

hxtorch.snn: Machine-learning-inspired Spiking Neural Network Modeling on BrainScaleS-2

Philipp Spilger *
Kirchhoff-Institute for Physics,
Heidelberg University
Heidelberg, Germany
pspilger@kip.uni-heidelberg.de

Elias Arnold*
Kirchhoff-Institute for Physics,
Heidelberg University
Heidelberg, Germany

Luca Blessing
Kirchhoff-Institute for Physics,
Heidelberg University
Heidelberg, Germany

Christian Mauch
Kirchhoff-Institute for Physics,
Heidelberg University
Heidelberg, Germany

Christian Pehle
Kirchhoff-Institute for Physics,
Heidelberg University
Heidelberg, Germany

Eric Müller
European Institute for Neuromorphic
Computing, Heidelberg University
Heidelberg, Germany

Johannes Schemmel
Kirchhoff-Institute for Physics,
Heidelberg University
Heidelberg, Germany

ABSTRACT

Neuromorphic systems require user-friendly software to support the design and optimization of experiments. In this work, we address this need by presenting our development of a machine learning-based modeling framework for the BrainScaleS-2 neuromorphic system. This work represents an improvement over previous efforts, which either focused on the matrix-multiplication mode of BrainScaleS-2 or lacked full automation. Our framework, called `hxtorch.snn`, enables the hardware-in-the-loop training of spiking neural networks within PyTorch, including support for auto differentiation in a fully-automated hardware experiment workflow. In addition, `hxtorch.snn` facilitates seamless transitions between emulating on hardware and simulating in software. We demonstrate the capabilities of `hxtorch.snn` on a classification task using the Yin-Yang dataset employing a gradient-based approach with surrogate gradients and densely sampled membrane observations from the BrainScaleS-2 hardware system.

KEYWORDS

hardware abstraction, modeling, accelerator, analog computing, neuromorphic

1 INTRODUCTION

Modern high-performance computing (HPC) environments speed up individual workloads with domain-specific hardware accelerators [9]. However, successfully establishing novel computing paradigms depends as much on the availability of a computing substrate as it does on the software infrastructure that provides the means to work with that substrate. In particular, the ‘programming model’ of

neuromorphic hardware is a very different approach to data processing than conventional systems [22, 23]. Recently, machine-learning-inspired training methods, and in particular gradient-based optimization of spiking neural networks (SNNs), have become increasingly popular [4, 13, 14, 24, 32].

We introduce `hxtorch.snn`, a PyTorch [25] wrapper library for the accelerated mixed-signal neuromorphic hardware system BrainScaleS-2 (BSS-2) [26]. We discuss and demonstrate our implementation using a simple benchmark task trained with a gradient-based method and the BSS-2 hardware system in the loop. The implementation described here has also been used in a real-world application [2] and will serve as the basis for future experiments.

Here we use the gradient estimation method introduced in [8], which uses the unrolled computational graph determined by the forward pass, to perform gradient estimation by injecting hardware voltage measurements and spikes in the backward pass. The overall library, however, is agnostic to this particular choice of gradient estimation method. Other approaches can be implemented, such as using a closed-form analytical formula [16] or the EventProp algorithm [38]. To support this generic framework for gradient estimation, we implement utilities for interpolation and normalization of hardware measurements, conversion of spike observations into dense tensors, and procedures for weight quantization.

Existing approaches to interfacing with Neuromorphic hardware tend to favor abstractions familiar to neuro-scientists, such as populations and projections [3, 10, 20, 28, 29] and only a few supported gradient-based methods [17, 21]. Recently several libraries for machine learning with spiking neurons in PyTorch have appeared, among them, Norse [27], Lava [18], NxTF [30], `snnTorch` [12], and `SpikingJelly` [15].

The BSS-2 mixed-signal neuromorphic hardware emulates networks of spiking neurons time-continuously in analog circuits [26]. Each of the 512 adaptive exponential integrate-and-fire (AdEx) [7] neuron compartments on BSS-2 are individually parameterized and, in particular, can be configured to resemble leaky-integrate and

*Contributed equally.

fire (LIF) and leaky integrator (LI) neuron dynamics. They receive input stimuli from a column of 256 exponential synapses in the synapse matrix with adjustable 6-bit weights. Configuration of the on-chip routing mechanism and the synapse matrix allows addressing the neurons’ digital spike events to target synapses and, thus, the realization of arbitrary topologies. In addition, artificial neural networks (ANNs) can be emulated by using the neurons in a non-spiking mode, which allows implementing analog vector-matrix multiplications [34, 37]. The recordings of spike events and membrane potentials, sampled in parallel via a columnar ADC (CADC), are accessible by the host computer and can be incorporated, e.g., for weight update computation in a hardware-in-the-loop (ITL) [31] fashion. A field-programmable gate array (FPGA) serves as the real-time experiment master and as interconnect between the neuromorphic chip and a host computer.

