



Learning Efficient Backprojections Across Cortical Hierarchies in Real Time

Kevin Max^{1(✉)}, Laura Kriener¹, Garibaldi Pineda García², Thomas Nowotny²,
Walter Senn¹, and Mihai A. Petrovici¹

¹ Department of Physiology, University of Bern, Bern, Switzerland
kevin.max@unibe.ch

² School of Engineering and Informatics, University of Sussex, Brighton, UK

Abstract. Models of sensory processing and learning in physical substrates (such as the cortex) need to efficiently assign credit to synapses in all areas. In deep learning, a well-established solution is error backpropagation; this however carries several biologically implausible requirements, such as weight transport from feed-forward to feedback paths. We present Phaseless Alignment Learning (PAL), a biologically plausible approach for learning efficient feedback weights in layered cortical hierarchies. Our dynamical system enables the simultaneous learning of all weights with always-on plasticity, and exclusively utilizes information locally available at the synapses. PAL is entirely phase-free, avoiding the need for forward and backward passes or phased learning, and enables efficient error propagation across multi-layer cortical hierarchies, while maintaining bio-physically plausible signal transport and learning.

Keywords: Credit assignment · Physical computing · Network plasticity

1 Summary

Neural activity is modulated through learning, i.e., long-term adaptation of synaptic weights. However, it remains unresolved how weights are adapted across the cortex to effectively solve a given task. A key question is how to assign credit to synapses that are situated deep within a hierarchical network. In deep learning, backpropagation (BP) is the current state-of-the-art for solving this issue, and may potentially serve as an inspiration for neuroscience. Application of BP to cortical processing is however non-trivial, due to several biologically implausible requirements it entails. For example, it requires information to be buffered for use at different stages of processing. Additionally, error propagation occurs through weights that must be mirrored at synapses in different layers, resulting in the weight transport problem. Furthermore, artificial neural networks (ANNs) operate in separate forward and backward phases, with inference and learning alternating strictly.

We introduce Phaseless Alignment Learning (PAL) [4], a biologically plausible technique for learning effective top-down weights across layers in cortical hierarchies. We propose that cortical networks can learn useful backward weights by utilizing a ubiquitous resource of the brain: noise. Despite being usually treated as a disruptive factor, noise can be leveraged by the feedback pathway as an additional carrier of information for synaptic plasticity.

PAL describes a fully dynamic system that effectively addresses all of the aforementioned problems: it models the dynamics of biophysical substrates, and all computations are carried out using information **locally available** at the synapses; learning occurs in a **completely phase-less** manner; **plasticity is always-on for** all synapses, both forward and backward, at all times. Our approach is consistent with biological observations and facilitates efficient learning without the need for wake-sleep phases or other forms of phased plasticity found in many other models of cortical learning.

PAL can be applied to a broad range of models and represents an improvement over previously known biologically plausible methods of credit assignment. For instance, when compared to feedback alignment (FA), PAL can solve complex tasks with fewer neurons and more effectively learn useful latent representations. We illustrate this by conducting experiments on various classification tasks using a cortical dendrite microcircuit model [7], which leverages the complexity of neuronal morphology and is capable of prospective coding [2].

2 Theory

PAL utilises the noise found in physical neurons, as information is sent across the cortical hierarchy, see Fig. 1 (a). Neuronal dynamics are described in a rate-based coding scheme of a network with $\ell = 1 \dots N$ layers,

$$\tau \dot{\mathbf{u}}_\ell = -\mathbf{u}_\ell + \mathbf{W}_{\ell,\ell-1} \mathbf{r}_{\ell-1} + \mathbf{e}_\ell + \boldsymbol{\xi}_\ell, \quad (1)$$

with bottom-up input $\mathbf{W}_{\ell,\ell-1} \mathbf{r}_{\ell-1}$, and noise $\boldsymbol{\xi}_\ell$; the local error signal \mathbf{e}_ℓ is used to update forward weights through $\dot{\mathbf{W}}_{\ell,\ell-1} \propto \mathbf{e}_\ell \mathbf{r}_{\ell-1}^T$. Errors are passed down from higher layers through top-down synapses $\mathbf{B}_{\ell,\ell+1}$ via $\mathbf{e}_\ell = \varphi' \cdot \mathbf{B}_{\ell,\ell+1} \mathbf{e}_{\ell+1}$.

