

D9.4.1 CDP5-Biological deep learning: Results for SGA2 Year 1 (M12)

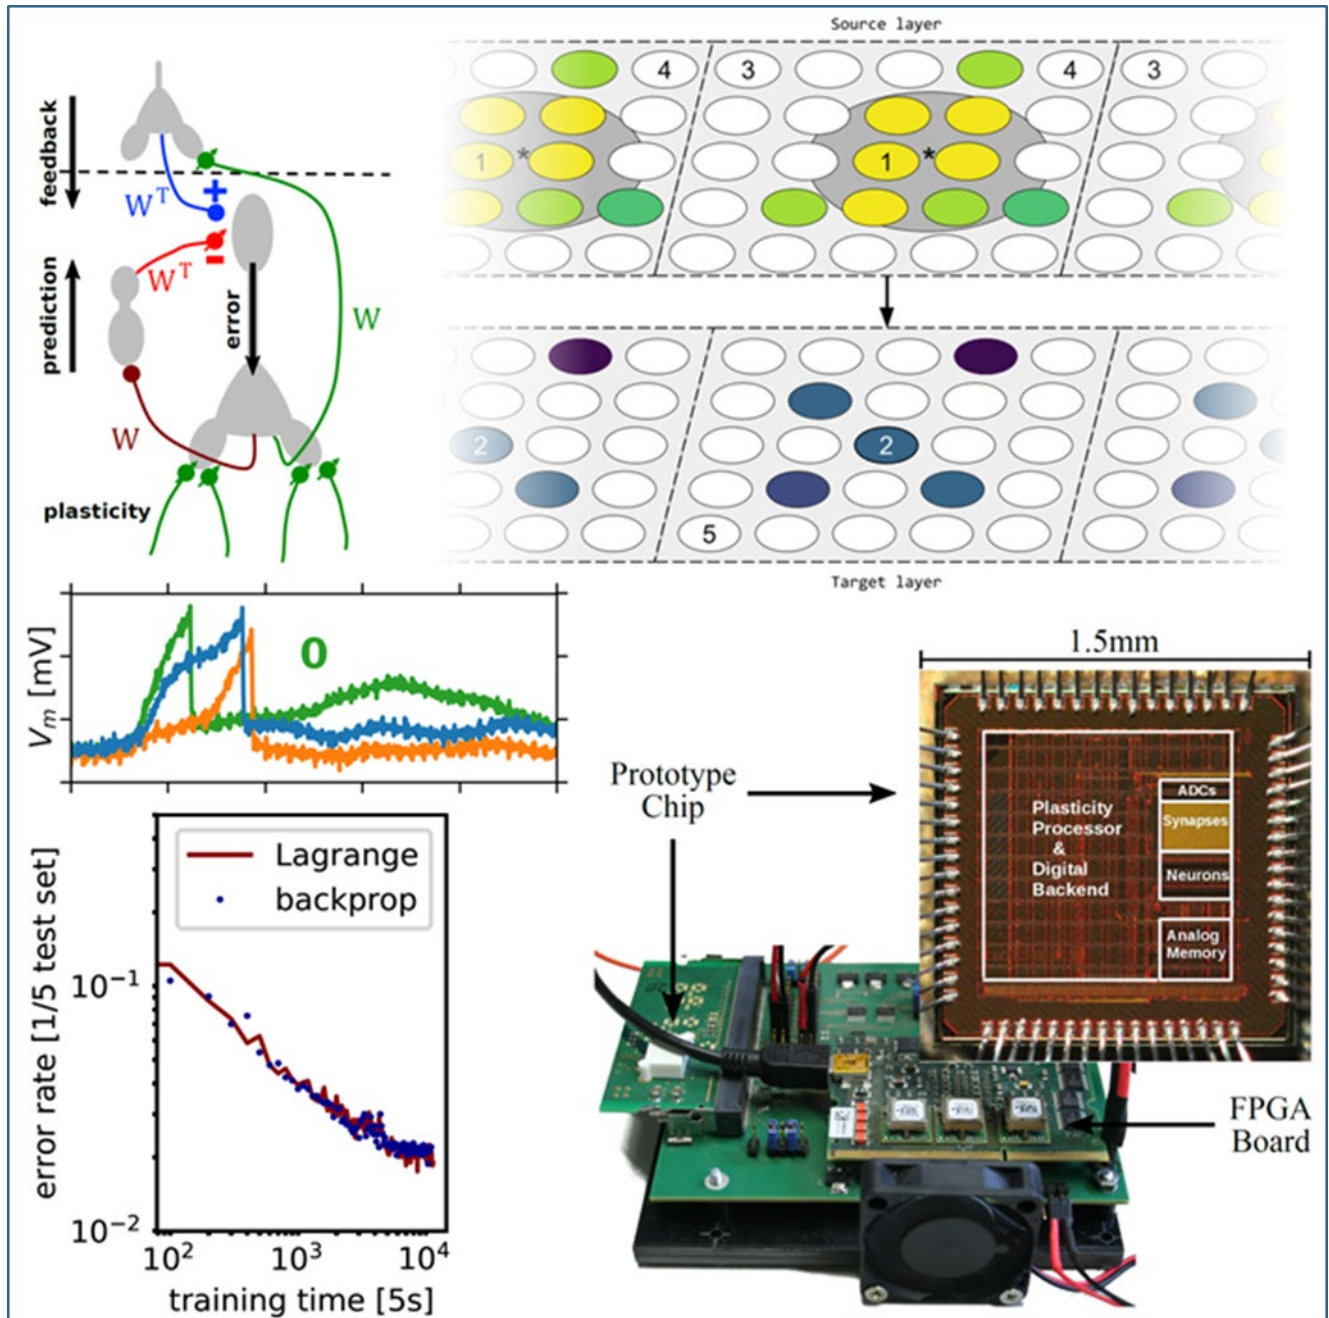


Figure 1: Biological learning for artificial substrates.

CDP5 covers and links a large spectrum of topics, including abstract mathematical models of learning (e.g., Lagrangian neurodynamics, see KRc5.3), network modeling (e.g., hierarchical sampling networks, KRc5.2) and neuromorphic implementation (e.g., structural plasticity on SpiNNaker, KRc5.5).

Project Number:	785907	Project Title:	Human Brain Project SGA2
Document Title:	D9.4.1 CDP5-Biological deep learning: Results for SGA2 Year 1 (M12)		
Document Filename:	D9.4.1 (D60.1 D27) SGA2 M12 ACCEPTED 190723		
Deliverable Number:	SGA2 D9.4.1 (D60.1, D27)		
Deliverable Type:	Report		
Work Package(s):	WP3.2, WP3.3, WP3.5, WP4.3, WP6.4, WP9.1, WP9.2, WP9.3, WP9.4, WP10.2, WP10.3		
Dissemination Level:	PU = Public		
Planned Delivery Date:	SGA2 M12 / 31 Mar 2019		
Actual Delivery Date:	SGA2 M13 / 16 Apr 2019; ACCEPTED 23 Jul 2019		
Authors:	Petrut BOGDAN, UMAN (P63), Oliver BREITWIESER, UHEI (P47), Dominik DOLD, UHEI (P47), Shirin DORA, UvA (P98), André GRÜNING, SURREY (P102), Jakob JORDAN, UBERN (P71), Jacques KAISER, FZI (P52), Elena KREUTZER, UBERN (P71), Pier Stanislao PAOLUCCI, INFN (P52), Mihai PETROVICI, UHEI (P47), Walter SENN, UBERN (P71), Dominik SPICHER, UBERN (P71), J. Camilo VASQUES-TIECK, FZI (P52), Kristin VÖLKL, UBERN (P71), Timo WUNDERLICH, UHEI (P47)		
Compiling Editors:	Mihai PETROVICI, UHEI (P47), Martina SCHMALHOLZ, UHEI (P47)		
Contributors:	<p>Mihai PETROVICI, UHEI (P47), contributed to each KRc, chapter and section</p> <p>Akos KUNGL, UHEI (P47), KRc5.1, KRc5.3, Ch. 3, 5.2; Eric MÜLLER, UHEI (P47), KRc5.1, Ch. 3, KRc5.3, Ch. 5.6, KRc5.5, Ch. 7.1; Johannes SCHEMMEL, UHEI (P47), KRc5.1, Ch. 3, KRc5.3, Ch. 5.5, 5.6, KRc5.4, Ch. 6.2; Timo WUNDERLICH, UHEI (P47), KRc5.1, Ch. 3; Dominik DOLD, UHEI (P47), KRc5.2, KRc5.3, Ch. 4.1, 5.2; Maximilian ZENK, UHEI (P47), KRc5.2, Ch. 4.1; Kristin VÖLKL, UBERN (P71), KRc5.2, Ch. 4.2; Walter SENN, UBERN (P71), KRc5.2, KRc5.3, Ch. 4.2, 5.1, 5.2, 5.4, 6.1; Elena KREUTZER, UBERN (P71), KRc5.3, Ch. 5.1; João SACRAMENTO, UBERN (P71), KRc5.3, Ch. 5.2, 5.3; Jakob JORDAN, UBERN (P71), KRc5.3, Ch. 5.3; Dominik SPICHER, UBERN (P71), KRc5.3, Ch. 5.4; Claudia CLOPATH, Imperial College London (external to HBP), KRc5.3, Ch. 5.4; Julian GÖLTZ, UHEI (P47), KRc5.3, Ch. 5.5; Oliver BREITWIESER, UHEI (P47), KRc5.3, Ch. 5.5; Sebastian SCHMITT, UHEI (P47), KRc5.3, Ch. 5.5; Brian GARDNER, SURREY (P102), KRc5.3, Ch. 5.6; Joseph CHROL-CANNON, SURREY (P102), KRc5.3, Ch. 5.6; André GRÜNING, SURREY (P102), KRc5.3, Ch. 5.6; Steve FURBER, UMAN, (P63), KRc5.3, Ch. 5.6, KRc5.5, Ch. 7.1; Shirin DORA, UvA, (P98), KRc5.4, Ch. 6.1; Sander BOHTE, UvA, (P98), KRc5.4, Ch. 6.1; Martin PEARSON, UWE, (P101), KRc5.4, Ch. 6.1; Cyriel PENNARTZ, UvA, (P98), KRc5.4, Ch. 6.1; Cristiano CAPONE, INFN (P92), KRc5.4, Ch. 6.2; Elena PASTORELLI, INFN (P92), KRc5.4, Ch. 6.2; Bruno GOLOSIO, INFN (P92), KRc5.4, Ch. 6.2; Pier Stanislao PAOLUCCI, INFN (P92), KRc5.4, Ch. 6.2; Maurizio MATTIA, ISS (P96), KRc5.4, Ch. 6.2; Sandra DIAZ-PIER, JUELICH (P20), KRc5.5, Ch. 7.1; Wouter KLIJN, JUELICH (P20), KRc5.5, Ch. 7.1; Petrut BOGDAN, UMAN (P63), KRc5.5, Ch. 7.1; Michel HOPKINS, UMAN (P63), KRc5.5, Ch. 7.1; Garibaldi PINEDA-GARCIA, SURREY (P102), KRc5.5, Ch. 7.1; Andrew ROWLEY, UMAN (P63), KRc5.5, Ch. 7.1; Oliver RHODES, UMAN (P63), KRc5.5, Ch. 7.1; Simon DAVIDSON, UMAN (P63), KRc5.5, Ch. 7.1; Robert JAMES, UMAN (P63), KRc5.5, Ch. 7.1; Christian PEHLE, UHEI (P47), KRc5.5, Ch. 7.1; Sebastian BILLAUDELLE, UHEI (P47), KRc5.5, Ch. 7.1; Benjamin CRAMER, UHEI (P47), KRc5.5, Ch. 7.1; Andrew DAVISON, CNRS (P10), KRc5.5, Ch. 7.1; Jacques KAISER, FZI (P52), KRc5.6, Ch. 8.1, 8.2; J. Camilo VASQUES-TIECK, FZI (P52), KRc5.6, Ch. 8.1, 8.2; Daniel REICHARD, FZI (P52), KRc5.6, Ch. 8.1, 8.2; Guillaume BELLEC, TUGRAZ (P55), KRc5.6, Ch. 8.1; Oliver RHODES, UMAN (P63), KRc5.6, Ch. 8.1; Emre NEFTCI</p>		

	(University of California, Irvine, UC, Department of Cognitive Sciences, external collaboration), KRc5.6, Ch. 8.1; Rüdiger DILLMANN, FZI (P52), KRc5.6, Ch. 8.1, 8.2; David KAPPEL, TUGRAZ (P55), KRc5.6, Ch. 8.2; Anand SUBRAMONEY, TUGRAZ (P55), KRc5.6, Ch. 8.2; Robert LEGENSTEIN, TUGRAZ (P55), KRc5.6, Ch. 8.2; Wolfgang MAASS, TUGRAZ (P55), KRc5.6, Ch. 8.2
SciTechCoord Review:	
Editorial Review:	
Description in GA:	CDP5 Annual compound deliverable Year 1. Overview of Key Results and Impact achieved in M1 to M12, tailored for presentation to the relevant audiences (research/industry/public) References to HBP/SP/CDP objectives and Use Cases for navigation between multiple Sub-Projects Linkage of Results to components/a set of component fact sheets (lower level information: Component ownership, Technology Readiness level, performed Quality Control checks, etc.). This deliverable will be a public report including theoretical research on learning principles, their transfer to the HBP computing infrastructures in SP7 and SP9, results of initial experiments and plans for M13 to M24 in SGA2.
Abstract:	In CDP5, HBP researchers from several SPs collaborate on the topic of biological and artificial intelligence. A key requirement for being classified as “intelligent” is a system’s ability to learn. In the course of the last year, important insights have been gained into the computational principles underlying biological learning and their application to the neuromorphic and neurobotic platforms developed in the HBP.
Keywords:	Learning, plasticity, (deep) hierarchical networks, spiking neurons, neuromorphic computation, neurorobotics
Target Users/Readers:	Scientists, Companies and other potential users of HBP results, Reviewers.

Table of Contents

1. Introduction	6
1.1 Summary of Key Results and Outputs	6
2. Key Result KRc5.1: Closed-loop navigation	7
2.1 Output 1: Adaptive closed-loop control in a virtual environment using local reinforcement learning on a BrainScaleS-2 prototype	7
3. Key Result KRc5.2: Deep spatio-temporal prediction	9
3.1 Output 1: Spatio-Temporal Predictions with Spiking Neural Networks.....	9
3.2 Output 2: Sequence learning by shaping hidden connectivity	12
4. Key Result KRc5.3: Deep learning with compartmental neurons	14
4.1 Output 1: Natural gradient for spiking neurons	14
4.2 Output 2: Lagrangian neurodynamics for real-time error-backpropagation across cortical areas ..	17
4.3 Output 3: Error-driven learning supports Bayes-optimal multisensory integration via conductance-based dendrites.....	20
4.4 Output 4: Robust predictive plasticity in dendrites: From a computational principle to experimental data ..	21
4.5 Output 5: Training deep networks with time-to-first-spike coding on the BrainScaleS wafer-scale system ..	23
4.6 Output 6: Spike-based learning on neuromorphic substrates	24
5. Key Result KRc5.4: Deep episodic learning	26
5.1 Output 1: Reconstruction of stimuli across multiple modalities.....	26
5.2 Output 2: Interaction between sleep and memory in a thalamo-cortical model performing visual classification (MNIST).....	28
6. Key Result KRc5.5: Learning to learn with reconfigurable networks	30
6.1 Output 1: Structural plasticity in the form of synaptic rewiring on the SpiNNaker digital neuromorphic system.....	31
7. Key Result KRc5.6: Deep closed-loop sensorimotor learning	32
7.1 Output 1: Real-world embodiment for neuromorphic learning research	32
7.2 Output 2: Embodied Synaptic Plasticity with Online REinforcement (SPORE) learning	35
8. Conclusion and Outlook	38
Annex A: Component Details	40

Table of Figures

Figure 1: Biological learning for artificial substrates.	1
Figure 2: Comparing the neuromorphic emulation on the prototype chip to a software simulation using the NEST simulator.....	8
Figure 3: Set-up of spatio-temporal predictions with spiking neural networks.	10
Figure 4: Effect of innervation on network performance.	12
Figure 5: Effects of different modifications to the network model.	13
Figure 6: Heterosynaptic plasticity counterbalances homosynaptic plasticity.	15
Figure 7: Lagrangian dynamics of cortical microcircuits.	18
Figure 8: Learning multisensory integration of visual and tactile information in conductance-based neurons with dendrites.	20
Figure 9: Predictive plasticity in dendrites.	22
Figure 10: Learning to classify reduced 7x7 MNIST digits with single spikes on the BrainScaleS-1 system...23	23
Figure 11: Training and test loss with the number of training epochs, using three-fold cross validation on Iris and Wisconsin (averaged over 20 runs).	25
Figure 12: Histogram of kurtosis for neurons in different layers of the model.	27
Figure 13: Effects of deep-sleep like Slow Oscillation on synaptic structure in a thalamo-cortical model. ..	29
Figure 14: Structural plasticity on SpiNNaker.	31
Figure 15: Our robotic setup embodying the event-driven random backpropagation synaptic learning rule.33	33
Figure 16: Accuracy on the DvsGestures dataset.	34
Figure 17: Our asynchronous framework based on open-source software components to evaluate SPORE. ..	36
Figure 18: Visualization of the setup for the two experiments.	36

Figure 19: Results for the reaching task.	36
Figure 20: Results for the lane following task with a medium prior ($cp = 0.25$).	37

1. Introduction

Whenever we learn a new skill, something changes in the physical structure of our brain. These structural changes, generally described as “plasticity”, enable the brain to acquire new functionality. It is this ability of biological neural systems to learn a huge variety of possible skills that remains unparalleled in the realm of AI.

