Letter

Physical learning beyond the quasistatic limit

Menachem Stern,[1,*] Sam Dillavou,[1] Marc Z. Miskin,[2] Douglas J. Durian,[1] and Andrea J. Liu[1]
[1]*Department of Physics and Astronomy, University of Pennsylvania, Philadelphia, Pennsylvania 19104, USA*
[2]*Department of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, Pennsylvania 19104, USA*

(Received 10 January 2022; accepted 28 April 2022; published 16 May 2022)

Physical networks, such as biological neural networks, can learn desired functions without a central processor, using *local* learning rules in space and time to learn in a fully distributed manner. Learning approaches such as equilibrium propagation, directed aging, and coupled learning similarly exploit local rules to accomplish learning in physical networks such as mechanical, flow, or electrical networks. In contrast to certain natural neural networks, however, such approaches have so far been restricted to the quasistatic limit, where they learn on timescales slow compared to their physical relaxation. This quasistatic constraint slows down learning, limiting the use of these methods as machine learning algorithms, and potentially restricting physical networks that could be used as learning platforms. Here we explore learning in an electrical resistor network that implements coupled learning, both in the laboratory and on the computer, at rates that range from slow to far above the quasistatic limit. We find that up to a critical threshold in the ratio of the learning rate to the physical rate of relaxation, learning speeds up without much change of behavior or error. Beyond the critical threshold, the error exhibits oscillatory dynamics but the networks still learn successfully.

DOI: 10.1103/PhysRevResearch.4.L022037

There is a class of physically realizable learning approaches that do not require a central processor and can in principle be implemented in real networks such as biological neuron, mechanical, flow, or electrical resistor networks. In order for such physical networks to learn on their own, they cannot minimize an arbitrary cost function by gradient descent since that is a global process that requires knowing all the microscopic details at once, carrying out the global computation of gradient descent, and then manipulating networks at the microscopic (node or edge) level. Rather, approaches such as contrastive learning [1–4], equilibrium propagation [5,6], directed aging [7–12], and coupled learning [13–15] use local rules, in which learning degrees of freedom (e.g., the conductances of edges in electrical networks of variable resistors) respond to physical degrees of freedom (e.g., node voltages) which automatically adjust themselves to minimize a physical cost function (e.g., difference in dissipated power), automatically performing an approximation to gradient descent on the cost function. Until now, the use of local rules has been restricted to the limit in which the physical degrees of freedom equilibrate rapidly compared to the learning degrees of freedom. This quasistatic or near-quasistatic condition of slow learning has also limited the utility of local learning approaches such as equilibrium propagation as machine learning algorithms on the computer, by rendering them too slow to compete with global approaches such as backpropagation [16].

It was shown recently that learning can occur beyond the strict quasistatic limit, and that continual updating of the learning degrees of freedom is possible for sufficiently low learning rates [17–19]. However, certain biological networks are capable of learning at rates not just close to the quasistatic limit but far above it. For instance, some natural neural networks may reach a steady spiking state on a similar timescale to synaptic plasticity processes [20–23]. If a system must adapt to a rapidly fluctuating environment supplying the training signals, learning must occur on timescales similar to that of physical adjustment. This may be the case in the humoral immune response, where antibodies are formed rapidly after infection, even before the onset of symptoms [24–26].

In this work, we explore learning as a function of the scaled learning rate $\mathcal{R}$, the ratio of the rates of evolution of the learning and physical degrees of freedom. We specifically consider the task of allosteric response on electrical resistor networks that implement coupled learning, but our approach can be applied to other tasks on any physical learning network. We find that local learning rules can lead to effective learning not only in the quasistatic limit as previously assumed ($\mathcal{R} \ll 1$), but also for values of $\mathcal{R}$ near and even far above unity. We show experimentally and numerically that up to a critical threshold $\mathcal{R}_c \sim 1$, learning speeds up with essentially no change in the error or behavior. Beyond $\mathcal{R}_c$ the system still learns but at the cost of a higher error and oscillatory learning dynamics. At still higher learning rates $\mathcal{R} \gg 1$, the achievable error drops once more, vanishing in the limit of infinitely fast learning $\mathcal{R} \to \infty$. It is notable that even for systems *trained* far from equilibrium, the error in *performing* the allosteric task in equilibrium is minimized as well.

*Corresponding author: nachis@sas.upenn.edu

We analyze the system theoretically to understand these results. Overall, our results suggest that learning in physical systems is achievable far from equilibrium, well above the quasistatic limit. This potentially allows for massive reduction in training times, relaxing limitations on local learning approaches as machine learning algorithms and opening up the use of coupled learning for training physical networks with slow physical dynamics.

