

Training Coupled Phase Oscillators as a Neuromorphic Platform using Equilibrium Propagation

Qingshan Wang and Clara C. Wanjura

Max Planck Institute for the Science of Light, Staudtstraße 2, 91058 Erlangen, Germany

Florian Marquardt

Max Planck Institute for the Science of Light, Staudtstraße 2, 91058 Erlangen, Germany and
Department of Physics, University of Erlangen-Nuremberg, 91058 Erlangen, Germany

(Dated: February 14, 2024)

Given the rapidly growing scale and resource requirements of machine learning applications, the idea of building more efficient learning machines much closer to the laws of physics is an attractive proposition. One central question for identifying promising candidates for such neuromorphic platforms is whether not only inference but also training can exploit the physical dynamics. In this work, we show that it is possible to successfully train a system of coupled phase oscillators—one of the most widely investigated nonlinear dynamical systems with a multitude of physical implementations, comprising laser arrays, coupled mechanical limit cycles, superfluids, and exciton-polaritons. To this end, we apply the approach of equilibrium propagation, which permits to extract training gradients via a physical realization of backpropagation, based only on local interactions. The complex energy landscape of the XY/ Kuramoto model leads to multistability, and we show how to address this challenge. Our study identifies coupled phase oscillators as a new general-purpose neuromorphic platform and opens the door towards future experimental implementations.

I. INTRODUCTION

We are witnessing the widespread adoption of machine learning in all branches of science and technology. At the same time, it becomes clear that the energy costs for both training and inference are exploding in an unsustainable way. This is especially apparent for the most powerful recent innovations like large-language models and diffusion models for image generation, which contain billions of parameters. Even the impressive advances in deploying specialized hardware like graphical processing units or tensor processing units are not able to keep up with the rapidly rising resource requirements. This makes it all the more urgent to explore alternative means for computation in this domain. The hope that such alternatives exist rests on several observations. First, machine learning is inherently stochastic and it therefore seems wasteful to run implementations on digital hardware that was originally conceived for implementing mathematical algorithms where precise answers matter. In addition, one can exploit the specific structure of machine learning problems to design new hardware platforms that are no longer general purpose. In particular, one can overcome the inefficient separation between processing and memory, the so-called von-Neumann bottleneck. Third, many physical platforms naturally allow for highly parallel processing.

These are the motivations for the rapidly growing field of neuromorphic computing [1, 2], which explores a large variety of different physical platforms, with the aim of designing more energy-efficient parallel processing devices, operating much closer to the underlying physics than a typical digital processor. Neuromorphic implementations do include alternative digital electronic platforms that overcome the von-Neumann bottleneck, but there is even more fundamental research to be done in the domain of analog, more directly physics-based approaches. These range from solid-state physics (building on components like memristors [3], Josephson junctions [4, 5], or spin oscillators [6]) and optics [7–12]

(free space or integrated photonics) to other domains, even including soft-matter physics (e.g. [13–16]).

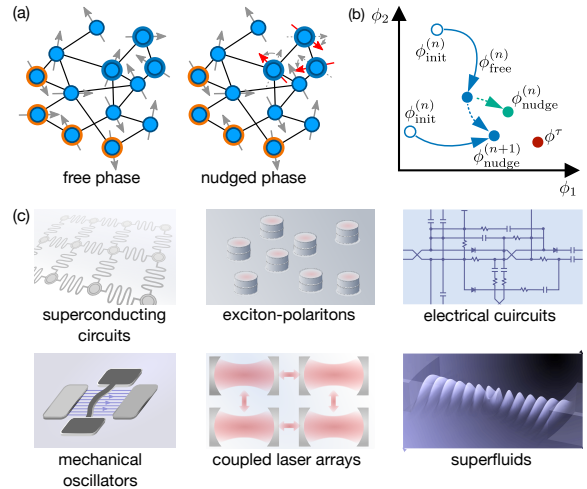


FIG. 1. Equilibrium propagation for coupled phase oscillators. (a) We train a neuromorphic system of coupled phase oscillators (units), here represented by blue circles. In the free phase (left), the input units (orange) are fixed and the system relaxes to its equilibrium via the standard overdamped dynamics of coupled phase oscillators. Output units are highlighted by a blue rim. In the nudge phase (right), the output units are forced weakly towards the desired target output (red arrows) and the system evolves to a new, slightly shifted equilibrium. The difference between these two equilibria allows to infer the gradients needed to train the system, according to the fundamental rule of equilibrium propagation, see Eq. (7). (b) Schematic representation of the evolution in the space of oscillator phases ϕ_j , with open circles depicting initial conditions and full circles depicting equilibria (ϕ^τ is the desired target configuration). (c) Possible experimental platforms that can give rise to the dynamics of the XY/ Kuramoto model of coupled phase oscillators include superconducting circuits, nonlinear electrical circuits, exciton-polariton systems, mechanical-oscillator limit cycles, coupled laser arrays and superfluids.

A neuromorphic platform usually has to exhibit nonlinear dynamics to be powerful enough in terms of its expressivity, i.e., its ability to approximate arbitrary input-output relations. Arguably one of the strongest generic classical nonlinearities is found in phase oscillators. These are systems where a combination of energy pump and nonlinearity establishes a limit cycle, with the phase of the oscillation as the resulting continuous degree of freedom. The laser is the most prominent example. When several such phase oscillators are coupled, they exhibit the dynamics of phase locking and synchronization. Theoretically, these systems are described by the so-called Kuramoto model [17], which is equivalent to the relaxation dynamics of an XY model [18], one of the most widely studied models in statistical physics. Coupled phase oscillators have been explored in a variety of platforms, including laser arrays [19, 20], coupled mechanical limit cycles in optomechanical [21, 22] and nano-electromechanical [23] systems, cold atoms [24], Josephson arrays in superconducting circuits [25], spin oscillators [26, 27], and exciton-polaritons [28–30].

In this article, we explore how a system of coupled phase oscillators can serve as a neuromorphic platform. Most importantly, we will investigate how it can be trained in a physics-based approach.

An array of coupled phase oscillators was first considered as a neuromorphic platform already in [31, 32]. In that case, the system was proposed as an associative Hopfield memory, where couplings can be determined in a single step based on a simple Hebbian rule. In contrast, in our case, we will be interested in the more general setting of supervised learning, where the goal is to learn a prescribed input-output relation, for regression or classification tasks – the most wide-spread application in machine learning. That such a goal might be achievable is indicated by the results in [33], where certain small architectures of coupled phase oscillators were hand-crafted to address simple tasks. More generally, coupled oscillators have been proposed in various forms as computing platforms (see the review in [34]), and the training of coupled digital oscillators has also been investigated recently in simulations [35]. From the neuromorphic perspective, oscillators seem like a natural choice for machine learning, since the brain is based on spiking neurons that can also show behaviour such as synchronization.

