

Leveraging chaotic transients in the training of artificial neural networks

Pedro Jiménez-González , Miguel C. Soriano , and Lucas Lacasa ^{*}
*Institute for Cross-Disciplinary Physics and Complex Systems (IFISC, CSIC-UIB),
 Campus UIB, 07122 Palma de Mallorca, Spain*



(Received 3 June 2025; accepted 6 March 2026; published 18 May 2026)

Traditional algorithms to optimize artificial neural networks when confronted with a supervised learning task are usually exploitation-type relaxational dynamics such as gradient descent (GD). Here, we explore the dynamics of the neural network trajectory along training for unconventionally large learning rates. We show that for a region of values of the learning rate, the GD optimization shifts away from purely exploitationlike algorithm into a regime of exploration-exploitation balance, as the neural network is still capable of learning but the trajectory shows sensitive dependence on initial conditions—as characterized by positive network maximum Lyapunov exponent. Interestingly, the characteristic training time required to reach an acceptable accuracy in the test set reaches a minimum precisely in such learning rate region, further suggesting that one can accelerate the training of artificial neural networks by locating at the onset of chaos. Our results—initially illustrated for the MNIST classification task—qualitatively hold for a range of supervised learning tasks, learning architectures (including both shallow and deep multilayer perceptrons and convolutional neural networks) and other hyperparameters (different activation functions and weight regularization), and showcase the emergent, constructive role of transient chaotic dynamics in the training of artificial neural networks.

DOI: [10.1103/t5p9-kv5w](https://doi.org/10.1103/t5p9-kv5w)

The multilayer perceptron (MLP) is an archetypal method in supervised machine learning [1,2], used to infer (regress or classify) complex input-output representations $x \rightarrow y$, where $x \in \mathbb{R}^m$ and y is usually another vector for regression tasks or an element from a discrete set in classification tasks. Accordingly, an MLP is mathematically an overparametrized nonlinear function $\mathcal{F}(x; \Omega)$, where $\Omega = \{w_k\}$ is the set of trainable parameters. Often represented as an (artificial neural) network [1,2], the MLP is visualized as a mathematical graph composed of various stacked layers of interconnected nodes, where the first layer represents the input vector $x \in \mathbb{R}^m$. The edges connecting nodes in adjacent layers represent (parametric) affine transformations of the vector elements of one layer to the next one (the coefficients of such transformations, usually called the edge weights, belong to the parameter set Ω), and the nodes integrate the incident linear compositions in the edges nonlinearly via what is called an activation function. The information thus flows from the input layer x onward until reaching the final, output layer whose nodes represent the elements of the output $y = \mathcal{F}(x; \Omega)$.

Optimizing the set of parameters Ω so that the mismatch between $\mathcal{F}(x; \Omega)$ and y —the so-called loss function $\mathcal{L}(x; \Omega)$ —is minimized is called the *training* process. In practice, training the MLP is usually done iteratively, with gradient descent (GD) in parameter space being the gold standard,

where

$$\omega_k(t+1) = \omega_k(t) - \eta \partial_{\omega_k} \mathcal{L}(x; \Omega(t)), \quad \forall \omega_k \in \Omega, \quad (1)$$

where η is the so-called learning rate and $\partial_x := \partial/\partial x$.

Observe that, along training, the MLP is effectively performing a trajectory in the (graph) space spanned by the set of parameters Ω . We contend that reinterpreting the training process as a (high-dimensional, latent) graph dynamics [3] allows to inquire the inner workings of learning algorithms following physics-inspired and complexity-based epistemics [4–6]. Instead of tracking the scalar projection of the full dynamics given by the time evolution of the loss function $\mathcal{L}(x; \Omega(t))$, let us consider the actual network trajectory [7]. We argue that this change of focus—which amounts to tracking full graph trajectories of the form $(\Omega(0), \Omega(1), \Omega(2), \dots)$, where the set $\Omega(t) = \{\omega_1(t), \dots, \omega_m(t)\}$ incorporates the updated values of the so-called weights and biases of the neural architecture at training epoch t —can provide valuable insights if at the same time we leverage principles and tools from network science [8–14] and dynamical systems [15,16].

