



HAL
open science

Supervised task learning via stimulation-induced plasticity in rate-based neural networks

Francesco Borra, Simona Cocco, Rémi Monasson

► **To cite this version:**

Francesco Borra, Simona Cocco, Rémi Monasson. Supervised task learning via stimulation-induced plasticity in rate-based neural networks. 2023. hal-04339067v1

HAL Id: hal-04339067

<https://hal.science/hal-04339067v1>

Preprint submitted on 12 Dec 2023 (v1), last revised 28 Nov 2024 (v2)

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 - NonCommercial - NoDerivatives 4.0 International License

Supervised task learning via stimulation-induced plasticity in rate-based neural networks

Francesco Borra^{1,*}, Simona Cocco¹, and Rémi Monasson¹

¹Laboratoire de Physique de l'Ecole Normale Supérieure, PSL Research, CNRS UMR8023, Sorbonne Université, 24 rue Lhomond, 75005 Paris, France

ABSTRACT

Synaptic plasticity dynamically shapes the connectivity of neural systems and is key to learning processes in the brain. To what extent plasticity mechanisms, which are intrinsically unsupervised, can be exploited to make a neural network achieve any computational task remains unknown. Here, we present a self-contained procedure which, through appropriate spatio-temporal stimulation control 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 tasks: a non-linear 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* biological computation, e.g. based on brain organoids.

Introduction

Recent progress makes now possible to grow, preserve, and study brain organoids¹, raising a surge of interest in *in vitro* biological computation^{2,3}. The capabilities of stimulating and recording neural populations⁴⁻⁶ 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 challenges raised by organoid technologies⁷, it remains conceptually unclear how they could be, in practice, best used and controlled to achieve high-level computation.

Computation with neural cultures has been largely based so far on the framework of reservoir computing^{8,9}, in which the readout of the high-dimensional neural activity is trained to perform the desired task. 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 circuitry¹⁰. These mechanisms can potentially be exploited to shape, through appropriate stimulations, the activity and the connectivity of networks¹¹. In this respect, a key question emerges: what spatio-temporal stimulation patterns are capable of remodeling a plastic neural network and make it achieve a desired computational task (Fig. 1A)? Deriving guiding principles and practical tools to answer this question is crucial to future progress in biological computation.

From a machine-learning point of view, supervised learning of a task is generally cast as an optimization problem in very high dimensions. One looks for the minimum of the cost (loss) function expressing the mismatch between the target and actual computation carried out by the network. In artificial nets, the cost 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 connectivity space, in particular the one along the gradient of the cost, 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 are functions of the firing rates of the neurons. As the number of neurons is generally much smaller than the number of synaptic interactions, plastic changes of different connections are highly inter-dependent, and cannot be individually tuned to match the gradient components. This fundamental limitation makes learning in biologically-plausible networks conceptually much more intricate than in their artificial, unconstrained counterparts. 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 work, we show how to bypass these intrinsic limitations, and how to train plastic networks to accomplish

desired computations in the case of neural rate-based models. We determine time and neuron dependent control stimulations capable of inducing adequate neural activity, progressively driving the network connectivity through its intrinsic Hebbian-like plasticity towards a desired state (Fig. 1A). Figure 1B 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 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.

Modeling framework: neural activity, connectivity, target, and stimulation protocol

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 pre-synaptic neuron j to the post-synaptic 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 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} (Methods and SI, Section 1). Control stimulations f_i are expressed in units of r_{max} , while connections J_{ij} are dimensionless (SI, Section 1).

Synaptic plasticity

Plasticity induces activity-dependent changes in the interactions. We assume that these changes derive from a Hebbian-like covariance rule with a post-synaptic threshold depending on the neuron type, $\varepsilon = E$ or I ¹², and are stabilized by two homeostatic feedbacks^{13–16}. The first one biases the activity of the post-synaptic neuron towards a baseline activity θ_0 , while the second one imposes a soft clipping of synaptic amplitudes outside the range $[-\bar{J}; +\bar{J}]$. The resulting equation for the synaptic dynamics is

$$\tau_s \frac{dJ_{ij}}{dt}(t) = \underbrace{\eta(\varepsilon_j) (r_i - \theta(\varepsilon_j)) r_j}_{\text{hebbian}} - \underbrace{\beta_1 |J_{ij}| (r_i^2 - \theta_0(\varepsilon_j)^2)}_{\text{homeostasis 1}} - \underbrace{\beta_2 \text{sign}(J_{ij}) h_2(|J_{ij}| - \bar{J})}_{\text{homeostasis 2}} \quad (2)$$

where ε_j is the pre-synaptic neuron type, and $h_2(u) = u^2$ if $u \geq 0$, 0 if $u < 0$. Plasticity for excitatory and inhibitory connections are assumed to have different learning rates $\eta(E)$ and $\eta(I)$ and are associated to 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 can describe Hebbian as well as anti-Hebbian learning. The former is obtained when $\eta(\varepsilon_j)$ is positive, independently of the pre-synaptic neuron type ε_j . To accommodate for anti-Hebbian inhibition¹⁷, 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.* $\tau_n \ll \tau_s$ (Methods and SI, Section 2). Due to this assumption only the slow variations (on the τ_s time scale) of the stimulations $\mathbf{f}(t) = \{f_i(t)\}$ matter. Furthermore, the neural activities r_i are quasi-stationary and locked to the slowly varying inputs. In mathematical terms, the r.h.s. of Eq. (1) vanish.

Target state of the network

Our goal is to drive the neural network connectivity $\mathbf{J}(t)$ in order to achieve some target, either structural or functional (Fig. 1A). In the former case the network connectivity is asked to reach some value, say, \mathbf{J}^{target} . The control stimulation should ensure that the loss

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

decays towards zero over time. We use weights $w_{\varepsilon_i, \varepsilon_j}$ 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 $\mu = 1, \dots, n_{pairs}$. A possible loss associated to this association task reads

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

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.

Control of network connectivity through stimulations

Training protocol. Driving the connectivity $\mathbf{J}(t)$ is done by applying piecewise constant external control stimulations $\mathbf{f}^*(t)$ to the neurons over successive intervals of duration Δt (Fig. 2). Informally speaking the control stimulation $\mathbf{f}^*(t)$ is such that the synaptic changes induced by plasticity decrease the loss as much as possible, that is,

$$\mathbf{f}^*(t) = \underset{\mathbf{f}}{\operatorname{argmin}} U(\mathbf{J}(t) + \Delta\mathbf{J}(\mathbf{f})), \quad (5)$$

where the change $\Delta\mathbf{J}$ in the connectivity depends on the applied control stimulation \mathbf{f} through their Hebbian-like dynamics in Eq. (2) over the period Δt . \mathbf{f}^* in Eq. (5) is such that the change in connectivity is best aligned along the gradient of the loss, see Fig. 3.

The determination of \mathbf{f}^* is a non trivial computational problem due to the non-linearities in the neural and synaptic dynamics, see Methods. Choosing $\Delta t \ll \tau_s$ (SI, Section 1) and additional regularization terms (Methods) ensure that the changes in \mathbf{J} remain small at the end of the training cycle. The learning process stops when the value of the loss U is considered small enough and the target has been reached.

Recordings of the activity and estimation of the network connectivity. Determining and updating the optimal control stimulation $\mathbf{f}^*(t)$ requires to know, to some degree of accuracy, the connectivity \mathbf{J} ¹⁸. To mimic realistic conditions, \mathbf{J} is not accessible through direct measurements, but is inferred through a simple procedure requiring multiple recordings of the activity (Fig. 2). We probe the network stationary activity states $\mathbf{r}^v = \{r_i^v\}$ corresponding to various random stimulation patterns $\mathbf{f}^v = \{f_i^v\}$, where $v = 1, \dots, n_{probes}$. The connectivity matrix is then estimated so that the stationary equations

$$r_i^v = \Phi\left(\sum_j J_{ij} r_j^v + f_i^v\right) \quad (6)$$

are fulfilled for all $i = 1, \dots, N$ and $v = 1, \dots, n_{probes}$, 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} : 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} (Methods). In practice, the connectivity varies little after each control stimulation period, and few recordings are needed to update our estimate.

