

Evolving Neuronal Plasticity Rules using Cartesian Genetic Programming

Henrik D. Mettler

henrik.mettler@unibe.ch, University of Bern

Maximilian Schmidt

RIKEN Center for Brain Science, Tokyo, Japan

Walter Senn

University of Bern

Mihai A. Petrovici*

University of Bern, Heidelberg University

Jakob Jordan*

University of Bern

ABSTRACT

We formulate the search for phenomenological models of synaptic plasticity as an optimization problem. We employ Cartesian genetic programming to evolve biologically plausible human-interpretable plasticity rules that allow a given network to successfully solve tasks from specific task families. While our evolving-to-learn approach can be applied to various learning paradigms, here we illustrate its power by evolving plasticity rules that allow a network to efficiently determine the first principal component of its input distribution. We demonstrate that the evolved rules perform competitively with known hand-designed solutions. We explore how the statistical properties of the datasets used during the evolutionary search influences the form of the plasticity rules and discover new rules which are adapted to the structure of the corresponding datasets.

CCS CONCEPTS

• **Computing methodologies** → **Heuristic function construction; Rule learning.**

KEYWORDS

Synaptic plasticity, metalearning, genetic programming

ACM Reference Format:

Henrik D. Mettler, Maximilian Schmidt, Walter Senn, Mihai A. Petrovici, and Jakob Jordan. 2021. Evolving Neuronal Plasticity Rules using Cartesian Genetic Programming. In *2021 Genetic and Evolutionary Computation Conference Companion (GECCO '21 Companion)*, July 10–14, 2021, Lille, France. ACM, New York, NY, USA, 2 pages. <https://doi.org/10.1145/3449726.3459420>

1 INTRODUCTION

Changes in coupling strength between neurons in the central nervous system are believed to be central for the acquisition of new skills and memories in humans and other animals. While the microscopic biochemical processes are extraordinarily complex, phenomenological models which describe changes in the postsynaptic response to presynaptic activity have long been explored and successfully related to experimental data [2]. Furthermore, modern approaches often provide a normative view on neuron and synapse dynamics [5]. Despite these successes, the construction of new phenomenological models remains a laborious, manual process. Here

we pursue an automated approach to constructing phenomenological models of synaptic plasticity by employing genetic programming to evolve rules for synaptic that learn efficiently. We refer to this approach as “Evolving to learn” (E2L).

A simple, but useful abstraction of information processing in cortical neurons is obtained by describing a neuron’s output y_i as a linear, weighted sum of presynaptic activities x_j , followed by the application of an activation function ρ : $y_i = \rho\left(\sum_{j=1}^n w_{ij}x_j\right)$. We consider plasticity rules f that determine changes in the coupling strength w_{ij} from neuron j to i : $\Delta w_{ij} \propto f(X_{ij})$. Here X_{ij} represents a set of local variables, such as pre- and postsynaptic activity traces or synaptic weights. We formulate the search for synaptic plasticity rules as an optimization problem [1]:

$$f^* = \operatorname{argmax}_f \mathcal{F}(f, \Omega). \quad (1)$$

Here $\mathcal{F}$ represents the fitness of rule f , and Ω represents the specific experimental conditions, for example the network model and task family. The fitness measures how well a given network with plasticity rule f solves tasks from the considered task family.

Recent work has defined f as parametric function, using evolutionary strategies to optimize parameter values [3]. While this approach allows the use continuous optimization methods, the choice of the parametric form severely constraints the search space. Other authors have encoded plasticity rules using artificial neural networks [8]. While this allows the plasticity rule to take, in principle, any computable form, the macroscopic computation by ANNs is notoriously difficult to understand, limiting the interpretability of the discovered rules. In contrast, we aim to discover interpretable synaptic plasticity rules in large search spaces. We employ Cartesian genetic programming (CGP) [6] to represent and evolve plasticity rules as compact symbolic expressions. Previous work has successfully demonstrated this approach on various learning paradigms for spiking neuronal networks [4]. Here we explore the application to rate-based models. As an example, we aim to discover plasticity rules that extract the first principal component of an input data set. We use the hand-designed “Oja’s rule” [7] as a competitive baseline.

