Grant Number:	720270	Grant Title:	Human Brain Project SGA1

Deliverable Title:	D4.7.3 (D25.3 D58) CDP5 Components Report for SGA1 M13-M24
Contractual Number and type:	SGA1 D4.7.3 - Report
Dissemination Level:	PU (= PUBLIC)
Version / Date:	Submitted: 08 May 2018; ACCEPTED 09 Jul 2018

Abstract:	The paramount role of a brain is to enable an individual to learn. A particular challenge for a project devoted to study the human brain is to understand this ability to learn. In fact, the subject of self-organisation, learning and plasticity on different time scales is one of the most important and multi-faceted questions of brain science. In the HBP, the confluence of work on plasticity, learning and development happens in CDP5. All Subprojects (SPs) contribute to this goal.
Keywords:	Synaptic plasticity mechanisms, Functional learning rules, network learning, porting learning to neuromorphic systems and cortical columns, big-systems learning demos.

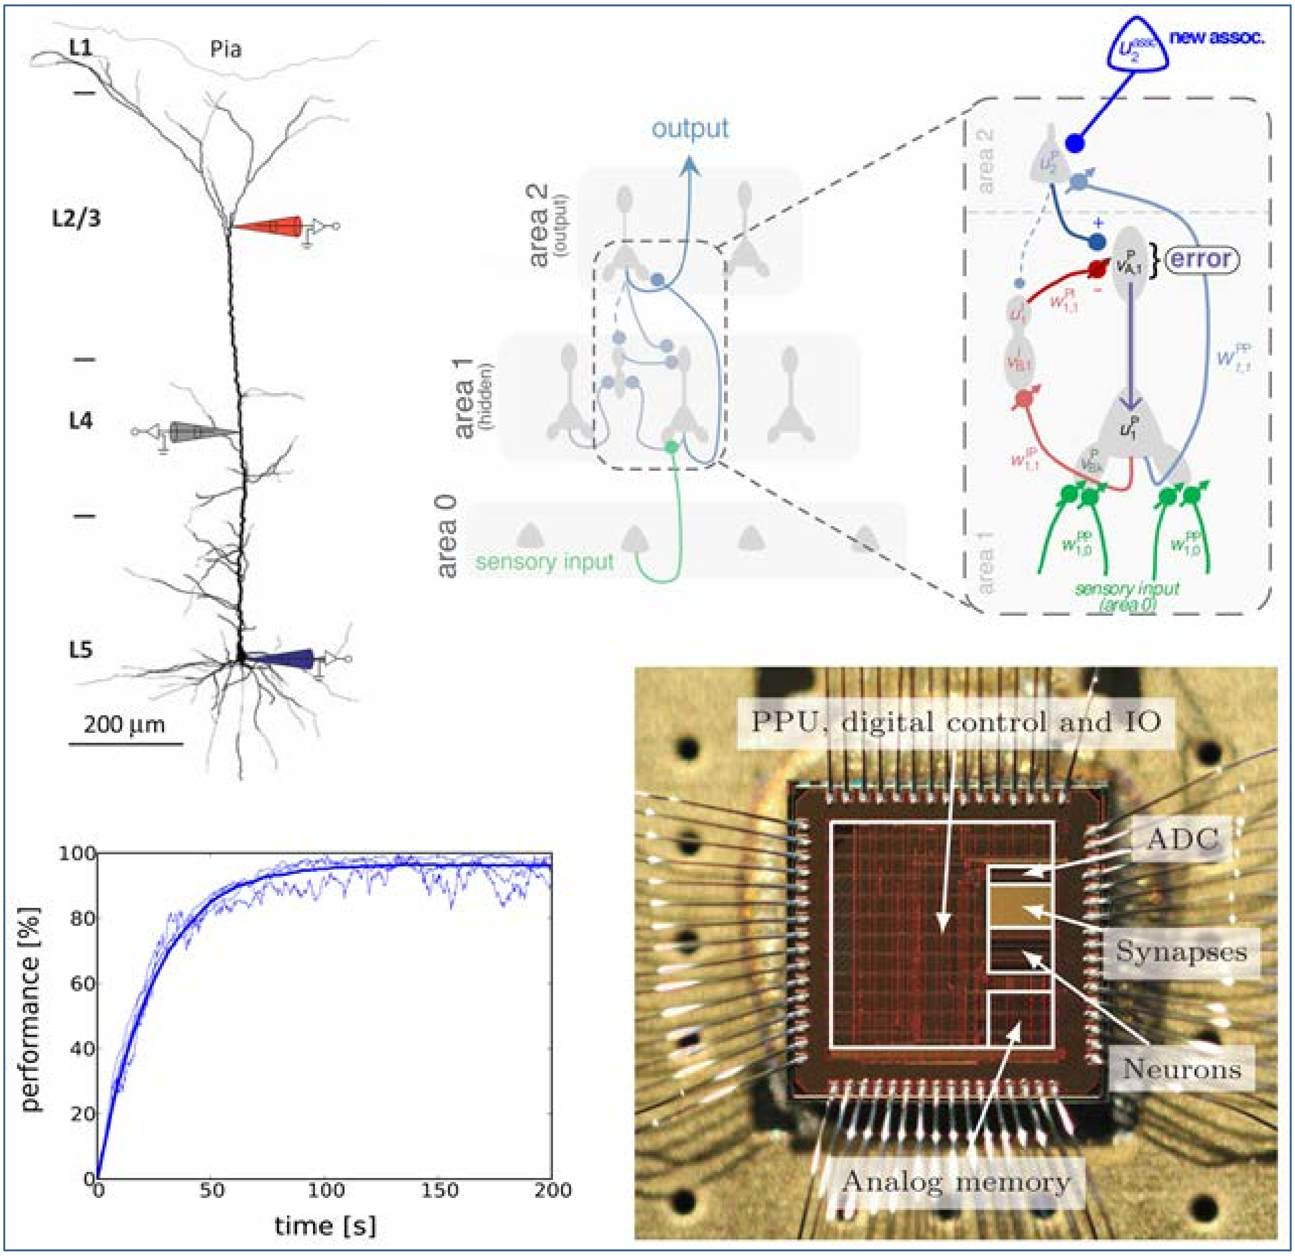

Plasticity in neural networks: from biology to silicon

Targeted users/readers	Scientific Community
Contributing Work-Package(s):	SGA1 WPs 3.2, 3.3, 3.4, 3.5, 3.6, 4.3, 4.6, 4.7, 9.1, 9.2, 9.3, 9.4, 9.5, 10.2, 10.3, 10.4, 10.5, 10.6 and 10.7
Initially Planned Delivery Date:	SGA1 M24 / 31 Mar 2018

Authors:	Authors and institutions are listed under each Key Result. They all contributed to writing the document.
Compiling Editors:	Mihai A. PETROVICI, Walter SENN, Martina SCHMALHOLZ
Contributors:	Authors and institutions are listed under each Key Result. They all contributed to writing the document.
SciTechCoord Review:	Science and Tech Coordination (SP11)
Editorial Review:	EPFL (P1): Guy WILLIS, Annemieke MICHELS

Table of Contents

List of Figures

1. Introduction

The paramount role of a brain is to enable an individual to learn. A particular challenge for a project devoted to study the human brain is to understand this ability to learn. In fact, the subject of self-organisation, learning and plasticity on different time scales is one of the most important and multi-faceted questions of brain science.

In the HBP, the confluence of work on plasticity, learning and development happens in CDP5. All Sub Projects (SPs) contribute to this goal:

- SP1, SP2 and SP5 provide and curate data on structural changes in mouse and human brain.
- Learning is the key element of cognitive neuroscience (SP3).
- The theory of learning principles represents a dedicated WP in SP4.
- The purpose of SP6 lies in the simulation of neural circuit dynamics, of which plasticity is a quintessential part.
- SP7 provides the HBP with infrastructure necessary for dynamic brain simulation.
- Failure to learn represents an important malfunction of the brain (SP8).
- A vital motivation for neuromorphic computing (SP9) lies in its promise to enable efficient learning in artificial systems.
- In order to efficiently cope with a complex, changing environment, robots must learn (SP10).

Besides brain science, another branch of science is interested in uncovering the principles of learning: artificial intelligence (AI). However, the field of AI has a different scope, as it is not confined by biological boundaries. Nevertheless, as a *Nature* editorial of 8 Feb 2018 on Hardware Upgrading stated, “Artificial intelligence is driving the next wave of semiconductor innovations”. The editorial concluded “We welcome papers that will enable computing architectures beyond von Neumann, such as components for neuromorphic chips and in memory processing. Scientists across many fields are waiting for the result….” By working on unravelling the learning ability of the mammalian brain, CDP5 implicitly addresses the advancement of AI and its desire to build fast, energy-efficient, massively parallel hardware. A key partner is therefore SP9, which is devoted to hardware implementations of learning circuitry, with the multi-core SpiNNaker system in Manchester and the physical-model BrainScaleS system in Heidelberg. The core of CDP5 connects the computational theories on learning and plasticity, developed in SP4 (WP4.3) and in part in SP9 itself (WP9.4), to the neuromorphic hardware developed in SP9.

2. Results

The CDP5 research of the past 24 months is reflected in three Key Results (KRs), each with several contributions, that have been achieved in a common effort across SPs as summarised below. These Key Results encompass theoretical studies on plasticity and learning that are envisaged in a next step for the hardware design (KR1.1-KR1.5) and/or robotics applications (KR1.6). Modelling work within SP3 on sleep and memory consolidation is being implemented in the software platform for spiking neurons (KR2.1 and KR2.2). Preparatory learning and plasticity experiments with neuromorphic hardware were combined with external devices (KR3.1) or were performed solely by the hardware (KR3.4). The emulation of networks with stochastic neurons that learn to store and recreate spatio-temporal patterns (KR3.2) represents an important step towards running similar large-scale experiments on the hardware platforms. Multi-compartmental neurons akin to cortical pyramidal neurons have been designed in the neuromorphic hardware under the guidance of experimental results in SP3 (KR3.3), and elements of spike-timing dependent plasticity were implemented in the hardware (KR3.5) and robotics platforms (KR3.6). For each Key Result, we have highlighted a selection of publications, education outreach activities that have contributed to the impact of our activities to the research community both within and outside the HBP.

Key Result 1: Brain-inspired learning algorithms

1.1 Natural gradient learning for spiking neurons (SP4 and SP9)

1.2 Sequence learning by shaping hidden connectivity (SP4 and SP9)

1.3 Error-backpropagation across cortical areas (SP4 and SP9)

1.4 Deep generative networks with spiking neurons (SP3 and SP4)

1.5 Motor learning with spiking neurons through adaptive control (SP4 and SP10)

1.6 Learning algorithms for neuromorphic hardware and robotics (SP9, SP3, SP6 and SP10)

Key Result 2: Plasticity and thalamocortical activity patterns

2.1 Slow-wave activity in a model of memory consolidation (SP3, SP6 and SP7)

2.2 Thalamocortical model simulating wakefulness, sleep, and state transitions (SP3, SP7 and SP9)

Key Result 3: Platform development and applications

3.1 In-the-loop training with neuromorphic hardware (SP9 and SP4)

3.2 Stochastic computing with spikes (SP9 and SP4)

3.3 Compartmental neurons in neuromorphic hardware (SP9, SP4 and SP3)

3.4 An embedded plasticity processor for the BrainScaleS-2 system (SP9 and SP4)

3.5 Plasticity rules in neuromorphic hardware (SP4 and SP9)

3.6 Robots learning sensory-motor loops with spiking neurons (SP10, SP9 and SP4)

2.1 Key Result 1: Brain-inspired learning algorithms

2.1.1 Natural gradient learning for spiking neurons

Elena Kreutzer, Mihai A. Petrovici, Walter Senn

CDP5 collaboration between SP4 (U Bern) and SP9 (U Heidelberg)

We suggest a model for supervised learning with spiking neurons based on the natural gradient algorithm that yields a consistent description of synaptic plasticity in the brain and is robust to fixed-pattern distortions in neuromorphic hardware.

Synaptic plasticity is known to be a key mechanism for learning in the brain. However, there are many equivalent ways to describe the strength of a synaptic connection in the brain, such as in terms of an EPSP slope, an EPSP amplitude, or the number of receptors at the synaptic cleft. In neuromorphic hardware, on the other hand, two synaptic weights that are intended to be equal are often represented in slightly different ways due to variations in construction.

In both cases, the specific choice of a synaptic weight parameterisation should not influence the learning enabled by the correspondingly transformed synaptic plasticity rule. However, classical error learning rules based on gradient descent disobey this requirement: a different parameterisation of the same model will result in a different gradient rule. U Bern suggests an alternative model for synaptic plasticity in supervised learning with spiking neurons that is based on the parameterisation-invariant natural gradient algorithm.

The latter has been successfully used in machine learning, exhibiting faster convergence by taking a more direct path to the learning target. We apply the natural gradient algorithm for the biological setting of learning with spiking neurons, and show that it allows consistent learning in spatially extended neurons, while converging faster than a classical gradient-based learning rule. Crucially, the framework also predicts biological phenomena such as heterosynaptic plasticity and a scaling of the learning rate by the variance of the presynaptic activity. Furthermore, due to its robustness against reparameterisations, the rule is well suited for implementation in analogue neuromorphic hardware (see KR3.3-KR3.5).

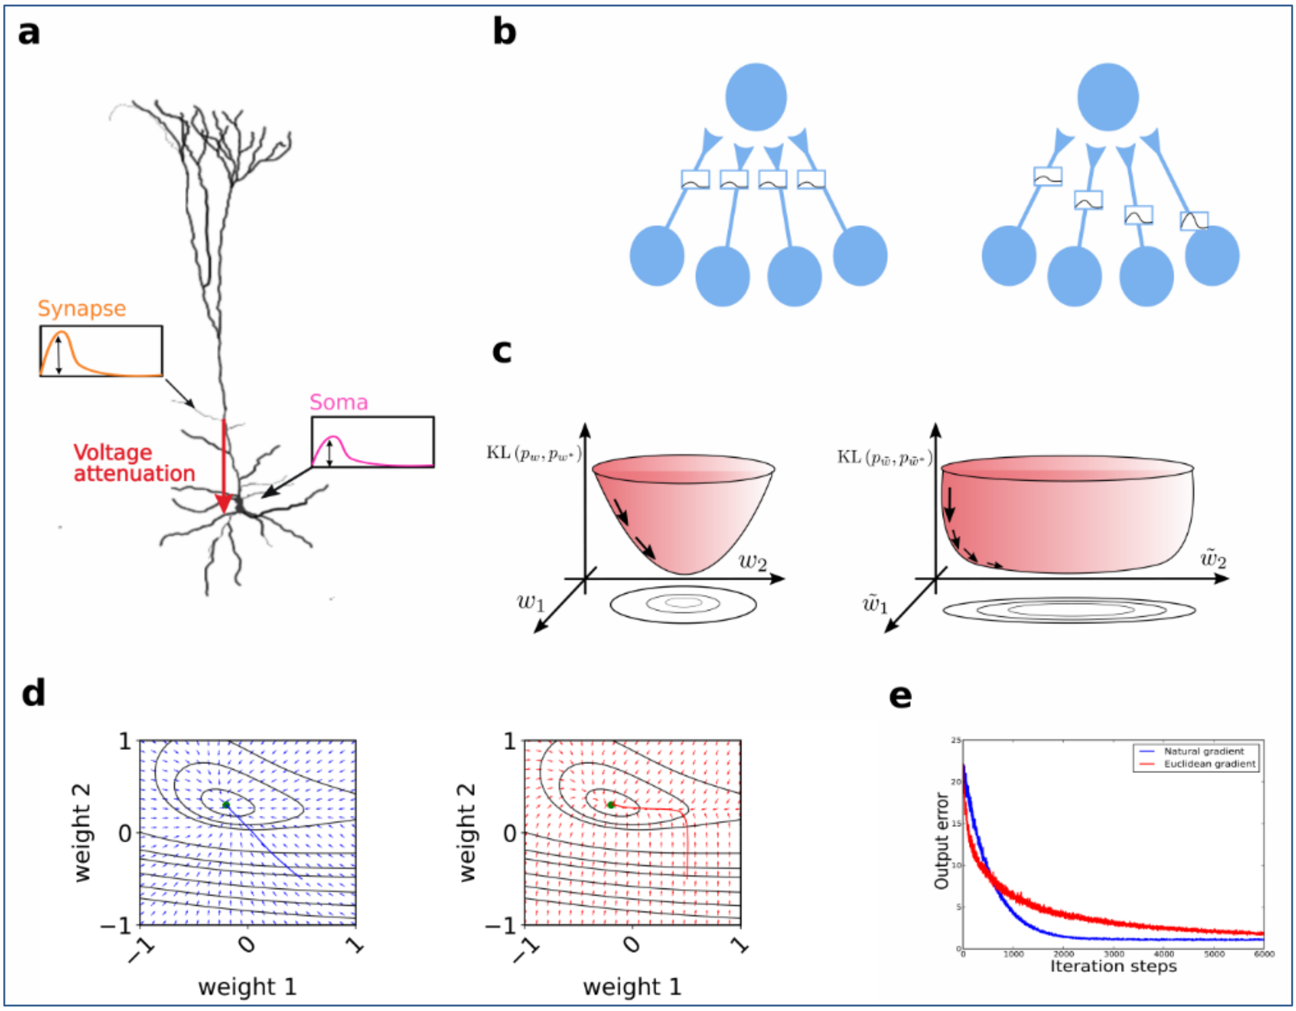

Figure 1: Synaptic strength may be parameterised in various ways.

(a, b) We show an example (a) where the synaptic weights are either parameterised as EPSP amplitudes at the soma, or as EPSP amplitude directly at the synapse. Synapses that exhibit the same EPSP amplitude at the soma (b, left) exhibit different amplitudes at the synapses (b, right) due to voltage attenuation in the dendritic tree. While classical gradient based learning will adapt somatic EPSP amplitudes equally strong in the first case, a derivation of the same rule in terms of synaptic EPSP amplitudes will change somatic amplitudes stronger for proximal than for distal synapses. The natural gradient learning rule resolves this inconsistency by adapting distal synapses stronger than proximal ones, leading to an equal change in somatic EPSP amplitudes. (c-e) Natural gradient learning is also more efficient than classical error-based learning. A non-optimal choice of parameterisation can turn a fairly symmetric error landscape (c, left), into a non-isotropic landscape with a shallow valley (c, right), a situation in which the standard gradient rule converges slowly. Whereas Euclidean gradient learning in the latter case follows the contour lines of the error function (d, left), taking a detour to the target state, natural gradient learning takes a more direct path (d, right). The learning curves (e) show that this results, on average, in faster convergence of the natural gradient rule.

