

Equilibrium Propagation for Complete Directed Neural Networks

Matilde Tristany Farinha,^{1,3} Sérgio Pequito,²
Pedro A. Santos,¹ Mário A. T. Figueiredo³

1 - INESC-ID & Dept. of Mathematics, IST, University of Lisbon, Portugal

2 - Dept. of Industrial and Systems Engineering, RPI, Troy (NY), USA

3 - Instituto de Telecomunicações, IST, University of Lisbon, Portugal

Abstract. Artificial neural networks, one of the most successful approaches to supervised learning, were originally inspired by their biological counterparts. However, the most successful learning algorithm for artificial neural networks, backpropagation, is considered biologically implausible. We contribute to the topic of biologically plausible neuronal learning by building upon and extending the equilibrium propagation learning framework. Specifically, we introduce: a new neuronal dynamics and learning rule for arbitrary network architectures; a sparsity-inducing method able to prune irrelevant connections; a dynamical-systems characterization of the models, using Lyapunov theory.

1 Introduction

Nowadays, many state-of-the-art approaches to *supervised learning* rely on *artificial neural networks* (ANNs). *Backpropagation* (BP) [1], the most successful algorithm to train ANNs [2], is considered bio-implausible since: (i) it lacks local error representation; (ii) it uses distinct forward and backward information passes; (iii) it requires symmetric synaptic weights [3]. To bridge the gap between biological and machine learning, it is thus crucial to find alternatives to BP-based algorithms that encompass properties of biological neural networks, a goal to which much research effort has been devoted.

The *equilibrium propagation* (EP [4]) model adopts a local learning rule and uses just one kind of computation both for the forward and backward information passes, therefore being a more bio-plausible alternative to BP. However, EP has some bio-unrealistic aspects: (i) it assumes symmetric synaptic weights; (ii) it has only been tested on layered architectures; and (iii) it does not promote sparse networks [5]. We tackle these problems by introducing the *DirEcted EP* (DEEP) learning framework, which: (i) allows for asymmetric feedback weights; (ii) allows for arbitrary network architectures; and (iii) promotes sparse networks. DEEP assumes the network is an arbitrary complete directed graph and its training algorithm actively removes the presumably expendable connections by inducing sparsity. We also establish sufficient conditions for convergence of the neuronal dynamics of DEEP's inference phase.

¹This work was supported by national funds through FCT, Fundação para a Ciência e a Tecnologia, under projects UIDB/50021/2020 and SLICE PTDC/CCI-COM/30787/2017, and the first author has conducted part of this research as a visiting research scholar in the Department of Industrial Systems and Engineering, Rensselaer Polytechnic Institute.

The remainder of this paper is organized as follows: in Section 2, state-of-the-art EP-like models are discussed; in Section 3, the new DEEP model is introduced; in Section 4, DEEP is experimentally evaluated and the results are analyzed; and in Section 5, conclusions and possible future work are discussed.

2 Related Work

DEEP was inspired by the original EP, which is an energy-based model described as a multi-layered continuous Hopfield network of recurrently connected neurons with symmetric weights [4]. The idea is to mimic a central aspect of the behaviour of the brain, where neurons perform local computations and evolve collectively towards an equilibrium state of the corresponding dynamical system.

Several extensions and adaptations of the original EP have been proposed. For instance, an asymmetric version of EP (with asymmetric synaptic weights) was proposed [6]; however, with a complete graph architecture, we observed experimentally that the asymmetric EP model is sometimes unable to learn. A bidirectional-EP working both as a generative and a discriminative model was also proposed; however, although it provides an insightful extension of EP [7], it does not solve the weight-transport problem [8]. Another EP adaptation considers spiking neurons, therefore taking an important step towards bio-plausibility [9]. Additionally, EP has been extended to convolutional architectures, for which the lowest error rate among EP-like models on the MNIST classification task [10] has been reported – “approximately 1%” [11]. However, this did not improve the model’s bio-plausibility, since convolutional architectures are considered bio-implausible due to their extensive weight sharing [12].

3 DirEcted Equilibrium Propagation (DEEP)

The DEEP model, for an architecture with a total of N neurons (from which P are input neurons), is completely specified by the following elements:

- a state vector $\mathbf{s}(t) = [s_j(t)]_{j=1}^N \in [0, 1]^N$, containing the neuronal activities, *i.e.*, “firing rates”; this state vector includes sub-vectors that correspond to (fixed) input $\mathbf{x} = [s_j(t)]_{j=1}^P$, hidden $\mathbf{h}(t)$, and output $\hat{\mathbf{y}}(t)$ neurons;
- a matrix $\mathbf{W} = [W_{ij}]_{i,j=1}^N$, where W_{ij} is the weight associated to the connection from neuron i to j ;
- a bias vector $\mathbf{b} = [b_j]_{j=1}^N$, where $b_j = 0$ for every $j : s_j(t) \in \mathbf{x}$; and
- a set of continuous-time differential equations defining its dynamics.