Quasistatic and nonequilibrium coupled learning. Recent learning methods in physical systems generally assume that the learning process is decoupled from physical processes. For example, the approach of directed aging for mechanical networks assumes that the physical degrees of freedom (node positions) equilibrate to minimize the physical cost function (elastic energy) on fast timescales compared to the evolution of the learning degrees of freedom (spring constants, equilibrium lengths).

In such frameworks [4–9,13] the physical degrees of freedom v_a on the nodes, indexed by a, first equilibrate to values v_a^*, and the learning degrees of freedom G_j on the edges, indexed by j, are then modified according to v_a^*. The physical degrees of freedom then equilibrate again, spurring further change to the learning ones. In particular, both equilibrium propagation [19] and coupled learning [13] compare two equilibrium states (the free state v_a^{*F} and the clamped state v_a^{*C}) to derive local learning rules that adjust the learning degrees of freedom. Likewise, approaches applied in the laboratory to mechanical networks [4] and electrical resistor networks [14] use learning rules comparing two different equilibrium states. In what follows, our notation is specific for electric resistor networks (with voltage values v_a on nodes indexed by a and conductances G_j on edges indexed by j), but can easily be generalized to other types of physical networks [13].

Throughout this Letter, we will treat the case where the electrical network is to learn a desired "allosteric" task in which a set of inputs V_S are applied to designated source nodes, and the network responds by producing desired outputs V_T at a designated set of target nodes. The free state is characterized by a vector v_a^{*F}, which is the set of all node voltages (physical degrees of freedom) when a set of input voltages V_S are applied to the source nodes, and the network equilibrates to produce the outputs voltages v_T^{*F} at the targets. Here the * notation indicates that the node voltages are equilibrated so that the physical cost function (dissipated power) is minimized; this is the quasistatic condition. The clamped state is defined by applying the inputs V_S to the source nodes *and* a set of voltages v_T^C at the target nodes. The values v_T^C are chosen to be nudged slightly from v_T^{*F} toward the desired output values, $v_T^C = v_T^{*F} + \eta[V_T - v_T^{*F}]$, with $\eta \ll 1$. The coupled learning rule is given by

$$\dot{G}_j = \gamma_\ell \eta^{-1} \frac{\partial}{\partial G_j} \left\{ \mathcal{P}^{*F}\left(V_S, v_T^{*F}\right) - \mathcal{P}^{*C}\left(V_S, v_T^C\right) \right\}, \quad (1)$$

where the learning degrees of freedom are the edge conductances G_j, indexed by j, $\mathcal{P}^{*F}$ and $\mathcal{P}^{*C}$ are the power dissipated by the system in the equilibrium free and clamped states, respectively, and γ_ℓ is a learning rate. Note that since the power $\mathcal{P}$ can be written as a sum over edges, Eq. (1) is spatially local. An approximation to this learning rule has been implemented in the laboratory with an electronic circuit of variable resistors, which can be trained successfully to perform diverse tasks [14]. The function $\mathcal{C}^* \equiv \eta^{-1}[\mathcal{P}^{*C} - \mathcal{P}^{*F}]$ is known as the *contrastive* function, and has been shown to approximate the mean-squared-error cost function $C^* \sim (v_T^{*F} - V_T)^2$ in the limit $\eta \to 0$ [5,13].

Suppose instead that the network learns (updates edge conductances) while the node voltages are still equilibrating. For simplicity we assume that the physical dynamics are overdamped, with decay rate γ_p:

$$\begin{aligned} \dot{v}_a^F &= -\gamma_p \partial_v \mathcal{P}^F, \\ \dot{G}_j &= -\gamma_\ell \partial_G \mathcal{C}. \end{aligned} \quad (2)$$

Here, the nonequilibrium contrastive function is $\mathcal{C} = \eta^{-1}[\mathcal{P}^C - \mathcal{P}^F]$. Note that the power (and contrastive function $\mathcal{C}$) of such physical networks is expressed as a sum over edges. Moreover, integrating Eq. (2) for some interval involves a number of iterations that scales with its length (fixed number of iterations for a forward Euler scheme). Therefore, the computational efficiency of learning using Eq. (2) is linear in network size, comparable to backpropagation in feed-forward neural networks. The scaled learning rate $\mathcal{R} \equiv \gamma_\ell/\gamma_p$ is the key quantity we vary here. The limit $\mathcal{R} \to 0$ is the quasistatic limit of coupled learning, in which the node voltages v_a evolve infinitely rapidly compared to the edge conductances G_j.

