DISCOVERING HETEROGENEOUS SYNAPTIC PLASTICITY RULES VIA LARGE-SCALE NEURAL EVOLUTION

Ziyuan Ye[1*], **Beichen Huang**[1*], **Yujie Wu**[2†], **Guozhang Chen**[3†], **Jibin Wu**[1,2†]
[1]Department of Data Science and Artificial Intelligence, The Hong Kong Polytechnic University
[2]Department of Computing, The Hong Kong Polytechnic University
[3]National Key Lab. for Multimedia Information Processing, School of Computer Science, Peking University
{ziyuan.ye, beichen.huang, yu-jie.wu, jibin.wu}@polyu.edu.hk,
guozhang.chen@pku.edu.hk

ABSTRACT

Synaptic plasticity is a fundamental substrate for learning and memory, where different synapse types exhibit distinct plasticity mechanisms. However, how functional behaviors emerge from heterogeneous synaptic plasticity mechanisms remains poorly understood. Here, we introduce a computational framework that harnesses Darwinian evolutionary principles to discover biologically plausible, heterogeneous synaptic plasticity rules within a biologically realistic model of the mouse primary visual cortex. Specifically, we parameterize several key factors related to synaptic plasticity, including presynaptic and postsynaptic spikes, their associated eligibility traces, and neuromodulatory signals. By integrating these factors via a truncated Taylor expansion, we construct a large-scale search space of candidate plasticity rules, with each rule containing over 2.6k optimizable parameters. Each rule is subsequently evaluated on both cross-domain visual task performance and biological validity. Leveraging a multi-objective evolutionary algorithm, we effectively navigate this high-dimensional search space to identify plasticity rules that are both biologically plausible and yield high task performance. We uncover diverse families of high-performing plasticity rules that achieve similar behavioral outcomes despite markedly different mathematical formulations, suggesting that real-world synaptic learning mechanisms may exhibit computational degeneracy. We further show that these biologically plausible rules are not only robust across network scales but also enable few-shot learning, offering a computational explanation for the emergence of innate ability.

1 INTRODUCTION

A long-standing ambition in both neuroscience and artificial intelligence has been to uncover the mechanisms by which the brain learns to generate intelligent behavior. Synaptic plasticity, as one of the most fundamental mechanisms underlying learning and memory (Bliss & Collingridge, 1993; Markram et al., 2011; McFarlan et al., 2023), has been primarily characterized through experimental studies, revealing a rich repertoire of plasticity mechanisms associated with different types of synapses (Markram et al., 1997; Bi & Poo, 1998; 2001; Woodin et al., 2003; Froemke & Dan, 2002; Kullmann & Lamsa, 2007; Kullmann et al., 2012; D'amour & Froemke, 2015). Nevertheless, the technical challenges inherent in large-scale in

[*]Equal contribution
[†]Corresponding authors

vivo experiments have limited our understanding of how these heterogeneous synaptic plasticity mechanisms cooperatively contribute to the emergence of behaviors (McFarlan et al., 2023).

In complement to experimental studies, theoretical modeling approaches aim to reveal the principles underlying synaptic plasticity mechanisms by developing mathematical frameworks that explain how neural activity patterns drive synaptic changes (Hebb, 1949; Bienenstock et al., 1982; Oja, 1982; Song et al., 2000; Pfister & Gerstner, 2006; Clopath et al., 2010; Lagzi & Fairhall, 2024; Agnes & Vogels, 2024). Data-driven inference approaches, on the other hand, seek to extract plasticity rules directly from neural recordings or behavioral data using machine learning or statistical methods (Bengio et al., 1991; Stevenson & Koerding, 2011; Robinson et al., 2016; Confavreux et al., 2020; Chen et al., 2023; Mehta et al., 2024; Bell et al., 2024; Kaleb et al., 2024; Confavreux et al., 2025a). However, none of these studies have explored the extensive landscape of heterogeneous plasticity rules in biologically grounded networks, while simultaneously accounting for biological constraints and functional objectives.

In this paper, we investigate the landscape of heterogeneous plasticity rules by asking:

> **Scientific Question**: What mathematical structure and computational principles allow heterogeneous plasticity rules to achieve functional efficacy in realistic cortical circuits?

Our approach leverages evolutionary algorithms to search a broad but interpretable candidate space of heterogeneous plasticity rules on a biologically realistic mouse primary visual cortex (V1) model, where each type of synapse is allowed to employ a distinct learning mechanism. The exploration process is evaluated under multiple objectives that balance task performance with biological plausibility. By simultaneously evaluating a diverse set of candidates within the *population*, our framework enables the study of not just a single "optimal" rule, but a family of rules that span various trade-offs. This includes rules that are potentially very simple yet surprisingly effective, as well as more complex and highly performant ones. As a result, our approach offers a way to explore the rich diversity of plasticity rules that may exist in the brain. The key contributions are:

- **Methodological advances:** (i) We construct an interpretable candidate space of plasticity rules through truncated Taylor expansion, enabling comprehensive enumeration of plasticity rules ranging from simple to complex forms, while maintaining biological interpretability (see Sec. 2.1). (ii) We introduce an evolutionary framework that allows multi-objective optimization on the constructed candidate space for discovering diverse families of plasticity rules (see Sec. 2.2). (iii) We develop several metrics covering task performance and biological constraints, enabling the discovery and systematic evaluation of plasticity rules with varying trade-offs among simplicity, efficiency, task effectiveness, and biological plausibility (see Sec. 2.3).
- **Neuroscientific insights:** (i) Using the proposed framework, we uncovered families of high-performing plasticity rules capable of producing similar memory behavioral outcomes, despite having significantly different mathematical forms (see Sec. 3.2, 4.1 and 4.2). (ii) We show that the evolved plasticity rules discovered by our framework enable few-shot learning, suggesting a potential mechanistic basis for innate abilities (see Sec. 3.3 and 4.3). (iii) We further demonstrate that the explored plasticity rules generalize well beyond their evolutionary settings, maintaining their effectiveness across networks of varying scales (see Sec. 3.4).

2 Synaptic plasticity rule exploration framework

In this section, we present the proposed multi-objective evolutionary optimization framework, which enables the discovery of biologically plausible synaptic plasticity rules by evolving a large-scale population of candidate rules, as illustrated in Fig. 1.

2.1 CONSTRUCTION OF SYNAPTIC PLASTICITY RULE CANDIDATES

Studies have shown that cortical circuits comprise diverse synaptic connections between neuronal populations, which may be governed by distinct plasticity mechanisms (McFarlan et al., 2023). To address this complexity, we build on the mouse V1 cortical model (Billeh et al., 2020; Chen et al., 2022), a biologically grounded spiking recurrent neural network comprising 17 neuron types: four excitatory subtypes from a single class (Exc) and 13 inhibitory subtypes across three classes (Pvalb, Sst, and Htr3a), distributed across six cortical layers (Fig. 1a). Specifically, as shown in Fig. 1b, we consider 289 possible synaptic connection types, corresponding to all pairwise combinations of the 17 neuron types. For each of these synaptic types, we employ a set of commonly identified neural signals in constructing synaptic plasticity rules:

$$\mathcal{T} = \{S_{\text{pre}}, S_{\text{post}}, X_{\text{pre}}, X_{\text{post}}, R\}, \tag{1}$$

where S_{pre} and S_{post} represent pre- and postsynaptic spikes, X_{pre} and X_{post} denote pre- and postsynaptic eligibility traces that capture the history of neural activity, and R corresponds to the reward prediction error trace that encodes neuromodulatory signals. These signals serve as fundamental building blocks, enabling the systematic construction of diverse plasticity rules through their additive and multiplicative combinations. It is important to note that the diversity of neuron types can significantly influence the temporal dynamics of both reward prediction error traces (Mohebi et al., 2024) and eligibility traces (Kerlin et al., 2010; He et al., 2015), resulting in heterogeneous plasticity rules across the network. To capture this biological heterogeneity, we use distinct trace decay parameters for different neuron types, as detailed below.