Along training, it is often assumed that in a typical neural network trajectory there is an arrow of time induced by the relaxational nature of Eq. (1), which is inherited in graph space: We often (perhaps misleadingly) assume that the dynamics always reaches equilibrium—some minimum of $\mathcal{L}(x; \Omega)$ —an intuition that parallels the behavior of Langevin dynamics—the continuous and stochastic counterpart to gradient descent—which, under broad conditions, relaxes to an equilibrium distribution (typically a Boltzmann measure) over the energy landscape. Accordingly, GD [17] is often seen as an *exploitation* search algorithm that iteratively performs small improvements of an initial solution, so that some fitness

^{*}Contact author: lucas@ifisc.uib-csic.es

function (here the loss function $\mathcal{L}$) is gradually (e.g., monotonically) decreased. This intuition is, at the end of the day, intimately related to the convergence properties of the GD scheme. Now, discrete-time GD convergence is not always fulfilled, specially for large enough learning rates, when GD can display more exotic behavior [18,19] (see Figs. 6 and 7, where we present a toy model of a gradient descent scheme displaying a period-doubling bifurcation cascade onto chaotic behavior, for a nonconvex loss function and a sufficiently large learning rate). At the same time, convergence to local minima is an asymptotic behavior, and interesting dynamics often emerge in transient times [20]. Altogether, it is interesting to consider whether, led by qualitative dynamical changes in the behavior of GD-type maps, not just the loss function but the whole MLP network trajectories *transition* from following a pure exploitation strategy to other search strategies—such as *exploration* [21]—when the learning rate is large enough. Incidentally, note that a similar balance of strategies [22] is known to yield optimal searching behaviors in animal foraging [23–29], transport [30], and a range of decision-making contexts [31,32], and is an explicit cornerstone of the reinforcement learning paradigm [33].

Our contention in this work is that, indeed, such transition takes place, and the emergence of explorationlike dynamics, induced by the onset of sensitive dependence on initial conditions—the hallmark of chaotic dynamics—in the dynamics of Eq. (1) yields a sweet spot where the balance of exploitation and exploration dynamics has a constructive, beneficial role for neural network learning. Let us clarify at this point that in this work we are neither considering intrinsically dynamic neural networks whose neurons show chaotic behavior [34–39] nor considering the stability of the input-to-output trajectory in deep networks [40]; instead, we are considering the training (optimization) dynamics in a neural network whose output is not dynamic. Moreover, we argue that (1) for an often unexplored region of large-but-not-too-large values of η , the map is optimally interpolating exploitation and exploration search strategies, and (2) that the average training time required for the network to learn a representation capable of generalizing—i.e., capable to reach a good performance on the test set— $\text{Test} = \{x_i, y_i\}_{i=1}^{N_{\text{test}}}$ is minimized in such *sweet spot*. Interestingly, this coincides with a similar type of optimality emerging when the loss function’s Hessian asymptotically evolves over training toward its so-called edge of stability, i.e., when its largest eigenvalue approaches $2/\eta$ [41]. In a nutshell, we argue that the transition to an exploitation-exploration balance is achieved by leveraging the onset of transient chaotic mixing [20], and at such transition the system efficiently minimizes training time, resulting in a possible demonstration of Langton’s hypothesis [42].

To actually demonstrate our hypothesis, we initially consider a vanilla classification task: MNIST image classification [43]. For parsimony, we initially choose a shallow MLP (one hidden layer with 64 neurons) and a tanh activation function (the results for other supervised learning tasks, activation functions deeper architectures, and other architectures with inductive biases are discussed in Figs. 8–13). After train/test split, this MLP is trained on a set $\text{Train} = \{(x_i, y_i)\}_{i=1}^N$ of $N = 6 \times 10^4$ labeled handwritten images (y_i is the label of




FIG. 1. Training loss trajectory of a neural network on the MNIST dataset for three different learning rates: $\eta = 0.01$, $\eta = 7.5$, and $\eta = 20$.

the i th image), and we use a cross-entropy loss function

$$\mathcal{L}(x; \Omega) = -\frac{1}{N} \sum_{i=1}^N y_i \ln \mathcal{F}(x_i; \Omega). \quad (2)$$

Training takes place by using Eq. (1). For the sake of simplicity, no regularization is initially added to the loss function (results with L_2 weight regularization are discussed in Fig. 11). Additionally, we use traditional GD schemes (no minibatch or stochastic gradient descent (SGD)) and discard dropout, to remove any source of stochasticity to the network dynamics [18]. Finally, to assess the performance for different learning rates, a constant learning rate η is fixed throughout training.

Figure 1 plots the time evolution of the (training) loss function $\mathcal{L}(x; \Omega(t)) \equiv \mathcal{L}(t)$ for three different learning rates η . Interestingly, such a loss is only monotonically decreasing for the standard range of small values of the learning rate. For larger values of η , other dynamical behaviors are found: converging loss functions with nonmonotonic, irregular transients, and other dynamical attractors for extremely large η , where the MLP does not seem to make any useful learning.

