



# Accelerated Physical Emulation of Bayesian Inference in Spiking Neural Networks

Akos F. Kungl<sup>1\*</sup>, Sebastian Schmitt<sup>1</sup>, Johann Klähn<sup>1</sup>, Paul Müller<sup>1</sup>, Andreas Baumbach<sup>1</sup>, Dominik Dold<sup>1</sup>, Alexander Kugele<sup>1</sup>, Eric Müller<sup>1</sup>, Christoph Koke<sup>1</sup>, Mitja Kleider<sup>1</sup>, Christian Mauch<sup>1</sup>, Oliver Breitwieser<sup>1</sup>, Luziwei Leng<sup>1</sup>, Nico Gürtler<sup>1</sup>, Maurice Güttler<sup>1</sup>, Dan Husmann<sup>1</sup>, Kai Husmann<sup>1</sup>, Andreas Hartel<sup>1</sup>, Vitali Karasenko<sup>1</sup>, Andreas Grübl<sup>1</sup>, Johannes Schemmel<sup>1</sup>, Karlheinz Meier<sup>1</sup> and Mihai A. Petrovici<sup>1,2</sup>

<sup>1</sup> Kirchoff-Institute for Physics, Heidelberg University, Heidelberg, Germany, <sup>2</sup> Department of Physiology, University of Bern, Bern, Switzerland

## OPEN ACCESS

### Edited by:

Themis Prodromakis,  
University of Southampton,  
United Kingdom

### Reviewed by:

Alexanrou Serb,  
University of Southampton,  
United Kingdom  
Adnan Mehonic,  
University College London,  
United Kingdom

### \*Correspondence:

Akos F. Kungl  
fkungl@kip.uni-heidelberg.de

### Specialty section:

This article was submitted to  
Neuromorphic Engineering,  
a section of the journal  
Frontiers in Neuroscience

**Received:** 10 July 2019

**Accepted:** 23 October 2019

**Published:** 14 November 2019

### Citation:

Kungl AF, Schmitt S, Klähn J, Müller P, Baumbach A, Dold D, Kugele A, Müller E, Koke C, Kleider M, Mauch C, Breitwieser O, Leng L, Gürtler N, Güttler M, Husmann D, Husmann K, Hartel A, Karasenko V, Grübl A, Schemmel J, Meier K and Petrovici MA (2019) Accelerated Physical Emulation of Bayesian Inference in Spiking Neural Networks. *Front. Neurosci.* 13:1201. doi: 10.3389/fnins.2019.01201

The massively parallel nature of biological information processing plays an important role due to its superiority in comparison to human-engineered computing devices. In particular, it may hold the key to overcoming the von Neumann bottleneck that limits contemporary computer architectures. Physical-model neuromorphic devices seek to replicate not only this inherent parallelism, but also aspects of its microscopic dynamics in analog circuits emulating neurons and synapses. However, these machines require network models that are not only adept at solving particular tasks, but that can also cope with the inherent imperfections of analog substrates. We present a spiking network model that performs Bayesian inference through sampling on the BrainScaleS neuromorphic platform, where we use it for generative and discriminative computations on visual data. By illustrating its functionality on this platform, we implicitly demonstrate its robustness to various substrate-specific distortive effects, as well as its accelerated capability for computation. These results showcase the advantages of brain-inspired physical computation and provide important building blocks for large-scale neuromorphic applications.

**Keywords:** physical models, neuromorphic engineering, massively parallel computing, spiking neurons, recurrent neural networks, neural sampling, probabilistic inference

## 1. INTRODUCTION

The aggressive pursuit of Moore's law in conventional computing architectures is slowly but surely nearing its end (Waldrop, 2016), with difficult-to-overcome physical effects, such as heat production and quantum uncertainty, representing the main limiting factors. The so-called von Neumann bottleneck between processing and memory units represents the main cause, as it effectively limits the speed of these largely serial computation devices. The most promising solutions come in the form of massively parallel devices, many of which are based on brain-inspired computing paradigms (Indiveri et al., 2011; Furber, 2016), each with its own advantages and drawbacks.

Among the various approaches to such neuromorphic computing, one class of devices is dedicated to the physical emulation of cortical circuits; not only do they instantiate neurons and synapses that operate in parallel and independently of each other, but these units are actually

represented by distinct circuits that emulate the dynamics of their biological archetypes (Mead, 1990; Indiveri et al., 2006; Jo et al., 2010; Schemmel et al., 2010; Pfeil et al., 2013; Qiao et al., 2015; Chang et al., 2016; Wunderlich et al., 2019). Some important advantages of this approach lie in their reduced power consumption and enhanced speed compared to conventional simulations of biological neuronal networks, which represent direct payoffs of replacing the resource-intensive numerical calculation of neuro-synaptic dynamics with the physics of the devices themselves.

However, such computation with analog dynamics, without the convenience of binarization, as used in digital devices, has a downside of its own: variability in the manufacturing process (fixed pattern noise) and temporal noise both lead to reduced controllability of the circuit dynamics. Additionally, one relinquishes much of the freedom permitted by conventional algorithms and simulations, as one is confined by the dynamics and parameter ranges cast into the silicon substrate. The main challenge of exploiting these systems, therefore, lies in designing performance network models using the available components while maintaining a degree of robustness toward the substrate-induced distortions. Just like for the devices themselves, inspiration for such models often comes from neuroscience, as the brain needs to meet similar demands.

With accumulating experimental evidence (Berkes et al., 2011; Pouget et al., 2013; Haefner et al., 2016; Orbán et al., 2016), the view of the brain itself as an analytical computation device has shifted. The stochastic nature of neural activity *in vivo* is being increasingly regarded as an explicit computational resource rather than a nuisance that needs to be dealt with by sophisticated error-correcting mechanisms or by averaging over populations. Under the assumption that stochastic brain dynamics reflect an ongoing process of Bayesian inference in continuous time, the output variability of single neurons can be interpreted as a representation of uncertainty. Theories of neural sampling (Buesing et al., 2011; Hennequin et al., 2014; Aitchison and Lengyel, 2016; Petrovici et al., 2016; Kutschireiter et al., 2017) provide an analytical framework for embedding this type of computation in spiking neural networks.

In this paper we describe the realization of neural sampling with networks of leaky integrate-and-fire neurons (Petrovici et al., 2016) on the BrainScaleS accelerated neuromorphic platform (Schemmel et al., 2010). With appropriate training, the variability of the analog components can be naturally compensated and incorporated into a functional network structure, while the network's ongoing dynamics make explicit use of the analog substrate's intrinsic acceleration for Bayesian inference (section 2.3). We demonstrate sampling from low-dimensional target probability distributions with randomly chosen parameters (section 3.1) as well as inference in high-dimensional spaces constrained by real-world data, by solving associated classification and constraint satisfaction problems (pattern completion, section 3.2). All network components are fully contained on the neuromorphic substrate, with external inputs only used for sensory evidence (visual data). Our work thereby contributes to the search for novel paradigms of

information processing that can directly benefit from the features of neuro-inspired physical model systems.

## 2. METHODS

### 2.1. The BrainScaleS System

BrainScaleS (Schemmel et al., 2010) is a mixed-signal neuromorphic system, realized in 180 nm CMOS technology, that emulates networks of spiking neurons. Each BrainScaleS wafer module consists of a 20 cm silicon wafer with 384 HICANN (High Input Count Analog Neural Network) chips, see **Figure 1A**. On each chip, 512 analog circuits emulate the adaptive exponential integrate-and-fire (AdEx) model (Brette and Gerstner, 2005; Millner et al., 2010) of spiking neurons with conductance-based synapses. The dynamics evolve with an acceleration factor of  $10^4$  with respect to biological time, i.e., all specific time constants (synaptic, membrane, adaptation) are  $\sim 10^4$  times smaller than typical corresponding values found in biology (Schemmel et al., 2010; Petrovici et al., 2014). To preserve compatibility with related literature (Petrovici et al., 2016; Schmitt et al., 2017; Leng et al., 2018; Dold et al., 2019), we refer to system parameters in the biological domain unless otherwise specified, e.g., a membrane time constant given as 10 ms is actually accelerated to 1  $\mu$ s on the chip.

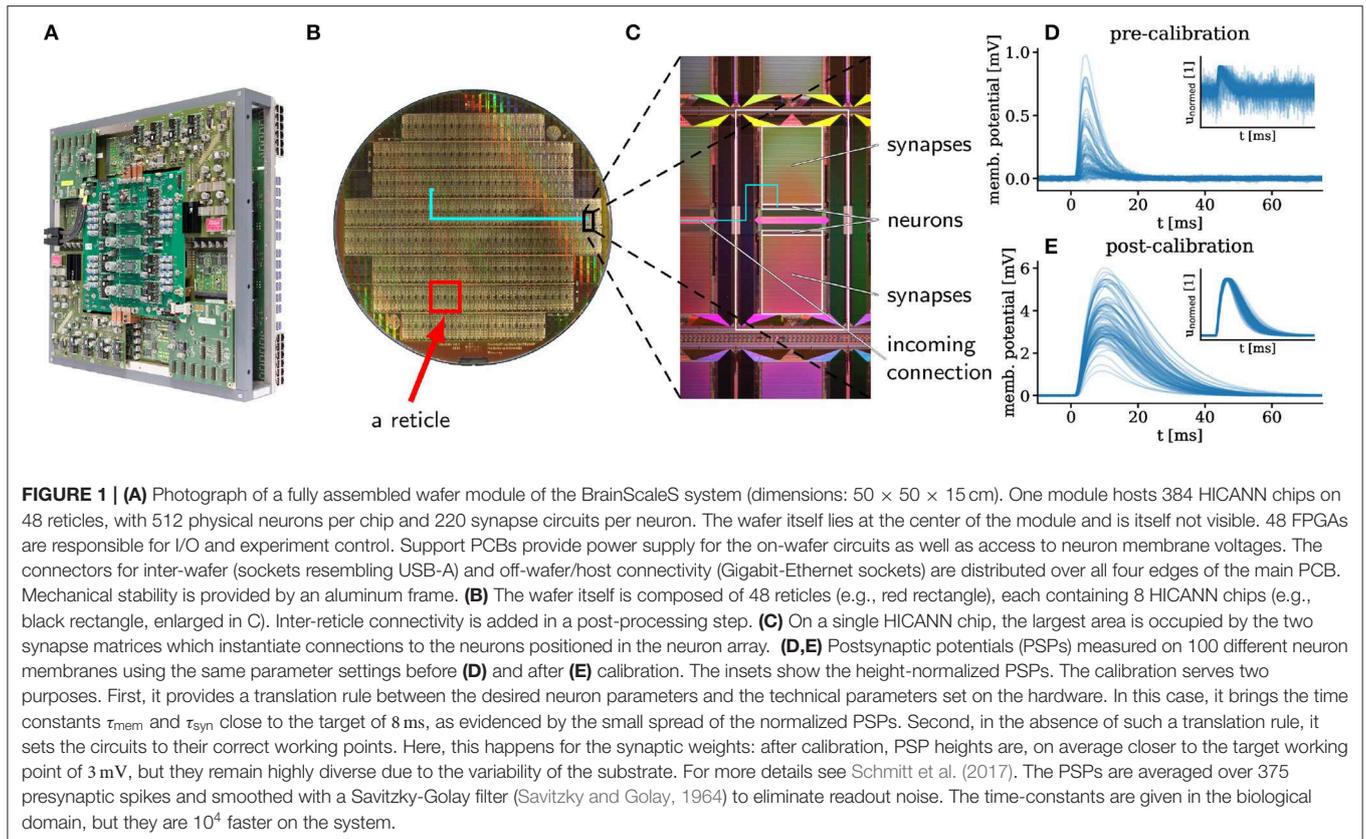
The parameters of the neuron circuits are stored in analog memory cells (floating gates) with 10 bit resolution, and the synaptic weights are stored in 4 bit SRAM (Schemmel et al., 2010). The analog memory cells are similar to the ones in Lande et al. (1996), and they are described in Looch (2006) and Millner (2012).