One particular challenge in neuromorphic physical platforms is the question of efficient training. The conceptually easiest, general approach that can be adopted in every case is the parameter shift method, where internal trainable parameters of the physical device are adjusted in a feedback loop based on the distance to the desired output (e.g. [36, 37] and many others). However, this scales very unfavourably, requiring for a single training step a number of evaluations rising linearly in the number of parameters. Another popular approach is reservoir computing [38–41], where the fixed (non-adjustable) nonlinear dynamics of a system is used to map the input to a higher-dimensional space which is then turned into the output via a trainable simple (possibly even linear) neural network, similar to kernel methods in machine learning. This has been very successful, but still relies on a digital neural network. One major goal in the field is to find ways to impli-

ment variants of backpropagation for neuromorphic devices, since this efficient method to find the gradients needed for training is the cornerstone of virtually all machine learning. One possibility consists in doing this via simulations (“digital twins”). This method suffers from the lower speed of simulations and the potential inaccuracy of the adopted model, although the latter constraint has recently been overcome via an ingenious hybrid method [42], where backpropagation happens in the simulation but the forward pass is performed on the real device. Physical implementations of backpropagation, exploiting the physical dynamics itself, are a conceptually very attractive goal. They have been occasionally put forward starting even in the 80s. The first ideas involved approaches that were not general purpose but relied on specific properties of specially engineered platforms, especially in optics [7, 43–45]. A more general type of physical backpropagation for an integrated-photonics platform was introduced and eventually implemented recently in [12, 46]. In certain cases, gradients can also be extracted directly from scattering experiments, as shown for the recently introduced nonlinear neuromorphic processing in linear wave scattering platforms [47]. However, only two generic methods exist for physically performing backpropagation in a wide class of systems. One of them is Hamiltonian Echo Backpropagation, which relies on the Hamiltonian dynamics of arbitrary time-reversal-invariant systems [48]. The other is Equilibrium Propagation [49], which encompasses arbitrary energy-based systems that relax to some equilibrium. Since coupled phase oscillators display relaxation-type phase dynamics, in this work we will employ Equilibrium Propagation.

We start by explaining how Equilibrium Propagation (EP) can be applied to systems of coupled phase oscillators. We then explore in numerical experiments the EP-based training of networks of oscillators in two illustrative examples (XOR and handwritten-digit recognition). We specifically point out the effects of multistability on training and how to account for them. Our results establish coupled phase oscillators as a viable neuromorphic platform for supervised training, although the choice of network architecture plays an important role.

II. THE XY MODEL AND EQUILIBRIUM PROPAGATION

In this section, we first describe our model and then briefly review the theory of EP and how to apply it in our setting. We address the problem of multistability, which arises especially in the very nonconvex energy landscape of coupled oscillators.

A. XY model of coupled phase oscillators

The dynamics of coupled phase oscillators can be described by some variant of the Kuramoto model, and this can be understood (in many cases) as overdamped relaxation dynamics in the energy landscape of an XY model. An XY model consists of interacting classical 2D spins (in our case, the phase oscillators), which are described by a set of phase angles ϕ_i .

Anticipating the neuromorphic application, we refer to the coupling between the spins as **weights** and assume that all weights are trainable. We adopt the following energy function

$$E(\phi; W, h, \psi) = -\frac{1}{2} \sum_{ij} W_{ij} \cos(\phi_i - \phi_j) - \sum_i h_i \cos(\phi_i - \psi_i). \quad (1)$$

The first term describes coupling between arbitrary oscillators, with coupling constants (weights) W_{ij} . The second term describes a tendency towards phase locking to some external oscillatory drives, where h_i is set by the amplitude of the respective drive and ψ_i is the phase of the drive. In the machine-learning context, these could be understood in analogy to the bias terms of a neural network. Beyond that, we remark that, in the synchronization context, one would normally expect to have another set of terms, of the type $\Omega_i \phi_i$. This would denote the fact that different oscillators are running at different intrinsic frequencies. When these terms are large, a transition towards desynchronization will take place. However, for the neuromorphic application we are interested in, we find that such terms – and the associated tendency towards desynchronization – are not helpful, as they may lead to the absence of stable fixed points. We therefore envisage a situation where all the oscillators’ intrinsic frequencies are the same. Small deviations from this ideal scenario can be tolerated, since the coupling terms would help to generate stable fixed points.

In this article, we focus on the deterministic dynamics of the system. We obtain the equilibrium by solving the following ODE, assuming that the phase angles are driven by the negative energy gradient

$$\begin{aligned} \dot{\phi}_i &= -\frac{\partial E}{\partial \phi_i}(\phi; \theta) \\ &= -\sum_j W_{ij} \sin(\phi_i - \phi_j) - h_i \sin(\phi_i - \psi_i). \end{aligned} \quad (2)$$

It is obvious from this equation that we took E (as well as the couplings and the biases) to have the dimensions of a frequency, to avoid having to introduce a separate mobility constant in front of the gradient. These are the equations of motion for a standard Kuramoto model [17] modified by the bias term $\sum_i h_i \cos(\phi_i - \psi_i)$. In this sense, Eq. (2) describes the dynamics of a perfectly synchronized network of oscillators with the same intrinsic frequency. Here and in the following, θ denotes the set of all the trainable parameters—in our case, $\theta = (W, h, \psi)$. As explained in the introduction, such a system can be realized in many physical platforms.

B. Equilibrium Propagation

EP was introduced in [49, 50] as a way to perform physical backpropagation, i.e., to extract training gradients based on the physical dynamics of a system undergoing relaxation towards a minimum-energy state. In a nutshell, it involves relaxing the system under two different choices of boundary conditions (“free” and “nudge”), comparing the resulting states locally and performing parameter updates via some feedback

mechanism (see below). As such, EP stands in the tradition of contrastive learning approaches which have been a well-known part of machine learning since the 80s [51], but importantly EP offers an implementation of those general mathematical ideas in terms of local physical interactions. In recent years, EP has been explored further in various directions. A variant termed “coupled learning” has been introduced [52] (see [53] for a very useful comparison of the different versions of such energy-based training approaches). EP has been proposed for the training of nonlinear resistor networks [54], and it has also been extended to deal with spiking networks [55]. The effects of the EP/ coupled-learning training dynamics on the response behaviour of a neuromorphic system was analyzed carefully in [56]. In addition, for EP, it was realized that it is possible to implement continual parameter updates already during the second phase of the equilibration dynamics, leading to an approach that is local in time [57]. Going an important step further, recently, a dynamical version of EP has been proposed [58] (see also [59]), where in addition to gradient extraction via physical dynamics there is also a way to perform the parameter update itself via physical dynamics, rather than via measurement and feedback. This development can be seen in the conceptual tradition of Hamiltonian Echo Backpropagation [48], where a general way to perform training updates via physical dynamics in neuromorphic systems was first introduced.

On the experimental side, EP and its variants have been implemented so far in only a few devices—almost all of them electronic systems. One set of demonstrations investigated linear and nonlinear resistor networks in a platform using two copies of the network working in tandem, trained using coupled learning [60–63]. An experimental application to elastic networks was recently presented in [16]. Furthermore, a classical Ising model has been trained using the ideas of EP and making use of a quantum annealer to efficiently reach equilibrium [64]. Probably the most technically advanced platform consists of a memristor crossbar array trained using a contrastive learning approach closely related to EP [65, 66].

In general, EP is applied to energy-based models, i.e., systems with an energy function. We will now describe its workings, already keeping in mind the system of coupled phase oscillators. During evaluation (inference), the input is injected by fixing the phases of some selected oscillators to the components of the input vector, e.g., representing the color values of a picture. This could as well be done by (sufficiently strong) external driving, which will generate the same tendency towards phase locking already present in the bias terms. The system then equilibrates via its natural dynamics, i.e., it relaxes to the lowest-energy state. The output is finally read off by observing a selected set of oscillators and taking their phases to represent the output vector.

