CDP5-Biological deep learning - Results for SGA2 Year 2 (D9.4.2 - SGA2)

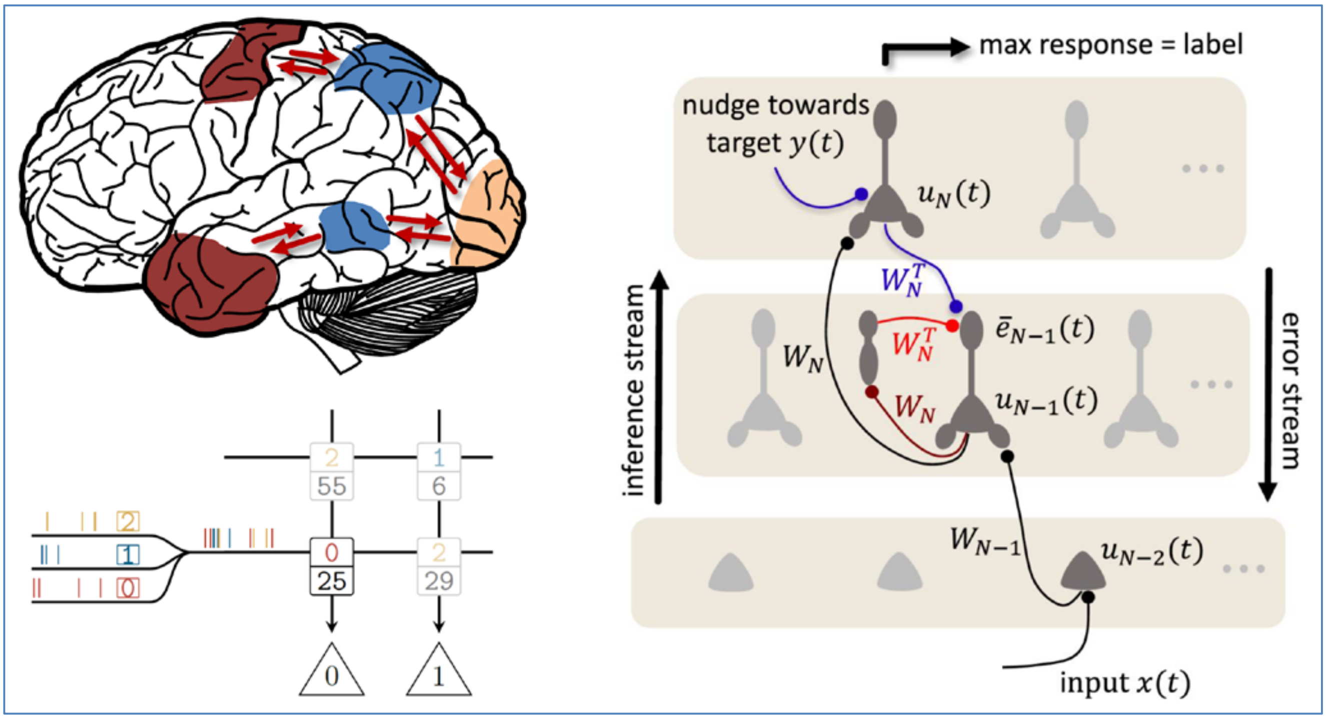

Fig. 1: Biological deep learning - neurophysiology, mathematics and silicon

How can the brain learn a functional, efficient, hierarchical representation of information?

Project Number:	785907	Project Title:	Human Brain Project SGA2

Document Title:	CDP5-Biological deep learning: Results for SGA2 Year 2
Document Filename:	D9.4.2 (D60.2 D469) SGA2 M24 ACCEPTED 201006.docx
Deliverable Number:	SGA2 D9.4.2 (D60.2, D46)
Deliverable Type:	Report
Work Packages:	3.2, 3.3, 3.4, 6.3, 6.4, 7.3, 7.4, 7.5, 9.1, 9.2, 9.3, 9.4, 10.2
Key Result(s):	KRc 5.1, 5.2, 5.3, 5.4, 5.5, 5.6
Dissemination Level:	PU = Public
Planned Delivery Date:	SGA2 M24 / 31 March 2020
Actual Delivery Date:	SGA2 M25 / 27 April 2020, resubmitted 6 Aug 2020; accepted 6 Oct 2020
Author(s):	Mihai A. PETROVICI, UBERN (P71) and UHEI (P47); Martina SCHMALHOLZ, UHEI (P47); Andreas BAUMBACH, UHEI (P47); Sebastian BILLAUDELLE, UHEI (P47); Nicolas DEPERROIS, UBERN (P71); Dominik DOLD, UHEI (P47); Julian GÖLTZ, UHEI (P47); Jakob JORDAN, UBERN (P71); Jacques KAISER, FZI (P52); Laura KRIENER, UBERN (P71); Kwangjun LEE, UVA (P98); Akos KUNGL, UHEI (P47); Pier S. PAOLUCCI, INFN (P92); Cyriel M.A. PENNARTZ, UVA (P98); Korbinian SCHREIBER, UHEI (P47); Jonas STAPMANNS, JUELICH (P20); Kristin VÖLK, UBERN (P71)
Compiled by:	Mihai A. PETROVICI, UBERN (P71) and UHEI (P47); Martina SCHMALHOLZ, UHEI (P47)
Contributor(s):	AALTO (P2): Oliver STRUCKMEIER; JUELICH (P20): David DAHMEN, Markus DIESMANN, Agnes KORCSAK-GORZO, Jonas STAPMANNS, Tom TETZLAFF, Sacha J. VAN ALBADA; UHEI (P47): Andreas BAUMBACH, Sebastian BILLAUDELLE, Oliver J. BREITWIESER, Ilja BYTSCHOK, Benjamin CRAMER, Dominik DOLD, Julian GÖLTZ, Andreas GRÜBL, Nico GÜRTLER, Maurice GÜTTLER, Andreas HARTEL, Arthur HEIMBRECHT, Dan HUSMANN, Kai HUSMANN, Vitali KARASENKO, Gerd KIENE, Johann KLÄHN, Mitja KLEIDER, Christoph KOKE, Alexander KUGELE, Akos F. KUNGL, Luziwei LENG, Christian MAUCH, Karlheinz MEIER, Eric MÜLLER, Paul MÜLLER, Christian PEHLE, Mihai A. PETROVICI, Oskar RIEDLER, Johannes SCHEMMEL, Sebastian SCHMITT, Korbinian SCHREIBER, Philipp SPILGER, David STÖCKEL, Yannik STRADMANN, Timo C. WUNDERLICH; FZI (P52): Rüdiger DILLMANN, Alexander FRIEDRICH, Jacques KAISER, Daniel REICHARD, Arne ROENNAU, Camilo V. TIECK; TUGRAZ (P55): Guillaume BELLEC, Elias HAJEK, Robert LEGENSTEIN, Wolfgang MAASS, Darjan SALAJ, Franz SCHERR; UMAN (P63): Steve FURBER, Oliver RHODES; UBERN (P71): Nicolas DEPERROIS, Jakob JORDAN, Laura KRIENER, Mihai A. PETROVICI, Walter SENN, Kristin VÖLK, Willem WYBO; INFN (P92): Cristiano CAPONE, Bruno GOLOSIO, Pier-Stanislao PAOLUCCI, Elena PASTORELLI; UvA (P98): Shirin DORA, Kwangjun LEE, Sander M. BOHTE, Cyriel M.A. PENNARTZ; UWE (P101): Martin PEARSON; UZ/ETHZ (external partners): Matteo CARTIGLIA, Germain HAESSIG, Giacomo INDIVERI, João SACRAMENTO; UGO (external partner): David KAPPEL; RIKEN (external partner): Maximilian SCHMIDT; UCI (external partner): Emre NEFTCI
SciTechCoord Review:	Martin TELEFONT, EPFL (P1)
Editorial Review:	Guy WILLIS, EPFL (P1)
Description in GA:	CDP5 Annual compound deliverable Year 2 Overview of Key Results and Impact achieved in M13 to M24, tailored for presentation to the relevant audiences (research/industry/public) References to HBP/SP/CDP objectives and Use Cases for navigation between multiple Sub-Projects Linkage of Results to components/a set of component fact sheets (lower level information: Component ownership, Technology Readiness level, performed Quality Control checks, etc.) This deliverable will be a public report including theoretical research on learning principles, their transfer to the HBP computing infrastructures in SP7 and SP9, final experimental results and applications, an evaluation of the developments outside the HBP and a planning outlook to SGA3.
Abstract:	In CDP5, HBP researchers from several SPs collaborate on the topic of biological and artificial intelligence. A key requirement for being classified as “intelligent” is a system’s ability to learn. In the course of the last year, important insights have been gained into the computational principles underlying biological learning and their application to the neuromorphic and neurorobotic platforms developed in the HBP.
Keywords:	Learning, plasticity, (deep) hierarchical networks, spiking neurons, neuromorphic computation, neurorobotics.
Target Users/Readers:	Scientists, companies and other potential users of HBP results.

NOTE:

Components: a full list of Components will be included in the SGA2 Periodic Report.

Dissemination: dissemination actions to promote specific Key Results and the Outputs that contribute to them will be documented in the SGA2 Periodic Report.

Table of Contents

Table of Tables

Table of Figures

Date	Change Requested / Change Made / Other Action
	Deliverable submitted to EC
200729	Resubmission with specified changes requested in Review Report Main changes requested: • Change 1 (Annex 1 - Expert opinion on deliverables: D9.4.2 (D60.2 D469) from Ref. Ares(2020)3990890)
200803	Revised draft sent by CDP5 (edited by Virginie Sabado) to PCO. Main changes made, with indication where each change was made: • Change 1 (see Section 4.5.2) • Change 2 (see Section 5.3.4) • Change 2 (see Section 7.2.2)
200806	Revised version resubmitted to EC by PCO via SyGMa

1 Introduction

To successfully deal with novel situations, adaptation is necessary. This nearly tautological statement holds irrespective of whether we consider biological or artificial problem-solving machines. When the machines in question are neuronal networks, this adaptation usually refers to gradual modifications in the interaction between neurons that lead to a better performance of the network on relevant tasks. As inter-neuron interactions are mediated by synapses, learning in neuronal network models thus boils down to synaptic plasticity. Since the ability to adapt to novel situations in order to achieve a given goal is, arguably, the defining characteristic of intelligence, the study of synaptic plasticity is one of the most interesting and prominent topics in both neuroscience and machine learning.

Spanning both these realms, CDP5 addresses the question of plasticity under two investigative paradigms. Firstly, we operate under a so-called normative approach. This means that we start from what a network "should" do (find an optimal statistical representation of sensory data, assign objects to categories, move in a way that minimizes energy expenditure and/or maximizes reward) and derive biologically plausible plasticity rules that enable the network to achieve its goals given specific natural constraints. Secondly, we assume that it is helpful for networks to represent information hierarchically, that is, different layers (or modules) of a network represent different aspects of the information they are manipulating. A classic example would be the explicit representation of objects in one part of the network and the categories to which they belong in another, separate part. Such an information hierarchy lies behind some of the most exciting recent advances in machine learning, while also matching well-established knowledge about cortical anatomy. However, efficient learning in such networks is not easily reconciled with biological constraints such as locality (neurons and synapses can essentially only access information from their immediate vicinity), which raises the fundamental research question addressed by our project: how can biologically constrained neuronal networks learn to solve tasks that profit from a hierarchical organization of information within them?

2 Overview

Interestingly, the recent explosion in deep-learning-related research was fuelled less by theoretical breakthroughs rather than, to a much larger extent, by advances in computing hardware and the increasing availability of large sets of labelled data. Many of the core underlying architectures (hierarchical networks) and algorithms (backpropagation of errors) are decades-old and, to some extent, inspired by neuroanatomy and physiology. However, deep learning has quickly left behind the realm of biological plausibility in favour of raw performance, leading many to believe that what started out as a brain-inspired computational paradigm has since become fundamentally incompatible with its biological archetype (see, e.g., Crick, Nature 1989).

It is only very recently that computational neuroscience has begun reassessing this assumption, with CDP5 researchers spearheading this endeavour and successfully solving many of the issues pointed out by Crick thirty years earlier (Richards *et al.*, Nature 2019). New theoretical insights developed in the HBP have made significant strides in reuniting the previously diverging worlds of biological and artificial neural networks, with important implications for understanding cortical processing, while bolstering the range of applications for neuromorphic hardware.

CDP5 spans a broad range of topics related to biologically plausible learning across information hierarchies in neuronal[1] networks. Many important contributions have been carried over and developed from CDP5 research in SGA1 and, naturally, the first half of SGA2. One of our core lines of research concerns the development of cortical microcircuit models, theories and EBRAINS implementations capable of learning hierarchical representations of visual data (Outputs 8, 9, 11 and 14). The underlying neuron and synapse models carry further computational properties useful

---

[1] We use "neuronal" as opposed to "neural" to emphasize biological plausibility.

for spatiotemporal learning and Bayesian computation (Outputs 2 and 10); importantly, they were also found to represent local optima of interpretable metalearning (Output 16). Continuing work that goes back as far as the RUP, we have made significant developments in the field of spike-based sampling (Outputs 3, 4, 5, and 6), as well as within the context of computation during sleep, with interesting convergence between initially separate lines of research (Outputs 3, 12, 13). Last but not least, the collaborations fostered by CDP5 have enabled the realization of many of our models within neuromorphic and neurorobotic platforms (Outputs 1, 5, 6, 7, 14, 15, 17).

In this document, we provide a brief description of our most recent results, organized along the five Key Results defined in our research plan for SGA2. We address the impact and significance of our research and provide a selection of publications for some of our more mature lines of work. Beyond the advancement of our scientific understanding, the collaborative work performed under the aegis of CDP5 has also allowed important insights about how funding body policies and incentives interact with the modus operandi and outcomes of science in large but diverse collaborations, which we address at the end of the document.

Table 1: Summary of Output-related links

C No.	Component Name	Link to	URL
C0001	SP9 BrainScaleS-1 Neuromorphic Computing System (version 1 = NM-PM1)	Technical Documentation	https://flagship.kip.uni-heidelberg.de/jss/FileExchange/D9.7.1_Neuromorphic_Platform_Specification_-_public_version.pdf?fID=1887&s=qqdXDg6HuX3&uID=65
		User Documentation	http://electronicvisions.github.io/hbp-sp9-guidebook/
C0209	NEST - The Neural Simulation Tool	Software Repository	https://github.com/nest/nest-simulator and P2517
		User Documentation	https://nest-simulator.readthedocs.io/en/stable/#how-the-documentation-is-organized
C0349	PyNN	Technical Documentation	http://neuralensemble.org/docs/PyNN/developers_guide.html
		User Documentation	http://neuralensemble.org/docs/PyNN/
C0457	BrainScaleS 2 Neuromorphic Computing System	User Documentation	https://electronicvisions.github.io/hbp-sp9-guidebook/
C1032	Plasticity: dendritic predictive plasticity that reproduces STDP data (Algo STDPpredictive)	User Documentation	N/A (manuscript in preparation)
C1788	T3.2.1 (2) Multipurpose simplified neuronal network model of different cortical areas matching SWA/wake transitions	User Documentation	https://collab.humanbrainproject.eu/#/collab/67068/nav/455631
C2060	T3.3.3 (1) Multi-area ensemble mechanisms of multi-feature detection in rodents	Documentation	Partially under Embargo, related publications: P2345, P2129
C2061	T3.3.3 (2) Multi-area ensemble mechanisms of object recognition in rodents	Documentation	Partially under Embargo, related publications: P2345, P2129
C2193	T3.5.2 (1) Cortical spiking model of the interplay between sleep and plasticity	Documentation	https://github.com/PierStanislaoPaolucci/2019thalCort-SNN-SO-AW-mem
C2226	T3.5.3 (1) Generative model of sensory cortical hierarchy and corticohippocampal network	Documentation	Partially under Embargo, related publications: P843, P2241, P1360, P1447, P2345, P2129 Related code: https://zenodo.org/record/3675212

C No.	Component Name	Link to	URL
C2228	T3.5.3 (3) Comparison with physiological data and related models	Documentation	Partially under Embargo, related publications: P2345, P2129
C2321	T3.5.3 (4) Integration of corticohippocampal network in Cognitive Architecture	Documentation	Partially under Embargo, related publications: P2345, P2129
C2419	SP4 SGA2 Plasticity: algorithms for multi-compartment models	Documentation	N/A (Journal paper in preparation)
C2420	Prototype implementations of rules and testing within and without the SP9 platforms	Documentation	Related publications: P2241, P1360, P1447, P2020, P2241, P2239
C2439	SP9 Methods for hierarchical neural sampling in networks of spiking neurons	Documentation	Related publication: P843, P2241, P1360, P1447
C2547	SP9 Computing with structured neurons	Documentation	Related publications: P253, P1905, P2536, P1999, P23584
C2696	Prototype implementation of framework and numerics for phenomenological neuron model with compartmentalized third factor plasticity	Technical Documentation	https://github.com/nest/nest-simulator/pull/1095 and https://github.com/nest/nest-simulator/pull/1257
		User Documentation	https://nest-simulator.readthedocs.io/
C2704	SGA2-C10.2.1.4 Scene Representation and Understanding	Documentation	Related publications: P2519, P1836
C2719	SGA2-C10.2.1.3 Learning object affordances	Documentation	Related publications: P2519, P1836
C2722	T3.5.4 (2) Multicompartmental pyramidal neuron models	Documentation	Related publication: P2534, P1905, P2536, P1999, P2538

3 Key Result KRc5.1: Closed-loop navigation

3.1 Overview of Outputs

- Output 1: Insect-inspired navigation

3.2 Insect-inspired navigation

Korbinian SCHREIBER, Sebastian BILLAUDELLE, Yannik STRADMANN, Benjamin CRAMER, Andreas BAUMBACH, Dominik DOLD, Julian GÖLTZ, Akos F KUNGL, Timo C WUNDERLICH, Andreas HARTEL, Eric MÜLLER, Oliver BREITWIESER, Christian MAUCH, Mitja KLEIDER, Andreas GRÜBL, David STÖCKEL, Christian PEHLE, Arthur HEIMBRECHT, Philipp SPILGER, Gerd KIENE, Vitali KARASENKO, Walter SENN, Mihai A PETROVICI, Johannes SCHEMMEL, Karlheinz MEIER

CDP5 collaboration between SP9 (UHEI, P47) and SP4 (UBERN, P71). Both partners have jointly developed exploitation strategies of the BSS-2 platform, in particular regarding experiments on the BSS-2 prototype, which led to a joint publication. UHEI has developed the specific implementation of the insect experiment on the BSS-2 prototype and has performed the associated emulations.

