

Effective theory of collective deep learning

Lluís Arola-Fernández* and Lucas Lacasa†

Instituto de Física Interdisciplinar y Sistemas Complejos IFISC (CSIC-UIB), Campus UIB, 07122 Palma de Mallorca, Spain

(Received 9 November 2023; accepted 16 October 2024; published 12 November 2024)

Unraveling the emergence of collective learning in systems of coupled artificial neural networks points to broader implications for physics, machine learning, neuroscience, and society. Here we introduce a minimal model of interacting deep neural nets that condenses several recent decentralized algorithms by considering a competition between two terms: the local learning dynamics in the parameters of each neural network unit, and a diffusive coupling among units that tends to homogenize the parameters of the ensemble. We derive an effective theory for linear networks to show that the coarse-grained behavior of our system is equivalent to a deformed Ginzburg-Landau model with quenched disorder. This framework predicts depth-dependent disorder-order-disorder phase transitions in the parameters' solutions that reveal a depth-delayed onset of a collective learning phase and a low-rank microscopic learning path. We validate the theory in coupled ensembles of realistic neural networks trained on MNIST and CIFAR-10 datasets under privacy constraints. Interestingly, experiments confirm that individual networks—trained on private data—can fully generalize to unseen data classes when the collective learning phase emerges. Our work elucidates the physics of collective learning and contributes to the mechanistic interpretability of deep learning in decentralized settings.

DOI: [10.1103/PhysRevResearch.6.L042040](https://doi.org/10.1103/PhysRevResearch.6.L042040)

Introduction. Collective behavior emerging from the dynamics of many interacting units is the flagship of complexity [1–3] and a common feature across natural and artificial complex systems [4–6]. Learning, a paradigmatic example of collective behavior [7,8], is a fascinating ability of the brain that also occurs in a collective manner among interacting agents that learn from each other, and at different scales, from ant colonies [5] to robot swarms [9] or social communities [10]. Such generality and multiscale nature of learning has attracted great interest across fields and major advances in the problem have benefited from cross-disciplinary research. In fact, ideas from complex systems and statistical physics are behind the early models of computation in neural networks [11] and in classical machine learning tools like particle swarm optimization [12], among other recent notable contributions [13–18].

The many-particle approach has proved successful in deep learning [19], a collection of algorithms involving large and overparametrized neural networks, which has recently shown astonishing results in many challenging tasks [20,21] and intriguing emerging behavior in large language models [22]. Interestingly, industry-related constraints posed by data privacy and power consumption during training [23,24], along with the perpetual quest for finding architectures with enhanced performance, have driven the field of deep learning

to explore solutions that capitalize on some form of collective behavior [17,19]. These solutions range from ensemble approaches [23,25] to federated learning [24,26,27], cooperative learning [28], or transfer learning [29]. This engineering success provides a strong motivation to understand, from a principled complex-systems viewpoint, whether interacting ‘brains’—rather than interacting ‘neurons’—develop collective behavior in the physical sense [1,5,17]. Does collective learning emerge when neural networks are put in interaction? Does the collective learning phase emerge abruptly as in the theory of phase transitions? And do deep learning architectures play nontrivial roles in such phenomenology? In this article we provide affirmative answers to all these questions. We study a minimal mathematical model of collective learning where we show that local brains trained for isolated tasks (private data) can generalize outside their training set when coupled. This happens via a collective learning phase transition—which is delayed by neural depth—as the coupling strength between the brains increases.

The discovery of collective behavior and critical phenomena in our work is triggered by an approximate correspondence that we find between the coarse-grained dynamics of learning parameters in currently popular decentralized algorithms and a paradigmatic model of statistical physics developed decades ago to understand emergent phenomena in systems of spins under interaction. This connection provides a rather simple effective theory that is amenable to mechanistic interpretability [13–16] at a collective scale, allowing us to make several qualitative and quantitative predictions that are subsequently confirmed in a range of realistic experiments on the MNIST and CIFAR-10 datasets.

Collective learning model. Our proof of concept considers a classification task to be solved by a coupled ensemble of vanilla neural nets, where each neural network is trained on

*Contact author: lluissarolaf@gmail.com†Contact author: lucas@ifisc.uib-csic.es




FIG. 1. Collective learning in coupled neural networks: (a) Small mean-field network of $N = 4$ coupled neural units trained on private data classes of MNIST, interacting with coupling strength σ . (b) Simplified deep feed-forward architecture of a neural unit (trained on digit-3 images). A test point outside the training set (a digit-0 image) is flattened as an input and processed towards the output layer. The unit predicts well an unseen class during training due to a collective learning mechanism emerging from the interactions.

data from a single class and evaluation is performed on the whole (multiclass) test set. However, the framework introduced below is flexible and extends to a generic supervised task or neural architecture.

Let $\mathcal{D} = \{(\mathbf{x}, \mathbf{y})\} \subseteq \mathbb{R}^{n_0} \times \mathbb{Z}^{n_{D+1}}$ denote the training set, where $\mathcal{X} = \{\mathbf{x} : (\mathbf{x}, \mathbf{y}) \in \mathcal{D}\}$ and $\mathcal{Y} = \{\mathbf{y} : (\mathbf{x}, \mathbf{y}) \in \mathcal{D}\}$ denote the input and output (labels) vectors with dimension n_0 and n_{D+1} , respectively. The training data is partitioned and distributed across N neural units, such that $\mathcal{D} = \cup_{i=1}^N \mathcal{D}_i$. Each unit is a fully-connected feed-forward deep network with D hidden layers with widths n_d , for $d = 1, \dots, D$, a readout (output) layer with n_{D+1} neurons, and nonlinear activation functions in the hidden neurons. We define $\bar{\mathbf{y}} = f_i(\mathbf{x}, \boldsymbol{\theta}_i)$ as the predicted output given the input $\mathbf{x}$ and the set of trainable parameters $\boldsymbol{\theta}_i = \{\theta_i^\alpha\}$, where α runs over all parameters of the i -th unit. We focus on a learning process under privacy constraints, such that each neural unit seeks to minimize the local empirical loss function

$$\mathcal{L}_i = \sum_{(\mathbf{x}_i, \mathbf{y}_i) \in \mathcal{D}_i} \ell[f(\mathbf{x}_i, \boldsymbol{\theta}_i), \mathbf{y}_i] + \gamma \|\boldsymbol{\theta}_i\|_2, \quad (1)$$

where the sum runs over data associated to the i -th unit. The individual loss function $\ell(\bar{\mathbf{y}}_i, \mathbf{y}_i) : \mathbb{R}^{n_{D+1}} \times \mathbb{R}^{n_{D+1}} \rightarrow \mathbb{R}$, measures the error between the prediction on a data point and the corresponding true value (or label) and the right-hand term is a standard weight decay (L_2 regularization) with strength γ that encourages lower weights' magnitudes, reducing the empirical risk of overfitting and improving generalization [17,19]. To minimize Eq. (1), the Backpropagation algorithm tunes the parameters of the units efficiently via Stochastic Gradient Descent (SGD) [19].

