Implementing Backpropagation for Learning on Neuromorphic Spiking Hardware

Renner, Alpha ; Sheldon, Forrest ; Zlotnik, Anatoly ; Tao, Louis ; Sornborger, Andrew

DOI: https://doi.org/10.1145/3381755.3381768

Originally published at:
Renner, Alpha; Sheldon, Forrest; Zlotnik, Anatoly; Tao, Louis; Sornborger, Andrew (2020). Implementing Backpropagation for Learning on Neuromorphic Spiking Hardware. In: NICE '20: Neuro-inspired Computational Elements Workshop, Heidelberg Germany, 2020, ACM.
DOI: https://doi.org/10.1145/3381755.3381768

Implementing Backpropagation for Learning on Neuromorphic Spiking Hardware

Alpha Renner
alpren@ini.uzh.ch
Institute of Neuroinformatics
University of Zurich and ETH Zurich
Zurich, Switzerland

Forrest Sheldon*
fsheldon@lanl.gov
Physics of Condensed Matter &
Complex Systems, T-4
Los Alamos National Laboratory
Los Alamos, New Mexico, USA

Anatoly Zlotnik
azlotnik@lanl.gov
Applied Mathematics & Plasma
Physics, T-5
Los Alamos National Laboratory
Los Alamos, New Mexico, USA

Louis Tao
taolt@mail.cbi.pku.edu.cn
Center for Quantitative Biology
Peking University
Beijing, China

Andrew Sornborger*
sornborg@lanl.gov
Information Sciences, CCS-3
Los Alamos National Laboratory
Los Alamos, New Mexico, USA

ACM Reference Format:
Alpha Renner, Forrest Sheldon, Anatoly Zlotnik, Louis Tao, and Andrew Sornborger. 2020. Implementing Backpropagation for Learning on Neuromorphic Spiking Hardware. In *Neuro-inspired Computational Elements Workshop (NICE '20), March 17–20, 2020, Heidelberg, Germany.* ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/3381755.3381768

1 SUMMARY

Many contemporary advances in the theory and practice of neural networks are inspired by our understanding of how information is processed by natural neural systems. However, the basis of modern deep neural networks remains the error backpropagation algorithm [1], which though founded in rigorous mathematical optimization theory, has not been successfully demonstrated in a neurophysiologically realistic circuit. In a recent study, we proposed a neuromorphic architecture for learning that tunes the propagation of information forward and backwards through network layers using an endogenous timing mechanism controlled by thresholding of intensities [2]. This mechanism was demonstrated in simulation of analog currents, which represent the mean fields of spiking neuron populations. In this follow-on study, we present a modified architecture that includes several new mechanisms that enable implementation of the backpropagation algorithm using neuromorphic spiking units. We demonstrate the function of this architecture in learning mapping examples, both in event-based simulation as well as a true hardware implementation.

2 BACKGROUND

There has been a rapid growth of interest in the re-formulation of classical algorithms for learning, optimization, and control using event-based information processing mechanisms inspired by the function of biophysiological neural systems [3]. The trend is driven by the advent of flexible computing architectures that enable experimentation with such algorithms in hardware [4]. Modern deep learning relies on a layered, feedforward network similar to the early layers of the visual cortex, with threshold nonlinearities at each layer that resemble mean-field approximations of neuronal integrate-and-fire models. However, the unnatural structure of backpropagation has made the algorithm notoriously difficult to implement in a neural circuit [5, 6]. A feasible neural implementation of the backpropagation algorithm has become more compelling with the rise of new neuromorphic computational architectures that feature local synaptic plasticity [4, 7, 8]. Neuromorphic systems have relied to date on conventional off-chip learning, and used on-chip computing only for inference [9, 10]. It has been a long-standing challenge to develop learning systems whose function is affected exclusively using neuromorphic mechanisms.

