Computational Neuroscience Group
Department of Physiology
University of Bern

Computers like brains

Models for computation with bio-inspired neuronal substrates

Mihai A. Petrovici

Electronic Vision(s)
Kirchhoff Institute for Physics
University of Heidelberg

Physics vs. neuroscience

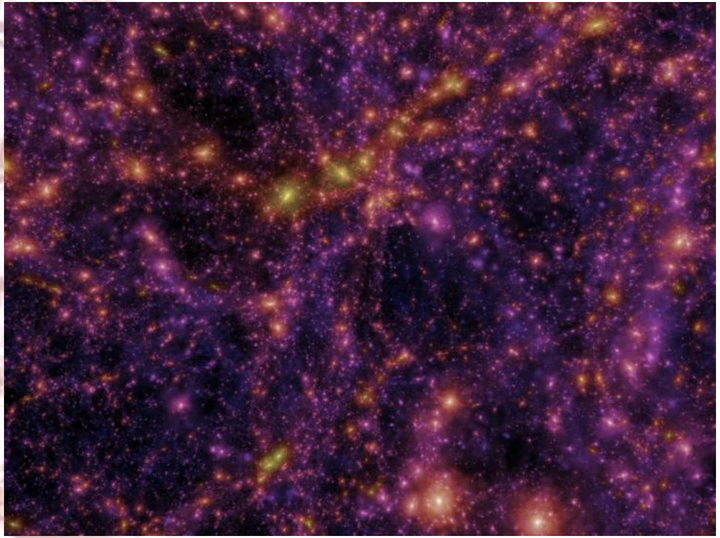

- huge particle numbers (>> N_A) → thermodynamics !
- largely identical particles
- relatively simple interaction, often short-range
- amenable to mean-field approaches

Physics vs. neuroscience

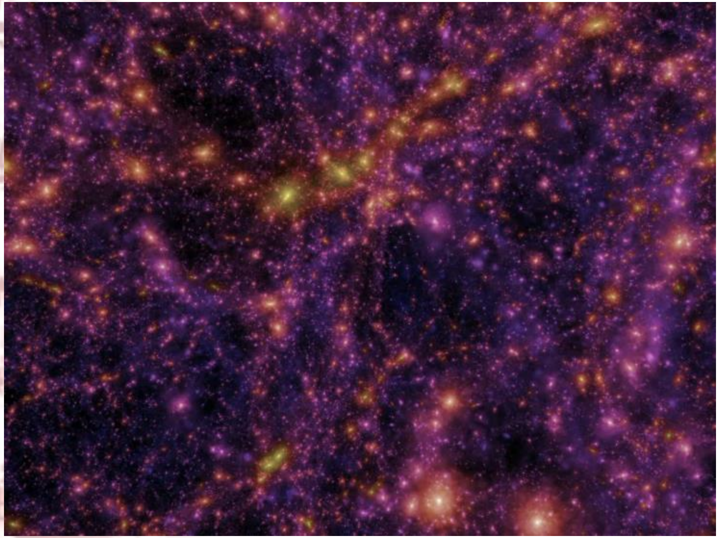

- huge particle numbers ($\gg N_A$) $\rightarrow$ thermodynamics !
- largely identical particles
- relatively simple interaction, often short-range
- amenable to mean-field approaches

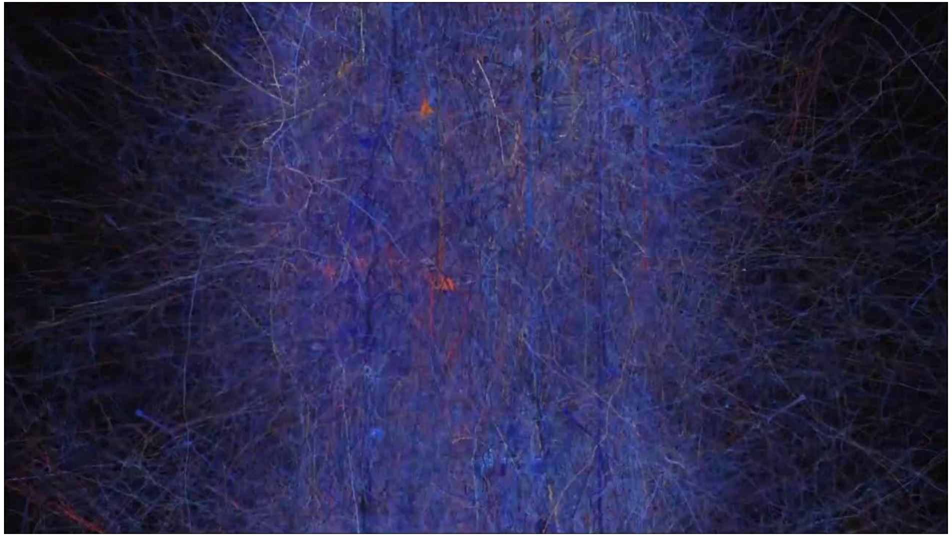

- not necessarily many „particles“ (neurons), but
- highly diverse, with
- long-range, time-dependent, complex interaction
- mean-field methods hide most of the interesting stuff

Simulating the brain

"The rule of simulation that I would like to have is that the number of computer elements required to simulate a large physical system is only to be proportional to the space-time volume of the physical system."

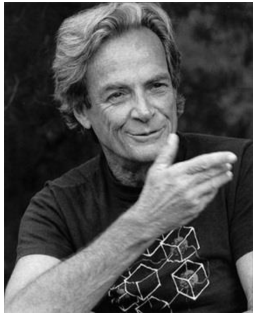

Simulating the brain

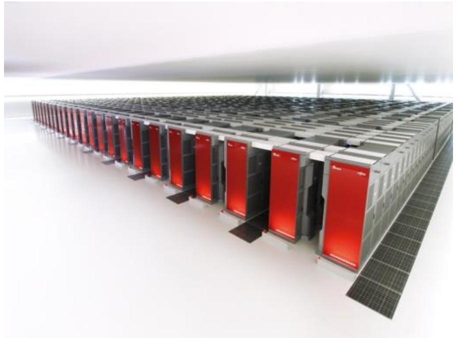

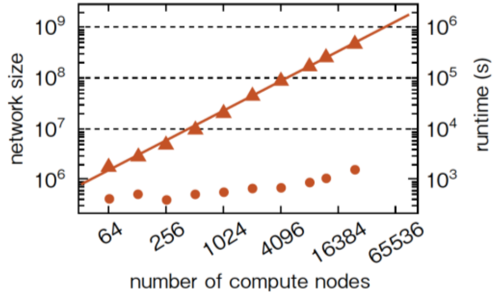

Diesmann (2012)

simulation time vs. biological real time: 1520:1

	nature	simulation
synaptic plasticity	seconds	hours
learning	days	years
development	years	millennia
evolution	> millennia	> millions of years

"The rule of simulation that I would like to have is that the number of computer elements required to simulate a large physical system is only to be proportional to the space-time volume of the physical system."

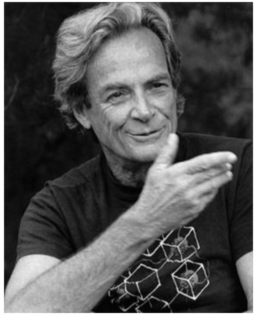

Simulating the brain

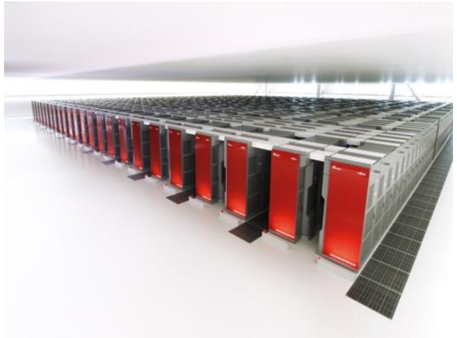

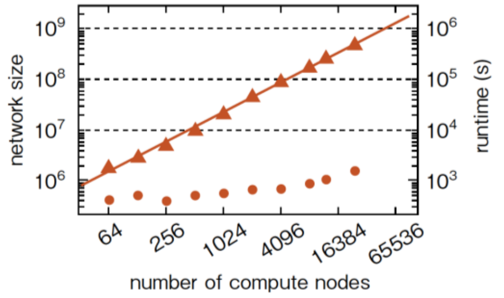

Diesmann (2012)

simulation time vs. biological real time: 1520:1

	nature	simulation
synaptic plasticity	seconds	hours
learning	days	years
development	years	millennia
evolution	> millennia	> millions of years

"The rule of simulation that I would like to have is that the number of computer elements required to simulate a large physical system is only to be proportional to the space-time volume of the physical system."

