

RESEARCH ARTICLE

Balancing complexity, performance and plausibility to meta learn plasticity rules in recurrent spiking networks

Basile Confavreux ^{1,2*}, Everton J. Agnes ³, Friedemann Zenke ⁴, Henning Sprekeler ⁵, Tim P. Vogels ¹

1 Institute of Science and Technology Austria, Klosterneuburg, Austria, **2** Gatsby Computational Neuroscience Unit, University College London, London, United Kingdom, **3** Biozentrum, University of Basel, Basel, Switzerland, **4** Friedrich Miescher Institute for Biomedical Research, Basel, Switzerland, **5** Technische Universität Berlin, Berlin, Germany

* basile.confavreux@gmail.com



OPEN ACCESS

Citation: Confavreux B, Agnes EJ, Zenke F, Sprekeler H, Vogels TP (2025) Balancing complexity, performance and plausibility to meta learn plasticity rules in recurrent spiking networks. *PLoS Comput Biol* 21(4): e1012910. <https://doi.org/10.1371/journal.pcbi.1012910>

Editor: Jonathan David Touboul, Brandeis University, UNITED STATES OF AMERICA

Received: July 08, 2024

Accepted: February 25, 2025

Published: April 24, 2025

Peer Review History: PLOS recognizes the benefits of transparency in the peer review process; therefore, we enable the publication of all of the content of peer review and author responses alongside final, published articles. The editorial history of this article is available here: <https://doi.org/10.1371/journal.pcbi.1012910>

Copyright: © 2025 Confavreux et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data availability statement: The full data and code is publicly available at <https://github.com/VogelsLab/SpikES>.

Abstract

Synaptic plasticity is a key player in the brain's life-long learning abilities. However, due to experimental limitations, the mechanistic link between synaptic plasticity rules and the network-level computations they enable remain opaque. Here we use evolutionary strategies (ES) to meta learn local co-active plasticity rules in large recurrent spiking networks with excitatory (E) and inhibitory (I) neurons, using parameterizations of increasing complexity. We discover rules that robustly stabilize network dynamics for all four synapse types acting in isolation (E-to-E, E-to-I, I-to-E and I-to-I). More complex functions such as familiarity detection can also be included in the search constraints. However, our meta learning strategy begins to fail for co-active rules of increasing complexity, as it is challenging to devise loss functions that effectively constrain network dynamics to plausible solutions *a priori*. Moreover, in line with previous work, we can find multiple degenerate solutions with identical network behaviour. As a local optimization strategy, ES provides one solution at a time and makes exploration of this degeneracy cumbersome. Regardless, we can glean the interdependencies of various plasticity parameters by considering the covariance matrix learned alongside the optimal rule with ES. Our work provides a proof of principle for the success of machine-learning-guided discovery of plasticity rules in large spiking networks, and points at the necessity of more elaborate search strategies going forward.

Author summary

Synapses between neurons in the brain change continuously throughout life. This phenomenon, called synaptic plasticity, is believed to be crucial for the brain to learn from and remember past experiences. However, the exact nature of these synaptic changes remains unclear, partly because they are hard to observe experimentally. Theorists have thus long tried to predict these synaptic changes and how they contribute to learning and

Funding: This project has received funding from the HORIZON EUROPE European Research Council (ERC) consolidator grant (SYNAPSEEK, awarded to TV), a Wellcome Trust Sir Henry Dale Research Fellowship (WT100000, awarded to TV), a Wellcome Trust Senior Research Fellowship (214316/Z/18/Z, awarded to TV), and a Sir Henry Wellcome Fellowship (110124/Z/15/Z, awarded to FZ). The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Competing interests: The authors have declared that no competing interests exist.

memory, using abstraction called plasticity rules. Although many plasticity rules have been proposed, there are many different synapse types in the brain and many more possible rules to test. A recent approach has thus been to automate this screening of possible plasticity rules, using modern Machine Learning tools. This idea, called meta learning plasticity rules, has so far only been applied to very simple models of brain synapses. Here, we scale up this idea to more complex and more faithful models. We optimize plasticity rules based on their ability to make model brain circuits solve some basic memory tasks. We find several different yet equally good plasticity rules (degeneracy). However, our method drops in performance when considering more complex rules or tasks.

Introduction

Synaptic plasticity is thought to be the cornerstone of learning and memory. *In silico*, the evolution of synaptic efficacies is modeled with plasticity rules [1–20] typically derived from *ex vivo* experiments in single synapses [7,21–25]. Even though such rules recapitulate the data gathered at the single neuron level, they often fail to elicit the observed functions or architectures at the network level, in part because of the enormous parameter space that must be trawled to elicit functions such as memory formation in spiking neuronal networks (SNNs) [12,13,16].

Instead of tuning the values of the parameters governing plasticity by hand (hand-tuning), an emerging approach dubbed “meta learning synaptic plasticity” consists in performing numerical optimization on the plasticity rules themselves so that candidate plasticity rules with desired network-level behaviors can be found automatically [26–32]. This approach has been successful in rate networks, both in elucidating the learning rules implemented in brains and proposing alternatives to back-propagation [33–38]. However, in the case of spiking neuronal networks, this meta learning approach has been restricted to two-layer feedforward networks performing simple tasks [31,39]. This dearth is partly owed to the fact that the parameterization of spike-based plasticity rules either involves high-dimensional expressions [6], ill-suited to numerical optimization in spiking networks, or search spaces so simple that they don’t contain truly novel rules [31] when used in isolation. Additionally, the non-differentiability of spiking network models and their large compute requirements contribute to the lack of results in meta-learning plasticity rules at the level of large recurrent spiking networks.

Here, we solve some of the above-mentioned difficulties to meta learn biologically plausible plasticity rules in large recurrent spiking networks with excitatory and inhibitory populations in a two-loop meta learning paradigm. In an inner loop, parameterized plasticity rules are embedded in spiking networks performing a given task, while in an outer loop, a Covariance Matrix Adaptation-Evolution Strategy (CMA-ES) [40] adjusts the parameters of the plasticity rules so that the spiking network in the inner loop performs better on the task at hand (Fig 1, see also [31]).

We compare several rule parameterizations; low-dimensional polynomial rules can provide interesting, easily interpretable solutions for multiple co-active rules; plasticity rules parameterized with neural networks (MLPs) allow us to include a richer set of potentially relevant effectors. We show that we can successfully extract suitable rules for a given function, such as stabilizing network dynamics or performing familiarity detection, as long as we focus on a single connection type (e.g., excitatory-excitatory) at a time. When we turn to more complex search spaces, such as co-active rules, or more elaborate tasks, the flexibility of the system makes it very difficult to craft successful loss functions for biologically plausible solutions that

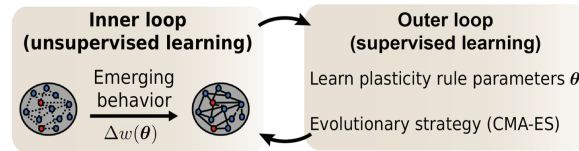


Fig 1. Meta learning approach to discover plasticity rules with a desired network function. Plasticity rules that change synaptic weights depending on some biologically plausible synaptic variables—e.g., pre- and postsynaptic spike times—are parameterized with parameters θ . These parameters are optimized using evolutionary strategies to find a plasticity rule that minimizes a loss function quantifying a desired network behavior.

<https://doi.org/10.1371/journal.pcbi.1012910.g001>

can be learned in finite time. Interestingly, when we find a suitable rule for a given task, we can usually find multiple others, confirming previous results on degenerate solution spaces of plasticity rules [31,41–43].

Results

The rules governing changes of neuronal connections across time remain an open question in neuroscience, despite decades of efforts. An emerging *in silico* method to propose interesting candidate plasticity rules from network-level constraints—meta learning—has been successful in small feedforward systems [31,39], using low-dimensional plasticity parameterizations [31]. However, it is unclear if this method can scale to larger recurrent spiking networks with more complex plasticity rules. Here, we show that we can meta learn plasticity rules with basic memory functions for a range of plasticity parameterizations. In fact, many different plasticity rules can perform a given computation, i.e., the solution space is degenerate.

Meta learning procedure

Here, we turned to a previously devised meta learning pipeline (Fig 1): We used CMA-ES [40] to iteratively improve upon an initial plasticity rule, parameterized with plasticity parameters θ . At every meta iteration, CMA-ES generated a collection of plasticity rules which were evaluated in individual spiking networks according to their fitness (loss function). CMA-ES then updated its internal model of parameter interdependencies and its guess for the best rule, i.e., the mean and covariance matrix of a Gaussian distribution in plasticity parameter space θ (see Methods). To help with the stability of the meta optimization, the initial plasticity rule was chosen such that it elicited no or few weight changes in the network.

Network stability with a small polynomial search space

We began with a simple search space for plasticity rules encompassing first-order (i.e., containing no square-terms, Methods) spike-timing-dependent plasticity (STDP) rules, which we refer to as the “small polynomial search space.” For this search space, similar to previous work [31], the weight from presynaptic neuron i to postsynaptic neuron j , $w_{ij}(t)$, evolved as

$$\frac{d}{dt}w_{ij}(t) = \alpha S_i(t) + \beta S_j(t) + \gamma S_j(t)x_i(t) + \kappa S_i(t)y_j(t) \quad (1)$$

where $S_i(t)$ is the spike train of neuron i , $x_i(t)$ is a low pass filters of the spike train of the pre-synaptic neuron i with time constant τ_{pre} , and $y_j(t)$ is a low pass filters of the spike train of the post-synaptic neuron j with time constant τ_{post} . The spike train is defined as

$S_i(t) = \sum_k \delta(t - t_k^i)$, where t_k^i is the time of the k -th spike of neuron i and δ is the Dirac delta. In total, this search space comprised six tunable plasticity parameters: $\theta = [\alpha, \beta, \gamma, \kappa, \tau_{pre}, \tau_{post}]$.

Previous work uncovered inhibitory to excitatory (I-to-E) rules that would enforce network stability in a feedforward setting [31]. To discover such rules in recurrent networks, we meta-learned I-to-E plasticity rules that enforced a target population firing rate of 10 Hz, quantified with a loss function on the network activity (see Methods, Fig 2A, “stability task”).