As suggested in [7], the different terms in Eq. (1) correspond to the different compartments of a pyramidal neuron, and the error is transported as the difference in firing rates of pairs of pyramidal and interneurons.

PAL learns from the noise $\boldsymbol{\xi}_\ell$ accumulated on top of a stimulus signal as it passes through the network. Backprojections are learned using high-pass-filtered rates $\hat{\mathbf{r}}_{\ell+1}$ through the rule

$$\dot{\mathbf{B}}_{\ell,\ell+1} \propto \boldsymbol{\xi}_\ell (\hat{\mathbf{r}}_{\ell+1})^T - \alpha \mathbf{B}_{\ell,\ell+1}. \quad (2)$$

By exploiting the autocorrelation properties of neuronal noise, this learning rule dynamically achieves approximate alignment $\mathbf{B}_{\ell,\ell+1} \parallel \mathbf{W}_{\ell+1,\ell}^T$ for all layers **simultaneously**, and without interrupting the learning of forward weights (see Fig. 1 (b,c)). This allows networks which implement PAL to efficiently learn

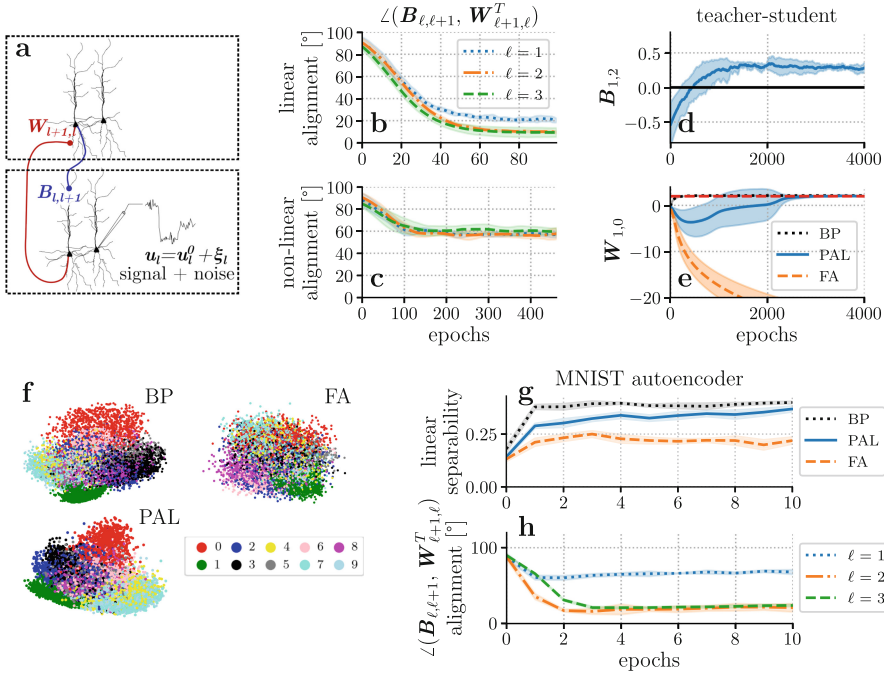


Fig. 1. PAL aligns weight updates with backpropagation in hierarchical cortical networks. (a) Cortical pyramidal cells as functional units of sensory processing and credit assignment. Bottom-up ($\mathbf{W}_{\ell+1,\ell}$) and top-down ($\mathbf{B}_{\ell,\ell+1}$) projections preferentially target different dendrites. Due to stochastic dynamics of individual neurons, noise is added to the signal. (b) We train the backward projections in a deep, dendritic microcircuit network of multi-compartment neurons with layer sizes [5-20-10-20-5] using our method PAL. All backward weights $\mathbf{B}_{\ell,\ell+1}$ are learned simultaneously, while forward weights are fixed. Forward weights are initialised s.t. neurons are activated in their linear regime. (c) Same as b, but with weights initialised in non-linear regime. (d) In a simple teacher-student task with a neuron chain [1-1-1] of dendritic microcircuits, PAL is able to flip the sign of backwards weights, which is crucial for successful reproduction of the teaching signal. (e) PAL solves teacher-student task, where feedback alignment fails. The teaching signal (red dashed) requires positive forward weights, whereas all student networks are initialised with negative $\mathbf{W}_{1,0}$. Note that PAL only learns the correct forward weights once the backwards weights have flipped sign (at epoch ~ 500). (f-h) PAL learns useful latent representations on the MNIST autoencoder task, whereas FA leads to poor feature separation. We train a network [784-200-2-200-784] using leaky-integrator neurons on the MNIST autoencoder task: (f) Shown are the activations after training in the two-neuron layer for all samples in the test set; colors encode the corresponding label. BP and PAL show improved feature separation compared to FA. (g) Linear separability of latent activation. (h) Alignment angle of top-down weights to all layers for networks trained with PAL. PAL is able adapt top-down weights while forward weights are learned at the same time. All curves show mean and standard deviation over 5 seeds.

