Physical models of the brain
From theory to neural substrates

D. Dold[1,2], K. Meier[†], A. Mielke[1], W. Senn[2], M. A. Petrovici[1,2]
as well as J. Sacramento, I. Bytschok, A. F. Kungl, A. Baumbach, R. P. Costa, O. Breitwieser, J. Schemmel, K. Schindler, J. Binas & Y. Bengio
[1]Heidelberg University, Kirchhoff-Institute for Physics, [2]University of Bern, Department of Physiology

The Manfred Stärk Foundation

Abstract

The underlying **computational principles** allowing **the brain** to deal with unreliable, high-dimensional and often incomplete data while having a power consumption on the order of a few Watt are still mostly unknown.
Here, I present ideas on how structures and mechanisms found in the cerebral cortex might be employed to perform **Bayesian computing with spiking neurons** and to implement the widely used **error backpropagation algorithms in cortical networks**. Such models are ideal candidates for hardware mimicking the vastly parallel structure of the brain (so-called **"neuromorphic" hardware**), promising a strongly accelerated and power-efficient implementation of powerful learning and inference algorithms.

A few building blocks of the brain

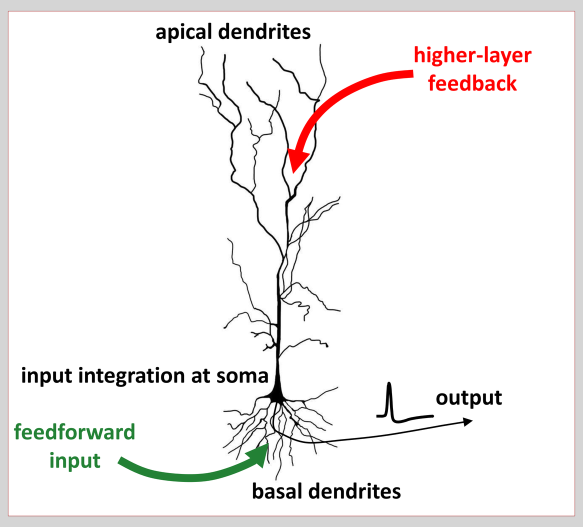

I Different from abstract neurons used in deep learning, biological neurons come in different types and shapes.
For instance, the main neuron type found in the cortex (pyramidal neurons) possesses tree-like structures (dendrites) for input integration.

What is the function of such neuron diversity?

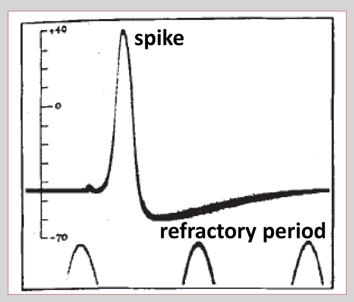

II Biological neurons communicate via all-or-nothing events called action potentials (spikes). The simplest model for this is a leaky integrator *(simplified for convenience)*

$$\tau \dot{u} = -u + \text{input}\,, \qquad (1)$$

which emits a spike when the membrane potential passes a threshold value. Afterwards, it cannot be excited again for some time (refractory period).

Spikes are energy efficient, but are there more benefits?

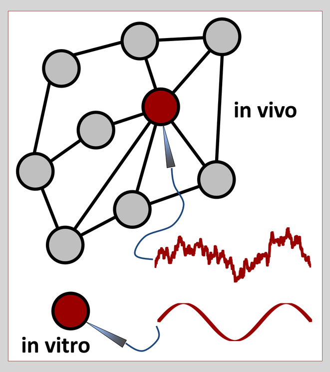

III Biological neurons behave deterministically in vitro, but are noisy in vivo due to a bombardement with spikes coming from approx. 10,000 adjacent cortical neurons.

Is this noise harmful or can it be utilized somehow?