A virtual insectoid agent uses path integration to navigate back home after spreading out randomly. Both the navigating neural network emulations and the environment and agent simulations are carried out on a BrainScaleS-2 prototype chip. This experiment highlights the versatility of the plasticity processing unit implemented in the BrainScaleS-2 architecture.

3.2.1 Research and Infrastructure

Based on physiological data from the bee's central complex, we emulated a network for path integration that reproduces bees' ability to return to their nest's location after exploring the environment for sources of food. Each experiment started with a spread-out phase, in which a virtual insect performed a random walk starting from a certain origin. In the second part, the return phase, the insect's motion was determined by the emulated neural network.

The plasticity processor handled multiple tasks: the processing of synaptic modulations for short-term memory neurons, the simulation of the 2D-environment, an emulation of all sensors including the corresponding spike stimuli, the translation of neuronal data into actions of motion, and the entire experiment control. The total flight duration was set to 200 ms on the hardware, which emulates 200 s in biology. In that time, sensory information and steering signals were exchanged between body and brain every 100 µs. During the first 50 ms, the insect performed its random outbound journey, after which it returned to the nest. The average spike rate of all neurons and spike generators was 300 kHz (30 Hz bio), which is in good agreement with experimental data from drosophila or locusts. Apart from the setup and readout phase, the experiment ran entirely self-contained on the BrainScaleS-2 system.

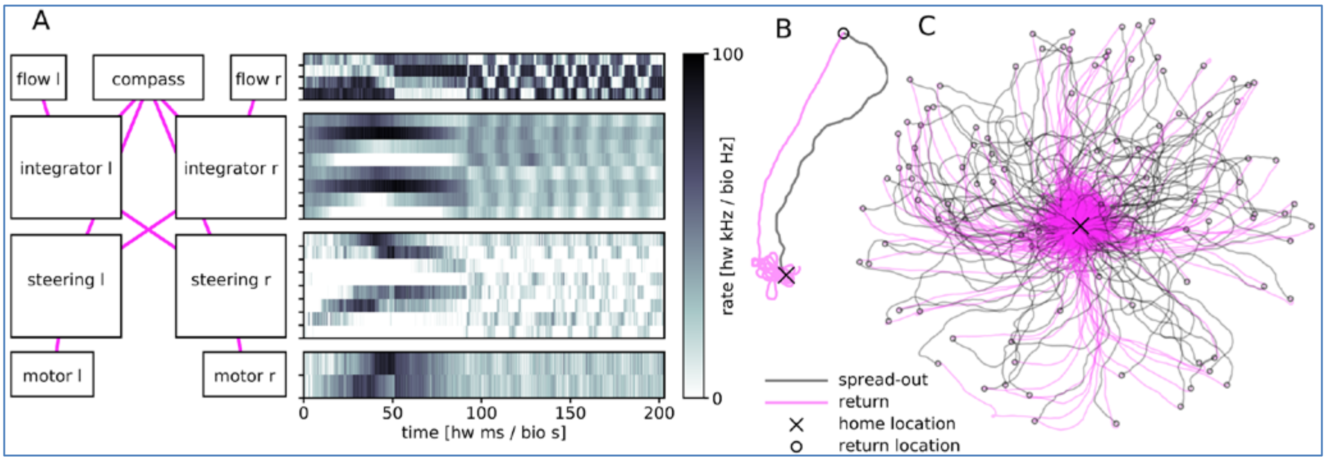

Fig. 2: Insect-inspired navigation

(A) Schematic network topology and neural activity of an entire experimental run. (B, C) Sample trajectories.

3.2.2 Impact and significance

Developing this experiment challenged the computational performance of the digital co-processor of the BrainScaleS-2 prototype chip to its current limit. Features like timed interrupts were used and tested for the first time and revealed certain aspects that led to further improvement of the cross-compiler. This experiment is the first real-time closed-loop robotic experiment to run entirely on a neuromorphic chip. Moreover, it is one of the first biologically plausible neural models of behaviour to be implemented on neuromorphic hardware.

3.2.3 Components

C0209, C0457, C2420

3.2.4 Publications

Sebastian Billaudelle, Yannik Stradmann, Korbinian Schreiber, Benjamin Cramer, Andreas Baumbach, Dominik Dold, Julian Göltz, Akos F Kungl, Timo C Wunderlich, Andreas Hartel, Eric Müller, Oliver Breitwieser, Christian Mauch, Mitja Kleider, Andreas Grübl, David Stöckel, Christian Pehle, Arthur Heimbrecht, Philipp Spilger, Gerd Kiene, Vitali Karasenko, Walter Senn, Mihai A Petrovici, Johannes Schemmel, Karlheinz Meier. *Versatile emulation of spiking neural networks on an accelerated* *neuromorphic substrate.* https://arxiv.org/abs/1912.12980, accepted at the International Symposium of Circuits and Systems 2020, to appear in IEEE Xplore (2019) (P2241).

4 Key Result KRc5.2: Deep spatiotemporal prediction

4.1 Overview of Outputs

- Output 2: Sequence learning by shaping hidden connectivity
- Output 3: Slow waves as tempering
- Output 4: Deterministic networks for probabilistic computing
- Output 5: Accelerated physical emulation of spike-based Bayesian inference
- Output 6: Stochasticity from function in spiking sampling networks
- Output 7: Fast and deep neuromorphic learning with time-to-first-spike coding

4.2 Sequence learning by shaping hidden connectivity

Kristin VÖLK, Walter SENN, Mihai A. PETROVICI

CDP5 collaboration between SP4 (UBERN, P71) and SP9 (UHEI, P47). Both partners contributed to the design of the model and simulation scenarios. The simulations were performed at UBERN.

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

4.2.1 Research and Infrastructure

In the previous CDP5 Deliverables, we showed how a sparse scaffold of somato-somatic connections, formed during network development, can guide the learning of dense somato-dendritic connections. We showed how this enables learning of non-Markovian patterns and illustrated some benefits of learning in the hidden pool by contrasting it with having only somato-somatic connections and no learning in the hidden population. We showed how this learning in the hidden population leads to the network being robust to “deletion” of partial sequences and, under certain circumstances, to “deletion” of the complete pattern. Importantly, the neuron and synapse model used here can be derived within a normative framework and carries relevance for several other outputs (8, 9, 10 and 11).

Since the last reporting period, we investigated the importance of learning in the hidden pool further. The first step here was to make the network even more biologically realistic by introducing random transmission delays between somato-somatic and somato-dendritic synapses. In the previous model, the somato-somatic delays were all 10ms and somato-dendritic transmission was instantaneous. Now the delay for each synapse is drawn randomly from a uniform distribution between 5ms and 15ms. With the introduction of these random delays, sequence states and somato-somatic delays are no longer matched. Furthermore, we introduced a more complex pattern, namely we took a snippet from a midi-file representing the piano piece for “*Für Elise*” (Fig 3(a-b)).

With this mismatch between sequence timing and synaptic transmission delays, the importance of learning in the hidden pool is even more apparent. Where there is no somato-dendritic learning (Fig 3 (a)) in the hidden population, the complex pattern is not even learned correctly and the performance of pattern replay decreases rapidly (Fig 3(a) and Fig 3(c)). However, once we introduce somato-dendritic learning among the neurons in the hidden population, the pattern can be learned correctly (Fig 3(b)) and pattern replay performance stays at a generally high level (Fig 3(b) and Fig 3(c)) over 55 cycles. Given these findings, we conclude somato-dendritic learning is not only vital for achieving robustness to pattern disruptions, but also for general learning when pattern delays and synaptic delays are mismatched, which is the case in most realistic learning scenarios. It also shows that somato-dendritic learning is capable of reordering the given somato-somatic scaffold in a way that aids the learning of the desired pattern.

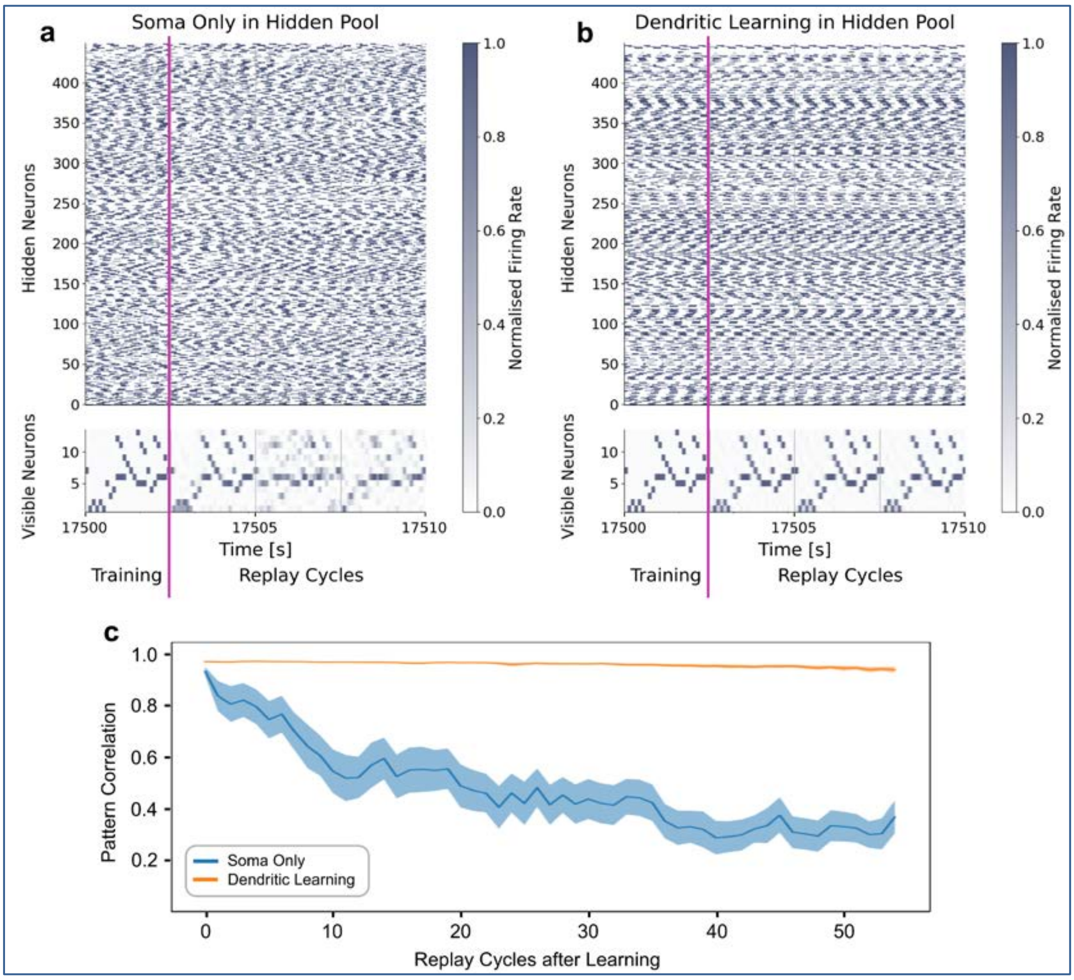

Fig. 3: Sequence learning by shaping hidden connectivity

(a) Raster-rate plot of network with random dendritic delays (both somato-somatic and somato-dendritic). The hidden neurons are simulated as point neurons only governed by their somatic inputs. Hence, there is no learning in the hidden pool. The visible neurons are two-compartment neurons with somato-dendritic learning. Consequently, the only learning happens among visible neurons and between hidden-to-visible connections. On the left side of the violet vertical line the last training cycle (after 7,000 cycles) is shown. On the right side three consecutive replay cycles after the end of training are shown. (b) The same as in (a) but here the hidden neurons are modelled as two-compartment neurons and hence somato-dendritic learning is activated among all neurons, including the hidden-to-hidden connections. (c) Comparison of the ongoing performance during pattern replay after learning measured as correlation between the desired pattern and the actual produced pattern. In orange the mean performance of networks with hidden somato-dendritic learning is shown (as in (b)). In blue the mean performance of networks without hidden somato-dendritic learning is shown (as in (a)). Shading represents the standard deviation of the mean.

4.2.2 Impact and significance

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

4.2.3 Components

C2226, C2419, C2420, C2547, C2722

4.2.4 Publications

Journal paper in preparation.

4.3 Slow waves as tempering

Agnes KORCSAK-GORZO, Luziwei LENG, Andreas BAUMBACH, Oliver J. BREITWIESER, Sacha J. VAN ALBADA, Walter SENN, Karlheinz MEIER, Mihai A. PETROVICI

CDP5 collaboration between SP4 (UBERN, P71, JUELICH, P20) and SP9 (UHEI, P47). All three partners contributed to the design of the model and simulation scenarios. The simulations were performed at UHEI and JUELICH.

It is well-known that slow waves appear during sleep and have thus been hypothesized to play a role in memory formation. We propose a mechanism by which slow waves can improve the generative properties of neural sampling networks, thereby enabling the replay of a more diverse set of memories.

4.3.1 Research and Infrastructure

Under the hypothesis of neural sampling, aspects of cortical spiking activity can be interpreted as sampling from some underlying distribution that represents the space of possible realities that are compatible with sensory information. However, all sampling methods suffer from the fundamental problem that the dynamics of the sampling system might not reach all relevant parts of the tested distribution within a reasonable amount of time, a difficulty often referred to as “the mixing problem”.

The dynamics in our hierarchical network models of LIF neurons performing Bayesian inference essentially implement such a sampling method (see also Outputs 4, 5 and 6). As an example of the mixing problem, we can see that, especially when trained to a high degree of classification accuracy, model networks tend to generate images from few or even only a single mode (for example, a single MNIST digit class). As reported previously (CDP5/SP9 in SGA1), mixing can be improved by employing short-term plasticity, but without the guarantee of conserving the target distribution.

Here, we turn our attention to another ubiquitous aspect of cortical dynamics: population-level oscillatory activity, which is particularly pronounced during slow-wave sleep (see also Outputs 12 and 13). We propose that these slow oscillations implement a tempering scheme, similar to the adaptive simulated tempering (AST) algorithm. Modulating the intensity of the background activity effectively changes the Boltzmann temperature at which the system operates. Increasing this temperature suppresses the correlations that keep the system in its old dynamic state and thereby limit the reach of the dynamics within the system’s state space. Doing the opposite, namely decreasing the temperature, freezes out variations, enabling the system to settle on a distinct mode. Slow waves naturally implement a sequence of these two phases, thereby allowing the networks to jump quickly between different imprinted memories.

Effective tempering hinges on a careful balance between background excitation and inhibition. In such a regime, the entropy of the resulting distribution matches the oscillatory background behaviour, increasing for high temperatures and decreasing for low temperatures. The resulting sampling mechanism exhibits significantly improved mixing behaviour, as indicated by the faster switching between modes.

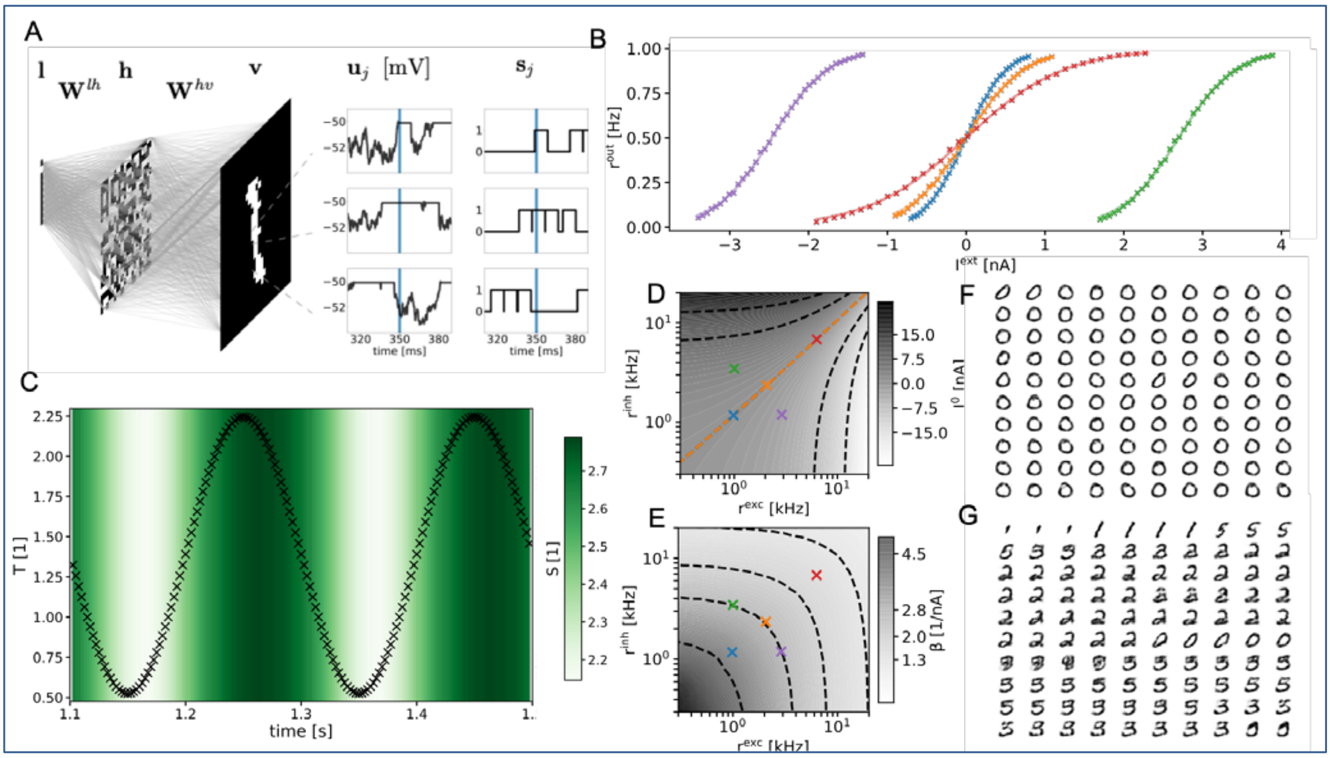

Fig. 4: Slow waves and tempering

(A) Hierarchical spiking network with label (left), hidden (middle) and visible (right) layer; exemplary membrane traces and associated state assignment (z=1 for refractory neurons, z=0 otherwise) for 3 visible neurons. Neurons are connected layer-wise and receive excitatory and inhibitory stochastic noise spikes. (B) Exemplary response functions of a LIF neuron for different background noise configurations. The red, orange and blue curves are obtained for well-balanced excitation and inhibition. (C) Tempering schedule (crosses) and resulting entropy of the sampled distribution (colour shade). Increased temperature (noise level) leads to higher-entropy distributions. (D) Mean and (E) width of the neuron response function. Coloured crosses correspond to the exemplary functions from B. Arbitrary changes can induce shifts of the response function (e.g. green and purple), while increased noise activity widens the response function (increased effective temperature). (F) Visible layer activation probabilities under static noise. Consecutive samples are drawn 10 s apart, the network is stuck in the “zero” mode. (G) Like (F) for tempered noise with a schedule period of 400 ms. Consecutive samples drawn at the cold-hot T=1 transition point, the network quickly switches through different modes (digits), while generating clearly discernible digits.

