Attractor Dynamics and Embodiment of Neural Computing

Yulia Sandamirskaya

Institute of Neuroinformatics (INI)
University of Zurich and ETH Zurich
Switzerland

@NICE 2019

What do brains “compute”?






- biological neural systems evolved to generate **movement**
- **goal-directed movement** requires
  - perception (state estimation)
  - calibration (internal and external alignment)
  - online adaptation (control)

  ➡ biological neural networks are intelligent **controllers**

What is required for intelligent control?

➡ working memory

- stabilisation of neuronal states

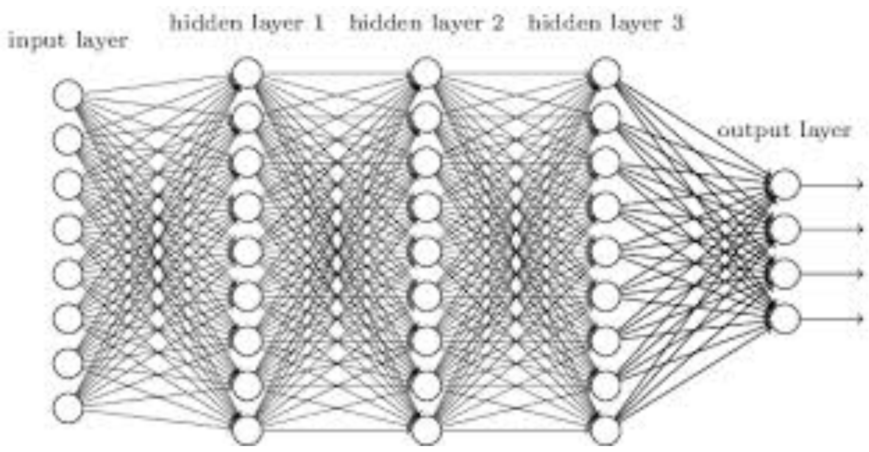

➡ decision making

- selection among alternatives
- “attention”

Not something artificial neural networks solve today!

What is required to enable “purely” neuronal computing / control?

➡ structure (autonomy)

➡ interfaces to sensors and motors (embodiment)

Towards neuromorphic SLAM

Simultaneous localisation and mapping

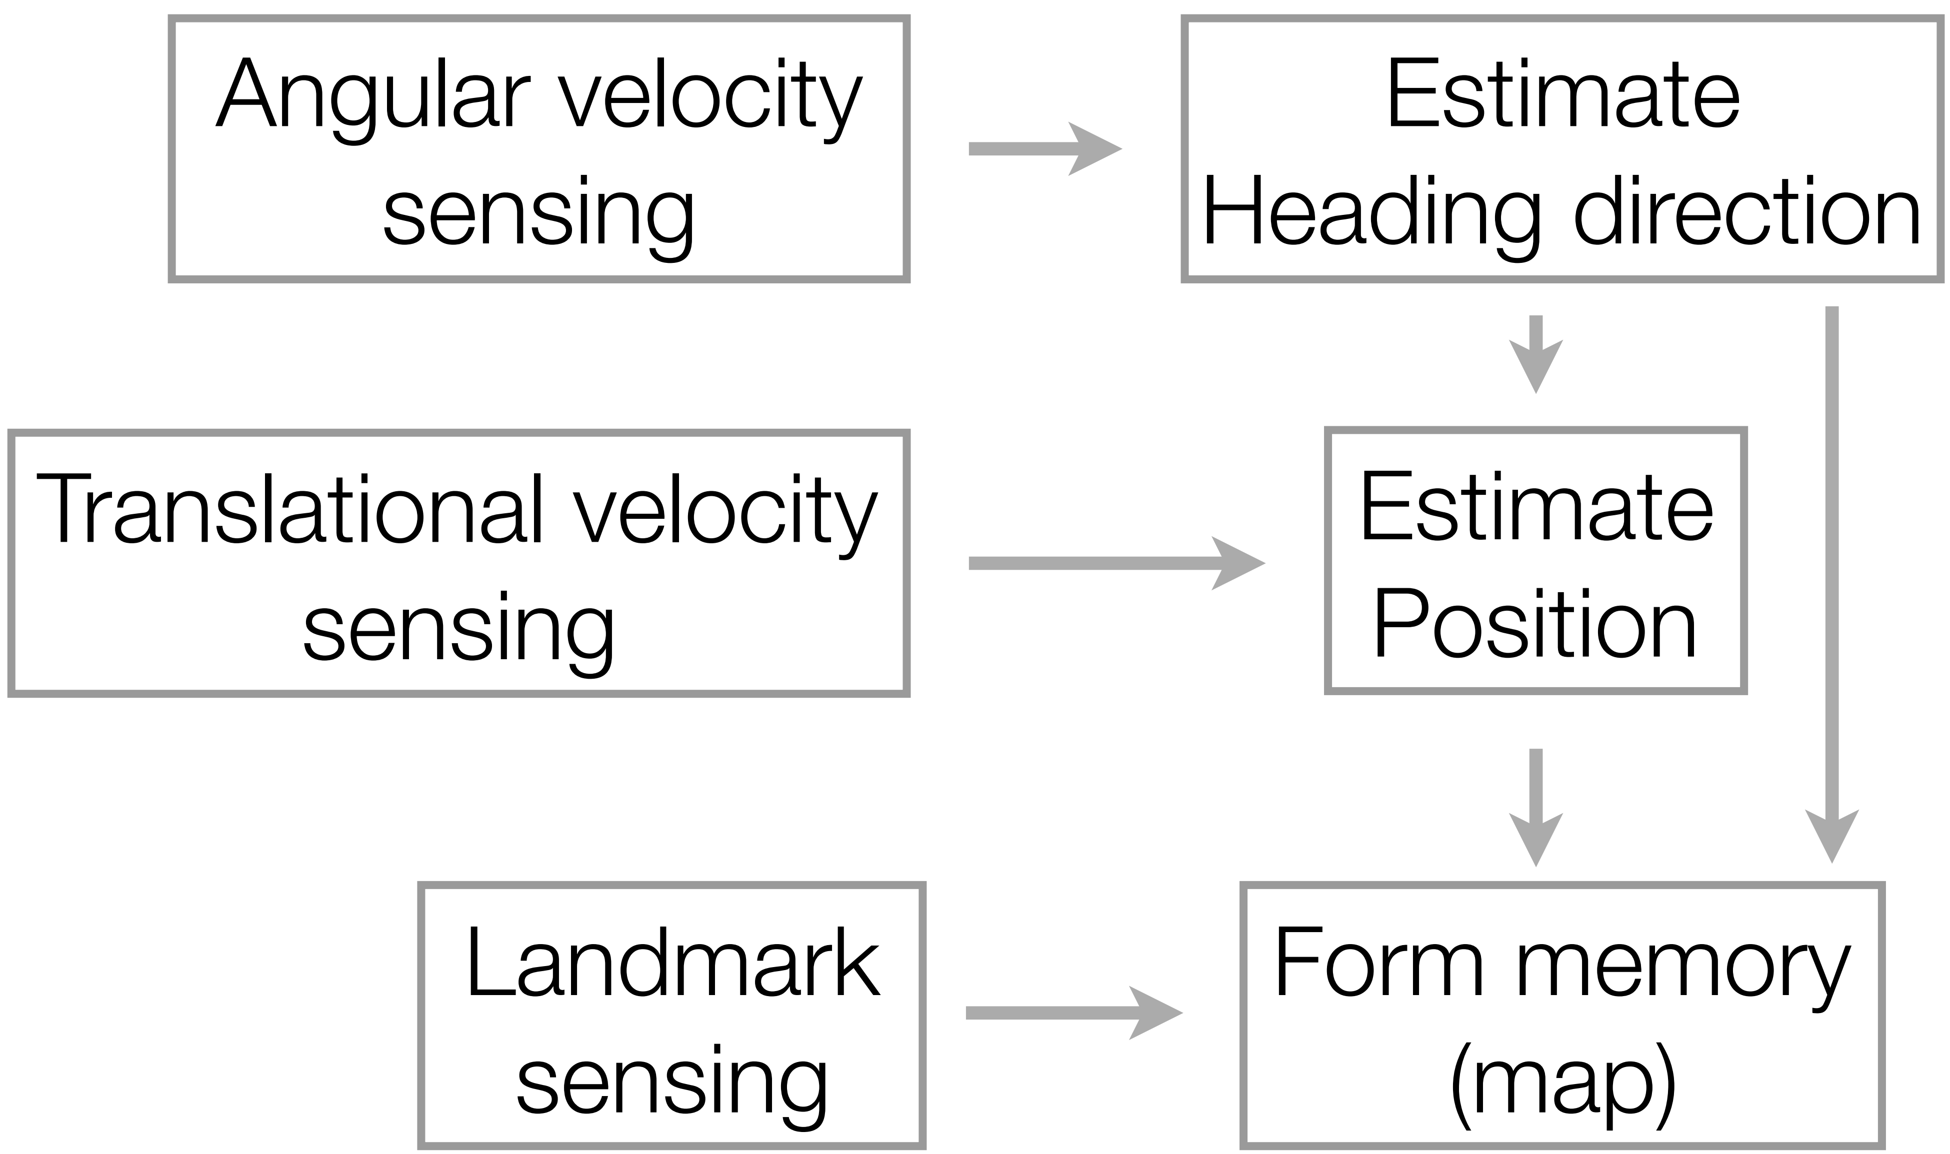

Neuromorphic hardware

ROLLS

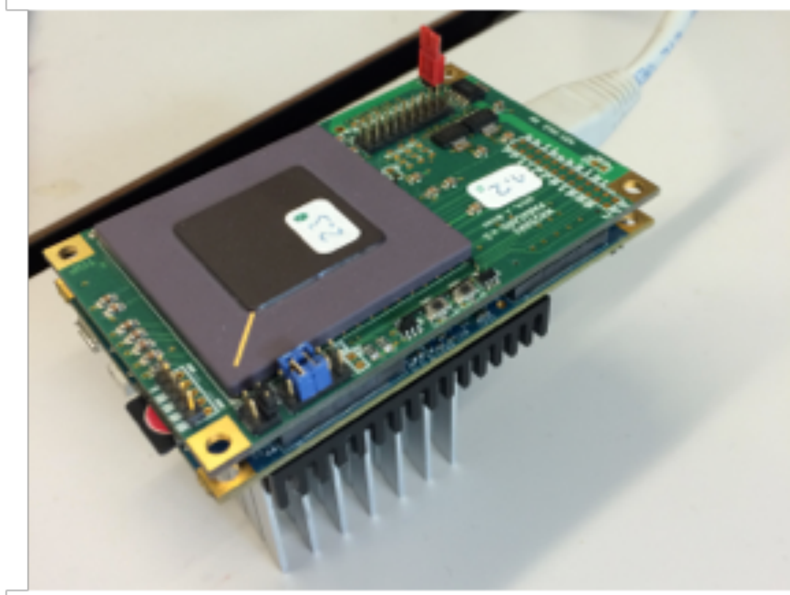

(Qiao, Indiveri, 2015)

- mixed-signal circuitry
- 256 artificial neurons
- 256 x 256 plastic synapse circuits
- 180nm process
- ultra-low power

DYNAP

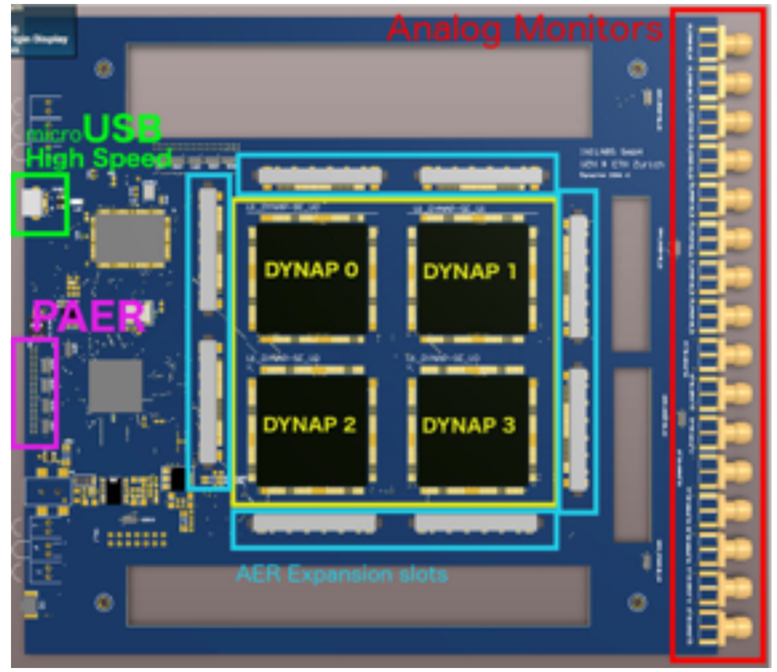

(Qiao, Indiveri, 2018)

- mixed-signal circuitry
- 4 x 1024 artificial neurons
- 64 x 1024 synapses
- 180nm process
- ultra-low power
- scalable

Loihi

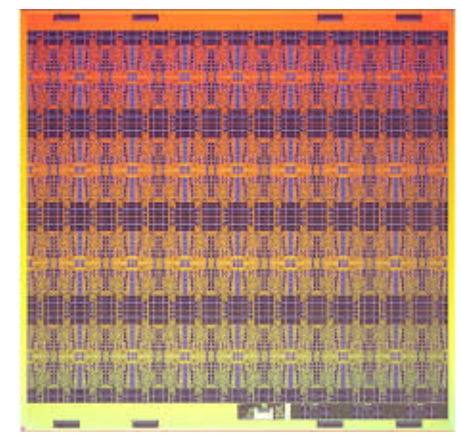

(Davies, 2018)

- fully digital circuitry
- 130,000 artificial neurons
- 130 million synapses
- programmable learning engine
- 14nm process
- low power and scalable

Reconfigurable OnLine Learning Spiking (ROLLS)

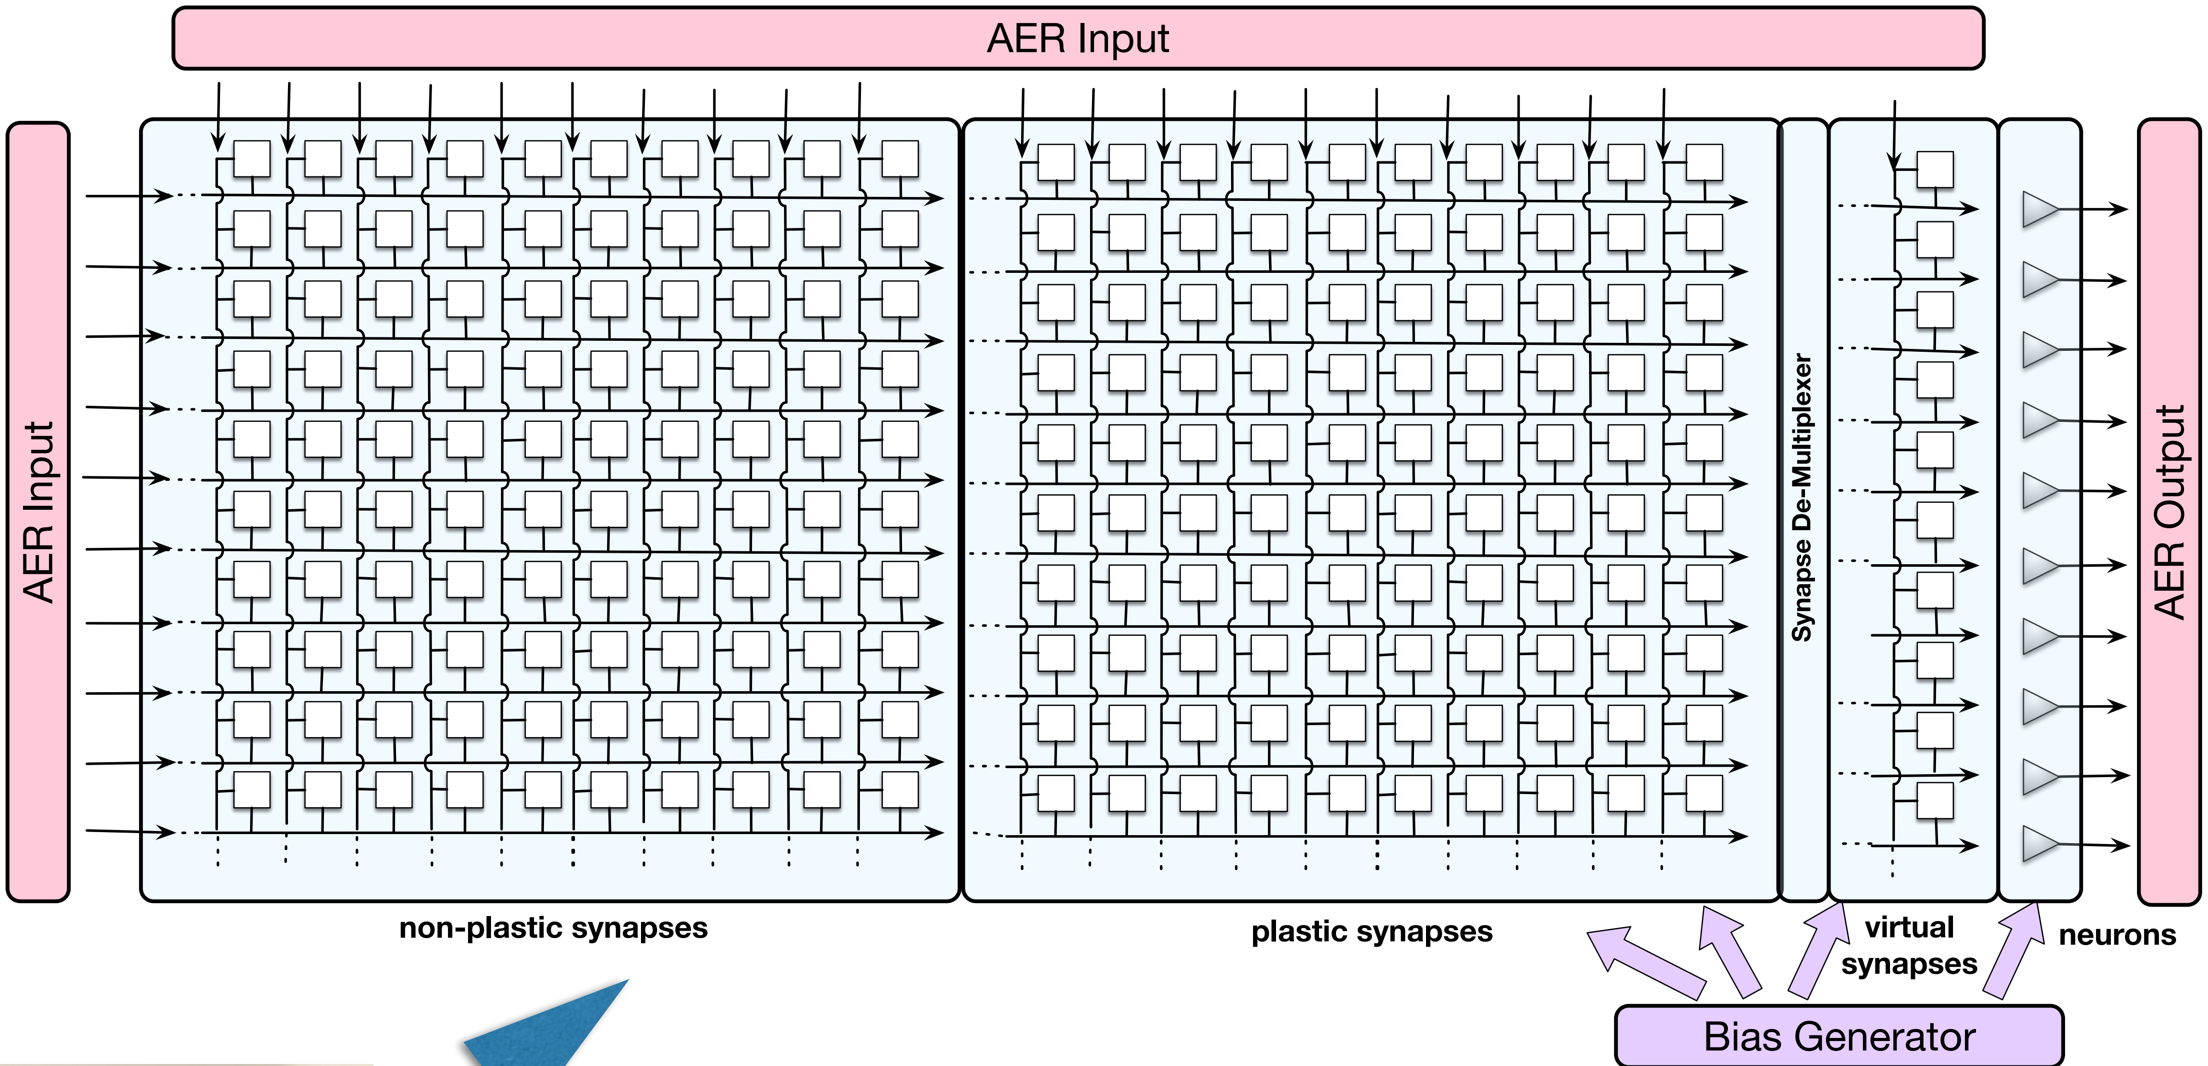

- analog circuits for neurons and synapses
- digital communication of spikes

➡“programming” = wiring-up and setting parameters

Qiao et al, 2015

Heading direction network

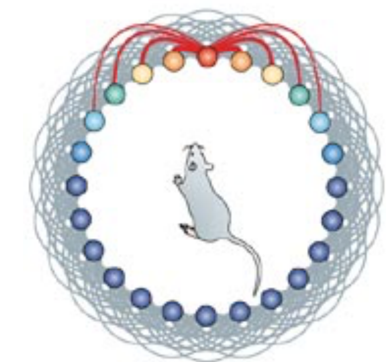

Heading
Direction

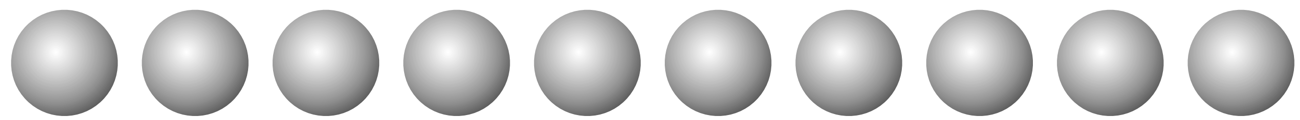

Heading direction network

Heading
Direction

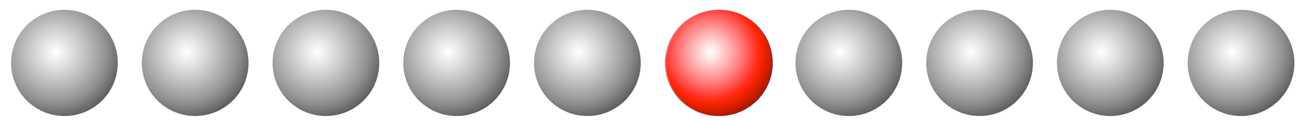

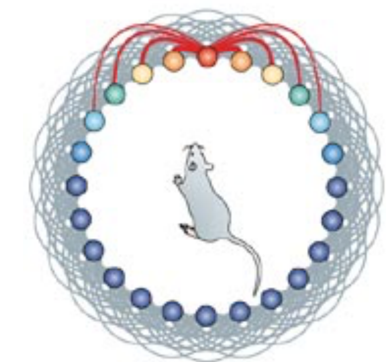

Heading direction network

Heading
Direction

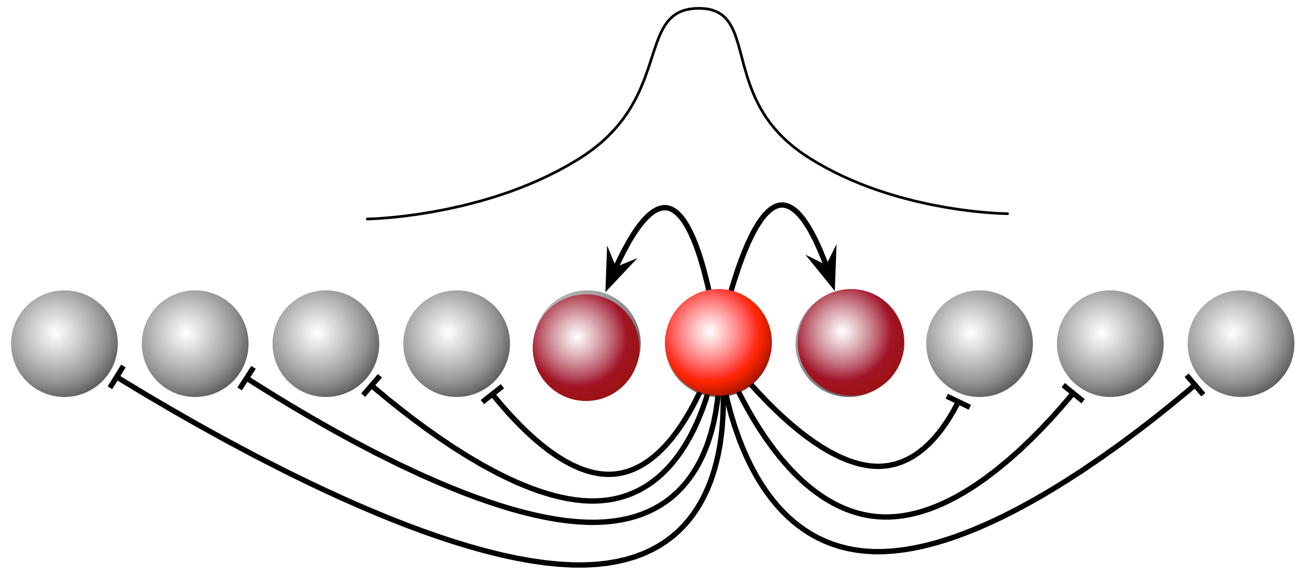

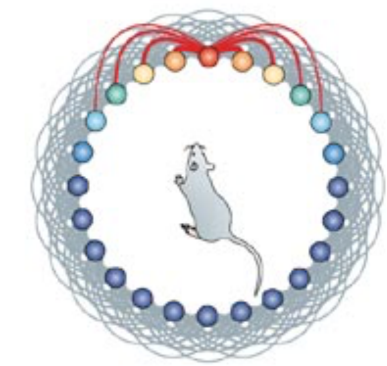

➡(soft) “winner take all”

➡dynamic neural field

- Population activity dynamics:

$$\tau \dot{u}(x,t) = -u(x,t) + h + \int f\big(u(x',t)\big)\omega(x-x')dx' + I(x,t)$$

Amari, S. **Dynamics of pattern formation in lateral-inhibition type neural fields**. Biological Cybernetics, **1977**, 27, 77-87

Wilson, H. R. & Cowan, J. D. **A mathematical theory of the functional dynamics of cortical and thalamic nervous tissue**. Kybernetik, **1973**, 13, 55-80

