A Framework for plasticity and meta-plasticity learning rules (D3.11 – SGA3)

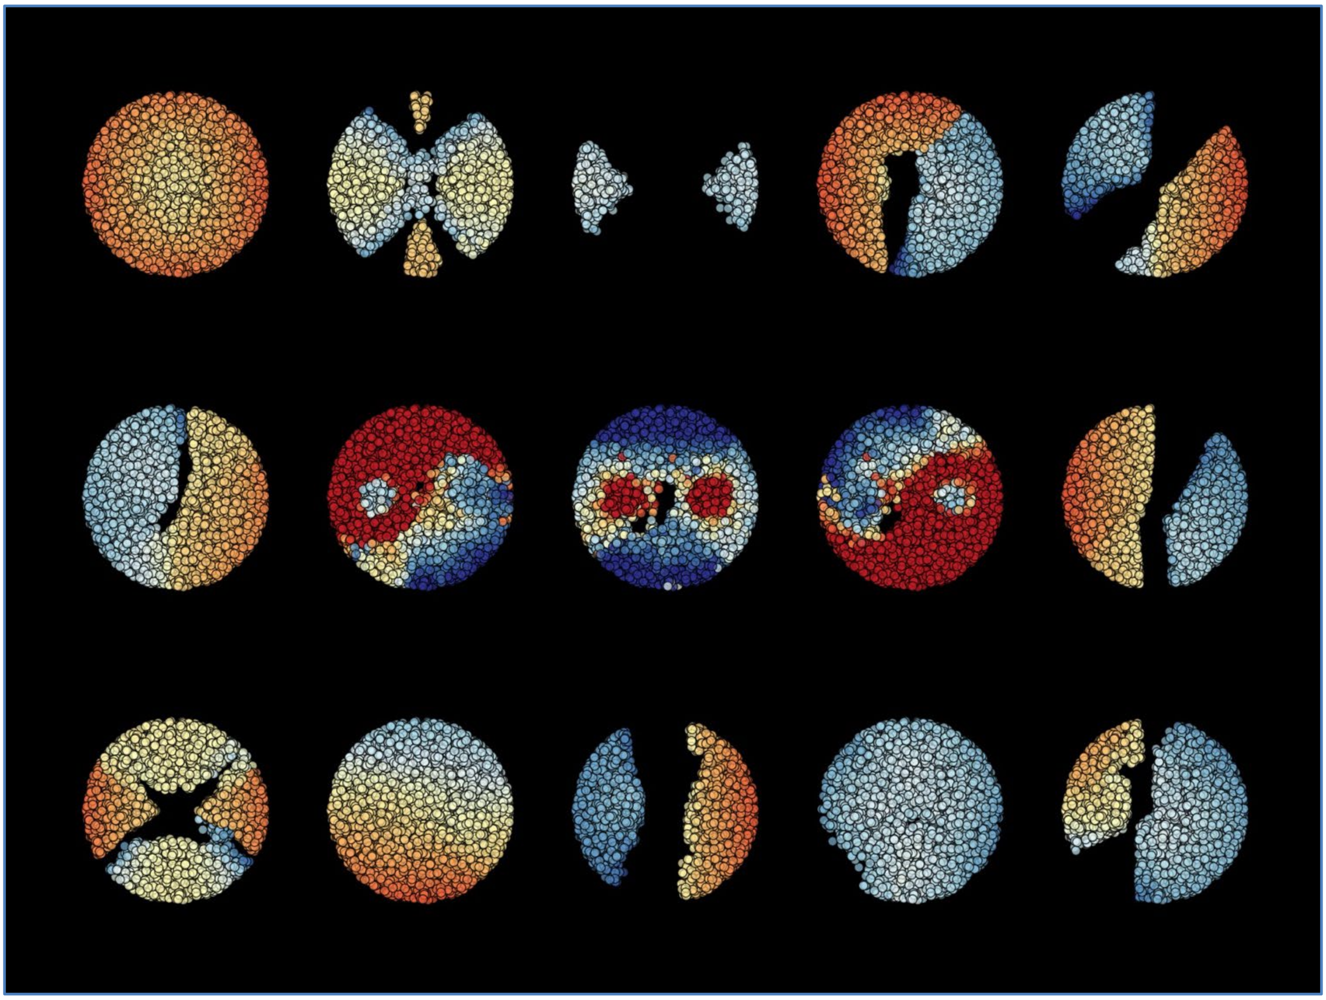

Figure 1: Illustration of the on-chip classification process with the Yin-Yang dataset

Adapted from Göltz et al., 2021 - P2917. This rendering was part of several art exhibitions, including one at Marseille main station and the Théâtre de l'Alcazar, during the final HBP Summit in 2023.

Project Number:	945539	Project Title:	HBP SGA3

Document Title:	A framework for plasticity and meta-plasticity learning rules
Document Filename:	D3.11 (D30) SGA3 M42 SUBMITTED 230922.docx
Deliverable Number:	SGA3 D3.11 (D30)
Deliverable Type:	Demonstrator
Dissemination Level:	PU = Public
Planned Delivery Date:	SGA3 M42 / 30 SEP 2023
Actual Delivery Date:	SGA3 M42 / 22 SEP 2023
Author(s):	Mihai A. PETROVICI, UBERN (P71) & UHEI (P47)
Compiled by:	Mihai A. PETROVICI, UBERN (P71) & UHEI (P47), Virginie SABADO, UBERN (P71)
Contributor(s):	Walter SENN, UBERN (P71), contributed to Section 2.1.1.1, 2.1.1.5, 2.1.2.7 Cyriel Pennartz, UvA (P98), contributed to Section 2.1.1.3 André Grüning, HOST (P132), contributed to Section 2.2.1.6 Thomas NOWOTNY, UoS (P106), contributed to Section 2.2.2.1
WP QC Review:	Jeannette BOSCHMA, UM (P117); Anita VAN OERS, UM (P117)
WP Leader / Deputy Leader Sign Off:	Rainer GOEBEL, UM (P117)
T7.4 QC Review:	N/A
Description in GA:	Large-scale deployment of the models and algorithms from OP3.13 on available EBRAINS platforms, as well as a hybrid demonstrator exploiting the advantages of available EBRAINS platforms for learning-to-learn setups involving both synaptic plasticity and meta-plasticity.
Abstract:	All neural networks, both biological and artificial, learn through plasticity. In biology, a plethora of data on plasticity exists, but a unifying theoretical framework is missing. Conversely, deep learning techniques can be rigorously derived, but require unbiological components. In SGA3, we have taken significant steps towards unification. Guided by normative principles and inspired by frameworks from machine learning and physics, we have developed a series of models that shed light on various aspects of this grand question. They encompass both rate- and spike-based models and are explicitly simultaneously aimed at explaining biological observations and at advancing the state of the art in machine intelligence, in particular through deployment on neuromorphic substrates.
Keywords:	Plasticity, learning, dendritic computation, Bayesian inference, evolutionary strategies, benchmarking, deep learning, spike timing, feedback learning, creativity, dreams, sleep, energy-based models, neuromorphic engineering, neuromorphic algorithms, quantum tomography, structural plasticity, error backpropagation, backpropagation through time, cortical microcircuits, teaching
Target Users/Readers:	Clinicians, computational/theoretical/systems/imaging neuroscientists, computer scientists, neuromorphic/AI experts consortium members, funders, general public, HPC community, platform users, policymakers, researchers, students

Table of Contents

Table of Figures

1. Introduction

Synaptic plasticity is the core mechanism underlying the brain's capability to learn and improve. Over the past decades, numerous studies of plasticity have been published, both experimental and theoretical. However, many questions remain, and the field is in dire need of rigorous theoretical frameworks with both demonstrated complex functionality and experimental predictive power.

In our work, we mostly follow a top-down approach: starting from normative frameworks inspired by well-established theories in physics and mathematics, we have derived mathematical expressions for synaptic learning rules that can be directly connected to electrophysiological data and exhibit enough computational power to rival state-of-the-art approaches in machine learning. This approach facilitates a rigorous and comprehensive understanding of neuron and synapse dynamics and how they synergise to yield particular types of functionality at the network level. Moreover, it allows us to establish common ground between different existing theories, thus enabling a unified perspective on cortical computation and learning. Complementing this analytical approach, we have also developed a novel evolutionary framework – an automated, heuristic search for powerful synaptic plasticity rules. Importantly, this method also produces explicit mathematical expressions for synapse dynamics, following the same general goals of interpretability as our top-down approaches, thus enabling deeper understanding of the underlying computational principles.

Our strong links to modern methods in machine learning match and bolster recent efforts to rekindle the connections between computational neuroscience and artificial intelligence – a field that has been directly inspired by key insights from neuroscience many decades ago but has since strayed away from biological plausibility in favour of pure functionality and performance. In particular, the core algorithm behind deep learning in neural networks – error backpropagation – has long been described as unbiological due to its reliance on spatiotemporally nonlocal computation. However, these arguments are receiving increasing pushback, and recent years have seen a veritable explosion of research on bio-plausible backpropagation, with applications in both basic (neuroscience) and applied research (neuromorphic engineering). In our work, we tackle some of the unsolved problems of credit assignment in deep cortical hierarchies and demonstrate powerful applications of our models to machine learning problems. Furthermore, we harness multiple hardware platforms (HPC, GPU, neuromorphic) to advance the state of the art in computational speed and efficiency.

Thus, our work is positioned at the interface between biological and artificial intelligence, with a particular focus on the physical substrate of neuronal computation. We simultaneously tackle open issues in our understanding of the brain and provide novel blueprints for in-silico computation that push the technological state of the art. This interdisciplinary nature of our R&D efforts and outputs addresses a broad spectrum of interests, especially in the communities working in cortical electrophysiology, computational neuroscience, neuromorphic computing and artificial intelligence.

In addition to our scientific work, we also undertook significant efforts in terms of education and teaching. We have developed and taught courses on various aspects of computational neuroscience and brain-inspired computing for multiple venues (conferences, workshops, summer schools etc.) and have also created a neuromorphic teaching platform for the general public but public but targeted particularly at schoolchildren and students.

Our research addresses neuro-synaptic information processing at multiple scales, from detailed models of neurons and synapses to cortical microcircuits and ultimately to multi-area models of cortex. In Task T3.5, our focus lies more on the detailed dynamics of neurons and synapses, but these feed directly into larger-scale models of cortical and cortex-inspired processing that are distributed among multiple Tasks in Work packages 2 and 3 of the Human Brain Project. Therefore, we provide many direct links to such associated lines of research and the corresponding documentation throughout this report. All of our work directly addresses WPO3.1: Enhanced real-world task performance through biologically plausible adaptive cognitive architectures running on neuromorphic hardware and closed-loop neuro-robotics platform.

Our results and workflows rely on several core components provided by the HBP consortium in general and the EBRAINS research infrastructure in particular. Many of them are rooted in close collaborations between partners in the HBP and make use of dedicated computational resources, either for simulation using the high-performance computing at the Swiss National Supercomputing

Centre (CSCS), provided through the ICEI/Fenix Infrastructure project, or for demonstrating applications in the realm of AI (BrainScaleS-2, GeNN, CSCS).

This document is structured as follows. The main results are described in Section 1, which is subdivided according to the two major groups of outputs: small-scale, detailed studies of synaptic plasticity in single neurons, neuronal microcircuits and small networks (Section 1.1) and large-scale deployment on the HBP/EBRAINS Platforms (Section 1.2). A succinct overview of the undertaken education and outreach is given in Section 1.3. Finally, given how our work combines both foundational and application aspects, we briefly point to potential continuations that go beyond the HBP, but that can profit greatly from the EBRAINS infrastructure (Section 2).

1. Outputs

1.1 Algorithms and software, first tests and applications, as well as small-scale hardware demonstrators of detailed synaptic plasticity and meta-plasticity rules for cortical hierarchies with structured neurons

1.1.1 Micro- and mesoscopic models

1.1.1.1 Neurophysiology and function of dendritic morphology

Contributing Partners: UBERN (P71)

Dendrites shape information flow in neurons. Yet, there is little consensus on the level of spatial complexity at which dendrites operate. We developed an automated Toolbox to reduce the dendritic morphology and compartmentalisation to a given number of branches, while trying to keep the nonlinear synaptic and dendritic processing of the original tree the same (Figure 2).

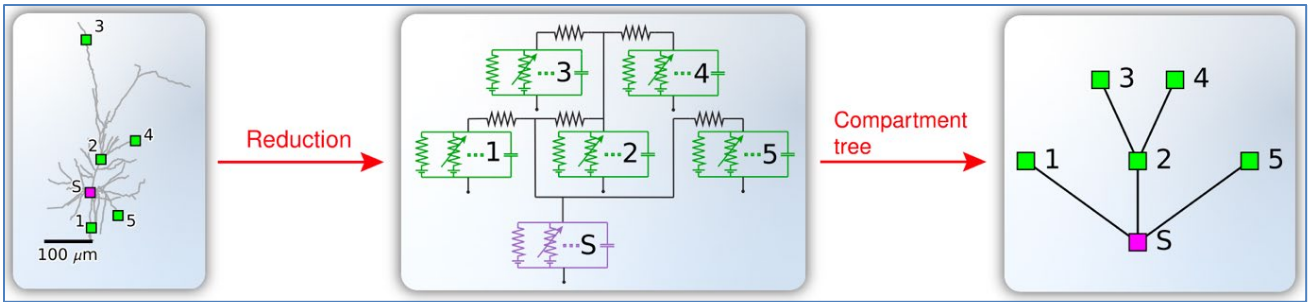

Figure 2: Automatised dendritic reduction method.

Adapted from Wybo et al., 2021 - P2760. For any set of locations on a given morphology (left, here an L2/3 pyramidal cell [Branco and Häusser, 2010]), a reduced compartmental model can be derived (middle), with an associated schematic representation (right).

Through carefully chosen parameter fits, we obtained accurate reduced compartmental models at any level of complexity. We showed that (back-propagating) action potentials, Ca^{2+} spikes, and N-methyl-D-aspartate (NMDA) spikes can all be reproduced with few compartments. We also investigated whether afferent spatial connectivity motifs admit simplification by ablating targeted branches and grouping affected synapses onto the next proximal dendrite. We found that the voltage in the remaining branches is reproduced if temporal conductance fluctuations stay below a limit that depends on the average difference in input resistance between the ablated branches and the next proximal dendrite. Furthermore, our methodology fits reduced models directly from experimental data, without requiring morphological reconstructions. We provided software that automatises the simplification, eliminating a common hurdle toward including dendritic computations in network models. This work was published in PNAS (Wybo et al., 2021 – P2760).

1.1.1.2 Conductance-based Bayesian dendritic computation

Contributing Partners: UBERN (P71)

To steer successful behaviour, our brains can leverage a plethora of information sources. We can rely on our various senses of the world, as well as on expectations shaped by prior experience (Figure 3a). Crucially, these sources are characterised by different levels of reliability. For example, low lighting conditions may make your eyes less reliable relative to your ears. Thus, our brain needs to be able to not only integrate these disparate sources of information, but additionally weigh them according to their (relative) reliability/(un)certainty. Such computations can be formalised in a Bayesian framework, but their implementation in the biological substrate remains unclear.

We have developed a new model of learning uncertainty-modulated integration of information in cortex. In contrast to previous approaches arguing for integration on the circuit level, our model proposes that individual neurons are naturally performing the required computations. In this view, probabilistic computations are deeply ingrained into the very substrate of our brains, rather than an emergent network phenomenon. An associated publication (Jordan et al., 2021 – P2837) is in review.

We propose that basal dendrites represent likelihood functions of a neuron's preferred feature, while apical dendrites encode priors. The soma acts to combine these two ingredients and thus computes a posterior distribution (see Figure 3b). Neuron models, which mimic synaptic coupling via conductance changes of ion channels in the membrane, implement the required computations naturally through their dynamics. Furthermore, we derived a local plasticity rule, which allows neurons to learn approximate posterior distributions from examples.

We demonstrated that our proposed models could learn to optimally combine uncertain information from different information sources. The behaviour of the model closely matches data from psychophysical experiments on mice performing a similar task (see Figure 3c, d). Simultaneously, the dynamics of the model, specifically its firing activity, closely resembles recordings from monkeys in a cross-modal suppression scenario: neuronal responses to a visual stimulus can be suppressed by simultaneously providing a mismatched vestibular stimulus.

Besides these, our model makes additional experimentally testable predictions. A characteristic feature of Bayesian computation is the reduction of uncertainty through new information. Since uncertainty in our model represents uncertainty via membrane potential variability, we would expect membrane potential variability to decrease after stimulus onset (see Section 1.1.1.3).

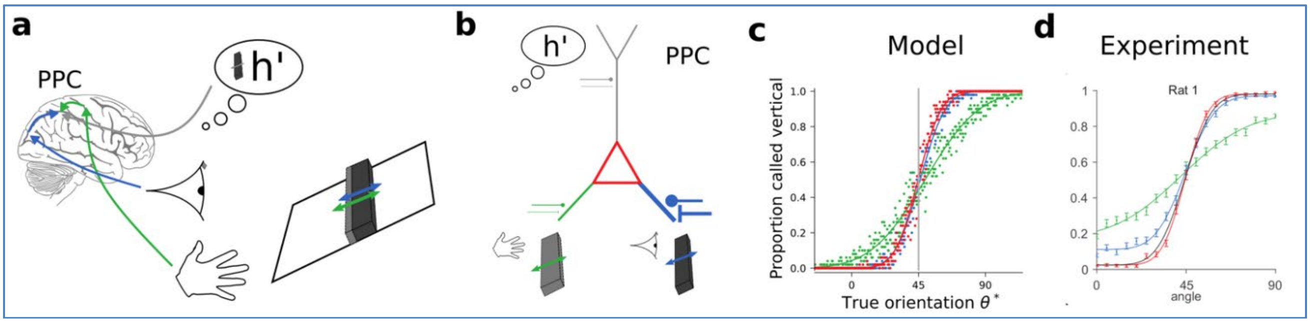

Figure 3: Dendritic opinion weighting.