2.1.1.1 Achieved Impact

1) Natural gradient for spiking neurons

Elena Kreutzer, Walter Senn

Poster at the Bernstein Conference, Berlin (Germany), 2016

2) Natural gradient for spiking neurons

Elena Kreutzer, Walter Senn

Talks and posters at the GCB Symposia, Bern (Switzerland), 2016, 2017, 2018

3) Natural gradient for spiking neurons

Elena Kreutzer, Walter Senn

Poster at the BMI Symposium "Neural Implementation of Learning Models", Lausanne (Switzerland), 2017

4) Natural gradient for spiking neurons

Elena Kreutzer, Walter Senn

Poster at the Bernese Clinical Neuroscience Day, Bern (Switzerland), 2016

5) Dendrites and plasticity: From first principles to structure and dynamics

Mihai A. Petrovici

Talk at the Dendritic integration and computation with active dendrites workshop, Paris (France), 2018

2.1.1.2 Component Dependencies

Summarised links to components this Key Result depends on:

Component ID	Component Name	HBP Internal	Comment
457	SP9: BrainScaleS 2 neuromorphic computing system (hardware)	Prototype accessible HBP internal	Hardware emulation of model in SGA2
468	SP9: Principles for brain-like computation (model)	No	CDP5-related results
969	SP4: Plasticity - Two-compartment neuron (model)	No	KR builds on and contributes to this model
1066	SP4: Plasticity - Synaptic plasticity and learning (model)	No	KR builds on and contributes to this model
2419	SP4: Plasticity - Algorithms for multi-compartment models (model)	Yes	KR builds on and contributes to this model

2.1.2 Sequence learning by shaping hidden connectivity

Kristin Völk, Mihai A. Petrovici, Walter Senn

CDP5 collaboration between SP4 (U Bern) and SP9 (U Heidelberg)

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

Behaviour can often be described as a temporal sequence of actions. These sequences are grounded in neural activity. In order for neural networks to learn complex sequential patterns, memories of past activities are required. These past activities need to be stored by hidden neurons in the network from which the 'visible' neurons can read out the memory.

We consider recurrently connected pools of visible and hidden neurons that are wired together with a small number of non-plastic somato-somatic connections and a large number of plastic somato-dendritic connections. We postulate that the few somato-somatic connections evolve stochastically during development to provide a characteristic scaffold of sequential and loopy connectivity patterns. Depending on the sequences imposed on the visible neurons during the learning process, different somato-somatic sequences and loops are activated in the hidden neurons. Based on dendritic predictive plasticity (Urbanczik & Senn, Neuron 2014), the many somato-dendritic connections evolve to stabilise the appropriated combination of the sequences and loops that support the sequential of activity imposed on the visible neurons. The stored sequential memories can be exploited by new visible sequences to be learned by slightly rewiring the somato-dendritic connections. The model provides hypotheses of how the cortical or hippocampal connectivity patterns evolve during development and the subsequent learning periods. Because all computations are local both in the development and in the learning phase, this model is also amenable to implementation in physical model hardware (see KR3.3-KR3.5).

The work is an extension of our previous model on prospective coding, Brea, Gaal, Urbanczik & Senn, Prospective Coding by Spiking Neurons, PLoS Comp Biol 2016 (http://dx.plos.org/10.1371/journal.pcbi.1005003) with code available at https://github.com/unibe-cns/prospectiveCoding

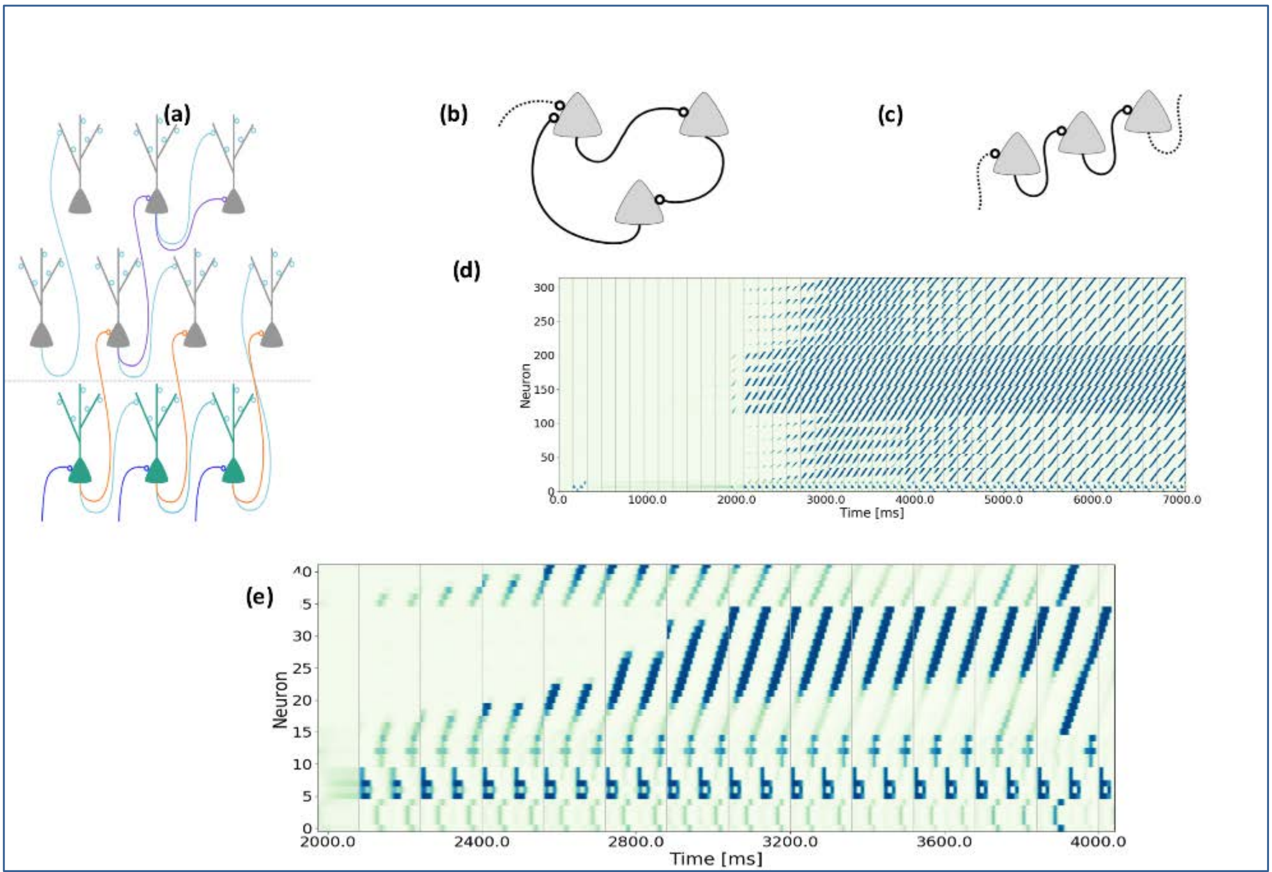

Figure 2: Developmental model of cortical pyramidal neurons for sequence learning.

(a) Schematic illustration of the framework. Some neurons (in the `visible' layer at the bottom) will receive external connections (blue), which represent the desired activation pattern. They will generate memory traces of past activities via somato-somatic connections to other neurons (orange and violet). (b, c) Schematic illustration of a delay line and a loop - two structures which enable the storage of memories over longer periods of time. (d) Whole network simulation of a network learning the sequence "b, c, b, f" in the visible population. Development of delay lines in the "hidden" neuron pool can be observed. (e) Zoomed-in version of (d), illustrating the learning phase during the development of delay lines.

2.1.2.1 Achieved Impact

6) Time series learning through hidden population shaping by somatic nudging

Kristin Völk, Walter Senn

Poster at the Bernstein Conference, Göttingen (Germany), 2017

5) Dendrites and plasticity: From first principles to structure and dynamics

Mihai A. Petrovici

Talk at the Dendritic integration and computation with active dendrites workshop, Paris (France), 2018

2.1.2.2 Component Dependencies

Summarised links to components this key result depends on:

Component ID	Component Name	HBP Internal	Comment
457	SP9: BrainScaleS 2 neuromorphic computing system (hardware)	Prototype accessible HBP internal	Hardware emulation of model in SGA2
468	SP9: Principles for brain-like computation (model)	No	CDP5-related results
969	SP4: Plasticity - Two-compartment neuron (model)	No	KR builds on and contributes to this model
1066	SP4: Plasticity - Synaptic plasticity and learning (model)	No	KR builds on and contributes to this model
2419	SP4: Plasticity - Algorithms for multicompartment models (model)	Yes	KR builds on and contributes to this model

2.1.3 Error backpropagation across cortical areas

João Sacramento, Rui Ponte Costa, Dominik Dold, Mihai A. Petrovici, Yoshua Bengio, Walter Senn

CDP5 collaboration between SP4 (U Bern), SP9 (U Heidelberg) and external partner (U Montreal, YB)

Inspired by a recurring connectivity motif found across brain areas, CDP5 researchers from U Bern and U Heidelberg in collaboration with U Montreal have developed a network mechanism to encode and propagate prediction errors. The proposed model approximates the backpropagation-of-errors algorithm widely used in artificial neural networks. Unlike its artificial counterpart, however, it operates continuously and uses only local synaptic plasticity rules, potentially uncovering a way of learning in deep biological and *in silico* neuronal networks.

Cortical networks of biological neurons, as well as state-of-the-art artificial neural models of pattern recognition and generation (termed deep neural networks), comprise multiple stages of interconnected processing elements. Learning in such networks involves determining changes to synaptic connections in order to achieve better performance in a given task, such as visual object recognition. This is the so-called credit assignment problem. A simple, but highly effective strategy employed in artificial neural network training is to define an appropriate cost function and then let synaptic changes follow gradient descent dynamics on this cost. Such gradients can be efficiently computed using the backpropagation-of-errors algorithm.

In neuroscience, few solutions to the credit assignment problem have been put forth so far, and computational models typically focus on solving simple tasks where learning is more straightforward. Drawing inspiration from cortical architecture, researchers from U Bern and U Heidelberg in collaboration with U Montreal have now developed a model for the computation and backpropagation of prediction errors which relies on a novel circuit motif. Such circuitry allows the network to perform credit assignment. Two key design elements are introduced: First, the network uses simplified compartmental neurons modelled after pyramidal cells, which continuously integrate bottom-up sensory input with top-down `semantic' input. The bottom-up input mainly projects to the basal tree of the model pyramidal neurons, and the top-down input, originating downstream processing areas, projects to the apical tree. Second, the top-down feedback signals are converted within the apical tree into neuron-specific prediction errors using lateral inhibition. The lateral inhibition (e.g. mediated by somatostatin interneurons) learns to explain away the top-down input, and the remaining apical voltage represents the prediction error that drives synaptic plasticity of the sensory input targeting the basal tree.

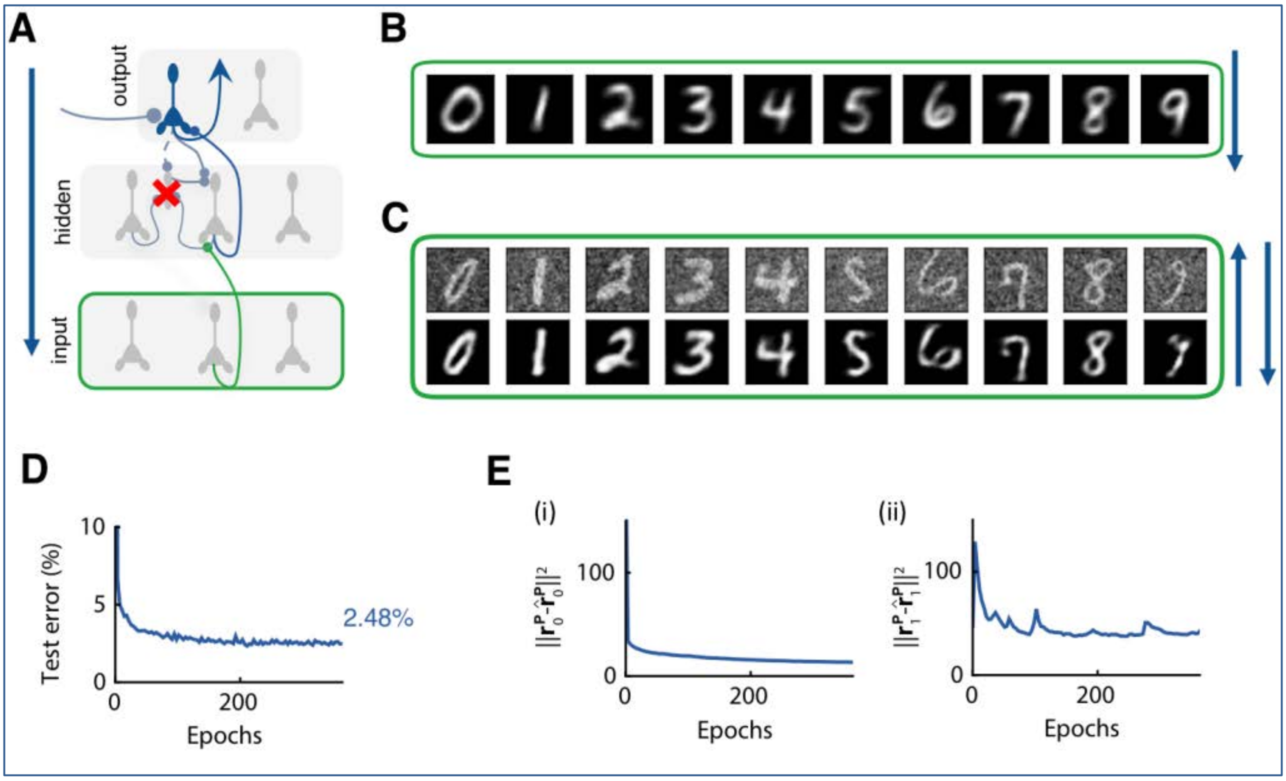

Figure 3: Top-down synapses can be adapted to simultaneously drive bottom-up learning, input construction and de-noising.

(A) Classification performance of a 784-1000-10 network exposed to MNIST images, with plastic top-down synapses that learn to predict lower-area activities. Top-down and forward weights co-adapt without pauses or phases. (B) Driving the network top-to-bottom (i.e. initialising the output area to a particular digit and turning off lateral and bottom-up inputs of both hidden and input areas) recreates class-specific image examples in the input area. The top-down connections can be tuned to encode a simple inverse visual model. (C) Such an inverse model yields image de-noising, which can be achieved by reconstructing corrupted inputs from hidden area activities. (D) The network also successfully learns to classify images. (E) Inverse reconstruction losses of original images (i) and hidden (ii) neuron activities. Top-down synapses connecting hidden pyramidal neurons back to the input area learn to reconstruct pixel arrangements given hidden neuron activities; synapses originating at the output area learn to predict hidden area activities given the current class label estimate.

The computational role proposed for the dendritic compartmentalisation of pyramidal cells, as well as the proposed distinct cell types, lead to a number of experimental predictions regarding learning and plasticity. Furthermore, being capable of operating in real-time and relying only on local plasticity rules (see KR3.3-KR3.5), the model suggests potential ways of learning in an online fashion deep neural networks that can be implemented in the HBP neuromorphic platforms (see

KR3.1 and KR3.2). The code for this model is available at https://github.com/unibe-cns/multiareaBackprop.

2.1.3.1 Achieved Impact

7) Dendritic error backpropagation in deep cortical microcircuits

João Sacramento, Rui Ponte Costa, Yoshua Bengio, Walter Senn

arXiv:1801.00062, 2017

8) Bayesian multisensory integration by dendrites

João Sacramento, Walter Senn

Talk at the COSYNE Conference, Salt Lake City (USA), 2016

9) Dendritic functions in learning

João Sacramento, Walter Senn

Selected young researcher talk at the DENDRITES Conference, Heraklion (Greece), 2016

10) Learning based on error representations in apical dendrites of L5 pyramidal neurons

João Sacramento, Walter Senn

Talk at the Champalimaud Neuroscience Symposium, Lisbon (Portugal), 2016

11) Error backpropagation in cortical circuits

João Sacramento, Yoshua Bengio, Walter Senn

Selected young researcher talk at the 5th Annual Human Brain Project Summit, Glasgow (UK), 2017

5) Dendrites and plasticity: From first principles to structure and dynamics

Mihai A. Petrovici

Talk at the Dendritic integration and computation with active dendrites workshop, Paris (France), 2018

2.1.3.2 Component Dependencies

Summarised links to components this Key Result depends on:

Component ID	Component Name	HBP Internal	Comment
457	SP9: BrainScaleS 2 neuromorphic computing system (hardware)	Prototype accessible HBP internal	Hardware emulation of model in SGA2
468	SP9: Principles for brain-like computation (model)	No	CDP5-related results
969	SP4: Plasticity - Two-compartment neuron (model)	No	KR builds on and contributes to this model
1066	SP4: Plasticity - Synaptic plasticity and learning (model)	No	KR builds on and contributes to this model
2622	SP4: Plasticity - Predictive plasticity for deep learning (model)	No	KR builds on and contributes to this model

2.1.4 Deep generative networks with spiking neurons

Shirin Dora, Sander Bohte, Cyriel Pennartz, Walter Senn

CDP5 collaboration between SP3 (U Amsterdam) and SP4 (U Bern).

Biological deep neural network models with spiking neurons are investigated that learn both to classify objects in a visual scene and to re-generate the visual input. The model will be extended towards multisensory integration and is inspired by data from SP3 (U Amsterdam). It provides a guideline for implementations in the HBP platforms.

Besides local plasticity rules, biological deep learning models also need to take into account the architectural constraints observed in the brain, such as the retinotopic organisation of the primary visual cortex, and replicate the response properties of biological neurons. In 12), we trained a hierarchical neural network model that employed neurons with overlapping receptive fields. The model utilises predictive coding to infer latent representations for a given input stimulus at each level in the hierarchy. These latent representations inferred by the model represent the hidden hypothetical causes in the external environment that generated a given stimulus. The latent representations at higher levels in the model denote more abstract information in comparison to lower levels in the model. We have shown that the latent representations at all levels in the model could be used to reconstruct the original stimulus.