Under the aegis of CDP5, a multidisciplinary team of HBP scientists are working to bridge the gap between learning in biological and artificial systems. Building on the current state of the art in experimental, theoretical and computational neuroscience, machine learning and neuro-inspired silicon substrates, we strive to answer two questions: What are the fundamental computational principles behind the brain’s ability to learn? And which properties of biological neural networks can we transfer to artificial systems in order to improve their learning capabilities?

Over the last year, our research in CDP5 has yielded a number of important results. We have demonstrated how purely local plasticity rules can endow hierarchical (deep) networks with the ability to learn complex, multimodal spatiotemporal signals. We have developed a theoretical framework for error backpropagation using only local information in time-continuous systems. Aspects of these models have been linked to experimental data. Models of learning using spikes have been implemented on the SpiNNaker and BrainScaleS platforms, harnessing their capabilities in terms of speed and energy efficiency. Significant progress has also been made on the development and application of spiking neural networks for robotic control.

With accumulating insight into synaptic plasticity, the gap between biological and artificial intelligence is constantly narrowing. Through CDP5, the HBP continues to occupy a unique position in this rapidly developing line of research.

1.1 Summary of Key Results and Outputs

- KRc5.1: Closed-loop navigation
 - Output 1: Adaptive closed-loop control in a virtual environment using local reinforcement learning on a BrainScaleS-2 prototype
- KRc5.2: Deep spatiotemporal prediction
 - Output 1: Spatio-temporal predictions with spiking neural networks
 - Output 2: Sequence learning by shaping hidden connectivity
- KRc5.3 Deep learning with compartmental neurons
 - Output 1: Natural gradient for spiking neurons
 - Output 2: Lagrangian neurodynamics for real-time error-backpropagation across cortical areas
 - Output 3: Error-driven learning supports Bayes-optimal multisensory integration via conductance-based dendrites
 - Output 4: Robust predictive plasticity in dendrites: From a computational principle to experimental data
 - Output 5: Training deep networks with time-to-first-spike coding on the BrainScaleS wafer-scale system
 - Output 6: Spike-based learning on neuromorphic substrates
- Key Result KRc5.4 Deep episodic learning
 - Output 1: Reconstruction of stimuli across multiple modalities

- Output 2: Interaction between sleep and memory in a thalamo-cortical model performing visual classification
- Key Result KRc5.5 Learning to learn with reconfigurable networks
 - Output 1: Structural plasticity in the form of synaptic rewiring on the SpiNNaker digital neuromorphic system.
- Key Result KRc5.6 Deep closed-loop sensorimotor learning
 - Output 1: Real-world embodiment for neuromorphic learning research
 - Output 2: Embodied Synaptic Plasticity with Online REinforcement (SPORE) learning

2. Key Result KRc5.1: Closed-loop navigation

Overview of Outputs

Output 1: Adaptive closed-loop control in a virtual environment using local reinforcement learning on a BrainScaleS-2 prototype

2.1 Output 1: Adaptive closed-loop control in a virtual environment using local reinforcement learning on a BrainScaleS-2 prototype

Timo WUNDERLICH, Akos KUNGL, Eric MÜLLER, Johannes SCHEMMEL, Mihai A. PETROVICI

CDP5 collaboration between SP9 (UHEI, P47) and SP4 (UBERN, P71)

For the first time, we demonstrated that a prototype of the BrainScaleS-2 neuromorphic system can be used to implement neuromodulated plasticity to solve a closed-loop learning task in a virtual environment. Our experiments show that the BrainScaleS hardware can be used to evaluate learning processes in an accelerated and energy-efficient fashion.

2.1.1 Research and Infrastructure

Neuromodulation of synaptic plasticity is known to be a key process in biological learning that, for example, mediates reward-based learning using the neuromodulator dopamine. It is therefore highly desirable for neuromorphic hardware to support plasticity rules that model neuromodulation. The BrainScaleS-2 neuromorphic platform will encompass embedded plasticity processors on each chip to allow researchers to implement a wide range of plasticity rules, including neuromodulated plasticity. Using a prototype chip that contains such a plasticity processor, we implemented a form of reinforcement learning using reward-modulated spike-timing dependent plasticity (R-STDP). The R-STDP learning rule can be viewed as a model of dopaminergic neuromodulation and in our case, allowed the agent embodied by the neural network to solve the commonly known Pong video game. The virtual environment that represents the Pong game was simulated by the plasticity processor, allowing us to implement the experiment fully on-chip, as all necessary observables (neuronal firing rates, synaptic eligibility traces) can be accessed from the processor.

By conservatively comparing our neuromorphic emulation on the prototype chip to a software simulation on an off-the-shelf processor using the NEST simulator, we are able to conclude that our emulation is at least an order of magnitude faster and three orders of magnitude more energy-efficient (Fig. 2). These advantages in terms of speed and energy efficiency will become more apparent in large neural networks on the full-scale BrainScaleS-2 system, as the speed-up of the neuromorphic emulation is constant while digital simulations are faced with slower simulation speeds

as the network size is scaled up. Besides this, the experiment shows that trial-to-trial variability of neuronal firing rates (due to the analog nature of the neurons) can be used as a computational resource, rather than being a nuisance, as it mediates action exploration and trial-and-error learning. Our results also demonstrate that fixed-pattern noise of neuron parameters, which does not change over time, can be implicitly compensated by the learning process, decreasing the need for calibration.

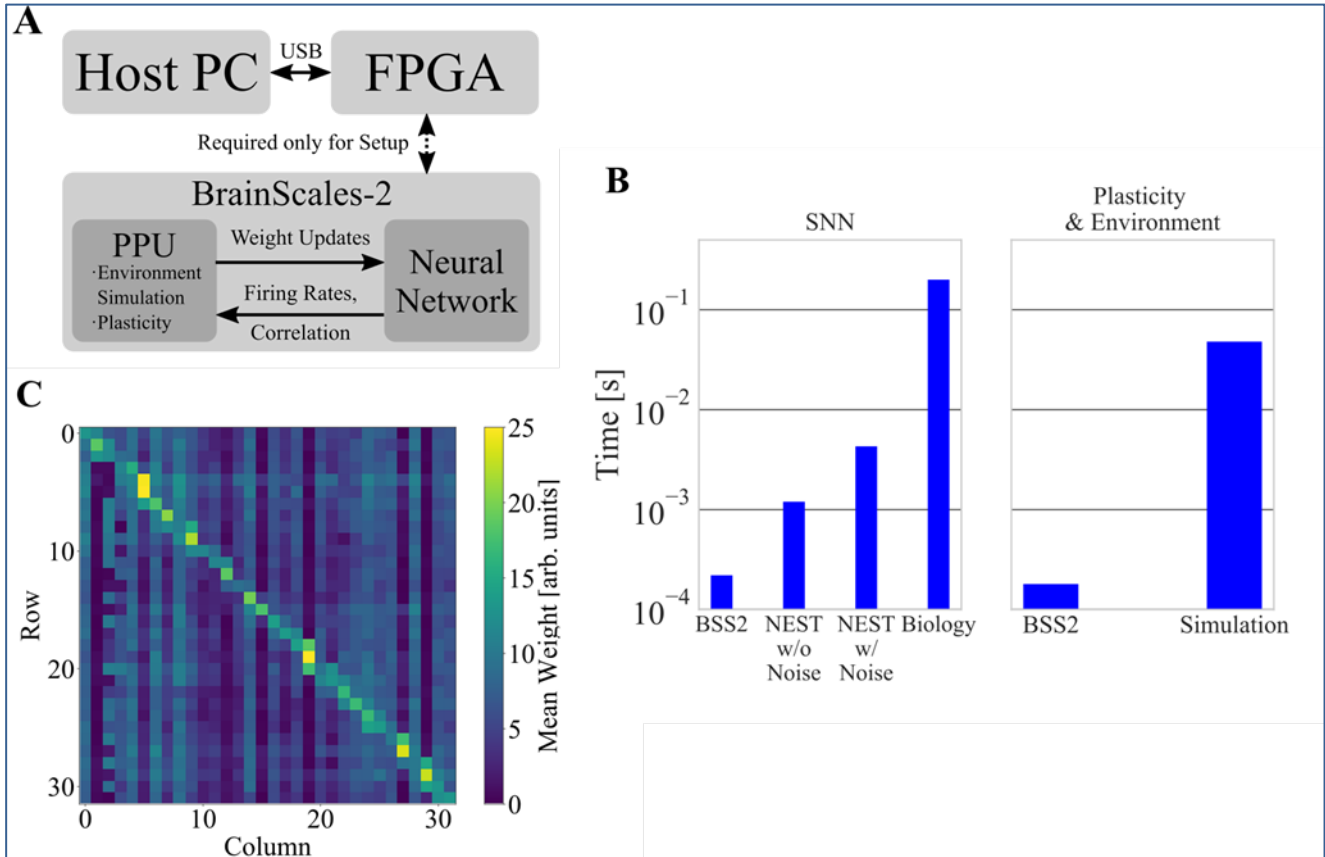


Figure 2: Comparing the neuromorphic emulation on the prototype chip to a software simulation using the NEST simulator

(A): The experiment takes place fully on-chip, the external FPGA is used only for configuration. The plasticity processing unit (PPU) simulates the environment (Pong) and calculates weight updates using a reward signal and neural network observables. (B): A software simulation on a digital processor is at least one order of magnitude slower than on BrainScaleS-2 (BSS2), when comparing the time required for spiking neural network (SNN) emulation/simulation only. (C): Weight matrix on the BrainScaleS-2 prototype after learning. The matrix is diagonally dominant which is the goal of the learning task. The vertical structure is due to the compensation of neuronal parameter variations.

2.1.2 Impact

In order to achieve these results, extensive work on the software environment of the BrainScaleS2 prototype was necessary. These developments benefit all present and future users of the hardware.

Within the greater neuromorphic community, these results have attracted considerable attention. The combination of an analog substrate with an embedded plasticity processor is unique and this was the first demonstration of its capabilities. The achieved combination of acceleration and energy efficiency represents an important point of reference on the neuromorphic map and thereby a key outcome of this line of research in the HBP.

2.1.3 Components

C209 NEST - The Neural Simulation Tool

C457 BrainScaleS 2 Neuromorphic Computing System

C2420 Prototype implementations of rules and testing within and without the SP9 platforms

2.1.4 Publications

- 1) Demonstrating Advantages of Neuromorphic Computation. Timo Wunderlich, Akos F. Kungl, Eric Müller, Andreas Hartel, Yannik Stradmann, Syed A. Aamir, Andreas Grübl, Arthur Heimbrecht, Korbinian Schreiber, David Stöckel, Christian Pehle, Sebastian Billaudelle, Gerd Kiene, Christian Mauch, Johannes Schemmel, Karlheinz Meier, Mihai A. Petrovici, *Frontiers In Neuromorphic Engineering*, 2019

2.1.5 Further measures for increasing impact

- 1) Form versus function: Theory and models for neuronal substrates. Mihai A. Petrovici, Talk at the Symposium on the Theory of Neuromorphic Computing, Heidelberg, December 13-14, 2018
- 2) Live Reinforcement Learning in an Accelerated Neuromorphic Spiking Neural Network, Video at <https://youtu.be/LW0Y5SSIQU4>
- 3) Demonstration at the Structures Review in Köln, June 26, 2018
- 4) Demonstration at the HBP Open Day in Maastricht, October 15, 2018
- 5) Demonstration at the ICT2018 in Wien, December 4-6, 2018
- 6) 6. Demonstration at the NICE (Neuro-Inspired Computational Elements) Workshop, Albany, March 26-29, 2019

3. Key Result KRc5.2: Deep spatio-temporal prediction

Overview of Outputs

- Output 1: Spatio-Temporal Predictions with Spiking Neural Networks
- Output 2: Sequence learning by shaping hidden connectivity

3.1 Output 1: Spatio-Temporal Predictions with Spiking Neural Networks

Maximilian Zenk, Dominik Dold, Mihai A. Petrovici

CDP5 collaboration between SP9 (UHEI, P47) and SP4 (UBERN, P71)

In previous work, it has been shown that stochastic spiking neural networks sample from posterior probability distributions parametrized by the network's synaptic connections. So far, this has been used for static inputs only, e.g., to perform classification after training such networks to approximate some underlying data distribution. In this work, we extended the approach to time-continuous problems, allowing the network to act and react in a time-dependent environment. The presented model is portable to the neuromorphic hardware currently developed in the HBP.

3.1.1 Research and Infrastructure

Building on their ability to perform pattern completion, we used stochastic spiking neural networks to solve the spatiotemporal problem of predicting the trajectory of a ball in a two-dimensional potential landscape (Fig. 3A). Since the ball's trajectory unfolds in time, the network can only predict future ball positions based on already available information, i.e., past ball positions. As a measure of performance, we let the network take control of a paddle (with finite velocity) to catch the ball in real-time, allowing the network to interact with its environment. To test the network's capabilities and limitations, we investigated its performance in software simulations for two different potentials (flat and with a Gaussian hill in the centre) and a scenario introducing additional random deflections of the ball, representing unpredictable external factors such as wind.

To obtain good performance on the spatiotemporal prediction task, we found that both long-term plasticity as well as short-term plasticity are indispensable. Long-term contrastive Hebbian plasticity was used to let the network learn important features of trajectories in two-dimensional potential landscapes (Fig. 3B). After initial training, the network was tested by letting it predict the complete trajectory of a rolling ball. We found that initially, when the ball has barely started rolling and only a small portion of its trajectory is known, the network becomes stuck in whatever initial trajectory prediction it happens to be residing in at that time, and is only able to leave this local mode once the ball has rolled forward along a significant portion of its trajectory. To shorten the duration of this initial phase, we utilized short-term plasticity, which promotes a stochastic, sampling-based exploration of different possible target trajectories by slowly pushing the network out of local attractors. This allowed the network to quickly access different solutions that are compatible with the current observations before making a final decision (Fig. 3C,D). This idea builds on previous work in CDP5 (Leng et al., Nat. Sci. Rep. 8 (1), 2018), which we reported in M24 of SGA1.

The presented model allows stochastic spiking neural networks to be used in agent-based tasks, like playing games (e.g., Pong), where the network's model of the environment has to be used to perform predictions not only over space but also in time.