Reward prediction error trace model. Neuromodulators, such as dopaminergic signals, typically influence synaptic plasticity in the form of reward prediction errors (Steinberg et al., 2013; Chang et al., 2016; Corkrum et al., 2020; Gershman et al., 2024). Motivated by this finding, we model the neuromodulator signal as a reward prediction error R in our plasticity formulation, with optimizable neuron-type-specific decay time constants τ_R to capture the potentially heterogeneous temporal profiles of neuromodulation across different neural populations.

In each trial, the neuron population maintains and shares a reward prediction signal, modeled as a simple moving average of recent rewards: $\mathbf{H}_{l_i} = [r_{l_i-1}, r_{l_i-2}, \ldots, r_{l_i-N_{win}}]$, where $N_{win} = 20$ is the window size. This approach reflects the fact that reward signals can persist and influence neuronal activity even after the immediate reward has been received (Hamid et al., 2016). Specifically, the population-shared reward prediction error $\delta_R(l_i)$ for trial l_i can be formulated as:

$$\bar{r}(l_i) = \frac{1}{N_{win}} \sum_{n=1}^{N_{win}} r_{l_i-n}, \tag{2}$$

$$\delta_R(l_i) = r(l_i) - \bar{r}(l_i). \tag{3}$$

Eqs. (2)-(3) drives δ_R close to 0 when rewards become predictable. Given that neuromodulators signals exhibit distinct dynamics across different synaptic types (Huang et al., 2024b; Mohebi et al., 2024), we model this biological heterogeneity by allowing different neuronal types to maintain distinct reward prediction error traces R^m:

$$\frac{dR^m}{dt} = \begin{cases} -\frac{R^m}{\tau_R^m} + \delta_R(l_i) & \text{at trial onset } l_i + 1 \\ -\frac{R^m}{\tau_R^m} & \text{during trial execution,} \end{cases} \tag{4}$$

where τ_R^m is the time constant for neuron type m that controls the temporal dynamics of the reward signal. Note that the reward prediction error $\delta_R(l_i)$ is only given at the onset of the next trial.

Eligibility trace model. Similar to neuromodulatory signals, the eligibility traces of pre- and postsynaptic activities also exhibit heterogeneous temporal dynamics across different neuron types (He et al., 2015). In

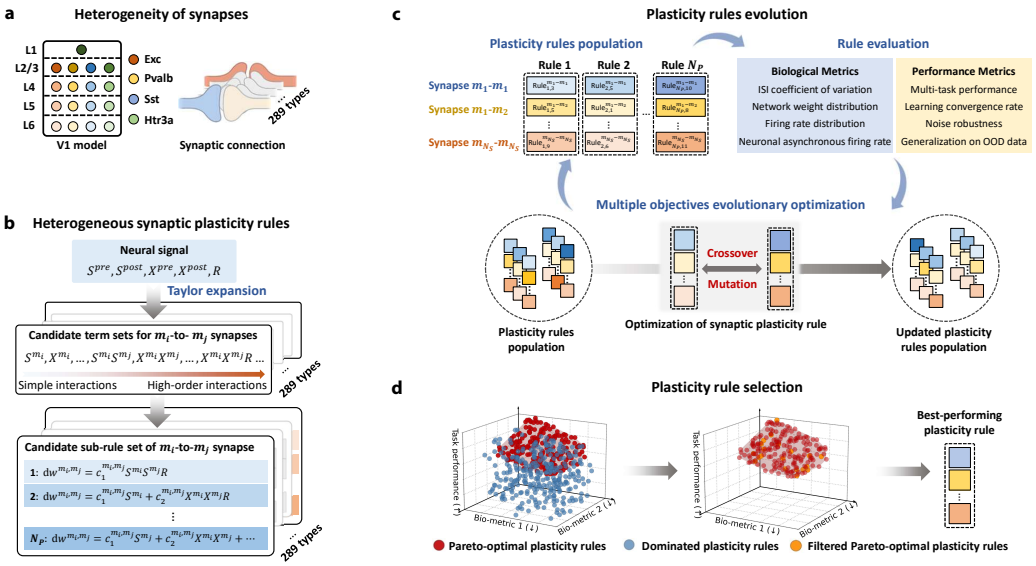

Figure 1: Overview of synaptic plasticity rule exploration framework. **a**. Our framework accounts for synaptic heterogeneity by considering one excitatory (Exc) and three classes of inhibitory synapses (Pvalb, Sst, Htr3a), each highlighted with a distinct color. Light and dark variants further indicate layer-specific subtypes within these classes. **b**. Plasticity rules are constructed from fundamental neural signals, including pre- and postsynaptic spikes, their associated eligibility traces, and the reward prediction error. These signals are systematically combined to generate candidate plasticity sub-rules for each synapse type. **c**. Plasticity rules are evolved through a multi-objective optimization process that simultaneously considers both biological and task performance metrics in a manner that parallels Darwinian evolutionary processes. **d**. The evolutionary search produces a Pareto-optimal subgroup of plasticity rules in the final generation, from which additional constraints are applied to select the filtered population and the overall best-performing rule.

our model, each neuron type maintains an eligibility trace that decays exponentially over time, with neuron-type-specific time constants to reflect this biological diversity. Specifically, the eligibility trace X^i for neuron i is denoted as

$$X^i(t+\Delta t) = X^i(t) - \frac{\Delta t}{\tau_E^m} X^i(t) + S^i(t), \tag{5}$$

where τ_E^m refers to the time constant of the eligibility trace for neuron type m, and $S^i(t)$ denotes spike occurrence at time t.

Candidate rule set. We further generate candidate synaptic plasticity terms that act as fundamental building blocks for plasticity rules. Leveraging Taylor expansion's ability to enumerate possible neural signal interactions while preserving biological interpretability, we derive these plasticity terms by expanding the basic neural signals defined in Eq. (1) up to third order. Formally:

$$\mathcal{P} = \left\{ \prod_{j=1}^{q} u_j \;\middle|\; u_j \in \{S_{\text{pre}}, S_{\text{post}}, X_{\text{pre}}, X_{\text{post}}, R\},\ q \le 3 \right\}. \tag{6}$$

In practice, redundant terms (e.g., $S_{\text{pre}}^2 = S_{\text{pre}}$ since S_{pre} is binary) and non-meaningful terms (e.g., R^2) are removed. This results in a candidate term set $\mathcal{S} = \{\mathcal{P}^k\}_{k=1}^{N_\mathcal{P}}$, where $N_\mathcal{P} = 25$, to be used in the subsequent

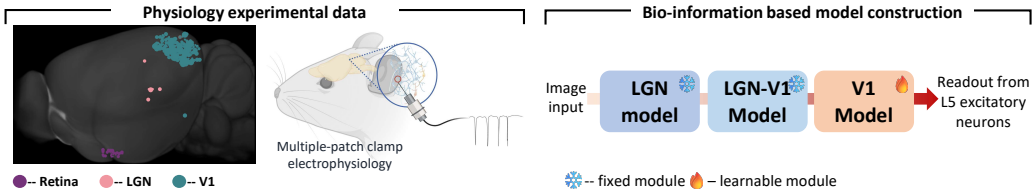

Figure 2: Overview of the biologically realistic mouse V1 model (Billeh et al., 2020; Chen et al., 2022) used in this work.

