All Models of Plasticity, Learning and Memory developed during SGA2 (D4.3.1 – SGA2)

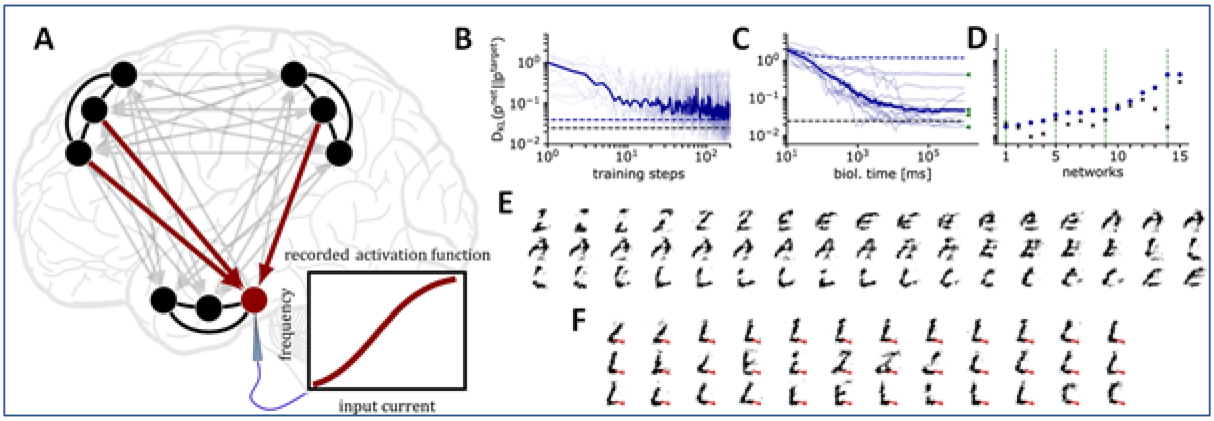

Figure 1: Learning and sampling from memories in the brain.

Project Number:	785907	Project Title:	Human Brain Project SGA2

Document Title:	All Models of Plasticity, Learning and Memory developed during SGA2
Document Filename:	D4.3.1 (D23.1 D69) SGA2 M24 ACCEPTED 201005.docx
Deliverable Number:	SGA2 D4.3.1 (D23.1, D69)
Deliverable Type:	Other
Work Packages:	WP4.3
Key Result(s):	KR4.2
Dissemination Level:	PU = Public
Planned Delivery Date:	SGA2 M24 / 31 Mar 2020
Actual Delivery Date:	SGA2 M26 / 20 May 2020, approved 30 Jul 2020; resubmitted 1 Oct 2020; accepted 5 Oct 2020
Author(s):	Wulfram GERSTNER (EPFL) (P1) Walter SENN (UBERN) (P71) Andre GRUNING (Surrey and Stralsund) (P102/P132) Misha TSODYKS (Weizmann) (P84)
Compiled by:	Andre GRUNING, SURREY/HOST (P102/P132)
Contributor(s):	
SciTechCoord Review:	Martin TELEFONT (EPFL), P1
Editorial Review:	N/A
Description in GA:	For consistent presentation of HBP results, SGA2 M24 Deliverables describing the accomplishments of an entire SP, WP or CDP have been prepared according to a standard template, which focuses on Key Results and the outputs that contribute to them. Project management elements such as Milestones and Risks will be covered, as per normal practice, in the SGA2 Project Periodic Report.
Abstract:	This report describes the progress made on the various models developed in WP4.3 Learning and memory in SGA2 together with related outputs and publications.
Keywords:	Models, plasticity, deep learning, networks.
Target Users/Readers:	Neuroscience scientific community; Computational neuroscience scientific community.

Table of Contents

Table of Figures

History of Changes made to this Deliverable (post Submission)

Date	Change Requested / Change Made / Other Action
20 May 2020	Deliverable submitted to EC
30 Jul 2020	Deliverable approved by EC
01 Oct 2020	Minor editorial change by PCO
01 Oct 2020	Revised version resubmitted to EC by PCO via SyGMa

1. Overview

The theoretical and computational models developed in SP4 occupy a central position in the HBP. One the one hand, they are derived from experimental data produced in the HBP. On the other hand, models are implemented in the HBP platforms, where they serve as "first users". These models also constitute the building blocks of work that will be continued in SGA3, such as bridging scales, network models, models of plasticity, models of cognitive processes and whole-brain models.

