Online learning in SNNs with e-prop and Neuromorphic Hardware

Adam Perrett
Sara Summerton
Andrew Gait
Oliver Rhodes

Overview

- Motivation
- Memory in neural networks
- E-prop
- SpiNNaker
- Wave form matching
- Temporal credit assignment
- Conclusions

Overview

- Motivation
- Memory in neural networks
- E-prop
- SpiNNaker
- Wave form matching
- Temporal credit assignment
- Conclusions

Motivation

- Online learning
  - No break between acting and learning
  - Local tailoring of a system
- Neuromorphic
  - Low power
  - Low latency
  - Event-based computing
  - Biologically inspired
- SpiNNaker
  - Real-time computing
  - Adaptable digital platform
  - Scale

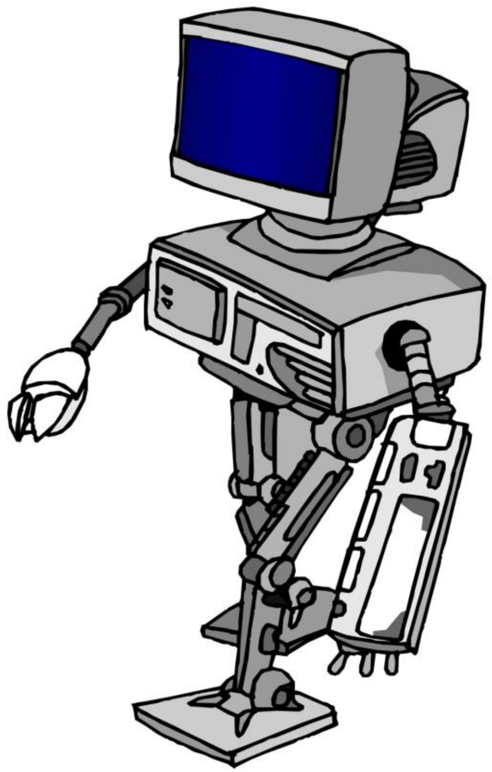

Overview

- Motivation
- Memory in neural networks
- E-prop
- SpiNNaker
- Wave form matching
- Temporal credit assignment
- Conclusions

Recurrent models with working memory

- Common practice in machine learning is to use Long Short Term Memory (LSTM) units
- Very successful in temporal tasks such as language processing and video prediction

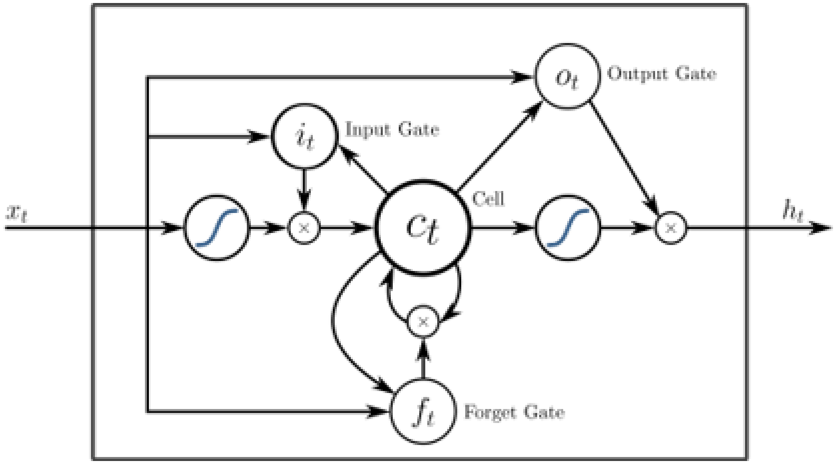

Back propagation through time (BPTT)

- Error is propagated backwards through the network and time
- Unravelling as another layer with shared weights for each time step
- All states must be recorded for all time steps you wish to learn over
- Not suited to neuromorphics

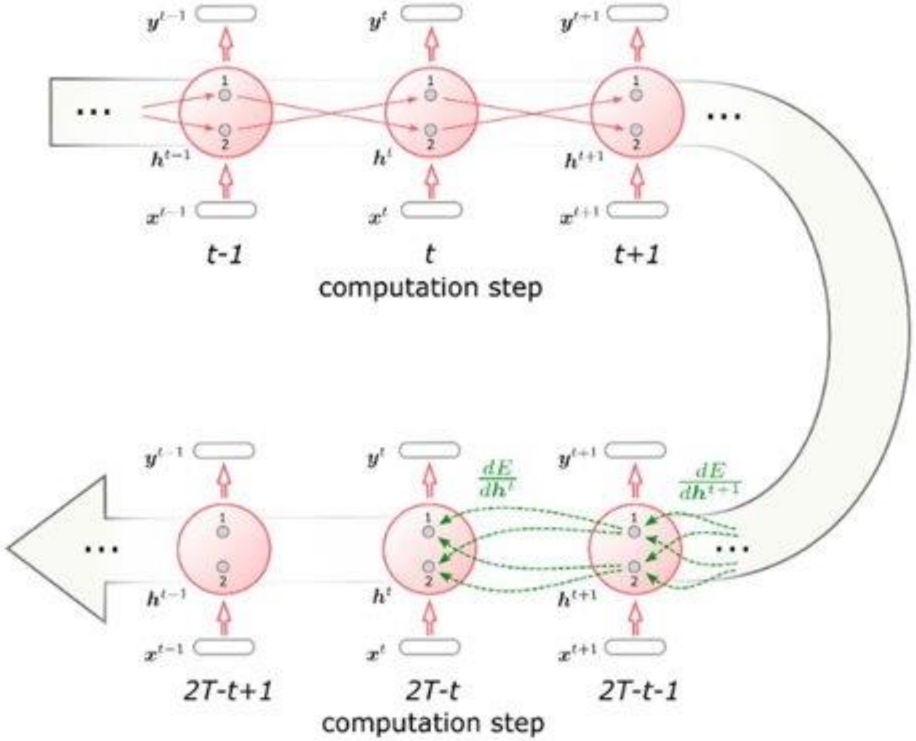

Bellec, G., Scherr, F., Subramoney, A., Hajek, E., Salaj, D., Legenstein, R., & Maass, W. (2020). A solution to the learning dilemma for recurrent networks of spiking neurons. *Nature communications*, *11*(1), 1-15.

Working memory in spiking neurons

- The firing threshold $A_j(t)$ of a leaky integrate-and-fire (LIF) neuron j contains a time-varying component $a_j(t)$
- It's temporarily increased by each spike $z_j(t)$ of the neuron:

$$A_j^t = v_{\mathrm{th}} + \beta a_j^t$$

$$a_j^{t+1} = \rho a_j^t + z_j^t$$

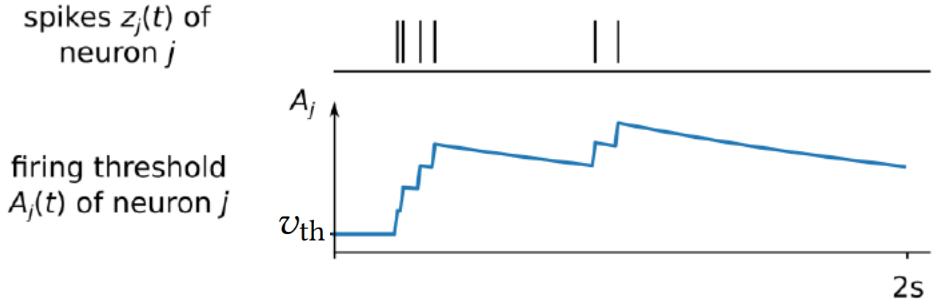

- Performance comparable with LSTMs when trained with BPTT!

