
Learning using switching synaptic plasticity rules

Denis Turcu

Allen Institute, Seattle, WA, USA
University of Washington, Seattle, WA, USA
denis.turcu@alleninstitute.org

Jonathan Cornford

University of Leeds, Leeds, UK
J.Cornford@leeds.ac.uk

Sven Dorkenwald

Allen Institute, Seattle, WA, USA
University of Washington, Seattle, WA, USA
sven.dorkenwald@alleninstitute.org

Stefan Mihalas

Allen Institute, Seattle, WA, USA
stefanm@alleninstitute.org

Abstract

We explore networks whose synapses can independently change their governing plasticity rule during training. The MICrONS connectome data revealed that cortical synapses are well described by two states, corresponding to strong synapses having developed the spine apparatus and weak synapses lacking it. The spine apparatus, a calcium reservoir affecting synaptic dynamics, plays a significant role in plasticity and learning, though its exact function is not fully understood. Although the connectome data is static, synapses can dynamically gain or lose the spine apparatus. Here, with simplifying assumptions, we model a network whose synapses can switch between one of two learning rules: a weak, pre-post rule governing weak synapses (Hebbian-like), and a strong, credit-assignment rule governing strong synapses (backpropagation, BP). We explore such a system using recurrent neural networks (RNN) and contrast our plasticity-switching RNNs with vanilla, BP-only, RNNs. Surprisingly, we found that switching RNNs learn faster, i.e. with fewer examples, than vanilla RNNs. We also found that the recurrent weights matrix of the trained plasticity-switching RNNs is significantly more antisymmetric than the vanilla RNNs' matrix. This surprising prediction, considering Hebbian updates are nearly symmetric, deserves further investigation to reconcile with connectomic graph analysis.

1 Introduction

Traditional artificial neural networks rely on BP, which requires global error signals at all synapses—a mechanism that lacks biological plausibility [Crick, 1989, Bengio et al., 2015]. While local plasticity rules like Hebbian learning are biologically realistic, they are typically insufficient for complex tasks requiring precise credit assignment [Gerstner, 2011]. This has motivated recent interest in hybrid approaches that combine local and global learning mechanisms [Lillicrap et al., 2016, Sacramento et al., 2018, Richards and Lillicrap, 2019].

Recent analysis of the MICrONS connectome revealed a striking bimodal distribution of synaptic strengths in cortical pyramidal neurons [Dorkenwald et al., 2022]. Synapses naturally segregate into weak synapses lacking specialized structures and strong synapses with spine apparatus—a calcium reservoir that enables sophisticated plasticity [Spacek, 1985, Jedlicka et al., 2008]. We hypothesize this reflects a functional division: weak synapses use local Hebbian plasticity, while strong synapses support credit assignment through calcium-dependent signaling [Shouval et al., 2002, Graupner and Brunel, 2012]. Additionally, recent theoretical work proved that, starting from unimodal distributions, single plasticity rules can only produce final unimodal weight distributions [Pogodin et al., 2023],

Cornford et al., 2024], providing more support for the hypotheses that the bimodal distribution of weights indicates two distinct learning rules.

Motivated by these experimental and theoretical findings, we propose neural networks with switching plasticity rules based on synaptic strength. Weak synapses use Hebbian learning while strong synapses use BP, providing a biologically motivated balance between local and global learning. This approach not only explains observed bimodal distributions but also presents a novel framework for biologically realistic learning algorithms.

2 Methods

Tasks We evaluated our switching plasticity models on cognitive tasks from [Yang et al., 2019], focusing on two tasks with different difficulty levels that test temporal memory and credit assignment capabilities. The *delaygo* task (moderate difficulty) requires maintaining a spatial location during a delay period and responding after a go cue. A stimulus appears at one of eight locations, followed by a delay, then a go signal prompting response at the remembered location. The *multidelaydm* task (high difficulty) extends this by requiring comparison between sequentially presented stimuli with varying delays, demanding enhanced working memory maintenance. These tasks' temporal structure provides natural opportunities for Hebbian plasticity to operate on correlated activity patterns, while requiring precise credit assignment for successful performance [Yang et al., 2019, Driscoll et al., 2024], making them ideal for investigating local-global learning interactions in our switching plasticity framework.