Results

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

Task 1: Non-linear association between all-or-nothing inputs and digit-like outputs

Definition of the task. A fundamental computation carried out by neural circuits in organisms is the association of diverse behavioural (or motor) responses to multiple sensory inputs. This task is particularly non trivial to learn when it is non additive, as it cannot be realized by linear networks¹⁹, and may require to recruit dedicated brain areas²⁰.

We consider here a toy version of a non-linear 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. 4A. For the sake of visual clarity, the output patterns are represented as digits, with grey/black pixels corresponding to neurons with low/high activities arranged on a 5×3 grid. To render this task harder, we impose that

- there is no direct connection from the input to the output neurons, see Fig. 4B. 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 arrangement mimics the idealized structures of a modular chip.
- the association is non-linear. For instance, adding the input stimulations associated to digits 0 and 1 is required to elicit digit 2, which is quite distinct from the superimposed 0 and 1 digits, see Fig. 4A. As a result of non-linearity the association task cannot be simply implemented through separate neural pathways, activated by one input and silent for the other ones.

Training: costs and dynamics. To learn this task we introduce a cost over the connectivity matrices favoring the desired input/output associations, see Methods, Eq. (19). An additional cost aiming at regularizing the connectivity and speeding up convergence towards stationary states of activity is introduced, see Methods.

We start from a network with weak, random connections. As our spatio-temporal control stimulation is applied to the neurons (Fig. 4C), the connectivity gets modified, and the costs associated to the task and to the regularization decrease, as shown in Fig. 4D. The activities of neurons progressively cluster into two categories, depending on their expected values in the output patterns (Fig. 4E). 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. 4F. We let the protocol proceed until the two groups (active, inactive) are clearly separated, and the output neuron activities match the target patterns, compare Fig. 4A&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 sub-populations 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. 5A that learning takes place through the strengthening of a 4-dimensional communication ‘channel’ linking the input and outputs.

From a circuit point of view, this functional channel is supported by a network of connections, whose histogram is reported in Fig. 5B. 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.

We stress that, within this low-dimensional space, the network computation is not linear as imposed by the non-additivity of the task. We show in Fig. 5C 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 behaviour 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 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 sub-populations. We see in Fig. 5D 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 to the association task can be found in SI, Figs. 5.

Task 2: Ring-like connectivity supporting continuous attractor dynamics

Definition of the task. We now aim at reshaping network connectivity into a target matrix \mathbf{J}^{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^{21,22}, and have recently been observed in the ellipsoid body of fly²³. Neuron connections are organized along two rings, see Fig. 6A. Excitatory neurons form the outer ring, with connections decays with their distance, while inhibitory neurons compose the smaller inner ring. The connections between the two rings are such that a bump of activity of one side of the outer ring induces inhibition on the diametrically opposite side (see Fig. 6A). The position of the bump is arbitrary in the absence of external input, and is otherwise attracted by a weak and localised input to the outer neurons.

Control stimulations are computed with the cost function given by Eq. (3). The time behaviour of the cost is shown in Fig. 6B for ten different random initial networks (see Fig. 6C, left, for a realization of the naive connectivity). After a fast initial decay the cost relaxes to very low values, signalling 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. 6C. The strong similarity between the network connectivity at the end of learning (compare Figs. 6A and

C, step (4)) confers the desired functional properties of continuous attractors to the network. We report in Fig. 6D 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 sub-population of neurons emerges, whose activity is supported by the recurrent connections. This sub-population varies with the angle associated to 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 network²⁴.

The time behavior of the control stimulations applied to neurons during learning are shown in Fig. 6E. 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. 6F. 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. 6C.

Discussion

We have shown how appropriate stimulations can induce plastic reconfiguration in a network and make it capable of accomplishing a prescribed computation. Our approach could successfully be applied to different tasks (see also SI, Section 6) and plasticity rules (Hebbian and anti-Hebbian, see SI, Sections 5.3 and 4.2). These results can be thus 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.

The idea of learning through stimulations is not new to biology, and 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 immunization²⁵, and can be designed to enhance the production of the so-called broadly neutralizing antibodies having large response spectra²⁶. By analogy, one may expect that stimulation protocols optimized to the learning dynamics of neural systems could drive them to desired operational states.

From this point of view, our work addresses a fundamental conceptual issue in the emerging field of neuro-biological computation. Reservoir computing^{8,9,27–29}, which exploits the capability of randomly connected neural populations to produce highly variable patterns of activity, has been linked to plausible neuro-biological mechanisms^{30,31}. Associative effects of plasticity can be exploited for the sake of learning, and reach performance beyond standard reservoir setups^{32,33}. More targeted approaches have been proposed to directly exploit plasticity for task-training, for instance by choosing familiar/unfamiliar inputs as rewards/punishment³ or by exploiting stimulation avoidance³⁴. 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.

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 SI, Sections 4&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 (SI, Fig. 3A). 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 control³⁵, 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.

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). This fundamental limitation is clearly present even in the case of very simple tasks trainable on small networks (SI, Section 6 and SI, 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 ³⁶. We observe, indeed, that Task 1 is easier to learn when the size N_{proc} of the processing area (Fig. 4B) increases, see SI, Section 4.1 and SI, Fig. 3. As a result, functional targets are easier to reach than structural ones, as the latter cannot benefit from

any multiplicity of solutions. The nature of the cost U has also 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. 6D, may be reached for vanishing U only (SI, Section 5.3 and SI, Fig. 6A&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. (19), rather than from the average over all pairs as in Eq. (4), gives better performance for Task 1.

Our work relies on a number of assumptions, whose technological and biological plausibility needs to be discussed. First, we have considered that stimulations could be done at the level of individual neurons. In practice, the resolution of optogenetic^{4,37} and electrophysiological³⁸ 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^{4,39,40}), or discrete (on/off) rather than continuous values.

An important hypothesis is the separation between a fast time scale τ_r associated to the relaxation of neural activity, and the duration of the control stimulation during a cycle, Δt . Under this assumption, the neural activity is stationary, and the estimate of the connectivity can be updated by probing the network response to few random stimuli without inducing synaptic changes (Methods, Eq. (11)). Physiological time scales 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. (12)). We have checked that τ_r remains comparable to τ_n throughout the learning process for all the tasks considered here (SI, Section 2). While it is difficult to estimate with precision the time scale τ_s associated to 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 (SI, Section 1). 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 few hours. We are well aware that these considerations are speculative and indicative of orders of magnitude at best.

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 & 2 are a few tens of r_{max} , see Figs. 4C and 6D; details about how to convert input currents into rates can be found in SI, Section 1.2. Ultra-fast optogenetics can elicit spikes in genetically-modified neurons with high probability (for sufficient light pulse intensity) up to hundreds of Hz⁴¹. Hence, ratios f_{max}/r_{max} close to unity are experimentally accessible. It remains to be seen whether these techniques can be applied to organoids with similar results.

From a modelling 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, both in terms of plasticity⁴²⁻⁴⁴ and of activity⁴⁵⁻⁴⁷, to the point that artificial spiking networks may be trained by proxy rate-based models⁴⁸. Plasticity at the spike level is a complex process involving multiple timescales, delayed effects and spike train patterns⁴⁹, possible co-existing with non-synaptic plasticity⁵⁰, different effects depending on the origin of the neuron⁵¹ and meta-plasticity⁴⁴. The coarse-grained plasticity rules we consider respect some general principles: associativity (covariance rule with a threshold, compatible with experiments⁵²), locality (dependence on the post and pre-synaptic neurons⁵¹), and homeostasis¹⁵. 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 SI, Section 5.3 for reverse choices of the rules for inhibitory synapses. In this context, Hebbian and anti-Hebbian learnings induce, respectively, the weakening and the strengthening of connections between inhibitory neurons with correlated activity, and result in increased or diminished¹⁴ activity. It would be interesting to exhaustively study the performances of our stimulation protocol over a wider class of plasticity rules.