To better characterize what particular change in dynamical behaviors in the network evolution is inducing these projections in the loss function, we now resort to recently introduced graph-theoretical extensions of the maximum Lyapunov exponent [44,45], designed to estimate sensitivity to initial conditions in network trajectories. The procedure consists of three steps: (1) For a fixed learning rate, we define a set of q different network initializations $\mathcal{S} = \{\Omega(0)\}$. (2) Around each concrete network initialization $\Omega(0)$, we build an ϵ -ball formed by a set of M small network “perturbations” of $\mathcal{B} = \{\Omega(0)^{(j)}\}_{j=1}^M$. To do that, we perturb each trainable parameter $w_k \in \Omega$ with uniform noise $w'_k = w_k + \xi$, $\xi \sim \text{UNIFORM}(-\epsilon, \epsilon)$. Defining the distance between two network initializations Ω, Ω' as the L_1 norm $d(\Omega, \Omega') = \sum_{\omega_k \in \Omega} |\omega_k - \omega'_k|$, it follows that each perturbed network $\Omega(0)^{(j)}$ is at most at distance $\text{CARD}(\Omega)\epsilon$ from $\Omega(0)$. (3) Then, following Ref. [44], we measure the expansion rate of closeby network trajectories by adequately averaging the divergence of the




FIG. 2. Semilog plot of the evolution (along training) of the network distance $d(t)$ for pairs of network trajectories with closely initialization $\Omega(0)$, as a function of the number of epochs t , for a shallow MLP with $\tanh()$ activation function trained on MNIST with a large learning rate. $d(t)$ displays a stylized exponential expansion followed by saturation. The slope of the exponential phase corresponds to the local network Lyapunov exponent Λ and is indicative of chaotic mixing.

M elements inside each ϵ -ball throughout the action of the training dynamics in Eq. (1):

$$\Lambda_{\Omega(0)} = \frac{1}{\tau} \ln \frac{M^{-1} \sum_{j=1}^M d_j(\tau)}{M^{-1} \sum_{j=1}^M d_j(0)}, \quad (3)$$

where $d_j(\tau) \equiv d(\Omega(\tau), \Omega(\tau)^{(j)})$. $\Lambda_{\Omega(0)}$ is indeed the network version of a finite, local Lyapunov exponent [44], where τ is the characteristic time required for the elements inside the ϵ -ball centered at $\Omega(0)$ to diverge to distances on the order of the phase space diameter. For an illustration, in Fig. 2 we depict the distance $d(t)$ between $M = 5$ network initializations close to an initial condition $\Omega(0)$ ($\epsilon = 10^{-8}$), for a large learning rate $\eta = 10$. We observe a clear exponential phase up to $\tau \approx 30$ epochs, the slope denoting the finite local network Lyapunov exponent $\Lambda_{\Omega(0)}$.

The network's maximum Lyapunov exponent averages the local exponents over different network initializations $\lambda_{\text{nMLE}} = \langle \Lambda_{\Omega(0)} \rangle_{\Omega(0) \in \mathcal{S}}$. For illustration, this is displayed for $\text{CARD}(\mathcal{S}) = 50$ different initial conditions in Fig. 3. The inset of that panel depicts the histogram of local exponents, whose average gives $\lambda_{\text{nMLE}} \approx 0.68$, i.e., the system shows sensitivity to initial conditions. λ_{nMLE} can thus be seen as an order parameter distinguishing two phases: A phase where the search strategy induced by the GD map [Eq. (1)] is an exploitation one and $\lambda_{\text{nMLE}} \leq 0$, and a phase where the search strategy is of an exploration type with sensitive dependence on initial conditions and $\lambda_{\text{nMLE}} > 0$. The transition between both phases interpolates both search strategies. As a complementary metric that will later assist us in our empirical analysis of the network training dynamics, we also define ρ as the percentage of MLP initializations of the set $\mathcal{S}$ for which the ϵ -ball expansion can be well approximated by an exponential function with high statistical significance ($R^2 > 0.9$, $\Lambda > 0.05$); i.e., this is the percentage of local network Lyapunov exponents that are positive.




FIG. 3. Same as Fig. 2, but for many different ϵ -balls centered at different initial conditions Ω . Each initial condition leads in principle to a different local network Lyapunov exponent $\Lambda(\Omega)$. In the inset, we display the histogram of local network Lyapunov exponents. The average of this distribution is the estimation of the network MLE, $\lambda_{\text{nMLE}} \approx 0.68$.

In Fig. 4, we report λ_{nMLE} as a function of the gradient descent's learning rate η . Results indicate that there is a clear transition between exploitation to exploration search, where the exploitation-exploration interpolation hovers in the range $\eta \in [1, 10]$. Figure 5 displays ρ as a function of the learning rate. This metric confirms a transition between a phase where no initial conditions display exponential expansion, to a phase where virtually all regions of the phase space display chaotic transients.

Finally, to assess the MLP's learning and training efficiency in the context of the abovementioned phenomenology, Fig. 5 also depicts the average number of epochs $\langle \tau \rangle$ needed for the MLP to reach an average classification accuracy of at least 0.9 in the test set (this is averaged over all different




FIG. 4. Estimation of the network maximum Lyapunov exponent λ_{nMLE} for MLP trajectories as a function of the learning rate η . Error bars denote $\pm$ one standard deviation of the population of finite local network Lyapunov exponents $\{\Lambda(\Omega)\}$. The onset of sensitivity to initial conditions $\lambda_{\text{nMLE}} > 0$ marks the change from a purely exploitation-type optimization to an exploration/exploitation type.