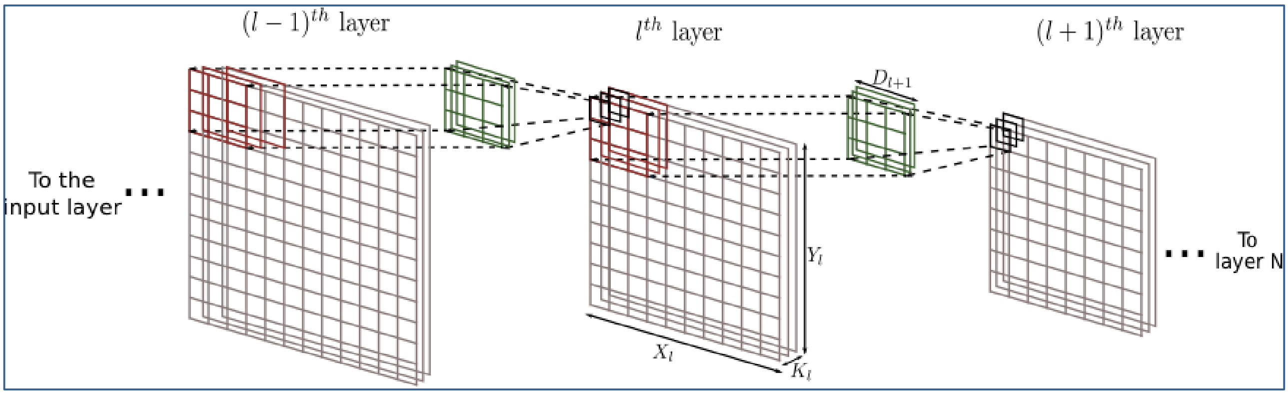

Figure 4: Architecture of a deep predictive coding network with retinotopic organisation.

As observed in the brain, the deeper layers in the model have larger receptive fields in comparison to the lower layers in the model. The model was trained on only 1000 images of horses and ships from the CIFAR-10 data set but we were able to infer latent representations for images of other objects in the data set that were not presented during training. This highlights the generalisation capabilities of the trained model across different classes of physical objects. Furthermore, we were also able to infer latent representations for translated versions of the images. This shows that the features learned by the model are translation-invariant.

Currently, the model is trained on unisensory stimuli (images from the CIFAR-10 data set). In SGA2, the goal will be to train and infer hierarchical representations in multisensory systems, simulating encoding and cross-modal recall and link sensory processing to episodic memory. A collaboration with the Senn lab at the University of Bern further targets the modelling of experimental data. An implementation in the physical-model BrainScaleS hardware has been expedited by small-scale implementations of hierarchical spiking networks (see KR3.1 and KR3.2).

2.1.4.1 Achieved Impact

12) Deep Predictive Coding Networks for Learning Latent Representations

Shirin Dora, Cyriel M.A. Pennartz, Sander M. Bohte

OpenReview.net, 2017

13) Predictive coding in deep neural networks

Shirin Dora, Sander M. Bohte and Cyriel M.A. Pennartz

Poster at the In-depth meeting on network learning, Fürberg (Austria), 2017

2.1.4.2 Component Dependencies

Summarised links to components this Key Result depends on:

Component ID	Component Name	HBP Internal	Comment
468	SP9: Principles for brain-like computation (model)	No	CDP5-related results
752	SP3: Analysis of network-level mechanisms constraining the in vivo implementation of learning rules and implementing integration, encoding and recall of multisensory memories (data)		KR constrained by this data
1066	SP4: Plasticity - Synaptic plasticity and learning (model)	No	KR builds on and contributes to this model
2622	SP4: Plasticity - Predictive plasticity for deep learning (model)	No	KR builds on and contributes to this model

2.1.5 Motor learning with spiking neurons through adaptive control

Aditya Gilra, Wulfram Gerstner, Walter Senn, J. Camilo Vasquez Tieck

CDP5 collaboration between SP4 (EPFL-LCN Lausanne, U Bern) and SP10 (FZI Karlsruhe)

For predicting and controlling non-linear dynamics using spiking neural networks, EPFL-LCN developed a local synaptic plasticity rule and network architecture, borrowing from adaptive control theory. The work hypothesises a neural basis for motor learning and control in the brain. Direct applicability to neuromorphic computing and neuro-robotics is envisaged. This forms the basis for a CDP5 collaboration within SP4 (EPFL-LCN, U Bern). The mode will be implemented in a closed-loop robotic experiment with partners from SP10 and serves as a guideline for the development of plasticity rules on the SpiNNaker hardware platform (SP9).

EPFL-LCN has proposed a learning scheme for neuronal networks to predict and control movement Gilra and Gerstner, 2017a, b), using spiking neurons adapting synaptic weights with local plasticity rules. The implementation of these “FOLLOW” learning networks on SpiNNaker and Neurorobotics Platform will contribute to several use cases of the SGA2 phase of CDP5. A brief overview of the learning scheme is provided below.

The brain needs to construct forward or inverse models of the non-linear dynamics of muscles, limbs and the external world for motor control and planning. How spiking neural networks can learn such models is an open problem, despite significant progress via reservoir computing, FORCE learning, and other methods Abbott *et al*., 2016, DeWolf *et al*., 2016, Denéve *et al*., 2017).

We proposed a self-supervised learning scheme (Figure 5): Feedback-based Online Local Learning Of Weights (FOLLOW) Gilra and Gerstner, 2017a), which especially draws from function and dynamics approximation theory Funahashi, 1989, Eliasmith and Anderson, 2004) and adaptive control theory Ioannou and Sun, 2012). Using our FOLLOW scheme, we enabled a recurrently connected network of heterogeneous spiking neurons to learn its feedforward and recurrent weights, so as to predict or control a low-dimensional non-linear dynamical system $dx/dt = f(x, u)$, where $u(t)$ is the control input and $x(t)$ are the state variables. We derived the learning rules showing global uniform (Lyapunov) stability with the error tending to zero asymptotically, under reasonable assumptions and approximations. The learning rules are synaptically local, involving the pre-synaptic firing rate and an error feedback current injected into the post-synaptic neuron.

Using a two-link arm as an example, we showed that our network learned a forward predictive model for motor planning i.e. it predicted the joint angles and velocities x(t) given joint torques u(t), or learned an inverse model for motor control i.e. it inferred the torque u(t) that would generate a desired state trajectory x(t). We further used the inverse model to control the arm to draw on a wall.

With FOLLOW learning, we proposed a more biologically plausible, specifically synaptically local, scheme of how the brain may learn forward and inverse models to perform motor planning and control. Extensions like incorporating Dale's law for further biological plausibility, hierarchical coding and control, semi-supervised learning, and applications to neuromorphic computing and neurorobotics are planned for future work (see KR3.6).

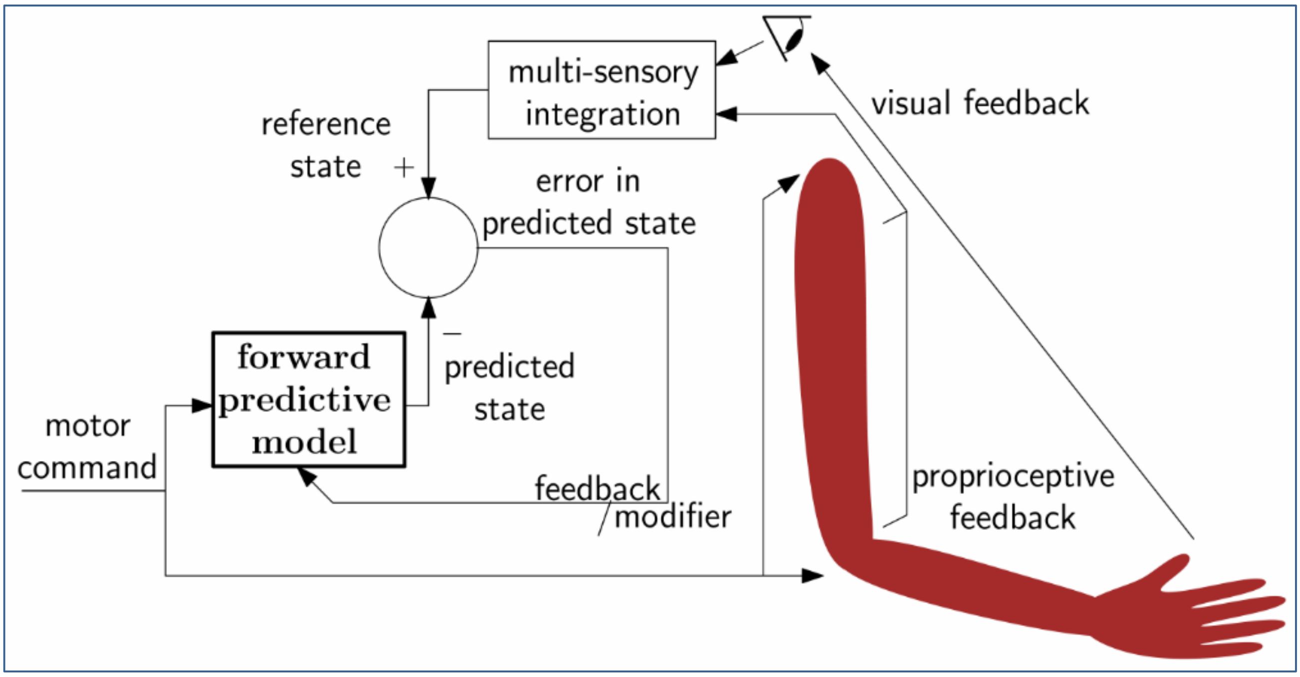

Figure 5: Schematic for self-supervised learning of forward model.

During learning, early in development, random motor commands (motor babbling) cause movements of the arm. An efference copy of the motor commands are also sent to the forward predictive model, which must learn to predict the positions and velocities (state variables) of the arm. The deviation of the predicted state from the reference state, obtained by visual and proprioceptive feedback, is used to learn the forward predictive model. The forward predictive model is a network of recurrently connected neurons whose connection strengths are adjusted by a synaptically local rule.

2.1.5.1 Achieved Impact

14) Predicting non-linear dynamics by stable local learning in a recurrent spiking neural network

Aditya Gilra and Wulfram Gerstner

eLife 6:e28295. DOI: 10.7554/eLife.28295, 2017

15) Non-linear motor control by local learning in spiking neural networks

Aditya Gilra and Wulfram Gerstner

arXiv:1712.10158 q-bio.NC)

16) A stable local learning scheme for recurrent spiking neural networks

Aditya Gilra and Wulfram Gerstner

Talk at the workshop "Machine Learning Meets Biology: Algorithms and Cortical Mechanisms", Geneva (Switzerland), 2017

14a) Predicting non-linear dynamics: a stable local learning scheme for recurrent spiking neural networks

Aditya Gilra and Wulfram Gerstner

Talk at the In-depth meeting on network learning, Fürberg (Austria), 2017

2.1.5.2 Component Dependencies

Summarised links to components this Key Result depends on:

Component ID	Component Name	HBP Internal	Comment
468	SP9: Principles for brain-like computation (model)	No	CDP5-related results
1066	SP4: Plasticity - Synaptic plasticity and learning (model)	No	KR builds on and contributes to this model

2.1.6 Learning algorithms for neuromorphic hardware and robotics

Guillaume Bellec, Zeno Jonke, David Kappel, Robert Legenstein, Wolfgang Maass, Dejan Pecevski, Christoph Pokorny, Arjun Rao, Anand Subramoney

CDP5 collaboration between SP9 (TU Graz, U Heidelberg, TU Dresden), SP3 (U Amsterdam), SP6 (KTH) and SP10 (FZI Karlsruhe)

Experimental data from neuroscience provided the basis for the design of new learning methods for networks of spiking neurons, in particular for neuromorphic hardware and neurorobotics.

Wolfgang Maass (SP9), in collaboration with Jeannette Hellgren-Kotaleski (SP6), organised an interdisciplinary CDP5 Brainstorming Workshop of SP1, SP2, SP4, SP6, SP9, SP10 on "Cellular Determinants of Functional Network Plasticity" in Fürberg, Austria, October 2016. In particular, results from molecular biology on intracellular processes in the postsynaptic density of synapses were examined from the systems perspective of network plasticity. A concrete outcome of these interdisciplinary discussions was a new perspective of the likely role of a key molecule in the postsynaptic density, CAMKII, for network plasticity 17). In addition, results and discussions on various long lasting processes on the molecular level in neurons and synapses provided new mechanisms for network learning (work in progress at TU Graz).

In addition, Wolfgang Maass (SP9) in collaboration with Cyriel Pennartz (SP3), organised a CDP5 Brainstorming Workshop on "Functional Network Plasticity" in Fürberg, Austria, May 2017. This workshop produced discussions and insights about network plasticity from several disciplines. Among the new results that were discussed at this workshop, and later results that benefited from discussions at this workshop, were 18), 19), 20), 21), 22), 23). Another collaboration between SP3 and SP9, which also started at this workshop, addresses the role of dendritic spikes for the interaction of bottom-up input and top-down feedback in cortical columns (work in progress).

The study 23) provides brain-derived principles for integrating synaptic rewiring into network plasticity. This forms the basis of a current collaboration between TU Graz and TU Dresden on an implementation of reward-based rewiring on the SANTOS chip for Spinnaker 2. It has also lead to a method for deep learning with sparse connectivity 24). Obviously, sparse connectivity is an essential prerequisite of efficient hardware or software implementations of deep learning. The relevance of this new method is highlighted by the fact that 24) has been accepted at a very competitive conference in Machine Learning (International Conference on Learning Representations 2018). The method from 23) is currently applied at FZI Karlsruhe for reward-based learning and rewiring in neural robot controllers. An already completed collaboration with FZI Karlsruhe is reported in 25).

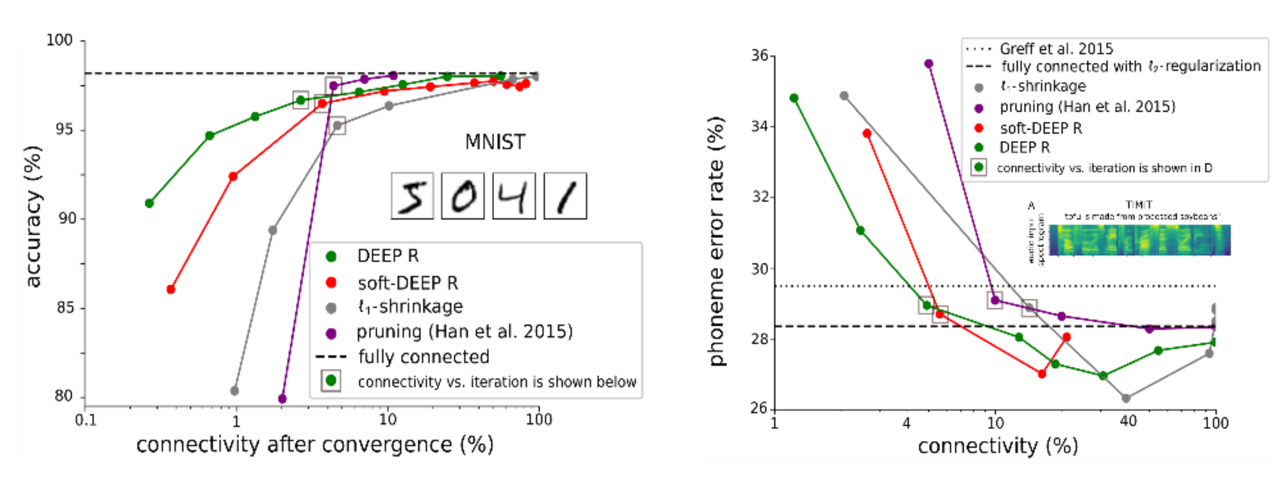

Figure 6: Deep learning for artificial neural networks

This can be performed with rewiring (DEEP R) in a sparsely connected network with just a small loss in performance; comparison with other algorithms. Left: Accuracy against connectivity for MNIST. Right: performance on the TIMIT dataset.

2.1.6.1 Achieved Impact

17) CaMKII activation supports reward-based neural network optimization through Hamiltonian sampling

Z. Yu, D. Kappel, R. Legenstein, S. Song, F. Chen, and W. Maass

arXiv:1606.00157 (2017)

18) Learning probabilistic inference through STDP

D. Pecevski and W. Maass

eNeuro (2016)

19) Feedback inhibition shapes emergent computational properties of cortical microcircuit motifs

Z. Jonke, R. Legenstein, S. Habenschuss, and W. Maass

Journal of Neuroscience, 37(35):8511-8523 (2017)

20) A probabilistic model for learning in cortical microcircuit motifs with data-based divisive inhibition

R. Legenstein, Z. Jonke, S. Habenschuss, and W. Maass

arXiv:1707.05182 (2017)

21) Associations between memory traces emerge in a generic neural circuit model through STDP

C. Pokorny, M. J. Ison, A. Rao, R. Legenstein, C. Papadimitriou, and W. Maass

bioRxiv:188938 (2017)

22) Assembly pointers for variable binding in networks of spiking neurons

R. Legenstein, C. H. Papadimitriou, S. Vempala, and W. Maass

arXiv preprint arXiv:1611.03698 (2017)

23) Reward-based stochastic self-configuration of neural circuits

D. Kappel, R. Legenstein, S. Habenschuss, M. Hsieh, and W. Maass

arXiv:1704.04238 (2017)

24) Deep rewiring: training very sparse deep networks

G. Bellec, D. Kappel, W. Maass, and R. Legenstein

International Conference on Learning Representations (ICLR), (2018)

25) Scaling up liquid state machines to predict over address events from dynamic vision sensors

Kaiser, J., Stal, R., Subramoney, A., Roennau, A. and Dillmann, R.

