. Whitelam, S. Approximating the dynamical evolution of systems of strongly interacting overdamped particles. *Mol. Simul.* **37**, 606–612 (2011).
32. Cybenko, G. Approximation by superpositions of a sigmoidal function. *Math. Control Signal. Syst.* **2**, 303–314 (1989).
33. Raisbeck, J. C., Allen, M., Weissleder, R., Im, H. & Lee, H., Evolution strategies converges to finite differences. Preprint at <https://arxiv.org/abs/2001.01684> (2019).
34. Staines, J. & Barber, D. Variational optimization. Preprint at <https://arxiv.org/abs/1212.4507> (2012).
35. Maheswaranathan, N., Metz, L., Tucker, G., Choi, D. & Sohl-Dickstein, J. Guided evolutionary strategies: Augmenting random search with surrogate gradients. In *International Conference on Machine Learning*. 4264–4273 (PMLR, 2019).
36. Whitelam, S. & Geissler, P. L. Avoiding unphysical kinetic traps in Monte Carlo simulations of strongly attractive particles. *J. Chem. Phys.* **127**, 154101 (2007).
37. Wilber, A. W. et al. Reversible self-assembly of patchy particles into monodisperse icosahedral clusters. *J. Chem. Phys.* **127**, 08B618 (2007).
38. Berthier, L. Revisiting the slow dynamics of a silica melt using Monte Carlo simulations. *Phys. Rev. E* **76**, 011507 (2007).
39. Sanz, E. & Marenduzzo, D. Dynamic Monte Carlo versus Brownian dynamics: a comparison for self-diffusion and crystallization in colloidal fluids. *J. Chem. Phys.* **132**, 194102 (2010).
40. Liu, X., Crocker, J. C. & Sinno, T. Coarse-grained Monte Carlo simulations of non-equilibrium systems. *J. Chem. Phys.* **138**, 244111 (2013).
41. Rovigatti, L., Russo, J. & Romano, F. How to simulate patchy particles. *Eur. Phys. J. E* **41**, 59 (2018).
42. Engel, A. & Van den Broeck, C., *Statistical Mechanics of Learning* (Cambridge University Press, 2001).
43. Bahri, Y. et al. Statistical mechanics of deep learning. *Ann. Rev. Condens Matter Phys.* **11**, 501–528 (2020).
44. Risken, H. Fokker-Planck Equation. In *The Fokker-Planck Equation*. Springer Series in Synergetics vol 18 (Springer, Berlin, Heidelberg, 1996). [https://doi.org/10.1007/978-3-642-61544-3\\_4](https://doi.org/10.1007/978-3-642-61544-3_4).
45. Van Kampen, N. G., *Stochastic Processes in Physics and Chemistry*, Vol. 1 (Elsevier, 1992).
46. Sinai, Y. B., <https://yohai.github.io/post/half-gaussian/> (2019).
47. Hansen, N. in *Towards A New Evolutionary Computation*, (eds Lozano J. A., Larrañaga P., Inza I. & Bengoetxea E.) 75–102 (Springer, 2006).

## Acknowledgements

This work was performed as part of a user project at the Molecular Foundry, Lawrence Berkeley National Laboratory, supported by the Office of Science, Office of Basic Energy Sciences, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231. This work used resources of the National Energy Research Scientific Computing Center (NERSC), a U.S. Department of Energy Office of Science User Facility operated under Contract No. DE-AC02-05CH11231. I.T. acknowledges funding from the National Science and Engineering Council of Canada.

## Author contributions

S.W. and I.T. initiated the study. S.W. did the analytic work and the shallow-net simulations, and V.S. did the deep-net simulations. S.-W.P. provided assistance with numerical simulations. All authors discussed the work and helped write the paper.

## 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/s41467-021-26568-2>.

**Correspondence** and requests for materials should be addressed to Stephen Whitelam or Isaac Tamblyn.

**Peer review information** *Nature Communications* thanks the anonymous reviewer(s) for their contribution to the peer review of this work.

**Reprints and permission information** is available at <http://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 license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license 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 license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2021