Network Architecture We implemented our switching plasticity model using a recurrent neural network (RNN) architecture (Fig 1A). The network consists of input weights (W_{in}), recurrent weights (W_{rec}), and output weights (W_{out}) with corresponding biases. All weights and biases are trained using BP, except recurrent weights which follow our switching plasticity rule where individual synapses dynamically switch between local Hebbian-like updates and global BP updates based on their strength. We simulated our models in discrete time:

$$h_t = (1 - \alpha)h_{t-1} + \alpha \phi(W_{in} \cdot i_t + W_{rec} \cdot h_{t-1} + b_{rec} + \eta_{rec}), \quad (1)$$

$$o_t = W_{out} \cdot h_t + b_{out}, \quad (2)$$

where $\phi(\cdot)$ is the activation function (tanh or ReLU), $\alpha = 0.2$, and we used standard hyperparameters from previous studies [Yang et al., 2019, Driscoll et al., 2024]. We analyzed networks with $N = 64, 128, 256$, and 512 hidden units, reporting results on larger sizes due to improved robustness.

The switching mechanism operates through strength-dependent thresholds: synapses below a lower threshold use Hebbian plasticity, while those above an upper threshold use BP. This creates a division of labor where weak synapses rely on local plasticity (consistent with simple synaptic connections) while strong synapses employ sophisticated credit assignment (analogous to spine apparatus-equipped synapses). The RNN architecture allows us to isolate switching plasticity effects without confounding factors from architectural innovations.

Training We trained vanilla (BP through time) and switching (custom) models using Adam optimizer [Kingma and Ba, 2017] with a learning rate of 0.0001 for 50 epochs of 500 batches, each containing 200 concatenated task trials. We initialized weights of both model types from a uniform distribution with a scale of $\sqrt{3/N}$ (Fig 1B), that produces an initial spectrum suitable for learning (Fig 1C).

Switching model The switching model used a custom approach where recurrent weights are computed as $W_{rec} = M \odot W_{BP} + (1 - M) \odot W_{Hebb}$ using a binary & dynamic mask M . At each batch end, we updated M based on synaptic weights crossing switching thresholds. Inspired by the long timescale of gaining or losing a spine apparatus, Hebbian weights approximate long-term plasticity, unlike typical Hebbian learning in the literature used for short-term, within-trial memory [Tyulmankov et al., 2022, Aitken and Mihalas, 2023]. We set a period $P = 100$ trials, stored hidden activity in $H_{P,T,N}$ (an array with three dimensions – the number of trials P , the number of time steps T , and the number of hidden units N), and computed:

$$\Delta W_{Hebb} = [H_{P,N,1:T-1} \times H_{P,2:T,N} / (T - 1)] .mean(axis = 0), \quad (3)$$