Last of all, a key point to be addressed in future works is noise and uncertainty. Noise could be taken into account, both in the dynamics of the neural activity and of synaptic interactions, 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. Model uncertainty is even more delicate: in practice, neither the equations describing the neural activity dynamics, nor the plasticity rules are exactly known. Any modelling error will be a source of correlated errors in the determination of the stimulation protocol. Controlling those errors and designing robust stimulation schemes will be crucial for future applications.

Methods

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 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 SI) and is chosen to be

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

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 (SI, Section 1).

Relaxation dynamics of the network activity

While single-neuron activities are associated to 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

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

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

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, $\varepsilon_i = E/I$, are initially unknown. To determine their values we start to probe the network activities r_i^v in response to n_{probe} randomly chosen stimulation patterns f_i^v , and rewrite the stationary relations in Eq. (6) as

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

for all i, v . 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 $n_{probe} \geq N$, 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^v are chosen to be large enough to avoid the presence of low firing rates ($r \simeq 0$), see SI, Section 1.

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

$$\varepsilon_i = \begin{cases} E & \text{if } \sum_j J_{ji} > 0 \\ I & \text{otherwise} \end{cases} \quad \text{and} \quad C_{ij} = \begin{cases} 1 & \text{if } |J_{ij}| > J_{min} \\ 0 & \text{otherwise} \end{cases}, \quad (10)$$

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 of 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 connections is \mathbf{J}'_i . To estimate this vector we probe the network with n_{probes} stimuli, and record the corresponding activities. We call \mathbf{R} the $(n_{probes} \times d_i)$ -dimensional matrix of activities r_j^v and \mathbf{t}_i the n_{probes} -dimensional vector of components t_i^v 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 obtain⁵³

$$\mathbf{J}'_i = \Pi_i^\perp \cdot \mathbf{J}_i + \mathbf{R}^+ \cdot \mathbf{t}_i, \quad (11)$$

where \mathbf{R}^+ denotes the pseudo-inverse of \mathbf{R} and $\Pi_i^\perp = \mathbf{Id} - \mathbf{R}^+ \cdot \mathbf{R}$ 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 .

Regularization of the connectivity matrix

A common practice in machine learning is to introduce a regularization cost over the model parameters, hereafter denoted as $U_{reg}(\mathbf{J})$, 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 SI, Section 2.2. For each singular value λ_k of \mathbf{J} , the regularization cost is approximately zero if $\lambda_k \ll 1$, and grows linearly as $g_2 \lambda_k$ for $\lambda_k \gg 1$. We implement this dependence through a softplus function:

$$U_{reg}(\mathbf{J}) = \frac{g_2}{g_1} \sum_{k=1}^{N^2} \log \left(1 + e^{g_1 (\lambda_k - 1)} \right) \quad \text{with} \quad g_2 = 2, \quad g_1 = 10. \quad (12)$$

This regularization cost can be explicitly differentiated with respect to the entries of the connectivity matrix, see below.

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 chain rule,

$$\frac{dU_{task}}{dJ_{ij}} = \frac{\partial U_{task}}{\partial J_{ij}} + \sum_k \frac{\partial U_{task}}{\partial r_k} \frac{\partial r_k}{\partial J_{ij}} \quad (13)$$

which requires knowledge of the derivatives of the neuron firing rates. These can be computed using the implicit function theorem⁵⁴ applied to the stationary equations for the dynamics, with the result:

$$\frac{\partial r_k}{\partial J_{ij}} = [(\mathbf{Id} - \mathbf{D} \cdot \mathbf{J})^{-1}]_{k,i} \mathbf{D}_{ii} r_j, \quad (14)$$

where \mathbf{D} was first used in Eq. (8). The gradient of the regularization cost, U_{reg} in Eq. (12), 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 to λ_k .

In addition, the gradient component associated to 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 pre-synaptic neuron type, *i.e.* $\frac{dU}{dJ_{ij}} > 0$ if $\varepsilon_j = E$ or $\frac{dU}{dJ_{ij}} < 0$ if $\varepsilon_j = I$. 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

$$\dot{U} = \sum_{i,j} \frac{\partial U}{\partial J_{ij}} \dot{J}_{ij}, \quad (15)$$

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 sum of the task and regularization costs, $U = U_{task} + U_{reg}$. To determine the best control stimulation \mathbf{f} over the next period we search for the minimum of the modified change

$$\Delta U(\mathbf{f}; \gamma) = U(\mathbf{J} + \Delta \mathbf{J}(\mathbf{f})) - U(\mathbf{J}) + \Delta t \gamma \|\Delta \mathbf{J}(\mathbf{f})\|^2 \quad (16)$$

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

$$\Delta \mathbf{J}(\mathbf{f}) = [\mathbf{J} + \Delta t \dot{\mathbf{J}}(\mathbf{f})]_{\varepsilon} - \mathbf{J} \quad (17)$$

where $[\cdot]_{\varepsilon}$ indicates that connections have been clipped to respect their excitatory/inhibitory types ε imposed by the pre-synaptic neurons. The last term in Eq. (16) constrains the change in the interactions to be small over the period of stimulation, effectively introducing a saturating non-linearity over the gradients.

Computing the gradient of ΔU in Eq. (16) 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. (13),(14), and with respect to time, see Eq. (15). In addition, the derivatives of the firing activities are needed. Using again the implicit function theorem we obtain, see Eq. (14),

$$\frac{\partial r_k}{\partial f_i} = [(\mathbf{Id} - \mathbf{D} \cdot \mathbf{J})^{-1}]_{k,i} \mathbf{D}_{ii}. \quad (18)$$

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 a iterative descent of ΔU in Eq. (16). 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 hyper-parameters involved can be found in SI, section 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 $\varepsilon_j = E$) and $1 - c_I$ and c_I for inhibitory ones ($\varepsilon_j = I$). 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 $\varepsilon_j > 0$ or $[-J_0, 0]$ if $\varepsilon_j < 0$. See SI, Section 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 $\boldsymbol{\sigma}^\mu$ (over the N_o output neurons), with $\mu = 1, \dots, n_{task}$. 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 to 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

$$U_{task}(\mathbf{J}) = \frac{\sum_{(\mu,i):\sigma_i^\mu=L} \sum_{(v,j):\sigma_j^v=H} \Delta(\mu, i; v, j) e^{\gamma \Delta(\mu, i; v, j)}}{\sum_{(\mu,i):\sigma_i^\mu=L} \sum_{(v,j):\sigma_j^v=H} e^{\gamma \Delta(\mu, i; v, j)}} \quad (19)$$

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

$$\Delta(\mu, i; v, j) = h_2(r_i^\mu - r_j^v + \delta r) \quad \text{with} \quad h_2(u) = u^2 \text{ if } u > 0, 0 \text{ otherwise.} \quad (20)$$

The exponential factors give more weights 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. (20) 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 $\Delta(\mu, i; v, j) = 0$.

Linear response analysis. To characterize how the association mechanism emerges across training, we consider the input-to-output response matrix $\mathbf{R} = \frac{d\mathbf{r}}{d\mathbf{f}}$ linearized around the average input $\bar{\mathbf{f}} = \sum_{\mu} \mathbf{f}^\mu / n_{task}$. We define the linear sub-spaces spanned by the inputs and outputs as, respectively, $In = \text{span}(\{\mathbf{f}^\mu - \bar{\mathbf{f}}\}_{\mu=1, \dots, n_{task}})$ and $Out = \text{span}(\{\boldsymbol{\sigma}^\mu - \bar{\boldsymbol{\sigma}}\}_{\mu=1, \dots, n_{task}})$, where we have converted $\sigma_i^\mu = \text{High/Low}$ to 1/0 and $\bar{\boldsymbol{\sigma}} = \sum_{\mu} \boldsymbol{\sigma}^\mu / n_{task}$ is the average output. We then partition \mathbf{R} in four blocks, through projections on In and its orthogonal complement In^\perp on one side, and Out and Out^\perp on the other side. The normalized (divided by $\sqrt{\dim(Out)}$ or $\sqrt{\dim(Out^\perp)}$) largest singular value of these four sub-matrices can be tracked across learning as a metric of the relative importance of each block.