II+III Sampling with deterministic spiking neurons

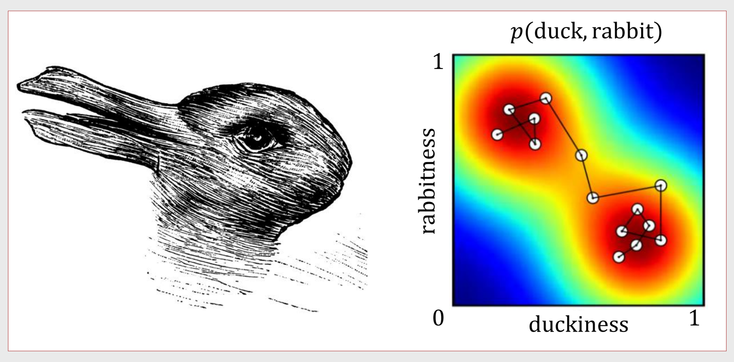

The noisy behavior of neurons is very likely the hallmark of a stochastic computation scheme. **Such a scheme explains how the brain deals with ambiguous input**, and how visual illusions like bi-stable images (duck/rabbit) might form, i.e., through sampling of modes.

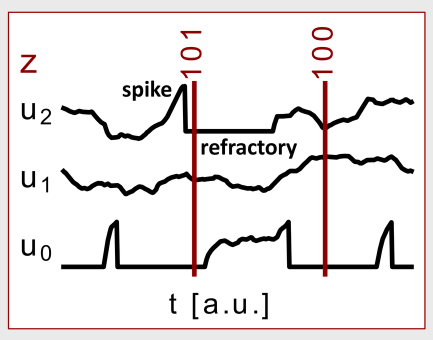

Stochastic comp. can naturally be implemented with spiking neurons by assigning refractory neurons the state $z = 1$ and 0 otherwise.

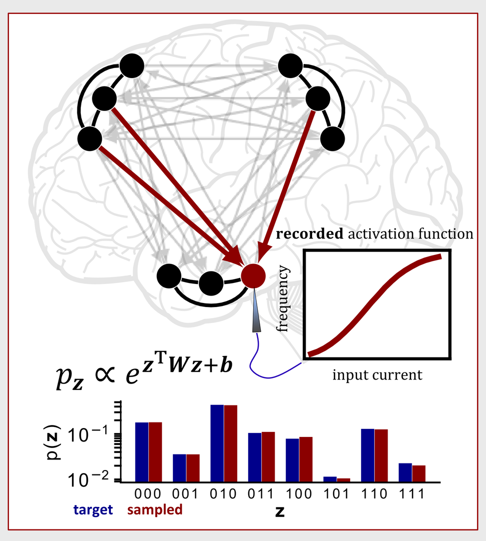

Neurons become stochastic when embedded in an ensemble of (functional) networks, like particles in a heat bath. In this scenario, **the spiking dynamics of every network in the ensemble sample from a probability distribution** parametrized by the respective synaptic weights, without any "true" source of stochasticity.

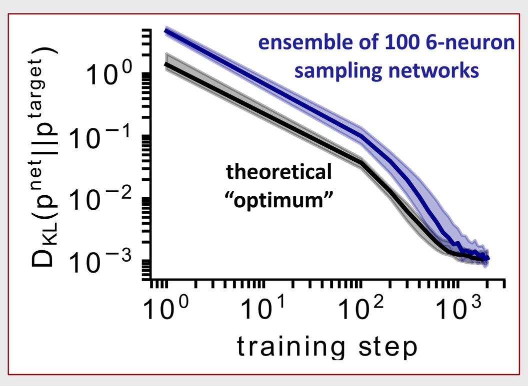

All networks of the ensemble can be trained with contrastive divergence to either sample from target distributions or model data, allowing them to perform pattern completion and classification tasks.