SP4 - Synaptic plasticity is the basis of learning and Memory Formation. Computational Models of Synaptic Plasticity describe changes of synaptic strength in the form of update rules either in form of differential equations in continuous time or algorithmically as steps in discrete time. Such models are sometimes also called learning rules, and they can be tested in different learning tasks. The rules of synaptic plasticity in this workpackage and the corresponding Key Result is a cornerstone for the link to cognition, to neuromorphic hardware implementations, and to compact large-scale simulations of learning.

In this deliverable we summarize the work performed as part of *KR4.2 Plausible biological models of plasticity for large networks with non-trivial functionality*.

Some of the learning rules and studies presented here have already been, or are about to be, tested in the NEST simulation environment of the HBP; or in the neuromorphic hardware of the HBP; or in cognitive experiments. For example, the voltage-dependent plasticity rule that was proposed by the HBP theoreticians in SGA1 has now been implemented in the NEST solftware package, under the lead of Mihai A. Petrovici (Bern), Walter Senn (Bern) and Markus Diesmann (Julich). The biologically plausible variant of learning in recurrent networks that avoids known problems of BackProp has been implemented in the hardware framework of HBP partner Uni Heidelberg. And learning rules using theoretically predicted eligibility traces have been tested in cognitive experiments by HBP partner Michael Herzog at EPFL.

2. Introduction

All the outputs in this deliverable are related to *KR4.2 Plausible biological models of plasticity for large networks with non-trivial functionality*.

The fundamental questions we addressed are:

First, how can we generalize the well-known Hebbian paradigm? An answer would be three-factor rules with eligibility traces, but can we design a cognitive experiment to test the model? Outputs 1 and 2 address this question.

Second, how can we avoid backpropagation of errors in the brain? Are the plausible alternatives? An answer to these questions is given in outputs 3 and 4.

Third, statistics needs sampling –but how can the brain sample? And answer to these questions is given in outputs 5 and 6.

Within the framework of KR4.2. several outputs have been achieved and are listed in the following pages. Details will be given in Section 3. These outputs include

1) Output 1: Theory and experiments on learning with eligibility traces
2) Output 2 Synaptic Eligibility Traces and NeoHebbian three-factor rules
3) Output 3 Biologically Plausible Deep Learning in Shallow networks
4) Output 4 Error-correcting learning with spiking neurons
5) Output 5 Deterministic networks for probabilistic computing
6) Output 6 Stochasticity from function in spiking sampling networks
7) Output 7: Spike-based learning on neuromorphic substrates
8) Output 8: Fundamental law of memory recall.

Note that most of the models will be available in the Knowledge Graph and are in the process of being transferred, the final list of the links will be given in due time.

3. Key Results KR4.2 Plausible biological models of plasticity for large networks with non-trivial functionality.

3.1 Outputs

3.1.1 Overview of Outputs

3.1.1.1 List of Outputs

- Output 1: Theory and experiments on learning with eligibility traces
- Output 2 Synaptic Eligibility Traces and NeoHebbian three-factor rules
- Output 3 Biologically Plausible Deep Learning in Shallow networks
- Output 4 Error-correcting learning with spiking neurons
- Output 5 Deterministic networks for probabilistic computing
- Output 6 Stochasticity from function in spiking sampling networks
- Output 7: Spike-based learning on neuromorphic substrates
- Output 8: Fundamental law of memory recall.

3.1.2 Output 1: Theory and experiments on learning with eligibility traces

A large class of algorithms in the field of reinforcement learning uses eligibility traces. In the machine learning theory, eligibility traces are decaying memories of past state-action pairs. In the neuroscience literature, these decaying state-action pairs can be reinterpreted as synaptic memory. Algorithms with eligibility traces are more efficient in learning tasks that involve a sparse and delayed reward. Also, algorithms with eligibility traces fit statistically human behaviour better than those without. However, nobody so far addressed the question whether there are clearly observable qualitative differences in human behaviour that would enable a distinction between the model class with eligibility traces from that without.

