

On-chip Few-shot Learning with Surrogate Gradient Descent on a Neuromorphic Processor

Kenneth Stewart
Computer Science Department
UC, Irvine
Irvine, CA 92697-2625 USA
kennetms@uci.edu

Garrick Orchard
Intel Labs
Intel Corporation
Santa Clara, California
garrick.orchard@intel.com

Sumit Bam Shrestha
Tamasek Laboratories @ NUS
National University of Singapore
Singapore
tslbs@nus.edu.sg

Emre Neftci
Cognitive Sciences Department
and
Computer Science Department
UC, Irvine
Irvine, CA 92697-2625 USA
eneftci@uci.edu

Abstract—Recent work suggests that synaptic plasticity dynamics in biological models of neurons and neuromorphic hardware are compatible with gradient-based learning [1]. Gradient-based learning requires iterating several times over a dataset, which is both time-consuming and constrains the training samples to be independently and identically distributed. This is incompatible with learning systems that do not have boundaries between training and inference, such as in neuromorphic hardware. One approach to overcome these constraints is transfer learning, where a portion of the network is pre-trained and mapped into hardware and the remaining portion is trained online. Transfer learning has the advantage that pre-training can be accelerated offline if the task domain is known, and few samples of each class are sufficient for learning the target task at reasonable accuracies. Here, we demonstrate on-line surrogate gradient few-shot learning on Intel’s Loihi neuromorphic research processor using features pre-trained with spike-based gradient backpropagation-through-time. Our experimental results show that the Loihi chip can learn gestures online using a small number of shots and achieve results that are comparable to the models simulated on a conventional processor.

Keywords— neuromorphic computing, spiking neural networks, on-chip learning, few-shot learning

I. INTRODUCTION

Understanding how the plasticity dynamics in multilayer biological neural networks are organized for efficient data-driven learning is a long-standing question in computational neuroscience [2], and elucidating it can open the door to scalable brain-inspired computers for continual “life-long” learning. Machine learning and artificial neural networks provide important hints for answering this question. Recent work has demonstrated that many ingredients of deep learning are compatible with biological neural networks and neuromorphic hardware [3], [1]. This compatibility paves the road towards neuromorphic learning machines with performances similar to deep learning, while being able to learn online, with input streams, and using a fraction of the energy.

However, neural network training is especially time-consuming in neuromorphic hardware since they generally operate in real-time and on a single stream of data (i.e. in deep learning terms, using a batchsize of 1). Furthermore, the training samples must be independently and identically distributed (iid). The failure to meet the iid constraint often results in catastrophic forgetting [4]. Methods such as transfer learning [5] and meta-learning [6] can mitigate these two problems by pre-training a network in an iid fashion on a known task domain, followed by a small number of “shots” to learn the task-specific output. Catastrophic forgetting is mitigated, because fewer data samples are needed to achieve a target accuracy compared to training from scratch, which translates into fewer synaptic weight updates.

Here, we leverage recent advances in surrogate gradient learning [1] and on-chip plasticity in neuromorphic hardware to demonstrate few-shot learning in an event-based gesture recognition task in Intel’s neuromorphic research processor, Loihi. A spiking convolutional neural network is trained offline using backpropagation through time [7] using an accurate behavioral model of the Loihi neuron. The classification layer is then learned on-line using a variation of a three-factor rule for surrogate gradient learning [8] adapted to overcome limitations in neuromorphic hardware such as locality and bit precision. We test the network learning performance on a gesture recognition task in a few-shot learning framework, and the performance is compared to that of an equivalent software model trained and executed on a conventional computer. Our results demonstrate that the Loihi hardware achieves similar results to that of the software simulations, and provided the building blocks for multilayer learning and transfer learning in neuromorphic hardware.