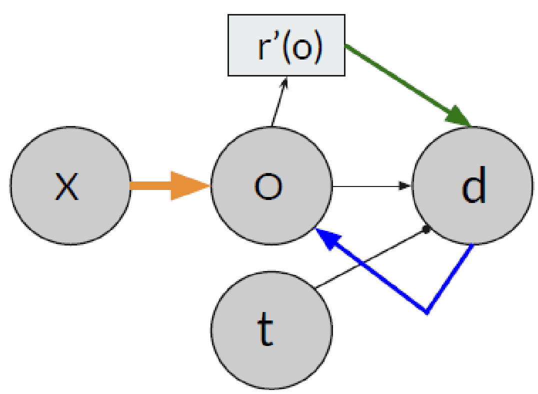

Figure 1: Simple outline of a (single layer) learning circuit, where x and o are circuit input and output, t is a target for learning, d is an error, and r' is the derivative of the activation function. Here the arrows represent feedforward copy (orange) and backpropagated error (blue), and derivative-weighted masking for error computation (green). For simplicity short-term memories are not shown.

*Corresponding authors.

NICE '20, March 17–20, 2020, Heidelberg, Germany

ACM ISBN 978-1-4503-7718-8/20/03.
https://doi.org/10.1145/3381755.3381768

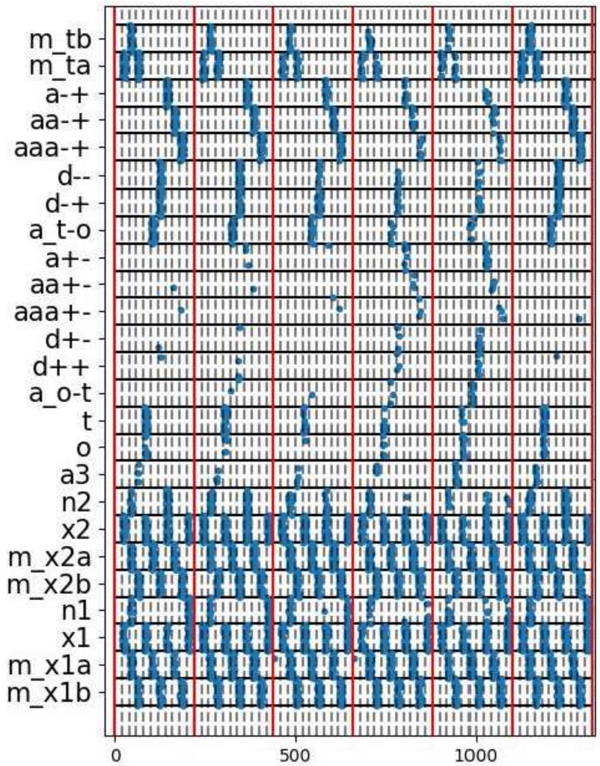

Figure 2: Spiking activity in one layer of the network, whose basic structure is given in Fig. 1. Neural sub-populations are identified on the ordinate, and time is shown on the abscissa. The time interval between two red lines is the time needed to propagate information through the circuit (see Fig. 3), and the time between two black notches on the abscissa is a single gating window.

3 CONTRIBUTIONS

In this study, we describe the hardware implementation of the backpropagation algorithm that makes use of mechanisms that have been developed and tested in simulation by the authors during the past decade, and synthesized in our recent study [2]. These proven neuronal and network features include propagation of graded information in a circuit composed of neural populations using synfire-gated synfire chains (SGSCs) [11–14], decision-making based on the interaction of synfire-chains [12], and regulation of Hebbian learning using pulse-gating [15, 16]. Furthermore, we introduce a number of recently discovered mechanisms that enable the precise regulation of the timing of information propagation using spiking intensities, rather than the approximating mean-field values that were previously demonstrated in simulation. Our approaches mathematically formalize the neurophysiological concepts of inhibition, excitation, potentiation, and depression. We then engineer the connectivity that affects these using event-based (spiking) neural computation, and integrate them into the proposed architecture to align the timing of unconditional and conditional gating, regularize the gap between layer activations caused by synaptic delay, and enable stable propagation of gradients in the evaluation of Hadamard products.