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 SI, Section 1 for parameter values.

Task and cost. The structural cost is defined in Eq. (3). We now specify the target connectivity. Neurons are associated angles on two rings, depending on their types $\varepsilon = E$ 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

$$J_{ij}^{target} = J_{\varepsilon_i, \varepsilon_j} \times \mathcal{M}(\theta_i^{\varepsilon_i} - \theta_j^{\varepsilon_j} - \pi \delta_{\varepsilon_i, I}; K_{\varepsilon_i, \varepsilon_j}), \quad (21)$$

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

$$\mathcal{M}(\Delta\theta; K) = e^{-K + K \cos(\Delta\theta)}, \quad (22)$$

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

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

References

1. Lancaster, M. *et al.* Cerebral organoids model human brain development and microcephaly. *Nature* **501**, 373–9 (2013).
2. Smirnova, L. *et al.* Organoid intelligence (oi): the new frontier in biocomputing and intelligence-in-a-dish. *Front. Sci.* **1** (2023).
3. Kagan, B. J. *et al.* In vitro neurons learn and exhibit sentience when embodied in a simulated game-world. *Neuron* **110**, 3952–3969 (2022).
4. Deisseroth, K. Optogenetics: 10 years of microbial opsins in neuroscience. *Nat. Neurosci.* **18**, 1213–1225 (2015).
5. Yang, W. & Yuste, R. Holographic imaging and photostimulation of neural activity. *Curr. Opin. Neurobiol.* **50**, 211–221 (2018).
6. Shin, H., Jeong, S., Lee, J. & al. 3d high-density microelectrode array with optical stimulation and drug delivery for investigating neural circuit dynamics. *Nat. Commun.* **12**, 492 (2021).
7. Qian, X., Song, H. & Ming, G.-l. Brain organoids: advances, applications and challenges. *Development* **146**, dev166074 (2019).
8. Tanaka, G. *et al.* Recent advances in physical reservoir computing: A review. *Neural Networks* **115**, 100–123 (2019).
9. Sumi, T. *et al.* Biological neurons act as generalization filters in reservoir computing. *Proc. Natl. Acad. Sci.* **120**, e2217008120 (2023).
10. Citri, A. & Malenka, R. Synaptic plasticity: Multiple forms, functions, and mechanisms. *Neuropsychopharmacol.* **33**, 18–41 (2007).
11. Carrillo-Reid, L., Yang, W., Kang Miller, J.-e., Peterka, D. S. & Yuste, R. Imaging and optically manipulating neuronal ensembles. *Annu. Rev. Biophys.* **46**, 271–293 (2017).
12. Vico, F. J. & Jerez, J. Stable neural attractors formation: Learning rules and network dynamics. *Neural Process. Lett.* **18**, 1–16 (2003).
13. Turrigiano, G. G. The self-tuning neuron: synaptic scaling of excitatory synapses. *Cell* **135**, 422–435 (2008).
14. Turrigiano, G. G. & Nelson, S. B. Hebb and homeostasis in neuronal plasticity. *Curr. Opin. Neurobiol.* **10**, 358–364 (2000).
15. Zenke, F., Gerstner, W. & Ganguli, S. The temporal paradox of hebbian learning and homeostatic plasticity. *Curr. opinion neurobiology* **43**, 166–176 (2017).
16. Zierenberg, J., Wilting, J. & Priesemann, V. Homeostatic plasticity and external input shape neural network dynamics. *Phys. Rev. X* **8**, 031018 (2018).
17. Földiák, P. Forming sparse representations by local anti-hebbian learning. *Biol. Cybern.* **64** (1990).
18. de Abril, I. M., Yoshimoto, J. & Doya, K. Connectivity inference from neural recording data: Challenges, mathematical bases and research directions. *Neural Networks* **102**, 120–137 (2018).

19. Minsky, M. & Papert, S. *Perceptrons: An Introduction to Computational Geometry* (MIT Press, Cambridge, MA, USA, 1969).
20. Devaud, J. *et al.* Neural substrate for higher-order learning in an insect: Mushroom bodies are necessary for configural discriminations. *Proc. Natl. Acad. Sci. (USA)* **112**, E5854–E5862 (2015).
21. Amari, S. Dynamics of pattern formation in lateral-inhibition type neural fields. *Biol. Cybern.* **27**, 77–87 (1977).
22. Ben-Yishai, R., Bar-Or, R. & Sompolinsky, H. Theory of orientation tuning in visual cortex. *Proc Natl Acad Sci USA* **25**, 3844–8 (1995).
23. Seelig, J. D. & Jayaraman, V. Neural dynamics for landmark orientation and angular path integration. *Nature* **521**, 186–191 (2015).
24. Kuehn, S. & Monasson, R. Information content in continuous attractor neural networks is preserved in the presence of moderate disordered background connectivity. *Phys. Rev. E* (2023).
25. Molari, M., Eyer, K., Baudry, J., Cocco, S. & Monasson, R. Quantitative modeling of the effect of antigen dosage on b-cell affinity distributions in maturing germinal centers. *eLife* **9**, e55678 (2020).
26. Sprenger, K., Louveau, J., Murugan, P. & Chakraborty, A. Optimizing immunization protocols to elicit broadly neutralizing antibodies. *Proc. Natl. Acad. Sci.* **117**, 20077–87 (2020).
27. Lukoševičius, M. & Jaeger, H. Reservoir computing approaches to recurrent neural network training. *Comput. science review* **3**, 127–149 (2009).
28. George, J. B. *et al.* Input coding for neuro-electronic hybrid systems. *Biosystems* **126**, 1–11 (2014).
29. Yada, Y., Yasuda, S. & Takahashi, H. Physical reservoir computing with force learning in a living neuronal culture. *Appl. Phys. Lett.* **119** (2021).
30. Rigotti, M., Rubin, D. B. D., Wang, X.-J. & Fusi, S. Internal representation of task rules by recurrent dynamics: the importance of the diversity of neural responses. *Front. computational neuroscience* **4**, 24 (2010).
31. Enel, P., Procyk, E., Quilodran, R. & Dominey, P. F. Reservoir computing properties of neural dynamics in prefrontal cortex. *PLoS computational biology* **12**, e1004967 (2016).
32. Morales, G. B., Mirasso, C. R. & Soriano, M. C. Unveiling the role of plasticity rules in reservoir computing. *Neurocomputing* **461**, 705–715 (2021).
33. Cai, H. *et al.* Brain organoid computing for artificial intelligence. *bioRxiv* 2023–02 (2023).
34. Masumori, A. *et al.* Emergence of sense-making behavior by the stimulus avoidance principle: Experiments on a robot behavior controlled by cultured neuronal cells. In *Artificial Life Conference Proceedings*, 373–380 (MIT Press, Cambridge, 2015).
35. Kirk, D. E. *Optimal control theory: an introduction* (Courier Corporation, 2004).
36. Marder, E. & Taylor, A. Multiple models to capture the variability in biological neurons and networks. *Nat Neurosci* **14**, 133–138 (2011).
37. Stamatakis, A. M. *et al.* Simultaneous optogenetics and cellular resolution calcium imaging during active behavior using a miniaturized microscope. *Front. neuroscience* 496 (2018).
38. Spira, M. E. & Hai, A. Multi-electrode array technologies for neuroscience and cardiology. *Nat. nanotechnology* **8**, 83–94 (2013).
39. Land, B. B., Brayton, C. E., Furman, K. E., LaPalombara, Z. & DiLeone, R. J. Optogenetic inhibition of neurons by internal light production. *Front. behavioral neuroscience* **8**, 108 (2014).
40. Han, X. & Boyden, E. S. Multiple-color optical activation, silencing, and desynchronization of neural activity, with single-spike temporal resolution. *PloS one* **2**, e299 (2007).
41. Mager, T., Lopez de la Morena, D., Senn, V. & al. High frequency neural spiking and auditory signaling by ultrafast red-shifted optogenetics. *Nat. Commun.* **9**, 1750 (2018).
42. Izhikevich, E. M. & Desai, N. S. Relating stdp to bcm. *Neural computation* **15**, 1511–1523 (2003).
43. Bush, D., Philippides, A., Husbands, P. & O’Shea, M. Reconciling the stdp and bcm models of synaptic plasticity in a spiking recurrent neural network. *Neural computation* **22**, 2059–2085 (2010).
44. Yger, P. & Gilson, M. Models of metaplasticity: a review of concepts. *Front. computational neuroscience* **9**, 138 (2015).

