

The Backpropagation Algorithm Implemented on Spiking Neuromorphic Hardware

Alpha Renner

University of Zurich and ETH Zurich <https://orcid.org/0000-0002-0724-4169>

Forrest Sheldon

London Institute for Mathematical Sciences

Anatoly Zlotnik

Los Alamos National Laboratory

Louis Tao

Center for Bioinformatics <https://orcid.org/0000-0003-1802-2438>

Andrew Sornborger (✉ sornborg@lanl.gov)

Los Alamos National Laboratory

Article

Keywords: Spiking Neural Network, backpropagation algorithm, three-layer circuit

Posted Date: August 18th, 2021

DOI: <https://doi.org/10.21203/rs.3.rs-701752/v1>

License:  This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

The Backpropagation Algorithm Implemented on Spiking Neuromorphic Hardware

Alpha Renner,^{1,*} Forrest Sheldon,^{2,†} Anatoly Zlotnik,^{3,‡} Louis Tao,^{4,5,§} and Andrew Sornborger^{6,¶}

¹*Institute of Neuroinformatics, University of Zurich and ETH Zurich, Zurich, Switzerland, 8057*

²*Physics of Condensed Matter & Complex Systems (T-4),*

Los Alamos National Laboratory, Los Alamos, New Mexico, USA, 87545

³*Applied Mathematics & Plasma Physics (T-5), Los Alamos National Laboratory, Los Alamos, New Mexico, USA, 87545*

⁴*Center for Bioinformatics, National Laboratory of Protein Engineering and Plant Genetic Engineering,*

School of Life Sciences, Peking University, Beijing, China, 100871

⁵*Center for Quantitative Biology, Academy for Advanced Interdisciplinary Studies, Peking University, Beijing, China, 100871*

⁶*Information Sciences (CCS-3), Los Alamos National Laboratory, Los Alamos, New Mexico, USA, 87545*

(Dated: July 9, 2021)

The capabilities of natural neural systems have inspired new generations of machine learning algorithms as well as neuromorphic very large-scale integrated (VLSI) circuits capable of fast, low-power information processing. However, it has been argued that most modern machine learning algorithms are not neurophysiologically plausible. In particular, the workhorse of modern deep learning, the backpropagation algorithm, has proven difficult to translate to neuromorphic hardware. In this study, we present a neuromorphic, spiking backpropagation algorithm based on synfire-gated dynamical information coordination and processing, implemented on Intel’s Loihi neuromorphic research processor. We demonstrate a proof-of-principle three-layer circuit that learns to classify digits from the MNIST dataset. To our knowledge, this is the first work to show a Spiking Neural Network (SNN) implementation of the backpropagation algorithm that is fully on-chip, without a computer in the loop. It is competitive in accuracy with off-chip trained SNNs and achieves an energy-delay product suitable for edge computing. This implementation shows a path for using in-memory, massively parallel neuromorphic processors for low-power, low-latency implementation of modern deep learning applications.

I. Introduction

1. Dynamical Coordination in Cognitive Processes.

Spike-based learning in plastic neuronal networks is playing increasingly key roles in both theoretical neuroscience and neuromorphic computing. The brain learns in part by modifying the synaptic strengths between neurons and neuronal populations. While specific synaptic plasticity or neuromodulatory mechanisms may vary in different brain regions, it is becoming clear that a significant level of dynamical coordination between disparate neuronal populations must exist, even within an individual neural circuit [1]. Classically, backpropagation (BP, and other learning algorithms) has been essential for supervised learning in artificial neural networks (ANNs). Although the question of whether or not BP operates in the brain is still an outstanding issue [2], BP does solve the problem of how a global objective function can be related to local synaptic modification in a network. It seems clear, however, that if BP is implemented in the brain, or if one wishes to implement BP in a neuromorphic circuit, some amount of dynamical information coordination is necessary to propagate the correct information to the correct location such that appropriate local synaptic modification may take place to enable learning.

2. The Grand Challenge of Spiking Backpropagation (sBP)

There has been a rapid growth of interest in the reformulation of classical algorithms for learning, optimization, and control using event-based information-processing mechanisms. Such spiking neural networks (SNNs) are inspired by the function of biophysiological neural systems [3], i.e., neuromorphic computing [4]. The trend is driven by the advent of flexible computing architectures such as Intel’s neuromorphic research processor, codenamed Loihi, that enable experimentation with such algorithms in hardware [5]. There is particular interest in deep learning, which is a central tool in modern machine learning. 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. While feedforward networks are readily translated to neuromorphic hardware [6–8], the far more computationally intensive training of these networks ‘on chip’ has proven elusive as the structure of backpropagation makes the algorithm notoriously difficult to implement in a neural circuit [9, 10]. A feasible neural implementation of the backpropagation algorithm has gained renewed scrutiny with the rise of new neuromorphic computational architectures that feature local synaptic plasticity [5, 11–13]. Because of the well-known difficulties, neuromorphic systems have relied to date almost entirely on conventional off-chip learning, and used on-chip computing only for inference [6–8]. It has been a long-standing challenge to develop learning systems whose function is realized exclusively using neuromorphic mechanisms. Backpropagation has been claimed to be biologically

* alpren@ini.uzh.ch

† fsheldon@lanl.gov

‡ azlotnik@lanl.gov

§ taolt@mail.cbi.pku.edu.cn

¶ sornborg@lanl.gov; Corresponding Author

implausible or difficult to implement on spiking chips because of several issues: (a) *Weight transport* – Usually, synapses in biology and on neuromorphic hardware cannot be used bidirectionally, therefore, separate synapses for the forward and backward pass are employed. However, correct credit assignment, i.e. knowing how a weight change affects the error, requires feedback weights to be the same as feedforward weights [14, 15]; (b) *Backwards computation* – Forward and backward passes implement different computations [15]. The forward pass requires only weighted summation of the inputs, while the backward pass operates in the opposite direction and additionally takes into account the derivative of the activation function; (c) *Gradient storage* – Error gradients must be computed and stored separately from activations; (d) *Differentiability* – For spiking networks, the issue of non-differentiability of spikes has been discussed and solutions have been proposed [16–19]; and (e) *Hardware constraints* – For the case of neuromorphic hardware, there are often constraints on plasticity mechanisms, which allow for adaptation of synaptic weights. On some hardware, no plasticity is offered at all, while in some cases only specific spike-timing dependent plasticity (STDP) rules are allowed. Additionally, in almost all available neuromorphic architectures, information must be local, i.e., information is only shared between neurons that are synaptically connected, in particular, to facilitate parallelization.

3. Previous Attempts at sBP

The most commonly used approach to avoiding the above issues is to use neuromorphic hardware only for inference using fixed weights, obtained by training of an identical network offline and off-chip [6, 7, 20–24]. This has recently achieved state-of-the-art performance [8], but it does not make use of neuromorphic hardware’s full potential, because offline training consumes significant power. Moreover, to function in most field applications, an inference algorithm should be able to learn adaptively after deployment, e.g., to adjust to a particular speaker in speech recognition, which would enable autonomy and privacy of edge computing devices. So far, only last layer training, without backpropagation, and using variants of the delta rule, has been achieved on spiking hardware [25–31]. Other on-chip learning approaches use alternatives to backpropagation [32, 33], bio-inspired non-gradient based methods [34], or hybrid systems with a conventional computer in the loop [35, 36]. Several promising alternative approaches for actual on-chip spiking backpropagation have been proposed recently [37–40], but have not yet been implemented in hardware.

To avoid the backwards computation issue (b) and because neuromorphic synapses are not bidirectional, a separate feedback network for the backpropagation of errors has been proposed [41, 42] (see Fig. 1). This leads to the weight transport problem (a) which has been solved by using symmetric learning rules to maintain weight symmetry [42–44] or with the Kolen-Pollack algorithm [44], which leads to symmetric weights automatically. It has

also been found that weights do not have to be perfectly symmetric, because backpropagation can still be effective with random feedback weights (random feedback alignment) [45], although symmetry in the sign between forward and backward weights matters [46].

The backwards computation issue (b) and the gradient storage issue (c) have been addressed by approaches that separate the function of the neuron into different compartments and use structures that resemble neuronal dendrites for the computation of backward propagating errors [2, 38, 40, 47]. The differentiability issue (d) has been circumvented by spiking rate-based approaches [7, 21, 48, 49] that use the ReLU activation function as done in ANNs. The differentiability issue has also been addressed more generally using surrogate gradient methods [6, 16–19, 22, 25, 50] and methods that use biologically-inspired STDP and reward modulated STDP mechanisms [51–54]. For a review of SNN-based deep learning, see [55]. For a review of backpropagation in the brain, see [2].

4. Our Contribution

In this study, we describe a hardware implementation of the backpropagation algorithm that addresses each of the issues (a)–(e) introduced above using a set of mechanisms that have been developed and tested in simulation by the authors during the past decade, synthesized in our recent study [56], and simplified and adapted here to the features and constraints of the Loihi chip. These neuronal and network features include propagation of graded information in a circuit composed of neural populations using synfire-gated synfire chains (SGSCs) [57–60], control flow based on the interaction of synfire-chains [58], and regulation of Hebbian learning using synfire-gating [61, 62]. We simplify our previously proposed network architecture [56], and streamline its function. We demonstrate our approach using a proof-of-principle implementation on Loihi [5], and examine the performance of the algorithm for learning and inference of the MNIST test data set [63]. The sBP implementation is shown to be competitive in clock time, sparsity, and power consumption with state-of-the-art algorithms for the same tasks.

II. Results

1. The Binarized sBP model

We extend our previous architecture [56] using several new algorithmic and hardware-specific mechanisms. Each unit of the neural network is implemented as a single spiking neuron, using the current-based leaky integrate-and-fire (CUBA) model (see Eq. (19) in Section IV) that is built into Loihi. The time constants of the CUBA model are set to 1, so that the neurons are memoryless. Rather than using rate coding, whereby spikes are counted over time, we consider neuron spikes at every algorithmic time step, so we can regard our implementation as a binary neural network. The feedforward

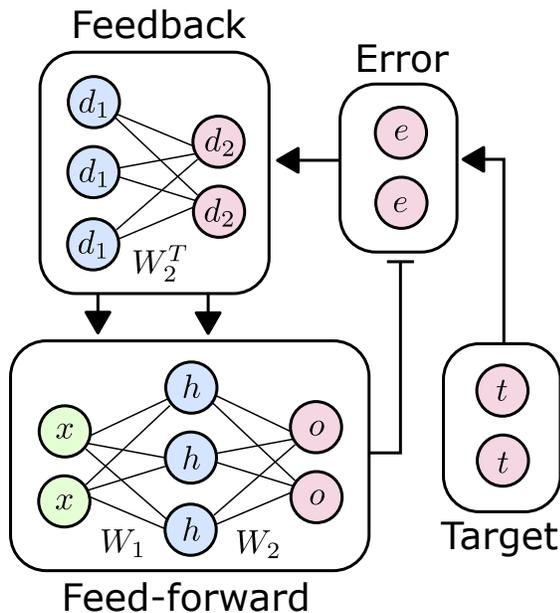


FIG. 1. **Overview of conceptual circuit architecture.** Feedforward activations of input (x), hidden (h) and output (o) layers are calculated by a feedforward module. Errors ($e = t - o$) are calculated from the output and the training signal (t). Errors are backpropagated through a feedback module with the same weights W_2 for synapses between h and o , but in the opposite direction (mathematically expressed as the transpose, W_2^T). Local gradients (d_1 , d_2) are gated back into the feedforward circuit at appropriate times to accomplish potentiation or depression of appropriate weights.

component of our network is a classic multilayer perceptron (MLP) with 3-layers, a binary activation function, and discrete (8 bit) weights. Our approach may, however, be extended to deeper networks and different encodings. In the following equations, each lowercase letter corresponds to a Boolean vector that represents a layer of spiking neurons on the chip (a spike corresponds to a 1). The inference (forward) pass through the network is computed as:

$$o = f(W_2 f(W_1 x)), \quad (1)$$

$$f(x) = H(x - 0.5), \quad (2)$$

$$H(x) = \begin{cases} 0, & x < 0, \\ 1, & x \geq 0, \end{cases} \quad (3)$$

where W_i is the weight matrix of the respective layer, f is a binary activation function with a threshold of 0.5, and H denotes the Heaviside function. The forward pass thus occurs in 3 time steps as spikes are propagated through layers. The degree to which the feedforward network's output (o) deviates from a target value (t) is quantified by the squared error, $E = \frac{1}{2} \|o - t\|^2$, which we would like to minimize. Performing backpropagation to achieve this requires the calculation of weight updates, which depend on the forward activations, and backward propagated lo-

cal gradients d_l , which represent the amount by which the loss changes when the activity of that neuron changes, as:

$$d_2 = (o - t) \circ f'(W_2 h), \quad (4)$$

$$d_1 = \text{sgn}(W_2^T d_2) \circ f'(W_1 x), \quad (5)$$

$$\frac{\partial E}{\partial W_l} = d_l (a_{l-1})^T, \quad (6)$$

$$W_l^{\text{new}} = W_l^{\text{old}} - \eta \frac{\partial E}{\partial W_l}, \quad l = 1, 2. \quad (7)$$

Here, \circ denotes a Hadamard product, i.e. the element-wise product of two vectors, T denotes the matrix transpose, $\text{sgn}(x)$ is the sign function, and a_l denotes the activation of the l th layer, $f(W_l a_{l-1})$, with $a_0 = x$, $a_1 = h$, $a_2 = o$. Here, η denotes the learning rate, and is the only hyperparameter of the model apart from the weight initialization. Here $'$ denotes the derivative, but because f is a binary thresholding function (Heaviside), the derivative would be the Dirac delta function, which is zero everywhere apart from at the threshold. Therefore, we use a common method [64], and represent the thresholding function using a truncated (between 0 and 1) ReLU (Eq. (8)) as a surrogate when back-propagating the error. The derivative of the surrogate is a box function (Eq. (9)):

$$f_{\text{surrogate}}(x) = \min(\max(x, 0), 1), \quad (8)$$

$$f'(x) = H(x) - H(x - 1). \quad (9)$$

The three functions (Equations (2), (8) and (9)) are plotted in the inset in Fig. 2.