Gerstner, Grossberg, Ermentrout, Coombes, **Schöner&Spencer, 2015**, Erlhagen…

Heading direction network

Heading
Direction

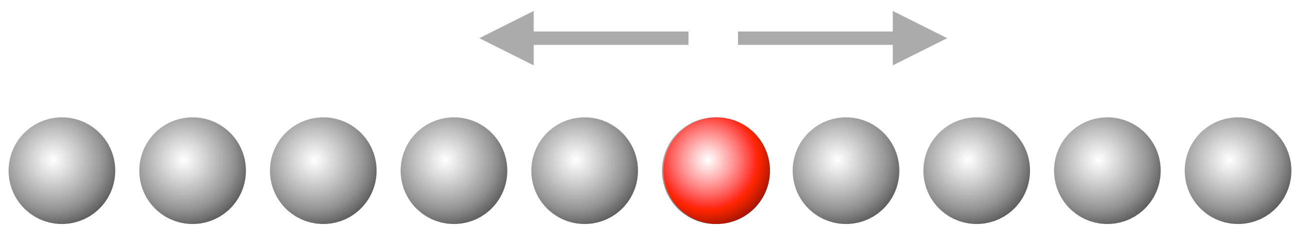

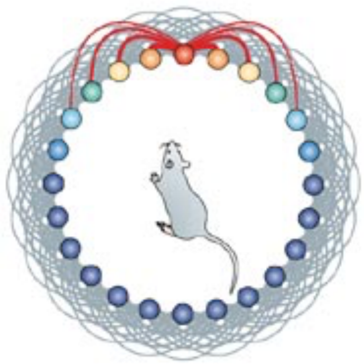

Heading direction network

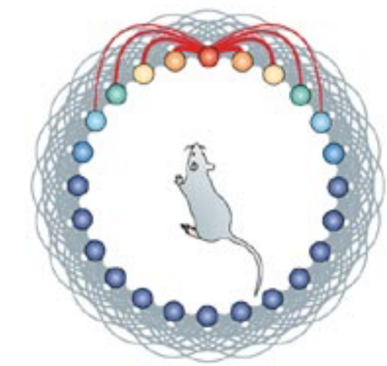

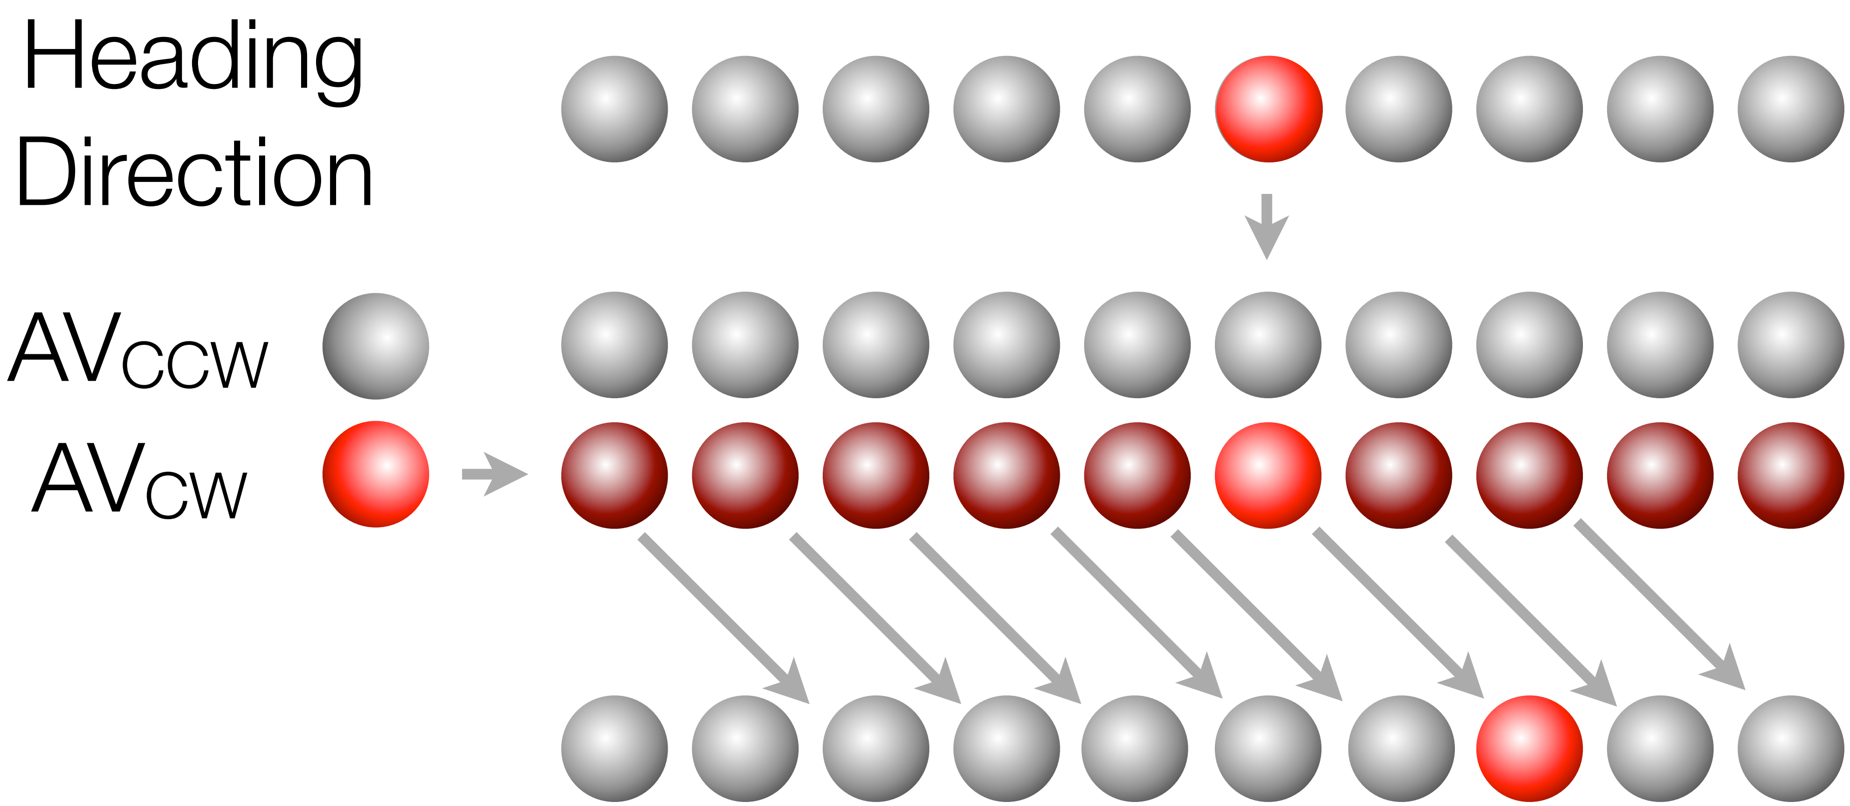

Heading direction network

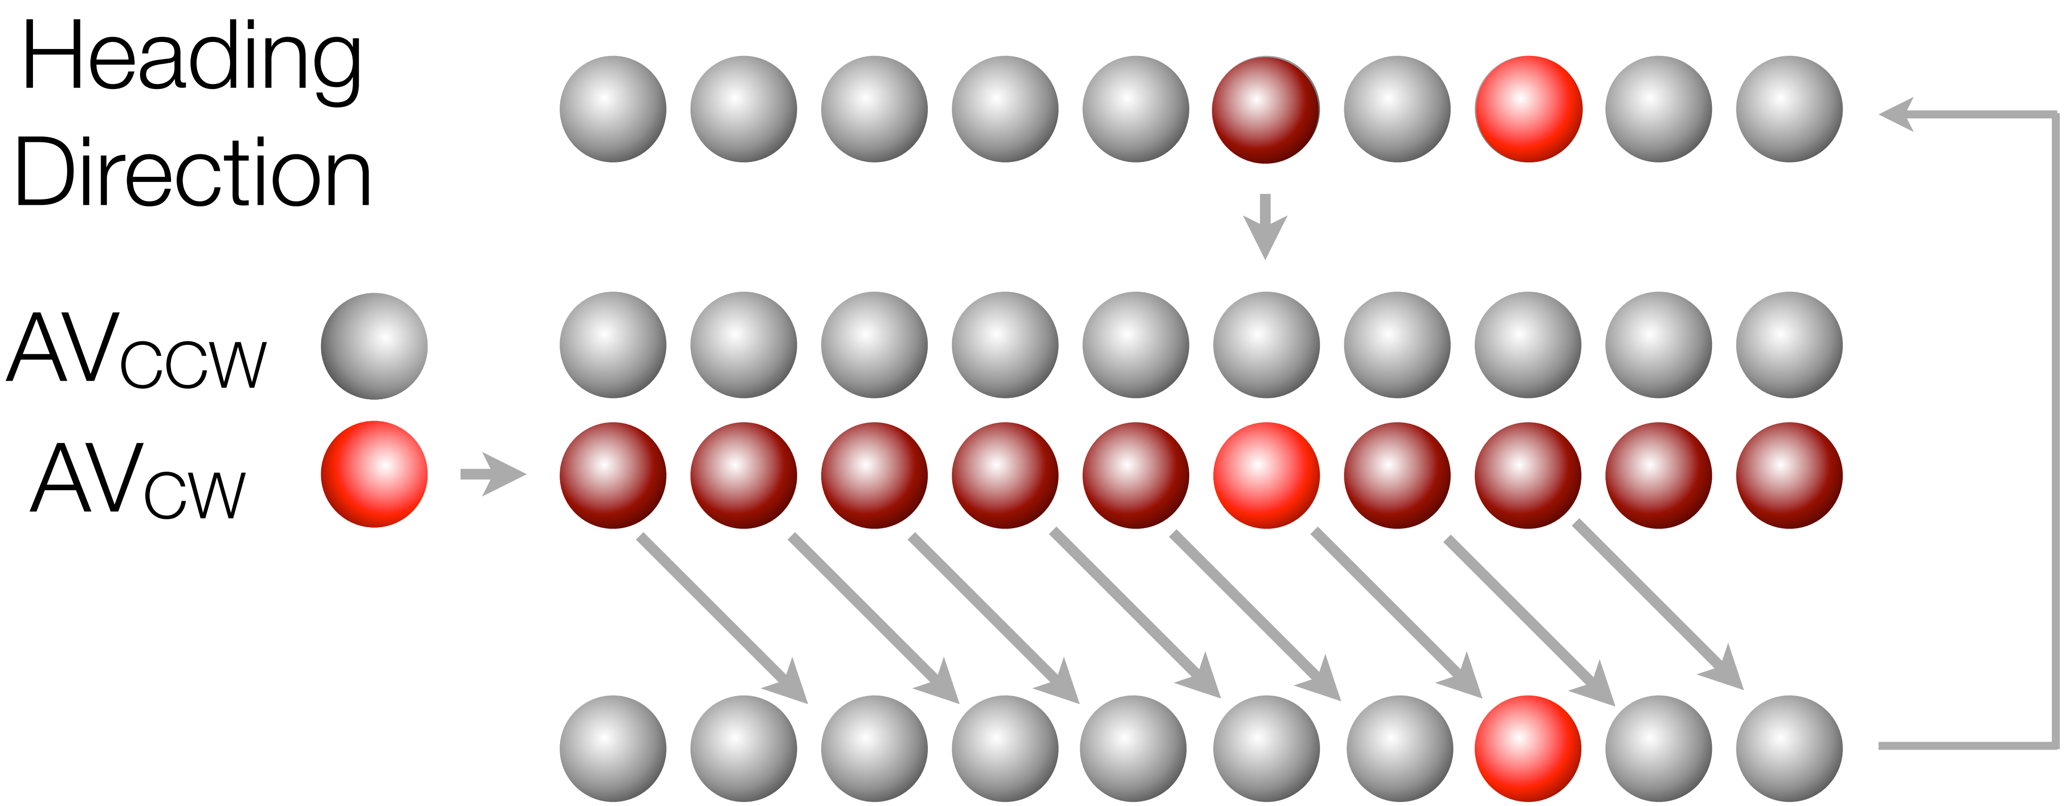

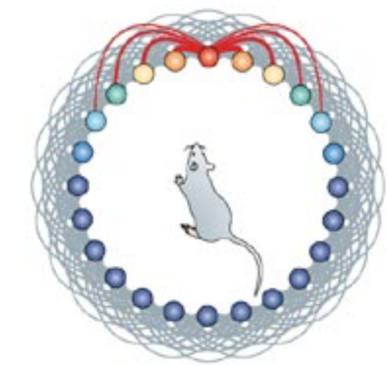

Heading direction network

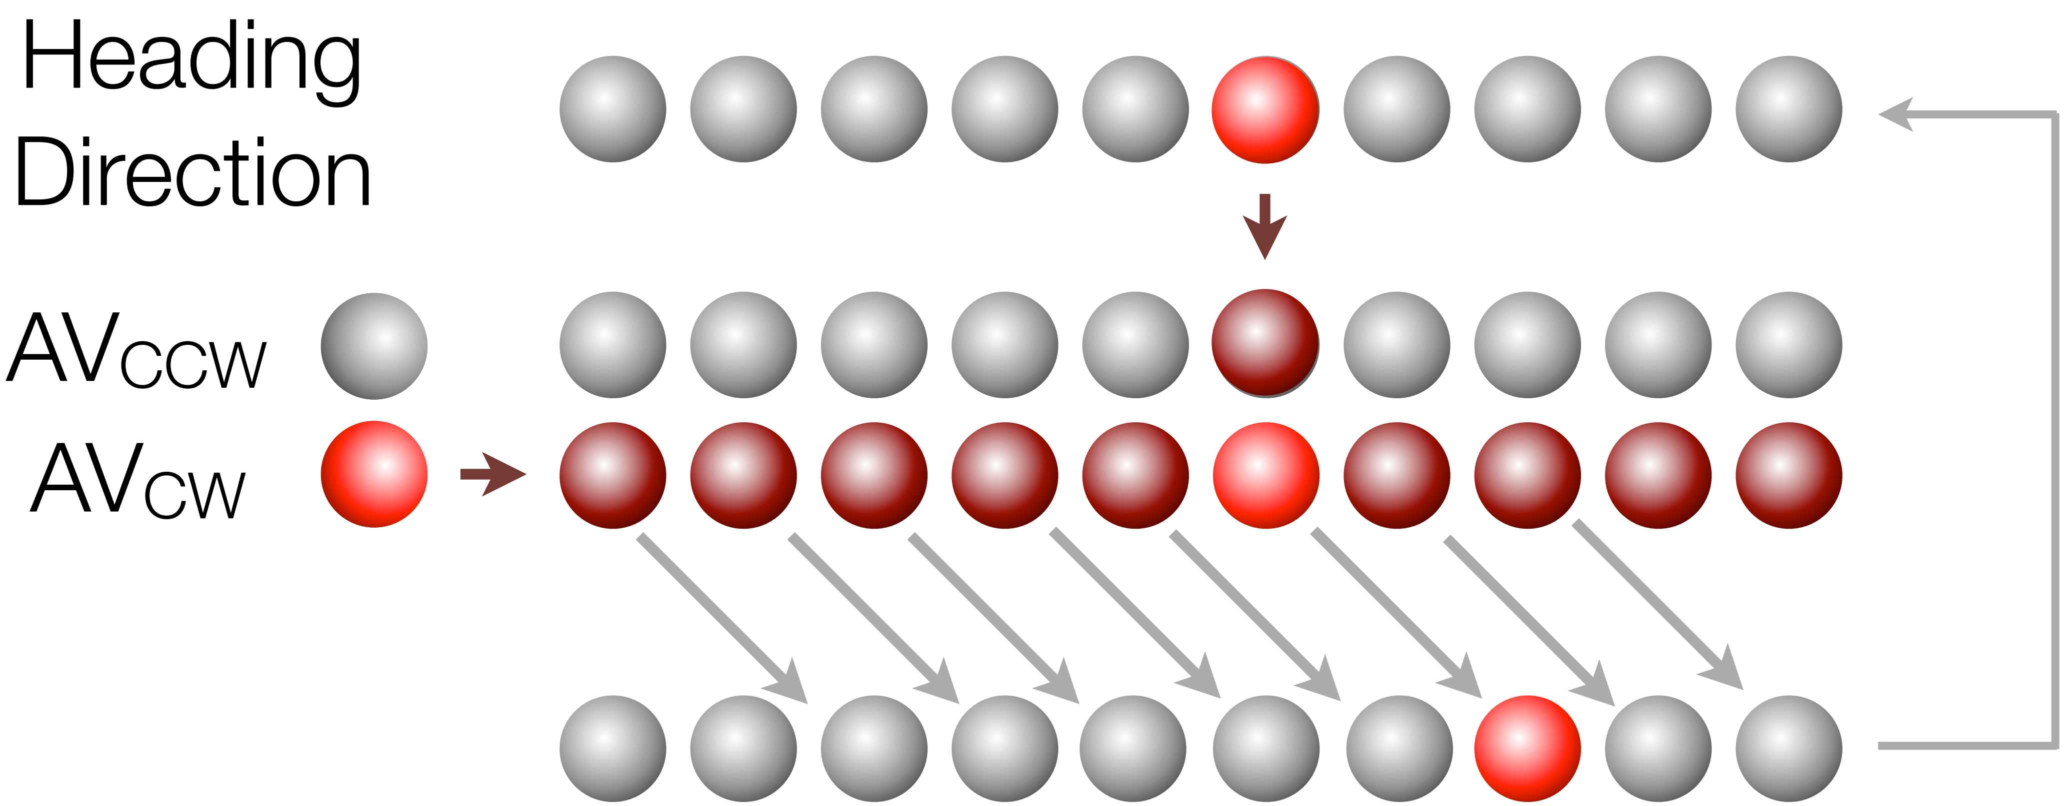

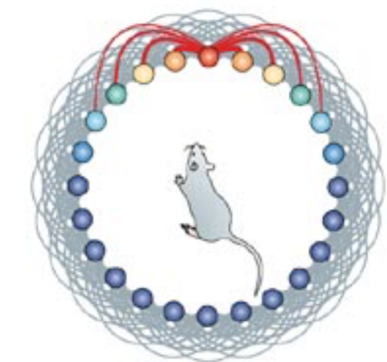

Analogue

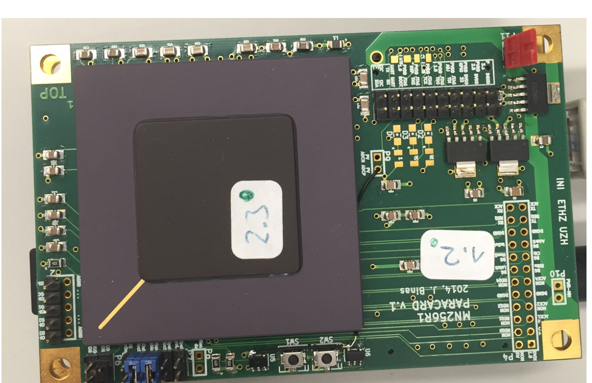

- missmatch
- variability
- low precision

Heading direction network

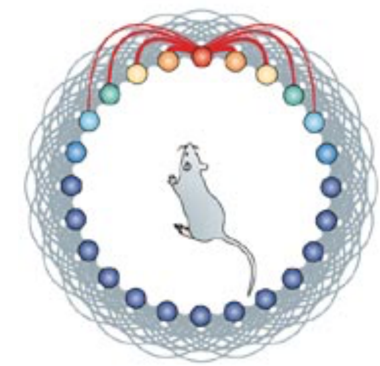

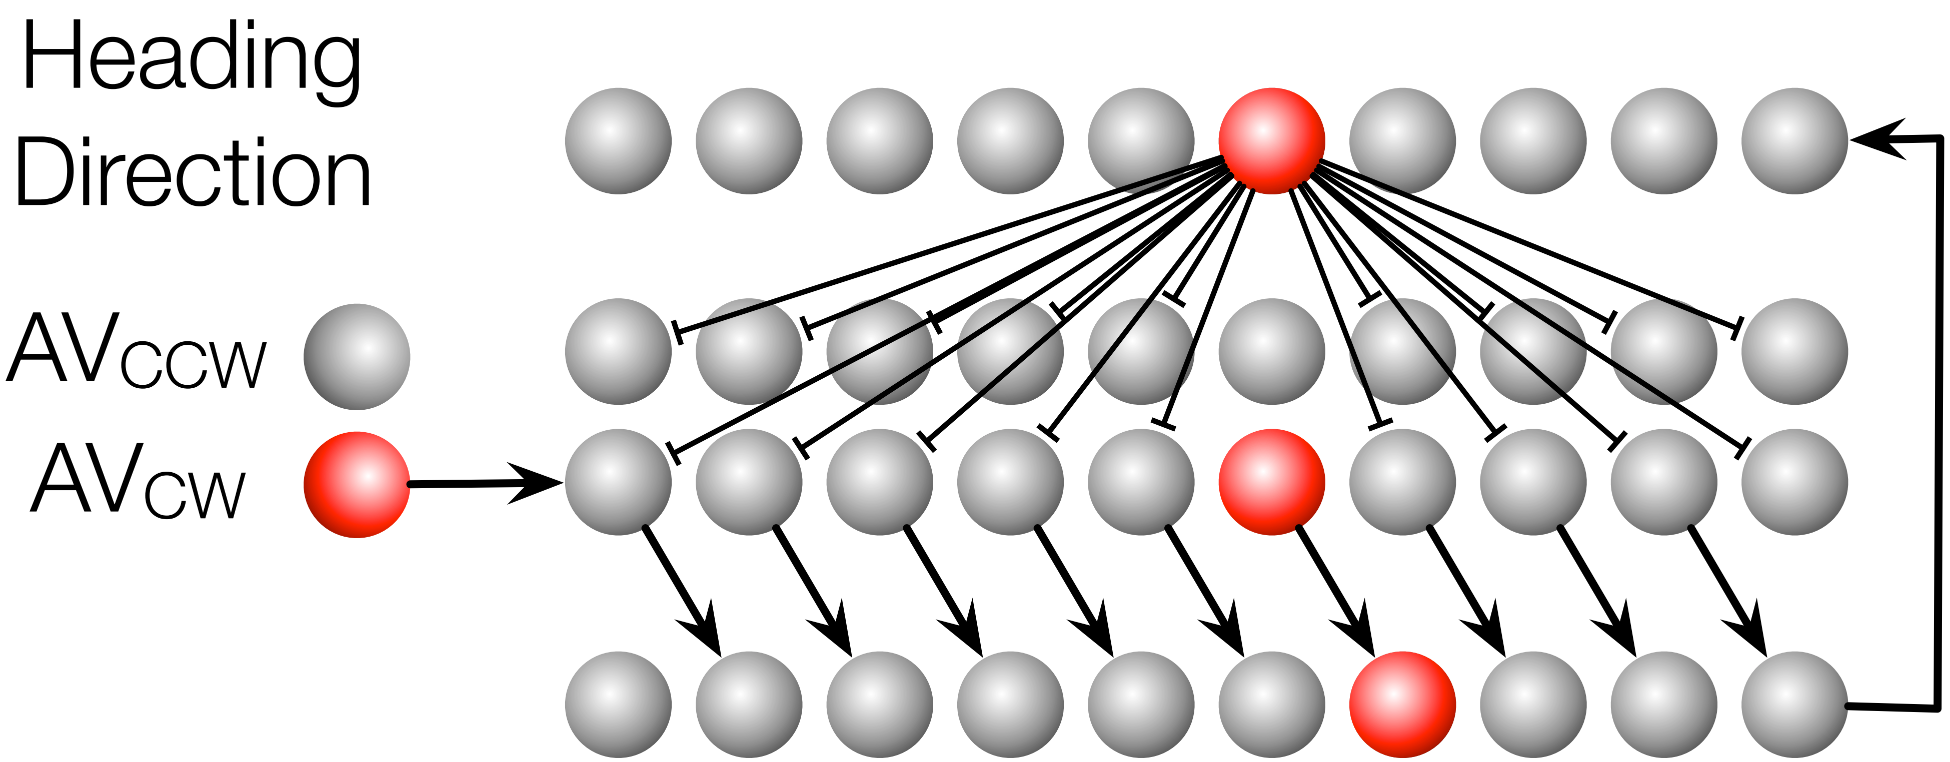

➡ More robust connectivity: desinhibition

Heading direction network

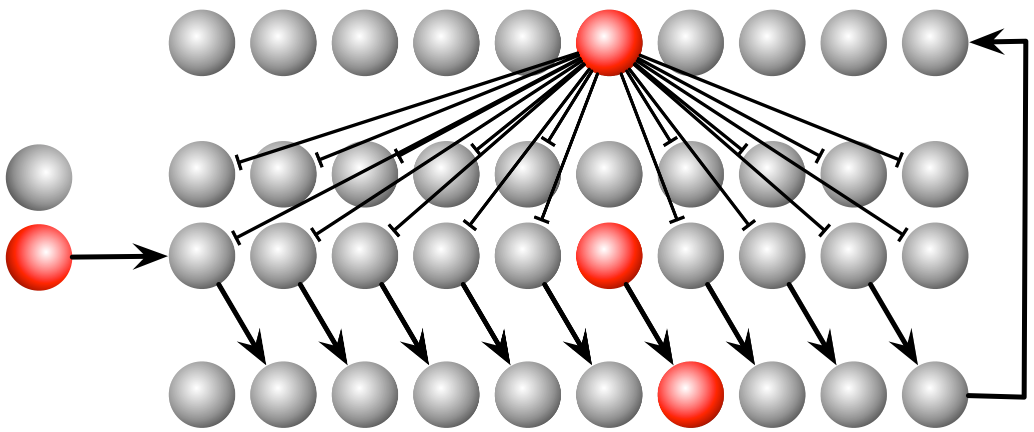

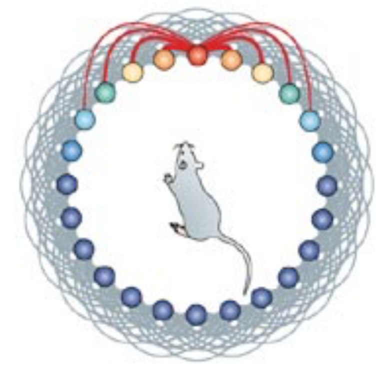

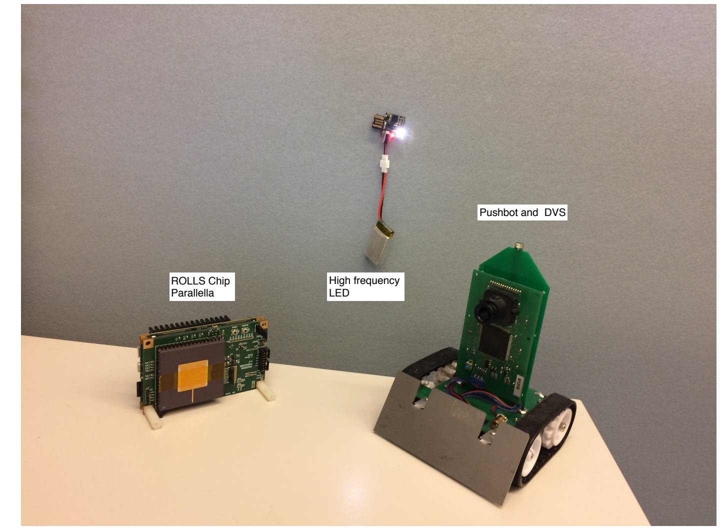

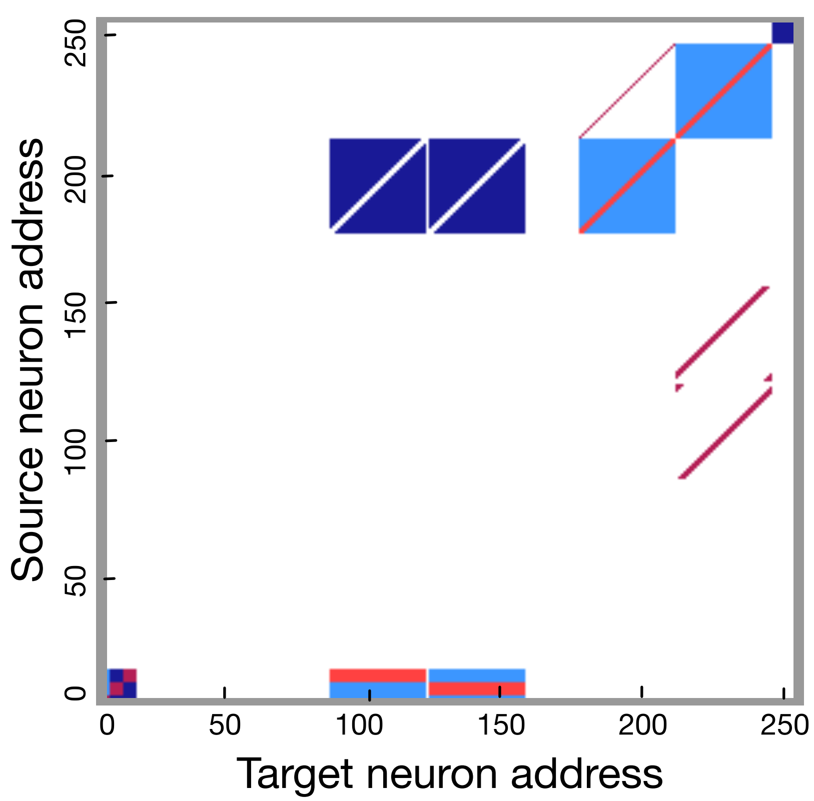

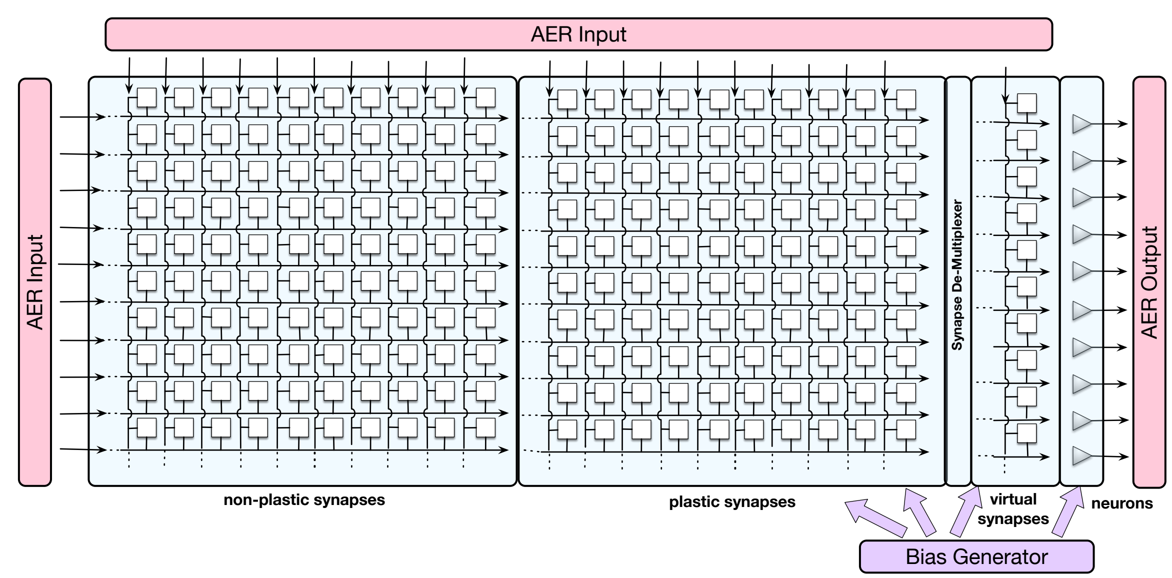

Kreiser, R.; Cartiglia, M. & Sandamirskaya, Y. A Neuromorphic approach to path integration: a head direction spiking neural network with visually-driven reset. IEEE Symposium for Circuits and Systems, ISCAS, **2018**

