



HAL
open science

Task learning through stimulation-induced plasticity in neural networks

Francesco Borra, Simona Cocco, Rémi Monasson

► **To cite this version:**

Francesco Borra, Simona Cocco, Rémi Monasson. Task learning through stimulation-induced plasticity in neural networks. *PRX Life*, 2024, 2 (4), pp.043014. hal-04339067v3

HAL Id: hal-04339067

<https://hal.science/hal-04339067v3>

Submitted on 12 Dec 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Task Learning through Stimulation-Induced Plasticity in Neural Networks

Francesco Borra , Simona Cocco , and Rémi Monasson 

Laboratoire de Physique de l'École Normale Supérieure, PSL Research, CNRS UMR8023, Sorbonne Université, 24 rue Lhomond, 75005 Paris, France



(Received 21 March 2024; accepted 22 October 2024; published 26 November 2024)

Synaptic plasticity dynamically shapes the connectivity of neural systems and is key to learning processes in the brain. To what extent the mechanisms of plasticity can be exploited to drive a neural network and make it perform some kind of computational task remains unclear. This question, relevant in a bioengineering context, can be formulated as a control problem on a high-dimensional system with strongly constrained and nonlinear dynamics. We present a self-contained procedure which, through appropriate spatiotemporal stimulations of the neurons, is able to drive rate-based neural networks with arbitrary initial connectivity towards a desired functional state. We illustrate our approach on two different computational tasks: a nonlinear association between multiple input stimulations and activity patterns (representing digit images), and the construction of a continuous attractor encoding a collective variable in a neural population. Our work thus provides a proof of principle for emerging paradigms of *in vitro* computation based on real neurons.

DOI: [10.1103/PRXLife.2.043014](https://doi.org/10.1103/PRXLife.2.043014)

I. INTRODUCTION

Natural neural networks have long been a key source of inspiration for machine learning and computing, starting from Rosenblatt's perceptron to neuromorphic computing [1] and nowadays to deep learning. Recently, it was suggested that biological neural structures might be directly exploited as a support for *vitro* computation [2,3]. Progress over the past decade has made it possible to grow, preserve, and study brain organoids [4]. The capabilities of stimulating and recording neural populations [5–7] allow for interfacing these systems, potentially turning brain organoids into miniature biological computers. Computing with organized assemblies of neurons could offer considerable advantages with respect to electronics-based devices, in particular in terms of low-energy consumption, massive parallel computation, and continuous learning. However, besides the practical challenges raised by these technologies [8], it remains unclear how they could best be used to achieve high-level computation. The aim of this paper is to address one conceptual question arising in this bioengineering context: how could circuits of biological neurons be trained to carry out desired computational tasks?

Computation with neural cultures has been largely based so far on the framework of reservoir computing [9,10], in which the readout of the high-dimensional neural activity is trained to perform the task of interest. However, reservoir computing falls short from fully realizing the computational potential of neural systems, as it does not exploit their capability for reconfiguration of biological connectivity across time. The

presence of plasticity mechanisms, involving multiple molecular and cellular processes, is crucial to the formation and maintenance of experience-related changes to neural function and circuitry [11]. These mechanisms can potentially be exploited to shape, through appropriate stimulations, the activity and the connectivity of networks [12], as has been shown in specific settings [13]. From a bioengineering point of view, the question is therefore to determine spatiotemporal stimulation patterns capable of remodeling a plastic neural network and make it achieve a desired computational task [Fig. 1(a)]. Establishing guiding principles and practical tools to obtain those stimulation patterns is essential to future progress in biological computation.

From a mathematical point of view, supervised learning of a task is generally cast as an optimization problem in very high dimensions. Given a population of N (artificial) neurons connected through S synaptic connections J_{ij} , one looks for the minimum of the cost (loss) function $U(\{J_{ij}\})$ expressing the mismatch between the target and actual computation carried out by the network. In artificial nets, the cost U can be gradually reduced through gradient descent (or one of its stochastic variants) in the space of connections, until a minimum with good performance is reached. A key point here is that all moves in the N^2 -dimensional interaction space, in particular the one along the gradient of the cost, $\frac{\partial U}{\partial J_{ij}}$, are allowed during learning.

The situation is much more constrained in the case of biological networks, in which the learning dynamics can obviously not be arbitrarily chosen, and plasticity mechanisms set strong limitations about feasible directions. As a concrete example, consider the case of Hebbian-like plasticity rules, in which the changes in the connections ΔJ_{ij} are functions of the firing rates r_i (over some appropriate timescale) of the neurons. As the number N of neurons is generally much smaller than the number N^2 of synaptic interactions, the plastic changes ΔJ_{ij} are highly interdependent, and cannot

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

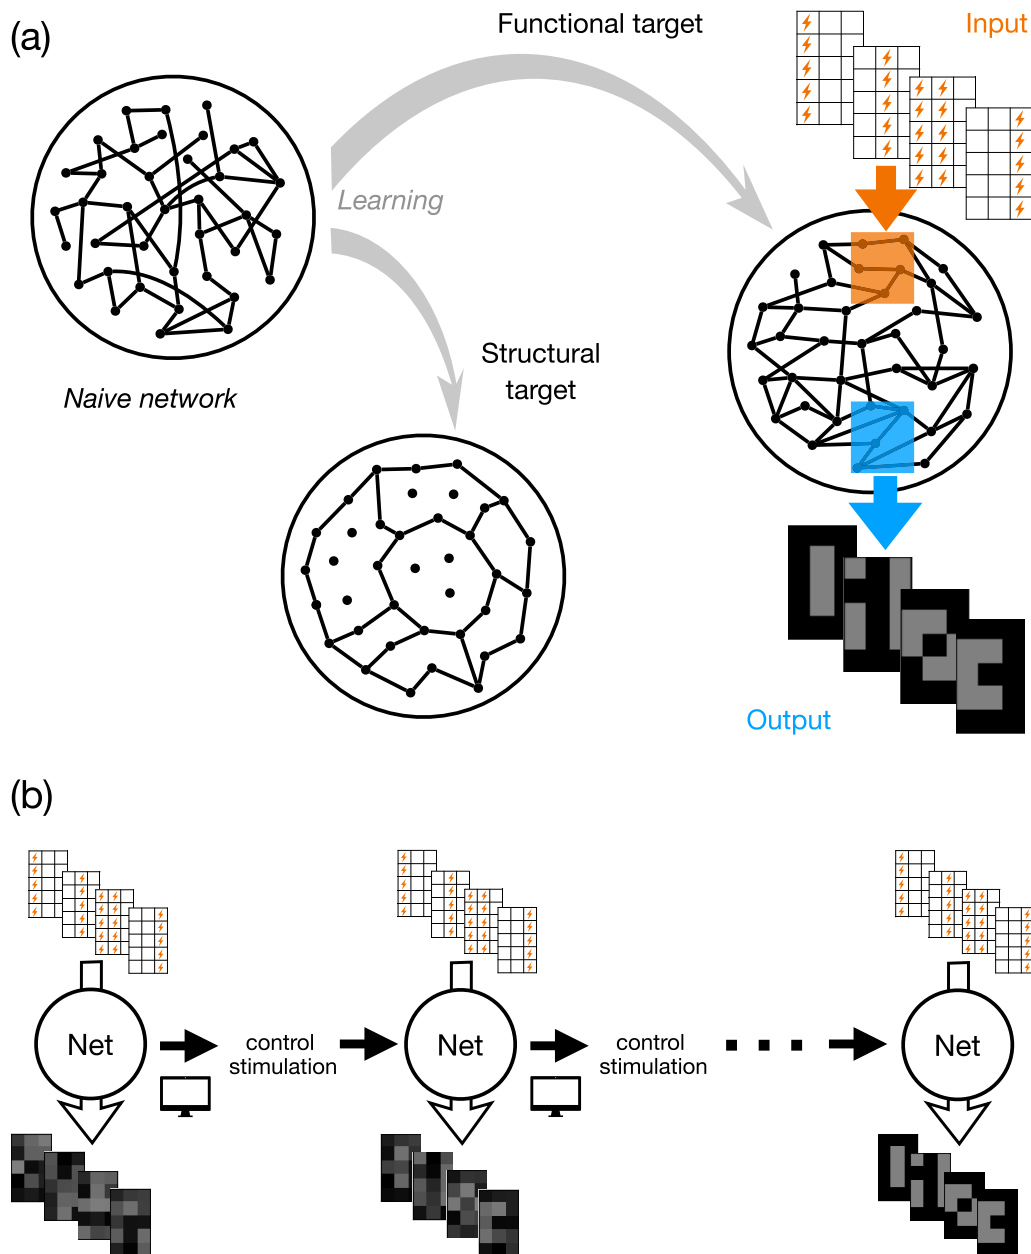


FIG. 1. Computational targets for plastic neural networks. (a) Our goal is to reconfigure a naive neural network and reach some target, either structural (a specific connectivity state) or functional (for instance, the network is required to implement some input-output associations). This reshaping is achieved through a learning process, in which appropriate spatiotemporal stimulations of the neurons exploit intrinsic plasticity mechanisms. (b) The computational target, here the functional task in panel (a), is reached through a learning cycle. Prior to the learning process, the naive network associates the inputs with random outputs (left). Our algorithm computes the best control to be applied to the network to modify its connectivity through plastic changes. Applying this control stimulation to the neural population results in an enhancement of the network performance (middle). The control is then reoptimized and applied during a new stimulation period. After multiple iterations, the correct input-to-output association is achieved (right).

be individually tuned to match the gradient components $\frac{\partial U}{\partial J_{ij}}$. This fundamental limitation makes supervised learning with biologically plausible rules conceptually much more intricate than with unconstrained dynamics. In addition, the magnitude of synaptic modifications is hard to control in biological networks, while the capability to implement adaptive learning rates is generally regarded as a key ingredient in machine learning.

In the present paper, we show that *in silico* neural networks that obey plausible plasticity mechanisms can be effectively trained, through the application of adequate control stimulations, to carry out very diverse computational tasks. These stimulations vary with time and from neuron to neuron, and they are obtained by solving a sequence of optimization problems. Intuitively, the control stimulations are capable of inducing adequate neural activity, progressively driving

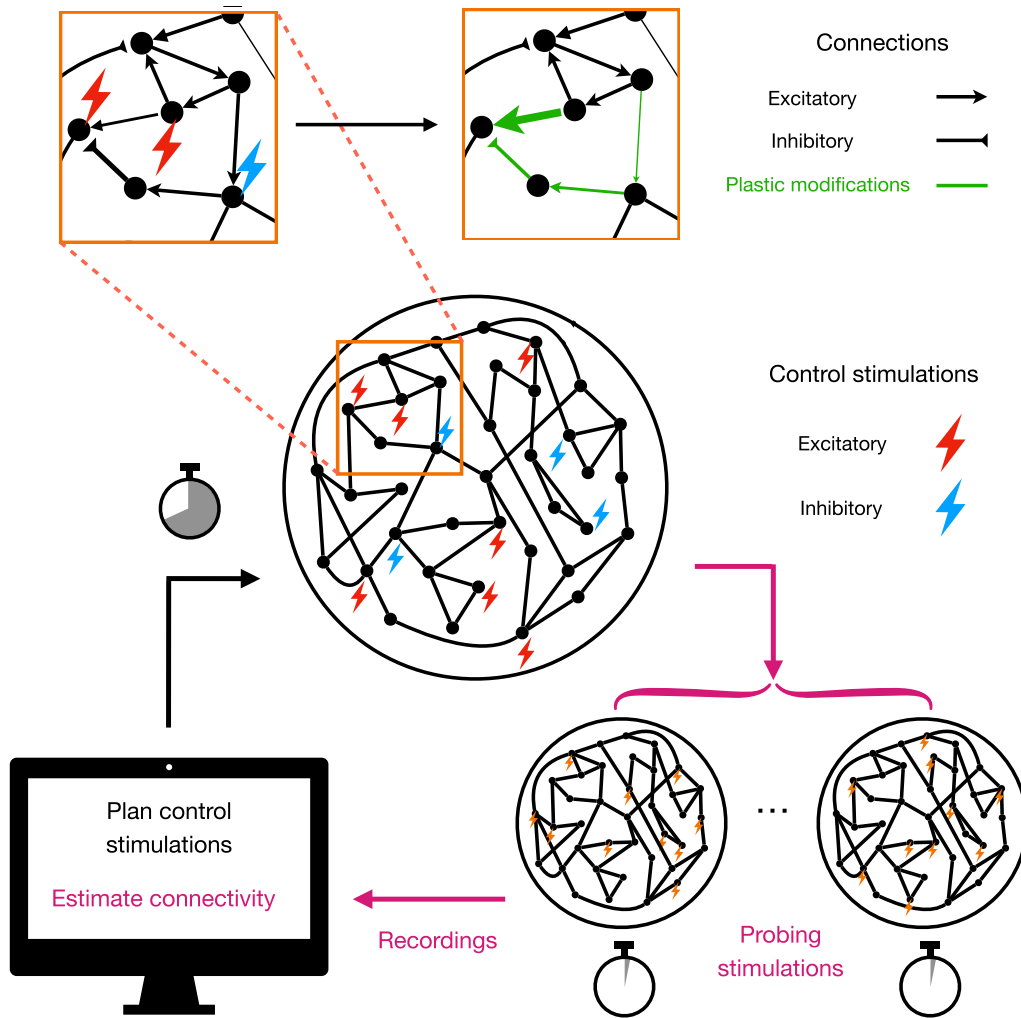


FIG. 2. Stages of the learning cycle: planning of control stimulations, reconfiguration under plasticity effects, and connectivity estimation. In each cycle, a control stimulation f_i^* is applied to the neurons i in the network, generating stationary firing rates r_i (middle). In turn this activity pattern leads, through the plasticity rule, to specific strengthening or weakening of the connections J_{ij} , indicated by the changes in the thicknesses of the connection arrows (top). Once the control period halts, few short and random stimulations are applied to the neurons, and the corresponding activities are recorded (bottom, right). These data are then used to update the estimate of the network connectivity, and to plan the optimal control for the next cycle (bottom, left).