When performed for each target (t) in the training set, the model may be thought of as a stochastic gradient descent algorithm with a fixed step size update for each weight in the direction of the sign of the gradient.

2. sBP on Neuromorphic Hardware

On the computational level, Equations (1)-(9) fully describe our model exactly as it is implemented on Loihi, excluding the handling of bit precision constraints that affect integer discreteness and value limits and ranges. In the following, we describe how these equations are translated from the computational to the algorithmic neural circuit level, thereby enabling implementation on neuromorphic hardware. Further details on the implementation can be found in Section IV.

a. Hebbian weight update Equation (7) effectively results in the following weight update per single synapse from presynaptic index i in layer $l - 1$ to postsynaptic index j in layer l :

$$\Delta w_{ij} = -\eta \cdot a_{l-1,i} \cdot d_{l,j}, \quad (10)$$

where η is the constant learning rate. To accomplish this update, we use a Hebbian learning rule [65] implementable on the on-chip microcode learning engine (for the exact implementation on Loihi, see Section IV B). Hebbian learning means that neurons that fire together, wire together, i.e., the weight update Δw is proportional

245 to the product of the simultaneous activity of the presynaptic (source) neuron and the postsynaptic (target) neuron. In our case, this means that the values of the 2 factors of Equation (10) have to be propagated simultaneously, in the same time step, to the pre- ($a_{l-1,i}$) and postsynaptic ($d_{l,j}$) neurons while the pre- and postsynaptic neurons are not allowed to fire simultaneously at any other time. For this purpose, a mechanism to control the information flow through the network is needed.

254 *b. Gating controls the information flow* As our information control mechanism, we use synfire gating [57–59, 66]. A closed chain of 12 neurons containing a single spike perpetually sent around the circle is the backbone of this flow control mechanism, which we call the gating chain. The gating chain controls information flow through the controlled network by selectively boosting layers to bring their neurons closer to the threshold and thereby making them receptive to input. By connecting particular layers to the gating neuron that fires in the respective time steps, we lay out a path that the activity through the network is allowed to take. For example, to create the feedforward pass, the input layer x is connected to the first gating neuron and therefore gated ‘on’ in time step 1, the hidden layer h is connected to the second gating neuron and gated ‘on’ in time step 2, and the output layer o is connected to the third gating neuron and gated ‘on’ in time step 3. A schematic of this path of the activity can be found in Fig. 2. To speak in neuroscience terms, we are using synfire gating to design functional connectivity through the network anatomy shown in Supplementary Fig. 4. Using synfire gating, the local gradient $d_{l,j}$ is brought to the postsynaptic neuron at the same time as the activity $a_{l-1,i}$ is brought back to the presynaptic neuron effecting a weight update. However, in addition to bringing activity at the right time to the right place for Hebbian learning, the gating chain also makes it possible to calculate and back-propagate the local gradient.

283 *c. Local gradient calculation* For the local gradient calculation, according to Equation (5), the error $o - t$ and the box function derivative of the surrogate activation function (Equation (9)) are needed. Because there are no negative (signed) spikes, the local gradient is calculated and propagated back twice for a Hebbian weight update in two phases with different signs. The error $o - t$ is calculated in time step 4 in a layer that receives excitatory (positive) input from the output layer o and inhibitory (negative) input from the target layer t , and vice versa for $t - o$.

294 The box function has the role of initiating learning when the presynaptic neuron receives a non-negative input and of terminating learning when the input exceeds 1, which is why we call the two conditions ‘start’ and ‘stop’ learning (inspired by the nomenclature of [67]). This inherent feature of backpropagation avoids weight updates that have no effect on the current output as the neuron is saturated by the nonlinearity with the current input. This regulates learning by protecting weights that

303 are trained and used for inference when given different inputs.

305 To implement these two terms of the box function (9), we use two copies of the output layer that receive the same input ($W_2 h$) as the output layer. Using the above-described gating mechanism, one of the copies (start learning, $o^<$) is brought exactly to its firing threshold when it receives the input, which means that it fires for any activity greater than 0 and the input is not in the lower saturation region of the ReLU. The other copy (stop learning, $o^>$) is brought further away from the threshold (to 0), which means that if it fires, the upper saturation region of the ReLU has been reached and learning should cease.

317 *d. Error backpropagation* Once the local gradient d_2 is calculated as described in the previous paragraph, it is sent to the output layer as well as to its copies to bring about the weight update of W_2 and its 4 copies in time steps 5 and 9. From there, the local gradient is propagated back through the transposed weight matrices W_2^T and $-W_2^T$, which are copies of W_2 connected in the opposite direction and, in the case of $-W_2^T$, with opposite sign. Once propagated backwards, the back-propagated error is also combined with the ‘start’ and ‘stop’ learning conditions, and then it is sent to the hidden layer and its copies in time steps 7 and 11.

3. Algorithm Performance

330 Our implementation of the sBP algorithm on Loihi achieves an inference accuracy of 95.7% after 60 epochs (best: 96.2%) on the MNIST test data set, which is comparable with other shallow, stochastic gradient descent (SGD) trained MLP models without additional allowances. In these computational experiments, the sBP model is distributed over 81 neuromorphic cores. Processing of a single sample takes 1.48 ms (0.169 ms for inference only) on the neuromorphic cores, including the time required to send the input spikes from the embedded processor, and consumes 0.653 mJ of energy on the neuromorphic cores (0.592 mJ of which is dynamic energy, i.e. energy used by our neural circuit in addition to the fixed background energy), resulting in an energy-delay product of 0.878 μ Js. All measurements were obtained using NxSDK version 0.9.9 on the Nahuku32 board ncl-ghrd-01. Table I shows a comparison of our results with published performance metrics for other neuromorphic learning architectures that were also tested on MNIST. Table III in the Supplementary Material shows a breakdown of energy consumption and a comparison of different conditions and against a GPU. Switching off the learning engine after training reduces the dynamic energy per inference to 0.0204 mJ, which reveals that the on-chip learning engine is responsible for most of the power consumption. Because the learning circuit is not necessary for performing inference, we also tested a reduced architecture that is able to do inference within four time steps using the previously trained weights. This architecture uses 0.00249 mJ of dynamic energy and 0.169 ms per inference.

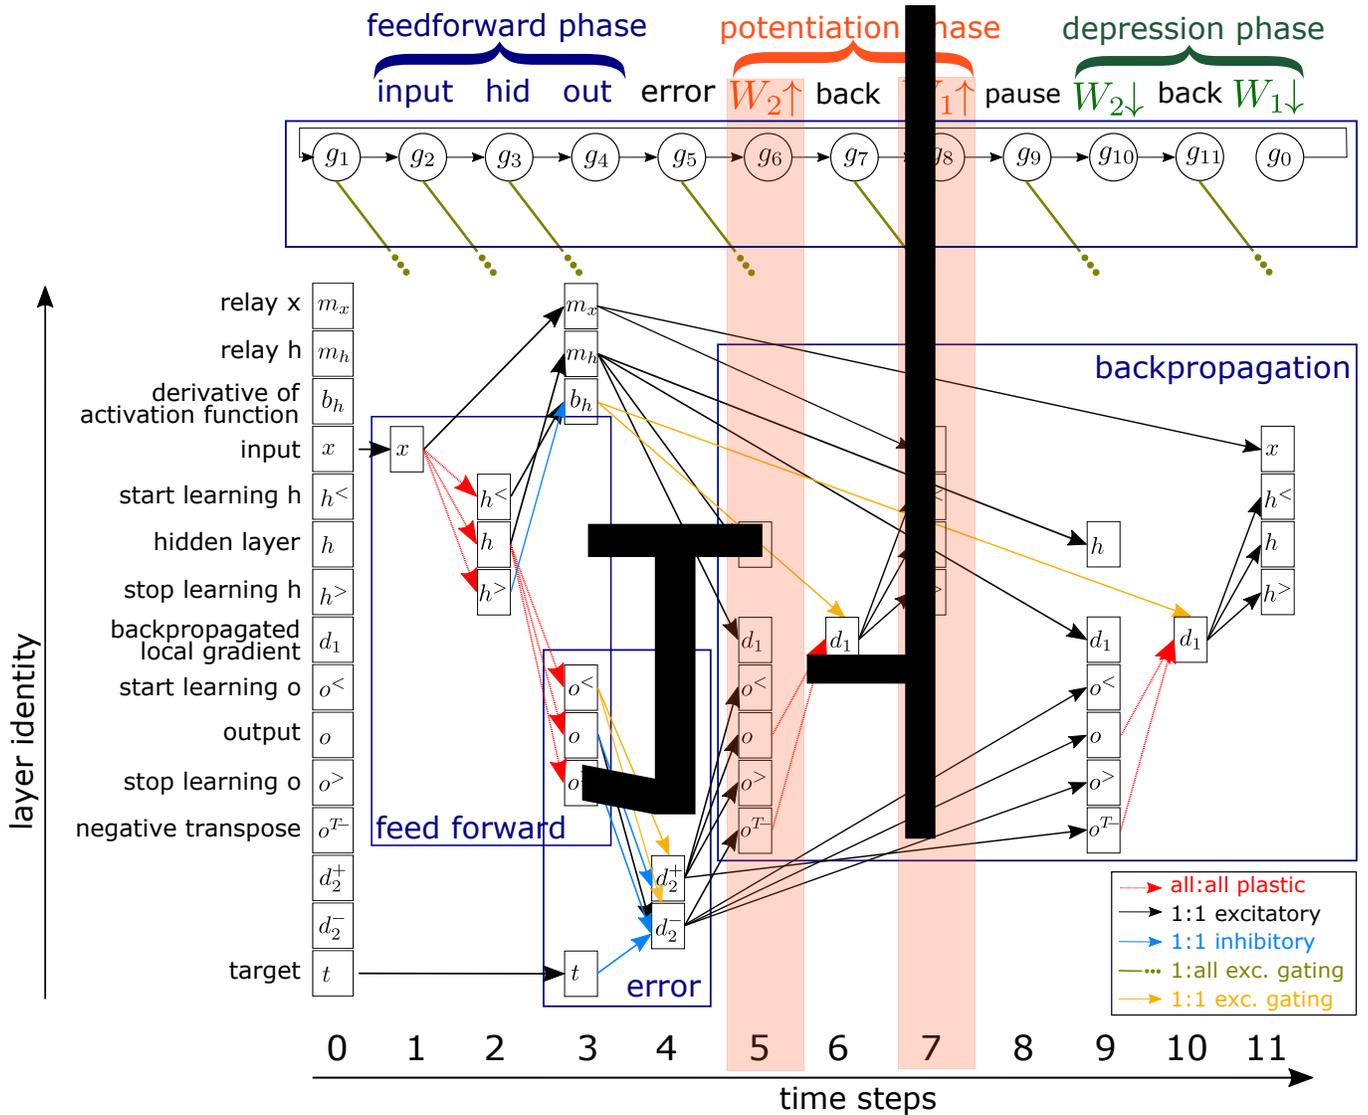


FIG. 2. **Functional connectivity of the 2 layer backpropagation circuit.** Layers are only shown when they are gated ‘on’ and synapses are only shown when their target is gated on. Plastic connections are all-to-all (fully connected), i.e. all neurons are connected to all neurons of the next layer. The gating connections from the gating chain are one-to-all, and all other connections are one-to-one, which means that a firing pattern is copied directly to the following layer. The names of the neuron layers are given on the left margin so that the row corresponds to layer identity. The columns correspond to time steps of the algorithm, which are the same as the time steps on Loihi. Table II shows the information contained in each layer in each respective time step. The red background in time steps 5 and 7 indicates that in these steps, the sign of the weight update is inverted (positive), as $r = 1$ in Eq. (21). A detailed step-by-step explanation of the algorithm is given in Section IV C 2 and in Table II in the supplementary material. The plot in the top left corner illustrates our approach to approximate the activation function f by a surrogate with the box function as derivative, $f_{\text{surr}} = H(x)H(1-x)$, where f' is the rectified linear map (ReLU) (see Equations 2, 8 and 9).