Input image

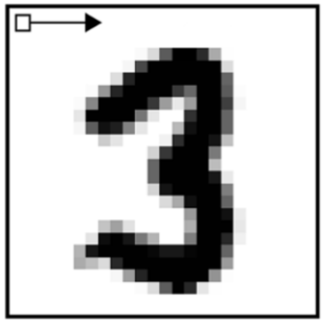

Performance comparison

for step sizes of 1ms (light color) and 2ms (saturated)

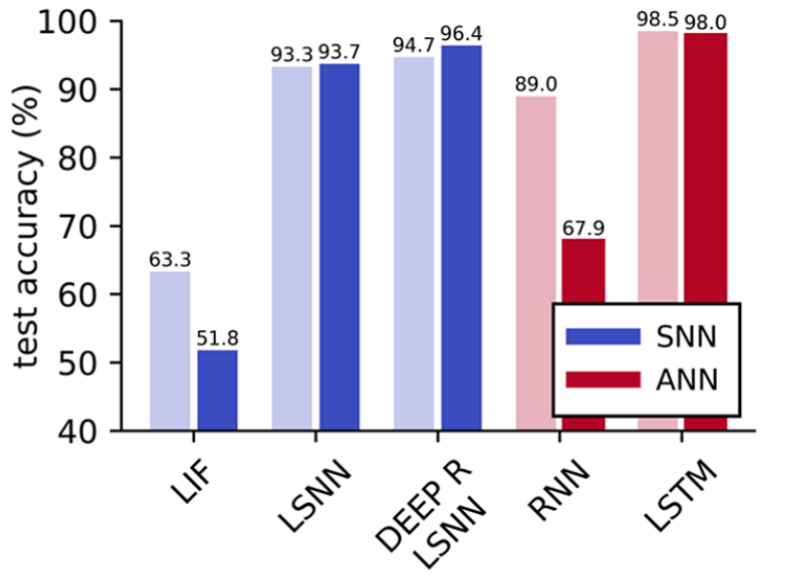

Bellec, G., Salaj, D., Subramoney, A., Legenstein, R., & Maass, W. (2018). Long short-term memory and learning-to-learn in networks of spiking neurons. *Advances in neural information processing systems, 31*.

Spiking neural network learning (LIF)

- Inherently recurrent as future states are influenced by past ones
- Non differentiable activation function
- Requires pseudo-derivative

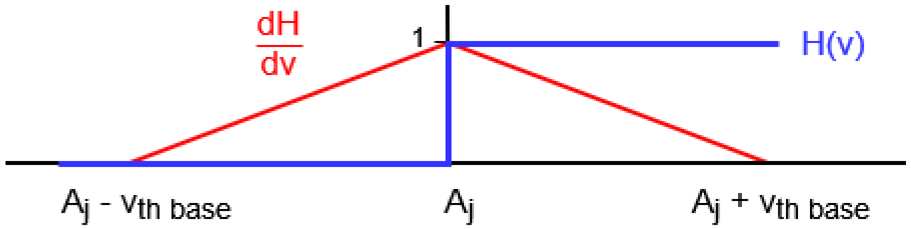

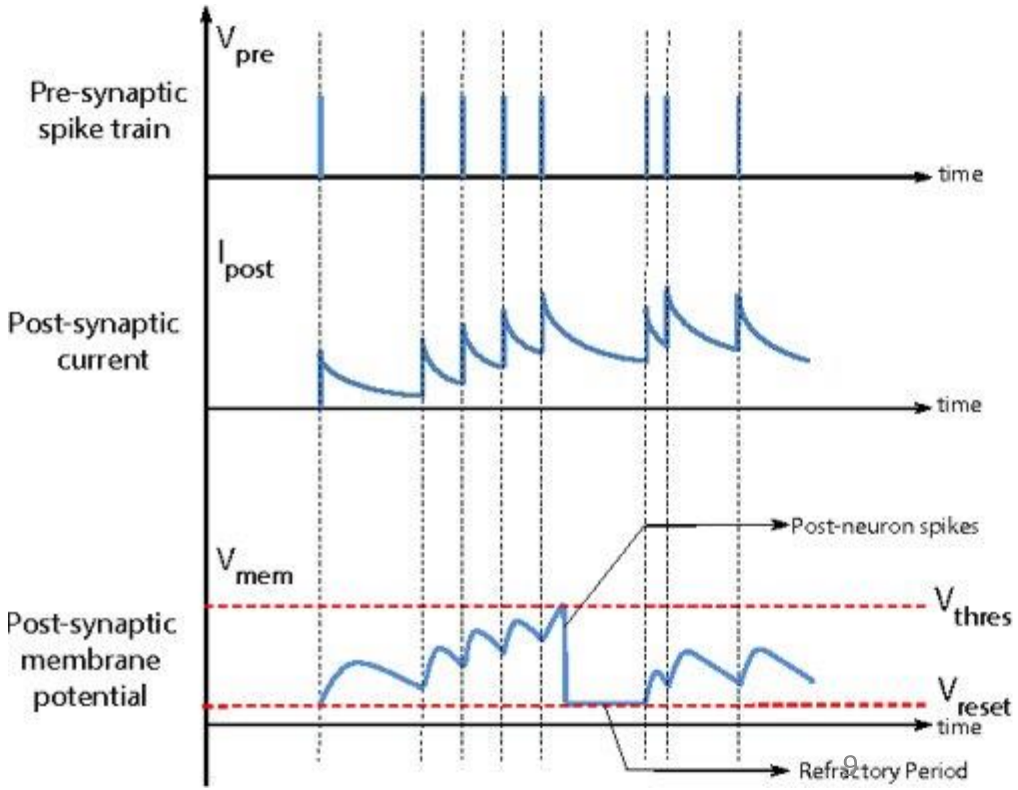

Overview

- Motivation
- Memory in neural networks
- E-prop
- SpiNNaker
- Wave form matching
- Temporal credit assignment
- Conclusions

Eligibility propagation (e-prop)

- Eligibility captures the history of behaviour
- It is lowpass filtered through time
- The error is broadcast along random feedback weights[1] creating the learning signal
- Eligibility and the learning signal create the weight updates