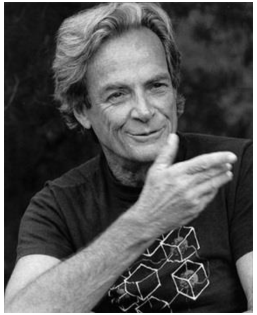

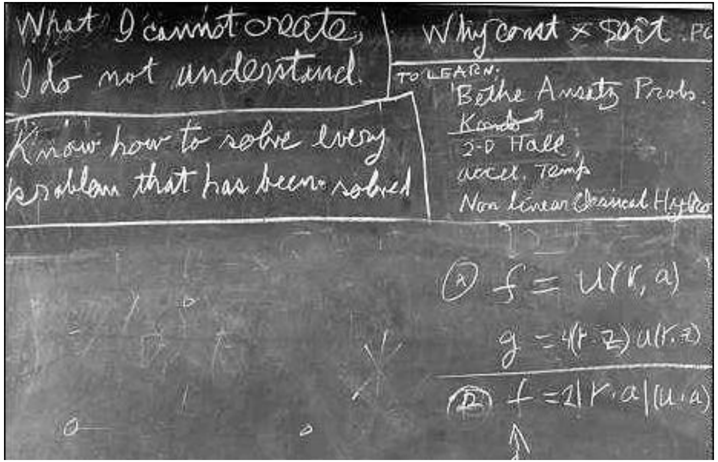

Caltech archives

Biological neurons

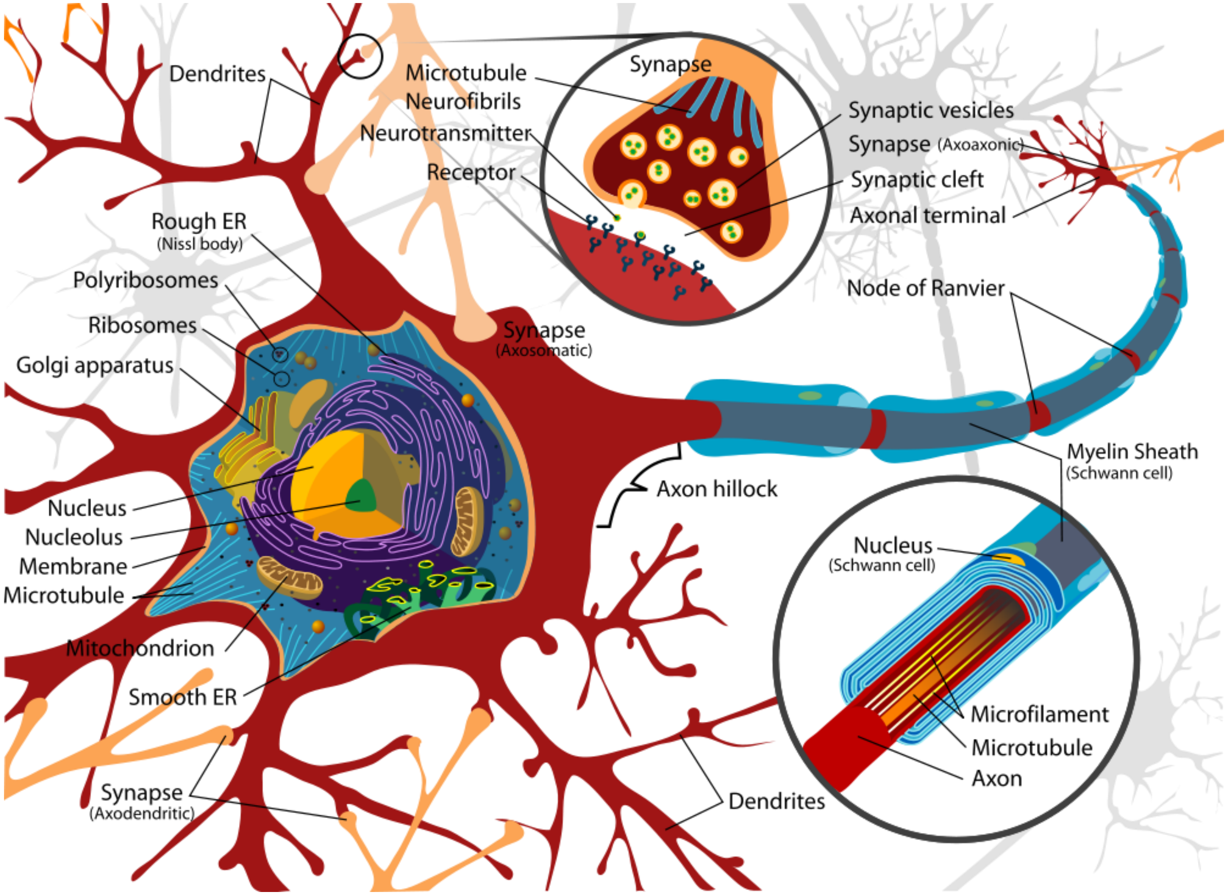

Wikipedia

From (too?) complex to (too?) simple

Biology

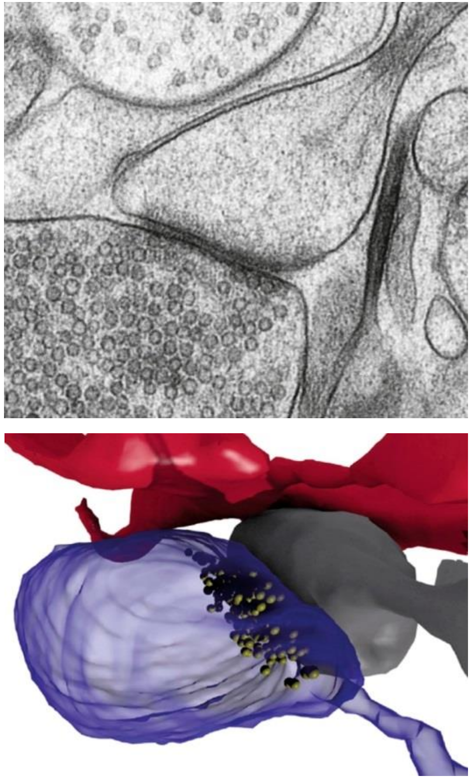

Korogod et al. (2015)

From (too?) complex to (too?) simple

Biology

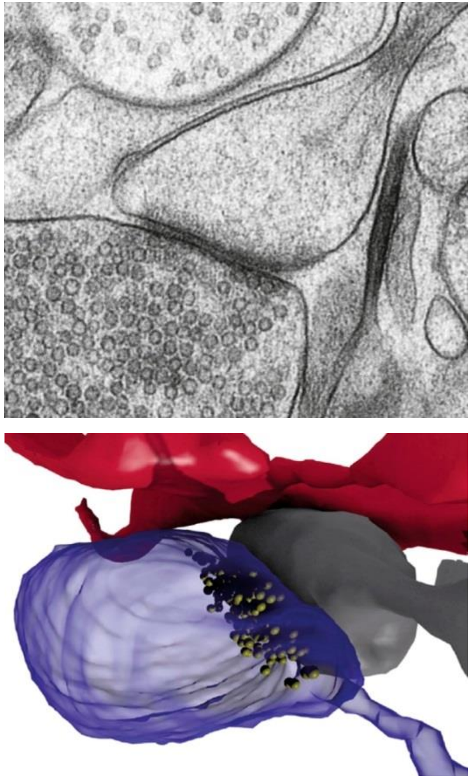

Korogod et al. (2015)

Theoretical neuroscience

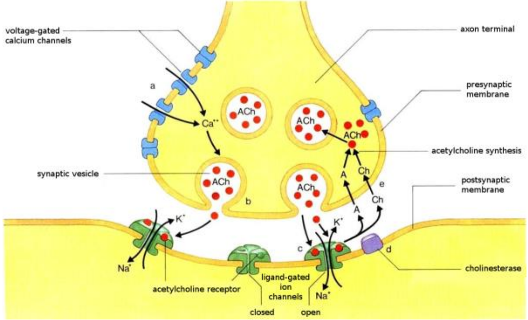

$$\frac{dR}{dt} = \frac{I}{\tau_{\text{rec}}} - \sum_{\text{spikes } s} UR\delta(t - t_s)$$

$$\frac{dE}{dt} = -\frac{E}{\tau_{\text{inact}}} + \sum_{\text{spikes } s} UR\delta(t - t_s)$$

$$\frac{dU}{dt} = \frac{U_0 - U}{\tau_{\text{facil}}} + \sum_{\text{spikes } s} U_0(1 - U)\delta(t - t_s)$$

From (too?) complex to (too?) simple

Biology

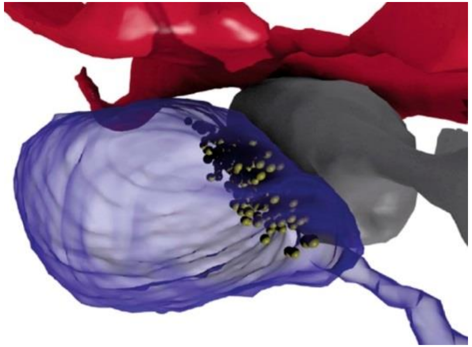

Korogod et al. (2015)

Theoretical neuroscience

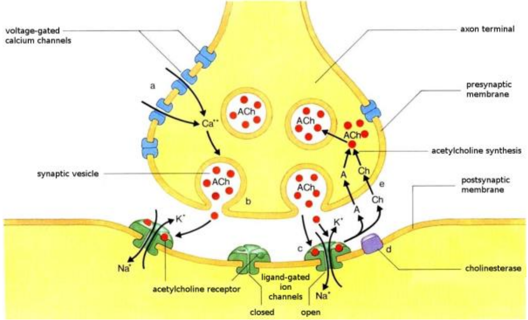

$$\frac{dR}{dt} = \frac{I}{\tau_{\text{rec}}} - \sum_{\text{spikes } s} UR\delta(t - t_s)$$

$$\frac{dE}{dt} = -\frac{E}{\tau_{\text{inact}}} + \sum_{\text{spikes } s} UR\delta(t - t_s)$$

$$\frac{dU}{dt} = \frac{U_0 - U}{\tau_{\text{facil}}} + \sum_{\text{spikes } s} U_0(1 - U)\delta(t - t_s)$$