Since local learning is restricted to the private data of each unit, we introduce an interaction mechanism to induce collective learning in the system, allowing units to learn from each other (see Fig. 1). As a minimal mechanism, we select a consensus-based model [28], where units are diffusively coupled in a relation 'parameter-to-parameter'. The interaction links (and their intensities) are captured by the weighted adjacency matrix $\mathcal{Q} \in \mathbb{R}^{N \times N} = \{q_{ij}\}$ of the supra-network,

i.e., the network that defines the interactions between the neural units. Combining both local (the minus gradient of the local loss) and interaction terms (the diffusive coupling), each of the α parameters of a unit is updated at each iteration with

$$\theta_i^\alpha(t+1) = \theta_i^\alpha(t) - \eta \nabla_{\theta_i^\alpha} \mathcal{L}_i + \frac{\eta \sigma}{N} \sum_{j=1}^N q_{ij} (\theta_j^\alpha(t) - \theta_i^\alpha(t)), \quad (2)$$

where η is the learning rate and σ the coupling strength. To assess collective learning, we let the parameter dynamics relax to a stationary state and then measure the mean test loss, i.e., the local loss evaluated on global (multiclass) test data $\mathcal{D}^{\text{test}}$, averaged over the units

$$\langle L \rangle = \frac{1}{N} \sum_{i=1}^N \sum_{(\mathbf{x}, \mathbf{y}) \in \mathcal{D}^{\text{test}}} \ell[f(\mathbf{x}, \boldsymbol{\theta}_i), \mathbf{y}], \quad (3)$$

which serves as our main metric of collective performance. We expect a low loss as a fingerprint of collective learning, indicating that units can generalize outside their training set. In classification tasks, the accuracy metric $\langle A \rangle$ measures the percentage of correct class predictions in the test set (again brackets indicate an average over neural units). In order to track microscopic information of the system beyond the ensemble (macroscopic) metrics, we use the cross-loss and cross-accuracy matrices $L_x, A_x \in \mathbb{R}^{N \times N}$, where the entry (i, j) determines the loss (accuracy) of the i -th unit when evaluated on data classes assigned to the j -th unit. Note that the dynamics in Eq. (2) lies in the cooperative learning framework [28], akin to the popular decentralized scheme of elastic averaging SGD [30] and also captures the spirit of standard federated learning [26], with the conceptual differences that here the self-organized dynamics of the units is independent of a centralized *master* model and that the diffusive interaction occurs simultaneously with the local training, instead of interacting periodically after several local updates. Previous results proved the convergence of cooperative and federated schemes, see for instance [27,28,30–34], unveiling that while heterogeneous and private data tend to slow down convergence [34], distributed schemes can find solutions with higher generalization properties [30,35,36]. The mechanisms underpinning the onset of emergent collective behavior in coupled neural networks under privacy constraints remain, however, poorly understood.

Coarse-grained theory. We leverage on learning scale separation and results on deep linear networks [16,37]—that shed light on the landscape of nonlinear networks and the existence of regularization-induced learning phase transitions—to derive a coarse-grained theory describing the dynamics in Eq. (2). Under suitable conditions (see Appendix A for full derivation), and considering a magnetization-like scalar order parameter m_i as defined in Appendix A (which coarse-grains the parameters of the i -th neural unit across layers), the dynamics of Eq. (2) approximately project onto a reduced system of N coupled differential equations

$$\dot{m}_i = \delta_i m_i^D - m_i^{2D+1} - \hat{\gamma} m_i + \frac{\hat{\sigma}}{N} \sum_{j=1}^N q_{ij} (m_j - m_i), \quad (4)$$

where $\dot{m}_i$ indicates the time derivative, $\hat{\sigma}$ and $\hat{\gamma}$ are effective hyper-parameters (coupling and regularization, respectively), q_{ij} are the entries of the supra-network adjacency matrix, and D is the neural depth of the units. The full information on the training set is encapsulated in δ , a sequence of scalars $\delta_i = \langle \hat{x}_i \hat{y}_i \rangle$, each of them being an effective one-dimensional representation of the expected input-output correlation, averaged over the data points assigned to the i -th unit. We build the effective test set by aggregating the private training distributions, such that $\langle \delta \rangle = (1/N) \sum_i \langle \hat{x}_i \hat{y}_i \rangle$. Interestingly, the mean effective loss—evaluated on the test set and averaged over the units—scales with the moments of the magnetization as

$$\langle \hat{L} \rangle \sim \langle m^{2(D+1)} \rangle - 2\langle \delta \rangle \langle m^{(D+1)} \rangle, \quad (5)$$

[see Appendix A for full derivation of Eq. (5)]. Our goal is to assess the onset of collective learning by solving Eq. (5) in the stationary regime of Eq. (4). From now on, we take an all-to-all unweighted supra-network (mean-field) with $q_{ij} = 1, \forall i \neq j$ and distinguish three scenarios:

Linear regression ($D = 0$). When units have no hidden layers, Eq. (4) becomes linear. Its stationary solution ($\dot{m}_i = 0$), found self-consistently, reads

$$m_i^* = \frac{\delta_i + \hat{\sigma} \langle \delta \rangle (1 + \hat{\gamma})^{-1}}{1 + \hat{\gamma} + \hat{\sigma}}. \quad (6)$$

The order parameter $\langle m \rangle = \langle \delta \rangle / (1 + \hat{\gamma})$ is independent of coupling $\hat{\sigma}$, as shown in Fig. 2(a) (light crosses). Increasing $\hat{\sigma}$ only makes the distribution of magnetization narrower, which translates into a monotonously decreasing loss function, as displayed in Fig. 2(b).