FIG. 5. Blue diamonds depict ρ , the percentage of MLP initializations Ω leading to training trajectories with positive local Lyapunov exponent $\Lambda(\Omega) > 0$ as a function of the learning rate η . In the same figure, we also plot (red dots) the average training time $\langle \tau \rangle$ (in number of gradient descent epochs) needed to reach an accuracy of 90% or larger in the test set. Training is found to be maximally efficient close to the onset of fully developed sensitivity to initial conditions [$\Lambda(\Omega) > 0 \forall \Omega$].

MLP initializations in $\mathcal{S}$), as a function of the learning rate η . We find that $\langle \tau \rangle$ displays a nonmonotonic shape, and indeed reaches a minimum in the exploitation-exploration interpolation region, at the learning rate $\eta \approx 7.5$, precisely marking the onset of fully developed sensitivity to initial conditions $\rho \approx 100\%$.

As advanced, this type of optimality in the training dynamics has been recently observed when the largest eigenvalue $\sigma_{\max}$ of the loss function's Hessian asymptotically converges to $2/\eta$, the so-called edge of stability [41]. In Fig. 14, we show how the time series of $\sigma_{\max}(t)$ over training indeed approaches $2/\eta$ when the learning rate hovers around the values for which $\langle \tau \rangle$ is minimized and suggests that asymptotically self-organizing around the Hessian's edge of stability is precursed by a chaotic transient.

Our results are robust against changes in the classification task (see Figs. 8 and 13 for further results on the Iris and CIFAR-10 classification, respectively), the type of activation function (see Fig. 9 for further results with sigmoid and rectified linear unit (ReLU) activation functions), the depth of the MLP (see Fig. 10 for a comparison of shallow vs deep network), the inclusion of weight regularization (see Fig. 11), and also hold for other neural network architectures with inductive biases (see Fig. 12 for analysis on a convolutional neural network) and overall highlight the relevance of leveraging transient chaotic mixing in the training of neural networks. In summary, we have found that, as the learning rate increases, the training dynamics transition from a regular, purely exploitation-type dynamics to a chaotic, purely exploration-type dynamics. The transition between both types is rather sharp and occurs in a region that trades off exploitation and exploration by the onset of a mechanism of chaotic transient—a chaotic fingerprint emerging in the first few dozen epochs of the training dynamics. This mechanism underpins an efficient search of graph space, eventually leading to faster learning. Evidence suggests that such early-stage, chaotic transient is

precursing later Hessian trajectories to eventually converge toward their edge of stability.

From a conceptual point of view, our findings suggest a demonstration of Langton's edge of chaos hypothesis [42], and a confirmation of Verschure's seminal idea of using chaos as a fast-search mechanism [46]. In hindsight, this is an example where instabilities emerging at the core of numerical schemes—often seen as a nuisance when it comes to, e.g., integrating differential equations or root finding—have in the case of searching for neural network representations an unexpected benefit. This benefit occurs because in neural network optimization one uses GD as a *search* algorithm rather than simply as a local minimizer, and as a consequence, exploitation benefits from exploration when both mechanisms are at play. Aside from this theoretical insight, forcing the training dynamics' learning rate to be located at the sweet spot could also yield practical benefits. This can be achieved, e.g., by using the bisection method to find the sweet spot, by iteratively refining an initial learning rate range $[\eta_{\min}, \eta_{\max}]$, such that $\rho(\eta_{\min}) \approx 0$ and $\rho(\eta_{\max}) \approx 100$ as a preprocessing step, before actually training the system.

Finally, further work is needed to clarify the modulating effect of stochastic sources in the network trajectories, such as using stochastic or minibatch as opposed to full-batch gradient descent [19,47] or adding dropout, and to address more sophisticated learning rate schedules. Likewise, investigating whether such instability-induced advantage holds in other optimization schemes, or whether other parameters such as the batch size [47] could play the role of a control parameter, and exploring if such phenomenology can be retrieved from low-dimensional projections of network trajectories [48,49], are all interesting open problems.

Acknowledgments. The authors thank K. Danovski for input in preliminary stages of this project, M. Matías for insightful discussions, A. Fernández Amil for pointing us to Ref. [46] after our work was submitted as a preprint, and referees for stimulating discussions. P.J. acknowledges funding from Maria de Maeztu (MdM) Seal of Excellence (CEX2021-001164-M) via the FPI programme (Grant No. PRE2022-104148), funded by the MICIU/AEI/10.13039/501100011033. L.L. acknowledges partial support from project CSxAI (PID2024-157526NB-I00) funded by MICIU/AEI/10.13039/501100011033/FEDER, UE. M.C.S. and L.L. acknowledge partial support from a Maria de Maeztu grant (CEX2021-001164-M) funded by the MICIU/AEI/10.13039/501100011033 and from the European Commission Chips Joint Undertaking Project No. 101194363 (NEHIL).