The physical cost function [$\mathcal{P}^F$ in Eq. (2)] is the total power dissipated by the network, and can be written as a sum over edges: $\mathcal{P} = \sum_j G_j \Delta v_j^2/2$, where the voltage drop Δv_j is the difference in voltage between the two nodes connected by that edge $\Delta v_j \equiv \Delta_{ja} v_a$ (Δ_{ja} is the incidence matrix). Equation (2) can therefore be simplified to

$$\begin{aligned} \dot{v}_a^{F,C} &= -\gamma_p \sum_j \Delta_{aj}^T G_j \Delta_{ja} v_a^{F,C}/2, \\ \dot{G}_j &= \gamma_\ell \eta^{-1} \left\{ \left[\Delta v_j^F\right]^2 - \left[\Delta v_j^C\right]^2 \right\}/2. \end{aligned} \quad (3)$$

Equations (1) and (3) are nontrivial to implement experimentally because it is impossible to apply simultaneously two different sets of boundary conditions (free and clamped). To circumvent this without storing information in any memory, the laboratory realization [14] of quasistatic coupled learning [Eq. (1)] introduced two identical networks subjected to the different boundary conditions. In the nonequilibrium case, Eq. (3) can likewise be implemented using two identical networks, a free network $v_a^F(t)$ and a clamped network $v_a^C(t)$, both evolving under the same overdamped physical dynamics and having the same edge conductances $G_j(t)$. As in Ref. [14], the difference between the two networks is the boundary conditions (applied external constraints): While both networks are constrained at the source(s) by $v_S^{F,C} = V_S$, the clamped network is further constrained at the target(s) by $v_T^C(t) = v_T^F(t) + \eta[V_T - v_T^F(t)]$.

Experimental results. To study nonequilibrium learning in the laboratory, we adapt the learning network introduced in Ref. [14] that approximates quasistatic coupled learning [13]. Here we lift the quasistatic condition by slowing down relaxation of the node voltages, introducing capacitors in parallel to each variable resistor to realize the overdamped physical dynamics of Eq. (3). Edges of the network with attached capacitors are shown in Fig. 1(a). The charging of

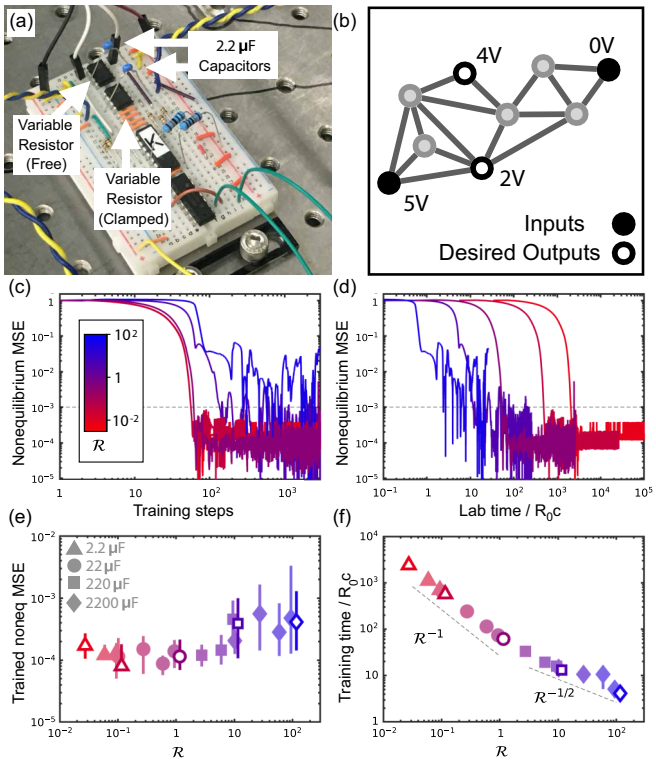

FIG. 1. Nonequilibrium learning in a physical learning network. (a) A single edge (in both free and clamped networks) with parallel 2.2 μF capacitors highlighted with arrows. (b) Network structure with inputs and outputs for the allosteric task used in (c)–(f). (c) Typical nonequilibrium (instantaneous) mean-squared-error (MSE) traces, divided by initial error value, for an allosteric task as a function of training steps and (d) real laboratory time. Colors indicate the scaled learning rate $\mathcal{R}$, and dotted line shows error threshold for (f). (e) Nonequilibrium MSE for networks trained for 3×10^3 steps as a function of $\mathcal{R}$. Error bars indicate first and third quartiles. Shapes indicate capacitor values used, with hollow points corresponding to the traces shown in (c) and (d). (f) Average training times τ when the system achieves a nonequilibrium MSE below a threshold of $\mathrm{MSE}(\tau) = 10^{-3}\mathrm{MSE}(0)$. Dotted lines are power laws of -1 and $-1/2$, respectively.

capacitors increases the time required to reach steady current in the network, but in other respects the experiment is almost [27] as described in Ref. [14]. The network is initialized by imposing the inputs and allowing the free network to reach equilibrium. The clamped boundary condition is applied by imposing both the desired input and output values and the clamped network is allowed to reach equilibrium as well. The training process then commences as in Ref. [14], but with a pause $1/\gamma_\ell$ prior to each update of the clamped voltages as they are adjusted toward the desired target values (see [28] for more information).