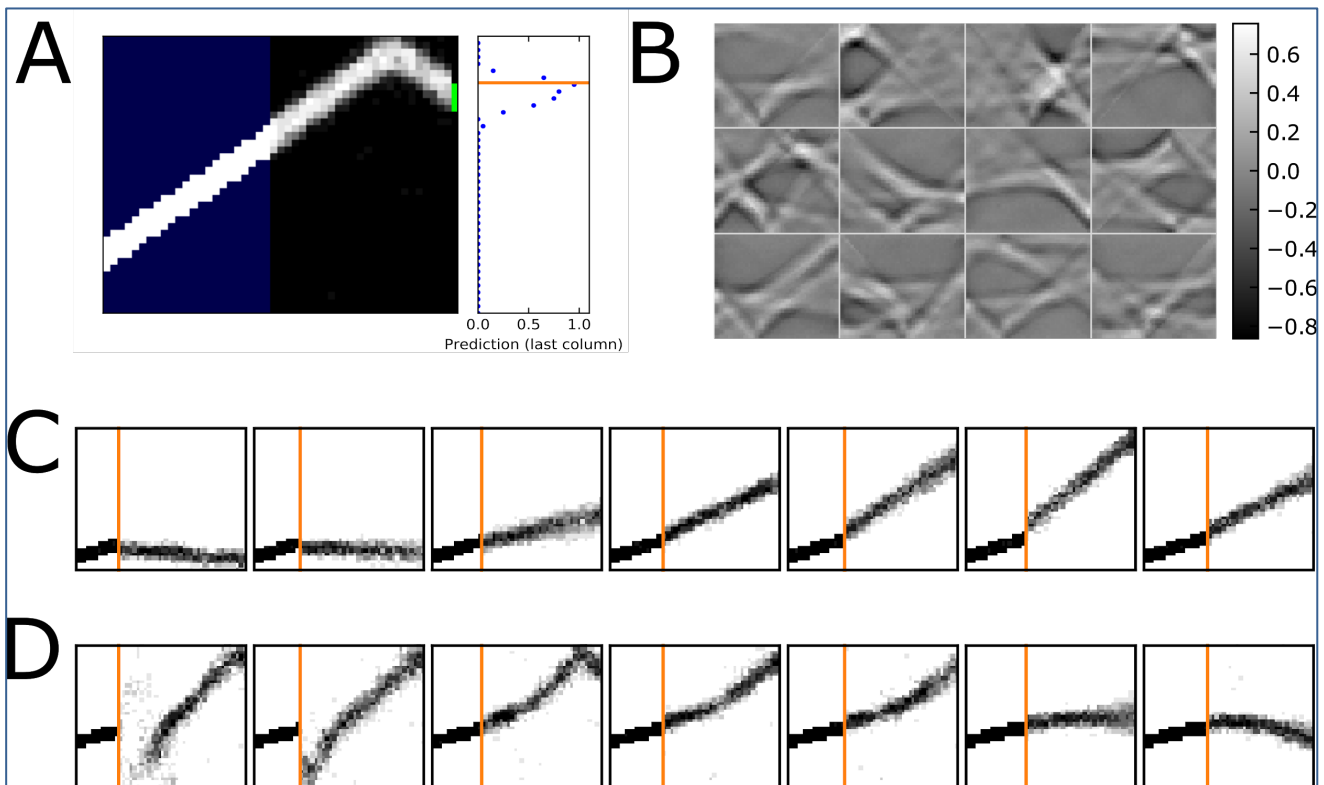


Figure 3: Set-up of spatio-temporal predictions with spiking neural networks.

(A): Illustration of the setup. The network receives a slowly unfolding trajectory as input (blue background) and predicts the remaining trajectory based on this information (black background). Furthermore, the network can control a paddle (green), which moves toward the position of the most likely predicted trajectory endpoint

(histogram, orange line). (B): Receptive fields learned by the network. (C): While the ball is moving forward (from left to right), short-term plasticity allows the network to effectively search the space of compatible solutions (predicted trajectories are right of the moving orange line, network input is to the left). This is shown here for the case of a flat potential (i.e., straight trajectories). (D): Same as (C), but for a ball traversing a Gaussian potential, leading to curved trajectories.

3.1.2 *Impact*

This work bridges the gap between the original models of spike-based inference on static inputs (see CDP5 SGA1 M24 Deliverable) and the need for Bayesian predictions in a dynamically changing world. Furthermore, it makes specific use of spike-triggered short-term plasticity to enable exploration/exploitation; this represents a unique computational perspective on the persisting question of why biology uses spikes, which was met with manifest interest by the computational neuroscience community (see Dissemination).

3.1.3 *Components*

C209 NEST - The Neural Simulation Tool

C349 PyNN

C2439 Methods for hierarchical neural sampling in networks of spiking neurons

3.1.4 *Publications*

- 1) Spatio-Temporal Predictions with Spiking Neural Networks. Maximilian Zenk, Master's Thesis, University of Heidelberg, 2018

3.1.5 *Further measures for increasing impact*

- 1) Spiking neuron ensembles and probabilistic inference. Mihai A. Petrovici, Talk at the "Beyond digital computing" conference, Heidelberg, March 18-21, 2018
- 2) Perspektiven. Mihai A. Petrovici, Talk at the Manfred Stärk Symposium, Bern, May 2-3, 2018
- 3) Form versus function: Theory and models for neuronal substrates. Mihai A. Petrovici, Talk at the Symposium on the Theory of Neuromorphic Computing, Heidelberg, December 13-14, 2018
- 4) Computers like brains. Mihai A. Petrovici, Talk at the "At the crossroad of physics and machine learning" conference, Santa Barbara, February 11-15, 2019
- 5) <http://online.kitp.ucsb.edu/online/machine-c19/petrovici/>
- 6) Why spikes? Signal detection, sampling, timing. Walter Senn, Talk at the Cosyne conference, Lisbon, February 28 - March 5, 2019
- 7) Mihai A. Petrovici, Interview with Euronews, Bern, March 11, 2019
- 8) <https://www.euronews.com/2019/04/01/the-human-brain-project-slicing-brains-and-thinking-machines>
<https://www.euronews.com/2019/04/01/can-we-make-a-computer-like-the-human-brain>
- 9) Computing with physics: From biological to artificial intelligence and back again. Mihai A. Petrovici, Talk at the Neuro-Inspired Computational Elements conference, Albany, March 26-29, 2019

3.2 Output 2: Sequence learning by shaping hidden connectivity

Kristin Völk, Mihai A. Petrovici, Walter Senn

CDP5 collaboration between SP4 (UBERN, P71) and SP9 (UHEI, P47)

A cortical developmental model is suggested for learning spatio-temporal patterns based on 2-compartment neuron models and dendritic plasticity. The model shapes an appropriate connectivity pattern in a pool of hidden neurons that allows the memorization of non-Markovian sequences in visible neurons. The model is portable to the neuromorphic hardware currently developed in the HBP.

3.2.1 Research and Infrastructure

In the CDP5 SGA1 M24 Deliverable, we showed how a sparse scaffold of somato-somatic connections, formed during network development, can guide the learning of dense somato-dendritic connections. This enabled learning of non-Markovian patterns. In our continued work on this model, we want to emphasize the benefits of having learning in the somato-dendritic connections compared to only utilizing the sparse somato-somatic connections.

While the somato-somatic connections already provide a blueprint of general structures (delay lines, loops etc.), the learning of the somato-dendritic connections recombines these into new structures that are better suited for the sequence to be learned. On one hand, this enables the hidden pool to form longer memories and hence learn more complex sequences; on the other hand, it also increases the robustness of the network to pattern disruptions. Figure 5(b) shows how we can “delete” the letter ‘c’ from the sequence, by clamping the associated neurons to the resting potential, without disrupting the correct continuation of the remaining sequence. In the case of having only somato-somatic connections and no somato-dendritic learning (Fig 5(a)), this cannot be done. This is because the shorter delay lines of the scaffold have not been recombined into longer ones during learning.

Furthermore, we can even delete a whole pattern (b,c,b,f) (Fig 5(c)) and ensure correct replay of the following sequence. Interestingly, the network activity following a deletion of a whole pattern is most stable when we have no somato-dendritic connections from the visible to the hidden pool (Fig. 4(a)). Not having these connections makes the hidden pool more robust against disruptions of the visible pool. However, it also increases learning time (Fig. 4(b)). These results illustrate how the somato-somatic connections facilitate the learning of appropriate memory structures, whereas somato-dendritic learning utilizes and adapts these structures to the learning problem.

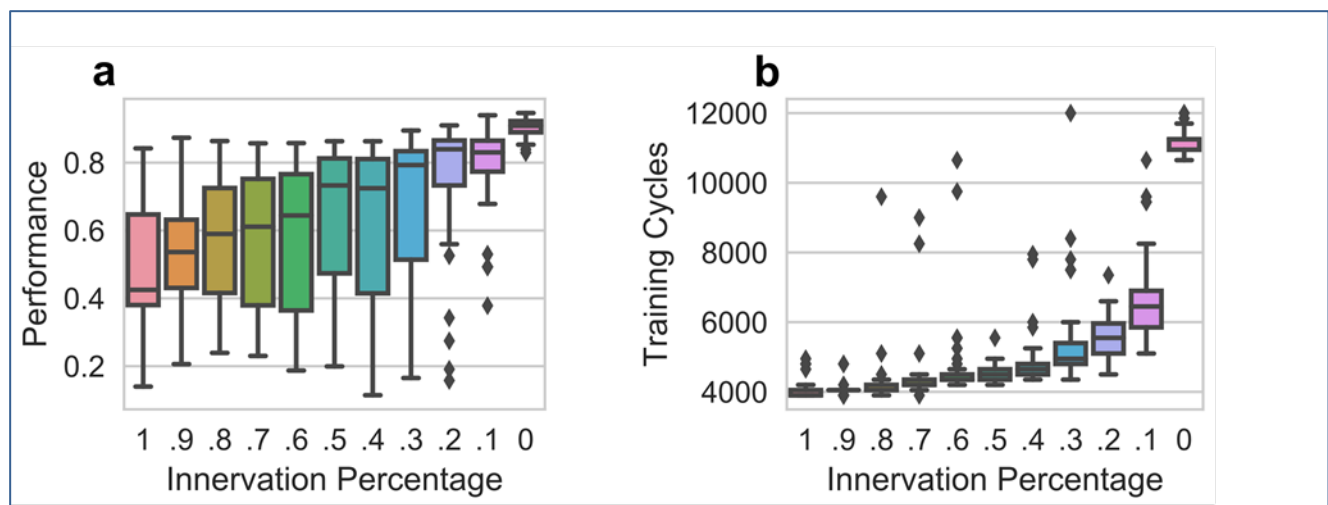


Figure 4: Effect of innervation on network performance.

Key: (a) Correlation performance for networks trained with a different percentage (from 1.0 to 0.0) of somato-dendritic visible to hidden innervations. For each percentage 50 networks are trained, visible to hidden connections are chosen randomly. (b) Number of training cycles needed to train the networks shown in (a). Again the results are shown for the different innervation percentages.

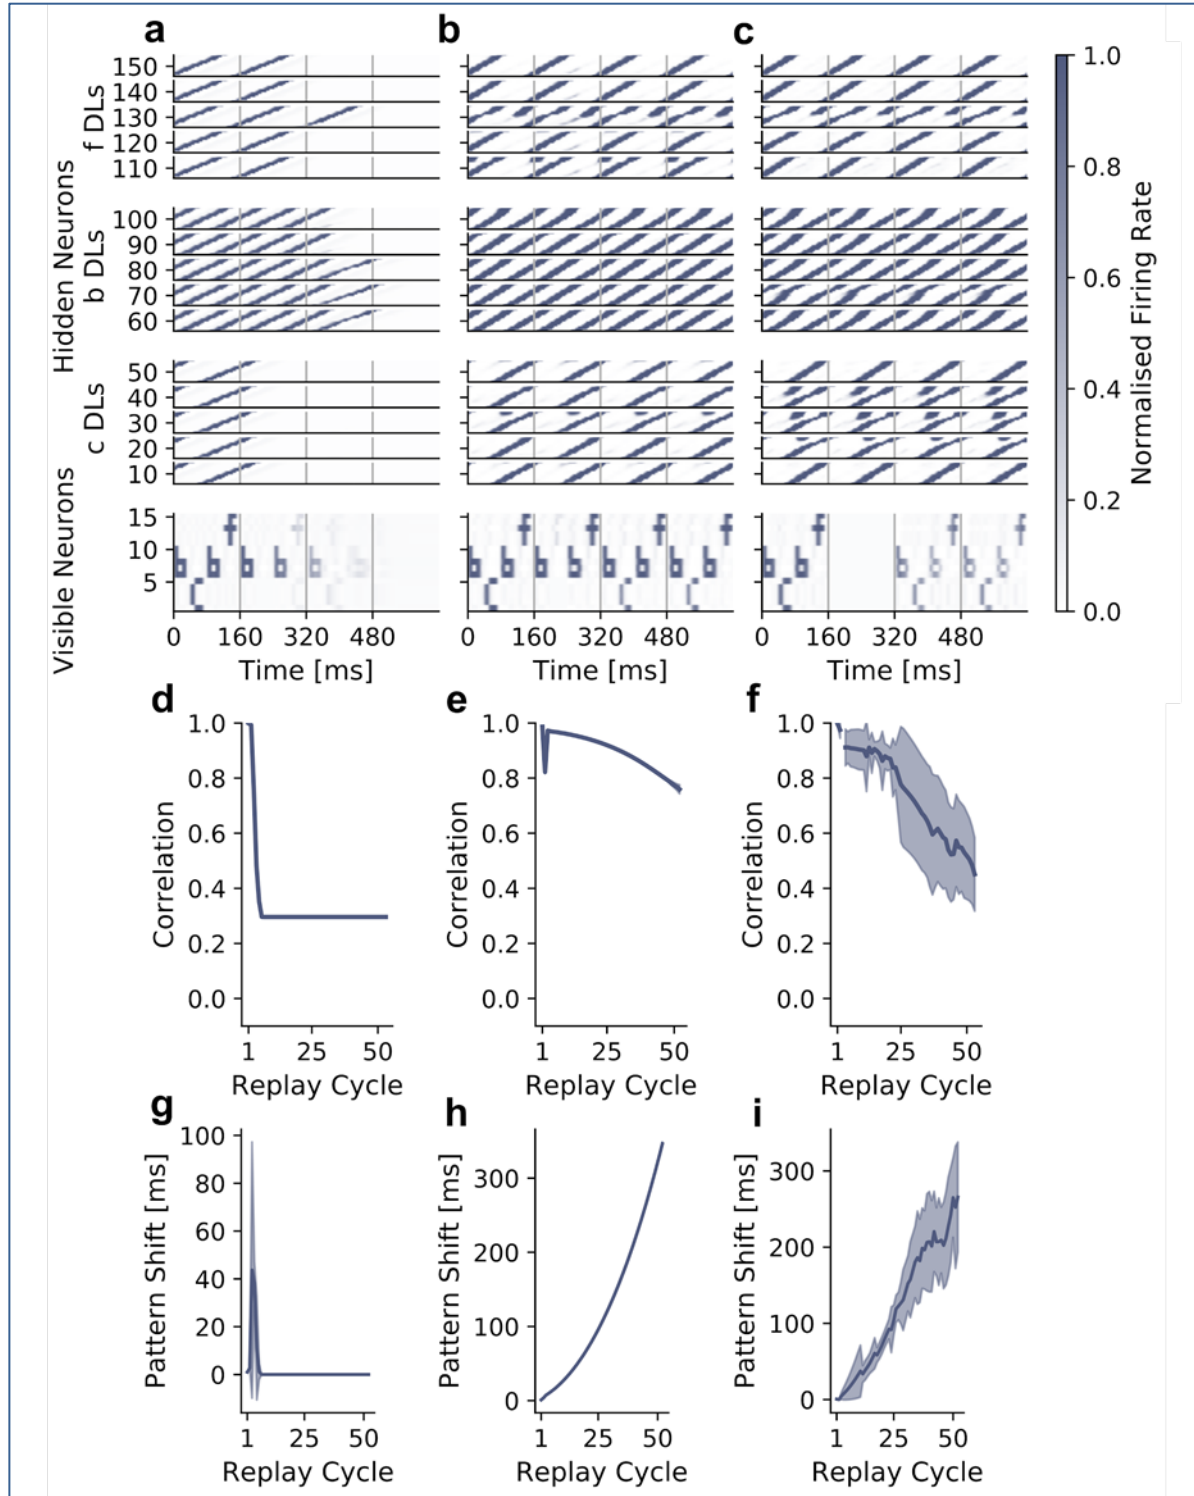


Figure 5: Effects of different modifications to the network model.

Key: (a-b) Raster Plot of simulated network: first replay cycle (after training), followed by the altered (c or whole pattern deletion) cycle and two consecutive pattern replays. DLs stands for delay lines and shows the delay lines present in the somato-somatic connections for each letter/neuron.