Bioinspiration & biomimetics, 12(5), p.055001 (2017)

2.1.6.2 Component Dependencies

Summarised links to components this Key Result depends on:

Component ID	Component Name	HBP Internal	Comment
1	SP9: BrainScaleS 1 neuromorphic computing system (hardware)	Publicly accessible via Collab	Hardware emulation environment
2	SP9: SpiNNaker neuromorphic computing system (hardware)	Publicly accessible via Collab	Hardware emulation environment
450	SP9: SpiNNaker 48-chips standalone system (hardware)	Boards publicly available for purchase or loan	Hardware emulation environment
451	SP9: SpiNNaker 4-chips standalone system (hardware)	Boards publicly available for purchase or loan	Hardware emulation environment
457	SP9: BrainScaleS 2 neuromorphic computing system (hardware)	Prototype accessible HBP internal	Hardware emulation environment
467	SP9: SpiNNaker 2 small-scale NM-MC system (hardware)	Prototype accessible HBP internal	Hardware emulation environment
468	SP9: Principles for brain-like computation (model)	No	CDP5-related results
1066	SP4: Plasticity - Synaptic plasticity and learning (model)	No	KR builds on and contributes to this model
1342	CDP5: Guiding platform design on functional plasticity (model)	No	CDP5 result

2.2 Key Result 2: Plasticity and thalamocortical activity patterns

2.2.1 Slow-wave activity in a model of memory consolidation

Pier Stanislao Paolucci, Maurizio Mattia, Fabrizio Capuani, Elena Pastorelli, Markus Diesmann, Hans Ekkehard Plesser

CDP5 collaboration between SP3 (INFN Rome, ISS Rome) and SP6/SP7 (FZ Jülich)

Preparing the study of the interplay between slow waves and plasticity.

Sleep is a physiological state in which periodically the brain falls expressing at its early stage a rhythmic activity at around 1 Hz. This is a rather stereotypical activity, a default activity pattern (Sanchez-Vives *et al*., Neuron 2017), which is invariantly expressed across animal species with widely different phylogenetic roots (Hobson *et al*., Nat Rev Neurosci 2009). Such invariance across evolution highlights the importance of the sleep function, which in turn is manifold ranging from neocortical maintenance, energy conservation and memory consolidation (Siegel, Nature 2005). In particular, the latter implies an impact of sleep slow rhythms in shaping/rewiring the cortical network synaptic matrix. More specifically, synapses appear to be downscaled after a sleep cycle with the exception of the largest 20% (De Vivo *et al*., Science 2017). Such global phenomena may in turn have a causal role in redistributing spontaneous spiking activity expressed during the whole sleep period (Watson *et al*., Neuron 2016). Thus, understanding the impact of synaptic plasticity during sleep-like brain states is of fundamental importance to understand its capability to contribute both to the self-consistency of this rhythm and to the impact on reinforcing memory engrams.

During SGA1, INFN and ISS improved both the simulation engine and the cortical models by showing that i) large–scale simulations of SWA are possible when synaptic plasticity is incorporated; ii) the conditions under which SWA is self-consistently expressed have been identified; and iii) the simulations are `affordable in terms' in terms of computational cost.

We started by refining a procedure, based on mean-field theory, to generate a cortical model with memories capable of sustaining SWA matching experimental *in vitro* measures of the phenomenon (Capone *et al*., Cereb Cortex 2017). The rhythmic alternation between Up and Down local activities levels gives rise to travelling waves (SWA), when coupled to a near neighbourhood lateral connectivity. Together this cortical network model and the simulation engine were improved to enable fast, scalable and efficient simulation of models expressing both SWA and irregular asynchronous activity. Tested model networks included up to tens of billions of synapses interconnecting up to tens of millions of neurons simulated on platforms including up to thousands of hardware cores and software processes. Neurons have been arranged in spatial grids (composed of up to 96x96 cortical modules, grid step 400 μm). The modules have been interconnected assuming a connection probability decaying with the distance with a decay length λ = 240 μm compatible with biological values (Pastorelli *et al*., PDP 2018). For this activity on such large systems a fast simulation engine as DPSNN was instrumental.

The preparation of the study of the interplay between SWA and memories requires plastic synapses. Therefore, we improved DPSNN by adding the classic spike-timing-dependent plasticity (STDP) described in (Gütig *et al*., J Neurosci, 2003). This STDP model was chosen as being a flexible two-factor STDP and being implemented also in NEST (Gewaltig and Diesmann, 2007). We observed that the computational cost of this simple plastic model is not prohibitive for the DPSNN simulator. Simulation speed on one thousand cores is about 1ns per synaptic /event for non-plastic models and preliminary measure show a decrease of about a factor 3 when plasticity is switched on. The simulations of the cortical model in SWA showed that, even in the presence of an evident effect of oscillations on synaptic weights (Figure 8), the rhythmic multiscale activity patterns generated by the cortical systems with plastic synapses is stable for a relatively long transient period (Figure 9).

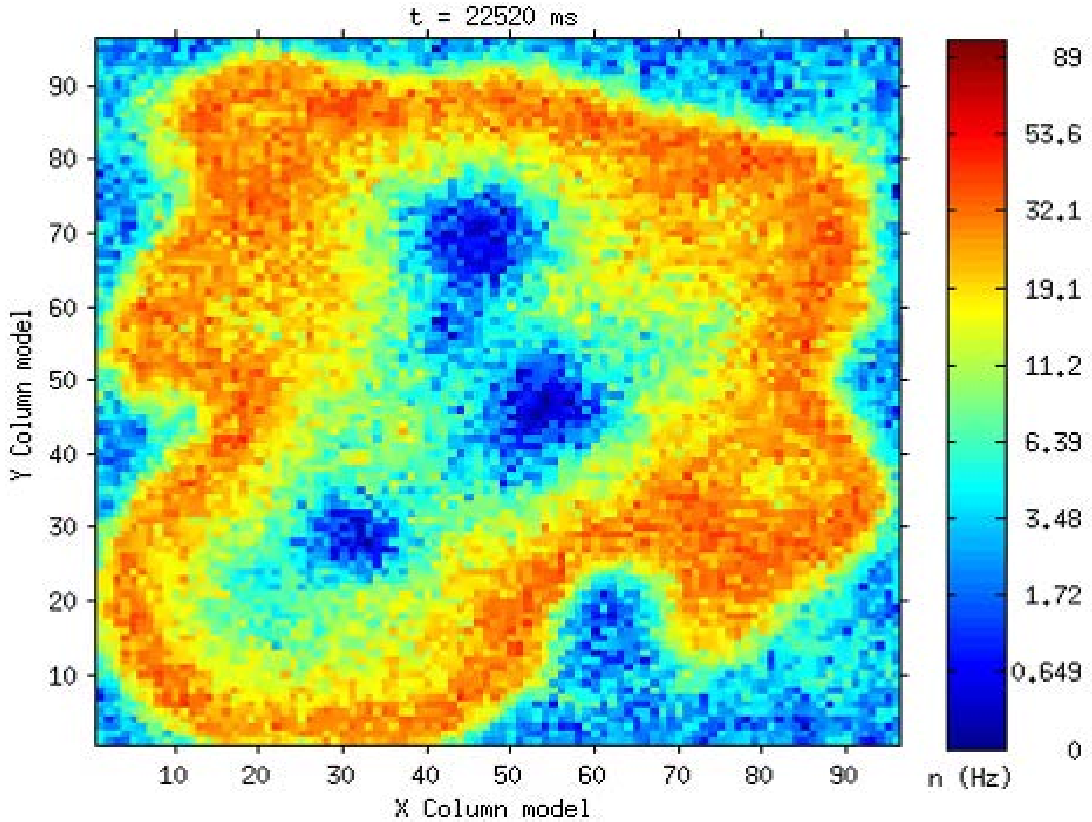

Figure 7: Snapshot of a slow wave propagating on a 96x96 grid of cortical modules

Grid step 400 µm. The model includes a total of 11.5M neurons and 17.5G.

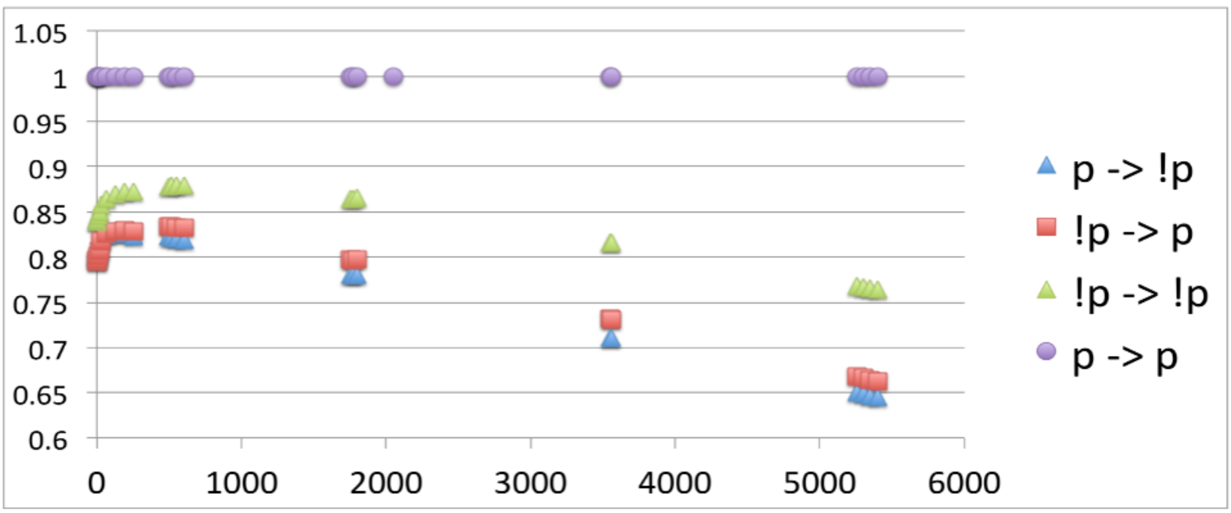

Figure 8: Time variation of relative synaptic strengths of a cortical column undergoing oscillatory activity.

Time variation of synapses connecting the population of neurons participating to the up-states (circles), neuron in population not participating to up-states (green triangles), or two neurons where the presynaptic population participates to the up-state, while the postsynaptic does not (blue triangles), or vice-versa (red squares). All synapses strengths are normalised by the strength of the synapses connecting two neurons participating to the oscillation. In the caption, p stands for participating in the up-state, while not p (!p) for not participating.

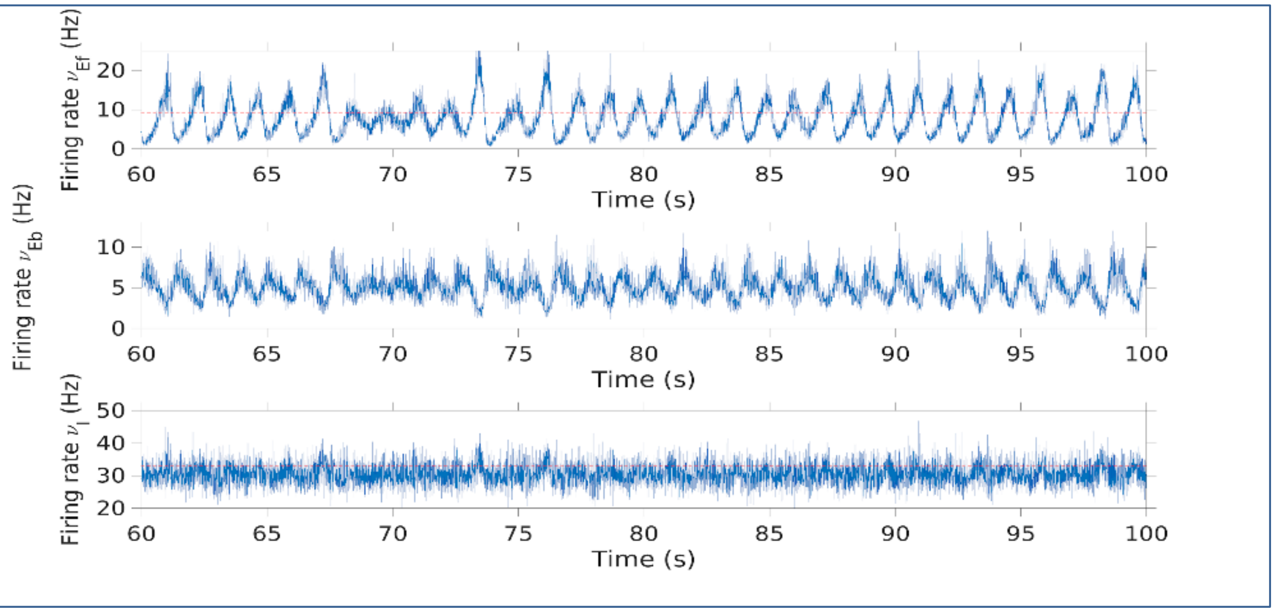

Figure 9: Time course of the slow-wave activity on a cortical module with plastic synapses.

Time variation of the firing rates of excitatory neurons participating to the slow-wave up-state activity (top panel), of excitatory neurons that do not participate in the slow-wave up-state (middle panel), and of inhibitory neurons (bottom panel). The figure shows that the interaction between the causal variant of STDP and SWA caused a progressive decrease of the amplitude of the waves, but that slow-wave activity persists for a relatively long transient period.

The porting of the SWA model without plasticity from DPSNN to NEST has been completed. In the final months of SGA1 we are testing the behaviour of NEST with the same STDP plasticity and preparing the delivery of the models in NEST format (SP6). Models and results will be shared with interested HBP partners using the HBP Collaboratory. In SGA2 perspective the provision of NEST models is expected to enable the run on the HBP Neuromorphic and Neurorobotic Platforms.

The activity of this SGA1 task is essential to prepare tools and models necessary for investigating the multiscale interplay between sleep and plasticity, which is a focus in SGA2. Key aspects of the SGA2 activity are: a) the evaluation of the stability properties of the SWA generated by a cortical network model incorporating plastic synapses and b) the study of the effect of SWA on the capability of recall memory engrams (distributed attractor states) stored in neocortex. Intuitively, we expect that the strong and coherent neuronal activation typical of SWA might strengthen (the relative weight of) synapses of functionally correlated neurons, while weakening the synapses connecting functionally uncorrelated neurons. To test this hypothesis, we must build a controlled simulation of a cortical module that (i) expresses SWA, (ii) incorporates synaptic plasticity, and (iii) is a realistic representation of mammalian cortical networks. This is possible as the collaboration within the HBP WP3.2 characterises itself by the effort to produce mesoscale models devised to bridge all brain scales when simulated on state-of-the art platforms with state-of-the art simulation engines.

2.2.1.1 Achieved Impact

26) Slow waves in cortical slices: how spontaneous activity is shaped by laminar structure

Cristiano Capone, Beatriz Rebollo, Alberto Munos, Xavi Illa, Paolo Del Giudice, Maria V Sanchez-Vives, Maurizio Mattia

Cerebral Cortex, 2017

27) Gaussian and exponential lateral connectivity on distributed spiking neural network simulation

Elena Pastorelli, Pier Stanislao Paolucci, Francesco Simula, Andrea Biagioni, Fabrizio Capuani, Paolo Cretaro, Giulia De Bonis, Francesca Lo Cicero, Alessandro Lonardo, Michele Martinelli, Luca Pontisso, Piero Vicini

Proceedings of PDP 2018, 26th Euromicro Intern. Conf. on Parallel, Distributed and Network-based Processing

28) (NEST implementation of the SWA model)

Distributed large scale simulation of synchronous slow-wave / asynchronous awake-like cortical activity

Elena Pastorelli, Cristiano Capone, Francesco Simula, Paolo Del Giudice, Maurizio Mattia, Pier Stanislao Paolucci

Poster presented at: 2017 NEST Conference, 19-20 December 2017, Haus Overbach, Julich, Germany

29) (DPSNN implementation of the SWA model)

Distributed large scale simulation of synchronous slow-wave / asynchronous awake-like cortical activity

Elena Pastorelli, Cristiano Capone, Francesco Simula, Paolo Del Giudice, Maurizio Mattia, Pier Stanislao Paolucci

Poster presented at: MSBDY 2017 (Brain Dynamics on Multiple Scales) Int. Workshop, 19-23 June 2017, Max Planck Inst. for the Physics of Complex Systems, Dresden, Germany

2.2.1.2 Component Dependencies

Summarised links to components this Key Result depends on:

Component ID	Component Name	HBP Internal	Comment
733	SP3: Single-area non-laminar model generating cortical slow waves - NEST flavour – WaveScalES (model)	No	KR builds on and contributes to this model
743	SP3: Multi-scale software model of cortical structures expressing slow waves and the transition to other consciousness states (model)	No	KR builds on and contributes to this model
749	SP3: Synaptic plasticity in slow wave activity simulations (model)	No	KR builds on and contributes to this model

2.2.2 Thalamocortical model simulating wakefulness, sleep, and state transitions

Johan Storm, Ricardo Murphy, Andre Nilsen, Bjørn Erik Juel, Hans Ekkehard Plesser

CDP5 collaboration between SP3 (U Oslo) SP7 (FZ Jülich) and SP9 (U Manchester)

We are developing a large multilayer thalamocortical network model consisting of conductance-based spiking neurons that can transition between wake- and sleep-like activity patterns when the neuronal properties are appropriately adjusted.

Cycling through distinct states of consciousness (such as wakefulness and different sleep stages) seems to be of importance for learning. However, the exact roles these states play in learning is not well understood. Some hypothesise that sleep is a price that must be paid for the ability to learn (such as in the synaptic rescaling suggested in the SHY hypothesis), while others suggest sleep is important for consolidation of memories formed during wakefulness (for example evidenced by pattern replay, see also KR2.1). Perhaps distinct sleep stages play different roles in synaptic plasticity and memory formation, consolidation, or other types of memory processing? In any case, models are useful for studying how naturally fluctuating brain states interact with plasticity and learning in the brain.

We are developing a large scale thalamocortical model in NEST, with the ability to cycle between wake-like and sleep-like stages, primarily to investigate clinically promising measures and biologically relevant mechanism of consciousness. The model also has the functionality to receive and respond to “retinal” input, with appropriate response patterns (across cortical regions) to structured visual stimuli in the wake-like state. Currently the connectivity required to induce such responses is hard-coded, but we wish to test how and to what extent the connectivity can be learned or modified by experience. Thus, the model can provide a testbed for investigating the effect of particular learning rules (e.g. KR1.1-KR1.3) once planned extensions such as 2/3-compartment neuron models are implemented. Thus, this model can become useful for investigating how learning and plasticity interacts with the sleep and wake states that mammalian brains naturally cycle through, in addition to being suitable for testing and generating predictions about consciousness in humans (which is the main purpose for the model).

The model is currently being prepared for publishing on the HBP Collaboratory and github to be available to the public. In addition, the model is being tested for compatibility with the SpiNNaker system in Manchester.