Machine learning

$$w_{ij}$$

From neurons to VLSI

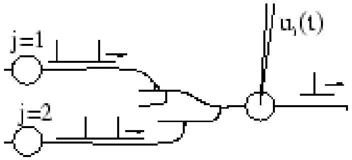

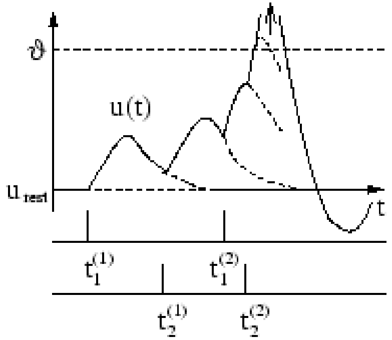

$$C_m\dot{u} = g_L(E_L - u) + g_{syn}\left(E_{syn} - u\right) + g_L\Delta_T \exp\left(\frac{V - V_T}{\Delta_T}\right) - w$$

$$\text{if } u = \vartheta \rightarrow \begin{cases} u \rightarrow E_r \\ w \rightarrow w + b \end{cases}$$

$$\tau_w \dot{w} = a(V - E_L) - w$$

From neurons to VLSI

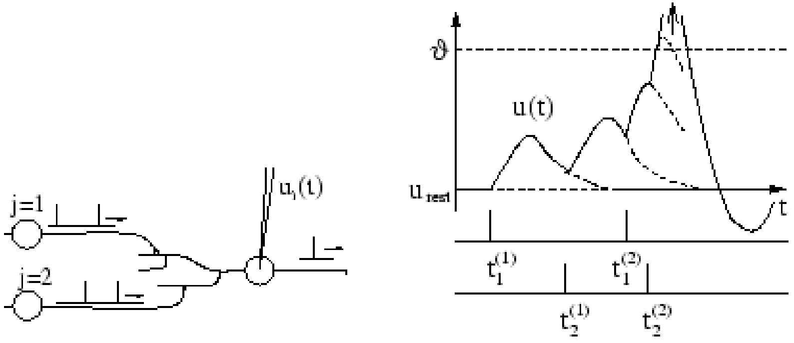

$$C_m \dot{u} = g_L(E_L - u) + g_{\mathrm{syn}}(E_{\mathrm{syn}} - u) + g_L \Delta_T \exp\left(\frac{V - V_T}{\Delta_T}\right) - w$$

$$\text{if } u = \vartheta \rightarrow \begin{cases} u \rightarrow E_r \\ w \rightarrow w + b \end{cases}$$

$$\tau_w \dot{w} = a(V - E_L) - w$$

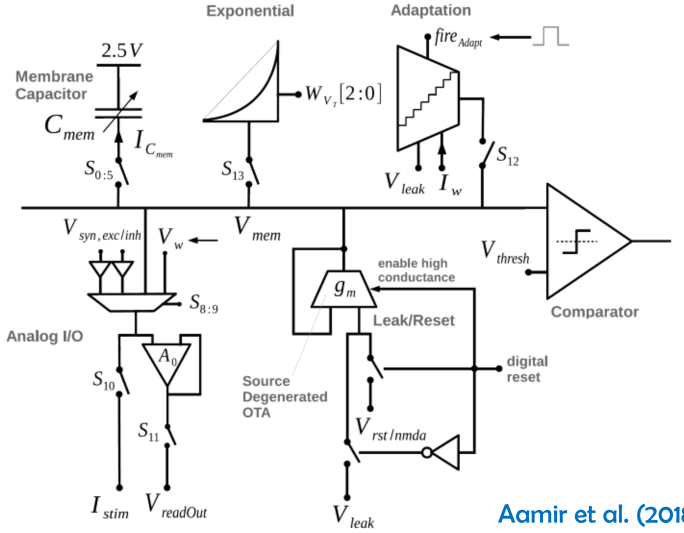

Aamir et al. (2018)

- mixed-signal VLSI:

  membrane → analog

  spikes → digital

- inherent speedup: $10^3 - 10^5$

Chip layout

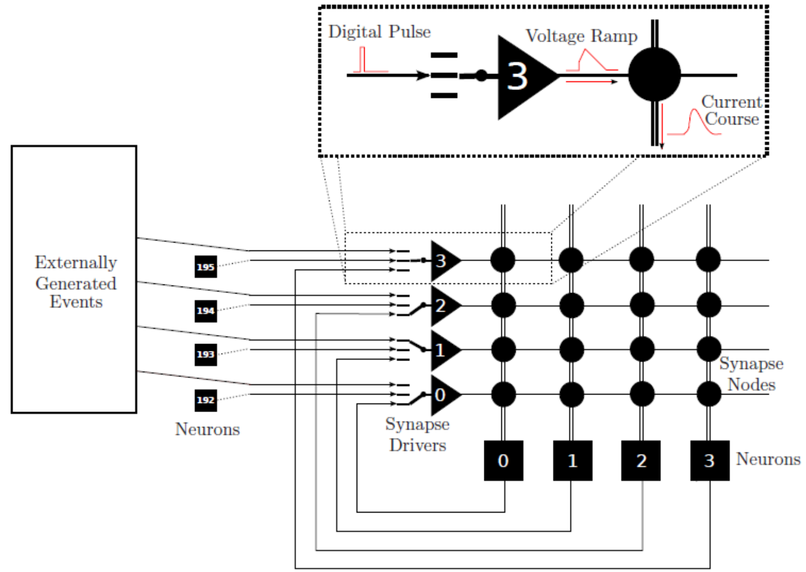

Pfeil et al. (2013)

Wafer-scale integration

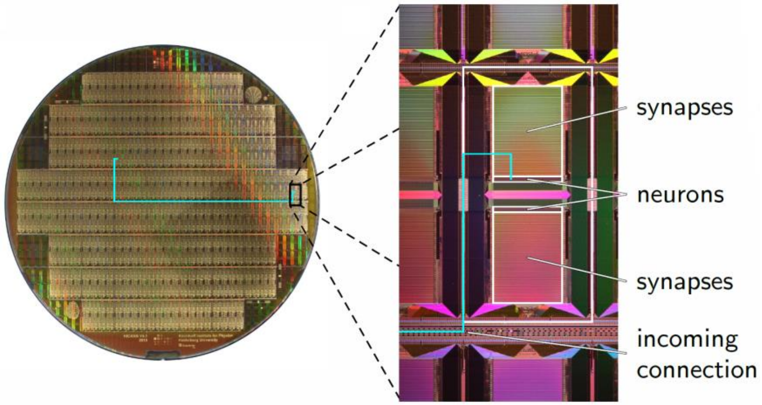

Schemmel et al. (2010)

Wafer-scale integration

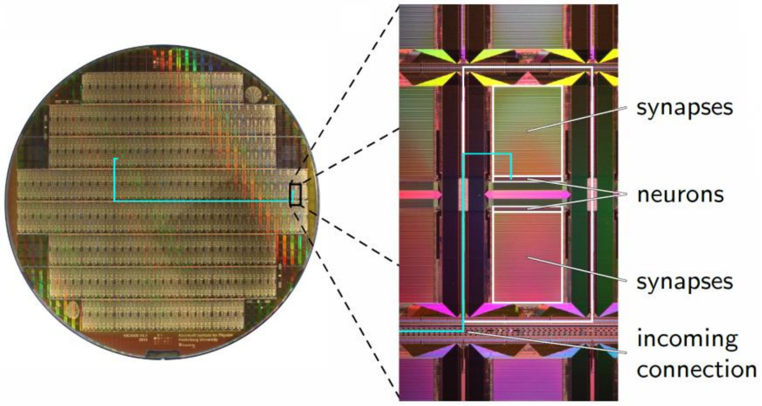

Schemmel et al. (2010)