II. SURROGATE GRADIENT DESCENT WITH THE LOIHI NEUROMORPHIC RESEARCH CHIP

A. Neuron and Plasticity Model

The dynamics of the Loihi neuron follow the current-based leaky integrate-and-fire model. The neuron uses two internal state variables: The synaptic response current $u_i[n]$ and the membrane potential $v_i[n]$. The current is computed as the sum of filtered input spike trains added to a constant bias current;

$$\begin{aligned} u_i[n] &= \sum_j w_{ij} Q_j[n] \\ Q_j[n+1] &= \alpha_Q Q_j[n] + \frac{1}{\tau_u} S_j[n]. \end{aligned} \quad (1)$$

The variable $S_j[n]$ is the input spike train, modeled as a sum of boxcar functions of width and height equal to 1, with n the discrete time step. $Q_j[n]$ here is a first-order low-pass filter that models the unweighted post-synaptic current. Upon an input spike S_j , u_i undergoes a jump of height w_{ij} , then decays exponentially with time constant τ_u , where $\alpha_Q = \exp(-\frac{1}{\tau_u})$. The connection weights w_{ij} are the trainable parameters of the neuron. This synapse model is commonly used in computational neuroscience modeling [9], [2]. In the equations above, inputs neurons are confounded with network neurons for simplicity of notation. Although recurrent connections can be learned using the same dynamics as described in this work, our experiments focus on training feed-forward weights only.

The membrane potential follows similar dynamics and is the integration of the synaptic current with an exponential decay. A neuron emits a spike when its membrane potential exceeds the firing threshold v_{th} .

$$\begin{aligned} S_i[n] &= \Theta(v_i[n] - v_{th}), \\ v_i[n] &= \sum_j w_{ij} P_j[n] + R_i[n], \end{aligned} \quad (2)$$

where,

$$\begin{aligned} P_j[n+1] &= \alpha_P P_j[n] + \frac{1}{\tau_v} Q_j[n], \\ R_i[n+1] &= \alpha_R R_i[n] - S_i v_{th}. \end{aligned} \quad (3)$$

The term $\alpha_P = \exp(-\frac{1}{\tau_v})$ captures the leak of state v . Following an input spike S_j , P follows an unweighted “alpha” post-synaptic potential with fast rise and slow decay. The term $R_i[n]$ captures the reset of the neuron after firing. The neuron model above is a discrete-time version of the Spike Response Model (SRM₀) [2] and is fully compatible with the Loihi neuron model.

The Loihi research processor offers a variety of local information to a programmable synaptic learning process including pre and post synaptic traces and reward traces corresponding to special reward spikes [10]. Synaptic weights can be updated via a learning rule expressed as a finite-difference equation with respect to a synaptic state variable that follows a sum-of-products form as follows [10]:

$$w_{ij}[n+1] = w_{ij}[n] + \sum_k C_k \prod_l F_{kl}, \quad (4)$$

where w_{ij} is the synaptic weight variable defined for the destination-source neuron pair being updated; C_k is a scaling constant; and F_{kl} may represent a synaptic state variable, a pre-synaptic trace, or a post-synaptic trace defined for the neuron pair i and j . The weights are stochastically rounded according to the programmed weight precision. All plastic weights in our work used 8 bit precision. Currently, the Loihi’s learning rules cannot scalably depend on an extrinsic factor such as modulation. (Although there exists an external reward modulation variable, these are limited to 4 per core but surrogate gradient learning requires one per neuron.) In the following paragraph, we describe a workaround for this constraint and the implementation of surrogate gradient learning.

B. Surrogate Gradient Learning in Spiking Neurons

The main contribution of this work is the implementation of a surrogate gradient-based learning rule in the Loihi research processor using three-factor rules. Three-factor rules are extensions of Hebbian learning and Spike Time Dependent Plasticity (STDP), which can be derived from a normative (top-down) approach [11]. Such rules are compatible with a wide number of unsupervised, supervised, and reinforcement learning paradigms [11], and implementations can have scaling properties comparable to that of STDP [12].