2 RESULTS

The neuronal network consists of n_{input} input units and a single output unit. Like previous work [7] we consider linear activation functions $\rho(x) = x$, hence $y = \sum_{j=1}^n w_j x_j$. A task is defined by a set $\mathcal{D}$ of M input vectors $\mathbf{x}$ sampled from a multi-dimensional Gaussian with zero mean and covariance matrix Σ . In every trial i we sample (without replacement) an input vector $\mathbf{x}^{(i)}$ from $\mathcal{D}$, compute the output activity y and update synaptic weights elementwise according to f : $\Delta w_j^{(i)} = \eta f(y^{(i)}, x_j^{(i)}, w_j^{(i-1)})$, where η is a fixed learning rate. Our goal is to discover rules which align the synaptic

*Shared senior authorship

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

GECCO '21 Companion, July 10–14, 2021, Lille, France

© 2021 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-8351-6/21/07.

<https://doi.org/10.1145/3449726.3459420>

weight vector $\mathbf{w}$ with the first principal component of the dataset ($\mathbf{PC}_0$). The set of all possible covariance matrices $\{\Sigma\}$ defines a task family $\mathcal{T}_0$. We further consider two additional task families: $\mathcal{T}_1$, where the components of $\mathbf{PC}_0$ are of approximately equal amplitude and $\mathcal{T}_2$, where $\mathbf{PC}_0$ is aligned with one of the axes. We define the fitness of a plasticity rule f for a dataset $\mathcal{D}$ as a sum of two terms, measuring the deviation of the weight vector from $\mathbf{PC}_0$, and a regularizer for its length, respectively, averaged over M trials:

$$\mathcal{F}(f, \mathcal{D}) = \frac{1}{M} \sum_{i=1}^M \left| \cos(\angle(\mathbf{w}_i, \mathbf{PC}_0)) - \alpha \left| \|\mathbf{w}_i\|_2 - 1 \right| \right|. \quad (2)$$

Here $\angle(\cdot, \cdot)$ denotes the angle between two vectors, and $\alpha > 0$ is a hyperparameter controlling the strength of the regularizer. To avoid overfitting plasticity rules to a single dataset, we define the fitness of a plasticity rule f for a task family $\mathcal{T}$ as the sampled average over K datasets from this family: $\mathcal{F}(f) = \mathbb{E}_{\mathcal{T}}[\mathcal{F}(f, \mathcal{D})]$.

When trained with tasks sampled from $\mathcal{T}_0$, 5 out of 6 evolutionary runs with different initial conditions evolve plasticity rules which allow the network to approximate $\mathbf{PC}_0$ of the respective dataset as good as or even slightly better than Oja's rule (Fig. 1a, b; $\Delta w_j^{\text{Oja}} = \eta y(x_j - w_j y)$, $\Delta w_j^{\text{lr}_1} = \eta(2y + 1 + w_j)(x_j - w_j y)$, $\Delta w_j^{\text{lr}_2} = \eta 2y(x_j - w_j y)$). These learning rules typically contain Oja's rule as a subexpression. Similarly to Oja's rule, learning rules evolved on datasets with random principle components generalize well to datasets with statistical structure (Fig. 1c, d). lr_2 slightly outperforms Oja across the investigated datasets due to a constant scaling factor which effectively increases its learning rate. These results demonstrate that our approach is able to robustly recover efficient hand-designed plasticity rules.

When evolved on structured data (task families $\mathcal{T}_1, \mathcal{T}_2$), learning rules tend to specialize and outperform their more general counterparts (Fig. 1c, $\Delta w_j^{\text{lr}_3} = \eta(-x_j)(x_j - w_j y)$; Fig. 1d, $\Delta w_j^{\text{lr}_4} = \eta(y + w_j x_j)(x_j - w_j y)$). However, evolved rules vary in their generalizability. For example, lr_3 rule does not generalize well to datasets with different statistical structure. The availability of plasticity rules as closed-form expressions helps us understand why. It is straightforward to derive the expected weight changes under lr_3 as $\mathbb{E}_D[\Delta w_j^{\text{lr}_3}] = \eta \left((w_j^2 - 1) \text{Var}[x_j] + w_j \sum_{i \neq j} w_i \text{Cov}[x_i x_j] \right)$. In two dimensions, this system of equations has only one stable fixed point with a wide basin of attraction that fully covers our assumed initialization space ($\|\mathbf{w}\|_2 = 1$) (Fig. 1e). For $\mathcal{D}$ from $\mathcal{T}_1$, the fixed point is close to $(-1, -1)$, thus approximately maximizing the fitness. For $\mathcal{D}$ from $\mathcal{T}_2$, the fixed point remains close to the diagonal, which is no longer aligned with $\mathbf{PC}_0$, thus prohibiting high fitness values (green dots in Fig. 1c, d). In contrast, learning rules evolved on datasets from $\mathcal{T}_2$, perform well on tasks from all task families (Fig. 1b, c, d), similar to Oja's rule.