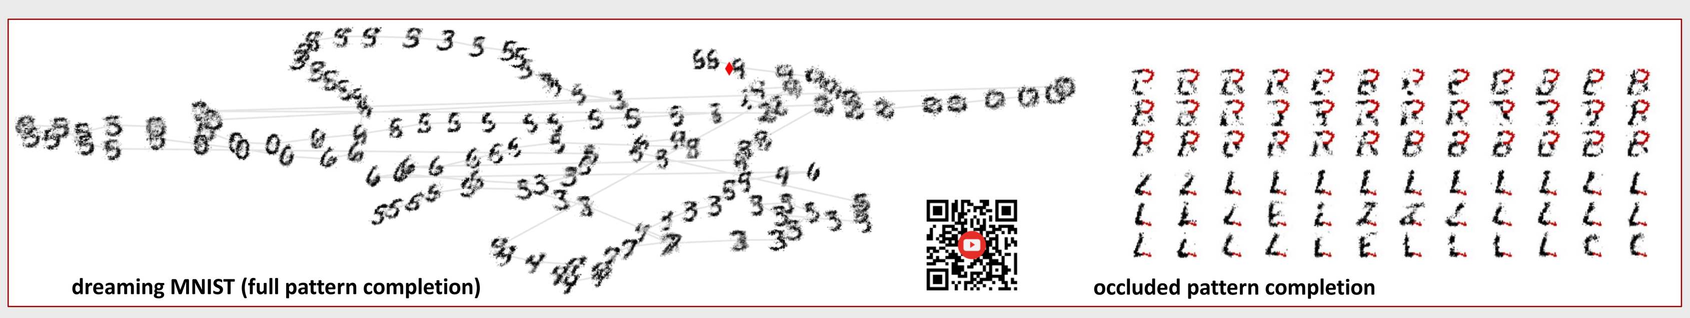

I Backprop in cortical networks

Whether the brain might use an optimization scheme like backprop to guide synaptic plasticity in deep hierarchical cortical areas is still an open question.
In our model, **backprop is approximated by cortical circuits combining different neuron types and extended neuron models.** Errors are calculated locally via lateral interneuron circuits that try to explain away feedback coming from higher areas. These errors nudge the soma, becoming accessible to a biologically plausible plasticity rule $\dot{W}_i \propto \underbrace{[u_i - W_i \varphi(u_{i-1})]}_{e_i} \varphi^T(u_{i-1})$.

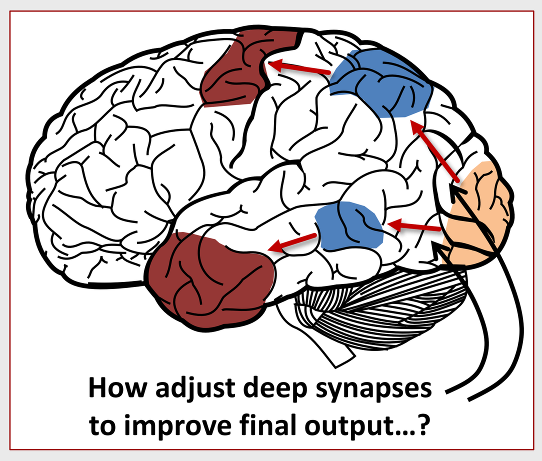

Errors are propagated backward through the network via feedback connections while sensory information is propagated forward. **Neurons minimize these local prediction errors e_i, which in turn reduces a global cost function.**