361 The sBP algorithm trains the network without explicit 370 neuron per inference). The spikes of the input and hid-
 362 sparsity constraints, and yet it exhibits sparsity because 371 den layer are repeated in the two learning phases (see
 363 of its binary (spiking) nature. After applying the binary 372 Fig. 2) independent of the training progress. However,
 364 threshold of 0.5 to the MNIST images, one image is 373 the error-induced activity from the local gradient layer
 365 encoded using on average 100 spikes in the input layer, 374 d_1 starts with 0.7 spikes per neuron per sample (during
 366 which corresponds to a sparsity of 0.25 spikes per neu- 375 0 spikes in the trained network as the error goes towards
 367 ron per inference. This leads to a typical number of 110 376
 368 spikes in the hidden layer (0.28 spikes per neuron per in- 377 0.
 369 ference) and 1 spike in the output layer (0.1 spikes per

Publication	Hardware	Learning Mode	Network Structure	Energy per Sample (mJ)	Latency per Sample (ms)	Test Accuracy (%)
On-chip backpropagation						
This study	Loihi	on-chip sBP	400-400-10 ^a	0.592	1.48	96.2
On-chip single layer training or BP alternatives						
[33] Shrestha et al. (2021)	Loihi	EMSTDP FA/DFA	CNN-CNN-100-10	8.4	20	94.7
[32] Frenkel et al. (2020)	SPOON	DRTP	CNN-10	0.000366 ^b	0.12	95.3
[30] Park et al. (2019)	unnamed	mod. SD	784-200-200-10	0.000253 ^b	0.01	98.1
[68] Chen et al. (2018)	unnamed	S-STDP	236-20 ^c	0.017	0.16	89
[27] Frenkel et al. (2018)	ODIN	SDSP	256-10	0.000015	-	84.5
[69] Lin et al. (2018)	Loihi	S-STDP	1920-10 ^c	0.553	-	96.4
[29] Buhler et al. (2017)	unnamed	LCA features	256-10	0.000050	0.001 ^b	88
On-chip inference only						
This study	Loihi	inference	400-400-10 ^a	0.00249	0.169	96.2
[33] Shrestha et al. (2021)	Loihi	inference	CNN-CNN-100-10	2.47	10	94.7
[32] Frenkel et al. (2020)	SPOON	inference	CNN-10	0.000313	0.12	97.5
[70] Göltz et al. (2019)	BrainScaleS-2	inference	256-246-10	0.0084	0.048	96.9
[69] Lin et al. (2018)	Loihi	inference	1920-10 ^c	0.0128 ^d	-	96.4
[68] Chen et al. (2018)	unnamed	inference	784-1024-512-10	0.0017	-	97.9
[72] Esser et al. (2015)	True North	inference	CNN (512 neurons)	0.00027	1	92.7
[72] Esser et al. (2015)	True North	inference	CNN (3840 neurons)	0.108	1	99.4
[73] Stromatias et al. (2015)	SpiNNaker	inference	784-500-500-10	3.3	11	95
Neuromorphic sBP in simulated SNN						
[74] Jin et al. (2018)	Simulation	BP	784-800-10	-	-	98.8
[75] Neftci et al. (2017)	Simulation	BP	784-500-10	-	-	97.7
[76] Shrestha et al. (2019)	Simulation	EM-STDP	784-500-10	-	-	97
[77] Tavanaei and Maida (2019)	Simulation	BP-STDP	784-500-150-10	-	-	97.2
[78] Mostafa (2017)	Simulation	BP	784-800-10	-	-	97.55
[79] Lee et al. (2016)	Simulation	BP	784-800-10	-	-	98.64
[80] O'Connor and Welling (2016)	Simulation	BP	784-300-300-10	-	-	96.4
[81] Diehl and Cook (2015)	Simulation	STDP	784-1600-10	-	-	95

^a 400 (20x20) corresponds to 784 (28x28) after cropping of the empty image margin of 4 pixels

^b Calculated from given values

^c Off-chip preprocessing

^d Dynamic energy reported in the Supplementary Material of [71]

TABLE I. Review of the MNIST Literature in SNN and on neuromorphic hardware. The table includes 4 relevant classes of literature. Studies that used ANN-SNN conversion purely in software are not reviewed here. Note that the energy-delay product may be computed from the Energy per Sample and Latency per Sample columns. Abbreviations: EMSTDP: Error-modulated spike-timing dependent plasticity; DFA: Direct feedback alignment; DRTP: Direct random target projection; SD: Segregated dendrites; SDSP: Spike-driven synaptic plasticity; LCA: Locally competitive algorithm.

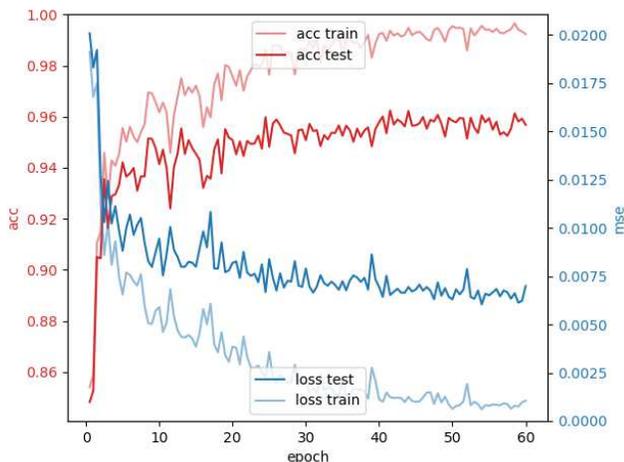


FIG. 3. Accuracy and loss (mean squared error) over epochs. Note separate axis scaling for accuracy (left) and loss (right).

III. Discussion

1. Summary

As we have demonstrated here, by using a well-defined set of neuronal and neural circuit mechanisms, it is possible to implement a spiking backpropagation algorithm on contemporary neuromorphic hardware. Previously proposed methods to address the issues outlined in Section I were not on their own able to offer a straightforward path to implement a variant of the sBP algorithm on current hardware. In this study, we avoided or solved these previously encountered issues with spiking backpropagation by combining known solutions with synfire-gated synfire chains (SGSC) as a dynamical information coordination scheme that was evaluated on the MNIST test data set on the Loihi VLSI hardware.

2. Solutions of Implementation Issues

The five issues (a)-(e) listed in Section I were addressed using the following solutions: (a) The *weight transport* issue was avoided via the use of a deterministic, symmetric learning rule for the parts of the network that im-

plement inference (feed-forward) and error propagation (feedback) as described by [82]. This approach is not biologically plausible because of a lack of developmental mechanisms to assure the equality of corresponding weights [83]. It would, however, without modifications to the architecture be feasible to employ weight decay as described by Kolen and Pollack [83] to achieve self-alignment of the backward weights to the forward weights or to use feedback alignment to approximately align the feedforward weights to random feedback weights [45]; (b) The *backwards computation* issue was solved by using a separate error propagation network through which activation is routed using an SGSC; (c) The *gradient storage* issue was solved by routing activity through the inference and error propagation circuits within the network in separate stages, thereby preventing the mixing of inference and error information. There are alternatives that would not require synfire gated routing, but are more challenging to implement on hardware [38, 40] as also described in a more comprehensive review [2]; (d) The *differentiability* issue was solved by representing the step activation function by a surrogate in the form of a (truncated) ReLU activation function with an easily implementable box function derivative; and (e) The *hardware constraint* issue was solved by the proposed mechanism’s straightforward implementation on Loihi because it only requires basic integrate-and-fire neurons and Hebbian learning that is modulated by a single factor which is the same for all synapses.

3. Encoding and Sparsity

While neural network algorithms on GPUs usually use operations on dense activation vectors and weight matrices, and therefore do not profit from sparsity, spiking neuromorphic hardware only performs an addition operation when a spike event occurs, i.e., adding the weight to the input current as in Equation (19). This means that the power consumption directly depends on the number of spikes. Therefore sparsity, which refers to the property of a vector to have mostly zero elements, is important for neuromorphic algorithms [71, 84], and it is also observed in biology [85]. Consequently, significant effort has been made to make SNNs sparse to overcome the classical rate-based approach based on counting spikes [70, 84, 86, 87]. The binary encoding used here could be seen as a limit case of the rate-based approach allowing only 0 or 1 spike. Even without regularization to promote sparse activity, it yields very sparse activation vectors that are even sparser than most timing codes. The achievable encoded information per spike is unquestionably lower, however. In a sense, we already employ spike timing to route spikes through the network because the location of a spike in time within the 12 time steps determines if and where it is sent and if the weights are potentiated or depressed. However, usage of a timing code for activations could be enabled by having more than one Loihi time step per algorithm time step. Therefore the use of SGSCs is not limited to this particular binary encoding, and in fact, SGSCs were originally designed for

a population rate code.

Similarly, the routing method we use in this work is not limited to backpropagation, but it could serve as a general method to route information in SNNs where autonomous activity (without interference from outside the chip) is needed. That is, our proposed architecture can act in a similar way as or even in combination with neural state machines [88, 89].

4. Algorithmically-Informed Neurophysiology

Although our implementation of sBP here was focused primarily on a particular hardware environment, we point out that the synfire-gated synfire chains and other network and neuronal structures that we employ could all potentially have relevance to the understanding of computation in neurophysiological systems. Many of these concepts that we use, such as spike coincidence, were originally inspired by neurophysiological experiments [57, 58, 90]. Experimental studies have shown recurring sequences of stereotypical neuronal activation in several species and brain regions [91–93] and particularly replay in hippocampus [94]. Recent studies also hypothesize [95, 96] and show [97] that a mechanism like gating by synfire chains may play a role in memory formation. Additional evidence [98] shows that large-scale cortical activity has a stereotypical, packet-like character that can convey information about the nature of a stimulus, or be ongoing or spontaneous. This type of apparently algorithmically-related activity has a very similar form to the SGSC controlled information flow found previously [57–60]. Additionally, this type of sequential activation of populations is evoked by the sBP learning architecture, as seen in the raster plot in Fig. 5 in the Supplementary Material.

Other algorithmic spiking features, such as the back-propagated local gradient layer activity decreasing from 0.7 spikes per neuron to 0 by the end of training, could be identified and used to generate qualitative and quantitative hypotheses concerning network activity in biological neural systems.

5. Future Directions

Although the accuracy we achieve is similar to early implementations of binary neural networks in software [64], subsequent approaches now reach 98.5% [99], and generally include binarized weights. However, networks that achieve such accuracy typically employ either a convolutional structure, or multiple larger hidden layers. Additional features such as dropout, softmax final layers, gain terms, and others could in principle be included in spiking hardware and may also account for this 3% gap. So, while we show that it is possible to efficiently implement backpropagation on neuromorphic hardware, several non-trivial steps are still required to make it usable in practical applications:

- 1) The algorithm needs to be scaled to deeper networks. While the present structure is in principle scalable to more layers without major adjustments, investigation is needed to determine whether gradient informa-

tion remains intact over many layers, and to what extent additional features such as alternatives to batch normalization may need to be developed.

2) Generalization to convolutional networks is compelling, in particular for application to image processing. The Loihi hardware presents an advantage in this setting because of its weight-sharing mechanisms.

3) Although our current implementation demonstrates on-chip learning, we train on the MNIST images in an offline fashion by iterating over the training set in epochs. Further research is required to develop truly continual learning mechanisms such that additional samples and classes can actually be learned without losing previously trained synaptic weights and without retraining on the whole dataset.

Additionally, the proposed algorithmic methodology can be used to inform hardware adjustments to promote efficiency for learning applications. Although our algorithm is highly efficient in terms of power usage, in particular for binary encoding, the Loihi hardware is not specifically designed for implementing standard deep learning models, but rather as general-purpose hardware for exploring different SNN applications [71].

This leads to a significant computational overhead for functions that are not needed in our model (e.g. neuronal dynamics), or that could have been realized more efficiently if integrated directly on the chip instead of using network mechanisms. Our results provide a potential framework to guide future hardware modifications to facilitate more efficient learning algorithm implementations. For example, in an upgraded version of the algorithm, it would be preferable to replace relay neurons with presynaptic (eligibility) traces to keep a memory of the presynaptic activity for the weight update.

6. Significance

To our knowledge, this work is the first to show an SNN implementation of the backpropagation algorithm that is fully on-chip, without a computer in the loop. Other on-chip learning approaches so far either use feed-back alignment [33], forward propagation of errors [32] or single layer training [25, 27, 29, 68, 69]. Compared to an equivalent implementation on a GPU, there is no loss in accuracy, but there are about two orders of magnitude power savings in the case of small batch sizes which are more realistic for edge computing settings. The network model we propose offers opportunities as a building block that can, e.g. be integrated into larger SNN architectures that could profit from a trainable on-chip processing stage.

IV. Methods

In this section, we describe our system on three different levels [100]. First, we describe the computational level by fully stating the equations that result in the intended computation. Second, we describe the spiking neural network (SNN) algorithm. Third, we describe the

details of our hardware implementation that are necessary for exact reproducibility.