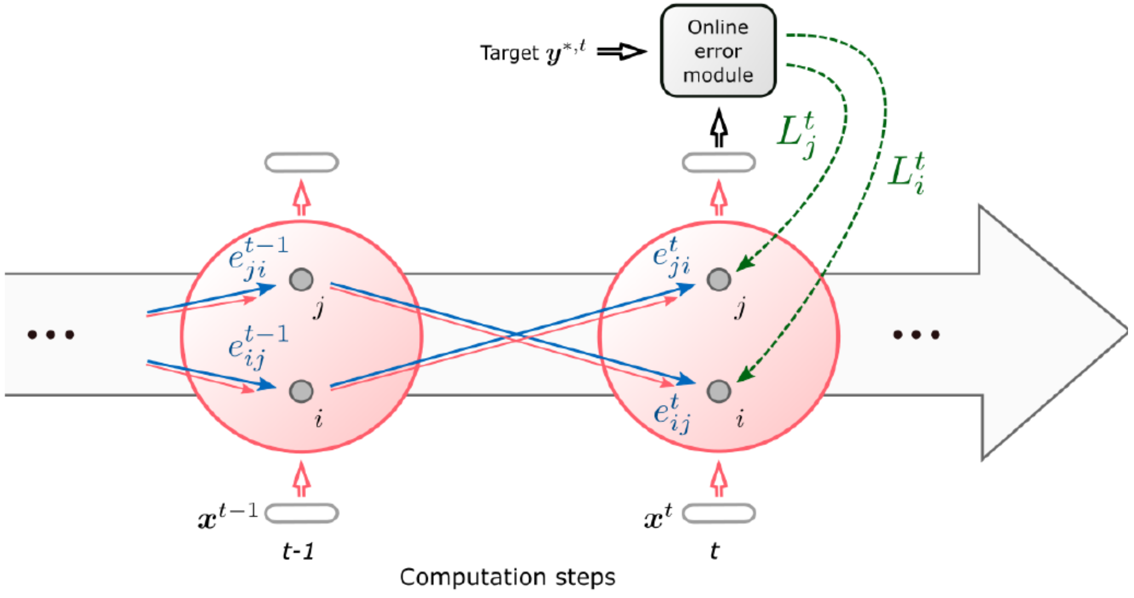

Bellec, G., Scherr, F., Subramoney, A., Hajek, E., Salaj, D., Legenstein, R., & Maass, W. (2020). A solution to the learning dilemma for recurrent networks of spiking neurons. *Nature communications*, *11*(1), 1-15.

[1] Lillicrap, T.P., Cownden, D., Tweed, D.B. and Akerman, C.J., 2014. Random feedback weights support learning in deep neural networks. *arXiv preprint arXiv:1411.0247*.

BPTT

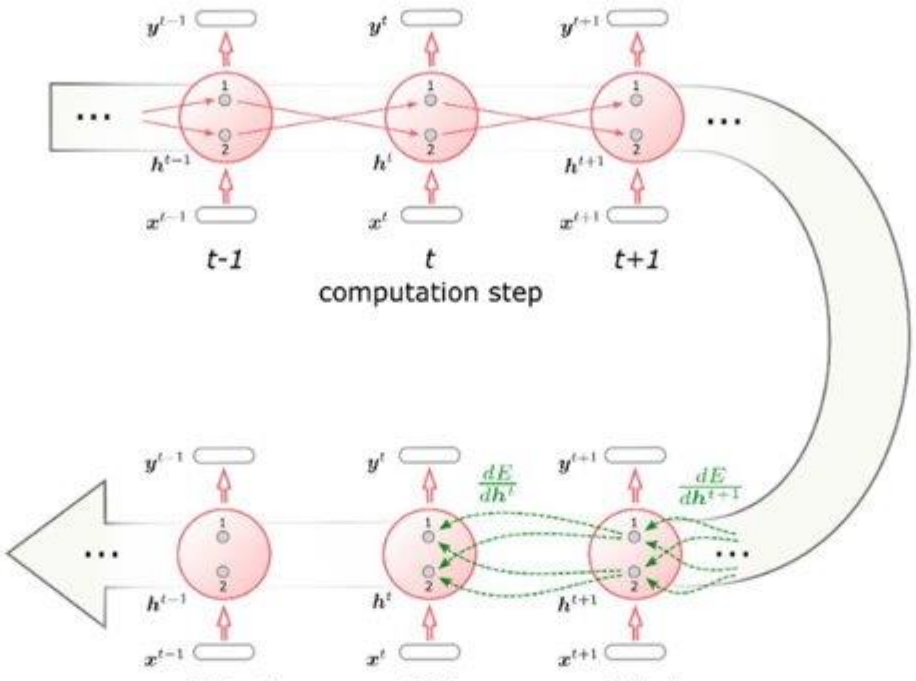

E-prop

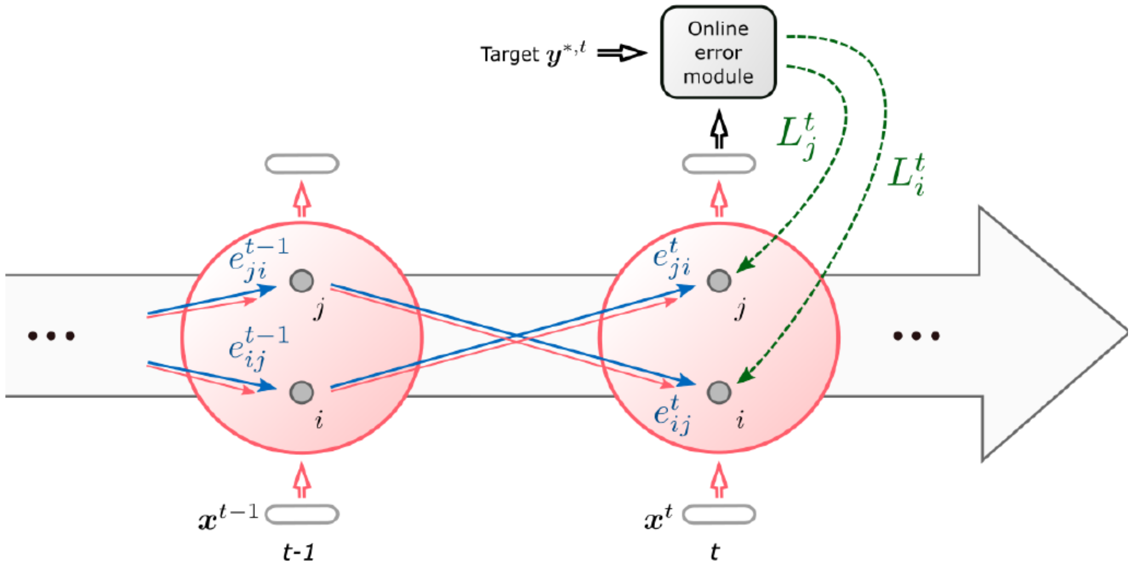

Bellec, G., Scherr, F., Subramoney, A., Hajek, E., Salaj, D., Legenstein, R., & Maass, W. (2020). A solution to the learning dilemma for recurrent networks of spiking neurons. *Nature communications*, *11*(1), 1-15.