the network connectivity through its intrinsic Hebbian-like plasticity towards a desired state [Fig. 1(a)]. Figure 1(b) summarizes our learning protocol, which relies on multiple cycles of stimulations and recordings of the neural population. At the beginning of each cycle, the responses of the network to a few short probing stimulations are recorded, and used to infer the connectivity of the network. Based on this estimate of the connectivity, we plan a control stimulation pattern, which is subsequently applied to the neurons. Under this control, plastic changes to the connections take place and enhance the network performance in achieving the desired computation. The procedure is iterated until the computational target is reached.

II. MODEL

A. Dynamics of neural activity

We consider a network of N neurons, characterized by their firing rates $\mathbf{r}(t) = \{r_i(t)\}$ at time t and the time-dependent

synaptic connectivity matrix $\mathbf{J}(t) = \{J_{ij}(t)\}$ (Fig. 2). By convention, J_{ij} refers to the coupling from the presynaptic neuron j to the postsynaptic neuron i . The neural population includes N_E excitatory (E) and N_I inhibitory (I) neurons, which constrains the signs of the corresponding synaptic interactions. The activities of the neurons obey the standard dynamical rate equations

$$\tau_n \frac{dr_i}{dt}(t) = -r_i(t) + \Phi \left(\sum_j J_{ij}(t) r_j(t) + f_i(t) \right), \quad (1)$$

where τ_n is the membrane relaxation time, and $f_i(t)$ is a time-dependent control stimulation on neuron i , which can be dynamically controlled at will in the range $[f_{\min}; f_{\max}]$. The input-to-rate transfer function Φ is a sigmoidal function, ranging between zero and maximal frequency r_{\max} [see Methods and the Supplemental Material (SM) [14], Sec. 1]. Control stimulations f_i are expressed in units of r_{\max} , while connections J_{ij} are dimensionless (SM, Sec. 1).

B. Synaptic plasticity

Plasticity induces activity-dependent changes in the interactions. While the precise rules describing plasticity are still debated in neuroscience, we use a simple and general model presenting three essential qualitative features: Hebbian associativity, regression towards a baseline (homeostasis), and bounded synaptic strengths. The equation for the synaptic dynamics is

$$\tau_s \frac{dJ_{ij}}{dt}(t) = \underbrace{\eta(\epsilon_j) [r_i - \theta(\epsilon_j)] r_j}_{\text{Hebbian}} - \underbrace{\beta_1 |J_{ij}| [r_i^2 - \theta_0(\epsilon_j)^2]}_{\text{homeostasis 1}} - \underbrace{\beta_2 \text{sgn}(J_{ij}) h_2(|J_{ij}| - \bar{J})}_{\text{homeostasis 2}}. \quad (2)$$

In the equation above, $\epsilon_j = E$ or I is the presynaptic neuron type, and $h_2(u) = u^2$ if $u \geq 0$, 0 if $u < 0$. The first term assumes that changes in the connections derive from a Hebbian-like covariance rule with a postsynaptic threshold depending on the neuron type [15]. The last two terms account for homeostatic feedback [16–19], biasing the activity of the postsynaptic neuron towards a baseline activity θ_0 , and imposing a soft clipping of synaptic amplitudes outside the range $[-\bar{J}; +\bar{J}]$. Plasticity for excitatory and inhibitory connections is assumed to have different learning rates $\eta(E)$ and $\eta(I)$ and is associated with different thresholds $\theta(E)$ and $\theta(I)$ to enhance the stability of the network activity states. The parameters β_1 and β_2 control the strengths of the homeostatic constraints.

The learning rule in Eq. (2) is flexible, and it can describe Hebbian as well as anti-Hebbian learning. The former is obtained when $\eta(\epsilon_j)$ is positive, independently of the presynaptic neuron type ϵ_j . To accommodate for anti-Hebbian inhibition [20], one chooses $\eta(E) > 0$, $\eta(I) < 0$. We report results with both choices of rules below. We stress that anti-Hebbian here refers to a negative feedback between correlations and synaptic strength change.

While the associative learning rule above can be easily modified, a key assumption we rely on is that synaptic changes are slow compared to the fast variations in the activity, i.e., $\tau_n \ll \tau_s$ (see Methods and the SM, Sec. 2). Due to this assumption, only the slow variations (on the τ_s timescale) of the stimulations $\mathbf{f}(t) = \{f_i(t)\}$ matter. Furthermore, the neural activities r_i are quasistationary and locked to the slowly varying inputs. In mathematical terms, the right-hand side of Eq. (1) vanishes.

C. Target state of the network

Our goal is to train the neural network connectivity \mathbf{J} to meet some target, either structural or functional [Fig. 1(a)]. In the former case, the network connectivity is asked to reach some value, say, $\mathbf{J}^{\text{target}}$. The training procedure should ensure that the loss

$$U_{\text{task}}(\mathbf{J}) = \sum_{i,j} w_{\epsilon_i, \epsilon_j} [J_{ij} - J_{ij}^{\text{target}}]^2 \quad (3)$$

decays towards zero over the training time. We introduce weights $w_{\epsilon_i, \epsilon_j}$ depending on the neuron types to ensure that

the four classes of connections $E, I \rightarrow E, I$ are equally contributing to U_{task} (Methods).

In the functional case, the target is not the connectivity itself as above, but the computation carried out by the network. The neurons i in the network are partitioned into three subpopulations, referred to as input (“in”), processing (“proc”), and output (“out”). The network is required to implement a set of n_{pairs} input/output mappings, that is, produce desired activities r_i^μ of the neurons i in the “out” subpopulation in response to specific input stimulations \mathbf{f}^μ applied to the neurons j in the “in” subpopulation, with $\mu = 1, \dots, n_{\text{pairs}}$. A possible loss associated with this association task reads

$$U_{\text{task}}(\mathbf{J}) = \sum_{\mu=1}^{n_{\text{pairs}}} \sum_{i \in \text{out}} [r_i(\mathbf{J}, \mathbf{f}^\mu) - r_i^\mu]^2, \quad (4)$$

where $r_i(\mathbf{J}, \mathbf{f})$ refers to the stationary solution of Eq. (1); see Methods for alternative choices of the loss. The task may be made harder by requiring that the “in” and “out” neurons share no direct connections, and communicate through the neurons in the “proc” subnetwork.

III. TRAINING LOOP: SUPERVISED CONTROL OF HEBBIAN PLASTICITY

Our training procedure is based on a sequence of learning cycles, labeled by $k = 0, 1, 2, \dots$ and sketched in Fig. 2. Each cycle k can be decomposed in three steps:

- (i) Estimation of the current connectivity, \mathbf{J}_k , through fast probing of the responses of the network to random stimuli.
- (ii) Calculation of the optimal control to be applied, \mathbf{f}_k^* , to shift the network connectivity state towards the desired target.
- (iii) Application of this control during the period Δt , leading to a reorganization of the network through its intrinsic plasticity mechanisms.

The learning process stops when the value of the loss U_k is considered small enough that the target has been reached. We now give more details about the steps above; technicalities can be found in Methods.

A. Estimation of the network connectivity

Determining and updating the optimal control stimulation requires knowing, to some degree of accuracy, the state of connectivity \mathbf{J}_k [21]. To mimic realistic conditions, \mathbf{J}_k cannot be accessed through direct measurements. We therefore infer the connectivity based on multiple recordings of the activity (Fig. 2). We probe the network stationary activity states $\mathbf{r}^\nu = \{r_i^\nu\}$ corresponding to various random stimulation patterns $\mathbf{f}^\nu = \{f_i^\nu\}$, where $\nu = 1, \dots, n_{\text{probes}}$. The connectivity matrix is then estimated so that the stationary equations

$$r_i^\nu = \Phi \left(\sum_j (J_k)_{ij} r_j^\nu + f_i^\nu \right) \quad (5)$$

are fulfilled for all $i = 1, \dots, N$ and $\nu = 1, \dots, n_{\text{probes}}$, see Eq. (1) and Methods. Increasing the number n_{probes} of recordings reduces the error over the estimated connectivity matrix, but the duration of the whole process should be sufficiently short not to induce any plastic modification to \mathbf{J}_k : hence n_{probes} should not exceed τ_s/τ_r , where τ_r is the relaxation time of

the network activity, which can be estimated from τ_n and \mathbf{J}_k (Methods). In practice, the connectivity varies little after each control stimulation period, and few recordings are needed to update our estimate.

B. Calculation of optimal control

Ideally, one would like the change of connectivity during one cycle, $\Delta\mathbf{J} = \mathbf{J}_{k+1} - \mathbf{J}_k$, to align along the direction of steepest descent of the loss U , i.e., $\Delta\mathbf{J} \propto -\partial U / \partial \mathbf{J}|_{\mathbf{J}_k}$. However, the connections obey the plasticity rule in Eq. (2), and their dynamics is not directly and arbitrarily controllable. From an informal point of view, the best we can hope for is to find a stimulation that will enhance neuronal activities in such a way as to move connections in a direction as aligned as possible with the gradient of U ; see Fig. 3(a).

Let us call $\Delta\mathbf{J}(\mathbf{f}, \Delta t, \mathbf{J}_k)$ the change in the connections produced by applying the stimulation \mathbf{f} for a period of time Δt to the network with initial connectivity state \mathbf{J}_k . In principle, $\Delta\mathbf{J}$ is obtained by integrating the Hebbian-like plasticity rule in Eq. (2) during the time interval (under fixed stimulation). Then, the best control is formally given by

$$\mathbf{f}_k^* = \underset{\mathbf{f}}{\operatorname{argmin}} U(\mathbf{J}_k + \Delta\mathbf{J}(\mathbf{f}, \Delta t, \mathbf{J}_k)). \quad (6)$$

Solving this optimization problem and finding the absolute minimum can be quite hard. We use a gradient descent procedure in the space of the controls \mathbf{f} , not to be confused with the training dynamics over \mathbf{J} , which is guaranteed to return a local minimum at least. We stress that this optimization step is an abstract computation, unrelated to any physical or biological process, and is done on a computer separate from the neural model; see Fig. 2. Informally speaking, each cycle can be thought of as an approximate gradient descent of U in the \mathbf{J} -space, whose precise direction is determined by another gradient descent in the \mathbf{f} -space. This double process is shown in Fig. 3(b).

The determination of \mathbf{f}_k^* in Eq. (6) via gradient descent is a nontrivial computational problem due to the nonlinearities in the neural and synaptic dynamics; see Methods. Choosing $\Delta t \ll \tau_s$ (SM, Sec. 1) and imposing additional regularization terms (Methods) ensure that the changes $\Delta\mathbf{J}$ remain small at the end of the learning cycle.

C. Plastic reorganization of the network under control stimulations

Once the optimal control \mathbf{f}_k^* has been computed (in a time we neglect compared to the timescales at play in the neural network), it is applied to the neurons for a period Δt . The activities of the neurons rapidly adapt to these external stimulations, and the connections start evolving under the plasticity dynamical rules.