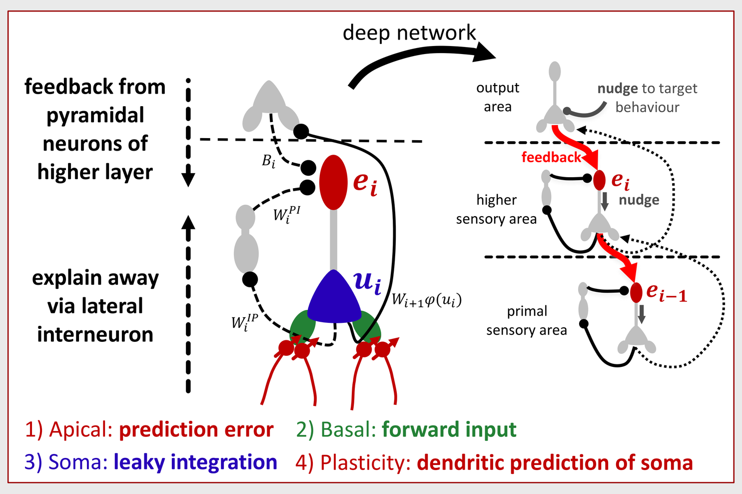

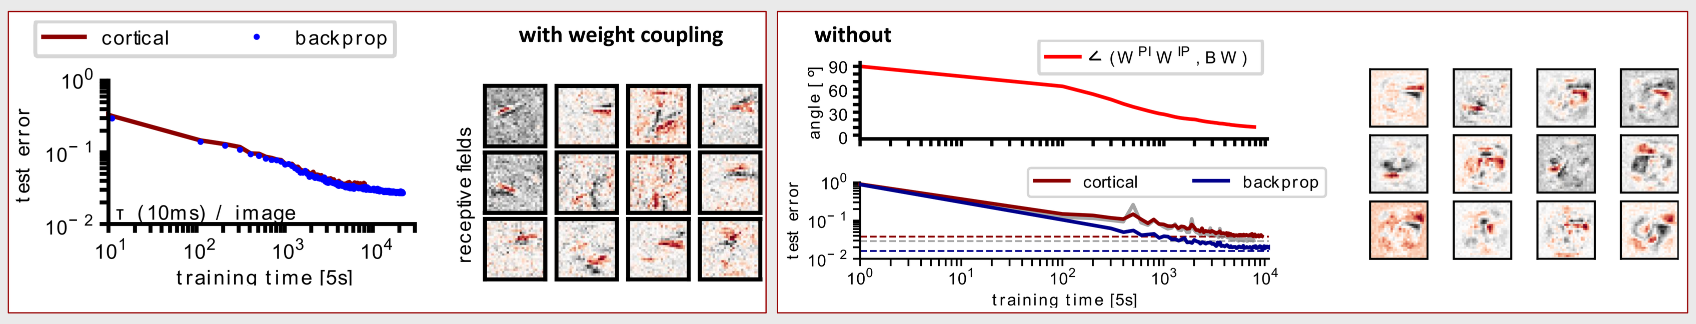

The full model can be derived from first principle by introducing a Lagrangian L that has to be stationary under neural dynamics and is minimized by synaptic dynamics.

The combined neurosynaptic dynamics lead to the emergence of backprop.

$$L(\boldsymbol{u}) = \sum_i \underbrace{\| \boldsymbol{u}_i - \boldsymbol{W}_i \boldsymbol{\varphi}(\boldsymbol{u}_{i-1}) \|^2}_{\text{prediction error } e_i} + \overset{\text{teacher strength}}{\beta} \cdot \boldsymbol{cost}$$

$$\boldsymbol{\delta L = 0} \quad \& \quad \dot{\boldsymbol{W}} = -\boldsymbol{\eta} \cdot \boldsymbol{\nabla_W L}$$

Theorem 1 *(real-time backprop)*

$$\boldsymbol{e_i} = W_{i+1}^{\mathrm{T}} \boldsymbol{e_{i+1}}\,, \quad \dot{\boldsymbol{W}}_i \propto \boldsymbol{e_i} \varphi(u_i)^{\mathrm{T}}$$

Theorem 2 *(real-time gradient descent)*

$$\frac{\mathbf{d}}{\mathbf{d}\boldsymbol{W_i}} \boldsymbol{cost} = \lim_{\beta \to 0} \frac{1}{\beta} \boldsymbol{e}_i^{\boldsymbol{\beta}} \varphi^{\beta}(u_i)^{\mathrm{T}}$$

Neuromorphic hardware: Towards silicon brains

Inspired by the brain, **non-Von Neumann architectures** are developed to explore novel computational paradigms. One such platform is the BrainScaleS physical model system in Heidelberg, **implementing neurons and synapses as analogue circuits with digital spike communication**. It promises great emulation speed (10^4 speed-up compared to biology) and low power demand.

On this system, **we achieved a physical realization of deterministic spiking sampling ensembles** (II+III).