Neuron model - (A)LIF

- Membrane voltage leaks
- Integrates spikes from the inputs and recurrent connections
- Reset by subtraction after spiking
- Produce a spike if above threshold
- The adaptive threshold increases following a spike then decays back down

$$v_j^{t+1} = \alpha v_j^t + \sum_{i \neq j} W_{ji}^{\mathrm{rec}} z_i^t + \sum_i W_{ji}^{\mathrm{in}} x_i^{t+1} - z_j^t A_j^t$$

$$z_j^t = \begin{cases} 1, & \text{if } (v_j^t - A_j^t) > 0 \\ 0, & \text{otherwise} \end{cases}$$

$$A_j^t = v_{\mathrm{th}} + \beta a_j^t$$

$$a_j^{t+1} = \rho a_j^t + z_j^t$$

Weight updates

- Eligibility of a neuron is the product of the pseudo derivative and the incoming spike trace

$$\psi_j^t = \frac{1}{v_{\text{th base}}}\gamma_{\text{pd}} \max\left(0, 1 - \left|\frac{v_j^t - A_j^t}{v_{\text{th base}}}\right|\right)$$

$$e_{ji}^t = \psi_j^t\left(\bar{z}_i^{t-1} - \beta\epsilon_{ji,a}^t\right)$$

- Weight updates are the product of the eligibility trace and the learning signal

$$\Delta W_{ji} = -\eta \sum_t \underbrace{\left(\sum_k B_{jk}\overbrace{\left(y_k^t - y_k^{*,t}\right)}^{=E^t}\right)}_{=L_j^t} \bar{e}_{ji}$$

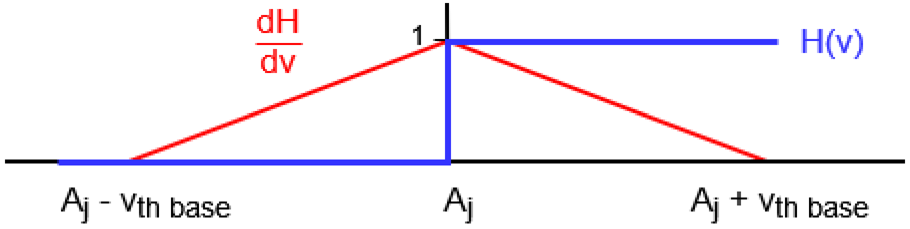

Overview

- Motivation
- Memory in neural networks
- E-prop
- SpiNNaker
- Wave form matching
- Temporal credit assignment
- Conclusions

SpiNNaker

- Memory constraints
  - 96 kB TCM per core (32kB ITCM, 64kB DTCM)
  - 128 MB SDRAM per chip
  - 18 cores/chip
- Real-time/online
- Weight precision (16-bit fixed point)
- Programmable

SpiNNaker board
(48 chips)

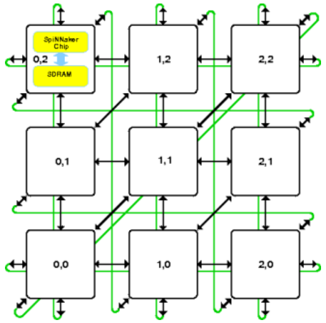

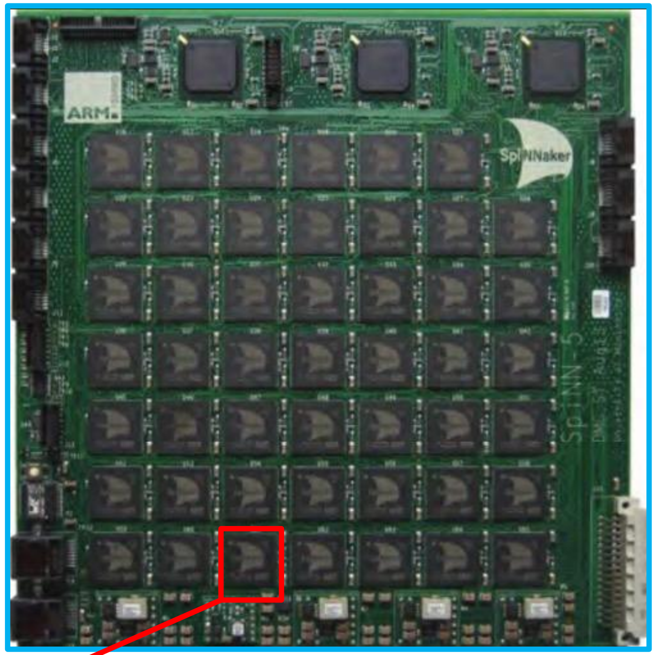

SpiNNaker chip
(18 ARM cores)