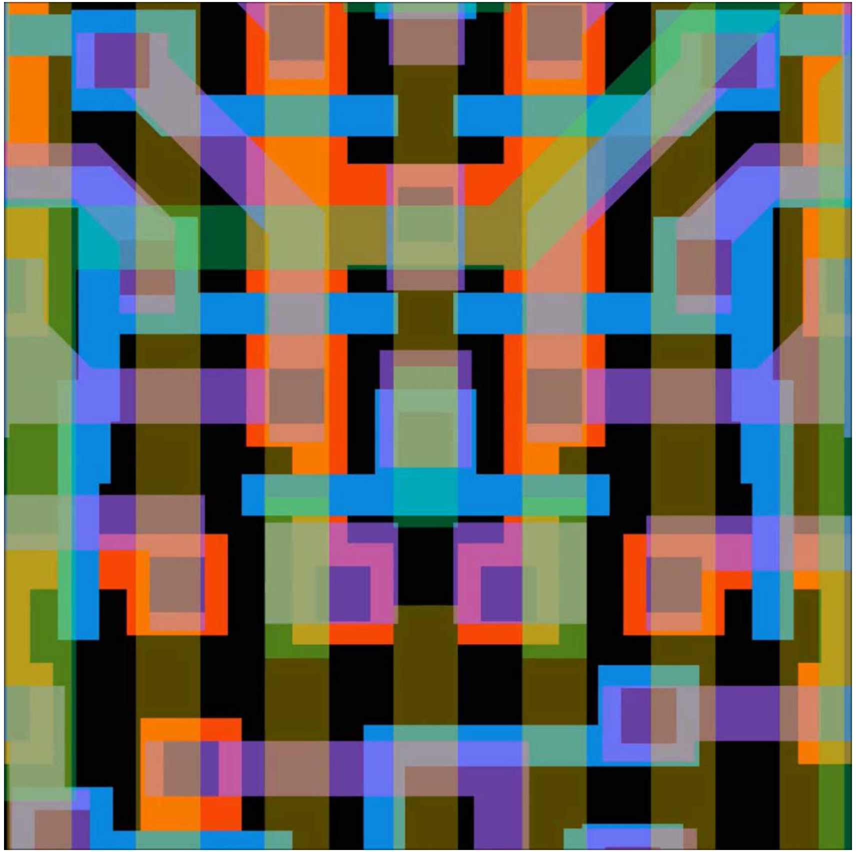

Wafer-scale integration

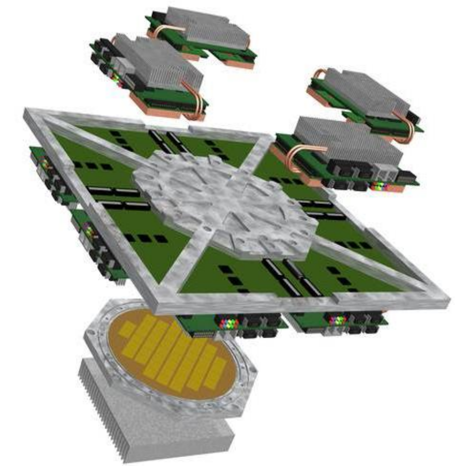

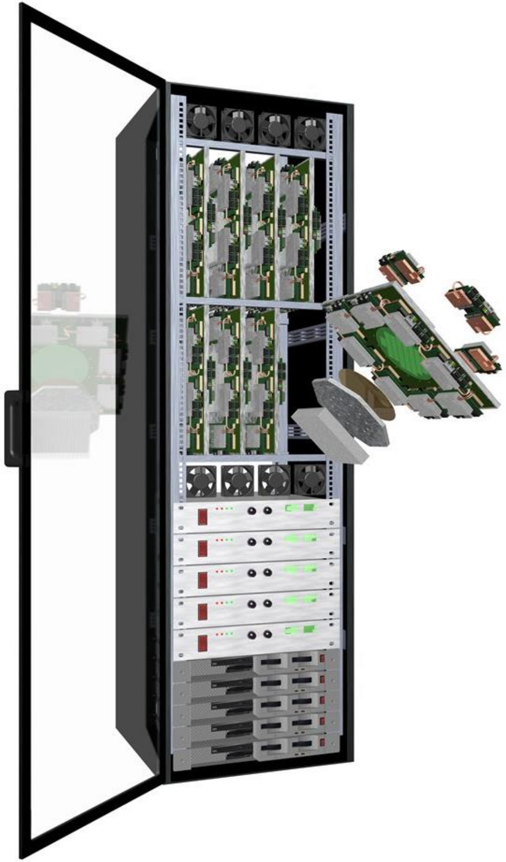

A back-end-independent high-level API

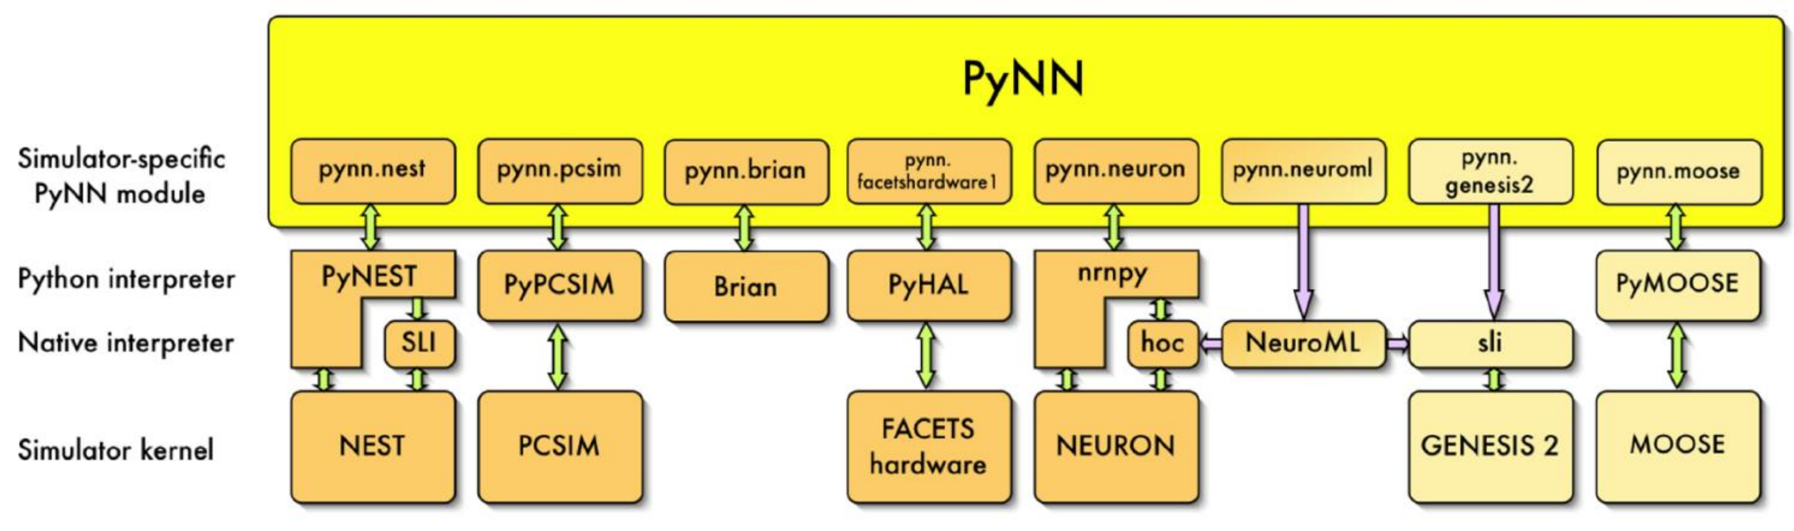


```
>>> p1 = Population(100, IF_curr_alpha, structure=space.Grid2D())
>>> p2 = Population(20, IF_curr_alpha, cellparams={'tau_m': 15.0, 'cm': 0.9})
>>> pulse = DCSource(amplitude=0.5, start=20.0, stop=80.0)
>>> pulse.inject_into(p1[3:7])
>>> prj2_1 = Projection(p2, p1, method=AllToAllConnector(), target='excitatory')
>>> run(1000.0)
>>> p1.printSpikes("spikefile.dat")
```

Davison et al. (2009)

Some examples

L2/3 cortical attractor memory

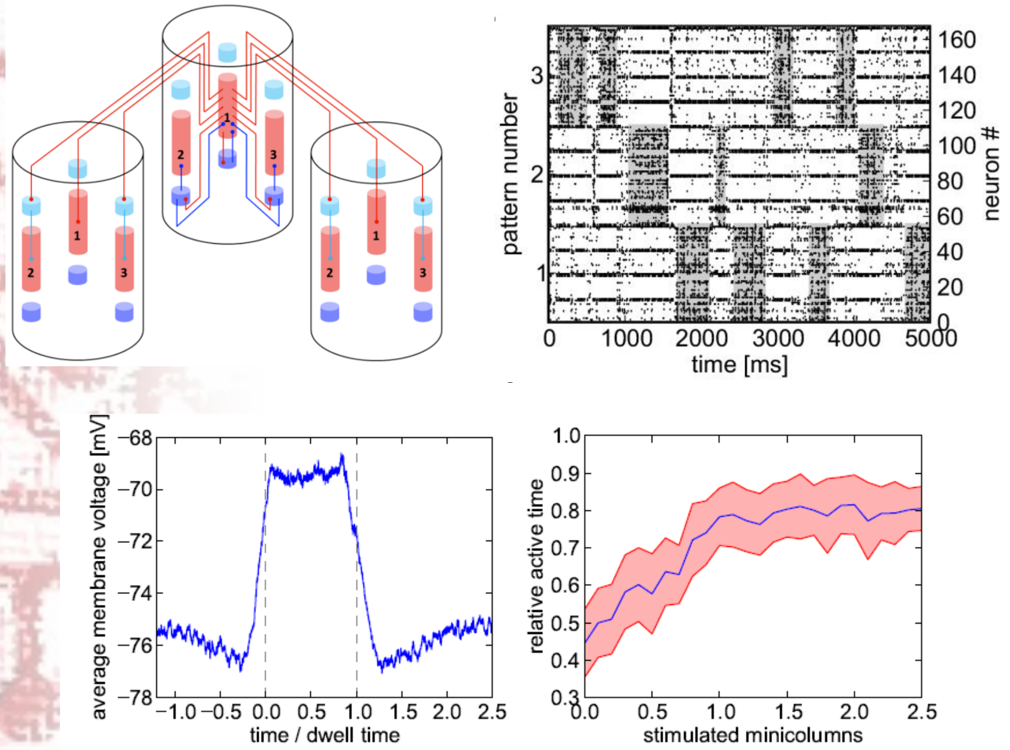

Classification with a spiking liquid

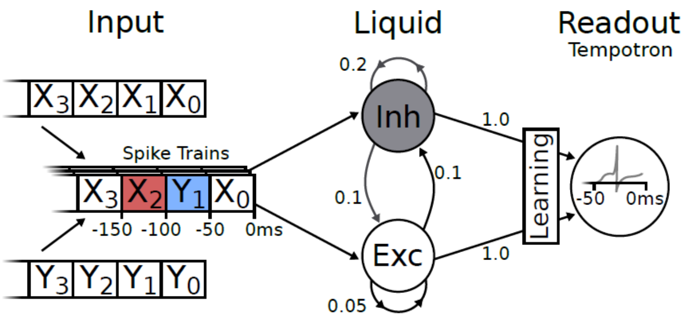

Pfeil & Grübl & Jeltsch & Müller & Petrovici & Schmuker et al. (2013)

Noise as a resource for computation

Bayesian inference & sampling

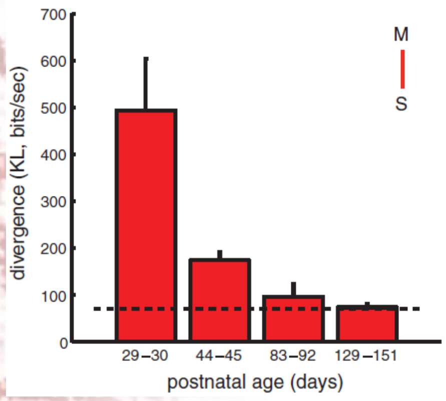

Berkes et al. (2011)

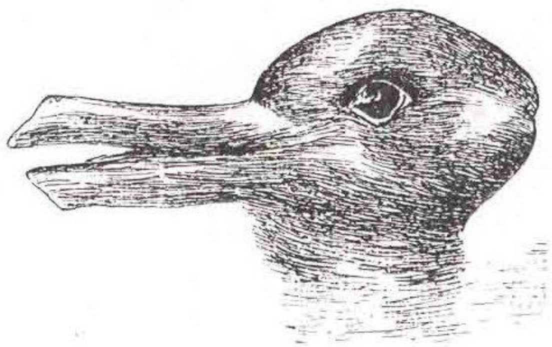

Bayesian inference & sampling

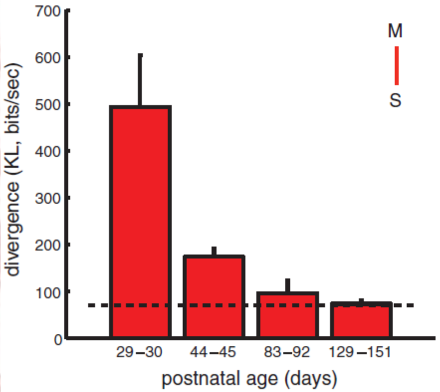

Berkes et al. (2011)

