Function from form
Two models of coding and learning in cortical circuits

D. Dold[1,2], M. A. Petrovici[1,2], W. Senn[2], K. Meier[1], A. Mielke[1]

as well as J. Sacramento, I. Bytschok, A. F. Kungl, A. Baumbach, 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

From a generic Bayesian perspective, cortical networks can be viewed as generators of target distributions. To enable such computation, models assume neurons to possess sources of perfect, well-behaved noise - an assumption that is both impractical and at odds with biology. We show how **local plasticity in an ensemble of spiking networks allows them to co-shape their activity towards a set of well-defined targets**, while reciprocally using the very same activity as a source of (pseudo-)stochasticity. This enables **purely deterministic networks** to simultaneously learn a variety of tasks, completely removing the need for true randomness by **using the available background activity of the whole ensemble as a resource to perform Bayesian computations**.

Dynamics: Leaky Integrate-and-Fire with COBA synapses

$$C_m\,\dot u = g_l(E_l - u) + I_{network} + I_{noise}$$

Coding: refractory **z=1**, else **z=0**

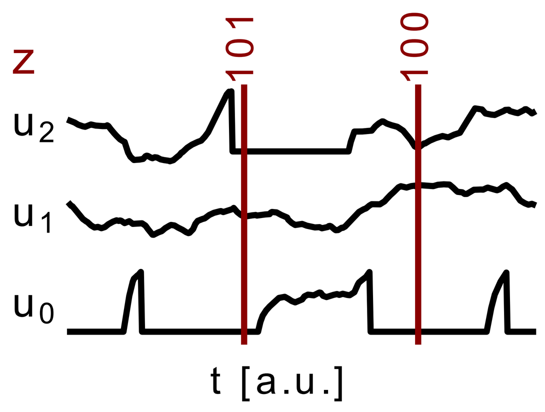

Network dynamics = **sampling** from Boltzmann distr.