Adapted from Jordan et al., 2021 – P2837. (a) Brains constantly need to combine information from different sources, here tactile and visual, to gather information about their environment. Typically, not all sources are equally reliable, thus requiring a probabilistic weighting of information. (b) We propose that individual cortical neurons are naturally equipped to perform the required computations. Specifically, we suggest that each dendritic "compartment", logical subunits of complex cell morphologies, represent likelihood functions and priors, while the soma combines these to compute a posterior distribution. (c) We trained our model to solve a typical multisensory integration task in which the two available modalities are characterised by different noise amplitudes. Similar to mice (d), the model learns to focus mostly on the more reliable modality while still being able to benefit from the complementary information in the less reliable stimulus.

1.1.1.3 Hallmarks of Bayesian computation in cortical data

Contributing Partners: UBERN (P71), UvA (P98)

A recently proposed theory (see Section 1.1.1.2) suggests that individual cortical neurons integrate incoming and prior information in a Bayesian way. The model predicts that neurons decrease their membrane potential variability in response to receiving additional information. More specifically, it predicts that this decrease is proportional to the reliability of this information.

We investigated these predictions in single neuron spiking data recorded in the mouse posterior parietal cortex (PPC). The PPC is known for receiving and integrating inputs from all sensory modalities. During measurement, the mice had to continuously monitor and report changes in two simultaneously presented stimuli; a tone that changed in pitch or a visual moving grating that changed its direction. To manipulate the reliability of the changes occurred either at perceptual threshold or at maximum intensity (see Figure 4Figure 4a).

Note that we were in interested in the intracellular membrane potential dynamics but that the data was an extracellular measurement of neuronal spiking. To obtain the former from the latter we used simulation-based inference (SBI). SBI allows us to infer the “hidden” parameters underlying an observable phenomenon. This is achieved by simulating a model neuron with a large number of parameters and then learning the mapping from the resulting observations (spiking activity) back to the parameters (membrane potential statistics).

We found that neurons responded to all stimulus changes with higher and more regular firing (see Figure 4b). Furthermore, we found that, in line with the predictions made by the normative neuron model (see Section 1.1.1.2), the inferred membrane potential becomes less variable post stimulus onset (see Figure 4c), and that this decrease is weaker for threshold stimuli. Moreover, the theory provides a normative explanation for the stronger decrease in audio-visual trials: since multiple stimuli convey more information than a single one, the decrease in uncertainty is larger, which, at the neuronal level, is represented by a stronger decrease in membrane potential variability.

This work represents a cross-WP collaboration with task T2.1 and is in the process of being published in the Live Paper 2.1b. It was also already featured at conferences such as CCCN and CNS 2023.

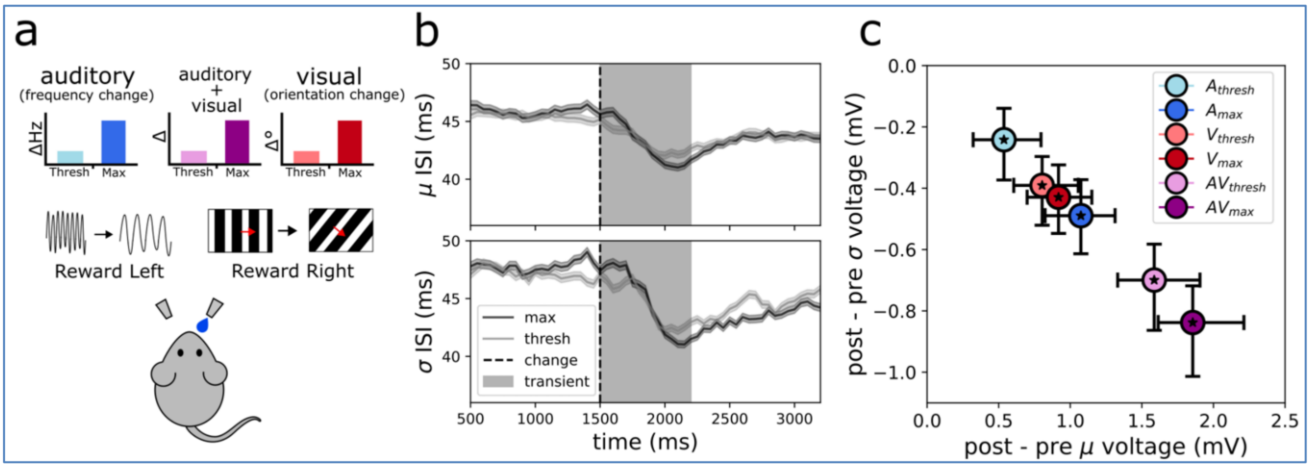

Figure 4: Bayesian inference in mouse data.

Adapted from the Live Paper 2.1b. (a) Audio-visual change detection task: mice were tasked with reporting changes in either of two simultaneously presented continuous stimuli. Changes in the pitch of the auditory stimulus and changes in the direction of the visual stimulus were reported with a lick to the left and right, respectively. (b) Neural activity surrounding stimulus change. Neuronal firing increases (decrease in mean ISI) and becomes more regular (decrease in standard deviation ISI) in response to stimulus change. (c) Inferred changes in membrane potential statistics. Neurons decrease their membrane potential variability post stimulus change. The decrease was the strongest in audio-visual change trials.

1.1.1.4 Learning to estimate the uncertainty of predictions

Contributing Partners: UBERN (P71)

To act in a complex world, our brain needs to correctly anticipate the near future. Prediction errors, the difference between predicted and observed circumstances, are the key to learning successful behavioural strategies. The real world is, however, highly variable and uncertain, such that the brain needs mechanisms to account for uncertainty when evaluating prediction errors. Cognitive studies indicate that uncertainty guides predictive learning, and neurons in cortical circuits that respond to unexpected stimuli could be identified. Yet, the neural mechanisms underlying the processing of prediction errors in uncertain environments are still largely unknown.

Based on normative theories, we have proposed that the identified prediction error neurons encode uncertainty-modulated prediction errors (UPEs). With computational modelling, we have studied circuit mechanisms for estimating uncertainty and computing UPEs.

Imagine a mouse has to predict a somatosensory stimulus (s) based on an auditory stimulus (Figure 5A). When the somatosensory stimulus is different than expected, layer 2/3 pyramidal cells in the posterior parietal cortex (PPC) become active. We hypothesise that these cells encode the UPE. When sampling somatosensory stimuli from a Gaussian distribution (Figure 5B), the activity of these prediction error neurons will be inversely modulated by the variance of the distribution.

In particular, we hypothesise that the layer 2/3 circuit calculates the UPE through the subtractive and divisive inhibition by different inhibitory cell types, such as somatostatin-positive (SST), and parvalbumin-positive (PV) interneurons (Figure 5C/D). By implementing the calculation of UPEs in a microcircuit model, we show that different cell types in cortical circuits can compute the means and variances of the stimulus distribution. With local activity-dependent plasticity rules, these computations can be learned context-dependently, and allow the prediction of upcoming stimuli and their distribution. Finally, we show that the resulting UPEs allow an organism to optimise its learning strategy with respect to external circumstances by enabling adaptive learning rates.

These mechanistic insights are an important step towards a better understanding of perception and learning in uncertain environments. A preprint (Wilmes et al., 2023 – P4014) is in review and has already received much attention, including invited talks, at multiple conferences and workshops (CNS, COSYNE 2023; Bernstein 2022).

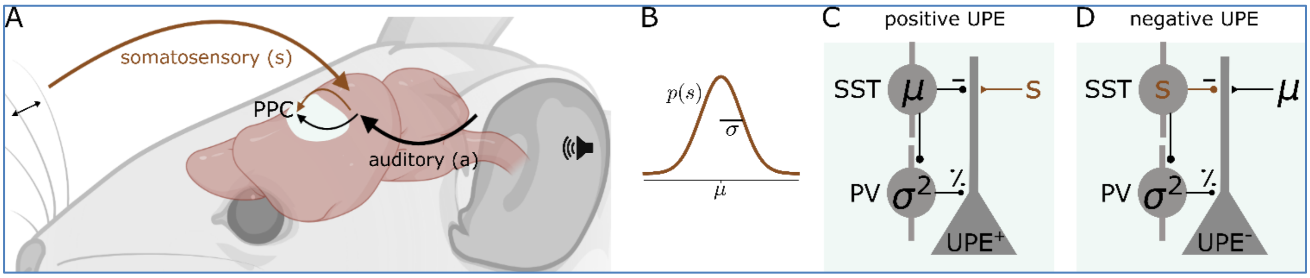

Figure 5: Uncertainty-modulated prediction errors.

Figure taken from Wilmes et al., 2023 - P4014. Computation of UPEs in cortical microcircuits. (A) a mouse learns the association between a sound (a) and a whisker deflection (s). The posterior parietal cortex (PPC) receives inputs from both somatosensory and auditory cortex. (B) The whisker stimulus intensities are drawn from a Gaussian distribution. (C) Positive prediction error circuit consisting of three cell types: layer 2/3 pyramidal cells (triangle), somatostatin-positive interneurons (SST, circle) and parvalbumin-positive interneurons (PV). Pyramidal cells represent the UPE, SSTs represent the mean prediction, and PVs the variance. (D) Negative prediction error circuit, similar to C, SST now represent the stimulus and the mean prediction is an excitatory input.

1.1.1.5 Predictive olfactory learning in the fruit fly

Contributing Partners: UBERN (P71) (external collaborators, U Fribourg conducting experiments)

Olfactory learning and conditioning in the fruit fly is typically modelled by correlation-based associative synaptic learning rules, such as Hebbian plasticity. While predictive coding has been suggested as a key to understanding cortical processing in more complex brains, it remains unclear how much these concepts can also be applied to the simple associative learning in simpler organisms. We considered olfactory learning in the fruit fly and showed that various data on behaviour and odour processing microcircuits can be interpreted in the framework of predictive coding.

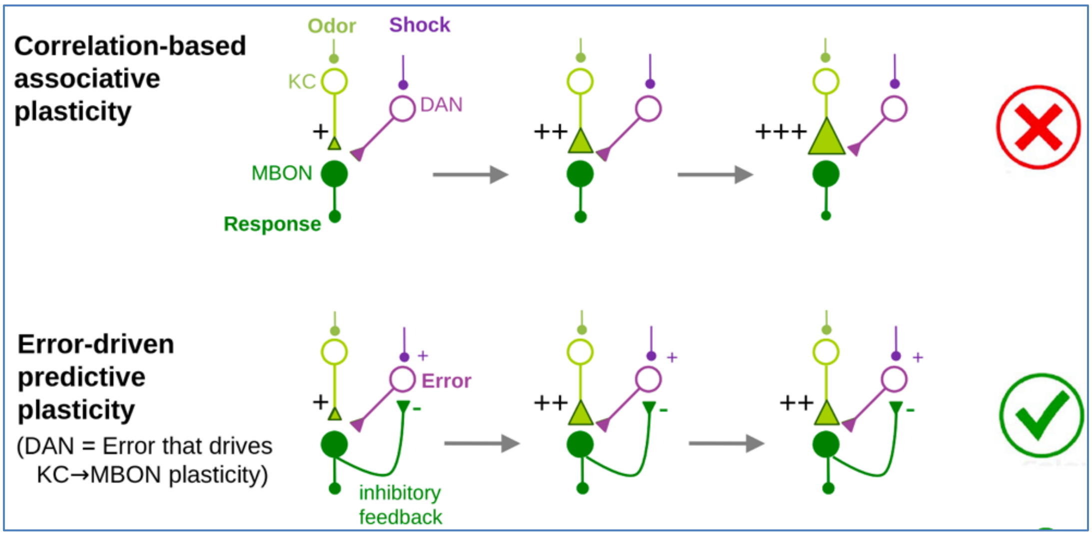

Figure 6: Associative versus error- or target-driven predictive plasticity.

Adapted from Zhao et al., 2021 – P2799. (Top) Pairing of an odour (CS) with a shock (US) is typically thought to induce correlation-based synaptic strengthening of the synapses mediating the conditioned response, here from the Kenyon cells (KCs) to the mushroom body output neurons (MBONs). Repeated pairing always leads to a stronger association strength. (Bottom) In predictive coding, plasticity stops when the strength of the US is correctly predicted by the CS. Plasticity of the KC-to-MBON synapses can be driven by the prediction error 'US-CS', formed by the dopaminergic neurons (DANs) that calculate the difference between the internal shock representation and the odour-evoked prediction by the MBONs.

It has been previously shown that the conditioning of an odour-evoked response by an electrical shock in fruit flies depends on the connections from Kenyon cells (KC) to mushroom body output neurons (MBONs). Although on the behavioural level, conditioning is recognised to be predictive, it remains unclear how MBONs form predictions of aversive or appetitive values (valences) of odours on the circuit level. We presented behavioural experiments that are not well explained by associative plasticity between conditioned and unconditioned stimuli, and we suggested an alternative model for how predictions can be formed (Figure 6). In error-driven predictive plasticity, dopaminergic neurons (DANs) represent the error between the predictive odour value and the shock strength. Predictive plasticity in KC-to-MBON synapses can also explain trace conditioning, the valence-dependent sign switch in plasticity, and the observed novelty-familiarity representation. The model offers a framework to dissect MBON circuits and interpret DAN activity during olfactory learning. It has been published in Nature Scientific Reports (Zhao et al., 2021 – P2799).

1.1.1.6 Natural gradient descent for complex spiking neurons

Contributing Partners: UBERN (P71), UHEI (P47)

In computational models of cortical activity, neuronal outputs are often approximated by scalar rates in order to simplify the associated mathematics. However, the spiking nature of neuronal outputs in vivo gives rise to unique properties that the brain can exploit for functional purposes. A relatively simple, but experimentally well-substantiated idea is to model neuronal outputs as Poisson spike trains, with an average rate equal to the scalar used in the simpler models. This, however, repositions the problem of synaptic plasticity in a probabilistic context.

The solution to credit assignment is often formulated as gradient descent on some cost function that measures the functional performance of a network. However, in a biologically plausible setting, the gradient with respect to synaptic weights is not invariant to various contingencies, such as the dendritic location of a synapse or the variance of its input. An elegant solution is offered by the natural (as opposed Euclidean) gradient, which inherently operates in the space of neuronal output probability distributions rather than the synaptic weights themselves. The plasticity rules predicted by natural gradient descent then also become parametrisation-invariant and automatically change synaptic weights in a way that optimally matches the desired change in neuronal outputs.

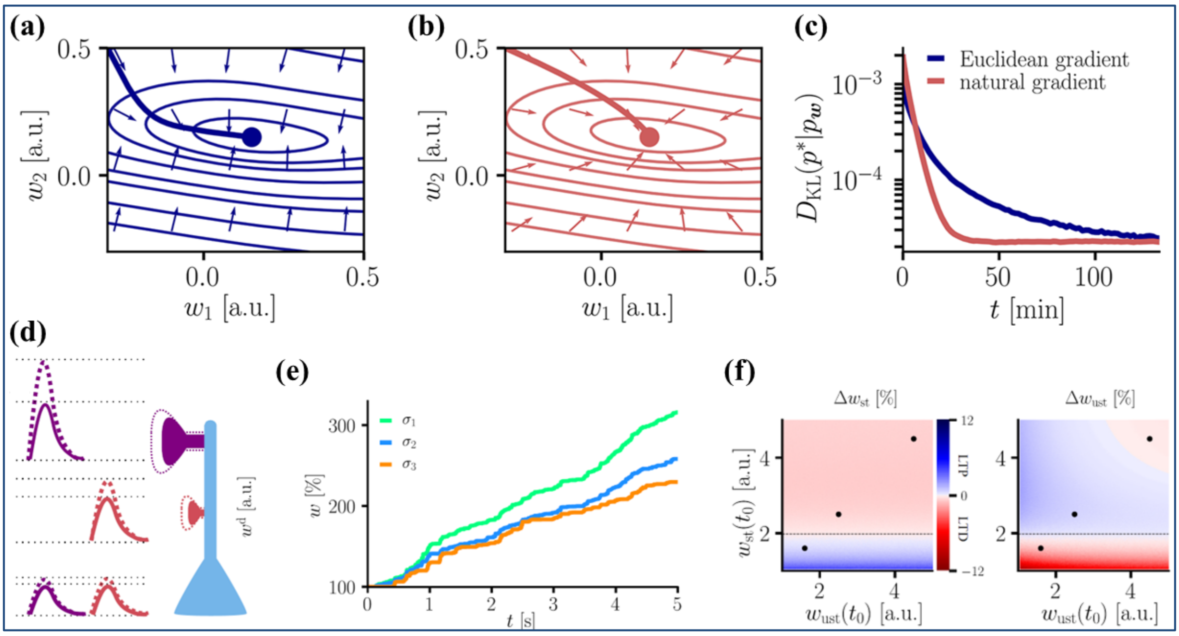

Figure 7: Synaptic plasticity as natural gradient descent.

Adapted from Kreutzer et al., 2022 - P2621. By making the cost landscape more isotropic (compare a to b), plasticity based on the natural, as opposed to the Euclidean gradient, can help synapses converge much faster towards their target values (c). This learning rule also predicts dendritic democracy (d), a scaling of the learning rate with input uncertainty (e) and a largely opposing effect of homo- and heterosynaptic plasticity (f).

We have derived such a plasticity rule for spiking neurons and demonstrated its functional advantages compared to classical gradient methods (see Figure 7a, b, c). This rule modifies classical error-correcting plasticity (ref US, LE, WNA, UPE, GeNN MCs/LE) to include terms that depend on neuronal morphology, presynaptic input statistics and heterosynaptic activity (inputs arriving from other synapses). This, in turn, provides direct links to experimental observations of synaptic plasticity that support our predictions (see Figure 7d, e, f). This work was published in eLife (Kreutzer et al., 2022 - P2621).