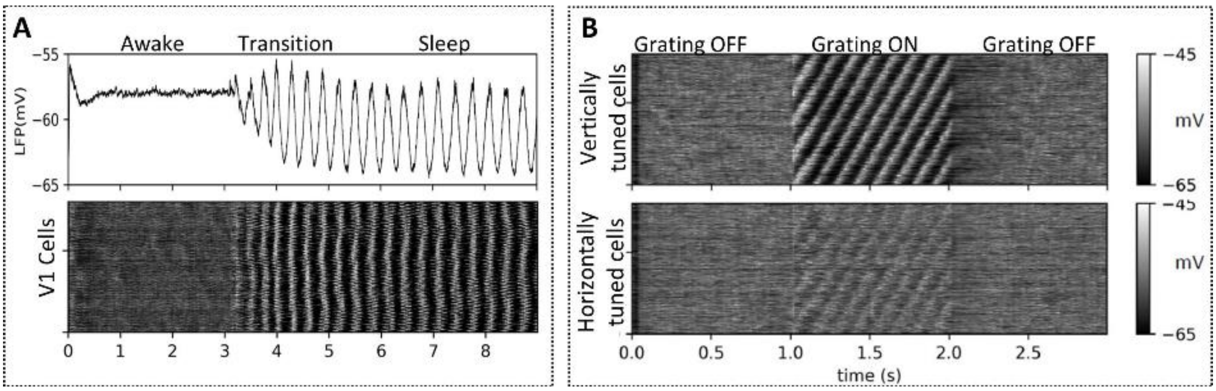

Figure 10: Initial results showing functionality of the thalamocortical model.

In A, a trace of the LFP (top) and a raster-like plot of membrane potential (bottom) of excitatory cells in V1 is shown as the relevant conductances are slowly changed from wake-like to sleep-like values. In B, the activity in two populations of cells in V1 (vertically and horizontally tuned cells in the top and bottom panels respectively) before, during, and after a moving vertical grating stimulus is fed into the primary thalamic nuclei in the model.

2.2.2.1 Achieved Impact

30) Simulating deep sleep and awake states in a mammalian thalamocortical model

Nilsen, Andre Sevenius; Murphy, Ricardo; Juel, Bjørn Erik; Storm, Johan Frederik.

Poster at the 2nd Nordic Neuroscience Meeting, Stockholm (Sweden), 2017

31) Implementation of the Hill-Tononi thalamocortical network model in the neural simulator NEST

Murphy, Ricardo; Nilsen, Andre Sevenius; Juel, Bjørn Erik; Storm, Johan Frederik.

Poster at the Human Brain Project Summit, Glasgow (UK), 2017

32) The Hill-Tononi thalamocortical network model implemented in NEST.

Nilsen, Andre Sevenius

Talk at Human Brain Project Summit, Glasgow (UK), 2017

2.2.2.2 Component Dependencies

Summarised links to components this Key Result depends on.

Component ID	Component Name	HBP Internal	Comment
2	SP9: SpiNNaker neuromorphic computing system (hardware)	Yes/No	Hardware emulation of model in SGA2
743	SP3: Multi-scale software model of cortical structures expressing slow waves and the transition to other consciousness states (model)	No	KR builds on and contributes to this model
969	SP4: Plasticity - Two-compartment neuron (model)	No	KR builds on and contributes to this model
749	SP3: Synaptic plasticity in slow wave activity simulations (model)	No	KR builds on and contributes to this model
1341	CDP5: Concept showcases in big systems (model)	No	CDP5 result

2.3 Key Result 3: Platform development and applications

2.3.1 In-the-loop training with neuromorphic hardware

Sebastian Schmitt, Johannes Schemmel, Karlheinz Meier, Mihai A. Petrovici, Walter Senn, Wolfgang Maass

CDP5 collaboration between SP9 (U Heidelberg, TU Graz) and SP4 (U Bern)

CDP5 partners have ported deep spiking networks to neuromorphic substrates, which were trained in software, with the hardware in the loop.

Partners from U Heidelberg, TU Graz and U Bern have successfully implemented a deep spiking network architecture on the BrainScaleS wafer-scale system. In 33, 34), a deep neural network trained in software (using Tensorflow) was converted to a spiking network on the BrainScaleS system. This process was followed by "in-the-loop" training, where in each training step, the network activity was first recorded in hardware and then used to compute the parameter updates in software via backpropagation of errors ("backprop"). An essential finding was that the parameter updates do not have to be precise, but only need to approximately follow the correct gradient, which simplifies the computation of updates. Using this approach, after only several tens of iterations, the spiking network showed an accuracy close to the ideal software-emulated prototype.

The presented techniques demonstrate that deep spiking networks emulated on analogue neuromorphic devices can attain good computational performance despite the inherent variations of the analogue substrate. The speedup compared to biological real time, including reconfiguration of the hardware and spike in and output, was found to be a factor of 1,000, i.e., the full training set of approx. 5,000 images was classified within 5 s of wall-clock time.

Furthermore, the ability to overcome hardware imperfections and a transition from continuous signals to spikes with an approximate learning rule akin to backprop paves the way for training larger and more complex hierarchical models on the physical-model BrainScaleS system (see KR1.3 and KR1.4).

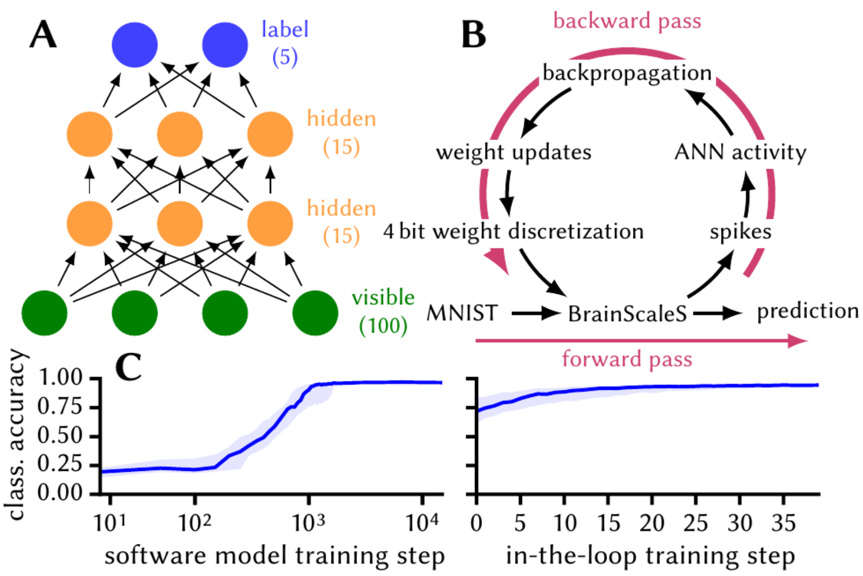

Figure 11: In-the-loop training.

(A) Structure of the feed-forward, rate-based deep spiking network. (B) Schematic of the training procedure with the hardware in the loop. (C) Classification accuracy over training step. Left: software training phase, right: hardware in-the-loop training phase.

2.3.1.1 Impact

33) Pattern representation and recognition with accelerated analog neuromorphic systems

Mihai A. Petrovici*, Sebastian Schmitt*, Johann Klähn*, David Stöckel*, Anna Schroeder*, Guillaume Bellec, Johannes Bill, Oliver Breitwieser, Ilja Bytschok, Andreas Grübl, Maurice Güttler, Andreas Hartel, Stephan Hartmann, Dan Husmann, Kai Husmann, Sebastian Jeltsch, Vitali Karasenko, Mitja Kleider, Christoph Koke, Alexander Kononov, Christian Mauch, Eric Müller, Paul Müller, Johannes Partzsch, Thomas Pfeil, Stefan Schiefer, Stefan Scholze, Anand Subramoney, Vasilis Thanasoulis, Bernhard Vogginger, Robert Legenstein, Wolfgang Maass, René Schüffny, Christian Mayr, Johannes Schemmel, Karlheinz Meier

Proceedings of the 2017 IEEE International Symposium on Circuits and Systems (ISCAS)

34) Neuromorphic Hardware In The Loop: Training a Deep Spiking Network on the BrainScaleS Wafer-Scale System

Sebastian Schmitt*, Johann Klaehn*, Guillaume Bellec, Andreas Grübl, Maurice Güttler, Andreas Hartel, Stephan Hartmann, Dan Husmann, Kai Husmann, Vitali Karasenko, Mitja Kleider, Christoph Koke, Christian Mauch, Eric Mueller, Paul Mueller, Johannes Partzsch, Mihai A. Petrovici, Stefan Schiefer, Stefan Scholze, Bernhard Vogginger, Robert Legenstein, Wolfgang Maass, Christian Mayr, Johannes Schemmel, Karlheinz Meier

Proceedings of the 2017 IEEE International Joint Conference on Neural Networks

35) Neuromorphic Computing: *Das Gehirn als Vorbild*

Sebastian Schmitt

Talk at the HBP Innovation Day, Munich (Germany), 2017

36) User Training

Sebastian Schmitt

HBP Hackathon 2018

37) Experiments on BrainScaleS

Sebastian Schmitt

Talk at the NICE Workshop, Portland (USA), 2018

2.3.1.2 Component Dependencies

Summarised links to components this Key Result depends on:

Component ID	Component Name	HBP Internal	Comment
1	SP9: BrainScaleS 1 neuromorphic computing system (hardware)	publicly accessible via Collab	Hardware emulation environment
1341	CDP5: Concept showcases in big systems (model)	No	CDP5 result

2.3.2 Stochastic computing with spikes

Akos Kungl, Johannes Schemmel, Karlheinz Meier, Mihai A. Petrovici, Walter Senn, Wolfgang Maass

CDP5 collaboration between SP9 (U Heidelberg, TU Graz) and SP4 (U Bern)

CDP5 partners have ported hierarchical sampling networks to neuromorphic substrates. After training, these networks are simultaneously discriminative and generative models of the learned data and can therefore be used both for classification and for pattern completion.

A major challenge in performing computing with neuromorphic substrates is the implementation of not only discriminative but also of generative models. Similar to the mammalian brain, such neural networks should be capable of recreating visual data based upon the training set, which enables them to also correct erroneous inputs or complete partially missing data based upon an internal Bayesian model evolved during training. In an extended neural sampling framework 38), we have demonstrated the ability of LIF neurons to sample from well-defined target distributions and perform Bayesian inference in the associated probability spaces. By endowing these networks with a hierarchical structure, it is then possible to train a model of complex visual data. Thereby, this work connects the concepts of probabilistic deep neural networks from machine learning to spiking neurons and hence to the neuromorphic systems developed in SP9.

Based on this model, in a collaboration between U Heidelberg, U Bern and TU Graz, such a network was implemented on the BrainScaleS neuromorphic platform 40,41). This was achieved by utilising the concept of a “sea of noise”, i.e. by replacing the source of randomness via the spiking activity of a randomly connected neural network - a concept developed in a collaboration between FZJ and U Heidelberg and described in 39).

We demonstrated that the implemented network is able to perform sampling from a targeted probability distribution. In order to test the discriminative and generative properties of the model we trained it on the MNIST handwritten digits dataset. The network reached a classification ratio comparable to that of an ideal software simulated model, and could successfully complete partially known digits and create recognisable ones based on its internal model. While the probabilistic nature of the model requires a fundamentally different kind of learning compared to the backpropagation of errors, the ability of the hardware to emulate hierarchical networks with both feedforward and feedback connectivity represents an important stepping stone towards embedding the network models presented in KR1.3 and KR1.4, thereby paving the way for multiple Use Cases in SGA2.

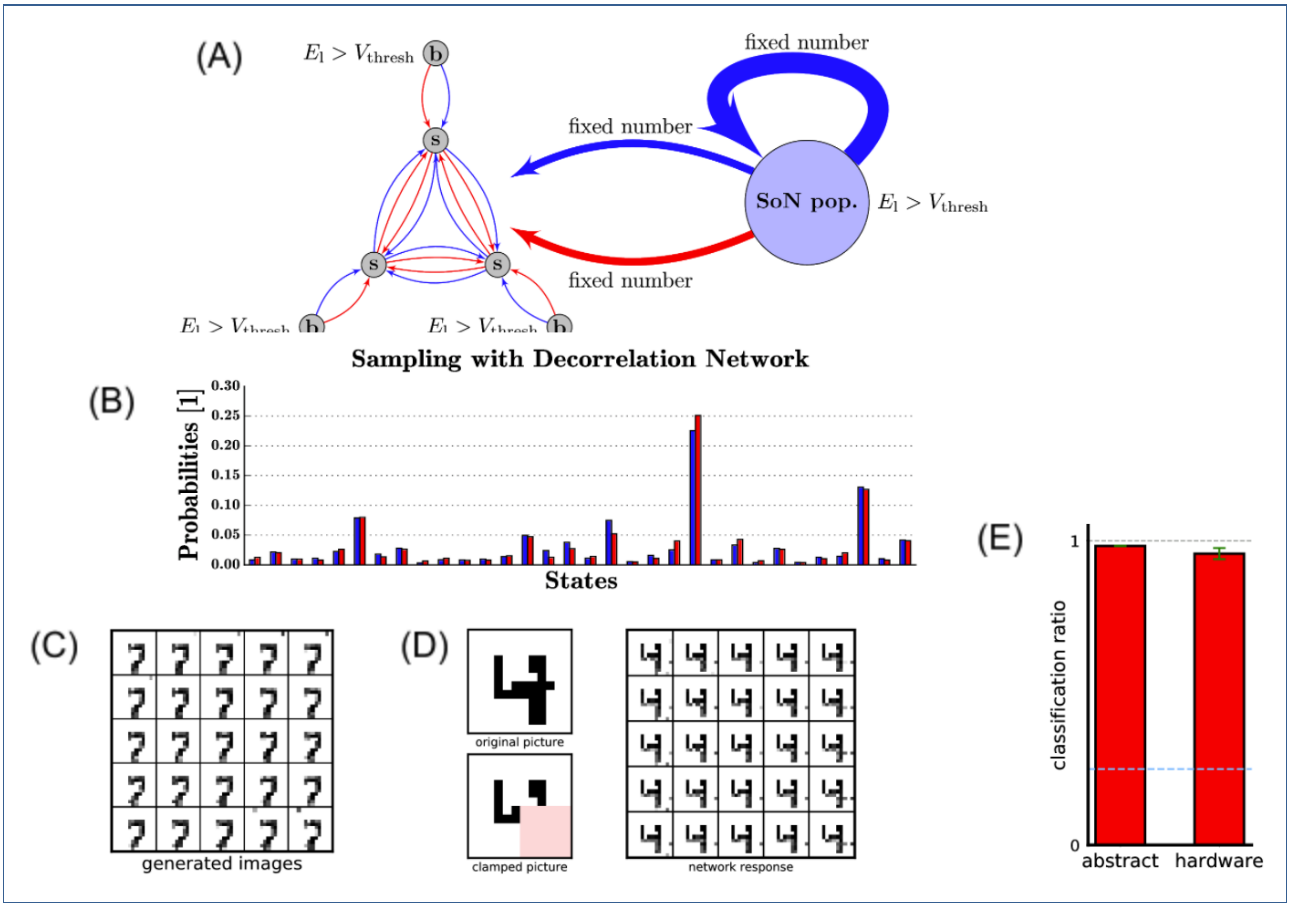

Figure 12: Stochastic computing with spikes

(A) Schematic of the implemented network. The sampling network receives its noise input from the spiking activity of a randomly connected spiking neural network. (B) Bar plots of a sampled probability distribution (red) compared to the target distribution (blue). (C) Example images generated by the model after being trained on a subset of the MNIST hand-written digits dataset. (D) The network is capable of completing partially available data based on its internal model it obtained during learning. (E) The reached classification ratio is comparable with that of an ideal software implementation.

2.3.2.1 Impact

38) Stochastic inference with spiking neurons in the high-conductance state

Mihai A. Petrovici*, Johannes Bill*, Ilja Bytschok, Johannes Schemmel, Karlheinz Meier

Physical Review E 94, 042312 (2016)

39) Stochastic neural computation without noise

Jakob Jordan, Mihai A. Petrovici, Oliver Breitwieser, Johannes Schemmel, Karlheinz Meier, Markus Diesmann, Tom Tetzlaff

arXiv:1710.04931

40) Robustness from structure: Inference with hierarchical spiking networks on analog neuromorphic hardware

Mihai A. Petrovici*, Anna Schroeder*, Oliver Breitwieser, Andreas Grübl, Johannes Schemmel, Karlheinz Meier

Proceedings of the 2017 IEEE International Joint Conference on Neural Networks

33) Pattern representation and recognition with accelerated analog neuromorphic systems

Mihai A. Petrovici*, Sebastian Schmitt*, Johann Klähn*, David Stöckel*, Anna Schroeder*, Guillaume Bellec, Johannes Bill, Oliver Breitwieser, Ilja Bytschok, Andreas Grübl, Maurice Güttler, Andreas Hartel, Stephan Hartmann, Dan Husmann, Kai Husmann, Sebastian Jeltsch, Vitali Karasenko, Mitja Kleider, Christoph Koke, Alexander Kononov, Christian Mauch, Eric Müller, Paul Müller, Johannes Partzsch, Thomas Pfeil, Stefan Schiefer, Stefan Scholze, Anand Subramoney, Vasilis Thanasoulis, Bernhard Vogginger, Robert Legenstein, Wolfgang Maass, René Schüffny, Christian Mayr, Johannes Schemmel, Karlheinz Meier

Proceedings of the 2017 IEEE International Symposium on Circuits and Systems (ISCAS)

41) Spike-based probabilistic inference with correlated noise

Ilja Bytschok, Dominik Dold, Johannes Schemmel, Karlheinz Meier, Mihai A. Petrovici

BMC Neuroscience 2017, 18 (Suppl 1):P200

42) Neural Sampling with Linear Feedback Shift Registers as a Source of Noise

Marcel Großkinsky

Bachelor Thesis, University of Heidelberg, 2016

43) Stochastic Computation in Spiking Neural Networks Without Noise

Dominik Dold

Master Thesis, University of Heidelberg, 2016

44) Magnetic Phenomena in Spiking Neural Networks

Andreas Baumbach

Master Thesis, University of Heidelberg, 2016

45) *Struktur schafft Robustheit: Eine Untersuchung hierarchischer neuronaler Netzwerke mit unpräzisen Komponenten*

Anna Schroeder

Bachelor Thesis, University of Heidelberg, 2016

46) Sampling with leaky integrate-and-fire neurons on the HICANNv4 neuromorphic chip

Akos Kungl

Master Thesis, University of Heidelberg, 2016

47) Accelerated Classification in Hierarchical Neural Networks on Neuromorphic Hardware

Carola Fischer

Bachelor Thesis, University of Heidelberg, 2017

48) Simulated Tempering in Spiking Neural Networks

Agnes Korcsak-Gorzo

Master Thesis, University of Heidelberg, 2017

49) Simulated Tempering in Biologically Inspired Neural Networks

Agnes Korcsak-Gorzo, Luziwei Leng, Oliver Julien Breitwieser, Johannes Schemmel, Karlheinz Meier, Mihai A. Petrovici

Distinguished poster at the Deutsche Physikerinnentagung, Oldenburg (Germany), 2018

50) Stochastic computation on spiking neuromorphic hardware

Dominik Dold, Ákos F. Kungl, Andreas Baumbach, Johann Klähn, Ilja Bytschok, Paul Müller, Oliver Breitwieser, Andreas Grübl, Maurice Güttler, Dan Husmann, Mitja Kleider, Christoph Koke, Alexander Kugele, Christian Mauch, Eric Müller, Sebastian Schmitt, Johannes Schemmel, Karlheinz Meier, Mihai A. Petrovici

Poster at the Bernstein Conference, Göttingen (Germany), 2017

51) Stochastic inference with deterministic spiking neurons

Mihai A. Petrovici, Luziwei Leng, David Stöckel, Oliver Breitwieser, Ilja Bytschok, Jakob Jordan, Roman Martel, Johannes Bill, Johannes Schemmel, Karlheinz Meier

Talk at the NICE workshop, Berkeley (USA), 2016

52) Stochastic computation in spiking neural networks

Mihai A. Petrovici

Invited Talk at the Neuroscience Seminar of the Department of Physiology, Bern (Switzerland), 2016

53) Stochastic computation in spiking neural networks

Mihai A. Petrovici

Talk at the STRUCTURES Initiative, Heidelberg (Germany), 2016

54) Exploring the substrate of intelligence

Mihai A. Petrovici

Lecture at the HBP Students' Conference, Vienna (Austria), 2017

55) Fast inference with spiking networks

Mihai A. Petrovici

Talk at the workshop "Machine Learning Meets Biology: Algorithms and Cortical Mechanisms", Geneva (Switzerland), 2017

56) A (deep) dive into probabilistic computing

Mihai A. Petrovici

Lecture at the HBP Summer School, Obergurgl (Austria), 2017

57) *Träumen Androiden von elektrischen Schafen*?

