



Neuromorphic Hardware Learns to Learn

Thomas Bohnstingl^{1*†‡}, Franz Scherr^{1‡}, Christian Pehle², Karlheinz Meier² and Wolfgang Maass¹

¹ Institute for Theoretical Computer Science, Graz University of Technology, Graz, Austria, ² Kirchhoff-Institute for Physics, Ruprecht-Karls-Universität Heidelberg, Heidelberg, Germany

OPEN ACCESS

Edited by:

Yansong Chua,
Institute for Infocomm Research
(A*STAR), Singapore

Reviewed by:

Sadique Sheik,
AiCTX AG, Switzerland
Garibaldi Pineda García,
University of Sussex, United Kingdom

*Correspondence:

Thomas Bohnstingl
boh@zurich.ibm.com

†Present Address:

Thomas Bohnstingl,
IBM Research - Zurich, Rüschlikon,
Switzerland

‡These authors have contributed
equally to this work

Specialty section:

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

Received: 31 January 2019

Accepted: 29 April 2019

Published: 21 May 2019

Citation:

Bohnstingl T, Scherr F, Pehle C,
Meier K and Maass W (2019)
Neuromorphic Hardware Learns to
Learn. *Front. Neurosci.* 13:483.
doi: 10.3389/fnins.2019.00483

Hyperparameters and learning algorithms for neuromorphic hardware are usually chosen by hand to suit a particular task. In contrast, networks of neurons in the brain were optimized through extensive evolutionary and developmental processes to work well on a range of computing and learning tasks. Occasionally this process has been emulated through genetic algorithms, but these require themselves hand-design of their details and tend to provide a limited range of improvements. We employ instead other powerful gradient-free optimization tools, such as cross-entropy methods and evolutionary strategies, in order to port the function of biological optimization processes to neuromorphic hardware. As an example, we show these optimization algorithms enable neuromorphic agents to learn very efficiently from rewards. In particular, meta-plasticity, i.e., the optimization of the learning rule which they use, substantially enhances reward-based learning capability of the hardware. In addition, we demonstrate for the first time Learning-to-Learn benefits from such hardware, in particular, the capability to extract abstract knowledge from prior learning experiences that speeds up the learning of new but related tasks. Learning-to-Learn is especially suited for accelerated neuromorphic hardware, since it makes it feasible to carry out the required very large number of network computations.

Keywords: spiking neural networks, learning-to-learn, markov decision processes, multi-armed bandits, neuromorphic hardware, HICANN-DLS, meta-plasticity, transfer learning

1. INTRODUCTION

The computational substrate that the human brain employs to carry out its computational functions, is given by networks of spiking neurons (SNNs). There appear to be numerous reasons for evolution to branch off toward such a design. For example, networks of such neurons facilitate a distributed scheme of computation, intertwined with memory entities, thereby overcoming known disadvantages in contemporary computer designs such as the von Neumann bottleneck. Importantly, the human brain serves as an inspiration for a power efficient learning machine, solving demanding computational tasks while consuming little resources. A characteristic property that makes energy efficient computation possible is the distinct communication among these neurons. In particular, neurons do not need to produce an output at all times. Instead, information is integrated over time and communicated sparsely using a format of discrete events, “spikes.”

The connectivity structure, the development of computational functions in specific brain regions, as well as the active learning algorithms are all subject to an evolutionary process. In particular, evolution has shaped the human brain and successfully formed a learning machine,

capable of carrying out a range of complex computations. In close connection to this, a characteristic property of learning processes in humans is the ability to take advantage of previous, related experiences and use them in novel tasks. Indeed, humans show both, the ability to quickly adapt to new challenges in various domains, and the ability to transfer prior acquired knowledge about different, but related tasks to new, potentially unseen ones (Taylor and Stone, 2009; Robert Canini et al., 2010; Wang and Zheng, 2015).

One strategy to investigate the benefit of a knowledge transfer between different, but related learning tasks is to impose a so-called Learning-to-Learn (L2L) optimization. L2L employs task-specific learning algorithms, but also tries to mimic the slow evolutionary and developmental processes that have prepared brains for the learning tasks humans have to face. In particular, L2L introduces a nested optimization procedure, consisting of an inner loop and an outer loop. In the inner loop, specific tasks are learned, while an additional outer loop aims to optimize the learning performance on a range of different tasks. This concept gave rise to an interesting body of work (Hochreiter et al., 2001; Wang et al., 2016; Finn et al., 2017) and showed that one can endow artificial learning systems with transfer learning capabilities. Recently, this concept was also extended to networks of spiking neurons. In a study by Bellec et al. (2018) it is shown that a biologically inspired circuit can encode prior assumptions about the tasks it will encounter.

Usually, one takes advantage of the availability of gradient information to facilitate optimization, here instead, we employ powerful gradient-free optimization algorithms in the outer loop that emulate the evolutionary process. In particular, we demonstrate the benefits of evolutionary strategies (ES) (Rechenberg, 1973) and cross entropy methods (CE) (Rubinstein, 1997), as they are able to deal with noisy function evaluations and perform in high-dimensional spaces. In the inner loop, on the other hand, we consider reinforcement learning problems (RL problems), such as Markov Decision Processes and Multi-armed bandits. Problems of this type appear quite often in general and therefore, a rich literature has emerged. However, it still remains that learning from rewards is particularly inefficient, as the feedback is given by a single scalar quantity, the reward. We show that by employing the concept of L2L we can produce agents that learn efficiently from rewards and exploit previous experiences on related, new tasks.

As another novelty, we implement the learning agent on a neuromorphic hardware (NM hardware). Specialized hardware of this type has emerged by taking inspiration of principles of brain computation, with the intent to port the advantages of distributed and power efficient computation to silicon chips (Mead, 1990). This holds the great promise to install artificial intelligence in devices without cloud connection and/or limited resource. Numerous architectures have been proposed that are either based on analog, digital or mixed-signal approaches: (Schemmel et al., 2010; Furber et al., 2014; Furber, 2016; Pantazi et al., 2016; Aamir et al., 2018; Ambrogio et al., 2018; Davies et al., 2018; Wunderlich et al., 2018). We refer to Schuman et al. (2017) for a survey on neuromorphic systems.

In order to further enhance the learning capabilities of NM hardware, we exploit the adjustability of the employed neuromorphic chip and consider the use of meta-plasticity. In other words, we evolve a highly configurable plasticity rule that is responsible for learning in the network of spiking neurons. To this end, we represent the plasticity rule as a multilayer perceptron (section 2.5.2) and demonstrate that this approach can significantly boost learning performance as compared to the level that is achieved by plasticity rules that we derive from general algorithms, see section 3.3.

NM hardware is especially well-suited for L2L because it renders the large number of simulations that need to be carried out feasible. Spiking neurons that are simulated on NM hardware typically exhibit accelerated dynamics as compared to their biological counterparts. In addition, the chosen neuromorphic hardware allows to emulate both, the RL environment as well as the learning algorithm at the same acceleration factor and hence, one unlocks the full potential of the specialized neuromorphic chip.

First, in section 2 we will discuss our approaches and methods, as well as the set of tools (https://github.com/bohnstingl/Neuromorphic_Hardware_learns_to_learn) that was used in our experiments. In particular, the employed NM hardware is discussed in section 2.3. Then, in section 3.1 we will exhibit the increase in performance and learning speed that we obtained on NM hardware for the conducted tasks and discuss which gradient-free algorithms worked best for our setting. Afterwards, we discuss in section 3.3 that performance can be further increased by the adoption of a highly customizable learning rule, i.e., meta-plasticity, that is shaped through L2L, and discuss its relevance in transfer learning. We also discuss the impact in terms of simulation time thanks to the underlying NM hardware. Finally, we conclude our findings and results in section 4.

2. METHODS AND MATERIALS

This section provides the technical details to the conducted experiments. First, we describe the background for L2L in section 2.1, and discuss the gradient-free optimization techniques that are employed. Subsequently, we provide details to the reinforcement learning tasks that we considered (section 2.2).

Since the agent that interacts with the RL environments is implemented on a NM hardware, we discuss the corresponding chip in section 2.3. We exhibit the network structure that we used throughout all our experiments in section 2.4. Subsequently, we provide details to the learning algorithms that we used in section 2.5 and discuss methods for analysis.

2.1. Learning-to-Learn and Gradient-free Optimization

The goal of Learning-to-Learn is to enhance a learning systems' capability to learn. In models of neural networks, learning performance can be enhanced by several methods. For example, one can optimize hyperparameters that affect the learning procedure or optimize the learning procedure as such. Often, this optimization is carried out manually and involves a lot of domain

knowledge. Here instead, we evolve suitable hyperparameters as well as learning algorithms automatically by the means of L2L.

In particular, L2L introduces a nested optimization that consists of two loops: an inner loop and an outer loop as displayed in **Figure 1**. In the inner loop, one considers a particular task C_i in which the model \mathcal{N} has to use its learning capabilities to succeed. The outer loop, on the other hand, is responsible to adapt the learning procedure that is used by \mathcal{N} such that it becomes better at learning tasks in a given family \mathcal{F} that share some similar concepts. To express the quality of the learning procedure, we introduce a learning fitness $f(C_i; \Theta)$ that measures how well the model \mathcal{N} can learn a task C_i , e.g., what is the cumulative reward that was achieved. This learning fitness depends on both the specific task that is being learnt, as well as the hyperparameters Θ that characterize the learning procedure. We write the goal of L2L is then as an optimization problem, where we want to find hyperparameters that yield the best learning procedure for tasks in the family \mathcal{F} :

$$\max_{\Theta'} \mathbb{E}_{C \sim \mathcal{F}} [f(C; \Theta')]. \quad (1)$$

In practice, the family of tasks could be comprised of infinite tasks and hence, the expectation in Equation (1) is approximated using batches of N different tasks: $\mathbb{E}_{C \sim \mathcal{F}} [f(C; \Theta)] \approx \frac{1}{N} \sum_{i=1}^N f(C_i; \Theta) = \hat{f}(\Theta)$. As a result of considering different tasks C_i in the inner loop each time, the hyperparameters can only assume task independent concepts that are shared throughout the family. In fact, one can consider L2L as an optimization that happens on two different timescales: fast learning of single tasks in the inner loop, and a slower learning process that adapts hyperparameters in order to boost learning on the entire family of learning tasks.