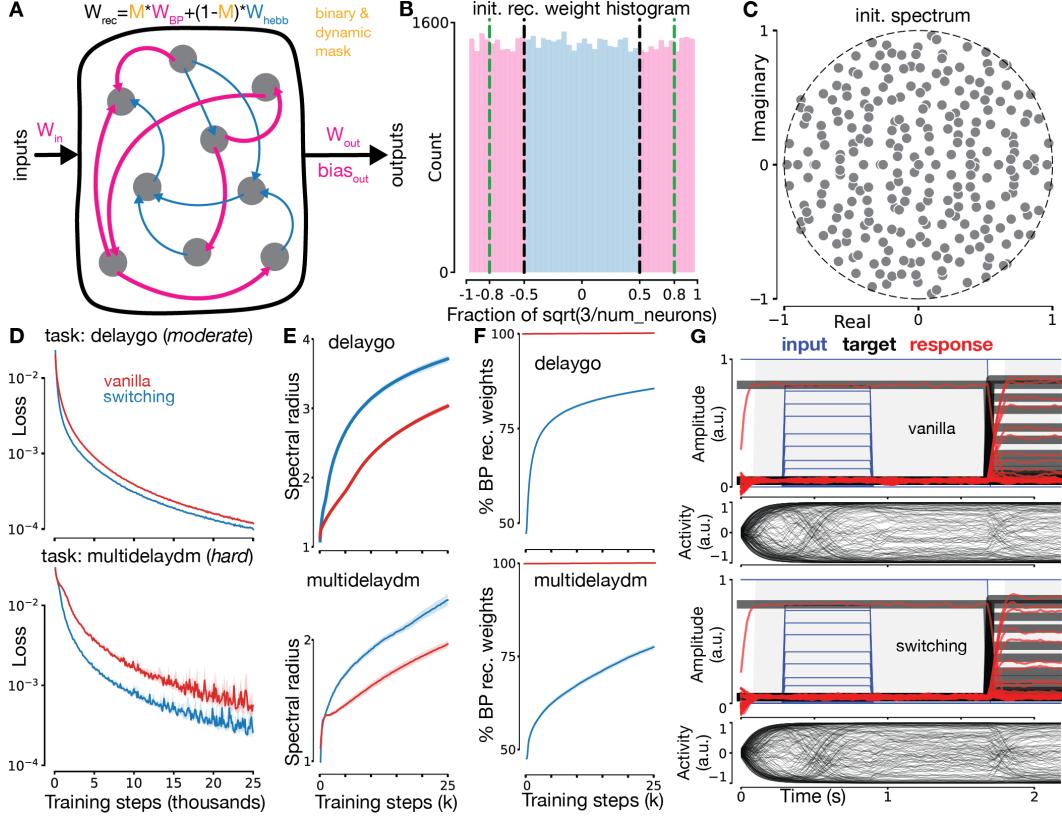


Figure 1: Model description and training metrics. **A** Schematic overview of the switching synaptic plasticity model. Input weights, recurrent biases, output weights, and output biases are trained using BP. Recurrent weights are trained using a dynamic, switching plasticity rule. The small recurrent weights are governed by local Hebbian-like learning (blue arrows), while the large recurrent weights are governed by BP (pink arrows). The binary, dynamic mask of the recurrent weights determines which synapses are in the weak or strong state at any moment in time using element-wise multiplication. **B** Example recurrent weights distribution for the switching model at initialization for a network with 256 recurrent neurons. The distribution is uniform with a scale that sets the spectral radius at initialization to 1. The switching thresholds are indicated by the dashed lines (green: hebb \rightarrow BP, black: BP \rightarrow hebb). The vanilla models start with the same distribution. **C** Spectrum of the recurrent weight matrix at initialization (both models). Dashed line indicates the unit circle. **D** Validation loss as a function of training steps across different models of size 512 (different colors), and as a function of task (top: delaygo, bottom: multidelaydm). Shaded areas represent 95% confidence interval across 20 different random initializations. Wilcoxon non-parametric one-sided test found p -value < 0.005 for the difference between the “vanilla” and “switching” between 500-16, 300 (600-13, 900) training steps for *delaygo* (*multidelaydm*). **E** Similar to **D**, but showing the spectral radius. P -value < 0.001 after 1,400 (1,300) training steps for *delaygo* (*multidelaydm*). **F** Similar to **D**, but showing the percentage of BP-trained weights. Vanilla models have 100% BP-trained weights at all times. P -value < 0.000001 at all training steps for both tasks. **G** Example trial (response amplitude and hidden activity) for *delaygo* task (top: vanilla, bottom: switching). The inputs (blue) dictate the targets (black) after the fixation period, and the model’s response (red) needs to match the targets. Gray boxes denote the time period where the loss function is computed.

followed by zeroing the diagonal and $W_{Hebb} \leftarrow \delta \cdot W_{Hebb} + \lambda \odot \Delta W_{Hebb}$ with $\delta = 0.9999$. “ \odot ” represents matrix multiplication along the last two axes. Subscripts denote appropriate transposes and time-shifts of the stored hidden activity. The learning rate λ was independently BP-trained for each synapse, concurrently with the BP recurrent weights, and was initialized from a pre-trained distribution optimized across network sizes and random initializations.