We call the vector of input phases ϕ_{in} , while the phase angles of the output units and the hidden units are denoted as ϕ_{out} and ϕ_{hidden} , respectively. The energy function can then be written as $E(\phi_{\text{in}}, \phi_{\text{hidden}}, \phi_{\text{out}}; \theta)$. In the context of EP, we call this the **internal energy**.

During training, we want to measure the deviation of the actual output to the desired output (the target). In the con-

text of phase oscillators, a simple possibility is the following **distance function**. It simply compares the normalized spin vectors \mathbf{s} , leading to a cosine-similarity measure, with the desired output values (target) denoted as ϕ_i^τ :

$$D(\phi, \phi^\tau) \equiv \sum_{i \in O} \frac{1}{2} |\mathbf{s}_i - \mathbf{s}_i^\tau|^2 = \sum_{i \in O} 1 - \cos(\phi_i - \phi_i^\tau). \quad (3)$$

Training via EP requires the system to have another weak contribution to the energy function which can be switched on when needed. This component describes the interaction between the output and the desired target. It plays the role of a cost function in supervised learning and is called **external energy**. The most straightforward ansatz for the cost function would be to employ the distance function introduced above. However, we found that this generates problems, since even incorrect solutions (for which phases differ by π) are fixed points—albeit unstable. Empirically, the following cost function works much better, as it eliminates these undesired unstable fixed points and quickly drives the system away from incorrect solutions

$$C(\phi, \phi^\tau) = \sum_{i \in S_{\text{out}}} -\ln(1 + \cos(\phi_i - \phi_i^\tau)). \quad (4)$$

We define S_{in} , S_{h} and S_{out} as the sets of input, hidden and output units, respectively. Although the logarithm appearing here was introduced empirically, it can be motivated by the connection to the categorical cross-entropy (Kullback-Leibler divergence) found for classification tasks. Interestingly, the cost function adopted here can also be related to fidelities occurring when comparing different quantum states, despite the classical setting we are working in, as shown in Appendix A.

Once a cost function has been assigned, EP introduces a small parameter β which quantifies the strength of the interaction between the output units and the target. This leads to the **total energy** of the system

$$F(\phi, \theta, \beta) = E(\phi_{\text{in}}, \phi_{\text{hidden}}, \phi_{\text{out}}; \theta) + \beta C(\phi_{\text{out}}, \phi^\tau). \quad (5)$$

The process of EP [49] contains the following steps:

1. Initialize the system such that the input units are fixed to the provided input values.
2. Set $\beta = 0$ and obtain a stable fixed point/ equilibrium by following the relaxation dynamics

$$\dot{\phi}_i = -\frac{\partial F}{\partial \phi_i}(\phi, \theta, \beta = 0) \quad (6)$$

towards the fixed point. This state is called the **free equilibrium** and this step is called the **free phase** of EP.

3. Switch on a small value of β and again follow the dynamics towards a new fixed point. This state is called **nudge equilibrium** and this step is called the **nudge phase** of EP.

4. Repeat the process 1-3 for all the input data from the given batch of training samples. The central insight of EP [49] is that the gradient of the cost function with respect to the trainable parameters can be approximated as

$$\begin{aligned} \frac{\partial C}{\partial \theta_\alpha} &= \frac{d}{d\beta} \left(\frac{\partial F}{\partial \theta_\alpha} \right)_{\beta=0} \\ &\approx \frac{1}{\beta} \left(\left\langle \frac{\partial E}{\partial \theta_\alpha} \right\rangle_{\phi_{\text{in}}}^{\text{nudge}} - \left\langle \frac{\partial E}{\partial \theta_\alpha} \right\rangle_{\phi_{\text{in}}}^{\text{free}} \right). \end{aligned} \quad (7)$$

in which, in our case, $\frac{\partial E}{\partial \theta_\alpha}$ are given by the following analytical expressions

$$\begin{aligned} \frac{\partial E}{\partial W_{ij}} &= -\cos(\phi_i - \phi_j) \\ \frac{\partial E}{\partial h_i} &= -\cos(\phi_i - \psi_i) \\ \frac{\partial E}{\partial \psi_i} &= -h_i \sin(\phi_i - \psi_i). \end{aligned} \quad (8)$$

It has been discovered that EP gradients may benefit from nudging the system away from the desired target state, i.e. $\beta < 0$, or even estimating the gradient in a symmetric way from positive and negative nudging [67]. However, it was not possible to use this trick in our work due to the fact that our cost function can become large quickly when nudged in the 'wrong' direction. We implement the dynamics of the equilibration numerically using the 4th order Runge-Kutta method with adaptive step size, implemented in python/ jax on a GPU.




FIG. 2. **The problem of multi-stability.** (a) A case with only one stable fixed point. The circles refer to the initial states and the arrows refer to the relaxation to the stable fixed points (equilibria; solid dots). The arrows with dashed lines refer the evolution of the equilibria during the training. The dark red dot refers to the target output, assuming that the two phase coordinates shown here both refer to output units, i.e., hidden phases and input phases are suppressed in this depiction. Different colors refer to different iterations during the training. (b) Multistability. Different equilibria can be reached at the same training iteration (same color), starting from different random initial conditions.

As we focus on the deterministic dynamics of a system with a complicated non-convex energy landscape, a discussion of multistability is inevitable. Apparently such a discussion was not needed in earlier applications of EP, due to the simpler nature of the energy landscape encountered there (e.g. [54]).

An XY model with random couplings can be viewed as a spin glass, for which [68] shows that the number of fixed points grows exponentially with the size of the network. Although only a small fraction of these fixed points are stable, the effect of possible multistability in this energy landscape cannot be neglected, as it can undermine the training if it is not properly handled. We emphasize that on top of this multistability of the energy landscape (peculiar to energy-based models) which hinders equilibration for each training step, there can also be multistability in the cost function training landscape (similar to neural networks), which is a distinct aspect.

The effect of multi-stability is shown schematically in Fig. 2. In panel (a), we assume that the system has only one stable fixed point. The fixed point then moves along the dashed arrows and approaches the desired result (the target) by adjusting the trainable parameters with EP. However, if there are multiple stable fixed points, the situation becomes more complicated, as shown in panel (b). Depending on the arbitrarily chosen initial phase configuration one may end up in different fixed points. Depending on which training updates are performed in which fixed points, this may even result in situations where some of the fixed points move away from the target during training, at least temporarily. Sometimes there are even more drastic cases as a stable fixed point may become unstable after the parameters have been adjusted. These cases can reduce the efficiency of the training.

In principle, the multistability problem can be solved by introducing noise into the system and rewriting Eq. (7) with a Boltzmann distribution (the theory is briefly discussed in Ref. [50]). However, we note here that the challenge of multistability can also be solved within the framework of deterministic dynamics. We propose to randomly initialize the hidden (and output) units for each new training step. The system will correspondingly end up in different fixed points, depending on the location of the randomly selected initial conditions with respect to the basins of attraction of those fixed points. Then, in the spirit of Stochastic Gradient Descent (SGD), we average the learning gradient over both the random initial configurations and the training samples. Based on the success of SGD in other contexts, we expect that all stable fixed points can be trained simultaneously (and our numerical experiments confirm this).