Shallow networks ($D = 1$). Eq. (4) becomes

$$\dot{m}_i = (\delta_i - \hat{\gamma})m_i - m_i^3 + \frac{\hat{\sigma}}{N} \sum_{j=1}^N (m_j - m_i). \quad (7)$$

Remarkably, this equation is formally identical to the mean-field, zero-temperature, quenched version of the Ginzburg-Landau (GL) model (ϕ^4 model) [38–42], traditionally used in condensed-matter physics to explore the critical behavior of systems with impurities under Landau's free energy approach. The local magnetization m_i or *spin* is the average parameter of our neural unit, whereas the heterogeneous allocation of training data plays the role of multiplicative quenched disorder, i.e., that of the material's random impurities [38]. Equation (7) displays a rich phenomenology which can be reinterpreted in the context of our neural system: For low coupling $\hat{\sigma}$, units trained on data $\delta_i > \hat{\gamma}$ relax in a double well local potential (with two symmetric stable points at $m_{i,\pm}^* = \pm \sqrt{\delta_i - \hat{\gamma}}$), and units with $\delta_i < \hat{\gamma}$ have a single equilibrium point at $m_i^* = 0$. This distribution produces a disordered state (with $|\langle m \rangle| \approx 0$). At a critical coupling, a *symmetry-breaking* mechanism induces a collective ordered phase (with $|\langle m \rangle| > 0$) via a second-order collective phase-transition. For even larger coupling, the system returns to a 'disordered' phase (via a reentrant phase-transition [39]), where $m_i^* \approx 0 \forall i$, by means of a *collective regularization* mechanism [when diffusion dominates in Eq. (4)]. These disorder-order-disorder phase transitions are visualized in Fig. 2(a) (triangles). Interestingly, the critical behavior of the local magnetization triggers a non-monotonous decay of the effective loss. As shown in Fig. 2(b)




FIG. 2. Theoretical predictions. (a) Mean magnetization depending on coupling, for increasing D (from light to dark). The theory predicts disorder-order-disorder transitions in $\langle m \rangle$ for $D > 0$. (b) Mean effective loss depending on coupling (inset shows $\langle \hat{A} \rangle \approx 1 - \langle \hat{L} \rangle$ in a reduced range). The theory predicts a depth-induced delay which translates into a delayed critical point of collective learning ($\langle \hat{A} \rangle > 0$). Markers denote mean values averaged over initial conditions and quenched disorders (shaded area represents one std.) and dashed lines the outcome of adiabatic protocols, averaged over quenched sequences. (c) Ranking of the $\hat{A}_x$ entries from Eq. (8) in the quenched 'data' disorder plane for no coupling, (d) Medium coupling, and (e) High coupling. See Appendix B for details on the numerical integration of the theory.

(and inset), the nonmonotonous shape delays the transition to the collective learning regime with respect to $D = 0$.

Deep networks ($D > 1$). In the deep learning realm, Eq. (4) represents an exotic deformation of the GL model— $a \phi^{2D+2}$ instead of a ϕ^4 model—with odd powers of the order parameter entering in Landau's free energy for even neural unit depths D , hence breaking rotational symmetry [42]. Figure 2(a) (dark circles) shows that disorder-order-disorder transitions are still found for $D = 2$, whereas the nonmonotonic behavior of the effective loss is enhanced with respect to the shallow case. This effect translates into a further delay of the transition to the collective learning regime, as observed in Fig. 2(b) and its inset, i.e., the transition occurs for a larger value of the coupling strength with respect to shallower models. Neural depth delays the emergence of the collective learning phase and it also induces an effective landscape with many local minima. A linear stability analysis reveals that $m_i^* = 0$ is always a local minimum, which triggers a first-order, regularization-induced, phase transition at the single unit level (see [16] for details). Indeed, the local bistability gets further amplified at the collective scale. Figures 2(a) and 2(b) show significant deviations between nonadiabatic (circles) protocol—where random initial conditions are drawn for each coupling value—, and adiabatic (dashed line) protocol—where the initial conditions are the stationary values of the parameters reached at the previous coupling value—, unveiling

that for $D > 1$ there is a strong sensitivity to initial conditions and a widespread presence of multistability.

Microscopic learning path. We leverage our theory to unveil, by analytical means, the order in which neural units learn from each other as the system enters the collective learning phase. We construct the cross-loss matrix $\hat{L}_x$, with $(\hat{L}_x)_{ij} \sim m_i^{D+1}(m_i^{D+1} - 2\delta_j)$. In the uncoupled regime of Eq. (4), we have $m_i(D) \sim m_i(0)^{1/(D+1)}$ and $\hat{L}_x$ can therefore be estimated from the $D = 0$ case. We extend this depth-independence as an *ansatz* to the coupled regime and compute $\hat{L}_x$ using Eq. (6). Using that a high accuracy requires a low loss [19] ($\hat{A}_x \sim -\hat{L}_x$) we get

$$(\hat{A}_x)_{ij} \sim \frac{\delta_i + \hat{\sigma} \langle \delta \rangle (1 + \hat{\gamma})^{-1}}{1 + \hat{\gamma} + \hat{\sigma}} \left(2\delta_j - \frac{\delta_i + \hat{\sigma} \langle \delta \rangle (1 + \hat{\gamma})^{-1}}{1 + \hat{\gamma} + \hat{\sigma}} \right). \quad (8)$$