$$p_z \propto e^{z^T W z + b}$$

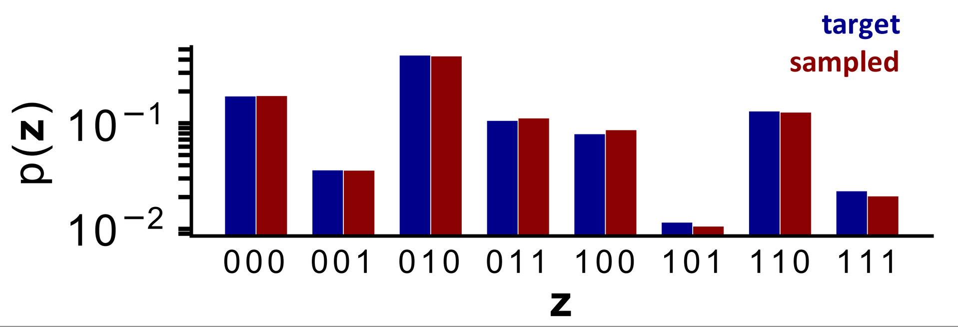

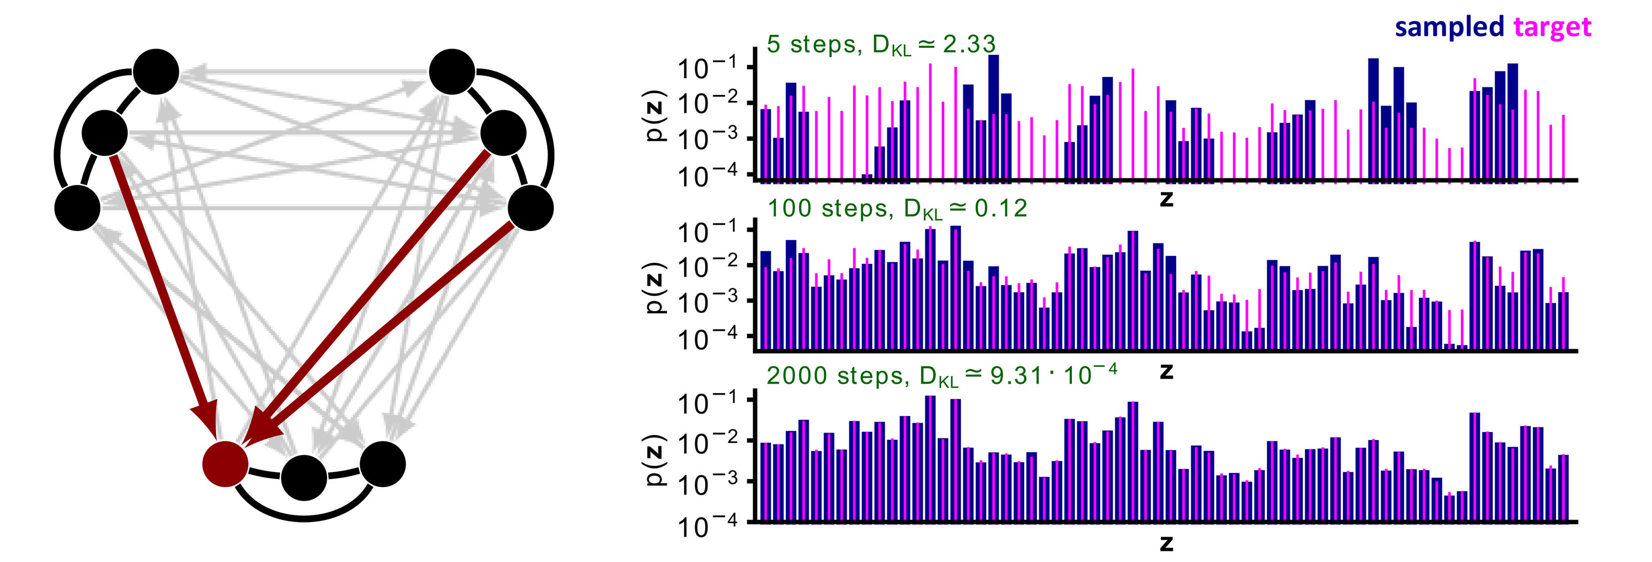

Noise source: Spikes from functionally disjunct subnetworks, no Poisson noise or any other pseudo-randomly generated noise.

Training: All networks simultaneously trained with Contrastive Divergence to sample from target distributions.

$$\Delta W_{ij} = p^{target}_{z_i=1,z_j=1} - p^{sampled}_{z_i=1,z_j=1} \qquad \Delta b_i = p^{target}_{z_i=1} - p^{sampled}_{z_i=1}$$

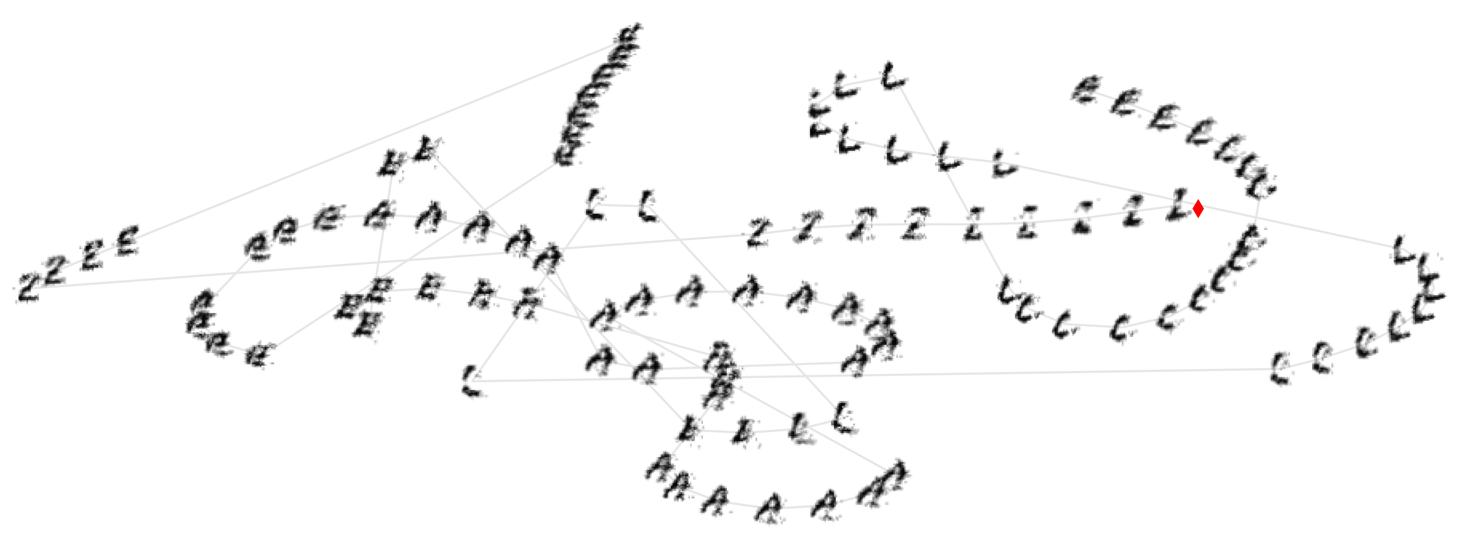

Hierarchical networks: Ensemble of (deterministic) networks trained on digits and letters are able to perform classification, pattern completion and sample generation (shown here).