The third factor in three-factor rules is motivated by neuromodulators, such as Dopamine, Acetylcholine, or Noradrenaline in reward-based learning [13]. Generally, the weight update of the three-factor learning rule can be written as follows:

$$\Delta w_{ij}^{3F} \propto f_{pre} f_{post} M_i \quad (5)$$

where f_{pre} and f_{post} correspond to functions over presynaptic and post synaptic variables, respectively (for example, but not restricted to traces P_i and P_j), and M_i is the modulating term of the postsynaptic neuron i . Note that, to reduce clutter, time indices will be omitted when there is no ambiguity (all terms in the same time step n). The modulating term is a task-dependent function, which can, for example, represent error, surprise, reward, or another variable relevant to the task at hand.

Three-factor rules can be derived from gradient descent on the spiking neural network [14], [1]. The significance of this compatibility is far-reaching because the techniques and tools of deep learning can be made compatible with training spiking neural networks. Such rules are often “local” in the sense that all the information necessary for computing the gradient is available at the post-synaptic neuron [3]. For this reason, three-factor rules are a candidate for implementation in neuromorphic hardware.

In our contribution, we start with a three-factor synaptic plasticity rule derived from surrogate gradient descent for classification or regression tasks [8]. For the output layer:

$$\Delta w_{ij} \propto Error_i \sigma'(v_i - v_{th}) P_j, \quad (6)$$

where the $Error_i$ is the instantaneous error (target _{i} – output _{i}), P_j is the pre-synaptic trace at the membrane potential, and $\sigma'(v_i[n] - v_{th})$ is the derivative of the smoothed output function with v_{th} the firing threshold. For the output layer of a classification or regression neural network, this learning rule is equivalent to the Delta rule in artificial neural networks. It is derived from the gradient of a Mean-Squared Error (MSE) loss with respect to the neuron parameters, ignoring the refractory variable R and replacing derivatives of Θ with those of the sigmoid function [1]. This rule is the building block of the backpropagation algorithm in multilayer networks, in which case $Error_i$ becomes the backpropagated errors (i.e. the “deltas”). Because this rule is derived from gradient descent on the spiking dynamics (i.e. there is no rate-based approximation), it naturally leverages spike timing dynamics. As a result, this rule and its variations have been shown to perform particularly well on spike-based classification tasks and outperform traditional synaptic plasticity rules such as STDP and reward/error modulated STDP [8], [15] in equivalent networks.

C. Implementation of Surrogate Gradient Learning in the Loihi Research Processor

Here, we map the Loihi learning dynamics to Eq. (6). By definition, the factor σ' is constrained to the range $[0, 1]$. It is non-zero when the membrane potential is close to 0. Removing this term will cause synaptic weights to endlessly grow, but does not prevent learning in the early stages of training. This mode of operation is compatible with the few-shot learning scenario experimented here. The two other terms f_{pre} and M_i , however, are critical for learning. Thus, we drop the σ' term in (6) and obtain provisionally:

$$\Delta w_{ij} \propto P_j P_i^{err}. \quad (7)$$

where P_i^{err} is an error trace as described below. To represent the error, we use dynamics:

$$\begin{aligned} S_i^{err} &= \Theta(v_i^{err} - v_{th}), \\ v_i^{err} &= \sum_j w_{ij} P_j - w_{tgt} \hat{P}_i + R_i^{err} + b^{err}, \end{aligned} \quad (8)$$

where P , $\hat{P}$ and R_i^{err} follow the same dynamics as in (2). $\hat{P}$ is the input due to the target spikes, and $w_{tgt} > 0$ is its weight. The constant b^{err} is a bias term used to offset the firing rate of the neuron, as explained in the following paragraphs.