At the end of this training cycle, the network connectivity may differ from the expected value $\mathbf{J}_k + \Delta\mathbf{J}(\mathbf{f}_k^*, \Delta t, \mathbf{J}_k)$ due to two main factors. First, the estimate \mathbf{J}_k , obtained prior to the stimulation, is not equal to the ground truth connectivity. The accuracy depends on the number of measurements, and is limited by the constraint of fast probing as explained above. Second, the true underlying plastic mechanisms controlling the network evolution can never be exactly modeled. The

unavoidable mismatch between the network dynamics and Eq. (2) induces, at the end of the stimulation period, a systematic bias in the value of the connectivity. To account for this bias and quantify its consequences, we consider below two variants of Eq. (2), with different sets of parameters, for the “true” plastic evolution of the network and for its model counterpart used to compute $\Delta\mathbf{J}(\mathbf{f}, \Delta t, \mathbf{J})$ and the optimal control \mathbf{f}^* .

IV. APPLICATIONS

We apply below our learning procedure to two computational tasks; a third task defined on a small network is detailed in the SM, Sec. 6, and in the Discussion section.

Task 1: Nonlinear association between all-or-nothing inputs and digitlike outputs

Definition of the task. A fundamental computation carried out by neural circuits in organisms is the association of diverse behavioral (or motor) responses to multiple sensory inputs. This task is particularly nontrivial to learn when it is non-additive, as it cannot be realized by linear networks [22], and it may require recruiting dedicated brain areas [23].

We consider here a toy version of a nonlinear associative task, in which a subset of N_{in} input neurons can be stimulated in n_{pairs} distinct ways. Each input stimulation is expected to elicit an associated activity pattern over N_{out} output neurons; see Fig. 4(a). For the sake of visual clarity, the output patterns are represented as digits, with gray/black pixels corresponding to neurons with low/high activities arranged on a rectangular grid. We make sure of the following:

(i) The association task is nonlinear by carefully choosing the input/output pairs. For instance, the sums of the input stimulations associated with the output digits 0 and 1 is the input required to elicit digit 2; however, 2 is quite distinct from the superimposition of 0 and 1; see Fig. 4(a). As a result of nonlinearity, the association task cannot be simply implemented through separate neural pathways, activated by one input and silent for the other ones.

(ii) There is no direct connection from the input to the output neurons; see Fig. 4(b). This arrangement mimics the idealized structures of a modular chip. Hence, the information about the input must be processed and conveyed by an intermediate processing (“proc”) region of the network, which contains most of the N neurons. This requirement also ensures that the decrease of the cost function in the initial part of the training is not due to a simple creation of a linear direct input-to-output connection, but rather to the creation of a complex structure.

Training: costs and dynamics. To learn the task, we introduce a cost over the connectivity matrices favoring the desired input/output associations; see Methods, Eq. (20). An additional cost aims at regularizing the connectivity and speeding up convergence towards stationary states of activity; see Methods.

We start from a network with weak, random connections. As our spatiotemporal control stimulation is applied to the neurons [Fig. 4(c)], the connectivity gets modified, and the costs associated with the task and the regularization decrease,

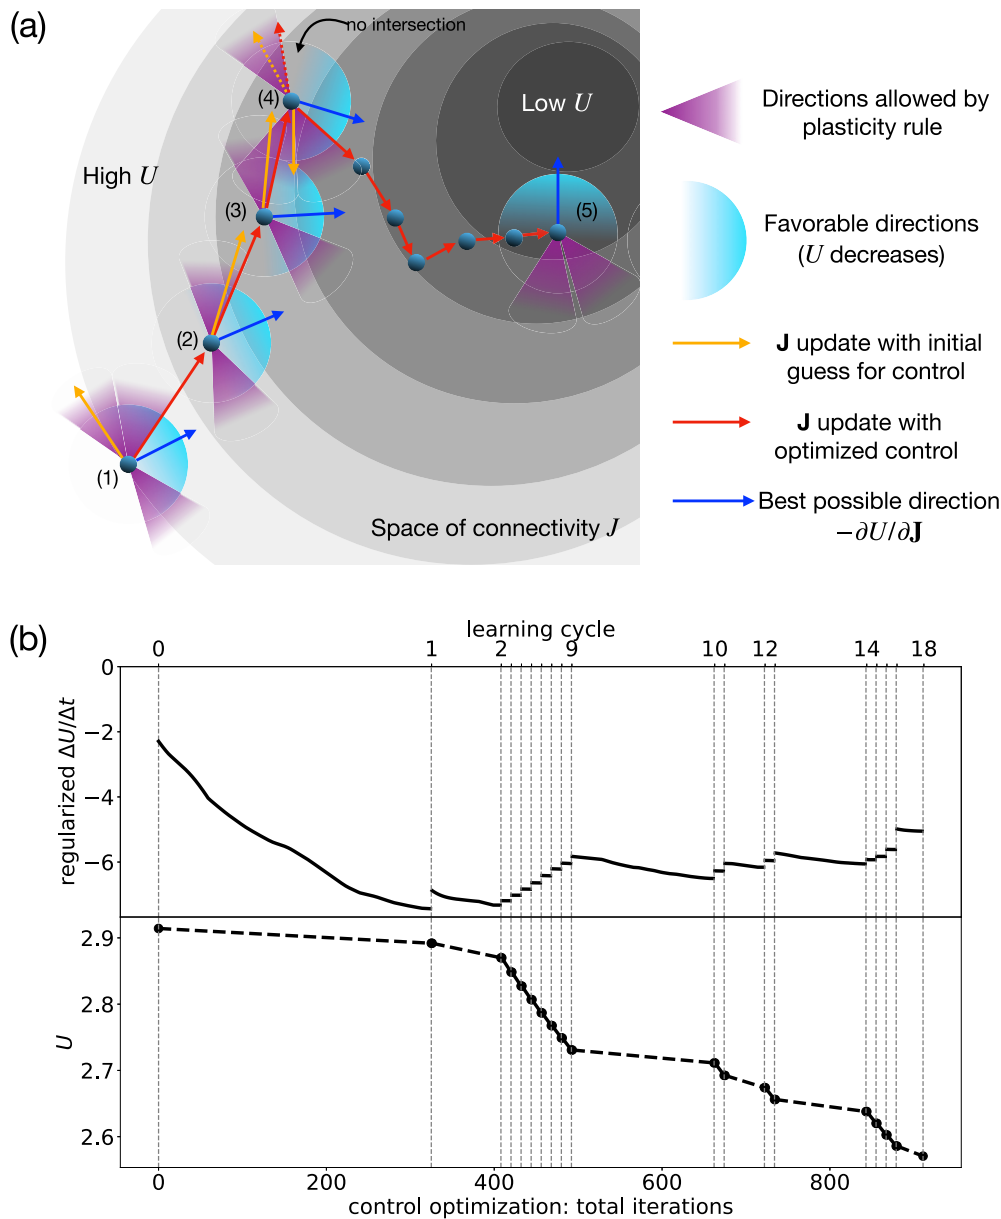


FIG. 3. Schematic dynamics of the network and feasible directions in the high-dimensional connectivity space during the training process. (a) Gray levels show values of the cost U , which quantifies the mismatch between the current connectivity \mathbf{J} of the network and the target. Dark blue balls locate the connectivity \mathbf{J} at the beginning of each control stimulation period. The set of all directions in the N^2 -dimensional connection space along which U decreases is shown by the cyan area, centered around the direction of steepest descent of U indicated by the blue arrows. The purple cone symbolically represents the set of all feasible directions for \mathbf{J} under the plasticity dynamics, i.e., which can be reached under any N -dimensional control stimulation \mathbf{f} . (1) Initial situation, prior to any control stimulation. The synaptic change corresponding to a random \mathbf{f} may point to a “bad” direction outside the cyan region (orange arrow), while the change associated with the optimal control stimulation \mathbf{f}^* lies on the edge of the purple cone (red arrow), as close as possible to the best direction (blue arrow). (2) After one control stimulation period. Repeating the same control stimulation as before would lead to a suboptimal change of \mathbf{J} (orange arrow), while the best stimulation (red arrow) yields a larger decrease in U . (3) Control stimulations are updated to ensure optimal synaptic changes until (4) all local updates of the previous \mathbf{f}^* lie outside the “good” (cyan) region. We then choose a new initial \mathbf{f} , and resume the search for the optimal \mathbf{f}^* as in cycle (1). This process is iterated, until the changes in U are very small and the optimization is completed, or the feasible and “good” regions do not overlap any longer [point (5)]. The performance of the network (final value of U) is then assessed. (b) Decay rate of the loss during the optimization over the control (top) and resulting loss U_k after each learning cycle k (bottom), for the setting corresponding to Task 1. Abscissa indicates the number n of optimization steps. Vertical dashed lines locate the beginnings of the training cycles k . Top: decay rate $[U(n) - U_k]/\Delta t$ at step n of the optimization algorithm, calculated with the regularized loss; see Eq. (17). For each cycle, the rate is decreased through gradient descent over \mathbf{f} until a plateau is reached. The optimization then halts, and the control defines \mathbf{f}_k^* . At the beginning of the new cycle, the network connectivity \mathbf{J} has changed, the previous control is not optimal any longer, and the optimization process resumes. The plateau values of the rate increase over cycles, making it progressively harder to find a good control to train the network.

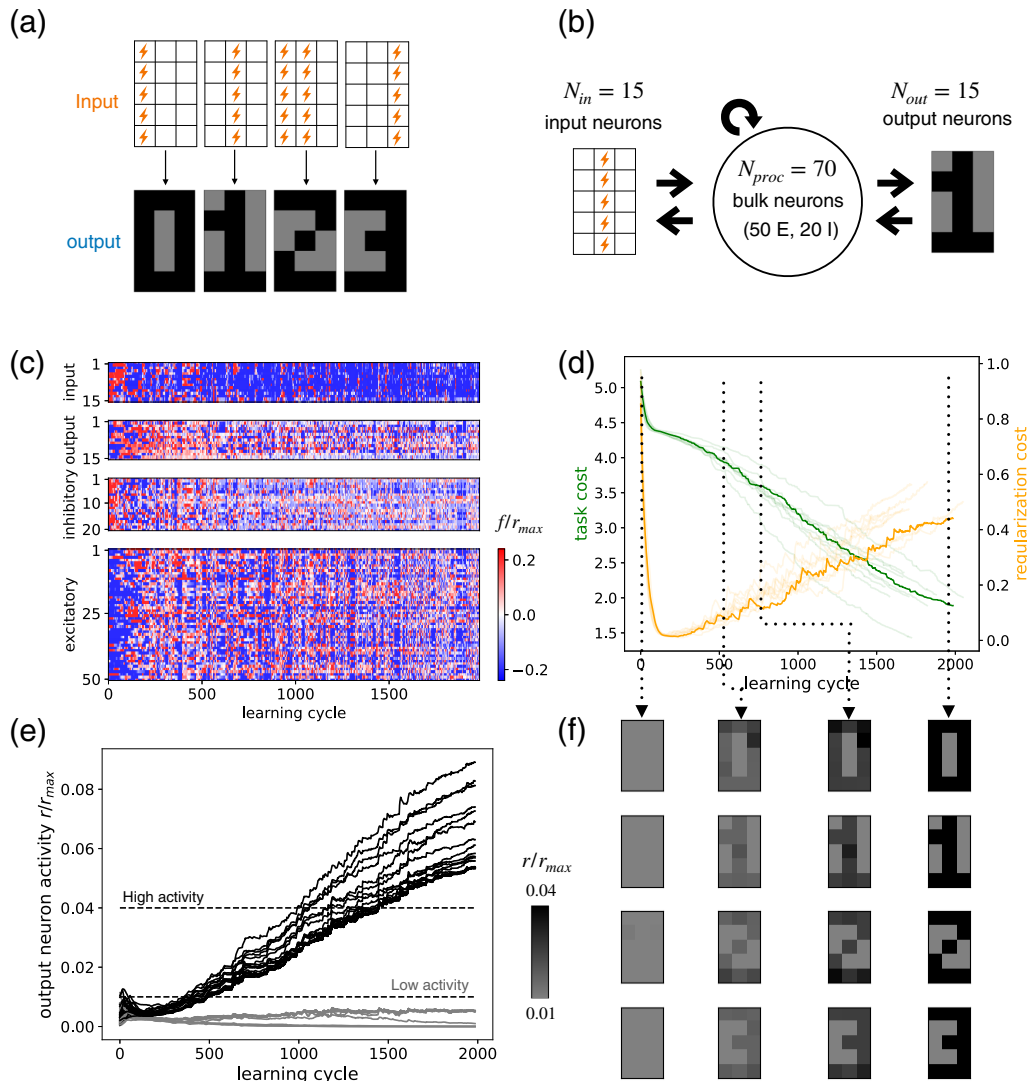


FIG. 4. Task 1: nonadditive input-output association. (a) The $n_{pairs} = 4$ pairs of input-output associations to be learned by the network. Input patterns: flashes locate the input neurons (out of $N_{in} = 15$) subject to a strong stimulus ($f_i = f_{max}$), while the other neurons (empty squares) receive no input ($f_i = 0$). Output patterns: target activities of the $N_{out} = 15$ output neurons (black: $r_i > 0.04r_{max}$, light gray: $r_i < 0.01r_{max}$). (b) Sketch of the network, with the bulk processing area including $N_{proc} = 70$ neurons, connecting the input and output neurons. The overall fractions of excitatory and inhibitory neurons in the network are equal to, respectively, 80% and 20%. (c) Control stimulations $f_i(t)$ (see color bar for values) applied to the neurons as a function of the learning step. From top to bottom: input, output, bulk inhibitory, and bulk excitatory neurons. (d) Costs associated with the task (green, left scale) and with regularization (orange, right scale) as functions of the learning cycle. (e) Activities of the output neurons in response to the input stimulation patterns during learning. The color (black or light gray) of each one of the $n_{pairs} \times N_{out} = 60$ curves indicates the level of activity requested for the corresponding output neuron [high or low, see panel (a)]. Training is successful when the black and light gray curves become well separated. (f) Average activities of the output neurons for the four time steps identified with the dashed line in panel (d), showing how the figures of digit emerge from the initial random image. Black and light gray activities are consistent with panel (e), with intermediate gray levels defined in the bar.