Spike events are transported digitally and can reach all other neurons on the wafer with the help of an additional redistribution layer that instantiates an on-wafer circuit-switched network (Zoschke et al., 2017) (**Figures 1B,C**).

Because of mismatch effects (fixed-pattern noise) inherent to the substrate, the response to incoming stimuli varies from neuron to neuron (**Figure 1D**). In order to bring all neurons into the desired regime and to reduce the neuron-to-neuron response variability, we employ a standard calibration procedure that is performed only once, during the commissioning of the system (Petrovici et al., 2017b; Schmitt et al., 2017). Nevertheless, even after calibration, a significant degree of diversity persists (**Figure 1E**). The emulation of functional networks that do not rely on population averaging therefore requires appropriate training algorithms (section 3.2).

### 2.2. Sampling With Leaky Integrate-and-Fire Neurons

The theory of sampling with leaky integrate-and-fire neurons (Petrovici et al., 2016) describes a mapping between the dynamics of a population of neurons with conductance-based synapses (equations given in **Table 1**) and a Markov-chain Monte Carlo sampling process from an underlying probability distribution over binary random variables (RVs). Each neuron in such a sampling network corresponds to one of these RVs: if the  $k$ -th neuron has spiked in the recent past and is currently refractory, then it is considered to be in the *on-state*  $z_k = 1$ , otherwise it is in



the *off-state*  $z_k = 0$  (Figures 2A,B). With appropriate synaptic parameters, such a network can approximately sample from a Boltzmann distribution defined by

$$p^*(z) = \frac{1}{Z} \exp\left(\frac{1}{2} z^T W z + z^T b\right), \quad (1)$$

where  $Z$  is the partition sum,  $W$  a symmetric, zero-diagonal effective weight matrix and  $b_i$  the effective bias of the  $i$ -th neuron (Figure 2D).

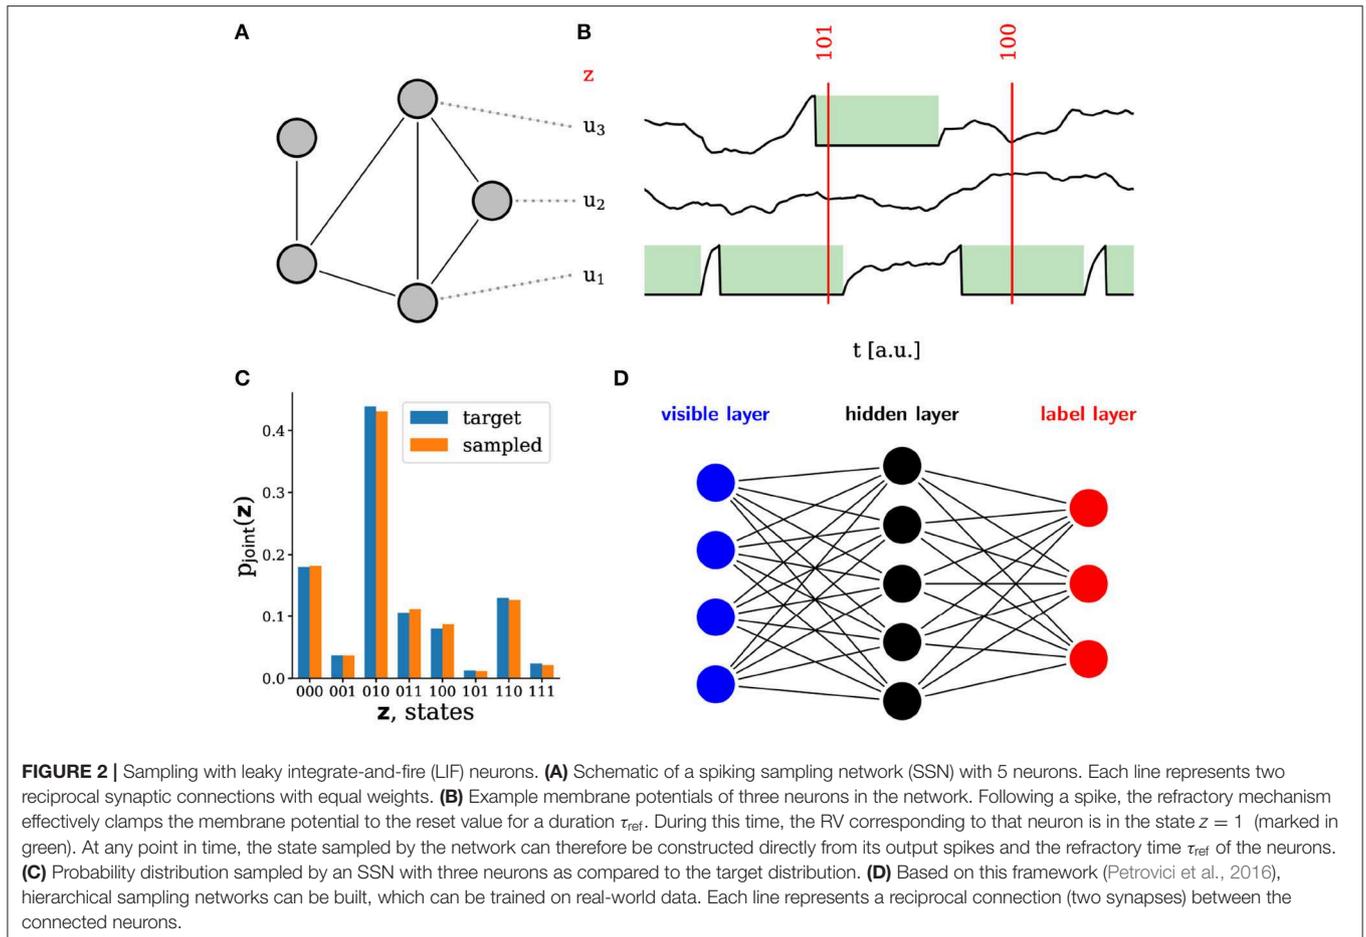
In the original model, each neuron receives excitatory and inhibitory Poisson input. This plays two important roles: it transforms a deterministic LIF neuron into a stochastic firing unit and induces a high-conductance state, with an effective membrane time constant that is much smaller than other time constants in the system:  $\tau_{eff} \ll \tau_{syn}, \tau_{ref}$  (see e.g., Destexhe et al., 2003; Petrovici, 2016), which symmetrizes the neural activation function, as explained in the following. The activation function of an LIF neuron without noise features a sharp onset, but only a slow converge to its maximum value, hence being highly asymmetric around the point of 50 % activity. Background Poisson noise smears out the onset of the activation function, while the reduced membrane time constant accelerates the convergence to the maximum, making the activation function more symmetric and thus more similar to a logistic function, which is a pre-requisite for this form of sampling. For the explicit derivation see Petrovici (2016) and Petrovici et al. (2016). A mapping of this activation function to the abovementioned

**TABLE 1 |** Description of the neuron and synapse model.

Type	Leaky integrate-and-fire (LIF), conductance-based synapse, exponential kernel
Subthreshold dynamics	Subthreshold dynamics [ $t \notin [t_{sp}, t_{sp} + \tau_{ref}]$ ]: $C_m(d/dt)u(t) = -g_l[u(t) - E_{leak}] - g_{syn}^{inh}(t)[u(t) - E_{inh}] - g_{syn}^{exc}(t)[u(t) - E_{exc}]$
Reset and refractoriness	[ $t \in [t_{sp}, t_{sp} + \tau_{ref}]$ ]: $u(t) = V_{reset}$
Spiking	This model was emulated on the BrainScaleS system (Schemmel et al., 2010) If $u(t)$ crosses $V_{thresh}$ from below at $t = t_{sp}$ , neuron emits a spike with timestamp $t_{sp}$
Synapse dynamics	For each presynaptic spike at $t_{sp}$ : $g_{syn}(t) = J \exp[-(t - t_{sp} - d)/(\tau_{syn})] \theta(t - t_{sp} - d)$ where $J$ is the synaptic weight, $d$ the synaptic delay and $\theta$ the Heaviside function This model was emulated on the BrainScaleS system (Schemmel et al., 2010)

The variables are described including their numerical values in the experiment in Table 2.

logistic function  $1/[1 + \exp(-x)]$  provides the translation from the dimensionless weights and biases of the target distribution to the corresponding biological parameters of the spiking network (Petrovici, 2016).



Although different in their dynamics, such sampling spiking networks (SSNs, **Figure 2D**) function similar to (deep) Boltzmann machines (Hinton et al., 1984), which makes them applicable to the same class of machine learning problems (Leng et al., 2018). Training can be done using an approximation of wake-sleep algorithm (Hinton et al., 1995; Hinton, 2012), which implements maximum-likelihood learning on the training set:

$$\Delta b_i = \eta(\langle z_i \rangle^* - \langle z_i \rangle), \quad (2)$$

$$\Delta W_{ij} = \eta(\langle z_i z_j \rangle^* - \langle z_i z_j \rangle), \quad (3)$$

where  $\langle \cdot \rangle$  and  $\langle \cdot \rangle^*$  represent averages over the sampled (model or sleep phase) and target (data or wake phase) distribution, respectively, and  $\eta$  is the learning rate.

In order to enable a fully-contained neuromorphic emulation on the BrainScaleS system, the original model had to be modified. The changes in the network structure, noise generation mechanism, and learning algorithm are described in section 2.3.

For low-dimensional, fully specified target distributions, we used the Kullback-Leibler divergence (DKL, Kullback and Leibler, 1951) as a measure of discrepancy between the sampled

( $p$ ) and the target ( $p^*$ ) distributions:

$$D_{\text{KL}}(p \parallel p^*) = - \sum_{z_i \in \Omega} p(z_i) \ln \left( \frac{p(z_i)}{p^*(z_i)} \right) \quad (4)$$

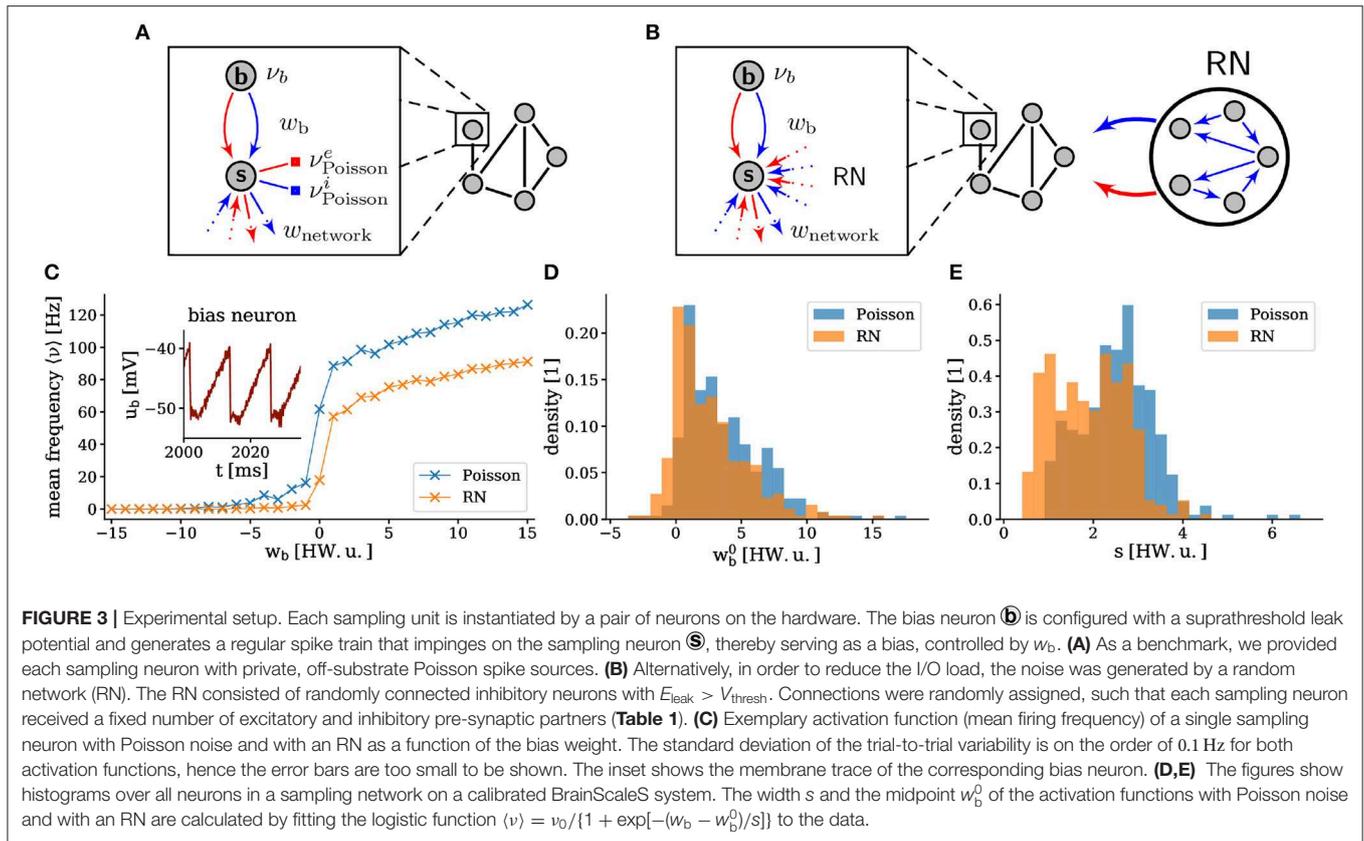
This was done in part to preserve comparability with previous studies (Buesing et al., 2011; Petrovici et al., 2015, 2016), but also because the DKL is the natural loss function for maximum likelihood learning. For visual datasets, we used the error rate (ratio of misclassified images in the test set) for discriminative tasks and the mean squared error (MSE) between reconstruction and original image for pattern completion tasks. The MSE is defined as