The system's physical relaxation time changes during training as the edge resistances evolve. Regardless of these changes, the relaxation time is proportional to the in-line capacitance F added across every resistor. In each experiment c takes one of four values, 2.2, 22, 220, or 2200 μF. We estimate a typical physical relaxation rate using resistance $R_0 \equiv 10$ KΩ (each resistor starts at 50 KΩ and can be varied through the range 781 Ω to 100 KΩ). Thus, the two rates in the system are the learning rate γ_ℓ and the physical relaxation rate $\gamma_p = (R_0 c)^{-1}$ so that the scaled relaxation rate, or ratio of the rate of learning to that of physical relaxation, is $\mathcal{R} = R_0 c \gamma_\ell$.

We train the network to learn a two-target, two-source allosteric task [Fig. 1(b)], varying the capacitance F and learning rate γ_ℓ to adjust $\mathcal{R}$. We define the instantaneous nonequilibrium error as the mean-squared difference between the free-state outputs and the desired outputs $C(t) \equiv \sum_T [V_T - v_T^F(t)]^2$, normalized by its initial value, as measured in real time so that the output voltages reflect the network's current nonequilibrium state. In the quasistatic regime, the scaled learning rate is low, $\mathcal{R} \ll 1$, and the system learns the task, as shown by the typical mean-squared-error (MSE) curves in red in Figs. 1(c) and 1(d), which eventually hover at the noise floor. In this regime, changing the ratio of relaxation rates, $\mathcal{R}$, does not affect the number of training steps required to learn the task τ [see collapse of the two reddest curves in (Fig. 1(c)] but changes the real time required to learn [see the same two curves in Fig. 1(d)]. As $\mathcal{R}$ increases, the system's behavior becomes qualitatively different, at first taking more training steps to reach low error, and then entering into a regime with perpetual oscillations ($\mathcal{R} \gtrsim 1$), as seen best in the blue curve in Fig. 1(d). These oscillations widen the distribution of errors observed, as shown in Fig. 1(e), and increase the number of training steps required for the system to fall below an arbitrary normalized nonequilibrium error threshold of $C(\tau) = 10^{-3}C(0)$ for 10 training steps (this 10-step requirement yields more consistent results by eliminating the effect of "lucky" experimental runs in the fast-learning regime, in which the system flies through a very low-error state but fails to maintain it). In Fig. 1(f) we see two regimes: for $\mathcal{R} \ll 1$, the scaled learning time is well described by $\tau \sim \mathcal{R}^{-1}$, while for $\mathcal{R} \gg 1$, the dimensionless learning time scales approximately as $\tau \sim \mathcal{R}^{-1/2}$.

Simulation results. To test the generality of these results, we turn to numerical simulations on a larger network with $N = 64$ nodes and $N_e = 143$ edges and use nonequilibrium coupled learning [Eq. (3)] to train a different allosteric task, involving $M_S = 5$ source nodes, with applied voltages sampled from a normal distribution $V_S \sim \mathcal{N}(0, 1)$, and $M_T = 3$ target nodes with desired voltages sampled from $V_T \sim \mathcal{N}(0, 0.2^2)$. We assess the success of learning with the equilibrium mean-squared-error (MSE) function $C^* \equiv \sum_T (V_T - v_T^{*F})^2$, normalized by its initial value, and calculated as learning degrees of freedom are evolving. Note that this cost function is in equilibrium (once the physical degrees of freedom reach steady state), compared to the experimentally computed nonequilibrium cost function C (for simulations of the nonequilibrium cost function, see [28]).