The L2L scheme allows separating the learning process in the inner loop from the optimization algorithms that work in the outer loop. We used Q-Learning and Meta-Plasticity to implement learning in the inner loop (discussed in section 2.5),

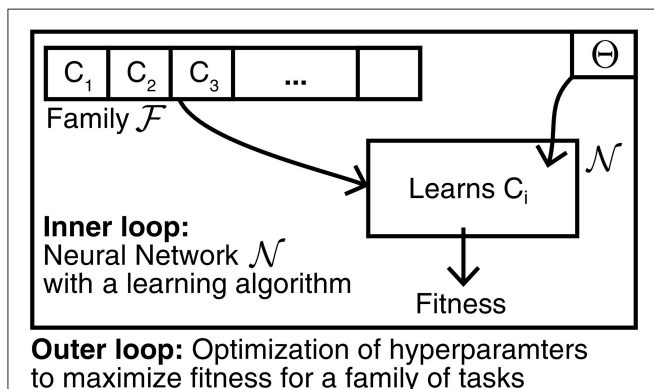


FIGURE 1 | Learning-to-Learn scheme. Learning-to-Learn introduces a nested optimization with two loops. In the inner loop a model is required to perform a learning task C_i from a family \mathcal{F} . The learning capabilities of the model are influenced by hyperparameters Θ that are optimized in the outer loop in order to maximize the average learning fitness over the entire task family \mathcal{F} .

while at the same time, we considered several gradient-free optimization techniques in the outer loop. The requirements for a well-suited optimization algorithm in the outer loop are the ability to operate in a high-dimensional parameter space, the ability to deal with noisy fitness evaluations, the ability to find a good final solution and the ability to do so using a small number of fitness evaluations. Due to this broad set of requirements, the choice of the outer loop algorithm is non-trivial and needs to be adjusted based on the task family that is considered in the inner loop. We selected a set of gradient-free optimization techniques such as cross-entropy methods, evolutionary strategies, numerical gradient-descent as well as a parallelized variation of simulated annealing. In the following, we provide a brief outline of the algorithms used and refer to the corresponding literature. For the concrete implementation, we employ a L2L software framework that provides several such optimization methods (Subramoney et al., 2019). In particular, the L2L optimization is carried out on a Linux-based host computer, whereas the inner loop is simulated in its entirety on the later discussed neuromorphic hardware, section 2.3.

2.1.1. Cross-entropy (CE) (Rubinstein, 1997)

In each iteration, this algorithm fits a parameterized distribution $p(\cdot; \phi)$ to the set of n best-performing hyperparameters in terms of maximum likelihood. In the subsequent step, new hyperparameters are sampled from this distribution and evaluated. Afterwards, the procedure starts over again until a stopping criterion is met. Through this process, the algorithm tries to find a region of individuals where the performance is high on average. We used a univariate Gaussian distribution with a dense covariance matrix.

2.1.2. Evolution Strategies (ES) (Rechenberg, 1973)

In each iteration, this algorithm maintains base hyperparameters Θ which are perturbed by random deviations ϵ to form a new set of n hyperparameters. This set is then evaluated and ranked by their fitness. In a subsequent step, the perturbations are weighted according to their rank to produce a direction of increasing fitness, which is used to update the base hyperparameters. Similar to Cross-entropy, ES also finds a region of hyperparameters with high fitness, rather than just a single one. Note that many variations of this algorithm have been proposed that differ for example in the way how the ranking or how the perturbations are computed (Salimans et al., 2017). In particular, we used Algorithm 1 from Salimans et al. (2017).

2.1.3. Simulated Annealing (SA) (Kirkpatrick et al., 1983)

In each iteration, the algorithm maintains hyperparameters Θ and a temperature T . The hyperparameters are perturbed with a random ϵ , whose size depends on the temperature T , and are evaluated later. The fitness of the unperturbed hyperparameters Θ is then compared with the perturbed hyperparameters Θ' . The Θ' replaces Θ with a probability of $\min\left(1, \exp\left(-\frac{\hat{f}(\Theta') - \hat{f}(\Theta)}{T}\right)\right)$. In the next step, the temperature is decreased following a predefined schedule and the new hyperparameters get

perturbed. In contrast to the other methods discussed before, a single set of hyperparameters is the result. In our experiments, we simultaneously perform a number of parallel SA optimizations, using a linear temperature decay.

2.1.4. Numerical Gradient-Descent (GD)

In each iteration, the algorithm maintains hyperparameters Θ which are perturbed randomly in many directions and then evaluated. Subsequently, the gradient is numerically estimated and an ascending step on the fitness landscape is performed.

2.2. Reinforcement Learning Problems

In all our experiments we considered reinforcement learning problems. Tasks of this type usually require many trials and sophisticated algorithms in order to produce a well-performing agent, since a teacher signal is only available in the form of a scalar quantity, the reward. To the worse, a reward does not arrive at every time step, but is often given very sparsely and only for certain events. **Figure 2A** depicts a generic reinforcement learning loop. The agent observes the current state $s(t)$ of the environment and has to decide on an action $a(t)$. In particular, the agent samples an action according to policy $\pi(a|s)$, which is a probability distribution over actions a given a state s . Upon executing the action, the environment will advance to a new state $s(t + 1)$ and the agent receives a reward $r(t)$. In all our experiments, the RL environment was simulated on the neuromorphic chip.

2.2.1. Markov Decision Process

Markov Decision Processes (MDPs) are a well-known and established model for decision making processes in literature. A MDP is defined by a five-tuple $(\mathbb{S}, \mathbb{A}, p, r, \gamma)$, with \mathbb{S} representing the state space, \mathbb{A} the action space, p the state transition function, r the reward function and γ a discount factor that weights future rewards differently from present ones. In particular, we

are concerned here with such MDPs that exhibit discrete and finite state and action spaces. In addition, rewards are given in the range of $[0, 1]$. **Figure 2B** shows a simple example of such a MDP with $\|\mathbb{A}\| = 2$ and $\|\mathbb{S}\| = 3$.

The goal of solving a MDP is to find a policy actions that yields the largest discounted cumulative reward R that is defined as:

$$R = \sum_{t=0}^T \gamma^t r(t) \tag{2}$$

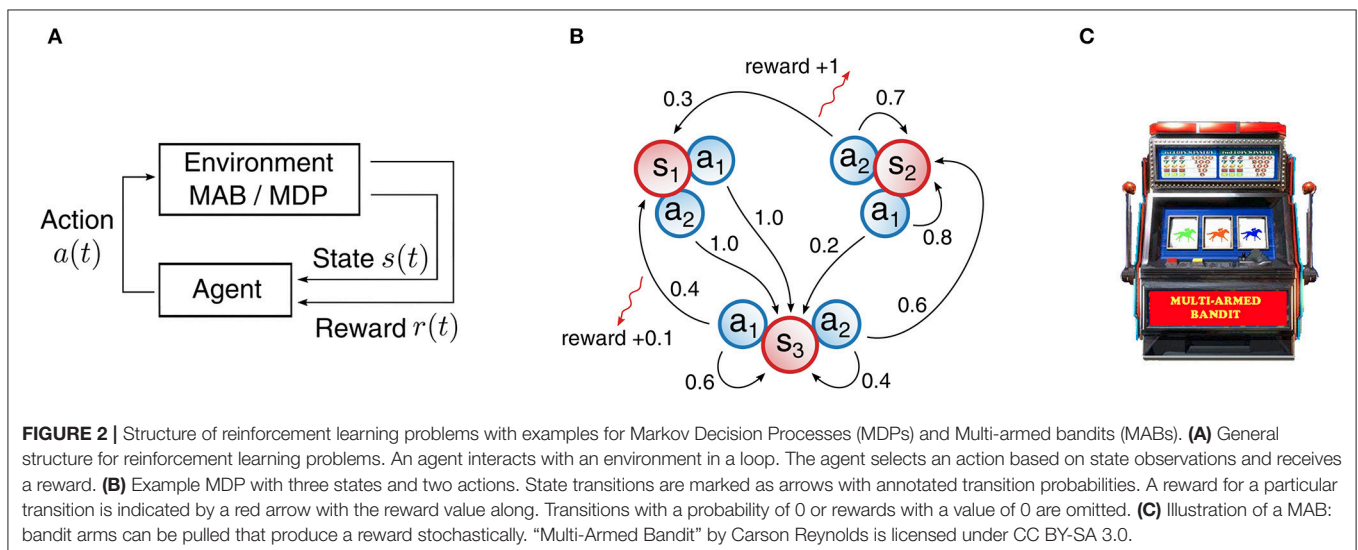
In order to perform well on MDPs, the agent has to keep track of the rewarding transitions and must therefore represent the transition probabilities. Furthermore, the agent has to make a trade-off between exploring new transitions and consolidating already known transitions. Such problems have been studied intensively in literature and a mathematical framework was developed to optimally solve them by Bellman et al. (1954). The so-called Value-Iteration (VI) algorithm emerged from this framework and yields an optimal policy. Therefore, this algorithm is considered as the optimal baseline in all following MDP results.

In order to apply the L2L scheme, we introduce a family of tasks consisting of MDPs with a fixed size of the action and the state space. MDPs of that family are generated according to the following sampling procedure: whenever a new task is required, the rewards r and the transition probabilities p are randomly sampled from the range $[0, 1]$. In addition, the elements of p are normalized such that the outgoing probabilities for all actions in each state sum up to 1.

We report our results in the form of a normalized discounted cumulative reward, where we scale between the performance of a random action selection and the performance of an optimal action selection, given by a policy produced by VI.