1.1.1.7 An evolutionary framework for interpretable plasticity

Contributing Partners: UBERN (P71)

Understanding how our brain makes sense of the outside world, how we learn from experiences and how we adapt to our environment are some of the key challenges in neuroscience. Changes in coupling strength between neurons in our brain- i.e., synaptic plasticity - are believed to be a key enabler of adaptive behaviour. The microscopic biochemical mechanisms of these changes are highly complex and mechanistic models are therefore challenging to connect to systems behaviour. One of the goals in theoretical neuroscience is thus to build phenomenological models of how these changes are orchestrated. So far, such models of synaptic plasticity, also often referred to as “learning rules”, have been developed to match either experimental data or derived mathematically from optimisation principles. However, both approaches are challenging and time-consuming.

We automate the search for new models of synaptic plasticity by leveraging genetic programming to discover free-form, interpretable learning rules. We have developed a software[1] for Cartesian Genetic Programming (CGP) in pure Python which evolves synaptic plasticity rules as symbolic expressions (Schmidt et al., 2020). Using this implementation, we demonstrated how we can discover plasticity rules for typical learning scenarios for networks of spiking and rate-based neurons. Our approach automatically (re)discovered previously hand-derived rules, as well as novel variants that improve upon these baselines and tailor plasticity to specific stimulus statistics.

---

[1] CGP software: https://github.com/Happy-Algorithms-League/hal-cgp/tree/0.3.0

Every experiment consists of a selected task family, a pre-defined network architecture, and a fitness function assessing the network's performance (Figure 8a). Crucially, this fitness depends on the plasticity rule currently being evaluated. In Mettler et al., 2021 - P2795, we sought rules that learn to map the weight vector of the network onto the first principal component of the distribution of the input data (principal component analysis). For different initialisations we find rules that perform on par or slightly better than a manually derived baseline rule (Oja's rule, see Figure 8b). When evolving rules only on datasets with specific statistics we discover rules that tailor well to the input statistics and thereby improve upon general solutions. A core strength of our approach is the availability of our rules as interpretable symbolic expressions. We can thus analytically study the weight dynamics generated by specific rules. For example, when inputs are typically positively correlated, we discover plasticity rules, which quickly converge to weights of similar magnitude, corresponding to the direction of the first principal component of the respective input distribution.

This work was published in eLife (Jordan et al, 2020 – P2556) and was featured at multiple conferences, including COSYNE, ACAIN, and GECCO (Mettler et al., 2021 - P2795). For a description of this model's deployment on the HBP/EBRAINS platforms, we refer to Sections 1.2.1.5 and 1.2.3.2.

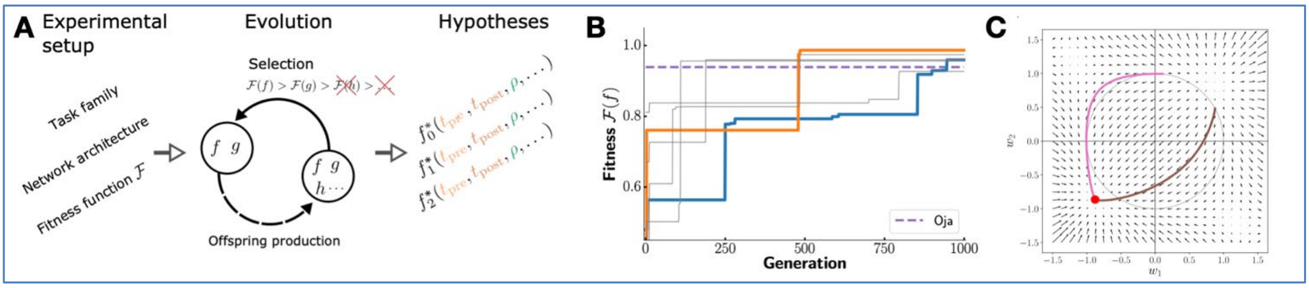

Figure 8: Evolutionary search finds interpretable learning rules.

Adapted from Jordan et al., 2020 – P2556 and Mettler et al., 2021 – P2795. (A) Setup of the automated search for learning rules. For a given experimental setup, the evolutionary algorithm generates learning rule candidates ("hypotheses"). (B) Evolutionary search finds learning rules that outperform a known manually derived baseline. Solid lines: evolved rules, dashed line: Oja's rule. (C) Weights dynamics for a learning rule evolved on specific input distributions. In the evolution of this rule, inputs were positively correlated, corresponding to a weight vector with equal magnitude for each input, or, equivalently, to a diagonal principal component.

1.1.2 Network models

1.1.2.1 The Yin-Yang dataset for challenging prototyping scenarios

Contributing Partners: UBERN (P71), UHEI (P47)

We have developed the Yin-Yang dataset for research on biologically plausible error backpropagation and deep learning in spiking neural networks. It is designed to be an alternative to classic deep learning datasets, especially in early-stage prototyping scenarios for both network models and hardware platforms. The Yin-Yang dataset provides several advantages (see Figure 9): it is smaller and therefore faster to learn, which makes it better suited for small-scale exploratory studies in both software simulations and hardware prototypes. Furthermore, it exhibits a clear gap between the accuracies achievable using shallow as compared to deep neural networks. Finally, it can be easily transferred between spatial and temporal input domains, thus ideal for multiple types of classification scenarios.

We used the EBRAINS CSCS service to develop the dataset and determine reference baseline performance values. The dataset was published in the ACM digital library (Kriener et al., 2022 - P2798), which was featured at the NICE conference and is available in its proceedings. It has already been used for a number of new developments in the field – see, e.g., Wunderlich et al., 2021; Göltz et al., 2021 – P2917 (see also Sections 1.1.2.2 and 1.2.1, as well as D3.10/T3.3); Müller et al., 2022; Max et al., 2022 – P3782 (see also Sections 1.1.2.2 and 1.2.2.1); Ma et al., 2023; Lee et al., 2023; Spilger et al. 2023.

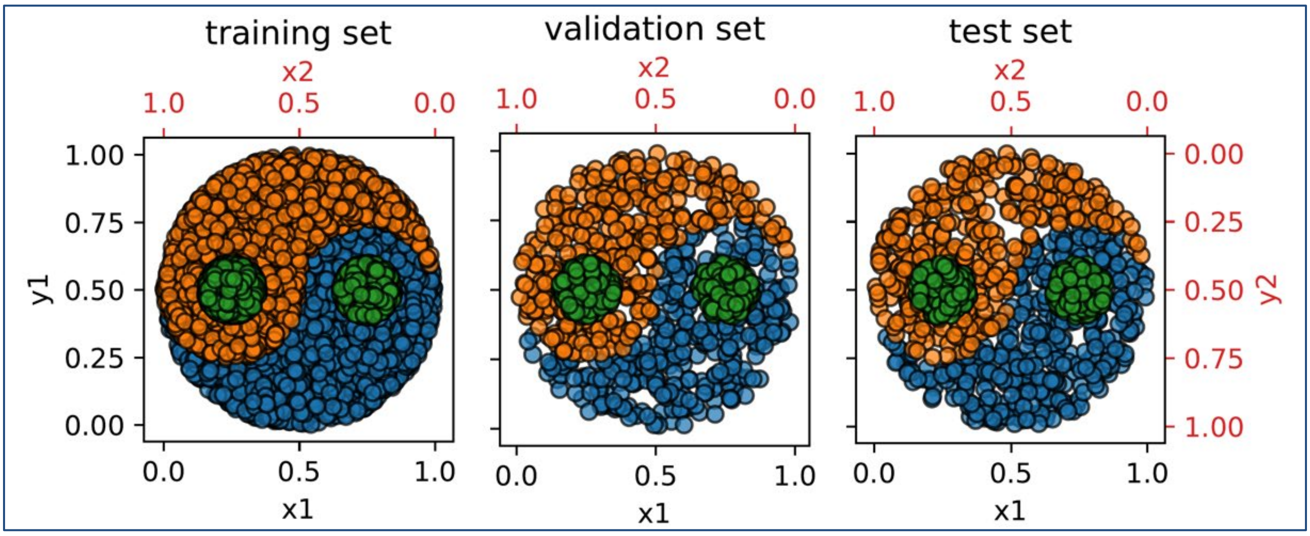

Figure 9: Training, validation, and test samples of the Yin-Yang dataset.

Figure taken from Kriener et al., 2022 - P2798. The data points of the training, validation, and test set are shown. The class distinction is highlighted by the colour of the dots. The shape resembles a Yin-Yang symbol. The dataset is publicly available on GitHub[2] (and soon on EBRAINS Knowledge graph).

1.1.2.2 Deep learning with first-spike times

Contributing Partners: UBERN (P71), UHEI (P47)

In biological substrates and also neuromorphic systems, the fundamental compute units are neurons that communicate via short electrical impulses called action potentials or spikes. Learning in networks of spiking neurons must explicitly take their temporal dynamics into account, and this has long relied on simple approximations. In contrast, by closely studying the dynamics of a particular model of spiking neurons (the LIF model), we have derived an exact relation between the precise timing of input and output spikes. This relation allows the optimisation of a loss function to satisfy a chosen goal in an exact manner. More precisely, we have shown how to precisely calculate the synaptic weight change required for a particular shift of neuronal spike times in arbitrary networks, including, in particular, deep hierarchical architectures (see Figure 10).

In practice, this is used for machine learning tasks such as image classification by providing the network with a set of spike-timing-encoded training samples and iteratively updating the connection weights to minimise the loss. For this, we have investigated multiple datasets, most importantly the Yin-Yang dataset (see Section 1.1.2.1) as well as the MNIST handwritten digit dataset. Since the application on a neuromorphic system was a core motivation for this work, a lot of care went into studying the robustness to typical imperfections found on neuromorphic substrates. While the derivation of the theory relies on mathematically exact parameters, the training mechanism is robust to imperfections, showing tolerance to levels far greater than seen on state-of-the art hardware.

This work is published in Nature Machine Intelligence (Göltz et al., 2021 - P2917) and was featured with invited talks, at multiple conferences and workshops (ICANN, CCNW 2022; ACAIN, NeuroFrance, NICE 2021; NAISys, ISCAS, SNUFA, COSYNE 2020). We have used compute time of the EBRAINS CSCS service to evaluate the theory on different tasks and investigate its robustness to substrate imperfections such as heterogeneous neuron parameters. For a description of this model's deployment on the HBP/EBRAINS platforms, we refer to Section 1.2.1.2.

---

[2] Yin-Yang dataset: https://github.com/lkriener/yin_yang_data_set

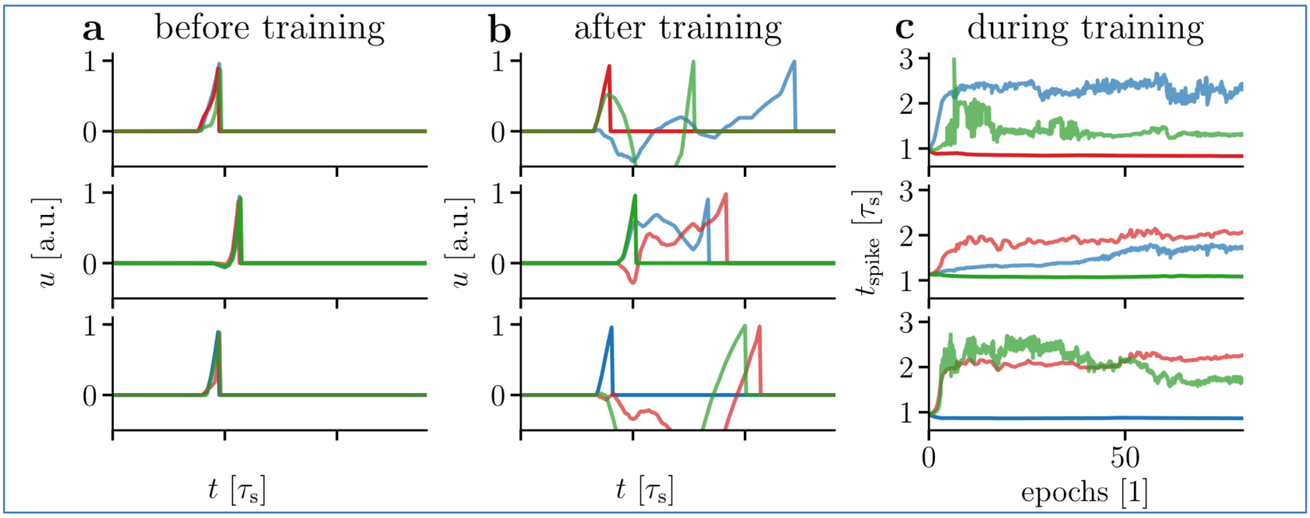

Figure 10: Dynamics in a spiking network before and after training on a classification task.

Figure adapted from Göltz et al., 2021 - P2917. (a,b) Voltage dynamics for three input samples before and after training. Before training, the responses of different neurons (shown by the three colours) to the samples cannot be distinguished. However, during training the weights are adapted such that the response of different neurons is clearly different, with the neuron coding for the class of the respective sample distinctly spiking first. (c) Evolution of the spike times during training. The spike time of the correct class (red, green, and blue for top, middle, and bottom row) is pushed to earlier times, while the other spike times are pushed to later times.

1.1.2.3 Latent Equilibrium

Contributing Partners: UBERN (P71)

Both animals and humans rely on fast information processing across the cortical hierarchy to act successfully in quickly changing environments. Simultaneously, in order to also learn from this experience, they need to solve the credit assignment problem and determine which synapses contributed to a successful outcome, continuously and in real-time. While there exists an elegant solution to this problem in the case of artificial neural networks, namely the error backpropagation (BP) algorithm, it is still an open question whether our brains use a similar algorithm.

While this question has led to the formulation of many approximate models of BP with local learning rules, most of them still suffer from what we refer to as the "relaxation problem" (Figure 11a), which arises because of the finite response times of computational elements: biological neurons need time to integrate their input with the characteristic time scale given by their membrane time constant and this property introduces a response lag which in consequence requires some sort of relaxation. This does not only delay information processing but in addition causes a timing mismatch between outputs and teaching signals that can even disrupt learning. Solutions so far circumvented this problem either by making use of very small learning rates in combination with unfeasible long training times or by turning off plasticity during the relaxation phase.

To address this problem, we have developed Latent Equilibrium (LE), a novel framework for fast computation and learning in networks of slow components that effectively decouples neuron and synapse dynamics. A core concept of our framework is prospective coding: by knowing how their current membrane potential changes, neurons can anticipate their future output. This allows neurons to overcome their intrinsic delays and enables, in principle, arbitrarily fast propagation of information (Figure 11b). The resulting model can be interpreted as a real-time, biologically plausible approximation of error backpropagation in deep cortical networks with continuous-time, leaky neuronal dynamics, and continuously active, local synaptic plasticity. Furthermore, the theory makes multiple specific predictions about neuronal morphology and cortical microcircuits (Figure 11c). In terms of performance, this model reproduces the state of the art on typical ML benchmarks such as the MNIST and CIFAR-10 datasets (see Section 1.2.3.1 for the large-scale deployment).

This work was published at the 35th NeurIPS conference 2021 (P2949) and also presented at multiple other conferences and workshops (including COSYNE, CNS, Bernstein, NAISys 2022 and CCNW 2023). A patent[3] application has been submitted in anticipation of deployment in novel edge devices.

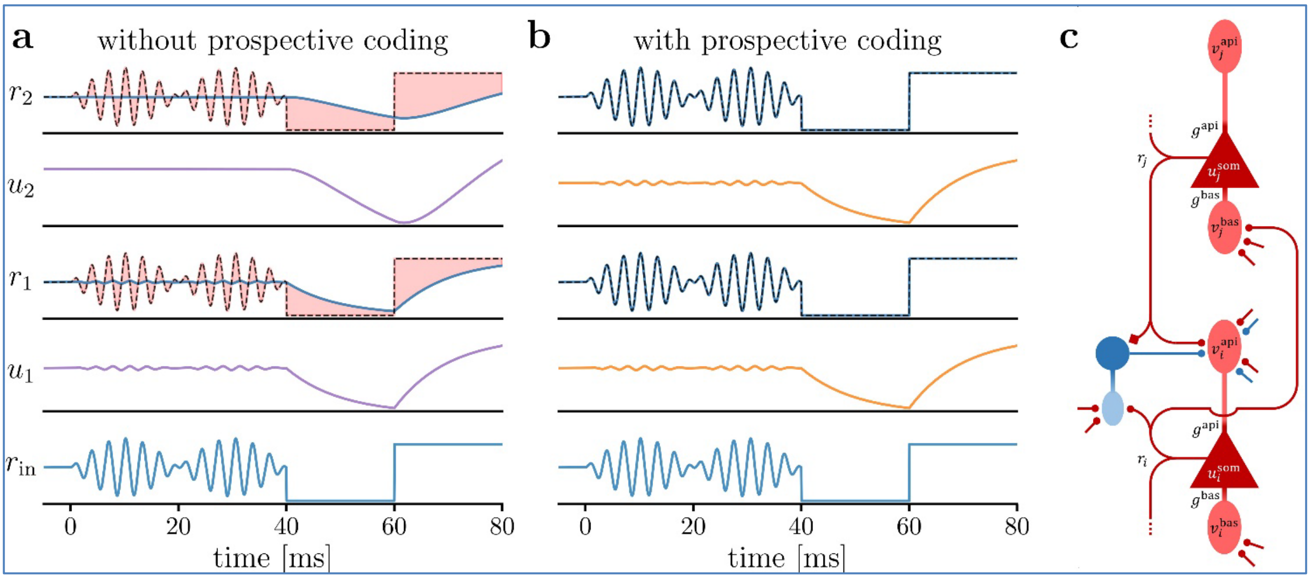

Figure 11: Prospectivity enables quasi-instantaneous computation and prevents unlearning.