We use a layered software approach to abstract hardware usage [22]. At the lowest level, a transport layer and an instruction set of commands to the FPGA, which execute sequentially, are employed. Above, individual hardware entity configuration, e.g., for each neuron, is abstracted into containers residing in a uniform address space for type safety and to remove requiring knowledge about the memory layout. Using this hardware abstraction, a data-flow-graph-based experiment notation, *grenade*, separates network topology specification, temporal evolution description, and execution to support the accelerated real-time nature of the neuromorphic hardware. On top, front ends for spiking and non-spiking experiments exist. Spiking neural networks can be described in the back-end-agnostic language PyNN [10] targeting computational neuroscience. For offloading ANNs, i.e. analog vector-matrix multiplications, a thin PyTorch machine-learning adapter, *hxtorch*, was developed [33], separating hardware interaction and machine-learning framework. Similarly, this work adds support for describing SNNs in this machine-learning framework yielding access to training, inference, and data alteration integrated into PyTorch’s ecosystem.

2 METHODS

For developing a machine-learning framework supporting the emulation of SNNs on BSS-2, we follow the same structure as for the developed ANN machine-learning adaptor *hxtorch* [33]. Due to its vast community, we continue to build our front end upon the PyTorch framework. Given the accelerated time-continuous nature of the neuromorphic hardware, network topology specification, temporal evolution and execution need to be separated in the machine-learning front end. In particular, individual numerical operations in a time-grid-based simulation cannot be trivially identified with dynamics on the neuromorphic substrate.

We use the data-flow-graph-based experiment notation [22] as the back end. The machine-learning front end then needs to handle data conversion to and from PyTorch tensors and wrapping of network topology descriptions in a PyTorch-compatible notation. Therefore, hardware interaction and machine-learning framework adaptation are separated.

In PyTorch, models are eagerly executed by the computation of a layer’s result when calling its forward method, which implies that the network builds up incrementally. However, the emulation of

SNNs in physical time requires knowledge of the complete network topology before experiment execution, to derive a corresponding hardware configuration. Therefore, an additional separation is required.

Since training neural networks benefits from the auto differentiation provided by the machine-learning framework, we aim to incorporate this into the network description. Within PyTorch, we train with the hardware in the loop by emulating the forward pass on BSS-2 and injecting its results in the backward pass on the host computer. Moreover, the convenience of describing models by a composition of modules, i.e. network layers, should be maintained and allows for defining network topologies with recurrently connected layers.

For developing a model for the neuromorphic hardware from scratch or a reference model, iterative development is beneficial. Exchanging a simulated model with emulation on BSS-2 requires parameter translations, e.g., weight scaling and adaptation to limited resolution on hardware or neuron model parameter translation to technical hardware parameters. A seamless comparison to and replacement of (parts of) the network by (custom) simulation facilitates this. It also enables hybrid networks, for example, by adding the possibility to introduce layers not representable on the neuromorphic hardware, e.g., neuron models not present on hardware or other arbitrary numerical operations on intermediate results, while retaining the neuromorphic hardware’s acceleration for the representable parts.

Lastly, the collection of hardware observables, such as membrane potential recordings or spike trains, shall be adjustable per layer depending on the training algorithms. Importantly, the transformation of acquired hardware observables into PyTorch data structures needs to be fully customizable to be best-tailored for the training algorithm at hand, e.g., the transition of sparse time-series data to a dense time grid, the interpolation of sparse membrane measurements, or the mapping of synapse weight parameters between PyTorch and BSS-2 can impact the network’s dynamics as well as the learning process.

3 RESULTS

The software library *hxtorch.snn* integrates modeling SNNs on BSS-2 into the PyTorch ecosystem. It extends the thin wrapper library *hxtorch*, which previously only targeted modeling ANNs.

Figure 1 shows the library structure, experiment description, and constitution. Similarly to ANN models in PyTorch, we describe constituents of a SNN model like synapse and neuron layers via custom modules derived from *HXModule*, containing parameters and representing the corresponding hardware entity. We want to retain PyTorch’s eager model construction and execution interface by successively applying modules’ `forward(input)` calls, but we need to separate the construction of the complete model from its execution on the hardware. Therefore, we only register each module’s invocation into an *Instance*, passed to each module on its construction. Instead of returning the actual results, we return *Handle*-typed promises of future results. The instance is then separately executable, and the promises are filled with the result data of the respective modules after the instance’s execution. An instance