Data availability. The code used in this work is publicly available in [50] and can be readily used or adapted to a range of classification tasks or network architectures. All data supporting the findings of this article were generated using this code.

Appendix A: Emergence of chaotic dynamics in toy gradient descent maps. Often, gradient descent optimization of neural networks weights generates monotonically decreasing loss function trajectories toward one of the minima, and common wisdom assumes that such monotonicity is also inherited in weight space. Here, we challenge this belief by presenting a




FIG. 6. Bifurcation diagram of the one-dimensional gradient map, displaying a period-doubling cascade to chaos.

toy model for the (gradient descent) optimization of neural networks, which display chaotic trajectories in weight space.

Let us consider a one-dimensional gradient dynamics map:

$$w_{t+1} = w_t - \eta \frac{dV(w)}{dw} \equiv f(w_t), \quad (\text{A1})$$

where the learning rate η is a map's parameter and $V(w)$ is a potential function that acts as our simplified loss function. The simplest case of a nonconvex loss with multiple minima is perhaps a double-well potential:

$$V(w) = \frac{1}{4}w^4 - \frac{1}{2}w^2, \quad (\text{A2})$$

which has two local minima for $w = -1, 1$. In this case, $f(w) = w[1 - \eta(w^2 - 1)]$ is thus a cubic map. We now show that even in such simple setting, the dynamics can become chaotic for sufficiently large learning rates.

While for low values of η the dynamics of Eq. (A1) converge toward one of the two minima, for larger values of this parameter, a bifurcation cascade takes place. In Fig. 6, we

plot the bifurcation diagram of the map, obtained numerically by iterating Eq. (A1) for 200 steps with a random initial conditions $w_0 \in N(0, 0.1)$ (after discarding a transient period) as a function of η . The figure shows a classical period-doubling bifurcation cascade toward a chaotic state with aperiodic dynamics, intertwined with some self-affine periodic windows. This is the first evidence.

Since this map is one-dimensional, it only has a single Lyapunov exponent λ , which can be approximated up to arbitrary precision with

$$\lambda = \lim_{t \rightarrow \infty} \frac{1}{t} \sum_{k=0}^{t-1} |\ln(f'(w_k))|. \quad (\text{A3})$$

In Fig. 7, we show λ as a function of the learning rate η , clearly finding regions for which $\lambda > 0$. All in all, this toy model illustrates that gradient descent scheme can yield chaotic dynamics (in the sense of Devaney) for nonconvex loss functions (such as the ones involved in neural network optimization)




FIG. 7. Lyapunov exponent of the gradient map in Eq. (A1) as a function of the learning rate η , showing regions with chaotic behavior.




FIG. 8. Lyapunov exponent (left column), percentage of valid exponents (ρ), and mean convergence time ($\langle \tau \rangle$) as functions of the learning rate (η) for each activation function. Results for the Iris task.

for sufficiently large learning rates. Since realistic (deep) neural network optimization just involves substantially more nonconvex loss functions of substantially higher-dimensional weight vectors, it is reasonable to expect that moving toward this more complex scenario makes the onset of chaotic behavior for large learning rates a generic phenomenon.

Appendix B: Onset of sensitivity to initial conditions and training efficiency: supplementary results. Below, we present the results of the key metrics (network maximum Lyapunov exponent λ_{nMLE} , percentage of positive finite

network Lyapunov exponents ρ , and average training time $\langle \tau \rangle$) for the additional set of configurations and tasks (Iris vs MNIST vs CIFAR-10, different activation functions, shallow vs deep architecture, effect of regularization, and Convolutional Neural Networks (CNNs))). These results provide additional evidence for the robustness of the phenomenology described in the main text. Unless otherwise stated, all the networks are initialized with weights drawn from $\mathcal{N}(0, 1)$ and biases initialized to zero, and are trained using the same protocol as in the main text. A baseline implementation of the




FIG. 9. Lyapunov exponent (left column), percentage of valid exponents (ρ), and mean convergence time ($\langle \tau \rangle$) as functions of the learning rate (η) for the remaining activation functions. Results for the MNIST task.

data-generation pipeline used in this work is publicly available in Ref. [50].

We start with the Iris classification task [51] (150 samples, 4 features, 3 classes; 120/30 train/test split). The network architecture is 4–10–3 (83 trainable parameters). Results are shown in Fig. 8. Despite the low dimensionality of the task, the same qualitative transition is

observed: The minimum in $\langle \tau \rangle$ coincides with the onset of positive λ_{nMLE} .