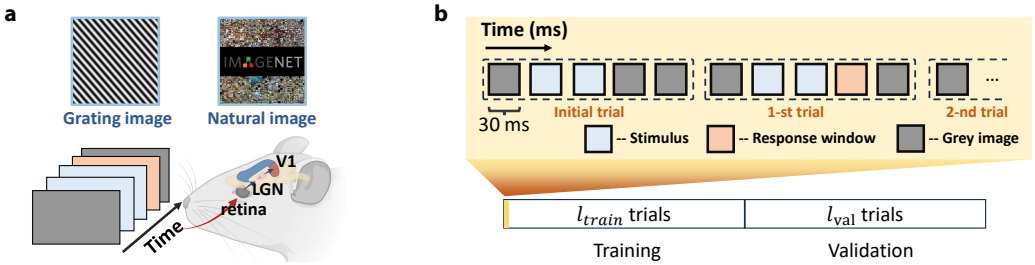

Figure 3: Overview of the multitask experimental setting. **a**. In each evaluation, stimuli are drawn either from gratings or from natural images in ImageNet (not mixed), and are presented to the V1 model. **b**. Each evaluation consists of training trials in the first half, with plasticity rules enabled, followed immediately by validation trials in the second half, with rules disabled.

construction of plasticity rules. Each plasticity rule is formulated as a weighted combination of a subset of these candidate terms. Specifically, for a given pair of presynaptic neuronal type m_{pre}, and postsynaptic neuronal type m_{post}, the candidate plasticity rules are defined as:

$$\Delta W^{(m_{\text{pre}},m_{\text{post}})} = \sum_{k=1}^{N_{\mathcal{P}}} g^k \, c^{k,(m_{\text{pre}},m_{\text{post}})} \, \mathcal{P}^k, \tag{7}$$

where $c^{k,(m_{\text{pre}},m_{\text{post}})}$ denotes the coefficient for the k-th candidate term specific to synapses connecting neuronal type m_{pre} to type m_{post}. The variable $g^k \in \{0,1\}$ is a binary selection gate that determines whether the k-th term is active, shared across all synaptic types. According to Eq. (7), the total number of optimizable parameters for each plasticity rule exceeds 2.6K, with the detailed calculation provided in Appendix D. During the weight optimization, Dale's law (Strata et al., 1999) is enforced and hard bounds on the maximum and minimum weight (Gerstner et al., 2014) are implemented to maintain biological plausibility and numerical stability during the weight update. The implementation details are also provided in Appendix D.

2.2 Plasticity rule exploration via multi-objective evolutionary optimization

All optimizable parameters of each synaptic plasticity rule are collected in a vector $\boldsymbol{\theta} = \{\boldsymbol{c}, \mathbf{g}, \boldsymbol{\tau}_E, \boldsymbol{\tau}_R\}$, which comprises the plasticity coefficients $\boldsymbol{c}$, binary gates $\mathbf{g}$, and the time constants $\boldsymbol{\tau}_E = \{\tau_E^m\}$ and $\boldsymbol{\tau}_R = \{\tau_R^m\}$ that govern the temporal dynamics of eligibility and reward traces across distinct neuron populations, respectively. This parameterization allows us to formulate synaptic plasticity rule exploration as a multi-objective optimization problem:

$$\text{Minimize } \mathcal{F}(\boldsymbol{\theta}) \in \mathbf{Y}, \boldsymbol{\theta} \in \Omega_\theta, \tag{8}$$

where $\mathcal{F}(\boldsymbol{\theta}) = (f_1(\boldsymbol{\theta}), \ldots, f_{N_o}(\boldsymbol{\theta}))$ represents N_o competing objective functions within the objective space $\mathbf{Y}$, while Ω_θ defines the feasible parameter search space. For each candidate synaptic plasticity rule defined

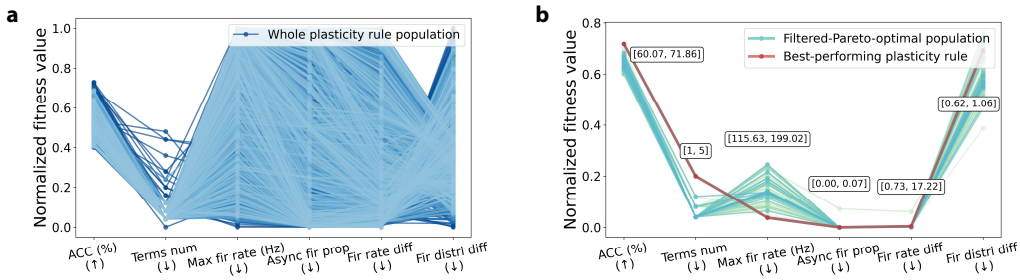

Figure 4: Performance of the filtered plasticity rules after convergence of the population. **a**. Normalized fitness value of 6 objectives for the whole plasticity rule population (N=4000). **b**. Normalized fitness value for the filtered Pareto-optimal population (N=70). The fitness values are normalized to $[0, 1]$, with original value ranges provided for reference.

by the parameter vector $\boldsymbol{\theta}$, we evaluate its performance by applying it to update the synaptic weights of the V1 model across different cognitive tasks.

As shown in Fig. 2, the input pathway to the V1 model is constructed based on biological data following Billeh et al. (2020). Visual stimuli are first processed by the fixed lateral geniculate nucleus (LGN) and LGN-V1 models before being passed to the learnable V1 model. Model performance is evaluated via readout from the layer 5 excitatory neurons in the V1 model. Model implementation is detailed in Appendix C.

Each plasticity rule undergoes a multi-objective evaluation based on its performance across different cognitive tasks. The evaluation framework incorporates biological plausibility metrics to assess whether both the V1 model's connectivity patterns and neural dynamics align with experimental observations after applying the plasticity rule. The evaluation framework also employs task performance metrics to quantify the plasticity rule's ability to solve the presented cognitive tasks. To explore this complex optimization landscape, we introduce a custom multi-objective evolutionary algorithm, whose design is described in Appendix H. During the evolutionary process, following principles analogous to Darwinian evolution, high-performing plasticity rules are retained and given opportunities to undergo crossover and mutation with other successful plasticity rules, thereby generating novel plasticity rules. Conversely, plasticity rules that exhibit poor performance on cognitive tasks or low biological plausibility are progressively eliminated from the population.

2.3 Metrics

To effectively evaluate each plasticity rule across the population, a comprehensive assessment comprising both task performance and biological metrics is employed, as illustrated in Fig. 1c. To maintain computational tractability and ensure a balanced evaluation, the following six objectives are considered in the experiments, including *cross-domain task performance, rule complexity, maximum firing rate, asynchronous firing proportion, firing rate difference and firing rate distribution difference*. We leave the implementation details in the Appendix E.

3 Experiments and Results

3.1 Experimental settings

Cross-domain task setting. Based on the mouse behavioral studies in two visual change detection experiments (Garrett et al., 2020; Siegle et al., 2021), we designed a cross-domain generalization learning (see

Mathematical formulation of plasticity rules in filtered population	Ratio	Obj-1 (↑)		Obj-2 (↓)		Obj-3 (↓)		Obj-4 (↓)		Obj-5 (↓)		Obj-6 (↓)	
		Avg	Best	Avg	Best	Avg	Best	Avg	Best	Avg	Best	Avg	Best
$\mathbf{\Delta w = S_{pre}}$	**48.57%**	65.18	70.29	0.04	0.04	157.22	127.83	0.00	0.00	1.25	0.86	0.83	**0.71**
$\mathbf{\Delta w = S_{pre} \cdot X_{post}}$	7.14%	64.54	68.43	0.04	0.04	165.10	125.96	0.00	0.00	1.16	0.87	0.91	0.86
$\Delta w = S_{post} \cdot X_{pre}$	5.71%	62.50	63.21	0.04	0.04	170.62	137.43	0.00	0.00	1.04	**0.84**	0.91	0.85
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮
$\mathbf{\Delta w = X_{post} + S_{pre} \cdot X_{pre} + S_{post} \cdot X_{pre} + X_{post}^2 + X_{post} \cdot R}$	1.43%	71.86	**71.86**	0.20	0.20	115.63	**115.63**	0.00	0.00	1.76	1.76	0.98	0.98

Table 1: Mathematical formulations of plasticity rules ranked by ratio of occurrence in the filtered population and their average performance on two visual change detection tasks. Each formulation may contain multiple rules with identical mathematical form but varying parameters. *Avg*: average fitness across individuals within each formulation; *Best*: optimal fitness within each formulation.

Fig. 3a). For the natural image change detection, all stimuli are selected from a set of 8 randomly chosen images from the ImageNet dataset (Deng et al., 2009). For the grating image change detection, static grating images are generated with orientations uniformly sampled from the range [60°, 120°] at 0.1° precision. The probability of stimulus change between consecutive presentations in two tasks is maintained at 50%. Change detection is signalled when the mean firing rate of the readout excitatory neurons during the response window exceeds a learnable, task-shared threshold φ. This experiment requires plasticity rules to endow the V1 model with working memory capabilities analogous to those required in 1-back working memory tasks (Owen et al., 2005).