Adapted from Haider et al., 2021 – P2949. (a) Firing rates r (blue) and membrane potentials u (purple) for a quickly changing input stimuli. The red shading highlights the mismatch between instantaneous and actual firing rates. (b) Incorporating prospective coding within our LE framework leads to firing rates (blue) that react as quickly as original input stimulus while membrane voltages (yellow) are still slow, as observed in biology. (c) Resulting microcircuit model with pyramidal cells (red) and interneurons (blue).

1.1.2.4 Phaseless Alignment Learning

Contributing Partners: UBERN (P71), UoS (P106)

Neural activity is optimised through long-term adaptation of synapses. However, it is still unclear how weights are adapted across the cortex to competently solve a task. One particular question is how to assign credit to synapses deep within a hierarchical network. The gold standard in machine learning is backpropagation (BP), which presents an interesting inspiration for neuroscience.

Many attempts have been made to reconcile BP with the properties and limitations of biophysical neurons. One of the main obstacles is how error signals are transported across cortical hierarchies: this implies the weight transport problem, where distant synapses in different cortical areas are required to have equal synaptic weights to enable efficient learning using a given error signal.

We have developed Phaseless Alignment Learning (PAL), an algorithm that is able to adapt weights to solve weight transport. PAL uses a signal, which is ubiquitous in cortex: noise. By modelling noise at every neuron, we do not only increase the bio-realism of our model, but also gain a type of additional information carrier. The high-frequency noise travels through the network, gathering information about bottom-up (to higher cortical areas) and top-down (from higher areas) weights. At a given top-down synapse, we extract the relevant signal using simple filtering, which we argue can easily be implemented by biological neurons. PAL can learn all error-transporting weights in a hierarchical network simultaneously and without phases, for which it also relies on Latent Equilibrium (Section 1.1.2.3).

We have evaluated PAL by integrating it into a fully dynamic, dendritic microcircuit model. The full model was tested in several small-scale setups, with and without teaching signals. During this evaluation, features and parameters of PAL as well as of the microcircuit model were developed

---

[3] Patent Signal processing method in a neural network: https://patents.google.com/patent/WO2023046504A1

iteratively. In these tests, PAL was shown to be a stable algorithm, providing quick convergence of error-transporting weights, without disrupting the learning of other weights (see Figure 12).

The simulation setups are described in a publication (Max et al., 2022 - P3782), which is currently under review and has been presented in several workshops and conferences such as CCNW, CNS and ICANN 2023.

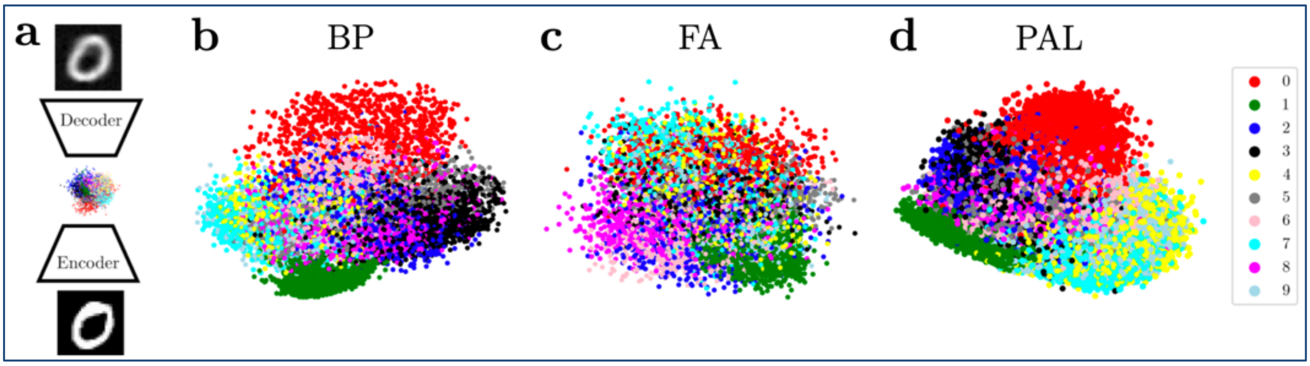

Figure 12: PAL learns useful latent representations, where feedback alignment fails to do so.

Figure taken from Max et al., 2022 - P3782. (a) Images are fed into an encoder network and a stacked decoder network aims to reproduce the original input. (b-d) Latent activations after training; colours encode the image class. Backpropagation and PAL show vastly improved feature separation compared to feedback alignment (FA). The simulation setups are publicly available on GitHub[4] (and soon on EBRAINS Knowledge Graph).

1.1.2.5 STDP enables bio-plausible weight transport

Contributing Partners: UBERN (P71), UHEI (P47)

To address the aforementioned weight transport problem in spiking neural networks, we developed a symmetrisation algorithm for reciprocal synapses purely based on the well-known STDP-algorithm. It finds its application amongst others in Bayesian inference with spiking sampling networks that fundamentally rely on a fully symmetric connectivity scheme. Similarly, spiking implementations of backpropagation in models of hierarchical cortical microcircuits benefit from aligning the top-down connections with their bottom-up counterparts.

Any physical neuronal system, be it the biological brain or an analogue neuromorphic hardware model thereof, exhibits substantial parameter variabilities across neurons and synapses. Therefore, establishing these symmetric in physical neuronal systems and maintaining it throughout the training requires an elaborate homeostatic mechanism.

Similar to PAL (see Section 1.1.2.2), our novel STDP-based symmetrisation algorithm exploits the presence of noise in cortex. The stochastic firing patterns of reciprocally connected neurons contain pre-post and post-pre correlations that represent a measure of the asymmetry of the involved weights. This observed asymmetry in the spike-timing statistics can then be used to even out the effects of the neurons on their post-synaptic partner by employing an anti-Hebbian STDP kernel.

We demonstrate the effectiveness of algorithm in a model of neural sampling, where an appropriate combination of Hebbian and anti-Hebbian STDP can significantly improve convergence to the target distribution as compared to Hebbian STDP alone when initialising the weights heterogeneously (Figure 13). This could explain how the brain is able to deal with synaptic heterogeneity, while providing a blueprint for neuromorphic substrates to do the same.

---

[4] Code repository for PAL: https://github.com/kma-code/Phaseless-Alignment-Learning

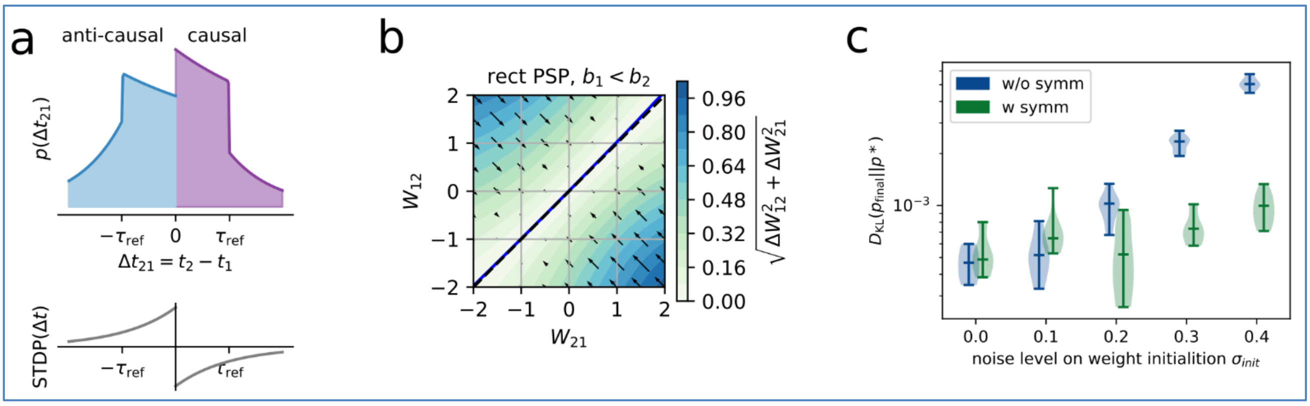

Figure 13: Fully local symmetrisation of reciprocal weights with STDP.

(a) Asymmetric connections lead to skewed patterns in the spike timing distribution (upper panel), which is exploited by the anti-Hebbian STDP kernel (lower panel). (b) Phase plane diagram of the weight updates as a function of weights for a fixed set of biases. The stable manifold (blue) coincides with symmetric weights. (c) Weight symmetrisation improves sampling accuracy: Kullback-Leibler-Divergence between the learned and the target distribution of a spiking sampling network with (green) and without (blue) symmetrisation.

1.1.2.6 Creative dreaming improves memory semantisation

Contributing Partners: UBERN (P71)

Throughout our lives, humans and other animals autonomously learn to recognise and respond to fuzzy object categories, for example tame vs. dangerous animals. This ability is believed to rely on cortical representations, i.e., neuronal activity patterns that are organised according to the semantics of the animal's environment. So far, the underlying mechanisms of this cortical organisation remain unclear. Previous models have attempted to explain semantic organisation through supervised learning, requiring millions of labelled examples. However, humans and animals only learn through sparse supervision. Brains should thus rather be considered self-supervised learning machines.

To characterise such an autonomous learning process, we introduced a cortical model with feedforward (FF) and feedback (FB) pathways, which supports learning of latent representations in higher cortical areas (Figure 14). The system can be either in a wake state, where external stimuli are perceived through FF pathways, or in a dream state, where sensory activity is driven by memory replay and FB pathways and the system is decoupled from its environment. During wakefulness (Figure 14), the cortex is exposed to diverse natural images that are processed along the cortical hierarchy through FF pathways until forming high-level neuronal representations, for example in the Inferior-Temporal (IT) cortex. During NREM sleep, a latent representation is replayed from the hippocampus, generating a NREM dream, re-enacting experiences from the previous day. This replayed dream is combined with additional perturbations, and higher cortical areas learn to map this perturbed dream to its original, replayed representation. Finally, during REM sleep, the representation stored from the previous day is replayed along with past, typically unrelated memories and noisy cortical background activity. In this state, the goal of FB pathways is to generate realistic (externally driven) low-level activity from it, by implementing an adversarial game against FF pathways that learn to classify this activity as unrealistic (internally driven).

We trained the model on benchmark datasets of natural images (CIFAR10, SVHN) and evaluated the quality of the learned representations. Our results suggest that REM dreams, generated through FB pathways via combinations of unrelated memories, become more realistic over learning, but remain different and novel compared to external sensory inputs and NREM dreams. These features are in line with well-studied dream phenomenology (Nir and Tononi, 2010; Scarpelli et al., 2019). Crucially, generating these virtual experiences through both memory combinations and adversarial learning improves the quality of the learned cortical representations, i.e., differences in neuronal responses are primarily driven by the semantic content of sensory inputs. NREM sleep, in contrast, serves a complementary function, by improving the robustness of neuronal responses to sensory perturbations.

Our model makes experimental predictions both at the neuronal and behavioural level, that can be tested in future neuroscientific studies. Moreover, it brings a new computational perspective of the brain that can learn in an unsupervised, autonomous manner, facilitated by dreams. Finally, it suggests a cortical implementation of adversarial learning, as the essence of creative dreams during REM sleep.

This work was published in eLife (Deperrois et al., 2022 – P2911) and featured at several important conferences in the field, including Bernstein and COSYNE 2021.

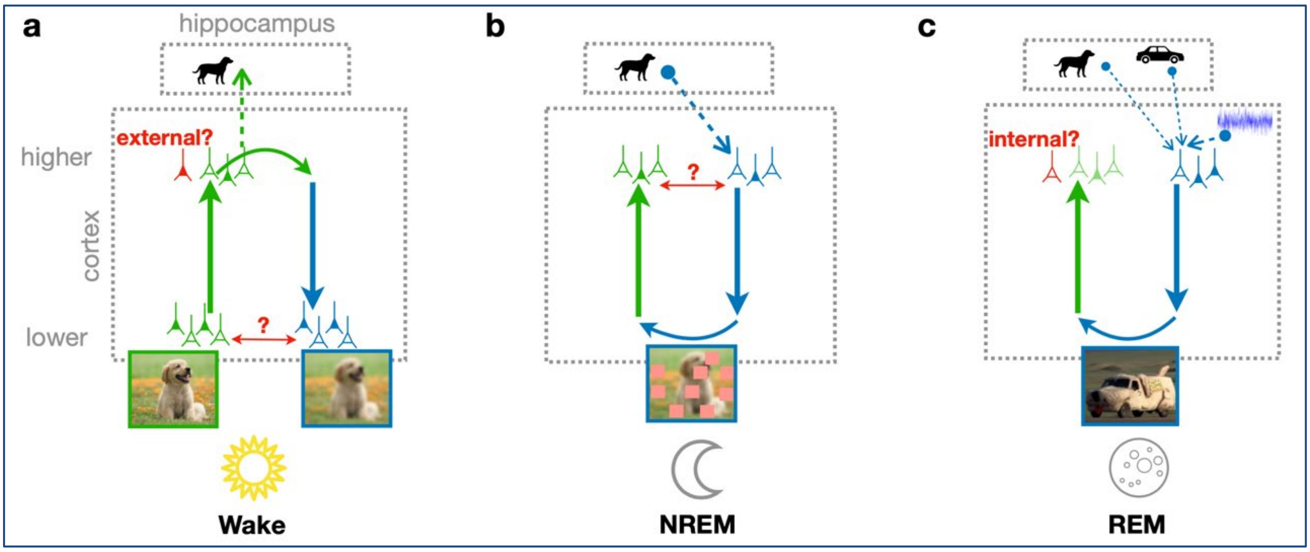

Figure 14: Cortical representation learning through Wake, NREM and REM sleep.

Figure taken from Deperrois et al., 2022 - P2911. (a) During wakefulness (Wake), cortical feedforward pathways learn to recognise that low-level activity is externally driven, and feedback pathways learn to reconstruct it from high-level neuronal representations. These high-level representations are stored in the hippocampus. (b) During non-rapid eye movement sleep (NREM), feedforward pathways learn to reconstruct high-level activity patterns replayed from the hippocampus affected by low-level perturbations, referred to as perturbed dreaming. (c) During rapid eye movement sleep (REM), feedforward and feedback pathways operate in an adversarial fashion, referred to as adversarial dreaming. Feedback pathways generate virtual low-level activity from combinations of multiple hippocampal memories and spontaneous cortical activity. While feedforward pathways learn to recognise low-level activity patterns as internally generated, feedback pathways learn to fool feedforward pathways.

1.1.2.7 Neuronal Least Action principle and cortical computation

Contributing Partners: UBERN (P71), UHEI (P47) (external collaborators, ETHZ and MILA, U Montreal)

One of the most fundamental laws of physics is the principle of least action. Motivated by its predictive power, we introduced a neuronal least-action (NLA) principle, saying that cortical activity is shaped by minimising the action of neurons and networks to produce a desired behaviour. The central notion of the principle is that of an error. It starts with a behavioural error produced by a neuronal network and shows that each neuron within the network did contribute with its own error to the output. The error within an individual neuron tells how much it could have changed its own activity to improve the behavioural outcome.

Apical dendrites of pyramidal neurons receive feedback information from downstream neurons closer to the behavioural output. Apical inputs represent neuron-specific errors that adapt the activity output to improve behaviour. The apical error also adapts the synaptic strengths of the driving input in the basal or somatic region of the dendritic tree. The basal synapses can read out the back propagated error in terms of a somato-dendritic mismatch error. This error adapts the synaptic strength to explain the mismatch away, and to eventually reduce the behavioural error (Figure 15).

We introduced the mathematics of the NLA and showed that it allows for predicting the network output and adapt the activities right in time. We showed how this principle can be applied for a motor task. We introduced the 'moving equilibrium hypothesis', dictating how deep neurons must instantaneously change to keep the trajectory of a movement on the desired track. The principle relates to the Latent Equilibrium (see Section 1.1.2.3) and formalises the classical equilibrium hypothesis in motor control.

This work was published in eLife (Senn et al., 2023 – P3937) and presented in numerous invited talks, workshops, and conferences. Given its general nature, we view this framework as a stepping-stone towards a comprehensive theory of the brain.

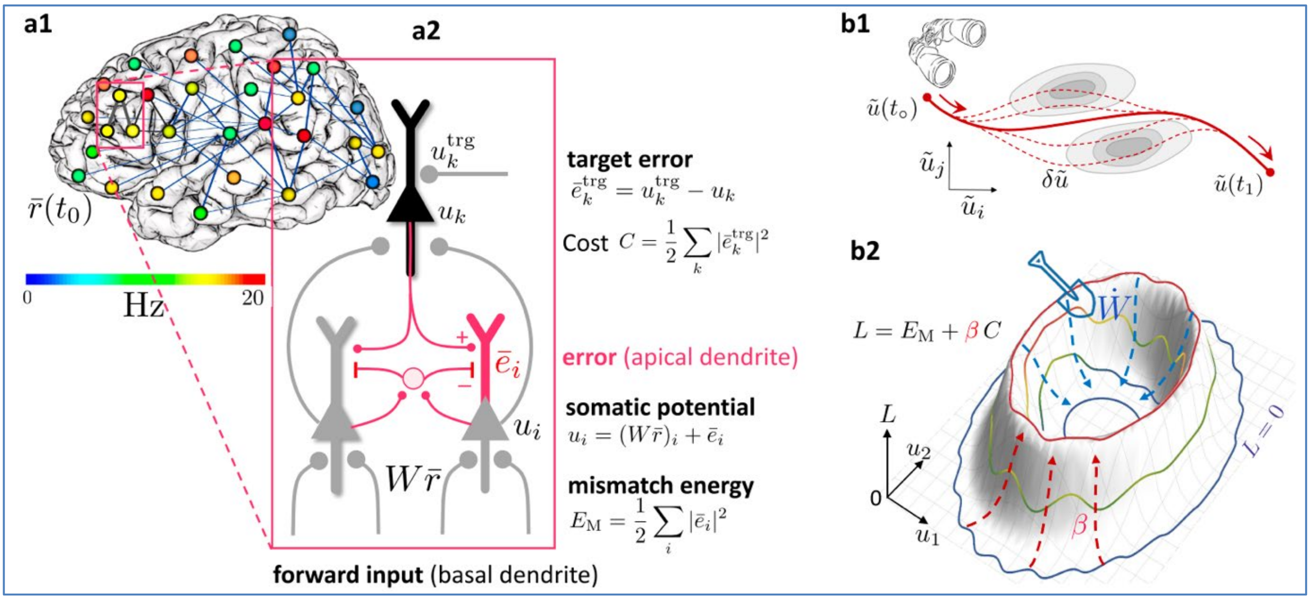

Figure 15: Somato-dendritic mismatch energies and the neuronal least-action principle.