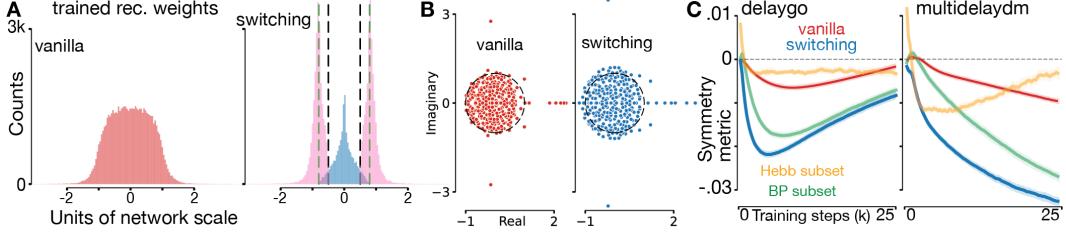


Figure 2: Post-training analysis. **A** Example synaptic weight distribution of trained models (left: vanilla, right: switching). The switching thresholds are shown on the switching model weights distribution. **B** Example spectrum of the recurrent weight matrix after training (left: vanilla, right: switching). **C** Symmetry metric as a function of training steps for different models (red, blue), and for different weight subsets of the switching models (green, yellow) for *delaygo*. The shaded area represents 95% confidence interval across 20 different random initializations. Wilcoxon non-parametric one-sided test found p -value < 0.0001 for the difference between the “vanilla” and “switching” after 750(500) training steps for *delaygo* (*multidelaydm*).

3 Results

Training and performance The switching model outperformed the vanilla in both training speed and final loss across tasks (Fig 1D), with more pronounced improvement on the harder task. Training the Hebbian learning rate λ was crucial for performance improvement, allowing adaptive adjustment of Hebbian update strength. Without training λ , the switching model did not outperform vanilla.

The switching model’s spectral radius increased faster during early training (Fig 1E), suggesting more effective learning across longer time scales. Switching models used 20 - 25% fewer BP-trained recurrent weights than vanilla models, consistent with more efficient learning through reduced sophisticated credit assignment requirements. Both models performed well qualitatively (Fig 1G). Decreasing the Hebbian update period P improved early training performance, suggesting that the interaction between weight subsets is crucial for flexible learning dynamics.

Qualitative examples show that both models performed well and had similar dynamics (Fig 1G). The switching model’s advantages stem from leveraging distinct functional weight subsets for more efficient information processing. We also found that decreasing the Hebbian update period P improved early training performance (data not shown), suggesting that interactions between the two weight subsets are crucial for flexible and adaptive learning dynamics.

Properties of switching models Post-training recurrent weight distributions reveal distinct patterns (Fig 2A). The vanilla model exhibits a unimodal weight distribution, as expected. The switching model shows a multi-modal distribution, indicating that the switching mechanism promotes formation of distinct functional weight subsets. The switching model also has a more expanded and uniform spectrum compared to vanilla models (Fig 2B).

The symmetry metric $\frac{\|W^{sym}\| - \|W^{skew}\|}{\|W^{sym}\| + \|W^{skew}\|} \in [-1, +1]$ [Hu et al., 2021] (sym and $skew$ are the symmetric and skew-symmetric parts of the matrix) shows that the switching model is significantly more anti-symmetric than the vanilla model (Fig 2C). This is surprising given that Hebbian updates are nearly symmetric ($H^T H$ in Eq. 3 is symmetric, and the minimal temporal shift only slightly affects that), and suggests that interaction between the two plasticity rules leads to more complex weight structure. This finding represents a new hypothesis about cortical connectivity patterns deserving further investigation.

4 Discussion

Our switching plasticity model demonstrates several key findings that bridge neuroscience and machine learning. Our model successfully reproduces biologically observed bimodal weight distributions while achieving better learning performance compared to vanilla models. Networks using switching plasticity showed faster learning and more efficient use of sophisticated credit assignment mechanisms, using 20-25% fewer BP-trained weights. The emergence of antisymmetric connectivity

patterns represents a novel prediction about cortical organization, while the model’s ability to leverage distinct functional weight subsets suggests fundamental advantages of hybrid learning mechanisms.

Future work should explore excitatory-inhibitory networks respecting Dale’s principle and more sophisticated spine apparatus dynamics with calcium-dependent switching. Key limitations include our binary switching mechanism that may not capture the continuous spectrum of synaptic states, the biologically implausible reliance on global error signals, and alternative explanations for the observed bimodal weight distributions, such as gating mechanisms. The efficiency of switching networks suggests promising applications in resource-constrained machine learning and neuromorphic computing, while our antisymmetric connectivity predictions provide testable hypotheses for future connectomic studies.