We next consider the MNIST dataset [43] (60 000/10 000 train/test samples; 28×28 grayscale images) using the 784–64–10 architecture (50 890 parameters). Figure 9 reports results for sigmoid and ReLU activation functions, complementing the tanh case analyzed in the main text. The




FIG. 10. Lyapunov exponent (left column), percentage of valid exponents (ρ), and mean convergence time ($\langle \tau \rangle$) as functions of the learning rate (η). Results for the MNIST task with an extra hidden layer in the architecture and tanh as the activation function.




FIG. 11. Lyapunov exponent (left column), percentage of valid exponents (ρ), and mean convergence time ($\langle \tau \rangle$) as functions of the learning rate (η). Results for the MNIST task with a single hidden layer and tanh activation function, now including L_2 regularization.

qualitative structure persists across nonlinearities. For the sigmoid case, the minimum of $\langle \tau \rangle$ extends over a broad range of learning rates, with the subsequent increase appearing only for $\eta \gtrsim 30$.

To assess depth effects, we replace the single hidden layer (64 neurons) by two hidden layers of 32 neurons each (784–32–32–10; 26 506 parameters), keeping total hidden width fixed. Results in Fig. 10 show that the phenomenology persists, although the minimum convergence time increases

by approximately a factor of 20 relative to the shallow configuration.

We further evaluate L_2 -regularized training with loss $\mathcal{L}_\lambda = \mathcal{L} + \frac{\lambda}{2} \|\mathcal{W}\|^2$, using $\lambda = 10^{-3}$ and 10^{-5} for the MNIST-tanh configuration. As shown in Fig. 11, the region of optimal training efficiency remains aligned with the onset of positive λ_{nMLE} .

Finally, we consider a convolutional architecture for MNIST consisting of two convolutional layers (8 and 16




FIG. 12. Lyapunov exponent (left column), percentage of valid exponents (ρ), and mean convergence time ($\langle \tau \rangle$) as functions of the learning rate (η). Results correspond to the CNN applied to the MNIST classification task.




FIG. 13. Lyapunov exponent (left column), percentage of valid exponents (ρ), and mean convergence time ($\langle\tau\rangle$) as functions of the learning rate (η) for each activation function. Results for the CIFAR-10 task with a target accuracy of 30%.

kernels, 3×3 , padding = 1), each follow by 2×2 max pooling, and a fully connected output layer (9098 trainable parameters). Convolutional layers use ReLU activations and weights are initialized from $\mathcal{N}(0, 0.01)$. Training is performed for up to 2000 epochs with 40 initial conditions and a target accuracy of 80%. As shown in Fig. 12, the characteristic sweet spot in learning rate remains visible.

We conclude with the CIFAR-10 dataset [52] (50 000/10 000 train/test samples; 32×32 RGB (red, green, blue) images) using the 3072–256–10 MLP architecture (789 258 parameters). Since MLPs typically achieve limited performance on this task, we set the target accuracy to 30%. As shown in Fig. 13, the same qualitative behavior is observed despite the increased task complexity.

Appendix C: Evolution of sharpness. Following the approach in Ref. [41], we measure the evolution of sharpness during training, defined as the maximum eigenvalue of the Hessian of the training loss with respect to the model parameters. This measure characterizes the local curvature of the loss surface and provides insight into the dynamics of the optimization process. The sharpness is tracked across epochs for the model trained with the MNIST dataset and tanh as activation function.

We use power iteration to approximate the top eigenvalue of the Hessian. This method is computationally efficient and does not require explicit computation and storage of the full Hessian matrix, so it is feasible to apply it during training. We track the evolution of the sharpness over




FIG. 14. Evolution of sharpness as a function of training epochs. The theoretical value $2/\eta$ is indicated by the dashed red line for reference. Results are shown for learning rate values $\eta = 7.5$ and $\eta = 10$.

training by measuring the top eigenvalue of the Hessian matrix of the loss function, as shown in Fig. 14. For learning rates that minimize $\langle \tau \rangle$, we find that, in line with theoretical considerations [41], sharpness remains bounded

once the training has stabilized and closely matches the theoretical limit of $2/\eta$, as indicated by the red dashed line. The results shown correspond to learning rates $\eta = 7.5$ and $\eta = 10$.