At EPFL, two HBP partners (Gerstner and Herzog) representing two different subprojects of the HBP collaborated to address the question of eligibility traces and designed a novel experimental paradigm to solve the above issue. The results of the project are unequivocal: humans use some form of eligibility traces. This clear conclusion stems from the fact that a standard reinforcement learning model of temporal difference learning WITHOUT eligibility traces needs several episodes until information about the location of the reward arrives at states that are a few steps away from the reward location. However, the experimental results clearly indicate that already after a SINGLE successful episode, humans bias their choices to actions toward the reward. This form of ‘one-shot-learning’ indicates that humans make use of eligibility traces. However, no statement could be made whether these eligibility traces are implemented on a synaptic level, or using some other transient memory mechanism.

In additional experiments it was found that pupil dilation correlates with reward-related information. Moreover, pupil dilation also reflects learning with eligibility traces. This result indicates that physiological correlates of learning also incorporate learning with eligibility traces.

Importantly, the experimental design would not have been possible with the rigorous theoretical analysis performed by the theory lab in SP4, WP4.3.

This result is an important step beyond the state-of-the art. In fact, earlier experimental designs were limited to experiments of a maximal sequence of two or three, whereas in the new HBP experiments much longer sequences are possible. Only the clever design made it possible to observe eligibility traces via one-shot learning.

This output is based on an interdisciplinary collaboration within HBP (Herzog + Gerstner) psychophysics and theory, spanning different subprojects within HBP.

3.1.3 Output 2: Synaptic Eligibility Traces and NeoHebbian three-factor rules

This review paper summarizes 50 years of research on biological implementation of reinforcement learning algorithms. The field probably started in the 60ies with first formulations of Crow and Barto and Sutton; then the field did a detour to more algorithmic reinforcement learning rules as summarized in the book of Sutton and Barto in the 90ies. But then the field went back to a computational neuroscience perspective and asked, in particular driven by the results of Wolfram Schultz and colleagues: how could reinforcement learning be implemented in biological wetware, i.e., what are the synaptic learning rules (at the level of synaptic contacts) that would allow reinforcement learning (at the level of behaviour).

This paper shows that theoretical work in the field of computational neuroscience predicted an extension of standard Hebbian learning rules to neoHebbian three-factor rules. Only after 2014, experimental results were published from different groups across the world (Japan, US, Europe) that checked these predictions. It was found that short and delayed pulses of neuromodulator indeed affect synaptic plasticity – as predicted by three-factor rules.

The basic idea of three-factor rules is as follows: joint activity of pre- and postsynaptic neurons in a Hebbian sense sets a flag and the synapse and marks the synapse as eligible for weight changes. However, weight changes do not occur immediately but only if, in parallel or during the next few seconds a neuromodulator is present. Thus, while in Hebbian learning two factors are needed, viz. pre- and postsynaptic activity, neoHebbian rules need three factors, viz. the two Hebbian factors AND a neuromodulator as a third factor.

Importantly such a framework provides a mathematical foundation not just for reward based learning, but also for modulation of learning by surprise, attention, or novelty.

This review paper has already started to influence the field, together with a review paper by HBP partner Pieter Roelfsema who also extends the framework of Hebbian two-factor rules to novel multi-factor rules from a slightly different perspective. This interaction shows that information on synaptic plasticity rules has circulated within the HBP between different subprojects and has now reached different subcommunities of neuroscience (outside HBP!)