In practice, in each iteration we randomly initialize the system many times and average the approximate gradient given by EP over both the training sample batch and the initial states. We denote the number of random initial values by M_{init} and the size of the training batch by M_{data} , and we update the trainable parameters accordingly

$$\delta\theta_\alpha = -\frac{\eta}{\beta} \frac{1}{M_{\text{data}}} \sum_{i=1}^{M_{\text{data}}} \frac{1}{M_{\text{init}}} \sum_{j=1}^{M_{\text{init}}} \left(\frac{\partial E}{\partial \theta_\alpha}(\phi^{i,j,\text{nudge}}; \theta) - \frac{\partial E}{\partial \theta_\alpha}(\phi^{i,j,\text{free}}; \theta) \right). \quad (9)$$

Here $\phi^{i,j,\text{nudge/free}}$ refers to the nudge/ free equilibrium for the i th input attained from the j th initial state, and η refers to the learning rate.

III. RESULTS

A. Learning the XOR function

We first test EP-based training of a coupled-oscillator network for the simple case of the XOR function. For artificial neural networks, the XOR problem is historically important, because it was the first task to require at least one hidden layer. Due to its simplicity, the XOR task allows us to investigate the dynamical properties of the network in great detail. We are able to explore the effects of the network size, the evolution of equilibria when the network is multi-stable, and the influence of random initial configurations.

We encode “True” and “False” with $\phi = \pi/2$ and $\phi = -\pi/2$, respectively. For XOR, if the two inputs are same, the output unit should give us a “False” and vice versa.

At the beginning of training, the weights are initialized independently according to a standard Gaussian distribution, while the strength and the angle of the bias are initialized according to the uniform distribution over $[-0.5, 0.5)$ and $[-\pi, \pi)$, respectively. During the training, we repeatedly update the parameters according to Eq. (9). In each iteration, the system is exposed to all 4 input-output pairs ($M_{\text{data}} = 4$). For each such pair, we repeat the following process for M_{init} times: the units of the network are initialized so that the phase angles of the hidden and output units independently follow the uniform distribution over $[-\pi, \pi)$; then we solve Eq. (2) for a fixed time interval T which is sufficiently long to ensure that the coupled phase oscillators approximately reach their equilibrium; finally, we calculate the average parameter gradient and update the trainable parameters using Eq. (9). In this section, we fix $T = 100$ and $\eta = 0.1$.

All our numerical results for the XOR task are collected in Fig. 3. We first set the number of initial configurations to $M_{\text{init}} = 1$ and train a 5-unit XY network with all-to-all connectivity for 1000 iterations. In each iteration, the mean distance $\langle D \rangle$ averaged over the 4 input-output pairs (computed with Eq. (3)) is recorded. In addition, we perform averages over different training runs with random initial weight configurations. The average distance is shown in panel (a) of Fig. 3. The white line depicts the average mean distance, which we define to be $\langle \overline{D} \rangle = \frac{1}{N_{\text{train}}} \sum_{i=1}^{N_{\text{train}}} \langle D(\phi, \phi^\tau) \rangle$, where the average is taken over training runs. We find that the distance almost always converges to 0, which proves the ability of XY networks to learn XOR via EP training.

The effect of size N on learning speed is shown in panel (b). For different N , we initialize and train an all-to-all network 100 times and plot the evolution of the mean distance function $\langle \overline{D} \rangle$. The results in (b) show that an increasing size initially accelerates the learning process. The speed of learning reaches its maximum at a certain number of phase oscillators and then decreases. The initial increase in speed can be explained by the increasing number of parameters, which makes the network more flexible. However, when the size becomes excessively large, the large number of parameters and multiple stable fixed points (all of which have to be trained) will reduce the efficiency of the training. According to the result




FIG. 3. Results of training an XY model on the XOR task. (a) Statistics of the training progress of 5-unit XY networks of all-to-all connectivity (the structure is shown in the upper right corner). Evolution of mean distance function $\langle D \rangle$ (white line, average over 1000 training runs). The density plot indicates the histogram of $\log_{10} D$ over these runs. (b) The dependence of the learning speed on the size N . For a given N , we randomly initialize the trainable parameters and train them for 1000 iterations, repeating this whole process 100 times. The curves show the evolution of $\langle D \rangle$ during the training. (c) Evolution of equilibria during a single training run. For each of the four input configurations, the possible stable fixed points at each training iteration are found by running the system with random initialization for 100 times. For each of these fixed points the state of the output unit (upper row) and one of the hidden units (lower row) are recorded. Output units for all fixed points converge to a unique, deterministic function of the input at the end of training. (d) Dependency of learning speed on the number of random initial configurations, M_{init} . For each N and M_{init} we show the evolution of $\langle D \rangle$ averaged over 100 training runs. (e) Upper panel: the learning speed of training measured by the average slope of the first 300 steps shown in (d) and its dependence on M_{init} . Lower panel: the “physical training speed”, measured by the slope divided by M_{init} .

in (b), the training becomes slower again when the number of units exceeds 8.

We now turn to a more detailed investigation of multistability [panel (c) of Fig.3]. We train a network of 30 units with 2000 iterations and record the parameters for each iteration. After training, for each iteration and each input-output pair, we search for all the possible equilibria by randomly initializing and running the network 100 times. We then inspect the phase angles for each of these equilibria. In the top row of panel (c), we show the evolution of the set of phase angles for an output unit. Importantly, training converges in the sense that the output unit phase angle becomes a deterministic function of the input, with no remaining scatter. At the same time however, multistability in the energy landscape remains. We can see this by observing the evolution of one selected hidden unit (bottom row), which shows scatter even after conclusion of the training. Overall, we confirm that despite the persistence of multistability, it is still possible to adjust all of the equilibria simultaneously to perfectly attain the training objective.

Finally, we study the effect of the number of random initial configurations, M_{init} , considered during each training iteration. First, for a fixed number of units N and M_{init} , we train an XY network with 100 different sets of initial trainable parameters and record the average mean distance function $\langle D \rangle$. The results are shown in panels (d) and (e). From panel (d), we observe that the learning speed increases with increasing M_{init} and saturates at large M_{init} . The effect of M_{init} is not obvious for networks of 5 units, as multi-stability is rare for small networks.

Another observation is that the average mean distance function $\langle D \rangle$ always decays approximately exponentially in the early stages of training. Therefore, we perform a linear regression to $\log \langle D \rangle$ on the first 300 iterations and use the negative slope to measure the speed of learning. The results are shown in the upper graph of panel (e). An increase in the learning speed and a saturation at large M_{init} are evident. However, a larger M_{init} physically means more equilibration runs. This results in larger time consumption in the physical world. To take this into account, we divide the learning speed extracted above by M_{init} and use this ratio to measure the physical speed of learning (“normalized convergence speed” in the figure). The results are shown in the bottom graph. We observe that the physical learning speed reaches its maximum at a relatively small value of M_{init} . This indicates that a small number of different initial configurations can already sufficiently accelerate the learning process in terms of reducing the physical time consumption.