as shown in Fig. 4(d). The activities of neurons progressively cluster into two categories, depending on their expected values in the output patterns [Fig. 4(e)]. The learning process is characterized by three stages. First, the activities of all output neurons approximately reach the same low activity level independently of the input pattern, as regularization tends to decrease the amplitudes of the connections and of the activities. Then activities start to separate according to the output patterns. At this stage, one can already guess the output digits; see Fig. 4(f). We let the protocol proceed until the two groups (active, inactive) are clearly separated, and the output neuron

activities match the target patterns; compare Figs. 4(a) and 4(f).

Transfer of information and neural representations. How the information is transferred from the input to the output neurons and represented in the “proc” area is investigated in Fig. 5. While the input and output subpopulations in our network both include $N_{in} = N_{out} = 15$ neurons, the associative task effectively takes place in a lower-dimensional space with $n_{pairs} = 4$ dimensions. We show in Fig. 5(a) that learning progressively concentrates most of the $N_{out} = 15$ -dimensional activity in the 4-dimensional subspace spanned by the output

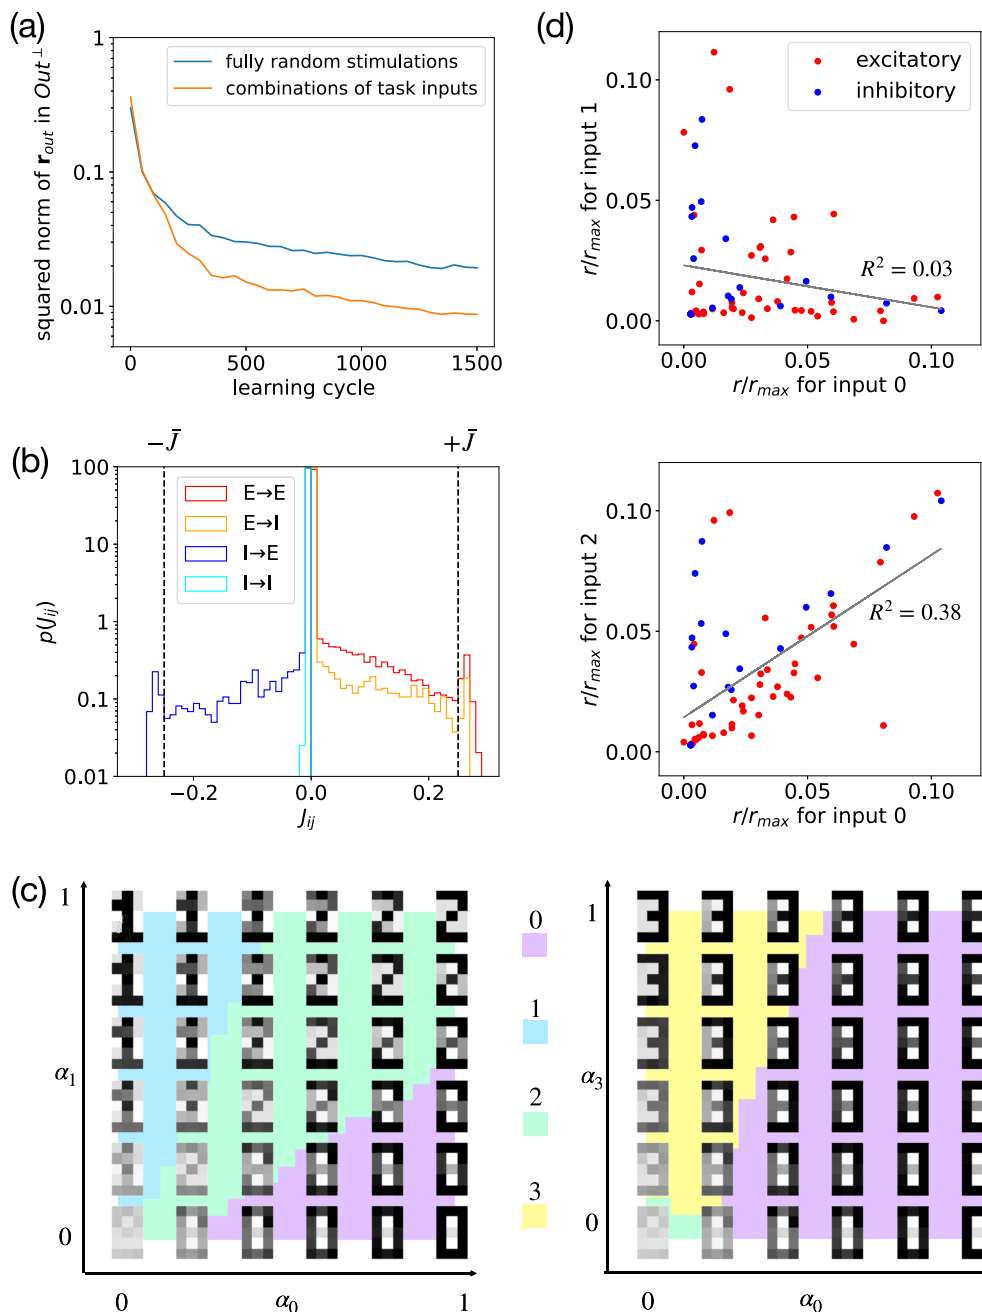


FIG. 5. Learning dynamics and internal representations for Task 1. (a) Squared norm of the projection of the output activity (over $N_{out} = 15$ neurons) orthogonal to the space spanned by the $n_{pairs} = 4$ digit patterns as a function of the number of training cycles. The total activity is normalized to unity. Blue curve: case of random 0-1 stimulations over the $N_{in} = 15$ input neurons; orange: stimulations are random convex combinations of the n_{pairs} inputs. Results were averaged over 10 training trials, and each point is averaged over 50 random input stimulations. (b) Histograms of interactions J_{ij} between the neurons in the E and I populations. Dashed lines show the soft bounds in $\pm \bar{J}$. (c) Output neuron activities (gray levels) in response to combinations of inputs $\alpha_\mu \mathbf{f}_\mu + \alpha_\nu \mathbf{f}_\nu$, with $0 \leq \alpha_\mu, \alpha_\nu \leq 1$. The bottom left corner corresponds to vanishing input and the bottom, right and top, left corners to pure inputs associated with, respectively, digits μ and ν . The background colors (middle squares) indicate the most resembling digits. (Left) case $\mu = 0, \nu = 1$: the output of the linear combination of inputs is not the linear combination of the outputs, as required by the association task. (Right) Case $\mu = 0, \nu = 3$: the response to linear combinations of \mathbf{f}_0 and \mathbf{f}_3 is approximately additive. (d) Scatter plot of the activities of the 70 bulk neurons for inputs 0,1 (top) and 0,2 (bottom) after training. The representations associated with patterns 0 and 1 are not correlated, as expected from the orthogonality of \mathbf{f}_0 and \mathbf{f}_1 [Fig. 4(a)]. Conversely, the representations of 0 and 2 are strongly correlated, reflecting the relation $\mathbf{f}_2 = \mathbf{f}_0 + \mathbf{f}_1$ [Fig. 4(a)].

digits. This dimensionality reduction phenomenon is even stronger when the N_{in} input neurons are stimulated by random combinations of the training inputs associated with the digits, implying the effective creation of a 4-to-4 dimensional channel upon learning. From a circuit point of view, this functional channel is supported by a network of connections, whose histogram is reported in Fig. 5(b). We observe that, during learning, the E and I subpopulations strengthen their connections, with the exception of recurrent interactions within the I neurons that remain weak. The creation of input-specific sub-networks of activity during the learning process can be seen in the animation provided in the SM, Sec. 7.1.

We stress that, within this low-dimensional space, the network computation is not linear as imposed by the nonadditivity of the task. We show in Fig. 5(c) how the network response to linear combinations of the inputs may largely differ from the superimposition of the responses to the single inputs; see case 0–1 (left). This result confirms that the network is capable of learning complex behavior beyond the superimposition of patterns in the low-dimensional subspace spans by the patterns. On the contrary, in the absence of any specific constraint imposed by the task, computation appears to be approximately linear; see the combinations of inputs 0 and 3 producing the 8 digit (right).

Furthermore, we may ask how the network, after training, is able to process and differentiate the patterns outside the input and output subpopulations. We see in Fig. 5(d) that inputs produce the activation of a significant fraction of the neurons in the “proc” area and that many of those neurons respond specifically to distinct inputs, with stronger correlations in the evoked activities for more correlated inputs. In particular, the role of inhibitory neurons is key to the differentiation of input 2 from both 0 and 1, and to depart from additivity as imposed by the task. Further information about how neurons in the “proc” region participates in the association task can be found in the SM, Fig. 6.

Task 2: Ringlike connectivity supporting continuous attractor dynamics

Definition of the task. We now aim at reshaping network connectivity into a target matrix $\mathbf{J}^{\text{target}}$ through learning. The target connectivity defines a ring attractor, capable of supporting a bump of activity coding for a continuously varying angle. Ring attractors were first theoretically hypothesized [24,25], and they have recently been observed in the ellipsoid body of fly [26]. Neuron connections are organized along two rings; see Fig. 6(a). Excitatory neurons form the outer ring, with connections that decay with their distance, while inhibitory neurons compose the smaller inner ring. The connections between the two rings are such that a bump of activity on one side of the outer ring induces inhibition on the diametrically opposite side [see Fig. 6(a)]. The position of the bump is arbitrary in the absence of external input, and is otherwise attracted by a weak and localized input to the outer neurons.

Control stimulations are computed with the cost function given by Eq. (3). The time behavior of the cost is shown in Fig. 6(b) for 10 different random initial networks [see Fig. 6(c), left, for a realization of the naive connectivity].

After a fast initial decay, the cost relaxes to very low values, signaling the success of the learning procedure.

Training: network connectivity and stimulations. Snapshots of the connectivity at different steps of the learning process are displayed in Fig. 6(c). The strong similarity between the network connectivity at the end of learning [compare Figs. 6(a) and 6(c), step (4)] confers the desired functional properties of continuous attractors to the network. We report in Fig. 6(d) the average firing activities of neurons when a weak input stimulation (intended to pin the bump of activity at a specific angle) is applied. At the initial state of learning (1), the network responds by a weak excess activity simply reflecting the localized pinning input stimulation. Through training, a subpopulation of neurons emerges, whose activity is supported by the recurrent connections. This subpopulation varies with the angle associated with the input stimulation; see (2) and (3). The emergence of receptive fields is improved with further training cycles (4). Notice that, even when the target structure is very well reconstructed, the direction of the polarization is not perfect, due to tiny correlated defects in the connectivity structure. The presence of these defects does not preclude the accuracy of the coding of angular information in the network [27].

The time behavior of the control stimulations applied to neurons during learning is shown in Fig. 6(e). We observe the presence of long-distance (large-angle) correlations at small times, which disappear at the end of the training. To better characterize this angular structure, we plot the Fourier transform (along neuron indices) of the control stimulations in Fig. 6(f). In the initial stage of training, the stimulation is essentially represented by cosine and sine waves, with periods matching the E and I ring extensions. As learning proceeds, higher and higher wave-number modes are recruited to refine the short scale structure of the synaptic matrix. This observation is in agreement with the progressive emergence of connectivity in Fig. 6(c).