3.1.4 Output 3. Biologically Plausible Deep Learning in Shallow networks

It is widely agreed that the standard backpropagation algorithm is not implementable as such in the brain. Several labs inside HBP (e.g., Roelsfema lab, Netherland Institute of Neuroscience and Walter Senn, Univ. of Berne Switzerland and outside HBP (e.g., Yoshua Bengio, Timothy Lillicrap, Blake Richards) are working on developing biologically more plausible versions of Deep Learning. In this HBP study in the Gerstner lab at the EPFL, a completely different approach was taken. Given that deep learning is difficult, the question was asked how much can be achieved with learning in shallow networks where only the output layer is learned in a supervised fashion, whereas the hidden layer is not. In fact, for the output layer, a gradient rule on the error surface does not pose any problem in terms of biological plausibility, because no backpropagation of errors is necessary. In the paper by Illing et al., it was found that many of the results that were taken as proof for deep learning with biologically plausible rules can in fact be repeated with shallow networks without any need of backpropagation of errors. This is a considerable advance over the state of the arts because it changes the perspective – and this question has brought together different labs within HBP

Stand-alone output, but influenced by discussions with HBP partners from different subprojects.

3.1.5 Output 4: Error-correcting learning with spiking neurons

In order to make learning algorithms in spiking neurons amenable to contemporary neuromorphic platforms, we have developed a version of neuronal backpropagation for networks of single-compartment neurons in a direct collaboration between theoreticians and hardware groups in HBP.

The HBP team has made significant progress beyond the state-of-the-art in implementing the general framework for learning from real-time data streams (as developed by the HBP theory group in Heidelberg) in low-power neuromorphic systems. The theoretical 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. 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. Through the choice of connectivity, neuron model and learning rule 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.

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.

3.1.6 **Output 5** *Deterministic networks for probabilistic computing*

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 neural networks can be used as sources of uncorrelated noise, exploiting the decorrelating effect of inhibitory feedback.

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

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.

3.1.7 Output 6: *Stochasticity from function in spiking sampling networks*

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.

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.

The HBP team 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"

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.

3.1.8 Output 7: Spike-based learning on neuromorphic substrates

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

We have worked on the implementation of biologically inspired but machine-learning relevant learning rules on the SP9 neuromorphic platforms. The neuromorphic platforms exhibit a number of constraints that make it difficult to translate fully-programmable software-based simulations to them in a one-to-one way. For example, the DLS system (SP9.2) has only one fixed type of pre- and postsynaptic traces.

SURREY (Brian Gardner and André Grüning) have converted standard-software simulations of larger spiking neural networks utilising HBP-internal (T433 RUP / T432 RUP/SGA1) and HBP-external plasticity rules into networks that only use a number of neurons commensurate with the current size of the currently available DLSv2 system (SP9.2) and learning rules that only require features already present in the HiCANN PPU systems: HBP-funded postdoc Brian Gardner has been successful in implementing learning on three standard ML-benchmark data set (Iris data set, Wisconsin breast cancer data set and MNIST) utilising only a number of neurons and synapses commensurate with the HiCANN-DLS platform. We currently need about 3-5 times more neurons than the DLSv2 offers, but since its larger successor

HiCANN-X is reaching maturity during SGA2, SGA3 will see the transfer of the simulation to the PPU of HICANN-X. An overview publication on computing with spiking neural networks has already spun off from this effort, and a publication on the core results is in preparation.

We have also reached out to the research community, and HBP-derived learning rules INST/FILT are now being explored by neuromorphic groups outside of HBP (with University of Minnesota, submitted).

3.1.9 Output 8: Fundamental law of memory recall.

Memory recall is a classical paradigm developed by cognitive psychologists to probe human memory. Usually it involves subjects to recall lists of randomly assembled words after a brief exposure. It was observed over the years that people cannot reliable recall even relatively short lists, and when lists become longer, the fraction of recall words steadily declines. Besides the number of words, recall performance depends on other experimental conditions, such as age of the subjects, their motivation, rate of presentation etc.. Thus, it appears that there can be no mathematical theory that would predict recall performance. Yet in this project we proposed a phenomenological model that predicts recall from a set of two fundamental principles concerning the encoding of memories in the brain and search algorithm that generates recall trajectories based on encoding. This model can be analytically solved to express the average number of recalled words, R, via the number of remembered words. By performing behavioural experiments, we confirmed that the relation between R and M is indeed predicted by this expression, and hence most of the variability in the performance is explained by the variability in the number of words in memory after presentation. This result demonstrates that some memory processes are universal for all the people, which opens a new direction in memory research. (Naim et al, 2020[1])

3.2 Validation and Impact

3.2.1 Actual and Potential Use of Output(s)

The actual and potential use of these Outputs is already described inside each Outputs.