B. Handwritten-Digits Recognition

In the previous section we showed the ability of XY networks to learn the simple XOR task via EP. In this section, we will analyse the performance on the recognition of written digits, a more complex task which has long been considered a benchmark for neural networks.

To keep the numerical effort in our simulations manageable,

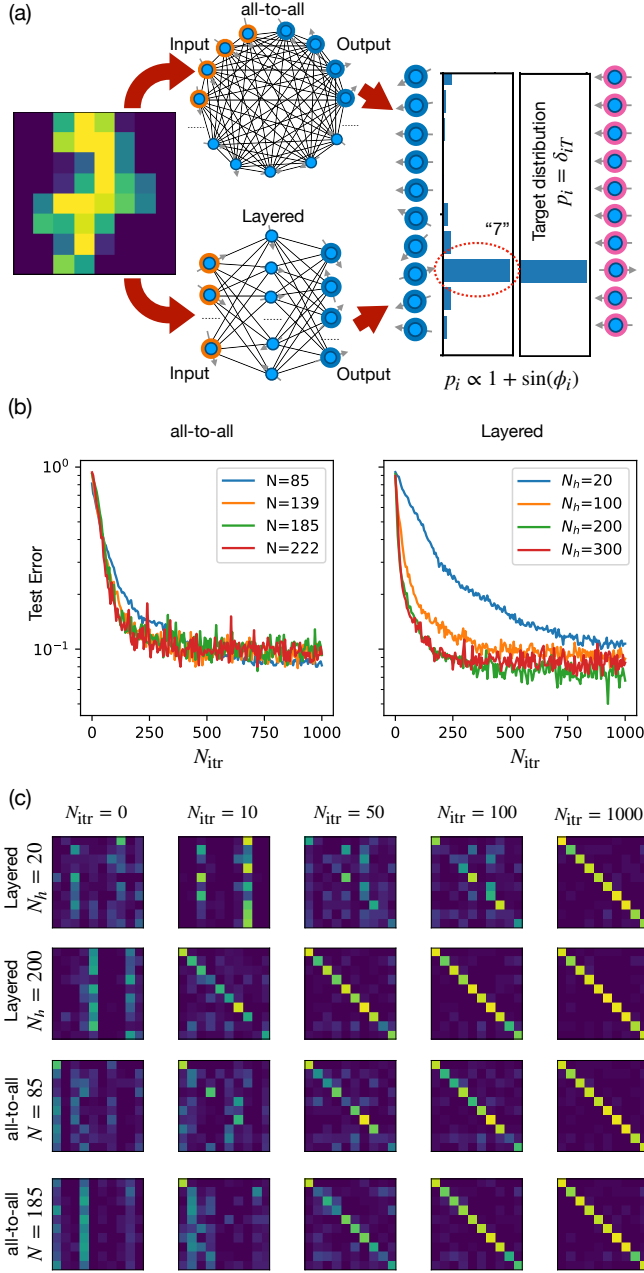


FIG. 4. **Handwritten-digit recognition in a network of coupled phase oscillators.** (a): We consider both all-to-all and layered connectivity networks, where the recognized digit is one-hot encoded in the phase configuration of the 10 output units. (b): Training curves for networks of different sizes with all-to-all connectivity and layer structure. The sizes of the networks are selected such that the number of trainable parameters for them are approximately the same (shown in I) (d): Training evolution of confusion matrices, for different structures and different network sizes.

we use a scaled-down version of the MNIST dataset. Specifically, we use the dataset provided by ScikitLearn which is a copy of the test set for the Optical Recognition of Handwritten Digits from the UCI ML hand-written digits datasets

[69]. It consists of 1797 8×8 -pixel greyscale images of written digits, approximately 180 for each. A sample figure is shown in the left of panel (a) of Fig 4. The size of the data set and the dimension of the pictures are all smaller than the commonly used MNIST, which allows us to do training via EP in smaller coupled-oscillator networks and test various architectures more easily. The structure of the networks and the results are shown in Fig 4.

As shown in Fig 4(a), we train XY networks with both all-to-all connectivity and a layered structure. In each such network, there are 64 input units and 10 output units. The input data (image pixels) are linearly rescaled into the range $[-\pi/2, \pi/2]$. As usual for classification tasks, the outputs are interpreted as a probabilities, such that $p_i \propto 1 + \sin \phi_i$, as shown in the right side of panel (a), where p_i denotes the probability that “the input digit is i ”. The inference result is given by taking the index of the highest probability. We note that in our setting these probabilities are not normalized, i.e. they do not sum up to one. This is different from usual artificial neural networks, where one uses a softmax activation function in the last layer to ensure normalization. The label for each input image is encoded in a “one-hot” configuration, giving a phase of $\pi/2$ at the output unit (phase oscillator) corresponding to the correct label of the digit seen in the image and $-\pi/2$ otherwise. We use the same cost function as for the XOR task.

For neural networks of large size, the influence of weight initialization on training can be crucial[70]. In this section, we initialize our networks as following:

1. For all-to-all networks, the weights are initialized according to $\mathcal{N}(0, 1/\sqrt{N})$, where \mathcal{N} refers to a normal Gaussian distribution and N is the total number of units in the network. The strength of the initial bias field is set to be 0 everywhere, while the initial bias directions are randomly set according to a uniform distribution over $[-\pi, \pi]$.
2. For layered networks, the initial weights connecting the i th and the $i+1$ th layers are initialized according to $\mathcal{N}(0, 1/\sqrt{N_i + N_{i+1}})$, where N_i and N_{i+1} are the size of the i th and the $i+1$ th layer. The biases are initialized in the same way as for all-to-all networks.

In order to evaluate the effect of the architecture, we compare the training process of all-to-all and layered XY networks with a similar number of parameters. For an all-to-all network of N units containing N_{in} input units and N_{out} output units, the number of trainable parameters (weights, strength and direction of bias) is $\frac{1}{2}N(N-1) - \frac{1}{2}N_{in}(N_{in}-1) + 2(N-N_{in})$. For a network of layered structure, the number of trainable parameters is $\sum_{i=0}^{L-1} N_i N_{i+1} + 2 \sum_{i=1}^L N_i$, where L here denotes the depth of the network (number of layers). The focus of our work is on the performance of all-to-all networks and layered networks with one hidden layer. The sizes and parameters of the tested networks are displayed in Tab. I.

The results are shown in panels (b) and (c) of Fig. 4. The data set is split into a training set and a test set. The training set consists of the first 100 images for each of the 10 digits

Layered		All-to-All	
Layer Structure	# Parameters	# Units	# Parameters
64, 20, 10	1540	85	1596
64, 100, 10	7620	139	7725
64, 200, 10	15220	185	15246
64, 300, 10	22820	222	22831

TABLE I. Sizes and parameters of tested networks

(1000 images in total), while the test set consists of the next 70 images for each digit. All networks are trained for 1000 iterations. In each iteration, we randomly select 30 images from the training set for each digit (so the total batch size is 300). We take the number of random initial phase configurations for each sample, M_{init} , to be one. This can be justified by our insights from the XOR task [panel (f) in Fig. 3] and the fact that a large batch size, such as the one adopted here, automatically leads to a variety of different initial configurations. Every 5 iterations, we test the accuracy of the inference with all the images in the test set.