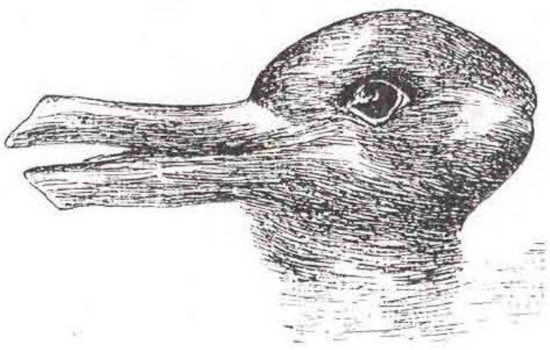

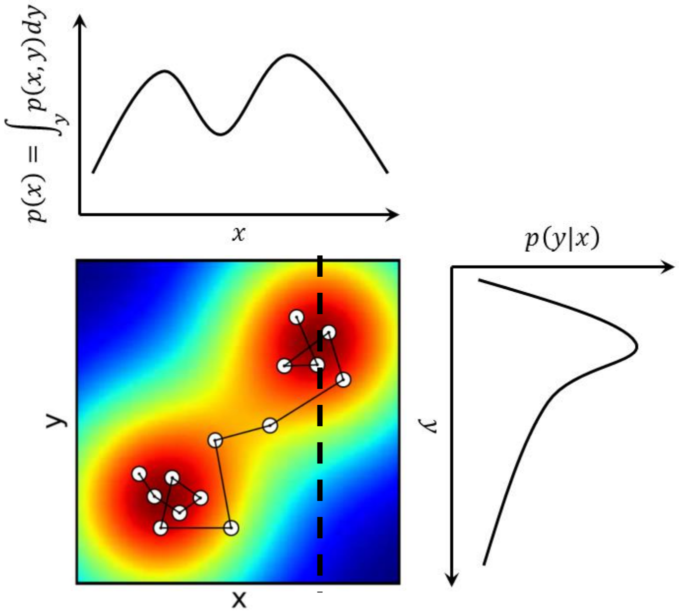

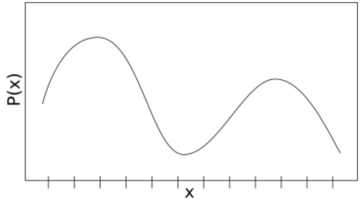

Sampling in spiking networks: single neuron dynamics

$$p(\boldsymbol{z}) = \frac{1}{Z} \exp\left[\frac{1}{2}\boldsymbol{z}^T \boldsymbol{W} \boldsymbol{z} + \boldsymbol{z}^T \boldsymbol{b}\right]$$

$$u_k = \log \frac{p(z_k = 1 \,|\boldsymbol{z}_{\setminus k})}{p(z_k = 0 \,|\boldsymbol{z}_{\setminus k})}$$

$$u_k = \sum_{i=1}^{K} W_{ki} z_i + b_k$$

$z_k = 1 \;\Leftrightarrow\;$ neuron has spiked in $[t - \tau, t)$

$\rightarrow$ spike pattern encodes states $\boldsymbol{z}^{(t)}$

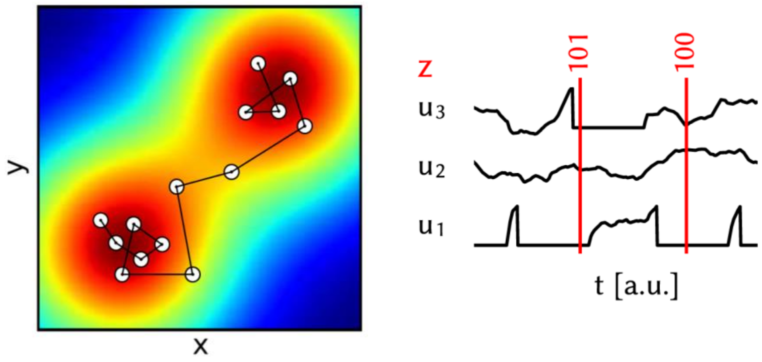

Petrovici & Bill et al. (2013, 2016), Petrovici et al. (2015)

Sampling in spiking networks: single neuron dynamics

$$p(\boldsymbol{z}) = \frac{1}{Z} \exp\left[\frac{1}{2}\boldsymbol{z}^T \boldsymbol{W} \boldsymbol{z} + \boldsymbol{z}^T \boldsymbol{b}\right]$$

$$u_k = \log \frac{p(z_k = 1 \,|\boldsymbol{z}_{\backslash k})}{p(z_k = 0 \,|\boldsymbol{z}_{\backslash k})}$$

$$u_k = \sum_{i=1}^{K} W_{ki} z_i + b_k$$

$z_k = 1 \;\Leftrightarrow\;$ neuron has spiked in $[t - \tau, t)$

$\rightarrow$ spike pattern encodes states $\boldsymbol{z}^{(t)}$

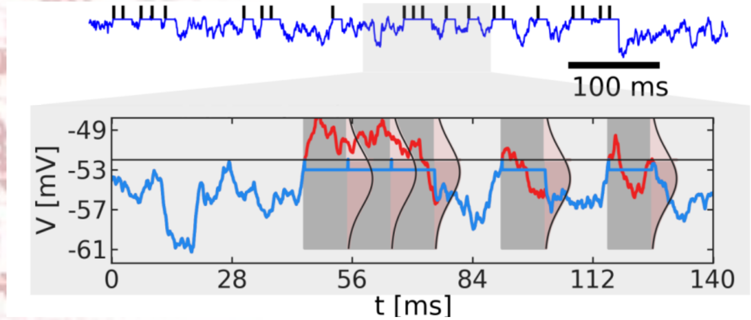

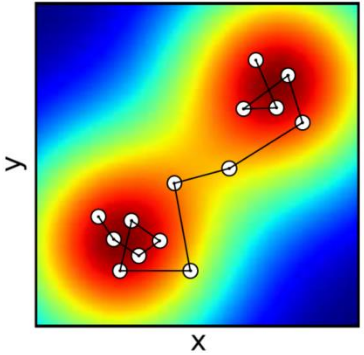

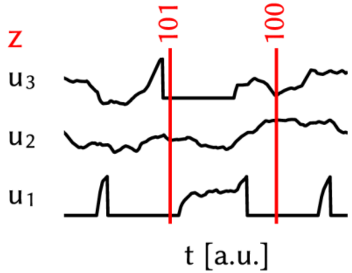

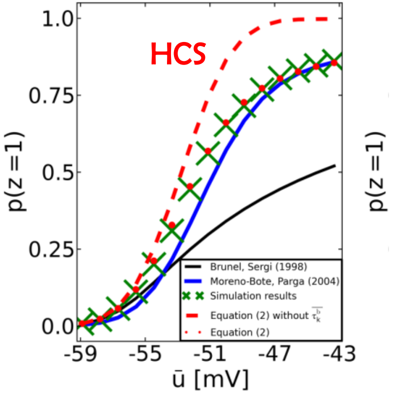

$$p(z_k = 1) = \frac{\sum_n P_n \, n \, \tau_{\text{ref}}}{\sum_n P_n \cdot \left(n \, \tau_{\text{ref}} + \sum_{k=1}^{n-1} \overline{\tau_k^b} + T_n\right)} \approx \frac{1}{1 + \exp(-u_k)}$$

Petrovici & Bill et al. (2013, 2016), Petrovici et al. (2015)

Generative & discriminative models of visual data

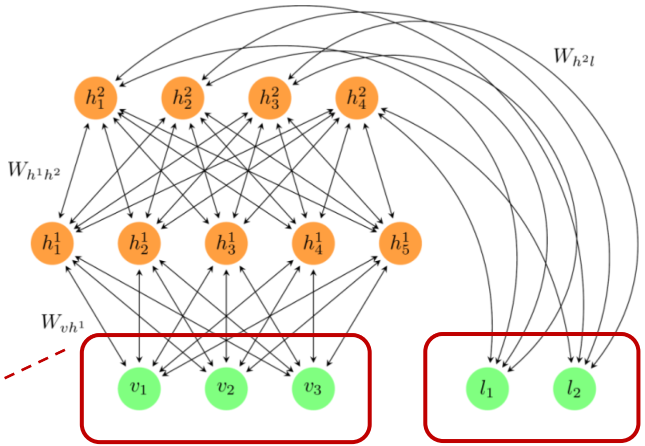

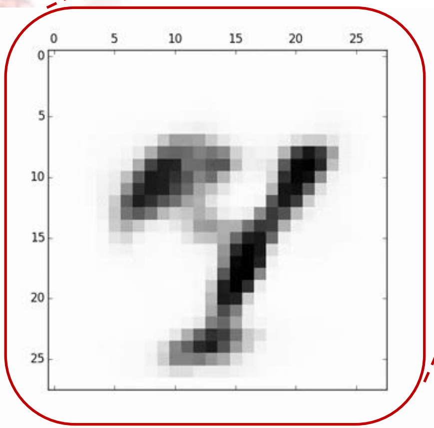

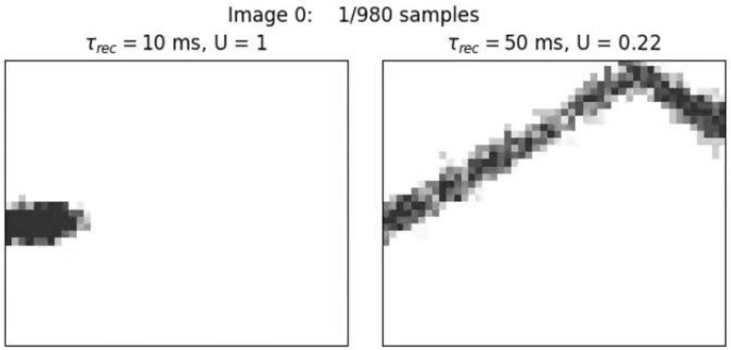

97% correct classification
with < 2000 neurons

$$\Delta w_{ij} \propto \langle z_i z_j \rangle_{\text{data}} - \langle z_i z_j \rangle_{\text{model}}$$

$$\Delta b_i \propto \langle z_i \rangle_{\text{data}} - \langle z_i \rangle_{\text{model}}$$

Leng & Petrovici et al. (2016, 2018), Zenk (2018), Kungl et al. (in prep.)