To analyze Eq. (8), we set $\gamma, \langle \delta \rangle \rightarrow 0^+$ and evaluate the normalized rankings of $(\hat{A}_x)_{ij}$ in the positive support of δ . First, in the uncoupled regime we have $(\hat{A}_x)_{ij} \sim 2\delta_i\delta_j - \delta_i^2$ [Fig. 2(c)]. The approximate diagonal shape indicates that units have the highest cross-accuracy when evaluated on data δ_j that is similar to their training data δ_i , as expected in the local learning regime. As coupling increases, $(\hat{A}_x)_{ij} \sim (2\delta_i\delta_j + \hat{\sigma} \langle \delta \rangle \delta_j)/(1 + \hat{\sigma})$. The first symmetric term [Fig. 2(d)] dominates for a wide range of couplings. Instead, for very large coupling [Fig. 2(e)], the cross-accuracy depends only on the data evaluated, not on the unit predicting it, as expected when the units become so similar (due to diffusive coupling) that they make the same predictions and errors. Last, from Eq. (8) we estimate the cumulative cross-accuracy matrix S_x to quantify the learning path. The integral $\hat{S}_x = \int_0^{\sigma'} \hat{A}_x d\sigma$ scales, in matrix form, as $\hat{S}_x \sim \log \sigma' \delta \delta^T + \sigma' \langle \delta \rangle \delta \mathbf{1}^T$ (where δ^T is the transpose of the quenched ‘data’ disorder vector and $\mathbf{1}$ is a vector of N ones). Consistent with the low-dimensional nature of our theory, the learning path between the units is approximately rank-one, i.e., driven by a single effective dimension (the vector δ). From the previous analysis, the interpretation of δ becomes clear: the higher the δ_j (the stronger the input-output correlation of the j -th data), the earlier a unit will predict it well (which is amplified if the unit is also trained on a high δ_i). These findings are reminiscent of how learning works in a single neural network trained on global data [37], which draws an interesting parallel between local and collective scales.

Empirical validation. To validate the effective theory, we tackle the standard MNIST image classification task with $N = 10$ coupled nonlinear feed-forward networks. Each unit is trained on images of a unique class (a digit from 0 to 9) and learns under the update rule in Eq. (2). Performance is evaluated on an independent, multiclass test set using Eq. (3) and related metrics. We refer to Appendix C for further experimental details. Figure 3(a) confirms that, for $D > 0$, disorder-order-disorder phase transitions in $\langle m \rangle$ appear as coupling increases. As predicted, the phase transitions induce a non-monotonous behavior of the mean loss [see Fig. 3(b)], along with a delay of the loss decay to zero (amplified with neural depth D). This effect triggers the appearance of a (depth-delayed) critical point where collective learning emerges,




FIG. 3. MNIST experiments. (a) Mean magnetization, (b) Mean test loss, and (c) Mean test accuracy depending on coupling (in log scale), for depths $D = 0, 1, 2$. Shaded area shows one std. over 10 independent runs. Inset shows results in a linear scale of coupling. Dashed line are the outcome of adiabatic protocols ($D = 2$). (d) Correlation between the entries of the empirical S_x matrix (averaged over 10 runs and shifted by the mean) for $D = 0, 1, 2$ against the entries of the rank-one truncation of S_x for $D = 0$. We show the R^2 scores and the best linear fit as a red line. (e) Empirical A_x for $D = 1$ at $\sigma = 0$, (f) $\sigma = 3$, and (g) $\sigma = 30$. Rows and columns have been ordered from lowest to highest cross-accuracy (as indicated by MNIST digits in (e) and the y-axis is inverted to guide the comparison with Fig. 2.

as observed in Fig. 3(c). The mean accuracy metric grows abruptly from $\langle A \rangle \approx 0.1$ (local learning phase) where units only predict well their private class, to $\langle A \rangle > 0.1$ at the critical point, up to $\langle A \rangle \approx 1$ (collective learning phase) for larger coupling, where units predict well all classes. Also note that for $D = 2$ the differences between nonadiabatic (circles) and adiabatic protocols (dashed line) confirm the presence of multistability in the deep case: Intriguingly, the collective learning transition is abruptly delayed when the system has ‘memory’ (i.e., when parameters are only initialized at $\sigma = 0$). Figure 3(d) displays the empirical correlation between the entries of the cumulative cross-accuracy matrix and its rank-one approximation for $D = 0$ (the predicted approximate learning path), computed using the singular value decomposition of S_x and keeping only the leading term. A considerable high R^2

score is sustained for varying depth, providing a quantitative validation of the coarse-grained approach at the microscopic level. Finally, Fig. 3 (bottom) shows snapshots of the cross-accuracy matrix A_x (for $D = 1$) at three coupling values. In (e) a local phase with an associated diagonal A_x for low coupling [as Fig. 2(c)], (f) a symmetric collective learning phase for medium coupling—above the critical point—[as Figs. 2(d) and 2(g)] a column-dependent phase for large coupling [43], where accuracy only depends on the evaluated class [as Fig. 2(e)], thus reproducing the patterns predicted by our analysis.

We refer to the Supplemental Material [44] File (SI) for additional experiments varying activation functions (Sec. A), the CIFAR-10 dataset and convolutional nets (Sec. B) and comparing the performance of the collective system against a single unit trained with global data (Sec. C). We test the robustness of the phenomena predicted by our theory against changes in dataset and neural architecture, and show that the collective performance of a system with private data is comparable (although with a slower convergence) to that of a single model trained with global data, validating our minimal setup to analyze realistic decentralized scenarios.

Discussion. While our effective-coarse-grained-theory is based on several simplifications and approximations, the predictions about the onset and properties of collective learning are well confirmed on realistic (high-dimensional, non-linear) experiments on standard datasets, which supports the mechanistic interpretability of the collective learning paradigm. Yet, much theoretical and experimental work is required to better understand its rich phenomenology and the relation of decentralized learning schemes with the statistical physics of (deformed) Ginzburg-Landau models. Our proof of concept did not cover the myriad of learning phases found in isolated neural networks [15,45–48], which can get more exotic in the collective case. Extending this theory to the complex plane may capture some of these nuances and even increase its analytical tractability (thanks to exact dimensional reductions and other techniques for coupled oscillators [49–52]). Other promising directions include exploring how the shape and difficulty of a dataset [48] (*effectively* encapsulated in the distribution of quenched disorders) and the changes in the supra-network topology [53] (beyond the mean-field, undirected, pair-wise and time-independent case) affect the phase transitions [54–56], and refining the diffusive coupling mechanism to model heterogeneous neural learners and indirect communication channels [24,57,58].