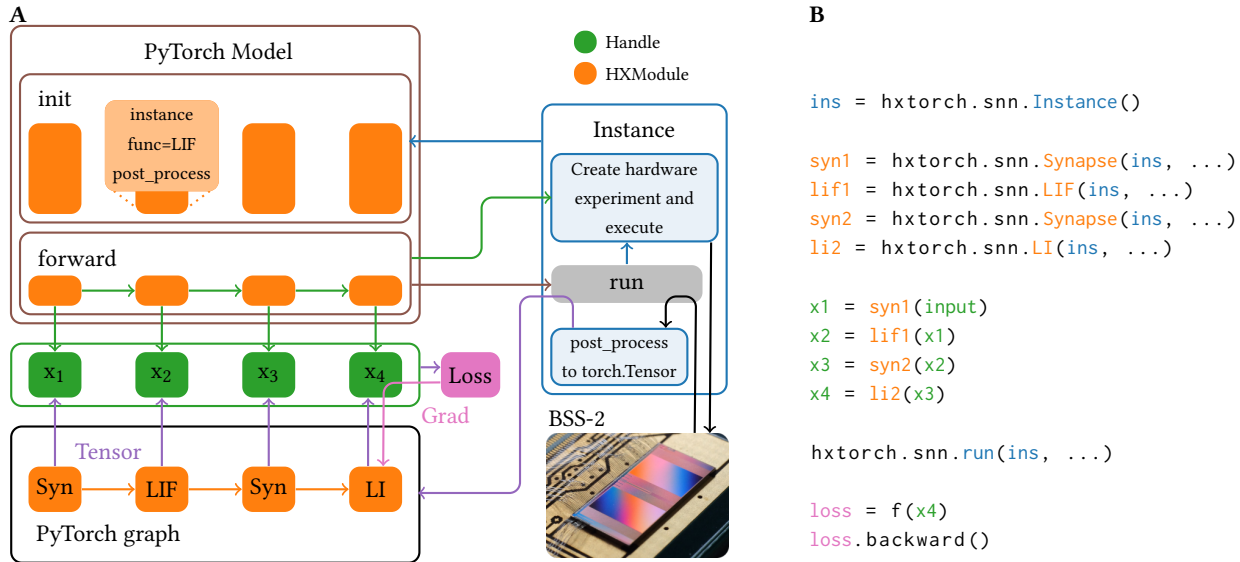


Figure 1: Software application programming interface (API) of `hxtorch.snn` shown exemplarily for the feed-forward network used in the classification of the Yin-Yang dataset [19], where A depicts the data flow in the model and with BSS-2 and B displays the corresponding source code. The network consists of two synapse layers and two LI/LIF neuron layers and is associated to an instance to be executed on hardware. Execution is triggered explicitly via the `run` function and the handles `x`, `y` and `loss` only afterwards carry their result data from the execution. This then in particular allows gradient calculation via backward functions attached to the handles’ data.

represents one hardware experiment execution. Upon explicit execution of an instance via `run`, the network topology is extracted from the registered modules and their invocations and converted into a graph-based description, which is then mapped to a hardware experiment description and executed on hardware. The resulting data is transformed into PyTorch tensors, post-processed via a customizable method, and finally annotated onto the corresponding data handles, thereby being accessible by the user. Since instances of `Instance` only require hardware allocations upon execution, multiple instances can coexist and be executed sequentially. This allows interleaving hardware-executed model parts and simulated parts as long as no inter-instance recurrence is required.

The model construction via network entity module invocations allows defining recurrent models by associating each module instance with a network entity and using invocations to connect them, cf. listing 1.

Listing 1: Construction of a recurrent network. When reusing the same network entity `nrn`, a recurrent connection is created.

```

nrn = hxtorch.snn.LIF(...)
syn1 = hxtorch.snn.Synapse(...)
syn2 = hxtorch.snn.Synapse(...)

x = syn1(input)
x = nrn(x) # feed-forward
x = syn2(x)
x = nrn(x) # recurrence

```

Upon construction, each module is equipped with a custom PyTorch-differentiable function, either directly as `Autograd.Function` or via a function defining both a simulated forward pass and an implicit backward pass. In the first case, `hxtorch.snn` injects the hardware data implicitly to be used in the backward function, and in the second case, hardware observables are provided as an additional function argument. This allows annotating a backward pass to the handle data tensors such that PyTorch’s auto differentiation can seamlessly backpropagate gradients based on hardware observations. When a simulated forward pass is provided, this module can also be used without hardware usage as long as it does not participate in a cyclic sub-network. This greatly aids network development and translation onto the hardware, since (parts of) a simulated reference network can gradually be interchanged with and compared to hardware entities within the same library.

Currently, `hxtorch.snn` supports LIF and LI neuron layers and has access to spike times and membrane voltage recordings of individual hardware neurons. To support the full extent of BSS-2 neuron dynamics, a forward and backward PyTorch implementation of the AdEx neuron model is planned.