Biological mechanisms for superior mixing

cortical oscillations

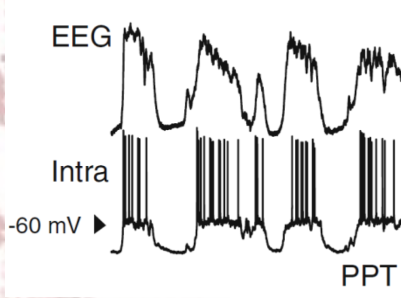

short-term synaptic plasticity

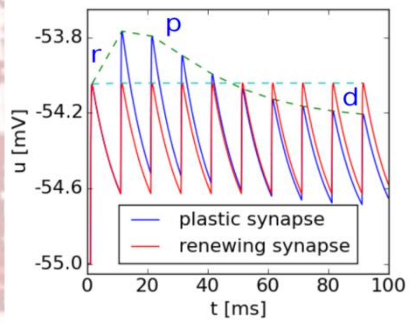

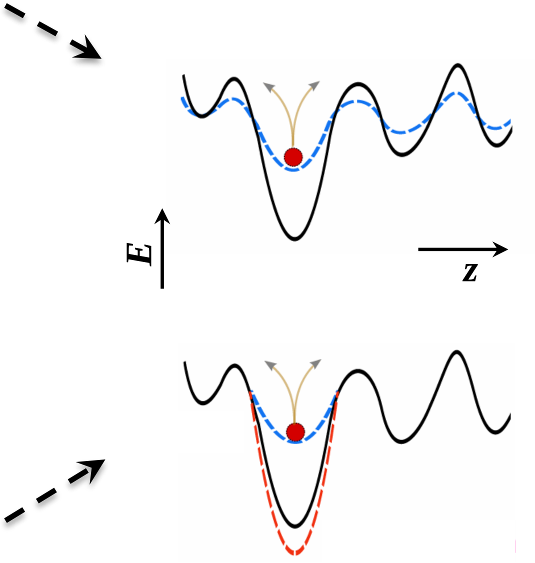

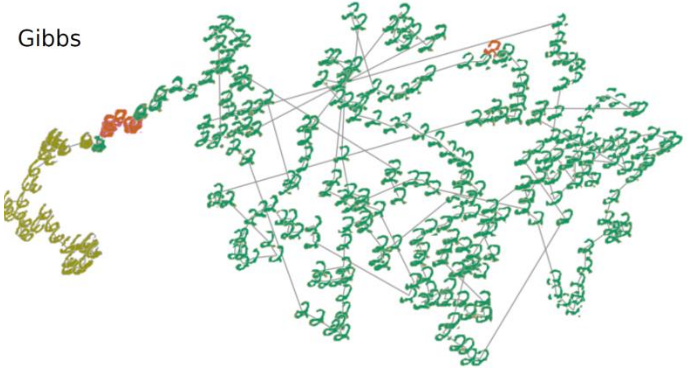

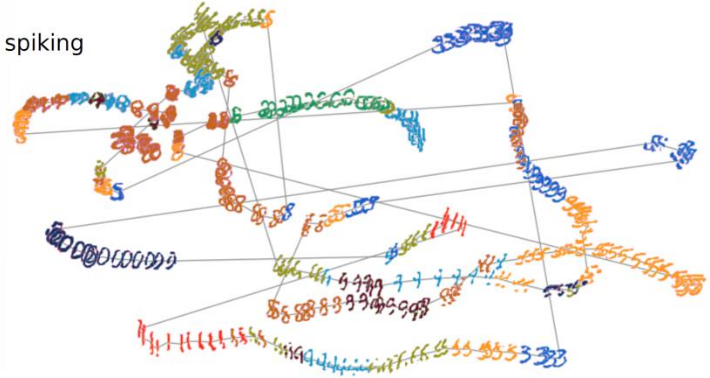

Leng & Petrovici et al. (2016, 2018), Korcsak-Gorzo & Baumbach & Petrovici et al., in prep.

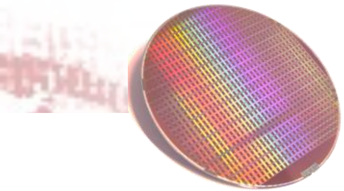

Physical stochastic computation without noise

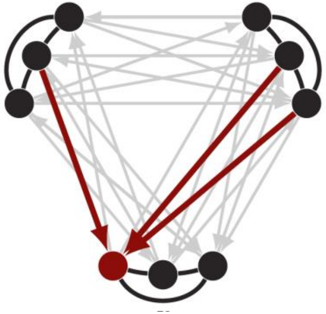

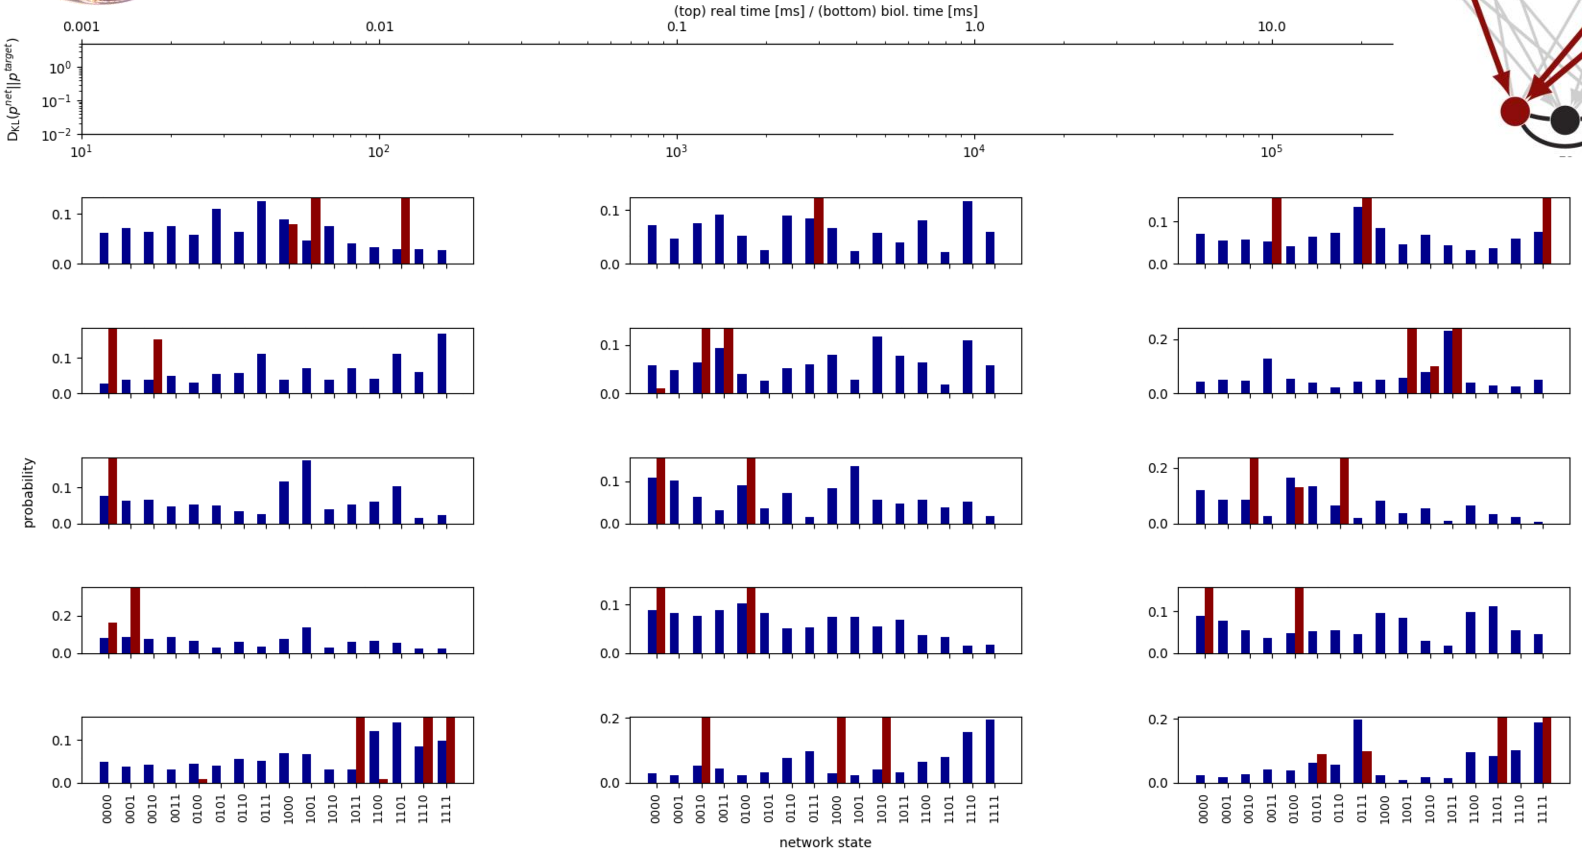

Petrovici & Stöckel et al. (2015), Jordan et al. (2017), Dold & Bytschok & Petrovici et al. (2018), Kungl et al. (in prep.)