$$\text{MSE} = \frac{1}{N_{\text{pixels}}} \sum_{k=1}^{N_{\text{pixels}}} \left( z_k^{\text{data}} - z_k^{\text{recon}} \right)^2, \quad (5)$$

where  $z_k^{\text{data}}$  is the reference data value,  $z_k^{\text{recon}}$  is the model reconstruction and the sum goes over the  $N_{\text{pixels}}$  pixels to be reconstructed by the SSN.

## 2.3. Experimental Setup

The physical emulation of a network model on an analog neuromorphic substrate is not as straightforward as a



software simulation, as it needs to comply with the constraints imposed by the emulating device. Often, it may be tempting to fine-tune the hardware to a specific configuration that fits one particular network, e.g., by selecting specific neuron and synapse circuits that operate optimally given a particular set of network parameters, or by manually tweaking individual hardware parameters after the network has been mapped and trained on the substrate. Here, we explicitly refrained from any such interventions in order to guarantee the robustness and scalability of our results.

All experiments were carried out on a single module of the BrainScaleS system using a subset of the available HICANN chips. The network setup was specified in the BrainScaleS-specific implementation of PyNN (Davison et al., 2009) and the standard calibration (Schmitt et al., 2017) was used to set the analog parameters. The full setup consisted of two main parts: the SSN and the source of stochasticity.

In the original sampling model (Petrovici et al., 2016), in order to affect biases, the wake-sleep algorithm (Equation 1) requires access to at least one reversal potential ( $E_l$ ,  $E_{exc}$ , or  $E_{inh}$ ), which are all controlled by analog memory cells. Given that rewriting analog memory cells is both less precise and slower than rewriting the SRAM cells controlling the synaptic weights, we modified our SSNs to implement biases by means of synaptic weights. To this end, we replaced individual sampling neurons by sampling units, each realized using two hardware neurons (Figures 3A,B). Like in the original model, a sampling neuron was set up to

encode the corresponding binary RV. Each sampling neuron was accompanied by a bias neuron set up with a suprathreshold leak potential that ensured regular firing (Figure 3C, inset). Each bias neuron projected to its target sampling neuron with both an excitatory and an inhibitory synapse (with independent weights), thus inducing a controllable offset of the sampling neuron's average membrane potential. Because excitatory and inhibitory inputs are routed through different circuits for each neuron, two types of synapses were required to allow the sign of the effective bias to change during training. For larger networks, in order to optimize the allocation of hardware resources, we shared the use of bias neurons among multiple sampling neurons (connected via distinct synapses). Similarly, in order to allow sign switches during training, connections between sampling neurons were implemented by pairs of synapses (one excitatory and one inhibitory) as well.

The dynamics of the sampling neurons were rendered stochastic in two different ways. The first setup served as a benchmark and represented a straightforward implementation of the theoretical model from (Petrovici et al., 2016), with Poisson noise generated on the host computer and fed in during the experiment (Figure 3A). In the second setup, we used the spiking activity of a sparse recurrent random network (RN) of inhibitory neurons, instantiated on the same wafer, as a source of noise (Figure 3B). For a more detailed study of sampling-based Bayesian inference with noise generated by deterministic networks, we refer to (Jordan et al., 2017). The mutual inhibition ensured a relatively constant (sub)population

firing rate with suitable random statistics that can replace the ideal Poisson noise in our application. Projections from the RN to the SSN were chosen as random and sparse; this resulted in weak, but non-zero shared-input correlations. The remaining correlations are compensated by appropriate training; the Hebbian learning rule (Equation 1) changes the weights and biases in the network such that they cancel the input correlations induced by the RN activity (Bytschok et al., 2017; Dold et al., 2019). Hence, the same plasticity rule simultaneously addresses three issues: the learning procedure itself, the compensation of analog variability in neuronal excitability, and the compensation of cross-correlations in the input coming from the background network. This allowed the hardware-emulated RN to replace the Poisson noise required by the theoretical model.

With these noise-generating mechanisms, the activation function of the neurons, defined by the firing rate as a function of the bias weight  $w_b$ , took on an approximately logistic shape, as required by the sampling model (Figure 3C). Due mainly to the variability of the hardware circuits and the precision of the analog parameters, the exact shape of this activation function varied significantly between neurons (Figures 3D,E). Effectively, this means that initial weights and biases were set randomly, but also that the effective learning rates were different for each neuron. However, as we show below, this did not prevent the training procedure from converging to a good solution. This robustness, with respect to substrate variability, represents an important result of this work. The used neuron parameters are shown in Table 2 and a summary of the used networks is given in Table 3. Our largest experiment, a network of 609 neurons with 208 sampling neurons, one bias neuron and 400 neurons in the RN (Table 3C) used hardware resources on 28 HICANN chips distributed over seven reticles. Each of these functional neurons was realized by combining four of the 512 neuronal compartments (“denmems”) available on each HICANN, in order to reduce variability in their leak potentials and membrane time constants; for details see (Schemmel et al., 2010).

To train the networks on a neuromorphic substrate without embedded plasticity, we used a training concept often referred to as in-the-loop training (Schmuker et al., 2014; Esser et al., 2016; Schmitt et al., 2017). With the setup discussed above, the only parameters changed during training were digital, namely the synaptic weights between sampling neurons and the weights between bias and sampling neurons. This allowed us to work with a fixed set of analog parameters, which significantly amplified the precision and speed of reconfiguration during learning, as compared to having used the analog storage instead. The updates of the digital parameters (synaptic weights) were calculated on the host computer based on the wake-sleep algorithm (Equation 1) but using the spiking activity measured on the hardware. During the iterative procedure, the values of the weights were saved and updated as a double precision floating point variable, followed by (deterministic) discretization in order to comply with the single-synapse weight resolution of 4 bits. The learning parameters are given in Table 4. Clamping (i.e., forcing neurons into state 1 or 0 with strong excitatory or inhibitory input) was done by injecting regular spike trains with a 100 Hz frequency

TABLE 2 | Neuron parameters.

(A) Sampling neuron

Name	Value	Description
$V_{\text{reset}}$	−35 mV	Reset potential
$E_{\text{leak}}$	−20 mV	Resting potential
$V_{\text{thresh}}$	−20 mV	Threshold potential
$E_{\text{inh}}$	−100 mV	Inhibitory reversal potential
$E_{\text{exc}}$	60 mV	Excitatory reversal potential
$\tau_{\text{ref}}$	4 ms	Refractory time
$\tau_{\text{mem}}$	ca. 7 ms	Membrane time constant*
$C_{\text{mem}}$	0.2 nF	Membrane capacity
$\tau_{\text{syn}}^{\text{exc}}$	8 ms	Excitatory synaptic time constant
$\tau_{\text{syn}}^{\text{inh}}$	8 ms	Inhibitory synaptic time constant

(B) Bias neuron

Name	Value	Description
$V_{\text{reset}}$	−30 mV	Reset potential
$E_{\text{leak}}$	60 mV	Resting potential
$V_{\text{thresh}}$	−20 mV	Threshold potential
$E_{\text{inh}}$	−100 mV	Inhibitory reversal potential
$E_{\text{exc}}$	60 mV	Excitatory reversal potential
$\tau_{\text{ref}}$	1.5 ms	Refractory time
$\tau_{\text{mem}}$	ca. 7 ms	Membrane time constant*
$C_{\text{mem}}$	0.2 nF	Membrane capacity
$\tau_{\text{syn}}^{\text{exc}}$	5 ms	Excitatory synaptic time constant
$\tau_{\text{syn}}^{\text{inh}}$	5 ms	Inhibitory synaptic time constant

(C) Neurons of the random network

Name	Value	Description (all analog)
$V_{\text{reset}}$	−60 mV	Reset potential
$E_{\text{leak}}$	−10 mV	Resting potential
$V_{\text{thresh}}$	−20 mV	Threshold potential
$E_{\text{inh}}$	−100 mV	Inhibitory reversal potential
$E_{\text{exc}}$	60 mV	Excitatory reversal potential
$\tau_{\text{ref}}$	4 ms	Refractory time
$\tau_{\text{mem}}$	ca. 7 ms	Membrane time constant*
$C_{\text{mem}}$	0.2 nF	Membrane capacity
$\tau_{\text{syn}}^{\text{exc}}$	8 ms	Excitatory synaptic time constant
$\tau_{\text{syn}}^{\text{inh}}$	8 ms	Inhibitory synaptic time constant

(D) Synapse

Name	Value	Description
$w_{\text{bias}}$	[0,15]	Synaptic bias weight in hardware values (digital)
$w_{\text{network}}$	[0,15]	Synaptic network weight in hardware values (digital)
$d$	On the order of 1 ms (uncalibrated)	Synaptic delay, estimated in Schemmel et al. (2010)

Parameters of the network setup specified in Table 1. The analog parameters are shown as specified in the software setup and not as realized on the hardware. For details on the calibration procedure see e.g., Schmitt et al. (2017). \*The calibration of the membrane time constant was not available at the time of this work, and the corresponding technical parameter was set to the smallest available value instead (fastest possible membrane dynamics for each neuron).

from the host through five synapses simultaneously, excitatory for  $z_k = 1$  and inhibitory for  $z_k = 0$ . These multapses (multiple synapses connecting two neurons) were needed to

**TABLE 3** | Network parameters.

<b>(A) Probability distribution with poisson noise</b>		
Name	Value	Description
$N_s$	5	Number of sampling neurons
$N_b$	1	Number of bias neurons
$N_r$	0	Number of random neurons
$K_{RN}$	–	Within-population in-degree of neurons in the random network
$K_{noise}$	–	In-degree of sampling neurons from the random network
$w_{RN}$	–	Synaptic weights in the random network in hardware units
$\nu_{Poisson}^{e/i}$	300Hz	Poisson frequency to sampling neurons per synapse type

**(B) Probability distribution with random network**

Name	Value	Description
$N_s$	5	Number of sampling neurons
$N_b$	1	Number of bias neurons
$N_r$	200	Number of random neurons
$K_{RN}$	20	Within-population in-degree of neurons in the random network
$K_{noise}$	15	In-degree of sampling neurons from the random network
$w_{RN}$	10	Synaptic weights in the random network in hardware units
$\nu_{Poisson}^{e/i}$	–	Poisson frequency to sampling neurons per synapse type

**(C) High-dimensional dataset**

Name	Value	Description
$N_s$	{207, 208}	Number of sampling neurons, { rFMNIST, rMNIST }
$N_b$	1	Number of bias neurons
$N_r$	400	Number of random neurons
$K_{RN}$	20	Within-population in-degree of neurons in the random network
$K_{noise}$	15	In-degree of sampling neurons from the random network
$w_{RN}$	10	Synaptic weights in the random network in hardware units
$\nu_{Poisson}^{e/i}$	–	Poisson frequency to sampling neurons per synapse type

Parameters are shown for the three different cases described in the manuscript: **(A)** Target Boltzmann distribution, Poisson noise. **(B)** Target Boltzmann distribution, random network for stochasticity. **(C)** Learning from data, random network for stochasticity. Note that the in-degree, sometimes also referred to as a fan-in factor, represents a neuron's number of pre-synaptic partners coming from some specific population.

exceed the upper limit of single synaptic weights and thus ensure proper clamping.