(a) Network without dendritic learning in the hidden population. Membrane potential of hidden neurons only governed by somatic inputs. (b) Network with full dendritic learning (visible and hidden layer). (c) Network without somato-dendritic visible to hidden connections. (d-f) Correlations between taught pattern and actually produced pattern

during replay phase for the network shown above. (g-i) Amount by which each pattern onset in the replay cycle is shifted from the ideal onset (this is cumulative) for the network shown above.

3.2.2 *Impact*

This work addresses the interplay between stereotypical (genetically imprinted) structure and learning in a reservoir computing scenario. The contrast between absence and existence of error-driven learning in different somato-dendritic projections (hidden-hidden and visible-hidden) within the network at the behavioral level is of general interest for research and applications of reservoir computing.

3.2.3 *Components*

C2419 Plasticity: algorithms for multi-compartment models

C2420 Plasticity: prototype implementations of rules and testing within and without the SP9 platforms

C2547 Computing with structured neurons

3.2.4 *Publications*

Journal paper in preparation.

3.2.5 *Further measures for increasing impact*

- 1) Shaping a Hidden Neuronal Pool for Time-Series Learning. Kristin Völk, Talk at the Fürberg Workshop, Fürberg, June 26-29, 2018
- 2) Time-series Learning through Hidden Population Shaping by Somatic Nudging. Kristin Völk, Talk at the Heidelberg-Bern Group Meeting, Triberg, July 18-20, 2018

4. Key Result KRc5.3: Deep learning with compartmental neurons

Overview of Outputs

Output 1: Natural gradient for spiking neurons

Output 2: Lagrangian neurodynamics for real-time error-backpropagation across cortical areas

Output 3: Error-driven learning supports Bayes-optimal multisensory integration via conductance-based dendrites

Output 4: Robust predictive plasticity in dendrites: From a computational principle to experimental data

Output 5: Training deep networks with time-to-first-spike coding on the BrainScaleS wafer-scale system

Output 6: Spike-based learning on neuromorphic substrates

4.1 Output 1: Natural gradient for spiking neurons

Elena Kreutzer, Mihai A. Petrovici, Walter Senn

CDP5 collaboration between SP4 (UBERN, P71) and SP9 (UHEI, P47)

We show that natural-gradient-based learning for spiking neurons predicts counterbalancing of homo- and heterosynaptic plasticity, thereby aligning models of error-correcting synaptic plasticity to experimental evidence. In addition, we demonstrate that in many cases, it can be approximated by a simpler rule, enhancing biological plausibility and facilitating implementation on neuromorphic hardware.

4.1.1 Research and Infrastructure

We investigated the predictions of the previously suggested supervised learning rule based on natural gradient descent, which leads to a consistent, parameter-invariant description of synaptic plasticity in the brain, with possible advantages for learning on neuromorphic hardware (see CDP5 SGA1 M24 Deliverable).

Whereas standard gradient-based rules can only account for homosynaptic plasticity, i.e. they only adapt connections to currently active afferents, we showed that the natural gradient rule also predicts changes to synapses that did not transmit input in the current time step, given a high overall activity of the neuron's afferents. This form of plasticity, called heterosynaptic, has been experimentally observed in both the hippocampus and the neocortex, and often been found to display the opposite sign compared to the homosynaptic changes.

In the context of a simple stimulation protocol, we demonstrated that qualitatively similarly to the experimental evidence, the natural gradient rule in general predicted the accompaniment of homosynaptic LTP with heterosynaptic LTD and vice versa, possibly enhancing competition during learning (Fig. 6). In the case of very large synaptic weights, we observed LTD in both stimulated and unstimulated pathways, keeping the synaptic weights in a suitable range.

Furthermore, we showed that under certain circumstances, our learning rule can be approximated by a simpler rule, making it easier for a synapse to implement. In addition, this increases the value of natural gradient learning as a possible training algorithm for neuromorphic hardware. First steps towards an implementation on the BrainScaleS2 platform have been discussed in a dedicated workgroup during the CDP5 meeting in Heidelberg.

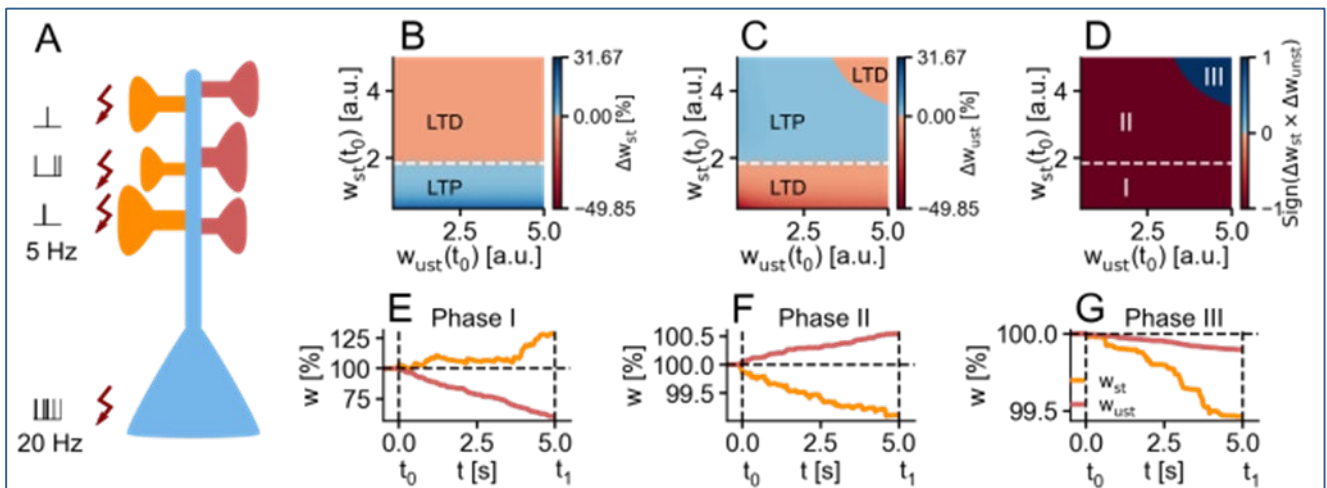


Figure 6: Heterosynaptic plasticity counterbalances homosynaptic plasticity.

Key: (A) We investigated the direction of homo- and heterosynaptic plasticity in a simple experiment with variable excitatory input and tonic inhibition. We stimulated half of a neuron's afferents providing excitatory Poisson input at 5Hz, while simultaneously delivering tonic inhibition and a teacher spike train at 20 Hz. This elicited homosynaptic plasticity at the stimulated and heterosynaptic plasticity at the unstimulated synapses. We investigated the direction of these components as a function of the magnitude of both the initial stimulated and initial unstimulated weights. For reasons of simplicity, the initial weights within each group were assumed to be equal. (B) Weight change of stimulated weights. While varying the size of the unstimulated weights has no effect on homosynaptic plasticity, increasing the initial weight of the simulated synapses results in a change from potentiation to depression, once the

postsynaptic error term becomes negative. (C) Weight change of the unstimulated weights. Heterosynaptic plasticity decreased the unstimulated weights in regions where the stimulated weights underwent LTP. Increasing the size of initial stimulated weights resulted in a change to potentiation at the same point where homosynaptic LTP turned into LTD. Further increase of either unstimulated or stimulated weights resulted eventually lead to depression of the unstimulated synapses. (D) Phase diagram comparing directions of homo- vs. heterosynaptic plasticity. Red regions indicate opposing signs, whereas regions where homo-and heterosynaptic plasticity pull in the same direction are marked in blue. (E-G) Example traces of synaptic weights of stimulated and unstimulated synapses during learning, with initial weights picked from the different phases indicated in (D).

4.1.2 *Impact*

On the technical side, the ability of a simplified synaptic plasticity rule to approximate the powerful natural gradient algorithm is essential for artificial systems with limited local bandwidth and computation capabilities. Regarding biological neural circuits, the new explanation of experimental data on heterosynaptic plasticity provides an elegant link between cortical dynamics and rigorous results in information geometry.

4.1.3 *Components*

C457 BrainScaleS 2 Neuromorphic Computing System

C2419 Plasticity: algorithms for multi-compartment models

C2420 Plasticity: prototype implementations of rules and testing within and without the SP9 platforms

C2547 Computing with structured neurons

C2722 Multicompartmental pyramidal neuron models

4.1.4 *Publications*

Journal paper in preparation.

4.1.5 *Further measures for increasing impact*

- 1) Neuromorphic computing - The physics of cognition. Mihai A. Petrovici, Talk at the HBP National Outreach Event, Amsterdam, March 2, 2018
- 2) Natural gradient for spiking neurons. Elena Kreutzer, Talk at the Manfred Stärk Symposium, Bern, May 2-3, 2018
- 3) Natural gradient for spiking neurons. Elena Kreutzer, Talk at the Fürberg Workshop, Fürberg, June 26-29, 2018
- 4) Natural gradient for spiking neurons. Elena Kreutzer, Talk at the Heidelberg-Bern Group Meeting, Triberg, July 18-20, 2018
- 5) Form versus function: Theory and models for neuronal substrates. Mihai A. Petrovici, Talk at the Symposium on the Theory of Neuromorphic Computing, Heidelberg, December 13-14, 2018
- 6) Mihai A. Petrovici, Interview with Euronews, Bern, March 11, 2019
<https://www.euronews.com/2019/04/01/the-human-brain-project-slicing-brains-and-thinking-machines>
<https://www.euronews.com/2019/04/01/can-we-make-a-computer-like-the-human-brain>
- 7) Computing with physics: From biological to artificial intelligence and back again. Mihai A. Petrovici, Talk at the Neuro-Inspired Computational Elements conference, Albany, March 26-29, 2019

4.2 Output 2: Lagrangian neurodynamics for real-time error-backpropagation across cortical areas

Dominik Dold, Akos F. Kungl, Joao Sacramento, Mihai A. Petrovici, Walter Senn

CDP5 collaboration between SP4 (UBERN, P71) and SP9 (UHEI, P47)

A major driving force behind the recent achievements of deep learning is the backpropagation-of-errors algorithm (backprop), which solves the credit assignment problem for deep neural networks. Its effectiveness in abstract neural networks notwithstanding, it remains unclear whether backprop represents a viable implementation of cortical plasticity. In CDP5, we developed a new theoretical framework that uses a least-action principle to derive a biologically plausible implementation of backprop.

4.2.1 Research and Infrastructure

In our model, neuronal dynamics are derived as Euler-Lagrange equations of a scalar function (the Lagrangian). The resulting dynamics can be interpreted as those of multi-compartment neurons with apical and basal dendrites, coupled with a Hodgkin-Huxley-like activation mechanism that undoes temporal delays introduced by finite membrane time constants. We suggest that a neuron's apical potential encodes a local prediction error arising from the difference between top-down feedback from higher cortical areas and the bottom-up prediction represented by activity in its home layer. This computation is enabled by a stereotypical cortical microcircuit, projecting from pyramidal neurons to interneurons back to the pyramidal neurons' apical compartments (Fig. 7A). When a subset of output neurons is slightly nudged towards a target behavior that cannot be explained away by bottom-up predictions, an error signal is induced that propagates back throughout the network through feedback connections. By defining synaptic dynamics as gradient descent on the Lagrangian, we obtain a biologically plausible plasticity rule that acts on the forward projections of pyramidal and interneurons in order to reduce this error.

The presented model incorporates several features of biological neurons that cooperate towards approximating a time-continuous version of backprop, where plasticity acts at all times to reduce an output error induced by mismatch between different information streams in the network. The model is not only restricted to supervised learning, but can also be applied to unsupervised and reinforcement learning schemes, as demonstrated in simulations (Fig. 7B-D).

To summarize, we present a formalized approach to deriving biologically plausible neurosynaptic dynamics integrating error backpropagation. Different from previous models, the derived dynamics allow a time-continuous implementation of error backpropagation, where none of the dynamics have to be stationary in order for plasticity to be able to reduce an output cost. Finally, we would like to stress two key points of the model's cortical implementation: (i) Learning is driven by a local and biologically plausible plasticity rule (the dendritic prediction of somatic activity) and (ii) by minimizing local prediction errors, each neuron contributes to reducing a global cost.

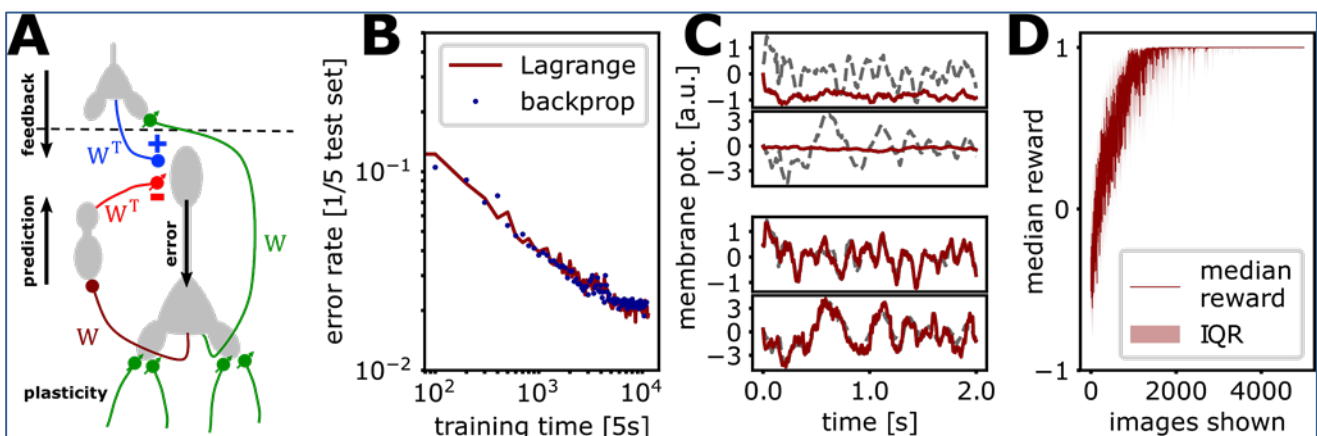


Figure 7: Lagrangian dynamics of cortical microcircuits.

Key: (A) Error coding scheme with compartmental model. A bottom-up prediction mediated via interneurons tries to explain away top-down feedback coming from higher cortical areas. If not all feedback can be explained away, this leads to a non-zero prediction error in the apical compartment, driving plasticity in the forward connections. Here, microcircuit, feedback and forward weights are coupled. However, these can also be learned independently, without requiring any weight transport. (B) Learning MNIST with a layered network (784-500-10). During training, every training image is only shown for a short amount of time (10ms) such that the network never reaches a stationary state. (C) Unsupervised learning of a time-continuous human intracortical EEG signal (56 electrodes, modelled by 56+40 fully recurrent neurons) before and after training. During test runs, the network only sees 46 of 56 inputs and reproduces the remainder (only 2 shown). (D) Classification of three images with reinforcement learning (reward is +1/-1). A winner-take-all like connectivity among the output neurons provides the necessary nudging when learning is based on scalar reward signals.

4.2.2 Impact

The derivation of neuronal and synaptic dynamics from a principled, unified mathematical framework (least-action principle) sheds new light on our understanding of cortical activity by highlighting a deep connection between the dynamics of biological and physical observables. This represents a novel and fundamental insight for computational neuroscience.

Two aspects of the theory are quintessential for learning in physical systems that are bound to some extent by information locality and the time-continuous nature of both their input and their internal dynamics: the local computation of errors and the prospective coding of neuronal layers. These unique properties, hitherto unavailable in neural networks in this combination, have attracted widespread interest within the neuroscientific, neuromorphic and machine learning communities, both within and outside the HBP, as evidenced by an abundance of invited and selected talks and collaborations with external partners such as MILA, CIFAR, Intel and DeepMind (see Dissemination).

4.2.3 Components

C2226 Generative model of sensory cortical hierarchy and corticohippocampal network

C2419 Plasticity: algorithms for multi-compartment models

C2420 Plasticity: prototype implementations of rules and testing within and without the SP9 platforms

C2547 Computing with structured neurons

C2722 Multicompartmental pyramidal neuron models

4.2.4 Publications

- 1) Dendritic cortical microcircuits approximate the backpropagation algorithm. Joao Sacramento, Rui Ponte Costa, Yoshua Bengio, Walter Senn. *Advances in Neural Information Processing Systems*, pp 8735-8746, 2018
- 2) Lagrangian dynamics of dendritic microcircuits enables real-time backpropagation of errors. Dominik Dold, Akos F. Kungl, João Sacramento, Mihai A. Petrovici, Kaspar Schindler, Jonathan Binas, Yoshua Bengio, Walter Senn. *Proceedings of the Cosyne conference*, Lisbon, 2019. Selected talk.