The ES converged to a rule with low loss values such that network activity remained stable at 10 Hz for twice the longest possible training duration (2 min, Fig 2C). Visualizing the meta-learned rule with a classic pre-post protocol (pairs of pre-/postsynaptic spikes with various delays, Fig 2B, left) revealed that it did not correspond to any known rules [44]. The

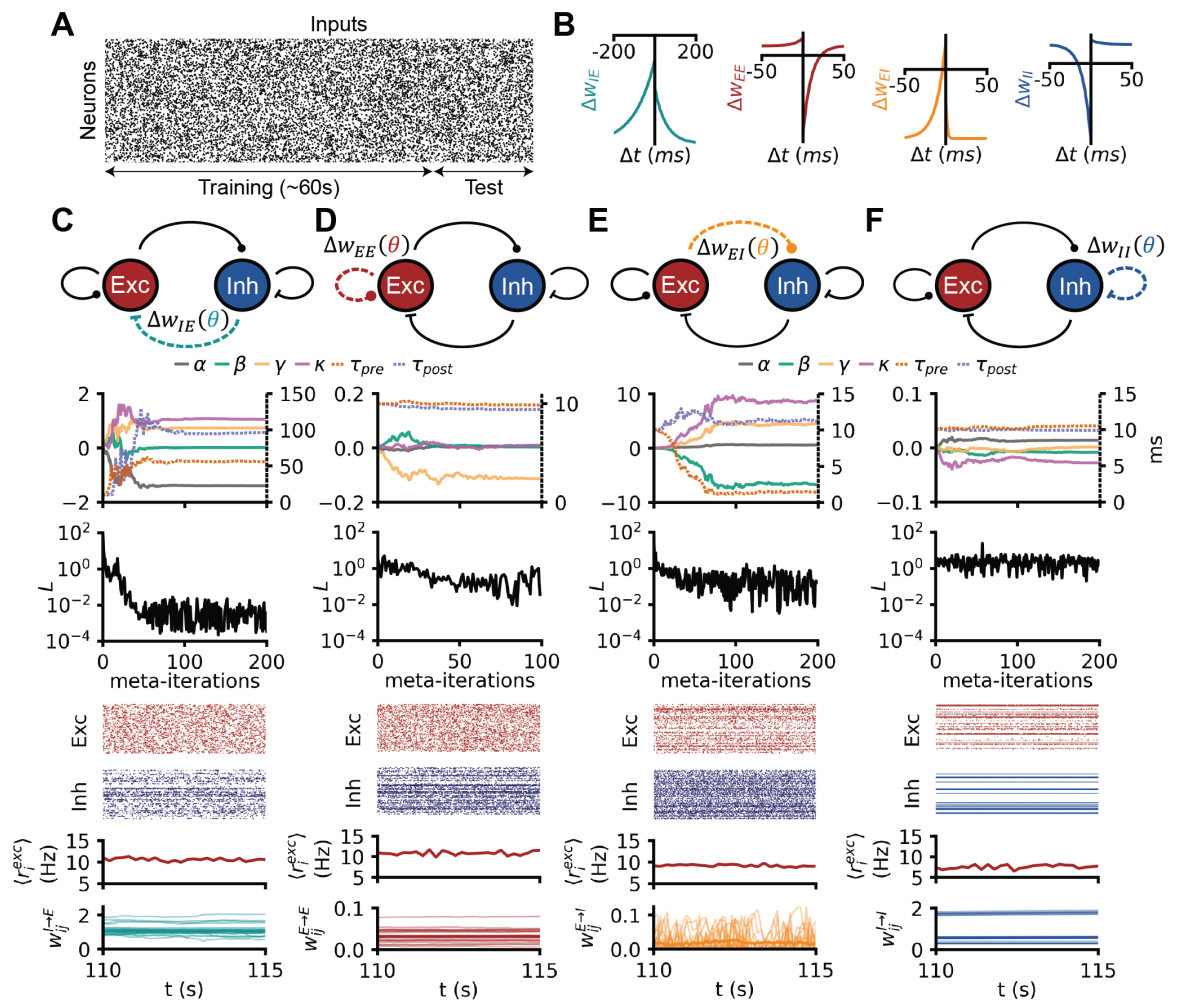


Fig 2. Network stabilization with simple polynomial rules in isolation. (A) Raster plot of inputs received by a recurrent spiking network undergoing the stability task. (B) Pre-post protocols of the four (separately) meta learned plasticity rules in C, D, E and F. (C) A spiking network received Poisson input at a random rate, the E-to-E synapses are plastic with a rule from the small polynomial search space. From top to bottom: (i) evolution of the 6 plasticity parameters during meta learning with CMA-ES. (ii) Evolution of the loss during meta-optimization. (iii) Raster plot of the 200 random excitatory neurons of a network evolving with the final meta learned I-to-E rule. (iv) same as (iii) for the inhibitory neurons. (v) evolution of the population firing rate of excitation (vi) evolution of E-to-E weights (thicker line: mean). D: Same as C, but for E-to-I plasticity. E: Same as C, but for I-to-E plasticity. F: Same as C, but for I-to-I plasticity.

<https://doi.org/10.1371/journal.pcbi.1012910.g002>

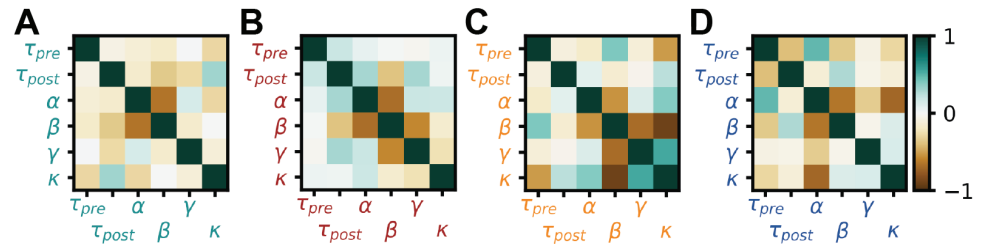


Fig 3. Interpretation of meta learned rules for stability. (A) Covariance matrix at meta-iteration 15 of the optimization in Fig 2C (See also Supplementary Materials and S2 Fig). (B) Same as A for the optimization shown in Fig 2D. (C) Same as A for the optimization shown in Fig 2E. (D) Same as A for the optimization shown in Fig 2F.

<https://doi.org/10.1371/journal.pcbi.1012910.g003>

rule resembled the symmetric inhibitory rules found in theory and experiments [10,24], but shifted downward such that all pairs of pre-post spikes elicited depression. Note that since in our simulations the number of pre- and post- synaptic spikes (the pre- and postsynaptic firing rates) do not have to be the same, this graph does not mean that the meta learned rule has no potentiation regions (Fig 2B). In the network simulation, the I-to-E weights do settle to intermediate values as a result of both potentiation and depression (Fig 2C, bottom).

Next, we used the same stability constraint to discover plasticity rules in other synapse types (E-to-E, E-to-I, or I-to-I, individually, Fig 2D–2F). In two scenarios—E-to-E and E-to-I—the optimization was able to find rules that established the target firing rate (Fig 2C and 2D), albeit not as effectively as with I-to-E plasticity, as evidenced by relatively high losses at the end of training.

Since all successful rules differed from previously observed rules (Fig 2B), we wanted to understand the inter-dependencies of the learned plasticity parameters. We thus plotted the covariance matrix between plasticity parameters as the meta learned rules emerged during the optimization. The covariance matrix is updated alongside the optimal rule in CMA-ES and contains information about the loss landscape, albeit heuristically. Our analysis revealed that the main structure in the rules was strongly anti-correlated non-Hebbian plasticity parameters, i.e., when α is increased β is likely decreased and vice versa (Fig 3), consistent with mean-field theory (see Mean-field analysis). For I-to-I plasticity, no plasticity rule could be found that improved meaningfully upon the initial (no plasticity) rule (Fig 3D), although we note that the absence of proof is not proof of the absence of a solution.

Familiarity detection with a small polynomial search space

Having established that rules from the small polynomial search space could stabilize recurrent spiking networks, we turned to a more complex, memory-related network function: familiarity detection. This ubiquitous form of memory [45,46] has already been the target of meta learning in rate networks [33], and has been shown to emerge in recurrent spiking networks with finely orchestrated, hand-tuned co-active synaptic plasticity rules [12,13,16].

To meta learn plasticity rules that would produce familiarity detection, we designed a protocol in which we stimulated a recurrent spiking network with the same stimulus multiple times, which we define as the “familiar” stimulus. After extensive stimulation with this familiar stimulus intended to induce strong changes in synapses, we compared network responses with a non-overlapping second stimulus that we define as “novel”. Successfully learning the familiar stimulus meant responding to it with a higher firing-rate compared to the novel stimulus after learning. To accomplish such a learning, we designed a loss function that

constrained plasticity rules such that the network would produce high firing rates to familiar, and low firing rates to novel stimuli (see Methods and S1 Fig).

We started by optimizing only I-to-E plasticity while all other rules were inactive. We refer to such scenarios as single-active rules. The I-to-E plasticity rule belonged to the small polynomial search space. Our ES algorithm converged to a rule that achieved low loss values (Fig 4C) and produced networks that responded more strongly to familiar than to novel

