BIOLOGICALLY PLAUSIBLE LOCAL SYNAPTIC LEARNING RULES IMPLEMENT CNNS AND DENOISING AUTOENCODERS

Masataka Konishi[1] **and Keiji Miura**[*1]
[1]Department of Biosciences, Kwansei Gakuin University, Sanda, Hyogo, JAPAN

ABSTRACT

Nowadays, deep learning with the backpropagation learning rule (BP) is so popular because of its high performance in scientific computing and others. Although the neuromorphic engineering attracts attention accordingly, BP is not necessarily biologically plausible because it requires too sophisticated information that propagates for a long distance across many layers to tune synaptic weights. Recently, a biologically plausible learning rules such as feedback alignment (FA) has been proposed in order to elucidate the possible synaptic learning rules used in the brain. Some previous benchmarks cautioned not to use FA for deeply layered networks. However, the benchmarks, that mostly aimed at image recognition tasks, only used the plain neural networks without convolutional layers or did not tune hyper-parameters solely specialized for FA. As the optimal values of hyper-parameters for BP and FA can be different, a principled training for FA, desirably with convolutional layers, is required. In fact, although the convolutional neural networks (CNNs), that need weight transport, are not necessarily biologically plausible in its original form, some form of convolutional layers seems indispensable for high performance. Here we show that biologically plausible CNNs and denoising autoencoders can be successfully implemented with FA. Our benchmark demonstrated that the both models, accelerated by GPUs for large scale datasets, can perform comparably with BP. These results suggest that the variety of tasks the neuromorphic computing can solve is enriched by various types of biologically plausible image transformations and their combination principally enables more sophisticated computation such as stable diffusion models.

Keywords feedback alignment · back propagation · biologically plausibility · autoencoders · CNN · deep learning

1 Introduction

It is known that the currently-prevailing backpropagation learning rule (BP) for deep learning is not necessarily biologically plausible and not implementable in the brain in its current form [6, 1, 2, 3, 4, 8, 10]. Then, what is the synaptic learning rule the brain adopts? The simplest candidates such as the extreme learning machine (ELM) that does not learn in the middle layers and the weight perturbation (WP) that randomly searches synaptic weights by trial-and-error without explicit gradient information showed only limited performances [3]. For example, they could not solve the reversal task or the XOR problem which mice and rats can solve. Although more sophisticated learning rules such as the predictive coding are proposed, they tend to require complicated network structures [7].

The candidates we consider in this paper is the recently proposed Feedback Alignment (FA) [6] and its variant FA_EX-100% [3]. In a word, FA updates synaptic weights in the middle layers by a modified backpropagation rule where the synaptic weights vector used to compute the synaptic update is simply replaced by a random vector B. Note that the difference between BP and FA resides in how to learn the synaptic weights in the middle layers, but the learning rule for the synaptic weights to the output layers is common for FA and BP. In FA_EX-100%, all the elements of B are simply set to 1. It has been reported that FA and FA_EX-100% are robust against rule perturbations and biologically inevitable noises [3], presumably because the correct prediction error (e.g. dopaminergic signals) can calibrate in the next step as a teaching signal if a perturbation creates a deviation.

*miura@kwansei.ac.jp

Table 1: classification performance of MNIST dataset trained with BP, FA and FA_EX-100% in which we tried both fully-connected feedforward neural networks and convolutional neural networks (CNNs).

MNIST classification		
Learning Rule	fully-connected NNs	CNNs
Backpropagation (BP)	98.10%	98.95%
Feedback Alignment (FA)	97.69%	98.63%
FA_EX-100%	88.24%	90.17%

Table 2: classification performance of CIFAR10 dataset trained with BP, FA and FA_EX-100% in which we tried both fully-connected feedforward neural networks and convolutional neural networks (CNNs).

CIFAR10 classification		
Learning Rule	fully-connected NNs	CNNs
Backpropagation (BP)	44.88%	60.46%
Feedback Alignment (FA)	41.24%	51.66%
FA_EX-100%	35.71%	27.48%

The candidates of the learning rule the brain implements can be narrowed down by comparing their performances. There, as the performances can heavily depend on the tasks imposed, it is very important to test different tasks. Although previous benchmarks of FA and its variants demonstrated fairly good performances for image recognition tasks, we believe that there are still two missing parts to attain biological plausibility as described next [6, 8].

First, the convolutional layer, that is essential for CNN's high performance, was not used [1, 2]. In a typical CNN, the same 3×3 filter is shared by different neurons in the next layer, which requires a weight transportation. Thus, apparently, the wiring for CNNs seems too complicated to be implemented in the brain [1, 2].

Second, the tasks that have high dimensional outputs, such as autoencoders, should also be tested. Such an "image transformation task" is more difficult and can be a good benchmark. Furthermore, the variety of tasks the brain can solve is enriched with biologically plausible image transformations. Apparently more difficult task can be occasionally decomposed into several image transformation tasks. For example, image generation by stable diffusion can be regarded as a repetition of denoising autocoders. Therefore it is important to include denoising autoencoders in a benchmark.