45. Ostojic, S. & Brunel, N. From spiking neuron models to linear-nonlinear models. *PLoS computational biology* **7**, e1001056 (2011).
46. Rueckauer, B. & Liu, S.-C. Conversion of analog to spiking neural networks using sparse temporal coding. In *2018 IEEE international symposium on circuits and systems (ISCAS)*, 1–5 (IEEE, 2018).
47. Sanzeni, A., Histed, M. H. & Brunel, N. Response nonlinearities in networks of spiking neurons. *PLoS computational biology* **16**, e1008165 (2020).
48. Kheradpisheh, S. R., Mirsadeghi, M. & Masquelier, T. Spiking neural networks trained via proxy. *IEEE Access* **10**, 70769–70778 (2022).
49. Froemke, R. C. & Dan, Y. Spike-timing-dependent synaptic modification induced by natural spike trains. *Nature* **416**, 433–438 (2002).
50. Mozzachiodi, R. & Byrne, J. H. More than synaptic plasticity: role of nonsynaptic plasticity in learning and memory. *Trends neurosciences* **33**, 17–26 (2010).
51. Caporale, N. & Dan, Y. Spike timing-dependent plasticity: a hebbian learning rule. *Annu. Rev. Neurosci.* **31**, 25–46 (2008).
52. Lim, S. *et al.* Inferring learning rules from distributions of firing rates in cortical neurons. *Nat. neuroscience* **18**, 1804–1810 (2015).
53. Montgomery, D. C., Peck, E. A. & Vining, G. G. *Introduction to linear regression analysis* (John Wiley & Sons, 2021).
54. Rudin, W. *Principles of mathematical analysis* (McGraw-Hill, 1953).

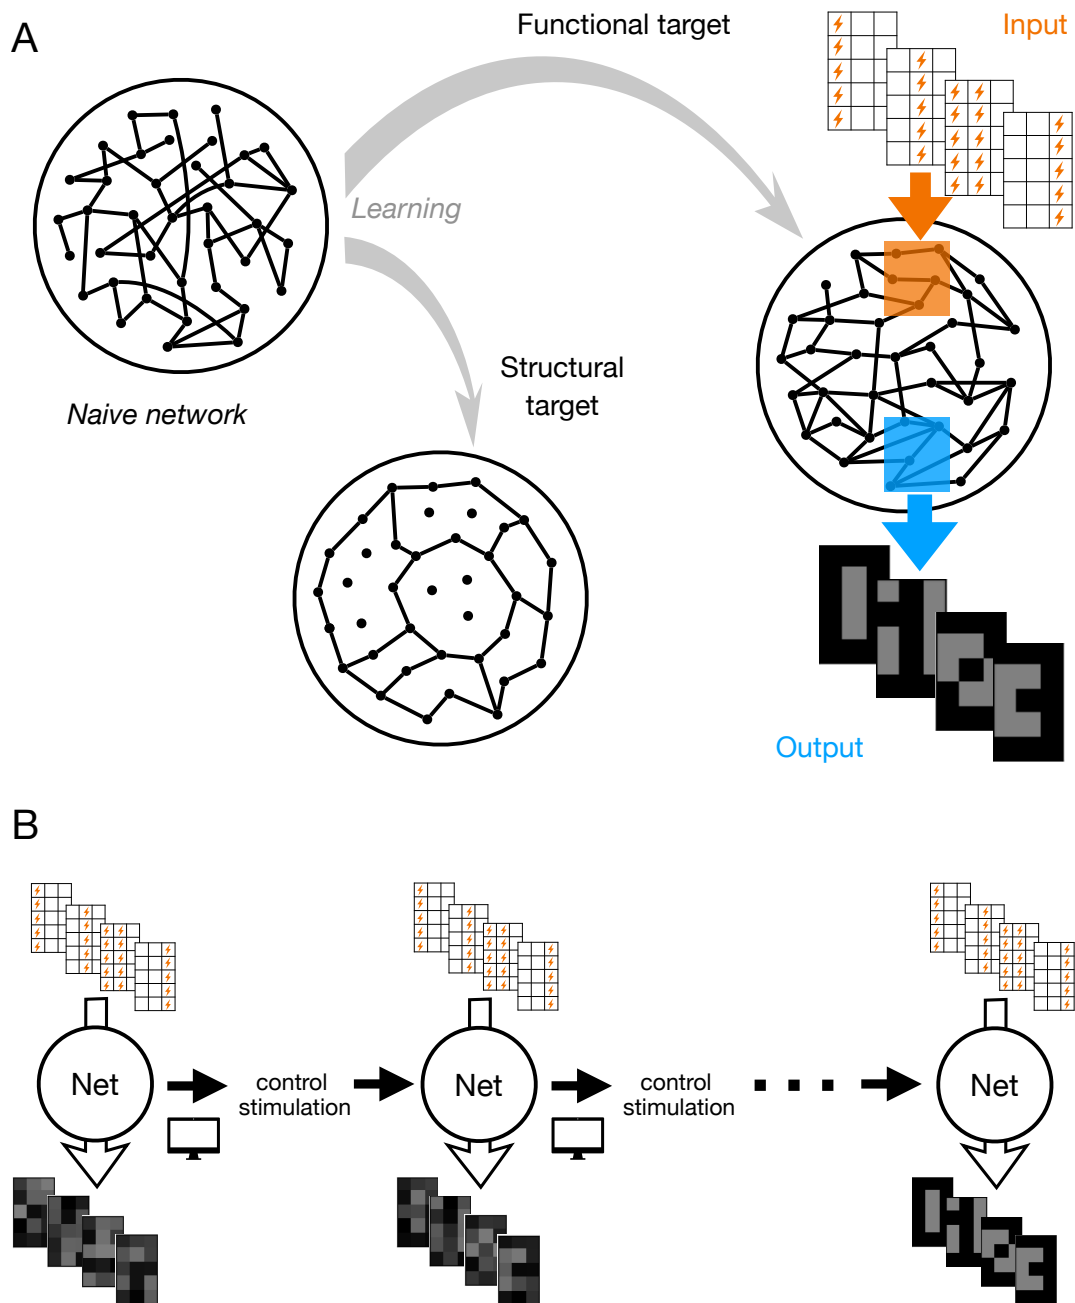


Figure 1

Figure 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 spatio-temporal stimulations of the neurons exploits 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 to 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 re-optimized and applied during a new stimulation period. After multiple iterations, the correct input-to-output association is achieved (right).