Stochasticity from function: Bayesian inference

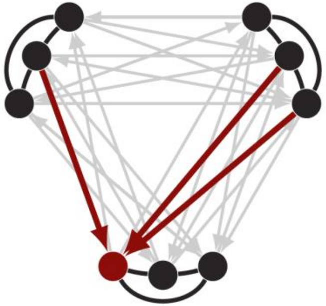

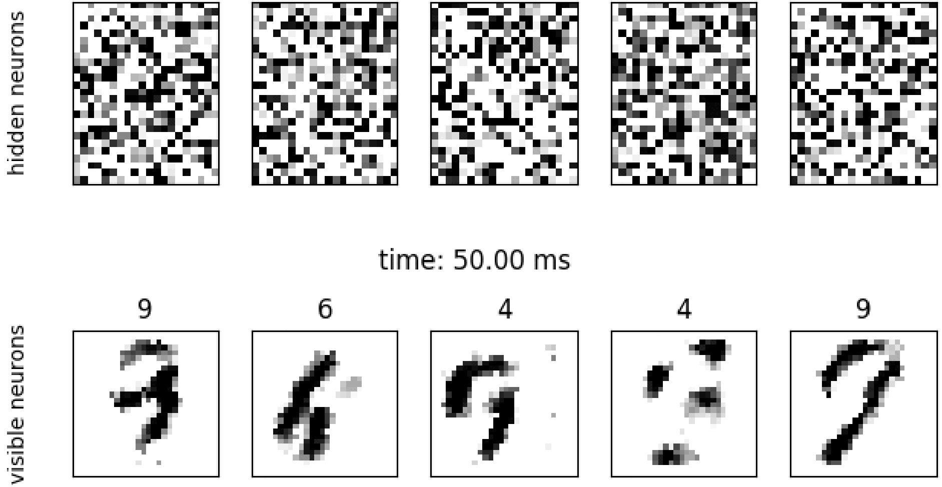

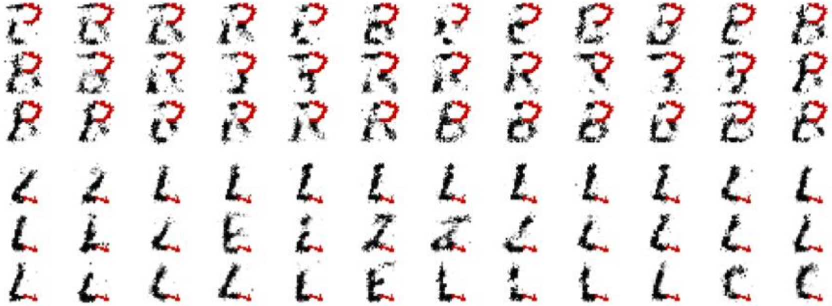

Dold & Bytschok & Petrovici et al. (2018)

Ensemble dynamics & statistics

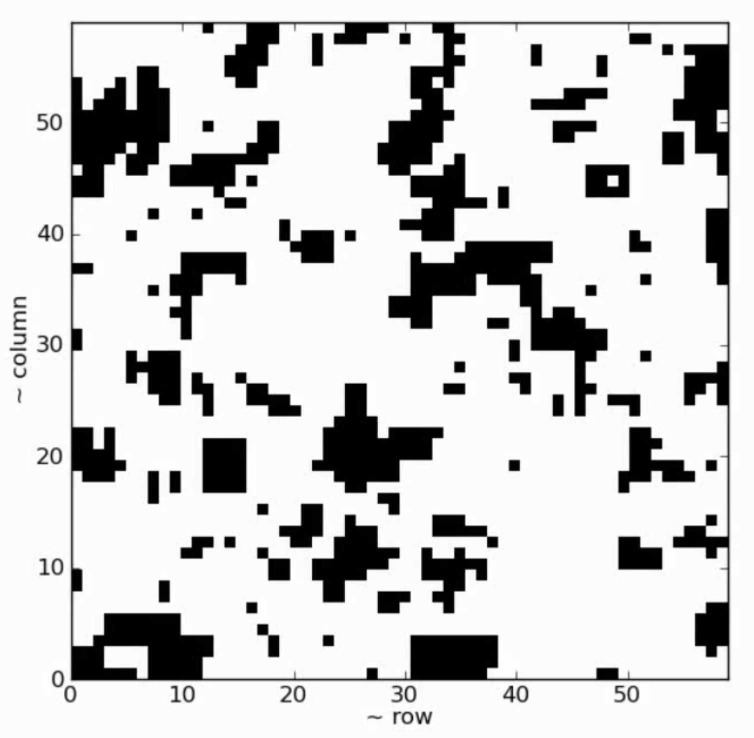

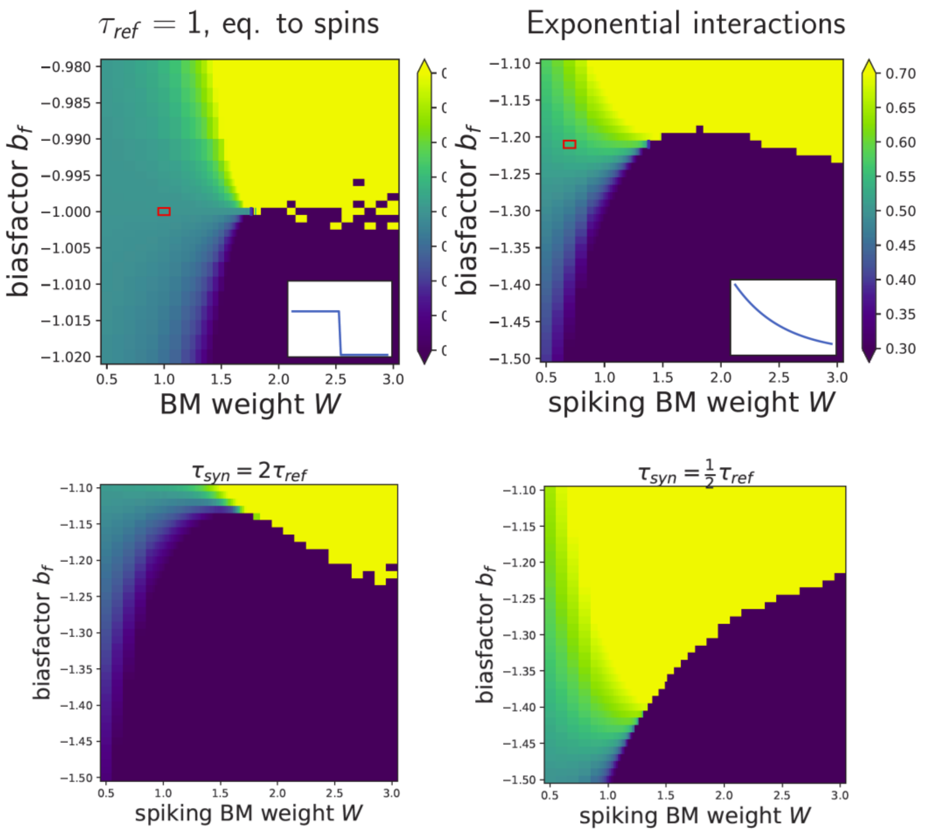

Baumbach et al., in prep.

Which way is down?

Lagrangian mechanics

principle of (least) stationary action

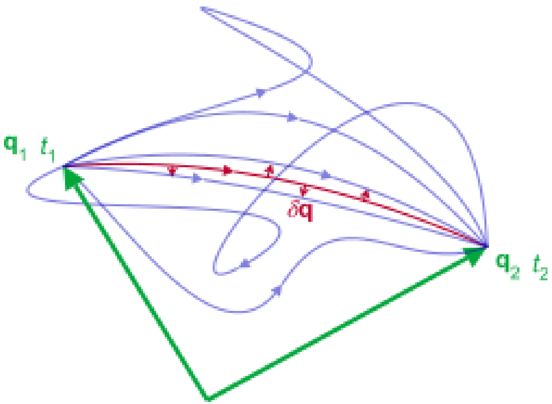

$$\delta\left(\int dt\, L(\boldsymbol{q},\dot{\boldsymbol{q}})\right) = 0$$

$$\Downarrow$$

$$\frac{\partial L}{\partial q_i} - \frac{d}{dt}\frac{\partial L}{\partial \dot{q}_i} = 0$$

Euler-Lagrange equations of motion

fundamental principle in

→ mechanics
→ geometrical optics
→ electrodynamics
→ quantum mechanics
→ ...
→ neurobiology?

Lagrangian mechanics for neuronal networks

$$E(\boldsymbol{u}) = \sum_i \frac{1}{2}\|\boldsymbol{u}_i - \boldsymbol{W}_i\bar{\boldsymbol{r}}_{i-1}\|^2 + \beta\frac{1}{2}\left\|\boldsymbol{u}_N - \boldsymbol{u}_N^{\mathrm{tgt}}\right\|^2 \xrightarrow{\boldsymbol{u} = \tilde{\boldsymbol{u}} - \tau\dot{\tilde{\boldsymbol{u}}}} L(\tilde{\boldsymbol{u}}, \dot{\tilde{\boldsymbol{u}}})$$

$$\frac{\partial L}{\partial \tilde{\boldsymbol{u}}_i} - \frac{d}{dt}\frac{\partial L}{\partial \dot{\tilde{\boldsymbol{u}}}_i} = 0$$

$$\dot{\boldsymbol{W}}_i = -\eta\frac{\partial E}{\partial \boldsymbol{W}_i}$$

$$\Downarrow$$

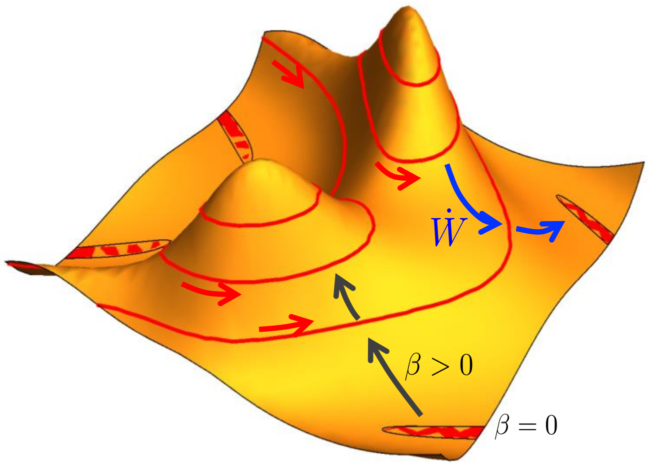

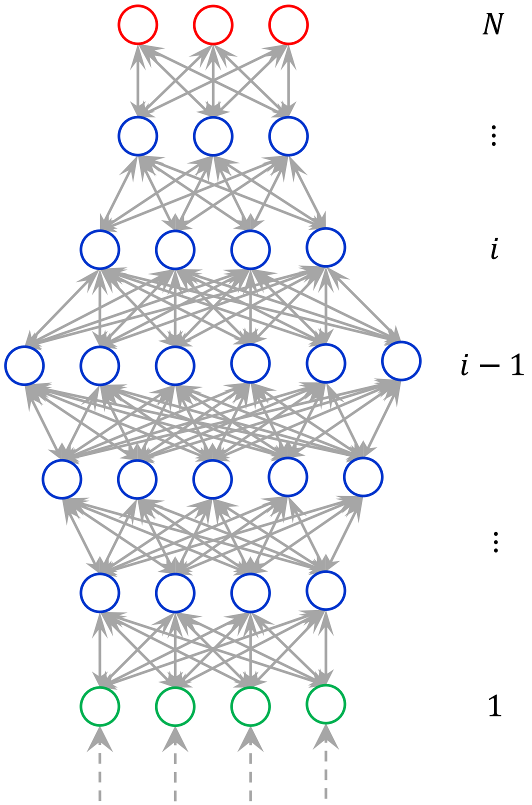

$\tau\dot{\boldsymbol{u}}_i = \boldsymbol{W}_i\boldsymbol{r}_{i-1} - \boldsymbol{u}_i + \boldsymbol{e}_i$ → neuron dynamics!