The evolution of the test errors is shown in Fig. 4 (b) and we list the best attained test accuracy for each architecture in Tab. II of Appendix B in comparison to results we obtained with artificial neural networks (ANNs) and linear classifiers. Increasing the number of parameters (by increasing the number of hidden units) seems to improve the test accuracy for the layer architecture although the accuracy again seems to drop as we increase the number of hidden units from 200 to 300. The best test accuracy we obtained with the layer architecture (at 200 hidden units) was 94.1 % which is above the accuracy of 90.7 % achieved by a linear classifier with the same number of parameters but below the accuracy of 95.0 % achieved by an ANN with one hidden layer (see Appendix B for further details). In contrast, the performance of a network with all-to-all connectivity remains approximately constant as we increase the number of parameters with a slight drop as the number of hidden units increases from 11 to 75. The best accuracy we achieved with all-to-all connectivity was 93.3 % with 11 hidden units. The comparable linear classifier achieves 90.4 % accuracy, while the ANN attains 94.3 %. This indicates that EP can be used successfully to train a network of coupled phase oscillators to perform standard classification tasks. The performance does not reach the state-of-the-art values of artificial neural networks, which have been highly optimized over the past two decades for the MNIST task (reaching accuracies close to 100% on the larger 28×28 images), but this was to be expected – further optimization and evolution of the neuromorphic platform (as well as the details of the training procedure) would certainly yield improved accuracy. In addition, we recall that the motivation for adopting neuromorphic platforms is to achieve gains in energy efficiency and other aspects like inference speed. Therefore, one may be willing to pay the price of reduced accuracy in certain applications.

The influence of the network architecture on convergence is also evident. For all-to-all networks, the curves depicting the training evolution of the test error almost overlap for the different choices of network size. In contrast, for layered networks, convergence speed increases significantly as the size

of the hidden layer increases. However, as the size of the hidden layer increases above 200, the speed of convergence tends to saturate. This difference suggests that the speed of learning is more sensitive to the size of the networks when they are layered.

Finally, we visualize the training evolution via the confusion matrix, which measures the probability that a given true digit is classified as one of the ten possible digits (correct classifications are on the diagonal). For different network sizes, we plot the confusion matrix after 0, 10, 50, 100 and 1000 iterations during training. Comparing the results of the same network architecture, we see a clear influence of size on the learning speed for layered networks, which is negligible for all-to-all networks. Comparing the the results for the a similar number of trainable parameters, we find that the learning process of layered networks is slower at smaller sizes, but faster when the size is sufficiently large.

In terms of network architectures, it is a natural question to ask whether a locally coupled network can also yield good training performance. This is an important question, since physical interactions are naturally local, so a neuromorphic device based on such an architecture would require less physical resources. We have explored this question and investigated nearest-neighbor-coupled square lattices of phase oscillators. In that setting, there are various choices to be made, e.g. as to how the input and output units are distributed over the lattice. In any case, despite our best efforts, we could not make that architecture train successfully on the MNIST task. Intuitively, we believe this is due to the difficulty for the system to find ways to effectively process different parts of the input that are spatially removed from each other (and from the output units). This difficulty may also be related to the physics of localization in disordered systems. More detailed investigations here would be needed in the future, since the performance of locally coupled neuromorphic devices in non-trivial tasks is an important question, even beyond the coupled phase oscillator platform discussed here.

IV. CONCLUSIONS

Our results confirm that it is possible to employ equilibrium propagation for training systems of coupled phase oscillators. As an important aspect, we observed that the complex energy landscape of the XY model leads to multistability. We showed that it is possible to address that challenge via stochastic initialization, leading to a simultaneous adjustment of all fixed points and eventually stable convergence of the training. In terms of physical implementations, we feel the most promising route might be based on coupled laser arrays of the type shown in [19], possibly combined with the flexibility of spatial light modulators or similar devices for setting the coupling weights, as shown for photonic neural networks in [9]. However, many other platforms are conceivable as well, with the main requirements that it is possible to read out phases, drive individual oscillators externally (for input injection and biases), as well as implement tuneable couplings between arbitrary oscillators.

V. ACKNOWLEDGEMENTS

We acknowledge funding via the German Research Foundation (DFG, Project-ID 429529648–TRR 306 QuCoLiMA, “Quantum Cooperativity of Light and Matter”).