### 3. RESULTS

#### 3.1. Learning to Approximate a Target Distribution

The experiments described in this section serve as a general benchmark for the ability of our hardware-emulated SSNs and the associated training algorithm to approximate fully specified target Boltzmann distributions. The viability of our proposal to simultaneously embed deterministic RNs as sources of pseudo-stochasticity is tested by comparing the sampling accuracy of RN-driven SSNs to the case where noise is injected from the host as perfectly uncorrelated Poisson spike trains.

**TABLE 4** | Parameters for learning.

Experiment	Learning rate	Momentum factor	Minibatch-size	Initial ( $W, b$ )
Target distribution, Poisson	1.0	0.6	–	$\mathcal{U}(-15, 15)$
Target distribution, random network	0.5	0.6	–	$\mathcal{U}(-15, 15)$
rMNIST	0.4	0.6	7/class	Pre-trained
rFMNIST	0.4	0.6	7/class	Pre-trained

We did not carry any systematic hyper-parameter optimization. Note that the used learning parameters in the experiments in section 3.1 are not directly comparable because the different statistics of the background noise (Poisson or random network) correspond to different effective learning rates.

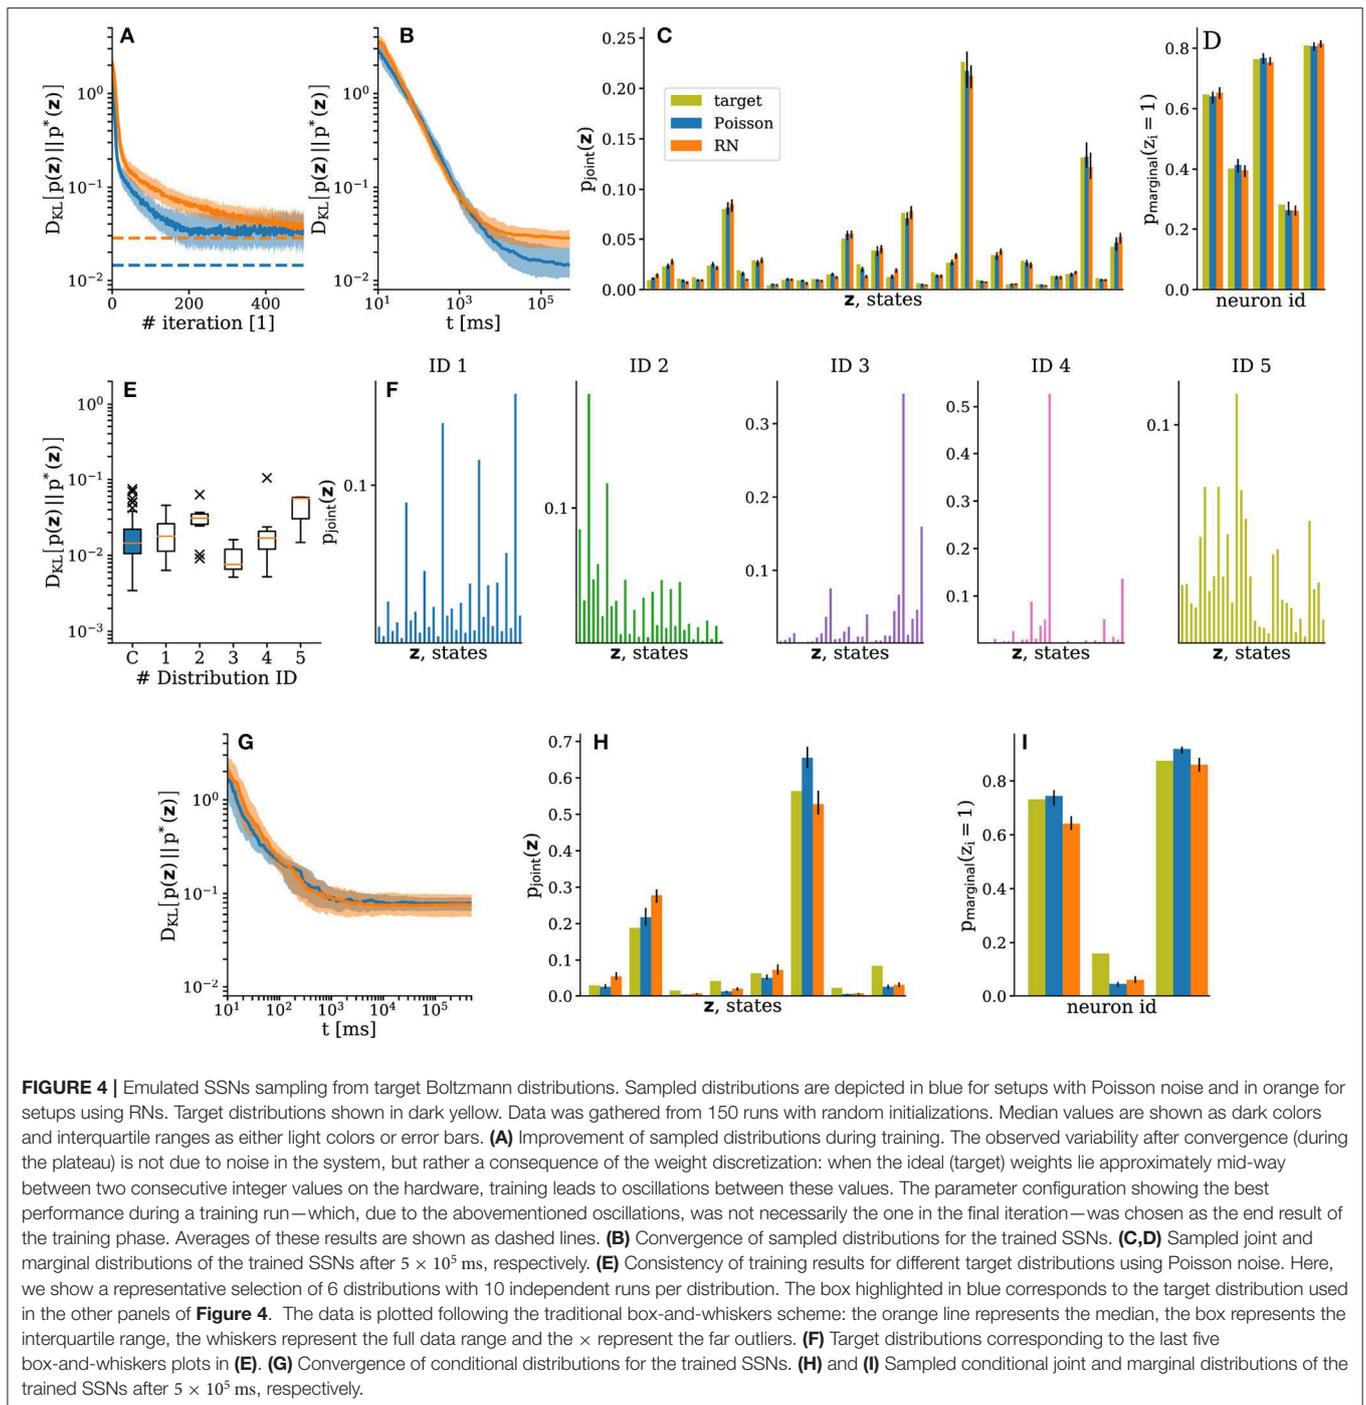
Target distributions  $p^*$  over 5 RVs were chosen by sampling weights and biases from a Beta distribution centered around zero:  $b_i, w_{ji} \sim 2[\text{Beta}(0.5, 0.5) - 0.5]$ . Similar to previous studies (Petrovici et al., 2016; Jordan et al., 2017), by giving preference to larger absolute values of the target distribution's parameters, we thereby increased the probability of instantiating rougher, more interesting energy landscapes. The initial weights and biases of the network were sampled from a uniform distribution over the possible hardware weights. Due to the small size of the state space, the “wake” component of the wake-sleep updates could be calculated analytically as  $\langle z_i z_j \rangle = p^*(z_i = 1, z_j = 1)$  and  $\langle z_i \rangle = p^*(z_i = 1)$  by explicit marginalization of the target distribution over non-relevant RVs.

For training, we used 500 iterations with  $1 \times 10^5$  ms sampling time per iteration. Afterwards, the parameter configuration that produced the lowest  $D_{KL}(p \parallel p^*)$  was tested in a longer ( $5 \times 10^5$  ms) experiment. To study the ability of the trained networks to perform Bayesian inference, we clamped two of the five neurons to fixed values  $(z_1, z_2) = (0, 1)$  and compared the sampled conditional distribution to the target conditional distribution. Results for one of these target distributions are shown in **Figure 4**.

On average, with Poisson noise, the training showed fast convergence during the first 20 iterations, followed by fine-tuning and full convergence within 200 iterations. As expected, the convergence of the setups using RNs was significantly slower due to the need to overcome the additional background correlations, but they were still able to achieve similar performance (**Figure 4A**).

In both setups, during the test run, the trained SSNs converged to the target distribution following an almost identical power law, which indicates similar mixing properties (**Figure 4B**). For longer sampling durations ( $\gg 10 \times 10^3$  ms), the systematic deviations from the target distributions become visible and the  $D_{KL}(p \parallel p^*)$  reaches the same plateau at approximately  $D_{KL}(p \parallel p^*) \approx 2 \times 10^{-2}$  as observed during training. **Figures 4C,D**, respectively show the sampled joint and marginal distributions after convergence (**Supplementary Video 1**). These observations remained consistent across a set of 20 different target distributions (see **Figure 4E** for a representative selection).