An additional journal paper is in preparation.

4.2.5 *Further measures for increasing impact*

- 1) Lagrangian mechanics for real-time error-backpropagation across multi-area cortical circuits. Dominik Dold, Talk at the EITN “From Neuroscience to Machine Learning” workshop, Paris, March 12-13, 2018
- 2) Lagrange-Mechanik neuronaler Netze. Dominik Dold, Talk at the Manfred Stärk Symposium, Bern, May 2-3, 2018
- 3) Continuous learning in dendritic cortical microcircuits using Lagrangian mechanics. Dominik Dold, Poster presentation at the EMBO Dendrites Workshop, Heraklion, June 17-20, 2018
Continuous error backpropagation in cortical microcircuits from Euler-Lagrange equations. Dominik Dold, Talk at the Fürberg Workshop, Fürberg, June 26-29, 2018
- 4) From Euler-Lagrange to time-continuous error backpropagation in cortical microcircuits. Dominik Dold, Talk at the “Neuroplasticity: From Bench to Machine Learning” workshop, Guildford, July 13-14, 2018
- 5) Continuous error backpropagation in dendritic cortical microcircuits using Lagrangian mechanics. Dominik Dold, Poster presentation at Bernstein Conference, Berlin, September 25-29, 2018
- 6) Real-time error backpropagation for deep cortical networks. Dominik Dold, Talk at the Intel INRC workshop, Reykjavík, September 28 - October 2, 2018
- 7) From Euler-Lagrange to error backpropagation in cortical circuits. Dominik Dold, Talk at ETH Institute of Neuroinformatics, Zurich, December 14, 2018
- 8) Dendritic cortical microcircuits approximate the backpropagation algorithm. Talk by João Sacramento at the NIPS Conference, Montreal, December 2-8, 2018
- 9) Form versus function: Theory and models for neuronal substrates. Mihai A. Petrovici, Talk at the Symposium on the Theory of Neuromorphic Computing, Heidelberg, December 13-14, 2018
- 10) Computers like brains. Mihai A. Petrovici, Talk at the “At the crossroad of physics and machine learning” conference, Santa Barbara, February 11-15, 2019
- 11) <http://online.kitp.ucsb.edu/online/machine-c19/petrovici/>
- 12) Lagrangian dynamics of dendritic microcircuits enables real-time backpropagation of errors. Dominik Dold, Selected talk at the Cosyne conference, Lisbon, February 28 - March 5, 2019
- 13) Google Deep Mind, London Feb. 5–7, Presentation by W. Senn: “Lagrangian neurodynamics and deep learning”.
- 14) Workshop Deep Learning, Barbados February 23 - February 28, 2019, Presentation by W. Senn: “Lagrangian neuro-dynamics for online error-backpropagation in the brain.”
- 15) Mihai A. Petrovici, Interview with Euronews, Bern, March 11, 2019
- 16) <https://www.euronews.com/2019/04/01/the-human-brain-project-slicing-brains-and-thinking-machines>
<https://www.euronews.com/2019/04/01/can-we-make-a-computer-like-the-human-brain>
- 17) Computing with physics: From biological to artificial intelligence and back again. Mihai A. Petrovici, Talk at the Neuro-Inspired Computational Elements conference, Albany, March 26-29, 2019
- 18) Gordon Conference “Dendrites”, Ventura CA, March 31 - April 4. Talk by W. Senn, “Cortical microcircuits and error-representations”

4.3 Output 3: Error-driven learning supports Bayes-optimal multisensory integration via conductance-based dendrites

Jakob Jordan, João Sacramento, Mihai A. Petrovici, Walter Senn

CDP5 collaboration between SP4 (UBERN, P71) and SP9 (UHEI, P47)

Animals receive information about their environment through a variety of senses that need to be integrated to form a coherent percept. To combine information from multiple senses meaningfully requires a representation of the reliability of each source. We developed a formal framework that maps such types of probabilistic computations to the biophysical dynamics of multi-compartment neurons with conductances-based synapses.

4.3.1 Research and Infrastructure

Starting from an energy function describing the macroscopic state of a network, we derived neuron and weight dynamics that minimize this energy via gradient descent. The resulting neuron dynamics resemble multi-compartment model with conductance-based synapses. Weight dynamics correspond to an error-driven plasticity rule for conductance-based synapses and contain an additional term that matches reliability of an input to the total conductance it generates. In this framework neurons learn to use membrane potentials and synaptic conductances to represent the first and second moments of probability distributions.

We demonstrated that the model can successfully perform a multisensory integration task for orientation estimation in which a layer of multisensory neurons receive inputs from visual and tactile feature detectors (Fig. 8). Each modality was presented with a one-dimensional stimulus obtained from a ground-truth orientation by adding Gaussian white noise with modality-specific amplitude. After learning, the network was able to combine these two noisy sources of information, reaching error levels comparable to a Bayes-optimal estimator.




Figure 8: Learning multisensory integration of visual and tactile information in conductance-based neurons with dendrites.

Corrupting noise on visual modality is 33% smaller than on the tactile modality. Left panel: Energy as a function of training step. Middle panel: root-mean-square deviation of true orientation and estimate reconstructed from output population activity. In the V and T conditions, input was only provided to the visual and tactile inputs respectively, while neurons of the other modality were clamped to their resting potential. Right panel: relative reliability of visual input.

4.3.2 Impact

The idea that individual neurons compute mainly by summing over their inputs represents a long-standing paradigm for neural networks. This work re-addresses the question of neural computation on a fundamental level by considering the effects of neuronal structure combined with conductance-based synapses. Our shift of perspective on the fundamental computational function of single neurons being the computation of a convex combination rather than a sum of their inputs allows us

to interpret their activity as Bayesian cue combination, which has the potential to explain a variety of experimental observations.

4.3.3 *Components*

C2419 Plasticity: algorithms for multi-compartment models

C2420 Plasticity: prototype implementations of rules and testing within and without the SP9 platforms

C2547 Computing with structured neurons

C2722 Multicompartmental pyramidal neuron models

4.3.4 *Publications*

- 1) Error-driven learning supports Bayes-optimal multisensory integration via conductance-based dendrites. Jakob Jordan, Joao Sacramento, Mihai Petrovici, Walter Senn, Proceedings of the Cosyne conference, Lisbon, 2019. Poster presentation.

4.3.5 *Further measures for increasing impact*

- 1) Certainty-weighted multisensory integration by dendrites. Jakob Jordan, Poster presentation at the EMBO Dendrites Workshop, Heraklion, June 17-20, 2018
- 2) Error-driven learning supports Bayes-optimal multisensory integration via conductance-based dendrites. Jakob Jordan, Poster presentation at the Cosyne conference, Lisbon, February 28 - March 5, 2019

4.4 **Output 4: Robust predictive plasticity in dendrites: From a computational principle to experimental data**

Dominik Spicher, Claudia Clopath, Walter Senn

CDP5 collaboration between SP4 (UBERN, P71) and Imperial College London (external to HBP)

We show how supervised learning in a single neuron can be sustained in the presence of strong intracellular voltage signals such as action potentials. Instead of the dendritic synapses being required to access the firing intensity function of the Poisson neuron in an exact manner, we postulate that different dendrites employ different approximations to give local estimates of the firing rate that convex-combine in the somatic compartment to a set of learnable voltage levels. We show how this can lead to stable learning and predict that voltages should cluster around different levels for different dendrites.

4.4.1 *Research and Infrastructure*

We have previously shown how a learning model that supports supervised learning in a single neuron (Urbanczik & Senn, 2014) can be brought into accordance with experimental evidence regarding the spike-timing dependent component of plasticity (STDP) in cortical pyramidal neurons. Besides the basic STDP phenomenon depicted in Fig. 9-1A, this further included frequency-dependent phenomena, as well as the impact of the synaptic distance to the soma on plasticity.

The depression part of the STDP phenomenon within this model depends crucially on the presence of a large voltage deflection caused by the action potential. This potentially poses a challenge to a supervised learning scheme where both the desired as well as the learned output are represented by continuous voltage levels. Indeed, the presence of a refractory period and the voltage deflection in the dendrite lead to deterioration of any learned voltage level once the teacher is removed (data not shown).

We thus studied a model where multiple dendrites are connected to a somatic compartment and synaptic plasticity serves to reproduce the somatic voltage in all dendrites individually. Depression in dendritic synapses is not driven by the somatic firing rate function, instead different dendrites employ different estimates thereof with step-like voltage-dependent processes (Fig. 9-2A). This leads to stable fixed points for the learning process that are robust to biases in the absence of a teacher. We demonstrate this behaviour in Fig. 9-2B where two dendrites convex-combine their fixed points after learning to achieve a somatic voltage trace close to the teaching signal.




Figure 9: Predictive plasticity in dendrites.

Key: 1) Spike-timing dependence of plasticity. Scatter plot shows data from Bi & Poo (1998), red curve shows the weight change in our model. 2A) Firing rate function of the Poisson unit (blue), as well as two dendritic estimates (green and orange) employed in the plasticity of synapses in the respective compartments. Dendritic estimates are step-like processes that establish fixed points. 2B) Learning a step-like teaching signal (red dashed) with the two dendrites whose firing rate estimates are depicted in Fig 2A. The first set of traces shows voltage levels before learning. On the right two consecutive cycles are depicted after learning when the teacher is removed. The two dendrites combine their fixed points such that the resulting convex combination approaches the teaching signal.

4.4.2 Impact

We have demonstrated how supervised learning in single neurons can be sustained in a robust manner in the presence of nuisance signals. When combining multiple such neurons in networks with a suitable wiring pattern, this learning rule can also perform unsupervised and reinforcement learning tasks. One particularly striking prediction of the model is that the voltages of different dendrites should cluster around distinct values on the subthreshold voltage spectrum.

4.4.3 Components

C1032 Plasticity: dendritic predictive plasticity that reproduces STDP data

C2419 Plasticity: algorithms for multi-compartment models

C2420 Plasticity: prototype implementations of rules and testing within and without the SP9 platforms

C2547 Computing with structured neurons

C2722 Multicompartmental pyramidal neuron models

4.4.4 Publications

Journal paper in preparation.

4.4.5 Further measures for increasing impact

- 1) Robust predictive plasticity in dendrites: From a computational principle to experimental data. Dominik Spicher, Talk at the Fürberg Workshop, Fürberg, June 26-29, 2018
- 2) Robust predictive plasticity in dendrites: From a computational principle to experimental data. Dominik Spicher, Talk at the Heidelberg-Bern Group Meeting, Triberg, July 18-20, 2018

4.5 Output 5: Training deep networks with time-to-first-spike coding on the BrainScaleS wafer-scale system

Julian Göltz, Oliver Breitwieser, Sebastian Schmitt, Johannes Schemmel, Mihai A. Petrovici
CDP5 collaboration between SP9 (UHEI, P47) and SP4 (UBERN, P71)

We established a framework that learns to recognize patterns on the BrainScaleS-1 using time-to-first-spike coding. Such learning with single spikes promises to be an energy efficient and fast approach to machine learning on neuromorphic hardware.

4.5.1 Research and Infrastructure

In classical neural networks, error backpropagation learning operates on neuronal outputs represented by (positive) real-valued numbers. In previous translations to spiking networks, these were represented as firing rates (see CDP5 SGA1 M24 Deliverable). However, single spikes can also encode such real values by means of the time of their occurrence.

To use these time-to-first-spike values for learning, exact and differentiable relations for the spike time of a leaky integrate-and-fire neuron under particular parametrisations were derived. These relations allow us to train a deep network on the neuromorphic hardware with backpropagation. In parallel, this computational framework was also studied in NEST simulations.

The training process for a reduced version of the MNIST dataset on the BrainScaleS-1 system is shown in Fig. 10. The input category was determined by the identity of the first neuron to spike in the label layer. After training, images could be classified within 40 ms of biological time (4 μ s wall-clock time) after the input was received. As classification occurs so quickly, most neurons spike at most once, drastically reducing the energy expenditure compared to the same computation being performed with rate coding. Experiments with larger networks and more complex datasets are in progress.




Figure 10: Learning to classify reduced 7x7 MNIST digits with single spikes on the BrainScaleS-1 system.

A network of shape 49>20>3 was trained with stochastic gradient descent using a differentiable function of spike times. We used the 0, 1, and 4 digits of the MNIST dataset averaged to 7x7 pixels. A) Accuracy increase during training.

The initial low accuracy increases within a few hundred training steps up to almost 100%. B) Confusion matrix after training. The matrix displays the classification results (columns) of the inputs (rows). Example images of the input 0, 1, and 4 are shown on the left. C) Voltage traces of the network after training. The rows show the reaction of the network to different example inputs (0, 1, 4 as written in the panels). Left column: six randomly selected hidden neurons (of 20 in total). Right column: neurons of the label layer. The colors of the traces in the label layer correspond to classes. Classifications are correct (color of first spiking neuron agrees with label) and fast (40 ms biological time).

4.5.2 *Impact*

Given the noisy and diverse nature of the analog components of the BrainScaleS systems, the training of precise spike times is a significant challenge. The successful realization of these small-scale experiments therefore represents an important milestone and points towards the feasibility of scaling them to larger networks and more difficult problems.

4.5.3 *Components*

- C1 BrainScaleS1 Neuromorphic Computing System
- C2 SpiNNaker Neuromorphic Computing System
- C209 NEST - The Neural Simulation Tool
- C349 PyNN
- C457 BrainScaleS2 Neuromorphic Computing System
- C2420 Plasticity: prototype implementations of rules and testing within and without the SP9 platforms

4.5.4 *Publications*

A Master's Thesis is in preparation.

4.5.5 *Further measures for increasing impact*

- 1) Mihai A. Petrovici, Interview with Euronews, Bern, March 11, 2019

<https://www.euronews.com/2019/04/01/the-human-brain-project-slicing-brains-and-thinking-machines>

<https://www.euronews.com/2019/04/01/can-we-make-a-computer-like-the-human-brain>

4.6 **Output 6: Spike-based learning on neuromorphic substrates**

Brian Gardner, Joseph Chrol-Cannon, Andre Grüning, Johannes Schemmel, Eric Müller, Steve Furber
CDP5 collaboration between SP9 (UHEI, P47 and UMAN, P63) and SP4 (SURREY, P102)

We adapted a number of classical benchmark learning tasks for use with spiking neural networks utilising learning rules that were developed through SGA1. We gave special consideration to the constraints and opportunities presented by the current releases of the neuromorphic systems within HBP.

4.6.1 Research and Infrastructure

It is important to find computing paradigm that go beyond classical von-Neumann architectures as these are reaching physical limits with respect to ever smaller integration and ever greater power-consumption. So far learning algorithms that match the performance of classical AI approaches are lacking for spike-based computation. We contribute to this quest for good learning algorithms on neuromorphic hardware in the following ways.

We implemented biologically inspired but machine-learning relevant learning rules on the SP9 neuromorphic systems. The neuromorphic systems exhibit a number of constraints that make it difficult to translate fully-programmable software-based simulations to them in a one-to-one way. For example, the DLSv2 system as currently 32 neuronal compartments (that can stand for neurons or biological compartments thereof) and one fixed type of pre- and postsynaptic traces.

We therefore converted standard-software simulations of larger spiking neural networks utilising HBP-internal (developed during the RUP and SGA1) and HBP-external plasticity rules into networks that only use a number of neurons commensurate with the current size of the currently available DLSv2 system and learning rules that only require features already present in the DLSv2.

We have successfully implemented learning on two standard ML-benchmark data set (Iris data set, <http://archive.ics.uci.edu/ml/datasets/Iris>, Wisconsin breast cancer data set, [https://archive.ics.uci.edu/ml/datasets/Breast+Cancer+Wisconsin+\(Diagnostic\)](https://archive.ics.uci.edu/ml/datasets/Breast+Cancer+Wisconsin+(Diagnostic))) utilising in the order of 32 neurons (commensurate with DLSv2) with performance at a competitive level to standard machine learning. Work on learning of the challenging MNIST data set is well underway, however in order to achieve competitive performance, we currently need about 3-5 times more neurons than the DLSv2 currently offers, and thus we are working to minimise the number of required neurons.