A. The Binarized Backpropagation Model

Network Model. Backpropagation is a means of optimizing synaptic weights in a multi-layer neural network. It solves the problem of credit assignment, i.e., attributing changes in the error to changes in upstream weights, by recursively using the chain rule. The inference (forward) pass through the network is computed as

$$o = f(W_N f(W_{N-1}(\dots f(W_1 x)))) , \quad (11)$$

where f is an element-wise nonlinearity and W_i is the weight matrix of the respective layer. The degree to which the network's output (o) deviates from the target values (t) is quantified by the squared error, $E = \frac{1}{2} \|o - t\|^2$, which we aim to minimize. The weight updates for each layer are computed recursively by

$$d_l = \begin{cases} (o - t) \circ f'(n_l), & l = N \\ W_{l+1}^T d_{l+1} \circ f'(n_l), & l < N \end{cases} , \quad (12)$$

$$\frac{\partial E}{\partial W_{l+1}} = d_{l+1} (a_l)^T , \quad (13)$$

$$W_l^{\text{new}} = W_l^{\text{old}} - \eta \frac{\partial E}{\partial W_l} , \quad (14)$$

where n_l is the network activity at layer l (i.e., $n_l = W_l f(W_{l-1} n_{l-1})$). Here, $'$ denotes the derivative, \circ denotes a Hadamard product, i.e. the element-wise product of two vectors, T denotes the matrix transpose, and a_l denotes $f(W_l a_{l-1})$, with $a_0 = x$. The parameter η denotes the learning rate. These general equations (12)-(14) are realized for two layers in our implementation as given by (5)-(7) in Section II. Below, we relate these equations to the neural Hebbian learning mechanism used in the neuromorphic implementation of sBP.

Although in theory the derivative f' of the activation function is applied in (12), in the case that f is a binary thresholding (Heaviside) function, the derivative is the Dirac delta function, which is zero everywhere apart from at the threshold. We use a common approach [64] and represent the activation function by a surrogate (or straight-through estimator [101]), in the form of a rectified linear map (ReLU) truncated between 0 and 1 (Eq. (16)) in the part of the circuit that affects error backpropagation. The derivative of this surrogate (f') is of box function form, i.e.

$$f(x) = H(x - 0.5) , \quad (15)$$

$$f_{\text{surrogate}}(x) = \min(\max(x, 0), 1) , \quad (16)$$

$$f'(x) = H(x) - H(x - 1) , \quad (17)$$

where H denotes the Heaviside function:

$$H(x) = \begin{cases} 0, & x < 0, \\ 1, & x \geq 0. \end{cases} \quad (18)$$

In the following section, we describe how we implement Equations (11)-(16) in a spiking neural network.

Weight Initialization. Plastic weights are initialized by sampling from a Gaussian distribution with mean of 0 and a standard deviation of $1/\sqrt{2/(N_{\text{fanin}} + N_{\text{fanout}})}$ (He initialization [102]). N_{fanin} denotes the number of neurons of the presynaptic layer and N_{fanout} the number of neurons of the postsynaptic layer.

Input Data. The images of the MNIST dataset [63] were cropped by a margin of 4 pixels on each side to remove pixels that are never active and avoid unused neurons and synapses on the chip. The pixel values were thresholded with 0.5 to get a black and white picture for use as input to the network. In the case of the 100 – 300 – 10 architecture, the input images were down-sampled by a factor of 2. The dataset was presented in a different random order in each epoch.

Accuracy Calculation. The reported accuracies are calculated on the full MNIST test data set. A sample was counted as correct when the index of the spiking neuron of the output layer in the output phase (time step 3 in Fig. 2) is equal to the correct label index of the presented image. In fewer than 1% of cases, there was more than one spike in the output layer, and in that case, the lowest spiking neuron index was compared.

B. The Spiking Backpropagation Algorithm

Spiking Neuron Model. For a generic spiking neuron element, we use the current-based linear leaky integrate-and-fire model (CUBA). This model is implemented on Intel’s neuromorphic research processor, codenamed Loihi [5]. Time evolution in the CUBA model as implemented on Loihi is described by the discrete-time dynamics with $t \in \mathbb{N}$, and time increment $\Delta t \equiv 1$:

$$\begin{aligned} V_i(t+1) &= V_i(t) - \frac{1}{\tau_V} V_i(t) + U_i(t) + I_{\text{const}} , \\ U_i(t+1) &= U_i(t) - \frac{1}{\tau_U} U_i(t) + \sum_j w_{ij} \delta_j(t) , \end{aligned} \quad (19)$$

where i identifies the neuron.

The membrane potential $V(t)$ is reset to 0 upon exceeding the threshold V_{thr} and remains at its reset value 0 for a refractory period, T_{ref} . Upon reset, a spike is sent to all connecting synapses. Here, $U(t)$ represents a neuronal current and δ represents time-dependent spiking input. I_{const} is a constant input current.

Parameters and Mapping. In our implementation of the backpropagation algorithm, we take $\tau_V = \tau_U = 1$, $T_{\text{ref}} = 0$, and $I_{\text{const}} = -8192$ (except in gating neurons, where $I_{\text{const}} = 0$). This leads to a memoryless point neuron that spikes whenever its input in the respective time step exceeds $V_{\text{thr}} = 1024$. This happens, when the neuron receives synaptic input larger than $0.5 \cdot V_{\text{thr}}$ and in the same timestep, a gating input overcomes the strong inhibition of the I_{const} , i.e. it is gated ‘on’. This is how the Heaviside function in Equation (15) is implemented. For

the other activation functions, a different gating input is applied.

There is a straightforward mapping between the weights and activations in the spiking neural network (SNN) described in this section and the corresponding artificial neural network (ANN) described in Section IV A:

$$\begin{aligned} w_{\text{SNN}} &= w_{\text{ANN}} \cdot V_{\text{thr}} , \\ a_{\text{SNN}} &= a_{\text{ANN}} \cdot V_{\text{thr}} . \end{aligned} \quad (20)$$

So, a value of $V_{\text{thr}} = 1024$ allows for a maximal ANN weight of 0.25, because the allowed weights on Loihi are the even integers from -256 to 254.

Feed-forward Pass. The feedforward pass can be seen as an independent circuit module that consists of 3 layers. An input layer x with 400 (20x20) neurons that spikes according to the binarized MNIST dataset, a hidden layer h of 400 neurons, and an output layer o of 10 neurons. The 3 layers are sequentially gated ‘on’ by the gating chain so that activity travels from the input layer to the hidden layer through the plastic weight matrix W_1 and then from the hidden to the output layer through the plastic weight matrix W_2 .

Learning Rule. The plastic weights follow the standard Hebbian learning rule with a global third factor to control the sign. Note, however, that here, unlike other work with Hebbian learning rules, due to the particular activity routed to the layers, the learning rule implements a supervised mechanism (backpropagation). Here we give the discrete update equation as implemented on Loihi:

$$\begin{aligned} \Delta w &= 4r(t) \circ x(t) \circ y(t) - 2x(t) \circ y(t) \\ &= (2r(t) - 1) \cdot 2x(t)y(t) , \\ r(t) &= \begin{cases} 1, & \text{if } (t \bmod T) = 5, 7 , \\ 0, & \text{otherwise} . \end{cases} \end{aligned} \quad (22)$$

Above, x , y , and r represent time series that are available at the synapse on the chip. The signals x and y are the pre- and postsynaptic neuron’s spike trains, i.e., they are equal to 1 in time steps when the respective neuron fires, and 0 otherwise. The signal r is a third factor that is provided to all synapses globally and determines in which phase (potentiation or depression) the algorithm is in. T denotes the number of phases per trial, which is 12 in this case. So, r is 0 in all time steps apart from the 5th and 7th of each iteration, where the potentiation of the weights happens. This regularly occurring r signal could thus be generated autonomously using the gating chain. On Loihi, r is provided using a so-called ‘reinforcement channel’. Note that the reinforcement channel can only provide a global modulatory signal that is the same for all synapses.

The above learning rule produces a positive weight update in time steps in which all three factors are 1, i.e., when both pre- and post-synaptic neurons fire and the reinforcement channel is active. It produces a negative update when only the pre- and post-synaptic neurons fire, and the weight stays the same in all other cases.

To achieve the correct weight update according to the backpropagation algorithm (see (10)), the spiking network has to be designed in a way that the presynaptic activity $a_{l-1,i}$ and the local gradient $d_{i,j}$ are present in neighboring neurons at the same time step. Furthermore, the sign of the local gradient has to determine if the simultaneous activity happens in a time step with active third factor r or not.

This requires control over the flow of spikes in the network, which is achieved by a mechanism called synfire gating [61, 62], which we adapt and simplify here.

Gating Chain. Gating neurons are a separate structure within the backpropagation circuit and are connected in a ring. That is, the gating chain is a circular chain of neurons that, once activated, controls the timing of information processing in the rest of the network. This allows information routing throughout the network to be autonomous to realize the benefits of neuromorphic hardware without the need for control by a classical sequential processor. Specifically, the neurons in the gating chain are connected to the relevant layers of the network, which allows them to control when and where information is propagated. All layers are inhibited far away from their firing threshold by default, as described in Section IV B, and can only transmit information, i.e., generate spikes, if their inhibition is neutralized via activation by the gating chain. Because a neuron only fires if it is gated ‘on’ AND gets sufficient input, such gating corresponds to a logical AND or coincidence detection with the input.

In our implementation, the gating chain consists of 12 neurons, which induce 12 algorithmic (Loihi) time steps that are needed to process one sample. Each neuron is connected to all layers that must be active in each respective time step. The network layers are connected in the manner shown in Supplementary Fig. 4, but the circuit structure, which is controlled by the gating chain, results in a functionally connected network as presented in Fig. 2, where the layers are shown according to the timing of when they become active during one iteration.

The weight of the standard gating synapses (from one of the gating neurons to each neuron in a target layer) is $w_{\text{gate}} = -I_{\text{const}} + 0.5V_{\text{thr}}$, i.e. each neuron that is gated ‘on’ is brought to half of its firing threshold, which effectively implements Eq. (15). In four cases, i.e., for the synapses to the start learning layers in time step 2 ($h^<$) and 3 ($o^<$) and to the backpropagated local gradient layer d_1 in time steps 6 and 10, the gating weight is $w_{\text{gate}} = -I_{\text{const}} + V_{\text{thr}}$. In two cases, i.e., for the synapses to the stop learning layer in time step 2 ($h^>$) and 3 ($o^>$), the gating weight is $w_{\text{gate}} = -I_{\text{const}}$. These different gating inputs lead to step activation functions with different thresholds, as required for the computations explained below, in Section IV B 3.

Backpropagation Network Modules. In the previous sections, we have explained how the weight update happens and how to bring the relevant values (a_{l-1} and d_l according to (10)) to the correct place at the correct time. In this section, we discuss how these values are actually

calculated. The signal a_{l-1} , which is the layer activity from the feedforward pass, does not need to be calculated but only remembered. This is done using a relay layer with synaptic delays, as explained in Section IV B 1. The signal d_2 , the last layer local gradient, consists of 2 parts according to (5). The difference between the output and the target $o - t$ (see Section IV B 2) and the box function f' must be calculated. We factorize the latter into two terms, a start and stop learning signal (see Section IV B 3). The signal d_1 , the backpropagated local gradient, also consists of 2 parts, according to Eq. (6). In addition to another ‘start’ and ‘stop’ learning signal, we need $\text{sgn}(W_2^T d_2)$, whose computation is explained in Section IV B 4.

In the following equation, the weight update is annotated with the number of the paragraph in which the calculating module is described:

$$\Delta W_2 = \eta \underbrace{(o - t)}_2 \circ \underbrace{f'(W_2 h)}_3 \underbrace{h^T}_1, \quad (23)$$

$$\Delta W_1 = \eta \underbrace{\text{sgn}(W_2^T d_2)}_4 \circ \underbrace{f'(W_1 x)}_3 \underbrace{x^T}_1. \quad (24)$$

1. Relay Neurons

The memory used to store the activity of the input and the hidden layer is a single relay layer that is connected both from and to the respective layer in a one-to-one manner with the proper delays. The input layer x sends its activity to the relay layer m_x so that the activity can be restored in the W_1 update phases in time steps 7 and 11. The hidden layer h sends its activity to the relay layer m_h so that the activity can be restored in the W_2 update phases in time steps 5 and 9.

2. Error Calculation

The error calculation requires a representation of signed quantities, which is not directly possible in a spiking network because there are no negative spikes. This is achieved here by splitting the error evaluation into two parts, $t - o$ and $o - t$, to yield the positive and negative components separately. Similarly, the calculation of back-propagated local gradients, d_1 , is performed using a negative copy of the transpose weight matrix, and it is done in 2 phases for the positive and negative local gradient, respectively. In the spiking neural network, $t - o$ is implemented by an excitatory synapse from t and an inhibitory synapse of the same strength from o , and vice versa for $o - t$. Like almost all other nonplastic synapses in the network, the absolute weight of the synapses is just above the value that makes the target neuron fire, when gated on. The difference between the output and the target is, however, just one part of the local gradient d_2 . The other part is the derivative of the activation function (box function).