Despite important differences to the experiment, the simulations capture all of the major features of the experiment (see [28] for details). In Fig. 2(a), error values C^* during training are shown for different values of the scaled learning rate $\mathcal{R}$. Here, we set the timescale so that a unit time corresponds to one training step. In the quasistatic regime where the physical dynamics are rapid compared to the learning dynamics, $\mathcal{R} \ll 1$, the number of required training steps is independent of $\mathcal{R}$ but the real training time speeds up linearly with $\mathcal{R}$

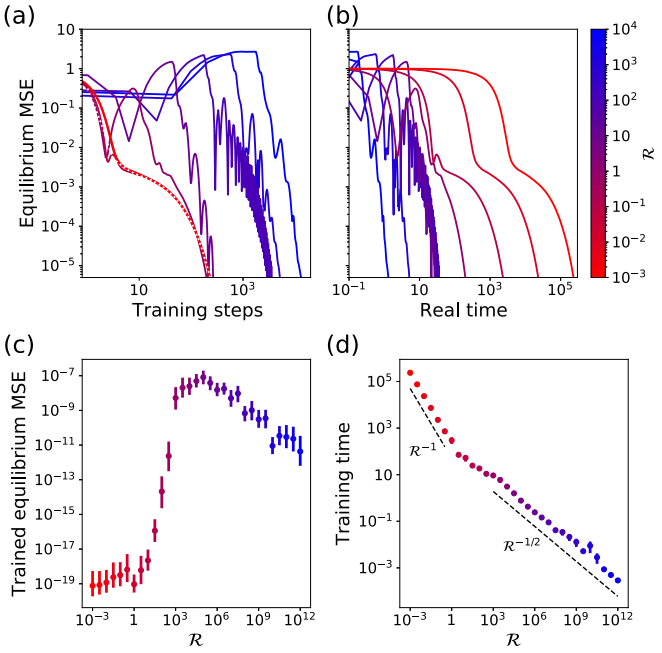

FIG. 2. Nonequilibrium *in silico* training of resistor networks. (a) Network of size $N = 64$ is trained for an allosteric task with $M_S = 5$ sources and $M_T = 3$ targets. We plot the equilibrium mean-squared error (MSE) during training at different scaled learning rates $\mathcal{R}$. When $\mathcal{R} \gtrsim 1$, more training steps are required to reduce the error (dotted line denotes the quasistatic limit $\mathcal{R} \to 0$). (b) However, training with a higher learning rate allows the system to learn more rapidly in real time. (c) The average MSE of trained networks ($N = 64, M_S = 10, M_T = 3$) shows that comparable success is achieved for learning rates up to $\mathcal{R} \approx 1$. While not as accurate, learning still substantially reduces errors even far from equilibrium, where the learning degrees of freedom relax rapidly relative to the physical ones $\mathcal{R} \gg 1$. (d) Training time τ until the equilibrium MSE reaches a threshold $\mathrm{MSE}^*(\tau) = 10^{-6}\mathrm{MSE}^*(0)$. As in the experiment, training time shortens linearly like $\mathcal{R}^{-1}$ for slow learning $\mathcal{R} \ll 1$, while for fast learning $\mathcal{R} \gg 1$ training times shrink as $\mathcal{R}^{-1/2}$.

[reddest curves in Figs. 2(a) and 2(b)]. As in the experiment, $\mathcal{R}$ can be increased to $\mathcal{R} \approx 1$ with little effect on learning. Once $\mathcal{R} > O(1)$, we observe that the error oscillates yet the network still learns successfully at longer times. The error oscillations strengthen and then weaken as the scale learning rate $\mathcal{R}$ increases further.

Figure 2(c) shows similar results for more complex allosteric tasks with 10 sources and 3 targets trained on a network with 64 nodes; results are averaged over 50 different choices of such tasks. As seen experimentally and computationally for the simpler tasks in Figs. 2(a) and 2(b), the system learns with nearly equal low error for all $\mathcal{R} < O(1)$. Moreover, the achievable error can reach low values even when the learning degrees of freedom relax rapidly relative to the physical degrees of freedom ($\mathcal{R} \gg 1$). In Fig. 2(d), we plot the real time τ taken for the network to be trained to a normalized error threshold $C^*(\tau) = 10^{-6}C^*(0)$. As in the experiment, we see a change of scaling in τ at $\mathcal{R} \sim 1$; for $\mathcal{R} \ll 1$ we see the expected linear improvement in the training time $\tau \sim \mathcal{R}^{-1}$. In the far-from-equilibrium regime, $\mathcal{R} \gg 1$, training becomes faster, but with slower scaling $\tau \sim \mathcal{R}^{-1/2}$.