Heading direction estimation: hardware results

Spiking activity on ROLLS

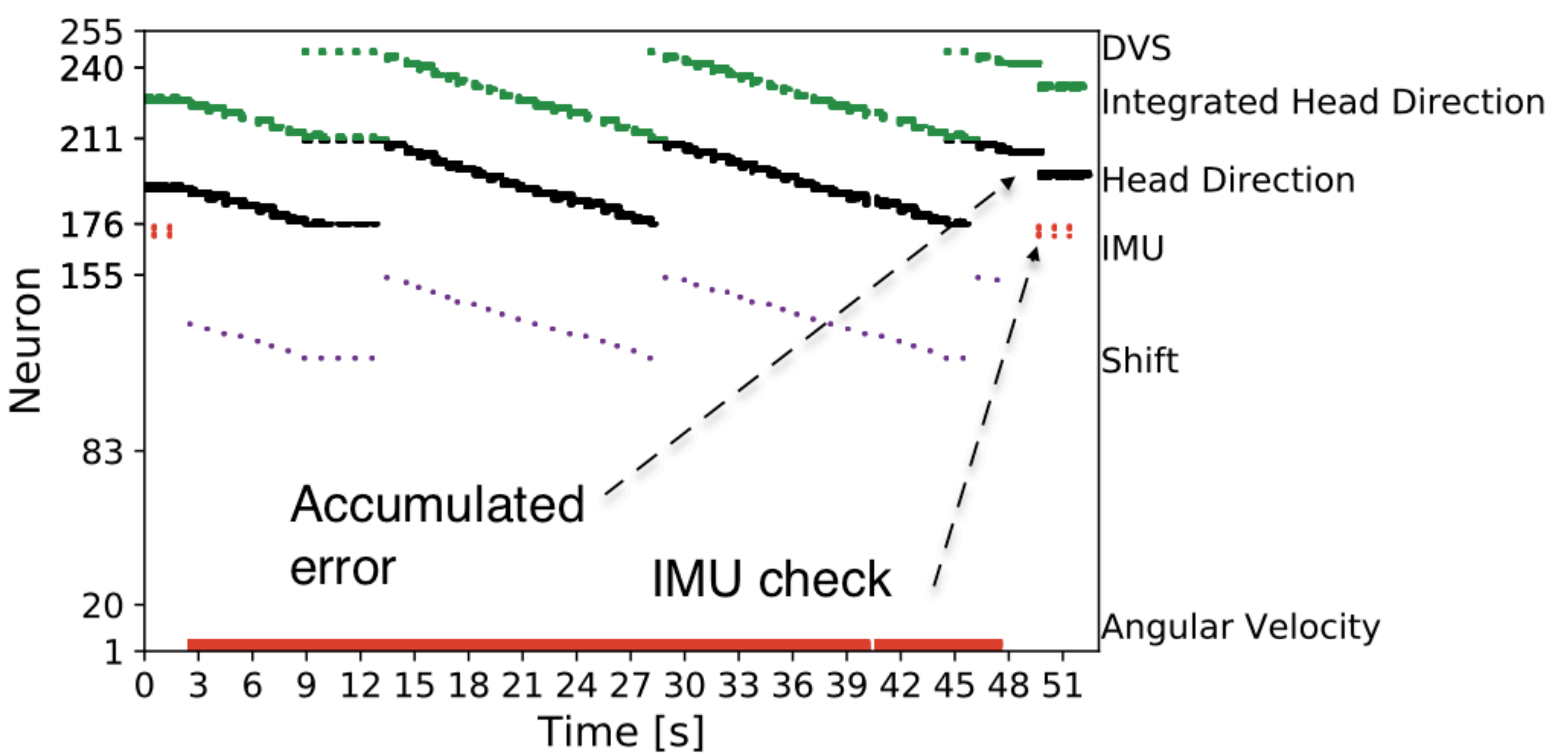

Error accumulation

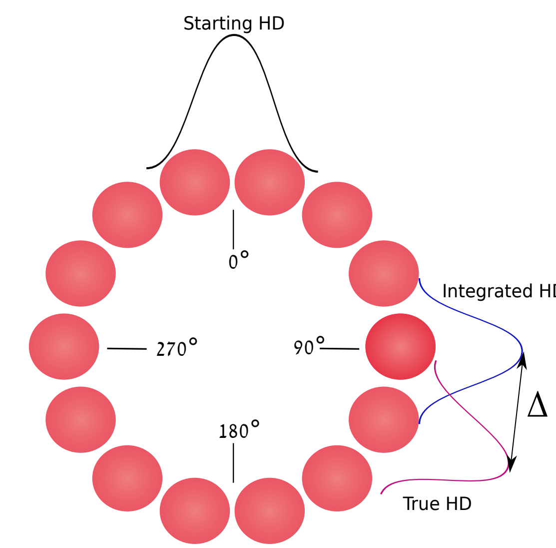

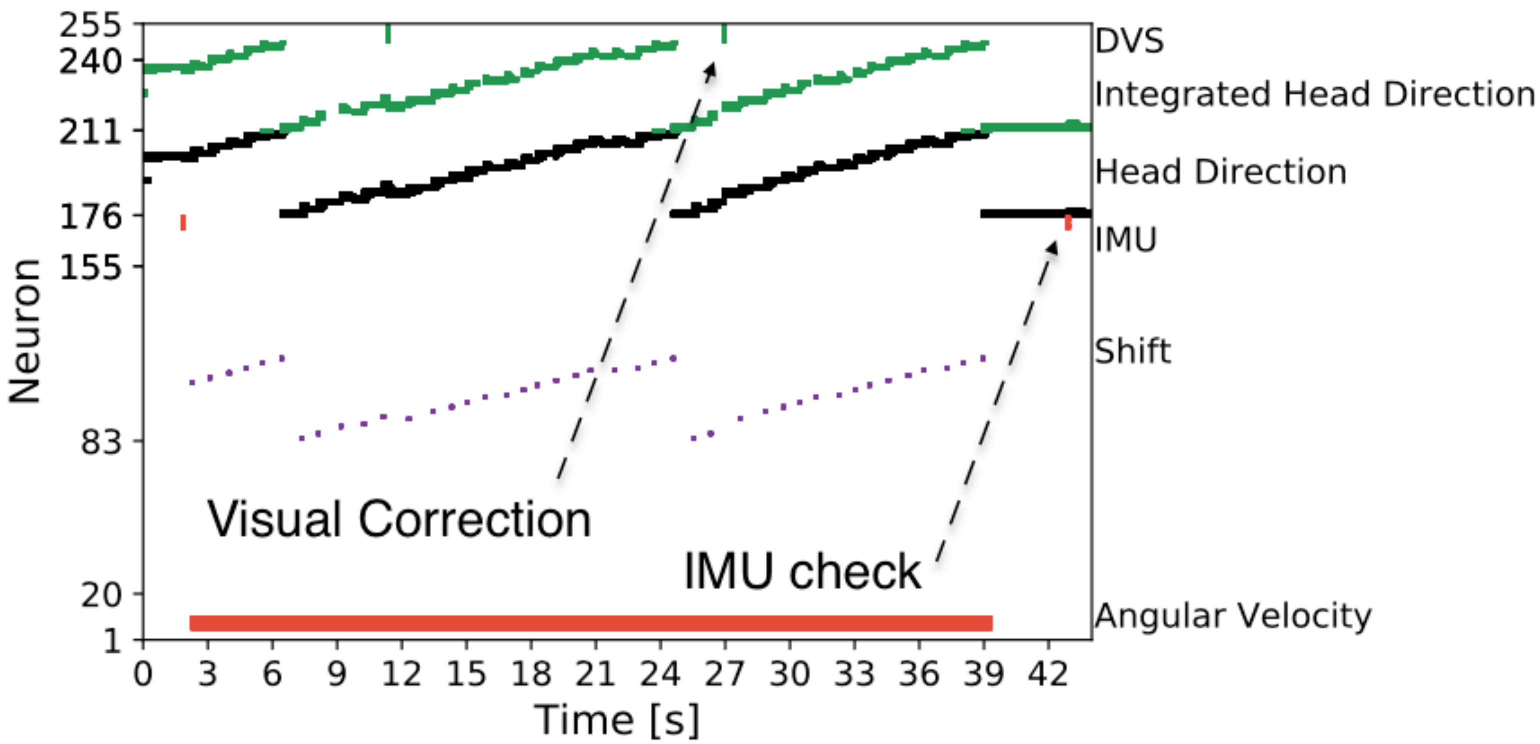

Correction using vision

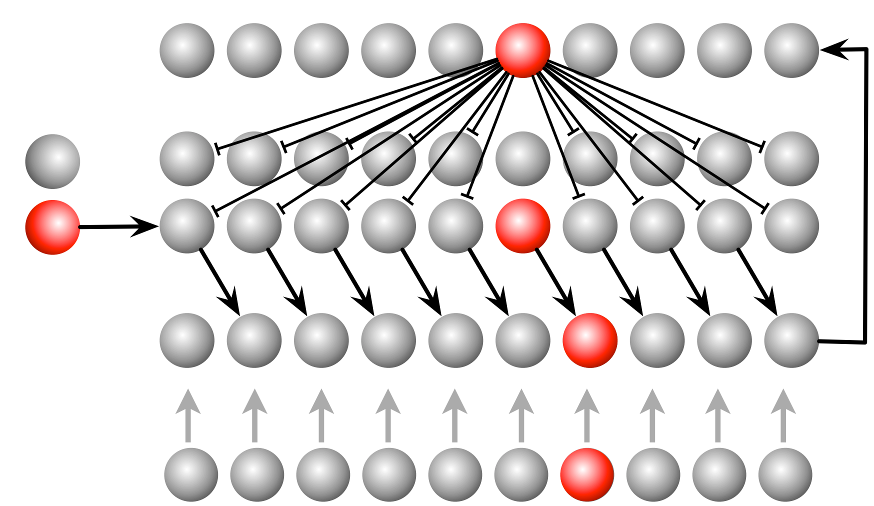

Reset

Real-time activity on the ROLLS chip

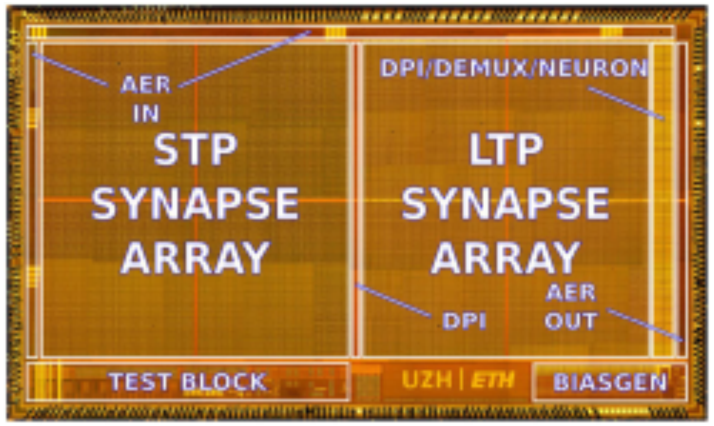

Qiao, Ning, et al, *Frontiers in neuroscience,* 2015

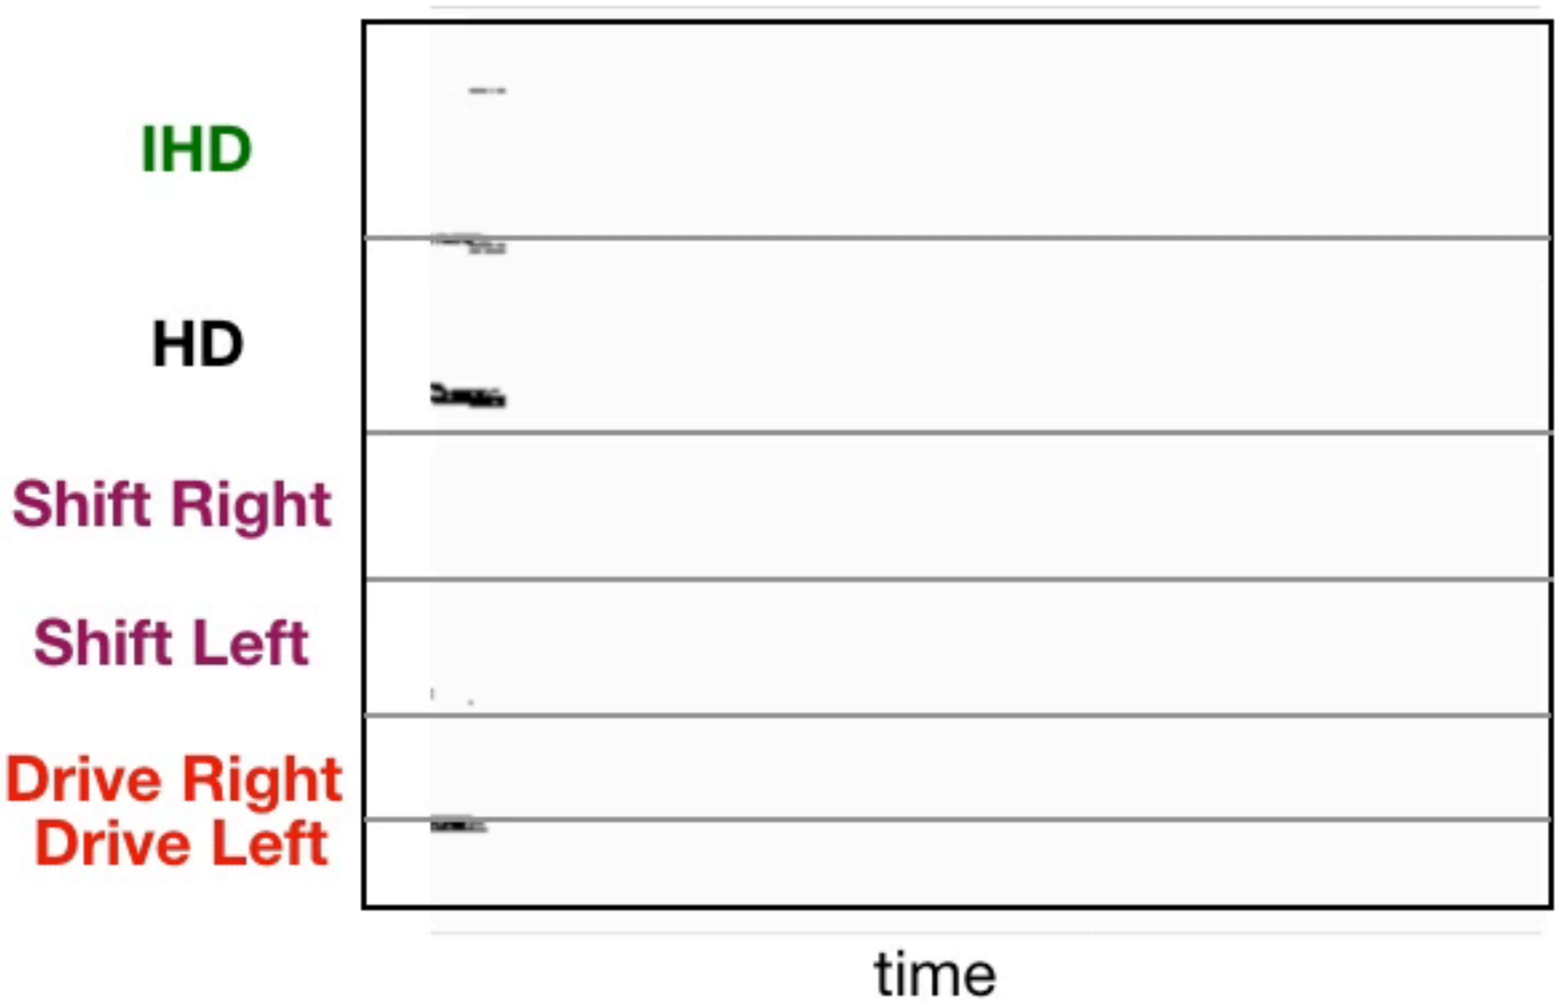

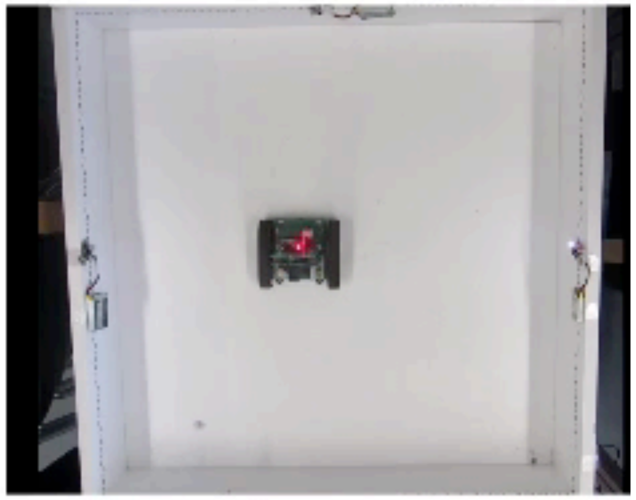

"Loop closure" and calibration

- How fast does the activity bump need to move?