In a nutshell, this work offers a minimal mathematical model of collective learning in natural and artificial systems. Our approach considers deep neural networks as the interacting units of a large complex system and leverages an effective theory at the collective scale, revealing that the emergence of collective learning occurs via phase transitions which are delayed by neural depth. This perspective enriches deep learning theories [13–18] and statistical physics approaches to interacting brains [5,6], and makes a first step towards a next-generation type of physical models [50,59] to describe emergent social behavior—such as collective learning—in populations of interacting agents. Last, it has not escaped our notice that this framework could contribute to tackle the so-called alignment problem [60] when independent learning models are put in interaction.

Python scripts that reproduce the results of this work are available on [44,61].

Acknowledgments. We thank many colleagues at IFISC, C. Granell, A. Arenas, and A. Raya for useful discussions. This work has been funded by the Spanish State Research Agency, through projects DYNDEEP, MISLAND and a María de Maeztu project (EUR2021-122007, PID2020-114324GB-C22 and CEX2021-001164-M, respectively), funded by the MICIU/AEI/10.13039/501100011033. The authors also acknowledge the support from the European Commission Chips Joint Undertaking Project No. 101194363 (NEHIL).

Appendix A: Analytical derivation of the effective system. To derive the coarse-grained dynamics, Eq. (4), we begin by assuming that the learning rate η in Eq. (2) is sufficiently small, such that we can approximate the discrete dynamics by the continuous version

$$\dot{\theta}_i^\alpha = -\nabla_{\theta_i^\alpha} \mathcal{L}_i(\mathbf{x}_i, \mathbf{y}_i, \boldsymbol{\theta}_i, \boldsymbol{\gamma}) + \frac{\sigma}{N} \sum_{j=1}^N q_{ij} (\theta_j^\alpha - \theta_i^\alpha), \quad (\text{A1})$$

where $\dot{\theta}_i^\alpha = d\theta_i^\alpha/dt$ is the time-derivative of a parameter, the time step dt in the numerical integration of Eq. (2) corresponds to the learning rate η and bold variables denote vectors (a notation used also in the main text).

Now we follow a common procedure in physics, i.e., to linearize the system and study the properties of the resulting linear approximation [16,17,37]. When the choice of $\ell(\bar{\mathbf{y}}, \mathbf{y})$ is the mean square error (MSE), the loss function of a deep linear network can be written as [16]

$$\hat{\mathcal{L}}_i = \mathbb{E}_{x_i} \left[\left(\sum_{d_0, \dots, d_{D+1}}^{n_0, \dots, n_{D+1}} [\Theta_{d_{D+1}d_D}^{D+1}]_i \cdots [\Theta_{d_1d_0}^1]_i [x_{d_0}]_i - y_i \right)^2 \right] + \gamma \sum_{d=1}^{D+1} \|\Theta_i^d\|_2^2, \quad (\text{A2})$$

where $\mathbb{E}_{x_i}[\cdot]$ denotes the expected value over the accessible data points and $\|\Theta_i^d\|_2$ is the squared L_2 norm of all elements in the parameters' matrix of the d -layer. Each sum in the first term of Eq. (A2) runs over all the neurons of a given layer with width n_d . The index i is kept to reflect that the loss, the parameters and the data belong to the i -th neural unit in the ensemble setting.

Here we follow the ‘mean-field’ analysis of [16], assuming that both the input and output are one-dimensional, and we approximate each matrix Θ_i^d by the mean value of its entries, a scalar $c_d \hat{\theta}_i^d$, where c_d is a layer-dependent constant. For the sake of simplicity, we set $c_d = 1$ in this study. The loss function of the i -th unit in Eq. (A2) then reads as

$$\hat{\mathcal{L}}_i \approx \mathbb{E}_{x_i} \left[\left(\hat{x}_i \prod_{d=1}^{D+1} \hat{\theta}_i^d - \hat{y}_i \right)^2 \right] + \gamma \sum_{d=1}^{D+1} (\hat{\theta}_i^d)^2, \quad (\text{A3})$$

where $\hat{x}_i$ and $\hat{y}_i$ are the one-dimensional projections of input and output for a data point assigned to the i -th unit. Under this mean-field approximation, we can reduce the number of coupled equations in Eq. (A1) from a very large number (α runs from 1 to the dimension of $\text{vec}(\bigcup_{d=1}^{D+1} \Theta^d)$) to just $(D+1) \times N$ equations, one for each layer of parameters in

each unit. In particular, we can now compute the gradient in Eq. (A1) with respect to $\hat{\theta}_i^d$, and using the loss function of Eq. (A3) we get

$$\hat{\theta}_i^d = -\frac{\partial \hat{\mathcal{L}}_i}{\partial \hat{\theta}_i^d} + \frac{\sigma}{N} \sum_{j=1}^N q_{ij} (\hat{\theta}_j^d - \hat{\theta}_i^d), \quad (\text{A4})$$

with $i \in [1, N]$, $d \in [1, D+1]$ and

$$\frac{\partial \hat{\mathcal{L}}_i}{\partial \hat{\theta}_i^d} = -2 \langle \hat{x}_i \hat{y}_i \rangle \prod_{k \neq d}^{D+1} \hat{\theta}_i^k + 2 \langle \hat{x}_i^2 \rangle \hat{\theta}_i^d \prod_{k \neq d}^{D+1} (\hat{\theta}_i^k)^2 + 2\gamma \hat{\theta}_i^d, \quad (\text{A5})$$

where we use the simplified notation $\langle \cdot \rangle$ to indicate an average (here over the data points assigned to a unit) as we do in the main text. Let us define the local magnetization $m_i = (D+1)^{-1} \sum_d \hat{\theta}_i^d$ as the average learning parameter of the i -th unit and compute the average value of the gradient in Eq. (A5), summing over the different layers, leading to

$$\frac{\partial \hat{\mathcal{L}}_i}{\partial m_i} \approx -2 \langle \hat{x}_i \hat{y}_i \rangle m_i^D + 2 \langle \hat{x}_i^2 \rangle m_i^{2D+1} + 2\gamma m_i, \quad (\text{A6})$$

where we have used the mean-field approximations

$$(D+1)^{-1} \sum_{d=1}^{D+1} (\partial \hat{\mathcal{L}}_i / \partial \hat{\theta}_i^d) \approx \frac{\partial \hat{\mathcal{L}}_i(m_i)}{\partial m_i}, \quad (\text{A7})$$

$$(D+1)^{-1} \sum_{d=1}^{D+1} \prod_{k \neq d}^{D+1} \hat{\theta}_i^k \approx m_i^D, \quad (\text{A8})$$

$$(D+1)^{-1} \sum_{d=1}^{D+1} \hat{\theta}_i^d \prod_{k \neq d}^{D+1} (\hat{\theta}_i^k)^2 \approx m_i^{2D+1}. \quad (\text{A9})$$