A previous work found that the "principled" hyper-parameter tuning is essential for solving more complicated tasks, although only direct feedback alighment (DFA), that is a variant of FA, was used as a learning rule there [5, 2, 4, 9]. In this paper, we tried similar principled hyper-parameter tuning for the original FA as a learning rule [6, 3]. We first demonstrate that the FA-based CNN can be successfully implemented in a biologically plausible way. We next show that the FA-based denoising autoencoder for MNIST dataset works comparably with BP. The both models were boosted by GPUs for large scale datasets.

2 Materials and Methods

All the numerical simulation for the multi-layer feedforward neural networks in this paper was performed by custom-made code using Python 3.9.13 and PyTorch 1.11.0+cu113, which implemented various learning rules for CNNs and auto-encoders. All the Python codes used for this paper are available publicly at `https://github.com/keiji-miura/FA-based_CNN_Autoencoder`.

2.1 Learning Rules: BP, FA and FA_EX-100%

BP is a conventional learning rule. BP updates synaptic weights by the usual backpropagation rule (both in the hidden and output layers). The weight vector is updated according to the gradient vector to minimize the loss function. The other learning rules are described in what follows, but, their difference only resides in the weight update rule in hidden layers. That is, the weight update rule in the output layer is common for all the learning rules and, thus, it is the same as that for BP.

FA updates synaptic weights in the hidden layers by the modified backpropagation rule in which the synaptic weights matrix W in the next layer, is replaced by a fixed ([-1,1]-uniformly) random vector B [6, 3].

FA should finish learning successfully, for example, if W gets close to B by the end of learning, consistently with the learning assumption ($W = B$).

FA_EX-100% updates synaptic weights in the same way as FA, but all the elements of B matrix is set to 1 [3]. In this learning rule, all the synapses are equal in the sense you do not need to assign random and heterogenous B values for different synapses. This simple rule may further enhance biological plausibility.

2.2 Implimenting FA with GPUs

FA was implemented by modifying the backward function of Linear class for linear layers in Pytorch. We believe that this rather minimal modification of original Pytorch function enables us the efficient usage of GPUs. We used the same codes for FA_EX-100%.

2.3 FA-based Convolutional Layers with Truly Local Connections

We have solely used 5×5 filtering layer for CNNs. The CNN with feedback alignment was implemented by representing the $M \times N$ sparse weight matrix by the dense $25 \times N$ matrix, where M and N represent the number of neurons in the pre and post layers, respectively. Note that in our feedback alignment-based CNN, the size of a single filter is $25 \times N$, instead of 25×1 for the conventional BP in which a 5×5 is shared by all the post layer neurons. In this paper, in the spirit of biological plausibility [1], we avoided to share the 5×5 weight and each post neuron has its own 5×5 filter. There, again, Pytorch Linear class was modified to create a hand-made class for the convolutional layer based on FA. Of course this implementation enabled us to utilize GPUs efficiently.

3 Results

3.1 FA-based CNNs with Biologically Plausible Local Connections

We compared the classification performance of MNIST or CIFAR10 dataset with backpropagation (BP), feedback alignment (FA) and FA_EX-100% in which we tried both fully-connected feedforward neural networks and convolutional neural networks (CNNs) (Table1 and 2). The difference in learning rules was summarized in Materials and Methods. The fully-connected feedforward neural network, used for both MNIST and CIFAR10, had a single hidden layer with 128 neurons followed by ReLU.

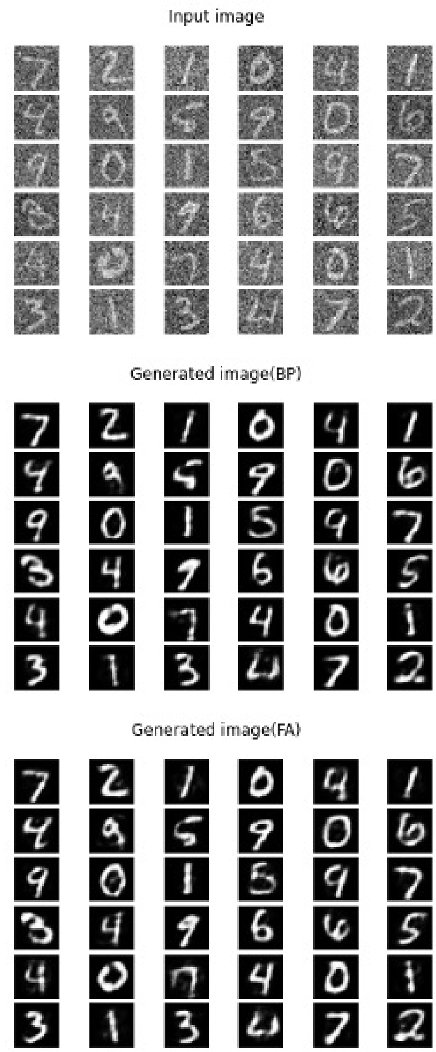

Figure 1: Reconstruction of MNIST images by the fully-connected feedforward neural network with a single hidden layer with 200 neurons trained with backpropagation (BP) or feedback alignment (FA).