all weights (feedforward and feedback) without phases, as opposed to many bio-inspired learning rules found in the literature (e.g., Difference Target Propagation and variants [1,3], AGREL [5,6], Equilibrium Propagation [8]).

3 Results

We have evaluated PAL on various tasks: for an excerpt of results, see Fig. 1 (b-h). Additionally, we benchmark PAL using standard tests such as the MNIST digit classification task, where the dendritic microcircuit model (of network size: [784-100-10]) achieves a final test error 3.9 ± 0.2 % using PAL and 4.7 ± 0.1 % with microcircuits with FA. We emphasize that our results were achieved through simulation of a fully dynamic, recurrent system that is biologically plausible. Weight and voltage updates were applied at every time step, and populations of multi-compartment neurons were used as a bio-plausible error transport mechanism. Our findings demonstrate that PAL can efficiently learn all weights and outperforms FA on tasks involving classification and latent space separation.

We argue that PAL can be realized both in biological and, more generally, physical components. Specifically, it capitalizes on the inherent noise present in physical systems and leverages simple filtering techniques to distinguish between signal and noise where necessary. A realization of PAL (or a variant) in physical form, whether in the cortex or on neuromorphic systems, constitutes an elegant solution to the weight transport problem, while enabling efficient learning with purely local computations.

References

1. Ernout, M., et al.: Towards scaling difference target propagation by learning back-prop targets. arXiv preprint [arXiv:2201.13415](https://arxiv.org/abs/2201.13415) (2022)
2. Haider, P., et al.: Latent equilibrium. *Adv. Neural Inf. Process. Syst.* **34**, 17839–17851 (2021)
3. Lee, D.-H., Zhang, S., Fischer, A., Bengio, Y.: Difference target propagation. In: Appice, A., Rodrigues, P.P., Santos Costa, V., Soares, C., Gama, J., Jorge, A. (eds.) *ECML PKDD 2015. LNCS (LNAI)*, vol. 9284, pp. 498–515. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-23528-8_31
4. Max, K., Kriener, L., García, G.P., Nowotny, T., Senn, W., Petrovici, M.A.: Learning efficient backprojections across cortical hierarchies in real time. arXiv preprint [arXiv:2212.10249](https://arxiv.org/abs/2212.10249) (2022)
5. Pozzi, I., et al.: A biologically plausible learning rule for deep learning in the brain. arXiv preprint [arXiv:1811.01768](https://arxiv.org/abs/1811.01768) (2018)
6. Roelfsema, P., Ooyen, A.: Attention-gated reinforcement learning of internal representations for classification. *Neural Comput.* **17**, 2176–2214 (2005)
7. Sacramento, J., et al.: Dendritic cortical microcircuits approximate the backpropagation algorithm. In: *Advances in Neural Information Processing Systems*. vol. 31 (2018)
8. Scellier, B., Bengio, Y.: Equilibrium propagation. *Front. Comput. Neurosci.* **11**, 24 (2017)