Homeostatic learning for classical artificial neural networks

Kristína Malinovská and Johanka Jakubove

Centre for Cognitie Science, DAI FMPI,
Comenius University in Bratislava, Slovakia
Mlynská dolina, 84248 Bratislava
Email: kristina.malinovska@fmph.uniba.sk

Abstract

Standard methods for training of the artificial neural networks (ANN) include various techniques for improving their success rates that are more mechanistic than brain-inspired. Brain-inspired learning makes use of elegant principles, such as homeostasis, for making biological neural networks learn more efficiently. The BCM rule is the most parsimonious model of neuronal learning in the brain today. An abstraction of this rule can potentially improve the efficiency of error-backpropagation-based learning in classical ANN. To test this hypothesis, we propose an adaptation of the classical multilayer perceptron to include a local learning rate component for each weight that decreases gradually through the training based on the activation of presynaptic and postsynaptic neurons.

1 Introduction

The most important principle in both biological and artificial neural networks is that they store their knowledge in patterns of synaptic weights, that modulated the efficiency of connections between the individual neurons (O'Reilly et al., 2024). Within artificial neural networks learning we distinguish two main paradigms based on what information is the network using to gain knowledge, the so called unsupervised learning - without any teaching signal - which is also implemented in the brain and the so called supervised learning which implements an input-output mapping function between the data samples and their labels (arbitrary categories given in the data).

1.1 Biologically motivated unsupervised learning

The general principle on which the synapses - weighted connections - between neurons are promoted or diminished has been postulated by famous Canadian psychologist Donald O. Hebb in 1940's (Hebb, 2005). The Hebb's rule states that "neurons that fire together, wire together". This also includes the causal link between the cells, the firing of the so called presynaptic cell has to precede (or trigger) the firing of the postsynaptic cell. This general principle indeed applies to learning in the brain, however, it has one particular cavity. If we formulate the change of the synaptic weight between these two neurons as:

$$\Delta w_{xy} = xy, \tag{1}$$

where x and y stand for activation (rate-coded approximation of the neuron's firing) of the presynaptic and the postsynaptic neuron respectively. We can observe that the weight w_{xy} can be increased until infinity just by correlated firing of these two cells.

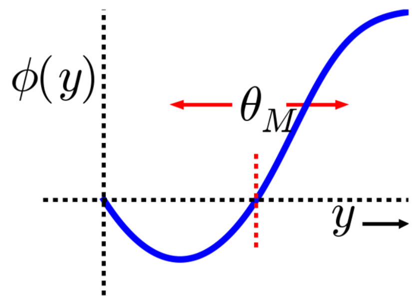

Fig. 1: The BCM modification function. Figure taken from (Blais and Cooper, 2008)

In reality, the weight adaptation needs to be controlled by a *homeostatic mechanism*, preserving balanced weight values and maintaining effective functioning of the synaptic weights. The best basic model of bio-realistic synaptic plasticity developed so far is the BCM model (Bienenstock–Cooper–Munro, Bienenstock et al., 1982). This model introduces a modification function of synaptic plasticity that is applied to Hebbian principle according to which in the extrema of the potentiation (strengthening) or depression of the synaptic weight is limited and thus modulated using a sliding modification threshold. The BCM theory proposes the BCM modification function ϕ of the synaptic plasticity update as follows:

$$\phi = y(y - \theta_M) \tag{2}$$

where θ_M is the sliding threshold adapted according to:

$$\theta_M = \lambda \langle y^2 \rangle_\tau, \tag{3}$$

where $\langle y^2 \rangle_\tau$ is the average of the square of neuron's past activity over the time interval τ and λ is a proportionality constant. The final weight update rule can then be formulated as follows:

$$\Delta w = \alpha \phi x_j, \tag{4}$$

where α is the learning rate. The modification function ϕ is illustrated in Fig. 1.

The threshold θ_m is the point of crossover between synaptic potentiation and depression. Considering that this synaptic modification is a function of postsynaptic response, if it goes below the modification threshold, active synapses depress. When postsynaptic activity surpasses θ_m, active synapses potentiate. The BCM function also compensates for the situation in which the synapses have not been potentiated for a long time and are actually more readily potentiated if some correlated firing appears. This is reflected in sliding of the threshold to the left in Fig. 1. Since the modification function of the BCM learning rule prevents the synaptic weight to potentiate extremely or endlessly, as well as lets the unused connection be potentiated even when the weight is very weak, it ensures *homeostasis* of the whole system. This is the property that we explore in the context of supervised learning, namely the classical multilayer perceptrons.