V. DISCUSSION

We have shown how appropriate stimulations can induce plastic reconfiguration in a network subject to Hebbian-like plasticity processes and make it capable of accomplishing a prescribed computation. Our approach could successfully be applied to different tasks (see also the SM, Sec. 6) and plasticity rules (Hebbian and anti-Hebbian, see the SM, Secs. 5.3 and 4.3). These results can thus be seen as a proof-of-concept for general-purpose neural training, which would compile a task into a neural network by externally guiding the intrinsic learning mechanisms. At its core, our training method addresses and solves a sort of inverse Hebbian problem: instead of looking for how synaptic interactions change under external inputs, we search for an appropriate set of inputs, or control stimulations implementing desired changes in the interactions through plasticity mechanisms. Notice that, strictly speaking, the issue of modifying connectivity with constrained control is not unique to biological neural networks: neuromorphic circuits, e.g., based on memristors, can undergo plastic changes [28], and computation in such networks could be among potential applications of our training protocol. We now discuss

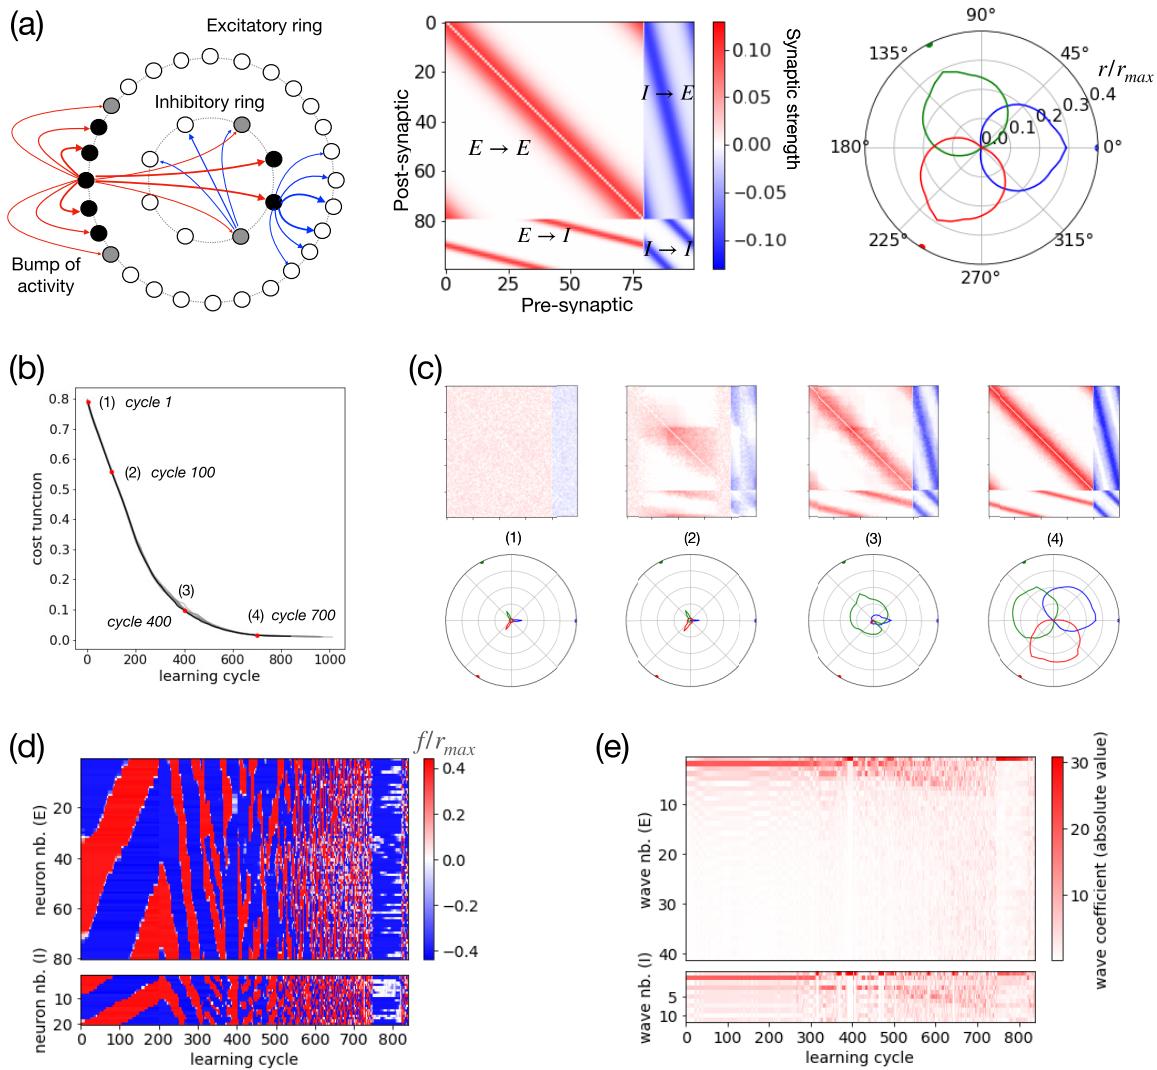


FIG. 6. Task 2: building a continuous attractor. (a) Target network (left) and connectivity matrix $\mathbf{J}^{\text{target}}$ (middle). Excitatory neurons (E) are arranged on the outer ring ($N_E = 80$), and inhibitory neurons (I) on the inner one ($N_I = 20$). Neurons on the E ring have strong excitatory connections with their neighbors, and project to inhibitory neurons diametrically opposed on the I ring. Inhibitory neurons repress neighboring neurons on the E ring and repress the opposite side of the I ring, leading to a localization of the activity (bump). Right: Radial plots of receptive fields of neurons associated with angles 0° (blue), 120° (green), and 240° (red) degrees on the E ring. A weak localized input stimulation is applied for each angle (Methods) and the activities (averaged over 100 trainings) of all neurons in the stationary state are shown. (b) Task cost evolution during learning for 10 random naive networks (gray curves); the back line shows the average cost. Four representative learning cycles, labeled (1), (2), (3), and (4), are referred to in panel (c). (c) Connectivity matrices (top) and receptive fields (bottom) at cycles (1), (2), (3), and (4), showing different stages of learning. Same color codes as in panel (a). (d) Control stimulations $f_i^\epsilon(t)$ as functions of the learning cycle. (e) Amplitude $\hat{f}_k^\epsilon(t) = \sum_{\ell=1}^{N_\epsilon} f_\ell(t) \cos(2\pi k \ell / N_\epsilon)$ of the Fourier modes $k = 0, 1, \dots, N_\epsilon/2$ associated with the control stimulations shown in panel (d), for excitatory ($\epsilon = E$) and inhibitory ($\epsilon = I$) neurons. Initially, the control stimulation mainly consists of large waves (low- k Fourier modes), while at the end of training, modes at large k are used to refine the synaptic structure on short angular scales. Notice the global suppression at cycle $\simeq 750$, after the receptive fields are formed.

our results, both from the computational and the bioengineering/biological points of view.

A. Computational aspects

Constrained learning and approximate gradient descent. As stressed in the Introduction, plasticity rules strongly constrain the manifold of possible changes in the connectivity matrix, and straightforward gradient descent of the loss is generally not feasible [Fig. 3(a)]. This fundamental limitation

is clearly present even in the case of very simple tasks trainable on small networks (SM, Sec. 6 and SM, Fig. 7). As the number N of neurons increases, the ratio between the number of control variables (N) and the dimension of the connectivity space (N^2 , or a fraction of N^2 depending on network sparsity) diminishes, and the situation ought to become less and less favorable. Fortunately, this effect is counterbalanced by the huge increase in the multiplicity of the networks realizing the task with N [29]. We observe, indeed, that Task 1 is easier to learn when the size N_{proc} of the processing area [Fig. 4(b)]

increases; see SM, Sec. 4.1 and SM, Fig. 3. As a result, functional targets are easier to reach than structural ones, as the latter cannot benefit from any multiplicity of solutions.

An additional difficulty is the impossibility of using an adaptive learning rate, a key ingredient in the convergence of most machine-learning algorithms. Though some variability in the duration Δt of the control can be considered, this is constrained by biological and experimental timescales and cannot vary over orders of magnitude. As a result, the magnitude of the changes $\Delta \mathbf{J} \propto \Delta t \mathbf{J}$ in the connectivity after each cycle cannot be tuned at will.

Choice of cost function: structural versus functional targets. The nature of the cost U has also a deep impact on the time course of performance throughout the learning process. Any task could, in principle, be given a structural cost, e.g., by running gradient descent of the functional cost on a computer and finding an adequate network \mathbf{J}^* implementing the target task. However, in the structural case, the value of U is not immediately informative about the computation carried out by the network. Good functional performance, such as the ability to create angle-specific receptive fields in Task 2, see Fig. 6(d), may be reached for vanishing U only [SM, Sec. 5.3 and SM, Figs. 8(a) and 8(b)], which is difficult to achieve in practice. For this reason, functional costs, for which low values of U directly imply good performances, are preferable. Accordingly, we find that defining U from the worst input-output mismatch, see Methods, Eq. (20), rather than from the average over all pairs as in Eq. (4), gives better performance for Task 1.

A consequence of the multiplicity of solutions in the functional case is that relatively little changes to the initial (and random) network are generally sufficient to meet the target, compared to what is needed to meet a structural target. This intuition is corroborated by the analysis reported in SM, Sec. 4.5 and by the animation provided in SM, Sec. 7.2. We observe that the internal activity of the network remains strongly self-correlated in time during function-based training, providing evidence for the ability of the learning procedure to reach a functional target with moderate changes to the connectivity.

Model uncertainty and robustness. A crucial point to be addressed in future works is the impact of noise and uncertainty. Dynamical noise, both in the neural activity and in the synaptic transmission, could be taken into account through a probabilistic formulation of the control problem. In addition, one should consider measurement errors in the activity recordings, affecting the estimation of the connectivity, or in the stimulation process, leading to inaccurate control of the network.

Taking into account model uncertainty is more delicate. As a matter of fact, neither the equations describing the neural activity dynamics nor the plasticity rules are exactly known. Any inaccuracy in the model is potentially a source of bias and of correlated errors in the determination of the stimulation protocol. To estimate the impact of such errors, we have studied the performance of our learning scheme in a mismatch setting, in which the plasticity rules governing the evolution of the connectivity and the ones used for the computation of the protocol differed. We show in SM, Sec. 4.2, that successful training is achieved even with 10% (relative) mismatches on the parameters entering the plasticity rules. This result is a good indication of the robustness of our approach to

uncertainties in the model. Needless to say, controlling errors in the modeling of the plasticity rules and designing robust stimulation schemes will be crucial for future developments.

Importance of the range of controls. At the implementation level, our training procedure involves high-dimensional optimization over the control stimulation \mathbf{f} , and requires careful numerical implementation. A number of parameters in the optimization loop have to be set, whose choice impacts not only the outcome but also the speed of learning; see the SM, Secs. 4 and 5. Of crucial importance is the range of variation of the control stimulation applicable to the neurons, $f_{\min} \leq f_i \leq f_{\max}$. If the range is wide, a lot of variability in the control stimulation schemes is possible, and learning is easy and fast. As $f_{\max} - f_{\min}$ decreases, the cost decays more slowly [SM, Fig. 3(a)]. Further investigations would be needed to better understand if a minimal value of $f_{\max} - f_{\min}$ exists below which the target task is not reachable any longer, and, more generally, how reachability depends on the task. Notice that the protocol we have implemented corresponds to a control strategy with zero time horizon, which optimizes the cost change ΔU at each cycle. Global protocols, derived in the framework of optimal control [30], could possibly achieve higher performance. However, our high-dimensional setting with incomplete knowledge (of connectivity) could suffer from serious error propagation over large time horizons, making implementation difficult.

Continual learning. The final configuration of the network (reached at the end of the control scheme) is in general not a stable state of the learning dynamics and would be lost over time. This is a general issue with systems undergoing continual learning [31]. From a computational point of view, one would either need to reapply stimulations to maintain the neural circuit in its state or use a cost function that ensures that the final configuration is stable. From a biological point of view, plasticity processes can be downregulated in many ways *in vivo* [11]. Diminishing plasticity is, for instance, essential at certain stages of brain development [32].