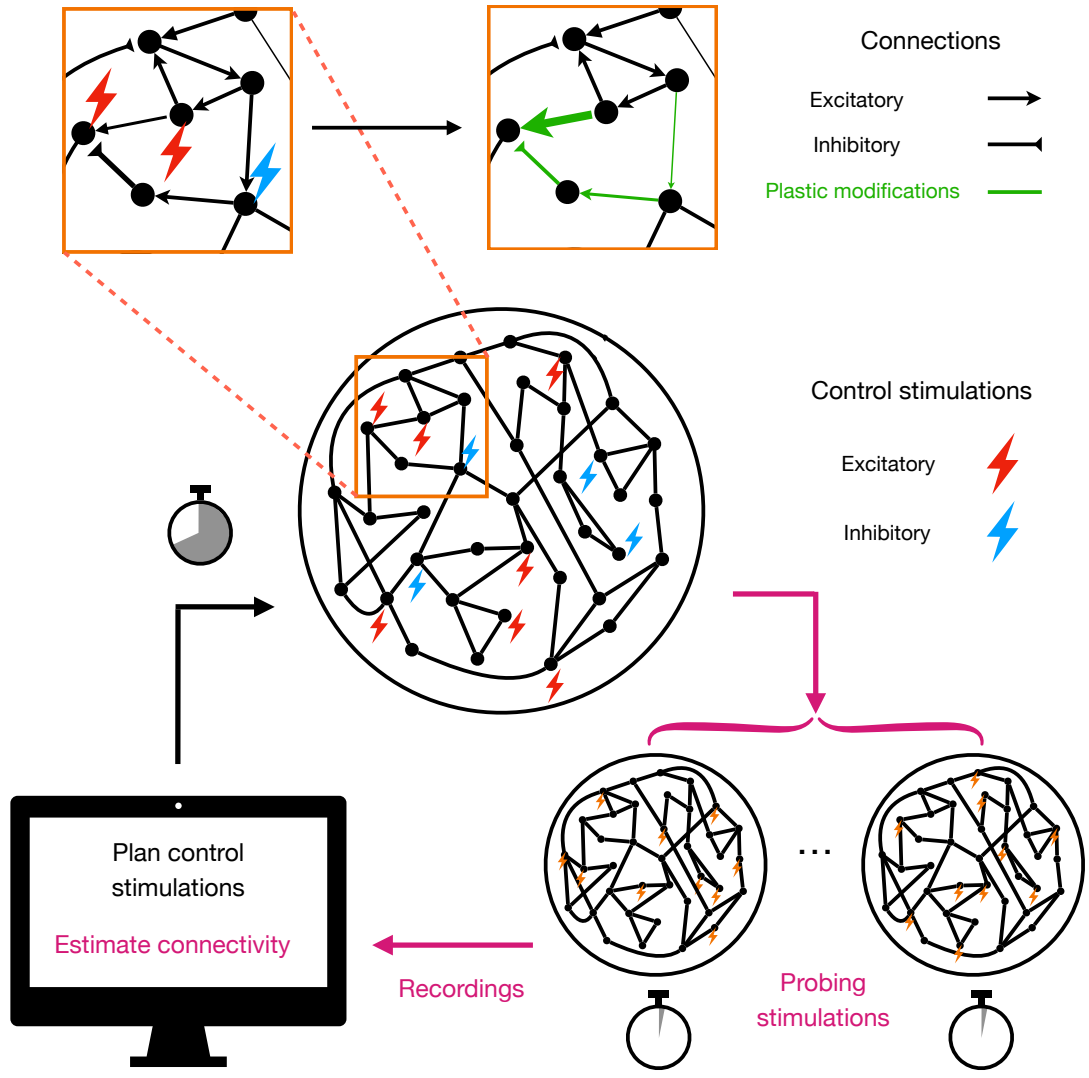


Figure 2

Figure 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).

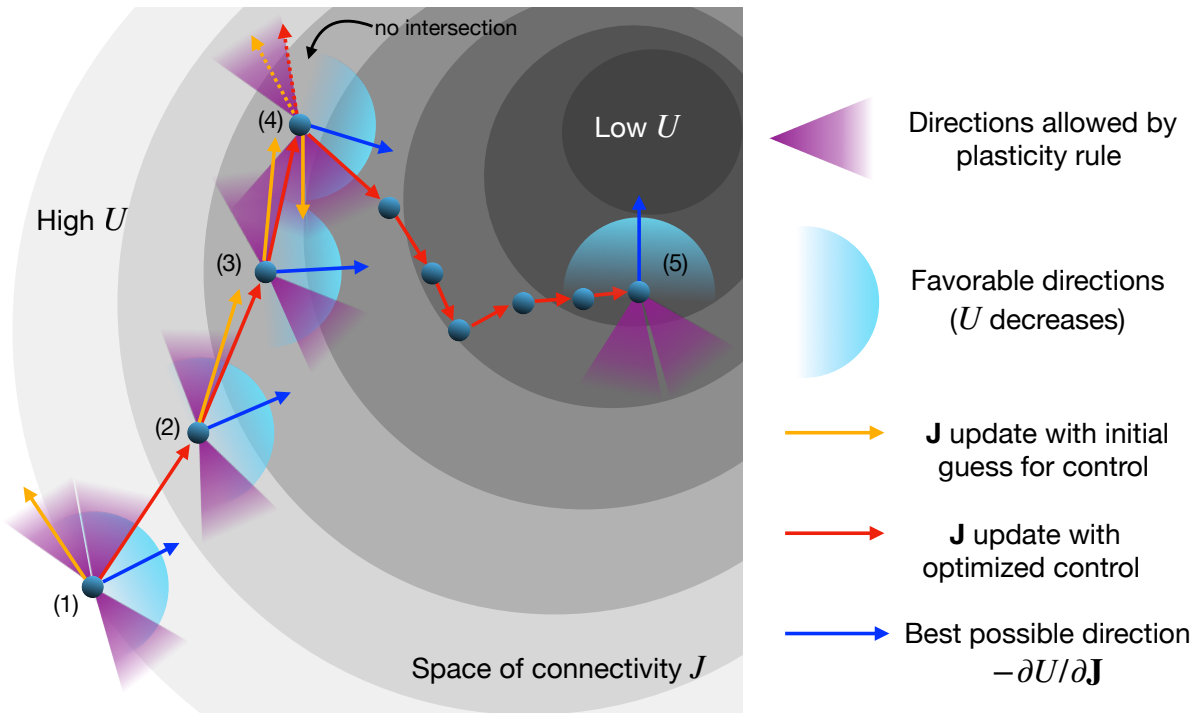


Figure 3

Figure 3: Schematic dynamics of the network and feasible directions in the high-dimensional connectivity space during the training process.

Grey 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 to 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 resumes the search for the optimal \mathbf{f}^* as in cycle (1). This process is iterated, until the change 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.