4.3.2 Impact and significance

In this work, we address a fundamental issue of the neural sampling hypothesis, namely the apparent adversity of precision and diversity, which are both required for agents to build good internal models of their surroundings based on prior experience. By linking the well-established phenomenon of cortical waves to classical tempering studied in physics and machine learning, we propose a functional role for slow waves during memory formation and retrieval. This work therefore carries relevance for both the neuroscience community and for the development of *in-silico* implementations of brain-inspired computing paradigms.

4.3.3 Components

C0209, C0349, C1788, C2193, C2226, C2439

4.3.4 Publications

Journal paper in preparation. Code released under:

- https://github.com/unibe-cns/spike-based-sampling/
- https://zenodo.org/record/3675212 (P2517)

4.4Deterministic networks for probabilistic computing

Jakob JORDAN, Mihai A. PETROVICI, Oliver BREITWIESER, Johannes SCHEMMEL, Karlheinz MEIER, Markus DIESMANN & Tom TETZLAFF

CDP5 collaboration between SP7 (JUELICH, P20), SP9 (UHEI, P47) and SP4 (UBERN, P71). All three partners contributed to the theoretical approach, the design of the model, the execution of simulations and the evaluation of data.

Neuronal network models of high-level brain functions such as memory recall and reasoning often rely on the presence of some form of noise. The majority of these models assumes that each neuron in the functional network is equipped with its own private source of randomness, often in the form of uncorrelated external noise. However, it was unclear what constitutes a suitable noise source for stochastic computations *in vivo*. We demonstrate how deterministic recurrent neuronal networks can be used as sources of uncorrelated noise, exploiting the decorrelating effect of inhibitory feedback.

4.4.1 Research and Infrastructure

In vivo, synaptic background input has been suggested to serve as the main source of noise in biological neuronal networks. However, the finiteness of the number of such noise sources constitutes a challenge to this idea. We demonstrated that shared-noise correlations resulting from a finite number of independent noise sources can substantially impair the performance of stochastic network models. We showed that this problem is naturally overcome by replacing the ensemble of independent noise sources by a deterministic recurrent neuronal network. By virtue of inhibitory feedback, such networks can generate small residual spatial correlations in their activity which, counter to intuition, suppress the detrimental effect of shared input. We exploited this mechanism to show that a single recurrent network of a few hundred neurons can serve as a natural noise source for a large ensemble of functional networks performing probabilistic computations, each comprising thousands of units (see also Outputs 3, 5, and 6).

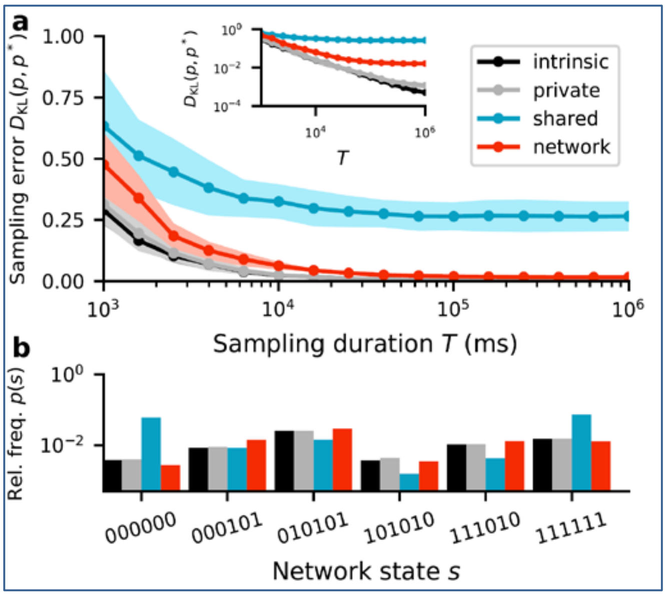

Fig. 5: Recurrent neuronal networks can serve as a suitable source of noise for stochastic computations.

(a) Sampling error as measured by Kullback-Leibler divergence between the empirical state distribution p of a sampling network and the state distribution p* generated by the corresponding Boltzmann machine as a function of the sampling duration T for different sources of noise (intrinsic: intrinsically stochastic units updating their states with a probability determined by their total synaptic input; private: deterministic units receiving private additive independent noise; shared: deterministic units receiving noise from a finite population of independent stochastic sources; network: deterministic units receiving quasi-random input generated by a finite recurrent network of deterministic units). Error bands indicate mean +/- SEM over 5 random network realizations. Inset: Same data as main panel in double-logarithmic representation. (b) Relative frequencies (vertical, log scale) of six exemplary states s (horizontal) for T = 10^6 ms.

4.4.2 Impact and significance

Our results demonstrate that background input from an active surrounding network can serve as a suitable source of noise for stochastic computations in neuronal networks, lending biological plausibility to many network models that rely on noisy background input. The proposed approach is sufficiently general to support a wide variety of network models without significant changes and provides a possibility for generating stochasticity in neuromorphic systems that often suffer from bandwidth limitations.

4.4.3 Components

C0001, C0209, C0349, C0457, C2226, C2439

4.4.4 Publications

Jakob Jordan, Mihai A. Petrovici, Oliver Breitwieser, Johannes Schemmel, Karlheinz Meier, Markus Diesmann, and Tom Tetzlaff. *Deterministic networks for probabilistic computing*. Scientific reports, 9(1), 1-17, (2019) (P843)

4.5 Accelerated physical emulation of spike-based Bayesian inference

Akos F. KUNGL, Sebastian SCHMITT, Johann KLÄHN, Paul MÜLLER, Andreas BAUMBACH, Dominik DOLD, Alexander KUGELE, Eric MÜLLER, Christoph KOKE, Mitja KLEIDER, Christian MAUCH, Oliver BREITWIESER, Luziwei LENG, Nico GÜRTLER, Maurice GÜTTLER, Dan HUSMANN, Kai HUSMANN, Andreas HARTEL, Vitali KARASENKO, Andreas GRÜBL, Johannes SCHEMMEL, Karlheinz MEIER, Mihai A PETROVICI

CDP5 collaboration between SP9 (UHEI, P47) and SP4 (UBERN, P71). UHEI and UBERN have jointly worked on the theory, network models and hardware implementation. UHEI was responsible for the commissioning of the system, providing the required hardware and software infrastructure.

Neuromorphic hardware seeks to achieve fast, parallel and robust computation by mimicking aspects of the nervous system. However, it needs appropriate algorithms and coding schemes, which can exploit the advantages of neuromorphic hardware. In previous works, we proposed a sampling-based framework for spiking neural networks (SNNs) both as models of cortical computation and as candidates for solving problems in machine learning. We adapted the framework and implemented it on the BrainScaleS-1 and BrainScaleS-2 neuromorphic systems. Of special interest for neuromorphic hardware with analogue components is whether and how the model can handle substrate-inherent difficulties, such as limited control over the parameters.

4.5.1 Research and Infrastructure

The proposed sampling framework and the accelerated spiking neuromorphic BrainScaleS systems excellently complement each other in many aspects. First, the sampling framework explicitly uses spikes as a computation paradigm, and not averaged quantities such as rates (see also Outputs 3, 4 and 6). Furthermore, sampling inherently benefits from the acceleration of the hardware, since many samples are needed to reach an appropriate estimate of the posterior distribution. Finally, local Hebbian-like learning has the potential to act as an on-chip learning procedure that exploits the acceleration of the system, since no global information (like global network states or weights) have to be shifted around from memory.

The noise necessary for sampling stems from an inhibition-dominated sparse random network to circumvent external bandwidth limitations between host computer and the chip. Such networks have been found to display chaotic dynamics in theory as well as neuromorphic implementations (Pfeil *et al*. 2016).

We show that the proposed sampling networks can cope with the challenges arising on analogue neuromorphic hardware, such as inhomogeneity (fixed-pattern noise) of the realized neurosynaptic circuits and limited weight resolution (4bit). This allows our model to exploit the fast emulation speeds (1,000 - 10,000 compared to biology) offered by the BrainScaleS systems, while experiencing only a negligible performance loss. We demonstrated this for three network architectures: firstly, we trained a fully-connected network on hardware with local contrastive Hebbian plasticity to sample from a predetermined target distribution and perform inference, i.e., calculate the correct conditional distributions. Secondly, we implemented a hierarchical network for classification, inference and data generation trained on a subset of the MNIST dataset. Finally, we trained an ensemble of deterministic sampling networks on neuromorphic hardware, realizing a closed physical implementation of sampling with deterministic networks shapeable by local contrastive Hebbian plasticity. We demonstrated the capabilities of the network in three different applications: Classification, pattern completion and guided data generation (Fig. 6). The hardware implementation reached a performance on a par with equivalent Restricted Boltzmann Machines run in software simulation.

Fig. 6: Accelerated physical emulation of spike-based Bayesian inference.

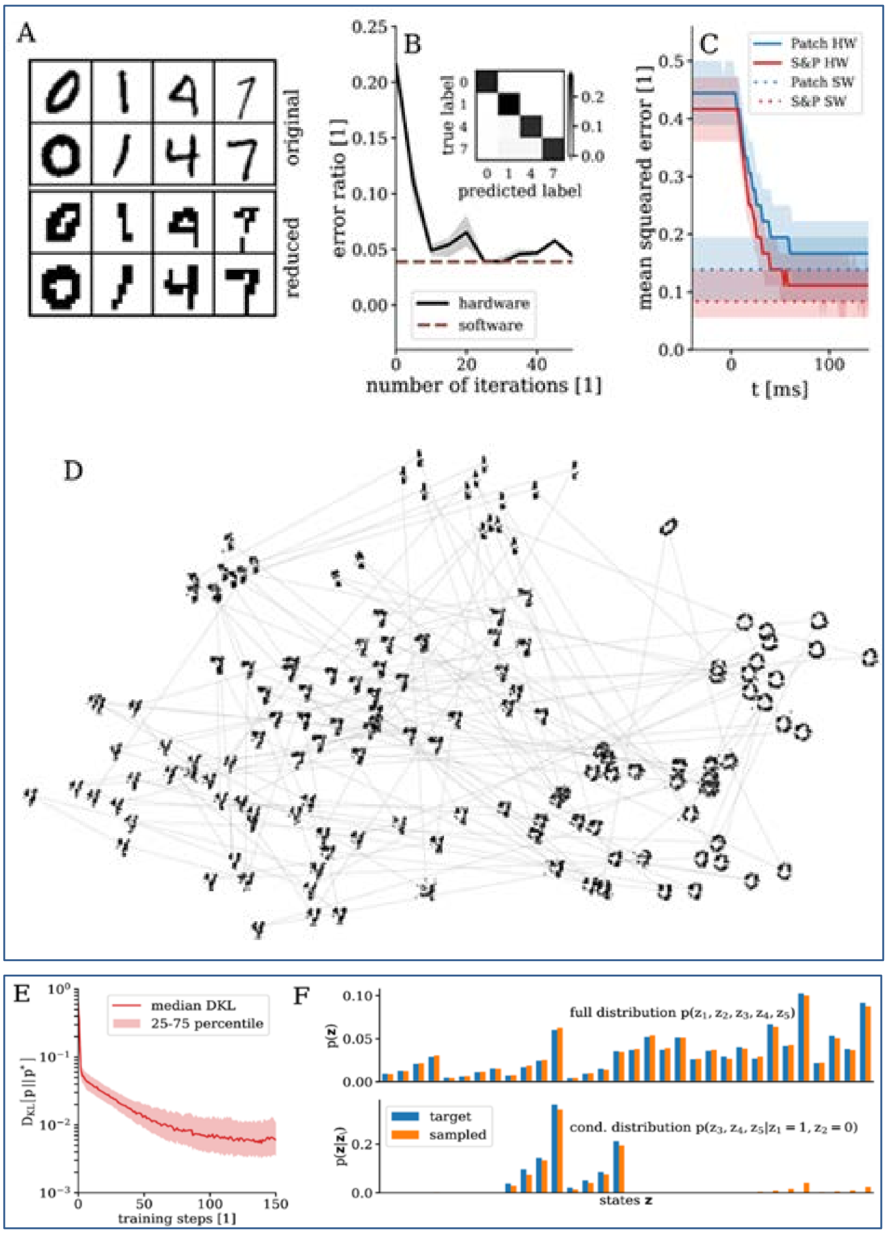

(A) We trained a hierarchical spiking sampling network on a reduced version of the MNIST dataset. We had to reduce the resolution of the dataset due to technical issues on the system limiting the size of implementable networks. (B) By training the system in the loop using a form of Hebbian contrastive plasticity we were able to restore the classification performance that was lost due to the translation from the abstract neural network domain to the spiking neuromorphic domain. (C) The network can perform pattern completion as measured by the mean squared error between the original image and network response. The neuromorphic realization is approximately on par with the original software solution and, in addition, strongly accelerated. (D) In dreaming mode, the network produced recognizable images from all classes, resembling the training set. We reached similar results on a reduced version of the Fashion MNIST dataset, see Kungl *et al*. (2019). (E) We successfully transferred the method to implement spike-based Bayesian inference on the DLSv2 prototype chip of the BSS-2 system. Sampling performance after training for 500 randomly generated target distributions. (F) Sampling from the learned (top) and an associated conditional distribution (bottom). Orange: sampled distribution. Blue: analytically calculated target distribution. Remaining error bars are too small to visualize. See also Billaudelle *et al*. (2019).

4.5.2 Impact and significance

The project has demonstrated the applicability of the spike-based sampling network on accelerated neuromorphic hardware. The implementation is unique in the sense that the underlying model explicitly uses the spiking nature of the neuron model and that the full implementation resulted in an autonomous sampling model. The established method proved to be robust enough that we could easily implement sampling on a prototype chip of the BrainScaleS-2 system (Billaudelle *et al*. 2019). Talk on neuromorphic computing at the NICE (Neuro-Inspired Computational Elements) Workshop, Albany, March 27, 2019 (recorded talk E2633).

4.5.3 Components

C0001, C0209, C0349, C0457, C2226, C2420, C2439

4.5.4 Publications

- Spike-based sampling on the BrainScaleS-2 prototype chip was published as part of the study:
  - Sebastian Billaudelle, Yannik Stradmann, Korbinian Schreiber, Benjamin Cramer, Andreas Baumbach, Dominik Dold, Julian Göltz, Akos F Kungl, Timo C Wunderlich, Andreas Hartel, Eric Müller, Oliver Breitwieser, Christian Mauch, Mitja Kleider, Andreas Grübl, David Stöckel, Christian Pehle, Arthur Heimbrecht, Philipp Spilger, Gerd Kiene, Vitali Karasenko, Walter Senn, Mihai A Petrovici*°, Johannes Schemmel, Karlheinz Meier. *Versatile emulation of spiking neural networks on an accelerated neuromorphic substrate*. arXiv:1912.12980, accepted for publication at the International Symposium of Circuits and Systems 2020, to appear in IEEE Xplore (2019) (P2241).
- Spike-based sampling on the BrainScaleS-1 system was published in:
  - Akos F Kungl, Sebastian Schmitt, Johann Klähn, Paul Müller, Andreas Baumbach, Dominik Dold, Alexander Kugele, Eric Müller, Christoph Koke, Mitja Kleider, Christian Mauch, Oliver Breitwieser, Luziwei Leng, Nico Gürtler, Maurice Güttler, Dan Husmann, Kai Husmann, Andreas Hartel, Vitali Karasenko, Andreas Grübl, Johannes Schemmel, Karlheinz Meier, Mihai A Petrovici. *Accelerated physical emulation of Bayesian inference in spiking neural networks.* Frontiers in neuroscience —- Neuromorphic Engineering Volume 13, Page 1201 (2019) (P1360)

4.6 Stochasticity from function in spiking sampling networks

Dominik DOLD, Ilja BYTSCHOK, Akos F. KUNGL, Andreas BAUMBACH, Oliver BREITWIESER, Walter SENN, Johannes SCHEMMEL, Karlheinz MEIER, Mihai A. PETROVICI

CDP5 collaboration between SP9 (UHEI, P47) and SP4 (UBERN, P71). UHEI and UBERN have jointly worked on the theory, network models and hardware implementation. UHEI was responsible for the commissioning of the system, providing the required hardware and software infrastructure.

The manner in which the brain could use a probabilistic computing scheme to process sensory information on the neuronal level is still an ongoing debate. Our work demonstrates how an ensemble of deterministic spiking networks, supplemented with weak and sparse inter-network connections, can be shaped by local plasticity to perform probabilistic computing in the form of sampling. This allows a self-sustained implementation of sampling-based computing in spiking networks, without the need for any explicit source of noise.