2.2.2. Multi-Armed Bandits

As a second category of RL problems, we consider multi-armed bandit (MAB) problems. A MAB is best described as a collection



of several one-armed bandits, each of which produces a reward stochastically when pulled. A depiction of which can be found in **Figure 2C**. In other words, one can view MAB problems as MDPs with a single state and multiple actions. Despite the deceptive simplicity of such problems, a great deal of effort was made in science to study these problems and the celebrated result of Gittins and Gittins (1979) showed that a learning strategy exists.

For the sake of brevity, we use the same notations for MABs as for MDPs. In particular, we say that the environment is always in one state s_1 and the agent is given the opportunity to pull several bandit arms i , which corresponds to actions a_i . In all experiments regarding MABs, we considered two-armed bandits, where each bandit produces a reward of either 0 or 1 with a fixed reward probability p_i . We investigate the impact of L2L on the basis of two different families of MAB tasks:

1. *unstructured bandits*: A task of this family is generated by sampling each of the two reward probabilities p_1, p_2 independently and uniform in $[0, 1]$.
2. *structured bandits*: A task of this family is generated by sampling the reward probability p_1 uniformly in $[0, 1]$ and compute $p_2 = 1 - p_1$.

Similar to MDPs, we report our results for MABs in the form of a normalized cumulative reward, where we scale between the performance of a random action selection and the performance of an oracle that always picks the best possible bandit arm. As a comparison baseline, we employ the Gittins index policy and note that the computation of the Gittins index value is calculated in the same way for both families. In particular, the Gittins index values are calculated assuming that the reward probabilities are independent (*unstructured bandits*).

2.3. Neuromorphic Hardware - HICANN DLSv2

Various approaches for specialized hardware systems implementing spiking neural networks emerged and fundamentally differ in their realizations, ranging from pure digital over pure analog solutions using optical fibers up to mixed-signal devices (Indiveri et al., 2011; Nawrocki et al., 2016; Schuman et al., 2017). Every NM hardware comes with certain advantages and limitations, one promising platform is the HICANN-DLS (Friedmann et al., 2017), herein it is used in the prototype version 2.

The hardware is a prototype of the second generation BrainScaleS-2 system currently under development as part of the Human Brain Project neuromorphic platform (Markram et al., 2011). It represents a scaled-down version of the future full-size chip and is used to evaluate and demonstrate new features as illustrated in this work.

Conceptually the chip is a mixed-signal design with analog circuits for neurons and synapses, spike-based, continuous time communication and an embedded microprocessor. The NM hardware is realized in a 65 nm CMOS process node by the company TSMC. It features 32 neurons of the leaky-integrate-and-fire (LIF) type connected by a 32x32 crossbar array of

synapses such that each neuron can receive inputs from a column of 32 synapses. Synaptic weights can be set with a precision of 6-bits and can be configured row-wise to deliver excitatory or inhibitory inputs. Synapses feature local short-term (STP) and long-term (STDP) plasticity, which is implemented by the embedded microprocessor described later. All analog time constants are scaled down by a factor of 1000 to represent an accelerated neuromorphic system compared to biological time-scales, a feature that is strongly exploited in this paper.

The embedded microprocessor is a 32-bit CPU implementing the Power-PC instruction set with custom vector extensions. It is used as a plasticity processing unit (PPU) to implement all synaptic weight changes. In particular, the PPU allows to devote memory to synapses in order to equip them with tagging mechanisms such as eligibility traces. As a general purpose processor, it can also act on any other on-chip data like neuron and synapse parameters as well as on the network connectivity. It can also send and receive off-chip signals like rewards or other control signals. Because of the large freedom in specifying programs for the PPU (written in C), we investigated different learning algorithms that are explained in section 2.5. They all exploit the proposed network structure from section 2.4 and have the commonality, that the reward information of the state transitions is encoded in the synaptic efficacy. In addition to learning algorithms, the plasticity processing unit also allows implementing environments for an agent. Since the system features a high speedup factor, any environment must also provide the same speedup factor in order to unlock full potential of the neuromorphic hardware, when using a closed-loop setup.

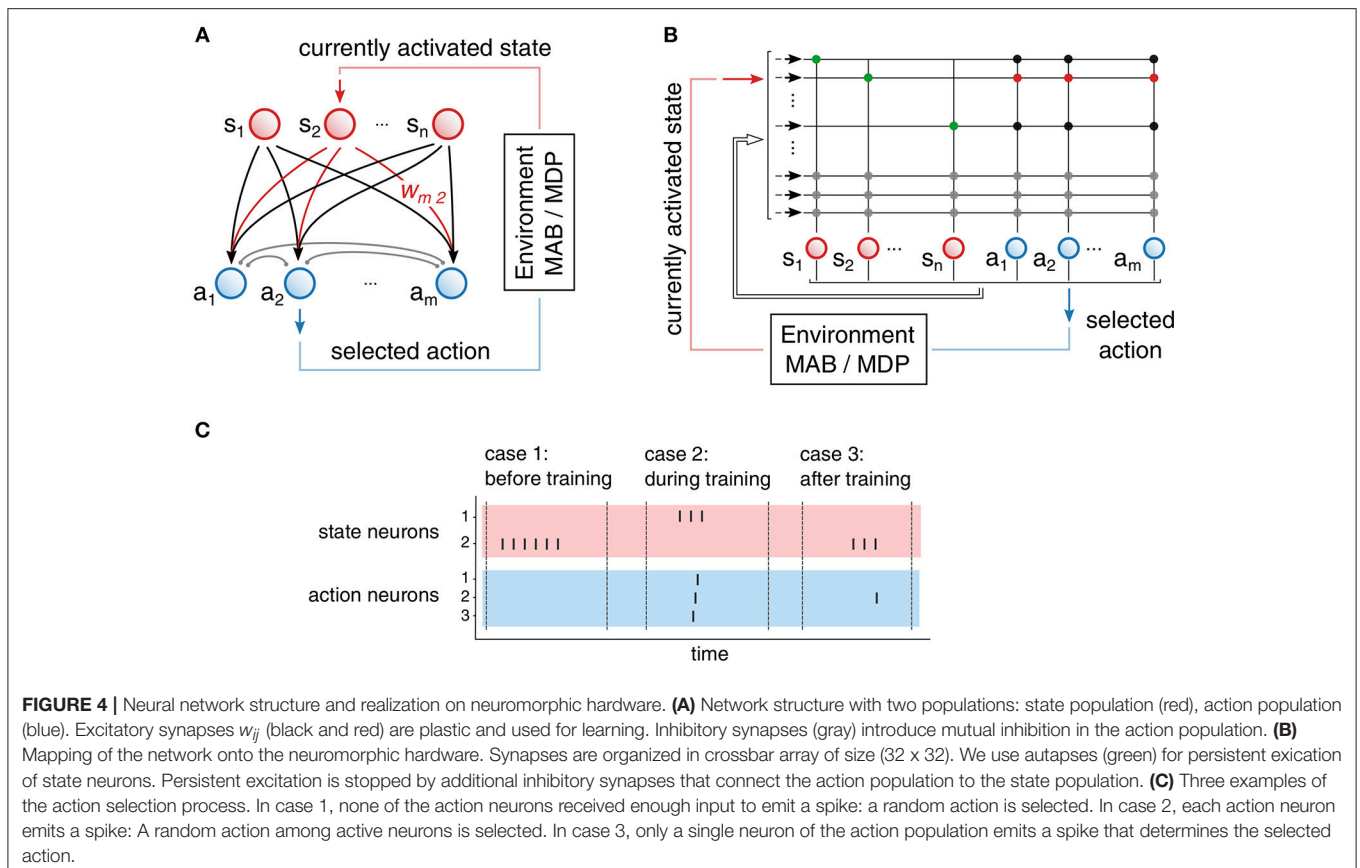
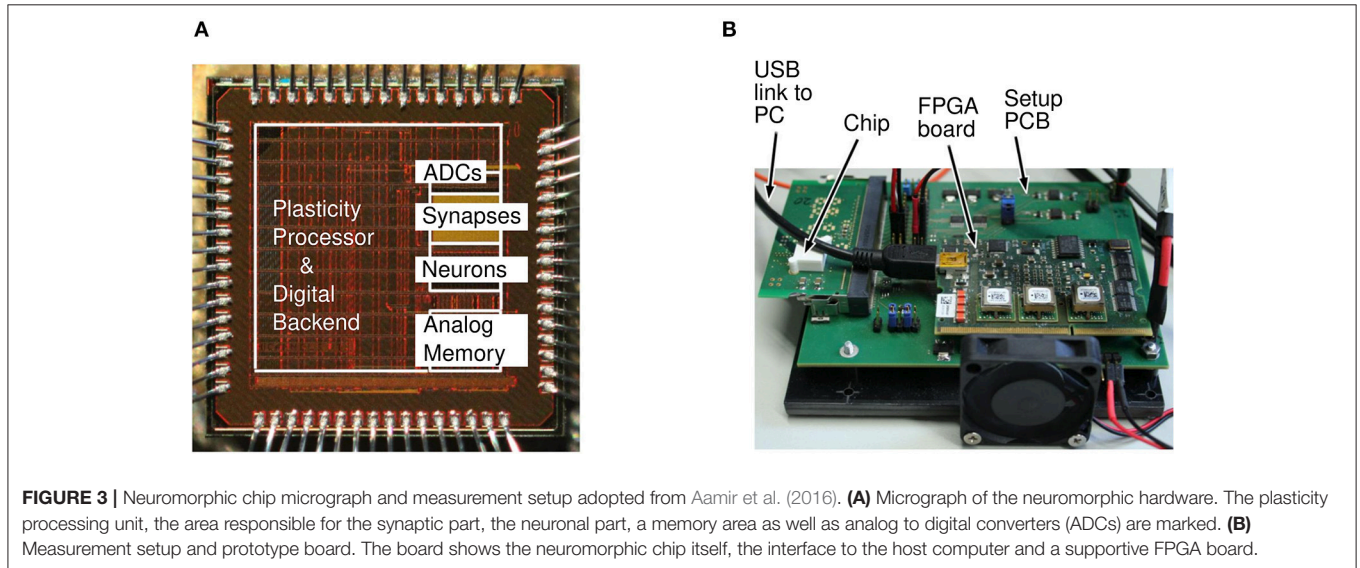
Some of the basic design rationales behind the second generation BrainScaleS-2 system with special emphasis on the PPU are described in Friedmann et al. (2017). **Figure 3A** shows the micrograph of the hardware and **Figure 3B** shows the measurement setup. In addition to other components, the measurement setup hosts the neuromorphic chip, a USB-Interface to connect the baseboard with a host computer as well as a separate FPGA board to control the experiments. The micrograph of the neuromorphic chip shows the different components and where they are located. A description of the actual prototype used in this work including details on the neuron implementation and the synaptic array can be found in Aamir et al. (2016).