Evaluation setting. Each plasticity rule is evaluated independently on each task using a two-phase protocol: 100 training trials (plasticity enabled) followed by 100 validation trials (plasticity disabled). After the population converges, rules undergo a final evaluation using 100 training and 200 testing trials per task to assess average performance across both tasks, with additional testing trials employed to reduce overfitting bias in performance assessment. Our evolution maintained a population of 4,000 plasticity rules across 150 generations, with each individual evaluated by applying its plasticity rule to a V1 model of 3,000 neurons, sampled following the strategy proposed in Chen et al. (2022). Details about the experimental setting and the search space Ω_θ can be found in Appendix E.

3.2 Discovered Plasticity Rules

As shown in Fig. 4, we identified and filtered 70 candidate plasticity rules from the Pareto-optimal population from the last generation based on overall task performance and biological validity criteria. The prevalent mathematical formulations, as demonstrated in Table 1, were ranked and selected according to their ratio of occurrence within the filtered population. We found that most plasticity rules exhibited performance consistent with biological observation. Specifically, the V1 model optimized using these plasticity rules achieves task accuracies comparable to those observed in mice (∼60% for grating change detection (Glickfeld et al., 2013) and ∼73% for natural image change detection (Garrett et al., 2020)). Results were consistent across seeds; one representative seed is shown. The full mathematical formulation table is provided in Appendix F.

In Table 1, the simple presynaptic dependent plasticity rule (*i.e.*, $\Delta w = S_{\text{pre}}$) emerged as the most prevalent formulation, comprising nearly half of the filtered rules. Several reward-free plasticity rules (*e.g.*, $\Delta w = S_{\text{pre}} \cdot X_{\text{post}}$) also succeeded in our reward-required task settings. To further analyze the explored plasticity rules, we selected the overall best-performing plasticity rule from the filtered population, which takes the mathematical form $\Delta w = X_{\text{post}} + S_{\text{pre}} \cdot X_{\text{pre}} + S_{\text{post}} \cdot X_{\text{pre}} + X_{\text{post}}^2 + X_{\text{post}} \cdot R$ (highlighted in red in

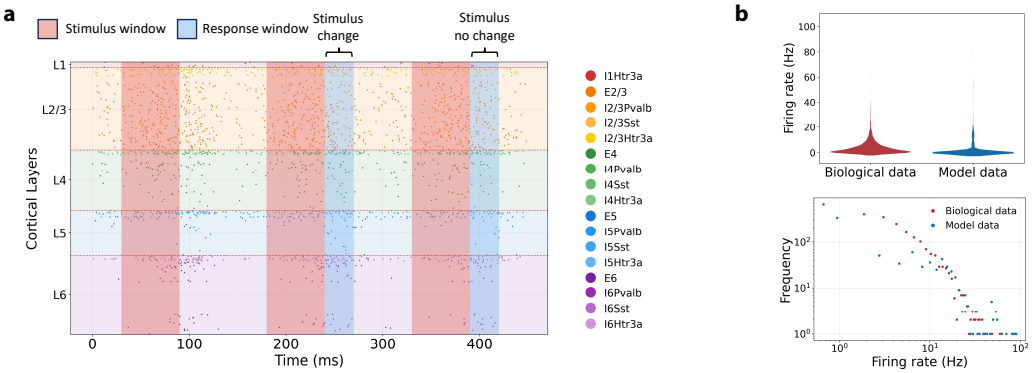

Figure 5: Visualization of spike raster plot and firing distribution of the V1 model optimized by the overall best-performing plasticity rule.

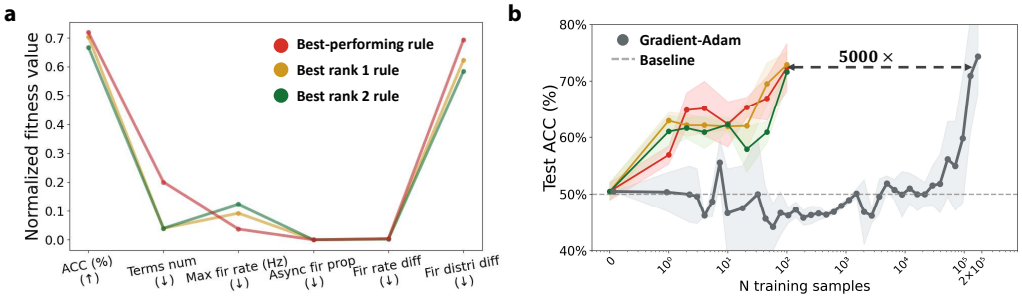

Figure 6: Comparison of performance on visual change detection of grating images between the best individual of three representative rules highlighted in Table 1 and the widely used gradient-based Adam optimizer. **a**. Performance overview. **b**. Mean test accuracy and standard deviation over 5 seeds (200 test trials each).

Fig. 4b and Table 1). The spike raster plot generated by this plasticity rule, as presented in Fig. 5a, illustrates the temporal dynamics of neural activity across cortical layers during both stimulus change and no-change conditions. Fig. 5b demonstrates the comparison of firing rate distributions between biological data (Garrett et al., 2020) and the optimized V1 model, both obtained during the grating change detection task. The analysis reveals that, consistent with the biological data, the optimized V1 model also demonstrates a heavy-tailed firing rate distribution during the grating change detection task.

3.3 Data efficiency analysis

To evaluate the efficiency of the discovered plasticity rules, we compared their performance against a conventional gradient descent (GD) baseline. Specifically, we trained the V1 model on the same visual change detection tasks using the Adam optimizer (Kingma & Ba, 2014) with a surrogate gradient function (Neftci et al., 2019a), serving as a baseline that utilizes global error signals but learns *without evolutionary priors*. To ensure a fair comparison, the biological validity metrics described in Sec. 2.3 were incorporated as regularization terms in the loss function. A detailed training setup is provided in Appendix G.

For benchmarking, we selected three representative rules from our discovered population: the best-performing rules from the top two prevalent formulations (shown in yellow and green in Table 1), and

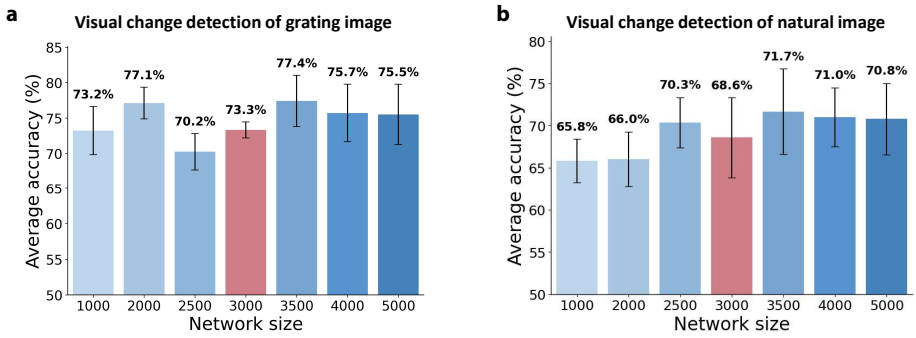

Figure 7: Evaluation of the scalability of the overall best-performing plasticity rule across network sizes. Red bars indicate test accuracy on the V1 model size used during the evolutionary search, while blue bars show test performance on different V1 model sizes. Results represent mean accuracy across 5 independent random seeds, with each evaluated with 200 testing trials. Error bars denote standard deviation across seeds.

the overall best-performing rule (shown in red). As illustrated in Fig. 6b, we observed a substantial disparity in data efficiency. All three plasticity rules achieved superior detection accuracy within the first 100 training trials, whereas the Adam baseline required nearly 5000× more samples to attain comparable performance levels. In practice, the wall-clock training time for the Adam optimizer is even greater than this factor suggests, due to the additional computational overhead introduced by backpropagation. A similar phenomenon is also observed in visual change detection on natural images (see Fig. S13). The detailed hyperparameter analysis of both Adam and Stochastic Gradient Descent (SGD) is provided in Appendix G.2.