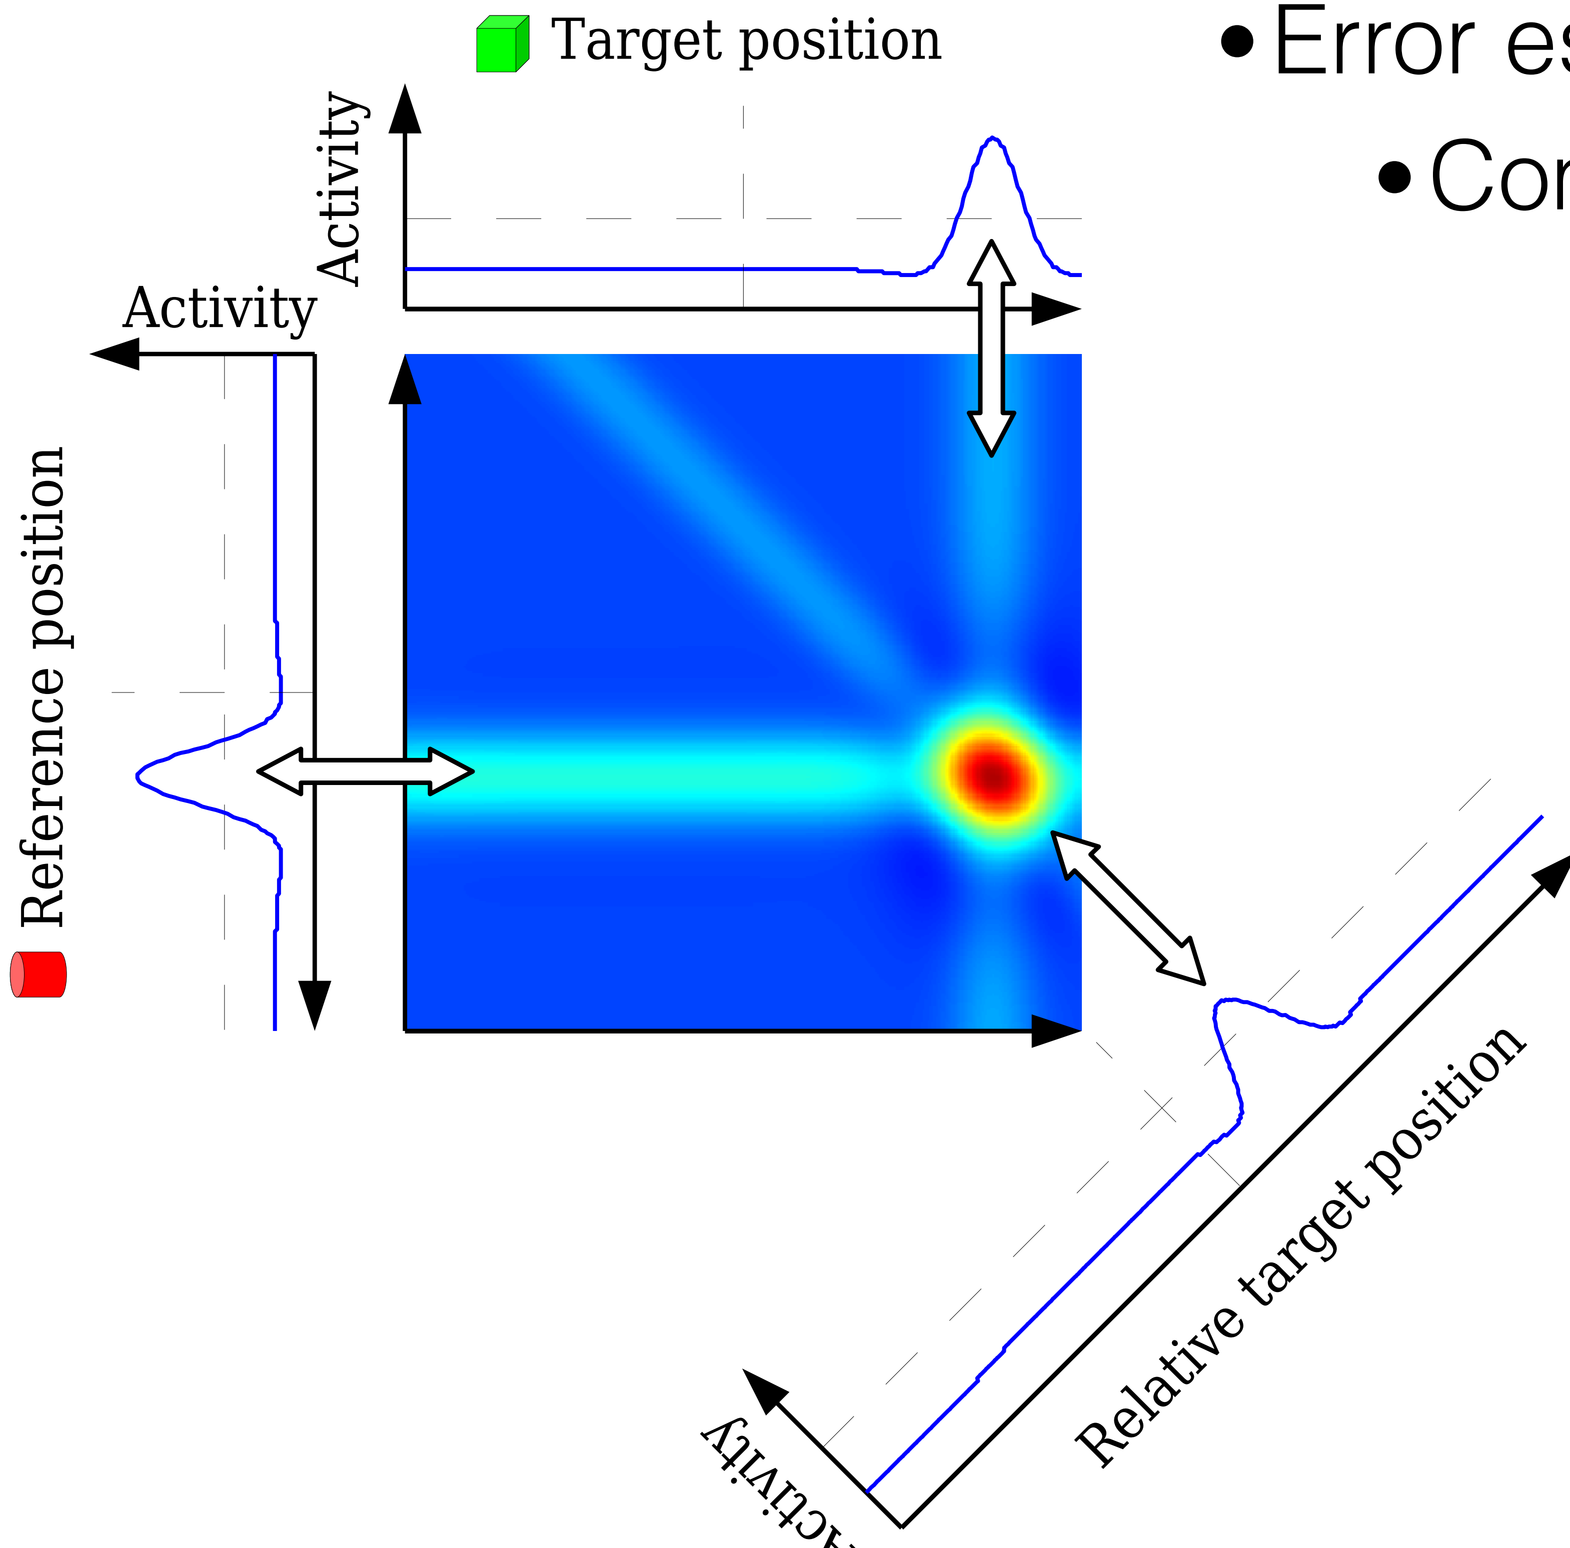

- Error estimation circuit
  - Computing differences

“Loop closure” and calibration

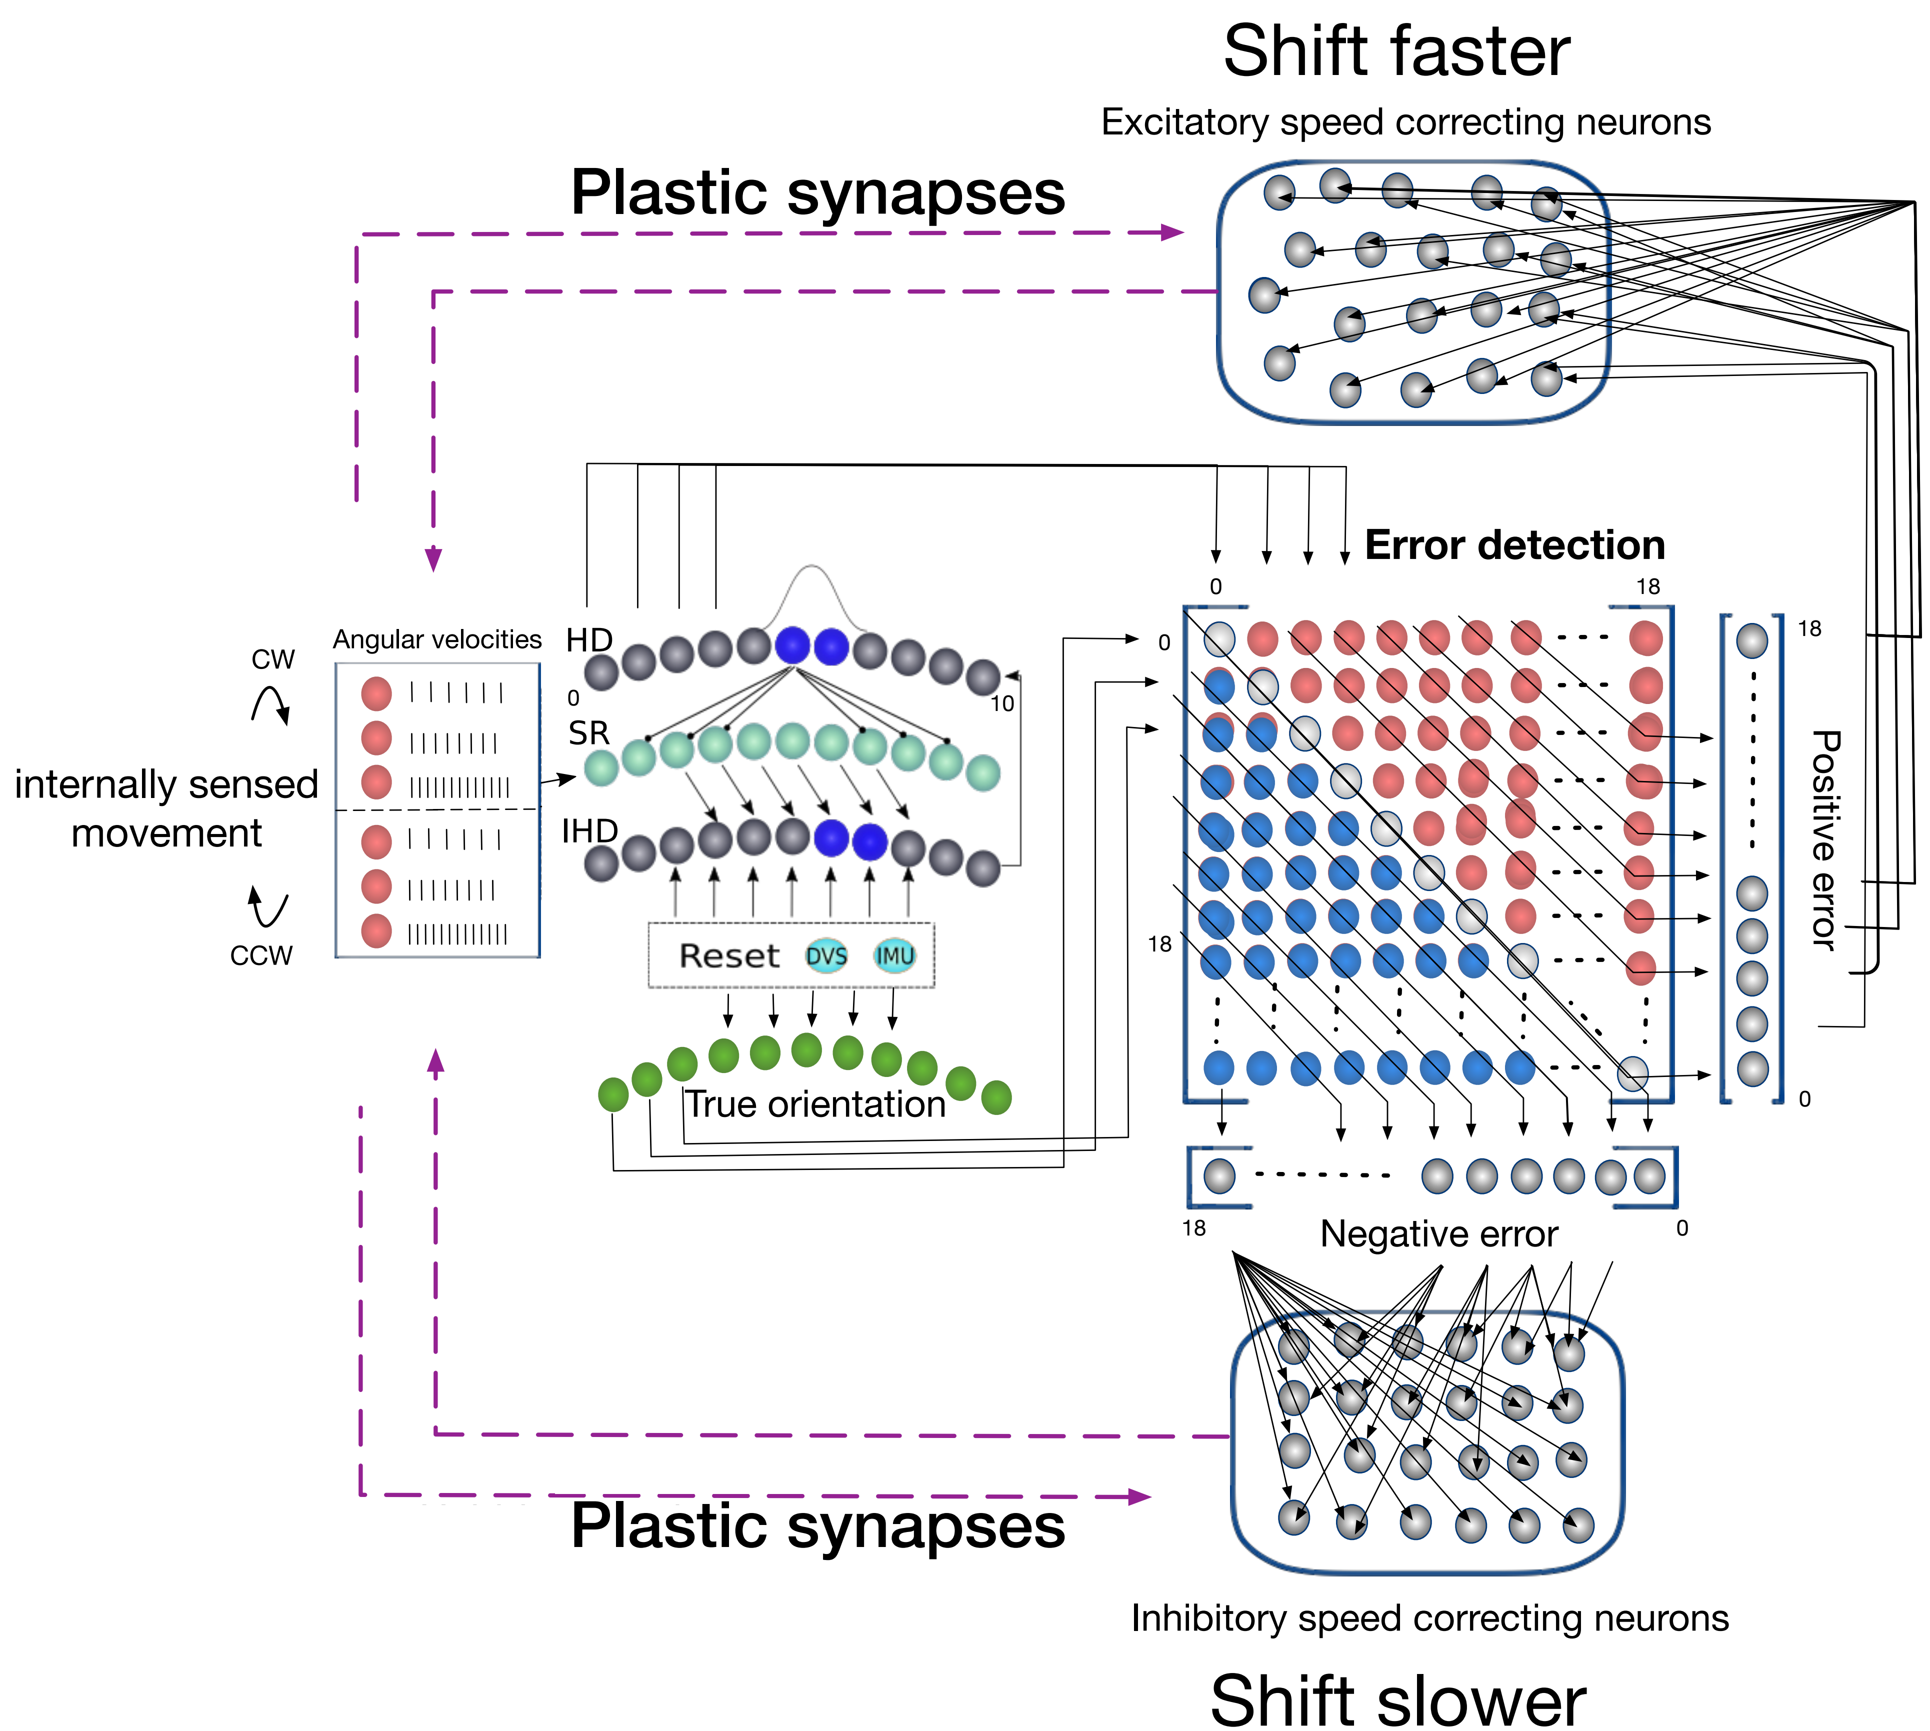

Matching activity shifting velocity to real velocity

Learning to shift faster

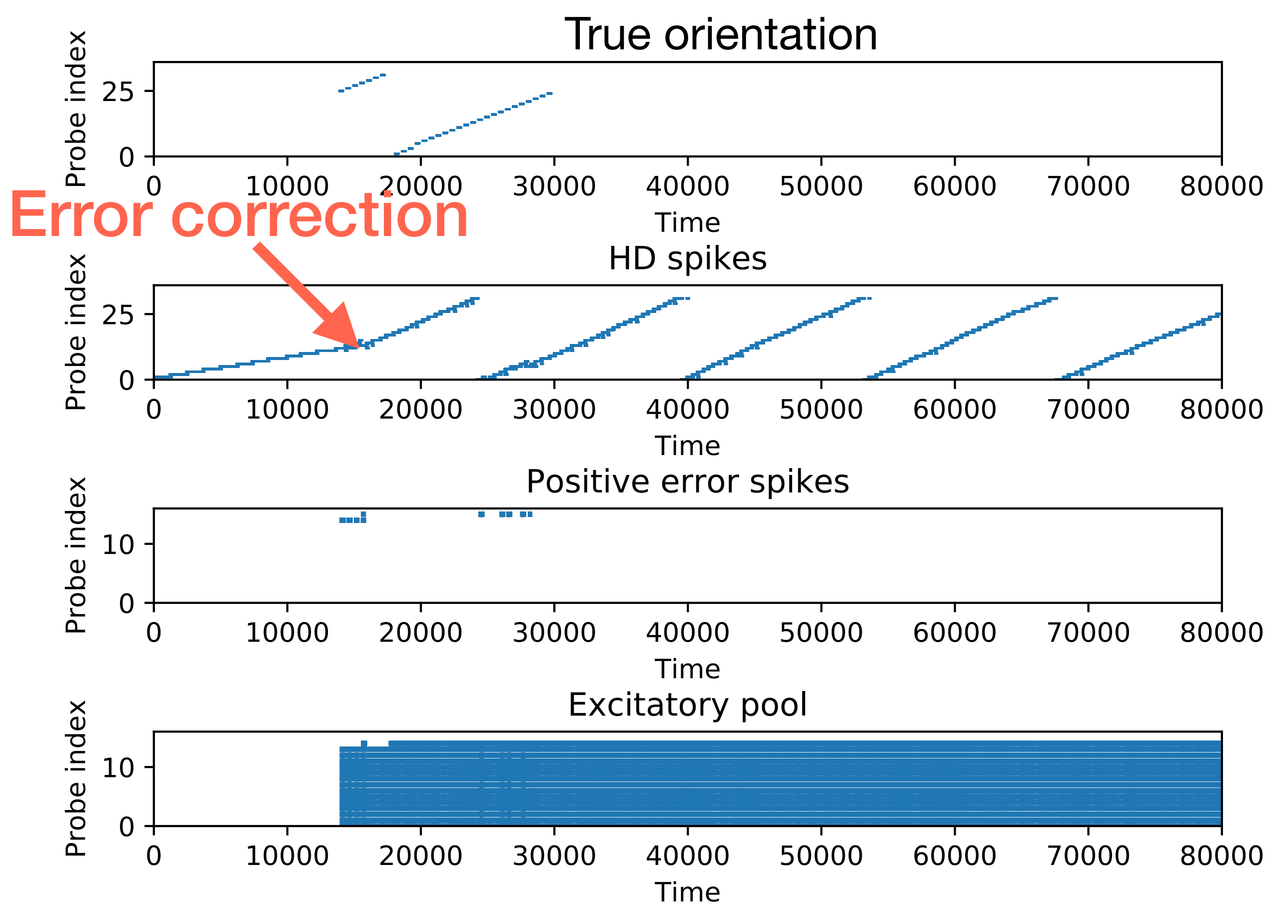

Learning to shift slower

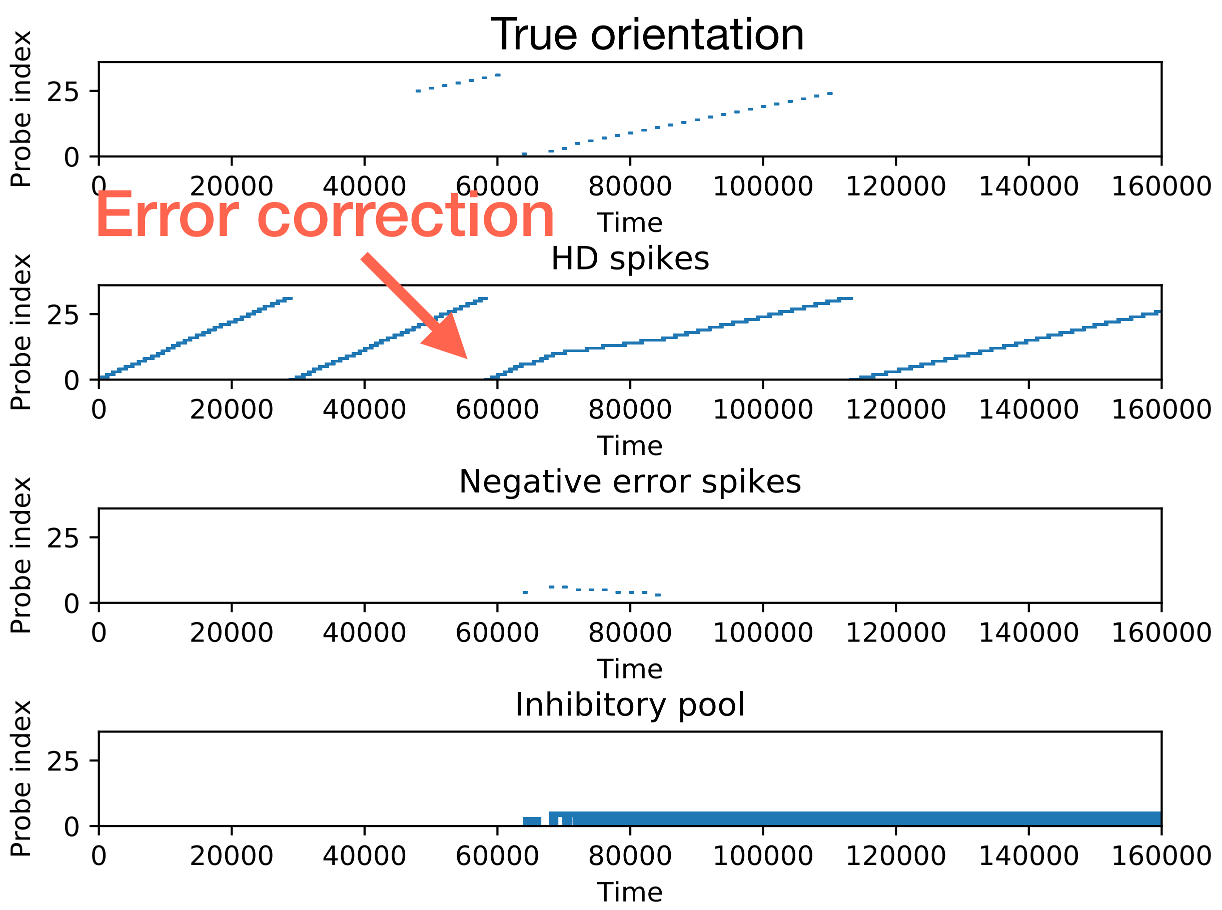

Position estimation

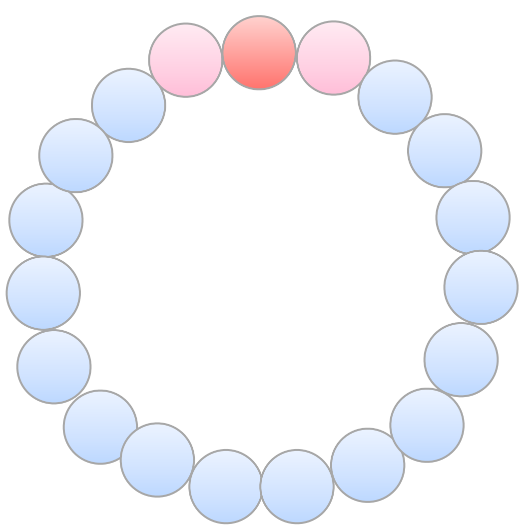

Heading direction estimation

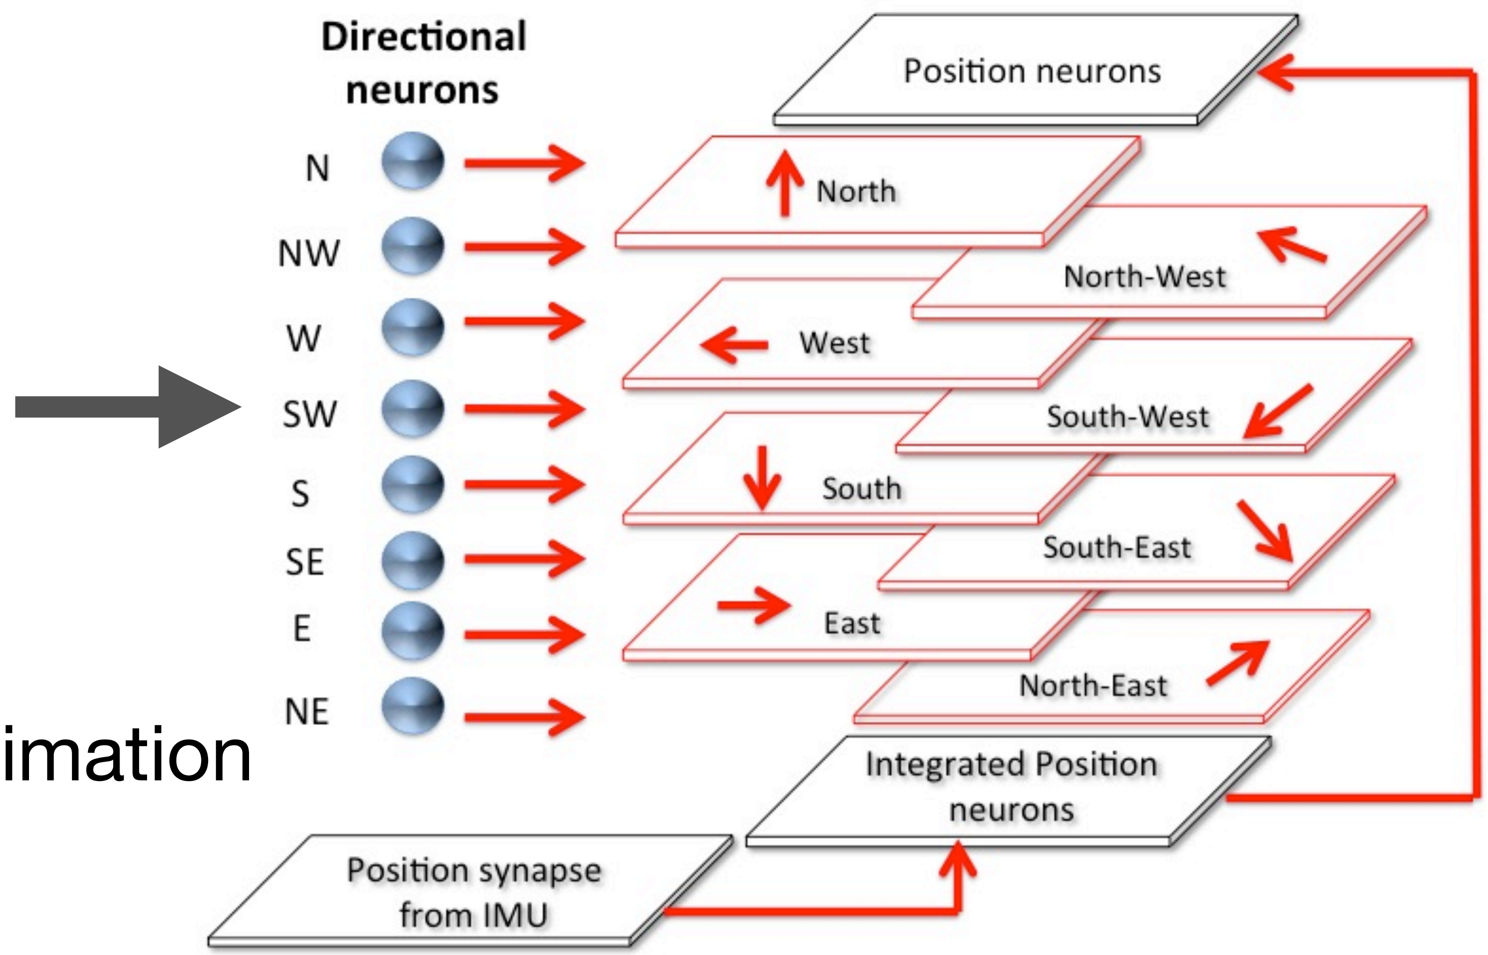

Position estimation network

Position estimation

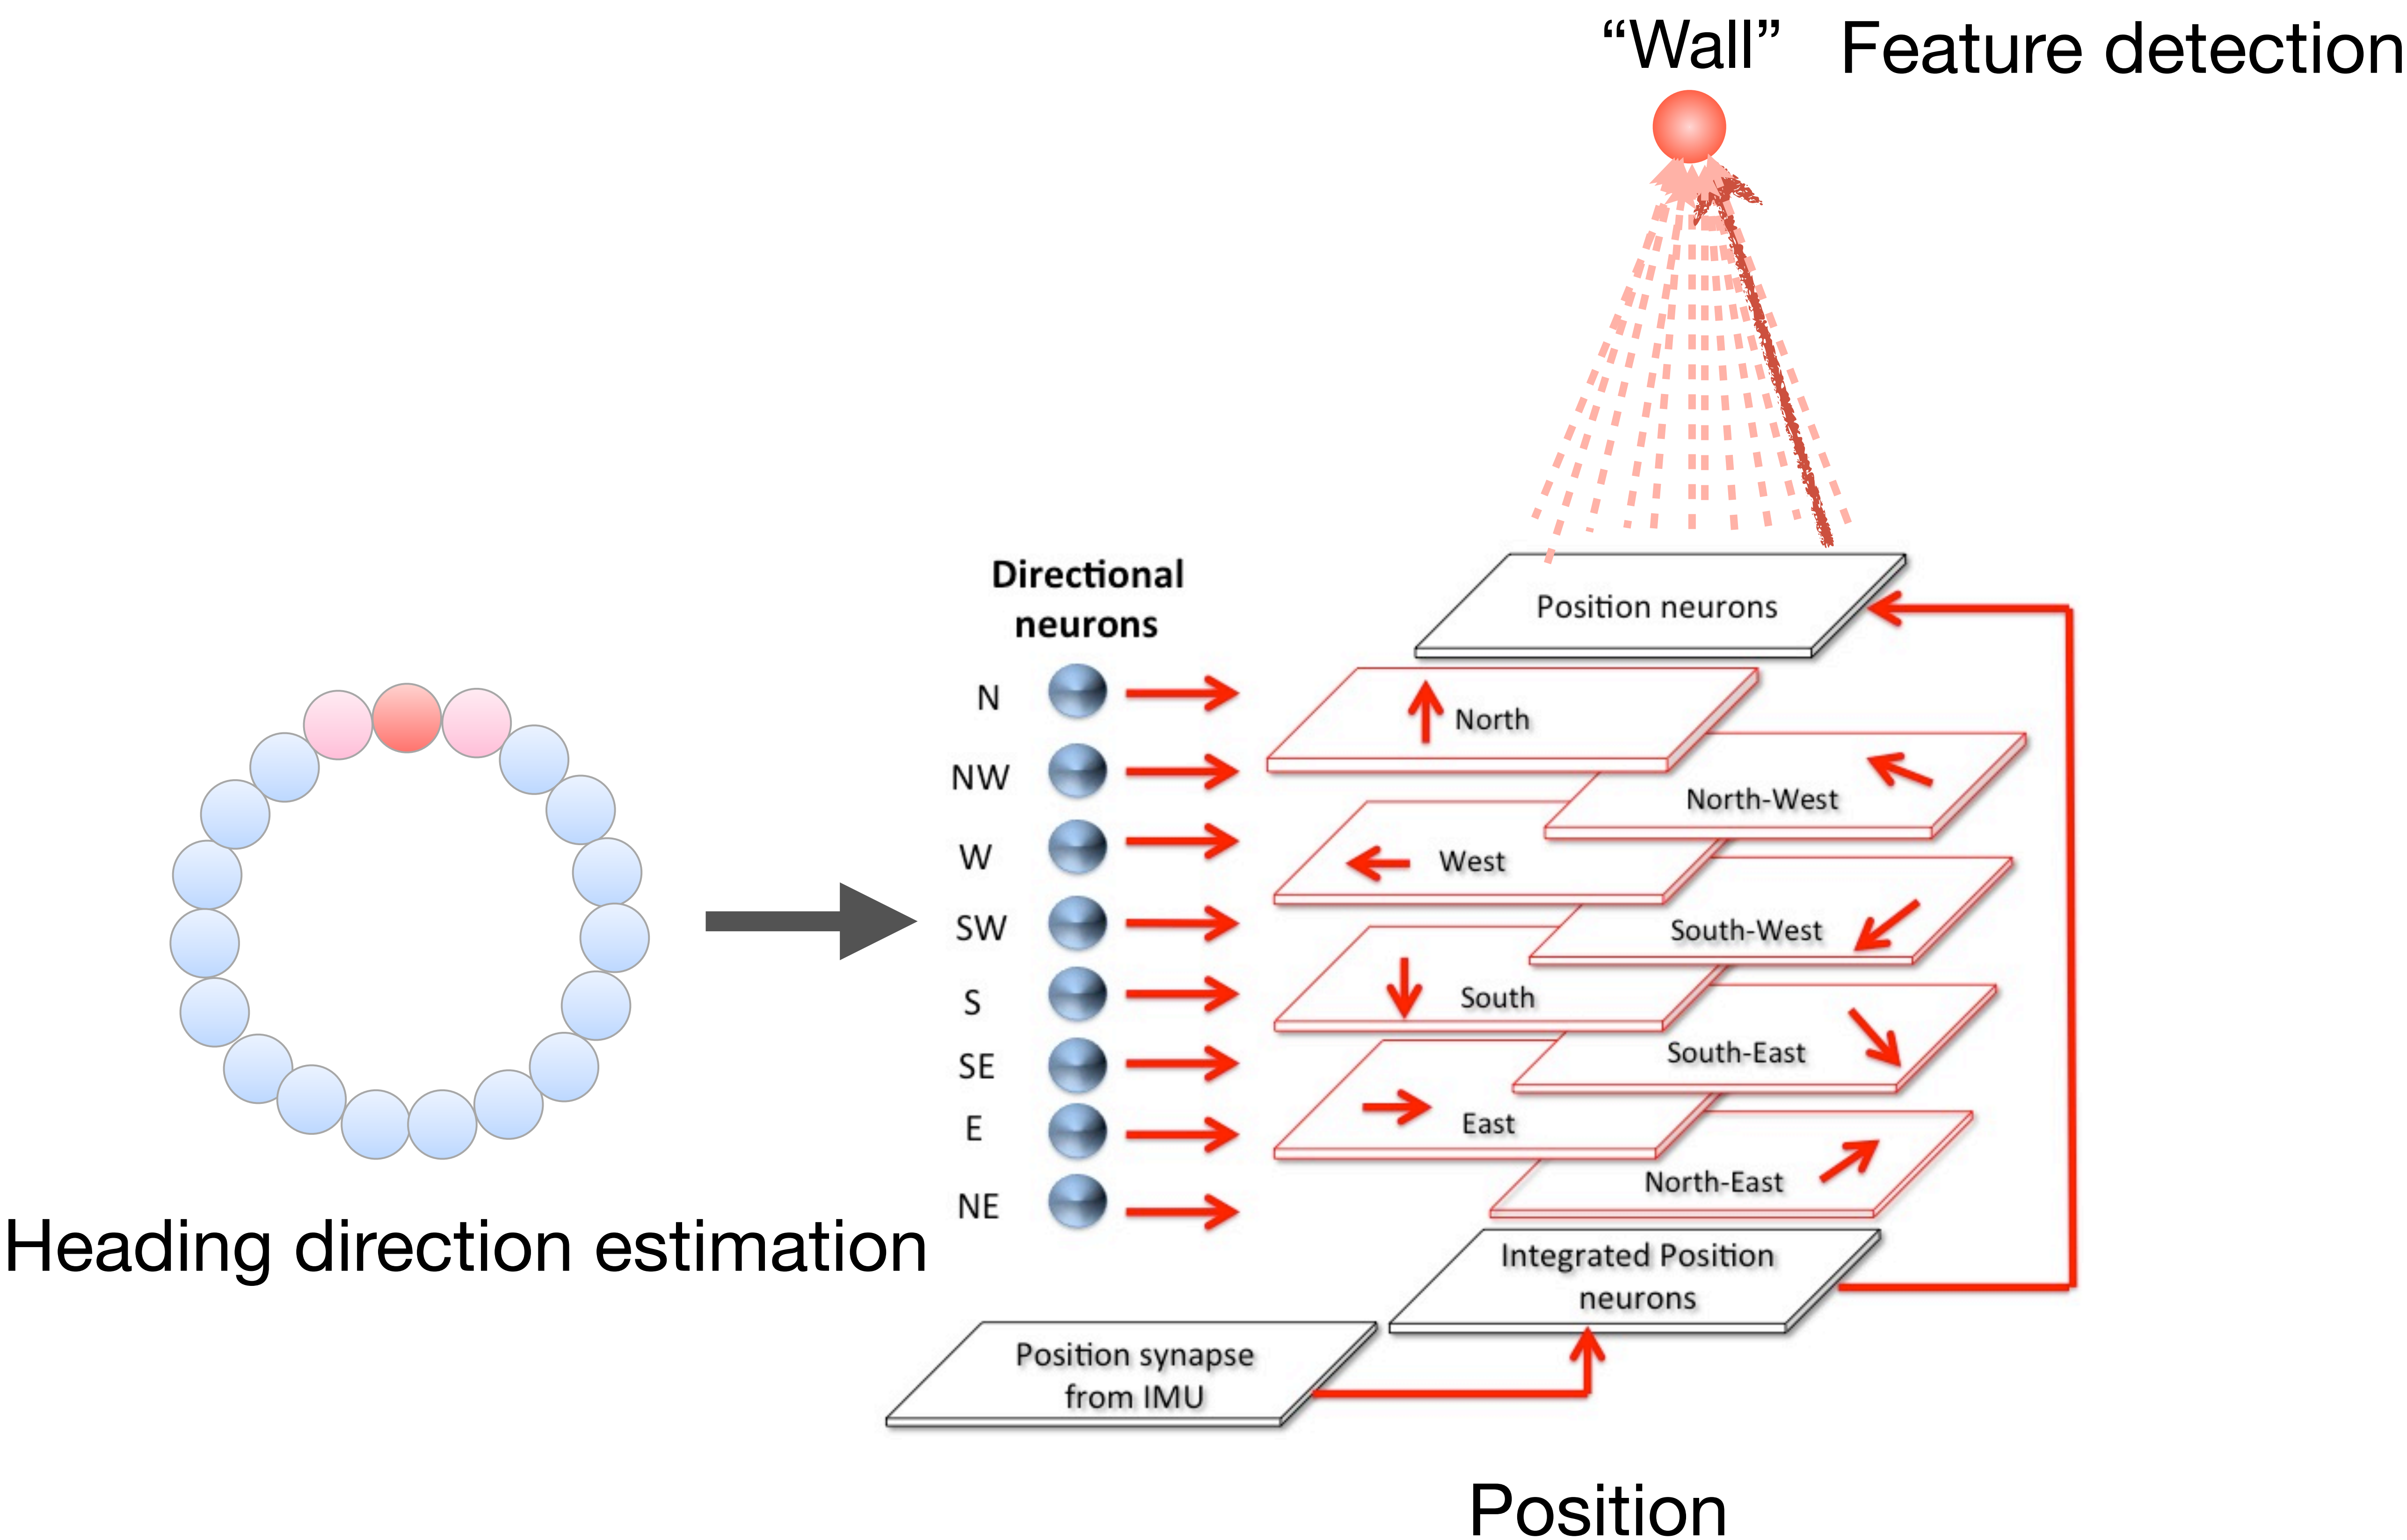

Map formation

Map (LTM) formation

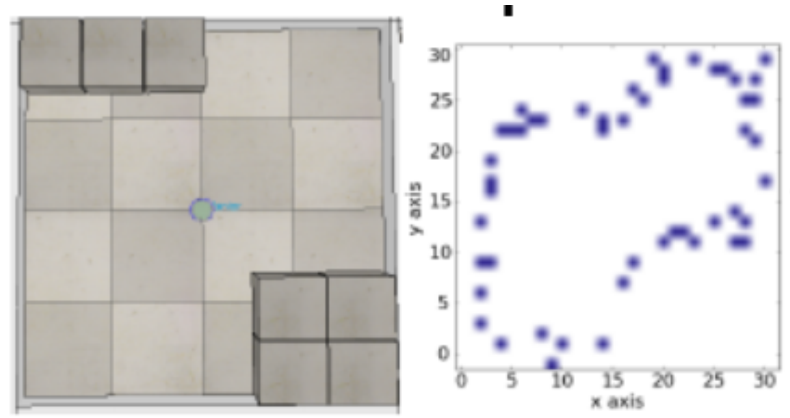

(d) Environment2 (e) Encoded map2

“Wall” Feature detection

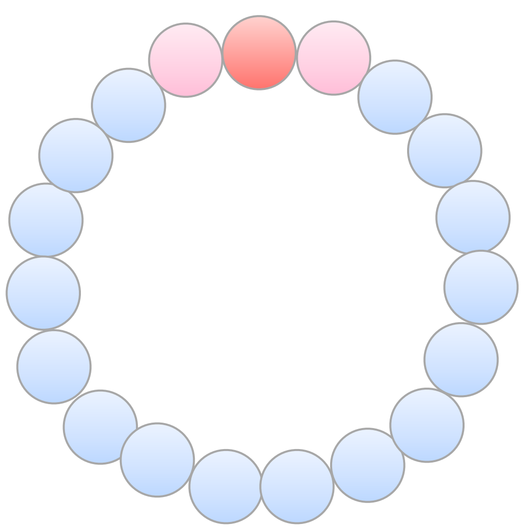

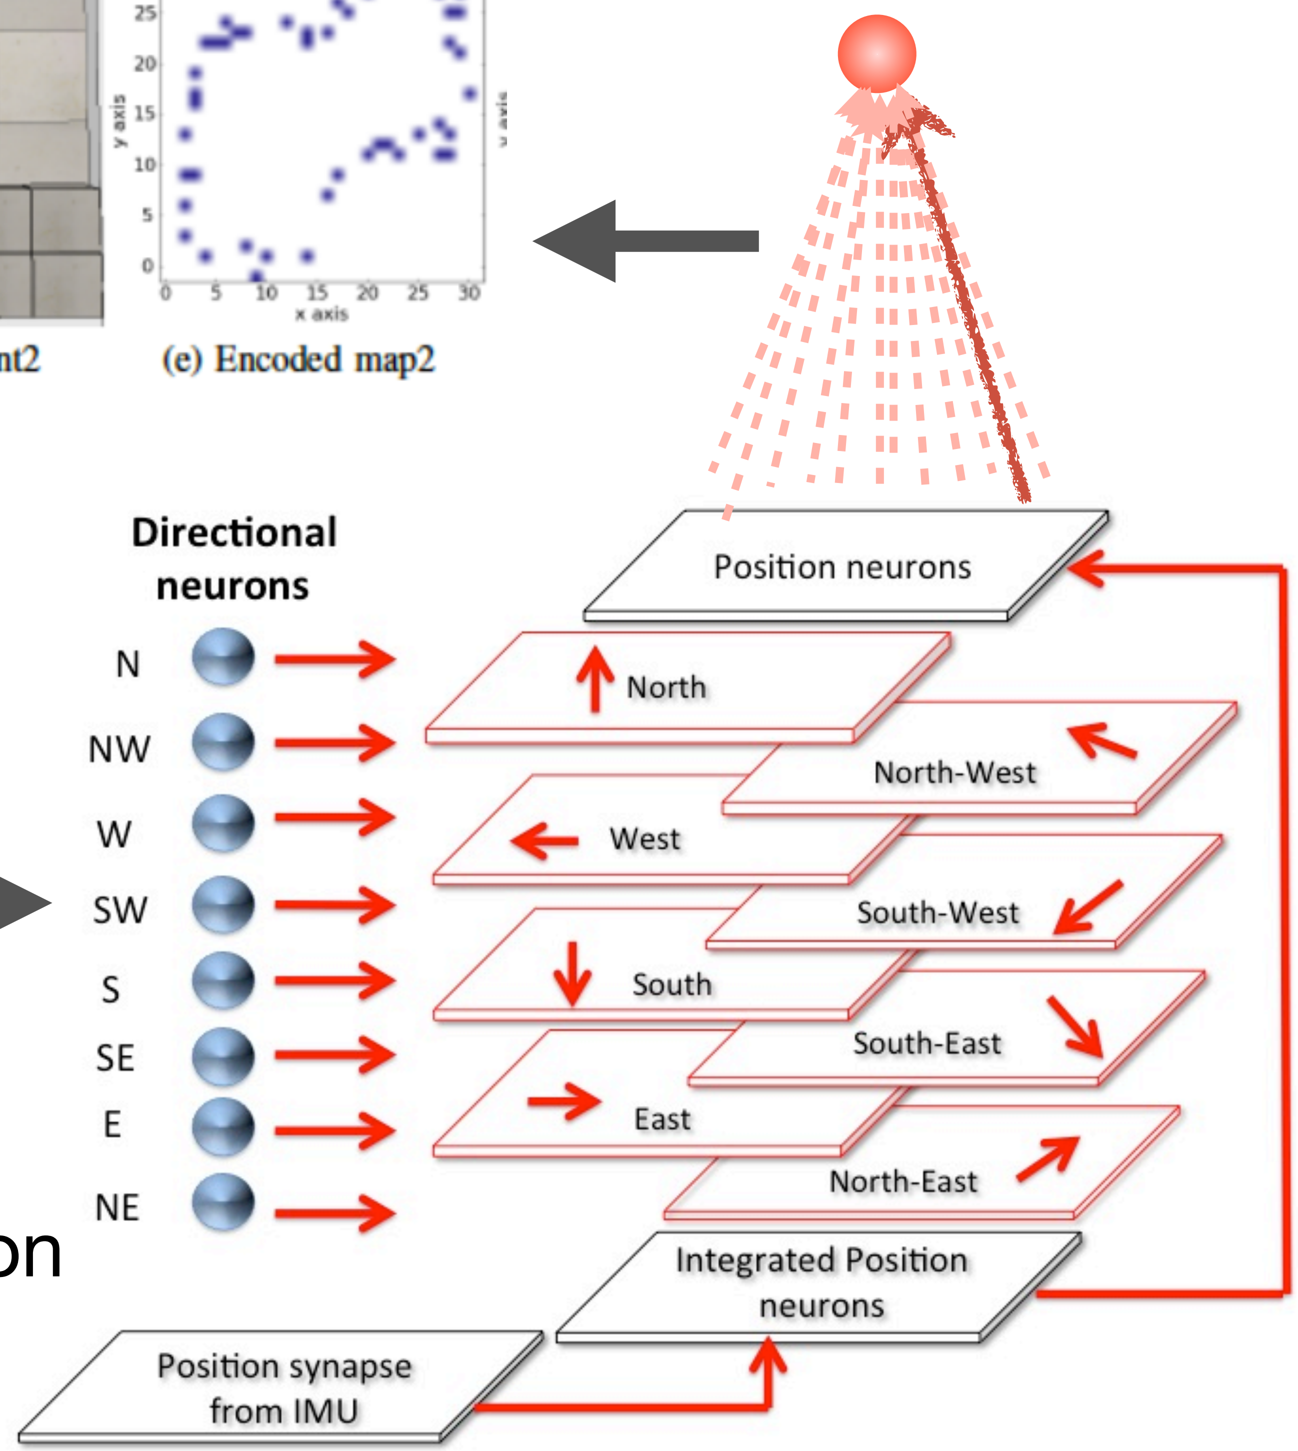

Heading direction estimation

Position

Map formation on the ROLLS chip

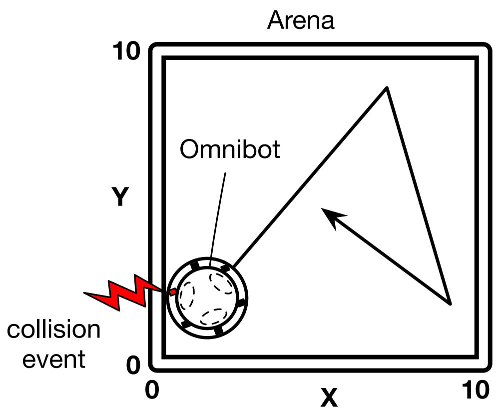

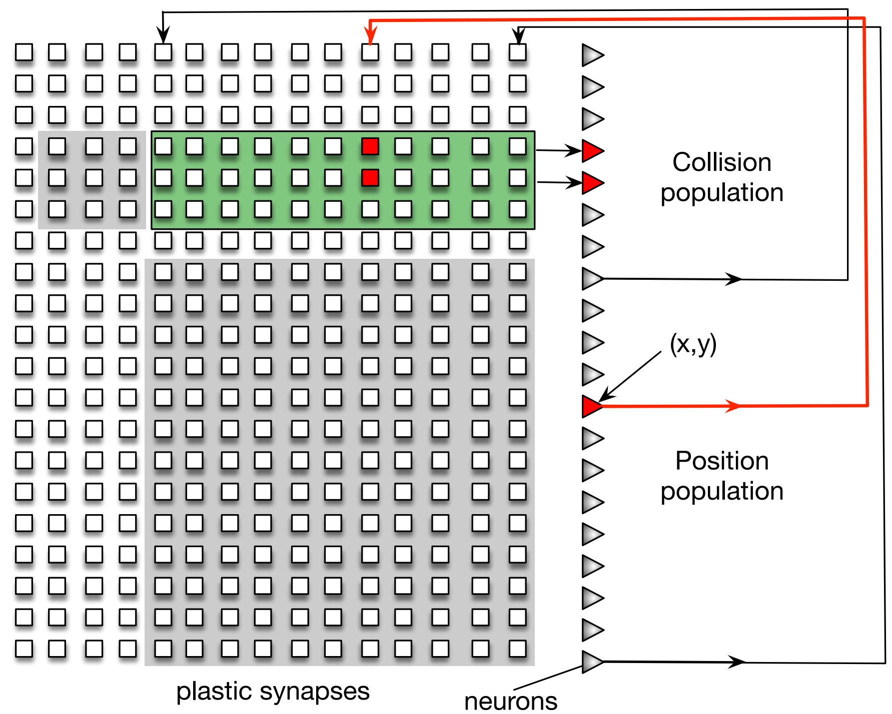

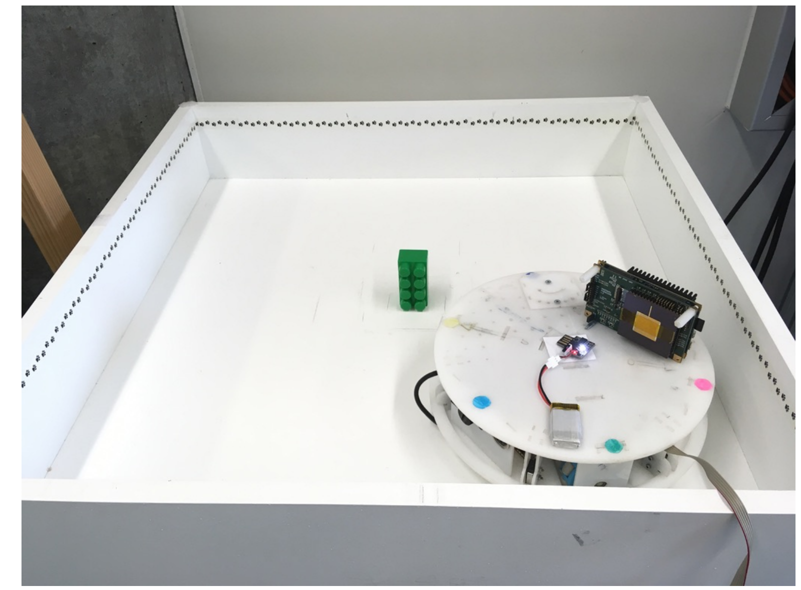

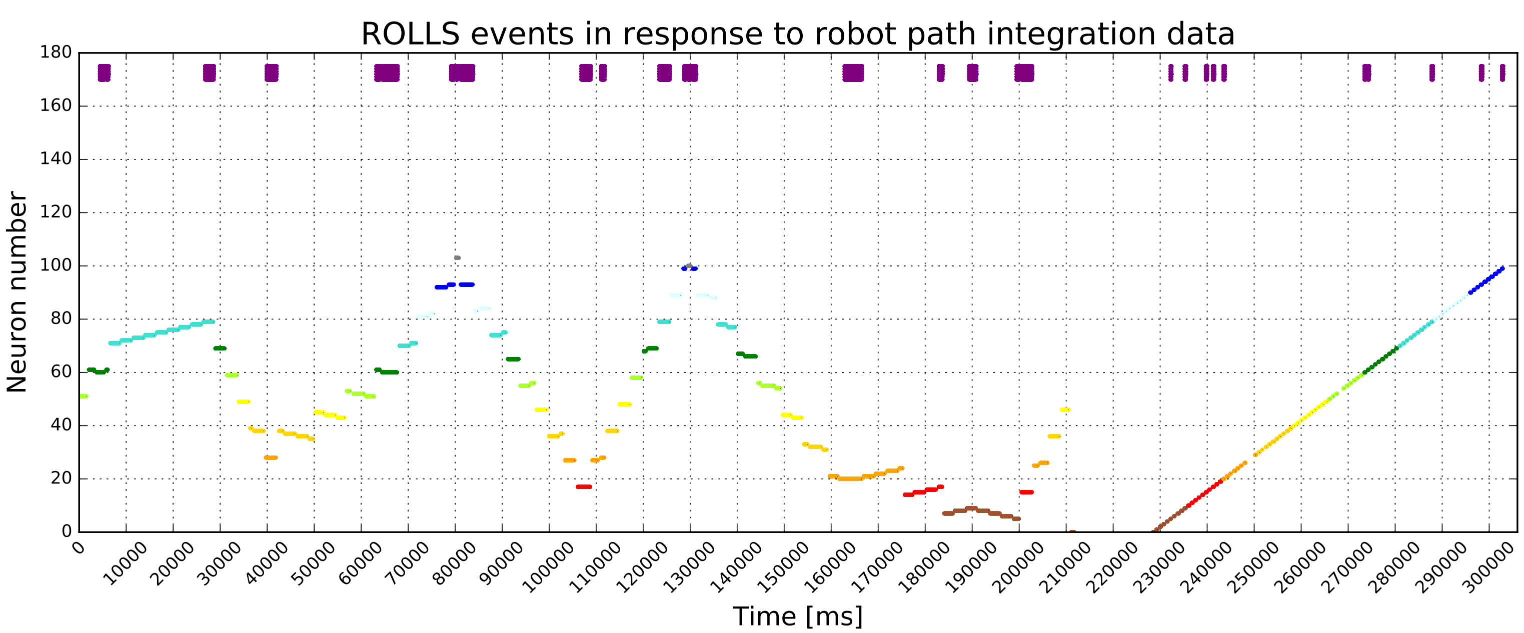

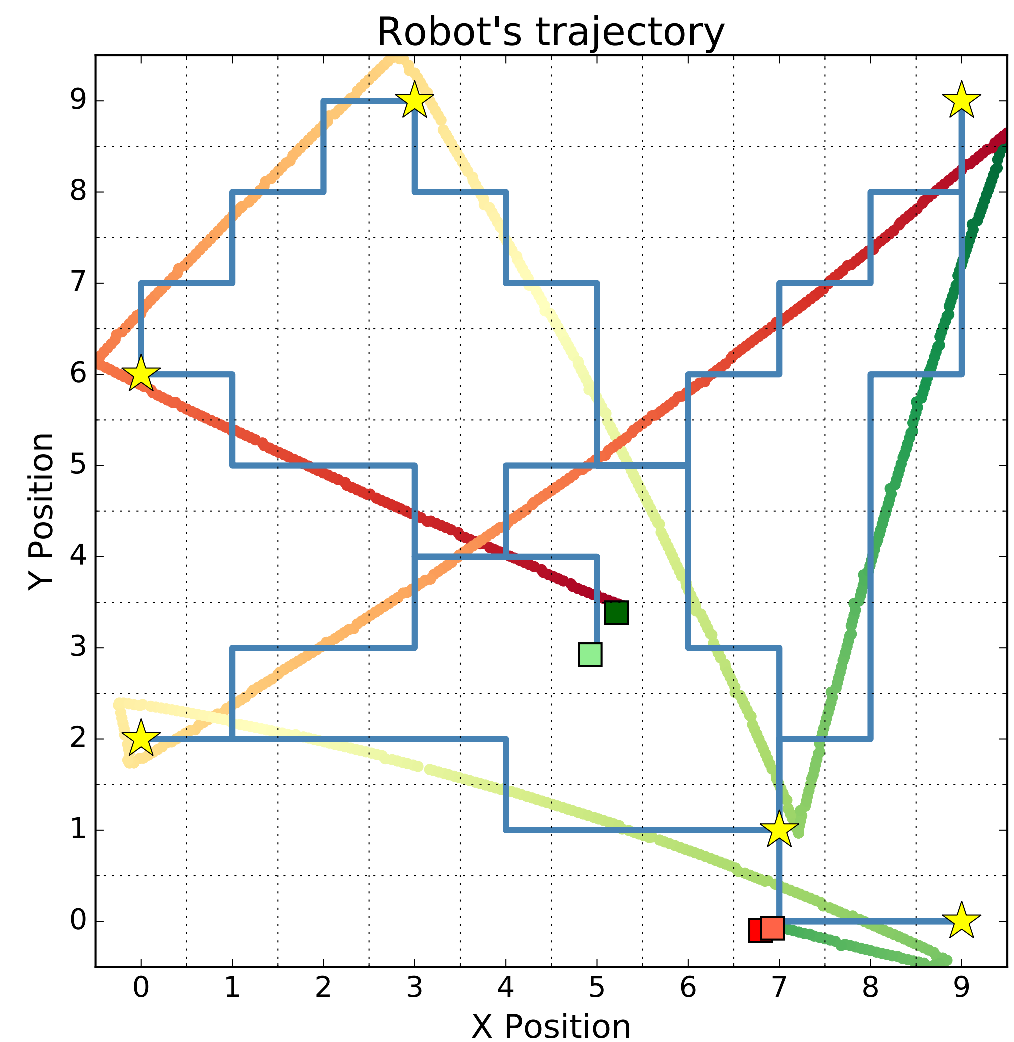

Map formation: Path integration in 2D

Learning different maps

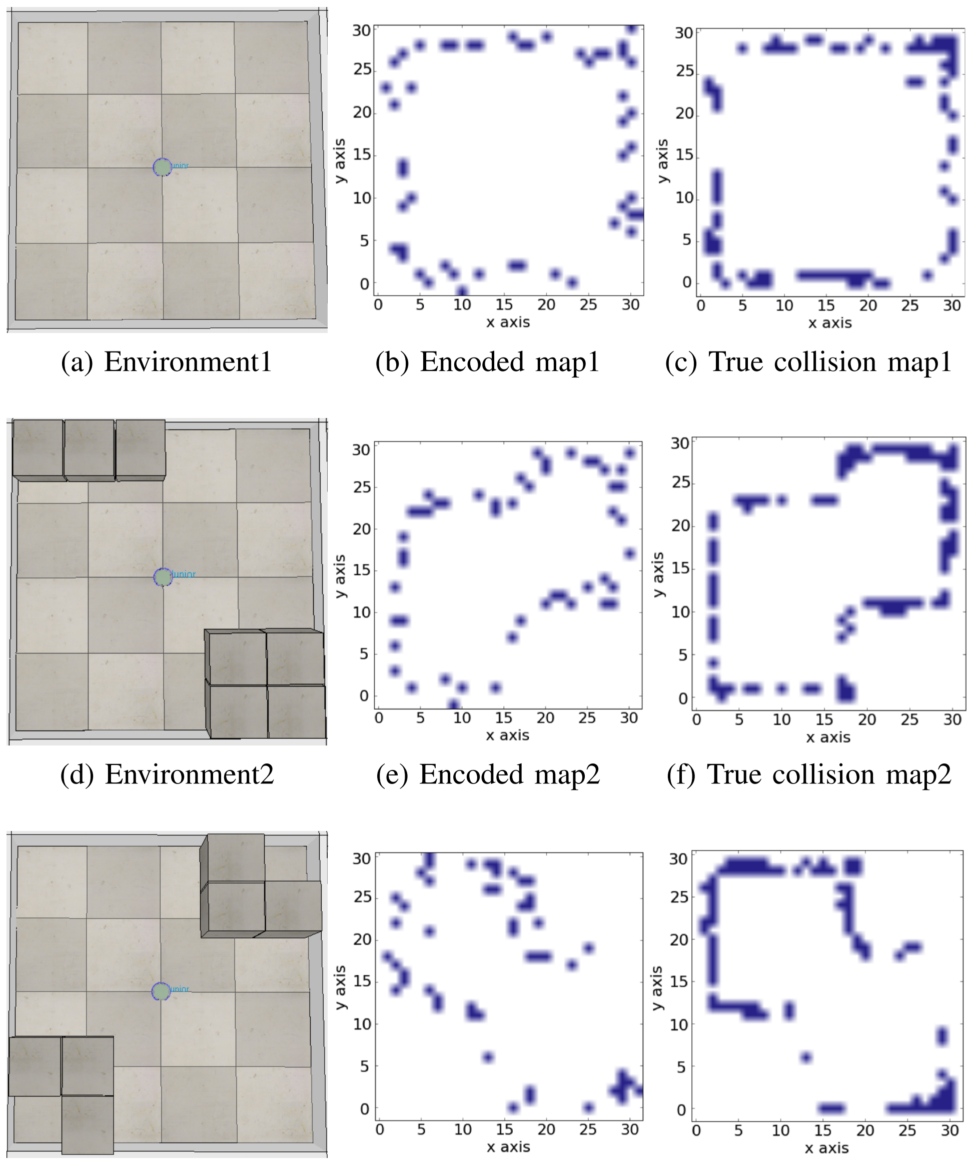

Unlearning a map

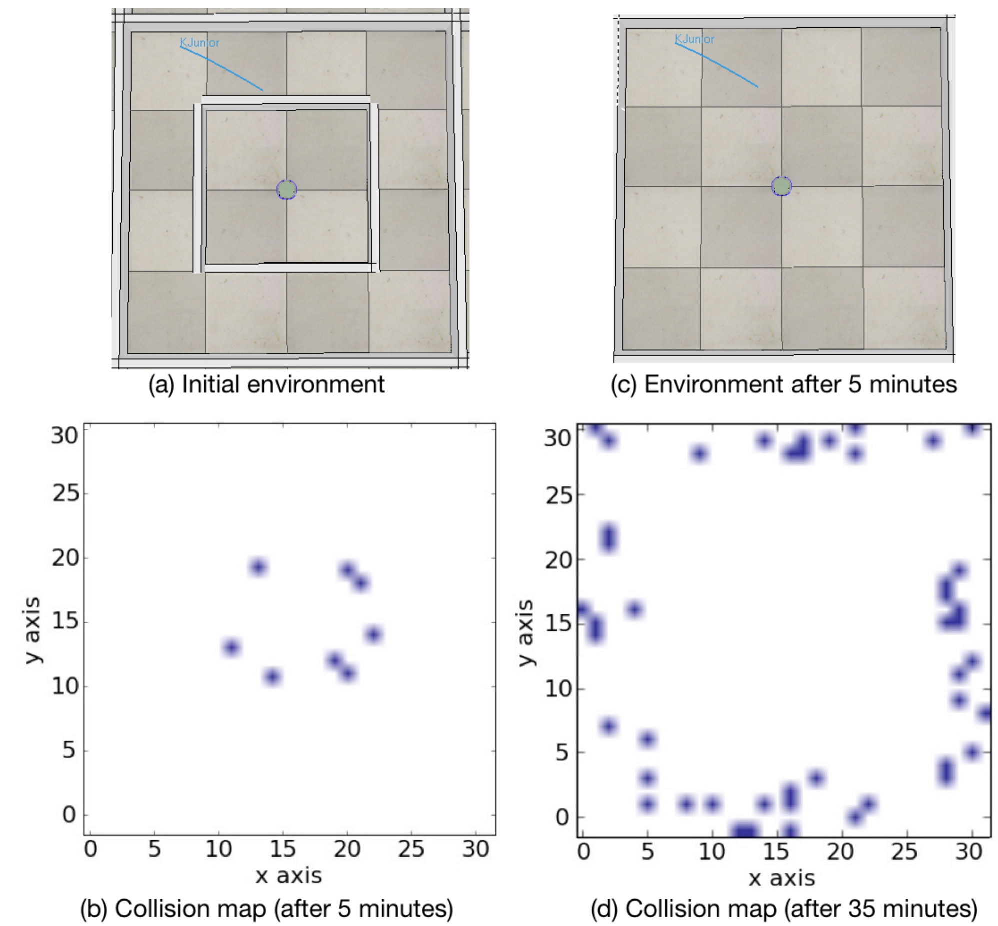

Kreiser, R.; Pienroj, P.; Renner, A. & Sandamirskaya, Y. Pose Estimation and Map Formation with Spiking Neural Networks: towards Neuromorphic SLAM. 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS, **2018**

How can we unlearn something?