Furthermore, we have been working on an application of HBP-developed multicompartment spike-based learning rules in imitation learning which is relevant for understanding biological learning behaviour. The learning rules have been transferred to ML tasks in robot learning to demonstrate the usefulness of HBP-learning rules on a large scale. We have now been successful in demonstrating the usefulness in a simplified imitation learning task and are working on expanding to a full complex imitation learning task of songbird song learning.




Figure 11: Training and test loss with the number of training epochs, using three-fold cross validation on Iris and Wisconsin (averaged over 20 runs).

Left: Iris, containing 150 samples split between three classes, classified by a 48x20x3 SNN. Right: Wisconsin, containing 699 samples split between two classes, classified by a 63x20x2 SNN.

4.6.2 Impact

We have demonstrated that 1. competitive performance in standard ML benchmarks is achievable with HBP-developed spiking-based learning rules and with a small number of neurons commensurate

with constraints of the neuromorphic systems, and that HBP-developed learning rules are implementable in principle on the prototypes of the neuromorphic systems available currently.

4.6.3 *Components*

- C1 BrainScaleS1 Neuromorphic Computing System
- C2 SpiNNaker Neuromorphic Computing System
- C2419 Plasticity: algorithms for multi-compartment models
- C2420 Plasticity: prototype implementations of rules and testing within and without the SP9 platforms

4.6.4 *Publications*

A journal publication is in preparation.

4.6.5 *Further measures for increasing impact*

- 1) Gardner B., Grüning A.: “Cognitive Computing -- Merging Concepts with Hardware”. Workshop in Hannover, Dec 2018.
- 2) Grüning A, Gardner B, Bulanov A, Furber S: “Plasticity and learning on neuromorphic platforms”, Workshop Abstract. Bristol Brain Research. 113(41). 11441-11446
- 3) Gardner B, Grüning A (2018): Pattern recognition using latency coding in multilayer spiking neural networks. Bernstein Conference 2018. doi: 10.12751/nncn.bc2018.0112
- 4) Mihai A. Petrovici, Interview with Euronews, Bern, March 11, 2019
<https://www.euronews.com/2019/04/01/the-human-brain-project-slicing-brains-and-thinking-machines>
<https://www.euronews.com/2019/04/01/can-we-make-a-computer-like-the-human-brain>

5. Key Result KRc5.4: Deep episodic learning

Overview of Outputs

- Output 1: Reconstruction of stimuli across multiple modalities
- Output 2: Interaction between sleep and memory in a thalamo-cortical model performing visual classification

5.1 Output 1: Reconstruction of stimuli across multiple modalities

Shirin Dora, Sander Bohte, Martin Pearson, Walter Senn, Cyriel Pennartz.

CDP5 collaboration between SP3 (UvA, P98 and UWE, P101) and SP4 (UBERN, P71)

We demonstrate that a multilayered neural network trained using predictive coding on unisensory stimuli results in neuronal activity that exhibits same properties as observed in experimental data. Furthermore, we extended the unisensory model of predictive coding to

handle multisensory data and showed that the multisensory representations inferred by the model can reconstruct stimuli in multiple modalities simultaneously.

5.1.1 Research and Infrastructure

We have shown that neural networks trained using predictive coding, exhibit neuronal properties (like selectivity, sparsity, etc.) similar to those observed in experimental data. Figure 11 shows histogram of selectivities for neurons in a network with 4 layers. Kurtosis was used as a measure of neuronal selectivity. These networks were trained on stimulus in a single modality. Subsequently, we have extended this approach to construct generative models on multisensory stimuli. The model can be used to infer a single multisensory representation that can simultaneously reconstruct stimuli across multiple modalities.

We also evaluated the trained model for the problem of crossmodal recall in the brain. For this purpose, we presented the trained model with stimulus in one of the modalities while the related stimulus in another modality was not presented to the trained model. The trained model was then used to infer representations using only the unisensory stimulus. We were able to reconstruct the stimulus that was not presented to the trained model based on the representations inferred using the partial stimulus.

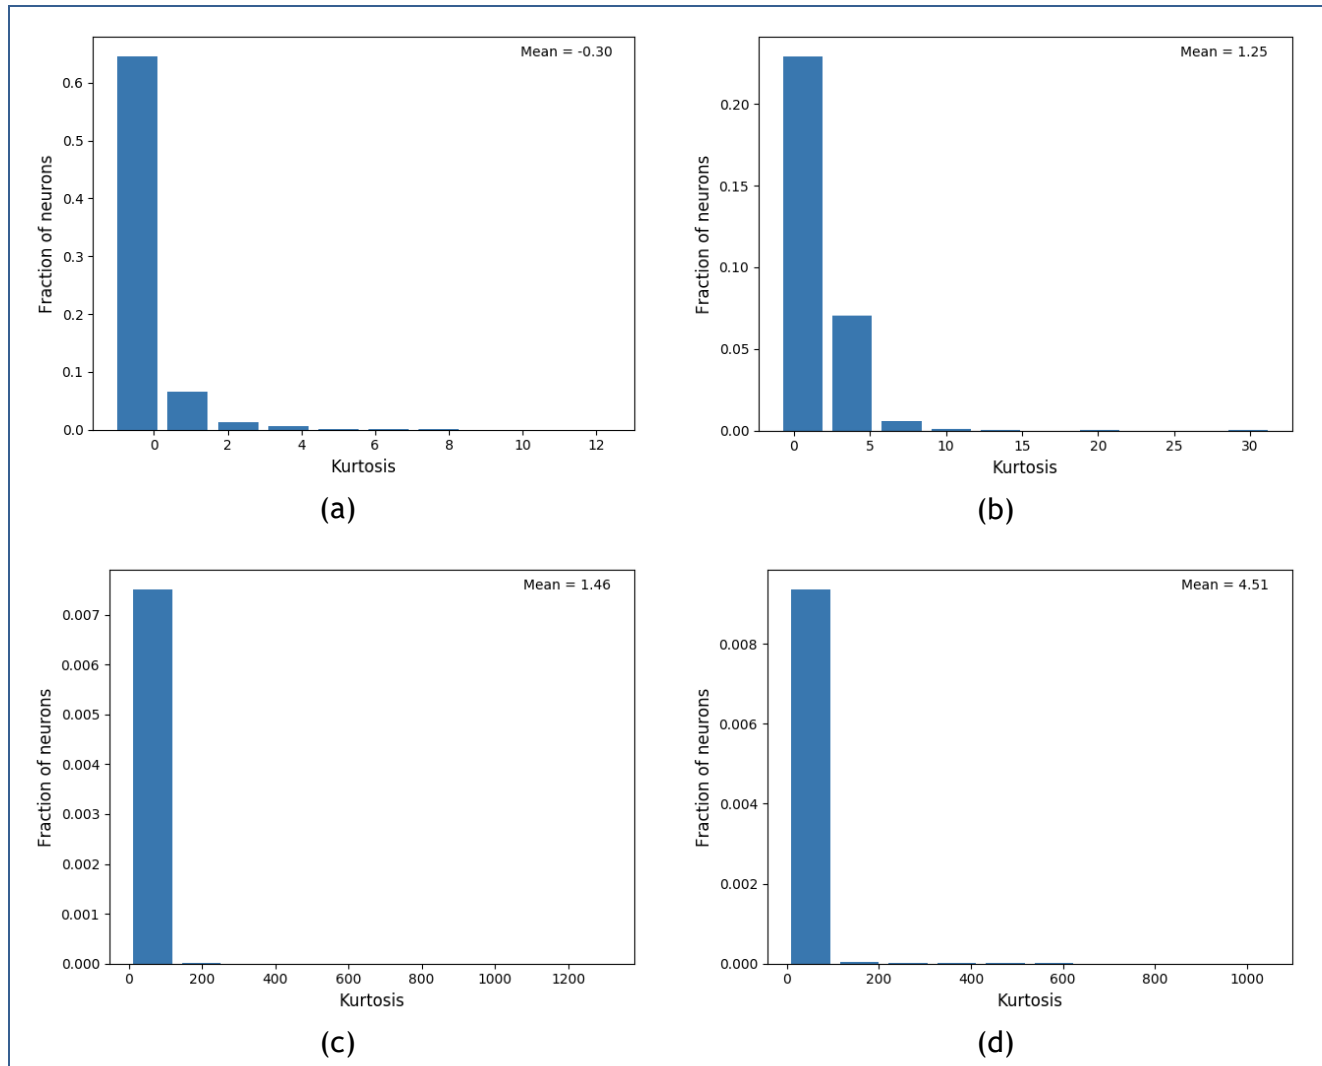


Figure 12: Histogram of kurtosis for neurons in different layers of the model.

Key: (a), (b), (c) and (d) show the histograms for neurons in layers 1, 2, 3 and 4 of the network, respectively. The mean kurtosis values for corresponding layers is shown top right corner of each plot.

5.1.2 *Impact*

Our model can be used to handle data in more than 2 modalities irrespective of the modalities used. It will be leveraged during the second year of SGA2 activity and will constitute a significant contribution for the multi-sensory cognitive architecture planned for the next HBP phase.

In relation to the model of multisensory integration and crossmodal recall using predictive coding, we want to study the feasibility of predictive coding as a mechanism that supports crossmodal recall in the brain (work in SP3). For this purpose, we are using this model together with experimental data from SP3 and sensor data from the Whiskey robot. The Whiskey robot will be used for gathering sensory input on objects used during experiments. The neuronal activities in the model will then be correlated with experimental data.

5.1.3 *Components*

C2061 Multi-area ensemble mechanisms of object recognition in rodents

C2226 Generative model of sensory cortical hierarchy and corticohippocampal network

C2228 Comparison with physiological data and related models

C2301 Instantiate model of WhiskEye into NRP

5.1.4 *Publications*

- 1) A deep predictive coding network for inferring hierarchical causes underlying sensory inputs. Shirin Dora, Cyriel Pennartz, Sander Bohte, International Conference on Artificial Neural Networks, pp 457-467, Springer, Cham, 2018
- 2) Deep predictive coding accounts for neural response properties along the visual cortical hierarchy, Shirin Dora, Sander Bohte and Cyriel Pennartz (in preparation)

5.1.5 *Further measures for increasing impact*

- 1) A Deep Predictive Coding Network for Inferring Hierarchical Causes Underlying Sensory Inputs. Shirin Dora, Talk at the International Conference for Artificial Neural Networks, Rhodes, October 4-7, 2018

5.2 **Output 2: Interaction between sleep and memory in a thalamo-cortical model performing visual classification (MNIST)**

Cristiano Capone, Elena Pastorelli, Bruno Golosio, Pier Stanislao Paolucci, Maurizio Mattia, Mihai Petrovici, Johannes Schemmel

CDP5 collaboration between SP3 (INFN, P92 and ISS, P96), SP4 (UBERN, P71) and SP9 (UHEI, P47)

For the first time, we demonstrated two beneficial effects of deep-sleep thalamo-cortical oscillations: the creation of categories from examples and a generalized homeostasis of synaptic weights, with an improvement in post-sleep classification rates and a normalization of firing rates during post-sleep wake activity. This mechanism could play an essential role in bio-inspired learning for neurorobotic applications based on neuromorphic computing.

5.2.1 Research and Infrastructure

Sleep is present in the brain of all animal species, notwithstanding the dangers associated to the disconnection from the environment (e.g. predation) and the reduction of productive time available for feeding and reproduction. It is known that sleep is beneficial to cognitive and mnemonic tasks, but a unified theoretical and computational approach demonstrating the underlying mechanisms is still lacking.

Using a simplified thalamo-cortical model, which is trained to encode, retrieve and classify images of handwritten digits, we demonstrated interesting effects of deep-sleep-like slow oscillation activity. Spike-timing-dependent-plasticity (STDP) produces a differential homeostatic process during deep-sleep like Slow Oscillations. This process is characterized by both a specific unsupervised enhancement of connections among groups of neurons associated to instances of the same class (digit) and a simultaneous down-regulation of stronger synapses created by the training (Fig. 12). This is reflected in a hierarchical organization of post-sleep internal representations. Such effects favour higher performances in retrieval and classification tasks and create hierarchies of categories in integrated representations.

The model leverages on the interaction between of top-down cortico-thalamic predictions and bottom-up thalamo-cortical projections during deep-sleep-like slow oscillations. Such a mechanism hints at possible applications to artificial learning systems. The model is written in NEST, using standard neuron (AdEX) and STDP models available in the NEST library, and is therefore directly compatible with the HBP platforms.

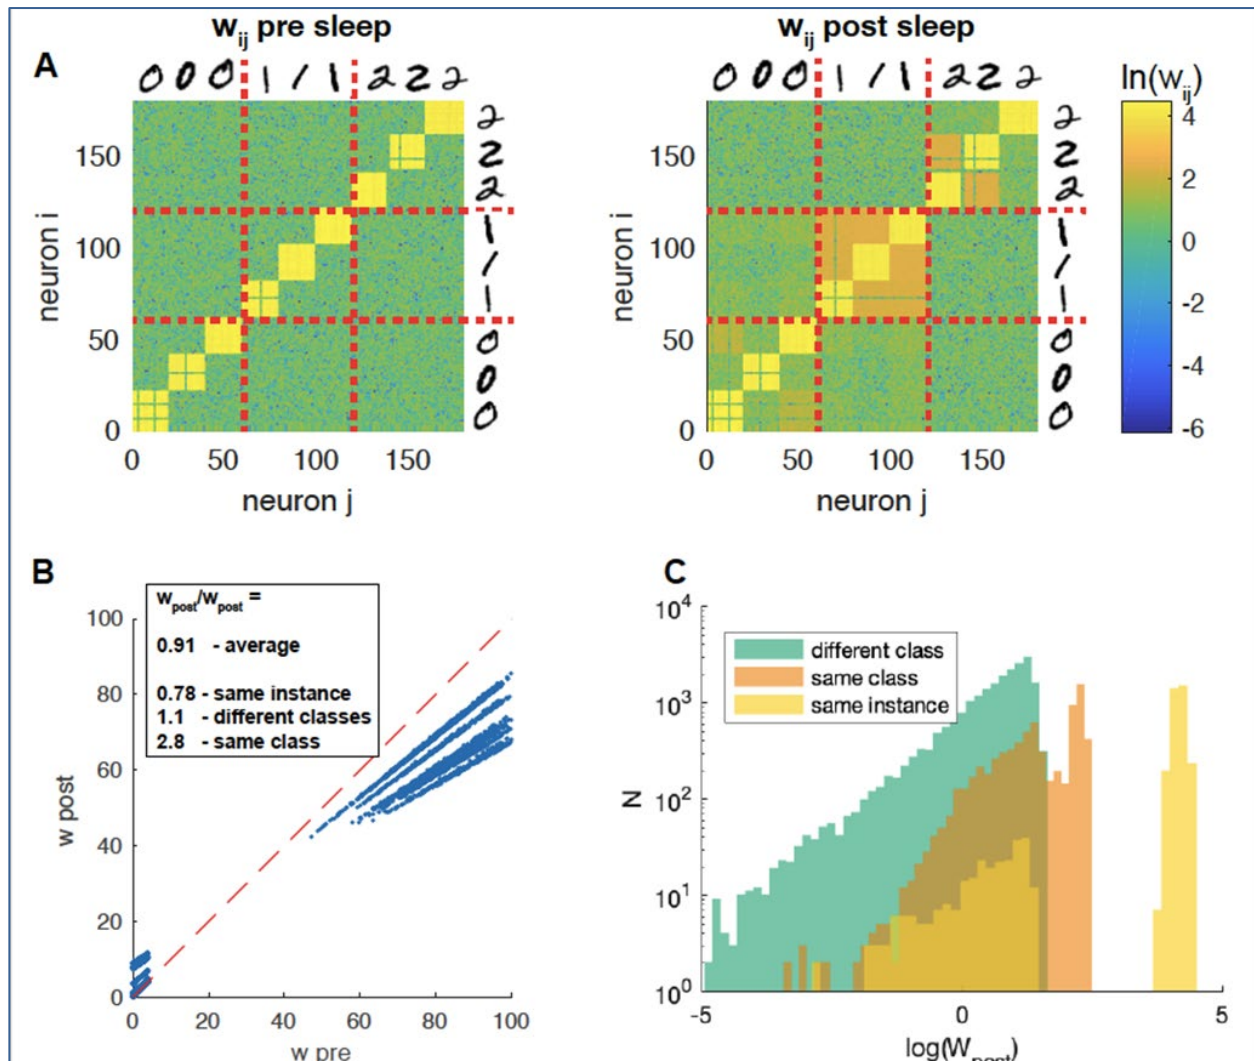


Figure 13: Effects of deep-sleep like Slow Oscillation on synaptic structure in a thalamo-cortical model.