A dropout module provides additional machine learning functionality by applying a batch-wise spiking mask to a preceding neuron layer and disabling the spike output on hardware accordingly.

Synaptic connections on BSS-2 are exposed to PyTorch as a `Synapse` module, representing a projection between neuron layers. This provides a framework for recording of additional synapse-related observables, such as spike-time correlation sensor recordings.

section	duration [s]	rel. duration [%]
network emulation duration	0.3	1.0
additional hardware runtime	7.5	26.3
additional back-end overhead	4.7	16.5
data transform to PyTorch	13.1	45.7
additional front-end overhead	2.5	8.6
gradient calculation	0.8	2.9
total duration	28.7	100.0

Table 1: Runtime performance of a training epoch of the Yin-Yang experiment consisting of 64 batches to 75 samples using a host computer with an AMD Ryzen 7 3800X central processing unit (CPU). While the emulated network duration falls two orders of magnitude behind the total training runtime, the majority of the additional hardware and back-end time consists of the analog-to-digital converter (ADC) neuron membrane potential readout. The data transformation to PyTorch tensors is dominated by interpolation of the sparse-in-time ADC data onto a dense time grid. The additional front-end overhead and gradient calculation don’t contribute significantly to the total training duration.

As a demonstration of our framework we consider the low-dimensional Yin-Yang classification task [19] visualized in Figure 2A. It consists of three classes *yin*, *yang*, and *dots*, each defining an area on the 2-dimensional xy -plane. Its i -th sample is a point $\mathbf{x}_i = (x_i, y_i)$, randomly drawn from the corresponding areas and labeled accordingly. As input to the SNN, the sample \mathbf{x}_i is translated to spike times of five input neurons [16, 19]. As shown in Figure 2B the point’s values, x_i and y_i , in each direction, as well as their inverse, $1 - x_i$ and $1 - y_i$, are scaled linearly to spike times in the time window $[t_{\text{early}}, t_{\text{late}}]$. In addition, a bias spike shortly after t_{early} is inserted to increase the network activity which is observed for the duration T . A hidden layer constitutes 120 LIF neurons in our SNN on BSS-2 integrates the input events and projects spikes itself onto a readout layer of 3 LI neurons, each corresponding to one class, as indicated by Figure 2A. This allows it to infer a class decision k_i for sample \mathbf{x}_i by interpreting the maximum membrane value of the output neurons v_k over time as a score for the corresponding class, i.e. $k = \operatorname{argmax}_k (\max_t v_k(t))$. For each sample, the SNN is emulated for $T = 60 \mu\text{s}$. We use the cross-entropy loss on the max-over-time values $s_k = \max_t v_k$ as the objective function, and SuperSpike [24] surrogate gradients for training. In `hxtorch.snn`, the PyTorch model is defined as outlined in Figure 1B. The upper plot in Figure 2D exemplifies the output traces while inferring a single sample. The membrane voltage of the output neuron corresponding to the sample’s class has the maximum value over time. The lower plot depicts the achieved accuracy of the SNN on BSS-2 over epochs; it reaches about $94.63 \pm 0.7\%$ (standard deviation over 15 seeds).

4 DISCUSSION

We presented the library `hxtorch.snn` to describe SNN models and their emulation on BSS-2 in the PyTorch ecosystem. It integrates

the auto-differentiation capabilities of PyTorch, providing access to the same training methodology as for ANNs. Separation of network construction and emulation allows taking full advantage of the emulation speed-up of BSS-2 and the description of feed-forward and recurrent network topologies. The API is designed for user-defined extension of neuron and synapse types, including customization of the backward pass while removing the need for expert knowledge of the hardware. Moreover, we showcased the library by using it to classify the YinYang dataset [19].

The Yin-Yang samples are encoded in spike events as proposed in [19] and fed into an SNN with one hidden LIF layer followed by a LI readout layer. When emulated on BSS-2, the SNN achieves a classification accuracy of $94.63 \pm 0.7\%$ when using a max-over-time loss. Compared to an accuracy of $95.0 \pm 0.9\%$ achieved on BSS-2 in [16], where the authors use spiking output neurons and an analytical solution for the gradient to optimize the model, our accuracy is only marginally lower. This verifies the design choices of our machine learning layer `hxtorch.snn` and demonstrates the implementation successfully.

The source code for the demonstration experiment, `hxtorch.snn`, and the underlying open-source software stack is available online [11, 36]. As part of the EBRAINS research infrastructure [1], we operate BSS-2 as a service allowing researchers to interactively conduct experiments and research on the neuromorphic platform. In [35] we provide and maintain a collection of experiments as interactive learning material for the platform.