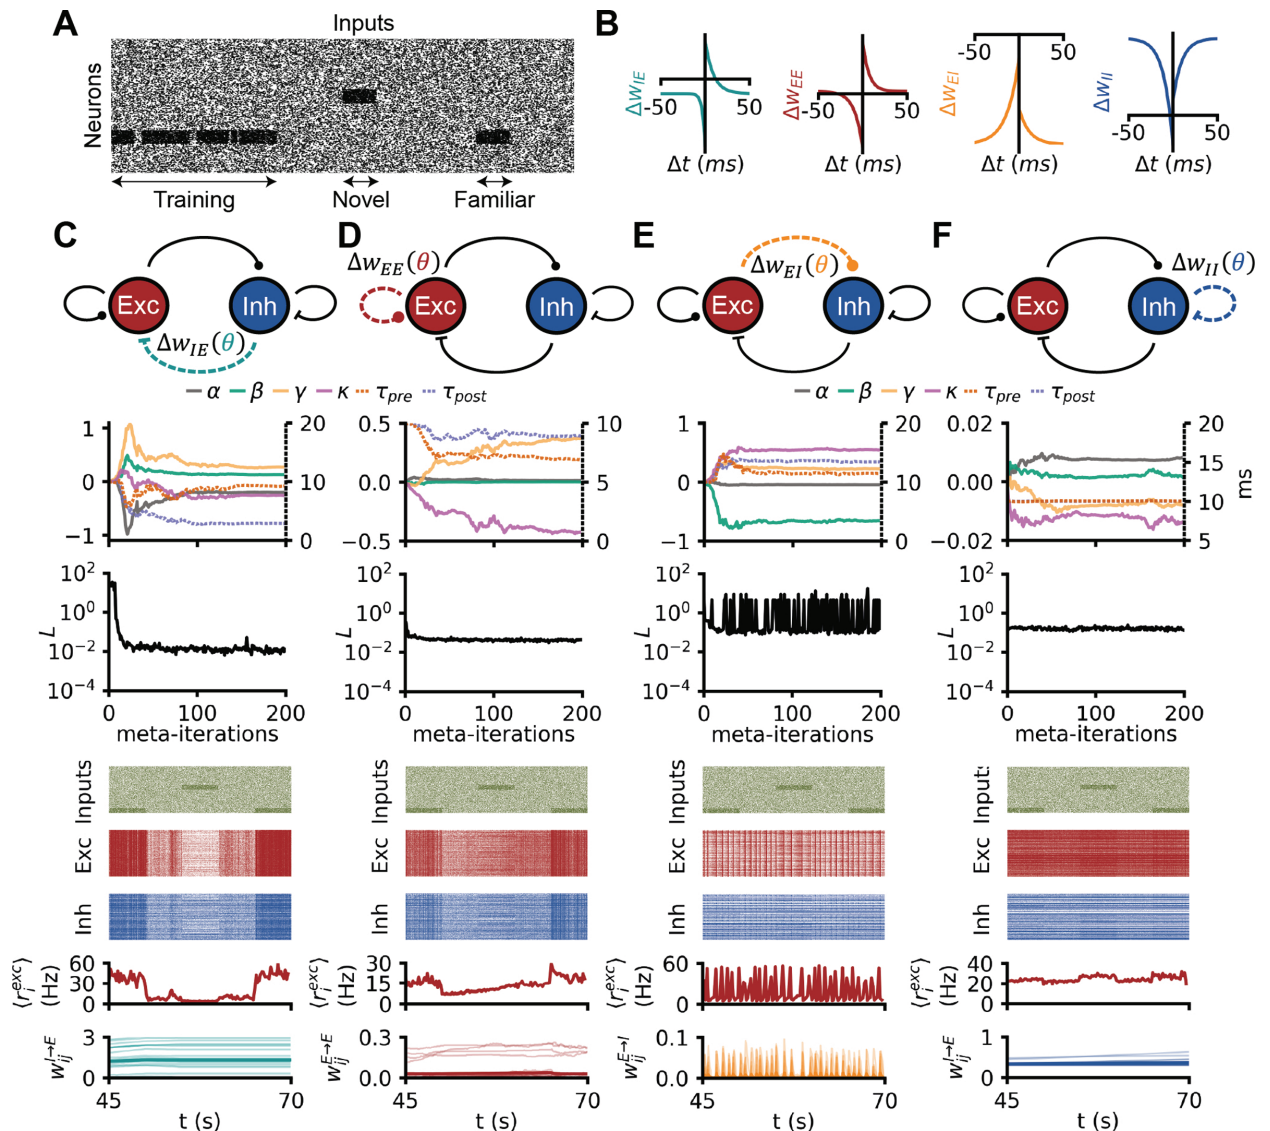


Fig 4. Familiarity detection with simple plasticity rules. (A) Raster plot of inputs received by a recurrent spiking network undergoing the familiarity task: the network is trained on a familiar stimulus, then after a break the network is shown a novel stimulus and the familiar stimulus. (B) Pre-post protocols of the four (separately) meta learned plasticity rules in C, D E and F. (C) A spiking network undergoing the familiarity task: the E-to-E synapses are plastic with a rule from the small polynomial search space. From top to bottom: (i) evolution of the 6 plasticity parameters during meta learning with CMA-ES. (ii) Evolution of the loss during meta optimization. (iii) Raster plot of the 200 random excitatory neurons of a network evolving with the final meta learned I-to-E rule. (iv) same as (iii) for the inhibitory neurons. (v) evolution of the population firing rate of excitation (vi) evolution of E-to-E weights (thicker line: mean). D: Same as C, but for E-to-I plasticity. E: Same as C, but for I-to-E plasticity. F: Same as C, but for I-to-I plasticity.

<https://doi.org/10.1371/journal.pcbi.1012910.g004>

stimuli. As a control, a network undergoing the same task without any plasticity was unable to exhibit asymmetric responses for novel versus familiar stimuli, confirming that the learned plasticity rule was responsible for this acquired behavior (S5 Fig). When we probed the plasticity rule with classical pre/post protocols, we found that the rule did not closely resemble any of the experimentally reported temporal relationships [44]. Notably, familiarity detection was achieved here with a single active I-to-E plasticity rule, contrary to previous work in which memory-related functions were always achieved *in tandem* with E-to-E plasticity [12, 13,16,47].

Next, we focused on E-to-E plasticity in isolation, using the same loss function. The ES found an E-to-E plasticity rule that was able to solve the familiarity task (Fig 4D), and its pre-post protocol resembled previously reported classical asymmetric STDP rules [21,22]. We also considered the other two synapse types (E-to-I, or I-to-I, individually active), but ES could not find satisfying solutions for either (Fig 4E and 4F). Note again, that this result does not prove that no solutions exist within the small polynomial search space for E-to-I or I-to-I rules.

The covariance matrices corresponding to the optimizations for the familiarity task were somewhat similar to the ones obtained on the stability task (α - β anti-correlations for Figs 3A, 3B and 5A, 5B; β - γ and β - κ anti-correlations as well as γ - κ correlation for Figs 3C and 5C; α - β , α - γ and α - κ anti-correlations for Figs 3D and 5D), suggesting similar structure in the relationships of the learned rules for both cases (strong anti-correlation between non-Hebbian parameters, Fig 5).

Familiarity detection with co-active simple polynomial rules

Inspired by previous work proposing co-active E-to-E and I-to-E rules for memory formation in spiking networks [12,13,16], we set out to meta learn jointly the E-to-E and the I-to-E plasticity rules for the familiarity detection task mentioned above (Fig 6). Since either rule (I-to-E or E-to-E) was shown to be able to solve this task individually, ES should succeed in finding at least one solution. As expected, ES converged on a solution that satisfied all constraints and displayed the hallmarks of cortical network dynamics. The learned E-to-E rule was similar to the above-described rule acting in isolation (Fig 4C), displaying a very similar shape as experimentally observed E-to-E rules [3,21,22]. The newly learned I-to-E rule, on the other hand, differed from previous experimental results [7,10,23,48] and also from the previous optimization above, showing an inverse, asymmetric “Bavarian Hat” with a tuft of potentiation. The covariance matrix revealed anti-correlation between non-Hebbian terms of each rule, and

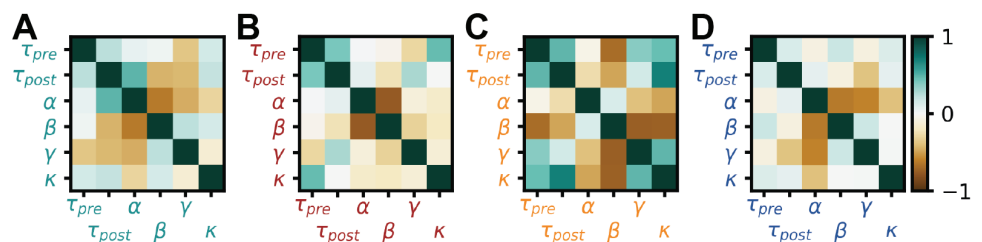


Fig 5. Interpretation of meta learned rules for familiarity detection. (A) Covariance matrix at meta-iteration 15 of the optimization in Fig 4C (B) Same as A for the optimization shown in Fig 4D. (C) Same as A for the optimization shown in Fig 4E. (D) Same as A for the optimization shown in Fig 4F.