B. Biological aspects

A new framework, in vitro biological computing. Our work is deeply motivated by the emerging field of *in vitro* computation with biological neurons. The main approach in this domain so far has been reservoir computing [9,10,33–35], which exploits the capability of randomly connected neural populations to produce highly variable patterns of activity, and has been linked to plausible neurobiological mechanisms [36,37]. Associative effects of plasticity can be exploited for the sake of learning, and reach performance beyond standard reservoir setups [38,39]. More targeted approaches have been proposed to directly exploit plasticity for task-training, for instance by choosing familiar/unfamiliar inputs as rewards/punishment [3] or by exploiting stimulation avoidance [40]. Our approach, on the contrary, aims at directly controlling and training the neural structure through the determination of adequate stimulations to be imposed to the neurons.

The idea of learning through stimulations is not new to biology, and it is the basis for the concept of vaccination. The immune system is endowed with learning mechanisms that

allow it to respond to pathogens after appropriate stimulations, e.g., following the administration of a pathogen protein or its coding RNA sequence. The vaccination scheme (timing and dosage) is crucial to the quality of the immunization [41], and it can be designed to enhance the production of the so-called broadly neutralizing antibodies having large response spectra [42]. By analogy, one may expect that stimulation protocols optimized to the learning dynamics of neural systems could drive them to desired operational states.

Limitations and plausibility. Our modeling approach is highly simplified and unlikely to capture detailed biological mechanisms, and it is primarily intended to explore conceptual challenges in the training of networks with constrained plasticity. Though the learning procedure is robust against mismatches between the “true” and modeled plasticity mechanisms (SM, Sec. 4.2), it relies on a number of assumptions that need to be discussed in terms of their technological and biological plausibility. First, we have assumed that stimulations could be done at the level of individual neurons: fine stimulation and recording capabilities are being developed [43] but, as of today, optogenetic [5,44] and electrophysiological [45] techniques rather allow one to target groups of nearby neurons. This limitation could be taken into account in our approach by imposing spatial correlations between the components of \mathbf{f} . Second, our protocol could be adapted to constrained controls \mathbf{f} , e.g., whose components have a given sign (though excitatory and inhibitory tools are available [5,46,47]), or discrete (on/off) rather than continuous values. Though these limitations could in principle be enforced, it remains to be seen how they would affect the learning procedure, both in terms of performances and training time.

Separation of neural and plasticity timescales. An important hypothesis is that the timescale τ_r associated with the relaxation of neural activity is much shorter than the duration of the control stimulation during a cycle, Δt , over which the connectivity changes. Under this assumption, the neural activity is stationary, and the estimate of the connectivity can be updated by probing the network response to a few random stimuli without inducing synaptic changes [Methods, Eq. (12)]. Physiological timescales related to single-neuron dynamics [τ_n in Eq. (1)] are expected to be of the order of tens of milliseconds. It is then crucial that network effects do not slow down the relaxation of the population activity on times $\tau_r \gg \tau_n$; see Methods, Eq. (8). This is ensured in practice in Task 1 through the introduction of a regularization cost on the connectivity, U_{reg} , which prevents the emergence of slow dynamics modes in the network [Methods, Eq. (13)]. We have checked that τ_r remains comparable to τ_n throughout the learning process for all the tasks considered here (SM, Sec. 2). While it is difficult to estimate with precision the timescale τ_s associated with plasticity effects in Eq. (2) and, consequently, the duration Δt of control stimulation, a biologically plausible order of magnitude for Δt is of a few tens of seconds (SM, Sec. 1). Note that this is loosely consistent with meaningful plastic modifications happening over a few minutes, such as those shown in [3,40]. We conclude that the hypothesis $\tau_r \ll \Delta t$ is satisfied. Furthermore, this estimate of Δt would imply that running hundreds of learning cycles would take a few hours. We are well aware that these considerations are speculative and indicative of orders of magnitude at best.

Intensities of stimulations. As emphasized above, having a large range of stimulation values at our disposal is crucial for reaching the target task at the end of learning. The values of f_{max} we have employed in Tasks 1 and 2 are a few tens of r_{max} , see Figs. 4(c) and 6(d); details about how to convert input currents into rates can be found in the SM, Sec. 1.2. Ultrafast optogenetics can elicit spikes in genetically modified neurons with high probability (for sufficient light pulse intensity) up to hundreds of Hz [48]. Hence, ratios $f_{\text{max}}/r_{\text{max}}$ close to unity are experimentally accessible. It remains to be seen whether these techniques could be applied to organoids with similar results.

Dependence on plasticity mechanism. From a modeling point of view, we have resorted to a coarse-grained description of activity at the level of firing rates, rather than of spiking events. Both descriptions are compatible with one another, at the plasticity [49–51] and activity [52–54] levels, to the point that artificial spiking networks may be trained by proxy rate-based models [55]. Plasticity at the spike level is a complex process involving multiple timescales, delayed effects, and spike train patterns [56], possibly coexisting with nonsynaptic plasticity [57], different effects depending on the origin of the neuron [58], and metaplasticity [51]. The coarse-grained plasticity rules we consider respect some general principles: associativity (covariance rule with a threshold, compatible with experiments [59]), locality (dependence on the post- and presynaptic neurons [58]), and homeostasis [18]. In addition, we test the robustness of our approach with two different plasticity scenarios, based on Hebbian and anti-Hebbian rules for inhibitory synapses in, respectively, Tasks 2 and 3; see the SM, Sec. 5.3 for reverse choices of the rules for inhibitory synapses. In this context, Hebbian and anti-Hebbian learning rules induce, respectively, the weakening and the strengthening of connections between inhibitory neurons with correlated activity, and they result in increased or diminished [17] activity. It would be interesting to exhaustively study the performances of our stimulation protocol over a wider class of plasticity rules.

Relevance for neuroscience. As emphasized above, our modeling approach, besides being an oversimplification of the processes taking place in neurons and in their connections, focuses on *in vitro* circuits. Such circuits share many features, at the molecular and cellular levels, with their counterparts in the brain of animals, but they also differ in fundamental ways. In particular, our learning protocol, based on control stimulations computed externally, was designed to meet bioengineering contexts and goals, but is *a priori* not directly relevant for neuroscience, e.g., to understand how animals learn in realistic situations.

From this point of view, the interpretation of the training stimulation sequence is an important point to consider. We have observed that, in the case of the structural Task 2, the structure of the training stimulation adapts to the ring symmetry of the target [Figs. 6(d) and 6(e)]. Briefly speaking, the stimulation protocol engraves Fourier modes on the ring at finer and finer scales as (training) times pass on. This finding shows the deep connection between the spatiotemporal pattern of stimulations and the structure of the target network. Interestingly, the possibility of shaping a network with structured activity patterns is reminiscent of what is speculated to happen during postnatal development, for instance in the visual cortex

[60]. It would be interesting to further investigate this analogy, in particular if the activity driving neural development could be interpreted as the result of some sort of control optimization.

VI. METHODS

A. Network structure and activation function

The fractions of excitatory and inhibitory subpopulations are equal to, respectively, 80% and 20%. The synaptic structure is encoded in the binary matrix C : neuron j is connected to neuron i if $C_{ij} = 1$, and it is not if $C_{ij} = 0$. The structure matrix C is random and fixed during the training, while the intensities J_{ij} of the existing connections ($C_{ij} = 1$) vary over time according to Eq. (2).

The neuron activation Φ is a monotonously increasing function of the input x (which is dimensionally a frequency, see the SM) and is chosen to be

$$\begin{aligned} \Phi(x) &= r_{\max} \frac{\psi(x)}{1 + \psi(x)}, \quad \text{where} \\ \psi(x) &= \frac{r_0}{r_{\max}} \log \left[1 + \exp \left(\frac{x}{r_0} \right) \right] \end{aligned} \quad (7)$$

is a differentiable and invertible sigmoid function. Here, r_{\max} is the maximal firing rate reached for large inputs, and r_0 is the magnitude of the spontaneous firing activity in the absence of input (SM, Sec. 1).

B. Relaxation dynamics of the network activity

While single-neuron activities are associated with the time constant τ_n in Eq. (1), the effective relaxation time τ_r can be longer due to network effects. We estimate τ_r by linearizing the dynamics of the rates around their stationary values, r_i . Denoting by \mathbf{D} the N -dimensional diagonal matrix with elements $D_{ii} = \Phi'(r_i)$ and zero outside the diagonal, we have

$$\begin{aligned} \tau_r &= \frac{\tau_n}{\rho_{\min}}, \\ \text{where } \rho_{\min} &= \begin{array}{l} \text{smallest real part} \\ \text{of the eigenvalues of } \mathbf{Id} - \mathbf{D} \cdot \mathbf{J}. \end{array} \end{aligned} \quad (8)$$

Estimates of τ_r are shown for the different tasks in the SM, Fig. 1.

C. Inference of connectivity

Initialization: estimate of the synaptic structure C and of the neuron types ϵ . The support of the synaptic interactions, $C_{ij} = 0, 1$, and the nature of neurons, $\epsilon_i = E/I$, are initially unknown. To determine their values, we start to probe the network activities r_i^ν in response to n_{probe} randomly chosen stimulation patterns f_i^ν , and we rewrite the stationary relations in Eq. (5) as

$$\sum_j J_{ij} r_j^\nu = t_i^\nu, \quad \text{where } t_i^\nu = \Phi^{-1}(r_i^\nu) - f_i^\nu \quad (9)$$

for all i, ν . As the initial number of probing stimulations does not have to be small (plastic changes would simply affect the starting state of the connectivity for the training phase), we

can choose $n_{\text{probe}} \geq N$, and the linear system above generically determines the matrix \mathbf{J} in a unique way. To avoid large numerical errors on currents through Φ^{-1} , the probing stimulations f_i^ν are chosen to be large enough to avoid the presence of low firing rates ($r \simeq 0$); see SM, Sec. 1.

Once J has been estimated, we estimate the identity of neurons and the support of the connectivity through, respectively,

$$\epsilon_i = \begin{cases} E & \text{if } \sum_j J_{ji} > 0, \\ I & \text{otherwise} \end{cases} \quad (10)$$

and

$$C_{ij} = \begin{cases} 1 & \text{if } |J_{ij}| > J_{\min}, \\ 0 & \text{otherwise,} \end{cases} \quad (11)$$

where J_{\min} is a small threshold, set to 10^{-6} . These estimates of ϵ and \mathbf{C} are left unchanged throughout the training process.

Updating the connectivity estimate after a training period. Let \mathbf{J}_i be the vector of connections incoming onto neuron i prior to a training stimulation period of duration Δt ; the dimension of \mathbf{J}_i is equal to $d_i = \sum_j C_{ij}$. After the training period, the vector connection is \mathbf{J}'_i . To estimate this vector, we probe the network with n_{probes} stimuli, and we record the corresponding activities. We call \mathbf{R} the $(n_{\text{probes}} \times d_i)$ -dimensional matrix of activities r_j^ν , and \mathbf{t}_i the n_{probes} -dimensional vector of components t_i^ν appearing in Eq. (9). The solution to the constraints $\mathbf{R} \cdot \mathbf{J}'_i = \mathbf{t}_i$ is therefore not unique. To lift this ambiguity, we select the solution \mathbf{J}'_i closest (in terms of L_2 norm) to the previous estimate \mathbf{J}_i . Formally, we obtain [61]

$$\mathbf{J}'_i = \Pi_i^\perp \cdot \mathbf{J}_i + \mathbf{R}^+ \cdot \mathbf{t}_i, \quad (12)$$

where \mathbf{R}^+ denotes the pseudoinverse of \mathbf{R} , and $\Pi_i^\perp = \text{Id} - \mathbf{R}^+ \cdot \mathbf{R}$ is the projector orthogonal to the n_{probes} -dimensional space spanned by the rows of \mathbf{R} . This updating step is then iterated over the N neurons i to obtain all the vectors \mathbf{J}'_i .

D. Regularization of the connectivity matrix

A common practice in machine learning is to introduce a regularization cost over the model parameters, hereafter denoted as $U_{\text{reg}}(\mathbf{J})$, to smooth out the training trajectories. In addition, regularization can help ensure that the neural activities reach their stationary values quickly (under fixed stimulation) through the network dynamics in Eq. (1), which is important for training and for connectivity estimation.

We impose regularization to bound the modulus $\sigma(\mathbf{J})$ of the largest singular value of the connectivity matrix. The rationale is that a sufficient condition for the network dynamics to have a unique activity stationary state and converge exponentially fast to this unique state is that $\sigma(\mathbf{J}) < 1$; see Eq. (8) and the SM, Sec. 2.2. For each singular value λ_k of \mathbf{J} , the regularization cost is approximately zero if $\lambda_k \ll 1$, and it grows linearly as $g_2 \lambda_k$ for $\lambda_k \gg 1$. We implement this dependence through a softplus function:

$$U_{\text{reg}}(\mathbf{J}) = \frac{g_2}{g_1} \sum_{k=1}^{N^2} \log(1 + e^{g_1 (\lambda_k - 1)}) \quad (13)$$

with $g_2 = 2$, $g_1 = 10$. This regularization cost can be explicitly differentiated with respect to the entries of the connectivity matrix; see below.