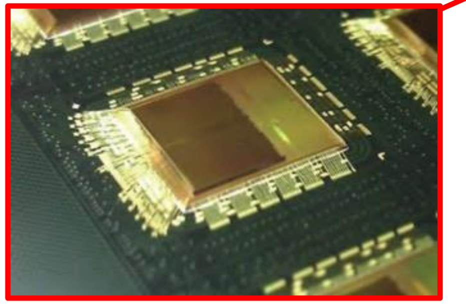

SpiNNaker implementation

- 8 neurons/core
- 250 synapses/neuron

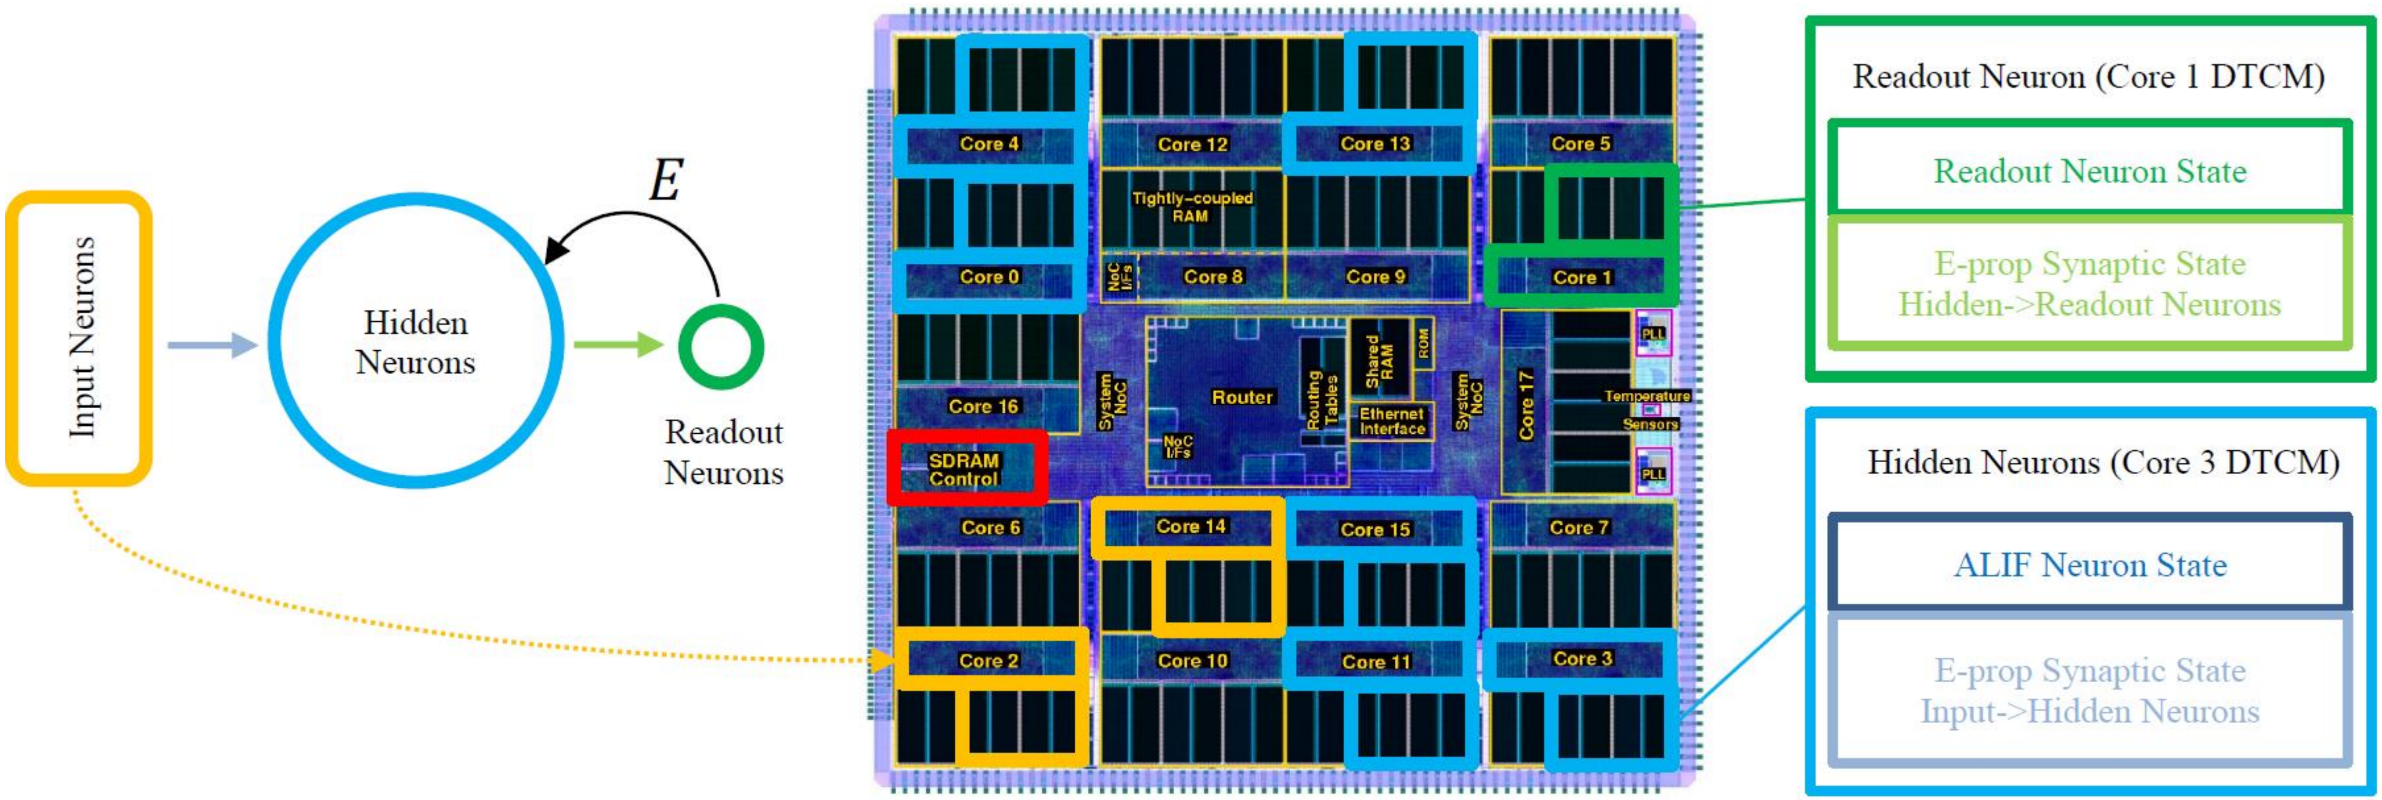

Overview

- Motivation
- Memory in neural networks
- E-prop
- SpiNNaker
- Wave form matching
- Temporal credit assignment
- Conclusions

Waveform matching

- Proof of concept to test matching a target output
- A repeated Poisson spike source is injected into a network
- A target waveform is compared with the output neuron's membrane potential to produce an error
- The learning signal is broadcast to the network to reduce the error