The synaptic plasticity processor operates on neural traces computed with pre-synaptic and post-synaptic spike times. In Loihi, traces are first-order linear filters. However, as discussed earlier P is a second-order linear filter. Second-order kernels can be implemented as a subtraction of two first-order kernels [2]. This subtraction is enabled by the sum-of-products formulation of the plasticity rule (4). The error can be implemented with the post-synaptic compartment using the following post-traces:

$$\begin{aligned} P_i^{err}[n+1] &= \alpha_v P_i^{err}[n] + Q_i^{err}[n], \\ Q_i^{err}[n+1] &= \alpha_u Q_i^{err}[n] + S_i^{err}[n]. \end{aligned}$$

Post-traces can only be positive but errors can be both positive and negative. This problem can be solved by offsetting the weight updates with a constant term b :

$$\Delta w_{ij} \propto P_j (P_i^{err} - b). \quad (9)$$

To enforce the constraint of zero error when output matches target, a constant firing rate is imposed on the error neurons via b^{err} . In the absence of inputs, this constant firing will create a repeating P_i^{err} trace with average $\langle P_i^{err} \rangle$. The bias b is then adjusted such that it is equal to the average of this trace: $\langle P_i^{err} \rangle = b$.

Two potential problems remain with this workaround: a) bounds on the error magnitude, b) Fluctuations in the weight updates even when the error is zero, since only the average update is zero with the baseline firing rate. Previously work suggest that a) does not prevent




Fig. 1. Setup and multi-compartment neuron used in our work. The distal compartment receives spikes from the spiking ConvNet with synaptic weights adjusted by the surrogate gradient rule. The distal compartment pushes its current to the proximal compartment for integration. During training, label spikes with weights of opposite signs are received by the distal compartment and proximal compartment. Error is calculated in the distal compartment as a subtraction of the label input and input from the previous layer.

convergence in spiking neural network classification tasks [16]. While (b) is currently unsolved, we observe that it does not prevent few-shot learning. The learning rule is implemented in Loihi as follows:

$$\Delta w_{ij} \propto C(Y_1(X_2 - X_1) + X_1 - X_2) \quad (10)$$

where C is a constant that captures the learning rate, X_1 and X_2 are two first-order pre-synaptic traces, Y_1 is a first-order post-synaptic trace. The $(X_2 - X_1)$ term corresponds to P_j in (9), and Y_1 corresponds to P_i^{err} in (9). Note that to simplify the rule, we approximated P_i^{err} using a single first-order filter.

D. Multi-compartment Neuron and Readout

In order to read out the class and to build multilayer learning in future work, we use a multi-compartment neuron. The neurons consist of two compartments, made of a distal (error) compartment that receives input from a pre-trained network, and a proximal (output) compartment whose voltage integrates the state u^{err} copied from the distal compartment, and does not directly receive spiking input from the convnet:

$$\begin{aligned} S_i^{out} &= \Theta(v_i^{out} - v_{th}) \\ P_i^{out} &= \alpha_v P_i^{out} + u^{err} \\ v_i^{out} &= P_i^{out} + R_i^{out} + w_{tgt} \hat{P}_i. \end{aligned} \quad (11)$$

During training, target spike signals are sent to both the distal compartment and the proximal compartment. The distal compartment receives positively weighted target spikes to adjust the weight of attached synaptic connections towards a desired class, while the proximal compartment receives equally negative weighted target spikes to cancel out the effect of the label spikes have on the current integrated from the attached distal compartment. Targets are not provided during the test phase. Error encoding and propagation to the proximal compartment takes place in the distal compartment without the need for separate phases. Classification accuracy is measured using the spike count in the proximal compartments of the output layer neurons.

III. EXPERIMENTAL EVALUATION

A. Setup

The IBM DvsGesture[17] dataset consists of recordings of 29 different individuals performing 10 different actions such as clapping and an unspecified gesture for a total of 11 classes. The actions are recorded using a DVS camera, an event-based neuromorphic sensor, under three different lighting conditions. The problem is to classify an action sequence video. Samples from the first 23 subjects were used for training and the last 6 subjects were used for testing. The test set contains 264 samples and will be used for the few-shot learning task. Each sample consists of the first 1.45 seconds of the gesture performed. The SNN model trained on the Loihi for DvsGesture classification is described in table I.