3.2.2 Publications

The significances are already explained inside the Outputs.

Output 1 : Theory and experiments on learning with eligibility traces

P2214 : M.P. Lehmann, H.A. Xu, V. Liakoni, M.H. Herzog, W. Gerstner, and K. Preuschoff (2019) One-shot learning and behavioral eligibility-traces in sequential decision making eLife 8:e47463 doi: 10.7554/eLife.47463

Output 2 : Synaptic Eligibility Traces and NeoHebbian three-factor rules

P1348 : W. Gerstner, M. Lehmann, V. Liakoni, and J. Brea (2018) Eligibility traces and plasticity on behavioral time scales: experimental support of NeoHebbian three-factor learning rules. Front. Neural Circuits, 12:53 doi: 10.3389/fncir.2018.00053

Output 3 : Biologically Plausible Deep Learning in Shallow networks

P2041 : Biologically plausible deep learning—But how far can we go with shallow networks? B Illing, W Gerstner, J Brea ; Neural Networks 118, 90-101

[1] Naim, Katkov, Romani, Tsodyks, (2020) Fundamental law of memory recall. Phys Rev. Letters 124:018101 https://journals.aps.org/prl/pdf/10.1103/PhysRevLett.124.018101

Output 4 : Error-correcting learning with spiking neurons

João Sacramento, Rui Ponte Costa, Yoshua Bengio, and Walter Senn. Dendritic cortical microcircuits approximate the backpropagation algorithm. In Advances in Neural Information Processing Systems, pages 8721–8732, 2018. (P1527) https://arxiv.org/pdf/1810.11393.pdf

Output 5 : Deterministic networks for probabilistic computing

P843 : Jordan, J., Petrovici, M. A., Breitwieser, O., Schemmel, J., Meier, K., Diesmann, M., & Tetzlaff, T. (2019). Deterministic networks for probabilistic computing. Scientific reports, 9(1), 1-17. https://www.nature.com/articles/s41598-019-54137-7

Output 6 : Stochasticity from function in spiking sampling networks

P1447 : Dold*, D., I., Bytschok*, Kungl, A. F., Baumbach, A., Breitwieser, W. Senn O., Schemmel, J., Meier, K. and Petrovici*, M. A. (2019). Stochasticity from function - why the Bayesian brain may need no noise. Neural Networks, 119, 200-213. doi: 10.1016/j.neunet.2019.08.002. https://boris.unibe.ch/136724/1/Dold2019Stochastic.pdf

Output 7 : Spike-based learning on neuromorphic substrates

P2221 : Jang, Simeone, Gardner, Grüning**:** "An Introduction to Probabilistic Spiking Neural Networks: Probabilistic Models, Learning Rules, and Applications ", IEEE Signal Processing Magazine 36(6), 64-77, 2019. DOI: https://doi.org/10.1109/MSP.2019.2935234 (P2221)

Output 8 : Fundamental law of memory recall.

N/A

4. Conclusion and Outlook

Achievements: KR4.2 planned to develop biologically plausible learning rules in functional networks. In SGA2 several of such rules have been developed that include a biological form of eligibility traces and local error representations. The eligibility-trace learning addressed the temporal credit assignment problem and was applied to temporal tasks in shallow networks. The error-driven plasticity addressed the spatial credit assignment problem and was applied to deep networks. Spiking versions of such plasticity rules were studied in recurrent networks. The spiking networks were simulated in standard software, and they were also partially implemented in the Heidelberg neuromorphic hardware. In addition, a theory of free memory recall was developed.

Significance: The plasticity rules yield insight into how biological neurons may solve the spatio-temporal credit assignment problem. The work helps to implement functionally relevant plasticity rules in neuromorphic hardware. The memory research reveals universal laws of memory recall in humans.

Outlook: In SGA3 we will combine the theories of eligibility trace plasticity and error-driven plasticity with the applications to spike-based computation. We expect this combination to boost the computational power of neuromorphic hardware when applied to learning problems.