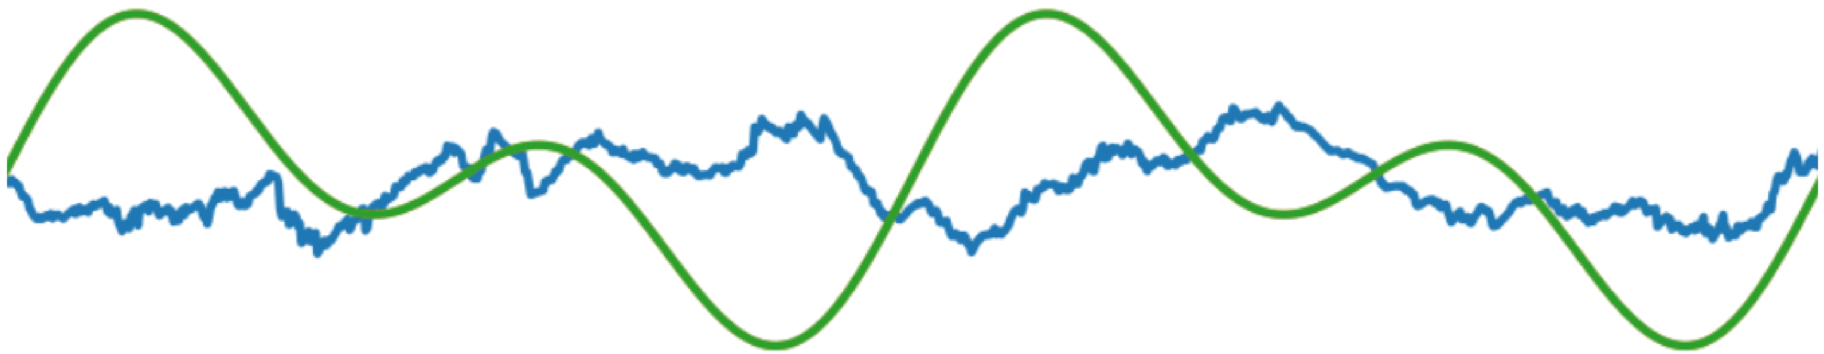

Waveform matching architecture

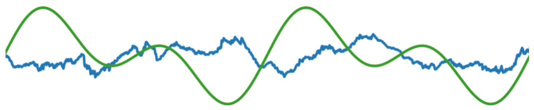

Repeating Poisson spike train

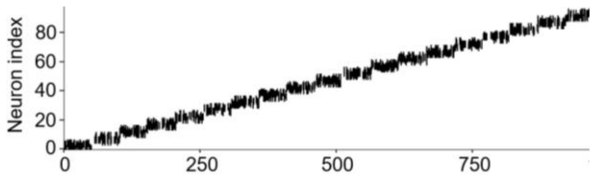

Hidden population
n=100

Only LIF
No adaptive threshold

Readout neuron compares desired output with V_{mem}

Learning signal

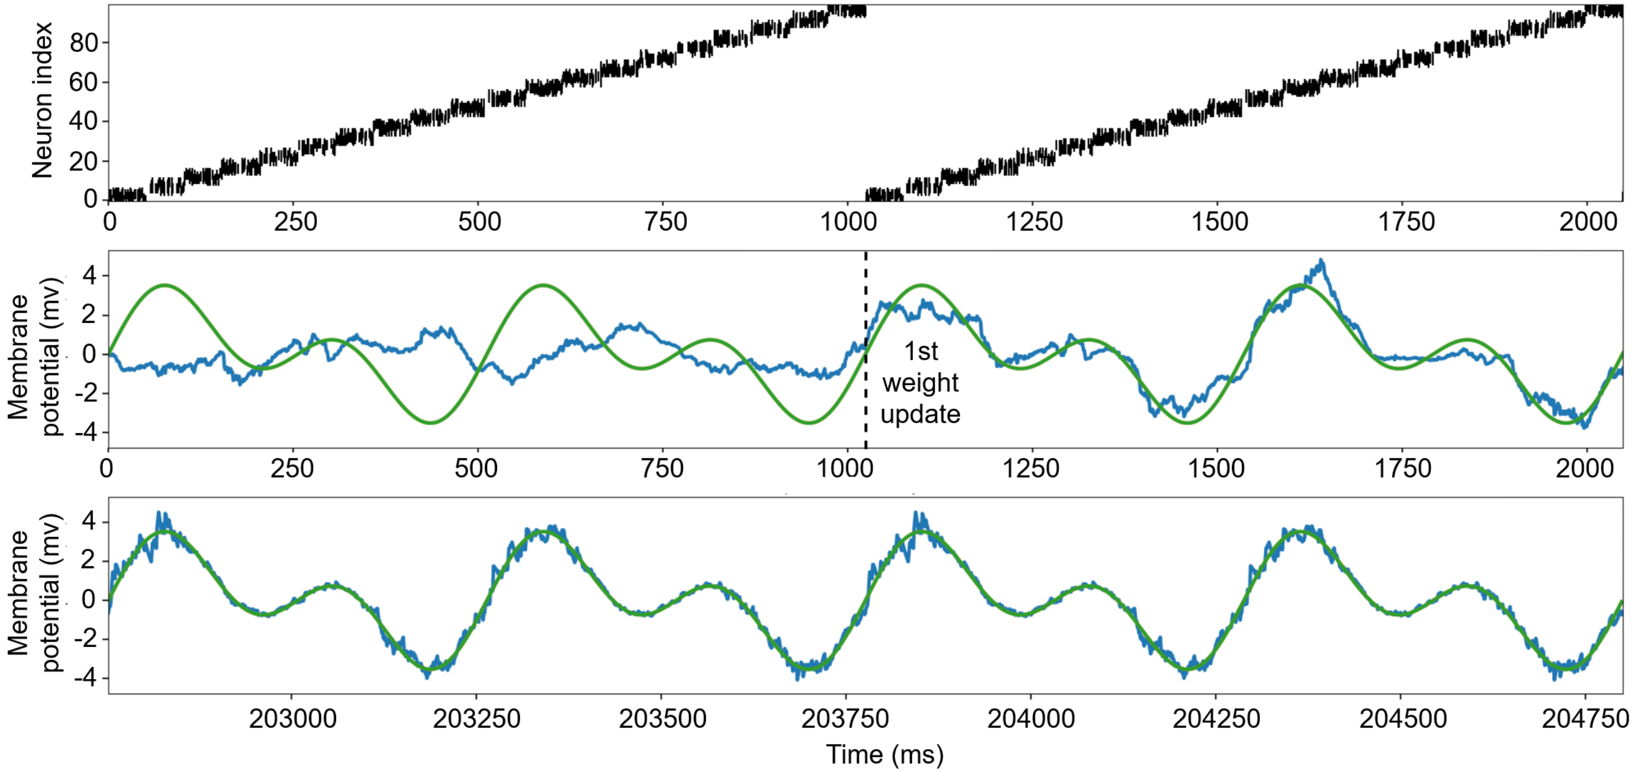

Neuron index
Membrane potential (mv)
1st weight update
Membrane potential (mv)
Time (ms)

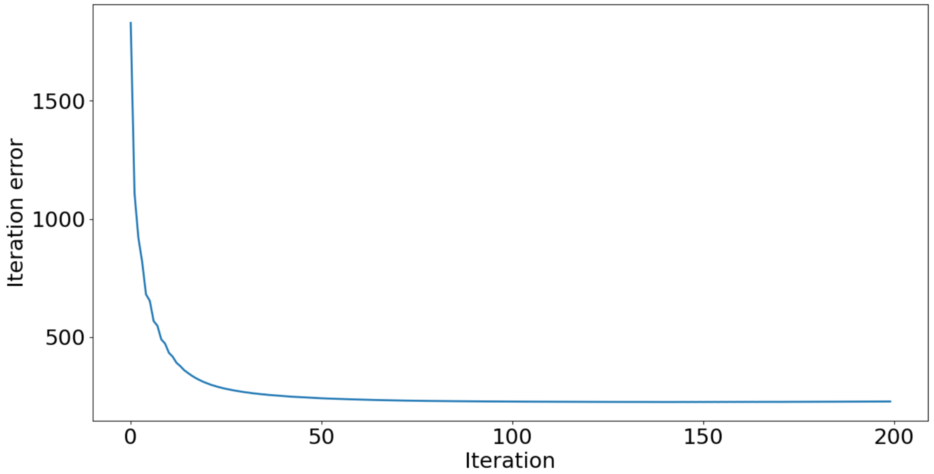

Iteration error
Iteration

Overview

- Motivation
- Memory in neural networks
- E-prop
- SpiNNaker
- Wave form matching
- Temporal credit assignment
- Conclusions

Temporal credit assignment

- Left and right signals are presented to a mouse as it walks down a hallway
- At the end of the hallway it should select the direction which presented the most cues