Henceforth, the time dependency of the state variable is omitted to shorten the notation. Defining $\theta = (\mathbf{W}, \mathbf{b})$ and with the input neurons fixed, the neuronal dynamics we propose is dictated by the vector fields (recall that $\dot{s}_j = ds_j/dt$)

$$\dot{s}_j = V_{\theta,j}^\beta(\mathbf{s}, \mathbf{y}, \beta) = \sum_{i=1}^N W_{ij}s_i + b_j - s_j \sum_{i=1}^N W_{ji} - \beta \frac{\partial C_\theta(\hat{\mathbf{y}}, \mathbf{y})}{\partial s_j} \mathbf{1}_{\hat{\mathbf{y}}(s_j)}, \quad (1)$$

where \mathbf{y} is the vector of target/desired outputs, $C_\theta(\hat{\mathbf{y}}, \mathbf{y})$ is the cost (e.g., mean squared error – MSE), β controls how much the cost influences the dynamics, and $\mathbb{1}_{\hat{\mathbf{y}}}(s_j) = 1$ if and only if $s_j \in \hat{\mathbf{y}}$ (i.e., if s_j is an output neuron), otherwise it is zero. Note that this new dynamics accounts for the weighted incoming and outgoing connections and the leakage of a neuron’s activity depends not only on its own activity (as in the original and asymmetric versions of the EP [4, 6]), but also on the weighted sum of outgoing connections.

The training algorithm has two distinct phases: the first phase, with $\beta = 0$; and the second phase, with $\beta \neq 0$. In the first and second phases, the network settles to equilibrium states, denoted \mathbf{s}^0 and \mathbf{s}^β , respectively². The loss function is defined as the cost when the network is at its first equilibrium state, $J_\theta(\mathbf{x}, \mathbf{y}) = C_\theta(\mathbf{s}^0, \mathbf{y})$. When used for inference, the activities of the input neurons are fixed and the network evolves to equilibrium \mathbf{s}^0 , from which the output is read at the corresponding output neurons $\hat{\mathbf{y}}$.

The learning rule we propose is obtained by numerical integration of the bio-inspired weight dynamics $\dot{W}_{ij} \propto s_i(t)\dot{s}_j(t)$ [13] in the path from \mathbf{s}^0 to \mathbf{s}^β in M_β -steps (the time derivative \dot{s}_j is approximated by a backward difference):

$$\Delta W_{ij} \propto \frac{1}{M_\beta} \sum_{m=M_0+1}^{M_0+M_\beta} s_i(m)(s_j(m) - s_j(m-1)), \quad (2)$$

where M_0 is the number of steps of the first phase. Interpreting the biases as the weighted outgoing connections from a neuron with activity fixed to one, their updates are also given by Equation (2). Note that the learning rule of the previously proposed asymmetric version of EP [6] corresponds to a one-step version of the numerical integration in Equation (2) with a forward difference approximation of \dot{s}_j , which yields simply $\Delta W_{ij} \propto s_i^0(s_j^\beta - s_j^0)$.

Sparsity-inducing ℓ_1 regularization is added to the proposed learning rule. Moreover, sparsity is actively enforced at each weight update when such weight is below a certain threshold. Specifically, each weight $|W_{ij}| < \lambda$, with $\lambda > 0$ small, is randomly removed from the network with probability p_{ij} , given by a Boltzmann distribution defined across the incoming weights of neuron j , i.e.,

$$p_{ij} = e^{-|W_{ij}|/T} / \sum_{k=1}^N e^{-|W_{kj}|/T}. \quad (3)$$

In this context, the temperature T represents how likely it is for stronger connections (in absolute value) to be deemed irrelevant.

3.1 Analytic Properties

3.1.1 Time-invariant Sum of Firing Rates

DEEP’s neuronal dynamics is bio-plausible as it has been reported that, in the absence of external sensory stimulus or motor activity, the grand mean firing

²To ensure that the firing rates are bounded by $[0, 1]$ and facilitate convergence, after each discrete update of the state variable \mathbf{s} , the state variable is bounded by the hard-sigmoid function.

rate of the hippocampal neurons remains constant [14].