We have tuned the structures of our FA-based CNNs and successfully improved their performances to the levels comparable with those of BP-based CNNs. During that process, we observed that the structures optimized for DFA in [5] also worked for FA basically, although those structures looked rather uncommon for the conventional BP-based CNNs. The CNN structure we used for MNIST was "Input - Conv2d(6 filters, 5x5, tanh, padding) - Maxpooling(2x2, strides=2) - Conv2d(16 filters, 5x5, tanh, padding) - Maxpooling(2x2, strides=2) - Fc(128, tanh) - Fc(10, softmax)". The CNN structure for CIFAR10 was "Input - Conv2d(32 filters, 5x5, tanh, padding) - Avgpooling(3x3, strides=2) - Conv2d(64 filters, 5x5, tanh, padding) - Avgpooling(3x3, strides=2) - Conv2d(64 filters, 5x5, tanh, padding) - Avgpooling(3x3, strides=2) - Fc(128, tanh) - Fc(10, softmax)".

3.2 FA-based Denoising Autoencoders

We next implemented the FA-based denoising autoencoder. Both the BP-based (=conventional) and FA-based autoencoders had a single hidden layer with 200 neurons followed by ReLU. One example of the recovered data for the noise amplitude $\sigma = 100$ shown in Figure 1 demonstrates that not only BP-based but also FA-based denoising autoencoder worked fairly well. Note that the original images had from 0 to 255 pixel values, on which the normal noise was added independently on each pixel. The reconstruction error in binary-crossentropy was summarized in Table 3. We did not show the performance of FA_EX-100% as its performance was much lower and it failed to reconstruct.

Table 3: Mean squared reconstruction error of MNIST dataset with backpropagation (BP), feedback alignment (FA). The fully-connected feedforward neural network had a single hidden layer with 200 neurons followed by ReLU.

MNIST reconstruction error			
Sigma (noise amplitude)	0	100	255
Backpropagation (BP)	0.0816	0.105	0.157
Feedback Alignment (FA)	0.0880	0.112	0.164

4 Discussion

Although the overall performance for BP is higher than that of FA as expected, the FA-based CNNs showed the performance that is comparable to the one for BP-based CNNs. The contribution of this paper resides in the GPU-accelerated large scale benchmarks of truly localized CNNs and denoising autoencoders, which are successfully trained with FA. These results suggest that the variety of tasks the neuromorphic computing can solve is enriched by various types of biologically plausible image transformations and their combination principally enables more sophisticated computation such as stable diffusion models.

Acknowledgments

K. Miura is partially supported by JSPS KAKENHI Grant Numbers JP22K19816 and JP22H02364.

References

[1] Sergey Bartunov, Adam Santoro, Blake Richards, Luke Marris, Geoffrey E Hinton, and Timothy Lillicrap. Assessing the scalability of biologically-motivated deep learning algorithms and architectures. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018.

[2] Brian Crafton, Abhinav Parihar, Evan Gebhardt, and Arijit Raychowdhury. Direct feedback alignment with sparse connections for local learning. *Front. Neurosci.*, 13:525, 2019.

[3] Masataka Konishi, Kei M Igarashi, and Keiji Miura. Biologically plausible local synaptic learning rules robustly implement deep supervised learning. *Front. Neurosci.*, 17:1160899, 2023.

[4] Julien Launay, Iacopo Poli, François Boniface, and Florent Krzakala. Direct feedback alignment scales to modern deep learning tasks and architectures. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 9346–9360. Curran Associates, Inc., 2020.

[5] Julien Launay, Iacopo Poli, and Florent Krzakala. Principled training of neural networks with direct feedback alignment. In *Preprint*, 2019.

[6] Timothy P. Lillicrap, Daniel Cownden, Douglas B. Tweed, and Colin J. Akerman. Random synaptic feedback weights support error backpropagation for deep learning. *Nat. Commun.*, 7:13276, 2016.

[7] Beren Millidge, Tommaso Salvatori, Yuhang Song, Rafal Bogacz, and Thomas Lukasiewicz. Predictive coding: Towards a future of deep learning beyond backpropagation? In Luc De Raedt, editor, *Proceedings of the 31st International Joint Conference on Artificial Intelligence and the 25th European Conference on Artificial Intelligence, IJCAI-ECAI 2022, Survey Track, Vienna, Austria, July 23-29, 2022*, pages 5538–5545. IJCAI/AAAI Press, 2022.

[8] Arild Nø kland. Direct feedback alignment provides learning in deep neural networks. In D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc., 2016.

[9] Ruben Ohana, Hamlet Jesse Medina Ruiz, Julien Launay, Alessandro Cappelli, Iacopo Poli, Liva Ralaivola, and Alain Rakotomamonjy. Photonic differential privacy with direct feedback alignment. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, 2021.

[10] Maria Refinetti, Stéphane d'Ascoli, Ruben Ohana, and Sebastian Goldt. Align, then memorise: the dynamics of learning with feedback alignment. *Journal of Statistical Mechanics: Theory and Experiment*, 2022(11):114002, 2022.