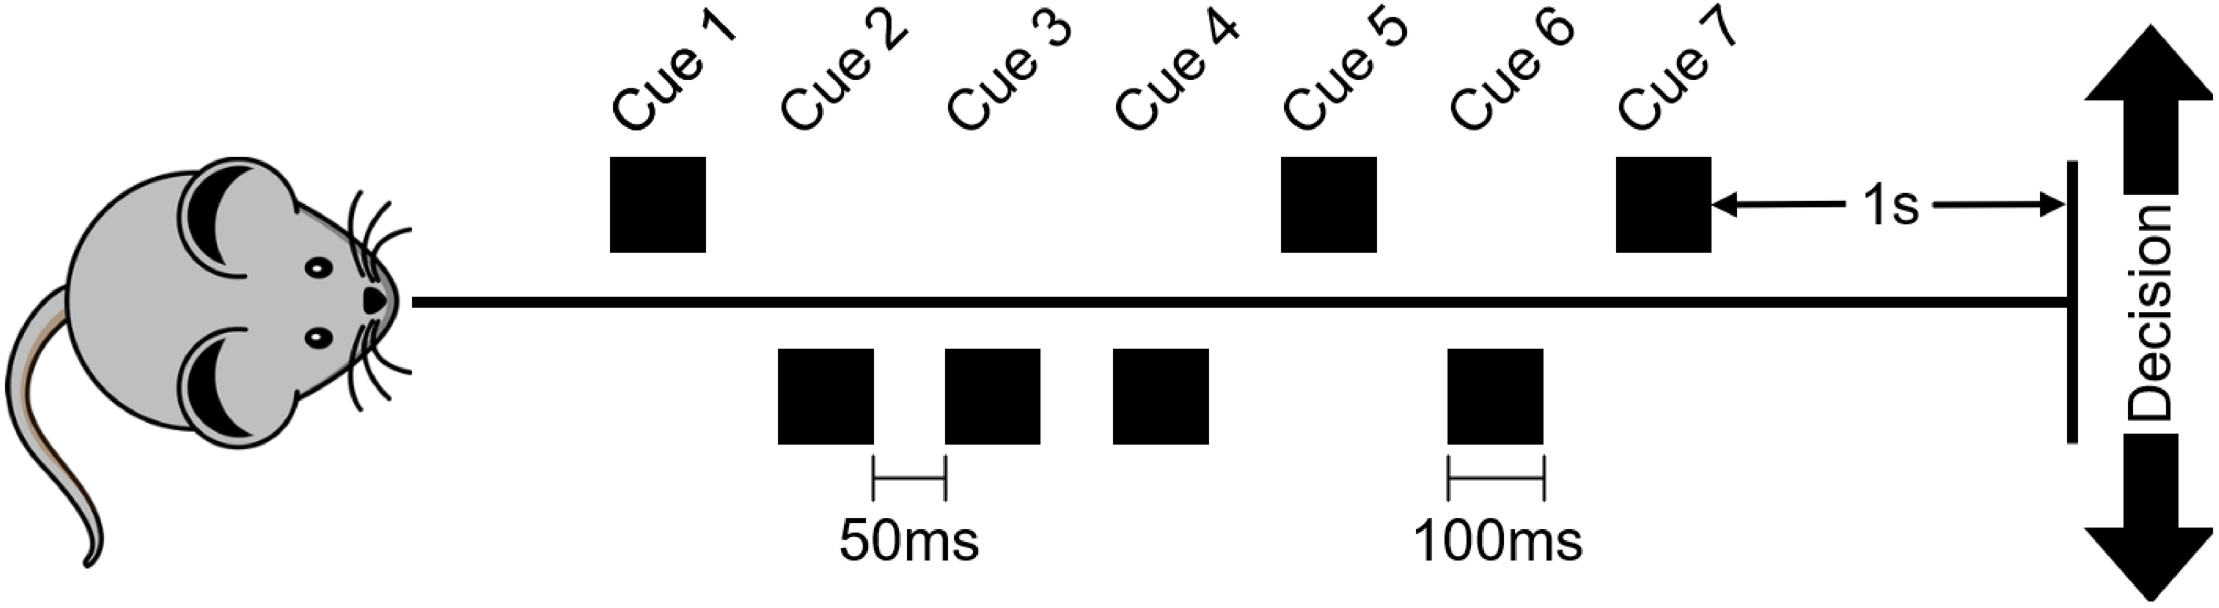

Temporal credit assignment

- input 0 = left cue
- input 1 = right cue
- input 2 = prompt signal
- input 3 = 10Hz noise

- ON = 100Hz
- OFF = 0Hz

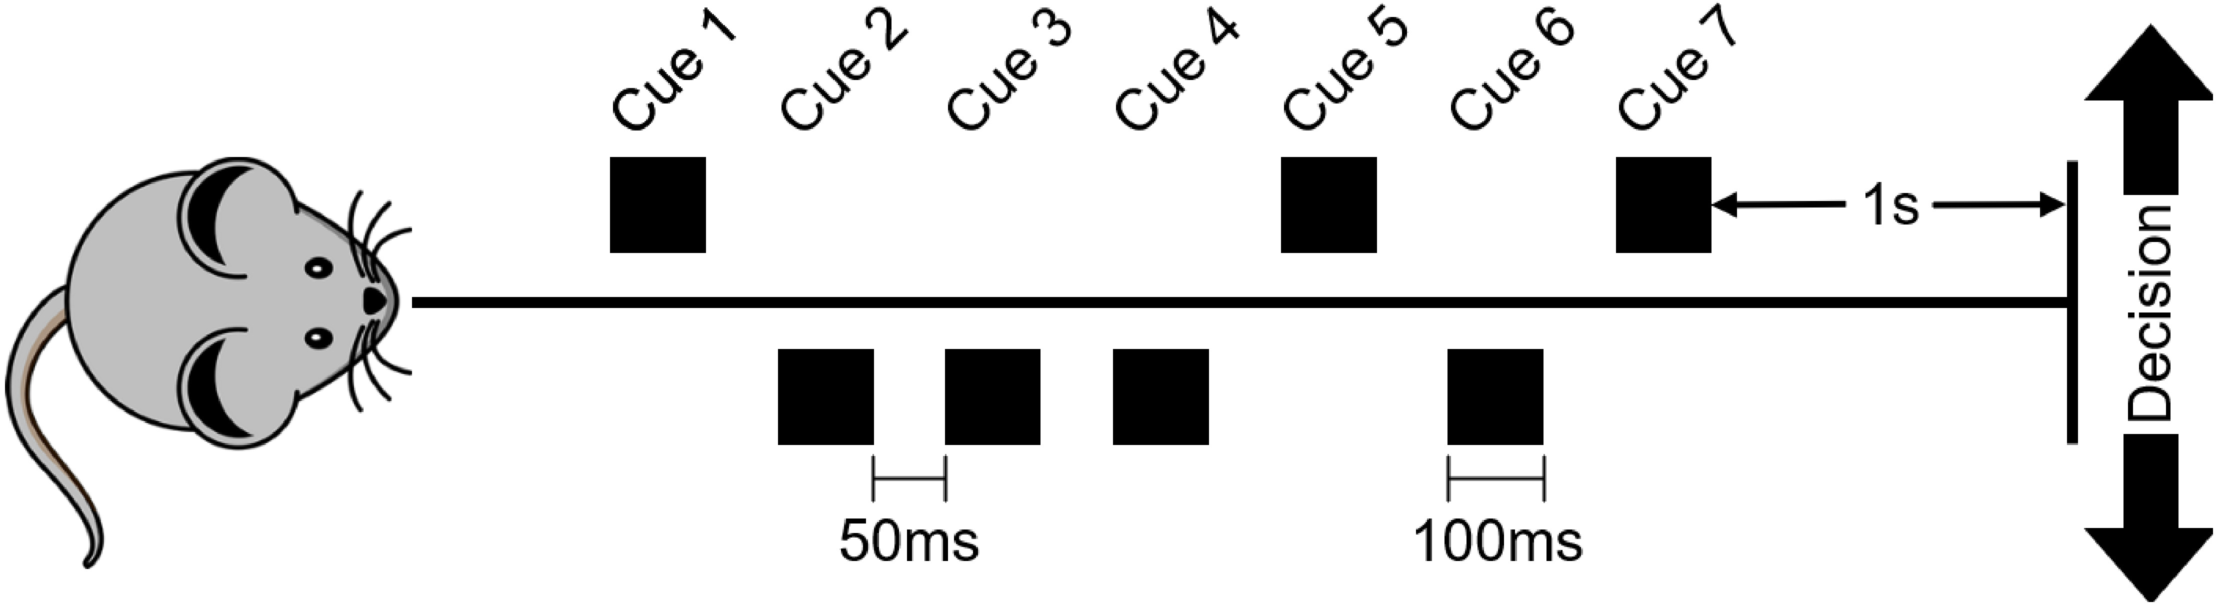

- output 0 = right decision
- output 1 = left decision

Temporal credit assignment

- Receives left and right cues
- After presentation there is 1s delay
- Prompt signal indicates to the network that a decision must be made
- A learning signal is only broadcast during the duration of the prompt
- Curriculum learning increases the number of cues after threshold performance (1, 3, 5, 7)