-
- [1] I. J. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning* (MIT Press, Cambridge, MA, USA, 2016), <http://www.deeplearningbook.org>.
 - [2] C. C. Aggarwal, *Neural Networks and Deep Learning* (Springer, Cham, 2018), Vol. 10.
 - [3] L. Lacasa, J. P. Rodríguez, and V. M. Eguiluz, Correlations of network trajectories, *Phys. Rev. Res.* **4**, L042008 (2022).
 - [4] L. A. N. Amaral, Artificial intelligence needs a scientific method-driven reset, *Nat. Phys.* **20**, 523 (2024).
 - [5] G. Bianconi, A. Arenas, J. Biamonte, L. D Carr, B. Kahng, J. Kertesz, J. Kurths, L. Lü, C. Masoller, A. E. Motter, *et al.*, Complex systems in the spotlight: Next steps after the 2021 Nobel Prize in Physics, *J. Phys.: Complexity* **4**, 010201 (2023).
 - [6] L. Arola-Fernández and L. Lacasa, Effective theory of collective deep learning, *Phys. Rev. Res.* **6**, L042040 (2024).
 - [7] K. Danovski, M. C. Soriano, and L. Lacasa, Dynamical stability and chaos in artificial neural network trajectories along training, *Front. Complex Syst.* **2**, 1367957 (2024).

- [8] V. Latora, V. Nicosia, and G. Russo, *Complex Networks: Principles, Methods and Applications* (Cambridge University Press, Cambridge, 2017).
- [9] P. Holme and J. Saramäki, Temporal networks, *Phys. Rep.* **519**, 97 (2012).
- [10] N. Masuda and R. Lambiotte, *A Guide to Temporal Networks* (World Scientific, Singapore, 2016).
- [11] O. E. Williams, L. Lacasa, A. P. Millán, and V. Latora, The shape of memory in temporal networks, *Nat. Commun.* **13**, 499 (2022).
- [12] A. Badie-Modiri, C. Boldrini, L. Valerio, J. Kertész, and M. Karsai, Initialisation and network effects in decentralised federated learning, *Appl. Network Sci.* **10**, 53 (2025).
- [13] E. L. Malfa, G. L. Malfa, G. Nicosia, and V. Latora, Deep neural networks via complex network theory: A perspective, in *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)* (2024), pp. 4361–4369.
- [14] Z. Zheng, H. Liang, V. Snasel, V. Latora, P. Pardalos, G. Nicosia, and V. Ojha, On learnable parameters of optimal and suboptimal deep learning models, in *Proceedings of the International Conference on Neural Information Processing* (Springer, 2024), pp. 126–140.
- [15] H. G. Schuster, W. Just, *Deterministic Chaos: An Introduction* (John Wiley & Sons, Weinheim, 2006).
- [16] H. Kantz and T. Schreiber, *Nonlinear Time Series Analysis* (Cambridge University Press, Cambridge, 2003).
- [17] D. Bertsekas, A. Nedic, and A. Ozdaglar, *Convex Analysis and Optimization* (Athena Scientific, Belmont, MA, 2003), Vol. 1.
- [18] L. Kong and M. Tao, Stochasticity of deterministic gradient descent: Large learning rate for multiscale objective function, *Adv. Neural Inf. Process. Syst.* **33**, 2625 (2020).
- [19] L. Herrmann, M. Granz, and T. Landgraf, Chaotic dynamics are intrinsic to neural network training with SGD, *Adv. Neural Inf. Process. Syst.* **35**, 5219 (2022).
- [20] Y.-C. Lai and T. Tél, *Transient Chaos: Complex Dynamics on Finite Time Scales* (Springer Science & Business Media, New York, 2011), Vol. 173.
- [21] M. Črepinšek, S.-H. Liu, and M. Mernik, Exploration and exploitation in evolutionary algorithms: A survey, *ACM Comput. Surv. (CSUR)* **45**, 35 (2013).
- [22] O. Berger-Tal, J. Nathan, E. Meron, and D. Saltz, The exploration-exploitation dilemma: A multidisciplinary framework, *PLoS One* **9**, e95693 (2014).
- [23] N. E. Humphries, H. Weimerskirch, N. Queiroz, E. J. Southall, and D. W. Sims, Foraging success of biological Lévy flights recorded in situ, *Proc. Natl. Acad. Sci. USA* **109**, 7169 (2012).
- [24] G. Ramos-Fernández, J. L. Mateos, O. Miramontes, G. Cocho, H. Larralde, and B. Ayala-Orozco, Lévy walk patterns in the foraging movements of spider monkeys (*Ateles geoffroyi*), *Behav. Ecol. Sociobiol.* **55**, 223 (2004).
- [25] S. Eliassen, C. Jørgensen, M. Mangel, and J. Giske, Exploration or exploitation: Life expectancy changes the value of learning in foraging strategies, *Oikos* **116**, 513 (2007).
- [26] A. Reynolds, E. Ceccon, C. Baldauf, T. K. Medeiros, and O. Miramontes, Lévy foraging patterns of rural humans, *PLoS One* **13**, e0199099 (2018).
