

Order from chaos: Interplay of development and learning in recurrent networks of structured neurons

Laura Kriener^{*,†,1}, Kristin Völk^{*,1,2},

Ben v. Hünenbein¹, Federico Benitez¹, Walter Senn¹, Mihai A. Petrovici^{†,1}

* These authors contributed equally

[†]Corresponding authors {laura.kriener, mihai.petrovici}@unibe.ch)

¹Department of Physiology, University of Bern, 3012 Bern, Switzerland.

²Catlab Engineering GmbH, 82284 Grafrath, Germany.

Behavior can be described as a temporal sequence of actions driven by neural activity. To learn complex sequential patterns in neural networks, memories of past activities need to persist on significantly longer timescales than relaxation times of single-neuron activity. While recurrent networks can produce such long transients, training these networks in a biologically plausible way is challenging. One approach has been reservoir computing, where only weights from a recurrent network to a readout are learned. Other models achieve learning of recurrent synaptic weights using propagated errors. However, their biological plausibility typically suffers from issues with locality, resource allocation or parameter scales and tuning. We suggest that many of these issues can be alleviated by considering dendritic information storage and computation. By applying a fully local, always-on plasticity rule we are able to learn complex sequences in a recurrent network comprised of two populations. Importantly, our model is resource-efficient, enabling the learning of complex sequences using only a small number of neurons. We demonstrate these features in a mock-up of birdsong learning, in which our networks first learn a long, non-Markovian sequence that they can then reproduce robustly despite external disturbances.

INTRODUCTION

Adaptation to a complex world requires to evolve complex and flexible behaviors, which are themselves grounded in neural activity. In order

for a neuronal network to learn a complex sequential pattern, memories of past activities need to persist on behavioral time scales—significantly longer than the characteristic relaxation times of single-neuron activity. Typically, network recurrence is assumed to produce such long transients, but learning in recurring networks remains a challenging problem. Commonly used approaches can be divided into two groups. First, one can simply forego recurrent plasticity and only learn the weights of a “readout” layer (reservoir computing [1]). Second, the recurrent weights can be learned as well, with the most successful algorithms using propagated errors to update synaptic weights. However, their biological plausibility suffers from issues with locality (BPTT [2]), resource allocation and scaling (RTRL and its approximations [3]), or parameter scales and tuning (FORCE [4], E-PROP [5]).

We suggest that many of these issues can be alleviated through dendritic information storage and computation. Specifically, we consider the interplay between two populations of cortical neurons (for example, in motor and premotor areas) that undergo two phases of network evolution. During early development, first, neurons form a sparse, random scaffold of somato-somatic connections. During later learning, this scaffold is then used to transmit errors that induce the growth and tuning of new synapses along dendritic trees. This fully local, always-on plasticity ultimately carves out strong attractors that generate robust activation sequences. Importantly, our model makes efficient

NETWORK MODEL

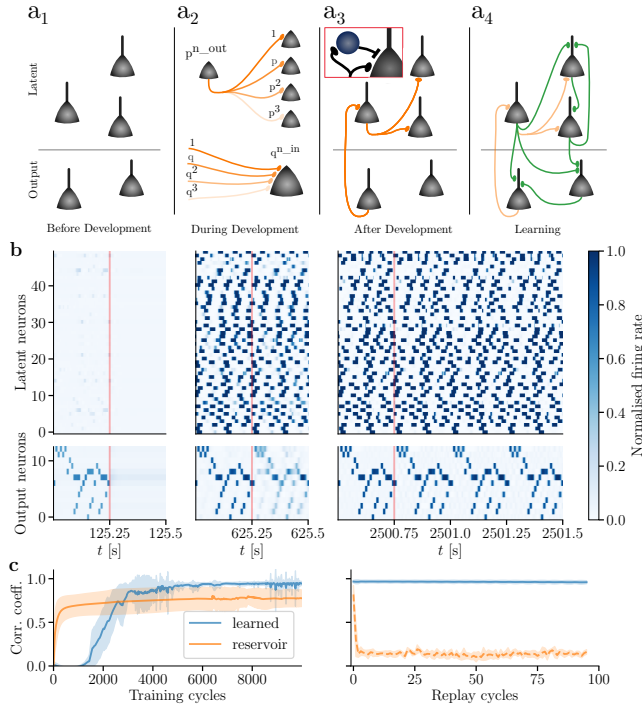


Fig. 1: **(a)** Development and learning in a two-population model of latent and output neurons. Following neurogenesis (a1), axons (orange) extend to form a sparse scaffold of somato-somatic connections between pyramidal neurons. Its structure (phenotype) is controlled by parameters p and q (genotype) (a2). Inhibition in this scaffold is mediated by interneurons (a3). Following the development phase, somato-dendritic synapses (green) evolve according to a three-factor plasticity rule (a4). **(b)** Evolution of network activity during learning. Output neurons are first nudged towards a particular target and then released in order to observe the network’s spontaneous activity (as demarcated by the red line). Snapshots during early (left), intermediate (middle) and final (right) stages of training. **(c)** Accuracy and stability of sequence replay, measured by the correlation between generated and target activity. Comparison between our model (capable of recurrent learning) and an equivalent reservoir (where only latent-to-output connections are learned). Left: Accuracy during training. Right: Accuracy after repeated replay cycles.