3. Start and Stop Learning Conditions

The box function (17) can be split in two conditions: a ‘start’ learning and a ‘stop’ learning condition. These two

805 conditions are calculated in parallel with the feedforward
 806 pass. The feedforward activation $f(x) = H(x - 0.5V_{\text{thr}})$
 807 corresponding to Eq. (1) is an application of the spiking
 808 threshold to the layer’s input with an offset of $0.5V_{\text{thr}}$,
 809 which is given by the additional input from the gating
 810 neurons. The first term of the box function (9), $H(x)$, is
 811 also an application of the spiking threshold, but this time
 812 with an offset equal to the firing threshold so that any
 813 input larger than 0 elicits a spike. We call this first term
 814 the ‘start’ learning condition, and it is represented in $h^<$
 815 for the hidden and in $o^<$ for the output layer. The second
 816 term of the box function in Eq. (9), $-H(x - 1V_{\text{thr}})$, is
 817 also an application of the spiking threshold, but this time
 818 without an offset so that only an input larger than the
 819 firing threshold elicits a spike. We call this second term
 820 the ‘stop’ learning condition, and it is represented in $h^>$
 821 and $o^>$ for the hidden and output layers, respectively.
 822 For the W_1 update, the two conditions are combined in a
 823 box function layer $b_h = h^< - h^>$ that then gates the d_1 lo-
 824 cal gradient layer. For the W_2 update, the two conditions
 825 are directly applied to the d_2 layers because an interme-
 826 diate b_o layer would waste one time step. The function
 827 is however the same: the stop learning $o^>$ inhibits the
 828 d_2 layers and the ‘start’ learning signal $o^<$ gates them.
 829 In our implementation, the two conditions are obtained
 830 in parallel with the feedforward pass, which requires two
 831 additional copies of each of the two weight matrices. An
 832 alternative method to avoid these copies but takes more
 833 time steps, would do this computation sequentially and
 834 use the synapses and layers of the feedforward pass three
 835 times per layer with different offsets, and then route the
 836 outputs to their respective destinations.

837 4. Error Backpropagation

838 Error calculation and gating by the start learning sig-
 839 nal and inhibition by the stop learning signal are com-
 840 bined in time step 4 to calculate the last layer local gra-
 841 dients d_2^+ and d_2^- . From there, the local gradients are
 842 routed into the output layer and its copies for the last
 843 layer weight update. This happens in 2 phases: The posi-
 844 tive local gradient d_2^+ is connected without delay so that
 845 it leads to potentiation of the forward and backward last
 846 layer weight matrices in time step 5. The negative local
 847 gradient is connected with a delay of 4 time steps so that
 848 it arrives in the depression phase in time step 9. For the
 849 connections to the o^{T-} layer which is connected to the
 850 negative copy $-W_2^T$, the opposite delays are used to get
 851 a weight update with the opposite sign. See Fig. 2 for a
 852 visualization of this mechanism. Effectively, this leads to
 853 the last layer weight update

$$\begin{aligned} \Delta W_2 = & \eta(H(t - o) \circ f'(W_2 h)) h^T \\ & - \eta(H(o - t) \circ f'(W_2 h)) h^T, \end{aligned} \quad (25)$$

854 where the first term on the right hand side is non-zero
 855 when the local gradient is positive, corresponding to an
 856 update happening in the potentiation phase, and the sec-
 857 ond term is nonzero when the local gradient is negative,

858 corresponding to an update happening in the depression
 859 phase. The functions f and f' are as described in Equa-
 860 tions (15) and (9).

861 The local gradient activation in the output layer does
 862 not only serve the purpose of updating the last layer
 863 weights, but it is also directly used to backpropagate the
 864 local gradients. Propagating the signed local gradients
 865 backwards through the network layers requires a positive
 866 and negative copy of the transposed weights, W_2^T and
 867 $-W_2^T$, which are the weight matrices of the synapses be-
 868 tween the output layer o and the back-propagated local
 869 gradient layer d_1 , and between o^{T-} and d_1 , respectively.
 870 Here o^{T-} is an intermediate layer that is created exclu-
 871 sively for this purpose. The local gradient is propagated
 872 backwards in two phases. The part of the local gradi-
 873 ent that leads to potentiation is propagated back in time
 874 steps 5 to 7, and the part of the local gradient that leads
 875 to depression of the W_1 weights is propagated back in
 876 time steps 9 to 11. In time step 6, the potentiating local
 877 gradient is calculated in layer d_1 as

$$d_1^+ = H(W_2^T d_2^+ + (-W_2^T) d_2^-) \circ b_h, \quad (26)$$

878 and in timestep 10 the depressing local gradient is calcu-
 879 lated in layer d_1 as

$$d_1^- = H((-W_2^T) d_2^+ + W_2^T d_2^-) \circ b_h. \quad (27)$$

880 Critically, this procedure does *not* simply separate the
 881 weights by sign, but rather maintains a complete copy
 882 of the weights that is used to associate appropriate sign
 883 information to the back-propagated local gradient val-
 884 ues. Note that here the Heaviside activation function
 885 $H(x)$ is used rather than the binary activation function
 886 $f = H(x - 0.5V_{\text{thr}})$, so that *any* positive gradient will
 887 induce an update of the weights. Any positive threshold
 888 in this activation will lead to poor performance by mak-
 889 ing the learning insensitive to small gradient values. The
 890 transposed weight copies must be kept in sync with their
 891 forward counterparts, so the updates in the potentiation
 892 and depression phases are also applied to the forward and
 893 backward weights concurrently.

So in total, after each trial, the actual first layer weight
 update is the sum of four different parts:

$$\begin{aligned} \Delta W_1 = & \eta((H(W_2^T d_2^+ + (-W_2^T) d_2^-) \circ b_h) x^T \\ & - \eta((H((-W_2^T) d_2^+ + W_2^T d_2^-) \circ b_h) x^T). \end{aligned} \quad (28)$$

894 These four terms are necessary because, e.g., a posi-
 895 tive error can also lead to depression if backpropagated
 896 through a negative weight matrix and the other way
 897 round.

898 C. The sBP Implementation on Loihi

899 **Partitioning on the Chip.** To distribute the spike
 900 load over cores, neurons of each layer are approximately
 901 equally distributed over 96 cores of a single chip. This

distribution is advantageous because only a few layers are active at each time step, and Loihi processes spikes within a core in a sequential manner. In total, the network as presented here needs $2N_{\text{in}}+6N_{\text{hid}}+7N_{\text{out}}+12N_{\text{gat}}$ neurons. With $N_{\text{in}} = 400$, $N_{\text{hid}} = 400$, $N_{\text{out}} = 10$, and $N_{\text{gat}} = 1$, these are 3282 neurons and about 200k synapses, most of which are synapses of the 3 plastic all-to-all connections between the input and the hidden layer.

Learning Implementation. The learning rule on Loihi is implemented as given in Eq. (21). Because the precision of the plastic weights on Loihi is maximally 8 bits with a weight range from -256 to 254 , we can only change the weight in steps of 2 without making the update non-deterministic. This is necessary for keeping the various copies of the weights in sync (hence the factor of 2 in Eq. (21)). With $V_{\text{thr}} = 1024$, this corresponds to a learning rate of $\eta = \frac{2}{1024} \approx 0.002$. The learning rate can be changed by increasing the factor in the learning rule, which leads to a reduction in usable weight precision, or by changing V_{thr} , which changes the range of possible weights according to Eq. (20). Several learning rates (settings of V_{thr}) were tested with the result that the final accuracy is not very sensitive to small changes. The learning rate that yielded the best accuracy is reported here. In the NxSDK API, the neuron traces that are used for the learning rule are $x0$, $y0$, and $r1$ for $x(t)$, $y(t)$ and $r(t)$ in 21 respectively. $r1$ was used with a decay time constant of 1, so that it is only active in the respective time step, effectively corresponding to $r0$. To provide the r signal, a single reinforcement channel was used and was activated by a regularly firing spike generator in time steps 5 and 7.

Weight Initialization. After the He initialization as described in Section IV A, the weights are mapped to Loihi weights according to (20). Then, the weights are truncated to $[-240, 240]$ and discretized to 8 bits resolution, i.e, steps of 2, by rounding them to the next valid number towards 0.

Power Measurements. All Loihi power measurements are obtained using NxSDK version 0.9.9 on the Nahuku32 board ncl-ghrd-01. Both software API and hardware were provided by Intel Labs. All other probes, including the output probes, are deactivated. For the inference measurements, we use a network that only consists of the three feedforward layers with non-plastic weights and the gating chain of four neurons. The power was measured for the first 10000 samples of the training set for the training measurements and all 10000 samples of the test set for the inference measurements.

References

- [1] Roelfsema, P. R. & Holtmaat, A. Control of synaptic plasticity in deep cortical networks. *Nat Rev Neurosci* **19**, 166–180 (2018).
- [2] Lillicrap, T. P., Santoro, A., Marris, L., Akerman, C. J. & Hinton, G. Backpropagation and the brain. *Nature Reviews Neuroscience* 1–12 (2020).
- [3] Yamins, D. L. & DiCarlo, J. J. Using goal-driven deep learning models to understand sensory cortex. *Nat. Neurosci.* **19**, 356–365 (2016).
- [4] Mead, C. Neuromorphic electronic systems. *Proceedings of the IEEE* **78**, 1629–1636 (1990).
- [5] Davies, M. *et al.* Loihi: A neuromorphic manycore processor with on-chip learning. *IEEE Micro* **38**, 82–99 (2018). URL <https://doi.org/10.1109/MM.2018.112130359>.
- [6] Esser, S. *et al.* Convolutional networks for fast, energy-efficient neuromorphic computing. *PNAS* **113**, 11441–11446 (2016).
- [7] Rueckauer, B., Lungu, I.-A., Hu, Y., Pfeiffer, M. & Liu, S.-C. Conversion of continuous-valued deep networks to efficient event-driven networks for image classification. *Frontiers in neuroscience* **11**, 682 (2017).
- [8] Severa, W., Vineyard, C. M., Dellana, R., Verzi, S. J. & Aimone, J. B. Training deep neural networks for binary communication with the whetstone method. *Nature Machine Intelligence* **1**, 86–94 (2019).
- [9] Grossberg, S. Competitive learning: From interactive activation to adaptive resonance. In Waltz, D. & Feldman, J. A. (eds.) *Connectionist Models and Their Implications: Readings from Cognitive Science*, 243–283 (Ablex Publishing Corp., Norwood, NJ, USA, 1988). URL <http://dl.acm.org/citation.cfm?id=54455.54464>.
- [10] Crick, F. The recent excitement about neural networks. *Nature* **337**, 129–132 (1989).
- [11] Painkras, E. *et al.* SpiNNaker: A multi-core system-on-chip for massively-parallel neural net simulation. In *Proceedings of the IEEE 2012 Custom Integrated Circuits Conference* (2012).
- [12] Schemmel, J. *et al.* 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).
- [13] Qiao, N. *et al.* A reconfigurable on-line learning spiking neuromorphic processor comprising 256 neurons and 128k synapses. *Frontiers in neuroscience* **9**, 141 (2015).
- [14] Grossberg, S. Competitive learning: From interactive activation to adaptive resonance. *Cognitive science* **11**, 23–63 (1987).
- [15] Liao, Q., Leibo, J. & Poggio, T. How important is weight symmetry in backpropagation? In *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 30 (2016).
- [16] Bohte, S. M., Kok, J. N. & La Poutré, H. Error-backpropagation in temporally encoded networks of spiking neurons. *Neurocomputing* **48**, 17–37 (2002). URL <http://linkinghub.elsevier.com/retrieve/pii/S0925231201006580>.
- [17] Pfister, J.-P., Toyozumi, T., Barber, D. & Gerstner, W. Optimal spike-timing-dependent plasticity for precise action potential firing in supervised learning. *Neural Computation* **18**, 1318–1348 (2006). URL <http://dx.doi.org/10.1162/neco.2006.18.6.1318>.
- [18] Zenke, F. & Ganguli, S. Superspike: Supervised learning in multilayer spiking neural networks. *Neural computation* **30**, 1514–1541 (2018).
- [19] Huh, D. & Sejnowski, T. J. Gradient descent for spiking neural networks. In *Advances in Neural Information*