Proposition 1. *In the absence of external stimulus (no fixed activity of the input or bias neurons), the dynamics in Equation (1), with $\beta = 0$, preserves the sum of firing-rates through time (i.e., $\sum_j \dot{s}_j(t) = 0$, $t \in \mathbb{R}_0^+$).*

Proof. It follows from $\sum_{j=1}^N \dot{s}_j = \sum_{j=1}^N V_{\theta,j}^0(\mathbf{s}) = \sum_{i,j=1}^N (W_{ij}s_i - W_{ij}s_i) = 0$. \square

3.1.2 Sufficient Conditions for Stability

Sufficient conditions for local asymptotic stability of \mathbf{s}^0 (i.e., the equilibrium state reached during inference) can be obtained by leveraging on Gergschorin’s circle theorem and nonlinear control analysis tools from Lyapunov theory [15]. Specifically, sufficient conditions for the stability of the inference phase of DEEP are given in the following proposition.

Proposition 2. *If $\sum_{i=1}^N W_{ji} > 0$ and $\sum_{i=P+1}^N |W_{ij}| < \left| \sum_{i=1}^N W_{ji} \right|$ is verified for $j \in \{P+1, \dots, N\}$, then \mathbf{s}^0 (equilibrium state with respect to Equation (1), with $\beta = 0$) is locally asymptotically stable.*

Proof. For a time-invariant nonlinear dynamical system $\dot{\mathbf{s}}(t) = f(\mathbf{s}(t))$, an equilibrium state \mathbf{s}^* is locally asymptotically stable if $J_{jj} < 0$ and $R_j < |J_{jj}|$ is satisfied for $j \in \{1, \dots, N\}$, where $\mathbf{J} = Df(\mathbf{s}^*) \in \mathbb{R}^{N \times N}$ (the Jacobian of f evaluated at \mathbf{s}^*), and $R_j = \sum_{i=1, i \neq j}^N |J_{ji}|$ [15]. These conditions can be particularized for the system given by Equation (1) with $\beta = 0$, where $\mathbf{J} = [\partial V_{\theta,j}^0(\mathbf{s}(t), \mathbf{y}, \beta) / \partial s_j]_j$, for $j \in \{P+1, \dots, N\}$, yielding the conditions stated in the proposition. \square

4 Experiments and Results

DEEP is constrained by the curse of dimensionality due to its high dimensional search space, so its performance is analyzed for simple tasks such as learning logical operations. For these tasks, the architecture considered is a 8-neuron complete directed graph. Although DEEP can learn XOR with 1 hidden neuron and AND and OR with none, 5 hidden neurons are used so that sparsity is perceivable when using the sparsity-inducing method.

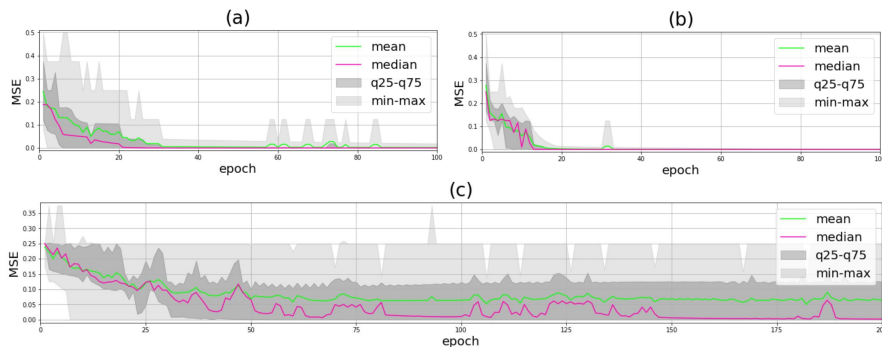


Fig. 1: MSE convergence during the first phase of 10 independently trained models while learning the logical operations: (a) AND, (b) OR, and (c) XOR.




Fig. 2: Sparse networks trained for learning the logical operations: (a) AND (93.75% sparse), (b) OR (95.83% sparse), and (c) XOR (62.50% sparse). The indexes refer to the neuron’s type: ①-① input, ②-⑥ hidden, ⑦ output; the connections’ opacity is proportional to their strength (in absolute value).

Figure 1³ illustrates how, when learning XOR, the MSE converges to zero most of the times but not always. Besides, it converges less frequently than when learning AND or OR, which highlights how XOR is harder to learn. Additionally, DEEP is also trained with the sparsity-inducing method mentioned above. We observed that, with this method, the most relevant connections are strengthened and, while for simple tasks all the expendable connections are removed, for more complex tasks only a few are removed (see Figure 2).

The performances of DEEP and the asymmetric version of EP [6] are compared with the same architecture (complete directed graph) and hyperparameters. Asymmetric EP fails to learn XOR and, for AND and OR, we observed that the MSE converged slower than DEEP, which took approximately half the number of epochs, and not always to zero.