use of its neuronal resources, allowing the learning of complex sequences with only a small number of neurons. We demonstrate these features in a mock-up of birdsong learning, in which our networks learn a long, non-Markovian sequence (a sample of Beethoven’s “Für Elise”) that they can reproduce robustly despite severe external disturbances.

The network is comprised of structured pyramidal neurons with a somatic (u) and a dendritic (v) compartment and leaky-integrator dynamics:

$$C_{m,u}\dot{u} = -g_l u + g_{\text{den}}(v - u) + \sum_j w_{ij}^{\text{som}} r_j (E^{\text{rev}} - u),$$

$$C_{m,v}\dot{v} = -g_l v + \sum_j w_{ij}^{\text{den}} r_j,$$

where C_m is the capacitance of the compartment, g_l is the leak conductance, w the synaptic weights, E^{rev} the reversal potential and g_{den} the coupling between the compartments. To account for the electrophysiology of signal propagation, somatic compartments receive weak conductance-based input (nudging), while the effect of dendritic synapses is current-based. The output of a neuron is its firing rate $r = \varphi(u)$, which is a non-linear function of the somatic potential u .

We model the interplay between an output and a latent population of neurons. While our model is more general, in a motor context these populations would respectively be found in motor and premotor areas of cortex. During development, a sparse scaffold of somato-somatic connections is formed (Fig. 1a). These connections are static and induce random delays. Their sparsity and uniformity is controlled by the affinity of pyramidal neurons to form new output connections p and to accept new input connections q . After development, the network learns an output pattern by adapting somato-dendritic synapses, which also induce transmission delays. A teacher nudges output neurons towards a target pattern, and the somato-somatic scaffold transports this nudging signal throughout the network. The somato-dendritic synapses use these signals to adapt based a local error-correcting learning rule [6]

$$\Delta w_{ij}^{\text{den}} = \eta \left[\varphi(u_i) - \varphi \left(\frac{g_L E_L + g_D v_i}{g_L + g_{\text{den}}} \right) \right] \bar{r}_j$$

where \bar{r}_j is the lowpass-filtered presynaptic firing rate. The error used for updating the weights is locally stored in the postsynaptic neuron as the difference between the somatic firing rate $\varphi(u)$ and the firing rate that the soma would produce if it were only driven by the dendrite. During the learning phase, the somato-dendritic connections form a new scaffold that extracts increasingly

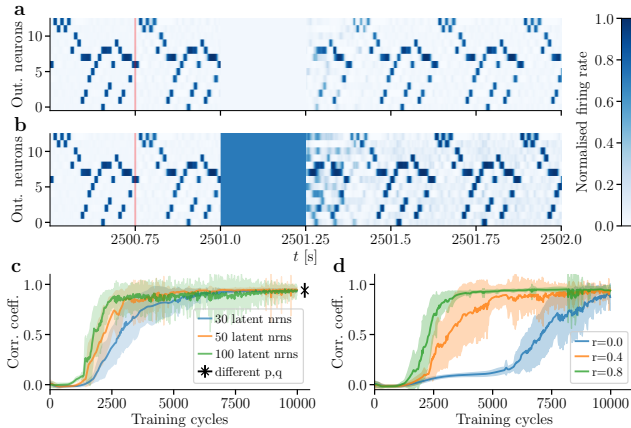


Fig. 2: **(a, b)** Robustness of pattern replay. As in Fig. 1, external nudging (teacher input) ends at the red line. A strong disturbance is introduced by voltage-clamping the output layer such that it produces very low (a) and very high (b) firing rates. After the disturbance, the network quickly recovers the correct sequence. **(c)** Efficient allocation of latent resources. Smaller latent populations require longer training times, but are still capable of producing the desired output. The results are robust across different levels of scaffold sparsity ($p \in [0.1, 0.5]$) and uniformity ($q \in [0.05, 0.25]$). **(d)** Robustness to different levels of output feedback. The strength of this feedback is determined by the output-to-latent somato-dendritic innervation ratio r (only somato-dendritic synapses can learn large weights during training). The ability of the network to learn this output feedback ($r > 0$) strongly influences its learning speed, but even in the complete absence of this capability ($r = 0$), the network can learn robust latent attractors that give rise to accurate sequence replays.

useful information from the nudging signals. This, in turn, gives rise to a robust dynamical attractor that learns to generate the correct output pattern even without the external teacher.