Now we plug Eq. (A6) into the sum of the $(D+1)$ equations in Eq. (A4) for each unit. Assuming orthogonal input representations [37,45] (which holds exactly for whitened input data [37]) we set $\langle \hat{x}_i^2 \rangle = 1$ without loss of generality. After rescaling time with $dt' = 2dt$ and absorbing any additional constants into the effective variables $\hat{\sigma}$, $\hat{\gamma}$ and $\delta_i = \langle \hat{x}_i \hat{y}_i \rangle$, we get Eq. (4) in the main text.

To derive the mean effective loss in Eq. (5), first note from Eq. (A3) that the effective output function of a neural unit takes the simple form $\bar{y}_i \approx m_i^{D+1} \hat{x}_i$. Then, using the MSE as the individual loss function $\ell(\bar{y}_i, \hat{y}_i)$ —where $\bar{y}_i$ and $\hat{y}_i$ are both scalars—, the effective loss of the i -th unit when evaluating data assigned to the j -th unit (i.e., the cross-loss entry $(\hat{\mathcal{L}}_x)_{ij}$) becomes

$$(\hat{\mathcal{L}}_x)_{ij} = (m_i^{D+1} \hat{x}_j - \hat{y}_j)^2. \quad (\text{A10})$$

Decomposing terms, neglecting the constant $\langle \hat{y}_j^2 \rangle$ that only depends on the dataset, and assuming again orthogonal input data, we have

$$(\hat{\mathcal{L}}_x)_{ij} \sim m_i^{D+1} (m_i^{D+1} - 2\delta_j), \quad (\text{A11})$$

which is used as the starting point to estimate the microscopic learning path in the main text. Finally, by summing all the entries of this matrix and using that in the test set $\langle \delta \rangle = N^{-1} \sum_j \delta_j$, we get Eq. (5) in the main text.

Appendix B: Numerical integration of the effective system. We integrate the dynamics of Eq. (4) using a RK45 method implemented with the SciPy Python solver. At each value

of the coupling strength, we iterate for a time span of 10s, averaging the metrics during the second half of the simulation to discard transient behavior. The number of nodes (neural units) is set to $N = 200$, regularization is $\hat{\gamma} = 10^{-3}$ and results are averaged over 200 realizations. The sequence of quenched disorders δ is drawn from a Gaussian distribution $\mathcal{N}(0, 2)$ and the initial condition for the magnetization values is the uniform support in $[-2 + \epsilon, 2 + \epsilon]$, where a small ϵ is used to break the symmetry towards positive solutions of $\langle m \rangle$ and reproduce what occurs in the experiments (where the nonlinear neuron activation function always produce $\langle m \rangle > 0$ in the ordered regime).

We consider two protocols in order to show the presence of multistability and the dependence on the initial conditions. In the adiabatic protocol (dashed lines in the main text) we use the distribution of magnetization values in the units from the previous coupling value as the initial condition of the following one, while in the nonadiabatic protocol (markers) we initialize the magnetization values at each coupling strength.

Appendix C: Experimental details. Here we describe details on the dataset, on the architecture of the neural units and on the parametrization of the learning algorithm used in our experiments. The code and data to reproduce the results of this work are available on github.com/mystic-/collective-learning.

i) Dataset. We use the standard MNIST dataset to validate the effective theory. It consists of 60000 labeled black-and-white images of 28×28 pixels corresponding to handwritten digits from zero to nine, and it is divided into train set (50000) and test set (10000) for cross-validation purposes. The images are flattened as vectors to feed as inputs into the neural network (see below). In the SI (Sec. B), we consider the slightly more complex CIFAR-10 dataset to further validate our predictions.

ii) Neural architecture. We use fully-connected feed-forward neural networks as the backbone of the neural units. We remark that the current version of the algorithm can be implemented on any other neural architectural (e.g., convolutional nets, recurrent nets or transformers) as long as the number of parameters is the same in all the units. Interestingly, in the SM (Sec. B), we show that the results predicted by the effective theory also hold for convolutional neural networks, even if the theory is derived assuming a feed-forward architecture.

For feed-forward nets, for each data point $x \in \mathbb{R}^{n_0}$, we use $h^d(x) \in \mathbb{R}^{n_d}$ and $x^d(x) \in \mathbb{R}^{n_d}$ for the pre- and post-activation functions. The recurrence relation for a layer of the network is defined as

$$h^{d+1} = x^d W^{d+1} + b^{d+1}, \quad x^{d+1} = \phi(h^{d+1}), \quad (\text{C1})$$

where $\phi()$ is a point-wise, nonlinear activation function. In the main text we use the standard ‘ReLU’ function $\phi(x) = \max(0, x)$ and in the SI (Sec. A) we consider the classical ‘Sigmoid’ and ‘Tanh’ functions as well [19]. $W^{d+1} \in \mathbb{R}^{n_d \times n_{d+1}}$ is the matrix of weights between layers and $b^{d+1} \in \mathbb{R}^{n_{d+1}}$ is the vector of biases. We define Θ^d as the matrix of layer d that includes both the weights and biases, thus $\Theta^d \in \mathbb{R}^{(n_d+1) \times (n_{d+1}+1)}$. We also define $\theta = \text{vec}(\cup_{d=1}^{D+1} \Theta^d)$, as the flattened vector of all the parameters of a single neural

unit. In the main experiments, we fix the number of hidden neurons to $n_d = 20$ for all hidden layers.

iii) *Training and evaluation.* To minimize the local term in Eq. (2), we use standard Backpropagation with vanilla mini-batch gradient descent (without momentum), with batch size $B = 2^5$ and Cross-Entropy as the loss function [19], suitable for categorical outputs. The batch size is set sufficiently large [48] such that we can neglect for simplicity an additive noise term that would emerge in Eq. (4) if

SGD was used to approximate the global loss function by taking the gradient of a single data point. The learning rate is fixed to $\eta = 0.005$ and weight decay to $\gamma = 10^{-3}$. Results for each coupling value σ are time-averaged during $2 \cdot 10^4$ batch iterations, after $2 \cdot 10^4$ more iterations to discard transient behavior. We run independent realizations of the process, where all parameters are initialized with $\mathcal{N}(0, 1)$ and consider both adiabatic and nonadiabatic protocols in the process.