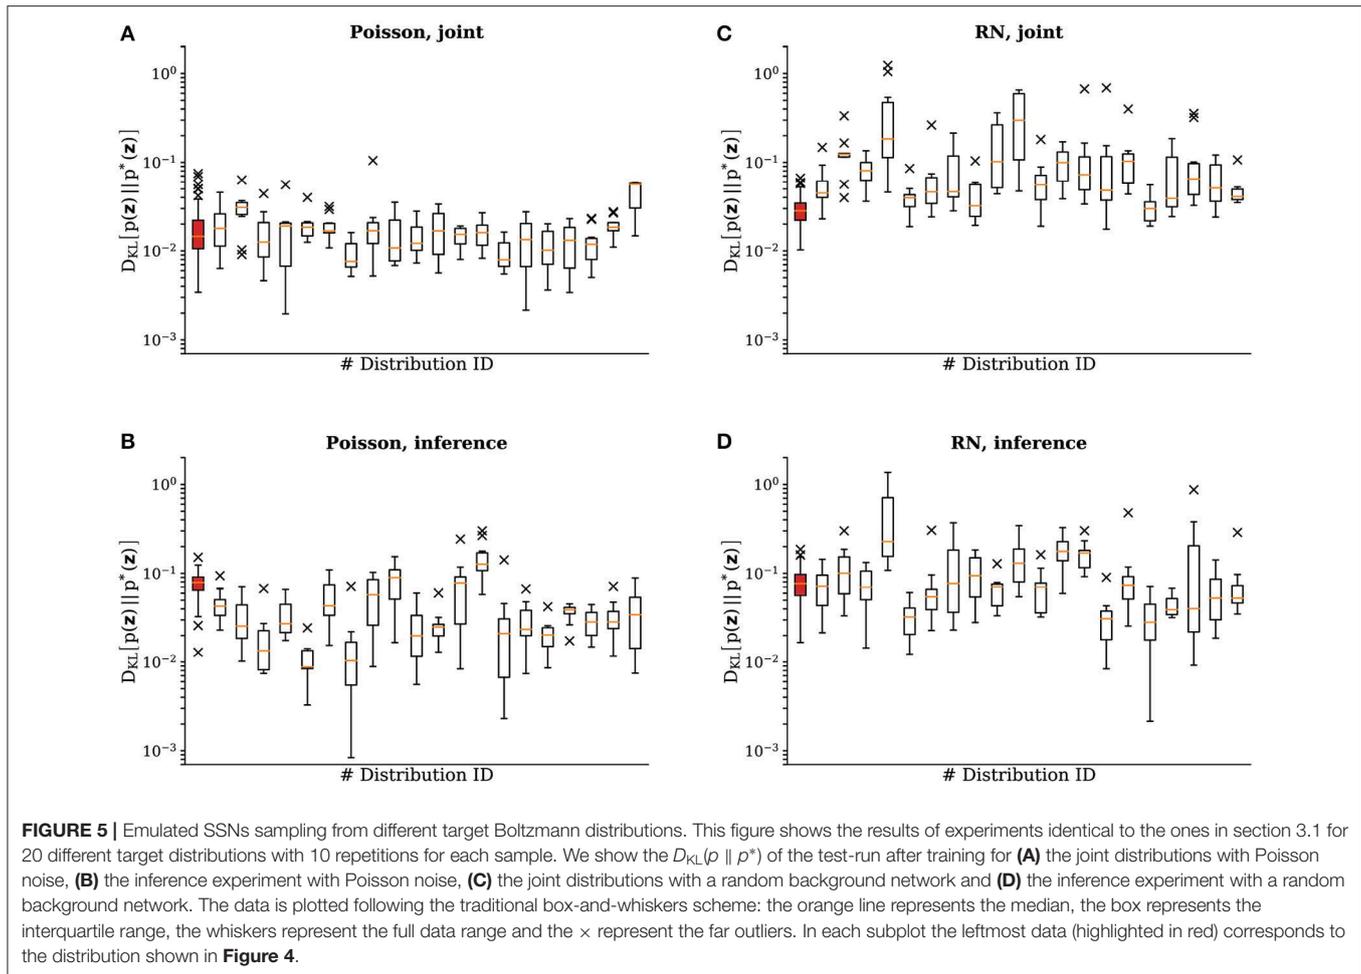
Similar observations were made for the inference experiments. Due to the smaller state space, convergence happened faster



(**Figure 4G**). The corresponding joint and marginal distributions are shown in **Figures 4H,I**, respectively. The lower accuracy of these distributions is mainly due to the asymmetry of the effective synaptic weights caused by the variability of the substrate, toward which the learning algorithm is agnostic. The training took  $5 \times 10^2$  s wall-clock time, including the pure experiment runtime, the initialization of the hardware and the calculation of the updates on the host computer (total turn-over time of the

training). This corresponds to a speed-up factor of 100 compared to the equivalent  $5 \times 10^4$  s of biological real time. While the nominal  $10^4$  speed-up remained intact for the emulation of network dynamics, the total speed-up factor was reduced due to the overhead imposed by network (re)configuration and I/O between the host and the neuromorphic substrate.

We carried out the same experiments as described previously with 20 different samples for the weights and the biases of the



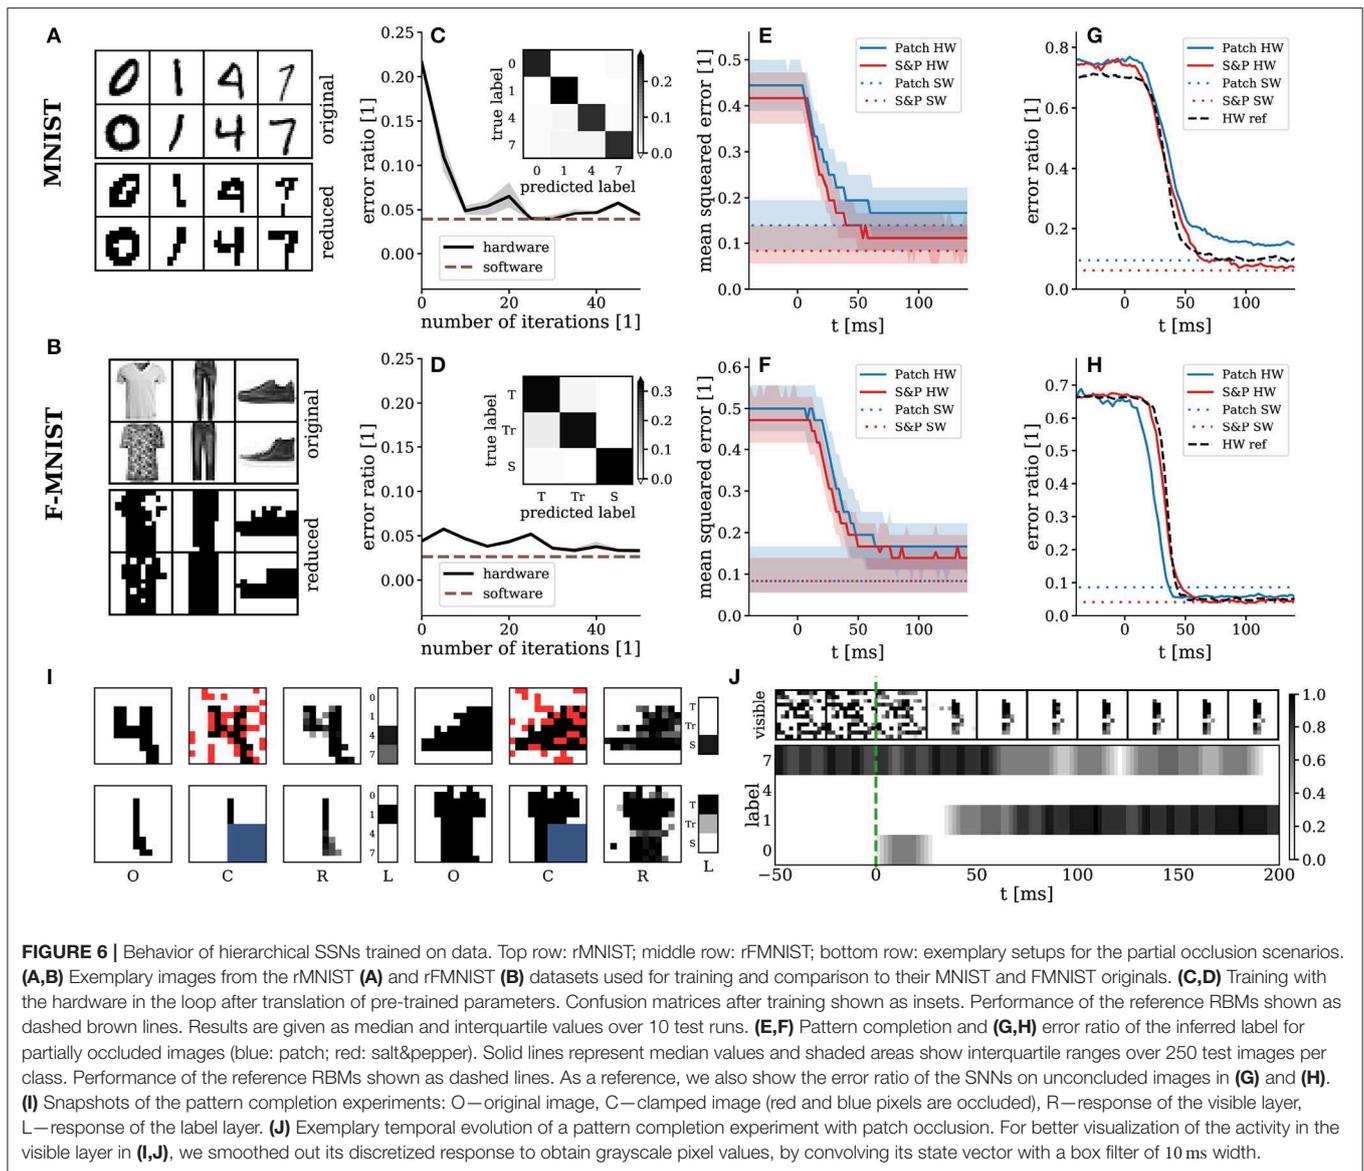
target distribution. In **Figure 5** we show the final DKLs after training to represent a target distribution both with Poisson noise and with the activity of a random network. The experiments were repeated 10 times for each sample. Median learning results remained consistent across target distributions, with the variability reflecting the difficulty of the problem (discrepancies between LIF and Glauber dynamics become more pronounced for larger weights and biases). Variability across trials for the same target distribution is due to the trial-to-trial variability of the analog parameter storage (floating gates), due to the inherent stochasticity in the learning procedure (sampling accuracy in an update step), as well as due to systematic discrepancies between the effective pre-post and post-pre interaction strengths between sampling units, which are themselves a consequence of the aforementioned floating gate variability.

### 3.2. Learning From Data

In order to obtain models of labeled data, we trained hierarchical SSNs analogously to restricted Boltzmann machines (RBMs). Here, we used two different datasets: a reduced version of the MNIST (LeCun et al., 1998) and the fashion MNIST (Xiao et al., 2017) datasets, which we abbreviate as rMNIST and rFMNIST

in the following. The images were first reduced with nearest-neighbor resampling [`misc.imresize` function in the SciPy library (Jones et al., 2001)] and then binarized around the median gray value over each image. We used all images from the original datasets (~6,000 per class) from four classes (0, 1, 4, 7) for rMNIST and three classes (T-shirts, Trousers, Sneakers) for rFMNIST (**Figures 6A,B**). The emulated SSNs consisted of three layers, with 144 visible, 60 hidden, and either four label units for rMNIST or three for rFMNIST.

Pre-training was done on simulated classical RBMs using the CAST algorithm (Salakhutdinov, 2010). The pre-training provided a starting point for training on the hardware in order to accelerate the convergence of the in-the-loop training procedure. We use the performance of these RBMs in software simulations using Gibbs sampling as a reference for the results obtained with the hardware-emulated SSNs. After pre-training, we mapped these RBMs to approximately equivalent SSNs on the hardware, using an empirical translation factor based on an average activation function (**Figure 3C**) to calculate the initial hardware synaptic weights from weights and biases of the RBMs. Especially for rMNIST, this resulted in a significant deterioration of the classification performance (**Figure 6C**). After mapping, we continued training using the wake-sleep algorithm, with the