Adapted from Senn et al., 2023 – P3937. (a1) Sketch of a cross-cortical network of pyramidal neurons described by NLA. (a2) Correspondence between elements of NLA and biological observables such as membrane voltages and synaptic weights. (b1) The NLA principle postulates that small variations *δu* (dashed) of the trajectories u (solid) leave the action, as defined by the temporal integral of a Lagrangian *L*, invariant. It is formulated in look-ahead coordinates (symbolised by the spyglass). (b2) In the absence of output nudging ($\beta = 0$), the trajectory u(t) is solely driven by the sensory input, and errors vanish (A = 0, outer blue trajectory at bottom). When nudging outputs towards a target ($\beta > 0$), somato-dendritic prediction errors appear, the energy increases (red dashed arrows) and the trajectory u(t) moves out of the *L* = 0 hyperplane, riding on top of the ‘volcano’ (red trajectory). Synaptic plasticity ($\dot{W}$) reduces the somato-dendritic mismatch along the trajectory by optimally ‘shovelling down the volcano’ (blue dashed arrows) while the trajectory settles in a new place on the *L* = 0 hyperplane (inner blue trajectory at bottom).

1.2 Large-scale deployment of the models and algorithms from Output 13 available on EBRAINS platforms; hybrid demonstrator exploiting the advantages of available EBRAINS platforms for learning-to- learn setups involving both synaptic plasticity and meta-plasticity

Note: For a more detailed description of associated publications and dissemination, we often refer to previous sections that describe the particular models.

1.2.1 BrainScaleS-2

1.2.1.1 Versatile neuromorphic emulation

Contributing Partners: UHEI (P47), UBERN (P71)

The field of neuromorphic hardware is a rapidly growing area of research. This expansion is driven on the one hand by the promise of efficiency gains by the incorporation of ideas from the brain, on the other hand by the expectation to better understand mechanisms in the brain. While this research is ongoing, the chips are often designed with some attempt at universality of their dynamics, i.e., suitability for more than just a single use case. One such neuromorphic substrate which aims for this kind of generality is the BrainScales-2 chip developed by the Electronic Vision(s) group in Heidelberg (Figure 16A). The chip is a mixed-signal neuromorphic hardware with an included plasticity-processing unit (PPU) that can be used to implement on-chip learning rules (Figure 16B). Due to the generality of the implemented neuron model, a range of different network models and spike-based algorithms have been realised on the chip.

In Billaudelle et al., 2020 (P2241), we summarised a set of five experiments that highlight the capabilities of this new class of neuromorphic hardware. One experiment implements deep learning in networks of spiking neurons, described in more detail in Section 1.2.1.2. Another experiment is closely related to structural plasticity, described in more detail in Section 1.2.1.4.

On the same chip, we have also emulated neural sampling in spiking networks. This exploits the binarity of the state of refractoriness and can be used to train networks to sample from a given arbitrary binary probability. When testing the fidelity of the sampled distribution and comparing it against the target distribution (Figure 16D) we see how the distributions become very similar. The unprecedented combination of sampling speed and precision enabled by this setup has led to a serendipitous application outside of classical computational neuroscience, which we describe in more detail in Section 1.2.1.3.

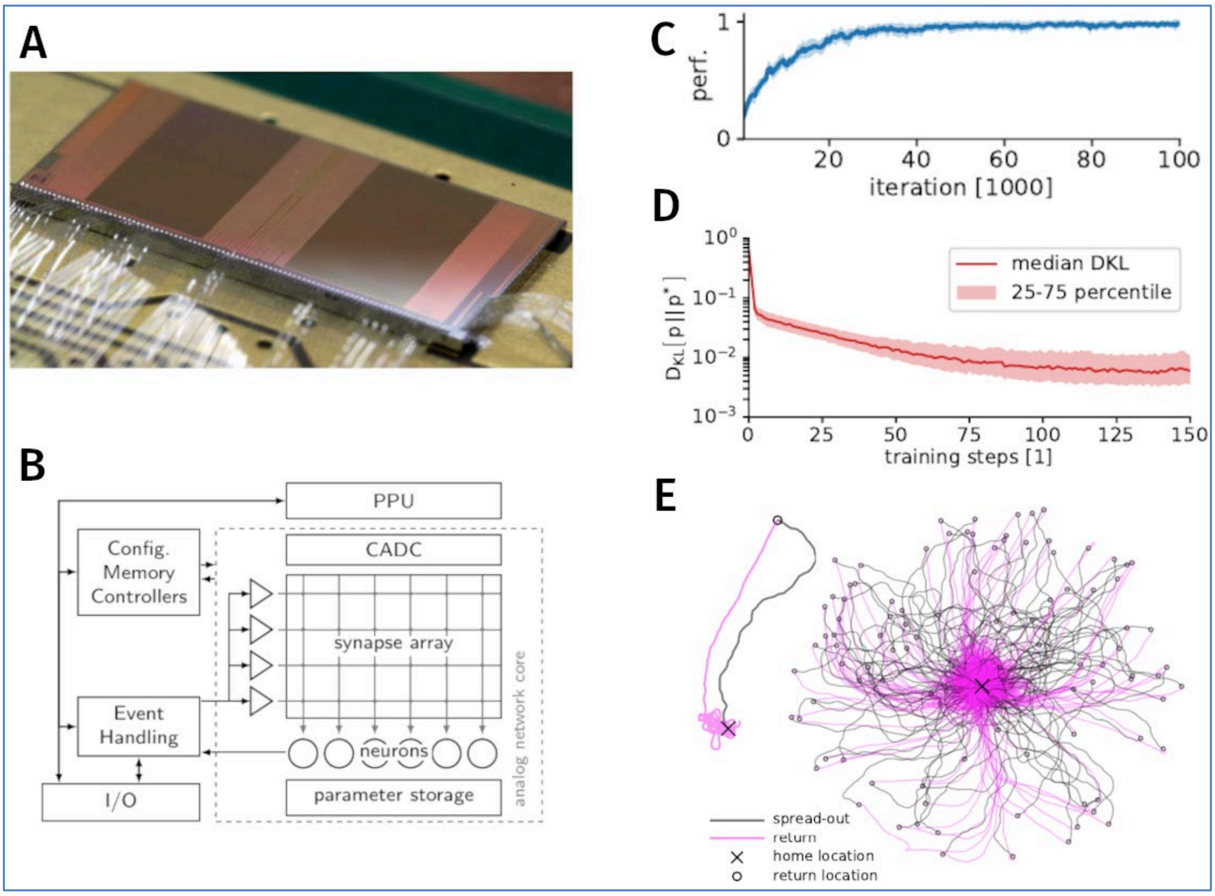

Figure 16: The BrainScaleS-2 neuromorphic hardware and its diverse applications.

Figure adapted from Billaudelle et al., 2020 - P2241. (A, B) Photograph and simplified schematic of the BrainScaleS-2 chip. (C) Reward over a training run on a reward-modulated spike-time-dependent plasticity (STDP) task. The convergence towards high performance (perfect performance corresponds to a value of 1) indicates successful learning of the task. D: Convergence of a learned probability distribution towards its target (Kullback-Leibler divergence) over time. E: A single (left) and collection of many trajectories (right) of an emulated insectoid that needs to return to the home location (x) after an explorative path (grey). These attempts to return home are shown in pink, and the relative straightness of these paths demonstrates the successful path integration performed by the system.

The PPU on the chip can be used to both implement learning rules and control the environment of experiments. In Wunderlich et al., 2019 (P1721) this was used to implement a reinforcement learning task. This task was solved with a reward-modulated STDP learning rule (Figure 16C), fully on-chip –

a pioneering achievement in the field. Environment control was demonstrated in a setup inspired by insect navigation. More specifically, we emulated the corresponding part of a bee's brain to perform path integration, allowing an agent to return to its nest after exploring for food (Figure 16E).

This large variety of experiments is highlighting the versatility of the BrainScales-2 system, one of the core computing platforms of the HBP/EBRAINS, as well as the various applications of spiking neuron models developed by several partners in the HBP, especially within WP3 during SGA3 and CDP5 during SGA2. Due to diversity of the presented results, it was prominently featured in multiple venues, starting with IEEE ISCAS 2020 and continuing with various other important conferences in the field, as highlighted in the sections describing the individual results.

1.2.1.2 Deep learning of precise spike timing

Contributing Partners: UBERN (P71), UHEI (P47)

Efficiently training networks of spiking neurons towards high functional performance on neuromorphic substrates represents a long-standing goal of the community. Our modelling efforts (see Section 1.2.1.1, 1.1.2.2, D3.10/T3.3 and Göltz et al., 2021 - P2917) have laid the theoretical groundwork for a new training method that enables error backpropagation in spiking substrates based purely on spike timing. Importantly, our method does not require the readout of other internal variables such as membrane potentials (Cramer et al., 2022 – P2756), nor does it need the separate simulation of a network copy (Esser et al., 2015). The practical realisation was done on the BrainScaleS-2 Platform (Figure 17a) in close collaboration with the hardware developers.

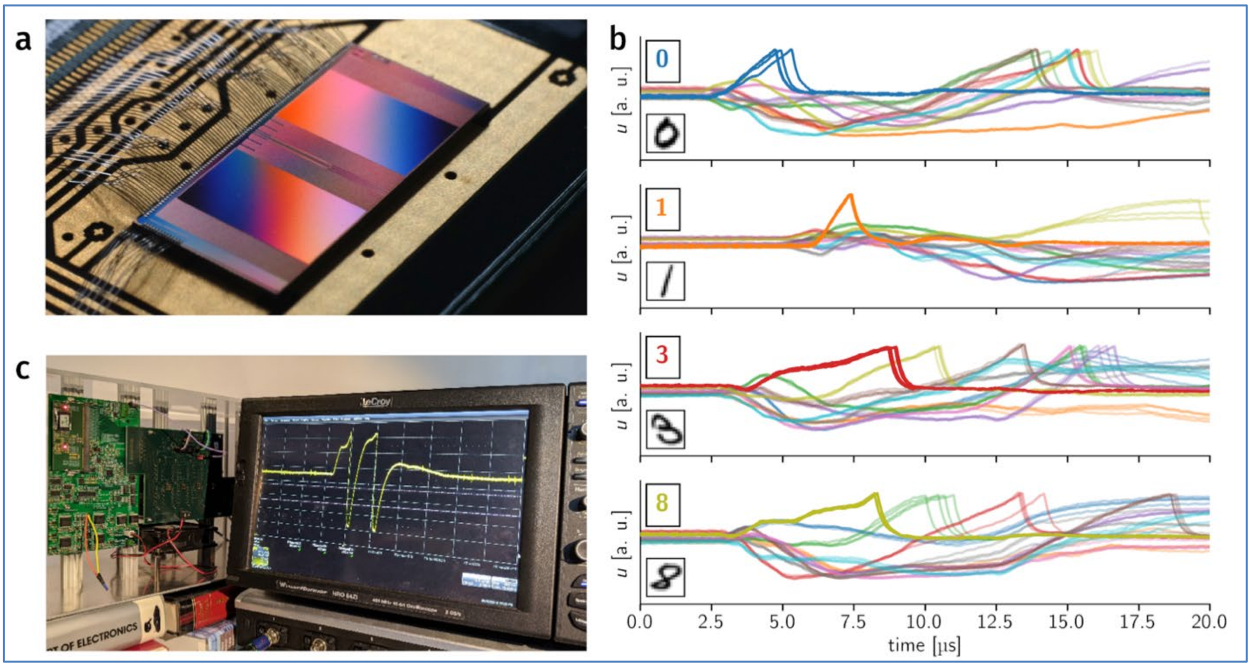

Figure 17: Neuromorphic emulation of spike-time learning on BrainScaleS-2.

Figure adapted from Göltz et al., 2021 - P2917. (a) Close up figure of the neuromorphic chip BrainScaleS-2. (b) After training, voltage traces of sample inputs are recorded. In all cases shown here, the correct neuron robustly fired first, as shown by the steep reset of the voltage curve. The same inputs were presented multiple times, and the variance in the curves is due to the analogue nature of the substrate. (c) Setup of a BrainScaleS-2 chip with supporting hardware. The membrane voltage of an analogue neuron is read out and displayed on an oscilloscope. For a compelling illustration of the chip's power consumption, the bright red lights on the left side are two control LEDs that consume roughly one fifth as much as the chip itself.

BrainScaleS-2 is a mixed-signal substrate, meaning that while the spikes are communicated digitally, the dynamics of the neurons as well as the synapses are emulated in an analogue manner. Consequently, the network will behave differently from an ideal simulation due to effects such as fixed-pattern noise on the parameters or jitter on the spike times. The successful training on the hardware (Figure 17b, see also Figure 1) exemplarily shows a clear separation of spike timing, corresponding to a correct classification for multiple classes) thus demonstrates the robustness of our approach, which we believe to be an essential ingredient for any realistic model of computation in physical spiking networks, be they biological or artificial.

Apart from the manifest robustness of our training method, a crucial result of this line of research is the efficiency of the resulting network: while the energy efficiency is often a motivation to work with neuromorphic hardware, we explicitly also measured the latency and energy consumption of our approach. The full time it took the setup to classify the full 10'000 test samples of the MNIST dataset was less than one second, including the preparation of data, communication, and evaluation. At the same time, the analogue chip only consumed a power of 175 mW, which is close to the consumption of simple light emitting diodes (LEDs). Indeed, the red control LEDs on the setup, seen in Figure 17c on the left side, initially influenced the power measurement despite serving no functional purpose. Altogether, the measured energy consumption per classification throughout a continuously running experiment comes down to less than 10 µJ.

This prototype result is on par with other state-of-the-art methods, even compared to dedicated chips and high-performance GPUs. Given that, the BrainScaleS-2 system is a general-purpose research chip and thus not optimised for any particular scenario, our results evince a clear impact on the current state of the art. This success is owed to the combination of a well-designed neuromorphic substrate, well-chosen data encoding and a rigorous and robust training algorithm – the result of a team effort that has greatly profited from the synergistically multidisciplinary approach of the HBP.

1.2.1.3 Quantum many-body states: A novel neuromorphic application

Contributing Partners: UHEI (P47), UBERN (P71)

The capability to perform probabilistic (Bayesian) inference represents an essential asset of the mammalian brain. A large body of neuroscientific work has been dedicated to this research topic, to which much of our own work in the HBP has contributed with novel insights (e.g., Petrovici et al., 2016 – P1213; Leng et al., 2018 – P1334; Kungl et al., 2019 – P1360; Dold et al., 2019 – P1447; Jordan et al., 2019 – P843; Korcsak-Gorzo et al., 2022 – P2806; see also D3.9 – cortical oscillations, as well as Sections 1.1.1.2, 1.1.1.3, 1.1.1.4, 1.1.2.5 and 1.2.1.1 above).

The ability to represent probability distributions (Figure 18a) endows spiking networks with a wide range of possible applications. In conjunction with efficient neuromorphic hardware (see also Section 1.2.1.1), these scenarios extend far beyond the field of (computational) neuroscience. In a serendipitous collaboration with quantum physicists from the University of Heidelberg, we have demonstrated how neuromorphically accelerated spiking sampling networks can efficiently solve certain classes of problems in quantum mechanics.

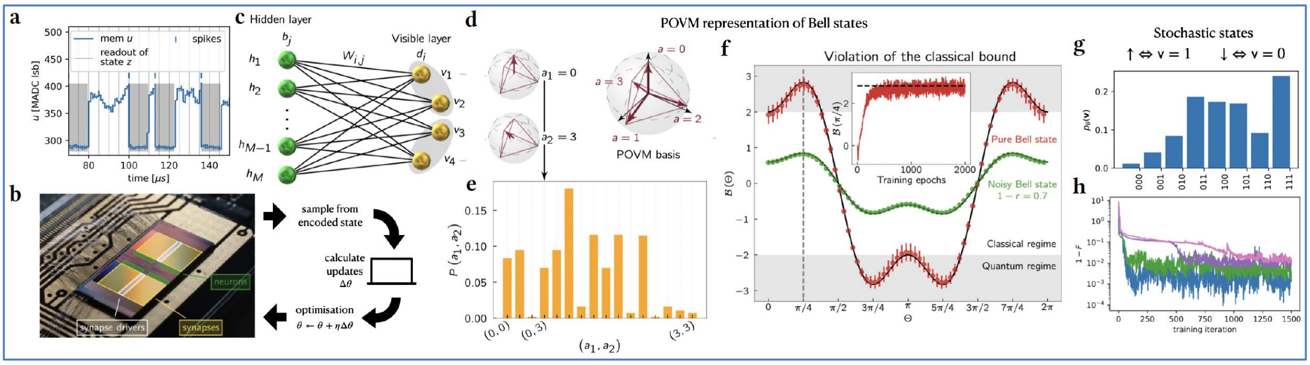

Figure 18: Neuromorphic emulation of quantum systems.

Figure adapted from Baumbach et al., 2022 – P3261. A spike-based representation of probability distributions (a) can be realised on the BrainScaleS-2 system (b) to train hierarchical sampling networks (c). Quantum spin ensembles, usually represented on a Bloch sphere (d) can be translated to a probabilistic representation (e), thereby becoming amenable to neuromorphic representation. We have demonstrated this both for the iconic Bell states (f) as well as for stochastic systems (g), achieving extremely fast times-to-solution (h) owing to the thousandfold neuromorphic acceleration factor compared to biological real-time.