4.6.1 Research and Infrastructure

The sensory input available to humans (as well as machines) is not only noisy, but also ambiguous and often incomplete due to the nature of our environment. For instance, when playing Poker, only a limited amount of definite information is available (i.e., the cards on the table and in our hand) and remaining information is either not accessible or uncertain (i.e., the cards of other players). In such a situation, it is best to choose actions with the highest probability of a favourable outcome. How to estimate and choose the action with the highest probability, and how to reduce uncertainty when accumulating new evidence is mathematically known as Bayesian inference.

Recently, it has been proposed that Bayesian inference might be realized in the brain as sampling, i.e., spikes represent samples from a probability distribution (Buesing *et al*., 2011; Probst *et al*., 2015; Petrovici *et al*., 2016, see also Outputs 3, 4 and 5). One key question that remains is how such systems attain their stochasticity in the first place to perform sampling-based computations. In cortical models, this is often realized using externally generated, high-frequency Poisson noise. However, this is rather problematic: for neuromorphic implementations, it demands additional architectural constraints and bandwidth considerations as external noise has to be induced into the system; and in case of the brain, such a well-behaved and artificial source of stochasticity is not plausible. Even though cortical neurons behave stochastically *in vivo*, this mainly originates from a strong background bombardment of up to 10,000 presynaptic partners per neuron *in vivo* (Ariele *et al.*, 1996); many of which are not involved in the current functionality of the postsynaptic neuron's local cortical area. Thus, inspired by the modular structure of the brain, i.e., strongly connected functional clusters that are weakly interconnected, we propose a similar architecture for spiking sampling networks: a weakly interconnected ensemble of functional subnetworks, where the interconnections are used to provide background noise to every subnetwork (Fig. 7A).

We found that ensembles of functional networks can be set up such that the output statistics generated by the ensemble equals the input statistics of the assumed background noise (self-consistency) - enabling self-sustained sampling. Such a self-consistent state can be found automatically via synaptic plasticity, independent of the underlying functionality each network has to implement. This enables a self-sufficient and parsimonious implementation of spike-based sampling, by allowing all neurons to take on a functional role and not dedicating any resources purely to the production of background stochasticity - which we demonstrate both on the neuromorphic hardware system "BrainScaleS-1" (Fig. 7B-D) and for hierarchical networks in software simulations (Fig. 7E, F). These results reduce the architectural constraints imposed on physical neural substrates required to perform spike-based probabilistic computations both in biology and neuromorphic hardware.

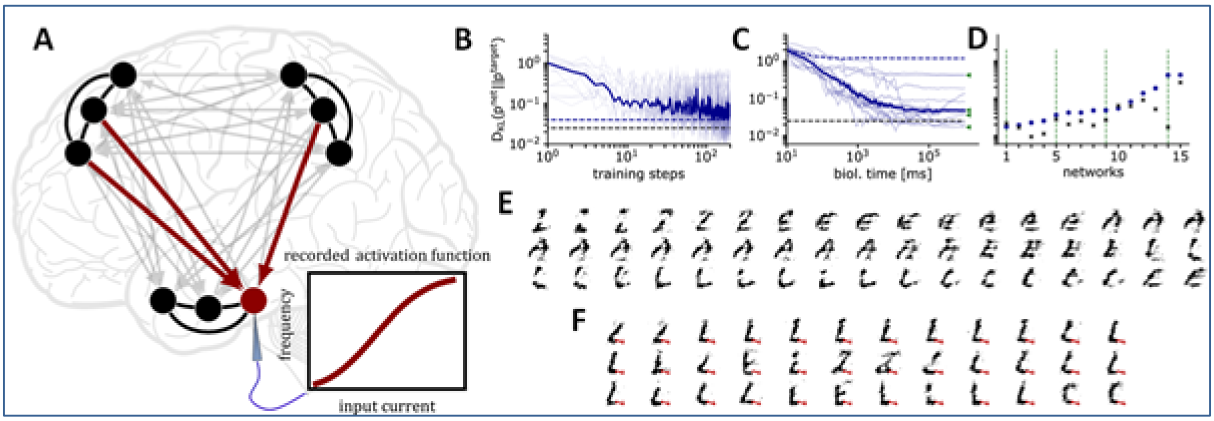

Fig. 7: Stochasticity from function in spiking sampling networks.

(A) Schematic visualisation of an ensemble of spiking sampling networks, inspired by the modular structure found in the mammalian cortex. In black, intra-network connections are shown, in grey inter-network connections. For instance, the red neuron receives spikes from neurons of the other two networks as background activity (red arrows), through which it attains the desired logistic activation function needed for spike-based sampling. (B-D) Sampling on BrainScaleS-1 with an ensemble of 15 deterministic functional networks with 4 neurons each. (B) Median performance of the ensemble (blue thick) and of individual networks (transparent blue) during training. The best performance reached by an equivalent implementation without interconnections using Poisson noise is shown in black. (C) Single run of the ensemble after training finished. By sampling longer, the networks approximate their target distributions better. The median performance pre-training is shown in (blue, dashed). (D) Final performance of the networks shown in (B), sorted from best to worst. Black crosses mark the performance of equivalent implementations with Poisson noise on the hardware. Green dashed lines correspond to the green triangles in (B). (E, F) An ensemble of four deterministic hierarchical networks provide each other with noise while performing Bayesian inference on the EMNIST dataset (software simulation using PyNEST). Each network was trained on a different subset of EMNIST ('a15YW', 'BRTXV', '13579', 'ACELZ'). (E) If no input is given to the networks, all of them generate patterns that match the data they were trained on (only one of four shown, we plot the spike activity of the visible layer averaged over a small time window here). (D) If only part of an image is shown to one of the networks (red), it is able to reconstruct the image and traverse through several plausible image interpretations (pattern completion / rivalry). The networks are also capable of classifying inputs (not shown here).

4.6.2 Impact and significance

The mammalian brain uses spikes to transmit information between neurons; however, there is still an ongoing debate on whether and what the functional purpose of spike-based coding might be. Here, we proposed that biological neurons use spikes for two specific functional purposes: (i) to draw samples from a posterior probability distribution constrained by sensory stimuli, offering an elegant algorithmic implementation of Bayesian inference and (ii) as sources of background irregularity needed to guarantee stochastic firing in the first place. This allows a resource-efficient and self-consistent realization of sampling-based Bayesian inference, both in the brain as well as neuromorphic hardware.

4.6.3 Components

C0001, C0209, C0349, C2226, C2420, C2439

4.6.4 Publications

Dominik DOLD, Ilja BYTSCHOK, Akos F. KUNGL, Andreas BAUMBACH, Oliver BREITWIESER, Johannes SCHEMMEL, Karlheinz MEIER, and Mihai A. PETROVICI. Stochasticity from function - why the Bayesian brain may need no noise. Neural Networks, 119, 200-213 (2019). https://arxiv.org/abs/1809.08045. doi: 10.1016/j.neunet.2019.08.002. (P1447)

4.7 Fast and deep neuromorphic learning with time-to-first-spike coding

Julian GÖLTZ, Andreas BAUMBACH, Sebastian BILLAUDELLE, Oliver BREITWIESER, Dominik DOLD, Laura KRIENER, Akos F. KUNGL, Walter SENN, Johannes SCHEMMEL, Karlheinz MEIER, Mihai A. PETROVICI

CDP5 collaboration between SP9 (UHEI, P47) and SP4 (UBERN, P71). UHEI and UBERN have jointly worked on the theory, network models and hardware implementation. UHEI was responsible for the commissioning of the system, providing the required hardware and software infrastructure.

The credit assignment problem represents a fundamental challenge for local synaptic plasticity mechanisms and is exacerbated by the use of non-smooth signals for communication, as is the case in spiking neural networks. We propose a mechanism for spike-based backpropagation of errors and demonstrate its functionality on an accelerated spiking neuromorphic substrate.

4.7.1 Research and Infrastructure

Training networks of spiking neurons to perform pattern recognition is a problem that has received increasing attention after recent successes of artificial neural networks, in the hope of producing similar functionality as machine learning and gaining insights about biological brains “on the go”. We have shown how the naive translation between artificial and spiking neural networks via rate coding (see CDP5 SGA1 M24 Deliverable) can be improved by using a more elaborate but sparse scheme, namely time-to-first-spike coding. We derived equations for the spike times of leaky integrate-and-fire (LIF) neurons in particular configurations (see also CDP5 SGA2 M12 Deliverable). Given a loss function, the equations allow us to train networks of spiking neurons with gradient descent and error backpropagation. In more recent work, we reformulated the equations to be more stable and thus perform better on analogue neuromorphic hardware.

The training of such a network on the BrainScaleS-2 system with a simple data set is shown in Figure 8. The coding favours early spikes, basically encoding stronger features by earlier spikes, and thus classification happens quickly (less than 5 µs wall-clock time). Learning of the MNIST data set in NEST-simulated networks is ongoing and translation to hardware is planned.

The algorithm was integrated into the Jenkins code development and monitoring framework. This allows automated, regular tests as well as manual testing of software changes to monitor the high-level state of the hardware and code base. Apart from setting a norm of standardisation and reproducibility for future experiments on BrainScaleS systems this has already helped debug complications in the past year.

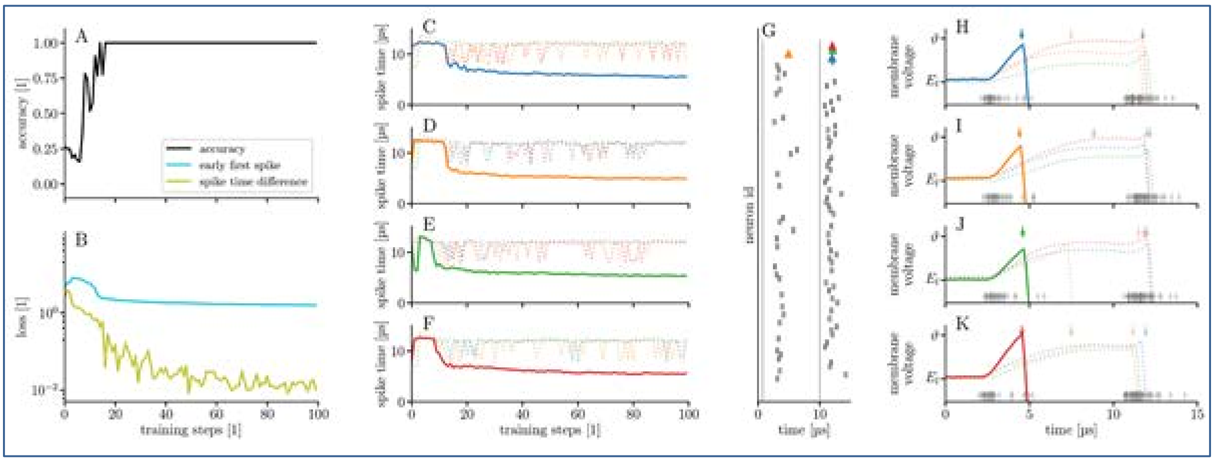

Fig. 8: Pattern recognition of a simple data set with time-to-first-spike coding on BrainScaleS-2

Evolution of key characteristics during training (A-F) and state after training (G-J). Training succeeds to 100% accuracy (A) and reduced loss (B) by effectively pushing correct neurons to spike earlier (C-F for the four classes of the data set, in each the correct neuron is highlighted over the other neurons). In the trained network (G), the classifying spike (▲) happens before most of the hidden neurons spike (█) and there is a large separation to the other label neurons' spikes. The voltage dynamics in the label layer (H-J) confirms the size of the separation.

4.7.2 Impact and significance

To harness the speed and parallelism of the neuromorphic system, one has to deal with the noisy and varying nature of analogue and mixed-signal neuromorphic systems, which represents a serious challenge. The successful training on both BrainScaleS generations and the use as a high-level monitoring tool on BSS-1, together with the analytical derivation of the spike times of LIF neurons demonstrate the feasibility of our approach and warrant further research into larger data sets, currently ongoing in software simulations.

4.7.3 Components

C0001, C0209, C0349, C0457, C2420

4.7.4 Publications

Julian Göltz. Training Deep Networks with Time-to-First-Spike Coding on the BrainScaleS Wafer-Scale System. Master's thesis. Universität Heidelberg (Apr. 2019). http://www.kip.uni-heidelberg.de/Veroeffentlichungen/details.php?id=3909. (P2020)

Sebastian Billaudelle, Yannik Stradmann, Korbinian Schreiber, Benjamin Cramer, Andreas Baumbach, Dominik Dold, and Eric Müller. Versatile emulation of spiking neural networks on an accelerated neuromorphic substrate. (2019) arXiv preprint arXiv:1912.12980. (P2241)

Julian Göltz, Andreas Baumbach, Sebastian Billaudelle, Oliver Breitwieser, Dominik Dold, Laura Kriener, and Mihai A. Petrovici. Fast and deep neuromorphic learning with time-to-first-spike coding. (2019) arXiv preprint https://arxiv.org/abs/1912.12980. (P2239)

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

5.1 Overview of Outputs

- Output 8: NEST implementation of three-factor error-correcting plasticity
- Output 9: Error-correcting learning with leaky integrate-and-fire point neurons

- Output 10: Bayes-optimal multisensory integration via conductance-based dendrites
- Output 11: Time-continuous deep reinforcement learning

5.2 NEST implementation of three-factor error-correcting plasticity

Jonas STAPMANNS, David DAHMEN, Markus DIESMANN, Laura KRIENER, Jakob JORDAN, Walter SENN, Mihai A. PETROVICI

CDP5 collaboration between SP7 (JUELICH, P20), SP4 (UBERN, P71) and SP9 (UHEI, P47). Contributions.

Due to the point-like nature of neuronal spiking, efficient neural network simulators often employ event-based simulation schemes for synapses. Yet many types of synaptic plasticity rely on the membrane potential of the postsynaptic cell as a third factor in addition to pre- and postsynaptic spike times. In order to update their strength, synapses therefore require continuous information which a priori necessitates a continuous update in a time-driven manner. Simulations of networks with this type of plasticity are therefore restricted to small network sizes. To overcome this problem, we implemented an efficient archiving of the history of postsynaptic membrane potentials in NEST to maintain an event-based update for synapses that minimizes the amount of stored data and communication between neurons.

5.2.1 Research and Infrastructure

Within component C2696, the storage and communication infrastructure for advanced plasticity mechanisms has been implemented in NEST by JUELICH. Besides the integration of the point-neuron rule by Clopath *et al*. (2010), a pull request has been made against NEST for the compartmentalized third-factor plasticity rule by Urbanczik and Senn (2014), which is the cornerstone of many learning paradigms studied in CDP5 (see also Outputs 2, 9, 10 and 11). The rule has been validated using simple examples, see e.g. (Fig. 9), provided by UBERN and UHEI.

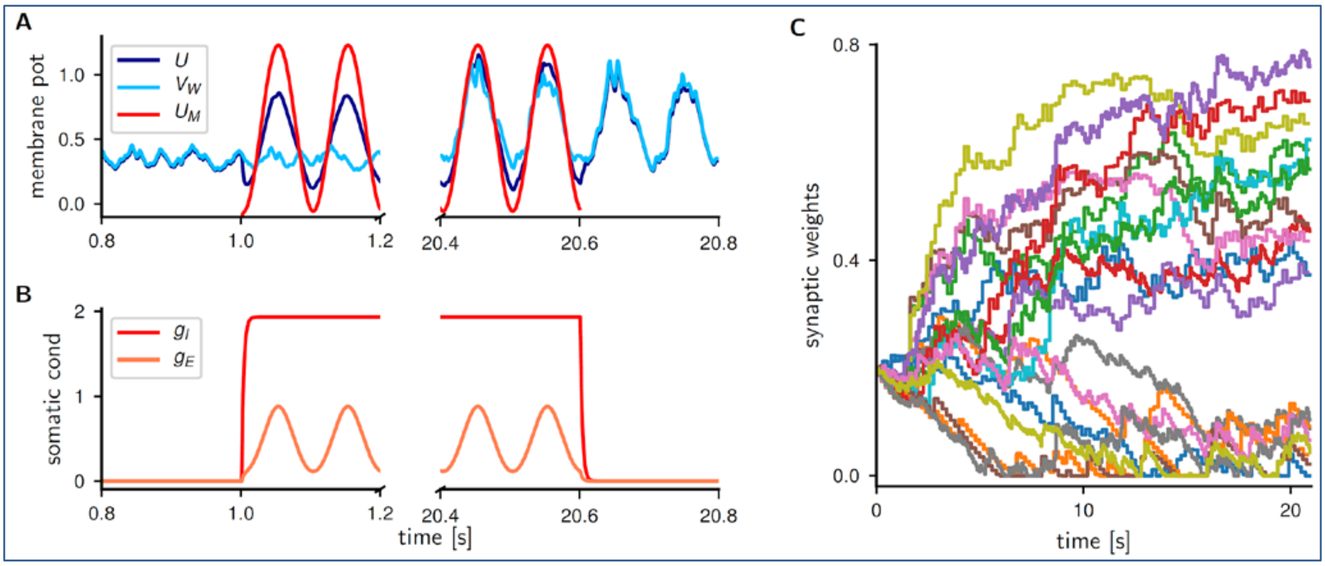

Fig. 9: Simple learning task using the Urbanczik-Senn plasticity rule for a 2-compartment neuron

The dendritic compartment receives a repeated spike pattern as an input via plastic synapses. During learning the synaptic weights are trained so that the membrane potential of the somatic compartment follows a teaching signal. A Membrane potential of the soma U (dark blue) and the dendrite V_W (light blue). The red curve denotes the nudging potential U_M resulting from somatic input (panel B). B Excitatory (g_E) and inhibitory (g_I) somatic conductances that produce the teaching signal. Panels A and B correspond to figure 1b in (R. Urbanczik, W. Senn (2014): Learning by the Dendritic Prediction of Somatic Spiking. Neuron, 81, 521 - 528). C Temporal evolution of the synaptic weights during learning.

5.2.2 Impact and significance

NEST is designed as a highly scalable neural network simulation engine that can be run on laptops as well as supercomputers. It provides well-documented, peer-reviewed and tested code to a large user community. Incorporating the infrastructure for third-factor plasticity rules makes this technology available to studies of functional networks. The possibility of simulating larger networks also allows investigation of learning paradigms in tasks of increasing complexity.