Physical modeling: The concept directly translates to analogue neuromorphic hardware (BrainScaleS), yielding an acceleration of 10^4 compared to biological time.

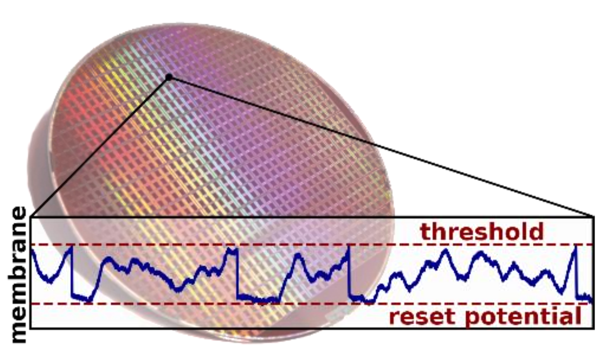

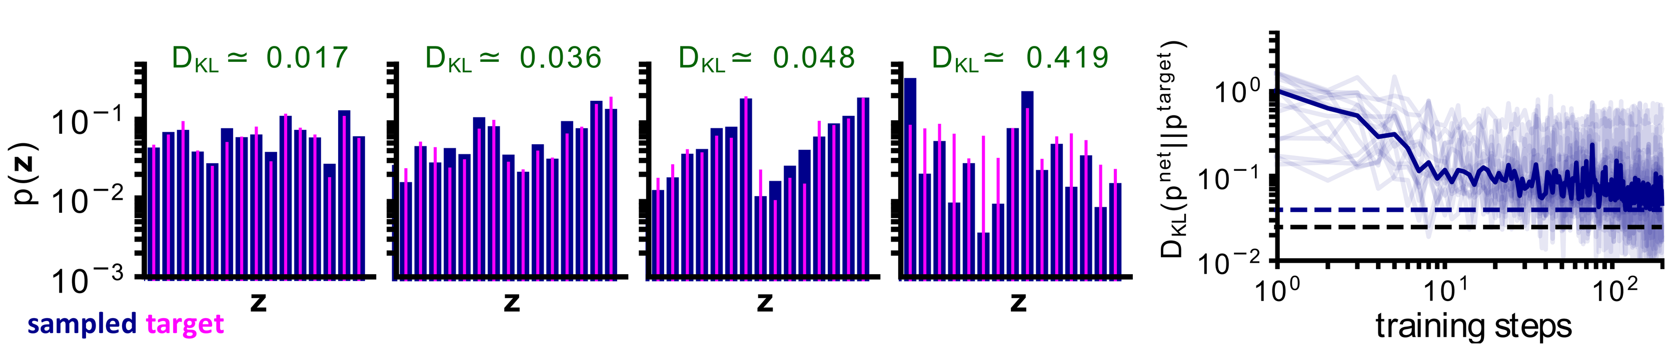

A major driving force behind the recent achievements of deep learning is the backpropagation-of-errors algorithm (backprop), which solves the credit assignment problem for deep neural networks. Its effectiveness in abstract neural networks notwithstanding, it remains unclear whether backprop represents a viable implementation of cortical plasticity. Here, we present a new theoretical framework that uses a **least-action principle to derive a biologically plausible implementation of backprop**. The presented model incorporates several features of biological neurons that cooperate towards approximating a **time-continuous version of backprop**, where **plasticity acts at all times to reduce an output error induced by mismatch between different information streams in the network**.

Energy function encodes network architecture and cost.

$$E = \frac{1}{2}\sum_{k=1}^{N} \underbrace{\|u_k - W_k \bar r_{k-1}\|^2}_{\text{prediction error}} + \underbrace{\beta C}_{\text{cost function}}$$

Dynamics: Euler-Lagrange eqs. on advanced potential.

$$\tilde u(t) = \frac{1}{\tau}\int_t^{\infty} dt'\, u(t')\, e^{-\frac{t'-t}{\tau}} \xrightarrow{L(\tilde u, \dot{\tilde u}) = -E(u)} \frac{\partial L}{\partial \tilde u} - \frac{d}{dt}\frac{\partial L}{\partial \dot{\tilde u}} = 0$$