Many emergent phenomena in condensed-matter physics, such as superconductivity or magnetism, are fundamentally rooted in quantum interactions. These phenomena often remain poorly understood, in part due to their mathematical complexity and the computational demands of their simulation. We were able to formulate two types of precise mathematical equivalence between such quantum systems and the probability distributions sampled by spiking networks. The first one uses a probabilistic representation of quantum density matrices, while the second one realises a direct representation of wave function coefficients for stochastic Hamiltonians.

This, in turn, allowed us to use the BrainScaleS-2 platform to implement such probabilistic representations of two particular types of quantum states (Figure 18b, c). First, we have successfully completed a hallmark test for "quantumness" by representing a quantum state that violates the classical bounds of the Bell inequality (Figure 18d-f). Second, we showed that we can represent the large class of stochastic quantum states with fidelities above 98% for moderate system sizes (Figure 18g, h). The physical nature of BrainScaleS-2, i.e., the implementation of the spiking neural system as electronic circuits, allowed extremely fast times-to-solution (on the scale of 1 second of wall-clock time), which, moreover, is independent of the size of the used/modelled system. This allows our method to scale to larger systems, thereby conclusively uncovering the emulation of quantum many-body states as a potential application for (spike-based) neuromorphic systems.

This work was presented at several conferences (SIAM, NICE, APS) and published in SciPost Physics (Czischek et al., 2022 – P2574) and iScience (Klassert et al., 2022 – P2932). It relies on a software resource developed within and curated by the HBP (Breitwieser et al., 2020 – P2517).

1.2.1.4 Structural plasticity on spiking neuromorphic hardware

Contributing Partners: UHEI (P47), UBERN (P71)

Synaptic plasticity represents the key physical process underlying learning in neural networks, whether biological or artificial. In most models, networks are pre-wired, such that the connections between neurons are fixed and only their weights are changed during learning. However, this is usually not the most efficient solution to the credit assignment problem, nor is it a faithful representation of biological learning processes. It is well known that synapses in the brain are subject to constant turnover, with many of them appearing and disappearing on a daily basis. This is also preferable from a hardware point of view, as it permits a redistribution of limited synaptic resources (bandwidth and energy) to where they are ultimately needed.

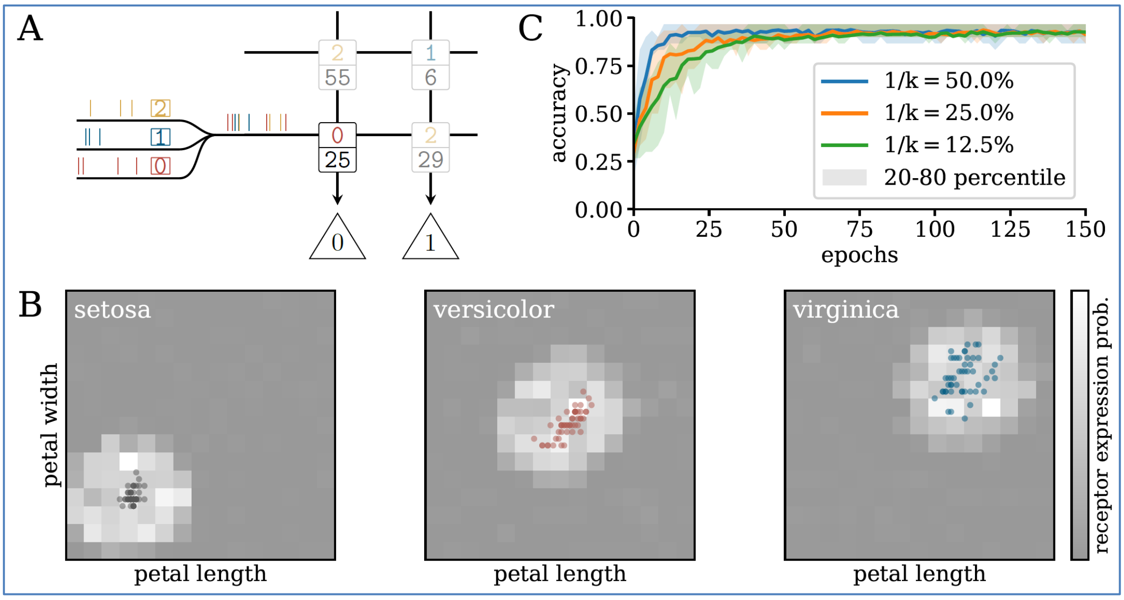

Figure 19: Self-organising receptive fields through structural plasticity on BrainScaleS-2.

Figure adapted from Billaudelle et al., 2020 – P2240. (A) The synaptic event communication infrastructure on BrainScaleS-2 allows a direct implementation of structural plasticity. (B) Despite limited resources, our structural plasticity rule allows an efficient assignment of neuronal receptive fields for data classification. (C) Reduced resources induce longer convergence times for structural plasticity, but as long as these resources suffice for solving the required task, optimal performance can ultimately be achieved.

Thus, we have proposed an efficient implementation of structural plasticity and demonstrated its functionality on a prototype of the BrainScaleS-2 system. To this end, we have exploited the address-event representation scheme and local event filtering in the synaptic array. The analog core of the platform is augmented by an embedded microprocessor incorporating an efficient vector processing unit, along with parallel access to the synapse array (Figure 19A). This plasticity-processing unit (PPU) provides the flexibility required for studying a wide range of learning algorithms and in this work serves as the substrate for structural reconfiguration.

The synaptic update policy enables neurons to dynamically select a set of suitable synapses out of a pool of potential connections. It is based on a weight update algorithm, which consists of three independent terms: An STDP term that potentiates correlated connections, a homeostatic regulariser that limits post-synaptic firing rates and encourages synaptic competition, and a stochastic component that induces exploration. The weight updates are accompanied by a pruning condition which, executed periodically, removes synapses with a weight below a certain threshold and randomly reassigns the freed resources.

We applied the above algorithm to a supervised learning task, where the network was trained to classify the Iris data set (Billaudelle et al., 2020 – P2240). The emulated plasticity rule led to self-organised reconfiguration of their receptive fields (Figure 19B), as the correlation between teacher signal and receptor proximity to the presented data drove the potentiation of associated synaptic weights. For higher degrees of enforced sparsity, convergence times were longer, as the search for relevant inputs in the feature space became statistically more challenging. Ultimately however, the learning rule enabled the network to achieve near-perfect classification in all three scenarios (Figure 19C), demonstrating its ability to ensure a better utilisation of synaptic resources without prior knowledge of the input data.

The core of our plasticity mechanism is compatible not only with the correlation-based weight dynamics used in the demonstrator task; it can instead be applied to all tasks and frameworks where a synapse's relevance is determined by its efficacy. Thus, our work provides a starting point for a multitude of resource-efficient learning algorithms on spiking neuromorphic hardware.

1.2.1.5 Evolutionary framework for interpretable plasticity

Contributing Partners: UBERN (P71), UHEI (P47)

Synaptic plasticity is believed to be a key process enabling us to understand, learn and adapt. The microscopic biochemical mechanisms of these changes are highly complex and detailed models are challenging to connect to system-level behaviour. One of the goals in theoretical neuroscience is thus to build phenomenological models of how these changes are orchestrated. So far, such models of synaptic plasticity, also termed "learning rules", have been developed to either match experimental data or implement various optimisation methods. However, both approaches are challenging and time-consuming.

We automate the search for new models of synaptic plasticity with evolutionary algorithms by leveraging genetic programming to discover free-form, interpretable learning rules (see Section 1.1.1.7). Based on this approach, we aim to implement the evaluation of candidate plasticity rules on the BrainScale-S2 (BSS2) neuromorphic system (see Figure 20a). Evolving neuronal learning rules on neuromorphic hardware offers us two major advantages: On one hand, the thousandfold acceleration of BSS2 compared to biological real-time allows us to accelerate the evaluation of learning rules and thus mitigates a crucial drawback of evolutionary algorithms, namely their relative slowness – in our case, due to the significant computational overhead of simulating each agent (spiking network) with its own synaptic learning rule. Furthermore, since the dynamics of neuron models on BSS2 mimic (some of) the dynamics of neurons in our brain, rules found through search on BSS2 should give us more insight on rules that might actually be useful for our brain.

For a prototype demonstration of this approach, we consider a "teacher-student" learning task (see Figure 20b). A stochastic spike source projects to two neurons, the student, and the teacher. The learning rules controls the dynamics of the synaptic weight to the student over a trial. Well-performing rules should enable the student to learn to mimic the teacher's membrane potential. And indeed, the rules evolved on the hardware learn to minimise the difference in membrane potentials between student and teacher and enable the plastic student weight to converge close to the fixed teacher weight over the duration of an experiment (Figure c). When looking at the evolved expressions, it becomes immediately apparent how the discovered rules are close in form to previously discovered rules derived from minimisation principles. This highlights an important aspect of our approach: the availability of the discovered rules as compact expressions allows for their mathematical analysis and hence an understanding of their underlying principles.

Building upon this work, we are currently implementing the evolutionary search for learning rules in hierarchical networks that should learn to classify the Yin-Yang dataset (see Section 1.1.2.1).

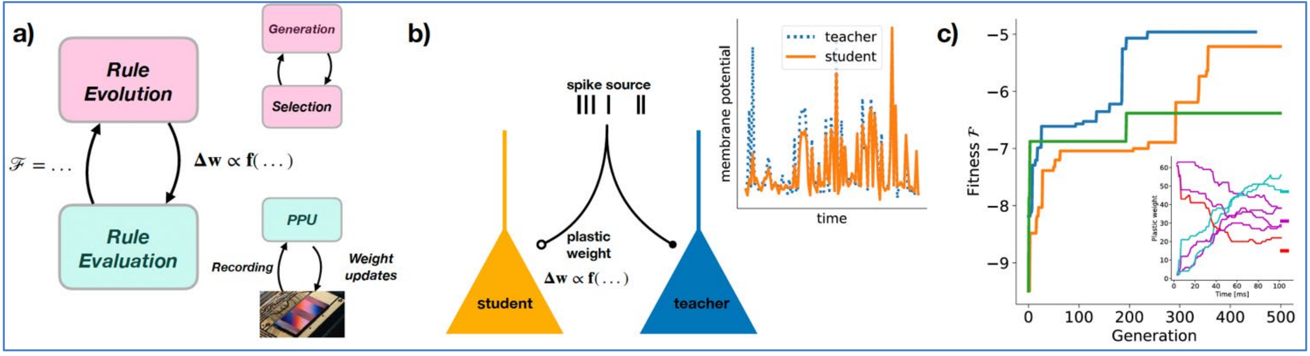

Figure 20: Evolving neuronal learning rules on neuromorphic hardware.

(a) Experimental setup: Evolutionary search is executed on a host computer and rule evaluation is performed on neuromorphic hardware, emulating the network dynamics, and sending back recorded observables. (b) Initial learning task: A spike source is connected to a student neuron via a plastic weight and to a teacher neuron via a fixed weight. Fitness over a single trial is the absolute difference between student and teacher membrane potential. (c) Fitness of the best rules over generations of the evolutionary algorithm for different initialisations. Inset: Weight traces over a trial in multiple experimental conditions.

1.2.1.6 New rules for approximating BPTT in spiking networks

Contributing Partners: HOST (P132)

Full backpropagation through time (BPTT) is not implementable on the current neuromorphic platforms (and in many respects not biologically plausible either). In order to still utilise spike timing-based encoding in input and output signals to spiking neural networks with hidden layers in supervised learning tasks, approximate rules were developed in earlier phases of the HBP (e.g., Gardner et al., 2021 – P3173, Göltz et al., 2021 – P2917, Bellec et al., 2019 – P1836, Zenke et al., 2021; see also Section 1.2.1.2 and D3.10/T3.3). Due to their simplicity, these rules have the promise of being implementable on neuromorphic platforms due to their lower demands with respect to available signals on the hardware. In preliminary simulations, we evaluated such approximate rules and full spiking BPTT on GPU servers, using frameworks for automatic differentiation (PyTorch, JaX).

Inspired by results in neuroscience (balanced networks) and deep learning (batch normalisation), we created a new class of regularisation for SNNs: regularisation of the statistical moments of the input spike trains weighted by their corresponding synaptic weights. The benefit of this approach is that there is no need to choose a "minimum" or "maximum" firing rate; one simply chooses a value for the second moment, which implicitly controls the average firing rate (thus ensuring gradient flow and effective learning). The first and third moments are kept at zero: the first to enforce balance, and the third to avoid degenerate solutions. With this approach, a variant of our proposed learning rule from Gardner et al., 2021 – P3173 was shown to be competitive with BPTT on neuromorphic MNIST, both in our own implementation and in the literature (Zenke et al., 2021).

- spike latency classifier + rate regularisation (88%, Gardner et al., 2021 – P3173)
- membrane classifier + input regularisation (97.4%, Mackwood & Grüning (in prep))
- membrane classifier + rate regularisation (97.8%, Zenke et al., 2021, BPTT)
- membrane classifier + input regularisation (98.8%, Mackwood & Grüning (in prep), BPTT)

1.2.2 GeNN

1.2.2.1 Cortical microcircuits

Contributing Partners: UoS (P106), UBERN (P71)

Inspired by some of our earlier work (Urbanczik et al., 2014), we have developed comprehensive theories for bio-inspired learning rules and circuits (Sacramento et al., 2018 – P1527; Haider et al., 2021 – P2949; see also Sections 1.1.2.3 and 1.1.2.7). In a close collaboration between UBERN and UoS, we have developed software to simulate circuits that employ these learning rules in the GPU-enhanced Neuronal Networks (GeNN) framework (Knight et al., 2021 – P2777). GeNN[5] is designed as a meta-compiler to generate efficient CUDA-C/ C++ code to run simulations of brain circuits. This allows for a high degree of flexibility where users of the GeNN framework can define their own neuron and synapse models "on the fly". Therefore, the bio-inspired microcircuits (Figure 21A) could be implemented in GeNN without substantial changes to the framework.

In most cases, GeNN is used for spiking neural network models where it offers the most benefits in terms of acceleration. In contrast, our microcircuits have developed within rate-based theories, where GeNN has traditionally had less advantage over other simulation methods. Nevertheless, we observed substantial speed-ups over the original Python-based implementations of the microcircuit models. The accelerated GeNN implementation has then enabled new research on a novel bio-plausible solution to the weight-transport problem, Phaseless Alignment Learning (PAL, see Max et al., 2022 - P3782 and see Section 1.1.2.2), which we describe below (see Section 1.2.2.2).

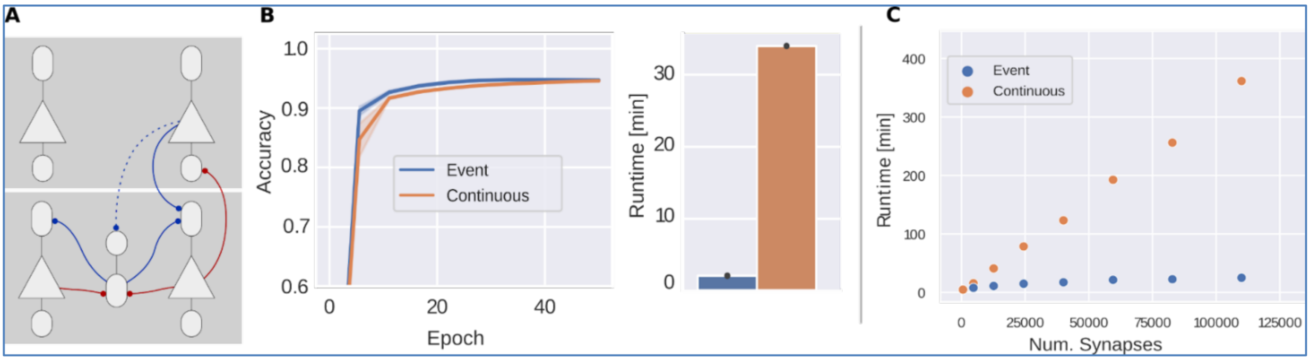

Figure 21: Cortical microcircuits in GeNN.

(A) Cortical microcircuit architecture as proposed by Sacramento et al., 2018 – P1527. (B) Test accuracy and total simulation time of a microcircuit network on the MNIST dataset. Using the event-based implementation led to an approximately 10-fold speedup. (C) Simulation time for training the RandMan classification task using varying network sizes / total number of synapses.

Finally, in order to take full advantage of the GPU acceleration of GeNN and approximate the spiking operation of biological microcircuits, we are now developing novel event-based implementations of the microcircuit and its associated learning rules. When testing on the popular MNIST benchmark for the recognition of handwritten digits, simulations with the event-based microcircuit are another order of magnitude faster than the rate-based GeNN implementation (Figure 21B) while achieving the same classification accuracy. The increased simulation speed is confirmed in scaling experiments using the so-called Random Manifolds (RandMan; Zenke et al., 2021) classification task (Figure 21C).

1.2.2.2 Phaseless Alignment Learning

Contributing Partners: UBERN (P71), UoS (P106)

Building on the results of our small-scale simulations of Phaseless Alignment Learning (PAL, see Section 1.1.2.2 and Max et al., 2022 – P3782), we expanded the scope of our experiments. We tested PAL against standard classification benchmarks (MNIST, Yin-Yang), greatly increasing the

---

[5] GeNN: https://github.com/genn-team/genn

computational demands. For this, we gratefully acknowledge the resources provided to us by the Swiss National Supercomputing Centre via the Fenix Infrastructure project.

We highlight that our results were obtained by simulating a fully dynamical, recurrent, and bio-plausible system with weight and voltage updates applied at every time step, and without any kind of steady-state approximation or biologically implausible information transfer (Figure 22). These simulations have only been made possible through the integration of PAL with microcircuits into the GPU-accelerated spiking neuron simulator GeNN (Knight et al., 2021 – P2777), and through the massive speed-up due to the use of prospective coding provided by our Latent Equilibrium framework (see Section 1.1.2.3 and Haider et al., 2021 - P2949).