2.4. Network Structure and Action Selection

As discussed in section 2.2, the agent is required to select an appropriate action $a(t)$ given a particular state $s(t)$ of the environment. We discuss in this Section how the agent can be implemented using a network of spiking neurons on neuromorphic hardware. Since our experiments were concerned with either Multi-armed bandits or Markov Decision Processes, we designed the network structure for the more general MDP problems. In particular, the design is based on the Markov Property of MDPs, using the fact that the next state $s(t + 1)$ solely depends on the chosen action $a(t)$ and the current state $s(t)$, similarly to Friedrich and Lengyel (2016).

Concretely, we make use of a feed-forward network of spiking neurons with two populations, as illustrated in **Figure 4A**. One population encodes the state of the environment (state population, marked in red) and the second population encodes all possible action choices (action population, marked in blue). We assume that all states exhibit the same number of possible

actions. Under this assumption, the resulting agent commits to specific actions by the following action selection protocol: Given that the agent finds itself in state s_j , then the corresponding state neuron receives stimulating input and produces output spikes that are transmitted to the neurons a_i of the action population by excitatory synapses w_{ij} . Eventually, this stimulation will trigger



a spike in the action population, depending on the synaptic strengths w_{ij} . The action $a(t)$ that will be taken is determined by the neuron of the action population that emits a spike first. In addition, neurons coding for actions are connected inhibitory among each other with synapses of strength ξ , through which a WTA-like network structure arises. Due to this mutual inhibition, mostly a single neuron of the action population will emit a spike and hence, trigger the corresponding action.

In practice, additional tricks are required to implement the proposed scheme on the neuromorphic device, see **Figure 4B**. To continually excite the active state neuron, we send a single spike that triggers a persistent firing through strong excitatory autapses (marked in green). If a neuron from the action population eventually emits a spike, the active state neuron needs to be prevented from further spiking. For this purpose, we use inhibitory synapses of strength ζ projecting from action neurons to state neurons. Due to synaptic delays, more than one action neuron may emit a spike. In such a case, an action is randomly selected among the set of active neurons. It is to be noted that smaller inhibition weights lead to more random exploration, because insufficient inhibition will not prevent spikes of other action neurons, in which case action selection becomes randomized.

One other implementation detail comes from the fact that the synaptic weights on the NM hardware yield a limited resolution of only 6 bit. This might cause that weights saturate at either 0 or the maximum weight value and prevent efficient learning. To avoid this problem, the weights w_{ij} are rescaled with a certain frequency $f_{rescale}$ according to:

$$k = \frac{W_{max} - W_{min}}{\max(w_{ij}) - \min(w_{ij})} \tag{3}$$

$$d = W_{max} - k \max(w_{ij}) \tag{4}$$

$$w'_{ij} = kw_{ij} + d \tag{5}$$

where W_{max} and W_{min} provide the upper and lower rescale boundary.

Figure 4C depicts typical examples of the action selection process for three common cases occurring throughout the learning process. In case 1 (usually before training), a state neuron, i.e., corresponding to state 2, is active and persistently emits a spike. However, none of the synapses connecting to the action neurons is strong enough to cause a spike. In such a case, after a predefined time, the state neuron is externally inhibited and a random action is selected by the implementation of the environment. In case 2 (likely during learning), another state neuron is active, but all synapses to the action neurons are strong enough to cause every action neuron to spike before the mutual inhibition sets in. In such a case, a random action among the active action neurons is selected (random selection is performed by the environment). Eventually, the system reaches case 3 (after learning), where a single action neuron is excited by a given state neuron.

Learning in this network structure is implemented by synaptic plasticity rules that act upon the excitatory weights w_{ij} projecting from the state to the action population. In particular, these

weights pin down which action has the highest priority for each state.

2.5. Learning Algorithms

2.5.1. Q-Learning

MDPs have been studied intensively in computer science and a rigorous framework on how to solve problems of this kind optimally was introduced by Bellman. An important quantity in MDPs is the so-called Q-Function, or Action-Value function. The Q-Function $Q^\pi(s, a)$ expresses the expected discounted cumulative reward, when the agent starts in state s , takes action a and subsequently proceeds according to its policy π . Formally, one writes this as:

$$Q^\pi(s_j, a_i) = \mathbb{E} \left[\sum_{k=0}^{\infty} \gamma^k r(t+k+1) | s(t) = s_j, a(t) = a_i \right] \tag{6}$$

where γ is the discount factor of the MDP and $r(t)$ is the immediate reward at time step t . As discussed before in section 2.2.1, we consider only discrete MDPs and the Q-Functions can therefore be represented in a tabular form. This property suits our network structure, since the synapses that project from the state population to the action population w_{ij} can represent all Q-values

$Q^\pi(s_j, a_i)$. Hence, we define $w_{ij} \stackrel{def}{=} Q^\pi(s_j, a_i)$.

To solve MDPs, the goal is to determine the optimal policy π^* . A common approach is to infer the Q-Function of an optimal policy Q^* and then reconstruct the policy according to:

$$\pi^*(a|s) = \begin{cases} 1 & \text{if } a = \arg \max_{a'} Q^*(s, a') \\ 0 & \text{else} \end{cases} \tag{7}$$

Indeed, as we aim to encode Q-values in synaptic weights w_{ij} , we emphasize that the argmax operation will be naturally carried out by the spiking neural network, as proposed in section 2.4. To infer the Q-values of the optimal policy, we derive rules of synaptic plasticity based on temporal difference algorithms as proposed by Sutton and Barto (1998).

2.5.1.1. TD(1)-Learning

Temporal Difference Learning (TD(1)-Learning) was developed as a method to obtain the optimal policy. The estimate of the optimal Q-Function is improved based on single interactions with the environment and TD(1)-Learning is guaranteed to converge to the correct solution (Watkins and Dayan, 1992; Dayan and Sejnowski, 1994). Based on TD(1), the synaptic weight updates take on the following form:

$$w_{ij}(t+1) = w_{ij}(t) + \alpha \left(r(t) + \gamma \max_k w_{kj}(t) - w_{ij}(t) \right) \tag{8}$$

for $s(t) = s_j, a(t) = a_i$

Where α denotes a learning rate.

2.5.1.2. TD(λ)-Learning

The convergence speed of TD(1)-Learning can be further improved if one uses additional eligibility traces $e_{ij}(t)$ per synapse. The resulting algorithm is then referred to as TD(λ)-Learning. In particular, the trace e_{ij} indicates to what extent a current reward makes the earlier visited state-action pair (s_j, a_i) more valuable and several convergence proofs of the resulting algorithm have been established (Dayan, 1992; Dayan and Sejnowski, 1994). To implement the algorithm, we update eligibility traces at every time step t according to the schedule

$$e_{ij}(t) = \begin{cases} \gamma \lambda e_{ij}(t-1) + 1 & \text{if } s(t) = s_j \text{ and } a(t) = a_i \\ \gamma \lambda e_{ij}(t-1) & \text{otherwise} \end{cases} \quad (9)$$

where the parameter $\lambda \in [0, 1]$ controls how many state transitions are taken into account. In the limit of $\lambda = 1$ one obtains TD(1)-Learning. In addition, we define an error $\delta(t)$ according to

$$\delta(t) = r(t) + \gamma \max_k w_{kj}(t) - w_{ij}(t) \quad \text{for } s(t) = s_j, a(t) = a_i \quad (10)$$

which enables us to express the resulting plasticity rule as a product of the eligibility trace and error $\delta(t)$. This update is carried out for every synapse w_{ij} :

$$w_{ij}(t+1) = w_{ij}(t) + \alpha \delta(t) e_{ij}(t) \quad \text{for all } i, j \quad (11)$$

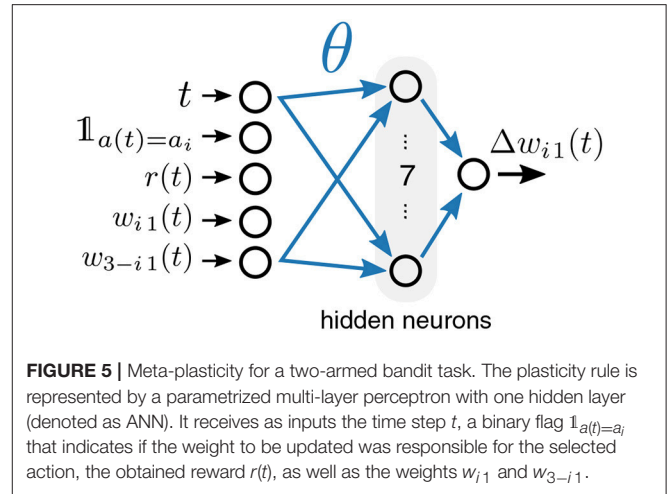
2.5.2. Meta-Plasticity

In order to tailor the specific update rule toward the actual task family at hand, we approached the problem also from the perspective of meta-plasticity. That is, we represent the synaptic weight update by a parameterized function approximator. We then optimize its parameters with L2L in such a way that a useful learning rule for a given task family emerges. We used a multilayer perceptron, the architecture of which is visualized in Figure 5. The perceptron receives five inputs, computes seven hidden units with sigmoidal activation and provides one output, the weight update Δw_{ij} . Effectively, the input to output mapping of this approximator is specified by a number of free parameters θ (weights of the multilayer perceptron) that are considered as hyperparameters and optimized as part of the L2L procedure. Since the multilayer perceptron is a type of an artificial neural network, this plasticity rule is referred to as ANN learning rule. The update of synaptic weights w_{ij} thus takes on the general form of:

$$w_{ij}(t+1) = w_{ij}(t) + f_{\text{ANN}}(\text{inputs}_{ij}(t); \theta) \quad (12)$$

The specific choice of inputs is salient for the possible set of learning rules that can emerge. In the case of the ANN learning rule, we only considered structured MAB, where each of the two synapses is updated at every time step. We set the inputs in this case to a vector

$$\text{inputs}_{i1}(t) = \begin{pmatrix} t \\ \mathbb{1}_{a(t)=a_i} \\ r(t) \\ w_{i1}(t) \\ w_{3-i1}(t) \end{pmatrix} \quad (13)$$



that is composed of the current time step t , the obtained reward $r(t)$, the weight $w_{i1}(t)$, and the weight of the synapse associated to the other bandit arm $w_{3-i1}(t)$. In addition, we included here a binary flag $\mathbb{1}_{a(t)=a_i}$ that is one iff the postsynaptic neuron caused the executed action at the last time step.