RESULTS

We demonstrate our model’s ability to learn complex, non-Markovian sequences by exposing it repeatedly to a sample of Beethoven’s “Für Elise” (Fig. 1b). Immediately after development, the latent population is only weakly active. During learning, dendritic connections evolve to shape the latent activity into an increasingly structured and regular attractor, allowing the network to replay the desired pattern independently of the teacher. Notably, our model is superior to a reservoir of neurons of the same size, both with regard to its ability to learn the pattern during training and to sustain it during replay (Fig. 1c).

For a biologically plausible model, robustness is of utmost importance. A model that only performs well in an idealized (disturbance-free) setting or for finely tuned and tightly controlled network parameters is not realistic. We therefore demonstrate our model’s resilience to multiple disturbances and its functionality in large ranges of the network parameter space. Despite severe temporary disruptions of the output population activity during pattern replay, the network is able to recover a correct replay of the learned pattern (Fig. 2a,b). Additionally, we demonstrate successful pattern learning across a range of network sizes and somatic scaffold parameters p and q (Fig. 2c). Finally, we demonstrate convergence to the target behavior even when the network is stripped of the capability to learn direct feedback from its output population, albeit at a slower pace (Fig. 2d).

CONCLUSION

In this work we propose a framework that enables fully local learning of sequences in recurrent networks of structured neurons. As a model of sequence learning in cortex, our method shows many advantages: It is resource-efficient, biologically plausible, and very robust. We demonstrate that our approach, compared to reservoir computation, is capable of both higher accuracy during training and sustained replay of the target pattern. Additionally, our model is not reliant on optimal initial conditions and retains its capability to learn and replay sequences across a wide range of parametrizations. Finally, we show that the attractor dynamics imprinted into our latent population via our local learning rule are able to withstand strong external disturbances. Our model starts with only a sparse, random connection scaffold generating weak and unstructured activity. We show that this is enough for local plasticity to extract useful information in order to imprint strong attractor dynamics. Thus, the combination of a developmental scaffold together with dendritic computation undergirds a robust and biologically plausible model of sequence learning in the brain.

REFERENCES

- [1] Mantas Lukoševičius and Herbert Jaeger. Reservoir computing approaches to recurrent neural

network training. *Computer science review*, 3(3):127–149, 2009.

- [2] Paul J Werbos. Backpropagation through time: what it does and how to do it. *Proceedings of the IEEE*, 78(10):1550–1560, 1990.
- [3] Owen Marschall, Kyunghyun Cho, and Cristina Savin. A unified framework of online learning algorithms for training recurrent neural networks. *The Journal of Machine Learning Research*, 21(1):5320–5353, 2020.
- [4] David Sussillo and Larry F Abbott. Generating coherent patterns of activity from chaotic neural networks. *Neuron*, 63(4):544–557, 2009.
- [5] Guillaume Bellec, Franz Scherr, Anand Subramoney, Elias Hajek, Darjan Salaj, Robert Legenstein, and Wolfgang Maass. A solution to the learning dilemma for recurrent networks of spiking neurons. *Nature communications*, 11(1):3625, 2020.
- [6] Robert Urbanczik and Walter Senn. Learning by the dendritic prediction of somatic spiking. *Neuron*, 81(3):521–528, 2014.

ACKNOWLEDGMENT

We gratefully acknowledge funding from the European Union for the Human Brain Project under the grant agreement 945539, the Manfred Stärk Foundation and the Swiss National Science Foundation (SNSF) for the “Prospective coding with pyramidal neurons”-grant (310030L_156863) and the Sinergia Grant “Neural Processing of Distinct Prediction Errors: Theory, Mechanisms & Interventions” (CR-SII5_180316). Our work has greatly benefitted from access to the Fenix Infrastructure resources, which are partially funded from the European Union’s Horizon 2020 research and innovation programme through the ICEI project under the grant agreement No. 800858.

AUTHOR CONTRIBUTIONS

KV, MAP and WS designed the model. KV wrote the initial code for the software simulation which was extended and adapted by LK. The experiments were jointly designed by LK, KV, BvH, FB, WS and MAP and conducted by LK. All authors contributed to the writing of the manuscript.

CODE AVAILABILITY

The code for the simulations is available at https://github.com/unibe-cns/temporal_sequence_learning.

COMPETING INTERESTS STATEMENT

The authors declare no competing interests.