<https://doi.org/10.1371/journal.pcbi.1012910.g005>

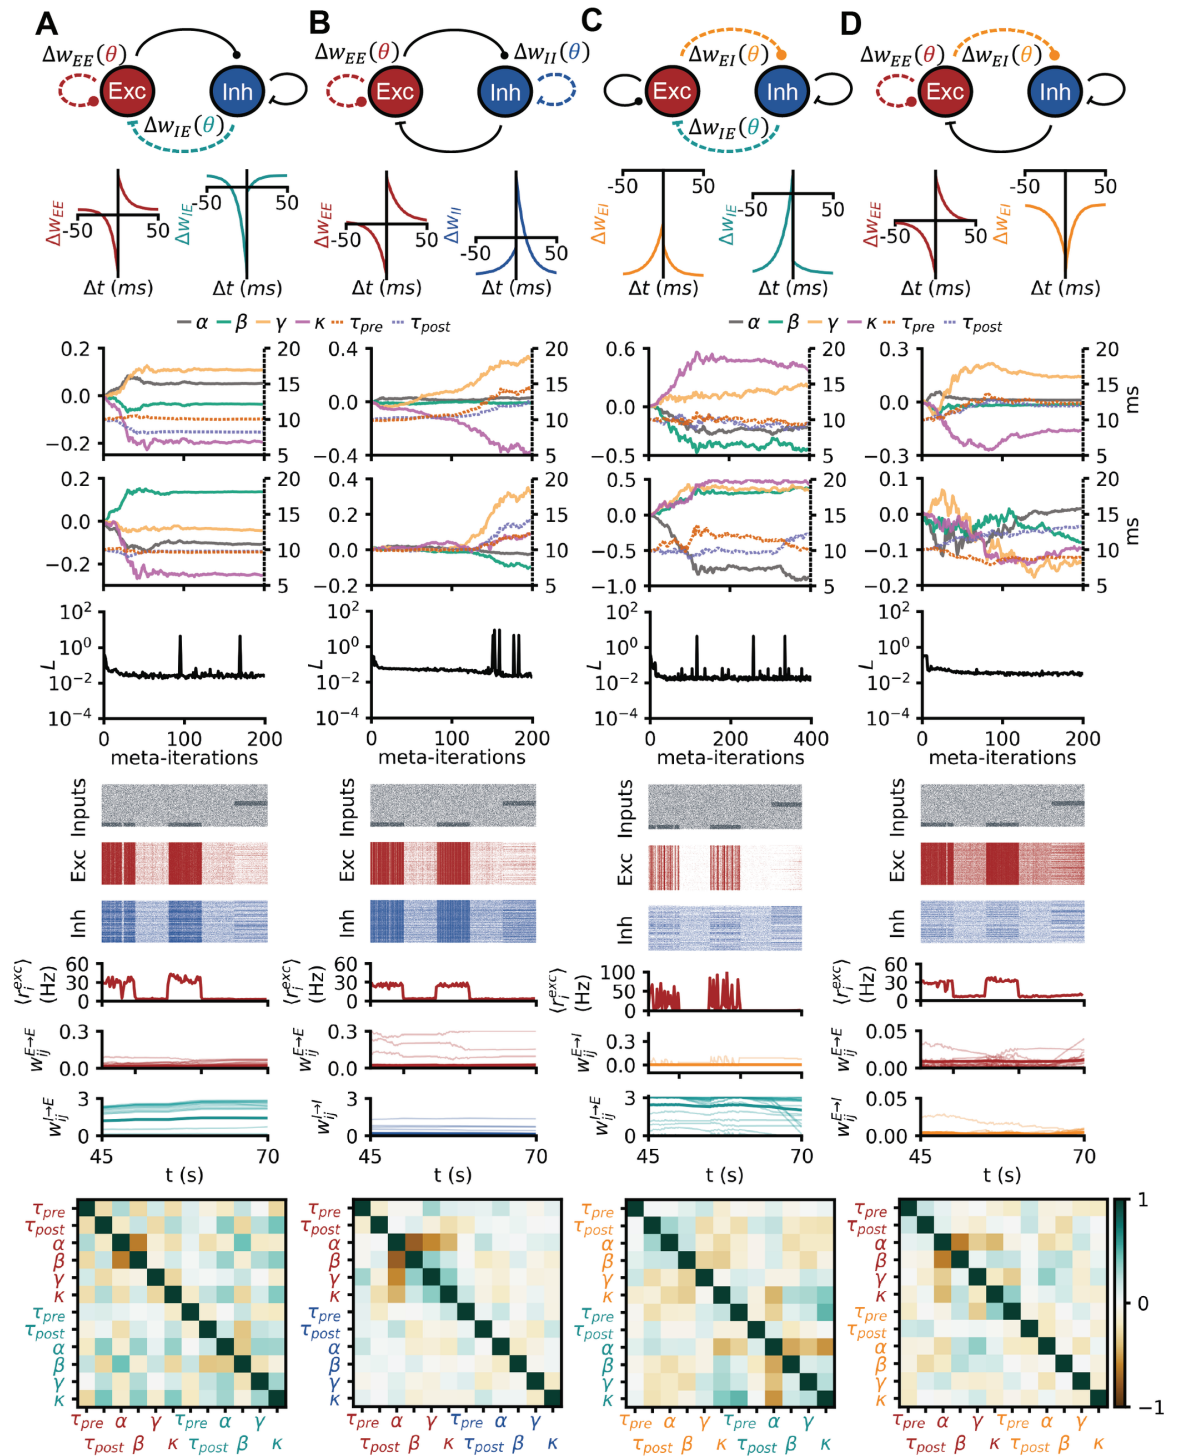


Fig 6. Familiarity detection with simple co-active rules. (A) Same network and familiarity task as in Fig 3, but both the E-to-E and I-to-I weights are plastic with rules from the simple polynomial search space. From top to bottom: network diagram, pre-post protocols of the 2 optimized co-active rules, evolution of the 6 parameters for both rules across the optimization, evolution of the loss across meta-training, covariance matrix at meta-iteration 20. (B) Same as A for a network with tunable E-to-I and I-to-I rule. (C) Same as A for a network with tunable E-to-E and E-to-I rule. (D) Same as A for a network with tunable E-to-E and I-to-E rule.

<https://doi.org/10.1371/journal.pcbi.1012910.g006>

some interactions between parameters of both rules, namely a inverse relationship between non-Hebbian parameters (Fig 6, bottom row).

We also tried other combinations of co-active rules on the same task (E-to-E and I-to-I, E-to-I and I-to-E, as well as E-to-E and E-to-I). In all cases, ES converged to solutions that elicited higher responses to familiar than to novel stimuli, but the network dynamics were biologically implausible (Fig 6), suggesting that E-to-E and I-to-E were the most useful synapse-type for the considered function. Alternatively, it could be that the plasticity rules we used were not flexible or broad enough to express biologically plausible solutions.

More complex plasticity rules

To capture more complex plasticity mechanisms we constructed two higher dimensional and more expressive plasticity rule parameterizations, i.e., (1) a polynomial with additional dependencies and (2) a neural network parameterization (“MLP”, a feedforward network that determines the synaptic changes of the recurrent spiking network, see Methods). We benchmarked these two new parameterizations on the same stability task as for the small search space (Fig 2).

First, we expanded the small polynomial search space, adding synaptic variables that contributed to weight updates, such as additional synaptic traces (triplets rules [8], bursts [17], voltage dependence [9], codependent plasticity [20] and weight dependence [5]),

$$\begin{aligned} \frac{d}{dt} w_{ij}(t) = & S_i(t) f_{pre} (x_{i,\tau_{short}}, x_{i,\tau_{long}}, x_{i,\tau_{long}}, x_{j,\tau_{short}}, w_{ij}, \langle V_j \rangle, C_{j,E}, C_{j,I}) \\ & + S_j(t) f_{post} (x_{i,\tau_{short}}, x_{i,\tau_{long}}, x_{i,\tau_{long}}, x_{j,\tau_{short}}, w_{ij}, \langle V_j \rangle, C_{j,E}, C_{j,I}), \end{aligned} \quad (2)$$

where $S_i(t)$ and $S_j(t)$ are the spike times of pre- and postsynaptic neurons, respectively, f_{pre} and f_{post} are polynomial functions with the following synaptic variables: $x_{i,\tau_{long}}(t)$ and $x_{i,\tau_{short}}(t)$ are low pass filters of the spike train of the pre-synaptic neuron i with time constants $\tau_{short} = 10$ ms and $\tau_{long} = 100$ ms (and similarly for post-synaptic neuron j); $C_{j,E}(t)$ and $C_{j,I}(t)$ are co-dependent terms representing the activity of neighboring synapses, i.e., low-pass filtered with fixed time constants $\tau_{CE} = 10$ ms and $\tau_{CI} = 100$ ms, as in previous work [20]; and $\langle V_j(t) \rangle$ is the low-pass filtered membrane potential, with a time-constant $\tau_{(V)} = 100$ ms [9]. We assumed separability of the synaptic variables, i.e., that the synaptic variables contributed independently to weight updates, which allowed us to incorporate additional dependencies to the weight updates without bloating the total number of plasticity parameters. Meta-learning I-to-E rules on the stability task in the larger polynomial search space resulted in solutions that achieved low losses (Fig 7A). However, when simulating the learned rule for longer than during training, we observed that the rule did not generalize as well as the rules from Fig 2A, with the excitatory activity showing large oscillations around the desired target of 10 Hz (Fig 7A). Additionally, some I-to-E weights reached the maximum weight (10, Fig 7A). The covariance matrix revealed a much sparser structure. The shape of the rule was ambiguous under the pre-post protocol, as it did not constrain the values of the additional synaptic variables.

Motivated by previous work proposing co-active rules that support memory formation and recall in spiking networks [12,13], we considered the same familiarity task and loss function as above, with E-to-E and I-to-E synapses plastic parameterized with the bigger polynomial search space (Fig 7G). Once again, we could meta learn rules that solved the task (Fig 7G). We verified that the learned co-active rules were able to elicit different population responses to the familiar and novel stimuli (Fig 7F). However, other aspects of network activity that were

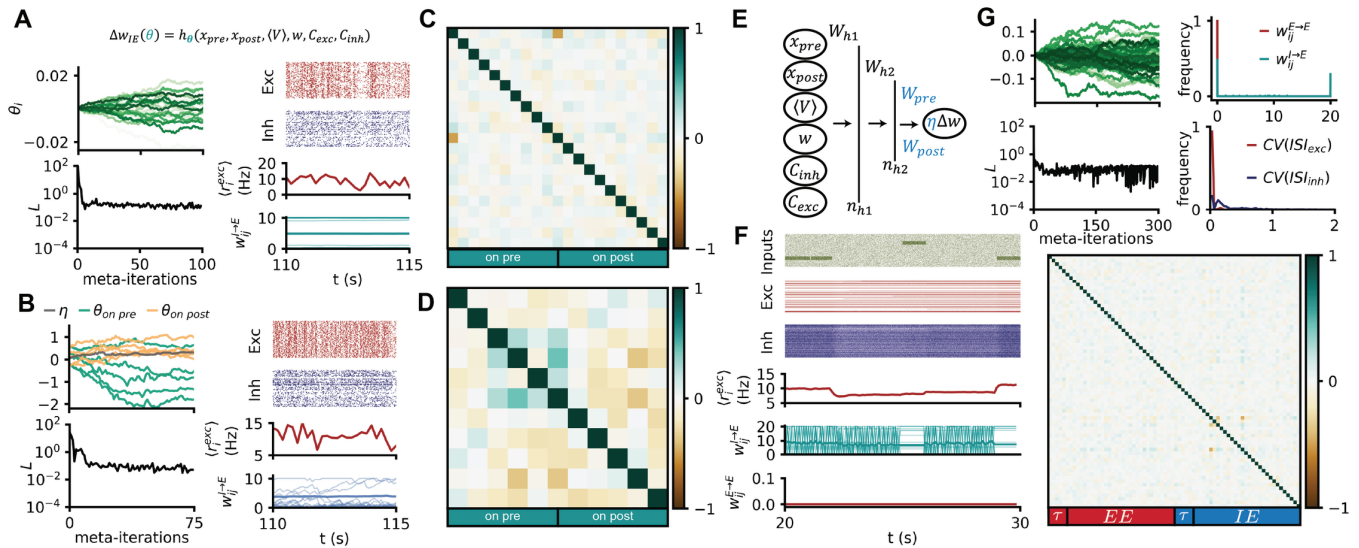


Fig 7. Meta learning complex plasticity rules with ES. (A) CMA-ES on I-to-E plasticity within the big polynomial search space for the stability task. Left: Evolution of the plasticity parameters and loss values along the optimization. Right: example network activity elicited by the meta learned I-to-E rule. (B) Same as A for I-to-E rules from the MLP search space. (C) Covariance-matrix of the plasticity parameters for the optimization shown in A. (D) Covariance-matrix of the plasticity parameters for the optimization shown in B. (E) Schematics of the MLP search-space: weight updates in the spiking network are computed by running forward an MLP with synaptic variables as inputs. (F,G) CMA-ES on E-to-E and I-to-E plasticity within the big polynomial search space for the familiarity detection task. G, left: Evolution of the plasticity parameters and loss values along the optimization. G, right and F: example network activity elicited by the meta learned I-to-E rule. F, right: Covariance-matrix of the plasticity parameters for the optimization shown in G.