- Processing Systems, 1433–1443 (2018). 1085
- [20] Rasmussen, D. Nengodl: Combining deep learning and neuromorphic modelling methods. *Neuroinformatics* **17**, 611–628 (2019). 1086
- [21] Sengupta, A., Ye, Y., Wang, R., Liu, C. & Roy, K. Going deeper in spiking neural networks: VGG and residual architectures. *Frontiers in Neuroscience* **13**, 95 (2019). URL <https://www.frontiersin.org/article/10.3389/fnins.2019.00095/full>. 1088
- [22] Shrestha, S. B. & Orchard, G. Slayer: Spike layer error reassignment in time. In *Advances in Neural Information Processing Systems*, 1412–1421 (2018). 1089
- [23] Nair, M. V. & Indiveri, G. Mapping high-performance rns to in-memory neuromorphic chips. *arXiv preprint arXiv:1905.10692* (2019). 1090
- [24] Rueckauer, B. *et al.* Nxtf: An api and compiler for deep spiking neural networks on intel loihi. *arXiv preprint arXiv:2101.04261* (2021). 1091
- [25] Stewart, K., Orchard, G., Shrestha, S. B. & Neftci, E. On-chip few-shot learning with surrogate gradient descent on a neuromorphic processor. In *2020 2nd IEEE International Conference on Artificial Intelligence Circuits and Systems (AICAS)*, 223–227 (IEEE, 2020). 1092
- [26] DeWolf, T., Jaworski, P. & Eliasmith, C. Nengo and low-power ai hardware for robust, embedded neuro-robotics. *Frontiers in Neurobotics* **14** (2020). 1093
- [27] Frenkel, C., Lefebvre, M., Legat, J.-D. & Bol, D. A 0.086-mm² 12.7-pj/sop 64k-synapse 256-neuron online-learning digital spiking neuromorphic processor in 28-nm cmos. *IEEE transactions on biomedical circuits and systems* **13**, 145–158 (2018). 1094
- [28] Kim, J. K., Knag, P., Chen, T. & Zhang, Z. A 640m pixel/s 3.65 mw sparse event-driven neuromorphic object recognition processor with on-chip learning. In *2015 Symposium on VLSI Circuits (VLSI Circuits)*, C50–C51 (IEEE, 2015). 1095
- [29] Buhler, F. N. *et al.* A 3.43 tops/w 48.9 pj/pixel 50.1 nj/classification 512 analog neuron sparse coding neural network with on-chip learning and classification in 40nm cmos. In *2017 Symposium on VLSI Circuits*, C30–C31 (IEEE, 2017). 1096
- [30] Park, J., Lee, J. & Jeon, D. 7.6 a 65nm 236.5 nj/classification neuromorphic processor with 7.5% energy overhead on-chip learning using direct spike-only feedback. In *2019 IEEE International Solid-State Circuits Conference-(ISSCC)*, 140–142 (IEEE, 2019). 1097
- [31] Nandakumar, S. *et al.* experimental demonstration of supervised learning in spiking neural networks with phase-change memory synapses. *Scientific reports* **10**, 1–11 (2020). 1098
- [32] Frenkel, C., Legat, J.-D. & Bol, D. A 28-nm convolutional neuromorphic processor enabling online learning with spike-based retinas. In *2020 IEEE International Symposium on Circuits and Systems (ISCAS)*, 1–5 (IEEE, 2020). 1099
- [33] Shrestha, A., Fang, H., Rider, D., Mei, Z. & Qui, Q. In-hardware learning of multilayer spiking neural networks on a neuromorphic processor. In *to appear in 2021 58th ACM/ESDA/IEEE Design Automation Conference (DAC)* (IEEE, 2021). 1100
- [34] Imam, N. & Cleland, T. A. Rapid online learning and robust recall in a neuromorphic olfactory circuit. *Nature Machine Intelligence* **2**, 181–191 (2020). 1101
- [35] Friedmann, S. *et al.* Demonstrating hybrid learning in a flexible neuromorphic hardware system. *IEEE transactions on biomedical circuits and systems* **11**, 128–142 (2016). 1102
- [36] Nandakumar, S. *et al.* Mixed-precision deep learning based on computational memory. *Frontiers in Neuroscience* **14**, 406 (2020). 1103
- [37] Payvand, M., Fouda, M. E., Kurdahi, F., Eltawil, A. M. & Neftci, E. O. On-chip error-triggered learning of multi-layer memristive spiking neural networks. *IEEE Journal on Emerging and Selected Topics in Circuits and Systems* **10**, 522–535 (2020). 1104
- [38] Payeur, A., Guerguiev, J., Zenke, F., Richards, B. & Naud, R. Burst-dependent synaptic plasticity can coordinate learning in hierarchical circuits. *bioRxiv* (2020). URL <https://doi.org/10.1101/2020.03.30.015511>. 1105
- [39] Bellec, G. *et al.* A solution to the learning dilemma for recurrent networks of spiking neurons. *bioRxiv* 738385 (2020). 1106
- [40] Sacramento, J., Costa, R. P., Bengio, Y. & Senn, W. Dendritic cortical microcircuits approximate the back-propagation algorithm. In *Advances in neural information processing systems*, 8721–8732 (2018). 1107
- [41] Stork, D. G. Is backpropagation biologically plausible. In *International Joint Conference on Neural Networks*, vol. 2, 241–246 (IEEE Washington, DC, 1989). 1108
- [42] Zipser, D. & Rumelhart, D. The neurobiological significance of the new learning models. In Schwartz, E. L. (ed.) *Computational Neuroscience*, 192–200 (The MIT Press, 1990). 1109
- [43] O’Reilly, R. C. Biologically plausible error-driven learning using local activation differences: The generalized recirculation algorithm. *Neural Computation* **8**, 895–938 (1996). URL <http://www.mitpressjournals.org/doi/10.1162/neco.1996.8.5.895>. 1110
- [44] Kolen, J. & Pollack, J. Backpropagation without weight transport. In *Proceedings of 1994 IEEE International Conference on Neural Networks (ICNN’94)*, 1375–1380 (IEEE, 1994). URL <http://ieeexplore.ieee.org/document/374486/>. 1111
- [45] Lillicrap, T. P., Cownden, D., Tweed, D. B. & Akerman, C. J. Random synaptic feedback weights support error backpropagation for deep learning. *Nature communications* **7**, 1–10 (2016). 1112
- [46] Liao, Q., Leibo, J. Z. & Poggio, T. How important is weight symmetry in backpropagation? In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, AAAI’16, 1837–1844 (AAAI Press, 2016). URL <http://dl.acm.org/citation.cfm?id=3016100.3016156>. 1113
- [47] Richards, B. A. & Lillicrap, T. P. Dendritic solutions to the credit assignment problem. *Current Opinion in Neurobiology* **54**, 28–36 (2019). URL <http://dx.doi.org/10.1016/j.conb.2018.08.003>. 1114
- [48] O’Connor, P., Neil, D., Liu, S.-C., Delbruck, T. & Pfeiffer, M. Real-time classification and sensor fusion with a spiking deep belief network. *Frontiers in Neuroscience* **7**, 178 (2013). URL <http://dx.doi.org/10.3389/fnins.2013.00178>. 1115
- [49] Kim, R., Li, Y. & Sejnowski, T. J. Simple framework for constructing functional spiking recurrent neural networks. *Proceedings of the national academy of sciences* **116**, 22811–22820 (2019). 1116
- [50] Neftci, E. O., Mostafa, H. & Zenke, F. Surrogate gradient learning in spiking neural networks. *IEEE Signal* 1117

- 1149 *Processing Magazine* **36**, 61–63 (2019). 1213
- 1150 [51] Izhikevich, E. M. Solving the distal reward prob- 1214
1151 lem through linkage of STDP and dopamine signal- 1215
1152 ing. *Cerebral Cortex* **17**, 2443–2452 (2007). URL 1216
1153 <http://dx.doi.org/10.1093/cercor/bhl1152>. 1217
- 1154 [52] Sporea, I. & Grünig, A. Supervised learning in multi- 1218
1155 layer spiking neural networks. *Neural Computation* **25**, 1219
1156 473–509 (2013). URL [http://dx.doi.org/10.1162/](http://dx.doi.org/10.1162/NECO_a_00396)
1157 [NECO_a_00396](http://dx.doi.org/10.1162/NECO_a_00396). 1220
- 1158 [53] Legenstein, R., Pecevski, D. & Maass, W. A learning 1221
1159 theory for reward-modulated spike-timing-dependent 1222
1160 plasticity with application to biofeedback. *PLoS Com- 1223*
1161 *putational Biology* **4**, e1000180 (2008). URL <http://dx.doi.org/10.1371/journal.pcbi.1000180>. 1224
1162 1225
- 1163 [54] Frémaux, N. & Gerstner, W. Neuromodulated spike- 1227
1164 timing-dependent plasticity, and theory of three-factor 1228
1165 learning rules. *Frontiers in Neural Circuits* **9**, 85 (2015). 1229
1166 URL <http://dx.doi.org/10.3389/fncir.2015.00085>. 1230
- 1167 [55] Tavanaei, A., Ghodrati, M., Kheradpisheh, S. R., 1231
1168 Masquelier, T. & Maida, A. Deep learning in spiking 1232
1169 neural networks. *Neural Networks* **111**, 47–63 (2019). 1233
- 1170 [56] Sornborger, A., Tao, L., Snyder, J. & Zlotnik, A. A 1234
1171 pulse-gated, neural implementation of the backpropaga- 1235
1172 tion algorithm. In *Proceedings of the 7th Annual Neuro- 1236*
1173 *inspired Computational Elements Workshop*, 10 (ACM, 1237
1174 2019). 1238
- 1175 [57] Sornborger, A., Wang, Z. & Tao, L. A mechanism 1239
1176 for graded, dynamically routable current propagation 1240
1177 in pulse-gated synfire chains and implications for infor- 1241
1178 mation coding. *J. Comput. Neurosci.* (2015). 1242
- 1179 [58] Wang, Z., Sornborger, A. & Tao, L. Graded, dynam- 1243
1180 ically routable information processing with synfire-gated 1244
1181 synfire chains. *PLoS Comp Biol* **12**, 6 (2016). 1245
- 1182 [59] Wang, C., Xiao, Z., Wang, Z., Sornborger, A. T. & 1246
1183 Tao, L. A fokker-planck approach to graded informa- 1247
1184 tion propagation in pulse-gated feedforward neuronal 1248
1185 networks. *arXiv preprint arXiv:1512.00520* (2015). 1249
- 1186 [60] Xiao, Z., Wang, B., Sornborger, A. & Tao, L. Mutual 1250
1187 information and information gating in synfire chains. 1251
1188 *Entropy* **20**, 102 (2018). 1252
- 1189 [61] Shao, Y., Sornborger, A. & Tao, L. A pulse-gated, pre- 1253
1190 dictive neural circuit. *Proceedings of the 50th Asilomar 1254*
1191 *Conference on Signals, Systems and Computers* (2016). 1255
- 1192 [62] Shao, Y., Wang, B., Sornborger, A. & Tao, L. 1256
1193 A mechanism for synaptic copy between neural cir- 1257
1194 cuits. *bioRxiv* (2018). URL [https://www.biorxiv.org/](https://www.biorxiv.org/content/early/2018/06/20/351114)
1195 [content/early/2018/06/20/351114](https://www.biorxiv.org/content/early/2018/06/20/351114). 1259
- 1196 [63] LeCun, Y., Bottou, L., Bengio, Y. & Haffner, P. 1260
1197 Gradient-based learning applied to document recogni- 1261
1198 tion. *Proceedings of the IEEE* **86**, 2278–2324 (1998). 1262
- 1199 [64] Hubara, I., Courbariaux, M., Soudry, D., El-Yaniv, 1263
1200 R. & Bengio, Y. Binarized neural networks. In *Ad- 1264*
1201 *vances in neural information processing systems*, 4107– 1265
1202 4115 (2016). 1266
- 1203 [65] Hebb, D. *The Organization of Behavior: A Neuropsy- 1267*
1204 *chological Approach* (John Wiley & Sons, 1949). 1268
- 1205 [66] Sornborger, A. & Tao, L. Exact, dynamically routable 1269
1206 current propagation in pulse-gated synfire chains. 1270
1207 *ArXiv:1410.1115* (2014). 1271
- 1208 [67] Senn, W. & Fusi, S. Learning only when necessary: 1272
1209 better memories of correlated patterns in networks with 1273
1210 bounded synapses. *Neural Computation* **17**, 2106–2138 1274
1211 (2005). 1275
- 1212 [68] Chen, G. K., Kumar, R., Sumbul, H. E., Knag, P. C. & 1276
Krishnamurthy, R. K. A 4096-neuron 1m-synapse 3.8- 1277
pj/sop spiking neural network with on-chip stdp learn- 1278
ing and sparse weights in 10-nm finfet cmos. *IEEE Jour- 1279*
nal of Solid-State Circuits **54**, 992–1002 (2018). 1280
- [69] Lin, C.-K. *et al.* Programming spiking neural networks 1281
on intel’s loihi. *Computer* **51**, 52–61 (2018). 1282
- [70] Göltz, J. *et al.* Fast and deep: energy-efficient neuro- 1283
morphic learning with first-spike times. *arXiv preprint*
arXiv:1912.11443v4 (2019). 1284
- [71] Davies, M. *et al.* Advancing neuromorphic computing 1285
with loihi: A survey of results and outlook. *Proceedings*
of the IEEE 1–24 (2021). 1286
- [72] Esser, S. K., Appuswamy, R., Merolla, P., Arthur, J. V. 1287
& Modha, D. S. Backpropagation for energy-efficient 1288
neuromorphic computing. *Advances in neural informa- 1289*
tion processing systems **28**, 1117–1125 (2015). 1290
- [73] Stromatias, E. *et al.* Scalable energy-efficient, low- 1291
latency implementations of trained spiking deep belief 1292
networks on spinnaker. In *2015 International Joint*
Conference on Neural Networks (IJCNN), 1–8 (IEEE, 1293
2015). 1294
- [74] Jin, Y., Zhang, W. & Li, P. Hybrid macro/micro level 1295
backpropagation for training deep spiking neural net- 1296
works. *arXiv preprint arXiv:1805.07866* (2018). 1297
- [75] Neftci, E. O., Augustine, C., Paul, S. & Detorakis, G. 1298
Event-driven random back-propagation: Enabling neu- 1299
romorphic deep learning machines. *Frontiers in neuro- 1300*
science **11**, 324 (2017). 1301
- [76] Shrestha, A., Fang, H., Wu, Q. & Qiu, Q. Approximat- 1302
ing back-propagation for a biologically plausible local 1303
learning rule in spiking neural networks. In *Proceed- 1304*
ings of the International Conference on Neuromorphic
Systems, 1–8 (2019). 1305
- [77] Tavanaei, A. & Maida, A. Bp-stdp: Approximating 1306
backpropagation using spike timing dependent plastic- 1307
ity. *Neurocomputing* **330**, 39–47 (2019). 1308
- [78] Mostafa, H. Supervised learning based on temporal 1309
coding in spiking neural networks. *IEEE transactions*
on neural networks and learning systems **29**, 3227–3235 1310
(2017). 1311
- [79] Lee, J. H., Delbruck, T. & Pfeiffer, M. Training deep 1312
spiking neural networks using backpropagation. *Front- 1313*
iers in neuroscience **10**, 508 (2016). 1314
- [80] O’Connor, P. & Welling, M. Deep spiking networks. 1315
arXiv preprint arXiv:1602.08323 (2016). 1316
- [81] Diehl, P. U. & Cook, M. Unsupervised learning of 1317
digit recognition using spike-timing-dependent plastic- 1318
ity. *Frontiers in computational neuroscience* **9**, 99 1319
(2015). 1320
- [82] Zipser, D. & Rumelhart, D. Neurobiological significance 1321
of new learning models. In *Computational Neuroscience*
(1990). 1322
- [83] Kolen, J. F. & Pollack, J. B. Backpropagation with- 1323
out weight transport. In *Proceedings of 1994 IEEE In- 1324*
ternational Conference on Neural Networks (ICNN’94), 1325
vol. 3, 1375–1380 (IEEE, 1994). 1326
- [84] Stöckl, C. & Maass, W. Optimized spiking neurons can 1327
classify images with high accuracy through temporal 1328
coding with two spikes. *Nature Machine Intelligence*
1–9 (2021). 1329
- [85] Baddeley, R. *et al.* Responses of neurons in primary 1330
and inferior temporal visual cortices to natural scenes. 1331
Proceedings of the Royal Society of London. Series B:
Biological Sciences **264**, 1775–1783 (1997). 1332