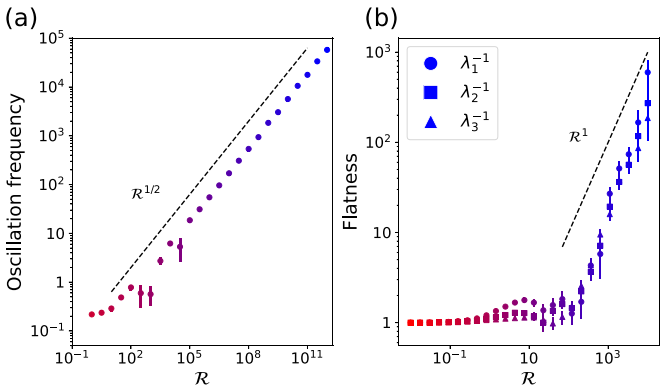

FIG. 3. Nonequilibrium dynamical effects for fast training of allosteric tasks in resistor networks. (a) Starting at a critical learning rate $\mathcal{R}_c \approx 1$, the training error exhibits oscillations whose frequency grows as $\mathcal{R}^{1/2}$. (b) Learning oscillations are suppressed as at higher learning rates, the network finds solutions of flatter cost, as indicated by the diminishing lead eigenvalues of the contrastive Hessian $\partial_G^2 \mathcal{C}$.

Theoretical analysis. To understand the observed behavior we take a second time derivative of the physical degrees of freedom in Eq. (2) to obtain a damped harmonic oscillator equation:

$$\gamma_p^{-2}\delta\ddot{v}_a^F + \gamma_p^{-1}H^F\delta\dot{v}_a^F + \mathcal{R}(D^F)^2\delta v_a^F = 0, \tag{4}$$

where we have defined the difference between the free and desired states $\delta v_T^F \equiv v_T^F - V_T$, and the physical Hessian and cross derivatives $H^F \equiv \partial_v^2 \mathcal{P}^F$, $D^F \equiv \partial_v \partial_G \mathcal{P}^F$. In addition to an inertial term there is an overdamped term for physical relaxation and a restoring one that scales with $\mathcal{R}$. As H^F, D^F do not depend on the rates γ_ℓ and γ_p, the restoring term becomes important only when learning is fast enough ($\mathcal{R} \gtrsim 1$). The solutions to this equation become oscillatory when the discriminant $(H^F)^2 - 4\mathcal{R}(D^F)^2$ becomes negative at a value $\mathcal{R}_c \approx 1$. For slow learning $\mathcal{R} \ll 1$, the decay time of the cost function $C \sim \delta v_a^2$ is dominated by the slowest decay mode of Eq. (4), scaling as $\tau \sim \mathcal{R}^{-1}$. Beyond the critical learning rate $\mathcal{R} \gg 1$, Eq. (4) shows that the oscillation frequency increases, scaling as $\mathcal{R}^{1/2}$. This frequency scaling dictates the training time $\tau \sim \mathcal{R}^{-1/2}$, which can be estimated by the first pass of the oscillatory dynamics through zero (see [28]). These predictions are verified numerically in Fig. 3(a) for a network with $N = 64$ nodes trained for a complex allosteric task with $M_S = 10, M_T = 3$. The oscillatory learning dynamics converge to solutions $\delta v_a \to 0$ as long as the physical derivatives H^F, D^F vary slowly relative to the physical and learning degrees of freedom v_a, G_j, i.e., if the physical cost function (dissipated power) is smooth enough with respect to the learning degrees of freedom. Rugged landscapes with narrow attractors would lead to physical dynamics that jump rapidly between basins, so that convergence cannot be guaranteed. The systems studied here have physical landscapes that are convex with a single basin so that they may be trained extremely fast.

These results also explain the suppression of learning oscillations. Expanding the contrastive function in series, we see that $\partial_G \mathcal{C} \approx -D^F \delta v_a^F$. If the learning dynamics pass

through a flat region, where $\mathcal{C}$ changes slowly compared to the learning parameters G_j, the oscillatory term in Eq. (4) is suppressed, and the dynamics may become overdamped. To test this, we train a $N = 64$ network for an allosteric task with $M_S = 10, M_T = 3$ at different $\mathcal{R}$. After training, we compute the Hessian of the contrastive function $\partial_G^2 \mathcal{C}$ and its top eigenvalues. Smaller eigenvalues indicate flatter contrastive landscapes around the solution. In Fig. 3(b), we show that the landscape becomes flatter linearly with increasing $\mathcal{R}$. In these networks, the learned solutions become flatter faster than the increase in oscillation amplitude $\sim\mathcal{R}^{1/2}$, explaining the suppression of oscillations at high $\mathcal{R}$ in Figs. 2(a) and 2(b).

Discussion. We have demonstrated that physical learning does not need to be restricted to the quasistatic limit in which physical degrees of freedom equilibrate rapidly compared to the learning degrees of freedom. Indeed, the learning process can be sped up by many orders of magnitude by updating learning degrees of freedom at a rate that is comparable to the relaxation rate of the physical degrees of freedom without any qualitative change of behavior or much quantitative change in the achievable error. It is notable that the simulations recover the experimentally observed phenomenology despite the differences between the simulated and experimental networks, suggesting that nonequilibrium coupled learning is robust to small alterations in the equations governing learning, as well as to noise and bias.