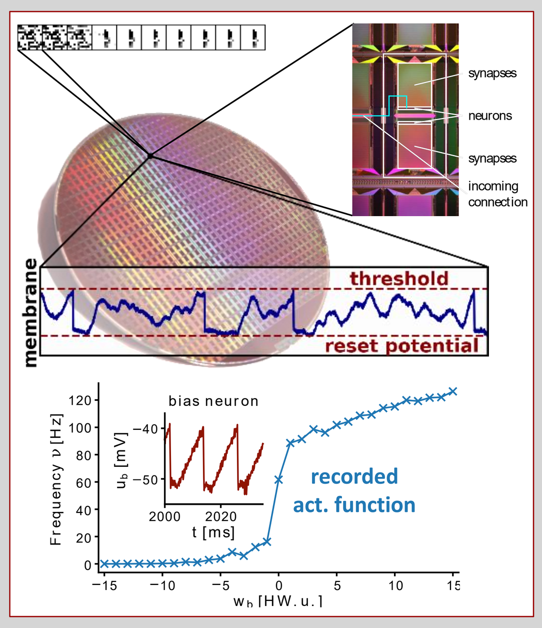

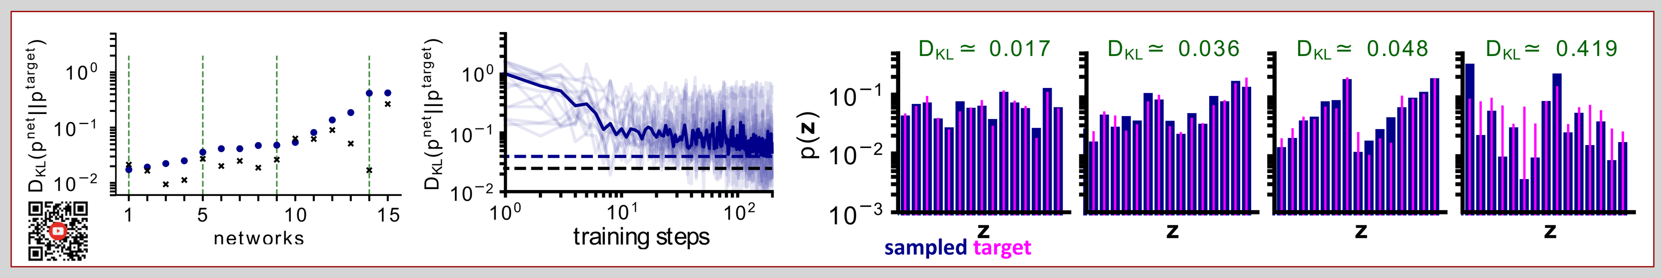

Currently developed systems - here shown: **HICANN-X (Heidelberg)** and **Loihi (Intel)** - feature **on-chip learning**, allowing an energy efficient and (possibly) accelerated implementation of local learning rules on neural substrates.

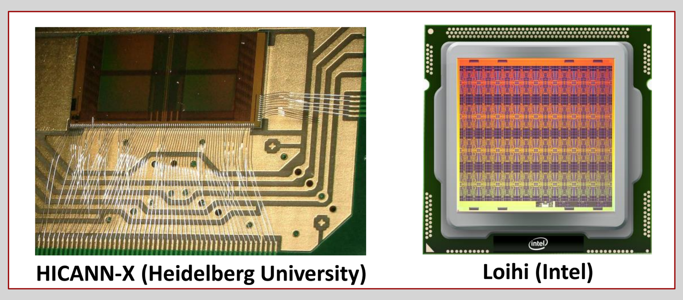

Conclusion

Why spikes?[1]
- energy efficient signal transmission
- possibly to encode samples

Stochasticity in the brain?[2]
- major source: background activity
- spiking activity as sampling from posterior distribution
- "deterministic" Bayesian computing

Can the brain do backprop?[3,4]
- possibly by employing dendrites, feedback and cortical circuitry
- learning rule itself local and biologically plausible / interpretable

Computers like brains?[5,6]
- utilizing spikes and finding efficient local learning rules are currently the main challenges to more "brain-like" hardware

[1]Petrovici, Mihai A., et al. **Stochastic inference with spiking neurons in the high-conductance state**. Physical Review E (2016).
[2]Dold, Dominik, et al. **Stochasticity from function-why the Bayesian brain may need no noise**. *arXiv preprint arXiv:1809.08045* (2018).
[3]Sacramento, João, et al. **Dendritic cortical microcircuits approximate the backpropagation algorithm**. NeurIPS (2018).
[4]Dold, Dominik, et al. **Lagrangian dynamics of dendritic microcircuits enables real-time backpropagation of errors**. Cosyne Abstracts (2019).
[5]Kungl, Akos F., et al. **Accelerated physical emulation of Bayesian inference in spiking neural networks**. *arXiv preprint arXiv:1807:02389* (2019).
[6]Wunderlich, Timo, et al. **Demonstrating Advantages of Neuromorphic Computation: A Pilot Study**. Frontiers in Neuroscience (2019).

This work has received funding from the **Manfred Stärk Foundation** and the **European Union Horizon 2020 Framework Programme** under grant agreement 720270 and 785907 (HBP). Calculations using GPUs were performed on **UBELIX** (http://www.id.unibe.ch/hpc), the HPC cluster of the University of Bern. Furthermore, this work was supported by the state of Baden-Württemberg through **bwHPC** and the **German Research Foundation (DFG)** through grant no INST 39/963-1 FUGG.

Email contact: dodo@kip.uni-heidelberg.de