TABLE I
NETWORK ARCHITECTURE

Layer	Kernel	Output	Training Method
input	1a	128x128x2	DVS128 (Sensor)
1	4a	32x32x2	SLAYER (BPTT)
2	16c5z	32x32x16	
3	2a	16x16x16	
4	32c3z	16x16x32	
5	2a	8x8x32	
6	-	512	
out	-	11	Loihi Plasticity Rule

The hidden layers 1 through 6 were trained offline using SLAYER [7] on the Loihi neuron model. The output layer synapse weights are trained using the three-factor surrogate gradient plasticity rule described in Sec. II-C.

We compare our model to the one used in [7]. To assess [7], the output layer weights are reset and retrained in few-shot learning conditions. Due to the limited bit precision of the synaptic weights in Loihi (256 levels), the effective learning rate is typically higher than that used in GPUs. To emulate this, we trained the output layer of the SLAYER model using a higher learning rate than the original

for a better comparison. In our few-shot learning experiments, using a higher learning rate achieves a higher accuracy because the training data is small and presented only once. We also compare our results to a model trained using SLAYER with hidden layers pre-trained SLAYER network but with the output layer replaced with a fully connected linear layer trained using the MSE of the membrane potentials of the neurons instead of classification with spike counts.

Because the hidden layers of each model are pre-trained on the DvsGesture training data, we train and test the models on the DvsGesture test dataset. Using the test dataset we set up three N-way classification tasks with 1, 5, and 14 shots. N-way classification involves selecting N unseen classes, provide the model with K different instances, also known as shots, of each of the N classes, and evaluate the model's ability to classify new instances within the N classes. The test data is split into separate training and testing sets. The training set consists of N-K samples taken from the first 154 samples of the DvsGesture Test dataset. During training, we present the data of the K shots for a single learning epoch. The N-way classification models are then tested on the last 110 samples of the DvsGesture Test dataset not used during training.

Additionally, we set up three M+N-way classification tasks with 1, 5, and 20 shots. For M+N-way classification, we pre-trained the SLAYER network offline using only M classes of the DvsGesture Training dataset. We then reset the output layer weights and retrain them using the surrogate gradient plasticity rule on Loihi on N unseen classes, provide the model with K shots of each of the N classes, and evaluate the model's ability to classify new instances within the N classes. The M+N-way classification models are trained using the DvsGesture Training dataset, and tested using the DvsGesture Test dataset.

B. Results

We present the N-way classification results in Table II, and the M + N-way classification results in Table III. Each model is trained with K shots presented a single time, and then tested on samples of each class held out from training. We use $N = 11$ for the number of classes for the N-way classification experiments, which is the number of classes in DvsGesture. For the M + N-way classification experiment, we use $M = 6$ for the number of classes pre-trained offline using SLAYER, and $N = 5$ for the number of classes trained on the network implemented on Loihi using the surrogate gradient plasticity rule. "Train" refers to classification accuracy on training samples using saved weights with plasticity disabled. "Test" is the classification accuracy on the samples held out of the training procedure. For each type of model, the same training and testing data samples were used. The SLAYER+Linear section of Table II shows the results of the model consisting of a pre-trained SLAYER hidden layer connected to an ANN Linear output layer that uses MSE loss, measuring the error of target versus output membrane potential quantities rather than a classification accuracy using the number of spikes over the presentation time.