3.4 Scalability of Discovered Plasticity Rules

To evaluate the robustness and scalability of our discovered plasticity rules, we tested the overall best-performing plasticity rule across different network scales. Specifically, we varied the V1 model size from 1,000 to 5,000 neurons using the sampling strategy described in Chen et al. (2022). As shown in Fig. 7, the overall best-performing plasticity rule maintained strong performance across different network sizes. This scalability suggests that the discovered plasticity rule captures scale-invariant principles underlying memory formation and retrieval, rather than being overfit to the 3,000-neuron V1 architecture used during the evolutionary optimization. This overall best-performing rule also demonstrates remarkable homeostatic properties under a significantly long time horizon modulation beyond its evolutionary setting (see Appendix F.7).

4 Discussion

4.1 Functionally Equivalent Plasticity Rules

Table 1 demonstrates that high task performance is not restricted to a single plasticity rule. In other words, multiple structurally different computational strategies can generate identical functional behavior, implying a degree of computational degeneracy similar to the recent findings in Confavreux et al. (2025b). This observation sheds light on why experimental synaptic plasticity studies focusing on the same type of synapse in the visual cortex sometimes report contradictory mechanisms (Lu et al., 2007; Sarihi et al., 2008; Huang et al., 2013), complementing previous explanations based on neuromodulation (McFarlan et al., 2023).

4.2 Explored plasticity rules support memory emergence

Fig. 5 and Fig. S10 reveal that sustained neural firing during delay periods between stimulus windows emerges naturally within the optimized V1 model, which is a canonical neural activity signature of working memory (Fuster & Alexander, 1971; Miyashita, 1988). Notably, *reward-free* plasticity rules relying exclusively on presynaptic activity constitute a significant proportion of the high-performing population (see Table 1) and also result in stable firing dynamics as the reward-required rule (see Fig. S11). This computational efficacy challenges the exclusivity of classical Hebbian coincidence detection and aligns with recent computational modeling in episode memory (Pang & Recanatesi, 2025). It also echoes experimental evidence from hippocampal mossy fiber synapses, where presynaptic activity has been identified as sufficient to support robust memory storage and recall (Vandael et al., 2020; Pelkey et al., 2023; Vandael & Jonas, 2024).

4.3 Evolutionary priors and the synaptic plasticity view of innate abilities

As illustrated in Fig. 6, our comparison with GD is not intended to establish algorithmic superiority, but to provide a possible explanation of the origins of biological learning efficiency. While GD operates as a general-purpose solver requiring extensive data to learn *from scratch*, evolution can embed task-related *inductive biases* into the structure and parameters of the plasticity rules. Consequently, this may enable the plasticity rules, even in a reward-free formulation, to achieve few-shot adaptation. This may explain why the sucking behavior of newborn mice requires maternal pheromone signals to be expressed (Logan et al., 2012). This perspective advances a '*synaptic plasticity view*' of innate behaviors beyond hardwired neural circuitry (Wilmer et al., 2010; Haimson & Mizrahi, 2025): innate capabilities may depend on pre-configured plasticity mechanisms that render specific behaviors accessible with minimal experience, thus bypassing the extensive trial-and-error learning required by acquired capabilities.

4.4 Limitations and Future Work

Building on these findings, our framework has several limitations that motivate future extensions. First, while several neural signals are used as building blocks of plasticity rules, incorporating additional variables such as synaptic weights would allow capturing a broader range of possible mechanisms. Second, by grounding our plasticity rules in mouse V1 data and visual processing pathways, we ensure biological fidelity to some extent but limit conclusions about generalization to non-visual sensory modalities, such as auditory data. Third, our framework captures only millisecond-scale plasticity driven by spike timing, without incorporating second-scale mechanisms such as behavioral timescale synaptic plasticity (BTSP) found in hippocampal CA1 (Bittner et al., 2017; Milstein et al., 2021). Fourth, the similarity between the optimized V1 model and biological data is evaluated from a neural representation-based aspect. Future work could incorporate recent advances in dynamics-based similarity metrics (Zhang et al., 2025) to constrain and accelerate the exploration. Addressing these limitations through expanded signal sets, cross-modal validation, multi-timescale integration, and advanced biological constraints represents important future directions.

5 Conclusion

In this paper, we present a computational framework that employs a multi-objective evolutionary algorithm to discover biologically plausible, heterogeneous plasticity rules within an experimentally grounded mouse V1 model. Our approach uncovers structurally distinct yet functionally equivalent rules, highlighting the role of computational degeneracy in neural robustness. Furthermore, our findings offer potential explanations for the origins of memory and innate abilities, suggesting that efficient learning can emerge from local synaptic dynamics. This work bridges the gap between evolutionary constraints and synaptic diversity, suggesting that the key to biological intelligence may lie in the degenerate yet robust landscape of learning rules.

ACKNOWLEDGEMENT

This work was partially supported by the National Natural Science Foundation of China (Grant No. 62306259), the Research Grants Council of the Hong Kong SAR (Grant No. C5052-23G, PolyU25216423, and PolyU15217424), The Hong Kong Polytechnic University (P0058445). This work was also supported by the National Natural Science Foundation of China (NSFC) under Grant No. 62576011.

ETHICS STATEMENT

This work adheres to the ICLR Code of Ethics and addresses several ethical considerations relevant to computational neuroscience and artificial intelligence research:

Biological data usage. Our research builds upon publicly available mouse V1 cortical data from the Allen Brain Atlas (Billeh et al., 2020), which was collected under appropriate institutional ethical approvals. We do not conduct new animal experiments and rely exclusively on established, ethically-sourced datasets.

Computational resource considerations. While our evolutionary framework requires substantial computational resources (8×A6000 GPUs), we have designed efficient algorithms based on EvoX (Huang et al., 2024a) to minimize environmental impact through batched evaluation and gradient-free optimization. The computational efficiency gained enables broader scientific exploration while reducing overall resource consumption compared to conventional approaches.

Potential applications and dual use. The discovered plasticity rules could inform the development of more efficient neuromorphic computing systems and brain-inspired AI. While these advances may have beneficial applications in assistive technologies and computational efficiency, we acknowledge that any powerful learning algorithm could potentially be misused. However, our work focuses on fundamental scientific understanding rather than specific applications, and the biological constraints inherent in our approach naturally limit potential for misuse.

Research integrity. All experimental protocols, data processing steps, and evaluation metrics are fully documented to ensure transparency and reproducibility. Our multi-objective evaluation framework explicitly incorporates biological plausibility constraints to prevent the discovery of biologically implausible but computationally expedient solutions.

REPRODUCIBILITY STATEMENT

We have made extensive efforts to ensure the reproducibility of our results:

Computational framework. Complete implementation details of our multi-objective evolutionary algorithm are provided in Appendix F, including the custom selection operators, reproduction mechanisms, and evaluation procedures. The framework is implemented using JAX and the EvoX library, with specific and hardware requirements detailed in Appendices E and H.

Model specifications. The V1 model architecture, connectivity patterns, and neuronal parameters are fully specified in Appendix A, building upon the publicly available Allen Brain Atlas data (Billeh et al., 2020) and following the architecture modification in Chen et al. (2022). All modifications to the original GLIF_3 model (Teeter et al., 2018) for plasticity compatibility are explicitly documented.

Experimental protocols. Detailed experimental settings, including stimulus generation, task protocols, evaluation metrics, and hyperparameter ranges, are provided in Appendices D and E. The visual change detection paradigm follows established behavioral protocols (Garrett et al., 2020; Siegle et al., 2021).

Baseline comparisons. The gradient-based Adam optimizer baseline implementation, including surrogate gradient functions and training procedures, is fully described in Appendix G to enable direct replication of comparative results.

Statistical analysis. All reported results include appropriate statistical measures (means, standard deviations, sample sizes) across 5 random seeds. The evolutionary search was conducted with a population of 4,000 rules over 150 generations, with final evaluations performed using 5 independent seeds and 200 test trials each.