5 Conclusions and Future Work

In this paper, we extended the *equilibrium propagation* (EP) learning framework towards more bio-plausible artificial neural networks by generalizing its architecture to a complete directed graph and introducing a new neuronal dynamics and learning rule. We also proposed a sparsity-inducing method to prune irrelevant connections. The resulting model was termed *DirEcted Equilibrium Propagation* (DEEP). Simulation results suggest that DEEP is able to learn logical operations that previous models are unable to learn. We supported our results with theoretical sufficient conditions to attain local asymptotic stability during inference.

As DEEP is defined by a continuous-time dynamical system, it provides an interesting line of research for algorithms that can be efficiently implemented with neuromorphic hardware [16]. Moreover, due to its unrestricted architecture, DEEP could be used as a network design tool: the optimized, possibly minimalist, structure of the trained networks could be used as an initial architecture for other learning algorithms.

³The min-max refers to the interval of values contained between the minimum and the maximum of the data, and the q25-q75 refers to the interval of values contained between the first and third quartiles of the data.

To accelerate the convergence of the first phase, a possible idea would be to initialize the neuronal activity as the solution of the neuronal dynamical system obtained when the non-linearity inserted by the hard-sigmoid is removed. Moreover, a detailed theoretical study should be made regarding the approximation scheme used to obtain the results in this discrete framework, in order to determine the influence of the hard-sigmoid in the stability of the equilibrium states reached with respect to this discrete dynamics.

It would also be interesting to study whether DEEP can be adapted to a spiking neural network (as in [9]) and, by leveraging on DEEP's recurrent nature, if it could be used for sequence prediction problems.

References

- [1] D.E. Rumelhart and J.L. McClelland. Learning internal representations by error propagation. In *Parallel Distributed Processing: Explorations in the Microstructure of Cognition: Foundations*, chapter 8, pages 318-362. MIT Press, 1st edition, 1987.
- [2] Y. Lecun, Y. Bengio, and G.E. Hinton. Deep learning. *Nature*, 521(7553):436-444, 2015.
- [3] J. Guerguiev, T. Lillicrap, and B. Richards. Towards deep learning with segregated dendrites. *eLife*, 6(e22901), 2017.
- [4] B. Scellier and Y. Bengio. Equilibrium propagation: bridging the gap between energy-based models and backpropagation. *Frontiers Computational Neuroscience*, 11(24), 2017.
- [5] R. O'Reilly. Six principles for biologically based computational models of cortical cognition. *Trends in Cognitive Sciences*, 2(11):455-462, 1998.
- [6] B. Scellier, A. Goyal, J. Binas, T. Mesnard, and Y. Bengio. Generalization of equilibrium propagation to vector field dynamics. Arxiv:1808.04873v1, 2018.
- [7] A. Khan. Bidirectional learning in recurrent neural networks using equilibrium propagation (master thesis), University of Waterloo, 2018.
- [8] S. Grossberg. Competitive learning: from interactive activation to adaptive resonance. *Cognitive Science*, 11(1):23-63, 1987.
- [9] P. O'Connor, E. Gavves, and M. Welling. Training a spiking neural network with equilibrium propagation. *JMLR*, 89:1516-1523, 2019.
- [10] Y. LeCun and C. Cortes. MNIST handwritten digit database, 2010.
- [11] M. Ernout, J. Grollier, D. Querlioz, Y. Bengio, and B. Scellier. Updates of equilibriumprop match gradients of backprop through time in an RNN with static input. In *NeurIPS*, 2019.
- [12] S. Bartunov, A. Santoro, B. Richards, L. Marris, G. Hinton, and T. Lillicrap. Assessing the scalability of biologically-motivated deep learning algorithms and architectures. In *NeurIPS*, 2018.
- [13] X. Xie and H. Seung. Spike-based learning rules and stabilization of persistent neural activity. In *NIPS*, 1999.
- [14] H. Hirase, X. Leinekugel, A. Czurkó, J. Csicsvari, and G. Buzsáki. Firing rates of hippocampal neurons are preserved during subsequent sleep episodes and modified by novel awake experience. *PNAS*, 98(16):9386-9390, 2001.
- [15] D. Bejarano, E.I. Mondragon, and E.G. Hernandez. A stability test for non linear systems of ordinary differential equations based on the Gershgorin circles. *Contemporary Engineering Sciences*, 11(91):4541-4548, 2018.
- [16] S. Ambrogio, P. Narayanan, H. Tsai, R.M. Shelby, I. Boybat, C. Di Nolfo, S. Sidler, M. Giordano, M. Bodini, N.C.P. Farinha, B. Killeen, C. Cheng, Y. Jaoudi, and W.G. Burr. Equivalent-accuracy accelerated neural-network training using analogue memory. *Nature*, 558(7708):60-67, 2018.