- LTD and LTP depends on both pre- and postsynaptic frequencies

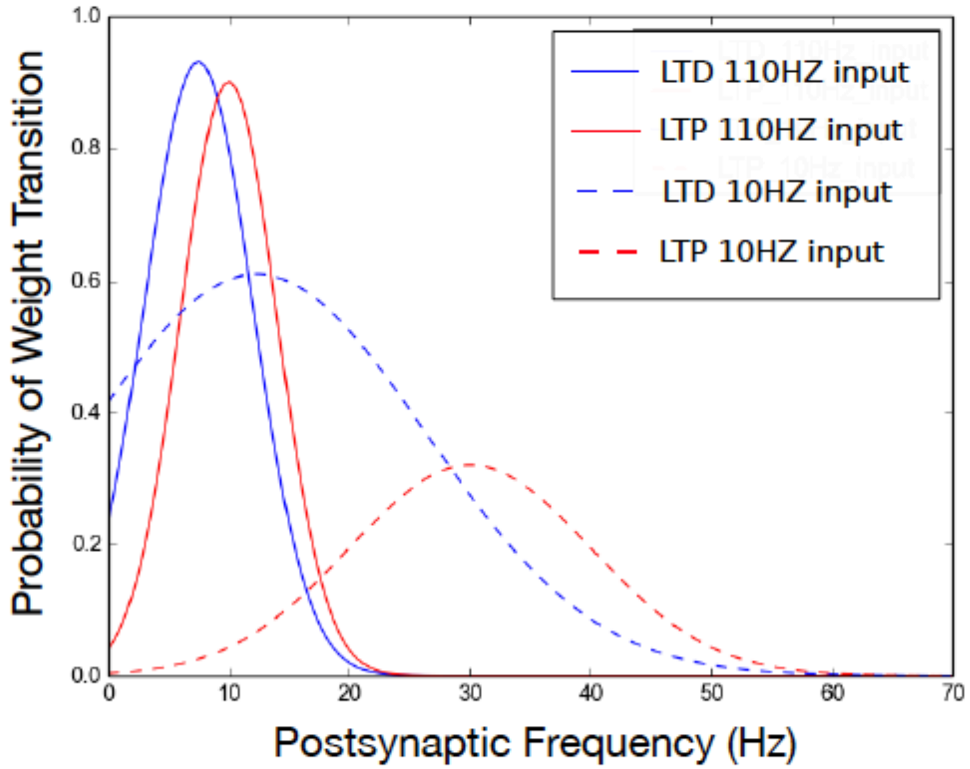

Stochastic weight update	$w_i = w_i + \Delta * w^+ \text{ if } V_{mem}(t_{pre}) > \theta_{mem} \text{ and } \theta_1 < Ca(t_{pre}) < \theta_3$
	$w_i = w_i - \Delta * w^- \text{ if } V_{mem}(t_{pre}) < \theta_{mem} \text{ and } \theta_1 < Ca(t_{pre}) < \theta_2$
Drift	$\frac{d}{dt} w_i = +C_{drift} \text{ if } w_i > \theta_w \text{ and } w_i < w_{max}$
	$\frac{d}{dt} w_i = -C_{drift} \text{ if } w_i < \theta_w \text{ and } w_i > w_{min}$
Binary weight of the synapse	$J_i = J_{max} f(w_i, q_J)$

(Brader, Senn, and Fusi, 2007)

Unlearning false collisions on chip

Ground truth recording

With obstacle

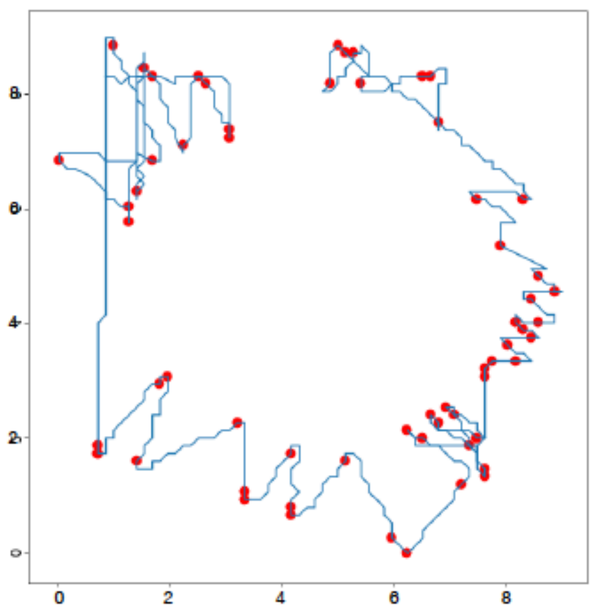

Without obstacle

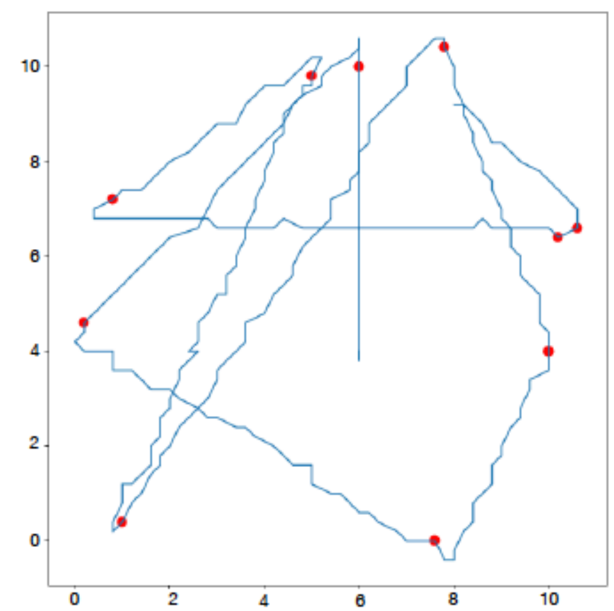

Learned collision maps

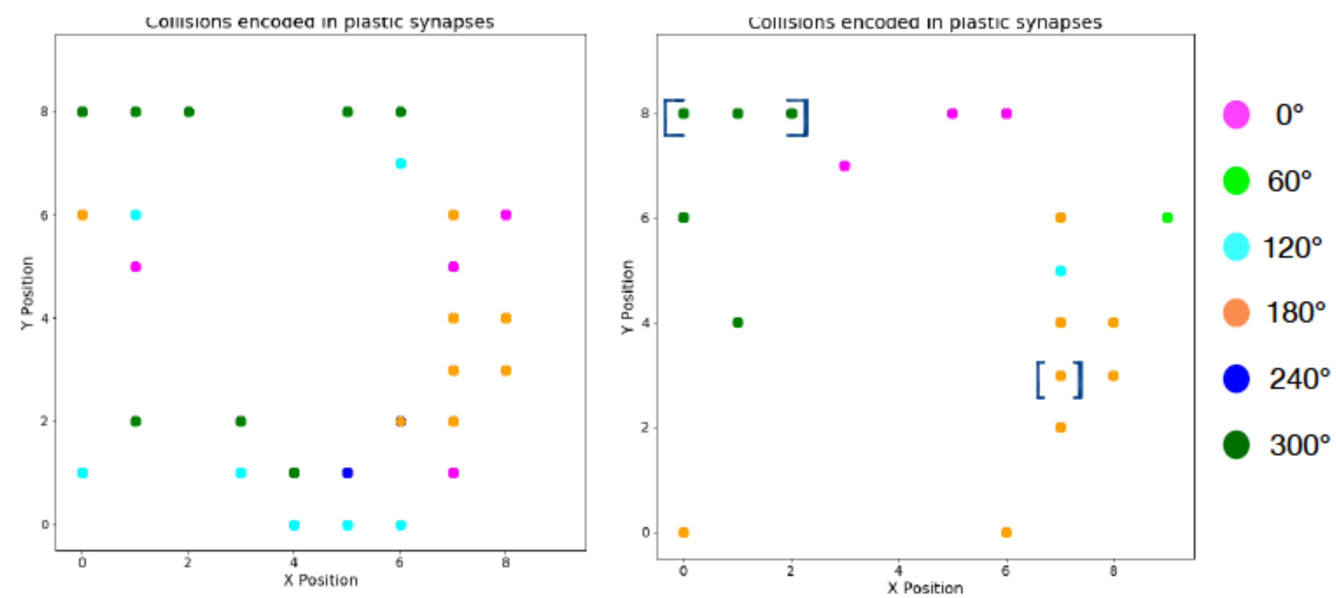

Trajectory recorded from firing activity

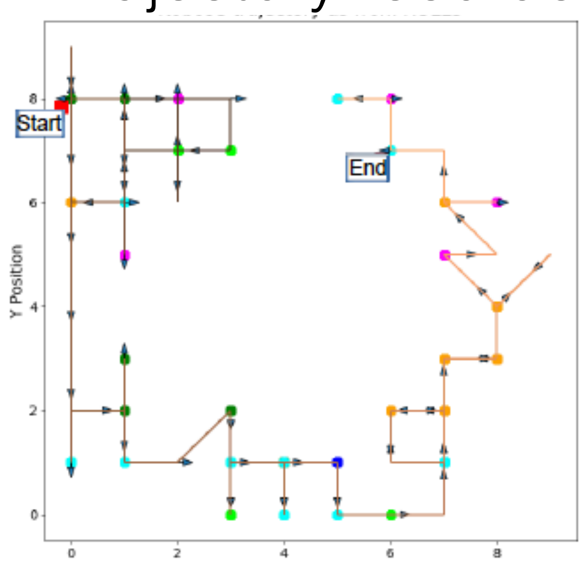

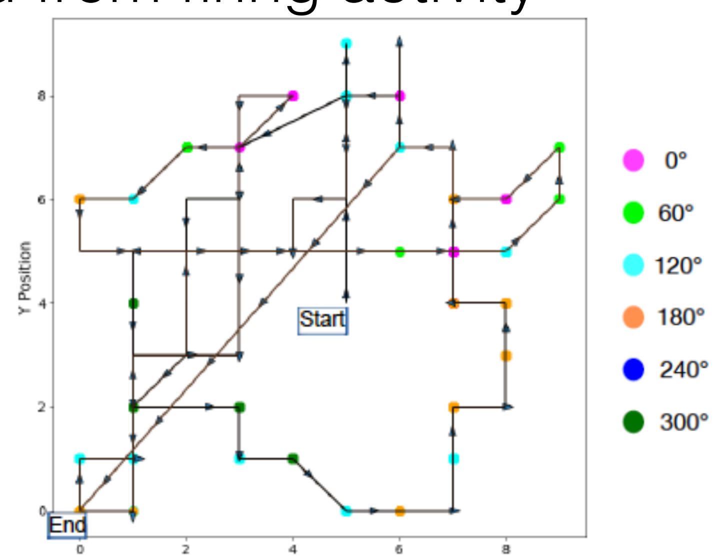

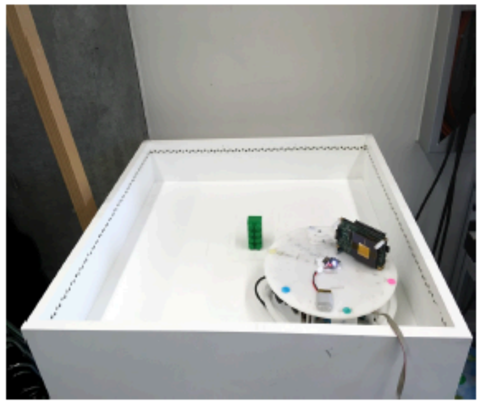

Robot arena with obstacle

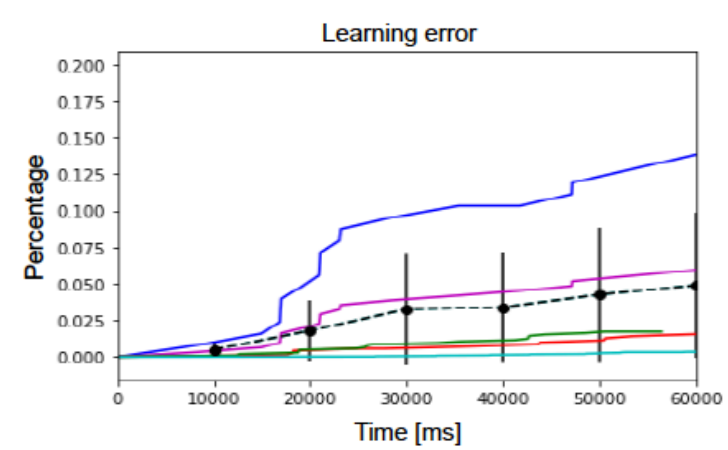

NICE?

➡ Simultaneous **localisation and mapping**: path integration, learning a map

- state estimation, building representations

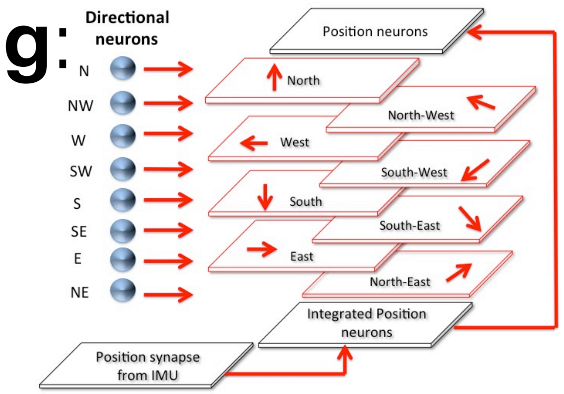

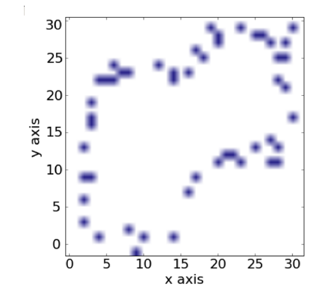

Kreiser et al 2018a, b

Blatter et al, ISCAS, under rev;

➡ **Braintenberg vehicle, sequences**

- attractors in a sensory-motor loop

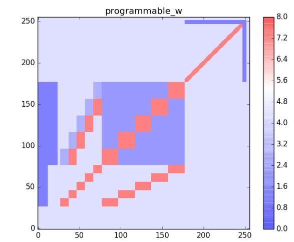

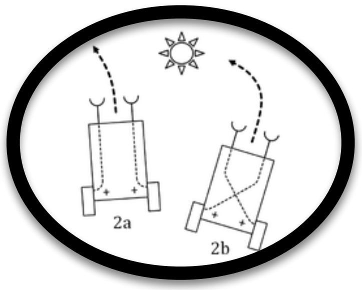

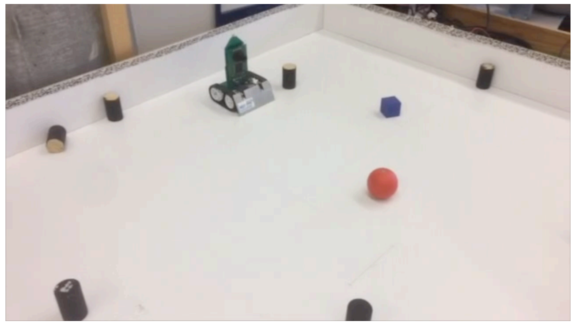

Milde et al 2017a,b; Kreiser et al 2018; Blum et al 2017

➡ **Reference frame** transformations

- key for linking modalities

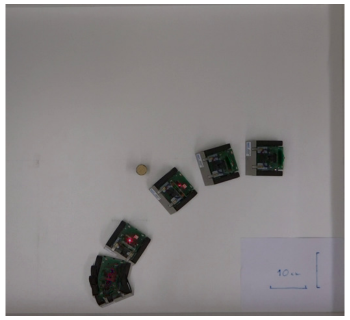

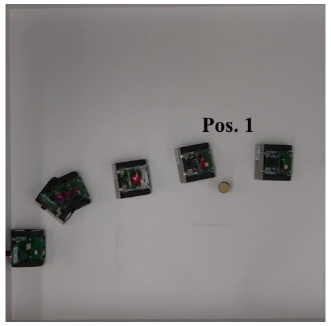

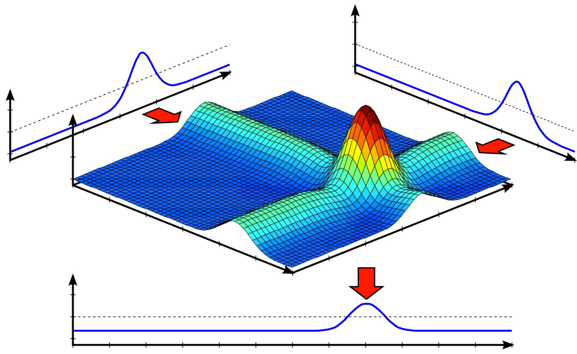

➡ **Adaptive motor control**

- key element for adaptive behavior

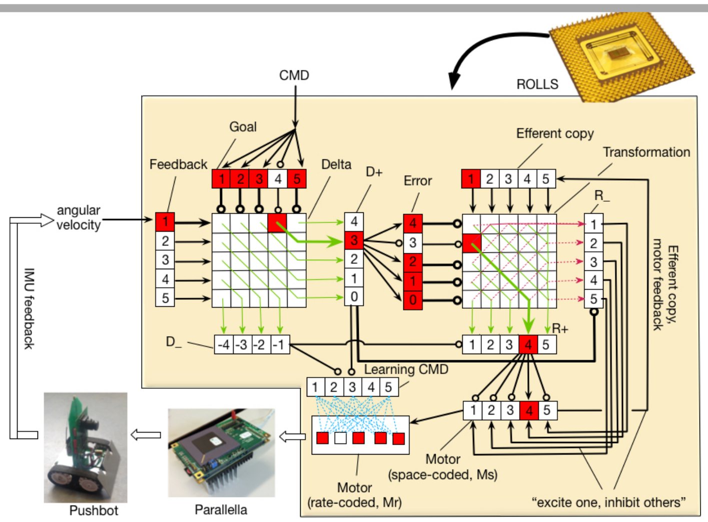

Glatz et al, arxiv, 2018

Conclusions

➡ lots of structure is needed to control behavior with neurons

- represent state with neuronal populations (“place code”)
- stabilise states and decision with recurrent connections (WTA)
- disinhibition for robustness
- adaptive couplings between sensed quantities and states
- error estimation and correction

➡ learning can then be very simple

- one-shot
- binary weight

➡ object representation as a map-formation problem, not a (just) pattern recognition

Thanks!

HORIZON 2020

- Marie Curie IF
- FET PROACT

FNSNF

- Ambizione
- Project coordination

Universität Zürich[UZH]

- Forschungskredit
- GRC Grant

ZNZ
Zentrum für Neurowissenschaften Zürich

- Junior Group fellowship

NEU COG ROB

PhD Students

Julien Martel
Alpha Renner*
Raphaela Kreiser*
Claudius Strub*
Moritz Milder
Dora Sumislawska

MSc, BSc theses

Gwendolyn English
Eloy Barrero
Llewyn Salt
Mathis Richter
Tobias Storck
Christian Bell
Claudia Rudolph
Jianlin Lu
Ammar Bitar
Jonathan Müller
Kay Müller
Sebastian Glatz
Valery Metry
Alpha Renner
David Niederberger
Raphaela Kreiser

Semester theses

Alexander Dietmüller
Mario Blatter
Frédéric Debraine
Lukas Blässig
Lennard de Graf
Michel Frising
Zahra Farsijani
Michael Purcell
Viviane Yang
Davide Plozza
Damiano Steger

Héctor Vazquez
Sebastian Glatz
Herman Blum
Matteo Cartiglia
Lin Jin
David Niederberger
Nicolas Känzig
Panin Pienroj
Paul Joseph
Nuria Armengol
Jozef Bucko
Balduim Dettling

Collaborators

Gregor Schöner
Giacomo Indiveri
Florentin Wörgötter
J.-C. Quinton
John Spencer
Piotr Dudek
Fatih Yanik
Jörg Conradt
Christian Faubel
Tobi Delbruck
John Lipinski
Richard Hahnloser
Matt Luciw
Jürgen Schmidhuber
Helge Ritter
Hajar Azgari
Abhishek Banerjee

Elementary module of neuronal control

"Reaching" task

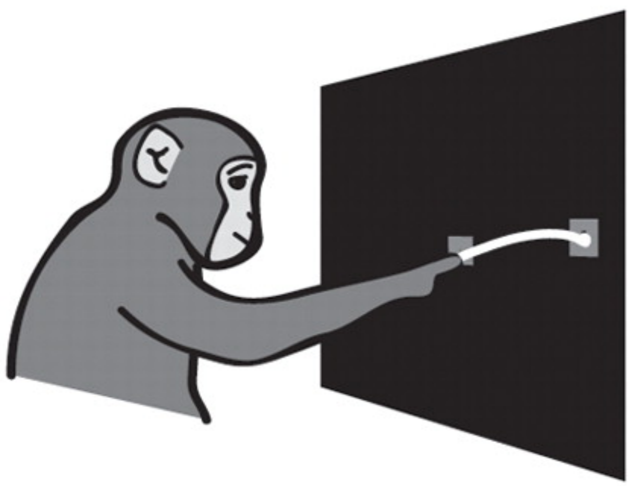

Distribution of population activation =
$\sum_{\text{neurons}}$ tuning curve * current firing rate

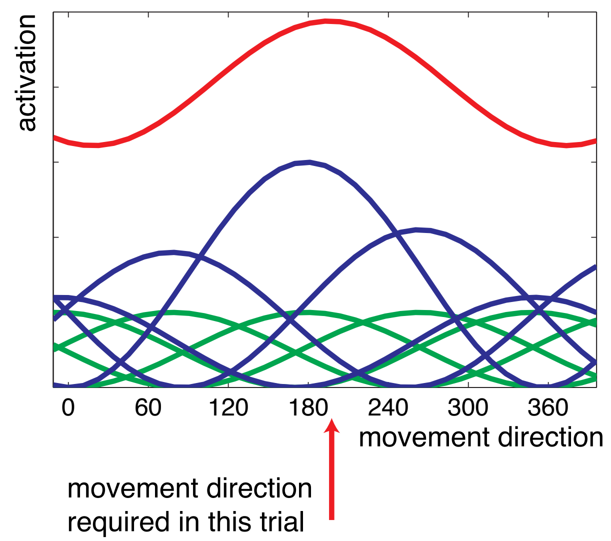

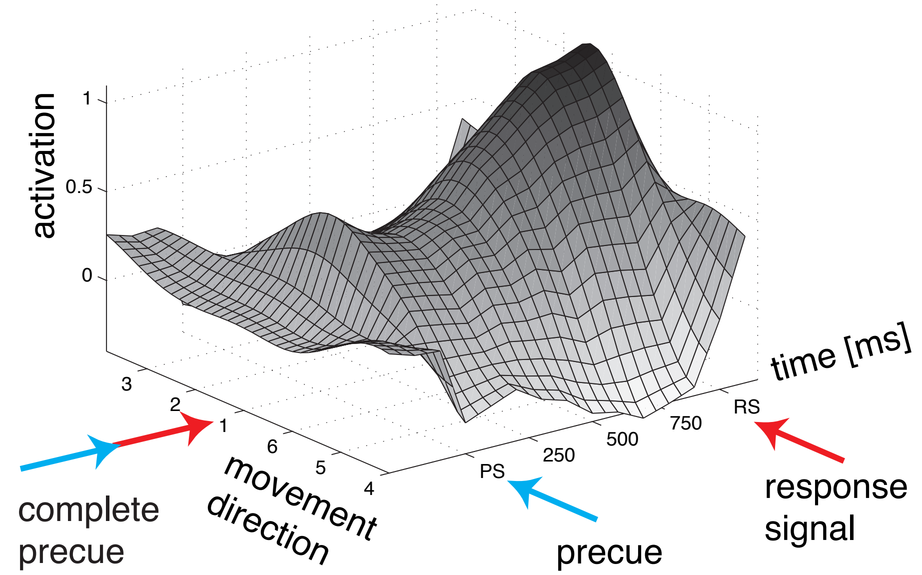

➡ Population activity dynamics:

$$\tau \dot{u}(x,t) = -u(x,t) + h + \int f\big(u(x',t)\big)\omega(x-x')dx' + I(x,t)$$

Amari, S. **Dynamics of pattern formation in lateral-inhibition type neural fields**. Biological Cybernetics, **1977**, 27, 77-87

Wilson, H. R. & Cowan, J. D. **A mathematical theory of the functional dynamics of cortical and thalamic nervous tissue**. Kybernetik, **1973**, 13, 55-80

Gerstner, Grossberg, Ermentrout, Coombes, Schöner&Spencer, Erlhagen…

Neural dynamics

Dynamic Neural Field, WTA, bump-attractor networks

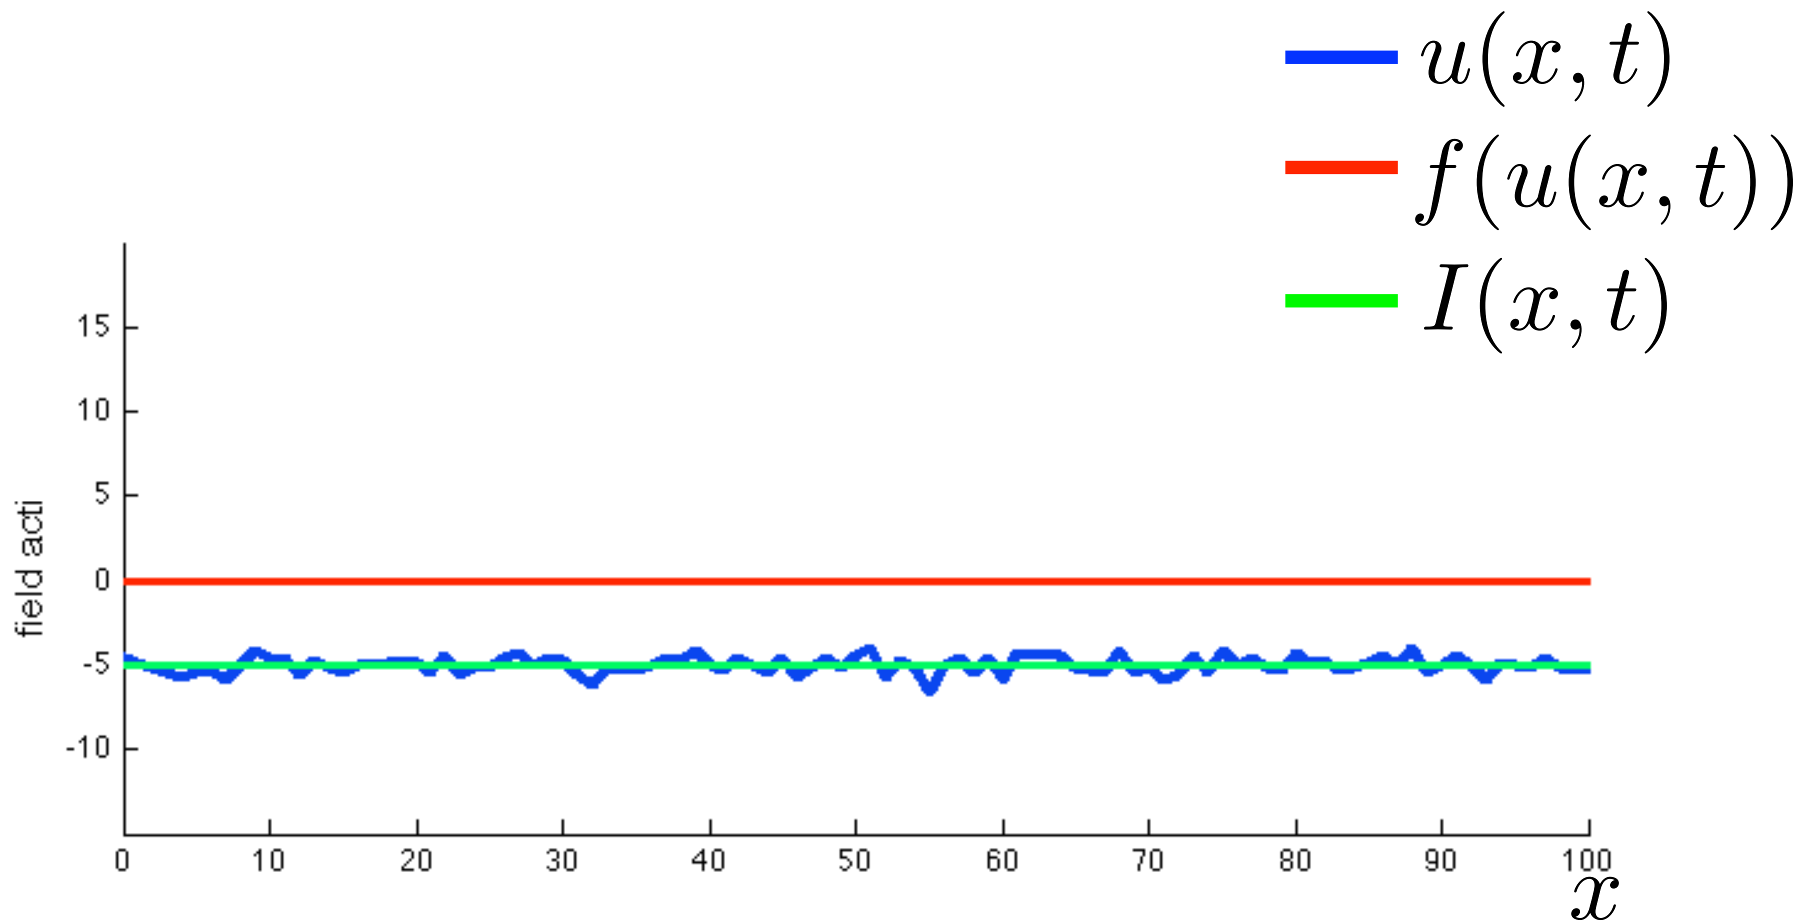

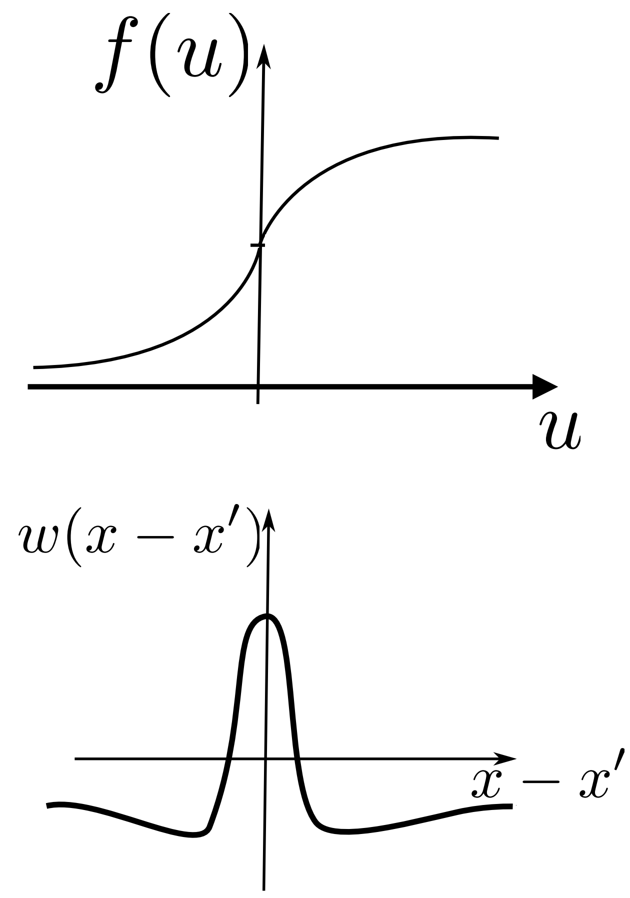

$$\tau\dot{u}(x,t) = -u(x,t) + h + \boxed{\int f(u(x',t))\omega(x-x')dx'} + I(x,t)$$