REFERENCES

Everton J Agnes and Tim P Vogels. Co-dependent excitatory and inhibitory plasticity accounts for quick, stable and long-lasting memories in biological networks. *Nature Neuroscience*, 27(5):964–974, 2024.

Zoe Ashwood, Nicholas A Roy, Ji Hyun Bak, and Jonathan W Pillow. Inferring learning rules from animal decision-making. *Advances in Neural Information Processing Systems*, 33:3442–3453, 2020.

David Bell, Alison Duffy, and Adrienne Fairhall. Discovering plasticity rules that organize and maintain neural circuits. *Advances in Neural Information Processing Systems*, 37:40732–40751, 2024.

Y Bengio, S Bengio, and J Cloutier. Learning a synaptic learning rule. In *IJCNN-91-Seattle International Joint Conference on Neural Networks*. IEEE, 1991.

Guo-qiang Bi and Mu-ming Poo. Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. *Journal of Neuroscience*, 18(24):10464–10472, 1998.

Guo-qiang Bi and Mu-ming Poo. Synaptic modification by correlated activity: Hebb's postulate revisited. *Annual Review of Neuroscience*, 24(1):139–166, 2001.

Elie L Bienenstock, Leon N Cooper, and Paul W Munro. Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex. *Journal of Neuroscience*, 2(1):32–48, 1982.

Yazan N Billeh, Binghuang Cai, Sergey L Gratiy, Kael Dai, Ramakrishnan Iyer, Nathan W Gouwens, Reza Abbasi-Asl, Xiaoxuan Jia, Joshua H Siegle, Shawn R Olsen, et al. Systematic integration of structural and functional data into multi-scale models of mouse primary visual cortex. *Neuron*, 106(3):388–403, 2020.

Katie C Bittner, Aaron D Milstein, Christine Grienberger, Sandro Romani, and Jeffrey C Magee. Behavioral time scale synaptic plasticity underlies ca1 place fields. *Science*, 357(6355):1033–1036, 2017.

Tim VP Bliss and Graham L Collingridge. A synaptic model of memory: long-term potentiation in the hippocampus. *Nature*, 361(6407):31–39, 1993.

James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, et al. Jax: composable transformations of python+ numpy programs. 2018.

Carlos Stein Naves de Brito and Wulfram Gerstner. Learning what matters: Synaptic plasticity with invariance to second-order input correlations. *PLoS Computational Biology*, 20(2):e1011844, 2024.

György Buzsáki and Kenji Mizuseki. The log-dynamic brain: how skewed distributions affect network operations. *Nature reviews neuroscience*, 15(4):264–278, 2014.

Chun Yun Chang, Guillem R Esber, Yasmin Marrero-Garcia, Hau-Jie Yau, Antonello Bonci, and Geoffrey Schoenbaum. Brief optogenetic inhibition of dopamine neurons mimics endogenous negative reward prediction errors. *Nature Neuroscience*, 19(1):111–116, 2016.

Guozhang Chen, Franz Scherr, and Wolfgang Maass. A data-based large-scale model for primary visual cortex enables brain-like robust and versatile visual processing. *Science Advances*, 8(44):eabq7592, 2022.

Shirui Chen, Qixin Yang, and Sukbin Lim. Efficient inference of synaptic plasticity rule with gaussian process regression. *Iscience*, 26(3), 2023.

Ran Cheng and Yaochu Jin. A Competitive Swarm Optimizer for Large Scale Optimization. *IEEE Transactions on Cybernetics*, 45(2):191–204, 2015. doi: 10.1109/TCYB.2014.2322602.

Claudia Clopath, Lars Büsing, Eleni Vasilaki, and Wulfram Gerstner. Connectivity reflects coding: a model of voltage-based stdp with homeostasis. *Nature Neuroscience*, 13(3):344–352, 2010.

Basile Confavreux, Friedemann Zenke, Everton Agnes, Timothy Lillicrap, and Tim Vogels. A meta-learning approach to (re) discover plasticity rules that carve a desired function into a neural network. *Advances in Neural Information Processing Systems*, 33:16398–16408, 2020.

Basile Confavreux, Poornima Ramesh, Pedro J Goncalves, Jakob H Macke, and Tim Vogels. Meta-learning families of plasticity rules in recurrent spiking networks using simulation-based inference. *Advances in Neural Information Processing Systems*, 36:13545–13558, 2023.

Basile Confavreux, Everton J Agnes, Friedemann Zenke, Henning Sprekeler, and Tim P Vogels. Balancing complexity, performance and plausibility to meta learn plasticity rules in recurrent spiking networks. *PLoS Computational Biology*, 21(4):e1012910, 2025a.

Basile Confavreux, Zoe PM Harrington, Maciej Kania, Poornima Ramesh, Anastasia N Krouglova, Panos A Bozelos, Jakob H Macke, Andrew M Saxe, Pedro J Gonçalves, and Tim P Vogels. Memory by a thousand rules: Automated discovery of multi-type plasticity rules reveals variety & degeneracy at the heart of learning. *bioRxiv*, pp. 2025–05, 2025b.

Michelle Corkrum, Ana Covelo, Justin Lines, Luigi Bellocchio, Marc Pisansky, Kelvin Loke, Ruth Quintana, Patrick E Rothwell, Rafael Lujan, Giovanni Marsicano, et al. Dopamine-evoked synaptic regulation in the nucleus accumbens requires astrocyte activity. *Neuron*, 105(6):1036–1047, 2020.

Saskia EJ de Vries, Joshua H Siegle, and Christof Koch. Sharing neurophysiology data from the allen brain observatory. *eLife*, 12:e85550, 2023.

K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan. A fast and elitist multiobjective genetic algorithm: Nsga-ii. *IEEE Transactions on Evolutionary Computation*, 6(2):182–197, 2002. doi: 10.1109/4235.996017.

Dominique Debanne, Beat H Gähwiler, and Scott M Thompson. Long-term synaptic plasticity between pairs of individual ca3 pyramidal cells in rat hippocampal slice cultures. *The Journal of Physiology*, 507 (Pt 1):237, 1998.

Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 248–255. Ieee, 2009.

Séverine Durand, Ramakrishnan Iyer, Kenji Mizuseki, Saskia De Vries, Stefan Mihalas, and R Clay Reid. A comparison of visual response properties in the lateral geniculate nucleus and primary visual cortex of awake and anesthetized mice. *Journal of Neuroscience*, 36(48):12144–12156, 2016.

James A D'amour and Robert C Froemke. Inhibitory and excitatory spike-timing-dependent plasticity in the auditory cortex. *Neuron*, 86(2):514–528, 2015.

Samuel Eckmann, Edward James Young, and Julijana Gjorgjieva. Synapse-type-specific competitive hebbian learning forms functional recurrent networks. *Proceedings of the National Academy of Sciences*, 121 (25):e2305326121, 2024.

Sami El-Boustani, Jacque PK Ip, Vincent Breton-Provencher, Graham W Knott, Hiroyuki Okuno, Haruhiko Bito, and Mriganka Sur. Locally coordinated synaptic plasticity of visual cortex neurons in vivo. *Science*, 360(6395):1349–1354, 2018.

J.E. Fieldsend and R.M. Everson. Multi-objective optimisation in the presence of uncertainty. In *2005 IEEE Congress on Evolutionary Computation*, volume 1, pp. 243–250, 2005. doi: 10.1109/CEC.2005.1554691.

Robert C Froemke and Yang Dan. Spike-timing-dependent synaptic modification induced by natural spike trains. *Nature*, 416(6879):433–438, 2002.

Joaquin M Fuster and Garrett E Alexander. Neuron activity related to short-term memory. *Science*, 173 (3997):652–654, 1971.

Marina Garrett, Sahar Manavi, Kate Roll, Douglas R Ollerenshaw, Peter A Groblewski, Nicholas D Ponvert, Justin T Kiggins, Linzy Casal, Kyla Mace, Ali Williford, et al. Experience shapes activity dynamics and stimulus coding of vip inhibitory cells. *eLife*, 9:e50340, 2020.