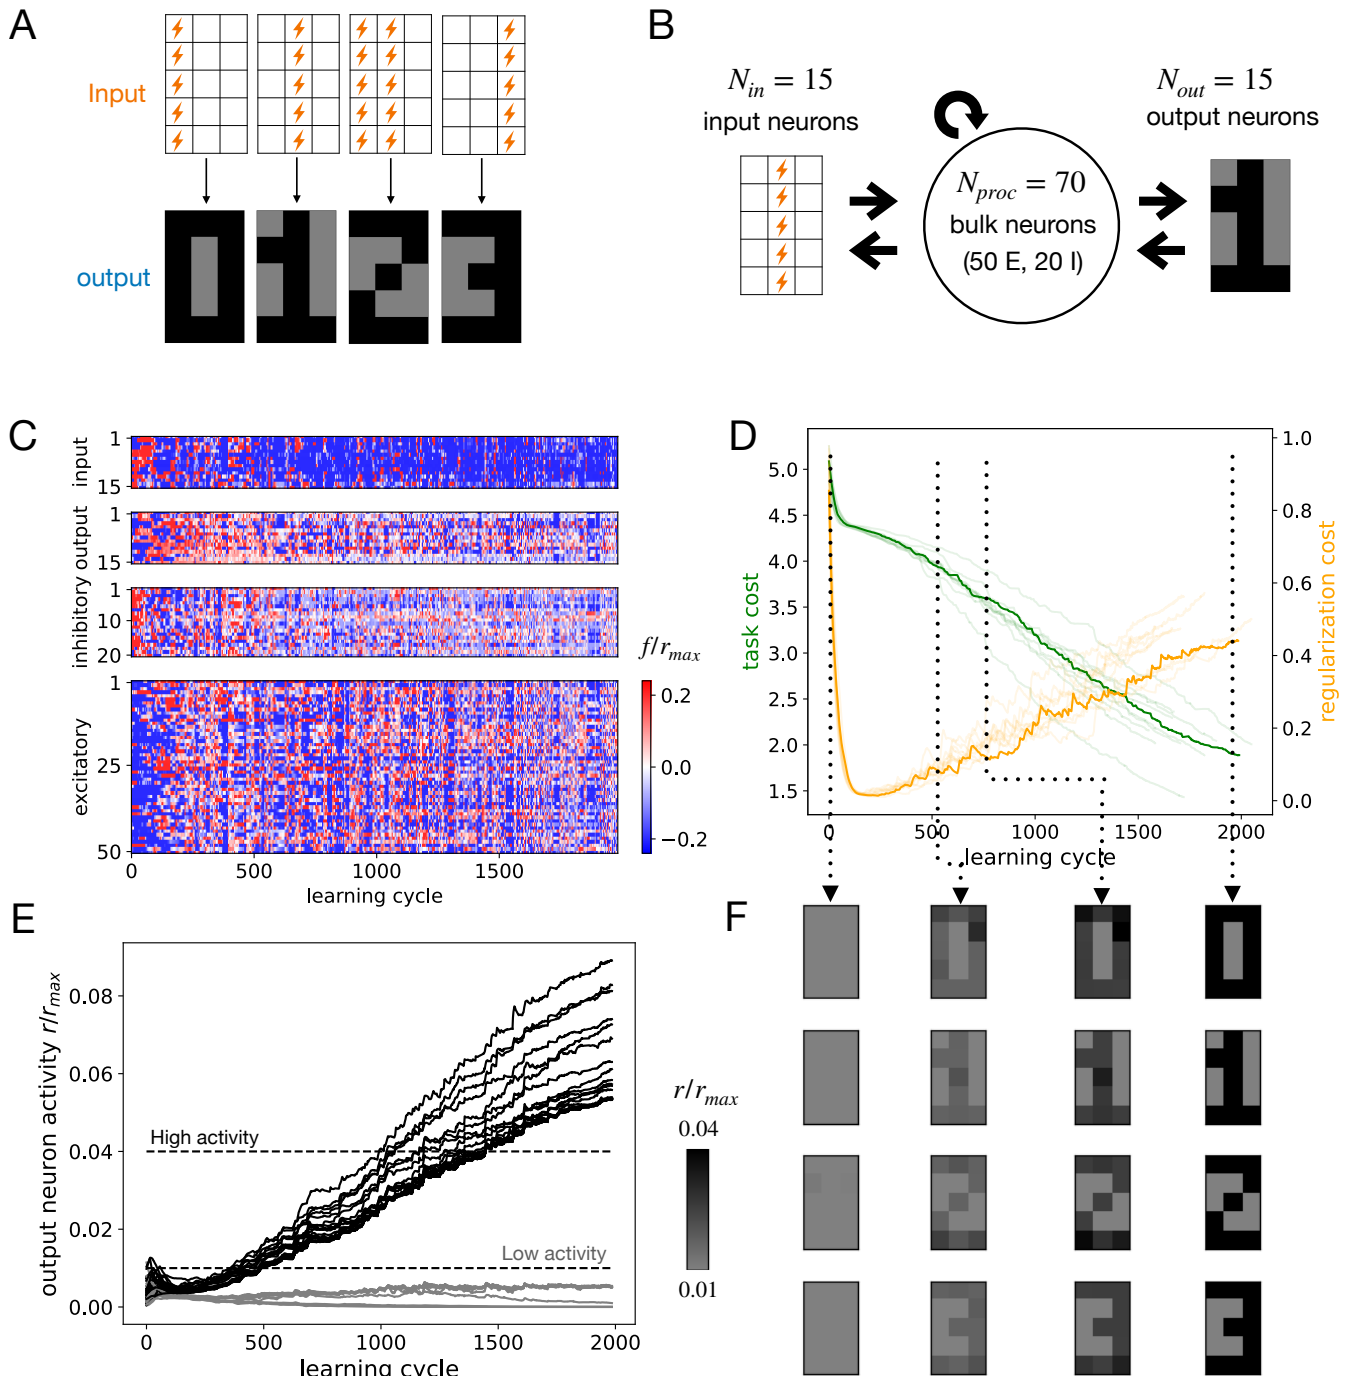


Figure 4

Figure 4: Task 1: non-additive 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 to the task (green, left scale) and to 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 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.