E. Determination of the optimal stimulation \mathbf{f}^*

Derivative of the costs with respect to connections and time. The task cost U_{task} depends on the connectivity \mathbf{J} , both directly and indirectly through the stationary activity \mathbf{r} . The expression for the total derivative of U_{task} with respect to an entry of \mathbf{J} is, according to the chain rule,

$$\frac{dU_{\text{task}}}{dJ_{ij}} = \frac{\partial U_{\text{task}}}{\partial J_{ij}} + \sum_k \frac{\partial U_{\text{task}}}{\partial r_k} \frac{\partial r_k}{\partial J_{ij}}, \quad (14)$$

which requires knowledge of the derivatives of the neuron firing rates. These can be computed using the implicit function theorem [62] applied to the stationary equations for the dynamics, with the result

$$\frac{\partial r_k}{\partial J_{ij}} = [(\mathbf{Id} - \mathbf{D} \cdot \mathbf{J})^{-1}]_{k,i} \mathbf{D}_{ii} r_j, \quad (15)$$

where \mathbf{D} was first used in Eq. (8). The gradient of the regularization cost, U_{reg} in Eq. (13), can be easily expressed using the derivatives of the singular values: $\frac{d\lambda_k}{dJ_{ij}} = u_{k,i} v_{k,j}$, where \mathbf{u}_k and \mathbf{v}_k denote, respectively, the left and right eigenvectors of \mathbf{J} associated with λ_k .

In addition, the gradient component associated with the connection J_{ij} is set to zero if $C_{ij} = 0$ (no connection is actually present), or if $J_{ij} = 0$ and descending the gradient would violate the constraint on the sign of the interaction resulting from the presynaptic neuron type, i.e., $\frac{dU}{dJ_{ij}} > 0$ if $\epsilon_j = E$ or $\frac{dU}{dJ_{ij}} < 0$ if $\epsilon_j = I$. Notice that knowledge of the gradients of U_{task} and U_{reg} with respect to J_{ij} gives access to the time derivatives of these costs (under fixed stimulation) through the generic formula

$$\dot{U} = \sum_{i,j} \frac{\partial U}{\partial J_{ij}} \dot{J}_{ij}, \quad (16)$$

where the latter term can be directly obtained from the dynamics over synapses described by Eq. (2).

Objective function and gradient over stimulations. The total cost U is defined as the sum of the task and regularization costs, $U = U_{\text{task}} + U_{\text{reg}}$. To determine the best control stimulation \mathbf{f} over the next period, we search for the minimum of the modified change

$$\Delta U(\mathbf{f}; \gamma) = U(\mathbf{J} + \Delta \mathbf{J}(\mathbf{f})) - U(\mathbf{J}) + \Delta t \gamma \|\Delta \mathbf{J}(\mathbf{f})\|^2, \quad (17)$$

where \mathbf{J} is the matrix of connectivity estimated through probing at the end of the previous control stimulation period, $\gamma = 1/[(\bar{J}/10)^2 c N^2]$, and

$$\Delta \mathbf{J}(\mathbf{f}) = [\mathbf{J} + \Delta t \dot{\mathbf{J}}(\mathbf{f})]_{\epsilon} - \mathbf{J}, \quad (18)$$

where $[\cdot]_{\epsilon}$ indicates that connections have been clipped to respect their excitatory/inhibitory types ϵ imposed by the presynaptic neurons. The last term in Eq. (17) constrains the change in the interactions to be small over the period of

stimulation, effectively introducing a saturating nonlinearity over the gradients.

Computing the gradient of ΔU in Eq. (17) with respect to the control stimulation components f_i requires the expressions for the derivatives of the costs with respect to the connections in Eqs. (14) and (15), and with respect to time; see Eq. (16). In addition, the derivatives of the firing activities are needed. Using again the implicit function theorem we obtain, see Eq. (15),

$$\frac{\partial r_k}{\partial f_i} = [(\mathbf{Id} - \mathbf{D} \cdot \mathbf{J})^{-1}]_{k,i} \mathbf{D}_{ii}. \quad (19)$$

To enforce the conditions $f_{\min} < f_i < f_{\max}$, we set the gradient components of ΔU_{tot} to zero when these boundaries are met.

Minimization of the cost. The optimal control \mathbf{f}^* can be found through an iterative descent of ΔU in Eq. (17). The control is updated along minus the gradient of $\Delta U(\mathbf{f}, \gamma)$, and the resulting f_i 's are clipped if the boundaries f_{\min} or f_{\max} are crossed. The iterative process halts if the decrease of $\Delta U(\mathbf{f}, \gamma)$ is smaller (in absolute value) than some threshold, or the expected decrease (based on a fit of the previous steps in the process) is too low. Then, if $\Delta U(\mathbf{f}, 0)$ is negative (i.e., the cost function is decreasing), the solution is accepted, while the gradient descent is repeated otherwise.

Details about the implementation, with the pseudocode for the optimization loop and the values of the hyperparameters involved, can be found in the SM, Sec. 3.

Task 1

Initialization and setting. Elements of the structure matrix C_{ij} are randomly set to 0 or 1 with respective probabilities $1 - c_E$ and c_E for excitatory columns (J such that $\epsilon_j = E$) and $1 - c_I$ and c_I for inhibitory ones ($\epsilon_j = I$). To ensure that the task is not easily implementable, no direct connection between the N_{in} input neurons and the N_{out} output neurons is present. Initial connections J_{ij} are drawn uniformly at random in the ranges $[0, J_0]$ if $\epsilon_j > 0$ or $[-J_0, 0]$ if $\epsilon_j < 0$. See the SM, Sec. 1 for parameter values.

Task-associated cost. We define a set of n_{task} input stimulations \mathbf{f}^{μ} (defined over the N_i input neurons) and output binary activation σ^{μ} (over the N_o output neurons), with $\mu = 1, \dots, n_{\text{task}}$. The goal of training is that output neuron i should display high or low firing activity in response to input μ when, respectively, $\sigma_i^{\mu} = H$ or L . In practice, we enforce that the firing rates associated with neurons with low ($\sigma = L$) and high ($\sigma = H$) activities should differ by more than a prescribed gap $\delta r = 0.12 r_{\max}$. The task cost reads

$$U_{\text{task}}(\mathbf{J}) = \frac{\sum_{(\mu,i):\sigma_i^{\mu}=L} \sum_{(v,j):\sigma_j^{\mu}=H} \Delta(\mu, i; v, j) e^{\gamma \Delta(\mu, i; v, j)}}{\sum_{(\mu,i):\sigma_i^{\mu}=L} \sum_{(v,j):\sigma_j^{\mu}=H} e^{\gamma \Delta(\mu, i; v, j)}}, \quad (20)$$

where $\gamma = 1/(\delta r/2)^2$ and we have introduced the mismatch

$$\Delta(\mu, i; v, j) = h_2(r_i^{\mu} - r_j^v + \delta r) \quad (21)$$

with

$$h_2(u) = u^2 \quad \text{if } u > 0, \quad 0 \text{ otherwise.} \quad (22)$$

The exponential factors give more weight to large mismatches Δ in the expression of the cost. The dependence of the cost on \mathbf{J} (and on the input stimulations \mathbf{f}^μ) in Eq. (21) is implicit through the firing rates r_i . Minimizing U_{task} is thus equivalent to separating the activities of the H and L neurons as required by the classification task, i.e., making all $\Delta(\mu, i; v, j) = 0$.

Analysis of input and output spaces. To characterize how the association mechanism emerges across training, we convert digits $\sigma_i^\mu = \text{High/Low}$ to 1/0 valued arrays. These digits span an $n_{\text{pairs}} = 4$ -dimensional space. Given a stimulation \mathbf{f} , we call $\mathbf{r}_{\text{out}}(\mathbf{f})$ the output activity and

$$\|\mathbf{r}_{\text{out}}^\perp(\mathbf{f})\|^2 = \frac{\min_{\beta} \left\| \sum_{\mu=1}^{n_{\text{task}}} \sigma_i^\mu \beta_\mu - \mathbf{r}_{\text{out}}(\mathbf{f}) \right\|^2}{\|\mathbf{r}_{\text{out}}(\mathbf{f})\|^2} \quad (23)$$

the normalized squared length of its projection orthogonal to the digit space. In Fig. 5(a), we report this quantity for (i) generic inputs \mathbf{f} with 0,1 entries with 50-50 probability, and (ii) random, uniform convex combinations of the inputs associated with the digits: $\mathbf{f} = \sum_{\mu=1}^{n_{\text{task}}} \alpha_\mu \mathbf{f}_\mu$ with $\alpha_\mu \geq 0$ and $\sum_\mu \alpha_\mu = 1$.

Task 2

Initialization and setting. The connectivity matrix is initialized as in the classification task above, but no *a priori* restriction is imposed to the structure, i.e., $c_E = c_I = 1$. See the SM, Sec. 1 for parameter values.

Task and cost. The structural cost is defined in Eq. (3). Weights are defined as $w_{\epsilon_i, \epsilon_j} = \left(\sum_{ab} [J_{\epsilon_i \epsilon_j}^{\text{target}}]_{ab}^2 \right)^{-1}$, where $J_{\epsilon_i \epsilon_j}^{\text{target}}$ is the submatrix of $\mathbf{J}^{\text{target}}$ connecting neurons of type $\epsilon_j = E/I$ to neurons of type $\epsilon_i = E/I$. This reweighting ensures that large, e.g., $E \rightarrow E$, and small, e.g., $I \rightarrow I$, blocks equally contribute to the cost, and are simultaneously and properly learned during the training process.

The target connectivity is defined as follows. Neurons are associated with angles on two rings, depending on their types $\epsilon = E$ or I . The angle attached to the excitatory neuron i is $\theta_i^E = 2\pi \frac{i}{N_E}$, where N_E is the number of excitatory neurons. A similar formula holds for the N_I inhibitory neurons: $\theta_i^I = 2\pi \frac{i}{N_I}$. The target synaptic connection from neuron j to i reads

$$J_{ij}^{\text{target}} = J_{\epsilon_i, \epsilon_j} \times \mathcal{M}(\theta_i^{\epsilon_i} - \theta_j^{\epsilon_j} - \pi \delta_{\epsilon_i, I}; K_{\epsilon_i, \epsilon_j}), \quad (24)$$

where δ_{\dots} denotes the Kronecker delta. The angular modulation is done through the von Mises function,

$$\mathcal{M}(\Delta\theta; K) = e^{-K+K \cos(\Delta\theta)}, \quad (25)$$

which is maximal for vanishing angular separation ($\Delta\theta = 0$) and decays over the width $\simeq K^{-1/2}$. Parameter values are listed in the SM, Sec. 1.

Receptive fields. We apply a weak local input stimulation $f_i = 0.06 r_{\text{max}}$, $f_{i\pm 1} = 0.028 r_{\text{max}}$, $f_{i\pm 2} = 0.012 r_{\text{max}}$, $f_{i\pm 3} = 0.004 r_{\text{max}}$ to induce polarization of neuron i , and we repeat the process for all i . Results for the activities are averaged over 100 trials.

ACKNOWLEDGMENTS

We thank S. Wolf for many useful discussions, and D. Chatenay for reading the manuscript. This research was supported by the European Union's Horizon 2020 research and innovation programme under Grant No. 964877 (project NEU-ChiP).

All authors planned and conceived the study. F.B. designed the algorithms, ran the simulations, and analyzed the results. F.B. and R.M. wrote the manuscript; all authors reviewed the manuscript.

The authors declare no competing interests.