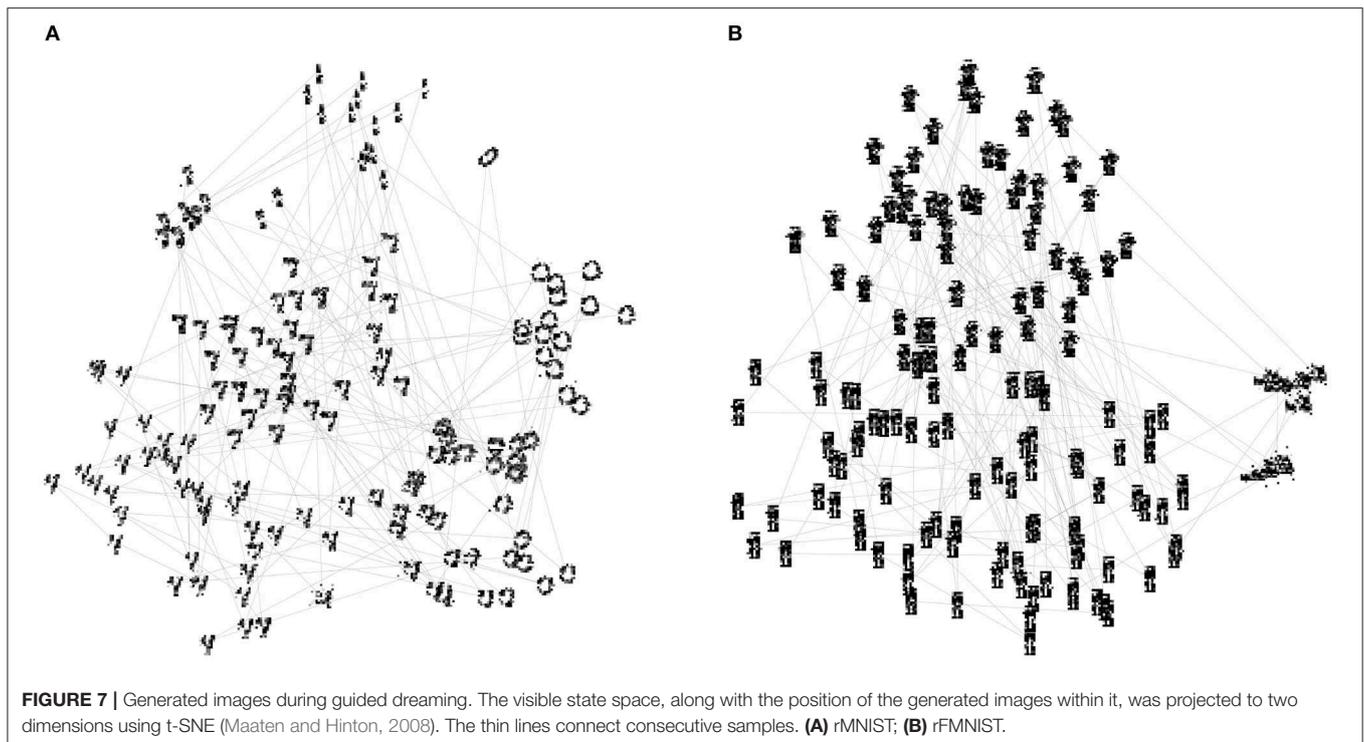


hardware in the loop. While in the previous task it was possible to calculate the data term explicitly, it now had to be sampled as well. In order to ensure proper clamping, the synapses from the hidden to the label layer and from the hidden layer to the visible layer were turned off during the wake phase.

The SSNs were tested for both their discriminative and their generative properties. For classification, the visible layer was clamped to images from the test set (black pixels correspond to  $z_k = 1$  and white pixels to  $z_k = 0$ ). Each image was presented for 500 biological milliseconds, which corresponds to 50  $\mu$ s wall-clock time. The neuron in the label layer with the highest firing rate was interpreted as the label predicted by the model. The spiking activity of the neurons was read out directly from the hardware, without additional off-chip post-processing. For both datasets, training was able to restore the performance lost in the translation of the abstract RBM to

the hardware-emulated SSN. The emulated SSNs achieved error rates of  $4.45^{+0.12}_{-0.36}$ % on rMNIST and  $3.32^{+0.27}_{-0.04}$ % on rFMNIST. These values are close to the ones obtained by the reference RBMs:  $3.89^{+0.10}_{-0.02}$ % on rMNIST and  $2.645^{+0.002}_{-0.010}$ % on rFMNIST (**Figures 6C,D**, confusion matrices shown as insets).

The gross wall-clock time needed to classify the 4125 images in the rMNIST test set was 10 s (2.4 ms per image,  $210\times$  speed-up). For the 3,000 images in the rFMNIST test set, the emulation ran for 9.4 s (3.1 ms per image;  $160\times$  speed-up). This subsumes the runtime of the BrainScaleS software stack, hardware configuration and the network emulation. The runtime of the software-stack includes the translation from a PyNN-based network description to a corresponding hardware configuration. As before, the difference between the nominal acceleration factor and the effective speed-up stems from the I/O and initialization overhead of the hardware system.



To test the generative properties of our emulated SSNs, we set up two scenarios requiring them to perform pattern completion (**Supplementary Video 2**). For each class, 250 incomplete images were presented as inputs to the visible layer. For each image, 25 % of visible neurons received no input, with the occlusion following two different schemes: salt&pepper (upper row in **Figure 6I**) and patch (lower row in **Figure 6I**). Each image was presented for 500 ms. In order to remove any initialization bias resulting from preceding images, random input was applied to the visible layer between consecutive images.

Reconstruction accuracy was measured using the mean squared error (MSE) between the reconstructed and original occluded pixels. For binary images, as in our case, the MSE reflects the average ratio of mis-reconstructed to total reconstructed pixels. Simultaneously, we also recorded the classification accuracy on the partially occluded images. After stimulus onset, the MSE converged from chance level ( $\approx 50\%$ ) to its minimum ( $\approx 10\%$ ) within 50 ms (**Figures 6E,F**). Given an average refractory period of  $\approx 10$  ms (**Figure 3C**), this suggests that the network was able to react to the input with no more than 5 spikes per neuron. For all studied scenarios, the reconstruction performance of the emulated SSNs closely matched the one achieved by the reference RBMs. Examples of image reconstruction are shown in **Figures 6I,J** for both datasets and occlusion scenarios. The classification performance deteriorated only slightly compared to non-occluded images and also remained close to the performance of the reference RBMs (**Figures 6G,H**). The temporal evolution of the classification error closely followed that of the MSE.

As a further test of the generative abilities of our hardware-emulated SSNs, we recorded the images produced by the visible

layer during guided dreaming. In this task, the visible and hidden layers of the SSN evolved freely without external input, while the label layer was periodically clamped with external input such that exactly one of the label neurons was active at any time (enforced one-hot coding). In a perfect model, this would cause the visible layer to sample only from configurations compatible with the hidden layer, i.e., from images corresponding to that particular class. Between the clamping of consecutive labels, we injected 100 ms random input to visible layer to facilitate the changing of the image. The SSNs were able to generate varied and recognizable pictures, within the limits imposed by the low resolution of the visible layer (**Figure 7**). For rMNIST, all used classes appeared in correct correspondence to the clamped label. For rFMNIST, images from the class “Sneakers” were not always triggered by the corresponding guidance from the label layer, suggesting that the learned modes in the energy landscape are too deep, and sneakers too dissimilar to T-shirts and Trousers, to allow good mixing during guided dreaming.

## 4. DISCUSSION

This article presents the first scalable demonstration of sampling-based probabilistic inference with spiking networks on a highly accelerated analog neuromorphic substrate. We trained fully connected spiking networks to sample from target distributions and hierarchical spiking networks as discriminative and generative models of higher-dimensional input data. Despite the inherent variability of the analog substrate, we were able to achieve performance levels comparable to those of software simulations in several benchmark tasks, while maintaining a

significant overall acceleration factor compared to systems that operate in biological real time. Importantly, by co-embedding the generation of stochasticity within the same substrate, we demonstrated the viability of a fully embedded neural sampling model with significantly reduced demands on off-substrate I/O bandwidth. Having a fully embedded implementation allows the runtime of the experiments to scale as  $\mathcal{O}(1)$  with the size of the emulated network; this is inherent to the nature of physical emulation, for which wall-clock runtime only depends on the emulated time in the biological reference frame. In the following sections, we address the limitations of our study, point out links to related work and discuss its implications within the greater context of computational neuroscience and bio-inspired AI.

#### 4.1. Limitations and Constraints

The most notable limitation imposed by the current commissioning state of the BrainScaleS system was on the size of the emulated SSNs. At the time of writing, due to limited software flexibility, system assembly and substrate yield, the usable hardware real-estate was reduced to a patchy and non-contiguous area, thereby strongly limiting the maximum connectivity between different locations within this area. In order to limit synapse loss to small values (below 2%), we restricted ourselves to using a small but contiguous functioning area of the wafer, which in turn limited the maximum size of our SSNs and noise-generating RNs. Ongoing improvements in post-production and assembly, as well as in the mapping and routing software, are expected to enhance on-wafer connectivity and thereby automatically increase the size of emulable networks, as the architecture of our SSNs scales naturally to such an increase in hardware resources.

To a lesser extent, the sampling accuracy was also affected by the limited precision of hardware parameter control. The writing of analog parameters exhibits significant trial-to-trial variability; in any given trial, this leads to a heterogeneous substrate, which is known to reduce the sampling accuracy (Probst et al., 2015). Most of this variability is compensated during learning, but the 4 bit resolution of the synaptic weights and the imperfect symmetry in the effective weight matrix due to analog variability of the synaptic circuits ultimately limit the ability of the SSN to approximate target distributions. This leads to the “jumping” behavior of the  $D_{\text{KL}}(p \parallel p^*)$  in the final stages of learning (Figure 4A). In smaller networks, synaptic weight resolution is a critical performance modifier (Petrovici et al., 2017b). However, the penalty imposed by a limited synaptic weight resolution is known to decrease for larger deep networks with more and larger hidden layers, both spiking and non-spiking (Courbariaux et al., 2015; Petrovici et al., 2017a). Furthermore, the successor system (BrainScaleS-2, Aamir et al., 2016) is designed with a 6-bit weight resolution.

In the setup we used shared bias neurons for several neurons in the sampling network. This helped us save hardware resources, thus allowing the emulation of larger functional networks. Such bias neuron sharing is expected to introduce some small amount of temporal correlations between the sampling neurons. However, this effect was too small to observe in our experiments for several reasons. First, the high firing

rate of the bias neurons helped smooth out the bias voltage induced into the sampling neurons. Second, the different delays and spike timing jitter on the hardware reduces such cross-correlations. Third, other dominant limitations overshadow the effect of shared bias neurons. In any case, the used training procedure inherently compensates for excess cross-correlations, thus effectively removing any distortions to the target distribution that this effect might introduce (Bytschok et al., 2017; Dold et al., 2019).

In the current setup, our SSNs displayed limited mixing abilities. During guided dreaming, images from one of the learned classes were more difficult to generate (Figure 7). Restricted mixing due to deep modes in the energy landscape carved out by contrastive learning is a well-known problem for classical Boltzmann machines, which is usually alleviated by computationally costly annealing techniques (Desjardins et al., 2010; Salakhutdinov, 2010; Bengio et al., 2013). However, the fully-commissioned BrainScaleS system will feature embedded short-term synaptic plasticity (Schemmel et al., 2010), which has been shown to promote mixing in spiking networks (Leng et al., 2018) while operating purely locally, at the level of individual synapses.

Currently, the execution speed of emulation runs is dominated by the I/O overhead, which in turn is mostly spent on setting up the experiment. This leads to the classification (section 3.2) of one image taking 2.4 to 3.9 ms, whereas the pure network runtime is merely 50  $\mu\text{s}$ . A streamlining of the software layer that performs this setup is expected to significantly reduce this discrepancy.

The synaptic learning rule was local and Hebbian, but updates were calculated on a host computer using an iterative in-the-loop training procedure, which required repeated stopping, evaluation and restart of the emulation, thereby reducing the nominal acceleration factor of  $10^4$  by two orders of magnitude. By utilizing on-chip plasticity, as available, for example, on the BrainScaleS-2 successor system (Friedmann et al., 2017; Wunderlich et al., 2019), this laborious procedure becomes obsolete and the accelerated nature of the substrate can be exploited to its fullest extent.

#### 4.2. Relation to Other Work

This study builds upon a series of theoretical and experimental studies of sampling-based probabilistic inference using the dynamics of biological neurons. The inclusion of refractory times was first considered in Buesing et al. (2011). An extension to networks of leaky integrate-and-fire neurons and a theoretical framework for their dynamics and statistics followed in Petrovici et al. (2013) and Petrovici et al. (2016). The compensation of shared-input correlations through inhibitory feedback and learning was discussed in Bytschok et al. (2017), Jordan et al. (2017), and Dold et al. (2019), inspired by the early study of asynchronous irregular firing in Brunel (2000) and by preceding correlation studies in theoretical (Tetzlaff et al., 2012) and experimental (Pfeil et al., 2016) work.