2.6. Analysis of Meta-Plasticity

After optimizing an artificial neural network in our meta-plasticity approach, we may have limited insight in what causes the emergent plasticity rule to work well. Therefore, we conduct in section 3.3 an analysis of the arising plasticity rule based on an approach called functional Analysis of Variance (fANOVA) which was presented by Hutter et al. (2014). This method originally aims to assess the importance of hyperparameters in the machine learning domain. It does so by fitting a random forest to the performance data of the machine learning model that was gathered using different hyperparameters.

We adopted this method but applied it to a slightly different, but related problem. Our goal is to assess the impact of each input of the ANN rule with respect to its output. To do so, the weights Θ of the plasticity network remain fixed, while the input values to the plasticity network as well as the output from the plasticity network are considered as inputs to the fANOVA framework. Based on this data, a random forest with 30 trees is fitted and the fraction of the explained variance of the output with respect to each input variable can be obtained.

3. RESULTS

This section presents the results of our approach implemented on the described neuromorphic hardware. First, we report how L2L can improve the performance and learning speed in section 3.1. Then, we investigate the impact of outer loop optimization algorithms in section 3.2 and demonstrate in section 3.3 that Meta-Plasticity yields competitive performance, while also enhancing transfer learning capabilities. Finally, we investigate the speedup gained from the neuromorphic hardware by comparing our implementation on the NM hardware

to a pure software implementation of the same model in section 3.4.

3.1. Learning-to-Learn Improves Learning Speed and Performance

Here, we first demonstrate the generality of our network structure when applied to Markov Decision Processes. Then, we examine the effects of an imposed task structure more closely by investigating Multi-armed Bandit problems. To efficiently train the network of spiking neurons, we employed Q-Learning and derived corresponding plasticity rules, as described in see section 2.5.1. The plasticity rule, as well as the concrete implementation on NM hardware, are influenced by hyperparameters Θ that we optimized by L2L, such that the cumulative discounted reward for a given family of tasks is improved on average, see section 2.1.

We implemented a neuromorphic agent that learns MDPs. In fact, the proposed network structure in section 2.4 is particularly designed for such tasks and we applied concretely TD(λ), see Equation (11). Hyperparameters included all occurring parameters of the employed TD(λ)-Learning rule $\alpha, \gamma, \lambda, \xi, \zeta, f_{\text{rescale}}, W_{\text{max}}, W_{\text{min}}$. Therefore, the complete hyperparameter vector was given as $\Theta = (\alpha, \gamma, \lambda, \xi, \zeta, f_{\text{rescale}}, W_{\text{max}}, W_{\text{min}})$. We used the discounted cumulative reward, Equation (2), as the fitness function $f(C; \Theta)$ and optimized Θ using CE. We used a batch size of $N = 20$.

The results for the MDP tasks are depicted in **Figure 6A** where we report the discounted cumulative reward for $T = 2,000$ steps. The discounted cumulative reward is normalized in such a way, that VI is scaled to 1 and the random policy is scaled to 0. To compare with, we used TD(λ)-Learning as a baseline, using the implementation from a software library¹ without a spiking neural network (green line).

We found that applying L2L improved the discounted cumulative reward (red solid line), compared to the case where the hyperparameters are randomly chosen (blue line). In addition, the learning speed was also increased, which can be seen in the zoom depicted in **Figure 6B**.

In the case of MABs, we focused on small networks and two arms in the bandit, which allowed us to complement the results that were obtained for general MDPs of larger size. We considered two families of MABs: *unstructured bandits* and *structured bandits* (2.2.2) which the neuromorphic agent had to learn using the TD(1)-Learning rule, see Equation (8), where we set $\gamma = 1$. In addition, we introduced here a learning rate schedule $\alpha(t) = \alpha_{\text{decay}}^t \cdot \alpha_0$ that decays a base learning rate α_0 at every time step by a constant decay factor of $\alpha_{\text{decay}} \in [0, 1]$. We then used L2L to carry out a hyperparameter optimization separately for both MAB families and optimized the parameters of the TD(1)-Learning rule α_0 and α_{decay} , the inhibition strength among action neurons ξ and the inhibitory

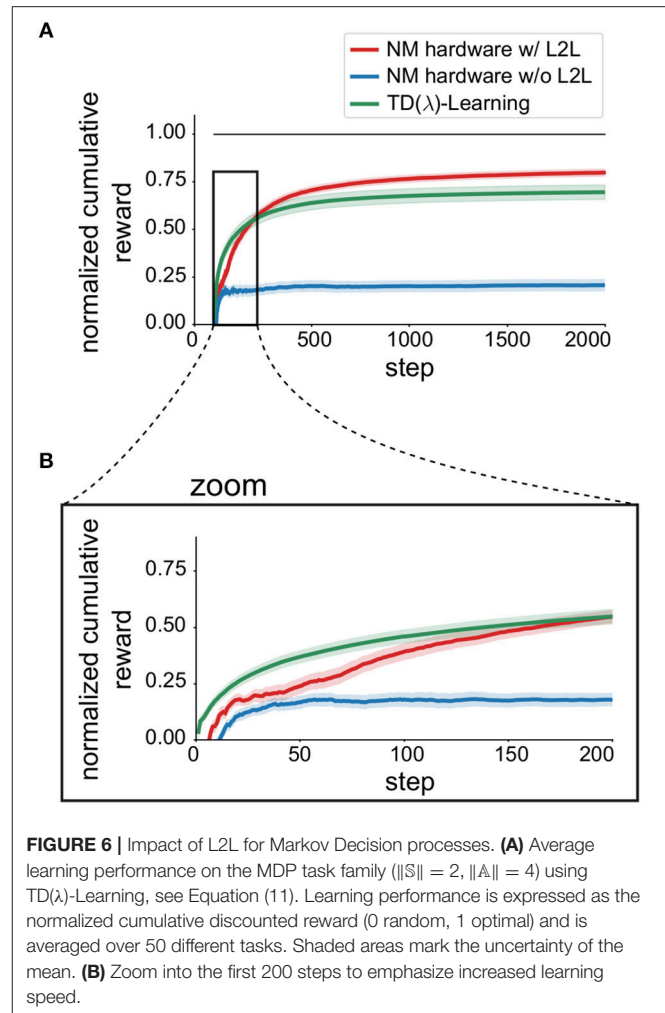
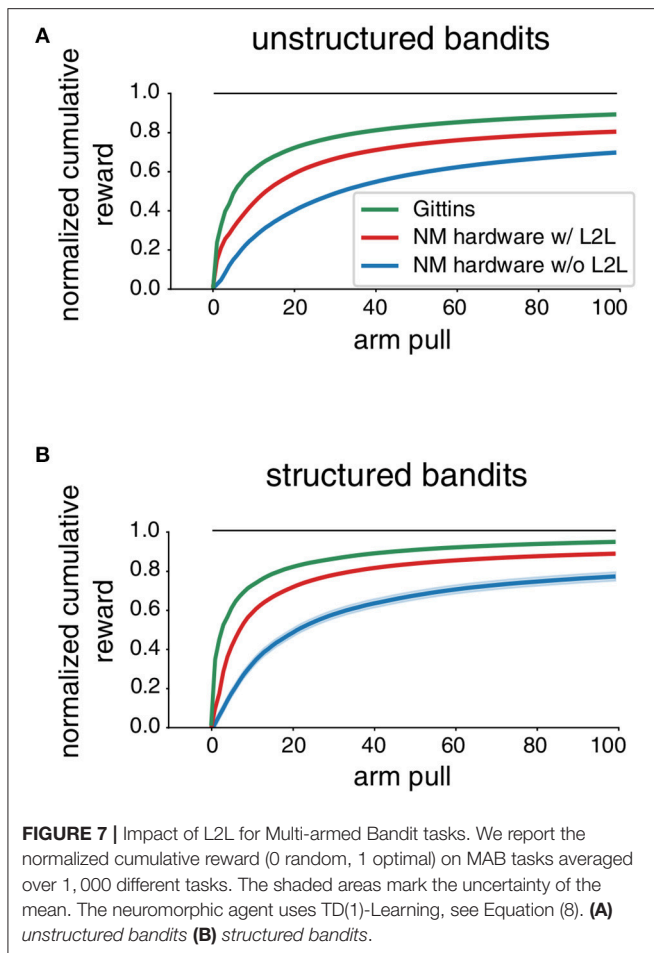


FIGURE 6 | Impact of L2L for Markov Decision processes. **(A)** Average learning performance on the MDP task family ($\|S\| = 2, \|A\| = 4$) using TD(λ)-Learning, see Equation (11). Learning performance is expressed as the normalized cumulative discounted reward (0 random, 1 optimal) and is averaged over 50 different tasks. Shaded areas mark the uncertainty of the mean. **(B)** Zoom into the first 200 steps to emphasize increased learning speed.

weights of synapses that connect the action population to the state population ζ . Hence, the hyperparameter vector was given as $\Theta = (\alpha_0, \alpha_{\text{decay}}, \xi, \zeta)$. We used the cumulative reward as the fitness function $f(C; \Theta)$ and optimized Θ using CE. We used a batch size of $N = 40$.