-
- [1] S. Furber, Large-scale neuromorphic computing systems, *J. Neural Eng.* **13**, 051001 (2016).
 - [2] L. Smirnova, B. S. Caffo, D. H. Gracias, Q. Huang, I. E. Morales Pantoja, B. Tang, D. J. Zack, C. A. Berlinicke, J. L. Boyd, T. D. Harris, E. C. Johnson, B. J. Kagan, J. Kahn, A. R. Muotri, B. L. Paulhamus, J. C. Schwamborn, J. Plotkin, A. S. Szalay, J. T. Vogelstein, P. F. Worley, and T. Hartung, Organoid intelligence (OI): The new frontier in biocomputing and intelligence-in-a-dish, *Front. Sci.* **1**, 1017235 (2023).
 - [3] B. J. Kagan, A. C. Kitchen, N. T. Tran, F. Habibollahi, M. Khajehnejad, B. J. Parker, A. Bhat, B. Rollo, A. Razi, and K. J. Friston, *In vitro* neurons learn and exhibit sentience when embodied in a simulated game-world, *Neuron* **110**, 3952 (2022).
 - [4] M. A. Lancaster, M. Renner, C.-A. Martin, D. Wenzel, L. S. Bicknell, M. E. Hurles, T. Homfray, J. M. Penninger, A. P. Jackson, and J. A. Knoblich, Cerebral organoids model human brain development and microcephaly, *Nature (London)* **501**, 373 (2013).
 - [5] K. Deisseroth, Optogenetics: 10 years of microbial opsins in neuroscience, *Nat. Neurosci.* **18**, 1213 (2015).
 - [6] W. Yang and R. Yuste, Holographic imaging and photostimulation of neural activity, *Curr. Opin. Neurobiol.* **50**, 211 (2018).
 - [7] H. Shin, S. Jeong, J.-H. Lee, W. Sun, N. Choi, and I.-J. Cho, 3D high-density microelectrode array with optical stimulation and drug delivery for investigating neural circuit dynamics, *Nat. Commun.* **12**, 492 (2021).
 - [8] X. Qian, H. Song, and G.-L. Ming, Brain organoids: Advances, applications and challenges, *Development* **146**, dev166074 (2019).
 - [9] G. Tanaka, T. Yamane, J. B. Héroux, R. Nakane, N. Kanazawa, S. Takeda, H. Numata, D. Nakano, and A. Hirose, Recent advances in physical reservoir computing: A review, *Neural Netw.* **115**, 100 (2019).
 - [10] T. Sumi, H. Yamamoto, Y. Katori, K. Ito, S. Moriya, T. Konno, S. Sato, and A. Hirano-Iwata, Biological neurons act as generalization filters in reservoir computing, *Proc. Natl. Acad. Sci. USA* **120**, e2217008120 (2023).
 - [11] A. Citri and R. C. Malenka, Synaptic plasticity: Multiple forms, functions, and mechanisms, *Neuropsychopharmacol.* **33**, 18 (2008).
 - [12] L. Carrillo-Reid, W. Yang, J.-E. Kang Miller, D. S. Peterka, and R. Yuste, Imaging and optically manipulating neuronal ensembles, *Annu. Rev. Biophys.* **46**, 271 (2017).

- [13] B. Liu, M. J. Seay, and D. V. Buonomano, Creation of neuronal ensembles and cell-specific homeostatic plasticity through chronic sparse optogenetic stimulation, *J. Neurosci.* **43**, 82 (2023).
- [14] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PRXLife.2.043014> for (1) further insight into the motivation and plausibility of biological modeling, and explains some mathematical aspects of the model (2) explores the robustness of the model by exploring different parameter choices (3) gives a detailed description of the algorithm implementation (4) provides additional technical analysis to provide further insight.
- [15] F. J. Vico and J. M. Jerez, Stable neural attractors formation: Learning rules and network dynamics, *Neural Proc. Lett.* **18**, 1 (2003).
- [16] G. G. Turrigiano, The self-tuning neuron: Synaptic scaling of excitatory synapses, *Cell* **135**, 422 (2008).
- [17] G. G. Turrigiano and S. B. Nelson, Hebb and homeostasis in neuronal plasticity, *Curr. Opin. Neurobiol.* **10**, 358 (2000).
- [18] F. Zenke, W. Gerstner, and S. Ganguli, The temporal paradox of Hebbian learning and homeostatic plasticity, *Curr. Opin. Neurobiol.* **43**, 166 (2017).
- [19] J. Zierenberg, J. Wilting, and V. Priesemann, Homeostatic plasticity and external input shape neural network dynamics, *Phys. Rev. X* **8**, 031018 (2018).
- [20] P. Földiák, Forming sparse representations by local anti-Hebbian learning, *Biol. Cybern.* **64**, 165 (1990).
- [21] I. M. de Abril, J. Yoshimoto, and K. Doya, Connectivity inference from neural recording data: Challenges, mathematical bases and research directions, *Neural Netw.* **102**, 120 (2018).
- [22] M. Minsky and S. Papert, *Perceptrons: An Introduction to Computational Geometry* (MIT Press, Cambridge, MA, 1969).
- [23] J. Devaud, T. Papouin, J. Carcaud, J. Sandoz, B. Grünewald, and M. Giurfa, Neural substrate for higher-order learning in an insect: Mushroom bodies are necessary for configural discriminations, *Proc. Natl. Acad. Sci. USA* **112**, E5854 (2015).
- [24] S.-I. Amari, Dynamics of pattern formation in lateral-inhibition type neural fields, *Biol. Cybern.* **27**, 77 (1977).
- [25] R. Ben-Yishai, R. Bar-Or, and H. Sompolinsky, Theory of orientation tuning in visual cortex, *Proc. Natl. Acad. Sci. USA* **92**, 3844 (1995).
- [26] J. D. Seelig and V. Jayaraman, Neural dynamics for landmark orientation and angular path integration, *Nature (London)* **521**, 186 (2015).
- [27] T. Kühn and R. Monasson, Information content in continuous attractor neural networks is preserved in the presence of moderate disordered background connectivity, *Phys. Rev. E* **108**, 064301 (2023).
- [28] T. Serrano-Gotarredona, T. Masquelier, T. Prodromakis, G. Indiveri, and B. Linares-Barranco, STDP and STDP variations with memristors for spiking neuromorphic learning systems, *Front. Neurosci.* **7**, 2 (2013).
- [29] E. Marder and A. Taylor, Multiple models to capture the variability in biological neurons and networks, *Nat Neurosci.* **14**, 133 (2011).
- [30] D. E. Kirk, *Optimal Control Theory: An Introduction* (Dover Publications, Mineola, NY, 2004).
- [31] L. Wang, X. Zhang, H. Su, and J. Zhu, A comprehensive survey of continual learning: Theory, method and application, *IEEE Trans. Pattern Anal. Mach. Intell.* **46**, 5362 (2024).
- [32] K. E. Carstens and S. M. Dudek, Regulation of synaptic plasticity in hippocampal area CA2, *Curr. Opin. Neurobiol.* **54**, 194 (2019).
- [33] M. Lukoševičius and H. Jaeger, Reservoir computing approaches to recurrent neural network training, *Comput. Sci. Rev.* **3**, 127 (2009).
- [34] J. B. George, G. M. Abraham, K. Singh, S. M. Ankolekar, B. Amrutur, and S. K. Sikdar, Input coding for neuro-electronic hybrid systems, *Biosystems* **126**, 1 (2014).
- [35] Y. Yada, S. Yasuda, and H. Takahashi, Physical reservoir computing with FORCE learning in a living neuronal culture, *Appl. Phys. Lett.* **119**, 173701 (2021).
- [36] M. Rigotti, D. B. D. Rubin, X.-J. Wang, and S. Fusi, Internal representation of task rules by recurrent dynamics: The importance of the diversity of neural responses, *Front. Comput. Neurosci.* **4**, 24 (2010).
- [37] P. Enel, E. Procyk, R. Quilodran, and P. F. Dominey, Reservoir computing properties of neural dynamics in prefrontal cortex, *PLoS Comput. Biol.* **12**, e1004967 (2016).
- [38] G. B. Morales, C. R. Mirasso, and M. C. Soriano, Unveiling the role of plasticity rules in reservoir computing, *Neurocomputing* **461**, 705 (2021).
- [39] H. Cai, Z. Ao, C. Tian, Z. Wu, H. Liu, J. Tchieu, M. Gu, K. Mackie, and F. Guo, Brain organoid computing for artificial intelligence, *bioRxiv* 2023.02.28.530502; doi: 10.1101/2023.02.28.530502.
- [40] A. Masumori, N. Maruyama, L. Sinapayen, T. Mita, U. Frey, D. Bakkum, H. Takahashi, and T. Ikegami, in *Artificial Life Conference Proceedings* (MIT Press, Cambridge, 2015), pp. 373–380.
- [41] M. Molari, K. Eyer, J. Baudry, S. Cocco, and R. Monasson, Quantitative modeling of the effect of antigen dosage on B-cell affinity distributions in maturing germinal centers, *eLife* **9**, e55678 (2020).
- [42] K. Sprenger, J. Louveau, P. Murugan, and A. Chakraborty, Optimizing immunization protocols to elicit broadly neutralizing antibodies, *Proc. Natl. Acad. Sci. USA* **117**, 20077 (2020).
- [43] H. Adesnik and L. Abdeladim, Probing neural codes with two-photon holographic optogenetics, *Nat. Neurosci.* **24**, 1356 (2021).
- [44] A. M. Stamatakis, M. J. Schachter, S. Gulati, K. T. Zitelli, S. Malanowski, A. Tajik, C. Fritz, M. Trulson, and S. L. Otte, Simultaneous optogenetics and cellular resolution calcium imaging during active behavior using a miniaturized microscope, *Front. Neurosci.* **12**, 496 (2018).
- [45] M. E. Spira and A. Hai, Multi-electrode array technologies for neuroscience and cardiology, *Nat. Nanotechnol.* **8**, 83 (2013).
- [46] B. B. Land, C. E. Brayton, K. E. Furman, Z. LaPalombara, and R. J. DiLeone, Optogenetic inhibition of neurons by internal light production, *Front. Behav. Neurosci.* **8**, 108 (2014).
- [47] X. Han and E. S. Boyden, Multiple-color optical activation, silencing, and desynchronization of neural activity, with single-spike temporal resolution, *PLoS ONE* **2**, e299 (2007).
- [48] T. Mager, D. Lopez de la Morena, V. Senn, J. Schlotte, A. D'Errico, K. Feldbauer, C. Wrobel, S. Jung, K. Bodensiek, V. Rankovi *et al.*, High frequency neural spiking and auditory signaling by ultrafast red-shifted optogenetics, *Nat. Commun.* **9**, 1750 (2018).
- [49] E. M. Izhikevich and N. S. Desai, Relating STDP to BCM, *Neural Comput.* **15**, 1511 (2003).

- [50] D. Bush, A. Philippides, P. Husbands, and M. O'Shea, Reconciling the STDP and BCM models of synaptic plasticity in a spiking recurrent neural network, *Neural Comput.* **22**, 2059 (2010).
- [51] P. Yger and M. Gilson, Models of metaplasticity: A review of concepts, *Front. Comput. Neurosci.* **9**, 138 (2015).
- [52] S. Ostojic and N. Brunel, From spiking neuron models to linear-nonlinear models, *PLoS Comput. Biol.* **7**, e1001056 (2011).
- [53] B. Rueckauer and S.-C. Liu, in *Proceedings of the 2018 IEEE International Symposium on Circuits and Systems (ISCAS)* (IEEE, Piscataway, NJ, 2018), pp. 1–5.
- [54] A. Sanzeni, M. H. Histed, and N. Brunel, Response nonlinearities in networks of spiking neurons, *PLoS Comput. Biol.* **16**, e1008165 (2020).
- [55] S. R. Kheradpisheh, M. Mirsadeghi, and T. Masquelier, Spiking neural networks trained via proxy, *IEEE Access* **10**, 70769 (2022).
- [56] R. C. Froemke and Y. Dan, Spike-timing-dependent synaptic modification induced by natural spike trains, *Nature (London)* **416**, 433 (2002).
- [57] R. Mozzachiodi and J. H. Byrne, More than synaptic plasticity: Role of nonsynaptic plasticity in learning and memory, *Trends Neurosci.* **33**, 17 (2010).
- [58] N. Caporale and Y. Dan, Spike timing-dependent plasticity: A Hebbian learning rule, *Annu. Rev. Neurosci.* **31**, 25 (2008).
- [59] S. Lim, J. L. McKee, L. Woloszyn, Y. Amit, D. J. Freedman, D. L. Sheinberg, and N. Brunel, Inferring learning rules from distributions of firing rates in cortical neurons, *Nat. Neurosci.* **18**, 1804 (2015).
- [60] C. P. Danka Mohammed and R. Khalil, Postnatal development of visual cortical function in the mammalian brain, *Front. Syst. Neurosci.* **14**, 29 (2020).
- [61] D. C. Montgomery, E. A. Peck, and G. G. Vining, *Introduction to Linear Regression Analysis* (Wiley, Hoboken, NJ, 2021).
- [62] W. Rudin, *Principles of Mathematical Analysis* (McGraw-Hill, New York, NY, 1953).