Our results may extend beyond contrastive learning in the electrical networks studied here. Similar results may arise in mechanical networks that learn using directed aging [7–9] or contrastive learning [4], studied so far in the quasistatic limit. There, the physical degrees of freedom are the node positions, the physical cost function is the elastic energy, and the learning degrees of freedom are edge stiffnesses, equilibrium lengths [7–9], or the presence or absence of edges [4]. The insight obtained here may shed light on certain forms of biological learning, where plasticity timescales of some neural circuits are similar to the rate of reaching steady state [22,23].

However, our results most directly apply in machine learning in artificial neural networks. There, the quasistatic condition for recurrent networks constitutes a computational bottleneck for equilibrium propagation [16]. Our results show that quasistaticity is not mandatory, reducing the computational effort required for training recurrent nets, making local rules possible competitors with backpropagation algorithms.

Our results also identify limits to how far the learning rate can be increased. We see that the learning rate can be sped up to the physical rate of relaxation, $\mathcal{R} \sim 1$, without penalty, and that the learning dynamics are overdamped up to that point. Beyond a critical rate $\mathcal{R}_c$ the learning dynamics develop underdamped oscillations, as learning degrees of freedom evolve too fast and overshoot desired solutions. These oscillations become faster and stronger at higher learning rates but can be suppressed when the cost function is sufficiently flat with respect to the learning degrees of freedom. This case is generic in the overparametrized regime of deep learning. In complex nonconvex landscapes underlying learning problems in the underparametrized regime, we expect fast learning to fail as the network would rapidly jump between basins in the physical cost function landscape due to the changing learning parameters. In the flow networks we trained for relatively few allosteric tasks, which lie in the convex, overparametrized regime, we observe that training is possible and successful far from equilibrium.

We trained physical networks for allosteric tasks, where a network learns a mapping between a single set of inputs and a single set of outputs. More complex learning problems such as classification, where networks are trained with multiple input and output examples [29–31], require further analysis. Training a network for multiple tasks introduces a new timescale corresponding to the rate at which the system is fed training examples, σ. The insight from our results here suggests that training for multiple tasks may only succeed if the task switching rate is lower than the rate of physical relaxation $\sigma \lesssim \gamma_p$. Otherwise, the system will not effectively evolve under the influence of the training example. We leave a more detailed account of the effect of task switching on far-from-equilibrium learning for future study.

Acknowledgments. We thank Sean Fancher, Eleni Katifori, and Vijay Balasubramanian for insightful discussions. This work was supported by the National Science Foundation via the DMR-2005749 (M.S.), Simons Foundation Investigator Grant No. 327939 to A.J.L (S.D.) and UPenn MRSEC/DMR-1720530 (S.D. and D.J.D.), Army Research Office Award No. W911NF-21-1-0076 and Air Force Office of Scientific Research Award No. FA9550-21-1-0313 (M.Z.M.), and DOE Office of Basic Energy Sciences, DE-SC0020963 (A.J.L.).

[1] J. R. Movellan, Contrastive Hebbian learning in the continuous Hopfield model, in *Connectionist Models* (Elsevier, San Diego, USA, 1991), pp. 10–17.

[2] V. Lopez-Pastor and F. Marquardt, Self-learning machines based on Hamiltonian echo backpropagation, arXiv:2103.04992.

[3] E. Martin, M. Ernoult, J. Laydevant, S. Li, D. Querlioz, T. Petrisor, and J. Grollier, EqSpike: Spike-driven equilibrium propagation for neuromorphic implementations, iScience **24**, 102222 (2021).

[4] N. Pashine, Local rules for fabricating allosteric networks, Phys. Rev. Materials **5**, 065607 (2021).

[5] B. Scellier and Y. Bengio, Equilibrium propagation: Bridging the gap between energy-based models and backpropagation, Front. Comput. Neurosci. **11**, 24 (2017).

[6] J. Kendall, R. Pantone, K. Manickavasagam, Y. Bengio, and B. Scellier, Training end-to-end analog neural networks with equilibrium propagation, arXiv:2006.01981.

[7] N. Pashine, D. Hexner, A. J. Liu, and S. R. Nagel, Directed aging, memory, and natures greed, Sci. Adv. **5**, eaax4215 (2019).

[8] D. Hexner, N. Pashine, A. J. Liu, and S. R. Nagel, Effect of directed aging on nonlinear elasticity and memory formation in a material, Phys. Rev. Research **2**, 043231 (2020).