-
- [1] P. W. Anderson, More is different, *Science* **177**, 393 (1972).
 - [2] M. Mézard, G. Parisi, and M. A. Virasoro, *Spin glass theory and beyond: An Introduction to the Replica Method and Its Applications*, Vol. 9 (World Scientific Publishing Company, 1987).
 - [3] S. H. Strogatz, *Sync: The Emerging Science of Spontaneous Order* (Hyperion, New York, 2003).
 - [4] M. A. Muñoz, *Colloquium: Criticality and dynamical scaling in living systems*, *Rev. Mod. Phys.* **90**, 031001 (2018).
 - [5] R. Solé, M. Moses, and S. Forrest, Liquid brains, solid brains, *Phil. Trans. R. Soc. B* **374**, 20190040 (2019).
 - [6] J. Piñero and R. Solé, Statistical physics of liquid brains, *Phil. Trans. R. Soc. B* **374**, 20180376 (2019).
 - [7] D. O. Hebb, *The Organization of Behavior: A Neuropsychological Theory* (John Wiley, New York, 1949).
 - [8] M. J. Mataric, Designing emergent behaviors: From local interactions to collective intelligence, in *PSICFAA2* (MIT Press, Cambridge, MA, 1993), pp. 432–441.
 - [9] D. Ha and Y. Tang, Collective intelligence for deep learning: A survey of recent developments, *Collective Intelligence* **1**, 26339137221114874 (2022).
 - [10] A. Olsson, E. Knapska, and B. Lindström, The neural and computational systems of social learning, *Nat. Rev. Neurosci.* **21**, 197 (2020).
 - [11] J. J. Hopfield, Neural networks and physical systems with emergent collective computational abilities, *Proc. Natl. Acad. Sci. USA* **79**, 2554 (1982).
 - [12] J. Kennedy and R. Eberhart, Particle swarm optimization, in *Proceedings of ICNN'95*, Vol. 4 (IEEE, Perth, 1995), p. 1942.
 - [13] G. Carleo, I. Cirac, K. Cranmer, L. Daudet, M. Schuld, N. Tishby, L. Vogt-Maranto, and L. Zdeborova, Machine learning and the physical sciences, *Rev. Mod. Phys.* **91**, 045002 (2019).
 - [14] M. Mézard, Spin glass theory and its new challenge: Structured disorder, *Indian J. Phys.* **98**, 3757 (2024).
 - [15] Z. Liu and M. Ueda, Towards understanding grokking: An effective theory of representation learning, in *Advances in Neural Information Processing System*, Vol. 35 (NeurIPS, 2022), p. 34651.
 - [16] L. Ziyin and M. Ueda, Exact phase transitions in deep learning, *arXiv:2205.12510*.
 - [17] P. Mehta, M. Bukov, C.-H. Wang, A. G. R. Day, C. Richardson, C. K. Fisher, and D. J. Schwab, A high-bias, low-variance introduction to machine learning for physicists, *Phys. Rep.* **810**, 1 (2019).
 - [18] D. A. Roberts, S. Yaida, and B. Hanin, *The Principles of Deep Learning Theory* (Cambridge University Press, Cambridge, MA, 2022).
 - [19] I. J. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning* (MIT Press, Cambridge, MA, 2016).
 - [20] J. Jumper, R. Evans, A. Pritzel, T. Green, M. Figurnov, O. Ronneberger, K. Tunyasuvunakool, R. Bates, A. de A. Potapenko, A. Bridgland, C. Meyer, S. A. A. Kohl, A. J. Ballard, A. Cowie, B. Romera-Paredes, S. Nikolov, R. Jain, J. Adler, T. Back *et al.*, Highly accurate protein structure prediction for the human proteome, *Nature (London)* **596**, 583 (2021).
 - [21] N. Brown and T. Sandholm, Superhuman AI for multiplayer poker, *Science* **365**, 885 (2019).
 - [22] S. Bubeck *et al.*, Sparks of artificial general intelligence: Early experiments with GPT-4, *arXiv:2303.12712*.
 - [23] T. Ben-Nun and T. Hoefler, Demystifying parallel and distributed deep learning: An in-depth concurrency analysis, *ACM Comput. Surv.* **52**, 1 (2019).
 - [24] T. Li, A. K. Sahu, A. Talwalkar, and V. Smith, Federated learning: Challenges, methods, and future directions, *IEEE Sig. P. Mag.* **37**, 50 (2019).
 - [25] S. Lee, S. Purushwalkam, M. Cogswell, D. Crandall, and D. Batra, Why M heads are better than one: Training a diverse ensemble of deep networks, *arXiv:1511.06314*.
 - [26] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. Aguera y Arcas, Communication-efficient learning of deep networks from decentralized data, in *Artificial intelligence and statistics* (PMLR, 2017), pp. 1273–1282.
 - [27] P. P. Liang, T. Liu, L. Ziyin, N. B. Allen, R. P. Auerbach, D. Brent, R. Salakhutdinov, and L.-P. Morency, Think locally, act globally: Federated learning with local and global representations, *arXiv:2001.01523*.
 - [28] J. Wang and G. Joshi, Cooperative SGD: A unified framework for the design and analysis of local-update SGD algorithms, *J. Mach. Learn. Res.* **22**, 9709 (2021).
 - [29] F. Zhuang, Z. Qi, K. Duan, D. Xi, Y. Zhu, and H. Zhu, A comprehensive survey on transfer learning, *Proc. IEEE* **109**, 43 (2020).
 - [30] S. Zhang, A. E. Choromanska, and Y. LeCun, Deep learning with elastic averaging SGD, *Adv. neural inf. process. syst.* **1**, 685 (2015).
 - [31] B. E. Woodworth, K. K. Patel, and N. Srebro, Minibatch vs local SGD for heterogeneous distributed learning, *Adv. Neural Inf. Process. Syst.* **33**, 6281 (2020).
 - [32] X. Li *et al.*, A convergence theory for federated average: Beyond smoothness, in *IEEE International Conference on Big Data* (IEEE, Piscataway, NJ, 2022), pp. 1292–1297.
 - [33] B. Huang, X. Li, Z. Song, and X. Yang, FL-NKT: A neural tangent kernel-based framework for federated learning convergence analysis, *arXiv:2105.05001*.
 - [34] X. Li, K. Huang, W. Yang, S. Wang, and Z. Zhang, On the convergence of FedAvg on non-IID data, *arXiv:1907.02189*.
 - [35] C. Baldassi, C. Borgs, J. T. Chayes, A. Ingrosso, C. Lucibello, L. Saglietti, and R. Zecchina, Unreasonable effectiveness of