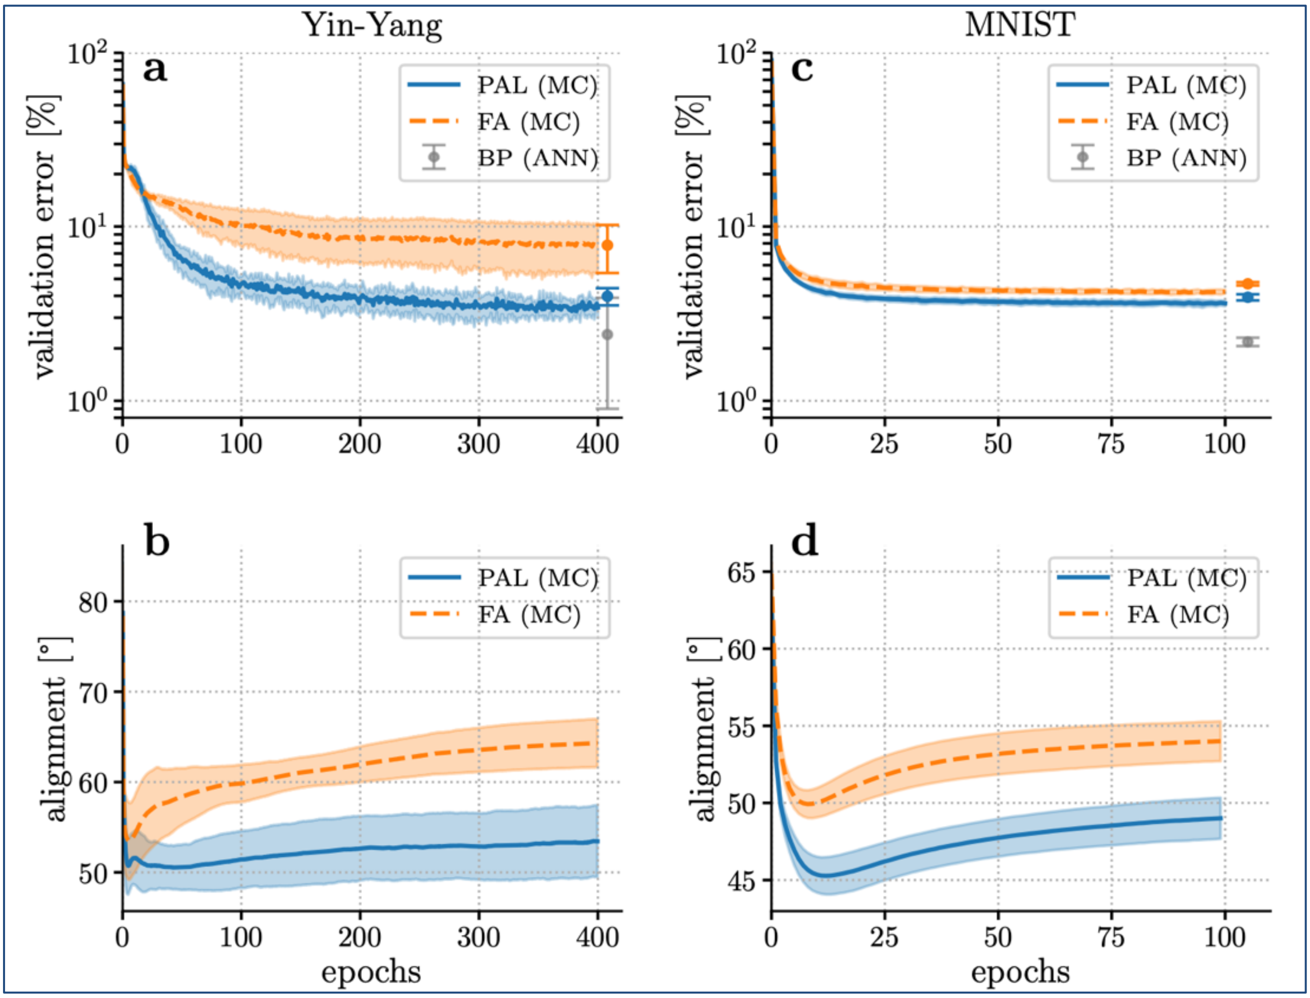

Figure 22: PAL is better than random synaptic feedback ("feedback alignment").

Figure taken from Max et al., 2022 - P3782. (a) & (c): A cortical microcircuit network trained with PAL learns to classify the Yin-Yang and MNIST digit dataset, respectively. Final accuracy is compared to an artificial neural network trained with backpropagation; in the context of bio-realistic learning, PAL outperforms FA (feedback alignment). (b) & (d): Alignment between backward and forward weights for the respective tasks; as anticipated by the theory, PAL enables better alignment, facilitating more efficient learning. We stress that these results have only been made achievable due to the integration of PAL and prospective coding with GeNN.

1.2.3 HPC

1.2.3.1 Latent Equilibrium

Contributing Partners: UBERN (P71)

Recent years have seen a surge of cortical learning models that address the problem of credit assignment in the brain. These models typically approximate the error backpropagation (BP) algorithm, which is the driving force behind most of modern machine learning. However, existing bio-plausible approximations to BP either require long relaxation phases following a change in sensory stimuli, which makes them unsuitable for tasks that require fast processing, or they impose some form of rapidly phased learning, which is not supported by experimental evidence. Fundamentally, this problem arises because of the finite response times of computational elements: biological neurons need time to integrate their input and thus introduce a response lag, which in turn requires some sort of relaxation. This does not only delay information processing, but further causes a timing mismatch between outputs and teaching signals that can even disrupt learning.

To address these limitations, we developed a novel computational framework called Latent Equilibrium (LE) that provides a unified perspective of cortical computation and learning (Section 1.1.2.3). LE allows neurons to phase-advance their output with respect to their membrane potential, which overcomes the intrinsic delays of biological neurons and enables in principle arbitrarily fast computation. The resulting model can be interpreted as a real-time, biologically plausible approximation of error backpropagation in deep cortical networks with continuous-time, leaky neuronal dynamics and continuously active, local synaptic plasticity.

To demonstrate the capabilities of our model, we tested our framework in simulations where we achieved competitive performance compared to vanilla backpropagation on standard machine learning benchmarks such as MNIST and CIFAR-10 datasets, but also on more exotic tasks such as finding Higgs bosons in LHC data, (Figure 23a-c). We further showed that we can easily integrate different deep network topologies such as multilayer perceptron and convolutional architectures to train networks on the aforementioned datasets. In addition, our model can be implemented within cortical microcircuitry resulting in a dynamical system that instantiates a bio-plausible version of BP with real-time dynamics and continuous local learning but with a significant increase in both inference and training speed compared to previous approaches such as the dendritic microcircuits by Sacramento et al., 2018 – P1527 (Figure 23d).

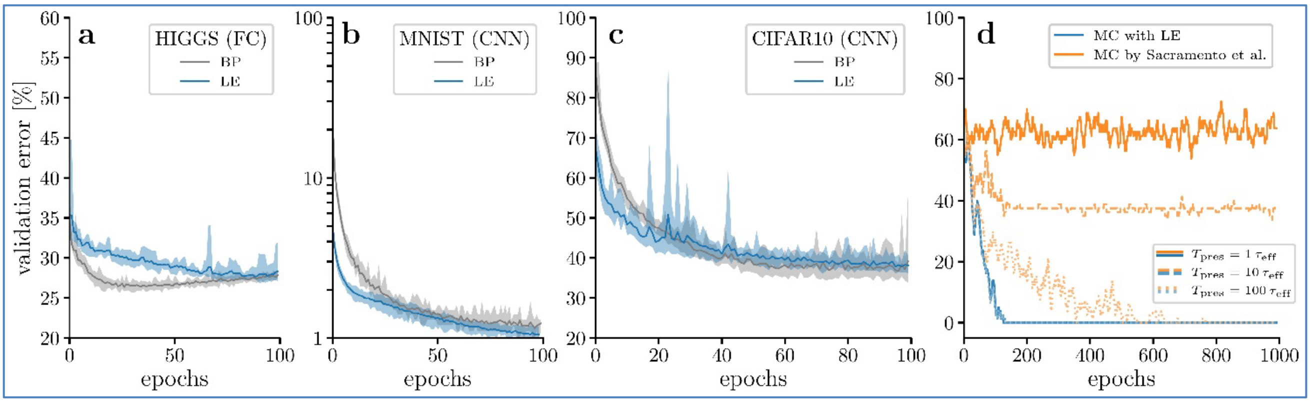

Figure 23: Performance of LE compared to vanilla BP and other microcircuit models.

Adapted from Haider et al., 2021 – P2949. (a-c) The LE framework allows the construction of cortical networks that perform on par with error backpropagation on various benchmark datasets. (d) LE enables functionality in various other models of deep learning in cortex (here: Sacramento et al., 2018 – P1527), which otherwise must rely on highly implausible and impractical relaxation times.

Importantly, we have simulated all networks with full dynamics, meaning there was no steady-state approximation and no phased plasticity involved. To do that, we heavily relied on and acknowledge computational resources from the Swiss National Supercomputing Centre (CSCS), provided through the Fenix Infrastructure project (Haider et al., 2021 - P2949).

1.2.3.2 Cartesian genetic programming framework for interpretable plasticity

Contributing Partners: UBERN (P71)

We aim to accelerate the search for phenomenological models of synaptic plasticity using evolutionary search (see Sections 1.1.1.7 and 1.2.1.5). While evolutionary algorithms are ideally suited for complex, non-differentiable search spaces such as the space of plasticity rules in spiking neuronal networks, they are computationally expensive. Resources provided by EBRAINS/ICEI on Piz Daint at the Swiss National Supercomputing Centre as well as on JUWELS at the Jülich Supercomputing Centre were essential to perform our experiments.

Our software library (Schmidt et al., 2020; see also Jordan et al., 2021 – P2556) exploits the potential for parallelism in evolutionary search on two different levels (Figure 24). First, we run several independent searches from different initial conditions. This results in different trajectories through the search space of possible learning rules, thereby mitigating the problem of local minima and increasing the chance of discovering (approximately) globally optimal solutions. Additionally, in each step (“generation”) of the evolutionary search, multiple simulations need to be performed and their results compared. Intuitively, this allows us to locally estimate the direction of fitness improvement. We parallelise these by evaluating multiple learning rules in parallel. This two-level parallelisation lends itself ideally to the architecture of supercomputing systems which allows for fine-grained parallelisation on the single node level to keep communication overhead small, while simultaneously being able to perform multiple independent parallel searches across different compute nodes.

Furthermore, using the resources provided on Piz Daint and JUWELS, we were able to perform multiple parameter scans for hyperparameters of our tasks, for example, learning rates and the number of training samples.

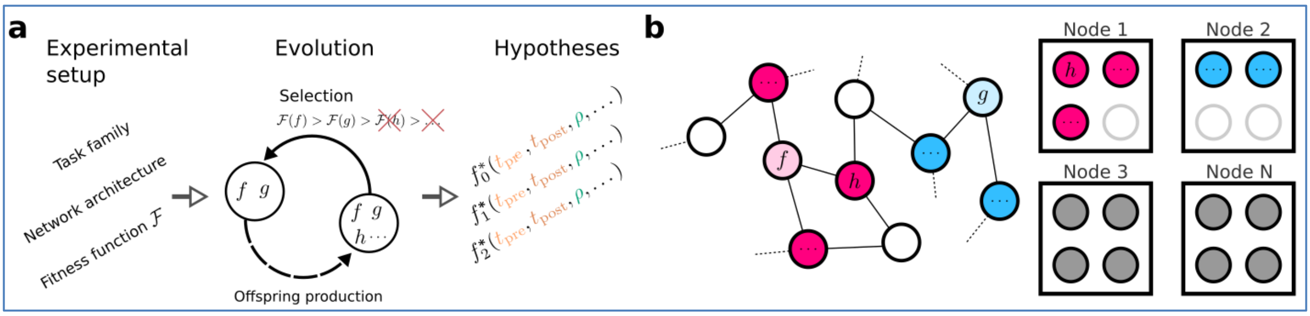

Figure 24: Evolutionary search maps well to modern supercomputer architectures.

(a) Rather than manually designing plasticity rules, our approach requires us to define fitness functions that quantify how well specific network architectures solve specific tasks provided a candidate plasticity rule. Evolutionary algorithms generate hypotheses for plasticity rules by creating offspring from previous candidates, for example changing the equation x + y to x * y. The quality of these offspring is quantified and compared by their fitness, typically measured across several simulations. (b) The search for plasticity rules can be abstracted as a graph search where each node represents an equation and links correspond to changing operators or operands. Several independent search processes can be parallelised across compute nodes (blue vs pink), while the evaluation of offspring (vivid colours) can be parallelised across cores within a node. Panel a adapted from Jordan et al., 2021 - P2556.

1.3 Education and teaching

1.3.1 A neuromorphic kit for simple experiments

Contributing Partners: UBERN (P71), UHEI (P47)

As researchers in the rapidly growing field of neuromorphic computing, we welcome and seek the opportunity to present out work to the public. One such occasion was the Nacht der Forschung (Night of Research) at the University of Bern in September 2022, where we sought to offer a playful, but evocative showcase of our approach towards understanding and replicating the human brain. To this end, we designed and produced a custom printed circuit board (PCB) that emulates a single Leaky Integrate-and-Fire (LIF) neuron. The neuron has tuneable parameters, and multiple neurons can be connected to create a network of neurons (Figure 25). We named the PCB Lu.i after its phonetic similarity to Louis Lapicque, the inventor of the LIF neuron model.

These PCBs are useful to explain multiple concepts that we are working with in a hands-on way, starting from idea of an individual spiking neuron. Furthermore, the ability to connect different neurons allows us to build simple spiking networks. Such demonstrations include simple delay lines, a network for sound localisation similar to the Jeffress model, and E-I networks. The setups are publicly available on GitHub.

Because the neurons can be used in a very accessible manner to explain our work, we have increasingly used the Lu.i PCBs not only for outreach, but also in a teaching context. This has been and continues to be explored on different levels, including worldwide activities with students, from early high school to university.

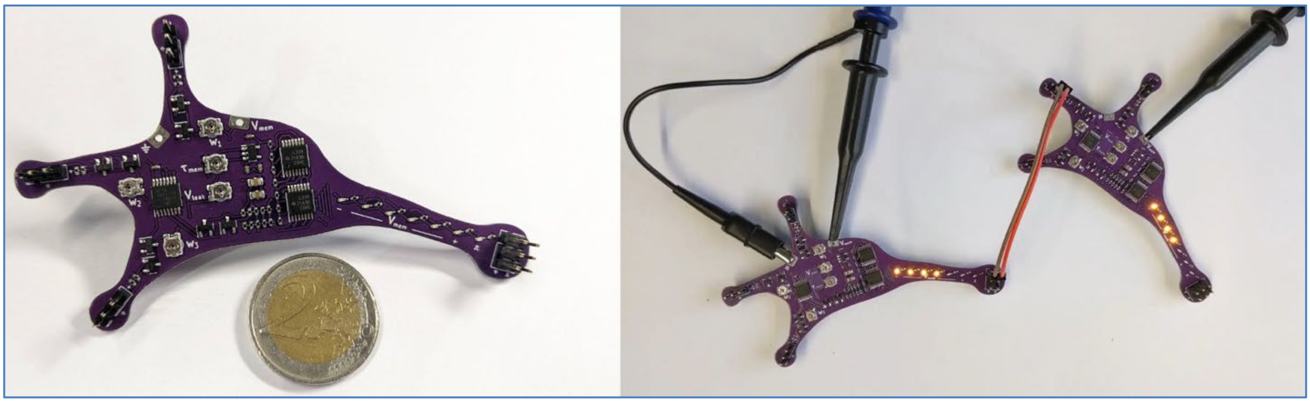

Figure 25: Photos of Lu.i PCBs

Left, photo of a single Lu.i PCB with a 2-Euro coin as a size comparison. Clearly visible are the axon terminal (lower right) and the three dendrites where neurons can be connected via jumper cables. Right, two Lu.i boards connected to each other. Turned on, the LEDs along the axon indicate the current membrane voltage and lead to a characteristic appearance. Given the flexibility of the parametrisation, a lot of different behaviours can be reproduced, which we are continuing to collect for the larger public in our open GitHub repository[6].

1.3.2 Other teaching

In addition to our many outreach activities, we have also incorporated our most recent scientific work, as well as the publicly available HBP/EBRAINS platforms, into our teaching curriculum.

At the 2021 Workshop of the Mediterranean Seminar for Consciousness (MESEC) – E3506, we gave an overview of our work on physical computing, in particular concerning our inspiration from theoretical physics (energy- and Lagrangian-based models), information coding and learning with spikes (both stochastic deterministic) and neuromorphic applications (all of them described in Section 1). While the workshop itself was mainly concerned with theories of consciousness, we underscored the idea that any change in conscious perception must be reflected by a corresponding change in the physical substrate from which it emerges, thus brining these questions into the realm of computational and theoretical neuroscience.

At the 2022 NeuroSpin Summer School – E5869, we gave an introductory lecture on computational neuroscience and spiking models. Throughout this lecture, we emphasised connections between neuroscience and machine learning, in particular on Bayesian inference and deep learning. With models developed by us in the HBP such as Spike-Based Sampling (see Section 1.2.1.3) and Latent Equilibrium (Section 1.1.2.3), we established possible applications for neuromorphic devices based on novel technologies, such as spintronics or magnonics.