Beyond the mathematical exposition of new neural information processing mechanisms and presentation of circuit architecture design, we demonstrate our approach with several examples, including an implementation of the classic XOR learning circuit using a rectified linear unit (ReLU) nonlinearity, as well as the MNIST test case. The function of the circuits used in the examples is first demonstrated and examined within the Brian2 simulation environment [17], with results as shown in Figures 2, 3 and 4. Then, significantly, algorithm performance is benchmarked in hardware using the Intel Loihi neuromorphic chip [4].

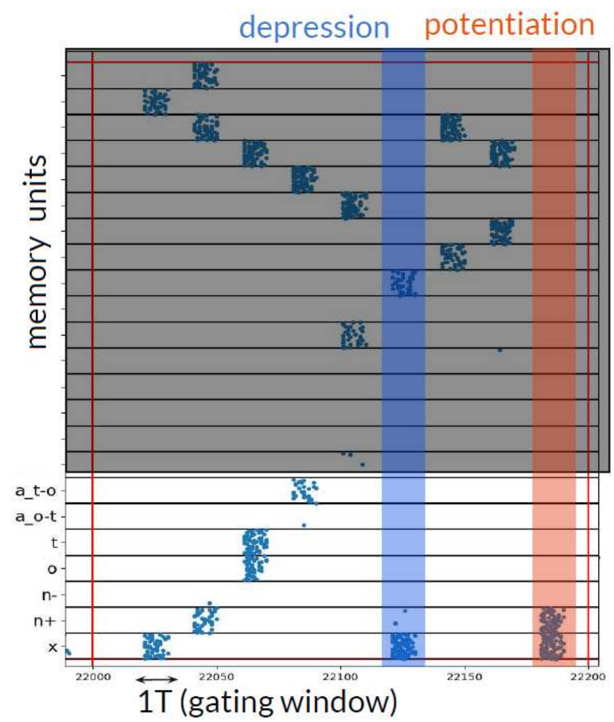

Figure 3: Spiking activity in one layer of the network shown in Fig. 2 during a single learning cycle. The neural sub-populations that perform memory functions are highlighted in grey on the y-axis, and time is shown on the x-axis. The event-based spiking activity is clearly seen on this time-scale.

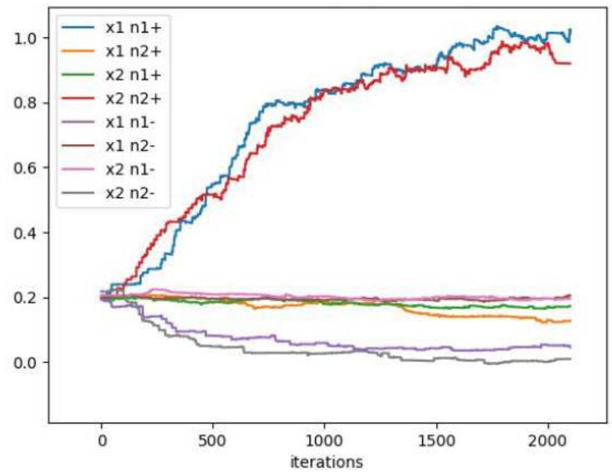

Figure 4: Convergence of synaptic weights in event-based (spiking) simulation of the circuit learning a linear map within the Brian2 environment.