Samuel J Gershman, John A Assad, Sandeep Robert Datta, Scott W Linderman, Bernardo L Sabatini, Naoshige Uchida, and Linda Wilbrecht. Explaining dopamine through prediction errors and beyond. *Nature Neuroscience*, 27(9):1645–1655, 2024.

Wulfram Gerstner, Werner M Kistler, Richard Naud, and Liam Paninski. *Neuronal dynamics: From single neurons to networks and models of cognition*. Cambridge University Press, 2014.

Abed Ghanbari, Aleksey Malyshev, Maxim Volgushev, and Ian H Stevenson. Estimating short-term synaptic plasticity from pre-and postsynaptic spiking. *PLoS Computational Biology*, 13(9):e1005738, 2017.

Lindsey L Glickfeld, Mark H Histed, and John HR Maunsell. Mouse primary visual cortex is used to detect both orientation and contrast changes. *Journal of Neuroscience*, 33(50):19416–19422, 2013.

Baruch Haimson and Adi Mizrahi. Integrating innate and learned behavior through brain circuits. *Trends in Neurosciences*, 2025.

Arif A Hamid, Jeffrey R Pettibone, Omar S Mabrouk, Vaughn L Hetrick, Robert Schmidt, Caitlin M Vander Weele, Robert T Kennedy, Brandon J Aragona, and Joshua D Berke. Mesolimbic dopamine signals the value of work. *Nature Neuroscience*, 19(1):117–126, 2016.

Kaiwen He, Marco Huertas, Su Z Hong, XiaoXiu Tie, Johannes W Hell, Harel Shouval, and Alfredo Kirkwood. Distinct eligibility traces for ltp and ltd in cortical synapses. *Neuron*, 88(3):528–538, 2015.

DO Hebb. The organization of behavior. a neuropsychological theory. 1949.

Beichen Huang, Ran Cheng, Zhuozhao Li, Yaochu Jin, and Kay Chen Tan. EvoX: A Distributed GPU-accelerated Framework for Scalable Evolutionary Computation. *IEEE Transactions on Evolutionary Computation*, 2024a. doi: 10.1109/TEVC.2024.3388550.

Jun Huang, Sylvain Crochet, Carmen Sandi, and Carl CH Petersen. Dopamine dynamics in nucleus accumbens across reward-based learning of goal-directed whisker-to-lick sensorimotor transformations in mice. *Heliyon*, 10(18), 2024b.

Shiyong Huang, Richard L Huganir, and Alfredo Kirkwood. Adrenergic gating of hebbian spike-timing-dependent plasticity in cortical interneurons. *Journal of Neuroscience*, 33(32):13171–13178, 2013.

Dongsung Huh and Terrence J Sejnowski. Gradient Descent for Spiking Neural Networks. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018. URL https://proceedings.neurips.cc/paper_files/paper/2018/file/185e65bc40581880c4f2c82958de8cfe-Paper.pdf.

Bernd Illing, Jean Ventura, Guillaume Bellec, and Wulfram Gerstner. Local plasticity rules can learn deep representations using self-supervised contrastive predictions. *Advances in Neural Information Processing Systems*, 34:30365–30379, 2021.

Yaochu Jin and J. Branke. Evolutionary optimization in uncertain environments-a survey. *IEEE Transactions on Evolutionary Computation*, 9(3):303–317, 2005. doi: 10.1109/TEVC.2005.846356.

Jakob Jordan, Maximilian Schmidt, Walter Senn, and Mihai A Petrovici. Evolving interpretable plasticity for spiking networks. *Elife*, 10:e66273, 2021.

Klara Kaleb, Barbara Feulner, Juan Gallego, and Claudia Clopath. Feedback control guides credit assignment in recurrent neural networks. *Advances in Neural Information Processing Systems*, 37:5122–5144, 2024.

Richard Kempter, Wulfram Gerstner, and J Leo Van Hemmen. Hebbian learning and spiking neurons. *Physical Review E*, 59(4):4498, 1999.

Aaron M Kerlin, Mark L Andermann, Vladimir K Berezovskii, and R Clay Reid. Broadly tuned response properties of diverse inhibitory neuron subtypes in mouse visual cortex. *Neuron*, 67(5):858–871, 2010.

Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

Dimitri M Kullmann and Karri P Lamsa. Long-term synaptic plasticity in hippocampal interneurons. *Nature Reviews Neuroscience*, 8(9):687–699, 2007.

Dimitri M Kullmann, Alexandre W Moreau, Yamina Bakiri, and Elizabeth Nicholson. Plasticity of inhibition. *Neuron*, 75(6):951–962, 2012.

Fereshteh Lagzi and Adrienne L Fairhall. Emergence of co-tuning in inhibitory neurons as a network phenomenon mediated by randomness, correlations, and homeostatic plasticity. *Science Advances*, 10(12): eadi4350, 2024.

Peter Lennie. The cost of cortical computation. *Current Biology*, 13(6):493–497, 2003.

Zhenyu Liang, Tao Jiang, Kebin Sun, and Ran Cheng. Gpu-accelerated Evolutionary Multiobjective Optimization Using Tensorized RVEA. In *Proceedings of the Genetic and Evolutionary Computation Conference*, GECCO '24, pp. 566–575, 2024. doi: 10.1145/3638529.3654223.

Sukbin Lim, Jillian L McKee, Luke Woloszyn, Yali Amit, David J Freedman, David L Sheinberg, and Nicolas Brunel. Inferring learning rules from distributions of firing rates in cortical neurons. *Nature Neuroscience*, 18(12):1804–1810, 2015.

Scott W Linderman, Christopher H Stock, and Ryan P Adams. A framework for studying synaptic plasticity with neural spike train data. *Advances in Neural Information Processing Systems*, 27, 2014.

Darren W Logan, Lisa J Brunet, William R Webb, Tyler Cutforth, John Ngai, and Lisa Stowers. Learned recognition of maternal signature odors mediates the first suckling episode in mice. *Current Biology*, 22 (21):1998–2007, 2012.

Jiang-teng Lu, Cheng-yu Li, Jian-Ping Zhao, Mu-ming Poo, and Xiao-hui Zhang. Spike-timing-dependent plasticity of neocortical excitatory synapses on inhibitory interneurons depends on target cell type. *Journal of Neuroscience*, 27(36):9711–9720, 2007.

Henry Markram, Joachim Lübke, Michael Frotscher, and Bert Sakmann. Regulation of synaptic efficacy by coincidence of postsynaptic aps and epsps. *Science*, 275(5297):213–215, 1997.

Henry Markram, Wulfram Gerstner, and Per Jesper Sjöström. A history of spike-timing-dependent plasticity. *Frontiers in Synaptic Neuroscience*, 3:4, 2011.

Amanda R McFarlan, Christina YC Chou, Airi Watanabe, Nicole Cherepacha, Maria Haddad, Hannah Owens, and P Jesper Sjöström. The plasticitome of cortical interneurons. *Nature Reviews Neuroscience*, 24(2):80–97, 2023.

Yash Mehta, Danil Tyulmankov, Adithya Rajagopalan, Glenn Turner, James Fitzgerald, and Jan Funke. Model based inference of synaptic plasticity rules. *Advances in Neural Information Processing Systems*, 37:48519–48540, 2024.

Thomas Miconi. Learning to acquire novel cognitive tasks with evolution, plasticity and meta-meta-learning. In *International Conference on Machine Learning*, pp. 24756–24774. PMLR, 2023.

Aaron D Milstein, Yiding Li, Katie C Bittner, Christine Grienberger, Ivan Soltesz, Jeffrey C Magee, and Sandro Romani. Bidirectional synaptic plasticity rapidly modifies hippocampal representations. *Elife*, 10: e73046, 2021.

Yasushi Miyashita. Neuronal correlate of visual associative long-term memory in the primate temporal cortex. *Nature*, 335(6193):817–820, 1988.