$\bar{\boldsymbol{e}}_i = \bar{\boldsymbol{r}}_i' \odot \left[\boldsymbol{W}_{i+1}^T(\boldsymbol{u}_{i+1} - \boldsymbol{W}_{i+1}\bar{\boldsymbol{r}}_i)\right]$ → error backprop!

$\dot{\boldsymbol{W}}_i = \eta(\boldsymbol{u}_i - \boldsymbol{W}_i\bar{\boldsymbol{r}}_{i-1})\bar{\boldsymbol{r}}_{i-1}^T$ → Urbanczik-Senn learning rule!

Dold et al., in prep.

Biophysical implementation

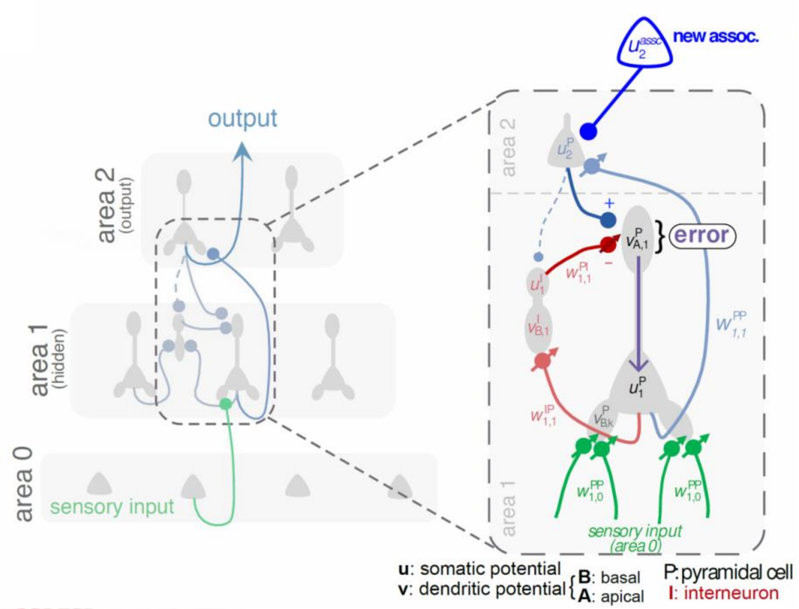

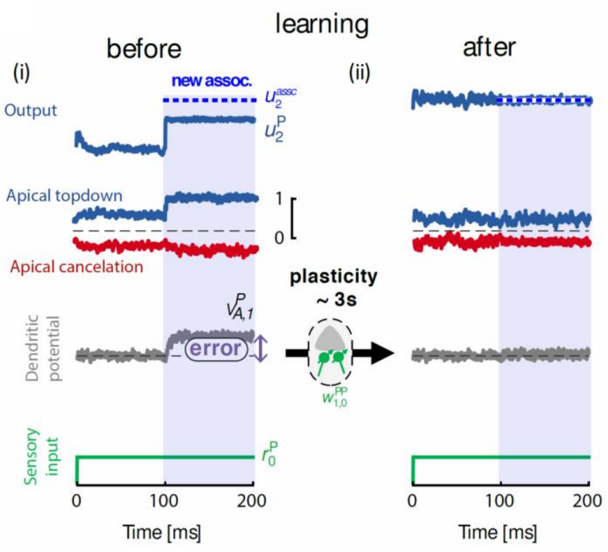

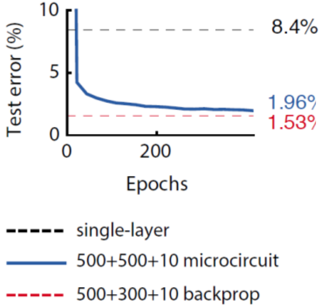

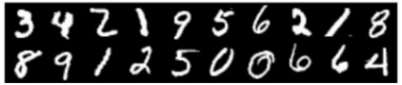

→ backpropagation of errors similar to feedback alignment

→ however, continuous dynamics! (no ff/fb phases)

Sacramento et al. (2017), Dold et al. (in prep.)

Outlook

Experiments

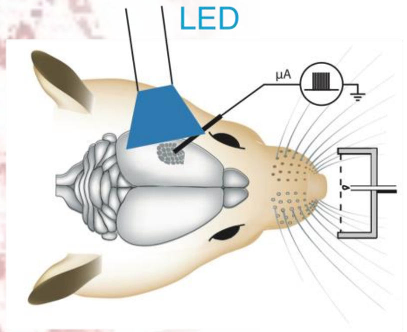

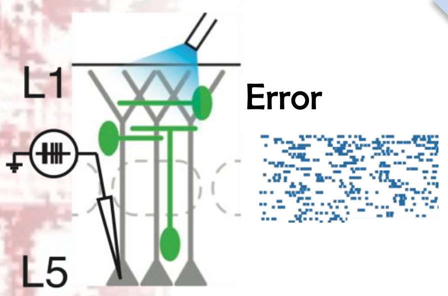

Theory

$$E(u,W)=\sum_{i=1}^{N}\underbrace{\frac{1}{2}(u_i - W_i\phi(u_{i-1}))^2}_{\text{Error}}$$

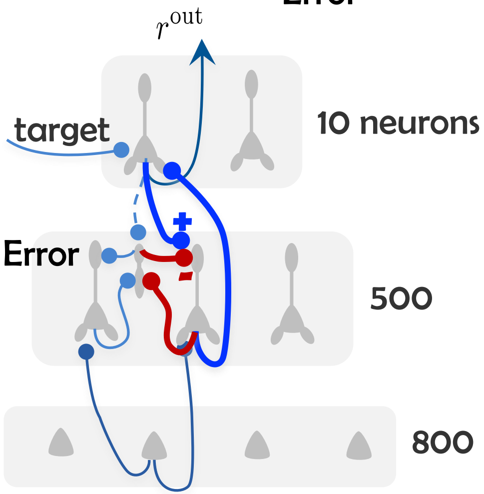

Hardware

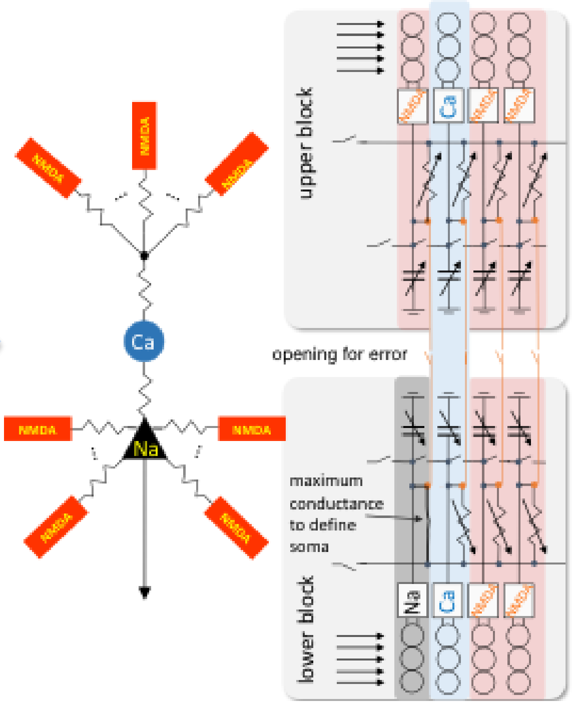

Some of the neural networks behind our neural networks

Dominik Dold
Ilja Bytschok
Ajos Kungl
Andreas Baumbach
Oliver Breitwieser
Walter Senn
Johannes Schemmel
Karlheinz Meier
Jakob Jordan
Markus Diesmann
Tom Tetzlaff
Sebastian Schmitt
Johann Klähn
David Stöckel
Anna Schröder
Johannes Bill
Andreas Grübl
Maurice Güttler
Andreas Hartel
Dan Husmann
Sebastian Jeltsch
Vitali Karasenko
Mitja Kleider
Christoph Koke
Alexander Kononov
Christian Mauch
Paul Müller
Thomas Pfeil
Bernhard Vogginger
Eric Müller
Dimitri Probst
Lyle Muller
Mikael Lundqvist
Alain Destexhe
Anders Lansner
Rene Schüffny
Michael Schmuker
Daniel Brüderle
Marc-Olivier Schwartz
Venelin Petkov
Roman Martel
Agnes Korcsak-Gorzo
Luziwei Leng
Nico Gürtler
Maximilian Zenk
Christian Weilbach
Marco Roth
Boris Rivkin
Elena Kreutzer
Joao Sacramento
Kristin Völk

Dominik Dold

Akos Kungl

Luziwei Leng

Oliver Breitwieser

Andi Baumbach

Walter Senn

Johannes Schemmel

Karlheinz Meier