In **Figure 7** we report the performance results that were obtained before and after applying L2L. The agent interacted for $T = 100$ steps with a single MAB and we compare with a baseline given by the Gittins index policy, as described in section 2.2.2. We found that after performing a L2L optimization the performance was enhanced, which was even more apparent for *structured bandits*. In particular, L2L endowed the agent with a better learning speed, which is exhibited by a faster rising of the performance curve. This can only be achieved when the hyperparameters of the learning system are well-tailored to the tasks that are likely to be encountered, which was the responsibility of L2L. We also observed that the agent could still learn a MAB task to a reasonable level even if no L2L optimization was carried out. This is implied by the fact that TD(1)-Learning is primed to learn RL tasks. However, this also raises the question of how well such a general plasticity rule can

¹<https://pymdptoolbox.readthedocs.io/en/latest/index.html>

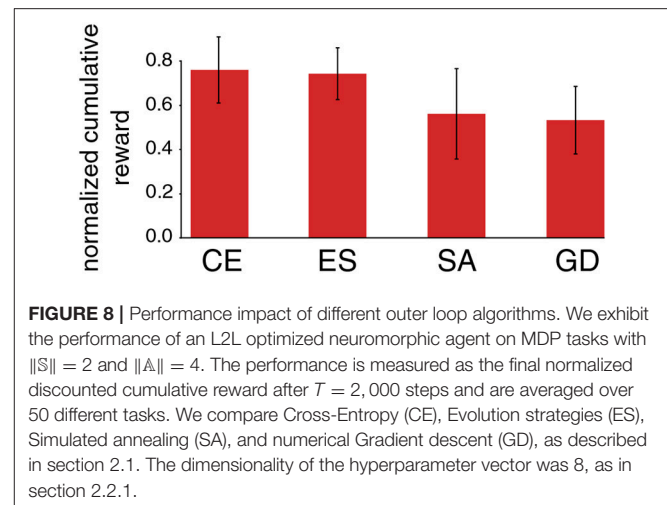


adapt to the level of variations exhibited by analog circuitry. We consider extensions in section 3.3.

3.2. Performance Comparison of Gradient-Free Optimization Algorithms in the Outer Loop

The results presented so far suggest that the concept of L2L can improve the overall performance and also lays the foundation that abstract knowledge about the task family at hand is integrated into an agent. However, the choice of a proper outer loop optimization algorithm is also crucial for this scheme to work well. The modular structure of the L2L approach used in this paper allows to interchange different types of optimization algorithms in the outer loop for the same inner loop task. To demonstrate the impact in terms of performance when using different optimization algorithms, several such algorithms were investigated for both general MDPs and also for specialized MAB tasks. **Figure 8** shows a comparison of the final discounted cumulative reward at the end of the tasks for different outer loop optimization algorithms.

Depending on the inner loop task considered, we found that the cross-entropy (CE) method, as well as evolution strategies (ES), work well because both aim to find a region

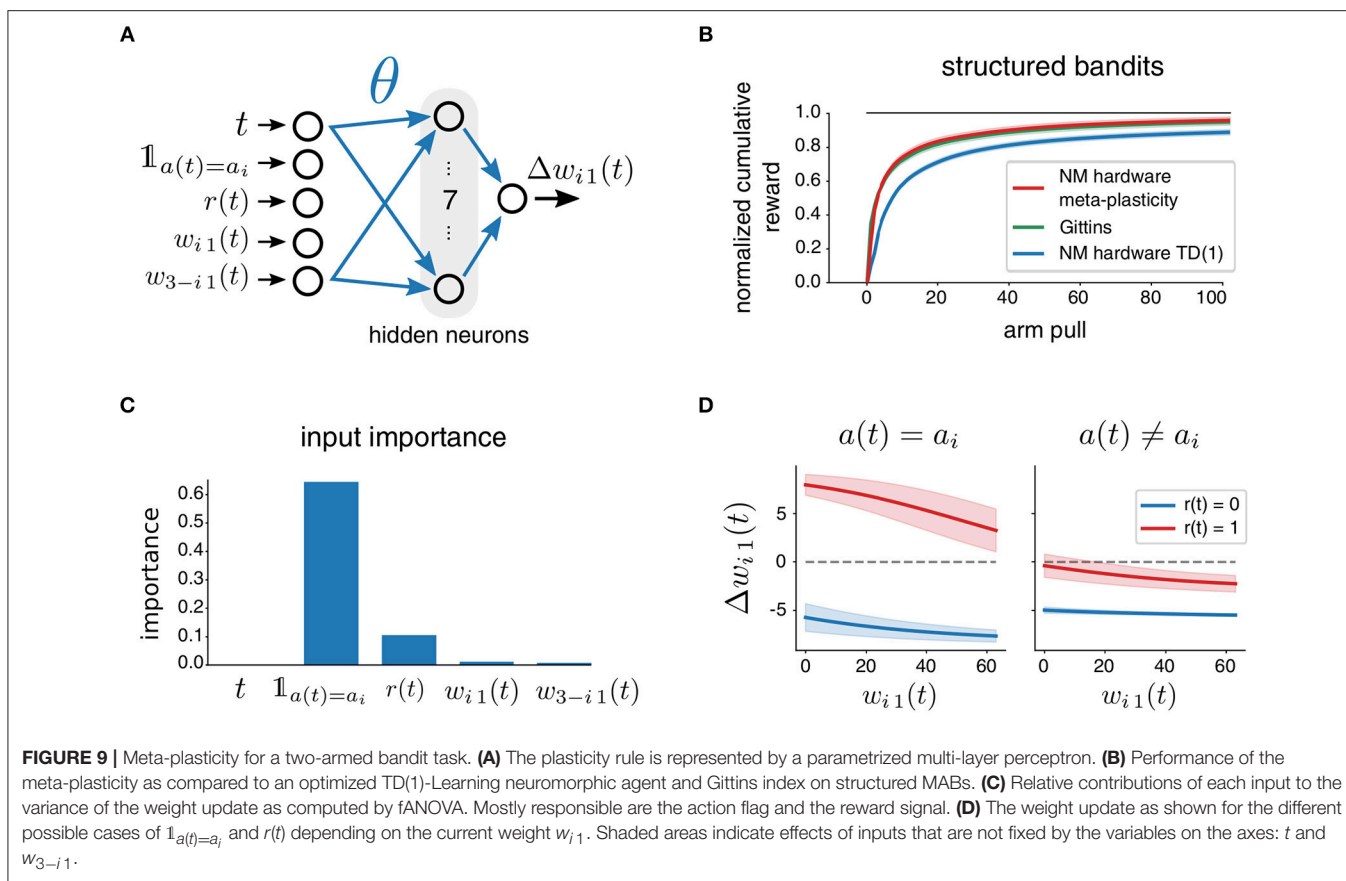


in the hyperparameter space, where the fitness is high. This property is particularly desired when it comes to noise in the fitness landscape due to imperfections of an underlying neuromorphic hardware. In addition, both can cope with noisy fitness evaluations and do not overestimate a single fitness evaluation which could easily lead to a wrong direction in the presence of high noise in the fitness landscape.

However, a simpler algorithm such as simulated annealing (SA) can also find a hyperparameter set with rather high fitness. Especially when running multiple separate annealing processes in parallel with different starting points, the results can almost compete with the ones found by CE or ES. However, SA does not aim at finding a good parameter region but just tries to find a single good set of working hyperparameters. This is prone to cause problems because a single good set of hyperparameters offers less robustness compared to an entire region of well-performing hyperparameters. A simple numerical gradient-based approach did not yield good results at all because of the noisy fitness landscape. In general, the developer is free to choose any optimization algorithm in the outer loop when using L2L. New algorithms can also be implemented which are specially tailored to a particular problem class, which can lead to a new research direction.

3.3. Performance Improvement Through Meta-Plasticity

Since the plasticity rule used so far is based on TD(1)-Learning, and also agnostic to the hardware being used, we raised the question if one could improve training on particular tasks by using an evolved plasticity rule, tailored specifically toward the neuromorphic device and task family at hand. We specified the plasticity rule by a multilayer perceptron with 7 hidden units (**Figure 9**) and considered the weights thereof as hyperparameters. This is apparently the first example of meta-plasticity on neuromorphic hardware, where a rule for synaptic plasticity is evolved through optimization by L2L.