Ali Mohebi, Wei Wei, Lilian Pelattini, Kyoungjun Kim, and Joshua D Berke. Dopamine transients follow a striatal gradient of reward time horizons. *Nature Neuroscience*, 27(4):737–746, 2024.

Elias Najarro and Sebastian Risi. Meta-learning through hebbian plasticity in random networks. *Advances in Neural Information Processing Systems*, 33:20719–20731, 2020.

Emre O. Neftci, Hesham Mostafa, and Friedemann Zenke. Surrogate gradient learning in spiking neural networks: Bringing the power of gradient-based optimization to spiking neural networks. *IEEE Signal Processing Magazine*, 36(6):51–63, 2019a. doi: 10.1109/MSP.2019.2931595.

Emre O. Neftci, Hesham Mostafa, and Friedemann Zenke. Surrogate Gradient Learning in Spiking Neural Networks, May 2019b. URL `http://arxiv.org/abs/1901.09948`. arXiv:1901.09948 [cs].

Cristopher M Niell and Michael P Stryker. Highly selective receptive fields in mouse visual cortex. *Journal of Neuroscience*, 28(30):7520–7536, 2008.

Erkki Oja. Simplified neuron model as a principal component analyzer. *Journal of Mathematical Biology*, 15(3):267–273, 1982.

Adrian M Owen, Kathryn M McMillan, Angela R Laird, and Ed Bullmore. N-back working memory paradigm: A meta-analysis of normative functional neuroimaging studies. *Human brain mapping*, 25(1): 46–59, 2005.

Rich Pang and Stefano Recanatesi. A non-hebbian code for episodic memory. *Science Advances*, 11(8): eado4112, 2025.

Eunsol Park, Dika A Kuljis, Rachel A Swindell, Ajit Ray, Mo Zhu, Joseph A Christian, and Alison L Barth. Somatostatin neurons detect stimulus-reward contingencies to reduce neocortical inhibition during learning. *Cell Reports*, 44(5), 2025.

Alexandre Payeur, Jordan Guerguiev, Friedemann Zenke, Blake A Richards, and Richard Naud. Burst-dependent synaptic plasticity can coordinate learning in hierarchical circuits. *Nature Neuroscience*, 24 (7):1010–1019, 2021.

Kenneth A Pelkey, Geoffrey A Vargish, Leonardo V Pellegrini, Daniela Calvigioni, Julio Chapeton, Xiaoqing Yuan, Steven Hunt, Alex C Cummins, Mark AG Eldridge, James Pickel, et al. Evolutionary conservation of hippocampal mossy fiber synapse properties. *Neuron*, 111(23):3802–3818, 2023.

Jean-Pascal Pfister and Wulfram Gerstner. Triplets of spikes in a model of spike timing-dependent plasticity. *Journal of Neuroscience*, 26(38):9673–9682, 2006.

Brian S Robinson, Theodore W Berger, and Dong Song. Identification of stable spike-timing-dependent plasticity from spiking activity with generalized multilinear modeling. *Neural Computation*, 28(11):2320–2351, 2016.

Abdolrahman Sarihi, Bin Jiang, Alireza Komaki, Kazuhiro Sohya, Yuchio Yanagawa, and Tadaharu Tsumoto. Metabotropic glutamate receptor type 5-dependent long-term potentiation of excitatory synapses on fast-spiking gabaergic neurons in mouse visual cortex. *Journal of Neuroscience*, 28(5): 1224–1235, 2008.

Geun Hee Seol, Jokubas Ziburkus, ShiYong Huang, Lihua Song, In Tae Kim, Kogo Takamiya, Richard L Huganir, Hey-Kyoung Lee, and Alfredo Kirkwood. Neuromodulators control the polarity of spike-timing-dependent synaptic plasticity. *Neuron*, 55(6):919–929, 2007.

Navid Shervani-Tabar and Robert Rosenbaum. Meta-learning biologically plausible plasticity rules with random feedback pathways. *Nature Communications*, 14(1):1805, 2023.

Joshua H Siegle, Xiaoxuan Jia, Séverine Durand, Sam Gale, Corbett Bennett, Nile Graddis, Greggory Heller, Tamina K Ramirez, Hannah Choi, Jennifer A Luviano, et al. Survey of spiking in the mouse visual system reveals functional hierarchy. *Nature*, 592(7852):86–92, 2021.

Sen Song, Kenneth D Miller, and Larry F Abbott. Competitive hebbian learning through spike-timing-dependent synaptic plasticity. *Nature Neuroscience*, 3(9):919–926, 2000.

Elizabeth E Steinberg, Ronald Keiflin, Josiah R Boivin, Ilana B Witten, Karl Deisseroth, and Patricia H Janak. A causal link between prediction errors, dopamine neurons and learning. *Nature Neuroscience*, 16 (7):966–973, 2013.

Ian Stevenson and Konrad Koerding. Inferring spike-timing-dependent plasticity from spike train data. *Advances in Neural Information Processing Systems*, 24, 2011.

Piergiorgio Strata, Robin Harvey, et al. Dale's principle. *Brain Research Bulletin*, 50(5):349–350, 1999.

Corinne Teeter, Ramakrishnan Iyer, Vilas Menon, Nathan Gouwens, David Feng, Jim Berg, Aaron Szafer, Nicholas Cain, Hongkui Zeng, Michael Hawrylycz, et al. Generalized leaky integrate-and-fire models classify multiple neuron types. *Nature Communications*, 9(1):709, 2018.

Danil Tyulmankov, Guangyu Robert Yang, and LF Abbott. Meta-learning synaptic plasticity and memory addressing for continual familiarity detection. *Neuron*, 110(3):544–557, 2022.

David Vandael and Peter Jonas. Structure, biophysics, and circuit function of a "giant" cortical presynaptic terminal. *Science*, 383(6687):eadg6757, 2024.

David Vandael, Carolina Borges-Merjane, Xiaomin Zhang, and Peter Jonas. Short-term plasticity at hippocampal mossy fiber synapses is induced by natural activity patterns and associated with vesicle pool engram formation. *Neuron*, 107(3):509–521, 2020.

Ganchao Wei and Ian H Stevenson. Tracking fast and slow changes in synaptic weights from simultaneously observed pre-and postsynaptic spiking. *Neural Computation*, 33(10):2682–2709, 2021.

Paul J Werbos. Backpropagation through time: what it does and how to do it. *Proceedings of the IEEE*, 78 (10):1550–1560, 1990. doi: 10.1109/5.58337.

Jeremy B Wilmer, Laura Germine, Christopher F Chabris, Garga Chatterjee, Mark Williams, Eric Loken, Ken Nakayama, and Bradley Duchaine. Human face recognition ability is specific and highly heritable. *Proceedings of the National Academy of Sciences*, 107(11):5238–5241, 2010.

Melanie A Woodin, Karunesh Ganguly, and Mu-ming Poo. Coincident pre-and postsynaptic activity modifies gabaergic synapses by postsynaptic changes in cl- transporter activity. *Neuron*, 39(5):807–820, 2003.

Friedemann Zenke and Surya Ganguli. SuperSpike: Supervised Learning in Multilayer Spiking Neural Networks. *Neural Computation*, 30(6):1514–1541, June 2018. ISSN 0899-7667. doi: 10.1162/neco_a_01086. URL https://www.ncbi.nlm.nih.gov/pmc/articles/PMC6118408/.

Friedemann Zenke, Everton J Agnes, and Wulfram Gerstner. Diverse synaptic plasticity mechanisms orchestrated to form and retrieve memories in spiking neural networks. *Nature Communications*, 6(1):6922, 2015.

Li I Zhang, Huizhong W Tao, Christine E Holt, William A Harris, and Mu-ming Poo. A critical window for cooperation and competition among developing retinotectal synapses. *Nature*, 395(6697):37–44, 1998.

Shimin Zhang, Ziyuan Ye, Yinsong Yan, Zeyang Song, Yujie Wu, and Jibin Wu. Koopstd: Reliable similarity analysis between dynamical systems via approximating koopman spectrum with timescale decoupling. In *Forty-second International Conference on Machine Learning*, 2025.