5.2.3 Components

C0209, C0349, C2419, C2420, C2547, C2696, C2722

5.2.4 Publications

Jordan, Jakob; Mørk, Håkon; Vennemo, Stine Brekke; Terhorst, Dennis; Peyser, Alexander; Ippen, Tammo; Deepu, Rajalekshmi; Eppler, Jochen Martin; van Meegen, Alexander; Kunkel, Susanne; Sinha, Ankur; Fardet, Tanguy; Diaz, Sandra; Morrison, Abigail; Schenck, Wolfram; Dahmen, David; Pronold, Jari; Stapmanns, Jonas; Trensch, Guido; Spreizer, Sebastian; Mitchell, Jessica; Graber, Steffen; Senk, Johanna; Linssen, Charl; Hahne, Jan; Serenko, Alexey; Naoumenko, Daniel; Thomson, Eric; Kitayama, Itaru; Berns, Sebastian; Plesser, Hans Ekkehard. NEST 2.18.0. https://doi.org/10.5281/zenodo.2605422 (P2534)

5.3 Error-correcting learning with leaky integrate-and-fire point neurons

Laura KRIENER, Matteo CARTIGLIA, Germain HAESSIG, Sebastian BILLAUDELLE, Benjamin CRAMER, Johannes SCHEMMEL, Giacomo INDIVERI, Walter SENN, Jakob JORDAN, Mihai A. PETROVICI

CDP5 collaboration between SP4 (UBERN, P71), SP9 (UHEI, P47) and external partners University of Zürich (UZ) and ETH Zürich (ETHZ). UBERN, UHEI, UZ and ETHZ have jointly worked on the theory and network models, as well as the implementation on different hardware platforms (BrainScaleS-2 and DynapSE). UBERN was responsible for the software simulations. UHEI, UZ and ETHZ were responsible for the commissioning of the system, providing the required hardware and software infrastructure.

In previous and ongoing work, we have shown how three-factor error-correcting learning rules enable powerful computation in hierarchical networks of structured neurons (see also CDP5 SGA1 M24 Deliverable and SGA2 M12 Deliverable). In order to make these learning algorithms amenable to contemporary neuromorphic platforms, we have developed a version of neuronal backpropagation for networks of single-compartment neurons.

5.3.1 Research and Infrastructure

Over the past year, we have made significant progress in implementing our general framework for learning from real-time data streams in low-power neuromorphic systems. The framework follows a first-principles approach to derive from a single objective function the neuron and weight dynamics in recurrent circuits of pyramidal neurons that extend across multiple cortical areas (see also Outputs 11 and 14). The network continuously learns to reduce its output error by forming local prediction errors from the combination of bottom-up and top-down neuronal activity represented in the lower and higher cortical area, respectively (Sacramento *et al*. 2018, Dold *et al*. 2019).

Originally described in Sacramento *et al*. 2018, the network consists of layers of dendritic microcircuits shown in Fig. 10 A and B. Through the choice of connectivity, neuron model and learning rule (see also Outputs 2, 8, 10, 11), this network natively implements a form of time continuous error-backpropagation, while relying only on locally available information for synaptic plasticity.

In Sacramento *et al*. 2018, dendritic microcircuits employ a multi-compartment neuron model. Most contemporary neuromorphic platforms do not support this model and, for those that do, the implementation details vary greatly. In order to stay platform-independent and not exclude platforms that do not feature multi-compartment neurons, we reformulated the dendritic microcircuit with point neurons. Figure 10 C shows our equivalent formulation of the original microcircuit. Figure 10 D illustrates the functionality of the point-neuron microcircuit using an instantaneous rate-based software simulation of the point neuron microcircuit. Early experiments of learning with point-neuron microcircuits on the neuromorphic hardware platforms DYNAPSE (ETHZ and UZ) and BrainScaleS-2 (UHEI) are shown in Fig. 11.

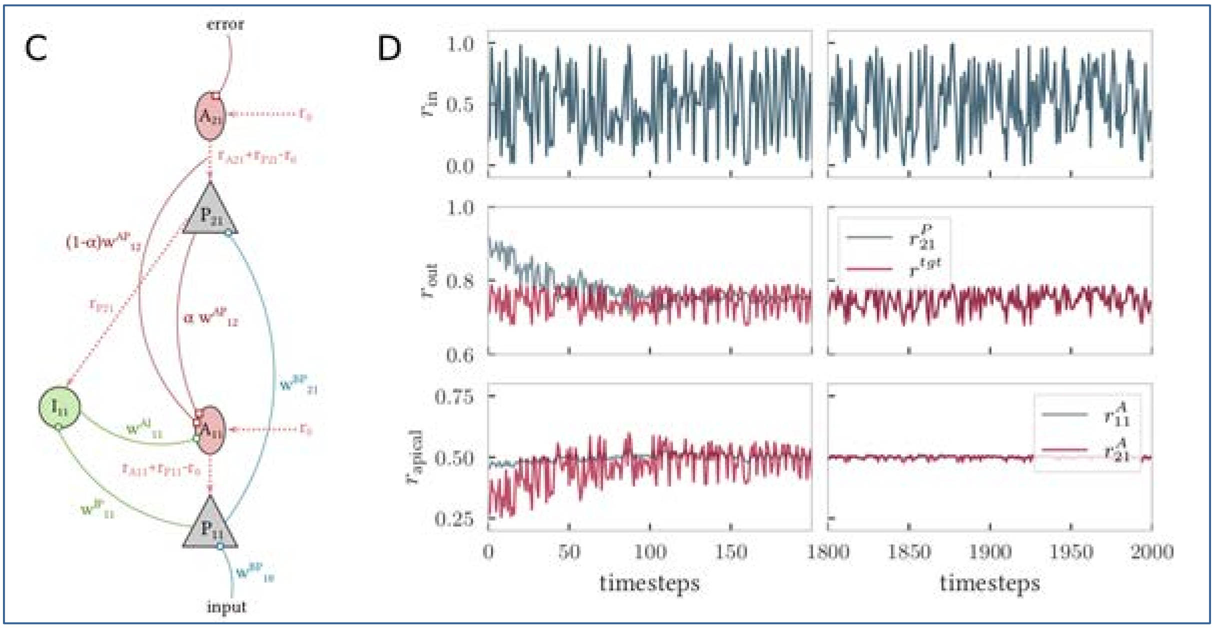

Fig. 10: Error backpropagation in cortical microcircuits

(A) Dendritic microcircuit using only point neurons. While the plasticity mechanism stays the same as for structured neurons, the multi-compartment pyramidal neurons are replaced by a combination of pyramidal point neurons and apical point neurons. The multi-compartment interneuron is replaced by a single point neuron. Dashed arrows represent the target firing rates for the neurons they are pointing towards. (D) Successful learning of an input/output-mapping with a point neuron microcircuit. The top panel shows the input, the middle panel shows the output in comparison to the target, the bottom panel shows the apical neurons which encode the error in the network (zero error corresponds to a rate of 0.5). Left panel: early phase of learning, right panel: late phase of learning.

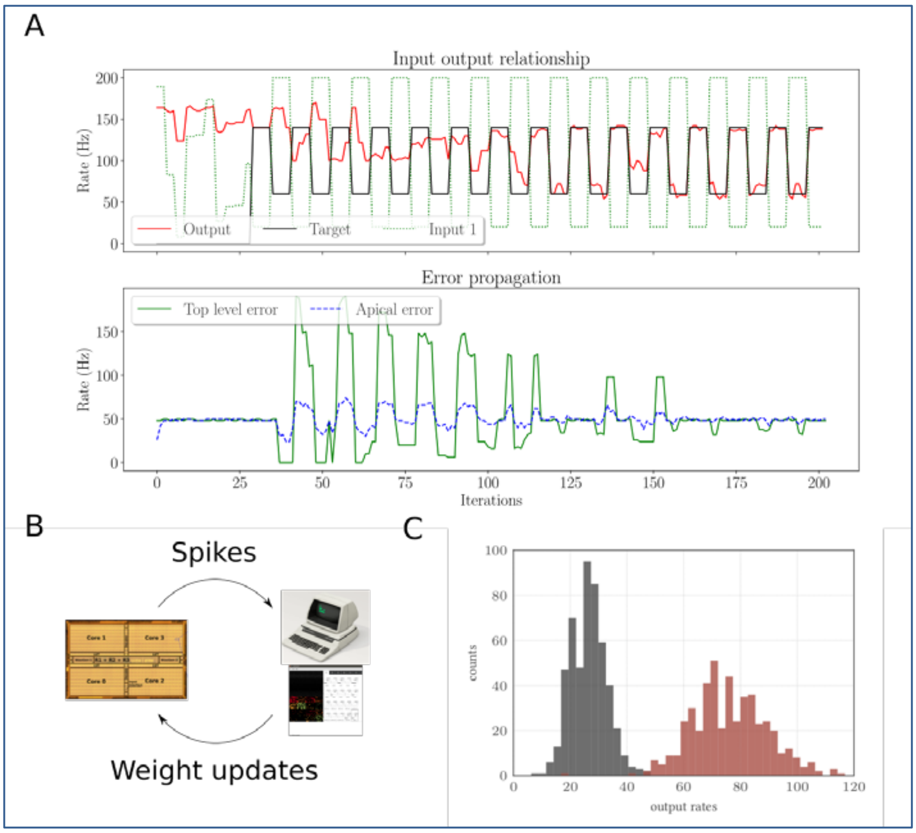

Fig. 11: Microcircuits on neuromorphic platforms.

(A) Successful learning of an input/output-mapping with a point neuron microcircuit on the DYNAPSE chip developed by ETHZ and UZ. The experimental setup was very similar to the simulation shown in Fig. 9 (D) and shows the portability of the model from software simulation to hardware emulation. (B) Setup used for learning in (A). Spikes of the microcircuit are recorded and sent from the chip to the host computer, where the weight updates are calculated. These are sent back to the chip and a new learning cycle begins. (C) Similar experiment on the BrainScaleS-2 chip. The microcircuit learns the correct weight configuration to produce to distinct target output rates for two given inputs.

5.3.2 Impact and significance

Our approach can be interpreted as a biophysically plausible implementation of real-time error-backpropagation amenable to the implementation in low-power neuromorphic hardware. Importantly, it does not require global scheduling for different phases of inference and learning. If successful, the framework can be extended to study time-continuous learning from ongoing stimulus streams in a mixed supervised, unsupervised and reinforcement-learning scenario, using spiking neural networks in neuromorphic systems.

5.3.3 Components

C0457, C2419, C2420, C2547, C2722

5.3.4 Publications

Dominik Dold, Akos F. Kungl, João Sacramento, Mihai A. Petrovici, Kaspar Schindler, Jonathan Binas, Yoshua Bengio, and Walter Senn. Lagrangian dynamics of dendritic microcircuits enables real-time backpropagation of errors. Cosyne abstracts, Lisbon (http://www.kip.uni-heidelberg.de/Veroeffentlichungen/details.php?id=3855), (2019) (P1905) (recorded talk: E2424)

5.4 Bayes-optimal multisensory integration via conductance-based dendrites

Jakob JORDAN, João SACRAMENTO, Willem WYBO, Mihai A. PETROVICI, Walter SENN

CDP5 collaboration between SP4 (UBERN, P71), SP9 (UHEI, P47) and external partners University of Zürich (UZ) and ETH Zürich (ETHZ). All partners contributed to the theoretical approach and the design of the model. The execution of simulations and the evaluation of data was done at UBERN.

Cue integration, the combination of different sources of information to reduce uncertainty, is a fundamental computational principle of brain function. Starting from a normative model, we show that the dynamics of multi-compartment neurons with conductance-based dendrites naturally implement the required probabilistic computations. The associated error-driven plasticity rule allows neurons to learn the relative reliability of different pathways from data samples, approximating Bayes-optimal observers in multisensory integration tasks.

5.4.1 Research and Infrastructure

Starting from a normative model of cue integration in single neurons, we derived membrane potential and weight dynamics in a principled fashion. The resulting neuron dynamics resembled leaky integrators with multiple dendritic compartments and the associated plasticity rule reduces output errors, while extracting the relative reliabilities of different sensory pathways from data samples (see also Outputs 2, 8, 9 and 11). We previously demonstrated that a model derived in this framework approximates Bayes-optimal observers in multisensory integration tasks assuming strong coupling of dendritic compartments to the soma. We now extended our mathematical analysis and numerical experiments to include realistic ratios of leak and transfer conductances extracted from biophysically detailed models. Since the relative weights of different pathways are learned from data, our model can account for the deviations introduced by finite transfer conductances and maintain a Bayes-optimal estimate, demonstrating the generality of our approach.

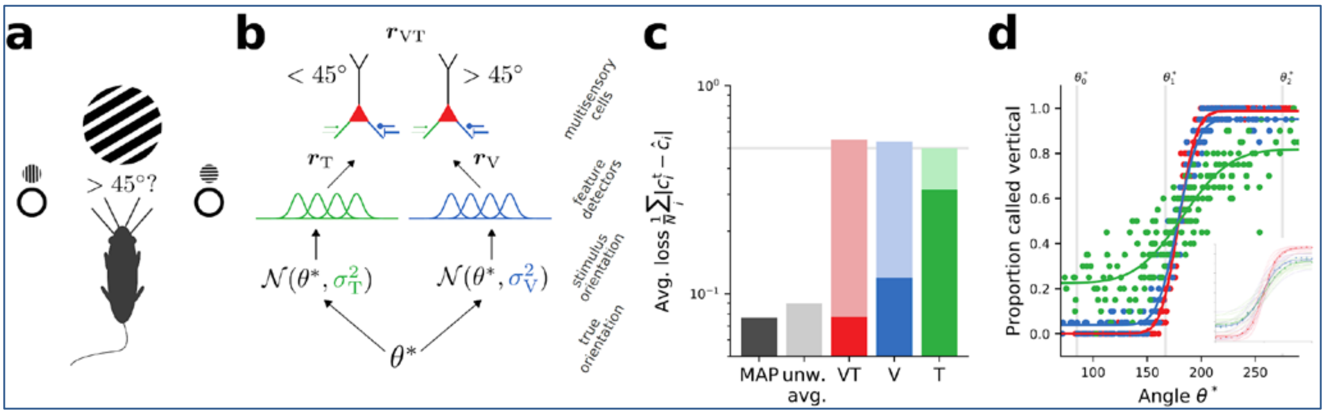

Fig. 12: Learning approx. Bayes-optimal inference of orientations from multi-modal stimuli

(a) Experimental setup; using visual and/or tactile information a rat classifies the orientation of a grating as either vertical or horizontal. (b) Network model. From a ground truth orientation θ visual and tactile stimuli are sampled with fixed, modality-specific noise amplitudes. These orientations are translated into firing rates using two populations of von-Mises feature detectors. All feature detectors project onto two multisensory neurons. (c) Loss (error) of the MAP estimate (dark grey), an unweighted estimate combining visual and tactile orientations equally (light grey), the trained model with bimodal cues (red), visual cues alone (blue) and tactile cues alone (green). Light coloured bars indicate loss before training, light grey line denotes chance level. (d) Psychometric curve, i.e., relative frequency of stimuli classified as vertical as a function of the true orientation of the grating, for bimodal (red) and unimodal conditions (blue and green). Dashed red line indicates psychometric curve estimated from unimodal measurements. Inset shows data from [Nikbakht *et al*., 2018].

5.4.2 Impact and significance

The classic view on neural computation abstracts integration of information in single neuron by a simple sum of their inputs. We suggest a fundamental change in perspective, considering membrane conductances as the central variables of neural computation. This new view exposes that single neurons are naturally equipped to perform a ubiquitous form of probabilistic computation: instead of simply summing their inputs, they compute a convex combination of stimulus-specific reversal potentials. Our framework provides parsimonious explanations for various experimental findings in cue integration tasks, connecting theories of probabilistic cue integration with neuron-level data. Furthermore, our approach has the potential to increase the robustness of artificial neural networks that need to merge different information streams.

5.4.3 Components

C2419, C2420, C2547, C2722

5.4.4 Publications

Jakob Jordan, Joao Sacramento, Mihai Petrovici, and Walter Senn. Conductance-based dendrites perform reliability-weighted opinion pooling. Proceedings of the NICE conference 2020, Heidelberg, Germany (2020, postponed). (P2536)

5.5Time-continuous deep reinforcement learning

Akos F. KUNGL, Dominik DOLD, Walter SENN, Oskar RIEDLER, Mihai A. PETROVICI

CDP5 collaboration between SP4 (UBERN, P71) and SP9 (UHEI, P47). Both partners contributed to the theoretical approach and the design of the model, the execution of simulations and the evaluation of data.

The hierarchical structure of the cortex raises the question of how plasticity in the brain is able to shape this complex network in order to solve the spatial and temporal credit assignment problems. Abstract neural networks - distant relatives of their biological archetypes - solve this problem with the backpropagation-of-errors algorithm. Despite its effectiveness in abstract neural networks, it remains an open question whether and how backprop might be implemented in biological neural networks. In M12 of SGA2, we reported the development of a theoretical framework that uses a least-action principle to derive a biologically plausible implementation of backpropagation. Here, we present an extension of the work to include not only supervised but also reinforcement learning.

5.5.1 Research and Infrastructure

In our model, the neurosynaptic dynamics are derived from an energy function using the variational principle. The resulting neuron and synapse models are closely related to those used in several other models (see Outputs 2, 8, 9 and 10). Errors are introduced to the network by nudging, and they are propagated to deeper layers via cortical microcircuits (Fig. 13 A; see also Outputs 9 and 14 for related architectures). In the resulting dynamics the phase-advanced firing of the neurons effectively undoes the network delay introduced by finite membrane time-constants (Fig. 13 B). Training in a supervised learning setup has been previously demonstrated on the MNIST dataset (Fig. 13 C).