- [86] Comsa, I. M. *et al.* Temporal coding in spiking neural networks with alpha synaptic function. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 8529–8533 (IEEE, 2020).
- [87] Rueckauer, B. & Liu, S.-C. Conversion of analog to spiking neural networks using sparse temporal coding. In *2018 IEEE International Symposium on Circuits and Systems (ISCAS)*, 1–5 (IEEE, 2018).
- [88] Baumgartner, S. *et al.* Visual pattern recognition with on-chip learning: towards a fully neuromorphic approach. In *Proceedings of the IEEE International Symposium on Circuits and Systems (ISCAS)* (IEEE, 2020).
- [89] Liang, D. *et al.* Neural state machines for robust learning and control of neuromorphic agents. *IEEE Journal on Emerging and Selected Topics in Circuits and Systems* **9**, 679–689 (2019).
- [90] Riehle, A., Grün, S., Diesmann, M. & Aertsen, A. Spike synchronization and rate modulation differentially involved in motor cortical function. *Science* **278**, 1950–1953 (1997).
- [91] Abeles, M., Bergman, H., Margalit, E. & Vaadia, E. Spatiotemporal firing patterns in the frontal cortex of behaving monkeys. *Journal of neurophysiology* **70**, 1629–1638 (1993).
- [92] Hahnloser, R. H., Kozhevnikov, A. A. & Fee, M. S. An ultra-sparse code underlies the generation of neural sequences in a songbird. *Nature* **419**, 65–70 (2002).
- [93] Ikegaya, Y. *et al.* Synfire chains and cortical songs: temporal modules of cortical activity. *Science* **304**, 559–564 (2004).
- [94] Foster, D. J. & Wilson, M. A. Reverse replay of behavioural sequences in hippocampal place cells during the awake state. *Nature* **440**, 680–683 (2006).
- [95] Rajan, K., Harvey, C. D. & Tank, D. W. Recurrent Network Models of Sequence Generation and Memory. *Neuron* **90**, 128–142 (2016).
- [96] Pang, R. & Fairhall, A. L. Fast and flexible sequence induction in spiking neural networks via rapid excitability changes. *Elife* **8**, e44324 (2019).
- [97] Malvache, A., Reichinnek, S., Villette, V., Haimerl, C. & Cossart, R. Awake hippocampal reactivations project onto orthogonal neuronal assemblies. *Science* **353**, 1280–1283 (2016).
- [98] Luczak, A., McNaughton, B. L. & Harris, K. D. Packet-based communication in the cortex. *Nature Reviews Neuroscience* **16**, 745–755 (2015).
- [99] Simons, T. & Lee, D.-J. A review of binarized neural networks. *Electronics* **8**, 661 (2019).
- [100] Marr, D. & Poggio, T. From understanding computation to understanding neural circuitry (1976).
- [101] Bengio, Y., Léonard, N. & Courville, A. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432* (2013).
- [102] He, K., Zhang, X., Ren, S. & Sun, J. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE international conference on computer vision*, 1026–1034 (2015).

Acknowledgements

This work was carried out at Los Alamos National Laboratory under the auspices of the National Nuclear Security Administration of the U.S. Department of Energy under Contract No. 89233218CNA000001 - SEEK: Scoping neuromorphic architecture impact enabling advanced sensing capabilities. Additional funding was provided by the LANL ASC Beyond Moore’s Law program and by LDRD Reserve Project 20180737ER - Neuromorphic Implementation of the Backpropagation Algorithm. L.T. acknowledges support from the Natural Science Foundation of China grants 31771147 (L.T.) and 91232715 (L.T.). L.T. thanks the Los Alamos National Laboratory for its hospitality. A.R. acknowledges support from the Swiss National Science Foundation (SNSF) grant Ambizione PZOOP2 168183 ELMA. F.S. carried out work under the kind support of the Center for Nonlinear Studies fellowship and final preparations at the London Institute for Mathematical Sciences. The Authors thank Intel Labs for providing support and access to the Loihi API and hardware.

Author Contributions

The authors contributed equally to the methodology presented here. Alpha Renner adapted and implemented the algorithm on Intel’s Loihi chip. Alpha Renner, Forrest Sheldon, and Anatoly Zlotnik formalized neuromorphic information-processing mechanisms, implemented the algorithm in simulation and hardware, and developed figures. All authors wrote the manuscript with Renner, Sheldon, and Zlotnik doing the bulk of the writing. Andrew Sornborger and Anatoly Zlotnik supervised the research and Sornborger and Louis Tao developed the concepts and algorithmic basis of the neuromorphic backpropagation algorithm and circuit structure.

Competing Financial Interests

The authors declare that they have no competing financial interests.

Supplementary Material

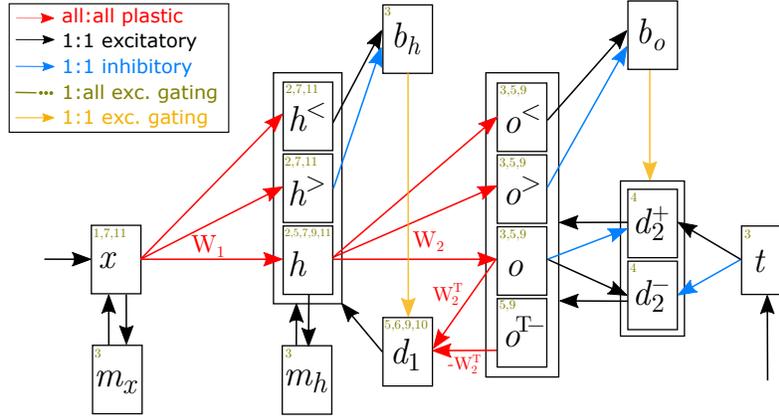


FIG. 4. **Anatomical connectivity of the 2 layer backpropagation circuit.** While in the Loihi implementation the o layers are connected directly go to the d_2 layers, here an intermediate fictional b_o layer is added for easier understanding. Arrows that end on the border of a box that encompasses several layers go to each of the layers. The gating chain is not shown, but the small numbers on top of each layer indicate when it is gated on. Colors are the same as in Fig. 2.

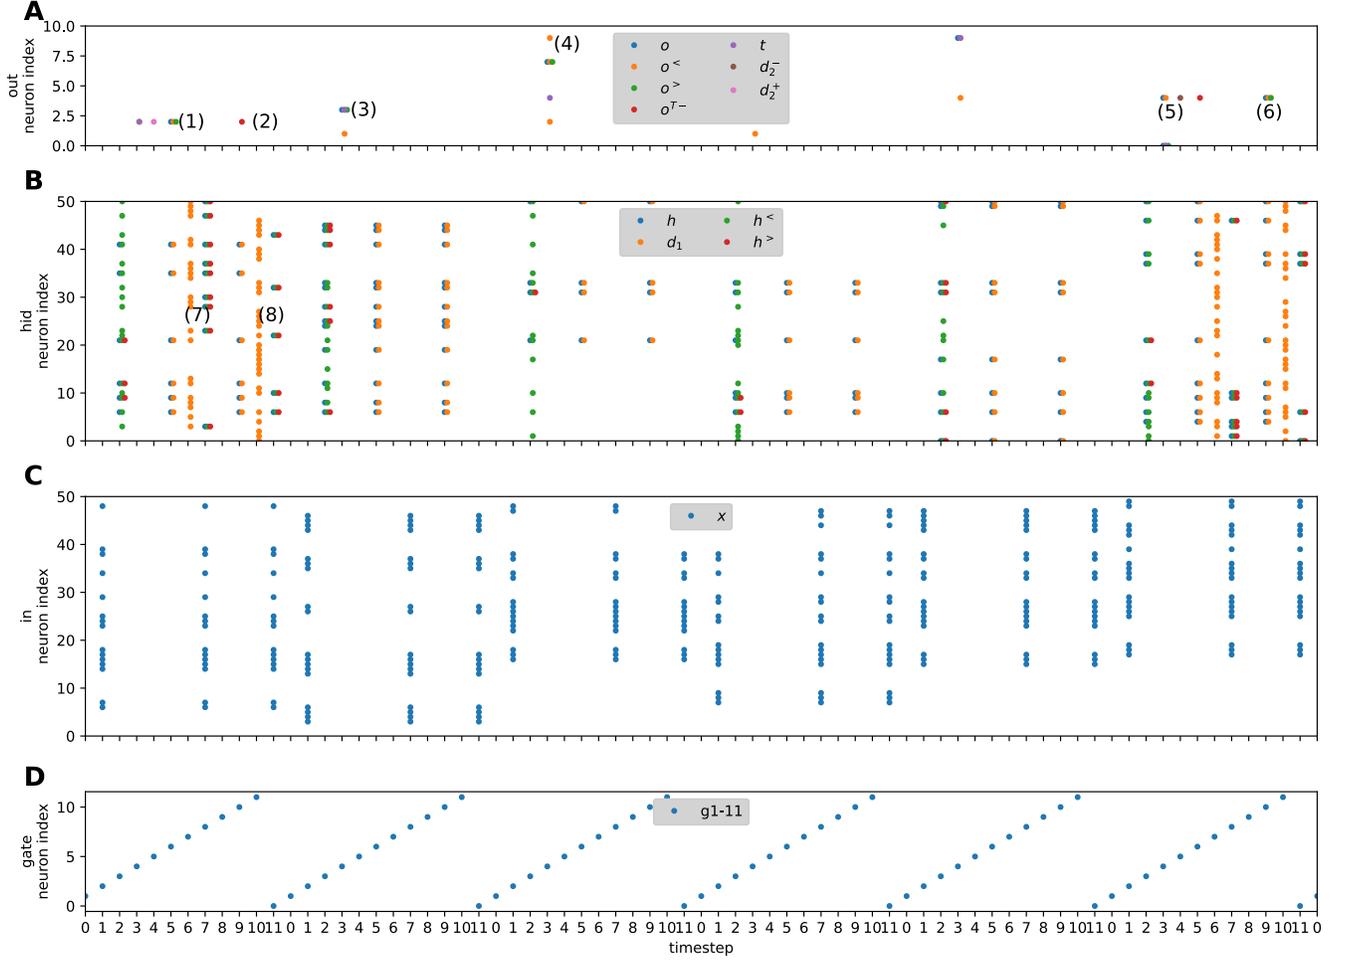


FIG. 5. Example raster plot of the spikes over six gating cycles. All populations of the same size are plotted in the same plot and only the first 50 neurons are plotted per layer. To avoid occlusion, a small offset in time is added to the time step of some layers. Refer to Tab. II for a detailed explanation of the spike propagation. (1) error (target but no output spike) leads to potentiation of the W_2 synaptic weight and the positive transpose; (2) the same error leads to depression of the negative transpose via activity of d_1 ; (3) no error because o and t fire at the same location, i.e. there is no update in this iteration; (4) there is an error (t fires at index 4, but o at index 7), but the local gradient is 0 because it is gated ‘off’ at index 7 because the derivative of the activation function is 0, i.e. both $o^<$ and $o^>$ fire. Also, it is not gated ‘on’ at index 4, because $o^<$ does not fire; (5) local gradient (output but not target), leads to potentiation of the weight of the synapses from o^{T-} to d_1 (red), and (6) depression of $h - o$ and $d_1 - o$ synaptic weights; (7) The orange spikes show the back-propagated local gradient from (1) which leads to potentiation of the $x - h$ weights. Note that for visualization purposes, the gating from b_h is applied one time step later directly to $h, h^<$ and $h^>$. That is, the orange spikes in time step 7 are the full backpropagated error, but only the neurons that are also gated ‘on’ by the combination of $h^<$ and $h^>$ are actually active in the potentiation phase in time step 8. (8) Same as (7), but the error from (2) leads to depression of the $x - h$ weights.