For over a decade, the yearly Capo Caccia Neuromorphic Workshop – E6448, E5134, E6528, E6530 and E6531, has represented one of the most important occasions for in-depth exchange between brain and computer scientists, with a strong emphasis on neuromorphic engineering and a large audience composed predominantly of young researchers. In 2023, we have presented an overview of energy-based approaches to computation and learning, a field that finds its beginnings in the 1980s with the Hopfield model, but that has grown considerably over the past decade. In particular, we focused on our combined theories of Latent Equilibrium (see Section 1.1.2.3) and Phaseless Alignment Learning (see Section 1.1.2.4).

---

[6] Setups for Lu.i PCB: https://github.com/giant-axon/lu.i-neuron-pcb

The summer schools of the Deutsche Schülerakademie (https://www.schuelerakademien.de/deutsche-schuelerakademie) are held yearly and are tailored towards talented and highly motivated students in their last two years at high school. The courses cover topics from all fields of science and art and are taught at an undergraduate university level. The main focus of our 2023 course was the field of neuromorphic computing. Following an introduction to neurobiology and computational neuroscience, we focused on the deployment of HBP/EBRAINS platforms for research in neuroscience and machine learning (see Section 1.2), in particular in HPC simulation (NEST, PyNN) and neuromorphic emulation (BrainScaleS-2, Lu.i).

2. Looking Forward

While much of our work was targeted at answering specific questions about credit assignment in animal or artificial brains, some of our models and frameworks retain significant untapped potential and provide detailed blueprints for future research. In particular, our more general computational frameworks for machine intelligence can be further extended to tackle new problems in the field with increased efficiency and performance. This delineates a continuing roadmap that extends beyond the aegis of the HBP, but that can profit greatly from several tools and platforms that this flagship has fostered, which will be part of the enduring EBRAINS infrastructure.

The ability to precisely manipulate individual spike times goes beyond the feedforward time-to-first-spike encoding used in our previous work. Our TTFS theory extends to spike trains of arbitrary structure and duration, thus enabling gradient-descent learning in recurrent networks for solving complex spatio-temporal tasks. Unlike alternative methods such as backpropagation through time, real-time recurrent learning and many of their numerous approximations, this framework promises a form of synaptic plasticity that depends purely on the timing of pre- and postsynaptic spikes. We believe that the demonstrated robustness with respect to substrate variations will also transfer to these more difficult settings, thus fostering a neuromorphic implementation that can, similarly to our existing results, challenge the state of the art in combined time- and energy-to-solution.

Our E2L evolutionary framework for meta-plasticity has already demonstrated its capabilities on conventional computing clusters, but we have also shown very promising results on the neuromorphic BrainScaleS-2 system. A continuation of this project after the end of the HBP is planned, but the potential for scaling beyond single chips is evident. While the massive acceleration of BrainScaleS-2 allows a uniquely rapid evaluation of an agent performance, the inherent parallelism of our framework suggests clear advantages gained from multi-chip systems. This presents a unique opportunity for exploiting the capabilities of the large-scale SpiNNaker-2 system, possibly in conjunction with BrainScaleS-2 for the simultaneous training of multiple agents of different size and computational capacity.

Finally, we have already shown (see also D3.12/T3.6) how our Latent Equilibrium framework can be generalised from spatial to spatiotemporal computing tasks. In conjunction with Phaseless Alignment Learning, this constitutes a complete theory for fully plastic, biologically plausible, and neuromorphically realisable neural networks, with the demonstrated ability to challenge the existing state of the art. In ongoing work using the EBRAINS GPU Platform (powered by GeNN), we aim to improve the robustness of the underlying microcircuit model and study the model’s amenability to an event-based form of information representation and exchange. In this context, our E2L framework will also be used to optimise the efficiency of these sparse codes. Altogether, the resulting architecture will not only lend itself to efficient deployment on the large-scale SpiNNaker-2 system but will also provide a blueprint for novel silicon circuits that will allow a highly accelerated and extremely power-efficient physical realisation, much like the BrainScaleS-2 platform has allowed for networks of leaky integrate-and-fire neurons.

3. Annex: Bibliography

HBP Publications

- S. Billaudelle et al., “Versatile Emulation of Spiking Neural Networks on an Accelerated Neuromorphic Substrate,” in 2020 IEEE International Symposium on Circuits and Systems (ISCAS), Oct. 2020, pp. 1–5. doi: 10.1109/ISCAS45731.2020.9180741 (P2241).
- S. Billaudelle et al., “Structural plasticity on an accelerated analog neuromorphic hardware system,” Neural Networks, vol. 133, pp. 11–20, Jan. 2021, doi: 10.1016/j.neunet.2020.09.024 (P2240).
- U. M. Mengual, W. A. M. Wybo, L. J. E. Spierenburg, M. Santello, W. Senn, and T. Nevian, “Efficient Low-Pass Dendro-Somatic Coupling in the Apical Dendrite of Layer 5 Pyramidal Neurons in the Anterior Cingulate Cortex,” J. Neurosci., vol. 40, no. 46, Art. no. 46, Nov. 2020, doi: 10.1523/JNEUROSCI.3028-19.2020 (P2759).
- W. A. Wybo, J. Jordan, B. Ellenberger, U. Marti Mengual, T. Nevian, and W. Senn, “Data-driven reduction of dendritic morphologies with preserved dendro-somatic responses,” eLife, vol. 10, p. e60936, Jan. 2021, doi: 10.7554/eLife.60936 (P2760).
- H. D. Mettler, M. Schmidt, W. Senn, M. A. Petrovici, and J. Jordan, “Evolving Neuronal Plasticity Rules using Cartesian Genetic Programming,” arXiv:2102.04312 [cs, q-bio], Feb. 2021, Accessed: Feb. 10, 2021. [Online]. Available: http://arxiv.org/abs/2102.04312 (P2795).
- C. Zhao, Y. F. Widmer, S. Diegelmann, M. A. Petrovici, S. G. Sprecher, and W. Senn, “Predictive olfactory learning in Drosophila,” Sci Rep, vol. 11, no. 1, Art. no. 1, Mar. 2021, doi: 10.1038/s41598-021-85841-y (P2799)
- J. Jordan, J. Sacramento, W. A. M. Wybo, M. A. Petrovici, and W. Senn, “Learning Bayes-optimal dendritic opinion pooling,” arXiv:2104.13238 [q-bio], Apr. 2021, Accessed: Apr. 28, 2021. [Online]. Available: http://arxiv.org/abs/2104.13238 (P2837)
- J. Göltz et al., “Fast and energy-efficient neuromorphic deep learning with first-spike times,” Nat Mach Intell, vol. 3, no. 9, Art. no. 9, Sep. 2021, doi: 10.1038/s42256-021-00388-x. (P2917)
- P. Haider, B. Ellenberger, L. Kriener, J. Jordan, W. Senn, and M. A. Petrovici, “Latent Equilibrium: A unified learning theory for arbitrarily fast computation with arbitrarily slow neurons,” in Advances in Neural Information Processing Systems, M. Ranzato, A. Beygelzimer, Y. Dauphin, P. S. Liang, and J. W. Vaughan, Eds., Curran Associates, Inc., 2021, pp. 17839–17851. [Online]. Available: https://proceedings.neurips.cc/paper/2021/file/94cdbdb84e8e1de8a725fa2ed61498a4-Paper.pdf (P2949)
- J. Jordan, M. Schmidt, W. Senn, and M. A. Petrovici, “Evolving interpretable plasticity for spiking networks,” eLife, vol. 10, p. e66273, Oct. 2021, doi: 10.7554/eLife.66273. (P2556)
- S. Czischek et al., “Spiking neuromorphic chip learns entangled quantum states,” SciPost Physics, vol. 12, no. 1, p. 039, Jan. 2022, doi: 10.21468/SciPostPhys.12.1.039. (P2574)
- K. Amunts et al., “The coming decade of digital brain research - A vision for neuroscience at the intersection of technology and computing,” Zenodo, Mar. 2023. doi: 10.5281/zenodo.7764003. (P3321)
- Korcsak-Gorzo et al., “Cortical oscillations support sampling-based computations in spiking neural networks,” PLOS Computational Biology, vol. 18, no. 3, p. e1009753, Mar. 2022, doi: 10.1371/journal.pcbi.1009753. (P2806)
- N. Deperrois, M. A. Petrovici, W. Senn, and J. Jordan, “Learning cortical representations through perturbed and adversarial dreaming,” eLife, vol. 11, p. e76384, Apr. 2022, doi: 10.7554/eLife.76384. (P2911)

- E. Kreutzer, W. Senn, and M. A. Petrovici, “Natural-gradient learning for spiking neurons,” eLife, vol. 11, p. e66526, Apr. 2022, doi: 10.7554/eLife.66526. (P2621)
- L. Kriener, J. Göltz, and M. A. Petrovici, “The Yin-Yang dataset,” in Neuro-Inspired Computational Elements Conference, in NICE 2022. New York, NY, USA: Association for Computing Machinery, Mar. 2022, pp. 107–111. doi: 10.1145/3517343.3517380. (P2798)
- A. Baumbach, R. Klassert, S. Czischek, M. Gärttner, and M. A. Petrovici, “Quantum many-body states: A novel neuromorphic application,” in Proceedings of the 2022 Annual Neuro-Inspired Computational Elements Conference, in NICE ’22. New York, NY, USA: Association for Computing Machinery, May 2022, pp. 104–106. doi: 10.1145/3517343.3517379. (P3261)
- R. Klassert, A. Baumbach, M. A. Petrovici, and M. Gärttner, “Variational learning of quantum ground states on spiking neuromorphic hardware,” iScience, vol. 25, no. 8, p. 104707, Aug. 2022, doi: 10.1016/j.isci.2022.104707. (P2932)
- K. Max, L. Kriener, G. P. García, T. Nowotny, W. Senn, and M. A. Petrovici, “Learning efficient backprojections across cortical hierarchies in real time.” arXiv, Dec. 20, 2022. doi: 10.48550/arXiv.2212.10249. (P3782)
- W. Senn et al., “A neuronal least-action principle for real-time learning in cortical circuits,” eLife, vol. 12, Aug. 2023, doi: 10.7554/eLife.89674. (P3937)
- M. Farisco et al., “A method for the ethical analysis of brain-inspired AI.” arXiv, May 18, 2023. doi: 10.48550/arXiv.2305.10938. (P4073)
- K. A. Wilmes, M. A. Petrovici, S. Sachidhanandam, and W. Senn, “Uncertainty-modulated prediction errors in cortical microcircuits.” bioRxiv, p. 2023.05.11.540393, Jul. 16, 2023. doi: 10.1101/2023.05.11.540393. (P4014)
- W. A. M. Wybo et al., “NMDA-driven dendritic modulation enables multitask representation learning in hierarchical sensory processing pathways,” Proceedings of the National Academy of Sciences, vol. 120, no. 32, p. e2300558120, Aug. 2023, doi: 10.1073/pnas.2300558120. (P3781)
- N. Deperrois, M. A. Petrovici, W. Senn, and J. Jordan, “Learning beyond sensations: how dreams organize neuronal representations.” arXiv, Aug. 03, 2023. doi: 10.48550/arXiv.2308.01830. (P4097)
- T. Wunderlich et al., “Demonstrating Advantages of Neuromorphic Computation: A Pilot Study,” Frontiers in Neuroscience, vol. 13, 2019, Accessed: Sep. 13, 2023. [Online]. Available: https://www.frontiersin.org/articles/10.3389/fnins.2019.00260 (P1721).
- B. Cramer et al., “Surrogate gradients for analog neuromorphic computing,” Proceedings of the National Academy of Sciences, vol. 119, no. 4, p. e2109194119, Jan. 2022, doi: 10.1073/pnas.2109194119 (P2756).
- M. A. Petrovici, J. Bill, I. Bytschok, J. Schemmel, and K. Meier, “Stochastic inference with spiking neurons in the high-conductance state,” Phys. Rev. E, vol. 94, no. 4, p. 042312, Oct. 2016, doi: 10.1103/PhysRevE.94.042312 (P1213).
- L. Leng et al., “Spiking neurons with short-term synaptic plasticity form superior generative networks,” Sci Rep, vol. 8, no. 1, Art. no. 1, Jul. 2018, doi: 10.1038/s41598-018-28999-2 (P1334).
- A. F. Kungl et al., “Accelerated Physical Emulation of Bayesian Inference in Spiking Neural Networks,” Frontiers in Neuroscience, vol. 13, 2019, Accessed: Sep. 13, 2023. [Online]. Available: https://www.frontiersin.org/articles/10.3389/fnins.2019.01201 (P1360).
- D. Dold et al., “Stochasticity from function — Why the Bayesian brain may need no noise,” Neural Networks, vol. 119, pp. 200–213, Nov. 2019, doi: 10.1016/j.neunet.2019.08.002 (P1447).
- J. Jordan et al., “Deterministic networks for probabilistic computing,” Sci Rep, vol. 9, no. 1, Art. no. 1, Dec. 2019, doi: 10.1038/s41598-019-54137-7 (P843).
- O. Breitwieser, A. Baumbach, A. Korcsak-Gorzo, J. Klähn, M. Brixner, and M. Petrovici, “sbs: Spike-based Sampling (v1.8.2).” Zenodo, Feb. 24, 2020. doi: 10.5281/zenodo.3686015 (P2517).

- B. Gardner and A. Grüning, “Supervised Learning With First-to-Spike Decoding in Multilayer Spiking Neural Networks,” Frontiers in Computational Neuroscience, vol. 15, 2021. Accessed: Sep.13, 2023. [Online]. Available: https://www.frontiersin.org/articles/10.3389/fncom.2021.617862 (P3173).
- J. Sacramento, R. Ponte Costa, Y. Bengio, and W. Senn, “Dendritic cortical microcircuits approximate the backpropagation algorithm,” in Advances in Neural Information Processing Systems, Curran Associates, Inc., 2018. Accessed: Sep. 13, 2023. [Online]. Available: https://papers.nips.cc/paper_files/paper/2018/hash/1dc3a89d0d440ba31729b0ba74b93a33-Abstract.html (P1527).
- J. C. Knight and T. Nowotny, “Larger GPU-accelerated brain simulations with procedural connectivity,” Nat Comput Sci, vol. 1, no. 2, Art. no. 2, Feb. 2021, doi: 10.1038/s43588-020-00022-7 (P2777).
- G. Bellec, F. Scherr, E. Hajek, D. Salaj, R. Legenstein, and W. Maass, “Biologically inspired alternatives to backpropagation through time for learning in recurrent neural nets,” Biologically inspired alternatives to backpropagation through time for learning in recurrent neural nets, vol. 2019, pp. 1–37, Jan. 2019 (P1836).

Non-HBP Publications

- F. Zenke et al., “Visualizing a joint future of neuroscience and neuromorphic engineering,” Neuron, vol. 109, no. 4, Art. no. 4, Feb. 2021, doi: 10.1016/j.neuron.2021.01.009.
- M. Schmidt and J. Jordan, “hal-cgp: Cartesian genetic programming in pure Python.” Zenodo, Jun. 11, 2020. doi: 10.5281/zenodo.3889163.
- R. Urbanczik and W. Senn, “Learning by the Dendritic Prediction of Somatic Spiking,” Neuron, vol. 81, no. 3, pp. 521–528, Feb. 2014, doi: 10.1016/j.neuron.2013.11.030.
- Y. Nir and G. Tononi, “Dreaming and the brain: from phenomenology to neurophysiology,” Trends in Cognitive Sciences, vol. 14, no. 2, pp. 88–100, Feb. 2010, doi: 10.1016/j.tics.2009.12.001.
- S. Scarpelli, C. Bartolacci, A. D’Atri, M. Gorgoni, and L. De Gennaro, “Mental Sleep Activity and Disturbing Dreams in the Lifespan,” International Journal of Environmental Research and Public Health, vol. 16, no. 19, Art. no. 19, Jan. 2019, doi: 10.3390/ijerph16193658.
- E. Müller et al., “A Scalable Approach to Modeling on Accelerated Neuromorphic Hardware,” Frontiers in Neuroscience, vol. 16, 2022, Accessed: Sep. 13, 2023. [Online]. Available: https://www.frontiersin.org/articles/10.3389/fnins.2022.884128
- J. Ma, G. Yu, W. Xiong, and X. Zhu, “Safe semi-supervised learning for pattern classification,” Engineering Applications of Artificial Intelligence, vol. 121, p. 106021, May 2023, doi: 10.1016/j.engappai.2023.106021.
- J. H. Lee, S. Haghighatshoar, and A. Karbasi, “Exact Gradient Computation for Spiking Neural Networks via Forward Propagation,” in Proceedings of The 26th International Conference on Artificial Intelligence and Statistics, PMLR, Apr. 2023, pp. 1812–1831. Accessed: Sep. 13, 2023. [Online]. Available: https://proceedings.mlr.press/v206/lee23b.html
- P. Spilger et al., “hxtorch.snn: Machine-learning-inspired Spiking Neural Network Modeling on BrainScaleS-2,” in Proceedings of the 2023 Annual Neuro-Inspired Computational Elements Conference, in NICE ’23. New York, NY, USA: Association for Computing Machinery, Apr. 2023, pp. 57–62. doi: 10.1145/3584954.3584993.
- S. K. Esser, R. Appuswamy, P. Merolla, J. V. Arthur, and D. S. Modha, “Backpropagation for Energy-Efficient Neuromorphic Computing,” in Advances in Neural Information Processing Systems, Curran Associates, Inc., 2015. Accessed: Sep. 13, 2023. [Online]. Available: https://papers.nips.cc/paper_files/paper/2015/hash/10a5ab2db37feedfdeaab192ead4ac0e-Abstract.html
- F. Zenke and E. O. Neftci, “Brain-Inspired Learning on Neuromorphic Substrates,” Proceedings of the IEEE, vol. 109, no. 5, pp. 935–950, May 2021, doi: 10.1109/JPROC.2020.3045625.

- W. A. M. Wybo et al., “Dendritic modulation enables multitask representation learning in hierarchical sensory processing pathways.” bioRxiv, p. 2022.11.25.517941, Nov. 26, 2022. doi: 10.1101/2022.11.25.517941.