REFERENCES

[1] D.E. Rumelhart, G.E. Hinton, and R.J. Williams. Learning representations by back-propagating errors. *Nature*, pages 533–536, 1986.
[2] Andrew Sornborger, Louis Tao, Jordan Snyder, and Anatoly Zlotnik. A pulse-gated, neural implementation of the backpropagation algorithm. In *Proceedings of the 7th Annual Neuro-inspired Computational Elements Workshop*, page 10. ACM, 2019.
[3] D. L. Yamins and J. J. DiCarlo. Using goal-driven deep learning models to understand sensory cortex. *Nat. Neurosci.*, 19(3):356–365, Mar 2016.
[4] Mike Davies, Narayan Srinivasa, Tsung-Han Lin, Gautham N. Chinya, Yongqiang Cao, Sri Harsha Choday, Georgios D. Dimou, Prasad Joshi, Nabil Imam, Shweta Jain, Yuyun Liao, Chit-Kwan Lin, Andrew Lines, Ruokun Liu, Deepak Mathaikutty, Steven McCoy, Arnab Paul, Jonathan Tse, Guruguhanathan Venkataramanan, Yi-Hsin Weng, Andreas Wild, Yoonseok Yang, and Hong Wang. Loihi: A neuromorphic manycore processor with on-chip learning. *IEEE Micro*, 38(1):82–99, 2018.
[5] Stephen Grossberg. Connectionist models and their implications: Readings from cognitive science. chapter Competitive Learning: From Interactive Activation to Adaptive Resonance, pages 243–283. Ablex Publishing Corp., Norwood, NJ, USA, 1988.
[6] F. Crick. The recent excitement about neural networks. *Nature*, 337:129–132, 1989.
[7] E. Painkras, L.A. Plana, J. Garside, S. Temple, S. Davidson, J. Pepper, D. Clark, C. Patterson, and S. Furber. SpiNNaker: A multi-core system-on-chip for massively-parallel neural net simulation. In *Proceedings of the IEEE 2012 Custom Integrated Circuits Conference*, 2012.
[8] J. Schemmel, D. Briiderle, A. Griible, M. Hock, K. Meier, and S. Millner. A wafer-scale neuromorphic hardware system for large-scale neural modeling. In *Proceedings of the IEEE 2010 IEEE International Symposium on Circuits and Systems*, 2010.
[9] S.K. Esser, P.A. Merolla, J.V. Arthur, A.S. Cassidy, R. Appuswamy, A. Andreopoulos, D.J. Berg, J.L. McKinstry, T. Melano, D.R. Barch, C. di Nolfo, Datta P., A. Amir, B. Taba, M.D. Flickner, and D.S. Modha. Convolutional networks for fast, energy-efficient neuromorphic computing. *PNAS*, 113:11441–11446, 2016.
[10] W. Severa, C.M. Vineyard, R. Delanna, S.J. Verzi, and J.B. Aimone. Whetstone: A method for training deep artificial neural networks for binary communication. 2018.
[11] A.T. Sornborger, Z. Wang, and L. Tao. A mechanism for graded, dynamically routable current propagation in pulse-gated synfire chains and implications for information coding. *J. Comput. Neurosci.*, August 2015.
[12] Z. Wang, A.T. Sornborger, and L. Tao. Graded, dynamically routable information processing with synfire-gated synfire chains. *PLoS Comp Biol*, 12:6, 2016.
[13] A.T. Sornborger and L. Tao. A Fokker-Planck approach to graded information propagation in pulse-gated feedforward neuronal networks. *ArXiv:1512.0052*, Oct 2015.
[14] Z.C. Xiao, B.X. Wang, A.T. Sornborger, and L. Tao. Mutual information and information gating in synfire chains. *Entropy*, 20:102, 2018.
[15] Y. Shao, A.T. Sornborger, and L. Tao. A pulse-gated, predictive neural circuit. *Proceedings of the 50th Asilomar Conference on Signals, Systems and Computers*, 2016.
[16] Yuxiu Shao, Binxu Wang, Andrew Sornborger, and Louis Tao. A mechanism for synaptic copy between neural circuits. *bioRxiv*, 2018.
[17] Dan FM Goodman, Marcel Stimberg, Pierre Yger, and Romain Brette. Brian 2: neural simulations on a variety of computational hardware. *BMC neuroscience*, 15(1):P199, 2014.