Step	1	2	3	4	5
	(x, x, g_1)	(m_x, x, g_2) $(h^<, f(W_1 x + 0.5), g_2)$ $(h^>, f(W_1 x - 0.5), g_2)$ $(h, f(W_1 x), g_2)$	(m_h, h, g_3) $(o^<, f(W_2 h + 0.5), g_3)$ $(o^>, f(W_2 h - 0.5), g_3)$ $(o, f(W_2 h), g_3)$ (t, t, g_3) $(b_h, h^< - h^>, g_3)$	$(d_2^+, f(t - o) \circ (o^< - o^>), o^< - o^>)$ $(d_2^-, f(o - t) \circ (o^< - o^>), o^< - o^>)$	(h, h, g_5) (h^T, h, g_5) (o, d_2^+, g_5) $(o^<, d_2^+, g_5)$ $(o^>, d_2^+, g_5)$ (o^{T-}, d_2^-, g_5)
	ff in	ff hid	ff out	error $\circ b_o$	potentiation $w_{hid,out}$

Step	6	7	8	9	10	11
	$(h^T, H(W_2^T(d_2^+) - W_2^T(d_2^-)))$ $\equiv d_1^+, g_6)$	(x, x, g_7) $(h, b_h \circ d_1^+, b_h)$ $(h^<, b_h \circ d_1^+, b_h)$ $(h^>, b_h \circ d_1^+, b_h)$		(h, h, g_9) (h^T, h, g_9) (o, d_2^-, g_9) $(o^<, d_2^-, g_9)$ $(o^>, d_2^-, g_9)$ (o^{T-}, d_2^+, g_9)	$(h^T, H(-W_2^T(d_2^-) + W_2^T(d_2^+)))$ $\equiv d_1^-, g_{10})$	(x, x, g_{11}) $(h, b_h \circ d_1^-, b_h)$ $(h^<, b_h \circ d_1^-, b_h)$ $(h^>, b_h \circ d_1^-, b_h)$
	backpropagation	potentiation $w_{in,hid}$		depression $w_{hid,out}$	backpropagation	depression $w_{in,hid}$

TABLE II. Information flow through the backpropagation network (see Fig. 2). Gating pulses ($g_1 - g_{11}$) are sent from the gating chain. The triplet notation (a, b, c) denotes information b in population a gated by population c .

	hardware	Power (mW)			Latency (ms)	Dyn. Energy per sample (mJ)	Energy delay product(μ Js)
		Static	Dynamic	Total			
Loihi training ds100-300-10	x86 cores	0.119	19.0	19.1	0.731	0.0139	0.0101
	neuron cores	35.7	193	229		0.141	0.103
	whole board	916	212	1128		0.155	0.113
Loihi training 400-300-10	x86 cores	0.133	18.9	19.0	1.18	0.0223	0.0264
	neuron cores	40.0	399	439		0.471	0.557
	whole board	914	418	1332		0.494	0.583
Loihi training 400-400-10 start	x86 cores	0.137	18.5	18.6	1.49	0.0275	0.0409
	neuron cores	37.6	398	435		0.591	0.88
	whole board	912	416	1328		0.619	0.92
Loihi training 400-400-10 end	x86 cores	0.15	18.5	18.7	1.48	0.0275	0.0408
	neuron cores	41.1	399	440		0.592	0.878
	whole board	914	418	1332		0.62	0.919
Loihi 400-400-10 end (learning engine off)	x86 cores	0.125	22.2	22.3	0.37	0.00822	0.00304
	neuron cores	34.1	55.1	89.2		0.0204	0.00755
	whole board	954	77.3	1031		0.0286	0.0106
Loihi inference net 400-400-10	x86 cores	0.143	23.4	23.6	0.169	0.00396	0.000669
	neuron cores	39.2	14.7	53.9		0.00249	0.000421
	whole board	980	38.2	1019		0.00645	0.00109
GPU (Nvidia GTX 1660S) training	batchsize 1	15000	19000	34000	1.0	19.0	19.0
	batchsize 10	15000	19000	34000	0.1	1.9	0.19
	batchsize 100	15000	19000	34000	0.01	0.19	0.0019

TABLE III. Breakdown of power consumption of the chip for different cases. ‘ds100-300-10’, means that the network is run using 100 input and 300 hidden layer neurons and with the input downsampled by 2. Here, ‘start’ means that the measurement is taken in the first 10000 iterations of the first training epoch, and ‘end’ means that it has been taken with the fully trained model. We use ‘learning engine off’ to say that the weights are set to nonplastic weights after training so that the on-chip circuitry that handles learning is not active. The inference net, which uses the same weights as the fully trained network, consists only of the 3 feedforward layers and the gating chain.

1. Power Measurements on GPU

1374
 1375 The power on the GPU was measured using `nvidia-smi` while running the TensorFlow implementation of our
 1376 algorithm. The dynamic power was calculated by taking the difference between the power reading before the start
 1377 of the run and during the run. The system is running Tensorflow version 2.3 on Windows 10 with an Intel i5-9600K
 1378 processor and 16 GB of RAM. The GPU is an NVIDIA GeForce GTX 1600S (driver version: 461.92, CUDA version:
 1379 11.2).

2. Detailed Description of the Algorithm

1380
 1381 Here we go through all steps of the algorithm as shown in Fig. 2:

- 1382 1. MNIST images are sent to the x layer in binarized form
- 1383 2. W_1 is applied three times with different offsets (received from the gating chain) before application of the nonlinear
 1384 activation function. This yields the hidden layer activity and start and stop learning conditions for the W_1 update
 1385 according to Eq. (9)
- 1386 3. The MNIST labels are sent as a one-hot vector to the target layer and W_2 is applied three times with different
 1387 offsets. This yields the output layer activity and start and stop learning conditions for the W_2 update according
 1388 to Eq. (9). Furthermore, the input and hidden layer activities are stored into the m_x and m_h relay layers for
 1389 later usage in the learning phases. In the box function layer (b_h), the start and stop learning conditions of W_1
 1390 are combined using an excitatory and inhibitory connection. This corresponds to the multiplication in Eq. (9)
 1391 as in this case, multiplication is the same as taking the difference with the inverted term.
- 1392 4. The positive and negative errors are calculated by excitation from the o layer and inhibition from the t layer (and
 1393 vice versa). The local gradient layers are not gated on by the gating chain but by the start learning condition
 1394 $o^<$, and furthermore, they are inhibited by the stop learning condition $o^>$. This way, the Hadamard product
 1395 from Eq. (5) is calculated.
- 1396 5. The positive local gradient is now sent to the output layer (and its two siblings) while at the same time the
 1397 hidden layer activity from time step 2 is reactivated from the relay layer. This leads to a positive weight update
 1398 of W_2 , and because also the d_1 layer is activated, to updates of W_2^T and W_2^{T-} according to the learning rule in
 1399 Eq. (21). Weights W_2^{T-} are however potentiated oppositely because layer o^{T-} receives input from layer d^- in
 1400 this time step. This time step implements the first term of Eq. (25).
- 1401 6. The actual backpropagation happens in this time step. By way of the transposed weight matrices, the d_1
 1402 layer receives activity from the output layer and the o^{T-} layer, which contains the positive and negative local
 1403 gradients. The layer is gated by the box function layer b_h , which calculates the Hadamard product from Eq. (6).
 1404 This gating includes an offset so that even the smallest possible nonzero input sum leads to a spike in the
 1405 neurons of the d_1 layer.
- 1406 7. This time step realizes the potentiation of the W_1 weights. The hidden layer and its two siblings receive the
 1407 local gradient that was propagated back in the previous time step, and the x layer restores its previous activity
 1408 from the relay layer. Because the third factor is switched on in this time step, potentiation happens according
 1409 to Eq. (21). This time step implements the first term of Eq. (28).
- 1410 8. To avoid interference of the input to the output layer from the previously active hidden layer (layer h is connected
 1411 to $o, o^<$ and $o^>$), no layer is gated on in this time step.
- 1412 9. The steps 9 to 11 are the same as 5 to 7, but with input of the opposite local gradient and without third factor
 1413 active and therefore the opposite sign of learning (depression).

3. Synapse Types

1414
 1415 There are five types of synapses (see different colors in Fig. 2 and Supplementary Fig. 4):

- 1416 • Plastic all-to-all synapses (red), that change weights according to a learning rule (Eq. (21)) and can be both
 1417 positive or negative ($w_p = [-256, 254]$).
- 1418 • Excitatory (black) one-to-one synapses with a fixed weight ($w_e = V_{\text{thr}} = 1024$) to copy the firing pattern of the
 1419 source layer to the target layer.
- 1420 • Inhibitory (blue) one-to-one synapses with a fixed weight ($w_i = -V_{\text{thr}} = -1024$) to always inhibit the target
 1421 neurons.

- 1422 • Excitatory one-to-all gating (green) synapses from a single gating chain neuron to all neurons of a layer with a
1423 fixed weight of usually $w_g = -I_{const} + 0.5V_{thr} = 8704$. The usual gating synapses are used to gate ‘on’ a layer
1424 that is supposed to have the feedforward activation function f . Different weights $w_{g<} = -I_{const} = 8192$ and
1425 $w_{g>} = -I_{const} + V_{thr} = 9216$ are used to gate ‘on’ the start and stop learning conditions to get the two step
1426 functions with which the box function is built.
- 1427 • Excitatory one-to-one gating (yellow) synapses from the start learning layer $o^<$ and the box function layer b_h
1428 to the respective local gradient layers. The w_g weight is used for gating of the d_2 layers in time step 4, and the
1429 $w_{g>}$ weight is used for gating of the d_1 layer in time steps 6 and 10.

1430 To a certain extent within the precision boundaries of the chip, the scaling of the weights is arbitrary. All plastic
1431 synapses and most other synapses have a synaptic delay of 0 time steps. Some synapses, specifically the ones originating
1432 from the relay layers, the d_2 layers, and the b_h layer, have delays up to 6 time steps, i.e. their output spikes affect the
1433 target neuron several time steps later. The delays, corresponding to the horizontal arrow length, can be read from
1434 Fig. 2.

1435 4. Downsampled MNIST

1436 Table III above refers to network structures with input sizes of 400 and 100. The 400-400-10 network used a set
1437 of trimmed images, where whitespace was removed to reduce the required size of the input layer. MNIST images for
1438 this had a border of width 4 removed. For the smaller 100-300-10 network, these images were then downsampled by
1439 averaging over 2×2 pixel grids and binarized by thresholding at 0.5. This gave a set of 10×10 binary images.

1440 5. Tensorflow Implementation

1441 For the feedforward network, our GPU tensorflow implementation is similar to the binary network presented in
1442 [64] but with continuous weights, and simplified to correspond with the neuromorphic design. The loss function was
1443 modified to mean squared error and minimized with stochastic gradient descent with a fixed learning rate. Dropout
1444 and batch normalization were removed and the batch size was set to 1.

1445 The network was initialized with Glorot parameters for each layer l , i.e., $\gamma_l = \sqrt{\frac{1.5}{N_{l-1} + N_l}}$ where N_l is the number
1446 of neurons in layer l . For each layer, weights were generated uniformly on the interval $[-\gamma_l, \gamma_l]$ and clipped to these
1447 values during learning. While these initializations do not take into account the binary activation functions of our
1448 network, tests with alternative initializations showed little or no improvement, likely due to the shallow depth of the
1449 networks considered.

1450 Gradients on the binary network in [64] were calculated with a ‘straight-through’ estimator in which the derivative
1451 of the binary activation function

$$\sigma(x) = \begin{cases} 1 & x \geq 0 \\ 0 & \text{otherwise} \end{cases} \quad (29)$$

1452 was approximated as

$$h_l(x) = \begin{cases} 0 & x \geq \gamma_l \\ \frac{1}{2\gamma_l} & -\gamma_l < x < \gamma_l \\ 0 & x < -\gamma_l \end{cases} \quad (30)$$

1453 This same gradient calculation is maintained, but to mimic the binary gradient signal in the neuromorphic implemen-
1454 tation, the final gradient with respect to each weight is thresholded. If the gradient calculated with respect to weight
1455 i is g_i , the applied gradient in the learning step is $\text{sgn}(g_i)\sigma(|g_i| - t)$ for a threshold t representing the threshold of
1456 signal propagation in the neuromorphic implementation and σ as defined above.