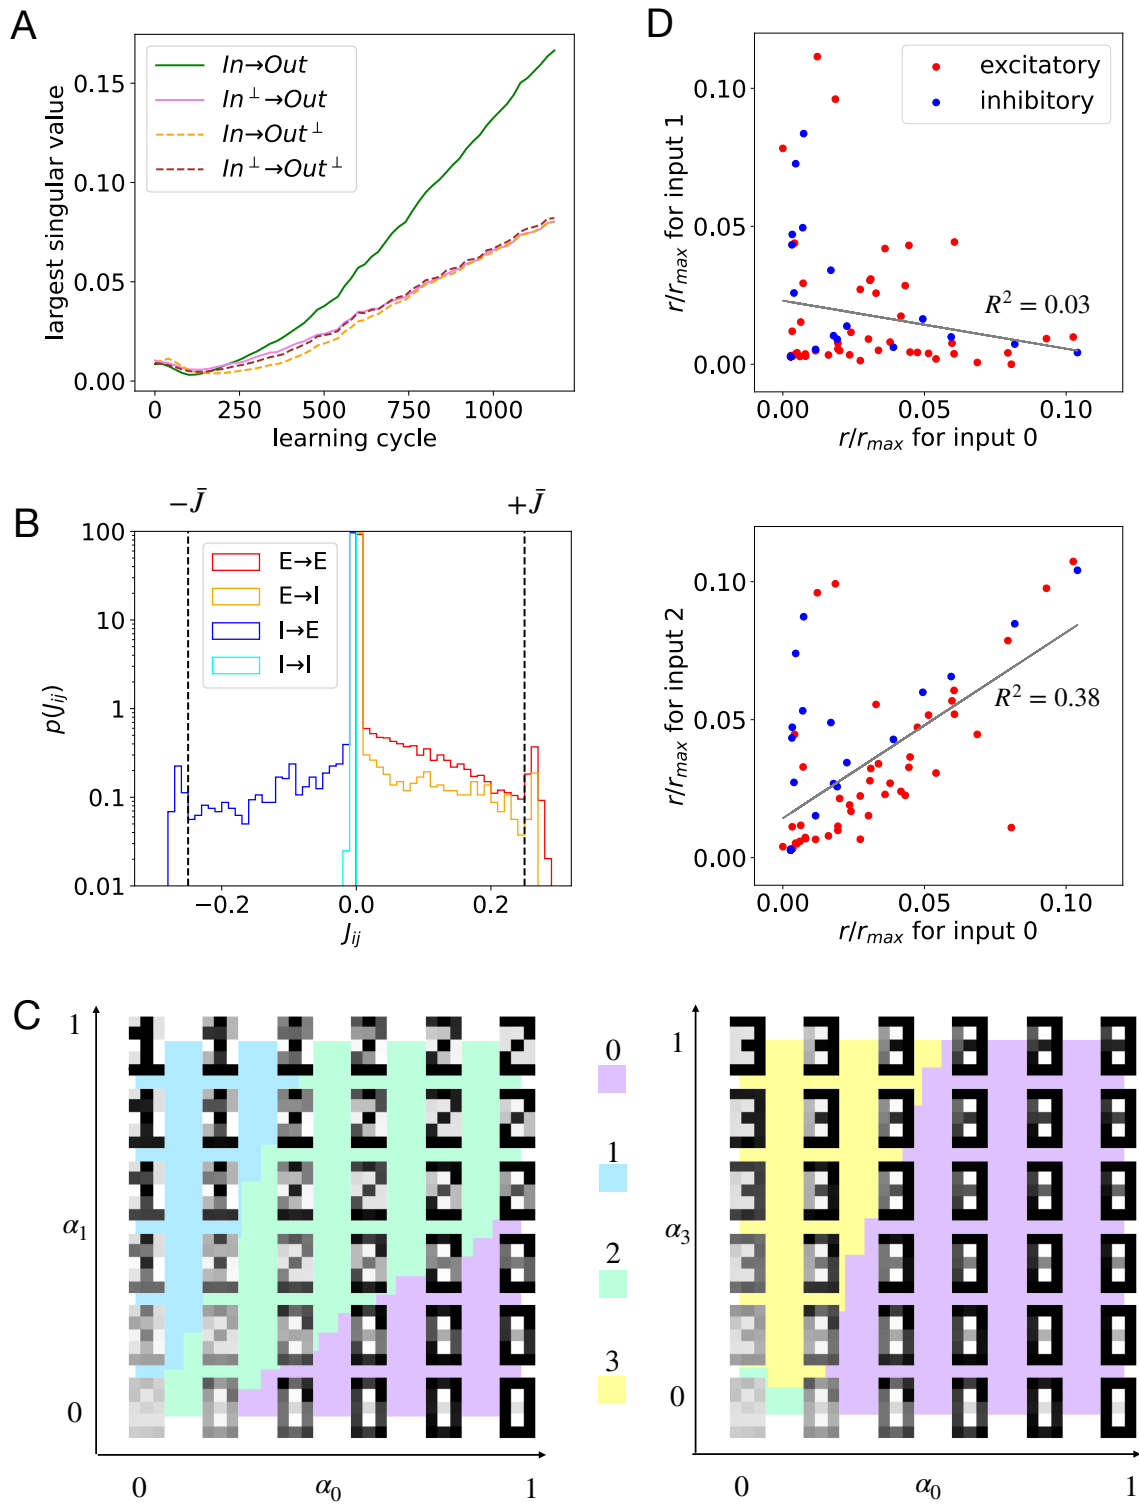


Figure 5

Figure 5: Learning dynamics and internal representations for Task 1.

A. Top singular values σ of the four blocks of the input-to-output matrix \mathbf{R} , see Methods. To allow for comparison the σ 's are normalized by the square root of the block dimensions (according to random matrix theory). At the beginning of training the four responses have similar amplitudes. All responses then decay due to regularization (Fig. 4C and Eq. (12) in Methods) until time step 400. As training proceeds, the response through the input-to-output channel (*In* and *Out*) becomes dominant.

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 to, 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 to patterns 0 and 1 are not correlated, as expected from the orthogonality of \mathbf{f}_0 and \mathbf{f}_1 (Fig. 4A). Conversely, the representations of 0 and 2 are strongly correlated, reflecting the relation $\mathbf{f}_2 = \mathbf{f}_0 + \mathbf{f}_1$ (Fig. 4A).

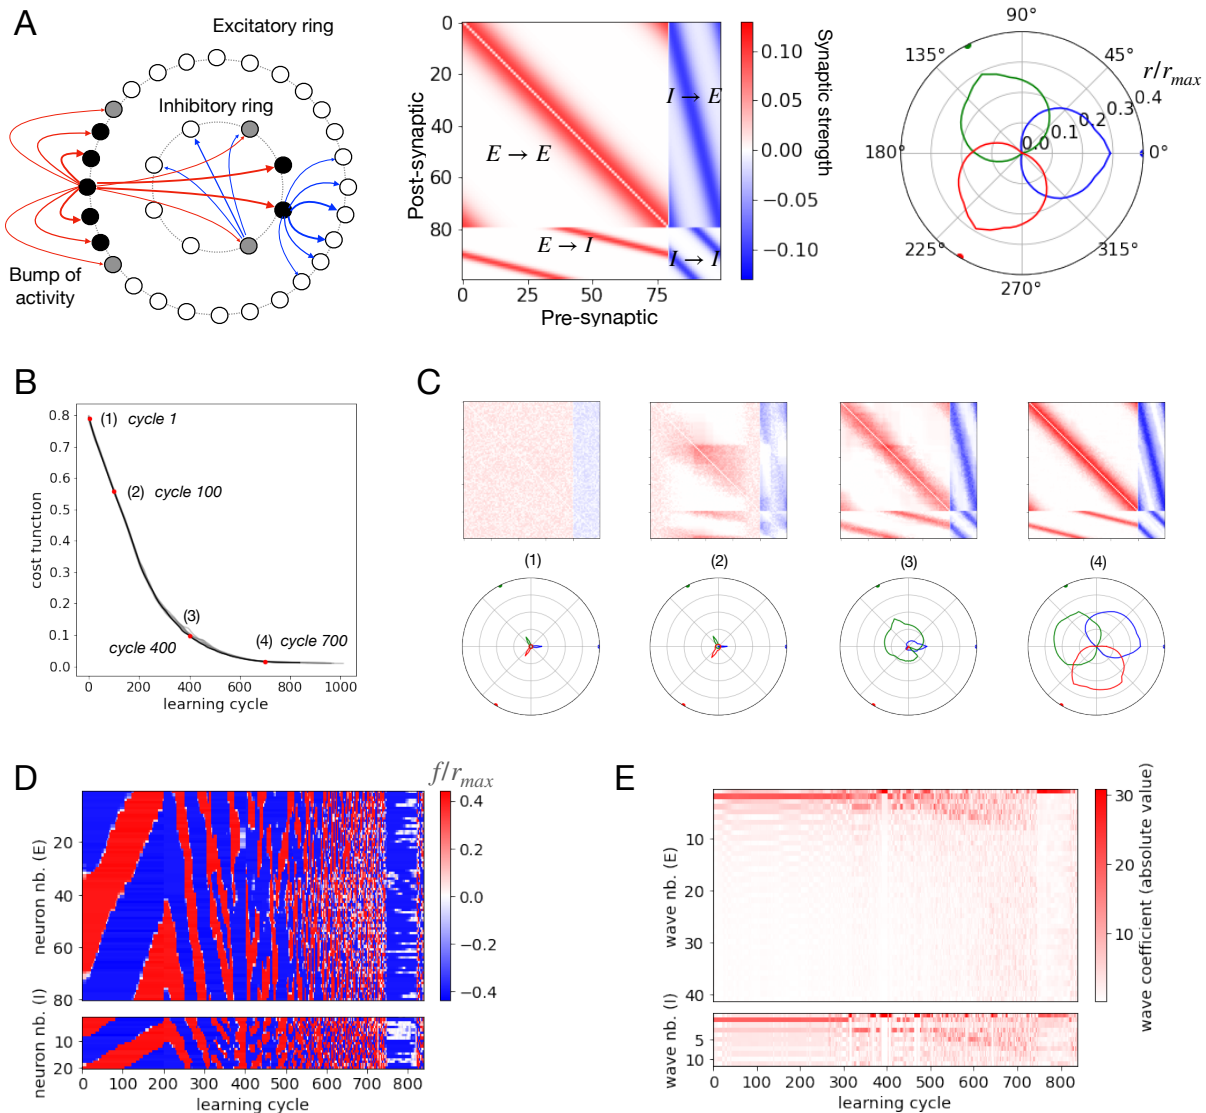


Figure 6

Figure 6: Task 2: building a continuous attractor.

A. Target network (left) and connectivity matrix \mathbf{J}^{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 neighbours, and project to inhibitory neurons diametrically opposed on the I ring. Inhibitory neurons repress neighbouring 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 to 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, labelled (1), (2), (3), (4) are referred to in panel C.

C. Connectivity matrices (top) and receptive fields (bottom) at cycles (1), (2), (3), (4), showing different stages of learning. Same color codes as in panel A.

D. Control stimulations $f_i(t)$ as functions of the learning cycle.

E. Amplitude $\hat{f}_k^\epsilon(t) = \sum_{\ell=1}^{N_E} f_\ell(t) \cos(2\pi k \ell / N_E)$ of the Fourier modes $k = 0, 1, \dots, N_E/2$ associated to the control stimulations shown in panel D, for excitatory ($\epsilon = E$) and inhibitory ($\epsilon = I$) neurons. Initially, the control stimulation stimulation mainly consist 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.

Acknowledgements

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 program under Grant No. 964977 (project NEU-ChiP).

Author contributions statement

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