- [27] J. M. Kembro, M. Lihoreau, J. Garriga, E. P. Raposo, and F. Bartumeus, Bumblebees learn foraging routes through exploitation–exploration cycles, *J. R. Soc. Interface* **16**, 20190103 (2019).
- [28] C. T. Monk, M. Barbier, P. Romanczuk, J. R. Watson, J. Alós, S. Nakayama, D. I. Rubenstein, S. A. Levin, and R. Arlinghaus, How ecology shapes exploitation: A framework to predict the behavioural response of human and animal foragers along exploration–exploitation trade-offs, *Ecol. Lett.* **21**, 779 (2018).
- [29] L. R. Paiva, S. G. Alves, L. Lacasa, O. DeSouza, and O. Miramontes, Visibility graphs of animal foraging trajectories, *J. Phys.: Complexity* **3**, 04LT03 (2022).
- [30] R. Klages, G. Radons, and I. M. Sokolov, *Anomalous Transport: Foundations and Applications* (Wiley-VCH, Weinheim, 2008).
- [31] T. T. Hills, P. M. Todd, D. Lazer, A. D. Redish, and I. D. Couzin, Exploration versus exploitation in space, mind, and society, *Trends Cognit. Sci.* **19**, 46 (2015).
- [32] M. A. Addicott, J. M. Pearson, M. M. Sweitzer, D. L. Barack, and M. L. Platt, A primer on foraging and the explore/exploit trade-off for psychiatry research, *Neuropsychopharmacology* **42**, 1931, (2017).
- [33] S. Ishii, W. Yoshida, and J. Yoshimoto, Control of exploitation–exploration meta-parameter in reinforcement learning, *Neural Networks* **15**, 665 (2002).
- [34] N. Bertschinger, T. Natschläger, and R. Legenstein, At the edge of chaos: Real-time computations and self-organized criticality in recurrent neural networks, *Adv. Neural Inf. Process. Syst.* **17**, 145 (2004).
- [35] J. Kadmon and H. Sompolinsky, Transition to chaos in random neuronal networks, *Phys. Rev. X* **5**, 041030 (2015).
- [36] D. Sussillo and L. F. Abbott, Generating coherent patterns of activity from chaotic neural networks, *Neuron* **63**, 544 (2009).
- [37] U. Pereira-Obilinovic, J. Aljadeff, and N. Brunel, Forgetting leads to chaos in attractor networks, *Phys. Rev. X* **13**, 011009 (2023).
- [38] D. Pazó, Discontinuous transition to chaos in a canonical random neural network, *Phys. Rev. E* **110**, 014201 (2024).
- [39] J. Luo, J. Chen, and H.-K. Zhang, The butterfly effect in neural networks: Unveiling hyperbolic chaos through parameter sensitivity, *Neural Networks* **189**, 107572 (2025).
- [40] L. Storm, H. Linander, J. Bec, K. Gustavsson, and B. Mehlig, Finite-time Lyapunov exponents of deep neural networks, *Phys. Rev. Lett.* **132**, 057301 (2024).
- [41] J. Cohen, S. Kaur, Y. Li, J. Z. Kolter, and A. Talwalkar, Gradient descent on neural networks typically occurs at the edge of stability, in *International Conference on Learning Representations* (2021).
- [42] C. G. Langton, Computation at the edge of chaos: Phase transitions and emergent computation, *Physica D* **42**, 12 (1990).
- [43] Y. LeCun, The MNIST database of handwritten digits (1998), <http://yann.lecun.com/exdb/mnist/>.
- [44] A. Caligiuri, V. M. Eguíluz, L. Di Gaetano, T. Galla, and L. Lacasa, Lyapunov exponents for temporal networks, *Phys. Rev. E* **107**, 044305 (2023).
- [45] A. Caligiuri, T. Galla, and L. Lacasa, Characterizing the dynamics of unlabeled temporal networks, *Chaos* **35**, 053122 (2025).

- [46] P. F. M. J. Verschure, Chaos-based learning, *Complex Syst.* **5**, 359 (1991).
- [47] Y. Dandi, F. Krzakala, B. Loureiro, L. Pesce, and L. Stephan, How two-layer neural networks learn, one (giant) step at a time, *J. Mach. Learn. Res.* **25**, 1 (2024).
- [48] L. Lacasa, F. J. Marín-Rodríguez, N. Masuda, and L. Arola-Fernández, Scalar embedding of temporal network trajectories, *Chaos, Solitons Fractals* **199**, 116599 (2025).
- [49] L. Lacasa, Fluid dynamics meet network science: Two cases of temporal network eigendecomposition, [arXiv:2509.03135](https://arxiv.org/abs/2509.03135).
- [50] <https://github.com/pedrojg8/chaotic-transients-in-ANNs>.
- [51] R. A. Fisher, The use of multiple measurements in taxonomic problems, *Ann. Eugen.* **7**, 179 (1936).
- [52] A. Krizhevsky and G. Hinton, Learning multiple layers of features from tiny images, Master's thesis, Department of Computer Science, University of Toronto, 2009.