[9] D. Hexner, A. J. Liu, and S. R. Nagel, Periodic training of creeping solids, Proc. Natl. Acad. Sci. USA **117**, 31690 (2020).

[10] M. Stern, V. Jayaram, and A. Murugan, Shaping the topology of folding pathways in mechanical systems, Nat. Commun. **9**, 4303 (2018).

[11] M. Stern, M. B. Pinson, and A. Murugan, Continual Learning of Multiple Memories in Mechanical Networks, Phys. Rev. X **10**, 031044 (2020).

[12] M. Stern, C. Arinze, L. Perez, S. E. Palmer, and A. Murugan, Supervised learning through physical changes in a mechanical system, Proc. Natl. Acad. Sci. USA **117**, 14843 (2020).

[13] M. Stern, D. Hexner, J. W. Rocks, and A. J. Liu, Supervised Learning in Physical Networks: From Machine Learning to Learning Machines, Phys. Rev. X **11**, 021045 (2021).

[14] S. Dillavou, M. Stern, A. J. Liu, and D. J. Durian, Demonstration of decentralized, physics-driven learning, arXiv:2108.00275.

[15] J. F. Wycoff, S. Dillavou, M. Stern, A. J. Liu, and D. J. Durian, Desynchronous learning in a physics-driven learning network, J. Chem. Phys. **156**, 144903 (2022).

[16] S. Bartunov, A. Santoro, B. Richards, L. Marris, G. E. Hinton, and T. Lillicrap, Assessing the scalability of biologically-motivated deep learning algorithms and architectures, in *Advances in Neural Information Processing Systems* (NeurIPS, Montreal, Canada, 2018), pp. 9368–9378.

[17] M. Ernoult, J. Grollier, D. Querlioz, Y. Bengio, and B. Scellier, Updates of equilibrium prop match gradients of backprop through time in an RNN with static input, *Advances in Neural Information Processing Systems 32* (NeurIPS, Montreal, Canada, 2019).

[18] M. Ernoult, J. Grollier, D. Querlioz, Y. Bengio, and B. Scellier, Equilibrium propagation with continual weight updates, arXiv:2005.04168.

[19] B. Scellier, A deep learning theory for neural networks grounded in physics, arXiv:2103.09985.

[20] R. S. Zucker and W. G. Regehr, Short-term synaptic plasticity, Annu. Rev. Physiol. **64**, 355 (2002).

[21] R. Kanai and F. A. J. Verstraten, Perceptual manifestations of fast neural plasticity: Motion priming, rapid motion aftereffect and perceptual sensitization, Vis. Res. **45**, 3109 (2005).

[22] S. Marom, Neural timescales or lack thereof, Prog. Neurobiol. **90**, 16 (2010).

[23] Y. Sagi, I. Tavor, S. Hofstetter, S. Tzur-Moryosef, T. Blumenfeld-Katzir, and Y. Assaf, Learning in the fast lane: New insights into neuroplasticity, Neuron **73**, 1195 (2012).

[24] S. M. Lemon and L. N. Binn, Serum neutralizing antibody response to hepatitis A virus, J. Infect. Dis. **148**, 1033 (1983).

[25] *WHO Guidelines on Hepatitis B and C Testing* (World Health Organization Press, Geneve, Switzerland, 2017).

[26] J. E. Banatvala and D. W. G. Brown, Rubella, Lancet **363**, 1127 (2004).

[27] The update rule is slightly different, now consisting of the evaluation of XOR($\Delta v_j^F > \Delta v_j^C$, $\Delta v_j^F > -\Delta v_j^C$) whereas before it was XOR($\Delta v_j^F > \Delta v_j^C$, $\Delta v_j^C > 0$). We have observed no effect from this change, but implemented it to adhere more closely to the original coupled learning rule.

[28] See Supplemental Material at http://link.aps.org/supplemental/10.1103/PhysRevResearch.4.L022037 for more details on the experiments, simulations, and accompanying theory.

[29] J. W. Rocks, H. Ronellenfitsch, A. J. Liu, S. R. Nagel, and E. Katifori, Limits of multifunctionality in tunable networks, Proc. Natl. Acad. Sci. USA **116**, 2506 (2019).

[30] M. Ruiz-García, A. J. Liu, and E. Katifori, Tuning and jamming reduced to their minima, Phys. Rev. E **100**, 052608 (2019).

[31] M. Ruiz-Garcia, G. Zhang, S. S. Schoenholz, and A. J. Liu, Tilting the playing field: Dynamical loss functions for machine learning, in *International Conference on Machine Learning* (PMLR, (ICML 2021) virtual conference, 2021), 9157–9167.