1.2 Classical supervised learning: the MLP

The multilayer perceptron (Rumelhart et al., 1986) has been proposed in 1980's as an extension of the single-layer perceptron that would overcome the problem of learning linearly inseparable problems. The MLP is also known to be a universal function approximator. A standard MLP consists of an input layer x, one or more hidden layers h, and an output layer y connected with the weight matrices v and w. Each projecting layer contains a trainable bias input fed with constant input -1, so when computing the layer activation the input vector has $k + 1$ nodes (where k is the actual size of the projecting layer). A generic MLP architecture is displayed in Fig. 2. Units in the network compute a weighted sum of activation from the previous layer modified by the activation function f (usually a sigmoid, e.g. logistic function) according to:

$$h_j = f(\sum_{i=1}^{n+1} v_{ij} x_i) \tag{5}$$

and

$$y_k = f(\sum_{j=1}^{q+1} w_{jk} h_j) \tag{6}$$

The MLPs are usually trained in a supervised manner using the error back-propagation (BP) (Rumelhart et al., 1986). After the network produces an estimate at the output layer y for a sample presented on the input layer x, the error computed as the difference between desired and estimated values on the output layer is propagated through the network in the backward direction (backward pass) and weights are updated according to:

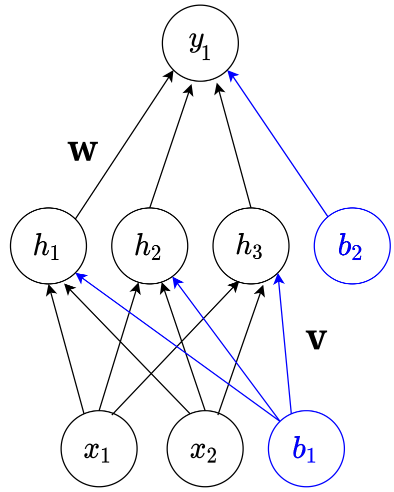

Fig. 2: Schematic depiction of a MLP.

$$\Delta w_{ik} = \alpha \delta_i h_k, \text{ where } \delta_i = (d_i - y_i) f'_i, \tag{7}$$

and

$$\Delta v_{kj} = \alpha \delta_k x_j, \text{ where } \delta_j = (\sum_i w_{ik} \delta_i) f'_k, \tag{8}$$

where $\alpha > 0$ is the learning rate.

The MLP performs supervised learning in a very efficient way. However, the saturation of the synaptic weights of the model is only controlled by the data presented to the network to learn. Our research question in this work is to explore whether regulation of the saturation of the synaptic weights with a mechanism preserving homeostasis inspired by the BCM could lead to improvement in its performance, namely via attributing each weighted connection a local learning rate that would be adapted based on the coactivation of the connected neurons.

2 Related Work

The concept of local learning rates and dynamically changing learning rates in the MLP has been proposed in combination with other regularization techniques (Jacobs, 1988). Various learning-rate adaptations have been studied (Magoulas et al., 1999; Riedmiller and Braun, 1993) and applied several times over the training of the network, as well as at each iteration (Magoulas et al., 2002). This is also known as the learning rate scheduling and has also been proposed as not decreasing, but cyclical function letting the learning rate to stay between reasonable values (Smith, 2017). It achieves improved classification accuracy without the need to manually find the most suitable value.

The BCM rule has been applied as modification of various models. Benuskova et al. (1994) proposed a single representative cell model that used a dynamic synaptic modification threshold in order to explain plasticity in the barrel cortex. Their results demonstrate generalization of plasticity under numerous conditions. The first application of BCM in the field of machine learning was proposed by Bachmann et al. (1994). It is used as an unsupervised learning algorithm using the modification threshold, which is compared with classical BP and literally inhibited BP using ISAR classification problem. Tino et al. (2000) applied a variation of BCM theory on a second-order recurrent neural network with the aim to investigate the unsupervised state space organization in the domain of predicting complex symbolic sequences. Meng et al. (2011) took a different approach with a spiking neural network and simulated BCM regulated by the gene regulatory network with results that demonstrate its effectiveness in human behavior pattern recognition, while acknowledging the need for more robust spatial feature extraction.

3 Our Model

We propose a novel modification of the classical multilayer perceptron to include local learning rate α_n for every weighted connection in the network such that each α_n is modified according to a modification function ψ related to the ϕ function in BCM. Our decision to manipulate the learning rate is based on the fact that the MLP is a machine learning model performing supervised learning in which the updates of the weight are driven by the data in such a way that the synaptic potentiation and depression are arising from the backpropagation of the errors from the data thus making it already possible for the model to drive the weights in both ways. In comparison, the BCM, as an unsupervised learning method, is driving the potentiation and depression based on the postsynaptic neuron q firing rate and does not use any error terms. Therefore our model is inspired by, yet different from the BCM.

The general formula of our new model's learning in a generalized form goes as follows:

$$\Delta w = \psi(\alpha, \theta_M)\delta x, \tag{9}$$

where w are the synaptic weights, the δ is computed in the same way as in equations 7 and 8 and the learning rate α is firstly set to initial value α_0 and subsequently modified by the function ψ throughout the network training.

Our experimentation is currently in progress around various mechanisms that we test separately and finally implement all mechanisms together to form the new model. We are testing several variations of this modification function to explore two crucial aspects of the model. In general, each local learning rate is kept constant until the point when the observed variable indicating the neuron's connection saturation is reached. Subsequently it will be gradually decreased until the end of training or until it reaches zero and thus the connection weight will stay fixed. For the shape of the function according to which the α_n is decreased we considered a linear decrease as well as a logarithmically decreasing function $-\ln(x) + 2$. For the model's basis for the threshold θ_M based on which α_n will start decreasing we have chosen to experiment with two different model's variables, the activation of the postsynaptic neuron q and the error term δ. In the first case the learning rate begins to decrease once the threshold value is reached and in the second case it works the opposite way, once the δ is smaller then the α_n starts to decrease. In our currently proposed model, the value of θ_M is set up experimentally and does not change over the course of the model training. Our proposed final modification function ψ is shown in Fig. 3.