Panel A: Synaptic weights matrix of the recurrent connectivity of cortical population, before (left) and after (right) the occurrence of sleep-like activity. The yellow squares represent high weights emerged between neurons encoding the visual input related to the same object (single instance of 0,1,2 ... image). Red solid lines separate the neurons encoding visual inputs related to different classes of objects (0,1,2 ...). Panel B: Scatter-plot of the same synaptic weights before and after sleep. Panel C: Synaptic weights after sleep, separated in three groups, synapses between neurons encoding the same object (yellow), the same class (but not the same object, orange) and different classes (green).

5.2.2 Impact

The thalamo-cortical model described in output 2 has been applied to MNIST handwritten characters, demonstrating learning after few examples, and simultaneous effects of synaptic homeostasis and memory association during the sleep phase, leading to three beneficial and biologically plausible results: higher post-sleep classification rates, reduction of firing frequencies of neurons spiking at higher frequency rates before the sleeping phase, emerging association between training examples sharing similarities. In perspective, the model can also be extended to multi-modal classification.

5.2.3 Components

C209 NEST - The Neural Simulation Tool

C2193 Cortical spiking model of the interplay between sleep and plasticity

5.2.4 Publications

1. Sleep-like slow oscillations improve visual classification through synaptic homeostasis and memory association in a thalamo-cortical model. C. Capone, E. Pastorelli, B. Golosio, P. S. Paolucci, arXiv:1810.10498, 2018 (submitted for publication)

5.2.5 Further measures for increasing impact

1. Real-time cortical simulation. Pier Stanislao Paolucci, Talk at PDP 2019 (Intern. Conf. on Parallel and Distributed Computing), special session on Neuroscience, Pavia, Italy, 13-15 February, 2019
2. Sleep, Cognition and Bio-Intelligence. Pier S. Paolucci, Talk at the World Brain Awareness Week Event, EBRI, Roma, March 11th, 2019
3. Sleep-Memory interaction. Cristiano Capone, Talk scheduled at the NEST Conference 2019, 24-25 June, Ås, Norway
4. Interaction between memories and cortical oscillations. Pier S. Paolucci, Talk at the Cognitive and System Neuroscience workshop, Berlin, September 13-14, 2018

6. Key Result KRc5.5: Learning to learn with reconfigurable networks

Overview of Outputs

Output 1: Structural plasticity in the form of synaptic rewiring on the SpiNNaker digital neuromorphic system

6.1 Output 1: Structural plasticity in the form of synaptic rewiring on the SpiNNaker digital neuromorphic system

Sandra Diaz-Pier, Wouter Klijn, Petrut Bogdan, Michel Hopkins, Garibaldi Pineda-Garcia, Andrew Rowley, Oliver Rhodes, Simon Davidson, Robert James, Steve Furber, Christian Pehle, Sebastian Billaudelle, Benjamin Cramer, Eric Müller, Andrew Davison

CDP5 collaboration between: SP7 (JUELICH, P20) and SP9 (UMAN, P63, UHEI, P47 and CNRS, P10)

The work demonstrates the value of synaptic rewiring in SNN applications, and their implementation on neuromorphic hardware. Topographic map formation and refinement is presented, together with applications in computer vision. It also makes synaptic rewiring functionality available to the international community of SpiNNaker users, along with its use in simulations run on the large-scale 1 million core SpiNNaker system at the University of Manchester.

6.1.1 Research and Infrastructure

Mammalian brains change their connectivity from early development and throughout adulthood. These changes are typically described in terms of a topographic map, particularly within sensory processing centres. A model of topographic map formation has therefore been implemented on the SpiNNaker neuromorphic system: running in real time, using point neurons, and making use of both synaptic rewiring and spike-timing dependent plasticity (STDP) [1]. It was demonstrated that synaptic rewiring refines an initially rough topographic map over and beyond the ability of STDP, and that input selectivity learnt through STDP is embedded into the network connectivity through rewiring (Fig. 13).

In subsequent work [2], this synaptic rewiring framework was applied within an event-based computer vision system. It was shown that maximum entropy sampling can be emulated in event-based processing systems, and techniques such as synaptic rewiring open the possibility of achieving online unsupervised learning in near-optimal ways - a result that it is difficult to deliver using frame-based approaches due to the very high computational cost of training such networks.

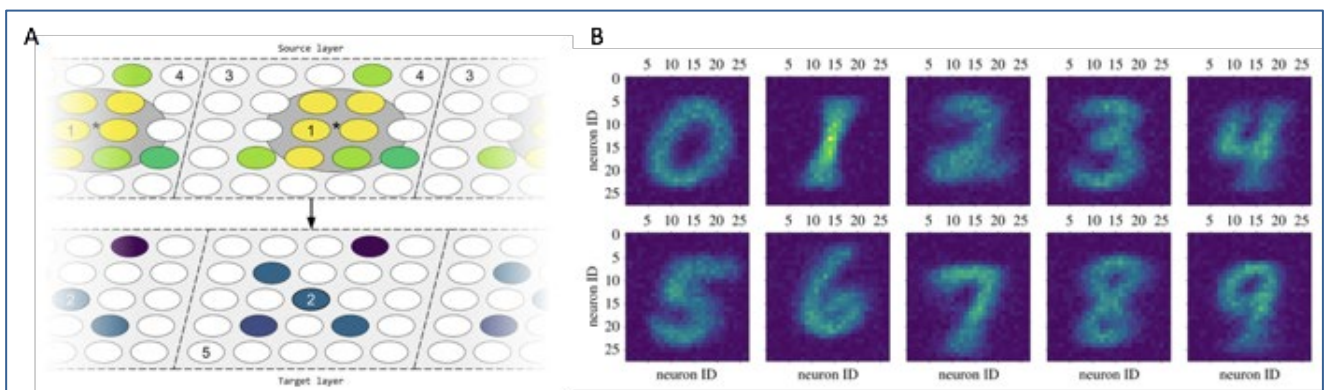


Figure 14: Structural plasticity on SpiNNaker.

Key: (A) Simplified network architecture for topographic map formation [1]. Neurons are placed on a rectangular mesh at discrete locations. Neuron 2 in the target layer has a receptive field formed by the connections coming from the source layer (feedforward), as well as connections coming from within the target layer (lateral), both of which can be updated through synaptic rewiring. (B) Average representation of each post-training MNIST input digit class reconstructed from the individual source population fan-out patterns when presented with non-filtered inputs [2].

6.1.2 *Impact*

This work makes synaptic rewiring functionality available to the international community of SpiNNaker users, along with its use in simulations run on the large-scale 1 million core SpiNNaker system at the University of Manchester.

6.1.3 *Components*

C0002 SpiNNaker Neuromorphic Computing System

C349 PyNN

C2420 Plasticity: prototype implementations of rules and testing within and without the SP9 platforms

6.1.4 *Publications*

- 1) Structural Plasticity on the SpiNNaker Many-Core Neuromorphic System. P. A. Bogdan, A. G. D. Rowley, O. Rhodes and S. B. Furber, *Frontiers in Neuroscience*, 12:434 (2018)
- 2) Spiking neural networks for computer vision. M. Hopkins, G. Pineda-Garcia, P. A. Bogdan, S. B. Furber, *Interface Focus* 8(4) (2018)
- 3) Event-based computation: Unsupervised elementary motion decomposition. P.A. Bogdan, G. Pineda-Garcia, S. Davidson, M. Hopkins, R. James, S. B. Furber, (in press) *Emerging Technology Conference* (2019)

6.1.5 *Further measures for increasing impact*

- 1) SpiNNaker: spiking neural networks for computer vision [2]. Steve Furber. Talk at Royal Society discussion meeting "Understanding Images in Biological and Computer Vision", February 19-20, 2018
- 2) <https://royalsociety.org/science-events-and-lectures/2018/02/understanding-images>
- 3) Structural plasticity for motion detection. Petrut Bogdan. Talk at internal seminar organised by the Machine Learning and Optimisation research group, Manchester, September 5, 2018

7. Key Result KRc5.6: Deep closed-loop sensorimotor learning

Overview of Outputs

Output 1: Real-world embodiment for neuromorphic learning research

Output 2: Embodied Synaptic Plasticity with Online REinforcement (SPORE) learning

7.1 Output 1: Real-world embodiment for neuromorphic learning research

Jacques Kaiser, J. Camilo Vasquez Tieck, Daniel Reichard, Guillaume Bellec, Oliver Rhodes, Emre Neftci, Rüdiger Dillmann

CDP5 collaboration between SP10 (FZI, P52), SP9 (TUGRAZ, P55 and UMAN, P63) and external partner University of California, Irvine, UC, · Department of Cognitive Sciences

We have set up a real-world robotics grasping experiment involving a Schunk arm, a five-finger hand and a robotic head with two Dynamic Vision Sensor. We integrated a synaptic learning rule - random-based backpropagation - trained to identify grasping affordances from microsaccadic eye movements. In the next phase, this setup will be used to benchmark newer learning rules implemented on SpiNNaker running in real-time.

7.1.1 Research and Infrastructure

In this experiment, we embodied the event-driven random backpropagation synaptic learning rule in the real-world grasping robotic setup depicted in Fig. 14. In this setup, the spiking network was trained to recognize four labels corresponding to four different types of grasp: ball-grasp, bottle-grasp, pen-grasp or do nothing.

During training, an object of a particular class is placed on a table at a specific position. The robotic head performs microsaccadic eye movements to extract visual information from the static object. The event streams are recorded together with the corresponding object affordance. During testing, a microsaccade is performed and the detected object affordance triggers the adequate predefined reaching and grasping motion on a Schunk LWA4P arm equipped with a five-finger Schunk SVH gripper (Fig. 14). This demonstrator was implemented by relying on the ROS Framework.

With a surprisingly small amount of training samples (70 per class) and training epochs (~30), the 4-layers network was capable of learning robustly the four discrete affordances (Fig. 15). Since the Dynamic Vision Sensor (DVS) does not sense colours, the network only relies on shape information, crucial for affordances. This allowed the network to moderately generalize despite the small amount of training samples. Indeed, a single object per affordance was used during training, but the network could recognize different objects of the same shape. Recognition also worked when the objects were slightly moved from the reference point used for grasping.




Figure 15: Our robotic setup embodying the event-driven random backpropagation synaptic learning rule.

The spiking network is trained in a supervised fashion to classify visual affordances from event streams provided by a Dynamic Vision Sensor (DVS). The DVS is mounted on a robotic head performing microsaccadic eye movements. The output of the network correspond to the four types of affordances: in the vertical row of ball-grasp, bottle-grasp,

pen-grasp or do nothing. At test time, a Schunk LWA4P arm equipped with a five-finger Schunk SVH gripper perform the detected reaching and grasping motion.




Figure 16: Accuracy on the DvsGestures dataset.

Classification accuracy over 30 epochs using an attention window of certain size. The average performance over all epochs was higher with an increasing attention window size, but the maximum classification accuracy was about 93% for all sizes. The inconsistency of the classification accuracy with a 32×32 attention window is mostly due to the same high learning rate for all three sizes and 32×32 having the smallest amount of parameters to learn. Another reason for this inconsistency is that the frame was too small for many samples to contain the complete motion.

7.1.2 Impact

This real-world task has the potential to enhance neuromorphic and neurorobotics research since many functional components (reaching, grasping, depth perception, eye movements, ...) can be easily segregated and implemented or learned with different brain models.

Neuromorphic technology enables the design of autonomous learning robots operating at high speed for a fraction of the energy budget of current solutions. Until recently, the advantages of this technology were limited due to the lack of efficient learning rules for spiking networks. This bottleneck has been addressed since the realization that gradient backpropagation can be approximated with spikes. We demonstrated the ability of eRBP to learn from real-world event streams provided by a DVS. First, with the addition of a simple covert attention mechanism, we have shown that eRBP achieved state-of-the-art accuracy on the DvsGesture benchmark. This attention mechanism improved performance by a large margin, by providing translation invariance at a low computational cost compared to convolutions. Second, we integrated eRBP in a real-world robotic grasping experiment, where affordances are detected from microsaccadic eye movements and conveyed to a robotic arm and hand setup for execution.

This would enable the setup to be extended to temporal sequence learning and reinforcement learning tasks. Other functional components to make the current experiment more generic, such as reaching motion, grasping motion, depth perception and others, can be easily segregated as new neural tasks. This work paves the way towards the integration of brain-inspired computational paradigms into the field of robotics.

7.1.3 Components

C2719 Learning object affordances

C2704 Scene Representation and Understanding

7.1.4 Publications

- 1) Embodied Event-Based Random Backpropagation, Jacques Kaiser, Alexander Friedrich, J. Camilo Vasquez Tieck, Daniel Reichard, Arne Roennau, Emre Neftci, Rüdiger Dillmann (in preparation)

7.1.5 Further measures for increasing impact

- 1) Sensori-motor learning in closed-loop neurorobotics scenarios and the Neurorobotics platform. J. Camilo Vasquez Tieck, Jacques Kaiser. Tutorial session at the Capocaccia Summer School, Alghero, 22 April - 03 May 2018

7.2 Output 2: Embodied Synaptic Plasticity with Online REinforcement (SPORE) learning

Jacques Kaiser, J. Camilo Vasquez Tieck, David Kappel, Daniel Reichard, Anand Subramoney, Robert Legenstein, Wolfgang Maass, Rüdiger Dillmann

CDP5 collaboration between SP10 (FZI, P52) and SP9 (TUGRAZ, P55)

We evaluate a biologically plausible reward-learning rule based on synaptic sampling on two closed-loop robotic tasks: reaching and lane following. We show that the method is capable of learning performing policies within the course of simulated hours for both tasks. Additionally, we carried out an empirical analysis of the important hyper-parameters of Synaptic Plasticity with Online REinforcement learning (SPORE) on the performance. This analysis identifies that the learning rate and the temperature driving the stochastic processes need to be regulated for performance improvements to be retained.

7.2.1 Research and Infrastructure

We evaluated the reward-learning rule SPORE on two closed-loop robotic tasks. SPORE is an instantiation of the synaptic sampling scheme. It incorporates a policy sampling method which models the growth of dendritic spines with respect to dopamine influx. Unlike current state-of-the-art reinforcement learning methods implemented with conventional neural networks, SPORE learns online from precise spike times and is entirely implemented with spiking neurons.

We evaluated this learning rule in a closed-loop reaching and a lane following setup (Fig. 16 and 17). In both tasks, an end-to-end visuomotor policy was learned, mapping visual input to motor commands. Learning control from visual input is in general challenging due to the high dimensionality of input images and the amount of data required for learning. In our setup, visual input was encoded in Address Event Representation with a Dynamic Vision Sensor (DVS) simulation.

We identified crucial techniques to regulate SPORE learning dynamics: a weak prior and a decreasing temperature (Fig. 18 and 19). With these techniques, we showed that SPORE is capable of learning policies online for moderately difficult embodied tasks within some simulated hours. On the other hand, SPORE was not capable to train hidden layers in deep neural architectures. In future work, we will swap SPORE for a more recent learning rule derived within the SGA2 phase of CDP5.

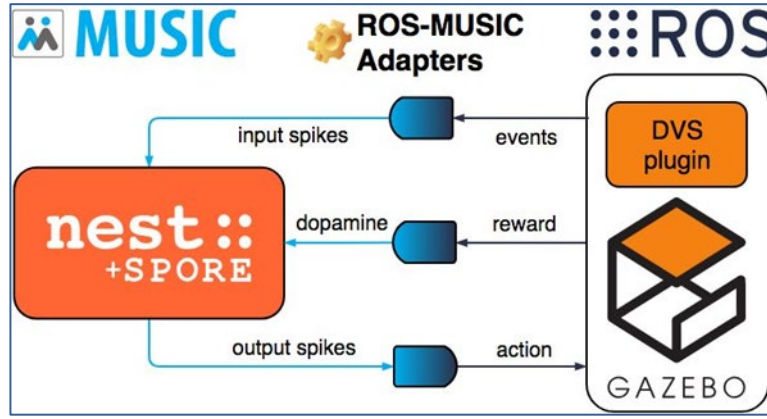


Figure 17: Our asynchronous framework based on open-source software components to evaluate SPORE.