Mihai A. Petrovici

Invited Talk at the NOVO Innovation Day, Bern (Switzerland), 2017

58) Neuromorphic computing - The physics of cognition

Mihai A. Petrovici

Talk at the HBP National Outreach Event, Amsterdam (Netherlands), 2018

59) Spiking neuron ensembles and probabilistic inference

Mihai A. Petrovici

Invited talk at the BDC Symposium, Heidelberg (Germany), 2018

60) Experiments on BrainScaleS

Sebastian Schmitt

Talk, NICE2018, http://niceworkshop.org/2018-nice-workshop/ February 27 to March 1

2.3.2.2 Component Dependencies

Summarised links to components this Key Result depends on.

Component ID	Component Name	HBP Internal	Comment
1	SP9: BrainScaleS 1 neuromorphic computing system (hardware)	Publicly accessible via Collab	Hardware emulation environment
468	SP9: Principles for brain-like computation (model)	No	CDP5-related results
1066	SP4: Plasticity - Synaptic plasticity and learning (model)	No	KR builds on and contributes to this model
1341	CDP5: Concept showcases in big systems (model)	No	CDP5 result
1342	CDP5: Guiding platform design on functional plasticity (model)	No	CDP5 result
2420	SP4: Plasticity - Prototype implementations of rules and testing within and without the SP9 platforms (model)	(currently in Fusi_plas-ticity branch of the official SpiNNaker github repository. Eventually this branch will be merged into master branch)	Towards Third-Factor learning rules on SpiNNaker

2.3.3 *Compartmental neurons in neuromorphic hardware*

Johannes Schemmel, Paul Müller, Karlheinz Meier, Walter Senn, Matthew Larkum

CDP5 collaboration between SP9 (U Heidelberg), SP4 (U Bern) and SP3 (HU Berlin)

We have developed and produced an accelerated mixed-signal neuromorphic system that incorporates inter-compartment connectivity and active dendrites supporting plateau potentials.

Biological neurons are complex information-processing units. Computational and hardware modelling often reduces the state of the neuron to a membrane voltage at a single point. The most common of these models is the Leaky Integrate-and-Fire (LIF) model, which abstracts a neuron to a leaky integrator that emits an action potential whenever its membrane potential reaches a fixed threshold. However, spatially extended neurons can support more complex dynamics. The local properties of the neuronal membrane, such as diameter and transmembrane protein concentration, enable behaviour that exceeds that of the simple LIF model. Image processing in the fly visual system, for example, relies in part on dendritic filtering of the input (for a review, see e.g. London & Häusser, Annu. Rev. Neurosci. 2005).

Parts of the dendritic tree can act as active components in information processing. It has been shown in Larkum, Trends in Neurosciences 2013, see also Figure 13, left), that neurons can act as nonlinear coincidence detectors for input arriving at dendrite and soma. While dendritic stimulus nearly vanishes at the soma (Figure 13A) and somatic stimulus evokes a single spike (Figure 13B), both stimuli combined interact nonlinearly, resulting in a burst at the soma (Figure 13C).

Building on input from SP3 and SP4, UHEI has developed and fabricated a single-chip prototype system that enables the configurable interconnection of multiple physical neuron models. The resulting circuit forms the equivalent of a multi-compartment neuron, which can be adjusted to contain active dendrites. A simulated configuration that mimics the biological use case is shown in Figure 13, panels D-F. Four compartments are used to emulate three neuronal compartments (inset): Dendritic (D) and somatic (E) stimulus interacts nonlinearly, causing a burst in the somatic compartment (F).

The implementation of multi-compartment neurons, along with the embedded plasticity processing unit (see KR3.4) will enable the BrainScaleS 2 system to emulate the two-compartment-based plasticity models described in KR1.1 - KR1.3.

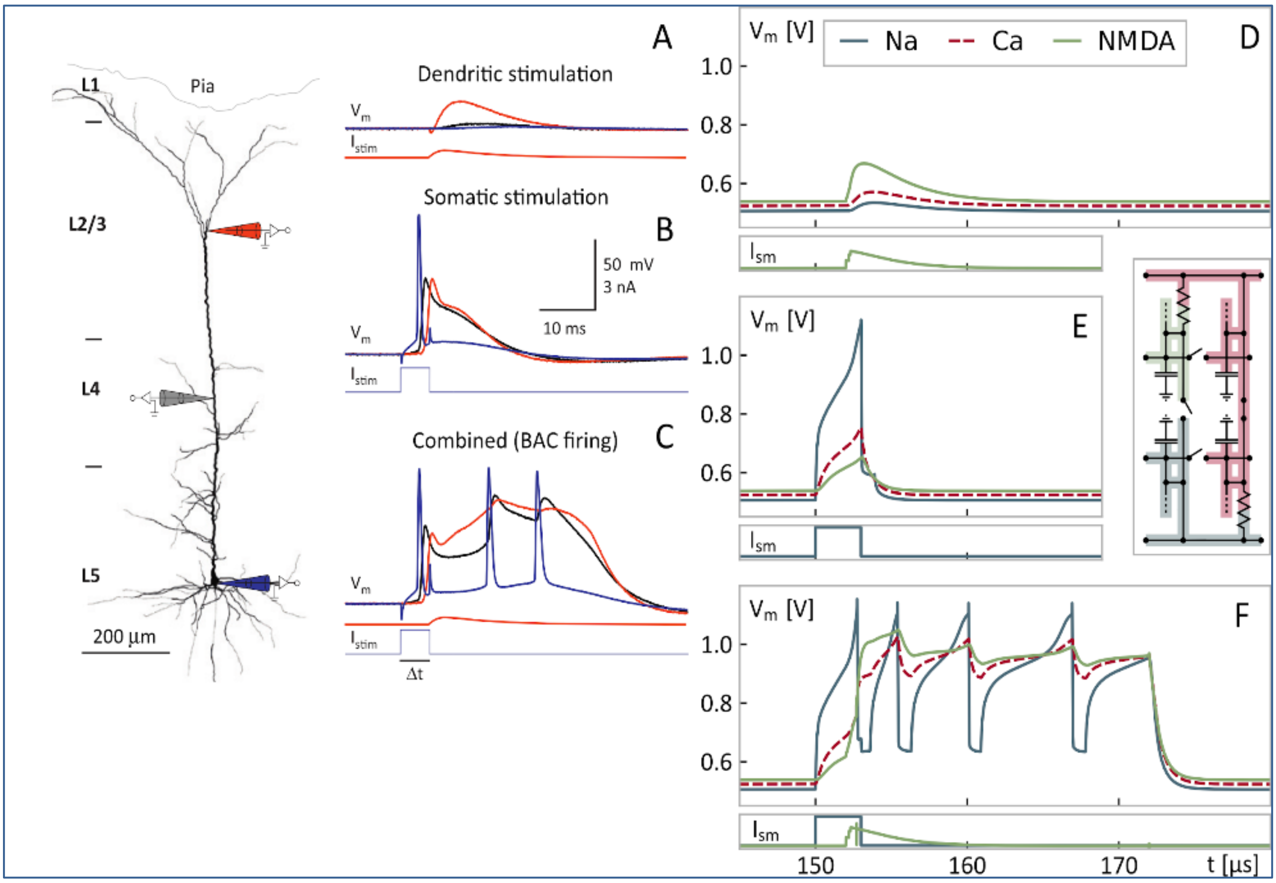

Figure 13: Simulation of multi-compartment functionality in the prototype chip system.

(A-C) Figure taken from Larkum, Trends in Neurosciences 2013). The response of a layer 5 pyramidal neuron to separate and correlated dendritic and somatic stimulus is shown. The traces show the membrane voltage at the Na-spike initiation zone (blue), the Ca-spike initiation zone (black) and the dendrites (red). (D-F) Simulation of the corresponding circuit configuration. (D) Spike input into the NMDA compartment (green, bottom) is chosen as not to cause firing in any compartment. An attenuated version of the post-synaptic potential is seen in the Ca and Na compartments. (E) Current stimulus into the Na compartment which is adjusted to initiate a single spike. The pictogram shows the switch and resistor configuration to which the simulations in E-G will correspond in the final chip. (F) Both inputs combined suffice to cause firing in the Ca and NMDA compartments. This, in turn, induces a burst in the Na compartment. Note that the time scale is milliseconds for the biological reference and microseconds for the circuit simulation due to the accelerated nature of the neuromorphic device.

2.3.3.1 Impact

61) An accelerated analog neuromorphic hardware system emulating NMDA-and calcium-based non-linear dendrites.

Johannes Schemmel, Laura Kriener, Paul Müller and Karlheinz Meier

Neural Networks (IJCNN), 2017 International Joint Conference on. IEEE, 2017.

62) Neuromorphic implementation of dendritic computing

Karlheinz Meier

Talk at the Dendritic integration and computation with active dendrites workshop, Paris (France), 2018

2.3.3.2 Component Dependencies

Summarised links to components this Key Result depends on.

Component ID	Component Name	HBP Internal	Comment
457	SP9: BrainScaleS 2 neuromorphic computing system (hardware)	Prototype accessible HBP internal	Hardware emulation environment
468	SP9: Principles for brain-like computation (model)	No	CDP5-related results
969	SP4: Plasticity - Two-compartment neuron (model)	No	KR builds on and contributes to this model
1033	SP4: Plasticity - STDP for a multi-compartment model with NMDA spikes (model)	Yes	KR builds on and contributes to this model
1066	SP4: Plasticity - Synaptic plasticity and learning (model)	No	KR builds on and contributes to this model
1341	CDP5: Concept showcases in big systems (model)	No	CDP5 result
1342	CDP5: Guiding platform design on functional plasticity (model)	No	CDP5 result
2419	SP4: Plasticity - Algorithms for multicompartment models (model)	Yes	KR builds on and contributes to this model

2.3.4 *An embedded plasticity processor for the BrainScaleS-2 system*

Johannes Schemmel, Benjamin Cramer, Karlheinz Meier, Nicolas Fremaux, Wulfram Gerstner, Mihai A. Petrovici, Walter Senn

CDP5 collaboration between SP9 (U Heidelberg) and SP4 (EPFL-LCN Lausanne, U Bern)

In computational neuroscience, new learning rules are being continuously suggested, either driven by empirical observations from biology or by functional requirements in machine-learning contexts. Across this ever-expanding plethora of plasticity mechanisms, there exists a huge disparity in terms of independent variables (such as spike times or membrane potentials) parameters and time scales. A purely analogue physical (neuromorphic) implementation would therefore be confined to choosing only a small subset of these plasticity rules. To increase the flexibility of the BrainScaleS system, an on-chip plasticity co-processor has been designed and built, and preliminary experiments have been successfully performed.

Former approaches to implement plasticity in accelerated, analogue neuromorphic hardware systems provide high efficiency by building dedicated electronic circuits, approximating the desired behaviour of a specific plasticity rule. This approach affords only little flexibility, since the implementation of new learning algorithms involves the design of specialised hardware circuits. However, in order to allow a large community of modellers to benefit from features of the BrainScaleS system such as the acceleration factor or the system size, flexibility in the implementable plasticity mechanisms should be provided. This is done by using a hybrid approach: a general-purpose processor, the plasticity processing unit (PPU), is combined with analogue neuromorphic hardware. The analogue elements provide a speed-up, necessary for performing long learning experiments and at the same time an energy efficient implementation. The PPU provides flexibility in the implementable learning algorithms.

With the PPU, plasticity algorithms beyond simple spike-timing-dependent plasticity (STDP) could be realised. Here, we show how it can be used to update synaptic weights iteratively to yield a stable state. These updates let the network evolve to a state where it becomes exquisitely sensitive to changes in its input, thereby increasing its computational power. This form of plasticity is only weakly sensitive to substrate imperfections and requires little tuning, which is of particular relevance for analogue neuromorphic substrates.

Together with the implementation of multi-compartment neurons (see KR3.3), the PPU endows the BrainScaleS 2 system with the necessary flexibility for implementing the learning models discussed in KR1.1-1.3.

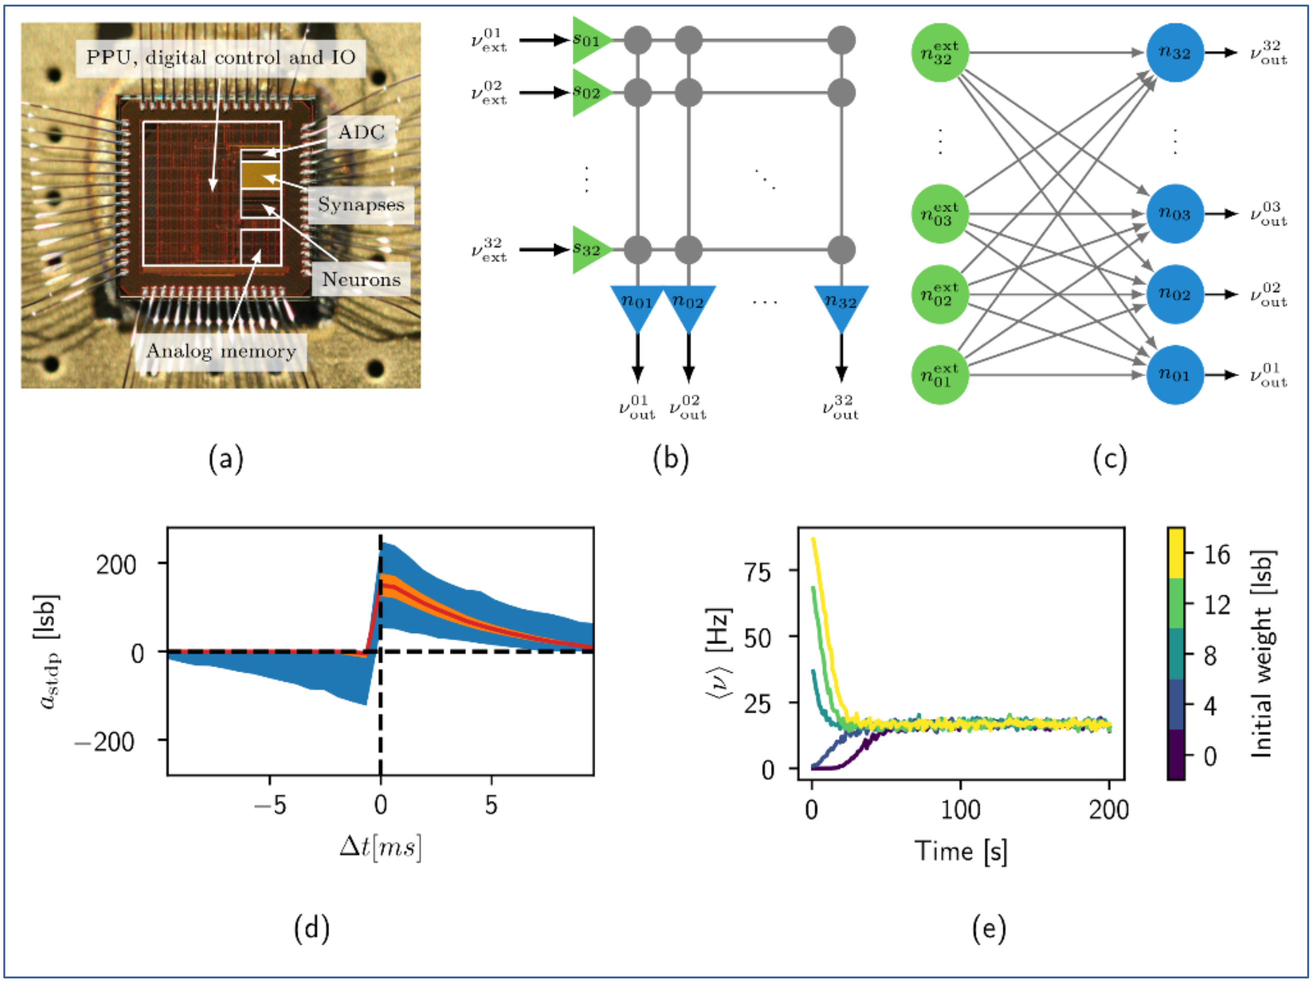

Figure 14: An embedded plasticity processor for the BrainScaleS-2 system

(a) Photograph of the HICANN-DLS die with the functional areas highlighted. The digital control and IO area handles the chip configuration and communication with the host. For in-the-loop reconfiguration during an experiment, there exists the PPU which is an embedded processor that can be programmed to perform various tasks, e.g. implement synaptic plasticity. (b) Abstract view of the on-chip network. There are 32 different inputs represented by the green triangles on the left. The blue triangles at the bottom represent the 32 neurons which can receive input from up to 32 synapses, show by the grey circles. (c) Abstraction of the on-chip network on the level of neurons and synapses. Every spike input to the chip is an abstract input-neuron which are represented by the green circles. The up to 32x32 synapses are represented by the grey arrows, targeting the on-chip neurons drawn as blue circles. (d) Every synapse has a correlation sensor that measures exponentially weighted pre-post spike time pairings. This figure shows the measured STDP kernel function for different spike-time differences. The blue range shows the variation in the correlation sensors sensitivities which varies due to analogue effects. (e) Average spike frequency of the on-chip neurons during a network emulation, where the PPU reconfigures the synaptic weights iteratively to yield a stable, moderate spiking activity. For each of the initial weights shown with different colours, the target frequency of 20Hz is reached after some time.