The surrogate gradient plasticity rule implemented on Loihi outperforms the purely SLAYER model on few-shot learning tasks in both the 11-way classification experiments, and the 6+5-way classification experiments. In the 11-way few-shot classification experiments, the plasticity rule achieved 32.1% better accuracy at test time than the SLAYER rule of the SLAYER+Spiking model for one-shot learning, with the plasticity rule achieving better accuracy on the 5 and 14 shot learning tasks as well. Interestingly, the SLAYER+Linear model achieved 21.7% better test accuracy than the SLAYER+Spiking model on one-shot learning, but performed poorer on the 5 and 14 shot learning tasks than the SLAYER+Spiking model. These results indicate that the plasticity rule augmented with the pre-trained features does not need large amounts of training data to generalize. The performance of the plasticity rule in the 6 + 5-way experiments was not as high as the 11-way experiments because the 5 unseen classes were unseen by the entire network and not just unseen by the output layer, but did outperform the 6 + 5-way SLAYER models. The plasticity rule achieved 24.5% better accuracy at test time than

TABLE II
11-WAY FEW-SHOT CLASSIFICATION ON THE DVSGesture DATASET

Dataset	Learning Method	Shots	Train	Test
DVSGesture	Loihi Plasticity Rule	1	100%	52.2%
		5	86.6%	56.8%
		14	72.2%	64.7%
	SLAYER+Spiking	1	9.1%	20.1%
		5	40%	50.1%
		14	56.5%	61.8%
	SLAYER+Linear	1	<1%	41.8%
		5	38.2%	42.7%
		14	53.9%	51.8%

TABLE III
6+5-WAY FEW-SHOT CLASSIFICATION ON THE DVSGesture DATASET

Dataset	Learning Method	Shots	Train	Test
DVSGesture	Loihi Plasticity Rule	1	40%	40%
		5	60%	43.3%
		20	73.5%	56.2%
	SLAYER+Spiking	1	40%	35%
		5	20%	41.7%
		20	37%	31.7%
	SLAYER+Linear	1	40%	33%
		5	36%	27.5%
		20	35%	33%

the SLAYER rule of the SLAYER+Spiking model and 23.2% better accuracy than the SLAYER+Linear model for 20 shot learning, with the plasticity rule achieving better accuracy on the 1 and 5 shot learning tasks as well. These results indicate that the plasticity rule augmented with features pre-trained on some classes but not all can generalize to completely new classes of data.

IV. CONCLUSION

The results achieved in the few-shot learning experiment demonstrate the building blocks for online multi-layer learning and transfer learning in neuromorphic hardware. Neuromorphic hardware that can achieve high accuracy on only a few observations is a long-sought feature in applications where large datasets are nonexistent, cannot be accessed or where energy is limited, such as in robotic and embedded IoT systems. However, there are still challenges for achieving highly accurate continual "life-long" learning in neuromorphic hardware such as catastrophic forgetting [4]. Future work will address these challenges by exploring synaptic consolidation dynamics [18], [19] and reducing the variance of the learning when no error is made (problem (b) in the text). The solution described here trains the final layer only and thus foregoes the deep credit assignment problem. Our ongoing and future work is addressing the deep credit assignment using the technique of local losses [15], thus allowing to train multiple layers simultaneously. As the plasticity rule is improved, the applicability is likely to extend beyond classification, namely unsupervised learning and reinforcement learning tasks.

V. ACKNOWLEDGMENTS

This work was supported by Intel Corporation (KMS); the National Science Foundation under grant 1652159 (EON); the National Science Foundation under grant 1640081 (EON), and the Nanoelectronics Research Corporation (NERC), a wholly-owned subsidiary of the Semiconductor Research Corporation (SRC), through Extremely Energy Efficient Collective Electronics (EXCEL), an SRC-NRI Nanoelectronics Research Initiative under Research Task ID 2698.003 (EON, KMS).

REFERENCES

- [1] E. O. Neftci, H. Mostafa, and F. Zenke, "Surrogate gradient learning in spiking neural networks: Bringing the power of gradient-based optimization to spiking neural networks," *IEEE Signal Processing Magazine*, vol. 36, no. 6, pp. 51–63, Nov 2019.