For reinforcement learning, instead of having an explicit teacher, the output neurons, which represent the actions, now form a soft winner-take-all network (Fig. 13 A). We propose that this winner-take-all structure evokes a nudging on the soma of the output neurons (Fig. 13 D), which is subsequently backpropagated through the network. We formally derive a reward prediction error $\delta = R - \langle R \rangle$ that modulates the plasticity multiplicatively as a global reward-specific neuromodulator. By construction, the learning rule approximates the policy gradient of the mean expected reward. Learning is not only possible, but the model also performs well under delayed reward (Fig. 13 E-F) and against noisy synaptic weights in the soft winner-take-all circuit.

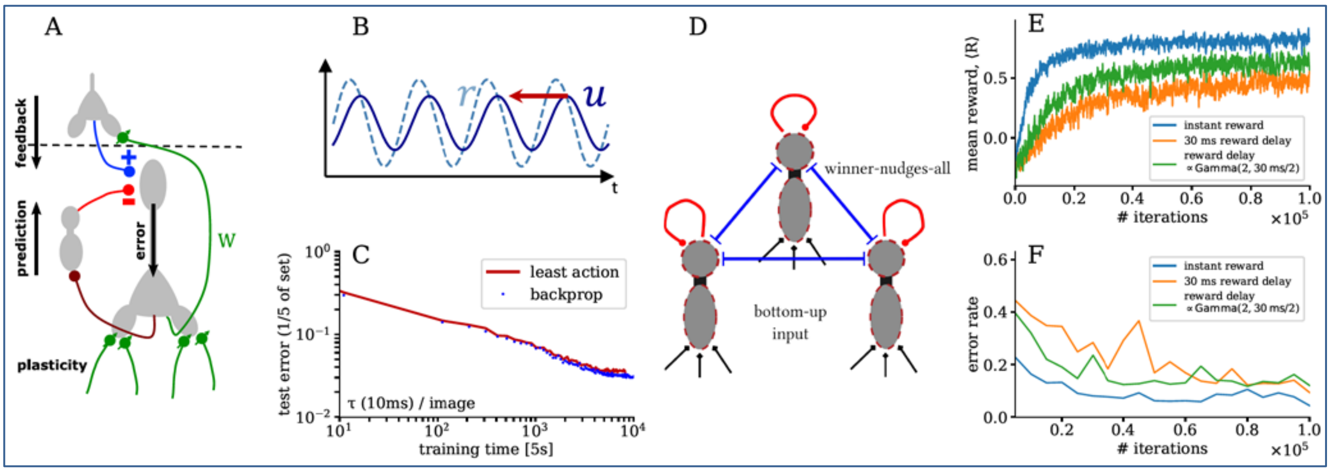

Fig. 13: Time-continuous deep reinforcement learning

(A) Cortical microcircuit used for error calculation and representation. Backpropagated errors are locally calculated by subtracting bottom-up prediction from top-down feedback. (B) The neuronal activity r is a non-linearly advanced version of the respective membrane potential that effectively undoes low-pass filtering caused by leaky integration. (C) A network of 786 - 800 - 10 neurons learns MNIST from a continuous movie, where each digit is only shown briefly (time scale of the membrane time constant). (D) Lateral somato-somatic interaction with self-excitation and mutual inhibition gives rise to an error nudging approximating policy gradient. (E-F) The network successfully learns on a small time-continuous classification problem based on 3 MNIST digits. Learning is robust although slower in the presence of delayed rewards, even if the reward delay is stochastic.

5.5.2 *Impact and significance*

Our work fills a gap between models of biological deep learning, which usually only concern supervised learning, and models of reinforcement learning, which are often restricted to shallow learning architectures. We connected soft winner-take-all architectures to models of backpropagation, by which we propose a novel role for winner-take-all networks: error-vector generation in the input layer. Finally, our studies regarding the robustness of the model are a prerequisite for future neuromorphic applications.

5.5.3 *Components*

C2419, C2420, C2547, C2722

5.5.4 *Publications*

Akos Ferenc Kungl, Dominik Dold, Oskar Riedler, Walter Senn, and Mihai A. Petrovici. Deep reinforcement learning in a time-continuous model. Poster on the Bernstein Conference (2019 Sept 18-20) - http://dx.doi.org/10.12751/nncn.bc2019.0168. (P1999).

Dominik Dold, João Sacramento, Akos F. Kungl, Walter Senn, and Mihai A. Petrovici. An energy-based model of folded autoencoders for unsupervised learning in cortical hierarchies. Poster at the Bernstein Conference (2019 Sept 18-20) (P2358).

6 Key Result KRc5.4: Deep episodic learning

6.1Overview of Outputs

- Output 12: Sleep-memory interaction in a thalamo-cortical spiking model
- Output 13: Visual memory consolidation during sleep
- Output 14: Deep predictive coding model of the visual cortical hierarchy

6.2 Sleep-memory interaction in a thalamo-cortical spiking model

Cristiano CAPONE, Elena PASTORELLI, Bruno GOLOSIO, Hans Ekkehard PLESSER, Johannes SCHEMMEL, Pier S. PAOLUCCI

CDP5 collaboration among SP3 (INFN, P92), SP7 (NMBU, P44), SP9 (UHEI, P47). INFN developed the model, defining with NMBU the necessary NEST features and with UHEI its porting to the BrainScalES architecture, planned for SGA3. The model has been released to the community through the EBRAINS Knowledge Graph.

While disconnected from external input and from the duties associated with wakefulness, animal brains are free to optimize internal representations, create novel association and plans and recover optimal working points for reduced energetic costs of activity. We propose a functional role of slow waves during sleep for the optimization and normalization of memory representations and the creation of novel associations, that improve post-sleep classification performance.

6.2.1 Research and Infrastructure

Human brains spend about one-third of their lifetime sleeping. Sleep is present in every animal species that has been studied. This happens despite two apparent drawbacks: the danger caused by sleep, that diminishes the capability to defend from predators and other threats, and the reduction of time available for activities targeting immediate rewards (e.g. hunting or gathering food). Sleep must therefore serve essential functions (see also Outputs 3 and 13).

In June 2019, Capone *et al*. demonstrated that sleep-like slow oscillations improve visual classification in a thalamo-cortical spiking model through synaptic homeostasis and memory association. The corresponding NEST model has been released (component C2193) through the EBRAINS Knowledge Graph. A network of spiking AdEx neurons is trained (STDP plasticity) to encode, retrieve and classify images of handwritten digits. Then, sleep-like oscillations are induced. A differential homeostatic process is observed. Slow oscillations induce both an unsupervised enhancement of connections among groups of neurons associated to instances of the same class (digit) and a simultaneous down-regulation of stronger synapses created by the training. This is reflected in a hierarchical organisation of post-sleep internal representations. This promotes higher performances in post-sleep retrieval and classification tasks and creates hierarchies of categories in integrated representations.

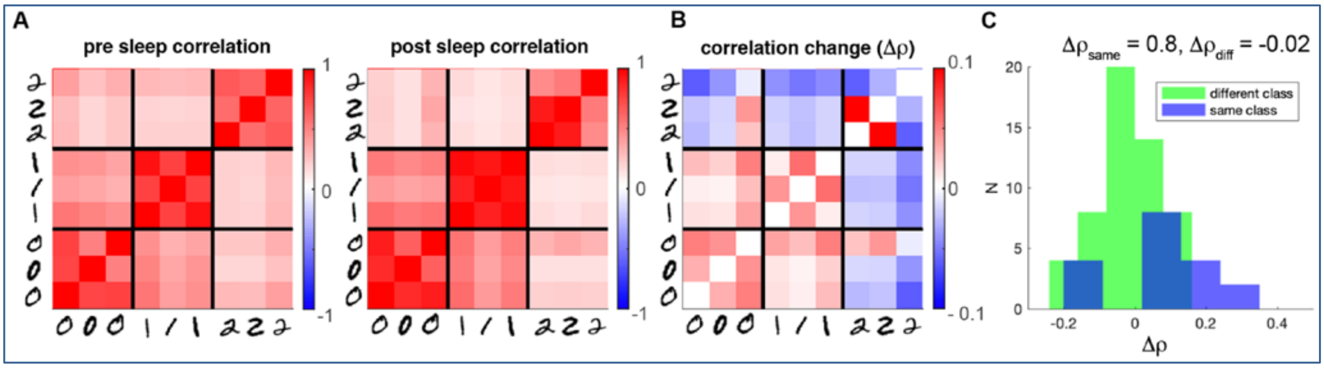

Fig. 14: Effects of slow oscillations on internal representation.

(A) Activity correlation between all pairs of populations representing the single images before (left) and after (right) sleep. (B) Correlation difference between after and before sleep. (C) Histogram of correlation differences for populations encoding the same class (blue) and different classes (green).

6.2.2 Impact and significance

The “Sleep-memory interaction in a thalamo-cortical spiking model” has been released through a dedicated GitHub (see link) and through the EBRAINS KG. It represents a contribution of WP3.5 to

CDP5 (Plasticity). The combination of this simplified model with the final SGA2 results of WP3.2 (Large-scale spiking simulations of slow-wave and asynchronous activity) prepares the ground for the planned SGA3 study of the interaction between brain states and memories in large-scale networks. Furthermore, it is conceivable that recurrent multi-area multi-level networks employed in future bio-inspired artificial-intelligence systems will need to enter specific brain-states to normalize, optimize and associate their internal representations.

6.2.3 Components

C0001, C0209, C1788, C2193

6.2.4 Publications

https://github.com/PierStanislaoPaolucci/2019thalCort-SNN-SO-AW-mem

https://kg.ebrains.eu/search/?facet_type[0]=Model&q=paolucci#Model/97670076281ccbdc38ea2c2d76a64e64

Cristiano Capone, Elena Pastorelli, Bruno Golosio, and Pier Stanislao Paolucci. Sleep-like slow oscillations improve visual classification through synaptic homeostasis and memory association in a thalamo-cortical model. Scientific Reports, Vol. 9, No. 1 (2019). https://www.nature.com/articles/s41598-019-45525-0. DOI: 10.1038/s41598-019-45525-0 (P2024)

6.3 Visual memory consolidation during sleep

Nicolas DEPERROIS, Jakob JORDAN, Mihai A. PETROVICI, Walter SENN

CDP5 collaboration between SP4 (UBERN, P71) and SP9 (UHEI, P47). Both partners contributed to the design of the model and simulation scenarios. The simulations were performed at UBERN.

The ability of the visual cortex to build semantic representations along the cortical hierarchy remains ill-understood. This process, known as visual memory consolidation, is not instantaneously achieved during wakefulness and persists as the brain is disconnected from the external world, such as during sleep. Moreover, the presence of creative dreams during rapid-eye-movement (REM) sleep suggests that the brain extracts the content of visual experiences by organizing them into semantic representations. Here, inspired by modern artificial intelligence, we argue that dreaming, especially during REM sleep, is a key process of visual memory consolidation allowing the emergence of semantic concepts.

6.3.1 Research and Infrastructure

From an early age, animals are able to recognize many objects in different situations without being explicitly taught how to achieve it. Thus, a natural role of the visual system is to build object-specific representations independently of different factors of variations such as position, distance, and illumination. Semantic memories are believed to arise when episodic memories stored in the hippocampus are reactivated during Slow-Wave Sleep (for further functional aspects of slow-wave sleep, see Outputs 3 and 12). What if, however, memory consolidation was additionally facilitated via creative dreaming mainly occurring during REM sleep?

A recent machine learning algorithm, called Generative Adversarial Networks (GANs) (Goodfellow *et al.*, 2014), known for its capability to generate high-quality image samples, could potentially explain how feedback connections learn to generate realistic visual representations during REM sleep. This generated imagery can in turn facilitate the formation of semantic representations. Here, we explore this idea by proposing a possible implementation of GANs in cortical microcircuits and considering the discriminator and the generator as part of the feedforward pathway and the feedback pathway of the visual cortex, respectively (Fig. 15). Notably, we show that REM sleep is an ideal stage where both networks can be trained in an adversarial way to improve the generation of visual representations.

We train the model with the benchmark dataset CIFAR-10 (50,000 colour images with 10 categories) using the DCGAN architecture introduced by Radford *et al*. (2015). At each epoch, we evaluated the quality of the latent representations by training a linear classifier on top of the 100-dimensional latent vectors obtained by passing the CIFAR-10 images through the discriminator network. As a baseline, we consider the classification performance achieved with a randomly initialized discriminator. Pure autoencoder training ("wake phase" only) does not improve the quality of the high-level representations from the baseline (Fig. 16C). Hippocampal replay of previously encoded latent states during sleep ("NREM phase") further improves the performance. Finally, a REM phase, in which random latent vectors generate previously unseen image samples, significantly improves the network's latent representations, leading to better classification performance. Together, these preliminary results show that hippocampal-replay during SWS and creative dreaming during REM sleep can be essential to build semantic representations in the visual cortex.

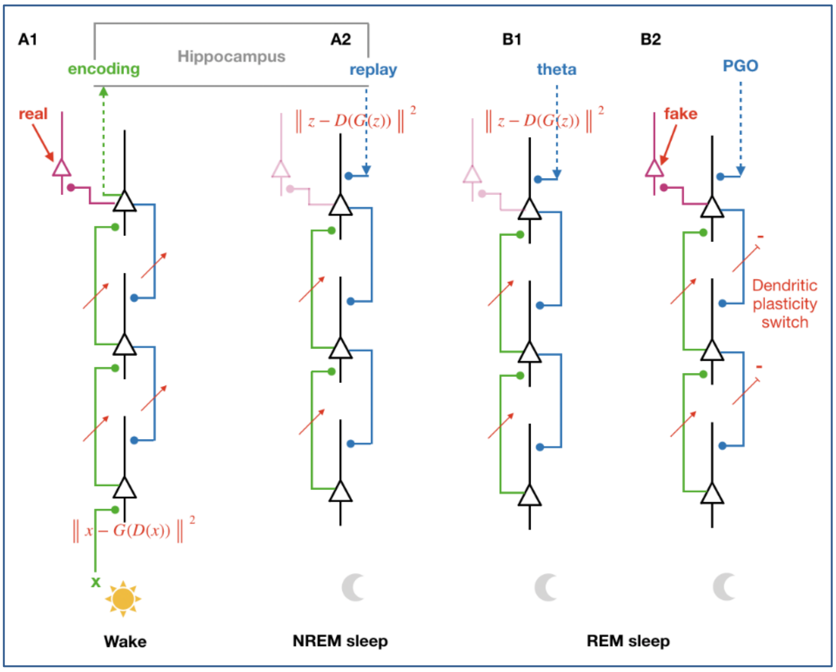

Fig. 15: Computational scheme of visual cortex for learning representations during wakefulness & sleep phases

x represents an image, z is a vector representing the latent layer (IT cortex). Each pyramidal neuron (triangle) represents a layer of the visual cortex. Discriminator (green) and generator (blue) connections are part of the DCGAN architecture and project to the basal and apical tree of the pyramidal neurons. Synaptic plasticity is symbolized by red arrows indicating that learning occurs on a given connection. (A) During wakefulness (A1), x is processed via the D connections and reconstructed via the G connections. D connections learn to classify sensory activity as "real" while G connections learn to minimize the reconstruction error |x-G(D(x))|. During NREM sleep (A2), activity z produced during wakefulness is replayed from the hippocampal memory and cortical synapses adapt to minimizes the reconstruction error |z - D(G(z))|. (B) During the postulated REM sleep, the network is trained to reconstruct random cortical activity z which generates new visual imagery (B1), and during PGO waves, D connections learn to classify the generated dreams as "fake" while G connections do the opposite. The reversal of plasticity mimics the adversarial action of GANs.

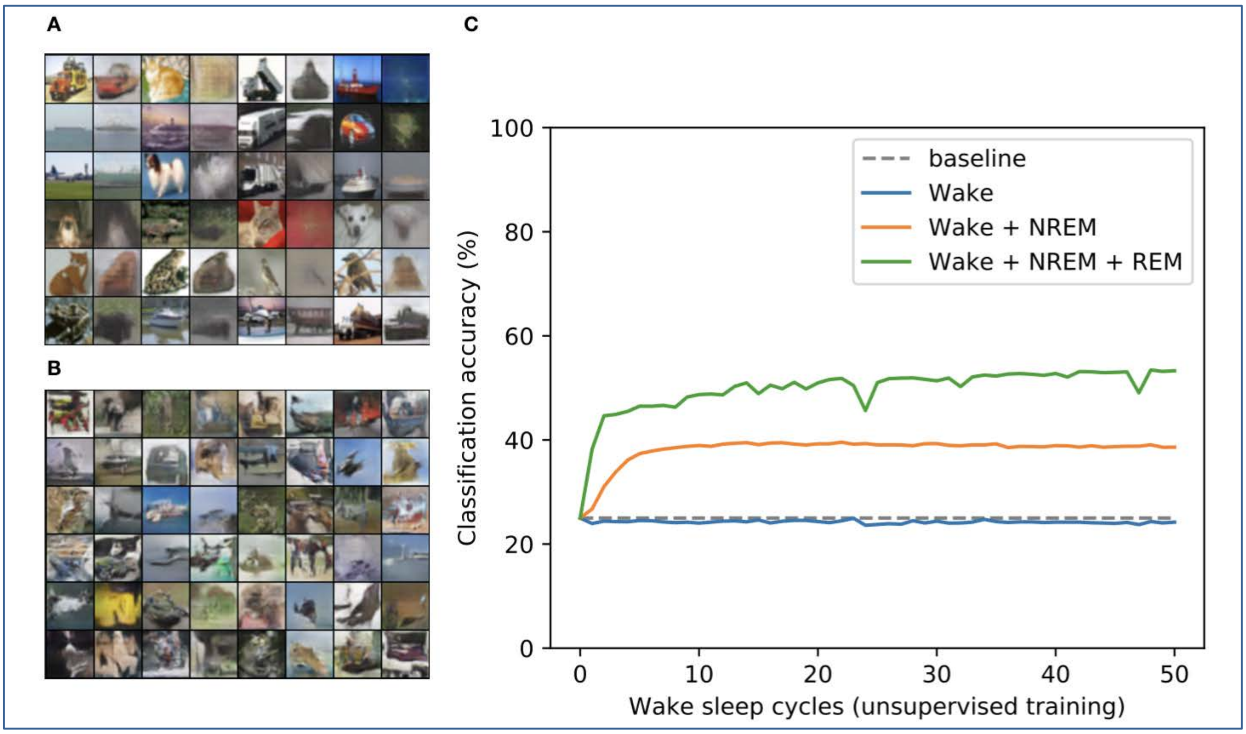

Fig. 16: Effect of sleep on high-level representations in IT

(A) Original image and reconstructions obtained via FF and FB connections. Reconstructed images match the original ones in terms of shape and semantic content, showing that cortex organizes its high-level representations so their feature represents objects. (B) Dreamed images G(z) used to consolidate IT representations during REM sleep. (C) Classification accuracy obtained by training a linear classifier on IT representations at each epoch of unsupervised learning for different sleeping conditions.