<https://doi.org/10.1371/journal.pcbi.1012910.g007>

not constrained by the loss function were unrealistic. For example, most neurons in the network were either silent or fired at unrealistic rates with in highly regular patterns (Fig 7F and 7G). The E-to-E connections mostly converged to zero weights (Fig 7F and 7G). The I-to-E connections underwent rapid switching between 0 and the maximum allowed weight at the millisecond scale, resulting in a bimodal distribution (Fig 7F and 7G).

Finally, we considered a neural-network-based search space for plasticity rules, in which the same synaptic variables as in the big polynomial were combined using a (generalized) multilayer perceptron (MLP, Fig 7B) [49]. Plastic synapses from this search-space underwent spike-triggered updates such that:

$$\frac{d}{dt} w_{ij}(t) = S_i(t) \text{MLP}_{\text{pre}}(x_i, \tau_{\text{long}}, x_j, \tau_{\text{short}}, w_{ij}, \langle V_j \rangle, C_{j,E}, C_{j,I}) + S_j(t) \text{MLP}_{\text{post}}(x_i, \tau_{\text{short}}, x_j, \tau_{\text{long}}, w_{ij}, \langle V_j \rangle, C_{j,E}, C_{j,I}), \tag{3}$$

using similar notations as for the big polynomial search-space (see [More complex plasticity rules](#)).

Similarly to the case with a single connection being plasticity, the MLP used to model synaptic changes taking place in the recurrent spiking network was a feedforward network with two hidden layers (50 and 4 hidden units, see Methods), in which only the final layer weights and bias were tunable, keeping all other layers fixed [49]. This design choice effectively decoupled the number of synaptic variables involved in the rule and the number of plasticity parameters to optimize, thus allowing for potentially highly non-linear dependencies on the synaptic variables while keeping the dimensionality of the search space as low as desired for the evolutionary strategy. This search space comprised a total of 11 parameters:

5 parameters for updates triggered by presynaptic spikes, 5 for postsynaptic updates, and a common learning rate (Fig 7E).

The evolutionary strategy was able to find plasticity rules that established the target firing rate in the MLP search space (Fig 7B). Similar to the big polynomial case, however, the learned rule was not as robust when tested on longer-than-training durations. In addition, the rule elicited biologically implausible network behaviours, for example, weights reaching the maximum allowed value and synchronous spiking patterns (Fig 7B).

Overall, all three parameterizations—small polynomial, big polynomial and MLP—led to rules that solved the task as it was quantified by the loss function. However, the meta learned rules from larger plasticity search spaces did not generalize as well as the simpler rules, and the resulting plastic networks exhibited implausible behaviors, such as synchronous regular firing patterns and bimodal distributions. In our hands, designing a loss function that constrained task performance alone was not sufficient to ensure that biologically relevant plasticity rules emerged.

Interpreting learned rules and degeneracy

Concerned by the impact of potential degeneracy on the rules proposed in this study, we set to test how reliable our rule predictions were on the familiarity task. Running two (intrinsically stochastic) evolutionary searches from the same starting point on the familiarity task with an I-to-E small polynomial rule converged to two plasticity rules with dissimilar pre-post protocols (Fig 8A). This shows, in agreement with previous work in rate networks [41], that at least two and probably many plasticity rules from the same search space can solve this task. This conclusion is not unique to I-to-E plasticity (Fig 8B). However, even though the plasticity rules differed across optimizations, the relationship between plasticity parameters appeared to be conserved. For example, we observed strong anti-correlations between non-Hebbian parameters in all simulations, as shown by the covariance matrices (Fig 8).

Discussion

In this study, we scaled up the automatic tuning of plasticity rules for homeostatic and memory-related tasks from single spiking neurons to large recurrent spiking networks. We used an evolutionary strategy to adjust flexibly parameterized plasticity rules in several search spaces and showed the potential and limitations of this gradient-free meta learning approach for *in silico* plasticity rule discovery.

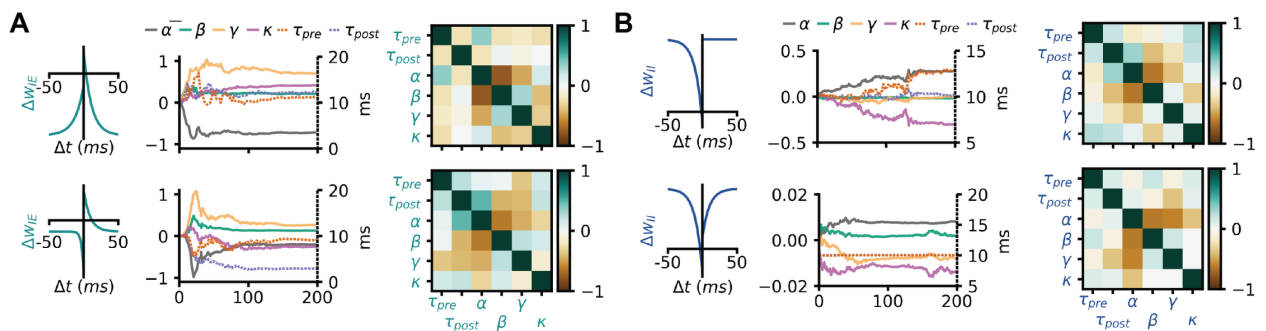


Fig 8. Degeneracy and solution manifolds. (A) Bottom: optimization for an I-to-E small polynomial rule on the familiarity task shown in Fig 4C. Top: pre-post protocol, parameter evolution and covariance matrix of another optimization for an I-to-E small polynomial rule on the familiarity task. (B) Same as A, for an I-to-I small polynomial rule on the familiarity task.

<https://doi.org/10.1371/journal.pcbi.1012910.g008>

As expected from previous work [10,11,15,31,49], we could find isolated I-to-E plasticity rules that enforced firing rate homeostasis (Fig 2). Similar homeostatic network effects could be achieved with isolated E-to-E and E-to-I rules (Fig 2). In our hands, no I-to-I plasticity rule could be found to serve rate homeostasis (Fig 2). We assumed that homeostasis by way of I-to-I synapses is more difficult to achieve because they can only affect the rates of excitatory neurons indirectly (see *Mean-field analysis*).

To get a better understanding of the meta learned plasticity rules, we made use of the covariance matrix as it emerged during meta learning with CMA-ES [40]. We interpreted deviations from zero in this matrix as acquired biases in the sampling of new plasticity rule candidates, i.e., the matrix unveiled likely interdependencies between parameters that led to successful plasticity action; strong anti-correlations between non-Hebbian parameters was a hallmark of all successful rules (Fig 3). We interpreted these inverse relationships to imply that even for Hebbian-seeming rules a substantial part of the weight changes were effected by pre-only and post-only—i.e., non-Hebbian—terms.

Next we searched for rules that could perform a computational task. We chose familiarity detection as a fundamental component of any memory function beyond mere activity homeostasis. Isolated E-to-E rules that were sufficient for this function could be found (Fig 4), in contrast to previous work that requires several co-active, finely orchestrated rules [12,13]. However, we did not check the long-term stability of the representations achieved by our rules, nor the emergence of attractor dynamics, which may require additional complexity, or additional types of rules. Moreover, isolated I-to-E plasticity rules could also be found to solve this memory task, hinting at a greater role for I-to-E plasticity, with a wide range of potential functions other than network stabilization [48,50]. Isolated E-to-I and I-to-I rules could not be found to store patterns in spiking networks, justifying *post hoc* the relative dearth of previous modeling studies on the function of these synapse types. Of course, the absence of proof does not prove the absence of solutions that could perform familiarity detection with isolated E-to-I or I-to-I plasticity. Remarkably, the parameter interdependencies for the familiarity task as revealed by the meta learned covariance matrix were similar to the ones for the stability task (Figs 3 and 5), suggesting that the same plasticity mechanisms could enforce homeostasis *and* support basic memory functions [49].

When we broadened our search spaces to successfully meta learn multiple co-active plasticity rules with similar basic memory functions (Fig 6), our joy in finding these sets of rules was somewhat tempered by the fact that isolated individual rules could solve the task at hand already, but our results provided a proof of principle for the possibility of discovering ensembles of co-active rules [49].

We also broadened the complexity of how we parametrized individual rules, going from simple polynomials [31] to expressions with more synaptic variables using either larger polynomials, or MLPs (Fig 7). Such an expansion of the search space did not scale well with regards to the compute requirements of the outer loop. We thus proposed a partially tunable MLP as a means to non-linearly mix multiple synaptic variables without bloating the parameter number, in line with previous work [49]. The added complexity in the rule space resulted in decreased robustness and generality of the meta learned rules (Fig 7), suggesting that the loss function was not constrained enough for these more flexible rules (Fig 7). Therefore, though the plasticity rules were automatically tuned by way of meta learning, the loss function now required extensive hand-tuning to effectively force network activity and weight dynamics into plausible regimes. To summarize, the pitfalls of the ES were, in our hands, (1) the known performance drop of CMA-ES in higher dimensions, (2) an exponential increase in compute time with increased dimensionality (more and more rules need to be tested per meta-iteration), and (3) the need to restart optimization from scratch upon defining a new

loss function. These three factors rendered hyperparameter optimization and debugging increasingly impractical. For meta learning to be able to scale to large plasticity search spaces and complex network models, we developed an approach elsewhere [49] that alleviates the problem of having to define *a priori* a loss function that controls both task performance and biological plausibility. Here, with ES, we can only guess the full loss function beforehand and subsequently identify flaws in these constraints by inspecting the “optimized” networks *post hoc*; refinement of the loss function then required us to restart the optimization from scratch, at great computational cost (Fig 7).

Moreover, parameters not directly related to synaptic plasticity also have an impact on the meta learned rules. For instance, the initial connectivity can render some plasticity rules stable or unstable [51]. Here, we evaluated each plasticity rule in several networks, with random initial mean weights, connectivity distributions, and input rates. The goal of this averaging was to ensure that we discovered rules that achieved low losses on a range of inputs and initial weight connectivities. However, such averaging could not explore exhaustively more fine-grained initial connectivity motifs [51]. On the other hand, systematically testing connectivity-rule pairs would bloat the parameter space substantially, and render meta learning computationally unfeasible.