TL;DR We propose a biologically-inspired neural network model where synapses dynamically switch between local Hebbian learning (weak synapses) and global backpropagation (strong synapses), successfully reproducing experimentally observed bimodal weight distributions while achieving superior learning performance compared to standard backpropagation-only networks.

Acknowledgements We thank Blake Richards and Roman Pogodin for helpful discussions.

References

- Francis Crick. The recent excitement about neural networks. *Nature*, 337(6203):129–132, 1989.
- Yoshua Bengio, Dong-Hyun Lee, Jorg Bornschein, Thomas Mesnard, and Zhouhan Lin. Towards biologically plausible deep learning. *arXiv preprint arXiv:1502.04156*, 2015.
- Wulfram Gerstner. Hebbian learning and plasticity. *From neuron to cognition via computational neuroscience*, pages 0–25, 2011.
- Timothy P Lillicrap, Daniel Cownden, Douglas B Tweed, and Colin J Akerman. Random synaptic feedback weights support error backpropagation for deep learning. *Nature communications*, 7(1):13276, 2016.
- João Sacramento, Rui Ponte Costa, Yoshua Bengio, and Walter Senn. Dendritic cortical microcircuits approximate the backpropagation algorithm. *Advances in neural information processing systems*, 31, 2018.
- Blake A Richards and Timothy P Lillicrap. Dendritic solutions to the credit assignment problem. *Current opinion in neurobiology*, 54:28–36, 2019.
- Sven Dorkenwald, Nicholas L Turner, Thomas Macrina, Kisuk Lee, Ran Lu, Jingpeng Wu, Agnes L Bodor, Adam A Bleckert, Derrick Brittain, Nico Kemnitz, et al. Binary and analog variation of synapses between cortical pyramidal neurons. *Elife*, 11:e76120, 2022.
- J Spacek. Three-dimensional analysis of dendritic spines. ii. spine apparatus and other cytoplasmic components. *Anatomy and embryology*, 171(2):235–243, 1985.
- Peter Jedlicka, Andreas Vlachos, Stephan W Schwarzacher, and Thomas Deller. A role for the spine apparatus in ltp and spatial learning. *Behavioural brain research*, 192(1):12–19, 2008.
- Harel Z Shouval, Mark F Bear, and Leon N Cooper. A unified model of nmda receptor-dependent bidirectional synaptic plasticity. *Proceedings of the National Academy of Sciences*, 99(16):10831–10836, 2002.
- Michael Graupner and Nicolas Brunel. Calcium-based plasticity model explains sensitivity of synaptic changes to spike pattern, rate, and dendritic location. *Proceedings of the National Academy of Sciences*, 109(10):3991–3996, 2012.
- Roman Pogodin, Jonathan Cornford, Arna Ghosh, Gauthier Gidel, Guillaume Lajoie, and Blake Richards. Synaptic weight distributions depend on the geometry of plasticity. *arXiv preprint arXiv:2305.19394*, 2023.
- Jonathan Cornford, Roman Pogodin, Arna Ghosh, Kaiwen Sheng, Brendan A Bicknell, Olivier Codol, Beverley A Clark, Guillaume Lajoie, and Blake A Richards. Brain-like learning with exponentiated gradients. *bioRxiv*, pages 2024–10, 2024.
- Guangyu Robert Yang, Madhura R Joglekar, H Francis Song, William T Newsome, and Xiao-Jing Wang. Task representations in neural networks trained to perform many cognitive tasks. *Nature neuroscience*, 22(2):297–306, 2019.
- Laura N Driscoll, Krishna Shenoy, and David Sussillo. Flexible multitask computation in recurrent networks utilizes shared dynamical motifs. *Nature Neuroscience*, 27(7):1349–1363, 2024.
- Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2017. URL <https://arxiv.org/abs/1412.6980>.
- 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.
- Kyle Aitken and Stefan Mihalas. Neural population dynamics of computing with synaptic modulations. *Elife*, 12:e83035, 2023.
- Brian Hu, Marina E Garrett, Peter A Groblewski, Douglas R Ollerenshaw, Jiaqi Shang, Kate Roll, Sahar Manavi, Christof Koch, Shawn R Olsen, and Stefan Mihalas. Adaptation supports short-term memory in a visual change detection task. *PLoS computational biology*, 17(9):e1009246, 2021.