2.3.4.1 Impact

63) Demonstrating Hybrid Learning in a Flexible Neuromorphic Hardware System

Friedmann, Schemmel, Grübl, Hartel, Hock, Meier

IEEE Transactions on Biomedical Circuits and Systems (2017)

64) From LIF to AdEx Neuron Models: Accelerated Analog 65 nm CMOS Implementation

Syed Ahmed Aamir*, Paul Müller*, Laura Kriener, Gerd Kiene, Johannes Schemmel and Karlheinz Meier

Accepted at 2017 IEEE Biomedical Circuits and Systems Conference (BioCAS)

65) Exploring Collective Neural Dynamics under Synaptic Plasticity on Neuromorphic Hardware

David Stöckel

Master Thesis, University of Heidelberg (2017)

66) STDP in a flexible neuromorphic hardware system

Benjamin Cramer and David Stöckel

Talk at the In-depth meeting on network learning, Fürberg (Austria), 2017

62) Neuromorphic implementation of dendritic computing

Karlheinz Meier

Talk at the Dendritic integration and computation with active dendrites workshop, Paris (France), 2018

2.3.4.2 Component Dependencies

Summarised links to components this Key Result depends on.

Component ID	Component Name	HBP Internal	Comment
457	SP9: BrainScaleS 2 neuromorphic computing system (hardware)	Prototype accessible HBP internal	Hardware emulation environment
468	SP9: Principles for brain-like computation (model)	No	CDP5-related results
1066	SP4: Plasticity - Synaptic plasticity and learning (model)	No	KR builds on and contributes to this model
1341	CDP5: Concept showcases in big systems (model)	No	CDP5 result
1342	CDP5: Guiding platform design on functional plasticity (model)	No	CDP5 result

2.3.5 Plasticity rules in neuromorphic hardware

André Grüning, Brian Gardner, Anna Bulanova, Andreas Hartel, Eric Müller, Eric Nichols, Oliver Rhodes, Steve Furber

CDP5 collaboration between SP9 (U Manchester, U Heidelberg) and SP4 (U Surrey, EITN Paris)

We co-develop biologically plausible learning rules and their implementation on the current neuromorphic platforms.

For future improved artificial intelligence and a better understanding of the brain, it is important to take neurobiological plasticity mechanisms seriously and explore their functional and computational properties on a larger scale. However, when bringing these rules together with next-generation computing hardware such as the neuromorphic platforms, these plasticity mechanisms cannot be implemented on the neuromorphic platforms straight away. In this task we therefore developed learning algorithms that are on the one hand biologically inspired, but adapted such that they are implementable on the neuromorphic platforms. Likewise, we also adapt the neuromorphic platforms such that they can support these biologically inspired learning rules. Our contribution consists of several efforts.

Gradient-descent learning on spike pattern likelihood

We developed a learning rule that strikes the balance between biological plausibility, technical performance, and suitability for implementation on the neuromorphic platforms in current state. This approach serves to eventually test both the computational function of the rule on a large scale as well as to test whether the neuromorphic platforms in their current state are capable to

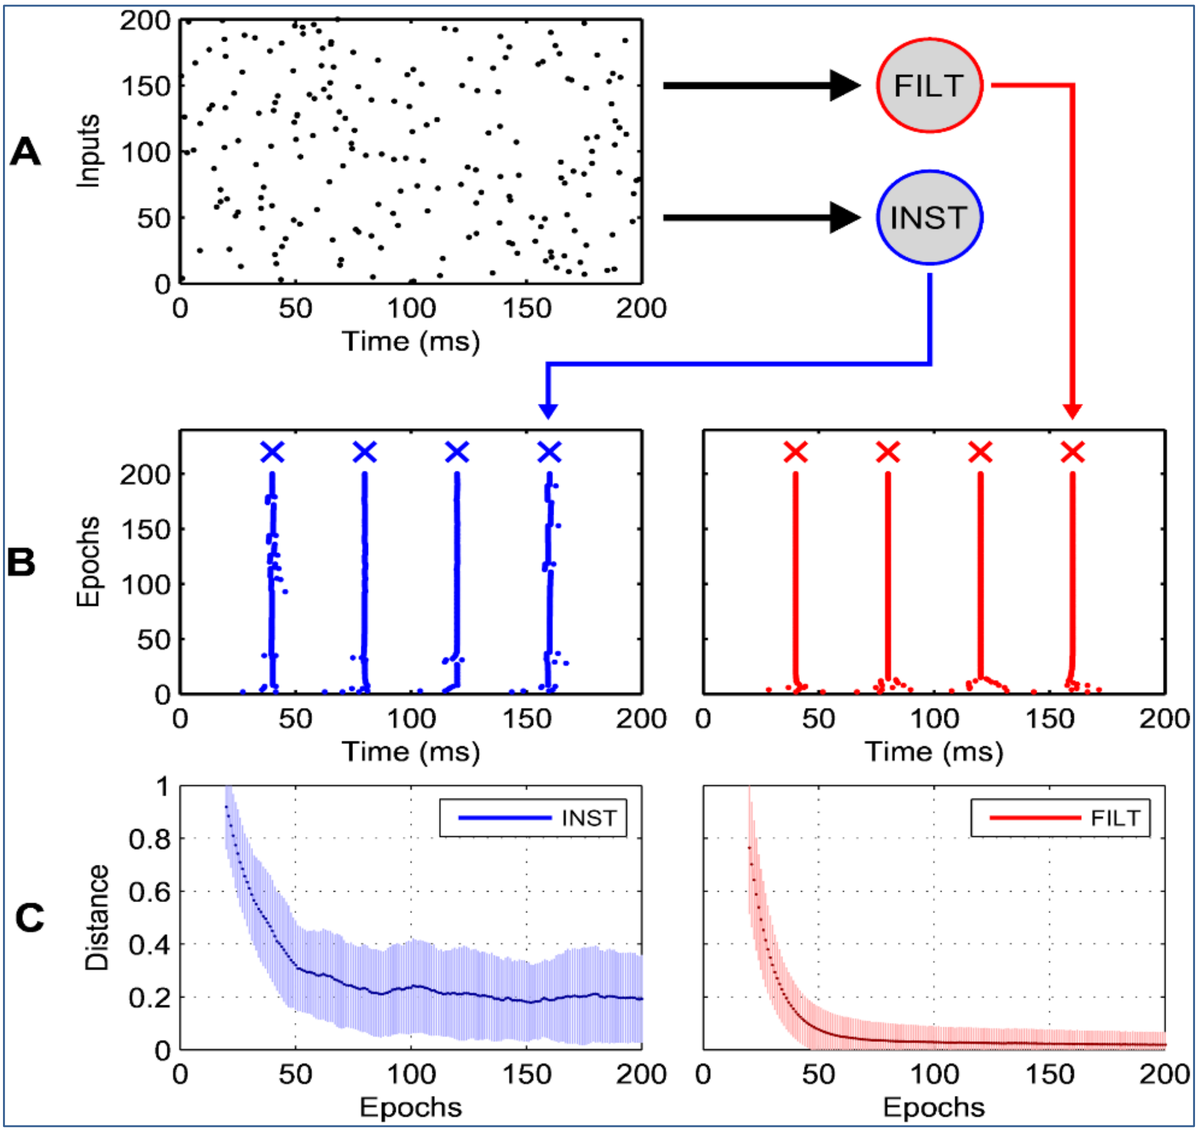

Figure 15: Two postsynaptic neurons trained under the proposed synaptic plasticity rules

Two postsynaptic neurons trained under the proposed synaptic plasticity rules that learned to map between a single, fixed input spike pattern and a four-spike target output train using the INST or FILT rule. (A): a frozen set of pseudo-Poisson spike trains serves as input to the neuron. Blue represents the use of the theoretically better founded INST rule, while red denotes the use of its heuristically improved FILT version. (B) A single learning run: The neuron was trained for 200 ms simulated time. For both rules actual output spikes times converge in several tens of training epochs to the target spike times denoted by crosses, with FILT showing better convergence. (C) Average van-Rossum distance of actual and target output spikes across 10 repetitions of the experiment: The error of the mean is lower for the heuristic FILT rule, demonstrating more reliable converge to the target spike times.

support these rules. This is by no means straight-forward as the neuromorphic platforms impose a number of constraints (based on their architecture) that are not transparent at first sight.

The supervised INST/FILT rule we developed utilises gradient descent on the likelihood of target spike frames and has only modest requirements for neuromorphic hardware compared to more biologically faithful approaches 67), Figure 15. More specifically, this rule only requires access to standard STDP traces, which, as of HICANN-DLS v2, are the only analogue hardware quantities that are accessible to the local plasticity processing unit (PPU). The implementation work on the HICANN-DLS v2 prototype has begun in SGA1 and will be carried over to SGA2 68) in a collaboration with U Heidelberg.

Evolving the SpiNNaker neuromorphic platform to support Third-Factor learning rules

Two years ago, the SpiNNaker platform did not provide the infrastructure/API to support learning rules that varied much from the standard Hebbian/STDP plasticity rules: The platform was optimised towards the efficient computation of downstream spike-based information flow in milliseconds of biological time. However, many rules that are of biological or conceptual interest deviate from this scheme and require lateral or downstream information flow (such as rewards or other third signals) which had not been a design consideration for the SpiNNaker platform.

Therefore, we carried out explorative experiments on the SpiNNaker platform to test the implementability of our plasticity rules (Figure 16, 69)). This included an effort to route a supervision signal (target spikes, Figure 17) to individual neurons and was successful as a prototype (Figure 18).

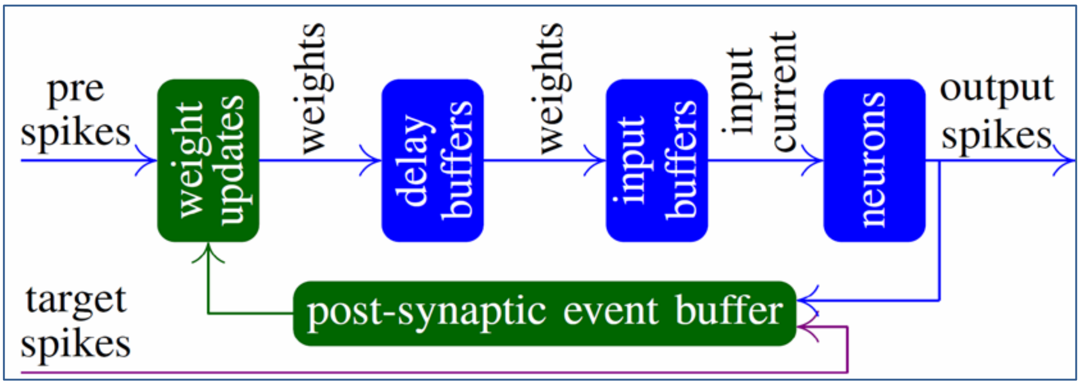

Figure 16: Overview of (intermediate-level) data flow in the Spinnaker platform

Dataflow on c-code level in blue is the standard SpiNNaker platform. To introduce Third-Factor supervised plasticity, code to enable data flow to include information on target spikes (the supervisory signal) had to be developed and integrated: Green illustrates the data flow that had to be modified to accommodate target spikes, and the purple arrow shows how the target spikes now enter the data flow within the SpiNNaker infrastructure code. Figure adapted on a drawing of the original SpiNNaker dataflow, courtesy of James Knight at partner UMAN.

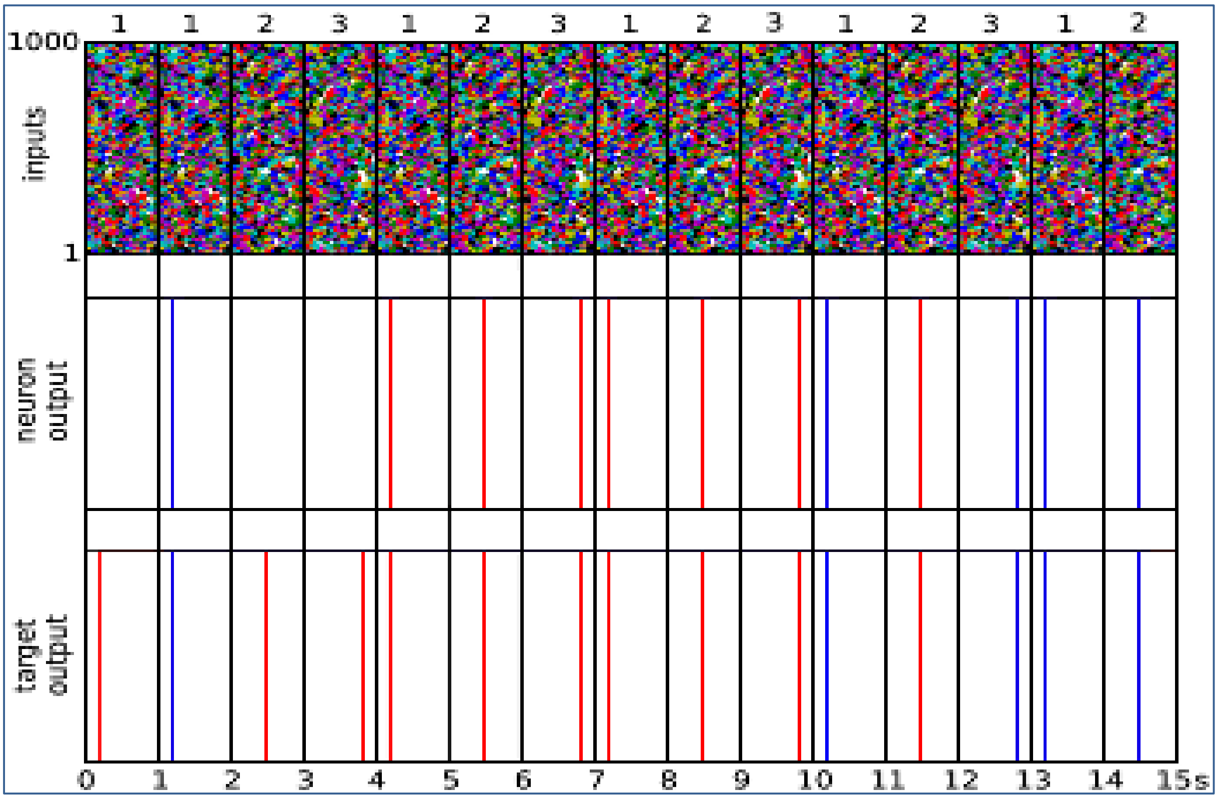

Figure 17: Supervised learning experiment with 3 frozen random input spike pattern, and 3 corresponding targets

With the new dataflow (Figure 16) included into the Spinnaker code, the INST/FILT rule was implemented on SpiNNaker and successfully used in a simple learning experiment. First row: 3 frozen (pseudo-Poisson) spike input spike patterns of 1s duration for 1000 incoming connections are applied to a single neuron in random order. The neuron's task is to produce a differently timed spike in response to each of the three input patterns. Middle and bottom row: actually produced and target spike times, respectively. Agreement of target and actual output spike coded in blue (within 1ms precision), deviation in red. After 12s the network has learned to produce all spike times correctly.

In the course of this work, it became apparent that a more systematic approach to implementation of such a learning rule was needed, and we therefore set up a cross-SP4-SP9 working group within CDP5 to facilitate the extension of the SpiNNaker API and underlying code infrastructure such that various custom third-factor plasticity rules can easily be implemented by third parties on SpiNNaker. Within this work group we prioritised existing plasticity rules with respect to their estimated implementation complexity within SpiNNaker, and are using in this sense simpler rules as stepping stones to progress to more complex rules. These rules are

- the Brader-Fusi-Senn Rule (Brader *et al*., Neural Computation 2007), see Figure 18.
- the Urbanczik-Senn Rule, in progress (see KR1.1-1.3).
- the INST/FILT rule, planned for SGA2

We deem that these rules represent a good cross-section of currently available rule types, and hence that developing the SpiNNaker API for these rules will facilitate future implementation of many other state-of-the art plasticity rules.