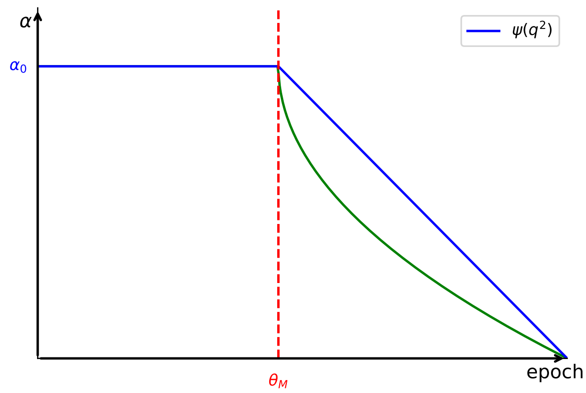

Fig. 3: Schematic depiction of two versions of our proposed modification function ψ applied individually to each local learning rate α_n in the MLP network with the linear (blue) and logarithmic (green) decrease as a function of training epoch (time).

4 Conclusion and Future Work

We have presented our novel adaptation of the classical MLP that defines local learning rates to maintain homeostasis of learning inspired by the BCM theory. There is a lot of experimentation ahead of us. Many aspects of the model shall be explored and studied, from the actual performance enrichment up to the traits such as sensitivity to over-training and others. The model is proposed in a way that would also allow us to implement various regularization techniques such as the momentum. Last, but not least, the computational and spatial complexity shall be analyzed thoroughly to have a conclusive idea on the model's possible use in the future.

Acknowledgement

This paper was written at the Centre for Cognitive Science at DAI FMPI Comenius University Bratislava, with the support of a grant VEGA 1/0373/23.We also thank for support to the Slovak Society for Cognitive SSKV[1].

References

Bachmann, C. M., Musman, S. A., Luong, D., and Schultz, A. (1994). Unsupervised bcm projection pursuit algorithms for classification of simulated radar presentations. *Neural Networks*, 7(4):709–728.

Benuskova, L., Diamond, M. E., and Ebner, F. F. (1994). Dynamic synaptic modification threshold: computational model of experience-dependent plasticity in adult rat barrel cortex. *Proceedings of the National Academy of Sciences*, 91(11):4791–4795.

Bienenstock, E. L., Cooper, L. N., and Munro, P. W. (1982). Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex. *Journal of Neuroscience*, 2(1):32–48.

Blais, B. S. and Cooper, L. N. (2008). BCM theory - Scholarpedia — scholarpedia.org. `http://www.scholarpedia.org/article/BCM_theory`. [Accessed 30-04-2025].

Hebb, D. O. (2005). *The organization of behavior: A neuropsychological theory*. Psychology press.

Jacobs, R. A. (1988). Increased rates of convergence through learning rate adaptation. *Neural networks*, 1(4):295–307.

Magoulas, G., Plagianakos, V., and Vrahatis, M. (2002). Globally convergent algorithms with local learning rates. *IEEE Transactions on Neural Networks*, 13(3):774–779.

Magoulas, G. D., Vrahatis, M. N., and Androulakis, G. S. (1999). Improving the convergence of the backpropagation algorithm using learning rate adaptation methods. *Neural Computation*, 11(7):1769–1796.

Meng, Y., Jin, Y., and Yin, J. (2011). Modeling activity-dependent plasticity in bcm spiking neural networks with application to human behavior recognition. *IEEE Transactions on Neural Networks*, 22(12):1952–1966.

O'Reilly, R. C., Munakata, Y., Frank, M. J., Hazy, T. E., and Contributors (2024). *Computational Cognitive Neuroscience*. Online Book, 5th Edition, URL: `https://compcogneuro.org`.

Riedmiller, M. and Braun, H. (1993). A direct adaptive method for faster backpropagation learning: the rprop algorithm. In *IEEE International Conference on Neural Networks*, pages 586–591 vol.1.

Rumelhart, D., Hinton, G., and Williams, R. (1986). *Learning internal representations by error propagation*, pages 318–362. Number 1. The MIT Press, Cambridge, MA.

Smith, L. N. (2017). Cyclical learning rates for training neural networks. In *2017 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 464–472.

Tino, P., Stancik, M., and Benuskova, L. (2000). Building predictive models on complex symbolic sequences with a second-order recurrent bcm network with lateral inhibition. In *Proceedings of the IEEE-INNS-ENNS International Joint Conference on Neural Networks. IJCNN 2000. Neural Computing: New Challenges and Perspectives for the New Millennium*, volume 2, pages 265–270 vol.2.

[1] https://cogsci.fmph.uniba.sk/sskv/