3 CONCLUSION

We demonstrated that E2L can successfully discover interpretable biophysically plausible plasticity rules allowing a neuronal network to solve a well-defined task. Not only did we recover Oja's rule, but by evolving rules on datasets with specific structure we obtained variations which are adapted to the corresponding task families.

This adaptation can be viewed as an example of "overfitting" that should be avoided. However, we believe this to be an important




Figure 1: E2L discovers plasticity rules which perform PCA. a) Fitness of the best-performing individual over generations for multiple evolutionary runs with different initial conditions with covariance matrix Σ sampled from $\mathcal{T}_0$. Random initial weights for each dataset, constant across generations to make individuals from different generations comparable. b-d) Fitness per dataset for $n = 100$ datasets not used in evolutionary run, with covariance Σ sampled from $\mathcal{T}_0$ (b), $\mathcal{T}_1$ (c) and $\mathcal{T}_2$ (d). Parameters: $n = 2$, $K = 10$, $M = 1000$, $\eta = 0.01$. f is constructed from the operator set $\{+, -, *\}$, with the input set $X_{ij} = \{\{w_{ij}, x_j, y\}\}$. For implementation details see [9]. e) Phase plane of lr_3 , trained on dataset, $\text{Var}[x_1] = 1.0$, $\text{Var}[x_2] = 0.9$, $\text{Cov}[x_1 x_2] = 0.3$, with two sample trajectories converging to the fixed point. Grey indicates possible initial weights.

feature of our approach: Evolving to learn from data with specific statistical structure and thus embedding empirical priors into plasticity rules could potentially explain some of the fascinating aspects of few-shot learning and quick adaptation to novel situations displayed by biological agents. For example, it seems reasonable to expect that plasticity mechanisms driving the organization of sensory cortices are adapted to the statistical structure of their inputs, reflecting an evolutionary specialization to the ecological niche of organisms.

REFERENCES

- [1] Samy Bengio, Yoshua Bengio, Jocelyn Cloutier, and Jan Gecsei. 1992. On the optimization of a synaptic learning rule. In *Preprints Conf. Optimality in Artificial and Biological Neural Networks*, Vol. 2. Univ. of Texas.
- [2] Guo-qiang Bi and Mu-ming Poo. 1998. Synaptic Modifications in Cultured Hippocampal Neurons: Dependence on Spike Timing, Synaptic Strength, and Postsynaptic Cell Type. *Journal of Neuroscience* 18, 24 (1998), 10464–10472.
- [3] Basile Confavreux, Everton J. Agnes, Friedemann Zenke, Timothy Lillicrap, and Tim P. Vogels. 2020. A meta-learning approach to (re)discover plasticity rules that carve a desired function into a neural network. *bioRxiv* (2020).
- [4] Jakob Jordan, Maximilian Schmidt, Walter Senn, and Mihai A. Petrovici. 2020. Evolving to learn: discovering interpretable plasticity rules for spiking networks. arXiv:q-bio.NC/2005.14149
- [5] Timothy P Lillicrap, Adam Santoro, Luke Marris, Colin J Akerman, and Geoffrey Hinton. 2020. Backpropagation and the brain. *Nature Reviews Neuroscience* 21, 6 (2020), 335–346.
- [6] Julian Miller. 2019. Cartesian genetic programming: its status and future. *Genetic Programming and Evolvable Machines* 21 (08 2019).
- [7] Erkki Oja. 1982. Simplified neuron model as a principal component analyzer. *Journal of Mathematical Biology* 15, 3 (1 Nov. 1982), 267–273.
- [8] Sebastian Risi and Kenneth O Stanley. 2010. Indirectly encoding neural plasticity as a pattern of local rules. In *International Conference on Simulation of Adaptive Behavior*. Springer, 533–543.
- [9] Maximilian Schmidt and Jakob Jordan. 2020. *hal-cgp: Cartesian genetic programming in pure Python*. <https://doi.org/10.5281/zenodo.3889163>