Previous work highlights degeneracy of mechanisms in neuroscience and more recently in synaptic plasticity [41–43,49,52]. We confirm these findings, although our ES approach is ill-suited to explore more thoroughly degeneracy due to its local search nature. Notably, we show that degeneracy emerges already in the simplest case (single small polynomial rule on the stability task, Fig 8), which may hint at degeneracy as a ubiquitous phenomenon on plasticity.

Understanding the meta learned rules is challenging, especially in high-dimensional search spaces. In the simpler case of stabilization with the small polynomial search space, we could rely on a predicted subspace of solutions using mean-field theory (see [Mean-field analysis](#), S2 FigC). In the bigger search spaces presented in this work, understanding the resulting learning rules and their relationships to other rules is important to be able to formulate experimental predictions. However, despite the use of an L1 regularization in all optimizations in this study, most meta learned rules still comprised many non-zero parameters (Fig 7) that made a direct comparison to experimental results challenging.

We proposed that the covariance matrix learned with CMA-ES alongside the best rule could help reveal structure in the meta learned parameters. The covariance matrix from an optimization on the stability task in the small polynomial search space is in agreement with insights from mean-field theory (see [Mean-field analysis](#)), in that the task is mainly solved by the non-Hebbian terms (Fig 3). The covariance matrix from optimizations of more complex parameterizations is much sparser than the meta learned solution (Fig 6 and 7) and suggests that a few terms are of special importance for this solution.

Overall, we believe that meta learning approaches for synaptic plasticity face a compromise: simple search spaces are easier to optimize, yet their simplicity makes them already amenable to theoretical analysis and often means no truly novel rules can be discovered. Here, meta learning with genetic search algorithms was successful albeit in very limited realms. On the other hand, it served as a first step and spawned a number of new approaches to automatically scan the uncharted depth of plasticity in the future.

Materials and methods

Neuron and network model

We considered recurrent networks of excitatory and inhibitory, conductance-based leaky-integrate-and-fire neurons. Two types of networks were implemented, emulating either the networks used by Vogels, Sprekeler et al. [10] or by Zenke et al. [13].

Networks following Vogels, Sprekeler et al. [10]. This network comprised 8000 excitatory and 2000 inhibitory neurons. The membrane potential dynamics of neuron j (excitatory or inhibitory) were given by

$$\tau_m \frac{d}{dt} V_j(t) = - (V_j(t) - V_{rest}) - g_j^{AMPA}(t) (V_j(t) - E_{AMPA}) - g_j^{GABA}(t) (V_j(t) - E_{GABA}), \tag{4}$$

with $\tau_m = 20$ ms, $V_{rest} = -60$ mV, $E_{AMPA} = 0$ mV and $E_{GABA} = -80$ mV. A postsynaptic spike was emitted whenever the membrane potential $V_j(t)$ crossed a threshold $V^{th} = -50$ mV, with an instantaneous reset to V_{rest} for the duration of the refractory period, $\tau_{ref} = 5$ ms.

The excitatory and inhibitory conductances, g^{AMPA} and g^{GABA} evolved such that:

$$\begin{aligned} \frac{d}{dt} g_j^{AMPA}(t) &= -\frac{g_j^{AMPA}(t)}{\tau_{AMPA}} + \sum_{i \in \text{Exc}} w_{ij}(t) S_i(t) \\ \frac{d}{dt} g_j^{GABA}(t) &= -\frac{g_j^{GABA}(t)}{\tau_{GABA}} + \sum_{i \in \text{Inh}} w_{ij}(t) S_i(t) \end{aligned} \tag{5}$$

with $\tau_{AMPA} = 5$ ms, $\tau_{GABA} = 10$ ms, $w_{ij}(t)$ the connection strength between neurons i and j (unitless), $S_i(t) = \sum \delta(t - t_i^*)$ the spike train of presynaptic neuron i , where t_i^* denotes the spike times of neuron i , and δ the Dirac delta. Unless mentioned otherwise, all neurons received input from 5000 Poisson neurons, with 5% random connectivity and constant rate $r_{ext} = 7$ Hz. The recurrent connectivity was instantiated with random sparse connectivity (2%).

This network was used for Figs 2 and 3. All other simulations used the network model described below.

Networks following Zenke et al. [13] This network comprised 4096 excitatory and 1024 inhibitory neurons. The membrane potential dynamics of neuron j (excitatory or inhibitory) followed:

$$\tau_m \frac{d}{dt} V_j(t) = - (V_j(t) - V_{rest}) - g_j^E(t) (V_j(t) - E_E) - g_j^I(t) (V_j(t) - E_I), \tag{6}$$

where E stands for excitation and I for inhibition, $\tau_m = 20$ ms, $V_{rest} = -70$ mV, $E_E = 0$ mV and $E_I = -80$ mV.

A postsynaptic spike was emitted whenever the membrane potential $V_j(t)$ crossed a threshold $V_j^{th}(t)$, with an instantaneous reset to $V_{reset} = -70$ mV. This threshold $V_j^{th}(t)$ was incremented by $V_{spike}^{th} = 100$ mV every time neuron j spiked and otherwise decayed following:

$$\tau_{th} \frac{d}{dt} V_j^{th}(t) = V_{base}^{th} - V_j^{th}(t), \tag{7}$$

with $V_{base}^{th} = -50$ mV. The excitatory and inhibitory conductances, g^E and g^I evolved such that

$$\begin{aligned} g_j^E(t) &= a g_j^{AMPA}(t) + (1 - a) g_j^{NMDA}(t) \quad \text{and} \\ \frac{d}{dt} g_j^I(t) &= -\frac{g_j^I(t)}{\tau_{GABA}} + \sum_{i \in \text{Inh}} w_{ij}(t) S_i(t) \end{aligned}$$

$$\begin{aligned} \text{with } \frac{d}{dt}g_j^{\text{AMPA}}(t) &= -\frac{g_j^{\text{AMPA}}(t)}{\tau_{\text{AMPA}}} + \sum_{i \in \text{Exc}} w_{ij}(t)S_i(t) \quad \text{and} \\ \frac{d}{dt}g_j^{\text{NMDA}}(t) &= \frac{g_j^{\text{AMPA}}(t) - g_j^{\text{NMDA}}(t)}{\tau_{\text{NMDA}}}, \end{aligned} \tag{8}$$

with $w_{ij}(t)$ the connection strength between neurons i and j (unitless), $a = 0.23$ (unitless), $\tau_{\text{GABA}} = 10$ ms, $\tau_{\text{AMPA}} = 5$ ms, $\tau_{\text{NMDA}} = 100$ ms, $S_i(t) = \sum \delta(t - t_i^*)$ the spike train of presynaptic neuron i , where t_i^* denotes the spike times of neuron k , and δ the Dirac delta. Unless mentioned otherwise, all neurons received input from 5000 Poisson neurons, with 5% recurrent connectivity and constant rate $r_{\text{ext}} = 7$ Hz. The recurrent connectivity was instantiated with random sparse connectivity (10%).

Plasticity rule parameterization

In this study, we considered three parameterizations for plasticity rules with various levels of complexity and expressivity.

“Small polynomial” parameterization. This polynomial search space, initially defined in [31], captured first order Hebbian spike-triggered updates:

$$\frac{d}{dt}w_{ij}(t) = \alpha S_i(t) + \beta S_j(t) + \gamma S_j(t)x_i(t) + \kappa S_i(t)x_j(t) \tag{9}$$

with $S_i(t) = \sum_k \delta(t - t_k^i)$ the spike train of neuron i , δ the Dirac delta function to denote the presence of a pre (post)-synaptic spike at time t . The synaptic traces x_i and x_j are low-pass filters of the activity of presynaptic neuron i and postsynaptic neuron j , with time constants τ_{pre} and τ_{post} , such that:

$$\frac{d}{dt}x_i(t) = -\frac{x_i(t)}{\tau_{\text{pre}}} + S_i(t) \quad \text{and} \quad \frac{d}{dt}x_j(t) = -\frac{x_j(t)}{\tau_{\text{post}}} + S_j(t), \tag{10}$$

Overall, this search space comprised 6 tunable plasticity parameters: $\theta = [\alpha, \beta, \gamma, \kappa, \tau_{\text{pre}}, \tau_{\text{post}}]$. Note that when these parameters were meta learned, the positivity constraints on the time constants τ_{pre} and τ_{post} were enforced by optimizing the natural logarithm of the time constants.

“Big polynomial” parameterization. Plasticity rules in this search space were parameterized such that:

$$\begin{aligned} \frac{d}{dt}w_{ij}(\theta, w_{ij}, x_i, \tau_{\text{short}}, x_i, \tau_{\text{long}}, x_j, \tau_{\text{short}}, x_j, \tau_{\text{long}}, \langle V_j \rangle, C_{j,E}, C_{j,I}) = \\ S_i(t)\theta_0 [1 + \theta_1 + \theta_2 w_{ij} + \theta_3 w_{ij}^2] [1 + \theta_4 \langle V_j \rangle] [1 + \theta_5 C_{j,E} + \theta_6 C_{j,E}^2] \\ [1 + \theta_7 C_{j,I}] [1 + \theta_8 x_i, \tau_{\text{long}}] [1 + \theta_9 x_j, \tau_{\text{short}}] + \\ S_j(t)\theta_{10} [1 + \theta_{11} + \theta_{12} w_{ij} + \theta_{13} w_{ij}^2] [1 + \theta_{14} \langle V_j \rangle] [1 + \theta_{15} C_{j,E} + \theta_{16} C_{j,E}^2] \\ [1 + \theta_{17} C_{j,I}] [1 + \theta_{18} x_i, \tau_{\text{short}}] [1 + \theta_{19} x_j, \tau_{\text{long}} + \theta_{20} x_j^3, \tau_{\text{long}}], \end{aligned} \tag{11}$$

with $C_{j,E}(t) = \langle g_j^E(t) (E_E - V_j(t)) \rangle$ and $C_{j,I}(t) = \langle g_j^I(t) (E_I - V_j(t)) \rangle$ co-dependent terms representing the activity of neighboring synapses, which were low-pass filtered with fixed time constants $\tau_{\text{CE}} = 10$ ms and $\tau_{\text{CI}} = 100$ ms, as in previous work [20]. $\langle V_j(t) \rangle$ the low-pass filtered membrane potential, with a time-constant $\tau_{\langle V \rangle} = 100$ ms [9]. Note that, unlike for the small

polynomial search space, all timescales in this search space were not learned, and fixed to values compatible with experimental data and previous studies [9,10,12,13]. The timescales for the synaptic traces were: $\tau_{EE}^{(1)} = \tau_{IE}^{(1)} = 10$ ms and $\tau_{EE}^{(2)} = \tau_{IE}^{(2)} = 100$ ms.