Table 1 evaluates the experiment runtime overhead of the software. The surrogate-gradient-based training is dominated by data transformations, mostly from membrane measurement data. However, more data-sparse training algorithms, such as EventProp [38], can avoid much of this conversion overhead. To further optimize performance, adding support for event-based numerical codes would also be beneficial, as it would allow us to avoid converting to a time grid. Here, for example, a port of `hxtorch.snn` to `jax` [6] could offer advantages based on increased flexibility yielding opportunities for more efficient data structures and codes.

The backward functions currently assigned to each module represent the entire backward dynamics of the module over time. Since this is a constraint for recurrent topologies, we aim to incorporate algorithms and interfaces that allow to either assign a single backward function to a defined sub-network (and forward in mock mode) or, for a defined time grid, assign each module a function representing one integration step and then build the backward graph by implicit time-unrolling in an outer loop (and forward in mock mode). Additionally, we plan to provide support for online learning rules like the e-prop learning method [5] involving code generation for the embedded single instruction, multiple data (SIMD) microprocessors on BSS-2.

While currently one Instance instance corresponds to a single hardware experiment, we want to support mixed networks that are only partially executed on hardware by implicitly identifying the sub-networks to run on BSS-2 and execute the remaining parts in software. In addition, we will implement the ability to execute layers partially in order to allow the use of layer sizes that cannot be mapped to BSS-2 due to limited neuron resources. This will enable the pipelined and parallel execution of large layers on multiple chips. Partial functionality of our `hxtorch.snn` will be consolidated




Figure 2: (A) Example samples from the Yin-Yang dataset [19]. The dataset consists of samples defined on the xy -plane, each assigned to one of three nonlinear separable classes yin (orange), yang (blue), and dots (green). (B) The spike encoding of the 2D point from the Yin-Yang dataset in A depicted as a red star. Each dimension of the point and an inverse of it is translated to a spike event. Additionally, a bias spike is added to increase network activity. (C) The SNN topology used to classify the Yin-Yang dataset. An input layer projects the spike-encoded point onto a hidden layer consisting of LIF neurons. A LI readout layer receives the spike events of the hidden layer. Their maximal membrane values over time are used to infer a decision. (D) The upper plot depicts the analog output membrane voltages while inferring the red sample in A. In the lower plot the classification accuracy is shown over the test epochs. On BSS-2, the SNN achieves an accuracy of $94.63 \pm 0.7\%$. The code is available as an interactive demo [36].

in Norse [27], and will serve as the mock backend, provide spike encoding and decoding schemes, and helper functions.

Overall, `hxtorch.snn` is a convenient and flexible tool for implementing and training SNN models on BSS-2, utilizing the auto-differentiation capabilities of PyTorch to enable machine-learning-inspired training methods. With further optimization and the support for the upcoming multi-chip systems, it has the potential to greatly facilitate and accelerate machine-learning-inspired SNN research and experimentation.

ACKNOWLEDGMENTS

The authors would like to thank the BrainScaleS software team for fruitful discussions on software design, the participants of the BrainScaleS-2 user meeting for their continuous support, especially Sebastian Billaudelle and Johannes Weis for sharing their hardware expertise, and all other current and former members of the Electronic Vision(s) research group who contributed to the BrainScaleS-2 neuromorphic platform.

This work has received funding from the EC Horizon 2020 Framework Programme under grant agreements 785907 (HBP SGA2) and 945539 (HBP SGA3), the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy EXC 2181/1-390900948 (the Heidelberg STRUCTURES Excellence Cluster), the German Federal Ministry of Education and Research under grant number 16ES1127 as part of the *Pilotinnovationswettbewerb ‘Energieeffizientes KI-System’*, from the Manfred Stärk Foundation, and from the Lautenschläger-Forschungspreis 2018 for Karlheinz Meier.

AUTHOR CONTRIBUTIONS

We give contributions in the *CRedit* (Contributor Roles Taxonomy) format: **PS & EA:** Conceptualization, visualization, methodology,

software, resources, writing – original draft, writing – reviewing & editing; **LB:** Investigation, validation, visualization, writing – original draft, writing – reviewing & editing; **CM:** Methodology, software, resources, writing – original draft, writing – reviewing & editing; **CP:** Conceptualization, methodology, writing – original draft, writing – reviewing & editing; **EM:** Conceptualization, methodology, software, resources, writing – original draft, writing – reviewing & editing, supervision; **JS:** Supervision, funding acquisition, writing – reviewing & editing.