Previous small-scale studies of sampling on accelerated mixed-signal neuromorphic hardware include (Petrovici et al., 2015, 2017a,b). An implementation of sampling with spiking neurons and its application to the MNIST dataset was shown in

Pedroni et al. (2016) using the fully digital, real-time TrueNorth neuromorphic chip (Merolla et al., 2014).

We stress two important differences between (Pedroni et al., 2016) and this work. First, the nature of the neuromorphic substrate: the TrueNorth system is fully digital and calculates neuronal state updates numerically, in contrast to the physical-model paradigm instantiated by BrainScaleS. In this sense, TrueNorth emulations are significantly closer to classical computer simulations on parallel machines: updates of dynamical variables are precise and robustness to variability is not an issue; however TrueNorth typically runs in biological real time (Merolla et al., 2014; Akopyan et al., 2015), which is 10,000 times slower than BrainScaleS. Second, the nature of neuron dynamics: the neuron model used in (Pedroni et al., 2016) is an intrinsically stochastic unit that sums its weighted inputs, thus remaining very close to classical Gibbs sampling and Boltzmann machines, while our approach considers multiple additional aspects of its biological archetype (exponential synaptic kernels, leaky membranes, deterministic firing, stochasticity through synaptic background, shared-input correlations etc.). Moreover, our approach uses fewer hardware neuron units to represent a sampling unit, enabling a more parsimonious utilization of the neuromorphic substrate.

### 4.3. Conclusion

In this work we showed how sampling-based Bayesian inference using hierarchical spiking networks can be robustly implemented on a physical model system despite inherent variability and imperfections. Underlying neuron and synapse dynamics are deterministic and close to their biological archetypes, but with much shorter time constants, hence the intrinsic acceleration factor of  $10^4$  with respect to biology. The entire architecture—sampling network plus background random network—was fully deterministic and entirely contained on the neuromorphic substrate, with external communication used only to represent input patterns and labels. Considering the deterministic nature of neurons *in vitro* (Mainen and Sejnowski, 1995; Reinagel and Reid, 2002; Troups et al., 2012), such an architecture also represents a plausible model for neural sampling in cortex (Jordan et al., 2017; Dold et al., 2019).

We demonstrated sampling from arbitrary Boltzmann distributions over binary random variables, as well as generative and discriminative properties of networks trained with visual data. The framework can be extended to sampling from arbitrary probability distributions over binary random variables, as it was shown in software simulations (Probst et al., 2015). For such networks, the two abovementioned computational tasks (pattern completion and classification) happen simultaneously, as they both require the calculation of conditional distributions, which is carried out implicitly by the network dynamics. Both during learning and for the subsequent inference tasks, the setup benefitted significantly from the fast-intrinsic dynamics of the substrate, achieving a net speedup of 100–210 compared to biology.

We view these results as a contribution to the nascent but expanding field of applications for biologically inspired

physical-model systems. They demonstrate the feasibility of such devices to solve problems in machine learning, as well as studying biological phenomena. Importantly, they explicitly address the search for robust computational models that are able to harness the strengths of these systems, most importantly their speed and energy efficiency. The proposed architecture scales naturally to substrates with more neuronal real-estate and can be used for a wide array of tasks that can be mapped to a Bayesian formulation, such as constraint satisfaction problems (Jonke et al., 2016; Fonseca Guerra and Furber, 2017), prediction of temporal sequences (Sutskever and Hinton, 2007), movement planning (Taylor and Hinton, 2009; Alemi et al., 2015), simulation of solid-state systems (Edwards and Anderson, 1975), and quantum many-body problems (Carleo and Troyer, 2017; Czischek et al., 2018).

## DATA AVAILABILITY STATEMENT

The datasets generated for this study are available on request to the corresponding author.

## AUTHOR CONTRIBUTIONS

AFK, AB, DD, LL, SS, PM, and MP designed the study. AFK conducted the experiments and the evaluations. NG contributed to the evaluations and provided software support for the evaluation. AFK wrote the initial manuscript. EM, CM, JK, SS, KH, and OB supported the experiment realization. EM coordinated the software development for the neuromorphic systems. AK, CK, and MK contributed with the characterization, calibration testing, and debugging of the system. AG, DH, and MG were responsible for system assembly. AG did the digital front- and back-end implementation. VK provided the FPGA firmware and supported system commissioning. JS is the architect and lead designer of the neuromorphic platform. MP, KM, JS, SS, and EM provided the conceptual and scientific advice. All authors contributed to the final manuscript.

## FUNDING

The work leading to these results received funding from the European Union Seventh Framework Programme (FP7) under grant agreement No. #604102, the EU's Horizon 2020 research and innovation programme under grant agreements No. #720270 and #785907 (Human Brain Project, HBP), the EU's research project BrainScaleS #269921 and the Heidelberg Graduate School of Fundamental Physics.

## ACKNOWLEDGMENTS

We thank Johannes Bill for many fruitful discussions. We acknowledge financial support by Deutsche Forschungsgemeinschaft within the funding programme

Open Access Publishing, by the Baden-Württemberg Ministry of Science, Research and the Arts and by Ruprecht-Karls-Universität Heidelberg. We owe particular gratitude to the sustained support of our research by the Manfred Stärk Foundation.

## REFERENCES

Aamir, S. A., Müller, P., Hartel, A., Schemmel, J., and Meier, K. (2016). "A highly tunable 65-nm cmos lif neuron for a large scale neuromorphic system," in *42nd European Solid-State Circuits Conference, ESSCIRC Conference 2016* (Lausanne: IEEE), 71–74.

Aitchison, L., and Lengyel, M. (2016). The hamiltonian brain: efficient probabilistic inference with excitatory-inhibitory neural circuit dynamics. *PLoS Comput. Biol.* 12:e1005186. doi: 10.1371/journal.pcbi.1005186

Akopyan, F., Sawada, J., Cassidy, A., Alvarez-Icaza, R., Arthur, J., Merolla, P., et al. (2015). Truenorth: design and tool flow of a 65 mw 1 million neuron programmable neurosynaptic chip. *IEEE Trans. Comput. Aided Des. Integr. Circuits Syst.* 34, 1537–1557. doi: 10.1109/TCAD.2015.2474396

Alemi, O., Li, W., and Pasquier, P. (2015). "Affect-expressive movement generation with factored conditional restricted Boltzmann machines," in *2015 International Conference on Affective Computing and Intelligent Interaction (ACII)* (Xian: IEEE), 442–448.

Bengio, Y., Mesnil, G., Dauphin, Y., and Rifai, S. (2013). "Better mixing via deep representations," in *International Conference on Machine Learning* (Atlanta, GA), 552–560.

Berkes, P., Orbán, G., Lengyel, M., and Fiser, J. (2011). Spontaneous cortical activity reveals hallmarks of an optimal internal model of the environment. *Science* 331, 83–87. doi: 10.1126/science.1195870

Brette, R., and Gerstner, W. (2005). Adaptive exponential integrate-and-fire model as an effective description of neuronal activity. *J. Neurophysiol.* 94, 3637–3642. doi: 10.1152/jn.00686.2005

Brunel, N. (2000). Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. *J. Comput. Neurosci.* 8, 183–208. doi: 10.1023/A:1008925309027

Buesing, L., Bill, J., Nessler, B., and Maass, W. (2011). Neural dynamics as sampling: a model for stochastic computation in recurrent networks of spiking neurons. *PLoS Comput. Biol.* 7:e1002211. doi: 10.1371/journal.pcbi.1002211

Bytschok, I., Dold, D., Schemmel, J., Meier, K., and Petrovici, M. A. (2017). Spike-based probabilistic inference with correlated noise. *BMC Neurosci.* 18:200. doi: 10.1186/s12868-017-0372-1

Carleo, G., and Troyer, M. (2017). Solving the quantum many-body problem with artificial neural networks. *Science* 355, 602–606. doi: 10.1126/science.aag2302

Chang, Y. F., Fowler, B., Chen, Y. C., Zhou, F., Pan, C. H., Chang, T. C., et al. (2016). Demonstration of synaptic behaviors and resistive switching characterizations by proton exchange reactions in silicon oxide. *Sci. Rep.* 6:21268. doi: 10.1038/srep21268

Courbariaux, M., Bengio, Y., and David, J.-P. (2015). "Binaryconnect: training deep neural networks with binary weights during propagations," in *Advances in Neural Information Processing Systems*, eds C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett (Montreal, QC: Neural Information Processing Systems Foundation, Inc.), 3123–3131.

Czischek, S., Gärtner, M., and Gasenzer, T. (2018). Quenches near ising quantum criticality as a challenge for artificial neural networks. *Phys. Rev. B* 98:024311. doi: 10.1103/PhysRevB.98.024311

Davison, A. P., Brüderle, D., Eppler, J. M., Kremkow, J., Müller, E., Pecevski, D., et al. (2009). Pynn: a common interface for neuronal network simulators. *Front. Neuroinformatics* 2:11. doi: 10.3389/neuro.11.011.2008

Desjardins, G., Courville, A., Bengio, Y., Vincent, P., and Delalleau, O. (2010). "Parallel tempering for training of restricted boltzmann machines," in *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics* (Cambridge, MA: MIT Press), 145–152.

Destexhe, A., Rudolph, M., and Paré, D. (2003). The high-conductance state of neocortical neurons *in vivo*. *Nat. Rev. Neurosci.* 4:739. doi: 10.1038/nrn1198

## SUPPLEMENTARY MATERIAL

The Supplementary Material for this article can be found online at: <https://www.frontiersin.org/articles/10.3389/fnins.2019.01201/full#supplementary-material>

Dold, D., Bytschok, I., Kungl, A. F., Baumbach, A., Breitwieser, O., Senn, W., et al. (2019). Stochasticity from function—why the bayesian brain may need no noise. *Neural Netw.* 119, 200–213. doi: 10.1016/j.neunet.2019.08.002

Edwards, S. F., and Anderson, P. W. (1975). Theory of spin glasses. *J. Phys. F Metal Phys.* 5:965.

Esser, S. K., Merolla, P. A., Arthur, J. V., Cassidy, A. S., Appuswamy, R., Andreopoulos, A., et al. (2016). From the cover: convolutional networks for fast, energy-efficient neuromorphic computing. *Proc. Natl. Acad. Sci. U.S.A.* 113:11441. doi: 10.1073/pnas.1604850113

Fonseca Guerra, G. A., and Furber, S. B. (2017). Using stochastic spiking neural networks on spinnaker to solve constraint satisfaction problems. *Front. Neurosci.* 11:714. doi: 10.3389/fnins.2017.00714

Friedmann, S., Schemmel, J., Grünbl, A., Hartel, A., Hock, M., and Meier, K. (2017). Demonstrating hybrid learning in a flexible neuromorphic hardware system. *IEEE Trans. Biomed. Circuits Syst.* 11, 128–142. doi: 10.1109/TBCAS.2016.2579164

Furber, S. (2016). Large-scale neuromorphic computing systems. *J. Neural Eng.* 13:051001. doi: 10.1088/1741-2560/13/5/051001

Haefner, R. M., Berkes, P., and Fiser, J. (2016). Perceptual decision-making as probabilistic inference by neural sampling. *Neuron* 90, 649–660. doi: 10.1016/j.neuron.2016.03.020

Hennequin, G., Aitchison, L., and Lengyel, M. (2014). Fast sampling for bayesian inference in neural circuits. *arXiv 1404.3521*.

Hinton, G. E. (2012). "A practical guide to training restricted boltzmann machines," in *Neural Networks: Tricks of the Trade*, eds G. Montavon, G. B. Orr, and K.-R. Müller (Berlin; Heidelberg: Springer), 599–619. doi: 10.1007/978-3-642-35289-8

Hinton, G. E., Dayan, P., Frey, B. J., and Neal, R. M. (1995). The "wake-sleep" algorithm for unsupervised neural networks. *Science* 268, 1158–1161.

Hinton, G. E., Sejnowski, T. J., and Ackley, D. H. (1984). *Boltzmann Machines: Constraint Satisfaction Networks That Learn*. Pittsburgh, PA: Carnegie-Mellon University, Department of Computer Science.