-
- [1] D. Marković, A. Mizrahi, D. Querlioz, and J. Grollier, *Nature Reviews Physics* **2**, 499 (2020).
- [2] D. V. Christensen, R. Dittmann, B. Linares-Barranco, A. Sebastian, M. Le Gallo, A. Redaelli, S. Slesazek, T. Mikolajick, S. Spiga, S. Menzel, *et al.*, *Neuromorphic Computing and Engineering* **2**, 022501 (2022).
- [3] M. Prezioso, F. Merrih-Bayat, B. D. Hoskins, G. C. Adam, K. K. Likharev, and D. B. Strukov, *Nature* **521**, 61 (2015).
- [4] M. L. Schneider, C. A. Donnelly, and S. E. Russek, *Journal of Applied Physics* **124** (2018).
- [5] J. M. Shainline, S. M. Buckley, R. P. Mirin, and S. W. Nam, *Physical Review Applied* **7**, 034013 (2017).
- [6] J. Torrejon, M. Riou, F. A. Araujo, S. Tsunegi, G. Khalsa, D. Querlioz, P. Bortolotti, V. Cros, A. Fukushima, H. Kubota, S. Yuasa, M. D. Stiles, and J. Grollier, *CoRR abs/1701.07715* (2017), [1701.07715](https://arxiv.org/abs/1701.07715).
- [7] K. Wagner and D. Psaltis, *Applied Optics* **26**, 5061 (1987).
- [8] Y. Shen, N. C. Harris, S. Skirlo, M. Prabhu, T. Baehr-Jones, M. Hochberg, X. Sun, S. Zhao, H. Larochelle, D. Englund, *et al.*, *Nature photonics* **11**, 441 (2017).
- [9] J. Bueno, S. Maktoobi, L. Froehly, I. Fischer, M. Jacquot, L. Larger, and D. Brunner, *Optica* **5**, 756 (2018).
- [10] J. Feldmann, N. Youngblood, C. D. Wright, H. Bhaskaran, and W. H. Pernice, *Nature* **569**, 208 (2019).
- [11] J. Feldmann, N. Youngblood, M. Karpov, H. Gehring, X. Li, M. Stappers, M. Le Gallo, X. Fu, A. Lukashchuk, A. S. Raja, *et al.*, *Nature* **589**, 52 (2021).
- [12] S. Pai, Z. Sun, T. W. Hughes, T. Park, B. Bartlett, I. A. Williamson, M. Minkov, M. Milanizadeh, N. Abebe, F. Morichetti, *et al.*, *Science* **380**, 398 (2023).
- [13] N. Pashine, *Physical Review Materials* **5**, 065607 (2021).
- [14] M. J. Falk, J. Wu, A. Matthews, V. Sachdeva, N. Pashine, M. L. Gardel, S. R. Nagel, and A. Murugan, *Proceedings of the National Academy of Sciences* **120**, e2219558120 (2023).
- [15] M. Stern and A. Murugan, *Annual Review of Condensed Matter Physics* **14**, 417 (2023).
- [16] L. E. Altman, M. Stern, A. J. Liu, and D. J. Durian, *arXiv preprint arXiv:2311.00170* (2023).
- [17] J. A. Acebrón, L. L. Bonilla, C. J. Pérez Vicente, F. Ritort, and R. Spigler, *Rev. Mod. Phys.* **77**, 137 (2005).
- [18] J. Kosterlitz, *Journal of Physics C: Solid State Physics* **7**, 1046 (1974).
- [19] M. Nixon, E. Ronen, A. A. Friesem, and N. Davidson, *Physical review letters* **110**, 184102 (2013).
- [20] Y. Takeda, S. Tamate, Y. Yamamoto, H. Takesue, T. Inagaki, and S. Utsunomiya, *Quantum Science and Technology* **3**, 014004 (2017).
- [21] G. Heinrich, M. Ludwig, J. Qian, B. Kubala, and F. Marquardt, *Phys. Rev. Lett.* **107**, 043603 (2011).
- [22] M. Zhang, S. Shah, J. Cardenas, and M. Lipson, *Physical review letters* **115**, 163902 (2015).
- [23] M. H. Matheny, M. Grau, L. G. Villanueva, R. B. Karabalin, M. Cross, and M. L. Roukes, *Physical review letters* **112**, 014101 (2014).
- [24] J. Struck, M. Weinberg, C. Ölschläger, P. Windpassinger, J. Simonet, K. Sengstock, R. Höppner, P. Hauke, A. Eckardt, M. Lewenstein, *et al.*, *Nature Physics* **9**, 738 (2013).
- [25] R. Cosmic, K. Kawabata, Y. Ashida, H. Ikegami, S. Furukawa, P. Patil, J. M. Taylor, and Y. Nakamura, *Phys. Rev. B* **102**, 094509 (2020).
- [26] J. Torrejon, M. Riou, F. A. Araujo, S. Tsunegi, G. Khalsa, D. Querlioz, P. Bortolotti, V. Cros, K. Yakushiji, A. Fukushima, H. Kubota, S. Yuasa, M. D. Stiles, and J. Grollier, *Nature* **547**, 428 (2017).
- [27] M. Romera, P. Talatchian, S. Tsunegi, F. Abreu Araujo, V. Cros, P. Bortolotti, J. Trastoy, K. Yakushiji, A. Fukushima, H. Kubota, *et al.*, *Nature* **563**, 230 (2018).
- [28] A. Baas, K. Lagoudakis, M. Richard, R. André, L. S. Dang, and B. Deveaud-Plédran, *Physical review letters* **100**, 170401 (2008).
- [29] N. G. Berloff, M. Silva, K. Kalinin, A. Askitopoulos, J. D. Töpfer, P. Cilibrizzi, W. Langbein, and P. G. Lagoudakis, *Nature Materials* **16**, 1120 (2017).
- [30] A. Kavokin, T. C. H. Liew, C. Schneider, P. G. Lagoudakis, S. Klembt, and S. Hoefling, *Nature Reviews Physics* **4**, 435 (2022).
- [31] F. C. Hoppensteadt and E. M. Izhikevich, *Physical Review Letters* **82**, 2983 (1999).
- [32] F. C. Hoppensteadt and E. M. Izhikevich, *Physical Review E* **62**, 4010 (2000).
- [33] N. Stroeve and N. G. Berloff, *Physical Review B* **104**, 205435 (2021).
- [34] G. Csaba and W. Porod, *Applied physics reviews* **7** (2020).
- [35] T. Rudner, W. Porod, and G. Csaba, *arXiv preprint arXiv:2309.02532* (2023).
- [36] M. J. Filipovich, Z. Guo, M. Al-Qadasi, B. A. Marquez, H. D. Morison, V. J. Sorger, P. R. Prucnal, S. Shekhar, and B. J. Shastri, *Optica* **9**, 1323 (2022).
- [37] S. Bandyopadhyay, A. Sludds, S. Krastanov, R. Hamerly, N. Harris, D. Bunandar, M. Streshinsky, M. Hochberg, and D. Englund, (2022), [arXiv:2208.01623 \[cs.ET\]](https://arxiv.org/abs/2208.01623).
- [38] F. Duport, B. Schneider, A. Smerieri, M. Haelterman, and S. Massar, *Optics express* **20**, 22783 (2012).
- [39] G. Tanaka, T. Yamane, J. B. Héroux, R. Nakane, N. Kanazawa, S. Takeda, H. Numata, D. Nakano, and A. Hirose, *Neural Networks* **115**, 100 (2019).
- [40] K. Nakajima, *Japanese Journal of Applied Physics* **59**, 060501 (2020).
- [41] G. Van der Sande, D. Brunner, and M. C. Soriano, *Nanophotonics* **6**, 561 (2017).
- [42] L. G. Wright, T. Onodera, M. M. Stein, T. Wang, D. T. Schachter, Z. Hu, and P. L. McMahon, *Nature* **601**, 549 (2022).
- [43] D. Psaltis, D. Brady, X.-G. Gu, and S. Lin, *Nature* **343**, 325 (1990).
- [44] X. Guo, T. D. Barrett, Z. M. Wang, and A. Lvovsky, *Photonics Research* **9**, B71 (2021).
- [45] J. Spall, X. Guo, and A. I. Lvovsky, “Training neural networks with end-to-end optical backpropagation,” (2023),

[arXiv:2308.05226 \[physics.optics\]](#).