REFERENCES

- [1] 2022. EBRAINS Research Infrastructure. <https://ebrains.eu>
- [2] Elias Arnold, Georg Böcherer, Eric Müller, Philipp Spilger, Johannes Schemmel, Stefano Calabrò, and Maxim Kuschnerov. 2022. Spiking Neural Network Equalization on Neuromorphic Hardware for IM/DD Optical Communication. In *European Conference on Optical Communication (ECOC) 2022*. Optica Publishing Group, Th1C.5. <https://opg.optica.org/abstract.cfm?URI=ECEOC-2022-Th1C.5>
- [3] Trevor Bekolay, James Bergstra, Eric Hunsberger, Travis DeWolf, Terrence C Stewart, Daniel Rasmussen, Xuan Choo, Aaron Voelker, and Chris Eliasmith. 2014. Nengo: a Python tool for building large-scale functional brain models. *Frontiers in Neuroinformatics* 7 (2014), 48.
- [4] Guillaume Bellec, Darjan Salaj, Anand Subramoney, Robert Legenstein, and Wolfgang Maass. 2018. Long short-term memory and learning-to-learn in networks of spiking neurons. *Advances in neural information processing systems* 31 (2018).
- [5] Guillaume Bellec, Franz Scherr, Anand Subramoney, Elias Hajek, Darjan Salaj, Robert Legenstein, and Wolfgang Maass. 2020. A solution to the learning dilemma for recurrent networks of spiking neurons. *Nature Communications* 11, 1 (2020), 3625. <https://doi.org/10.1038/s41467-020-17236-y>
- [6] James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang. 2022. *JAX: composable transformations of Python+NumPy programs*. <http://github.com/google/jax>
- [7] R. Brette and W. Gerstner. 2005. Adaptive Exponential Integrate-and-Fire Model as an Effective Description of Neuronal Activity. *J. Neurophysiol.* 94 (2005), 3637 – 3642. <https://doi.org/10.1152/jn.00686.2005>
- [8] Benjamin Cramer, Sebastian Billaudelle, Simeon Kanya, Aron Leibfried, Andreas Grübl, Vitali Karasenko, Christian Pehle, Korbinian Schreiber, Yannik Stradmann, Johannes Weis, et al. 2022. Surrogate gradients for analog neuromorphic computing. *Proceedings of the National Academy of Sciences* 119, 4 (2022).