6.3.2 Impact and significance

We hypothesize that, besides consolidation of experiences, sleep provides an ideal state for our brains to improve their ability to extract semantic structure from sensory data. In contrast to previous approaches we provide a concrete model for visual memory consolidation during sleep that can potentially be mapped to cortical microcircuits. This provides a unique opportunity to connect high-level theories of cortical reorganization during sleep with circuit level dynamics.

6.3.3 Components

C1788, C2061, C2193, C2226, C2419, C2420, C2547, C2722

6.3.4 Publications

Early research stage - no publications yet.

6.4 Deep predictive coding model of the visual cortical hierarchy

Shirin DORA, Kwangjun LEE, Sander M. BOHTE, Cyriel M.A. PENNARTZ, Oliver STRUCKMEIER, Martin J. PEARSON, Walter SENN, and Mihai A. PETROVICI

CDP5 collaboration between SP3 (UvA, P98; UWE, P101; AALTO, P2), SP4 (UBERN, P71), and SP9 (UHEI, P47). UvA worked on developing a deep predictive neural network model for the visual cortical hierarchy. UBERN and UHEI helped in formulating the model and discussion of its properties. UWE and AALTO implemented the model in a biomimetic robot performing place recognition tasks.

Predictive coding provides a computational paradigm for modelling perceptual processing as the construction of representations accounting for causes of sensory inputs. Using biologically plausible principles, we developed a deep predictive-coding network that mimics properties of the visual cortical hierarchy. Furthermore, a multi-sensory extension of the model is developed and implemented in a biomimetic navigating robot.

6.4.1 Research and Infrastructure

Understanding brain mechanisms of perception requires a computational approach based on neurobiological principles. Many deep-learning architectures are trained by supervised learning from large sets of labelled data (see also Outputs 9 and 11), but biological brains must also be able to learn from unlabelled sensory inputs. We developed a Predictive Coding methodology for building scalable networks that mimic deep sensory cortical hierarchies, perform inference on the causes of sensory inputs and are trained by unsupervised, Hebbian learning (Figure 17). The network models are well-behaved, in that they faithfully reproduce visual images based on high-level, latent representations. When ascending the sensory hierarchy, we find that many response properties associated with neurons along the sensory cortical hierarchy are also exhibited by these networks, like increasing image selectivity, improved object-level generalizability capabilities in higher areas.

Furthermore, we have extended this model to simultaneously handle data in multiple modalities. To study the inferential capacity of these multisensory networks, we used this network for a problem in robotics. More specifically, the representations inferred from a deep neural network trained using predictive coding on visual and tactile sensory inputs were used for the problem of place recognition in navigating robots (Figure 18). The results obtained from the model clearly show that the representations obtained from the model outperform the traditional models used for this problem. This work was done using simulation of the WhiskEye robot in Gazebo and work on application of the model in the actual platform (WhiskEye) is being carried out in collaboration with Martin Pearson at University of West England.

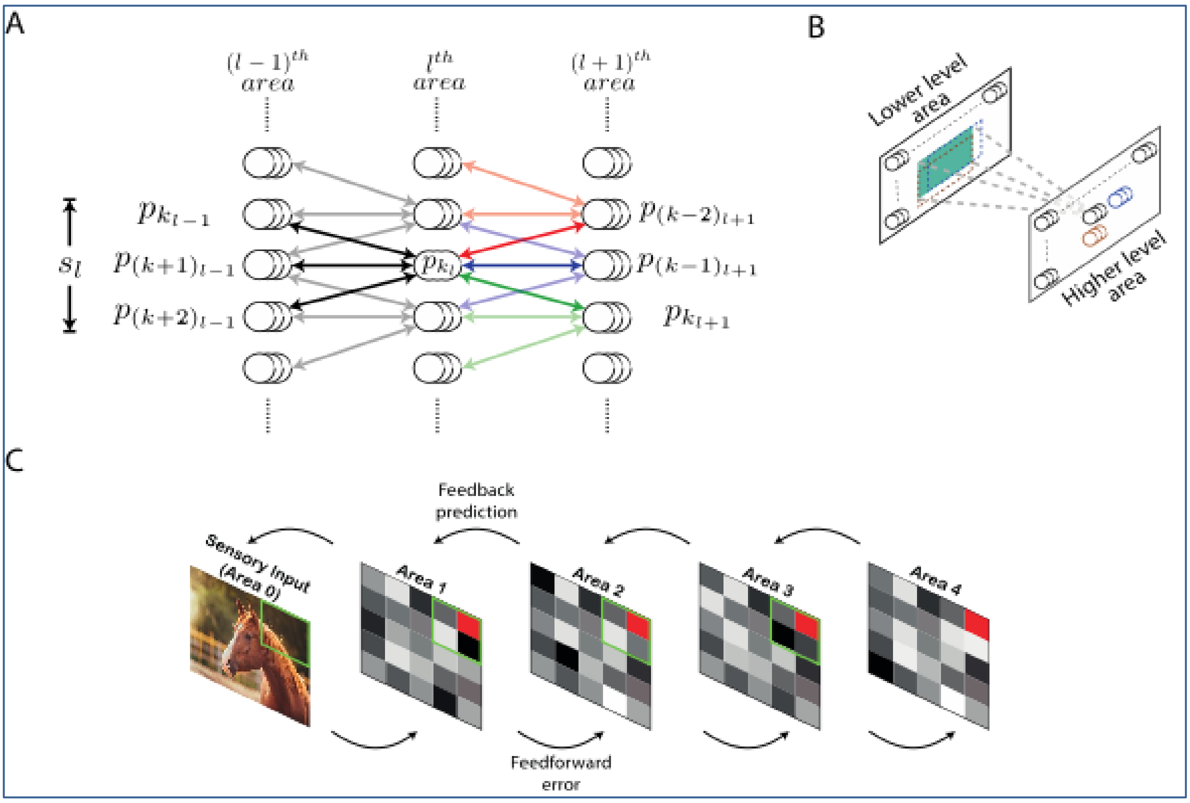

Fig. 17: Architecture of the deep predictive coding network with receptive fields

(A) A population of neurons having identical receptive fields is represented by three overlapping circles. p_{k_l} denotes the k_{th} population in the l_{th} area and s_l is the size of the receptive field of all populations in the l^{th} area. Both s_l and s_{l+1} have been set to 3 here. For this value of s_l, the populations $p_{k_{l-1}}$ through $p_{(k+2)_{l-1}}$ constitute the receptive field of the population p_{k_l} (their connections are represented by black lines). Similarly, for this value of

s_{l+1}, p_{k_l} will be present in the projective fields of populations $p_{(k-2)_{l+1}}$ through $p_{k_{l+1}}$. (B) For processing images, neuronal populations in each area can be visualized in a two-dimensional grid. Each population exhibits a two-dimensional receptive field (the receptive field of an example population in a higher-level area is shown in green). As a result, the receptive fields of two different populations can exhibit different overlaps horizontally and vertically. The receptive fields of two horizontally adjacent populations (black and blue) overlap completely in the vertical direction and partially in the horizontal direction. Similarly, the receptive fields of two vertically adjacent populations (black and brown) overlap completely in the horizontal direction and partially in the vertical direction. (C) An overview of the network with $n_l = 1$ for all areas. Sensory input is presented to the network through Area 0. Activity of neurons in areas 1-4 is represented by tiles in greyscale colours. The green square in a lower area denotes the receptive field of the population represented as a red tile in the higher area.

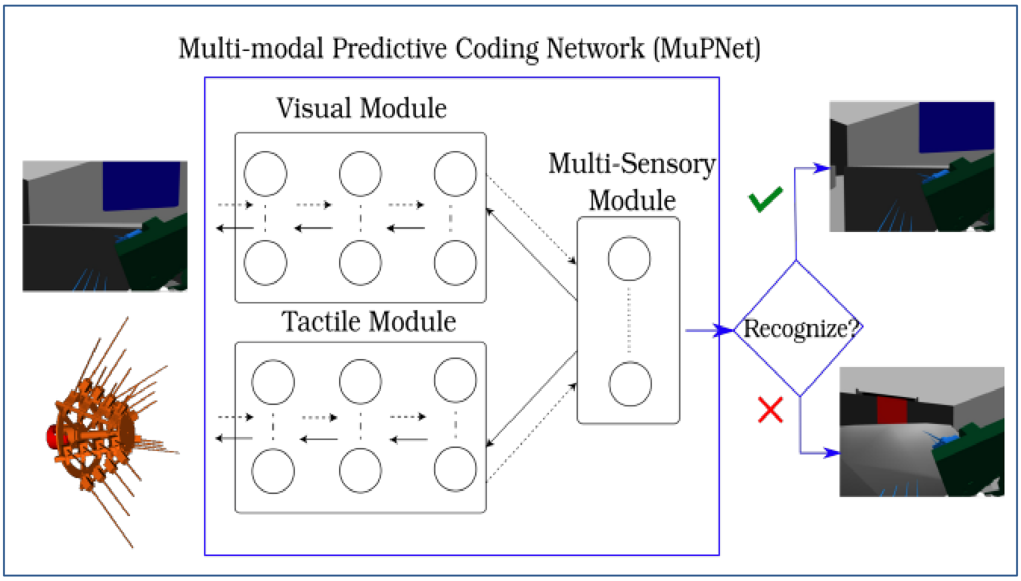

Fig. 18: High level overview of MuPNet architecture.

Visuo-tactile sensory data for a given scene is presented to the trained network and a latent representation of the multi-sensory stimuli is inferred. Note that the errors (dashed lines) are propagated forward and the predictions (solid lines) backwards. The resulting representation is used for place recognition.

6.4.2 Impact and significance

Our model shows how a complex neuronal phenomenology emerges from biologically plausible, deep networks for unsupervised perceptual representation. Based on the model, we developed MuPNet that extracts visuo-tactile latent representations for place recognition in a biomimetic rodent robot. The proposed predictive coding based approach for multimodal feature extraction is not limited to visuo-tactile processing, but also applicable to learning the joint latent representations from any co-incident multi-sensory inputs.

6.4.3 Components

C2060, C2061, C2226, C2228, C2321

6.4.4 Publications

Shirin Dora, Sander M. Bohte, Cyriel M.A. Pennartz. Deep predictive coding accounts for emergence of complex neural response properties along the visual cortical hierarchy. (2020). http://europepmc.org/article/PPR/PPR112163. bioRxiv doi: 10.1101/2020.02.07.937292. (P2345)

Oliver Struckmeier, Kshitij Tiwari, Shirin Dora, Martin J. Pearson, Sander M. Bohte, Cyriel M.A. Pennartz, and Ville Kyrki. MuPNet: Multi-modal Predictive Coding Network for Place Recognition by Unsupervised Learning of Joint Visuo-Tactile Latent Representations. (2019). https://arxiv.org/abs/1909.07201. (P2129)

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

7.1 Overview of Outputs

- Output 15: Structural plasticity on BrainScaleS-2
- Output 16: Evolving plasticity rules for spiking neural networks

7.2 Structural plasticity on BrainScaleS-2

Sebastian BILLAUDELLE, Benjamin CRAMER, Mihai A. PETROVICI, Korbinian SCHREIBER, David KAPPEL, Johannes SCHEMMEL, Karlheinz MEIER

CDP5 collaboration between SP9 (UHEI, P47), SP4 (UBERN, P71), and Uni Goettingen (UGO, external partner). UHEI, UBERN, and UGO have jointly worked on the theory, network models, experimental setups and data evaluation. UHEI was responsible for the hardware implementation, as well as the commissioning of the system, providing the required hardware and software infrastructure.

To be able to adapt to its inputs and learn new tasks, synaptic plasticity cannot be limited to only changing synaptic weights: the connectome itself undergoes continuous structural modifications during the lifetime of an individual. Since neuromorphic systems impose connectivity constraints on the implementable networks, one may again draw inspiration from their biological archetype. We propose an efficient implementation of structural plasticity and demonstrate its functionality on the BrainScaleS-2 system. The plasticity rule enables neurons to dynamically select a set of suitable synapses out of a pool of potential connections. This policy optimizes performance while at the same time maintaining a sparse connectome.

7.2.1 Research and Infrastructure

Synaptic plasticity is known to not only be limited to adjusting the strength of synapses; the connectome itself undergoes continuous change during the lifetime of an individual. By constraining the number of expressed synapses to enforce a certain level of sparsity, the nervous system appears to manage its spatial and energetic budget. Similar constraints apply to all physical information processing systems, with neuromorphic ones being no exception. In particular, the synaptic fan-in of silicon neurons is often limited.

We implemented a synaptic update policy that incorporates structural plasticity, enabling neurons to dynamically select a set of suitable synapses out of a pool of potential connections that optimizes performance for a chosen task, while maintaining a sparse connectome. An STDP term potentiates correlated connections, a homeostatic regulariser limits post-synaptic firing rates and encourages synaptic competition, and a stochastic component induces exploration. In addition, a pruning condition is executed periodically, removing synapses with a weight below a certain threshold and randomly reassigning them.

Structural plasticity is enabled by bundling k presynaptic sources and injecting them into a single synaptic row (Fig. 19A). Spikes are then distinguished by a label address identifying their respective source. Synapses hold a similar label in their local memory, which allows filtering of afferent spikes: only events with a matching address are forwarded to the home neuron. Pruning and reassignment of a synapse is simply implemented by changing this synaptic label. The reconfiguration is thereby fully local and, in particular, does not involve time-consuming sorting of routing tables or connectivity lists.

We applied the above algorithm to a supervised learning task, where the network was trained to classify the Iris data set. The emulated plasticity rule led to self-organized reconfiguration of receptive fields (Fig. 19B). Ultimately, however, the learning rule enabled the network to achieve near-perfect classification for different sparsity levels (Fig. 19C), demonstrating its ability to ensure a better utilization of synaptic resources without prior knowledge of the input data.

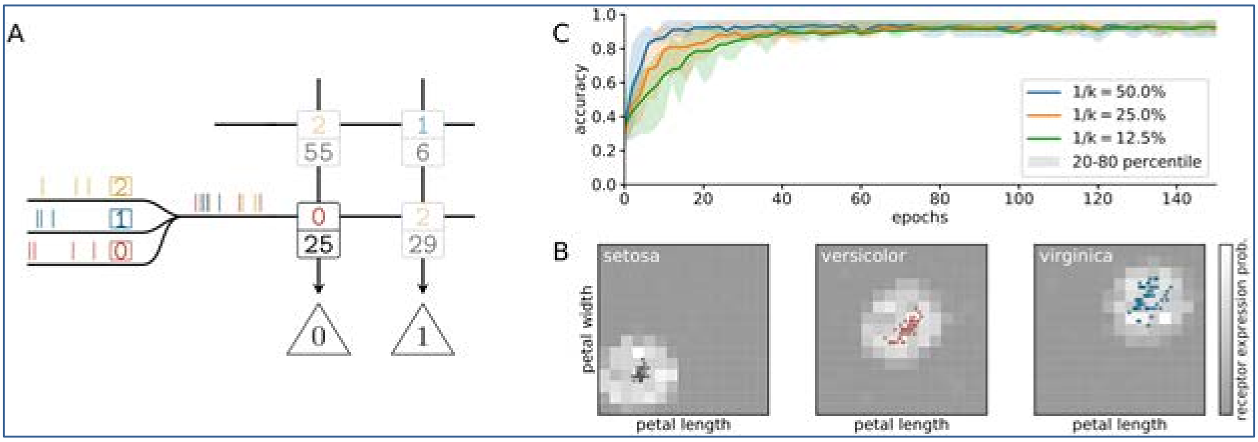

Fig. 19: Self-organizing receptive fields through structural plasticity

(A) Spike trains from different sources can be injected into a single synaptic row. Each synapse filters afferent spikes according to a locally stored label. (B) A network endowed with structural plasticity learns to discriminate between types of Iris flowers (dataset represented by coloured dots). The receptor distribution after training is adapted to the input data distribution. (C) Feature selection through structural plasticity allows the conservation of classification performance even for strongly enforced sparsity 1 − 1/k.

7.2.2 Impact and significance

The results demonstrate how to employ on-chip structural plasticity on BrainScaleS-2 and can be applied to other learning frameworks. This expands the set of experimental scenarios that are amenable to emulation on BrainScaleS-2 and also offers similar potential for other neuromorphic architectures.

We have contributed to two well received and widely viewed mini-documentaries featured in articles and on the Youtube platform: Euronews, April 2019 (E2564, article 1, article 2), and Bloomberg’s “Moonshot” series, November 2019 (E2327, article, Moonshot trailer).

7.2.3 Components

C0349, C0457, C2420

7.2.4 Publications

Sebastian Billaudelle, Benjamin Cramer, Mihai A. Petrovici, Korbinian Schreiber, David Kappel, Johannes Schemmel, and Karlheinz Meier. Structural plasticity on an accelerated analog neuromorphic hardware system. (2019). arXiv preprint https://arxiv.org/abs/1912.12047. (P2240)

Sebastian Billaudelle, Yannik Stradmann, Korbinian Schreiber, Benjamin Cramer, Andreas Baumbach, Dominik Dold, Julian Göltz, Akos F. Kungl, Timo C. Wunderlich, Andreas Hartel, Eric Müller, Oliver Breitwieser, Christian Mauch, Mitja Kleider, Andreas Grübl, David Stöckel, Christian Pehle, Arthur Heimbrecht, Philipp Spilger, Gerd Kiene, Vitali Karasenko, Walter Senn, Mihai A. Petrovici, Johannes Schemmel, Karlheinz Meier. Versatile emulation of spiking neural networks on an accelerated neuromorphic substrate. arXiv preprint https://arxiv.org/abs/1912.12980, accepted for ISCAS 2020, 2019 (P2241).