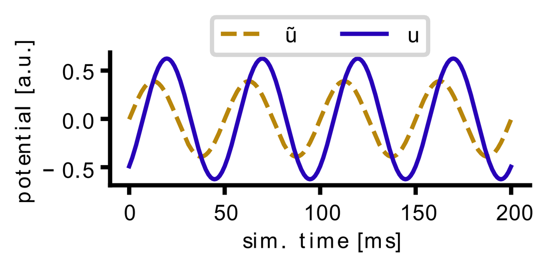

Equation of motion:

$$\underbrace{\tau \dot u_k = -u_k}_{\text{somatic integration}} + \underbrace{W_k r_{k-1} + e_k}_{\text{basal and apical compartment}}$$

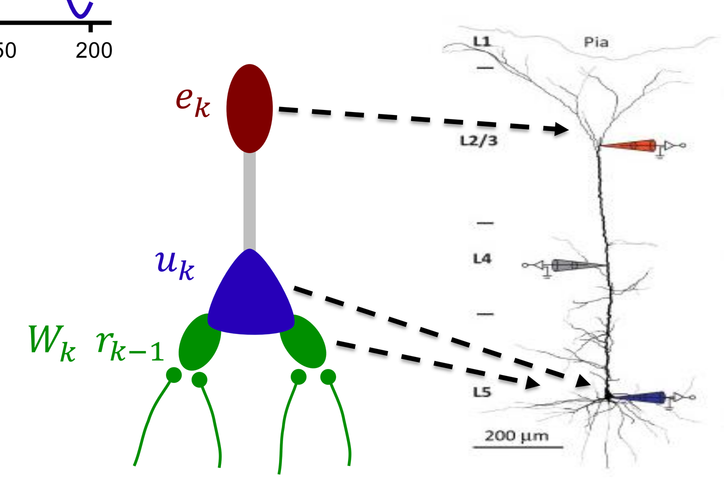

$$\dot W \sim -\nabla_W E \longrightarrow \dot W_k \sim \underbrace{(u_k - W_k \bar r_{k-1})\bar r_{k-1}^T}_{\text{Urbanczik-Senn (basal prediction of soma)}}$$

Look-ahead undoes low-pass filtering, allows time-continuous learning without phases.

$$r_k = \underbrace{\bar r_k + \tau \dot{\bar r}_k}_{\text{phase-advance of } \bar r_k = \varphi(u_k)}$$

$$\approx \bar r_k(t+\tau) \approx \underbrace{m_\infty^3 h(u, \dot u)}_{\text{sodium gating of HH neurons}}$$

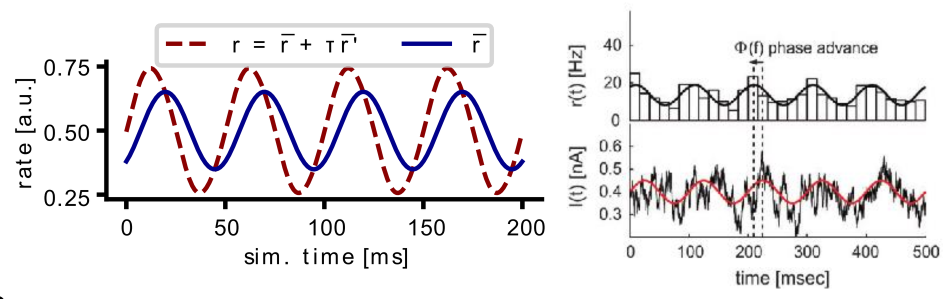

Prediction error encoded in apical dendrites:

$$\bar e_k = \bar r_k' \cdot W_{k+1}^T(u_{k+1} - W_{k+1}\bar r_k)$$

$$\sim W_{k+1}^T u_{k+1} - W_{k+1}^T W_{k+1} \bar r_k$$

$$= \underbrace{B_{k+1} u_{k+1}}_{\text{top-down feedback}} - \underbrace{W_k^{PI} W_k^{IP} \bar r_k}_{\text{bottom-up prediction}}$$

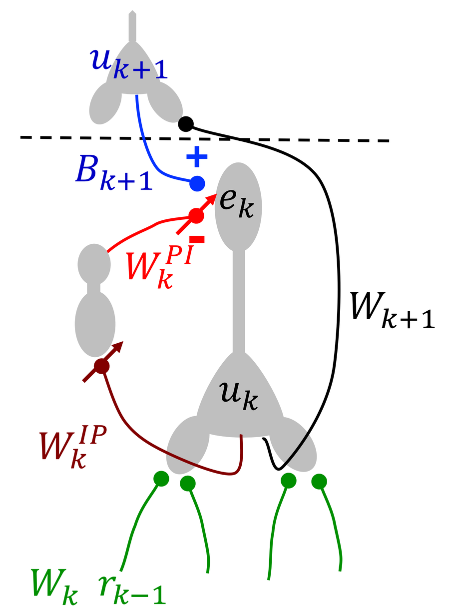

Combined neurosynaptic dynamics yields backprop

$$\dot W_k \sim \bar e_k \bar r_{k-1}^T$$

$$\bar e_k = \bar r_k' \cdot W_{k+1}^T \bar e_{k+1}$$

Illustration of error prop.