- [9] William J. Dally, Yatish Turakhia, and Song Han. 2020. Domain-Specific Hardware Accelerators. *Commun. ACM* 63, 7 (June 2020), 48–57. <https://doi.org/10.1145/3361682>
- [10] Andrew P. Davison, Daniel Bröderle, Jochen Eppler, Jens Kremkow, Eilif Müller, Dejan Pecevski, Laurent Perrinet, and Pierre Yger. 2009. PyNN: a common interface for neuronal network simulators. *Front. Neuroinform.* 2, 11 (2009). <https://doi.org/10.3389/fninf.2009.00111>
- [11] Electronic Vision(s) Group. 2022. Open-Source Software. <https://github.com/electronicvisions/>
- [12] Jason K Eshraghian, Max Ward, Emre Neftci, Xinxin Wang, Gregor Lenz, Girish Dwivedi, Mohammed Bannamoun, Doo Seok Jeong, and Wei D Lu. 2021. Training spiking neural networks using lessons from deep learning. *arXiv preprint* (2021). [arXiv:2109.12894](https://arxiv.org/abs/2109.12894) [cs.NE]
- [13] Steve K Esser, Rathinakumar Appuswamy, Paul Merolla, John V Arthur, and Dharmendra S Modha. 2015. Backpropagation for energy-efficient neuromorphic computing. *Advances in neural information processing systems* 28 (2015).
- [14] Steven K. Esser, Paul A. Merolla, John V. Arthur, Andrew S. Cassidy, Rathinakumar Appuswamy, Alexander Andreopoulos, David J. Berg, Jeffrey L. McKinstry, Timothy Melano, Davis R. Barch, Carmelo di Nolfo, Pallab Datta, Arnon Amir, Brian Taba, Myron D. Flickner, and Dharmendra S. Modha. 2016. Convolutional networks for fast, energy-efficient neuromorphic computing. *Proceedings of the National Academy of Sciences* 113, 41 (2016), 11441–11446. <https://doi.org/10.1073/pnas.1604850113> <https://www.pnas.org/content/113/41/11441.full.pdf>
- [15] Wei Fang, Yanqi Chen, Jianhao Ding, Ding Chen, Zhaofei Yu, Huihui Zhou, Timothée Masquelier, Yonghong Tian, and other contributors. 2020. SpikingJelly. <https://github.com/fangwei123456/spikingjelly>.
- [16] Julian Göltz, Laura Kriener, Andreas Baumbach, Sebastian Billaudelle, Oliver Breitwieser, Benjamin Cramer, Dominik Dold, Ákos Ferenc Kungl, Walter Senn, Johannes Schemmel, Karlheinz Meier, and Mihai A. Petrovici. 2021. Fast and energy-efficient neuromorphic deep learning with first-spike times. *Nature Machine Intelligence* 3, 9 (2021), 823–835. <https://doi.org/10.1038/s42256-021-00388-x>
- [17] Hananel Hazan, Daniel J. Saunders, Hassaan Khan, Devdhar Patel, Darpan T. Sanghavi, Hava T. Siegelmann, and Robert Kozma. 2018. BindsNET: A Machine Learning-Oriented Spiking Neural Networks Library in Python. *Frontiers in Neuroinformatics* 12 (2018), 89. <https://doi.org/10.3389/fninf.2018.00089>
- [18] Intel. 2021. Announcement of Loihi-2 and new software framework. <https://www.intel.com/content/www/us/en/newsroom/news/intel-unveils-neuromorphic-loihi-2-lava-software.html> <https://github.com/lava-nc/>
- [19] Laura Kriener, Julian Göltz, and Mihai A. Petrovici. 2022. The Yin-Yang Dataset. In *Neuro-Inspired Computational Elements Conference* (Virtual Event, USA) (NICE 2022). Association for Computing Machinery, New York, NY, USA, 107–111. <https://doi.org/10.1145/3517343.3517380>
- [20] Chit-Kwan Lin, Andreas Wild, Gautham N China, Yongqiang Cao, Mike Davies, Daniel M Lavery, and Hong Wang. 2018. Programming Spiking Neural Networks on Intel’s Loihi. *Computer* 51, 3 (2018), 52–61.
- [21] Dylan R. Muir, Felix Bauer, and Philipp Weidel. 2019. Rockpool Documentation. <https://doi.org/10.5281/zenodo.3773845>
- [22] Eric Müller, Elias Arnold, Oliver Breitwieser, Milena Czierlinski, Arne Emmel, Jakob Kaiser, Christian Mauch, Sebastian Schmitt, Philipp Spilger, Raphael Stock, Yannik Stradmann, Johannes Weis, Andreas Baumbach, Sebastian Billaudelle, Benjamin Cramer, Falk Ebert, Julian Göltz, Joscha Ilmberger, Vitali Karasenko, Mitja Kleider, Aron Leibfried, Christian Pehle, and Johannes Schemmel. 2022. A Scalable Approach to Modeling on Accelerated Neuromorphic Hardware. *Front. Neurosci.* 16 (2022). <https://doi.org/10.3389/fnins.2022.884128>
- [23] Eric Müller, Sebastian Schmitt, Christian Mauch, Sebastian Billaudelle, Andreas Grübl, Maurice Güttler, Dan Husmann, Joscha Ilmberger, Sebastian Jeltsch, Jakob Kaiser, Johann Klähn, Mitja Kleider, Christoph Koke, José Montes, Paul Müller, Johannes Partzsch, Felix Passenberg, Hartmut Schmidt, Bernhard Vogginger, Jonas Weidner, Christian Mayr, and Johannes Schemmel. 2022. The Operating System of the Neuromorphic BrainScaleS-1 System. *Neurocomputing* 501 (2022), 793–810. <https://doi.org/10.1016/j.neucom.2022.05.081> [arXiv:2003.13749](https://arxiv.org/abs/2003.13749) [cs.NE]
- [24] Emre O. Neftci, Hesham Mostafa, and Friedemann Zenke. 2019. Surrogate gradient learning in spiking neural networks: Bringing the power of gradient-based optimization to spiking neural networks. *IEEE Signal Processing Magazine* 36, 6 (2019), 51–63. <https://doi.org/10.1109/MSP.2019.2931595>