“Cognitive” properties of Neural Fields

- “Detection” and “forgetting” instabilities
  - continuous time ⟶ discrete “events”
- Localised “bumps”
  - continuous space ⟶ discrete “categories”
- “Selection” instability
  - stabilisation of selection decisions
- Sustained activation
  - modelling working memory

➡ DNF “Architectures”

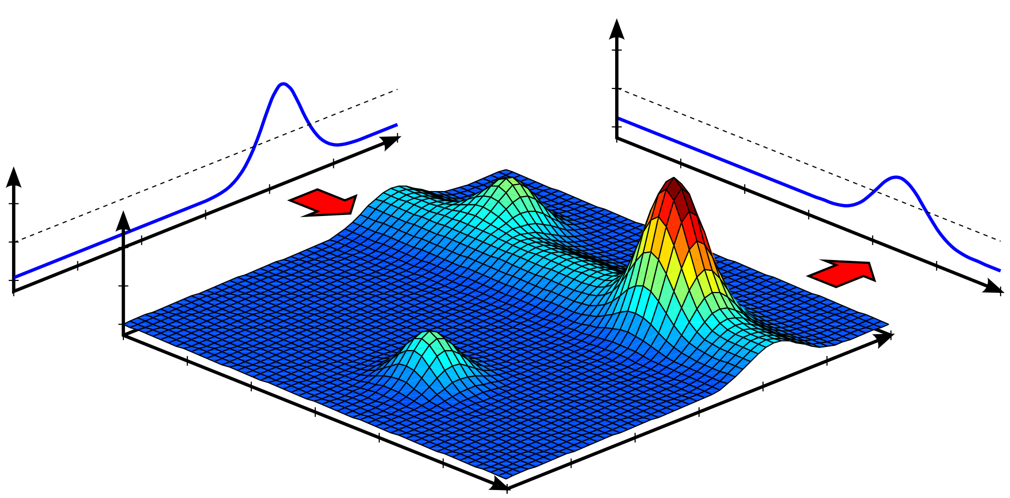

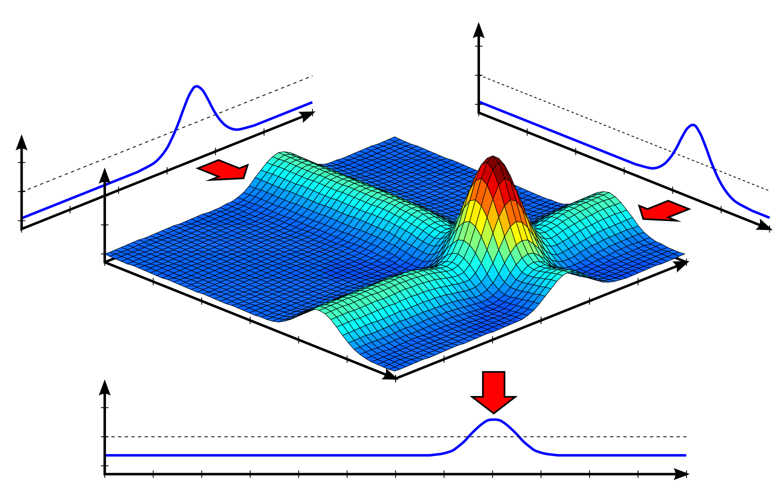

Embodied DNF architectures

Action selection

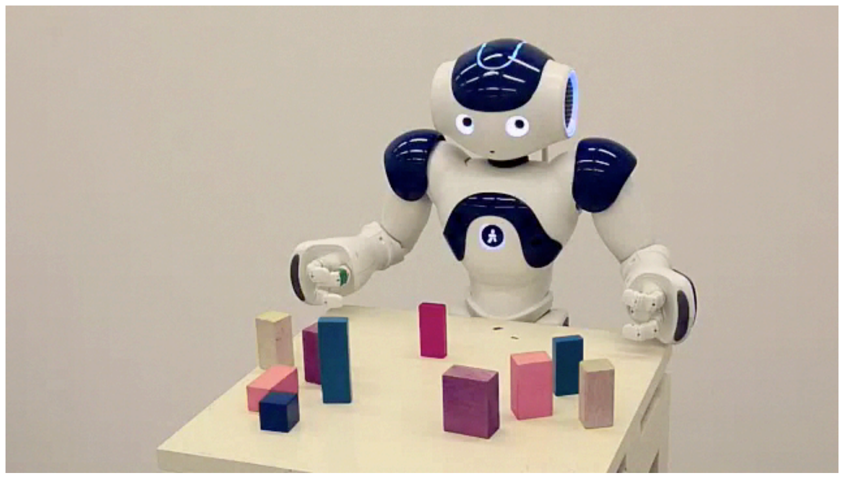

Learning to look

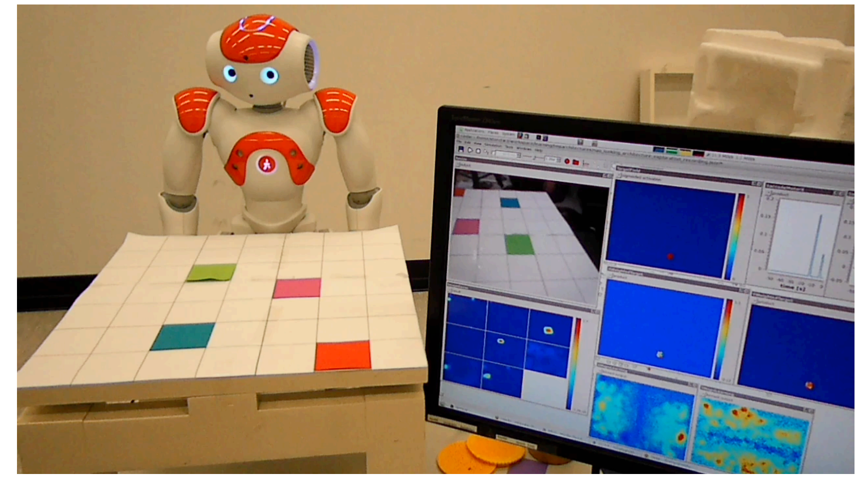

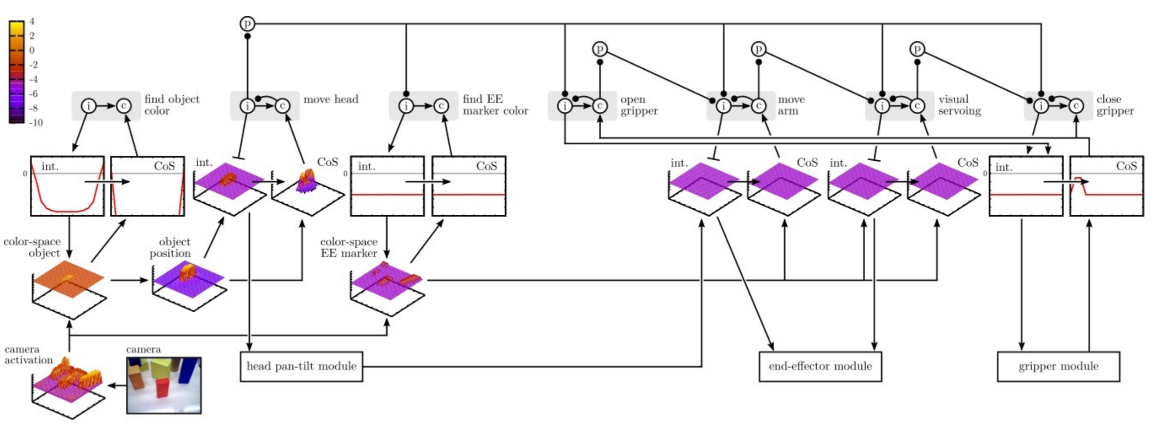

Sequence learning

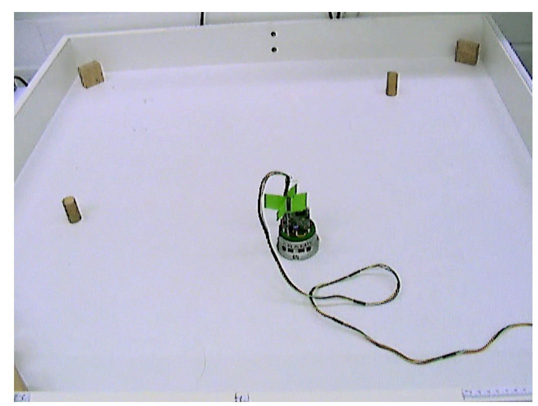

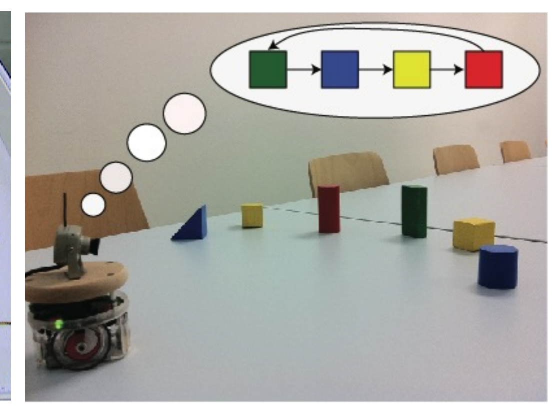

Haptic learning

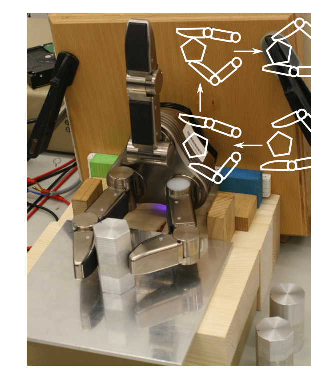

Planning & acting

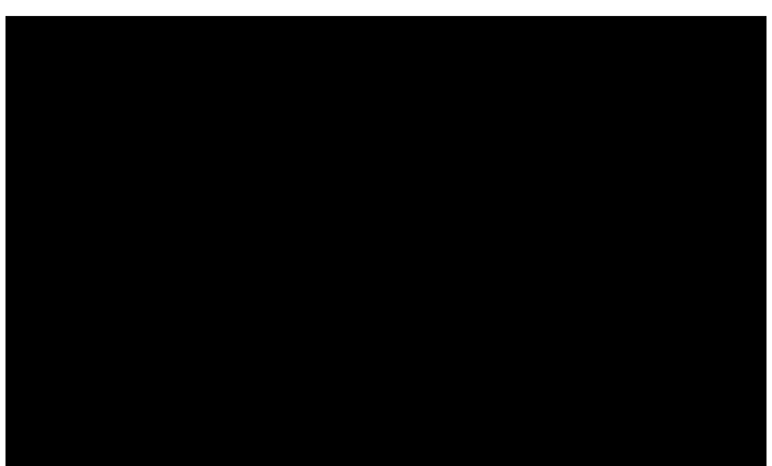

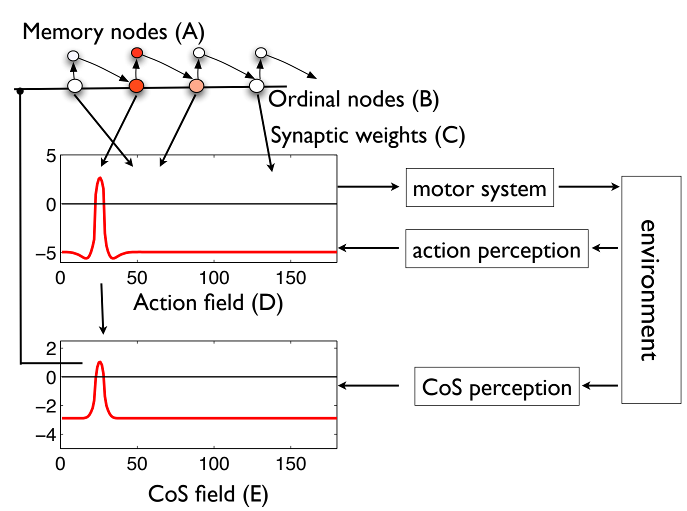

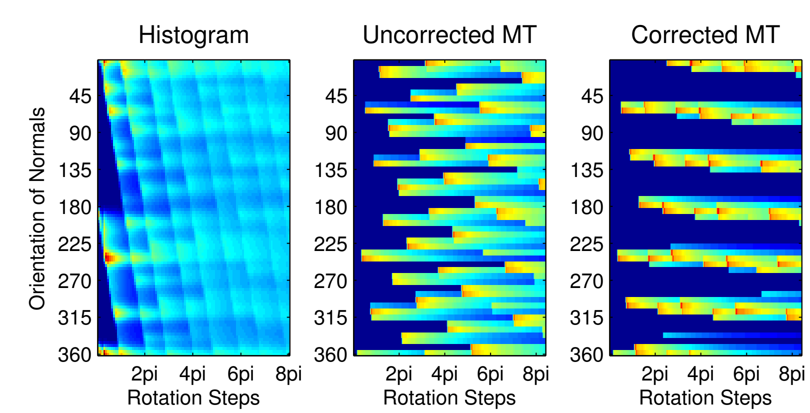

Sandamirskaya, 2010-2015

Why are these architectures fundamental?

➡ **Braintenberg vehicle, sequences**

- attractors in a sensory-motor loop

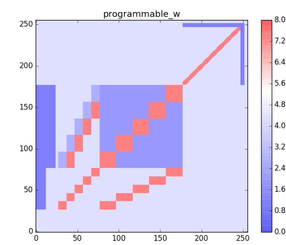

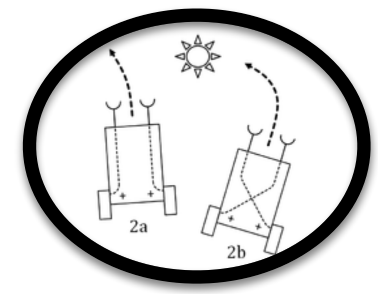

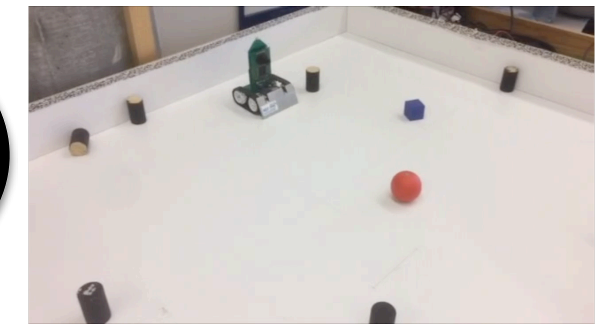

Milde et al 2017a,b; Kreiser et al 2018

➡ **Reference frame** transformations

- key for linking modalities

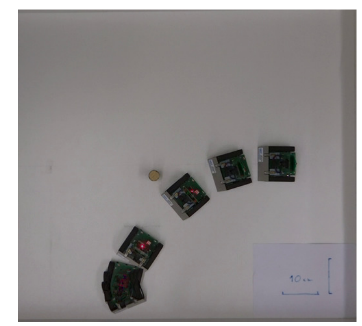

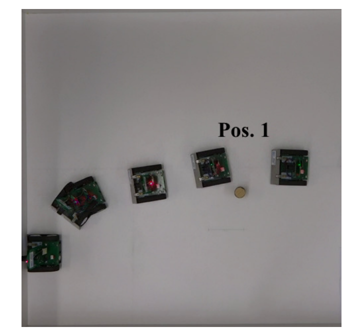

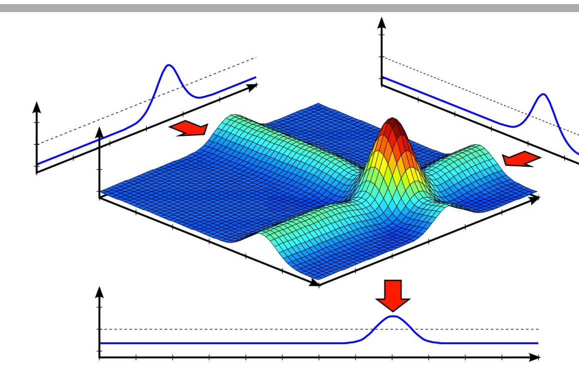

Blum et al 2017

➡ Simultaneous **localisation and mapping**: path integration, learning a map

- state estimation, building representations

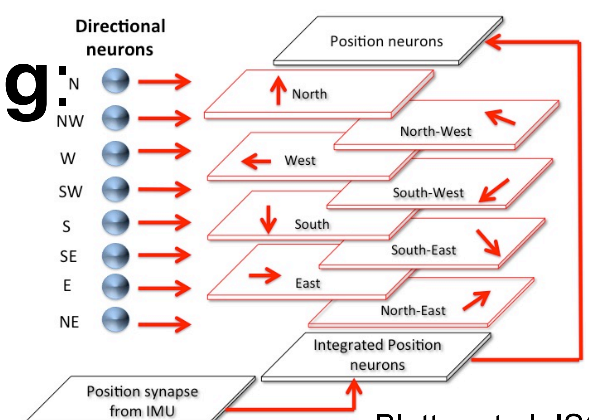

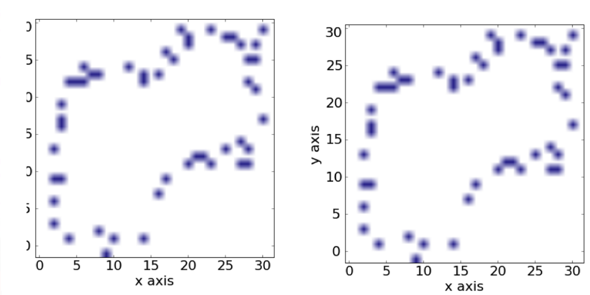

Blatter et al, ISCAS, under rev; Kreiser et al 2018a, b

➡ **Adaptive motor control**

- key element for adaptive behavior

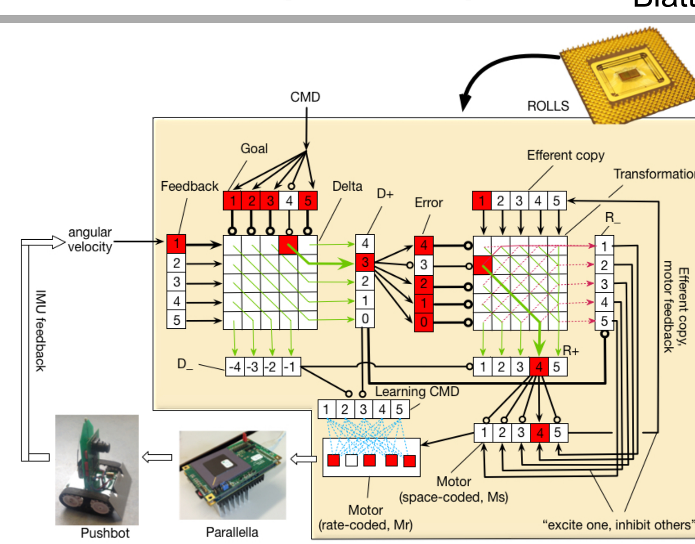

Glatz et al, arxiv, 2018

Reference frames transformation on chip

View-based target representation:

- target in view

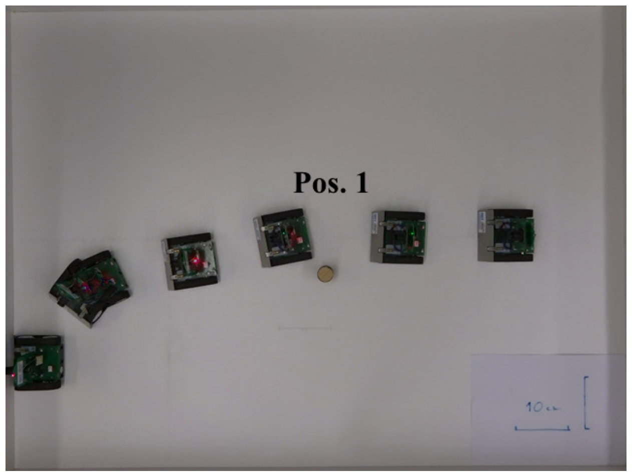

- target lost from view

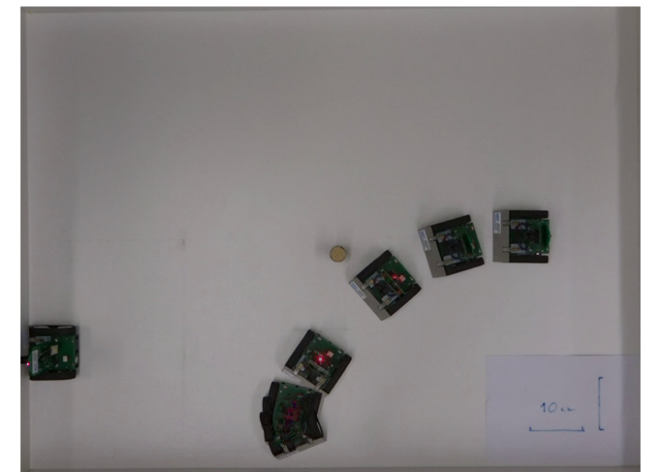

Allocentric target representation:

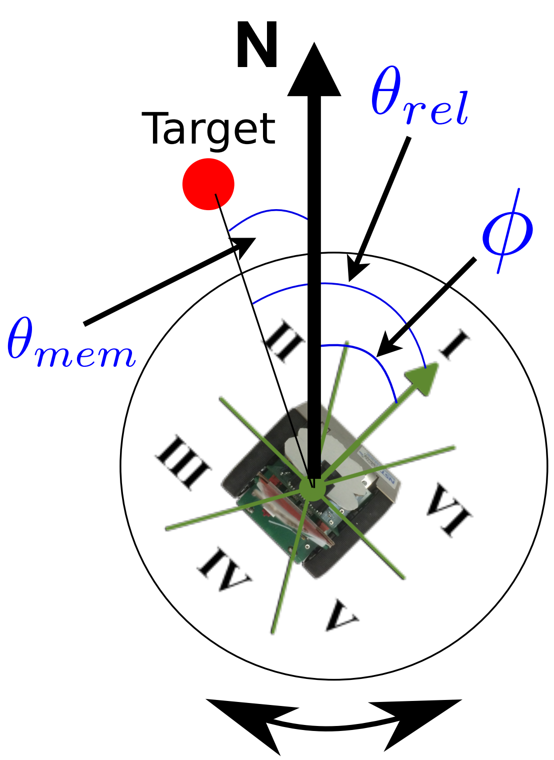

Neural ref. frame transformation:

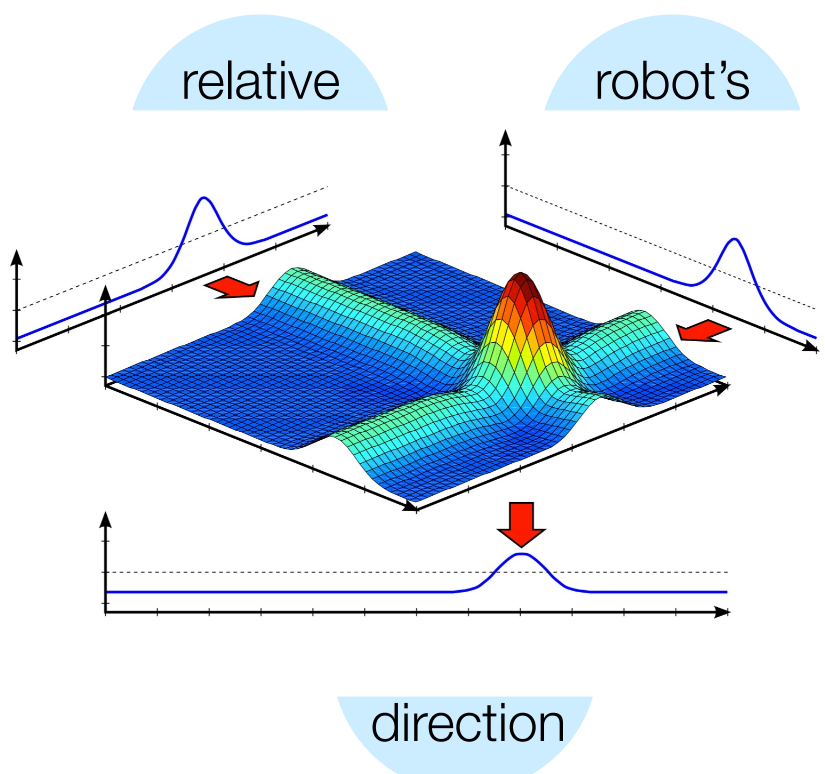

ROLLS device

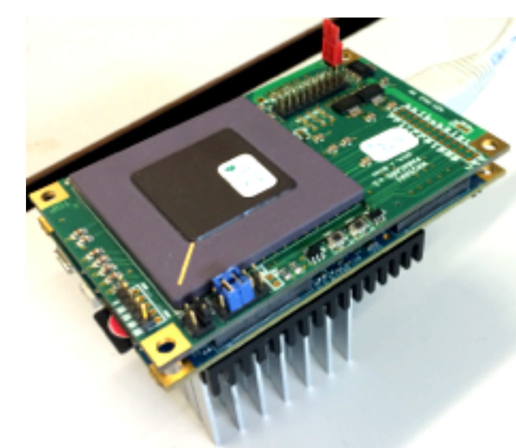

Blum, H.; Dietmüller, A.; Milde, M.; Conradt, J.; Indiveri, G. & Sandamirskaya, Y. A neuromorphic controller for a robotic vehicle equipped with a dynamic vision sensor. Robotics: Science and Systems (RSS), **2017**

What will be the bottleneck for purely neuronal control?

➡ interfaces

- sensors
  - neuromorphic SLAM as an approach to perception
- motor system
  - adaptive control

Motor control: results

Controller

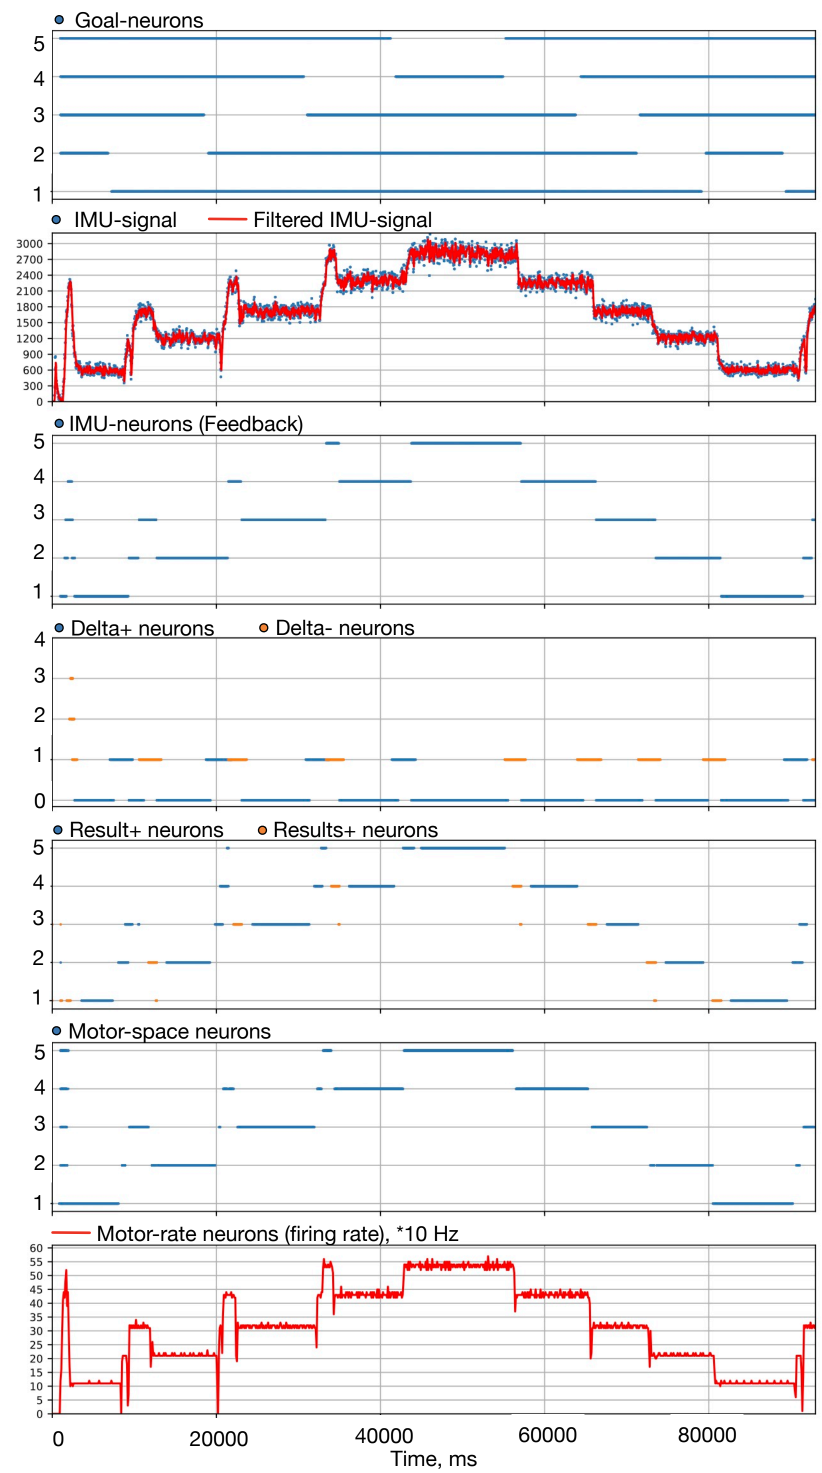

Learning the inverse mapping

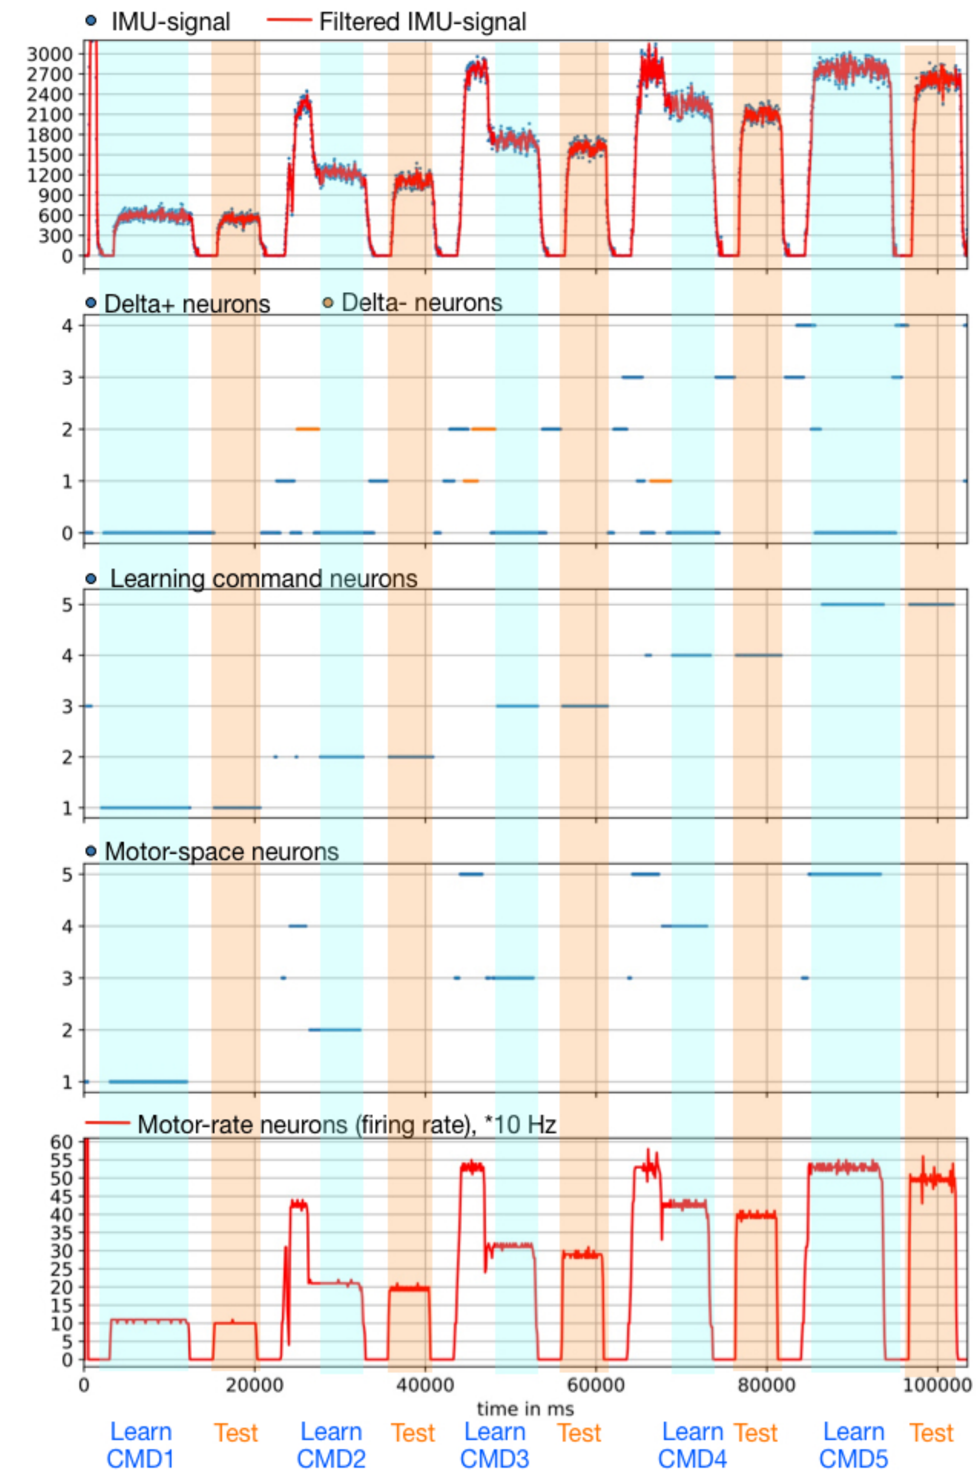

Glatz, S.; Kreiser, R.; Martel, J. N. P.; Qiao, N. & Sandamirskaya, Y. Adaptive motor control and learning in a spiking neural network, fully realised on a mixed-signal analog/digital neuromorphic processor. ICRA, arxiv, **2019**