- [46] T. W. Hughes, M. Minkov, Y. Shi, and S. Fan, *Optica* **5**, 864 (2018).
- [47] C. C. Wanjura and F. Marquardt, arXiv preprint arXiv:2308.16181 (2023).
- [48] V. Lopez-Pastor and F. Marquardt, arXiv preprint arXiv:2103.04992 (2021).
- [49] B. Scellier and Y. Bengio, *Frontiers in computational neuroscience* **11**, 24 (2017).
- [50] B. Scellier, arXiv preprint arXiv:2103.09985 (2021).
- [51] D. H. Ackley, G. E. Hinton, and T. J. Sejnowski, *Cognitive science* **9**, 147 (1985).
- [52] M. Stern, D. Hexner, J. W. Rocks, and A. J. Liu, *Physical Review X* **11**, 021045 (2021).
- [53] B. Scellier, M. Ernoult, J. Kendall, and S. Kumar, “Energy-based learning algorithms for analog computing: a comparative study,” (2023), [arXiv:2312.15103 \[cs.LG\]](#).
- [54] J. Kendall, R. Pantone, K. Manickavasagam, Y. Bengio, and B. Scellier, arXiv preprint arXiv:2006.01981 (2020).
- [55] E. Martin, M. Ernoult, J. Laydevant, S. Li, D. Querlioz, T. Petrisor, and J. Grollier, *Iscience* **24** (2021).
- [56] M. Stern, A. J. Liu, and V. Balasubramanian, “The physical effects of learning,” (2023), [arXiv:2306.12928 \[cond-mat.dis-nn\]](#).
- [57] M. Ernoult, J. Grollier, D. Querlioz, Y. Bengio, and B. Scellier, arXiv preprint arXiv:2005.04168 (2020).
- [58] B. Scellier, S. Mishra, Y. Bengio, and Y. Ollivier, “Agnostic physics-driven deep learning,” (2022), [arXiv:2205.15021 \[cs.LG\]](#).
- [59] M. Falk, A. Strupp, B. Scellier, and A. Murugan, “Contrastive learning through non-equilibrium memory,” (2023), [arXiv:2312.17723 \[cond-mat.dis-nn\]](#).
- [60] S. Dillavou, M. Stern, A. J. Liu, and D. J. Durian, *Physical Review Applied* **18**, 014040 (2022).
- [61] J. F. Wycoff, S. Dillavou, M. Stern, A. J. Liu, and D. J. Durian, *The Journal of Chemical Physics* **156** (2022).
- [62] M. Stern, S. Dillavou, M. Z. Miskin, D. J. Durian, and A. J. Liu, *Physical Review Research* **4**, L022037 (2022).
- [63] S. Dillavou, B. D. Beyer, M. Stern, M. Z. Miskin, A. J. Liu, and D. J. Durian, arXiv preprint arXiv:2311.00537 (2023).
- [64] J. Laydevant, D. Markovic, and J. Grollier, “Training an ising machine with equilibrium propagation,” (2023), [arXiv:2305.18321 \[cs.NE\]](#).
- [65] S.-i. Yi, J. D. Kendall, R. S. Williams, and S. Kumar, *Nature Electronics* **6**, 45 (2023).
- [66] S. Oh, J. An, S. Cho, R. Yoon, and K.-S. Min, *Micromachines* **14**, 1367 (2023).
- [67] A. Laborieux, M. Ernoult, B. Scellier, Y. Bengio, J. Grollier, and D. Querlioz, *Frontiers in neuroscience* **15**, 633674 (2021).
- [68] S. F. Edwards and P. W. Anderson, *Journal of Physics F: Metal Physics* **6**, 1927 (1976).
- [69] E. Alpaydin and C. Kaynak, “Optical Recognition of Handwritten Digits,” UCI Machine Learning Repository (1998), DOI: <https://doi.org/10.24432/C50P49>.
- [70] X. Glorot and Y. Bengio, in *Proceedings of the thirteenth international conference on artificial intelligence and statistics (JMLR Workshop and Conference Proceedings, 2010)* pp. 249–256.

Appendix A: Interpretation of the Cost Function

In Eq. (4) we introduced a cost function different from the straightforward distance function in Eq. (3). Here we comment further on the problems posed by the distance function. In addition, we point out an interesting connection between the adopted cost function and the overlap of quantum states.

From Eq. (3), we find that the training reaches its fixed points when

$$\sin(\phi_i - \phi_i^T) = 0 \quad (\text{A1})$$

holds for the output units. The distance function D reaches its minimum when $\phi_i = \phi_i^T$ and maximum when $\phi_i = \phi_i^T + \pi$. Theoretically, the maxima are unstable and therefore should be avoided. However, practically they can still cause a long stay on some plateaus during training (with $\phi_i = \phi_i^T + \pi$ for some output unit) and seriously increase the time cost for training. In contrast to this behaviour of D , the only fixed point for the adopted cost function C is the case when $\phi_i = \phi_i^T$ holds for all the output units. We find that this makes the training much more efficient.

Although C was initially chosen simply to provide a strong repulsion from the unstable fixed points (maxima), we can draw an interesting connection to quantum physics. Assume that we have an array of independent two-level quantum systems and prepare the state of the system based on the output units by letting

$$|\Phi_i\rangle = \frac{1}{\sqrt{2}}(|0\rangle + e^{i\phi_i}|1\rangle) \quad (\text{A2})$$

which is exactly an eigenstate of $\cos \phi_i X + \sin \phi_i Y$ of eigenvalue $+1$ (Here $X = |0\rangle\langle 1| + |1\rangle\langle 0|$ and $Y = i|0\rangle\langle 1| - i|1\rangle\langle 0|$).

We prepare another set of two-level systems based on the state of the target in the same way. Then we have two product states, corresponding to the actual output $|S\rangle$ and the target $|T\rangle$

$$\begin{aligned} |S\rangle &= \prod_{i \in S_{\text{out}}} |\Phi_i\rangle = \frac{1}{2^{|S_{\text{out}}|/2}} \prod_{i \in S_{\text{out}}} (|0\rangle + e^{i\phi_i}|1\rangle) \\ |T\rangle &= \prod_{i \in S_{\text{out}}} |\Phi_i^T\rangle = \frac{1}{2^{|S_{\text{out}}|/2}} \prod_{i \in S_{\text{out}}} (|0\rangle + e^{i\phi_i^T}|1\rangle). \end{aligned} \quad (\text{A3})$$

Then the task is to minimize the distance between $|S\rangle$ and $|T\rangle$ by maximizing $|\langle T|S\rangle|^2$. This is equivalent to minimizing

$$\begin{aligned} \mathcal{L} &= -\log |\langle T|S\rangle|^2 \\ &= -\sum_{i \in S_{\text{out}}} \log \left(\frac{1}{4} |1 + e^{i(\phi_i^T - \phi_i)}|^2 \right) \\ &= -\sum_{i \in S_{\text{out}}} \log \left(\frac{1}{2} (1 + \cos(\phi_i - \phi_i^T)) \right) \end{aligned} \quad (\text{A4})$$

which has the form of C expressed in Eq. (4). We note that $|\langle T|S\rangle|^2$ is the probability of measuring the quantum state $|T\rangle$ after preparing the state $|S\rangle$ (if measurements are performed with respect to a basis containing $|T\rangle$ as one of the

states). Thus, \mathcal{L} can also be seen as a categorical cross-entropy $-\sum P_j \ln Q_j$, for the case where $P_1 = 1$ is the desired probability of outcome ‘‘T’’ in such a measurement, and $Q_1 = |\langle T|S \rangle|^2$ is the probability of actually observing this outcome.

Appendix B: Digit recognition performance and benchmarks

We compared the performance of the neuromorphic system to artificial neural networks (ANNs) and linear classifiers which we implemented with tensorflow. The architecture of the models we implemented as well as the number of parameters matched those of the layered structure of our neuromorphic systems. We used a categorical cross-entropy loss function for the ANNs and a mean-squared-error loss function for the linear classifiers and trained both models using Adam optimization for 100 epochs dividing the training data into randomised mini-batches of 10 images. The maximum test accuracy is reported in Tab. II and compared to the performance of the corresponding neuromorphic systems with either layer structure or with all-to-all connectivity.

Layered			All-to-All			Benchmark	
# hidden units	# parameters	test accuracy	# hidden units	# parameters	test accuracy	ANN	linear classifier
20	1540	91.9 %	11	1596	93.3 %	94.3 %	90.4 %
100	7620	92.7 %	75	7725	92.7 %	95.0 %	90.3 %
200	15220	94.1 %	115	15246	92.4 %	95.0 %	90.7 %
300	22820	93.7 %	148	22831	92.9 %	95.1 %	90.7 %

TABLE II. **Digit recognition test accuracy and benchmarks.** Benchmarks were obtained with networks of the same structure as the layered neuromorphic system. For the linear classifier, we used linear activations and a mean-squared-error cost function; for the artificial neural network, we used a sigmoid activation function in the second layer, a softmax in the last layer and a categorical cross-entropy activation function.