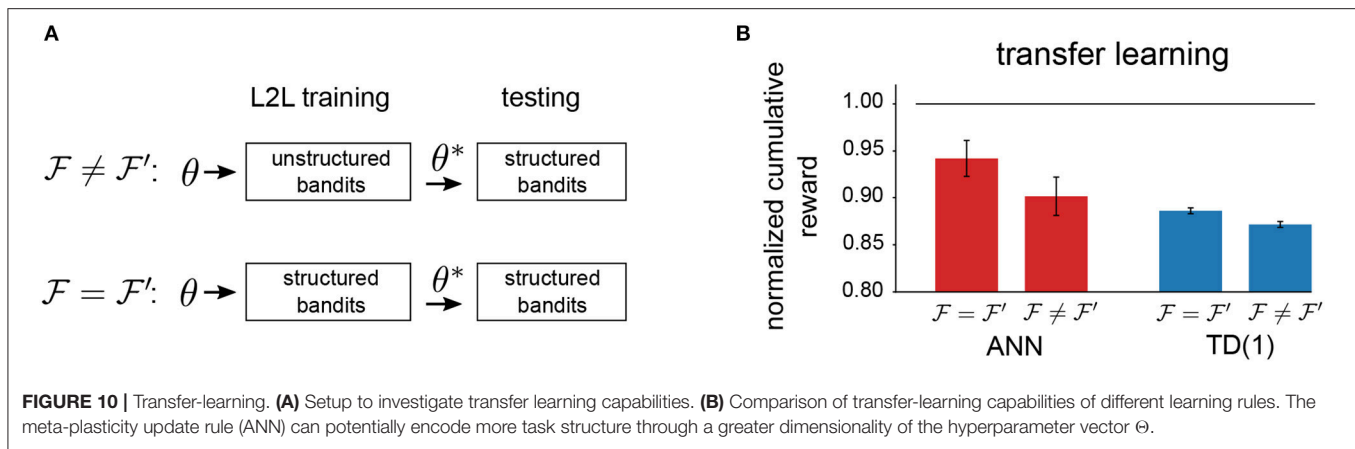
To test the approach, we used L2L to optimize all occurring hyperparameters on the task family of *structured bandits*. In particular, the hyperparameter vector was composed of the parameters of the plasticity rule θ and the inhibition strengths ξ and ζ : $\Theta = (\theta, \xi, \zeta)$. We used the cumulative reward, Equation (2), as the fitness function $f(C; \Theta)$ and optimized Θ using CE with a batch size of $N = 40$.

We summarize our results in **Figure 9** and observed a drastic increase in learning performance. Clearly, the use of meta-plasticity endowed the agent with better skill at learning *structured bandits*, as compared to the TD(1)-Learning rule. It also allows the agent to achieve a performance that is on the same level as the Gittins index policy. This highlights that the evolved plasticity rule can absorb task-structure, and counteract possible negative effects of imperfections in the neuromorphic hardware.

Even though the arising learning rule performs well on average on the family of tasks it has been trained on, there is no theoretical guarantee for that. Hence, an analysis of the optimized learning rule was conducted, where we examined the importance of the multiple inputs provided to the update rule for the resulting output, see **Figure 9C**. Apparently, the most important inputs are the flag that represents if the current weight was responsible for the last action and the obtained reward. Since both of the inputs can assume only two values, one can visualize the four different cases in four different curves. We report the expected weight change depending on the current weight, see **Figure 9D**,

where we average over other unspecified inputs. Updates for weights which were responsible for the previous action are in the direction of the obtained rewards. Hence, the meta-plasticity rule reinforced actions depending on the reward outcome, similarly to Q-learning rules. Interestingly however, the update of the synaptic weight which had not caused the last action was always negative independently of the reward. We believe that L2L simply found that it does not matter what happens to the weight that did not cause actions, because as long as it does not increase, it will not disturb the current belief of the best bandit arm.

To test if the reinforcement learning agent on the neuromorphic hardware has been optimized for a particular range of tasks, we carried out another experiment. We tried to answer if the agent can take advantage of the abstract task structure if it was present. To do so, we always tested learning performance on *structured bandits*, denoted as \mathcal{F}' . For optimization with L2L, we instead used either *unstructured bandits* or *structured bandits*, and we denote the family on which hyperparameter optimization was carried out by \mathcal{F} . This experimental protocol (**Figure 10A**) allowed us to determine to which extent abstract task structure can be encoded in hyperparameters. We report the results for neuromorphic agents in **Figure 10B**, where we considered the TD(1)-Learning rule and the meta-plasticity learning rule. Consistently, we observed that optimizing



hyperparameters for the appropriate task family enhances performance. However, we conjecture that the greater adjustability of the meta-plasticity learning rule renders it to be better suited for transfer learning as compared to TD(1)-Learning rule.

3.4. Exploiting the Benefit of Accelerated Hardware for L2L

One of the main features of neuromorphic hardware devices is the ability to simulate spiking neural networks very fast and efficiently. To make this more explicit for the MDP tasks, a software implementation with the same network structure and the same plasticity rule was conducted on a standard desktop PC using one single core of an IntelTM XeonTM CPU X5690 running at 3.47 GHz. The spiking neural network was implemented using the Neural Simulation Tool (NEST) (Gewaltig and Diesmann, 2007) with a Python interface and the plasticity rule as well as the environment were also implemented in Python. To have a better comparison, two families of MDP tasks with different sizes of $\|\mathcal{S}\|$ and $\|\mathcal{A}\|$ were defined. The first family is defined by $\|\mathcal{S}\| = 2$ and $\|\mathcal{A}\| = 4$ (small MDP) and the second family by $\|\mathcal{S}\| = 6$ and $\|\mathcal{A}\| = 8$ (large MDP).

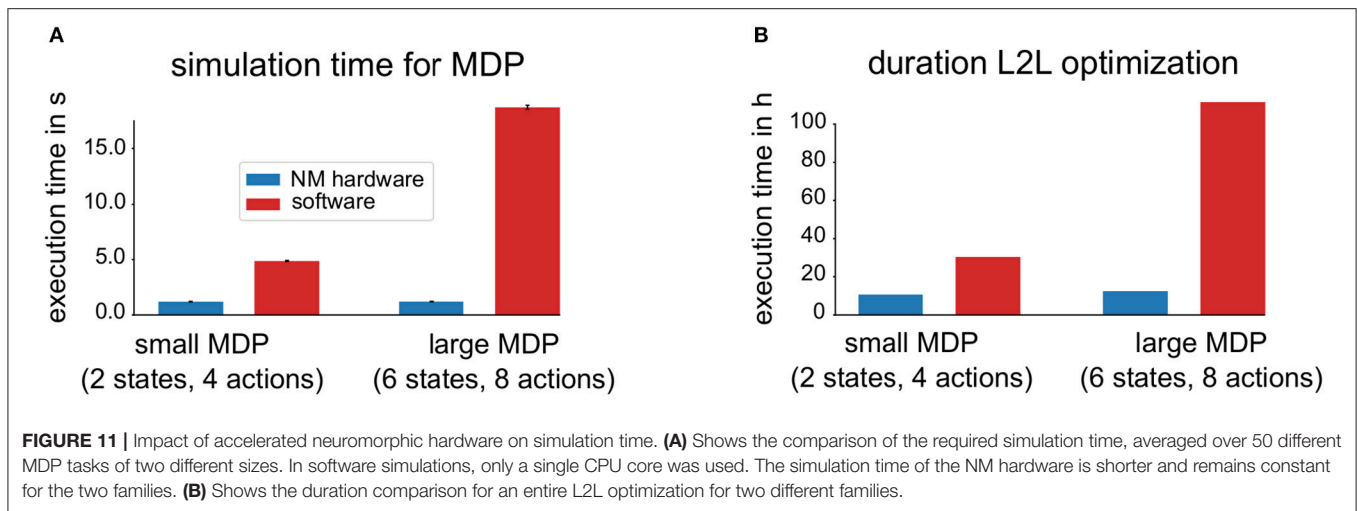
Figure 11A shows a comparison of the simulation time needed for a single randomly selected MDP tasks, averaged over 50 MDPs and for each of the two families. The simulation times include implementation specific overheads, for example, the communication overhead with the neuromorphic hardware. One can see that the simulation time needed for MDP tasks with both sizes are shorter using the neuromorphic hardware and in addition, the simulation time needed to solve the larger task does not increase. First, this indicates, that the neuromorphic hardware can carry out the simulation of the spiking neural network faster and second, that using a larger network structure does not yield an additional cost, as long as the network can fit on the NM hardware. In contrast to this, using more neurons requires longer simulation times in pure software. A similar key message can be found in **Figure 11B**, where instead of a single MDP run, an entire L2L run is evaluated on the neuromorphic hardware as well as with the software implementation. Both, the L2L run on neuromorphic hardware as well as the one in software

can in principle be easily parallelized when using more hardware systems or more CPU cores which would decrease the overall simulation time. Note that scheduler overheads are not taken into considerations.

4. DISCUSSION

Outstanding successes have been achieved in the field of deep learning, ranging from scientific theories and demonstrators to real-world applications. Despite impressive results, deep neural networks are not out of the box suitable for low-power or resource-limited applications. Instead, spiking neural networks are inspired by the brain, an arguably very power efficient computing machine. In this work we employ a neuromorphic hardware that was designed to port key aspects of the astounding properties of this biological circuitry to silicon devices.

The human brain has been prepared by a long evolutionary process with a set of hyperparameters and learning algorithms that can be used to cover a large variety of computing and learning tasks. Indeed, humans are able to generalize task concepts and port them to new, similar tasks, which provides them with a tremendous advantage as compared to most of the contemporary neural networks. In order to mimic this behavior, we employed gradient-free optimization techniques, such as the cross-entropy method or evolutionary strategies (see section 2.1), applied in a Learning-to-Learn setting. This two-looped scheme combines task-specific learning with a slower evolutionary-like process that results in a good set of hyperparameters as demonstrated in section 3.1. The approach is generic in the sense that both, the algorithms mimicking the slower evolutionary process and the learning agent can be exchanged. In principle, any agent with learning capabilities can be used as the learning agent and any optimization algorithms as the evolutionary process. We found that some outer loop optimization algorithm perform better than others and the optimization algorithms should ideally be chosen with the inner loop task in mind. Outer loop optimization algorithms need to operate in a high-dimensional parameter space, have the ability to deal with noisy result evaluations, have the ability to find a good final solution and also require a low number of parameter evaluations before



reaching a good solution. Algorithms that aim to find a region of hyperparameters with high performance such as evolution strategies or cross-entropy worked the best for us, see section 3.2.

L2L offers both, either to find optimal hyperparameters for a fixed individual task or to boost transfer learning capabilities of an agent when using a family of tasks. In addition, new optimization algorithms can be developed to further improve performance in the outer loop of L2L. In this work, we used reinforcement learning problems in connection with NM hardware to demonstrate the aforementioned benefits.