- learning neural networks: From accessible states and robust ensembles to basic algorithmic schemes, *Proc. Natl. Acad. Sci. USA* **113**, E7655 (2016).
- [36] C. Baldassi, F. Pittorino, and R. Zecchina, Shaping the learning landscape in neural networks around wide flat minima, *Proc. Natl. Acad. Sci. USA* **117**, 161 (2020).
- [37] A. M. Saxe, J. L. McClelland, and S. Ganguli, Exact solutions to the nonlinear dynamics of learning in deep linear neural networks, [arXiv:1312.6120](https://arxiv.org/abs/1312.6120).
- [38] C. Van den Broeck, J. M. R. Parrondo, J. Armero, and A. Hernandez-Machado, Mean field model for spatially extended systems in the presence of multiplicative noise, *Phys. Rev. E* **49**, 2639 (1994).
- [39] J. Buceta, J. M. R. Parrondo, and F. J. dela Rubia, Random Ginzburg-Landau model revisited: Reentrant phase transitions, *Phys. Rev. E* **63**, 031103 (2001).
- [40] R. Toral, C. J. Tessone, and J. M. V. P. Lopes, Collective effects induced by diversity in extended systems, *Eur. Phys. J.: Spec. Top.* **143**, 59 (2007).
- [41] N. Komin, L. Lacasa, and R. Toral, Critical behavior of a Ginzburg-Landau model with additive quenched noise, *J. Stat. Mech.* (2010) P12008.
- [42] P. Hohenberg and A. Krekhov, An introduction to the Ginzburg-Landau theory of phase transitions and nonequilibrium patterns, *Phys. Rep.* **572**, 1 (2015).
- [43] Too much coupling destroys the collective learning phase either by *collective regularization* or by an explosion of the diffusion term of Eq.(2) leading to divergences.
- [44] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevResearch.6.L042040> for the results on additional experiments with varying activation functions (Sec. A), the CIFAR-10 dataset and convolutional nets (Sec. B), and comparing the performance of the collective system against a single unit trained with global data (Sec. C).
- [45] A. Jacot, F. Gabriel, and C. Hongler, Neural tangent kernel: Convergence and generalization in neural networks, in *Advances in Neural Information Processing Systems*, Vol. 31 (NeurIPS, 2018).
- [46] C. H. Martin and M. W. Mahoney, Implicit self-regularization in deep neural networks: Evidence from random matrix theory and implications for learning, *J. Mach. Learn. Res.* **22**, 1 (2021).
- [47] A. Achille, M. Rovere, and S. Soatto, Critical learning periods in deep neural networks, [arXiv:1711.08856](https://arxiv.org/abs/1711.08856).
- [48] A. Sclocchi and M. Wyart, On the different regimes of stochastic gradient descent, *Proc. Natl. Acad. Sci. USA* **121**, e2316301121 (2024).
- [49] E. Ott and T. Antonsen, Low dimensional behavior of large systems of globally coupled oscillators, *Chaos* **18**, 037113 (2008).
- [50] C. Bick, M. Goodfellow, C. R. Laing, and E. A. Martens, Understanding the dynamics of biological and neural oscillator networks through exact mean-field reductions: A review, *J. Math. Neurosci.* **10**, 9 (2020).
- [51] G. Gottwald, Model reduction for networks of coupled oscillators, *Chaos* **25**, 053111 (2015).
- [52] L. Arola-Fernández, P. S. Skardal, and A. Arenas, Geometric unfolding of synchronization dynamics on networks, *Chaos* **31**, 061105 (2021).
- [53] M. Newman, *Networks: An Introduction* (Oxford University Press, Inc., New York, NY, 2010).
- [54] J. Gómez-Gardeñes, S. Gomez, A. Arenas, and Y. Moreno, Explosive synchronization transitions in scale-free networks, *Phys. Rev. Lett.* **106**, 128701 (2011).
- [55] P. S. Skardal, J. G. Restrepo, and E. Ott, Frequency assortativity can induce chaos in oscillator networks, *Phys. Rev. E* **91**, 060902(R) (2015).
- [56] L. Arola-Fernández, S. Faci-Lzaro, P. S. Skardal, E.-C. Boghiu, J. Gmez-Gardees, and A. Arenas, Emergence of explosive synchronization bombs in networks of oscillators, *Commun. Phys.* **5**, 264 (2022).
- [57] K. Chatterjee, D. Zufferey, and M. A. Nowak, Evolutionary game dynamics in populations with different learners, *J. Theor. Biol.* **301**, 161 (2012).
- [58] D. Gao, X. Yao, and Q. Yang, A survey on heterogeneous federated learning, [arXiv:2210.04505](https://arxiv.org/abs/2210.04505).
- [59] J. Gómez-Gardeñes, D. Soriano-Paños, and A. Arenas, Critical regimes driven by recurrent mobility patterns of reaction-diffusion processes in networks, *Nat. Phys.* **14**, 391 (2018).
- [60] R. Vinuesa, H. Azizpour, I. Leite, M. Balaam, V. Dignum, S. Domisch, A. Felländer, S. D. Langhans, M. Tegmark, and F. Nerini, The role of artificial intelligence in achieving the sustainable development goals, *Nat. Commun.* **11**, 233 (2020).
- [61] github.com/mystic-blue/collective-learning.