- [25] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *Advances in Neural Information Processing Systems* 32, H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett (Eds.), Curran Associates, Inc., 8024–8035. <http://papers.nips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>
- [26] Christian Pehle, Sebastian Billaudelle, Benjamin Cramer, Jakob Kaiser, Korbinian Schreiber, Yannik Stradmann, Johannes Weis, Aron Leibfried, Eric Müller, and Johannes Schemmel. 2022. The BrainScaleS-2 Accelerated Neuromorphic System with Hybrid Plasticity. *Frontiers in Neuroscience* 16 (2022). <https://doi.org/10.3389/fnins.2022.795876> [arXiv:2201.11063](https://arxiv.org/abs/2201.11063) [cs.NE]
- [27] Christian Pehle and Jens Egholm Pedersen. 2021. *Norse — A deep learning library for spiking neural networks*. <https://doi.org/10.5281/zenodo.4422025> Documentation: <https://norse.ai/docs/>.
- [28] Oliver Rhodes, Petruț A. Bogdan, Christian Brenninkmeijer, Simon Davidson, Donal Fellows, Andrew Gait, David R. Lester, Mantas Mikaitis, Luis A. Plana, Andrew G. D. Rowley, Alan B. Stokes, and Steve B. Furber. 2018. sPyNNaker: A Software Package for Running PyNN Simulations on SpiNNaker. *Frontiers in Neuroscience* 12 (2018), 816. <https://doi.org/10.3389/fnins.2018.00816>
- [29] Andrew G. D. Rowley, Christian Brenninkmeijer, Simon Davidson, Donal Fellows, Andrew Gait, David R. Lester, Luis A. Plana, Oliver Rhodes, Alan B. Stokes, and Steve B. Furber. 2019. SpiNNTools: The Execution Engine for the SpiNNaker Platform. *Frontiers in Neuroscience* 13 (2019), 231. <https://doi.org/10.3389/fnins.2019.00231>
- [30] Bodo Rueckauer, Connor Bybee, Ralf Goettsche, Yashwardhan Singh, Joyesh Mishra, and Andreas Wild. 2021. NxTF: An API and Compiler for Deep Spiking Neural Networks on Intel Loihi. *arXiv preprint* (Jan. 2021).
- [31] Sebastian Schmitt, Johann Klähn, Guillaume Bellec, Andreas Grübl, Maurice Güttler, Andreas Hartel, Stephan Hartmann, Dan Husmann, Kai Husmann, Sebastian Jeltsch, Mitja Kleider, Christoph Koke, Alexander Kononov, Christian Mauch, Eric Müller, Paul Müller, Johannes Partzsch, Mihai A. Petrovici, Bernhard Vogginger, Stefan Schiefer, Stefan Scholz, Vasilis Thanasoulis, Johannes Schemmel, Robert Legenstein, Wolfgang Maass, Christian Mayr, and Karlheinz Meier. 2017. Neuromorphic Hardware In The Loop: Training a Deep Spiking Network on the BrainScaleS Wafer-Scale System. *Proceedings of the 2017 IEEE International Joint Conference on Neural Networks (IJCNN)* (2017), 2227–2234. <https://doi.org/10.1109/IJCNN.2017.7966125>
- [32] Sumit Bam Shrestha and Garrick Orchard. 2018. SLAYER: Spike Layer Error Reassignment in Time. In *Advances in Neural Information Processing Systems*, S. Bengio, H. Wallach, H. Larochelle, K. Grauman, M. Cesa-Bianchi, and R. Garnett (Eds.), Vol. 31. Curran Associates, Inc. <https://proceedings.neurips.cc/paper/2018/file/82f2b308c3b01637c607ce05f52afed-Paper.pdf>
- [33] Philipp Spilger, Eric Müller, Arne Emmel, Aron Leibfried, Christian Mauch, Christian Pehle, Johannes Weis, Oliver Breitwieser, Sebastian Billaudelle, Sebastian Schmitt, Timo C. Wunderlich, Yannik Stradmann, and Johannes Schemmel. 2020. hxtorch: PyTorch for BrainScaleS-2 — Perceptrons on Analog Neuromorphic Hardware. In *IoT Streams for Data-Driven Predictive Maintenance and IoT, Edge, and Mobile for Embedded Machine Learning*. Springer International Publishing, Cham, 189–200. https://doi.org/10.1007/978-3-030-66770-2_14
- [34] Yannik Stradmann, Sebastian Billaudelle, Oliver Breitwieser, Falk Leonard Ebert, Arne Emmel, Dan Husmann, Joscha Ilmberger, Eric Müller, Philipp Spilger, Johannes Weis, and Johannes Schemmel. 2022. Demonstrating Analog Inference on the BrainScaleS-2 Mobile System. *IEEE Open Journal of Circuits and Systems* 3 (2022), 252–262. <https://doi.org/10.1109/OJCS.2022.3208413>
- [35] BrainScaleS-2 Team. 2022. hxtorch.snn: Yin-Yang Demo. <https://electronicvisions.github.io/documentation-brainscales2>
- [36] BrainScaleS-2 Team. 2022. Source code for the hxtorch.snn Yin-Yang Demo. https://github.com/electronicvisions/hxtorch/tree/nice2023_yydemo/src/pyhxtorch/hxtorch/lit_yinyang
- [37] Johannes Weis, Philipp Spilger, Sebastian Billaudelle, Yannik Stradmann, Arne Emmel, Eric Müller, Oliver Breitwieser, Andreas Grübl, Joscha Ilmberger, Vitali Karasenko, Mitja Kleider, Christian Mauch, Korbinian Schreiber, and Johannes Schemmel. 2020. Inference with Artificial Neural Networks on Analog Neuromorphic Hardware. In *IoT Streams for Data-Driven Predictive Maintenance and IoT, Edge, and Mobile for Embedded Machine Learning*. Springer International Publishing, Cham, 201–212. https://doi.org/10.1007/978-3-030-66770-2_15
- [38] Timo C Wunderlich and Christian Pehle. 2021. Event-based backpropagation can compute exact gradients for spiking neural networks. *Scientific Reports* 11, 1 (2021), 1–17. <https://doi.org/10.1038/s41598-021-91786-z>