Indiveri, G., Chicca, E., and Douglas, R. J. (2006). A VLSI array of low-power spiking neurons and bistable synapses with spike-timing dependent plasticity. *IEEE Trans. Neural Netw.* 17, 211–221. doi: 10.1109/TNN.2005.860850

Indiveri, G., Linares-Barranco, B., Hamilton, T. J., Van Schaik, A., Etienne-Cummings, R., Delbruck, T., et al. (2011). Neuromorphic silicon neuron circuits. *Front. Neurosci.* 5:73. doi: 10.3389/fnins.2011.00073

Jo, S. H., Chang, T., Ebong, I., Bhadviya, B. B., Mazumder, P., and Lu, W. (2010). Nanoscale memristor device as synapse in neuromorphic systems. *Nano Lett.* 10, 1297–1301. doi: 10.1021/nl904092h

Jones, E., Oliphant, T., and Peterson, P. (2001). *SciPy: Open Source Scientific Tools for Python*.

Jonke, Z., Habenschuss, S., and Maass, W. (2016). Solving constraint satisfaction problems with networks of spiking neurons. *Front. Neurosci.* 10:118. doi: 10.3389/fnins.2016.00118

Jordan, J., Petrovici, M. A., Breitwieser, O., Schemmel, J., Meier, K., Diesmann, M., et al. (2017). Stochastic neural computation without noise. *arXiv 1710.04931*.

Kullback, S., and Leibler, R. A. (1951). On information and sufficiency. *Ann. Math. Stat.* 22, 79–86.

Kutschireiter, A., Surace, S. C., Sprekeler, H., and Pfister, J.-P. (2017). Nonlinear bayesian filtering and learning: a neuronal dynamics for perception. *Sci. Rep.* 7:8722. doi: 10.1038/s41598-017-17246-9

Lande, T. S., Ranjbar, H., Ismail, M., and Berg, Y. (1996). "An analog floating-gate memory in a standard digital technology," in *Proceedings of Fifth International Conference on Microelectronics for Neural Networks* (Lausanne: IEEE), 271–276.

LeCun, Y., Bottou, L., Bengio, Y., and Haffner, P. (1998). Gradient-based learning applied to document recognition. *Proc. IEEE* 86, 2278–2324.

- Leng, L., Martel, R., Breitwieser, O., Bytschok, I., Senn, W., Schemmel, J., et al. (2018). Spiking neurons with short-term synaptic plasticity form superior generative networks. *Sci. Rep.* 8:10651. doi: 10.1038/s41598-018-28999-2
- Loock, J.-P. (2006). *Evaluierung eines floating gate analogspeichers für neuronale netze in single-poly umc 180nm cmos-prozess* (Diploma thesis), University of Heidelberg, Heidelberg, Germany.
- Maaten, L. V. D., and Hinton, G. (2008). Visualizing data using t-sne. *J. Mach. Learn. Res.* 9, 2579–2605. Available online at: <http://www.jmlr.org/papers/v9/vandermaaten08a.html>
- Mainen, Z. F., and Sejnowski, T. J. (1995). Reliability of spike timing in neocortical neurons. *Science* 268, 1503–1506.
- Mead, C. (1990). Neuromorphic electronic systems. *Proc. IEEE* 78, 1629–1636.
- Merolla, P. A., Arthur, J. V., Alvarez-Icaza, R., Cassidy, A. S., Sawada, J., Akopyan, F., et al. (2014). A million spiking-neuron integrated circuit with a scalable communication network and interface. *Science* 345, 668–673. doi: 10.1126/science.1254642
- Millner, S. (2012). *Development of a multi-compartment neuron model emulation* (PhD thesis) (Heidelberg: Heidelberg University, Faculty of Physics and Astronomy). doi: 10.11588/heidok.00013979
- Millner, S., Grübl, A., Meier, K., Schemmel, J., and Schwartz, M.-O. (2010). A VLSI implementation of the adaptive exponential integrate-and-fire neuron model. *Adv. Neural Inform. Process. Syst.* (Vancouver, QC) 23, 1642–1650. Available online at: <https://papers.nips.cc/paper/3995-a-vlsi-implementation-of-the-adaptive-exponential-integrate-and-fire-neuron-model>
- Orbán, G., Berkes, P., Fiser, J., and Lengyel, M. (2016). Neural variability and sampling-based probabilistic representations in the visual cortex. *Neuron* 92, 530–543. doi: 10.1016/j.neuron.2016.09.038
- Pedroni, B. U., Das, S., Arthur, J. V., Merolla, P. A., Jackson, B. L., Modha, D. S., et al. (2016). Mapping generative models onto a network of digital spiking neurons. *IEEE Trans. Biomed. Circuits Syst.* 10, 837–854. doi: 10.1109/TBCAS.2016.2539352
- Petrovici, M. A. (2016). *Form Versus Function: Theory and Models for Neuronal Substrates*. Springer International Publishing Switzerland. Available online at: <https://www.springer.com/gp/book/9783319395517>
- Petrovici, M. A., Bill, J., Bytschok, I., Schemmel, J., and Meier, K. (2013). Stochastic inference with deterministic spiking neurons. *arXiv 1311.3211*.
- Petrovici, M. A., Bill, J., Bytschok, I., Schemmel, J., and Meier, K. (2016). Stochastic inference with spiking neurons in the high-conductance state. *Phys. Rev. E* 94:042312. doi: 10.1103/PhysRevE.94.042312
- Petrovici, M. A., Schmitt, S., Klähn, J., Stöckel, D., Schroeder, A., Bellec, G., et al. (2017a). “Pattern representation and recognition with accelerated analog neuromorphic systems,” in *2017 IEEE International Symposium on Circuits and Systems (ISCAS)* (Baltimore, MD: IEEE), 1–4.
- Petrovici, M. A., Schroeder, A., Breitwieser, O., Grübl, A., Schemmel, J., and Meier, K. (2017b). “Robustness from structure: inference with hierarchical spiking networks on analog neuromorphic hardware,” in *2017 International Joint Conference on Neural Networks (IJCNN)* (Anchorage, AL: IEEE), 2209–2216.
- Petrovici, M. A., Stöckel, D., Bytschok, I., Bill, J., Pfeil, T., Schemmel, J., et al. (2015). “Fast sampling with neuromorphic hardware,” in *Advances in Neural Information Processing Systems (NIPS)*, eds C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, (Montreal, QC: Neural Information Processing Systems Foundation, Inc.) Vol. 28.
- Petrovici, M. A., Vogginger, B., Müller, P., Breitwieser, O., Lundqvist, M., Müller, L., et al. (2014). Characterization and compensation of network-level anomalies in mixed-signal neuromorphic modeling platforms. *PLoS ONE* 9:e108590. doi: 10.1371/journal.pone.0108590
- Pfeil, T., Grübl, A., Jeltsch, S., Müller, E., Müller, P., Petrovici, M. A., et al. (2013). Six networks on a universal neuromorphic computing substrate. *Front. Neurosci.* 7:11. doi: 10.3389/fnins.2013.00011
- Pfeil, T., Jordan, J., Tetzlaff, T., Grübl, A., Schemmel, J., Diesmann, M., et al. (2016). Effect of heterogeneity on decorrelation mechanisms in spiking neural networks: a neuromorphic-hardware study. *Phys. Rev. X* 6:021023. doi: 10.1103/PhysRevX.6.021023
- Pouget, A., Beck, J. M., Ma, W. J., and Latham, P. E. (2013). Probabilistic brains: knowns and unknowns. *Nat. Neurosci.* 16:1170. doi: 10.1038/nn.3495
- Probst, D., Petrovici, M. A., Bytschok, I., Bill, J., Pecevski, D., Schemmel, J., et al. (2015). Probabilistic inference in discrete spaces can be implemented into networks of LIF neurons. *Front. Comput. Neurosci.* 9:13. doi: 10.3389/fncom.2015.00013
- Qiao, N., Mostafa, H., Corradi, F., Osswald, M., Stefanini, F., Sumislawska, D., et al. (2015). A reconfigurable on-line learning spiking neuromorphic processor comprising 256 neurons and 128k synapses. *Front. Neurosci.* 9:141. doi: 10.3389/fnins.2015.00141
- Reinagel, P., and Reid, R. C. (2002). Precise firing events are conserved across neurons. *J. Neurosci.* 22, 6837–6841. doi: 10.1523/JNEUROSCI.22-16-06837.2002
- Salakhutdinov, R. (2010). “Learning deep boltzmann machines using adaptive MCMC,” in *Proceedings of the 27th International Conference on Machine Learning (ICML-10)*, (Haifa), 943–950.
- Savitzky, A., and Golay, M. J. (1964). Smoothing and differentiation of data by simplified least squares procedures. *Anal. Chem.* 36, 1627–1639.
- Schemmel, J., Brüderle, D., Grübl, A., Hock, M., Meier, K., and Millner, S. (2010). “A wafer-scale neuromorphic hardware system for large-scale neural modeling,” in *Proceedings of 2010 IEEE International Symposium on Circuits and Systems (ISCAS)* (IEEE), 1947–1950.
- Schmitt, S., Klähn, J., Bellec, G., Grübl, A., Guettler, M., Hartel, A., et al. (2017). “Neuromorphic hardware in the loop: training a deep spiking network on the brainscales wafer-scale system,” in *2017 International Joint Conference on Neural Networks (IJCNN)* (Anchorage, AL: IEEE), 2227–2234.
- Schmuker, M., Pfeil, T., and Nawrot, M. P. (2014). A neuromorphic network for generic multivariate data classification. *Proc. Natl. Acad. Sci. U.S.A.* 111, 2081–2086. doi: 10.1073/pnas.1303053111
- Sutskever, I., and Hinton, G. (2007). “Learning multilevel distributed representations for high-dimensional sequences,” in *Artificial Intelligence and Statistics*, eds M. Meila and X. Shen (San Juan), 548–555.
- Taylor, G. W., and Hinton, G. E. (2009). “Factored conditional restricted boltzmann machines for modeling motion style,” in *Proceedings of the 26th Annual International Conference on Machine Learning (ACM)*, 1025–1032.
- Tetzlaff, T., Helias, M., Einevoll, G. T., and Diesmann, M. (2012). Decorrelation of neural-network activity by inhibitory feedback. *PLoS Comput. Biol.* 8:e1002596. doi: 10.1371/journal.pcbi.1002596
- Toups, J. V., Fellous, J. M., Thomas, P. J., Sejnowski, T. J., and Tiesinga, P. H. (2012). Multiple spike time patterns occur at bifurcation points of membrane potential dynamics. *PLoS Comput. Biol.* 8:e1002615. doi: 10.1371/journal.pcbi.1002615
- Waldrop, M. M. (2016). The chips are down for Moore’s law. *Nat. News* 530:144. doi: 10.1038/530144a
- Wunderlich, T., Kungl, A. F., Müller, E., Hartel, A., Stradmann, Y., Aamir, S. A., et al. (2019). Demonstrating advantages of neuromorphic computation: a pilot study. *Front. Neurosci.* 13:260. doi: 10.3389/fnins.2019.00260
- Xiao, H., Rasul, K., and Vollgraf, R. (2017). *Fashion-MNIST: A Novel Image Dataset for Benchmarking Machine Learning Algorithms* (Zalando SE).
- Zoschke, K., Güttler, M., Böttcher, L., Grübl, A., Husmann, D., Schemmel, J., et al. (2017). “Full wafer redistribution and wafer embedding as key technologies for a multi-scale neuromorphic hardware cluster,” in *2017 IEEE 19th Electronics Packaging Technology Conference* (Singapore).

**Conflict of Interest:** The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Copyright © 2019 Kungl, Schmitt, Klähn, Müller, Baumbach, Dold, Kugele, Müller, Koke, Kleider, Mauch, Breitwieser, Leng, Gürtler, Güttler, Husmann, Husmann, Hartel, Karasenko, Grübl, Schemmel, Meier and Petrovici. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.