Overall, this search space amounted to 21 plasticity parameters per synapse type.

“MLP” parameterization. In line with previous work [49], we chose a two-hidden-layer fully-connected feedforward network (“MLP”), composed of 50 sigmoidal units in the first hidden layer and 4 in the second.

In this MLP search space, the same plasticity variables as for the big polynomial were combined such that:

$$\frac{d}{dt} w_{ij}(t) = S_i(t) \text{MLP}_{\text{pre}}(x_i, \tau_{\text{long}}, x_j, \tau_{\text{short}}, w_{ij}, \langle V_j \rangle, C_{j,E}, C_{j,I}) + S_j(t) \text{MLP}_{\text{post}}(x_i, \tau_{\text{short}}, x_j, \tau_{\text{long}}, w_{ij}, \langle V_j \rangle, C_{j,E}, C_{j,I}), \tag{12}$$

The input layer of the MLP was composed of 6 neurons, with the values of the relevant synaptic variables during a spike-triggered update. This layer was followed by a first fully connected hidden layer with 50 units and sigmoid non-linearity, then by another fully connected layer with 4 units and sigmoid non-linearity. The final layer was linear, fully connected. The weights of the 2 hidden layers were randomly initialized and fixed ($\sim \mathcal{U}(\frac{-1}{n_{\text{inp}}}, \frac{1}{n_{\text{inp}}})$, where n_{inp} is the number of input features at a given layer), with identical values for the on-pre and on-post MLPs. Only the weights, bias, and output learning rate of the final linear layer were trained, for a total of 4 weights + 1 bias for each MLP, as well as a common learning rate for a total of 11 plasticity parameters per plastic synapse type (see Fig 1B).

Mean-field analysis

Within the small polynomial search space, we performed mean-field analysis on the I-E connections to link the plasticity parameters and the population firing rates at steady state $r_{\text{exc}}^*, r_{\text{inh}}^*$, as done in previous work [10,15,31,49]:

$$\text{IE} : r_{\text{exc}}^* = - \frac{\alpha r_{\text{inh}}^*}{\beta + (\kappa \tau_{\text{post}} + \gamma \tau_{\text{pre}}) r_{\text{inh}}^*} \tag{13}$$

with the additional conditions $\alpha < 0$ and $\beta + (\kappa \tau_{\text{post}} + \gamma \tau_{\text{pre}}) r_{\text{inh}}^* > 0$ for stability.

We can perform similar derivations for the other three synapse types plastic in isolation:

$$\text{EE} : r_{\text{inh}}^* = - \frac{\alpha + \beta}{\kappa \tau_{\text{post}} + \gamma \tau_{\text{pre}}} \tag{14}$$

$$\text{EI} : r_{\text{inh}}^* = - \frac{\alpha r_{\text{exc}}^*}{\beta + (\kappa \tau_{\text{post}} + \gamma \tau_{\text{pre}}) r_{\text{exc}}^*} \tag{15}$$

$$\text{II} : r_{\text{inh}}^* = - \frac{\alpha + \beta}{\kappa \tau_{\text{post}} + \gamma \tau_{\text{pre}}} \tag{16}$$

Meta learning plasticity rules with evolutionary strategies

The parameters θ of the plasticity rules were optimized using an evolutionary strategy (CMA-ES [40]). It is difficult to compute usable gradients in long unrolled computational graphs, such as spiking networks, due to exploding or vanishing gradients [53]. In the case of spiking networks, simulation time-steps have to be small (0.1 ms), and total simulation times need to be long enough to give time for plasticity to carve the weights in the network at biologically

realistic timescales. Aware of the instability problems encountered in the training of learned optimizers in Machine Learning [35,53,54], we thus use evolutionary strategies instead, for their smoothing properties [53,54].

We chose CMA-ES for its robustness and low number of hyperparameters. Briefly, at meta-iteration $i + 1$, a set of n plasticity rules to evaluate is generated such that:

$$\{\theta_k\}_{1 \leq k \leq n} \sim \mathcal{N}(\theta_i^*, C_i) \quad (17)$$

with θ_i^* the current best "guess" at this stage of the optimization, and C_i a covariance matrix, both of which are updated at each meta-iteration based on the scores of the set of n rules tested at the current meta-iteration as well as their previous values [40].

Such gradient-free optimization strategies require the simulation of many plastic spiking networks. The generation size parameter of CMA-ES n was typically chosen to be twice the number of plasticity parameters, and the number of trials N_{trials} over which to evaluate a single meant that every meta-iteration required the simulation of nN_{trials} (values chosen between 4 and 10) separate recurrent spiking networks. Plastic networks were simulated in C++ using Auryn, a fast simulation software for spiking networks [55].

Covariance matrix: The covariance matrix is updated at every meta-iteration in CMA-ES. In Figs 3, 5 and 6, we only show the covariance matrix at one meta-iteration, before the loss plateaus (typically meta-iterations 10 to 20). Once the loss plateaus, we observed that most terms in the covariance became close to 1 or -1 (S2 Fig).

Supporting information

S1 Fig. Familiarity task loss function visualization. Visualization of the loss function used for the familiarity task.

(PNG)

S2 Fig. Interpretation of plasticity rules. (A) Covariance matrices at the last meta-iterations for the optimizations shown in Fig 2 and 3. (B) Optimization from Fig 2C, evolution of two plasticity parameters during the optimization trajectory. Dotted line is the mean-field theoretical prediction with the non-Hebbian terms only (same analysis as in [31]). This suggests that the task is being solved mainly via the two non-Hebbian parameters, an interpretation in line with the covariance matrix visualization.

(PNG)

S3 Fig. Another optimization with I-to-E plasticity. More details about the optimization shown in Fig 8A.

(PNG)

S4 Fig. Another familiarity detection optimization with I-to-I plasticity. More details about the optimization shown in Fig 8B.

(PNG)

S5 Fig. Familiarity detection without synaptic plasticity. From top to bottom, raster plot of input neurons to a network identical the ones used in Fig 4 and 6, with all connections static; raster plot of excitatory neurons; raster plot of inhibitory neurons; firing rate of the excitatory population.

(PNG)

S6 Fig. Delayed familiarity detection with flexible co-active rules. (A) Same network and plasticity search space as in Fig 6, but the task now involves a delay between stimulus presentation and measure of the population activity. (B) Top: evolution of the plasticity parameters across meta-training. The parameters are grouped according to whether they belong to the E-to-E or I-to-E rule, and whether they are part of the weight updates triggered by a presynaptic or by a postsynaptic spike. Bottom: corresponding evolution of the loss function. Right: Network simulated with the learned rule.
(PNG)

Acknowledgments

We would like to thank Chaitanya Chintaluri, Nicoleta Condruz and Douglas Feitosa Tomé for insightful discussions.

Author contributions

Conceptualization: Basile Confavreux, Everton J. Agnes, Friedemann Zenke, Tim P. Vogels.

Formal analysis: Basile Confavreux.

Methodology: Basile Confavreux, Everton J. Agnes, Friedemann Zenke, Tim P. Vogels.

Software: Basile Confavreux.

Supervision: Everton J. Agnes, Henning Sprekeler, Tim P. Vogels.

Visualization: Basile Confavreux.

Writing – original draft: Basile Confavreux, Tim P. Vogels.

Writing – review & editing: Basile Confavreux, Everton J. Agnes, Friedemann Zenke, Henning Sprekeler, Tim P. Vogels.

References

1. Bienenstock EL, Cooper LN, Munro PW. Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex. *J Neurosci.* 1982;2(1):32–48. <https://doi.org/10.1523/jneurosci.02-01-00032.1982> PMID: 7054394
2. Oja E. Simplified neuron model as a principal component analyzer. *J Math Biol.* 1982;15(3):267–73. <https://doi.org/10.1007/bf00275687> PMID: 7153672
3. Gerstner W, Kempter R, Van Hemmen JL, Wagner H. A neuronal learning rule for sub-millisecond temporal coding. *Nature.* 1996;383(6595):76–8. <https://doi.org/10.1038/383076a0> PMID: 8779718
4. Abbott LF, Nelson SB. Synaptic plasticity: taming the beast. *Nat Neurosci.* 2000;3:1178–83. <https://doi.org/10.1038/81453> PMID: 11127835
5. Van Rossum MC, Bi GQ, Turrigiano GG. Stable Hebbian learning from spike timing-dependent plasticity. *J Neurosci.* 2000;20(23):8812–21. <https://doi.org/10.1523/jneurosci.20-23-08812.2000> PMID: 11102489
6. Gerstner W, Kistler WM. Mathematical formulations of Hebbian learning. *Biol Cybern.* 2002;87(5):404–15. <https://doi.org/10.1007/s00422-002-0353-y> PMID: 12461630
7. Haas JS, Nowotny T, Abarbanel HD. Spike-timing-dependent plasticity of inhibitory synapses in the entorhinal cortex. *J Neurophysiol.* 2006;96(6):3305–13. <https://doi.org/10.1152/jn.00551.2006> PMID: 16928795
8. Pfister JP, Gerstner W. Triplets of spikes in a model of spike timing-dependent plasticity. *J Neurosci.* 2006;26:9673–9682. <https://doi.org/10.1523/jneurosci.1425-06.2006> PMID: 16988038
9. Clopath C, Büsing L, Vasilaki E, Gerstner W. Connectivity reflects coding: a model of voltage-based STDP with homeostasis. *Nat Neurosci.* 2010;13:344–352. <https://doi.org/10.1038/nn.2479> PMID: 20098420