The spiking network is implemented with the NEST neural simulator, which communicates spikes with MUSIC. The reward is streamed to all synapses in the spiking network learning with SPORE. Spikes are encoded from address events and decoded to motor commands with ROS-MUSIC tool-chain adapters. Address events are emitted by the DVS plugin within the simulated robotic environment Gazebo, which communicates with ROS.




Figure 18: Visualization of the setup for the two experiments.

Left: reaching experiment. The goal of the task is to control the ball to the center of the plane. Visual input is provided by a DVS simulation above the plane looking downward. The ball is controlled with Cartesian velocity vectors. Right: Lane following experiment. The goal of the task is to keep the vehicle on the right lane of the road. Visual input is provided by a DVS simulation attached to the vehicle looking forward to the road. The vehicle is controlled with steering angles.

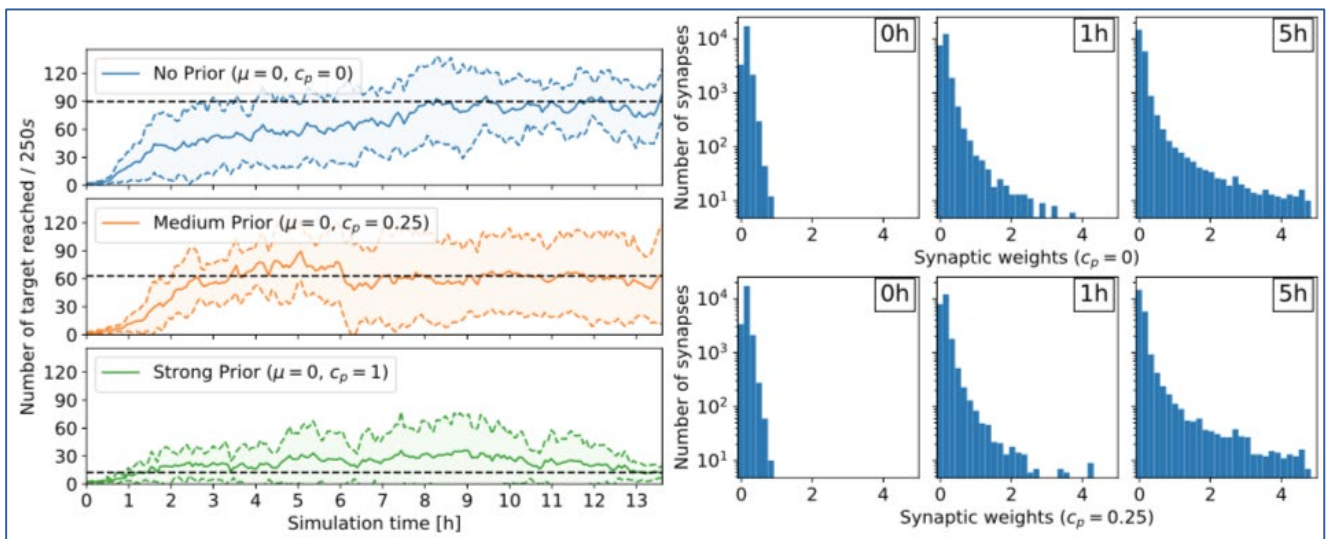


Figure 19: Results for the reaching task.

Left: comparing the effect of different prior configurations on the overall learning performance. The results were averaged over 8 trials. The performance is measured with the rate at which the target is reached (the ball moves to the centre and is reset at a random position). Right: Development of the synaptic weights over the course of learning for two trials: no prior ($c_p = 0.0$, top) and medium prior ($c_p = 0.25$, bottom). In both cases, the number of weak synaptic weights (below 0.07) increases significantly over time.

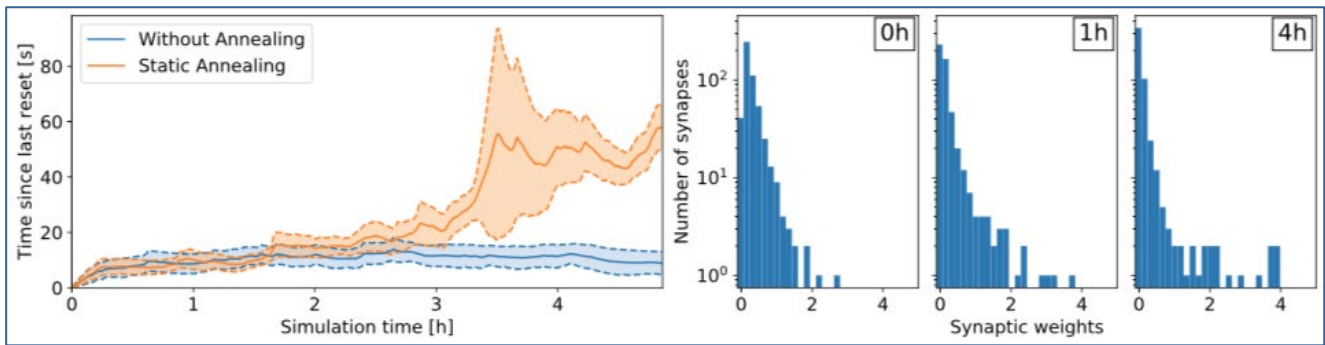


Figure 20: Results for the lane following task with a medium prior ($cp = 0.25$).

Left: comparing the effect of annealing on the overall learning performance. The results were averaged over 6 trials. Without annealing, performance improvements are not retained and the network does not learn to perform the task. With annealing, the learning rate β decreases over time and performance improvements are retained. Right: Development of the synaptic weights over the course of learning for a medium prior of $cp = 0.25$ with annealing. The number of weak synaptic weights (below 0.07) increases from 41 to 231 after 1h of learning to 342 after 4h of learning (out of 512 synapses in total).

7.2.2 Impact

The endeavour to understand the brain spans over multiple research fields. Collaborations allowing synaptic learning rules derived by theoretical neuroscientists to be evaluated in closed-loop embodiment are an important milestone of this endeavour. We evaluated the reward-learning rule SPORE (Kappel et al. (2018, 2015, 2014); Yu et al. (2016)) on two closed-loop robotic tasks. We successfully implemented a framework allowing this evaluation by relying on open-source software components for spiking network simulation Gewaltig and Diesmann (2007), Kappel et al. (2017), synchronization and communication Djurfeldt et al. (2010), Ekeberg and Djurfeldt (2008), Weidel et al. (2016), Quigley et al. (2009) and robotic simulation Koenig and Howard (2004), Kaiser et al. (2016a). Overall, we have shown that SPORE was capable of learning shallow feedforward policies online for moderately difficult embodied tasks within some simulated hours. This evaluation allowed us to characterize the influence of the prior distribution on the learned policy. Overall, deep learning methods still outperform biologically plausible learning such as SPORE on a functional scale. For future work, the performance gap between SPORE and deep learning methods should be tackled by taking inspiration from deep learning methods.

7.2.3 Components

C209 NEST - The Neural Simulation Tool

C2535 Reward-based learning for action adaption

7.2.4 Publications

1. Embodied Synaptic Plasticity with Online Reinforcement learning. Jacques Kaiser, Michael Hoff, Andreas Konle, Juan Camilo Vasquez Tieck, David Kappel, Daniel Reichard, Anand Subramoney, Robert Legenstein, Arne Roennau, Wolfgang Maass and Rüdiger Dillmann, Frontiers in Neurorobotics 2019 (submitted).
2. A dynamic connectome supports the emergence of stable computational function of neural circuits through reward-based learning. David Kappel, Robert Legenstein, Stefan Habenschuss, Michael Hsieh, Wolfgang Maass, eNeuro 5.2 (2018)

7.2.5 Further measures for increasing impact

- 1) Sensori-motor learning in closed-loop neurorobotics scenarios and the Neurorobotics platform. J. Camilo Vasquez Tieck, Jacques Kaiser. Tutorial session at the Capocaccia Summer School, Alghero, 22 April - 03 May 2018

8. Conclusion and Outlook

In SGA2, the organization of the HBP was changed significantly with respect to the Co-Design Projects (CDPs). The desire to promote co-design now stands somewhat in contrast to the absence of dedicated funding for this endeavour. However, despite such shortcomings, CDP5 remains fully on track thanks to the commitment of project partners to invest into cross-SP collaborations. The success of this collaborative undertaking is reflected by an ensemble of well-received publications, invited talks at key conferences in the field and collaborations with interested partners from science and industry outside the HBP.

In the transition from SGA1 to SGA2, CDP5 has designated biological deep learning as a focus topic. Efficient learning across neural population hierarchies represents a key prerequisite of biological intelligence. Analogously, the recent breakthroughs in machine learning have been based on a heuristic solution to the exact same problem, namely the backpropagation-of-errors algorithm. Over the past year, building on work from SGA1, CDP5 has made significant progress in bridging the gap between learning algorithms in biological and artificial hierarchical (deep) networks. A keystone of these insights is embodied by the idea of error-correcting, three-component plasticity enabled by neurons with dendritic structure. It is this morphological separation of neuronal compartments that in turn allows them to locally store multiple pieces of information required for the efficient computation of gradients at the level of individual synapses.

A further development that is of particular significance for a unified, rigorous understanding of error-gradient learning in physical neuronal systems is what we have called the Lagrangian theory of neurodynamics. Similar to the analogous approach in theoretical physics, this theory permits deriving equations of motion for plastic neural systems from a least-action postulate. The derived dynamics for synaptic weights and membrane potentials describe the behaviour of microcircuits that, on one hand, mirror known structures in cortex, while at the same time enabling error backpropagation using only local information.

These developments are further bolstered by parallel work on computational properties of experimentally well-established phenomena in biological neural networks. We have shown how spike-based computation coupled with short-term plasticity extends the predictive capabilities of hierarchical sampling networks to spatio-temporal problems. We have extended models of predictive coding to multisensory representations. Furthermore, we have studied the effect of ensemble phenomena (such as slow waves during sleep) on learning, which also creates a link to previous work on simulated tempering in spiking networks performed during SGA1.

Our efforts in theory and modeling are mirrored by work on implementing models of learning on the HBP neuromorphic and neurorobotic platforms. The results obtained with the physical model systems developed in Heidelberg demonstrate the unique capabilities of these devices, along with the surprising feasibility of precise computation based on single spikes despite substrate-inherent imperfections. The demonstration of structural plasticity on SpiNNaker also represents an important milestone by highlighting the value of resource-efficient learning for neuromorphic systems. Finally, the diversity of neurorobotics experiments achieved during the past year illustrates the maturation of interactions between roboticists and modelers, as well as the technology readiness level of the platform itself.

Throughout the remainder of SGA2, CDP5 will continue all of these multi-faceted, highly interactive and fruitful lines of research, laying the groundwork for a successful final phase of the HBP during SGA3. The work planned for the SGA3 phase has the ambition to achieve a measurable step forward in our understanding of how biological learning networks enable high-level human cognitive

functions. This will be pursued by emulating the architecture and operation of the brain that support these functions, and applying them to address cognitive problems in an embodied setting and on neuromorphic platforms. The groundwork laid in SGA2 by CDP5 will allow the development of open cognitive architectures with an unprecedented combination of biological realism and efficacy of performance for a number of selected cognitive tasks. State-of-the-art biological learning algorithms for recurrent networks will be refined to pursue human-level task performance with a reduced amount of training data. Cognitive architectures with selected functional modules will be implemented on neuromorphic hardware to illustrate both the merit of the developed architectures and advantages of NM computing on real-world physical systems such as improved energy efficiency, greater learning/execution speed, and smaller size. These goals, albeit ambitious, represent a direct continuation of our ongoing work and are only made possible by the unique, complementary expertise in the HBP, and by the growing synergies between the scientific roadmaps of its partners.

Annex A: Component Details

Component Field	Data
ID	C209
Title	NEST - The Neural Simulation Tool
Type	Software
Leader	Markus Diesmann
(Release) Name	NEST-2.16.0
Release URL	https://nest-simulator.readthedocs.io/en/latest/
Effective (Release) Date	

Component Field	Data
ID	C457
Title	BrainScaleS 2 Neuromorphic Computing System
Type	Hardware
Leader	Johannes Schemmel
(Release) Name	
Release URL	
Effective (Release) Date	

Component Field	Data
ID	C2420
Title	Plasticity: Prototype implementations of rules and testing within and without the SP9 platforms
Type	Model
Leader	André Grüning
(Release) Name	
Release URL	
Effective (Release) Date	

Component Field	Data
ID	C349
Title	PyNN
Type	Software
Leader	Andrew Davison
(Release) Name	0.9.4 - Added support for the SONATA format
Release URL	http://neuralensemble.org/docs/PyNN/
Effective (Release) Date	

Component Field	Data
ID	C2439
Title	Methods for hierarchical neural sampling in networks of spiking neurons
Type	Model
Leader	Mihai Petrovici
(Release) Name	
Release URL	
Effective (Release) Date	

Component Field	Data
ID	C2419
Title	Plasticity: algorithms for multi-compartment models
Type	Model
Leader	André Grüning
(Release) Name	
Release URL	
Effective (Release) Date	

Component Field	Data
ID	C2547
Title	Computing with structured neurons
Type	Model
Leader	Mihai Petrovici
(Release) Name	
Release URL	
Effective (Release) Date	

Component Field	Data
ID	C2722
Title	Multicompartmental pyramidal neuron models
Type	Model
Leader	Johan Storm
(Release) Name	
Release URL	
Effective (Release) Date	

Component Field	Data
ID	C2226
Title	Generative model of sensory cortical hierarchy and corticohippocampal network
Type	Model
Leader	Cyriel Pennartz
(Release) Name	v1 - The details of the model have been published in ICANN' 2018
Release URL	
Effective (Release) Date	

Component Field	Data
ID	C1032
Title	Plasticity: dendritic predictive plasticity that reproduces STDP data
Type	Model
Leader	Walter Senn
(Release) Name	
Release URL	
Effective (Release) Date	

Component Field	Data
ID	C1
Title	BrainScaleS1 Neuromorphic Computing System
Type	Hardware
Leader	Johannes Schemmel
(Release) Name	
Release URL	
Effective (Release) Date	

Component Field	Data
ID	C2
Title	SpiNNaker Neuromorphic Computing System
Type	Hardware
Leader	Andrew Rowley
(Release) Name	
Release URL	
Effective (Release) Date	

Component Field	Data
ID	C2061
Title	Multi-area ensemble mechanisms of object recognition in rodents
Type	Dataset
Leader	Cyriel Pennartz
(Release) Name	Electrophysiology_Rodent_Data_V1 - First preprocessed dataset acquired during multimodal object recognition task (V2, Bf1, Prh, Hippocampus)
Release URL	
Effective (Release) Date	03.09.2021

Component Field	Data
ID	C2228
Title	Comparison with physiological data and related models
Type	Dataset
Leader	Cyriel Pennartz
(Release) Name	
Release URL	
Effective (Release) Date	

Component Field	Data
ID	C2301
Title	Instantiate model of WhiskEye into NRP
Type	Software
Leader	Martin Pearson
(Release) Name	
Release URL	
Effective (Release) Date	

Component Field	Data
ID	C2193
Title	Cortical spiking model of the interplay between sleep and plasticity
Type	Model
Leader	Pier Stanislao Paolucci
(Release) Name	
Release URL	https://collab.humanbrainproject.eu/#/collab/45703/nav/314001
Effective (Release) Date	

Component Field	Data
ID	C2719
Title	Learning object affordances
Type	Model
Leader	Daniel Reichard
(Release) Name	
Release URL	
Effective (Release) Date	

Component Field	Data
ID	C2704
Title	Scene Representation and Understanding
Type	Model
Leader	Alois Knoll
(Release) Name	
Release URL	
Effective (Release) Date	

Component Field	Data
ID	C2535
Title	Reward-based learning for action adaption
Type	Model
Leader	Daniel Reichard
(Release) Name	
Release URL	
Effective (Release) Date	