- [2] W. Gerstner, W. M. Kistler, R. Naud, and L. Paninski, *Neuronal dynamics: From single neurons to networks and models of cognition*. Cambridge University Press, 2014.
- [3] E. O. Neftci, "Data and power efficient intelligence with neuromorphic learning machines," *iScience*, vol. 5, pp. 52–68, 2018. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S2589004218300865>
- [4] J. L. McClelland, B. L. McNaughton, and R. C. O'Reilly, "Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory." *Psychological review*, vol. 102, no. 3, p. 419, 1995.
- [5] J. Yosinski, J. Clune, Y. Bengio, and H. Lipson, "How transferable are features in deep neural networks?" in *Advances in neural information processing systems*, 2014, pp. 3320–3328.
- [6] C. Finn, P. Abbeel, and S. Levine, "Model-agnostic meta-learning for fast adaptation of deep networks," in *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. JMLR. org, 2017, pp. 1126–1135.
- [7] S. B. Shrestha and G. Orchard, "Slayer: Spike layer error reassignment in time," in *Advances in Neural Information Processing Systems 31*, S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, Eds. Curran Associates, Inc., 2018, pp. 1412–1421. [Online]. Available: <http://papers.nips.cc/paper/7415-slayer-spike-layer-error-reassignment-in-time.pdf>
- [8] F. Zenke and S. Ganguli, "Superspike: Supervised learning in multi-layer spiking neural networks," *arXiv preprint arXiv:1705.11146*, 2017.
- [9] A. Destexhe, Z. Mainen, and T. Sejnowski, *Methods in Neuronal Modelling, from ions to networks*. MIT Press, 1998, ch. Kinetic Models of Synaptic Transmission, pp. 1–25.
- [10] M. Davies, N. Srinivasa, T. H. Lin, G. Chinya, P. Joshi, A. Lines, A. Wild, and H. Wang, "Loihi: A neuromorphic manycore processor with on-chip learning," *IEEE Micro*, vol. PP, no. 99, pp. 1–1, 2018.
- [11] R. Urbanczik and W. Senn, "Learning by the dendritic prediction of somatic spiking," *Neuron*, vol. 81, no. 3, pp. 521–528, 2014.
- [12] G. Detorakis, S. Sheik, C. Augustine, S. Paul, B. U. Pedroni, N. Dutt, J. Krichmar, G. Cauwenberghs, and E. Neftci, "Neural and synaptic array transceiver: A brain-inspired computing framework for embedded learning," *Frontiers in Neuroscience*, vol. 12, p. 583, 2018. [Online]. Available: <https://www.frontiersin.org/article/10.3389/fnins.2018.00583>
- [13] W. Schultz, "Getting formal with dopamine and reward," *Neuron*, vol. 36, no. 2, pp. 241–263, 2002.
- [14] J.-P. Pfister, T. Toyozumi, D. Barber, and W. Gerstner, "Optimal spike-timing-dependent plasticity for precise action potential firing in supervised learning," *Neural computation*, vol. 18, no. 6, pp. 1318–1348, 2006.
- [15] J. Kaiser, H. Mostafa, and E. Neftci, "Synaptic plasticity for deep continuous local learning," *arXiv preprint arXiv:1812.10766*, 2018.
- [16] E. O. Neftci, C. Augustine, S. Paul, and G. Detorakis, "Event-driven random back-propagation: Enabling neuromorphic deep learning machines," *Frontiers in Neuroscience*, vol. 11, p. 324, 2017.
- [17] A. Amir, B. Taba, D. Berg, T. Melano, J. McKinstry, C. Di Nolfo, T. Nayak, A. Andreopoulos, G. Garreau, M. Mendoza *et al.*, "A low power, fully event-based gesture recognition system," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 7243–7252.
- [18] J. Kirkpatrick, R. Pascanu, N. Rabinowitz, J. Veness, G. Desjardins, A. A. Rusu, K. Milan, J. Quan, T. Ramalho, A. Grabska-Barwinska *et al.*, "Overcoming catastrophic forgetting in neural networks," *Proceedings of the national academy of sciences*, p. 201611835, 2017.
- [19] M. K. Benna and S. Fusi, "Computational principles of biological memory," *arXiv preprint arXiv:1507.07580*, 2015.