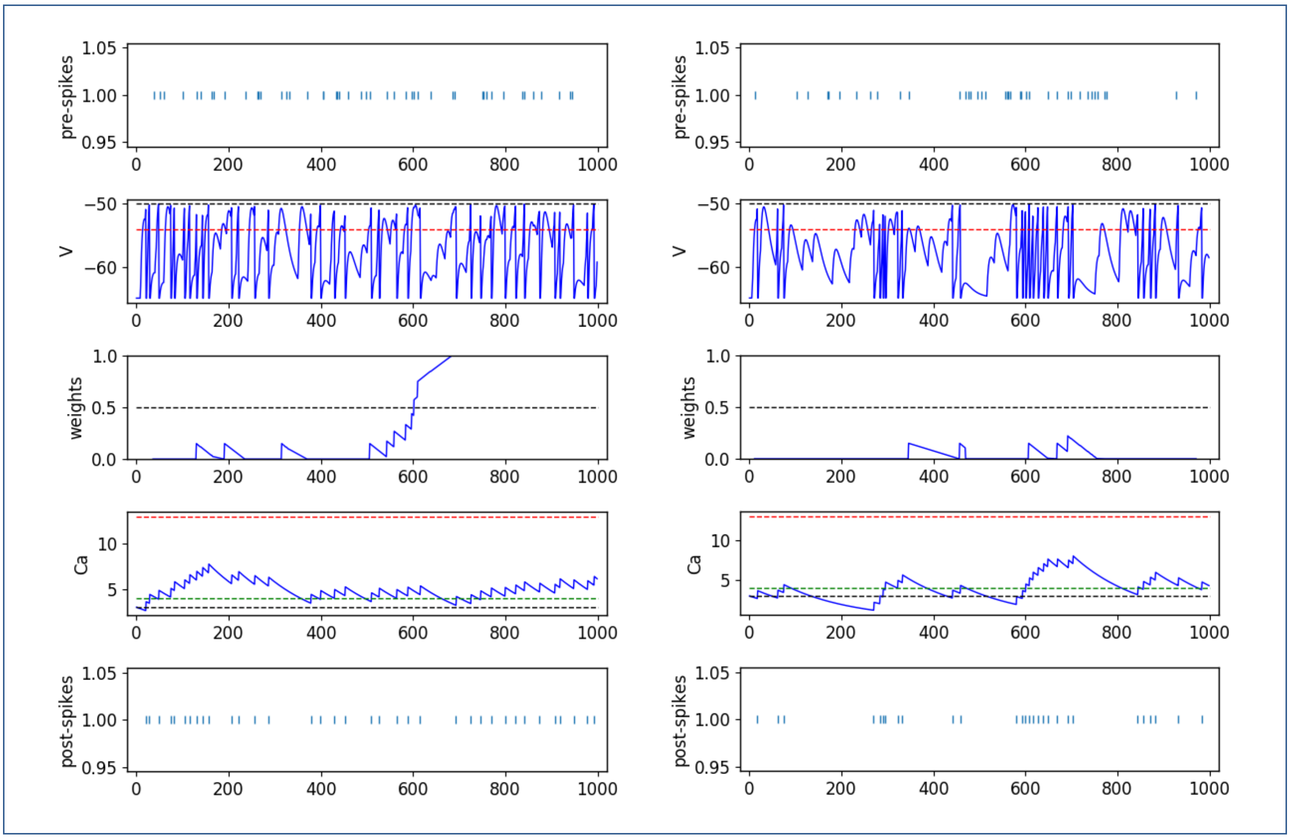

Figure 18: Replication of Figure 1 from (Brader *et al*., Neural Computation 2007)

Two SpiNNaker simulation runs (left/right) of the Brader-Fusi-Senn rule on a single neuron with the same parameters, inputs are Poisson trains with the same mean rates. Due to stochasticity of the input, one run results in a Long Term Potentiation (LTP) transition (left panel, middle subpanel), and the other one does not (right panel). In this rule the weights are binary, and plasticity changes are tracked using bistable internal weight variables with a slow drift towards maximum or minimum weight (third subpanel from the top). The plasticity updates occur at presynaptic spike times: potentiation if V is above the threshold (red line, second subpanel from the top) and Ca is in the interval between black and red lines (fourth subpanel); depression if V is below the threshold, and Ca is between black and green lines; no change otherwise.

2.3.5.1 Impact

67) Supervised learning in spiking neural networks for precise temporal encoding

B. Gardner and A. Grüning

PLoS ONE, 11(8), 1–28 (2016)

68) Towards utilising the dls v2 for supervised learning

B. Gardner and A. Grüning

doi: 10.5281/zenodo.999476 (Sep 2017)

69) Supervised Learning on the SpiNNaker Neuromorphic Hardware E. Nichols, B. Gardner and A. Grüning

doi: 10.5281/zenodo.574582 (July 2017)

2.3.5.2 Component Dependencies

Summarised links to components this Key Result depends on.

Component ID	Component Name	HBP Internal	Comment
457	SP9: BrainScaleS 2 neuromorphic computing system (hardware)	Prototype accessible HBP internal	Hardware emulation environment
467	SP9: SpiNNaker 2 small-scale NM-MC system (hardware)	Prototype accessible HBP internal	Hardware emulation environment
468	SP9: Principles for brain-like computation (model)	No	CDP5-related results
969	SP4: Plasticity - Two-compartment neuron (model)	No	KR builds on and contributes to this model
1066	SP4: Plasticity - Synaptic plasticity and learning (model)	No	KR builds on and contributes to this model
1203	SP4: Plasticity - INST/FILT rule (model)	No	KR builds on and contributes to this model
1341	CDP5: Concept showcases in big systems (model)	No	CDP5 result
1342	CDP5: Guiding platform design on functional plasticity (model)	No	CDP5 result
2420	SP4: Plasticity - Prototype implementations of rules and testing within and without the SP9 platforms (model)	(currently in Fusi_plasticity branch of the official SpiNNaker github repository. Eventually this branch will be merged into master branch)	Towards Third-Factor learning rules on SpiNNaker

2.3.6 Robots learning sensory-motor loops with spiking neurons

J. Camilo Vasquez Tieck, Jacques Kaiser, Rüdiger Dillmann, Wolfgang Maass, Walter Senn, Wulfram Gerstner

CDP5 collaboration between SP10 (FZI Karlsruhe), SP9 (TU Graz) and SP4 (EPFL-LCN Lausanne, U Bern)

In order to use spiking neural networks in closed-loop sensory-motor robotics experiments, we need to integrate models for learning, representation and processing of input and output motor information. Here, we highlight the collaborations to incorporate different models in our robotics scenarios, specifically: reinforcement learning with SPORE and SP9 (UGraz), a dynamic model for a robotic arm SP4 (EPFL-LCN), and one of the learning rules from SP4 (UBern).

An important prerequisite for cross-SP collaboration is the development of a full pipeline where models can be integrated and different experiments can be performed. As a demonstration, we show an initial experiment setup (see Figure 20) containing a ball that has to be steered in the centre of a plane and then balanced there. The vision system is event-based, using a simulation

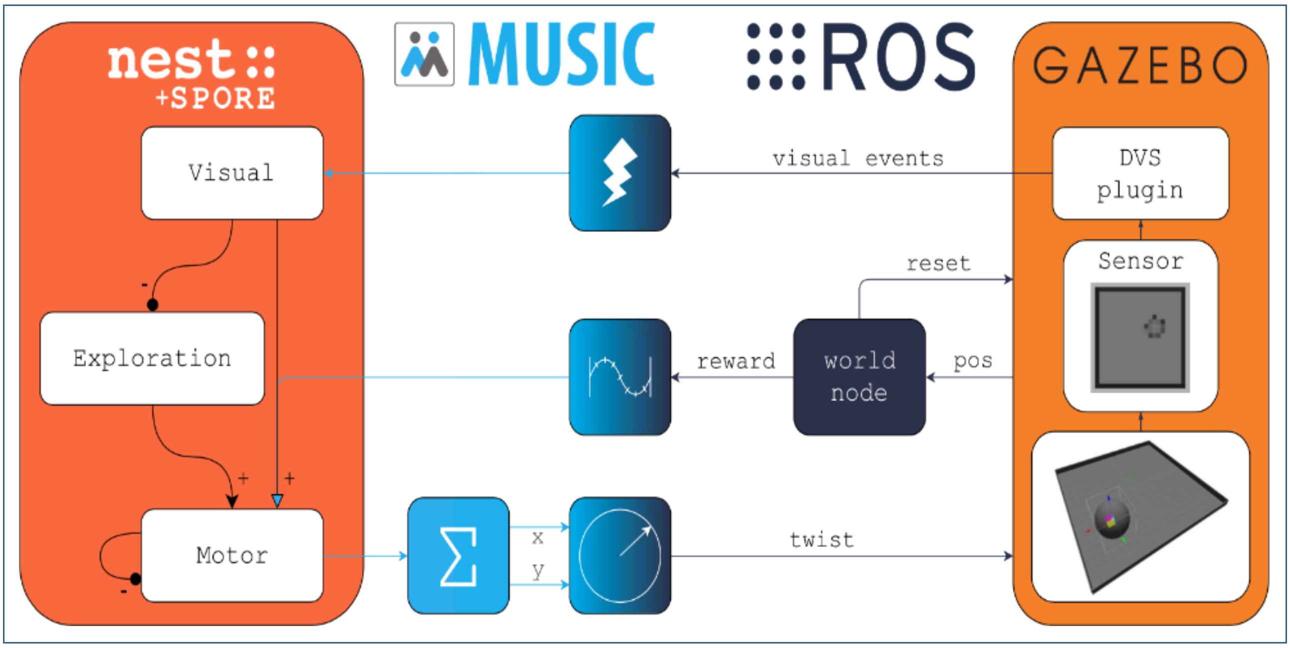

Figure 19: Closed-loop experiment setup.

The main components for communication and synchronization are presented. The SNN is simulated in Nest with the SPORE extension for reinforcement learning. The world and model are simulated in Gazebo. A simulation of a dynamic vision sensor (DVS) is used to get visual feedback.

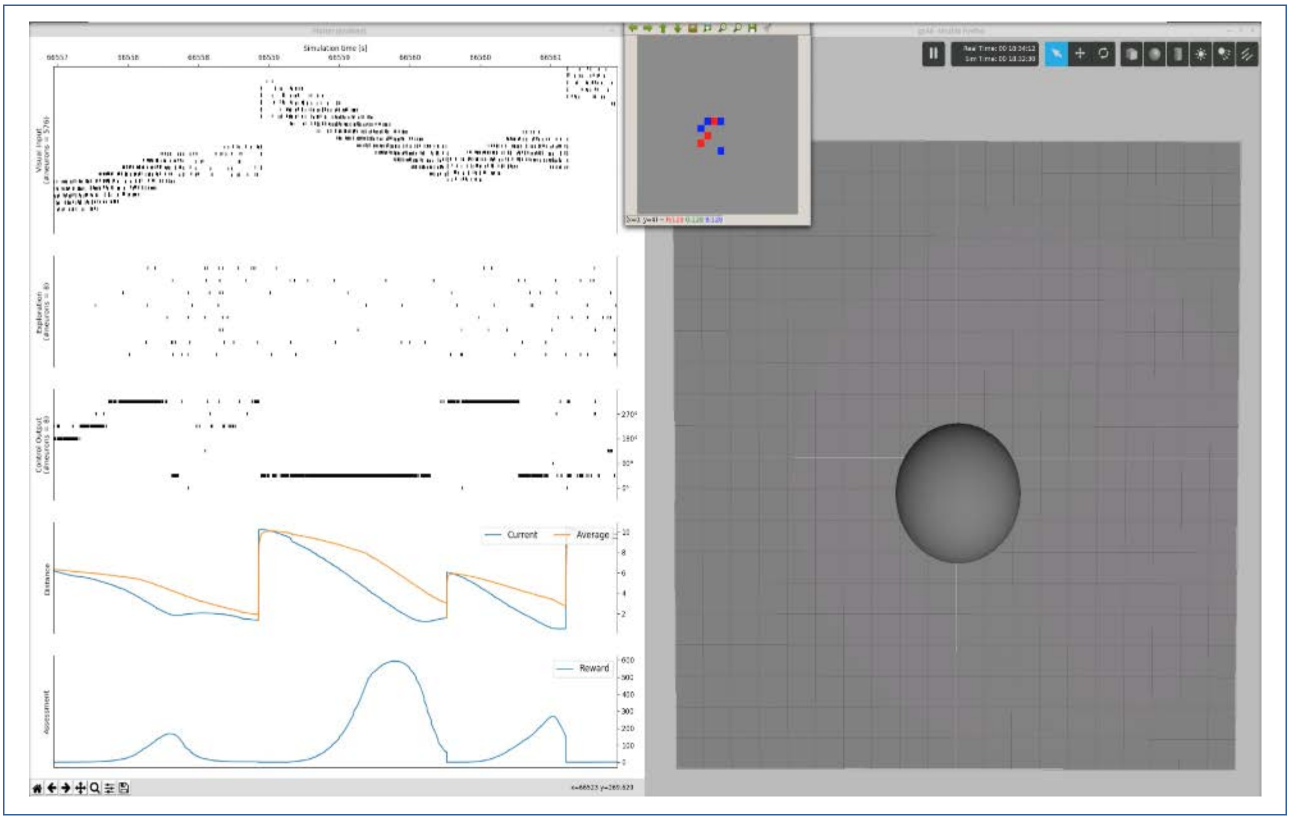

Figure 20: The running system: extract from the simulation after 20 hours of training.

The SNN learns to balance the ball in the centre of the plane. The plots on the left show the SNN activation, the reward signal and the encoded input. On the right, one can see the ball being balanced in the middle of the plane. The top middle panel shows the DVS simulation of the current scene.

for a DVS (Dynamic Vision Sensor) camera. The control of the different degrees of freedom is modelled after biology, using simple muscles for each direction. A reward is calculated based on the relative position of the ball to the centre of the plane. The SPORE (Synaptic Plasticity with Online Reinforcement) framework for reinforcement learning is used for training. Different populations are connected to the DVS input, the reward and the muscles respectively. This ball-balancing with a robotic platform experiment demonstrates the potential of using spiking neurons for closed-loop robot control. The network contains 550 spiking neurons and converges on average after 20 hours of training.

There is a great need of robot controllers that are adaptive, flexible and can learn. With the latest biological inspired models and techniques developed within HBP we were able to perform a validation on different robotics experiments to provide a real/simulated embodiment for the brain and a real application. In the SGA2 phase, the collaboration between SP10, SP4 and SP9 will move towards more complex scenarios and the integration of other models (see KR1.5).

2.3.6.1 Impact

70) Scaling up liquid state machines to predict over address events from dynamic vision sensors

Kaiser, Jacques and Stal, Rainer and Subramoney, Anand and Roennau, Arne and Dillmann, Rüdiger

Bioinspiration & biomimetics. https://doi.org/10.1088/1748-3190/aa7663 (2017)

71) Learning Closed-Loop Robot Control with Spiking Neurons and Event-Based Vision

Michael Hoff

Master Thesis, KIT Karlsruhe and FZI Karlsruhe, 2017

72) Using Spiking Neural Networks in Closed-Loop Sensory-Motor Robotics Experiments

J. Camilo Vasquez Tieck

Talk at Human Brain Project Summit, Glasgow (UK), 2017

2.3.6.2 Component Dependencies

Summarised links to components this Key Result depends on:

Component ID	Component Name	HBP Internal	Comment
924	SP10:NRP - Functional brain model for visual perception (model)	No (papers were accepted, to Bioinspiration & Biomimetics and ICANN 2017	KR builds on and contributes to this model
922	SP10: Functional brain model for humanoid grasping (model)	Yes	KR builds on and contributes to this model
921	SP10: NRP - Functional body and movement learning	Yes	KR builds on and contributes to this model

3. Component Details

The following is a list of the newly released internal Components for this Deliverable.

3.1 CDP5: Concept showcases in big systems (model)

Field Name	Field Content	Additional Information
ID	1341	
Component Type	Model	
Contact	GRÜNING, André	
Component Description	A collection of experiments run on the SP9 neuromorphic systems.	
Latest Release	Release date given by the publication date of the documents listed below.	
TRL	3	
Location	Manchester, Heidelberg	
Format	NA	
Curation Status	PLA registered	
Validation - QC	Unchecked	
Validation – Users	Yes	See publications below
Validation – Publications	Yes	Aamir et al. 2017), Friedmann et al. 2017), Gardner et al. 2017), Nichols et al. 2017), Petrovici et al. 2017a), Petrovici et al. 2017b), Schemmel et al. 2017), Schmitt et al. 2017), Stöckel 2017)
Privacy Constraints	No Privacy Constraint	
Sharing	consortium	
License	Release License Unspecified	
Component Access URL	NA	
Technical documentation URL	NA	
Usage documentation URL	NA	
Component Dissemination Material URL	NA	

3.2 CDP5: Guiding platform design on functional plasticity (model)

Field Name	Field Content	Additional Information
ID	1342	
Component Type	Model	
Contact	GRÜNING, André	
Component Description	SP9 developments with contributions from CDP5.	
Latest Release	Release date given by the publication date of the documents listed below.	
TRL	3	
Location	Manchester, Heidelberg	
Format	NA	
Curation Status	PLA registered	
Validation - QC	Unchecked	
Validation – Users	Yes	See publications below
Validation – Publications	Yes	Aamir *et al.* 2017), Friedmann *et al.* 2017), Gardner *et al.* 2016), Gardner *et al.* 2017), Nichols *et al.* 2017), Schemmel *et al.* 2017), Stöckel 2017)
Privacy Constraints	No Privacy Constraint	
Sharing	consortium	
License	Release License Unspecified	
Component Access URL	NA	
Technical documentation URL	NA	
Usage documentation URL	NA	
Component Dissemination Material URL	NA	

4. Summary and Outlook

The breadth of activities undertaken within the HBP offers unique scientific opportunities, but also represents a challenge for efficient collaboration. This circumstance is particularly relevant for CDP5, given how development and learning are core topics for all Sub Projects, as they represent *the* key for understanding intelligence, be it biological or artificial.

In the two-year period since its inception, CDP5 has become a terrain for true scientific collaborations on the topic of plasticity across the HBP SPs. The major challenge of the SGA1 period was to achieve a common language, of using common tools and of achieving common goals between specialists from areas as diverse as biology, theory, high-performance computing and neuromorphic engineering. As our selection of Key Results shows, this challenge has been met, thereby laying an auspicious groundwork for continuation in the following SGA2 period. The development of an organisational infrastructure consisting of regular discussions, workshops with HBP-internal and external participation and multiple lab exchanges has been instrumental in achieving this goal.

One aspect that is unique to the HBP is the research-driven advancement of novel, beyond von-Neumann, brain-inspired computing technologies, for which CDP5 plays an instrumental role. Guided by insights from experiments and theory, these neuromorphic platforms – including both the hardware and the software infrastructure – are being continuously enhanced, as espoused by several of our Key Results. Owing to these advances, the neuromorphic platforms, in turn, will foster insight into the computational principles of the brain on spatial and temporal scales that are otherwise prohibitive to classical simulation.

Among the achieved Key Results, we highlight the link from the theoretical work on computation in hierarchical (deep) networks to the first small-scale hardware implementation in the BrainScaleS system. On the theory side, significant progress has been made on the role of synaptic and neural dynamics – in particular, the interplay between synaptic plasticity and neuronal compartmentalisation – in shaping the flow of information in such networks; additionally, we are gaining an understanding of how emerging oscillations in large-scale networks may benefit this type of computation. On the emulation side, we have demonstrated the ability to implement these hierarchical networks on a highly accelerated substrate, while training them with a classical computer in the loop; at the same time, we have developed new hardware that is able to accelerate the learning process as well. Due to these advances in the SGA1 period, we have decided to emphasize this line of research in SGA2, within a correspondingly focused CDP5: "Biological deep learning".