Obstacle avoidance and target acquisition

Obstacle avoidance

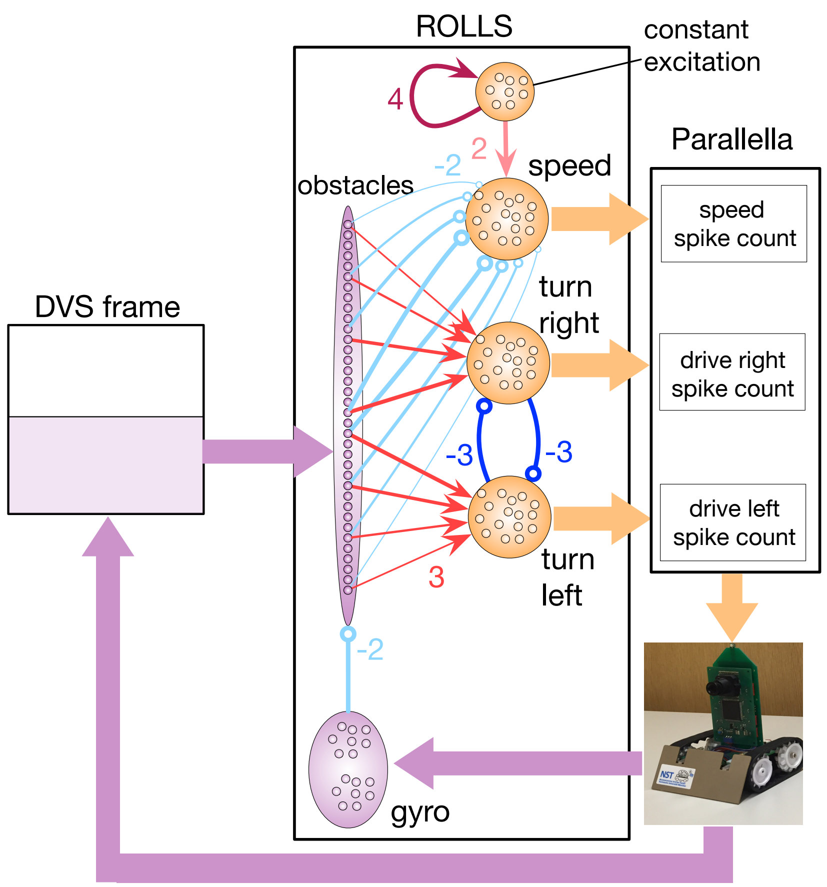

Braitenberg vehicle

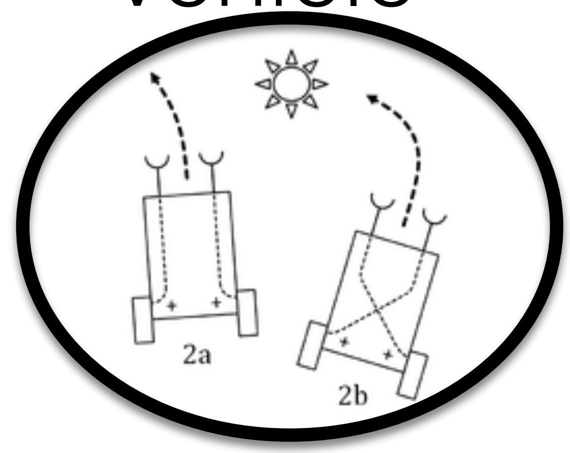

Target acquisition

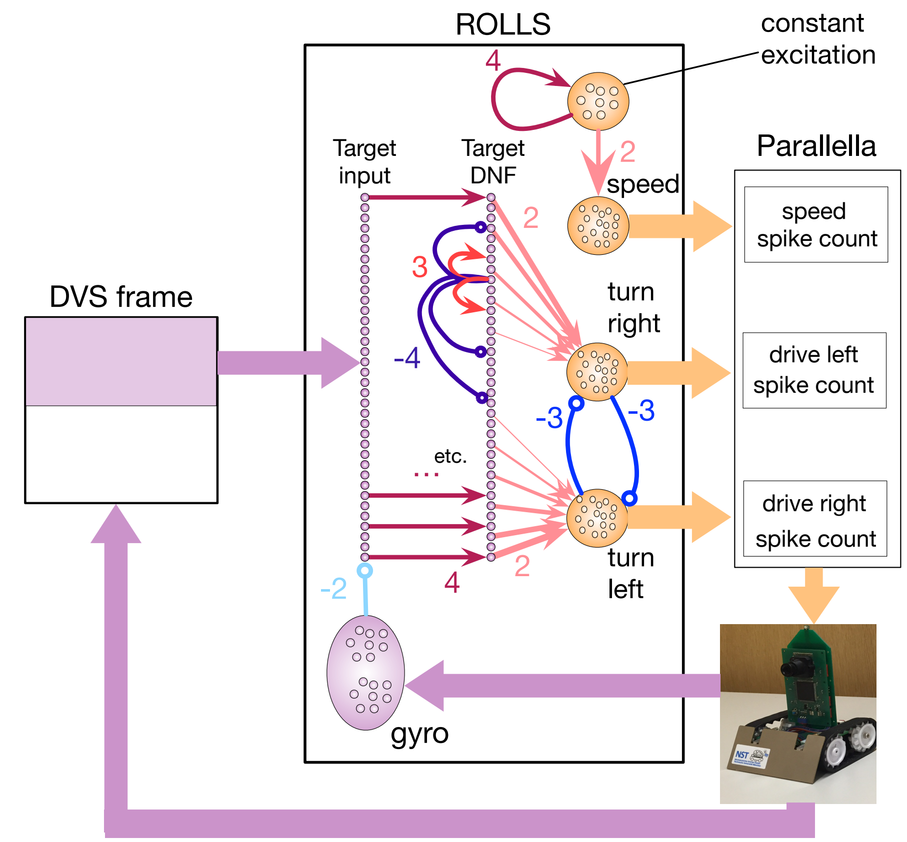

Connectivity

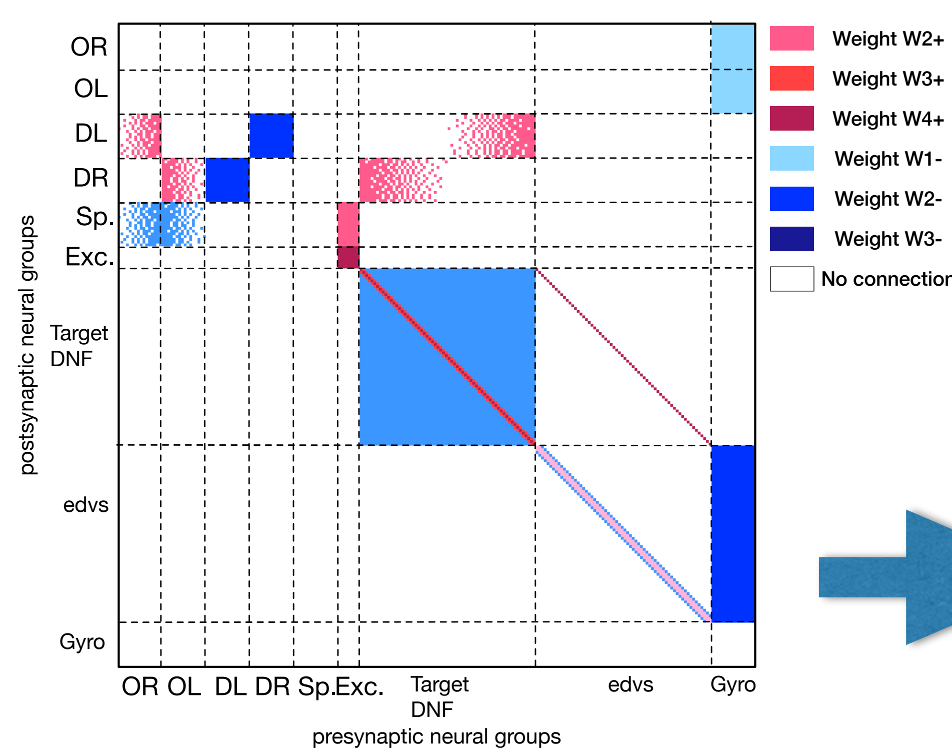

ROLLS

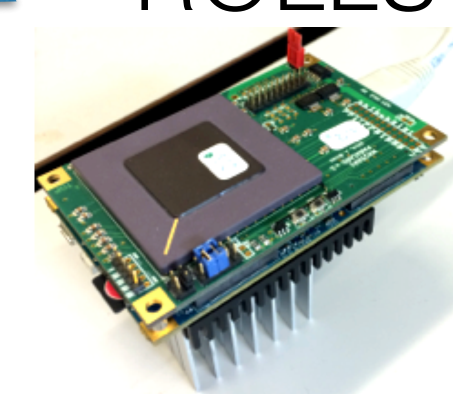

Sequence learning “program”

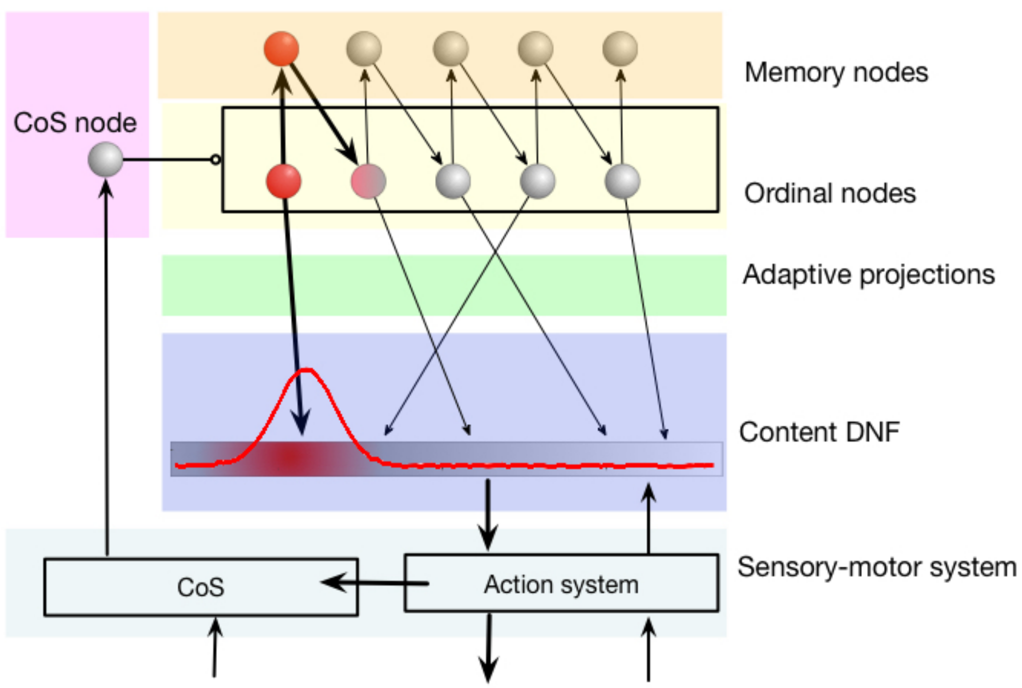

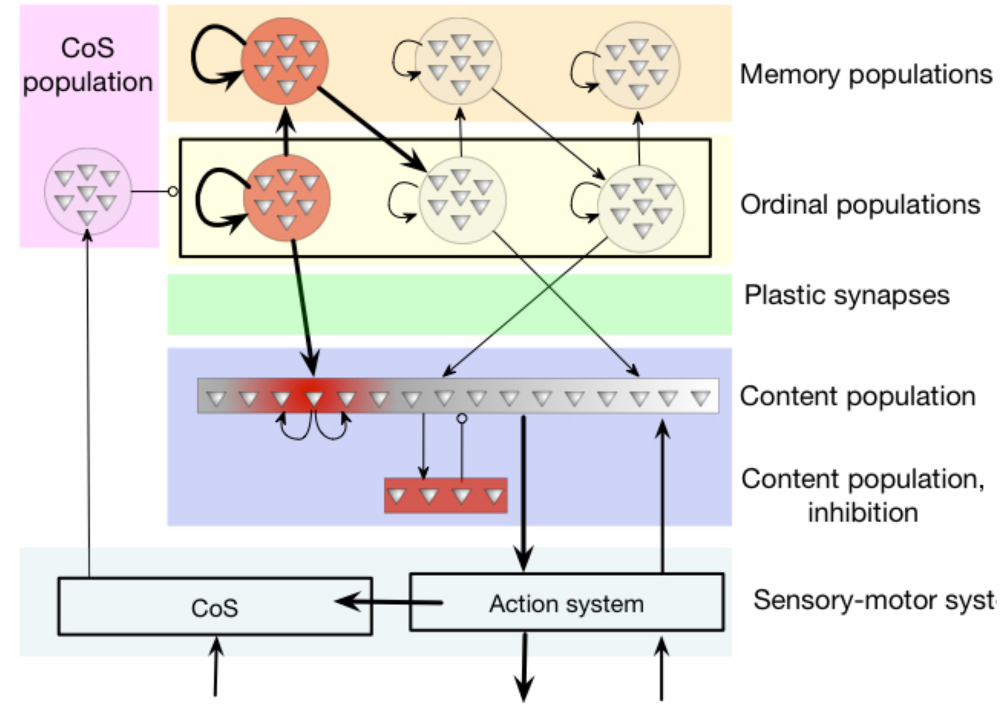

Connectivity matrix

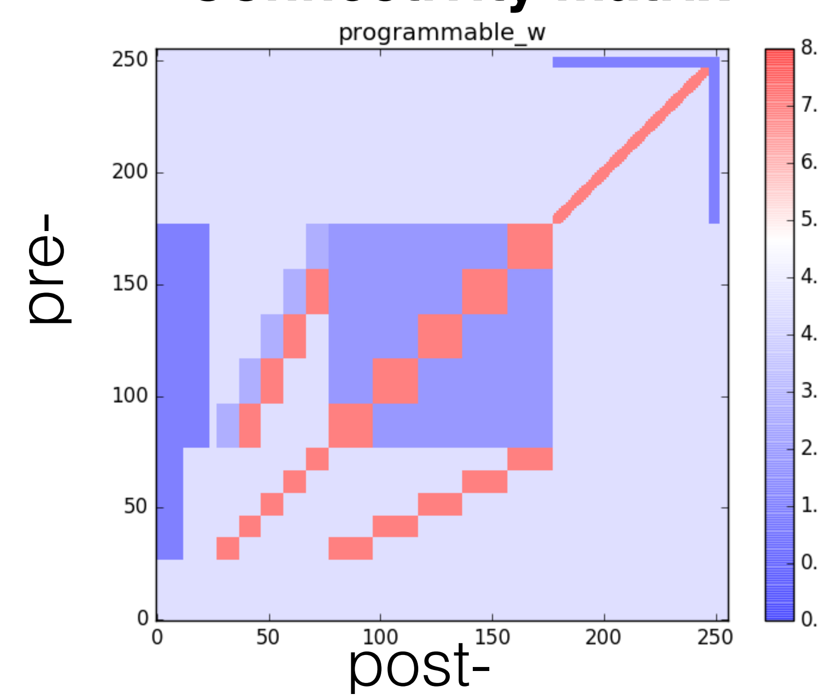

“Programm”

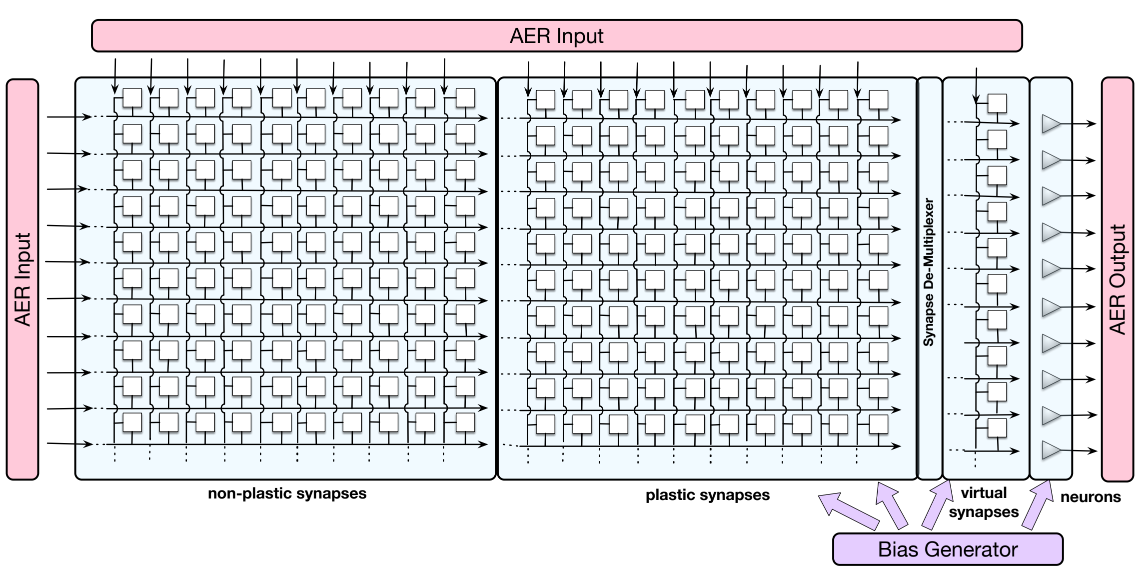

Kreiser, R.; Aathmani, D.; Qiao, N.; Indiveri, G. & Sandamirskaya, Y. Organising Sequential Memory in a Neuromorphic Device Using Dynamic Neural Fields. Frontiers in Neuromorphic Engineering, **2018**

Embodied experiment

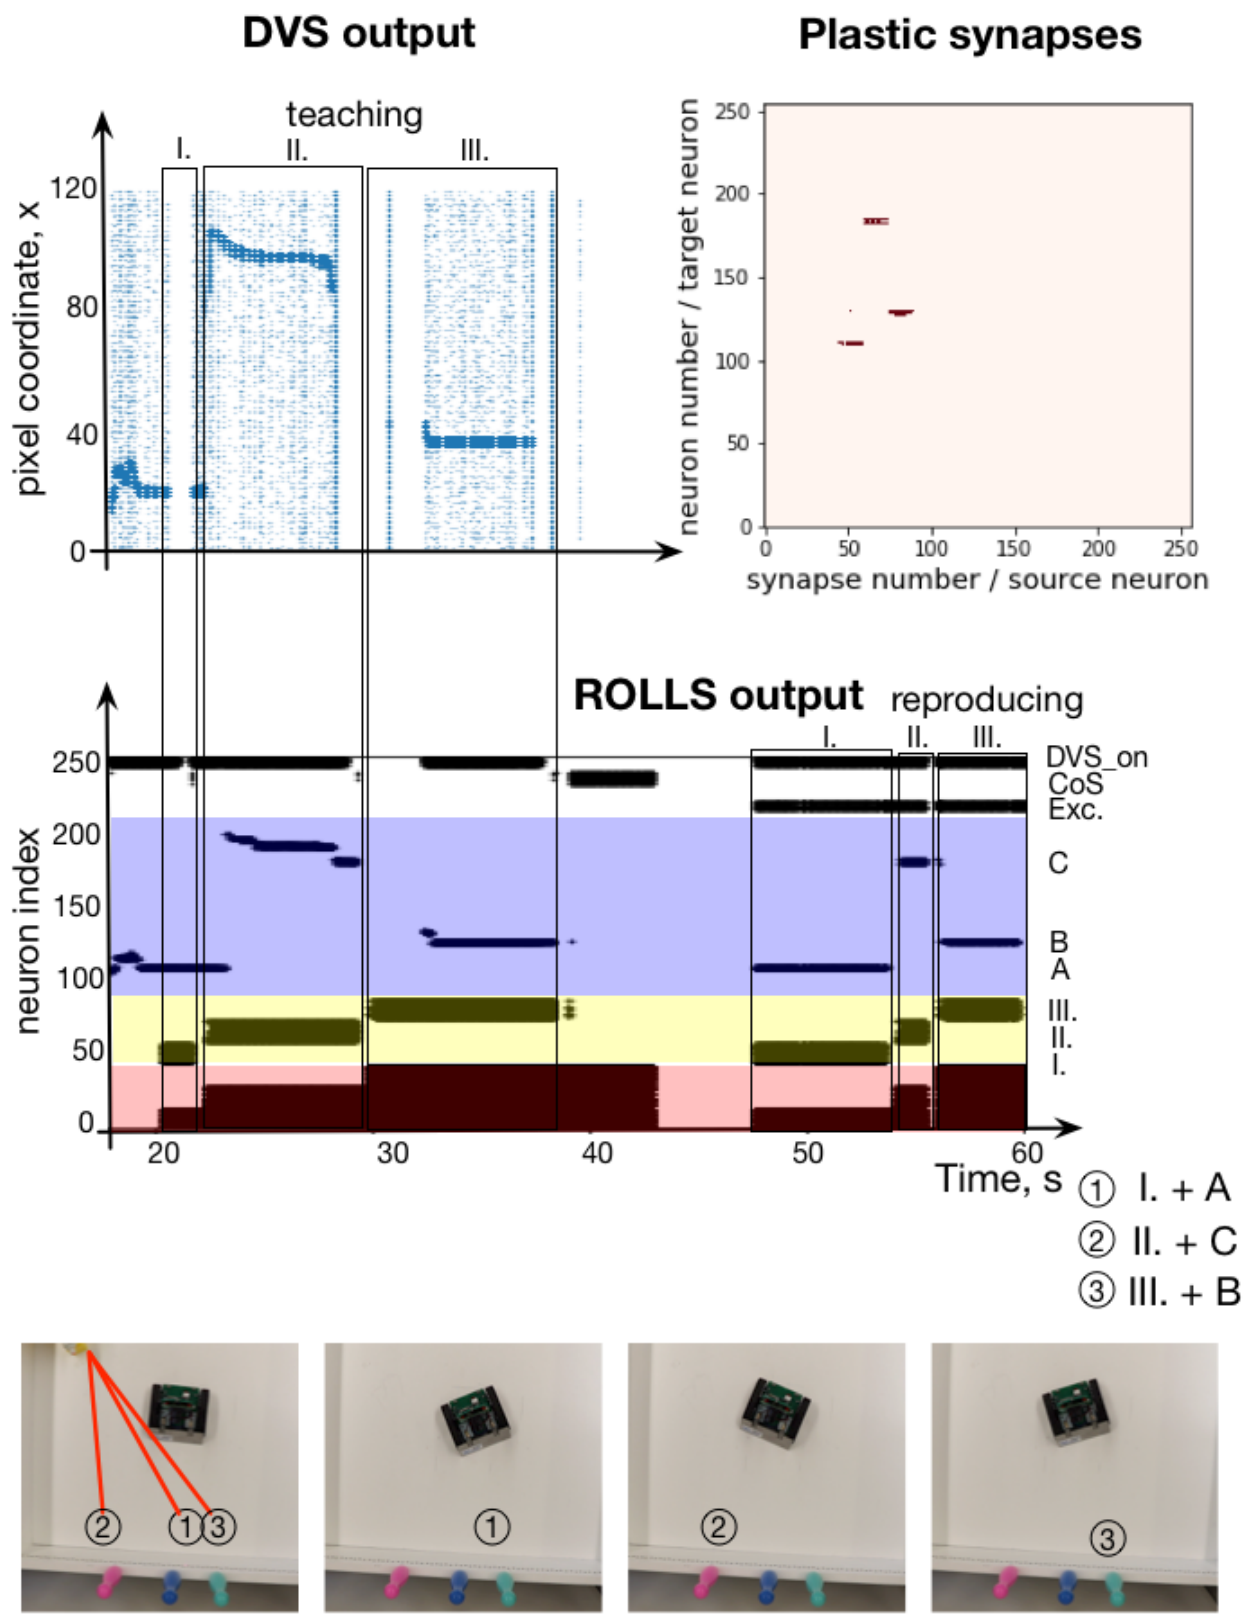

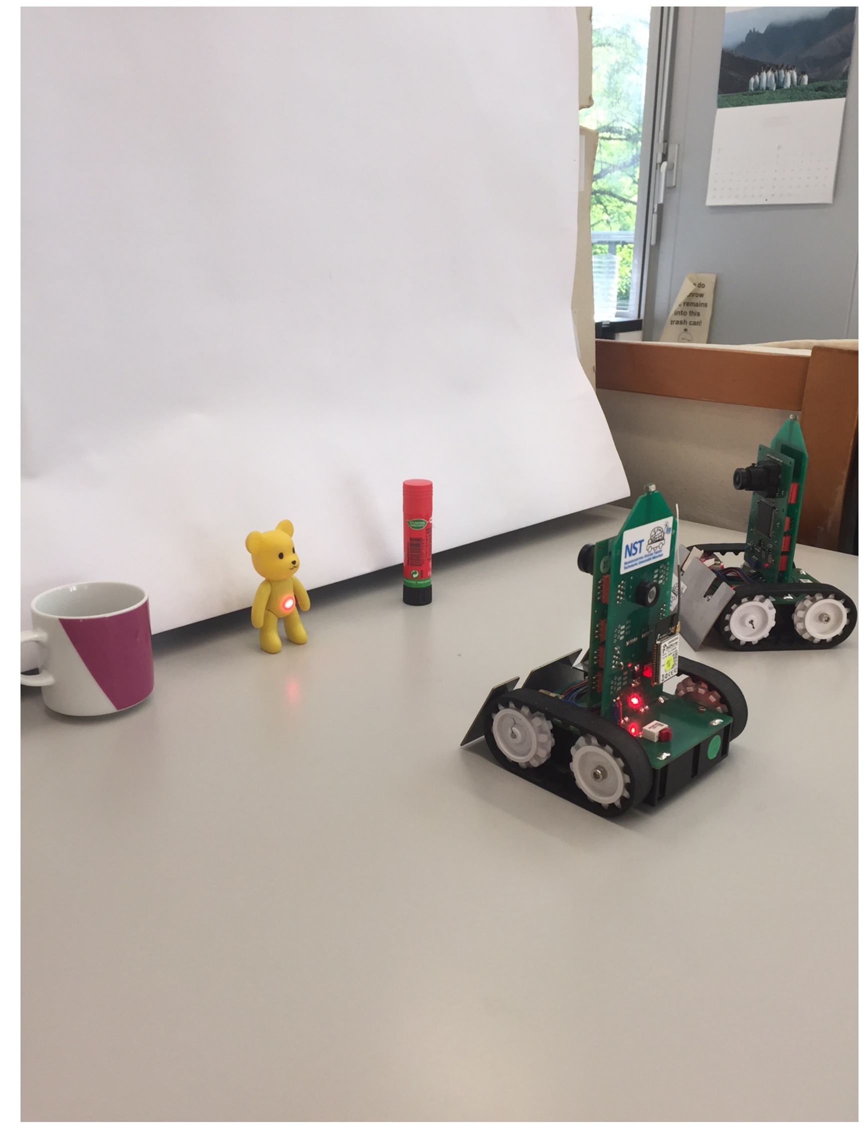

Kreiser, R.; Aathmani, D.; Qiao, N.; Indiveri, G. & Sandamirskaya, Y. Organising Sequential Memory in a Neuromorphic Device Using Dynamic Neural Fields. Frontiers in Neuromorphic Engineering, **2018**

Neuronal mechanisms: Braitenberg Vehicle

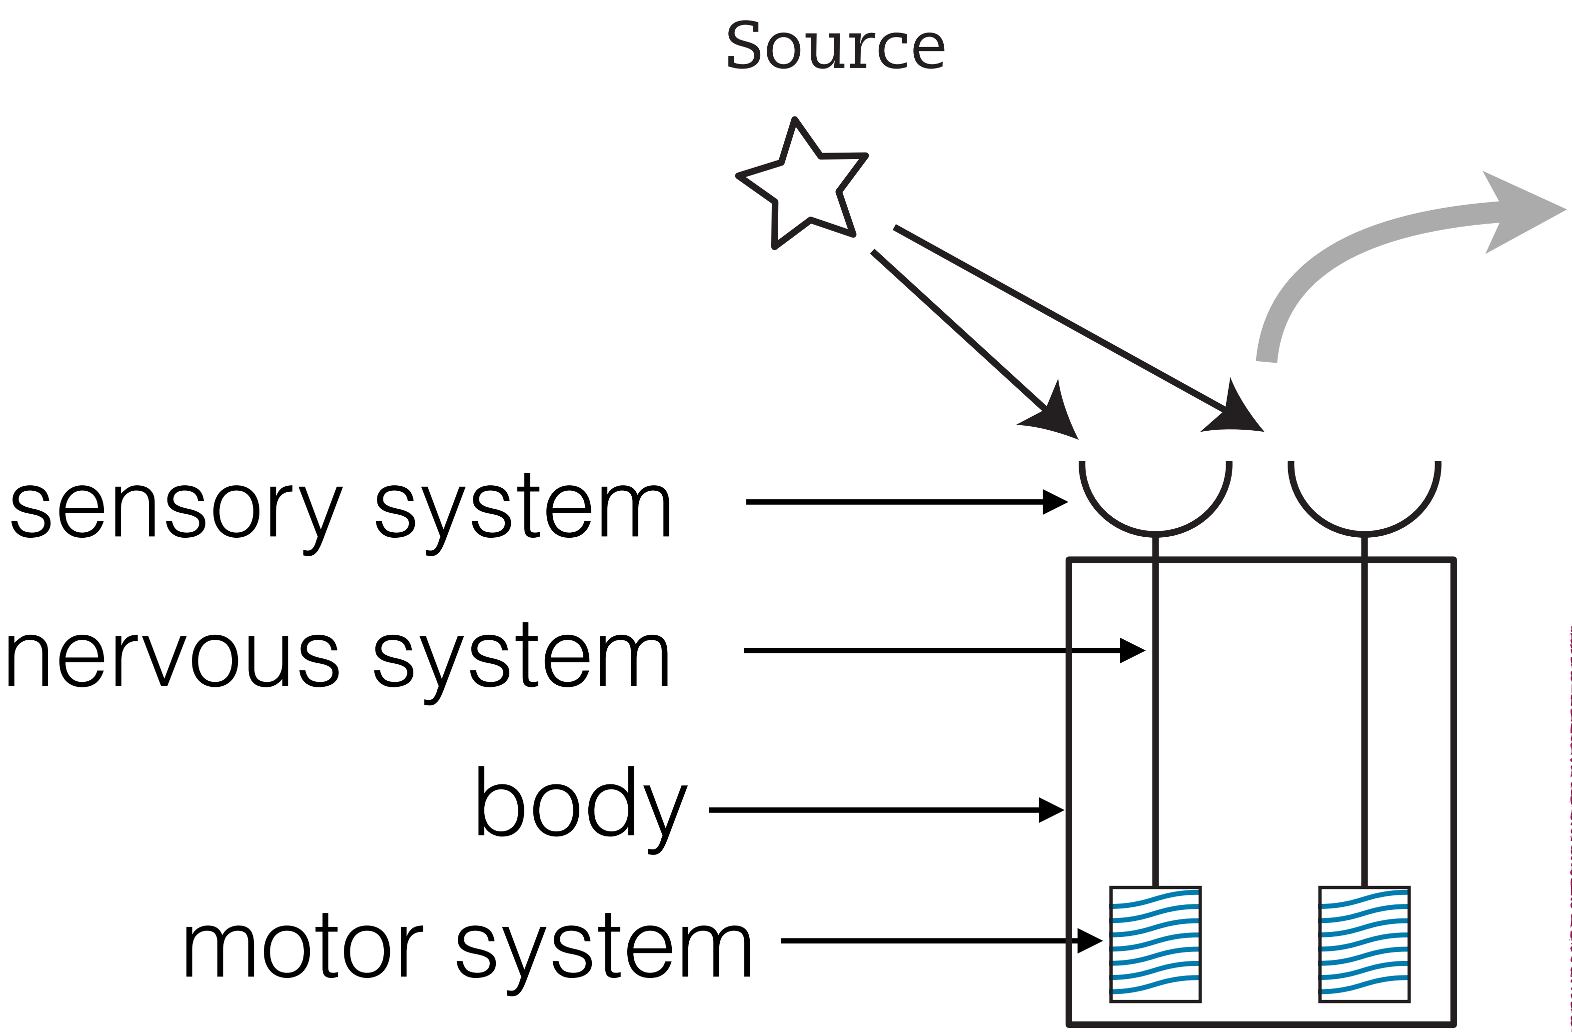

Different behaviours:

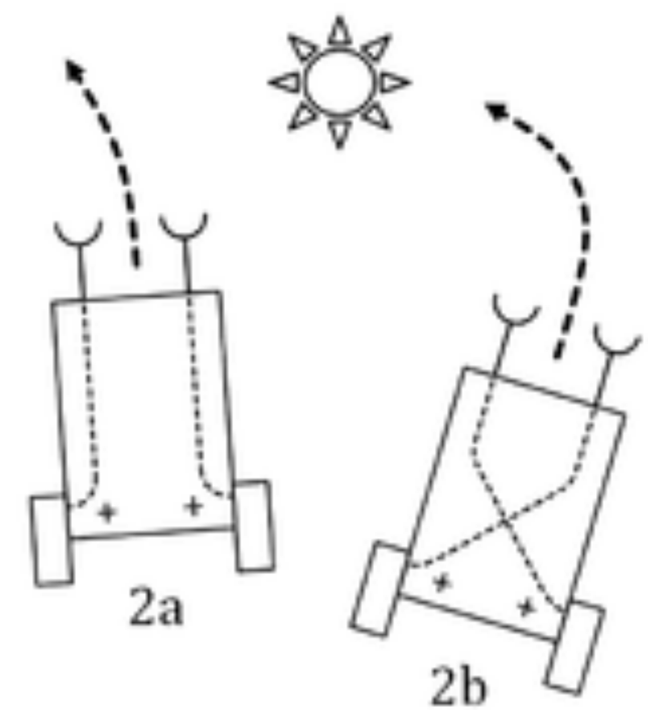

Mathematical formalisation: attractor dynamics

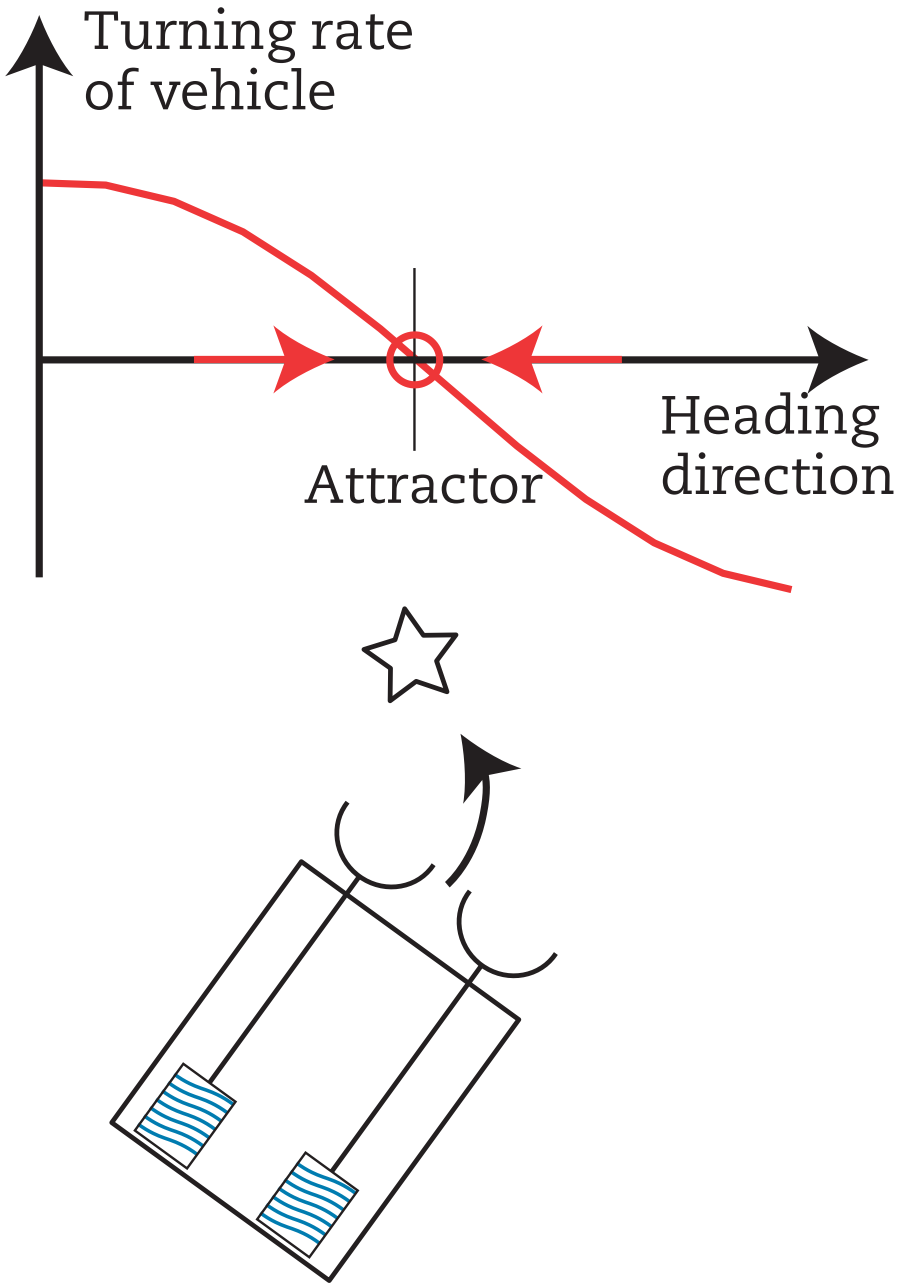

- "behavioral variable"
  - describes the behavior
- its rate of change:

$$\tau\dot{\phi}(t) = -\phi(t) + A(t)$$

  - determines its dynamics
- overt behavior corresponds to attractors
  - stability

Multiple targets

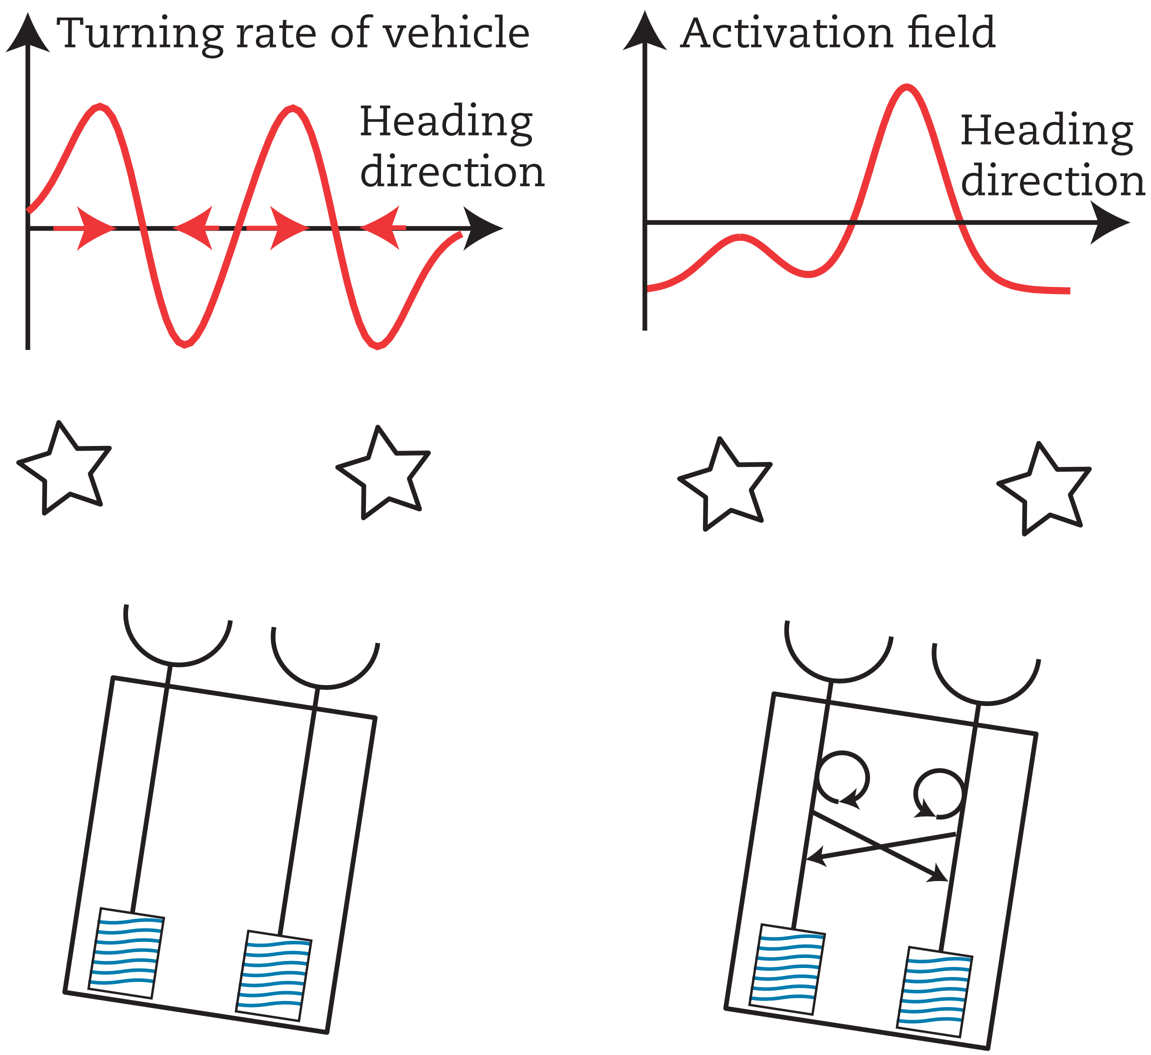

- represent "utility" of options
- stabilise decisions

$$\boxed{\dot{\phi}(t) \rightarrow \dot{u}(\phi, t)}$$

"activation" and its dynamics

→ Neural dynamics

Neuronal correlate of behavior: population activity

"Reaching" task

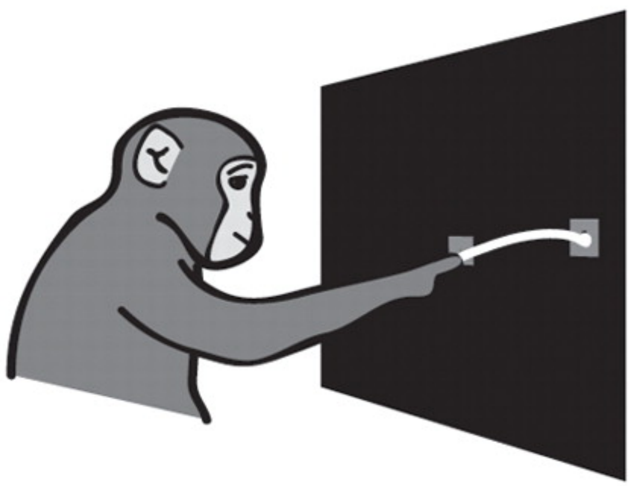

Distribution of population activation =
$\sum_{\text{neurons}}$ tuning curve * current firing rate

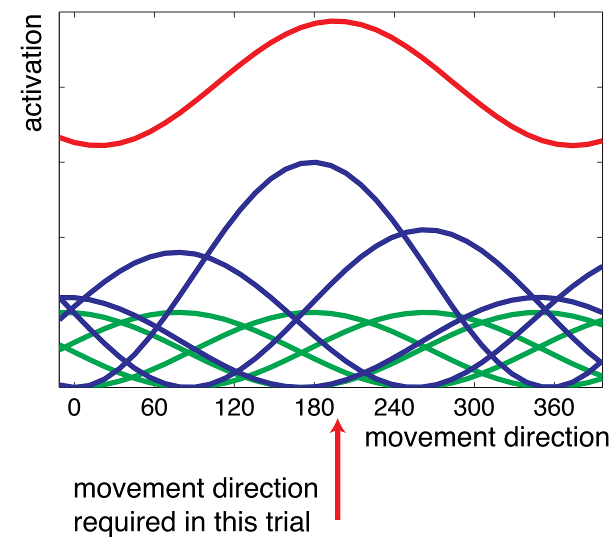

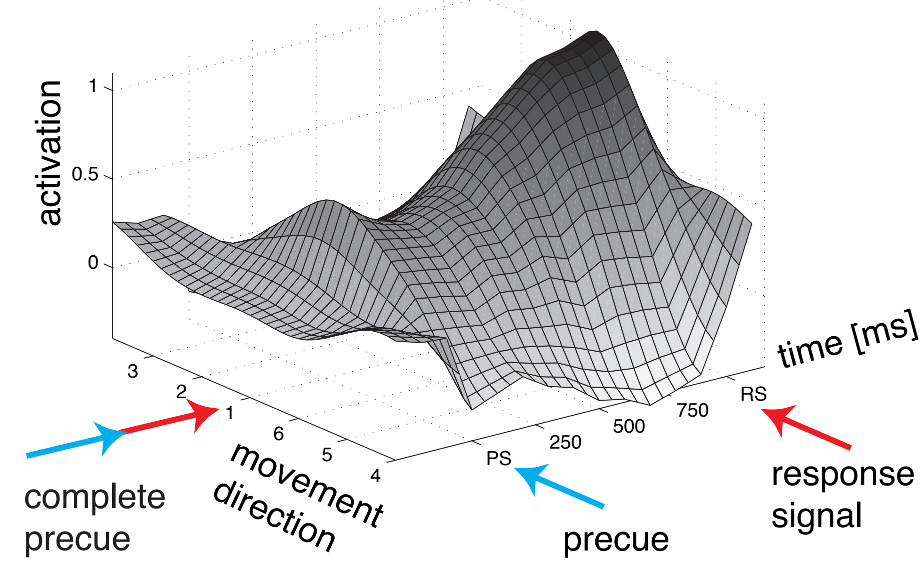

➡"Dynamic neural field" model

$$\tau \dot{u}(x,t) = -u(x,t) + h + \int f\big(u(x',t)\big)\omega(x-x')dx' + I(x,t)$$

Amari, S. **Dynamics of pattern formation in lateral-inhibition type neural fields**. Biological Cybernetics, **1977**, 27, 77-87

Wilson, H. R. & Cowan, J. D. **A mathematical theory of the functional dynamics of cortical and thalamic nervous tissue**. Kybernetik, **1973**, 13, 55-80

Gerstner, Grossberg, Ermentrout, Coombes, Schöner&Spencer, Erlhagen…

Dynamic Neural Fields explain behavior

- Habit formation

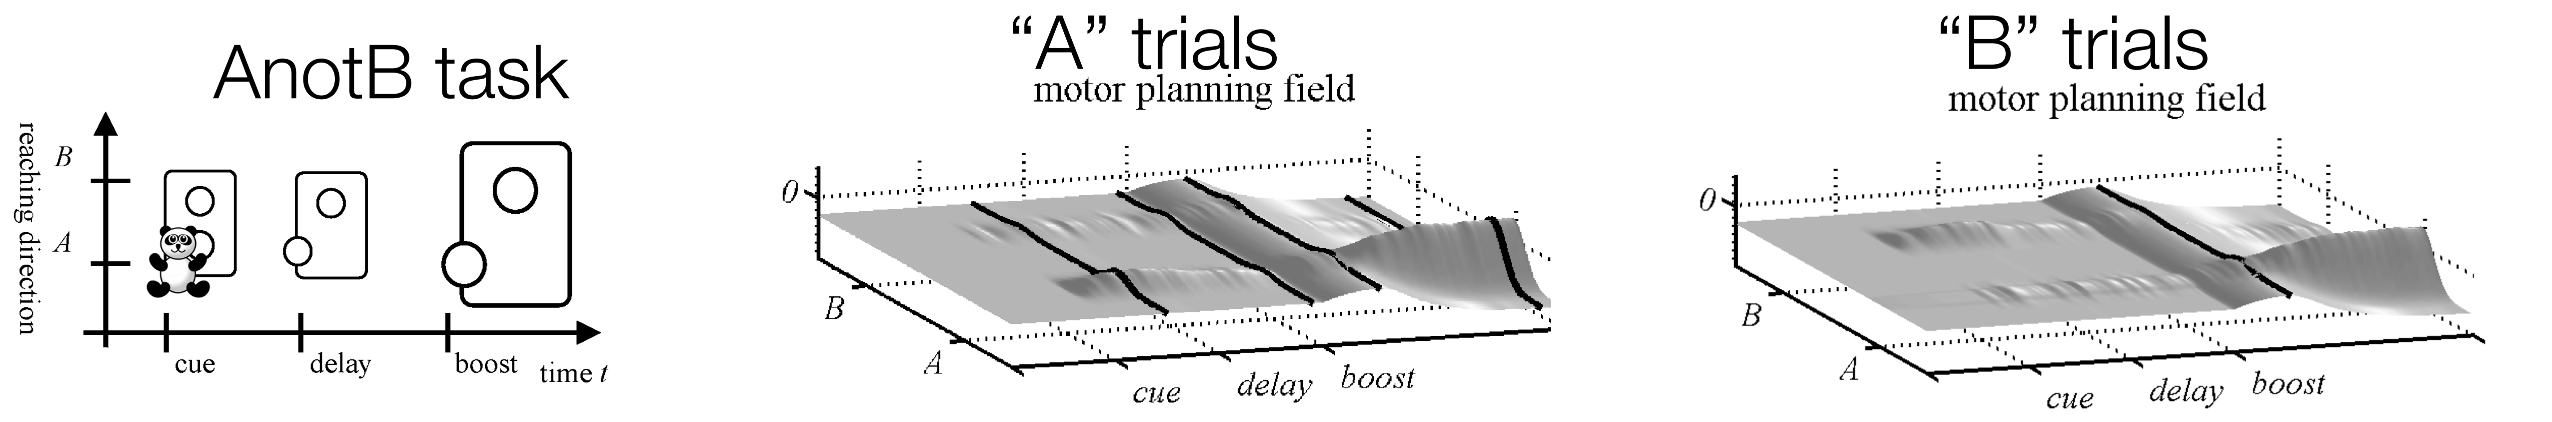

Dineva and Schöner, "How infants' reaches reveal principles of sensorimotor decision making", Connection Science, 2018

- Visual working memory

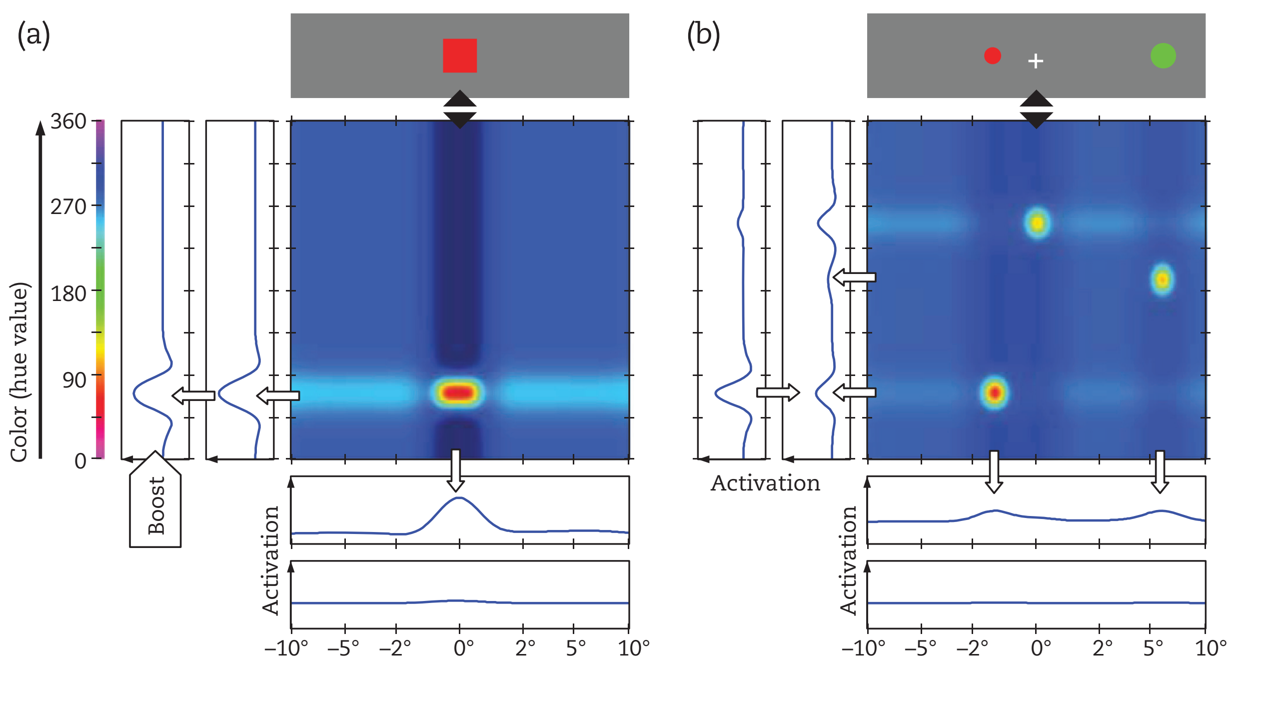

- Spatial working memory

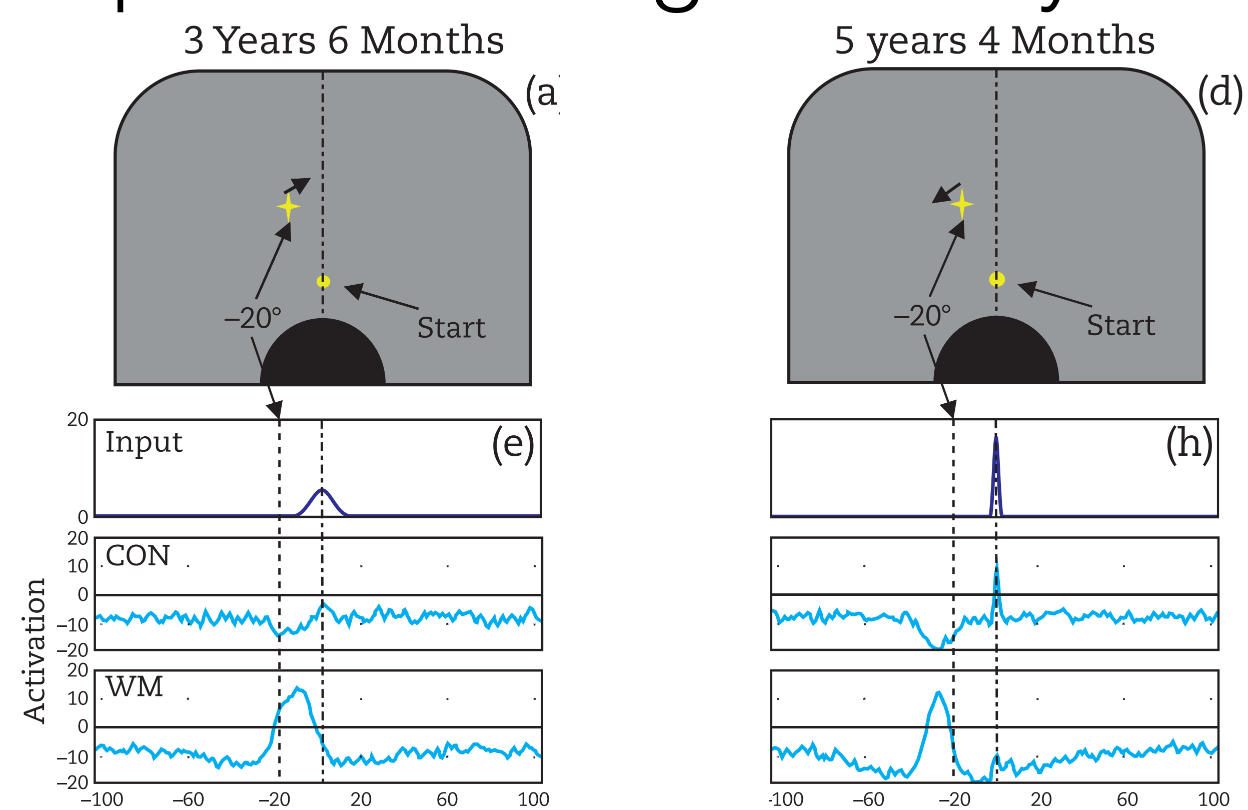

Schöner, Spencer, and the DFT group. "Dynamical Thinking: Primer to Dynamic Field Theory, Oxford Press, 2015

“Implementation issue”

Neuronal dynamics

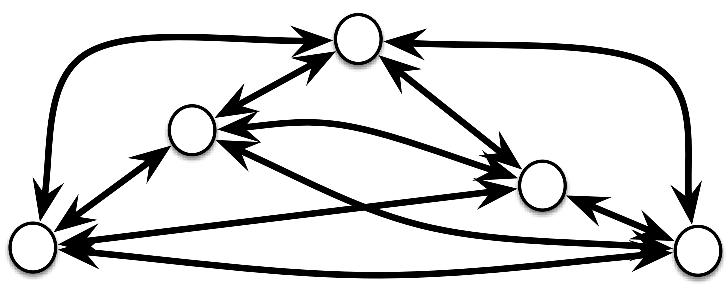

$$\tau \dot{u}(x,t) = -u(x,t) + h + \int f(u(x',t))\omega(x - x')dx' + I(x,t)$$

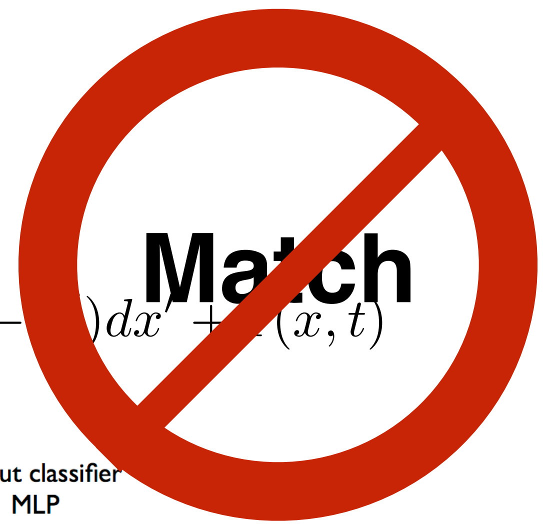

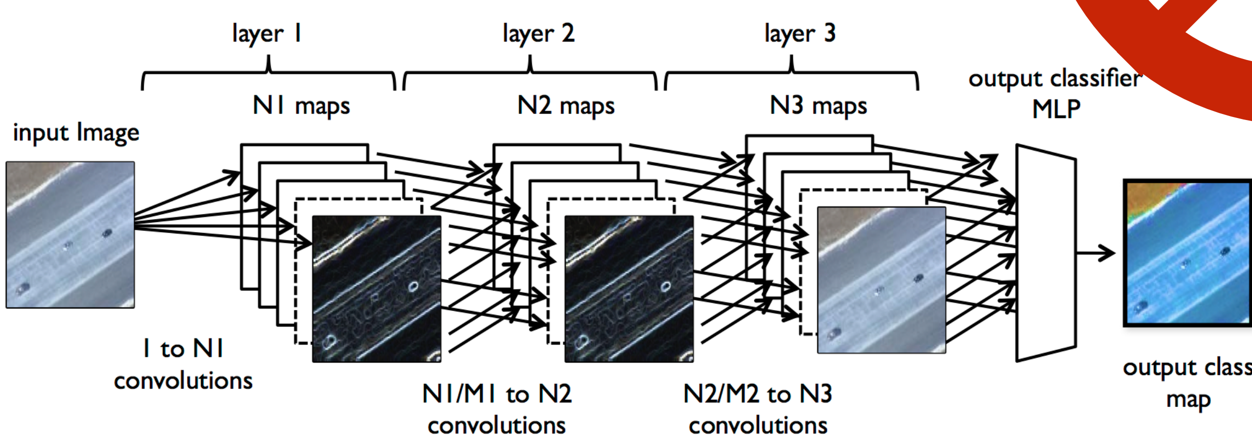

“Von Neumann” computer

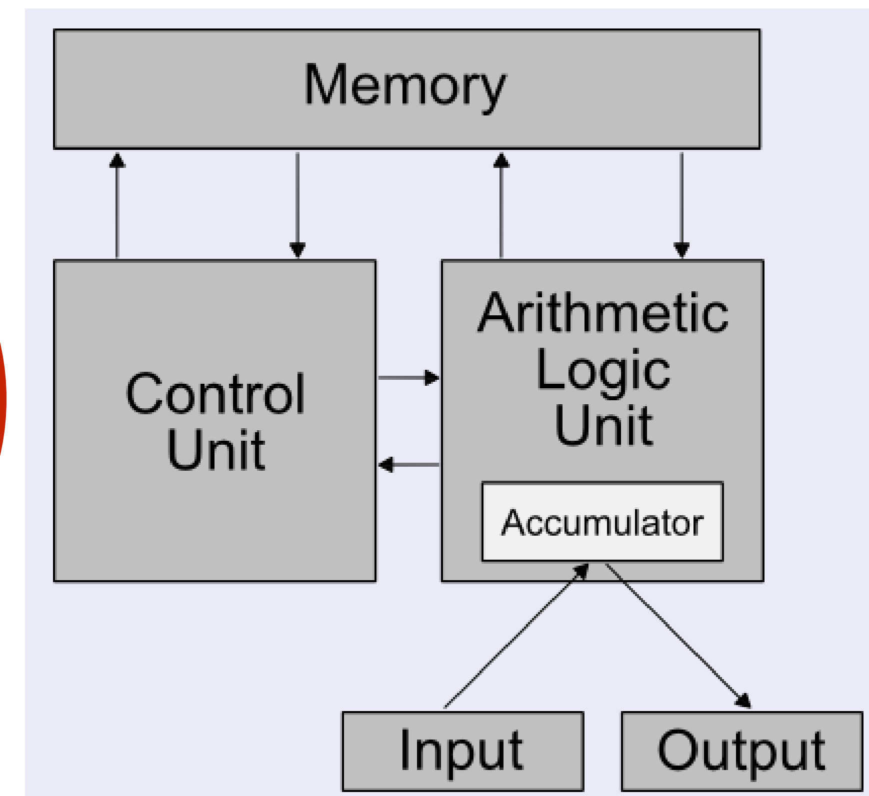

- analogue values
- parallel processing
- memory and computation interlinked

- digital representations
- sequential processing
- separate memory unit

Neuromorphic Hardware

➡ Brain-inspired computing or sensing devices that emulate activity of biological neurons and synapses

Analog

"BrainDrop" (Stanford)

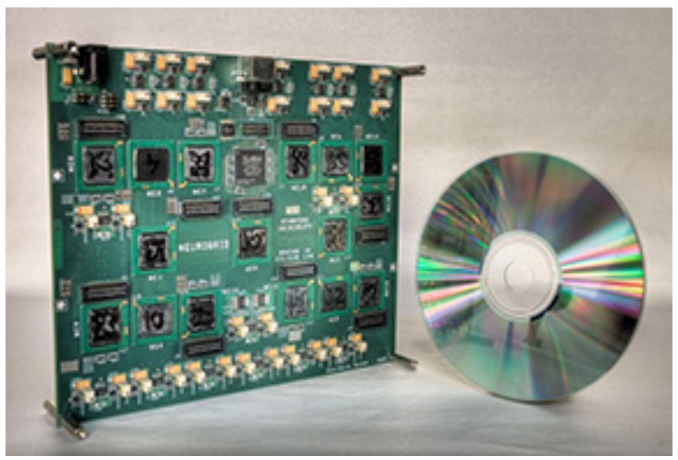

DYNAP (Zurich)

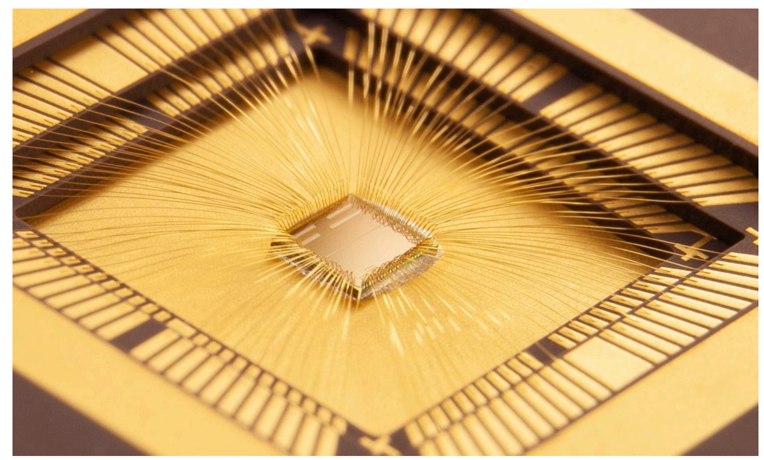

BrainScaleS (Heidelberg)

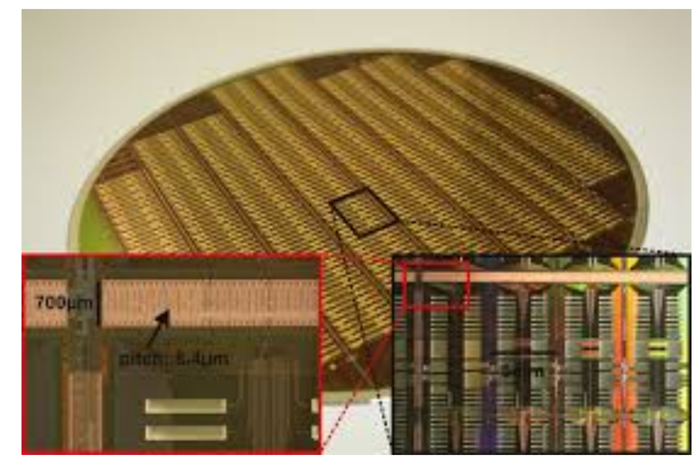

Digital

"TrueNorth" (IBM)

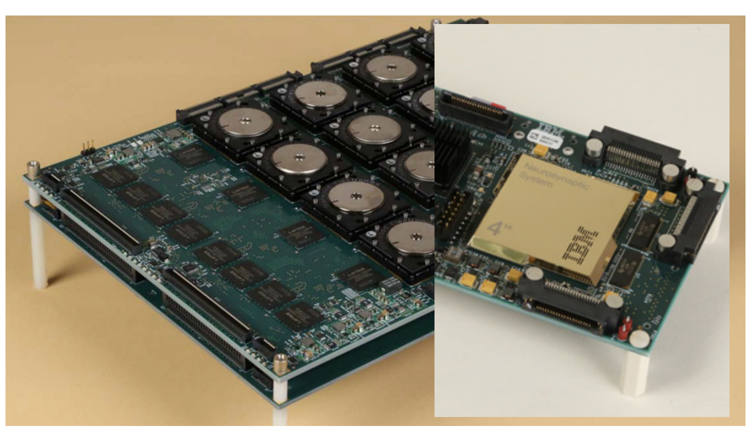

Loihi (Intel)

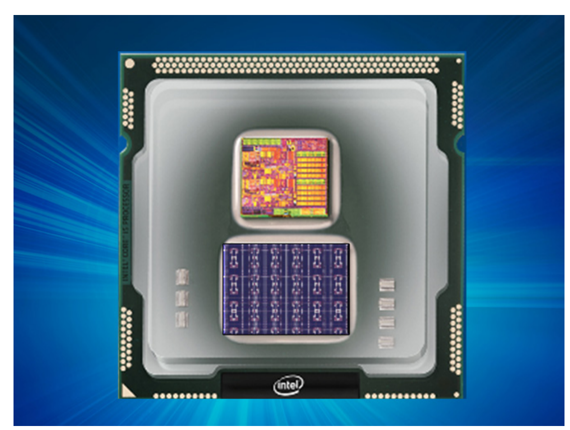

SpiNNaker (Manchester)

NEUROTECH

NEUROMORPHIC COMPUTING TECHNOLOGY LEADING TO AI REVOLUTION

Create and promote neuromorphic community in Europe: www.neurotechai.eu