7.3Evolving plasticity rules for spiking neural networks

Jakob JORDAN, Maximilian SCHMIDT, Walter SENN, and Mihai A. PETROVICI

CDP5 collaboration between SP4 (UBERN, P71), SP9 (UHEI, P47) and external partner RIKEN Center for Brain Science (RIKEN). All partners contributed to the design of the theoretical framework, *evolutionary model and simulation scenarios. The simulations and data analysis were performed at UBERN and RIKEN.*

The search for plasticity mechanisms in biological neural networks can be interpreted as an optimization problem, with researchers in various disciplines performing a distributed manual search based on functional performance, biophysical plausibility and experimental evidence. We propose to augment this manual search with automated optimization methods suitable for discovering biophysically plausible plasticity rules for spiking neural networks. In particular, we employ Genetic Programming to produce compact, interpretable, and generalizable plasticity rules.

7.3.1 Research and Infrastructure

Genetic Programming offers a unique opportunity to unite an automated search for plasticity rules with the desire to intuitively understand the plasticity rules discovered by the algorithm. We successfully applied a particular form of Genetic Programming to three different learning paradigms for spiking neural networks, error-driven, reward-driven and correlation-driven learning, evolving high-performance solutions. In particular, we reproduced previously known solutions for reward-driven learning, evolved approximations to well-known plasticity rules in error-driven learning (see also Outputs 2, 8, 9, 10 and 11) and generated new hypotheses for STDP rules solving correlation-driven learning tasks.

Our successful experiments demonstrate that Genetic Programming is a suitable method for generating interpretable hypotheses for plasticity rules in various learning paradigms. The analytical tractability of the resulting expressions allows one not only to understand the mechanism by which the evolved solution achieves a high task performance, but also to interpret the evolutionary steps by which the search progresses.

7.3.2 Impact and significance

Our results demonstrate the significant potential of Genetic Programming to support traditional research on plasticity in biological systems. The presented methods can be seen as a machinery for automatically generating, testing, or extending hypotheses on learning in spiking neural networks driven by problem instances and constrained by experimental evidence. This approach holds significant potential to accelerate progress towards deep insights into biological information processing.

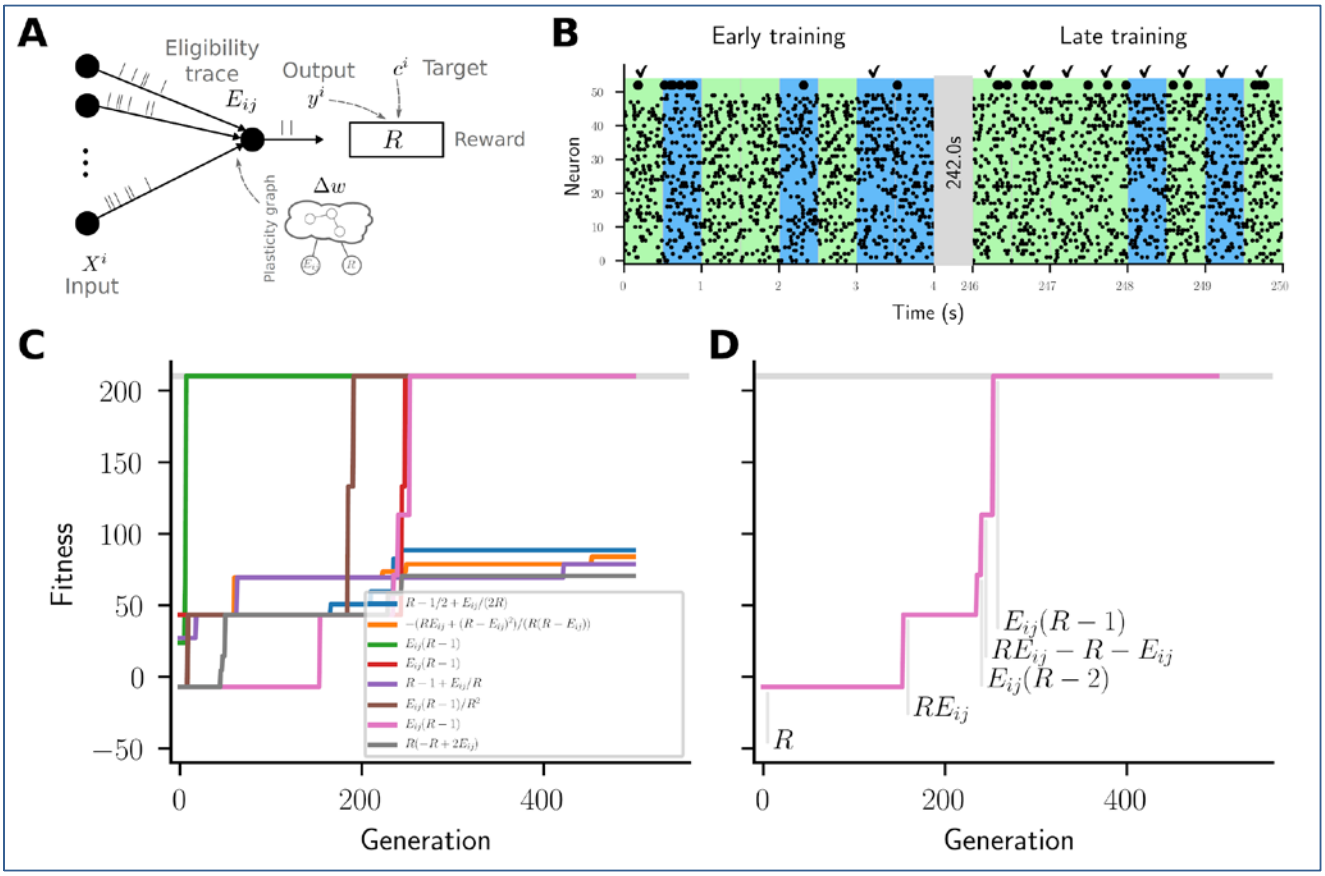

Fig. 20: Cartesian Genetic Programming evolves an efficient reward-driven learning rule

(A) Network sketch. A number of inputs project to a single output unit. Pre- and postsynaptic activity generate an eligibility trace in each synapse. Comparison between the output activity and the target activity generates a reward signal. Both of these are provided to the plasticity graph that returns a weight update. (B) Raster plot of the activity of input units (small black dots) and output unit (large black dots). Blue/green background indicate patterns for which the output should become in/active. Checkmarks indicate correct classifications. Shown are 8 trials at the beginning (left) and the end of training (right) using an evolved plasticity rule with high fitness. (C) Highest fitness per generation as a function of the generation number for multiple runs of the genetic algorithm. Labels are the rule at the end of the respective run. Light grey line indicates reward obtained with a gradient-descent rule. (D) Evolution of plasticity rules over one particular successful run.

7.3.3 Components

C0457, C1032, C2419, C2420, C2547, C2722

7.3.4 Publications

Manuscript in preparation.

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

8.1 Overview of Outputs

- Output 17: Learning event-based visual representations for grasping on SpiNNaker

8.2 Learning event-based visual representations for grasping on SpiNNaker

Jacques KAISER, Alexander FRIEDRICH, Camilo V. TIECK, Daniel REICHARD, Emre NEFTCI, Rüdiger DILLMANN, Guillaume BELLEC, Franz SCHERR, Elias HAJEK, Darjan SALAJ, Robert LEGENSTEIN, Wolfgang MAASS, Oliver RHODES, Steve FURBER

CDP5 collaboration between SP10 (FZI, P52), SP9 (TUGRAZ, P55; UMAN, P63) and external partner University of California, Irvine (UCI). TUGRAZ derived the learning rule e-prop and assisted UMAN for the SpiNNaker implementation of e-prop. FZI developed a pattern classification architecture and integrated it in the robotic demonstrator. UCI assisted with the neural classification architecture originally implemented with the learning rule eRBP on a CPU.

Three-factor synaptic plasticity rules approximating backpropagation-through-time for spiking neural networks were recently derived by CDP5 and the computational neuroscience community. Suited to neuromorphic hardware, these rules can learn online with asynchronous updates. We propose to evaluate the ability to efficiently learn spatio-temporal visual representations from an event-based vision sensor, using e-prop implemented on a SpiNNaker. This evaluation is integrated in a real robotic setup, where a gripper grasps objects according to detected visual affordances. This neuromorphic approach allows fast classification of the event stream while events are gathered by the sensor, thanks to the SpiNNaker live streaming feature.

8.2.1 Research and Infrastructure

The SpiNNaker implementation of the eligibility propagation (e-prop, Bellec *et al*., 2019) learning rule was integrated in a real-world robotic reaching and grasping experiment (Kaiser *et al*., 2019). The spiking network classifies event streams provided by a Dynamic Vision Sensor (DVS) into four classes: ball, bottle, pen and background. The DVS is mounted on a robotic head performing microsaccadic eye movements to perceive the static objects. The Schunk robotic arm and five-finger hand execute the corresponding predefined reaching and grasping motion. We show that e-prop learns quickly and accurately from the spatio-temporal event streams. Additionally, with its live spike streaming feature, SpiNNaker enables online classification as the events are emitted by the DVS.

The network is trained in a supervised fashion from a rosbag dataset of 200 samples. A sample consists of the address events obtained during a microsaccade, with the corresponding object label. Every sample lasts around 700ms. Labels are encoded as spikes and error signals are computed within the network with a population of positive and negative error neurons. Error spikes are integrated in a dedicated error compartment in the learning neurons. The accuracy during training is over 99%, when the network receives a learning signal. More experiments have to be conducted to stabilize learning once the learning signal is removed. A test script for generic pattern classification was provided in the SpiNNaker code base.

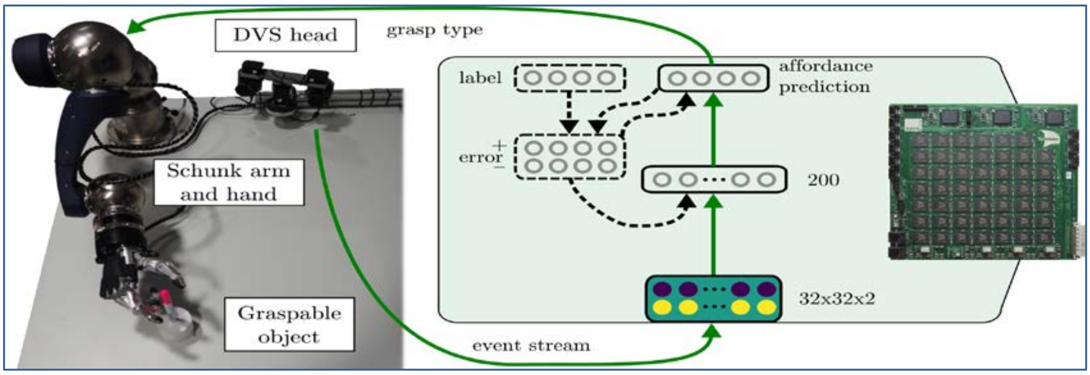

Fig. 21: The robotic closed-loop setup and the spiking neural network on SpiNNaker

The spiking neural network is trained with e-prop from a dataset of event stream in a supervised fashion on a pattern classification task. At test time, inferences about the object class are made online with live spike injection, and the corresponding affordance is sent to the robot arm. Event streams are gathered from static scenes by performing microsaccadic eye movements with a dynamic vision sensor.

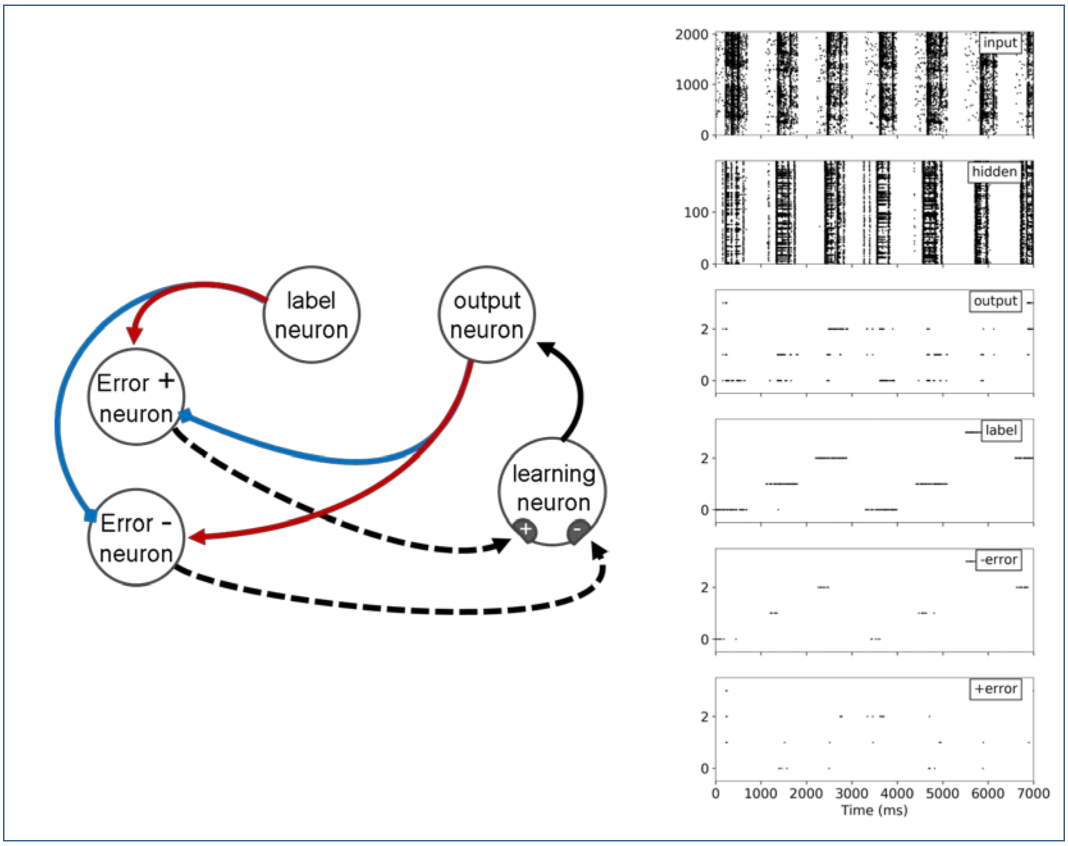

Fig. 22: Wiring of the dedicated learning synapses and network activity during learning

Left: Positive and negative error neurons enable LTP or LTD respectively on the synapses of the learning neurons. The activity of these neurons depend on the label neurons and the output neurons. Right: spike train of the neural populations while learning to classify event streams on the SpiNNaker with e-prop.

8.2.2 Impact and significance

This development paves the way towards the integration of neuromorphic computing technology into the field of robotics. The advantage of neuromorphic technology in this setup over traditional computer architecture is the speed at which inference can be performed, while sensory data from the dynamic vision sensor is streamed to the network.

8.2.3 Components

C2420, C2704, C2719

8.2.4 Publications

The e-prop/SpiNNaker evaluation in a robotic setup is still in the stage of preliminary results.

The corresponding evaluation of eRBP on CPU was published in a pre-print: Jacques Kaiser, Alexander Friedrich, Camillo V. Tieck, Daniel Reichard, Arne Roennau, Emre Neftci, and Rüdiger Dillmann. Embodied event-driven random backpropagation. (2019). arXiv preprint arXiv:1904.04805. (P2519)

The original e-prop pre-print: Guillaume Bellec, Franz Scherr, Elias Hajek, Darjan Salaj, Robert Legenstein, & and Wolfgang Maass. Biologically inspired alternatives to backpropagation through time for learning in recurrent neural nets. (2019). arXiv preprint https://arxiv.org/abs/1901.09049. (P1836)

9 Conclusion and Outlook

In order to further incentivise collaboration among HBP partners, several Co-Design Projects were launched at the beginning of SGA1. Their aim was to define explicitly collaborative objectives that would link research efforts across multiple SPs. Initially centred around the topic of plasticity, CDP5 was further refined and refocused onto the subject of biological deep learning during the transition from SGA1 to SGA2.

Over these four years, an organizational and scientific structure has emerged that reflects, on one hand, the scientific goals of the involved partners, and the incentives generated by the funding scheme (or lack thereof) on the other. CDP5 represented fertile soil for interactions between computational neuroscientists working on functional neuronal networks and developers of efficient neuromorphic substrates capable of implementing these models. Thus, our research was mainly driven by collaborations between SP4 and SP9, with important contributions from more computationally oriented groups in SP3, SP7 and SP10.

Unsurprisingly, the discontinuation of dedicated funding for CDPs in SGA2 had noticeable consequences. Organization and management became significantly more difficult, as it had to rely on SP-dedicated resources; the significant strain caused by the substantial reporting requirements and administration overhead, along with the difficulties surrounding the repeated reorganization of the project, left some partners unfortunately, though understandably, disappointed, leading to losses that will be felt in SGA3. With respect to science, the lack of funding led several other contributors to discontinue their efforts in CDP5 already during SGA2. This was, however, already foreseen during the planning phase for SGA2. While the loss of scientific expertise was certainly regrettable, it left behind a true coalition of the willing, capable of efficient teamwork within a more focused project.

Despite these difficulties, collaboration throughout SGA2 continued to be an exceptionally fruitful one and, in many aspects, even a more productive one, thanks to the foundations laid out in SGA1. Many of the previously initiated lines of research have led to important insights into biological and biologically-inspired learning, significantly advancing the state of the art, as evidenced by a number of well-received publications and a high visibility at international conferences and workshops. Even more importantly, new ideas have spawned new projects, which, even in their early stages, have already sparked interest in the community.

The work in CDP5 will be continued in SGA3 under the aegis of WP3, merging our efforts with those from CDP4. The topic of biological deep learning is deeply and explicitly embedded into the research plans of WP3. We welcome this renewed commitment of the EC to backing collaborative research in SGA3 and believe to have made good use of the associated restructuring process within the HBP in order to continue the successful pursuit of our common scientific goals.