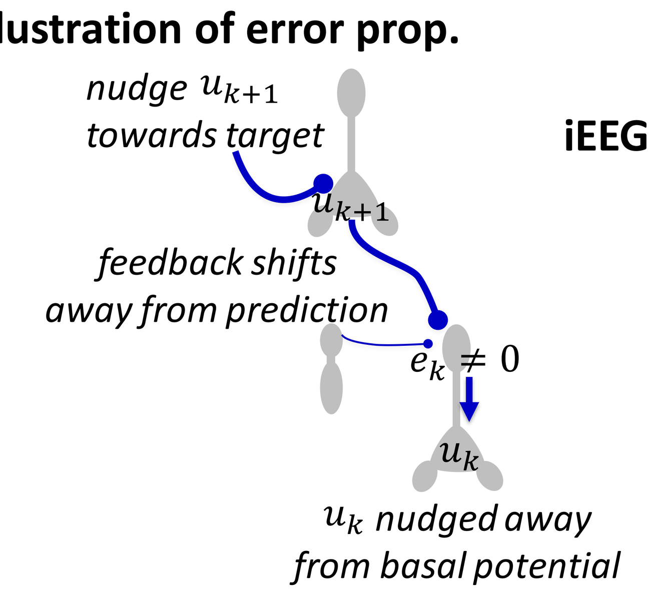

MNIST

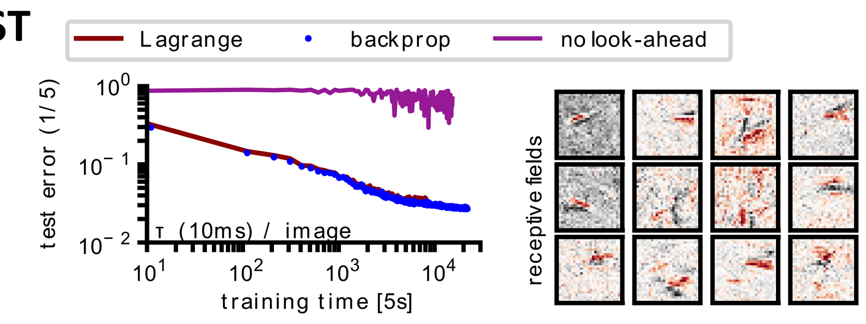

iEEG signals

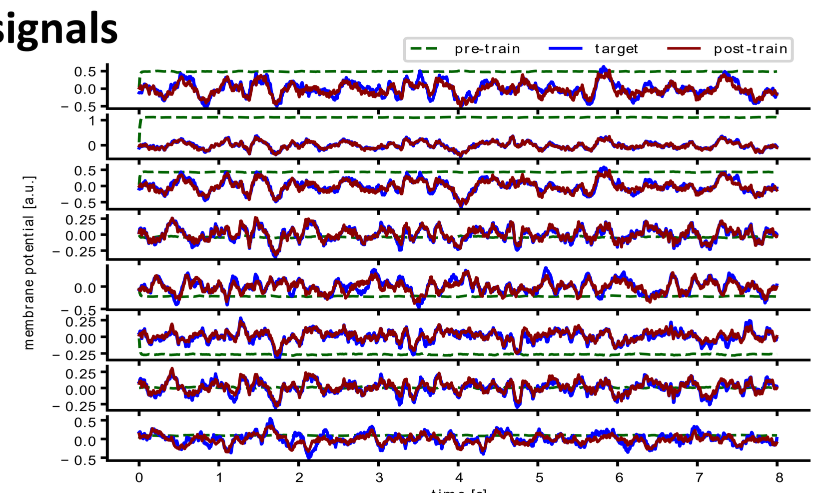

Dold, Dominik, et al. **Stochasticity from function-why the Bayesian brain may need no noise**. *arXiv preprint arXiv:1809.08045* (2018).

Dold, Dominik, et al. **Lagrangian dynamics of dendritic microcircuits enables real-time backpropagation of errors**. Cosyne Abstracts 2019, Lisbon, PT.

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

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