10. Vogels TP, Sprekeler H, Zenke F, Clopath C, Gerstner W. Inhibitory plasticity balances excitation and inhibition in sensory pathways and memory networks. *Science*. 2011;334:1569–1573. <https://doi.org/10.1126/science.1211095> PMID: 22075724
11. Luz Y, Shamir M. Balancing feed-forward excitation and inhibition via Hebbian inhibitory synaptic plasticity. *PLoS Comput Biol*. 2012;8:1–12. <https://doi.org/10.1371/journal.pcbi.1002334> PMID: 22291583
12. Litwin-Kumar A, Doiron B. Formation and maintenance of neuronal assemblies through synaptic plasticity. *Nat Commun*. 2014;5:5319. <https://doi.org/10.1038/ncomms6319> PMID: 25395015
13. Zenke F, Agnes EJ, Gerstner W. Diverse synaptic plasticity mechanisms orchestrated to form and retrieve memories in spiking neural networks. *Nat Commun*. 2015;6:6922. <https://doi.org/10.1038/ncomms7922> PMID: 25897632
14. Magee JC, Grienberger C. Synaptic plasticity forms and functions. *Annu Rev Neurosci*. 2020;43:95–117. <https://doi.org/10.1146/annurev-neuro-090919-022842> PMID: 32075520
15. Agnes EJ, Luppi AI, Vogels TP. Complementary inhibitory weight profiles emerge from plasticity and allow flexible switching of receptive fields. *J Neurosci*. 2020;40(50):9634–49. <https://doi.org/10.1523/jneurosci.0276-20.2020> PMID: 33168622
16. Schulz A, Miehl C, Berry II MJ, Gjorgjieva J. The generation of cortical novelty responses through inhibitory plasticity. *Elife*. 2021;10:e65309. <https://doi.org/10.7554/elife.65309> PMID: 34647889
17. Payeur A, Guerguiev J, Zenke F, Richards BA, Naud R. Burst-dependent synaptic plasticity can coordinate learning in hierarchical circuits. *Nat Neurosci*. 2021;24(7):1010–19. <https://doi.org/10.1038/s41593-021-00857-x> PMID: 33986551
18. Soldado-Magraner S, Seay MJ, Laje R, Buonomano DV. Paradoxical self-sustained dynamics emerge from orchestrated excitatory and inhibitory homeostatic plasticity rules. *Proc Natl Acad Sci U S A*. 2022;119(43):e2200621119. <https://doi.org/10.1073/pnas.2200621119> PMID: 36251988
19. Halvagal MS, Zenke F. The combination of Hebbian and predictive plasticity learns invariant object representations in deep sensory networks. *Nat Neurosci*. 2023;26(11):1906–15. <https://doi.org/10.1038/s41593-023-01460-y> PMID: 37828226
20. Agnes EJ, Vogels TP. Co-dependent excitatory and inhibitory plasticity accounts for quick, stable and long-lasting memories in biological networks. *Nat Neurosci*. 2024;27:964–74. <https://doi.org/10.1038/s41593-024-01597-4> PMID: 38509348
21. Markram H, Lübke J, Frotscher M, Sakmann B. Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science*. 1997;275(5297):213–5. <https://doi.org/10.1126/science.275.5297.213> PMID: 8985014
22. Bi Gq, Poo Mm. Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. *J Neurosci*. 1998;18(24):10464–72. <https://doi.org/10.1523/jneurosci.18-24-10464.1998> PMID: 9852584
23. Woodin MA, Ganguly K, Poo Mm. Coincident pre-and postsynaptic activity modifies GABAergic synapses by postsynaptic changes in Cl⁻ transporter activity. *Neuron*. 2003;39(5):807–20. [https://doi.org/10.1016/s0896-6273\(03\)00507-5](https://doi.org/10.1016/s0896-6273(03)00507-5) PMID: 12948447
24. D'amour JA, Froemke RC. Inhibitory and excitatory spike-timing-dependent plasticity in the auditory cortex. *Neuron*. 2015;86(2):514–28. <https://doi.org/10.1016/j.neuron.2015.03.014> PMID: 25843405
25. Lagzi F, Bustos MC, Oswald AM, Doiron B. Assembly formation is stabilized by Parvalbumin neurons and accelerated by Somatostatin neurons. *bioRxiv*. preprint. 2021; pp. 2021–9. <https://doi.org/10.1101/2021.09.06.459211>
26. Bengio Y, Bengio S, Cloutier J. Learning a synaptic learning rule. In: *IJCNN-91-Seattle International Joint Conference on Neural Networks*. 1991;2:969. <https://doi.org/10.1109/IJCNN.1991.155621>
27. Stevenson IH, Kording KP. Inferring spike-timing-dependent plasticity from spike train data. *Advances in Neural Information Processing Systems 27 (NeurIPS)*. 2014.
28. Linderman S, Stock C, Adams R. A framework for studying synaptic plasticity with neural spike train data. *Advances in Neural Information Processing Systems 27 (NeurIPS)*. 2014.
29. Robinson BS, Berger TW, Song D. Identification of Stable Spike-Timing-Dependent Plasticity from Spiking Activity with Generalized Multilinear Modeling. *Neural Comput*. 2016;28(11):2320–2351. https://doi.org/10.1162/neco_a_00883 PMID: 27557101
30. Lappalainen J, Herpich J, Tetzlaff C. A theoretical framework to derive simple, firing-rate-dependent mathematical models of synaptic plasticity. *Front Comput Neurosci*. 2019;13:26. <https://doi.org/10.3389/fncom.2019.00026> PMID: 31133837
31. Confavreux B, Zenke F, Agnes EJ, Lillicrap T, Vogels TP. A meta-learning approach to (re) discover plasticity rules that carve a desired function into a neural network. *Advances in Neural Information Processing Systems 34 (NeurIPS)*. 2020.

32. Lindsey J, Litwin-Kumar A. Learning to learn with feedback and local plasticity. *NeurIPS* 2020. <https://proceedings.neurips.cc/paper/2020/hash/f291e10ec3263bd7724556d62e70e25d-Abstract.html>
33. Tyulmankov D, Yang GR, Abbott L. Meta-learning synaptic plasticity and memory addressing for continual familiarity detection. *Neuron*. 2022;110(3):544–557. <https://doi.org/10.1016/j.neuron.2021.11.009>
34. Metz L, Maheswaranathan N, Cheung B, Sohl-Dickstein J. Learning unsupervised learning rules. *arXiv preprint*. 2018;1804.00222.
35. Metz L, Harrison J, Freeman CD, Merchant A, Beyer L, Bradbury J, et al. Velo: Training versatile learned optimizers by scaling up. *arXiv preprint*. *arXiv:221109760*. 2022.
36. Shervani-Tabar N, Rosenbaum R. Meta-learning biologically plausible plasticity rules with random feedback pathways. *Nat Commun*. 2023;14(1):1805. <https://doi.org/10.1038/s41467-023-37562-1> PMID: 37002222
37. Miconi T. Learning to acquire novel cognitive tasks with evolution, plasticity and meta-meta-learning. In: *International Conference on Machine Learning*. PMLR; 2023. pp. 24756–74.
38. Miconi T, Stanley K, Clune J. Differentiable plasticity: training plastic neural networks with backpropagation. In: *International Conference on Machine Learning*. PMLR; 2018. pp. 3559–68.
39. Jordan J, Schmidt M, Senn W, Petrovici MA. Evolving interpretable plasticity for spiking networks. *eLife*. 2021;10:e66273. <https://doi.org/10.7554/eLife.66273> PMID: 34709176
40. Hansen N. The CMA evolution strategy: a tutorial. *arXiv preprint*. 2016;1604.00772.
41. Ramesh P, Confavreux B, Goncalves PJ, Vogels TP, Macke JH. Indistinguishable network dynamics can emerge from unlike plasticity rules. *eLife*. 2024. <https://doi.org/10.7554/eLife.94411.1>
42. Nagaraj H, Narayanan R. Plasticity manifolds and degeneracy govern circadian oscillations of neuronal intrinsic properties in the suprachiasmatic nucleus. *iScience*. 2023;26(4):106503. <https://doi.org/10.1016/j.isci.2023.106503> PMID: 37123240
43. Anirudhan A, Narayanan R. Analogous synaptic plasticity profiles emerge from disparate channel combinations. *J Neurosci*. 2015;35(11):4691–705. <https://doi.org/10.1523/jneurosci.4223-14.2015> PMID: 25788686
44. Vogels T, Froemke R, Doyon N, Gilson M, Haas J, Liu R, et al. Inhibitory synaptic plasticity: spike timing-dependence and putative network function. *Front Neural Circuits*. 2013;7:119. <https://doi.org/10.3389/fncir.2013.00119> PMID: 23882186
45. Standing L. Learning 10000 pictures. *Q J Exp Psychol*. 1973;25(2):207–22. <https://doi.org/10.1080/14640747308400340>
46. Brady TF, Konkle T, Alvarez GA, Oliva A. Visual long-term memory has a massive storage capacity for object details. *Proc Natl Acad Sci U S A*. 2008;105(38):14325–14329. <https://doi.org/10.1073/pnas.0803390105> PMID: 18787113
47. Hendin O, Horn D, Tsodyks MV. The role of inhibition in an associative memory model of the olfactory bulb. *J Comput Neurosci*. 1997;4:173–82. <https://doi.org/10.1023/a:1008895429790> PMID: 9154523
48. Hennequin G, Agnes EJ, Vogels TP. Inhibitory plasticity: balance, control, and codependence. *Annu Rev Neurosci*. 2017;40:557–59. <https://doi.org/10.1146/annurev-neuro-072116-031005> PMID: 28598717
49. Confavreux B, Ramesh P, Goncalves PJ, Macke JH, Vogels TP. Meta-learning families of plasticity rules in recurrent spiking networks using simulation-based inference. In: *Thirty-seventh Conference on Neural Information Processing Systems*. 2023.
50. Gong Z, Brunel N. Inhibitory plasticity enhances sequence storage capacity and retrieval robustness. *bioRxiv preprint*. 2024; pp. 2024–4.
51. Ocker GK, Litwin-Kumar A, Doiron B. Self-organization of microcircuits in networks of spiking neurons with plastic synapses. *PLOS Comput Biol*. 2015;11(8):1–40. <https://doi.org/10.1371/journal.pcbi.1004458>
52. Prinz AA, Bucher D, Marder E. Similar network activity from disparate circuit parameters. *Nat Neurosci*. 2004;7(12):1345–52. <https://doi.org/10.1038/nn1352> PMID: 15558066
53. Metz L, Maheswaranathan N, Nixon J, Freeman D, Sohl-Dickstein J. Understanding and correcting pathologies in the training of learned optimizers. In: *International Conference on Machine Learning*. PMLR; 2019. pp. 4556–65.
54. Vicol P, Metz L, Sohl-Dickstein J. Unbiased gradient estimation in unrolled computation graphs with persistent evolution strategies. In: *International Conference on Machine Learning*. PMLR; 2021. pp. 10553–63.

55. Zenke F, Gerstner W. Limits to high-speed simulations of spiking neural networks using general-purpose computers. *Front Neuroinform.* 2014;8:76.
<https://doi.org/10.3389/fninf.2014.00076> PMID: 25309418