In particular, the concept of L2L allows to shape highly adjustable plasticity rules for specific task families. The usage is not only limited to spiking neural networks but can also be applied to artificial neural networks. This may yield potential for a future research direction. Apparently, this is the first time that the idea of L2L and Meta-Plasticity was applied to a NM hardware, see section 3.3. In addition, the NM hardware provides the possibility to implement advanced plasticity rule on a separate digital processor on-chip. This enables the search for new plasticity rules and might also enable new research directions.

A central role in the approaches explained in this paper is the used NM hardware. It allows to emulate a spiking neural network with a significant speedup compared to the biological equivalent, which makes a large number of computations, required in the L2L scheme feasible. To quantify the overall speedup of the accelerated NM hardware, a comparison with a pure software simulation on a conventional computer was carried out (see **Figure 11**). We conclude that the two-looped L2L scheme as well as the highly adjustable on-chip plasticity rule are especially suited for accelerated neuromorphic hardware.

REFERENCES

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

AUTHOR CONTRIBUTIONS

WM, TB, and FS developed the theory and experiments. TB implemented and conducted experiments with regard to MDPs, benchmarked performance impact of outer loop optimization algorithms and probed the performance benefit of NM hardware. FS implemented and conducted experiments with regard to MABs. FS, CP, and WM conceived meta-plasticity, FS and CP implemented it. FS tested the benefits in transfer learning. TB, FS, CP, WM, and KM wrote the paper.

FUNDING

This research/project was supported by the HBP Joint Platform, funded from the European Union's Horizon 2020 Framework Programme for Research and Innovation under the Specific Grant Agreement No. 785907 (Human Brain Project SGA2).

ACKNOWLEDGMENTS

We thank Anand Subramoney for his support and the contributions to the Learning-to-Learn framework. We are also grateful for the support during the experiments with the neuromorphic hardware. In particular, we like to thank David Stöckel, Benjamin Cramer, Aaron Leibfried, Timo Wunderlich, Yannik Stradmann, Christian Mauch, and Eric Müller. Furthermore, we also like to thank Elias Hajek for useful comments on earlier versions of the manuscript.

Aamir, S. A., Stradmann, Y., Müller, P., Pehle, C., Hartel, A., Gruebl, A., et al. (2018). An accelerated LIF neuronal network array for a large scale mixed-signal neuromorphic architecture. *IEEE Trans. Circuits Syst.* 12, 4299–4312. doi: 10.1109/TCSL.2018.2840718

Ambrogio, S., Narayanan, P., Tsai, H., Shelby, R. M., Boybat, I., di Nolfo, C., et al. (2018). Equivalent-accuracy accelerated neural-network training

- using analogue memory. *Nature* 558, 60–67. doi: 10.1038/s41586-018-0180-5
- Bellec, G., Salaj, D., Subramoney, A., Legenstein, R., and Maass, W. (2018). “Long short-term memory and Learning-to-learn in networks of spiking neurons,” in *Advances in Neural Information Processing Systems 31*, eds. S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (Curran Associates, Inc.), 787–797. Available online at: <http://papers.nips.cc/paper/7359-long-short-term-memory-and-learning-to-learn-in-networks-of-spiking-neurons.pdf>
- Bellman, R., Glicksberg, I., and Gross, O. (1954). The theory of dynamic programming as applied to a smoothing problem. *J. Soc. Indus. Appl. Math.* 2, 82–88. doi: 10.1137/0102007
- Davies, M., Srinivasa, N., Lin, T.-H., China, G., Cao, Y., Choday, S. H., et al. (2018). Loihi: a neuromorphic manycore processor with on-chip learning. *IEEE Micro* 38, 82–99. doi: 10.1109/MM.2018.112130359
- Dayan, P. (1992). The convergence of TD(X) for general X. *Mach. Learn.* 8, 341–362. doi: 10.1007/BF00992701
- Dayan, P., and Sejnowski, T. J. (1994). TD(λ) converges with probability 1. *Mach. Learn.* 14, 295–301. doi: 10.1007/BF00993978
- Finn, C., Abbeel, P., and Levine, S. (2017). “Model-agnostic meta-learning for fast adaptation of deep networks,” in *Proceedings of the 34th International Conference on Machine Learning, Proceedings of Machine Learning Research*, Vol. 70, eds. D. Precup and Y. Whye Teh (Sydney, NSW: International Convention Centre), 1126–1135.
- Friedmann, S., Schemmel, J., Grübl, 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
- Friedrich, J., and Lengyel, M. (2016). Goal-directed decision making with spiking neurons. *J. Neurosci.* 36, 1529–1546. doi: 10.1523/JNEUROSCI.2854-15.2016
- Furber, S. (2016). Large-scale neuromorphic computing systems. *J. Neural Eng.* 13:051001. doi: 10.1088/1741-2560/13/5/051001
- Furber, S. B., Galluppi, F., Temple, S., and Plana, L. A. (2014). The SpiNNaker project. *Proc. IEEE* 102, 652–665. doi: 10.1109/JPROC.2014.2304638
- Gewaltig, M.-O., and Diesmann, M. (2007). Nest (neural simulation tool). *Scholarpedia* 2:1430. doi: 10.4249/scholarpedia.1430
- Gittins, A. J. C., and Gittins, J. C. (1979). Bandit processes and dynamic allocation indices. *J. R. Stat. Soc. Ser. B*, 41, 148–177. doi: 10.1111/j.2517-6161.1979.tb01068.x
- Hochreiter, S., Younger, A. S., and Conwell, P. R. (2001). “Learning to learn using gradient descent,” in *ICANN*, volume 2130 of *Lecture Notes in Computer Science* (Springer), 87–94.
- Hutter, F., Hoos, H., and Leyton-Brown, K. (2014). “An efficient approach for assessing hyperparameter importance,” in *Proceedings of the 31st International Conference on Machine Learning*, volume 32 of *Proceedings of Machine Learning Research*, E. P. Xing and T. Jebara (Beijing: PMLR), 754–762.
- 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
- Kirkpatrick, S., Gelatt, C. D., and Vecchi, M. P. (1983). Optimization by simulated annealing. *Science* 220, 671–80. doi: 10.1126/science.220.4598.671
- Markram, H., Meier, K., Lippert, T., Grillner, S., Frackowiak, R., Dehaene, S., et al. (2011). Introducing the human brain project. *Proc. Comput. Sci.* 7, 39–42. doi: 10.1016/j.procs.2011.12.015
- Mead, C. (1990). Neuromorphic electronic systems. *Proc. IEEE* 78, 1629–1636. doi: 10.1109/5.58356
- Nawrocki, R. A., Voyles, R. M., and Shaheen, S. E. (2016). A mini review of neuromorphic architectures and implementations. *IEEE Trans. Electron Devices* 63, 3819–3829. doi: 10.1109/TED.2016.2598413
- Pantazi, A., Woźniak, S., Tuma, T., and Eleftheriou, E. (2016). All-memristive neuromorphic computing with level-tuned neurons. *Nanotechnology* 27:355205. doi: 10.1088/0957-4484/27/35/355205
- Rechenberg, I. (1973). *Evolutionsstrategie: Optimierung Technischer Systeme Nach Prinzipien der Biologischen Evolution*. Stuttgart: Frommann-Holzboog.
- Robert Canini, K., Shashkov, M. M., Griffiths, T. L. (2010). “Modeling transfer learning in human categorization with the hierarchical dirichlet process,” in *Proceedings of the 27th International Conference on Machine Learning (ICML-10)*, eds J. Fürnkranz and T. Joachims (Haifa: Omnipress), 151–158. Available online at: <https://icml.cc/Conferences/2010/papers/180.pdf>
- Rubinstein, R. Y. (1997). Optimization of computer simulation models with rare events. *Eur. J. Oper. Res.* 99, 89–112.
- Salimans, T., Ho, J., Chen, X., and Sutskever, I. (2017). Evolution strategies as a scalable alternative to reinforcement learning. *CoRR* abs:1703.03864
- Schemmel, J., Brüderle, D., Griibl, 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 (IEEE)*, 1947–1950.
- Schuman, C. D., Potok, T. E., Patton, R. M., Birdwell, J. D., Dean, M. E., Rose, G. S., et al. (2017). A survey of neuromorphic computing and neural networks in hardware. *CoRR* abs:1705.06963.
- Subramoney, A., Rao, A., Scherr, F., Bohnstingl, T., Jordan, J., Kopp, N., et al. (2019). IGITUGraz/L2L: v0.4.3.
- Sutton, R. S., and Barto, A. G. (1998). *Reinforcement Learning: An Introduction*. MIT Press.
- Taylor, M. E., and Stone, P. (2009). Transfer learning for reinforcement learning domains: a survey. *J. Mach. Learn. Res.* 10, 1633–1685. doi: 10.1007/978-3-642-01882-4
- Wang, D., and Zheng, T. F. (2015). “Transfer learning for speech and language processing” in *Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA)* (Hong Kong: IEEE), 1225–1237. doi: 10.1109/APSIPA.2015.7415532
- Wang, J. X., Kurth-Nelson, Z., Soyer, H., Leibo, J. Z., Tirumala, D., Munos, R., et al. (2016). “Learning to reinforcement learn,” in *Proceedings of the 39th Annual Meeting of the Cognitive Science Society*, eds. G. Gunzelmann, A. Howes, T. Tenbrink, and E. J. Davelaar (London: cognitivesciencesociety.org). Available online at: <https://mindmodeling.org/cogsci2017/papers/0252/index.html>
- Watkins, C. J. C. H., and Dayan, P. (1992). Q-learning. *Mach. Learn.* 8, 279–292. doi: 10.1023/A:1022676722315
- Wunderlich, T., Kungl, A. F., Müller, E., Hartel, A., Stradmann, Y., Aamir, S. A., et al. (2018). Demonstrating advantages of neuromorphic computation: a pilot study. *arxiv:1811.03618*. doi: 10.3389/fnins.2019.00260

Conflict of Interest Statement: 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 Bohnstingl, Scherr, Pehle, Meier and Maass. 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.