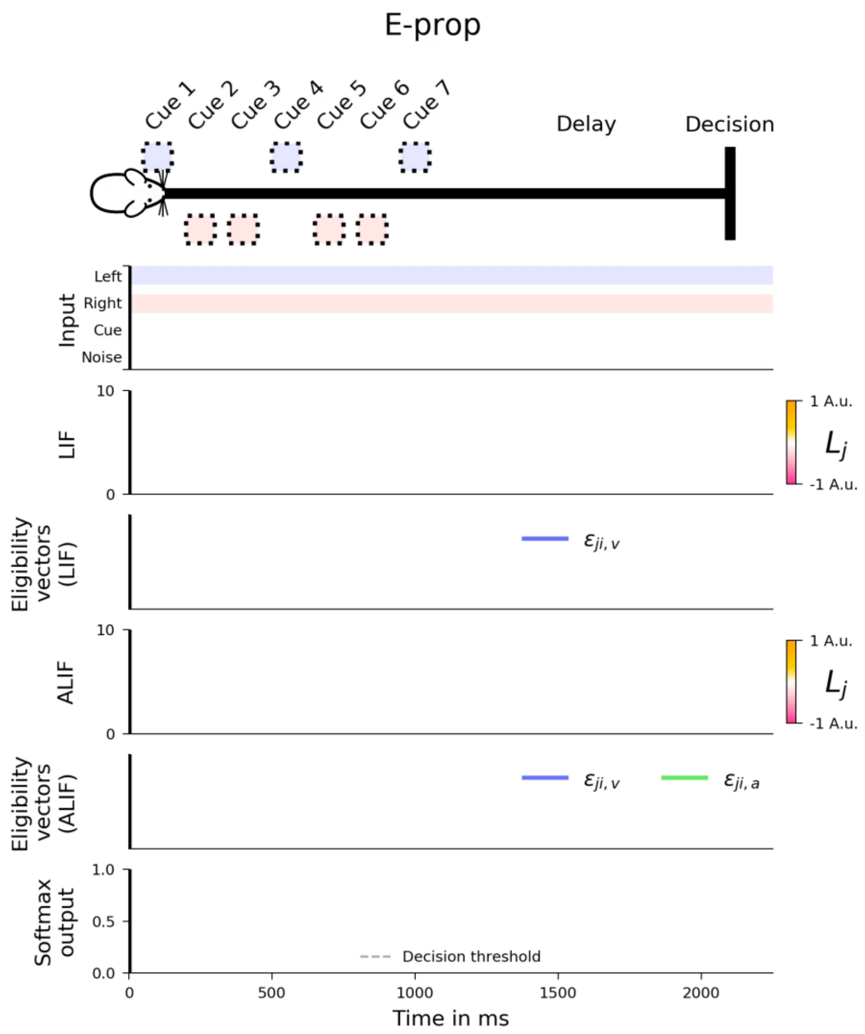

Bellec, G., Scherr, F., Subramoney, A., Hajek, E., Salaj, D., Legenstein, R., & Maass, W. (2020). A solution to the learning dilemma for recurrent networks of spiking neurons. *Nature communications*, *11*(1), 1-15.

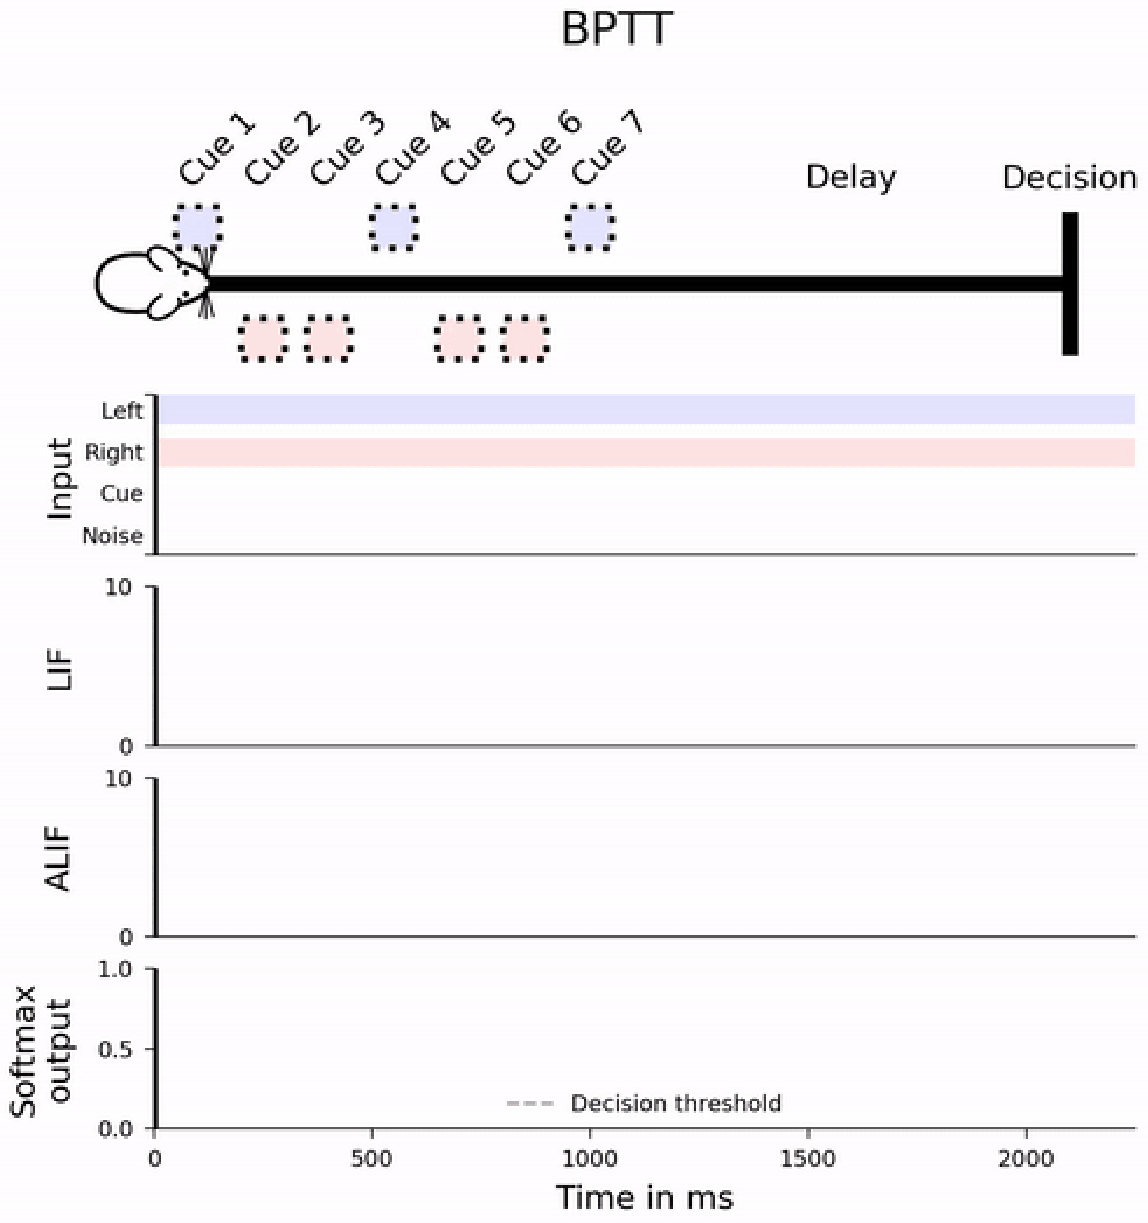

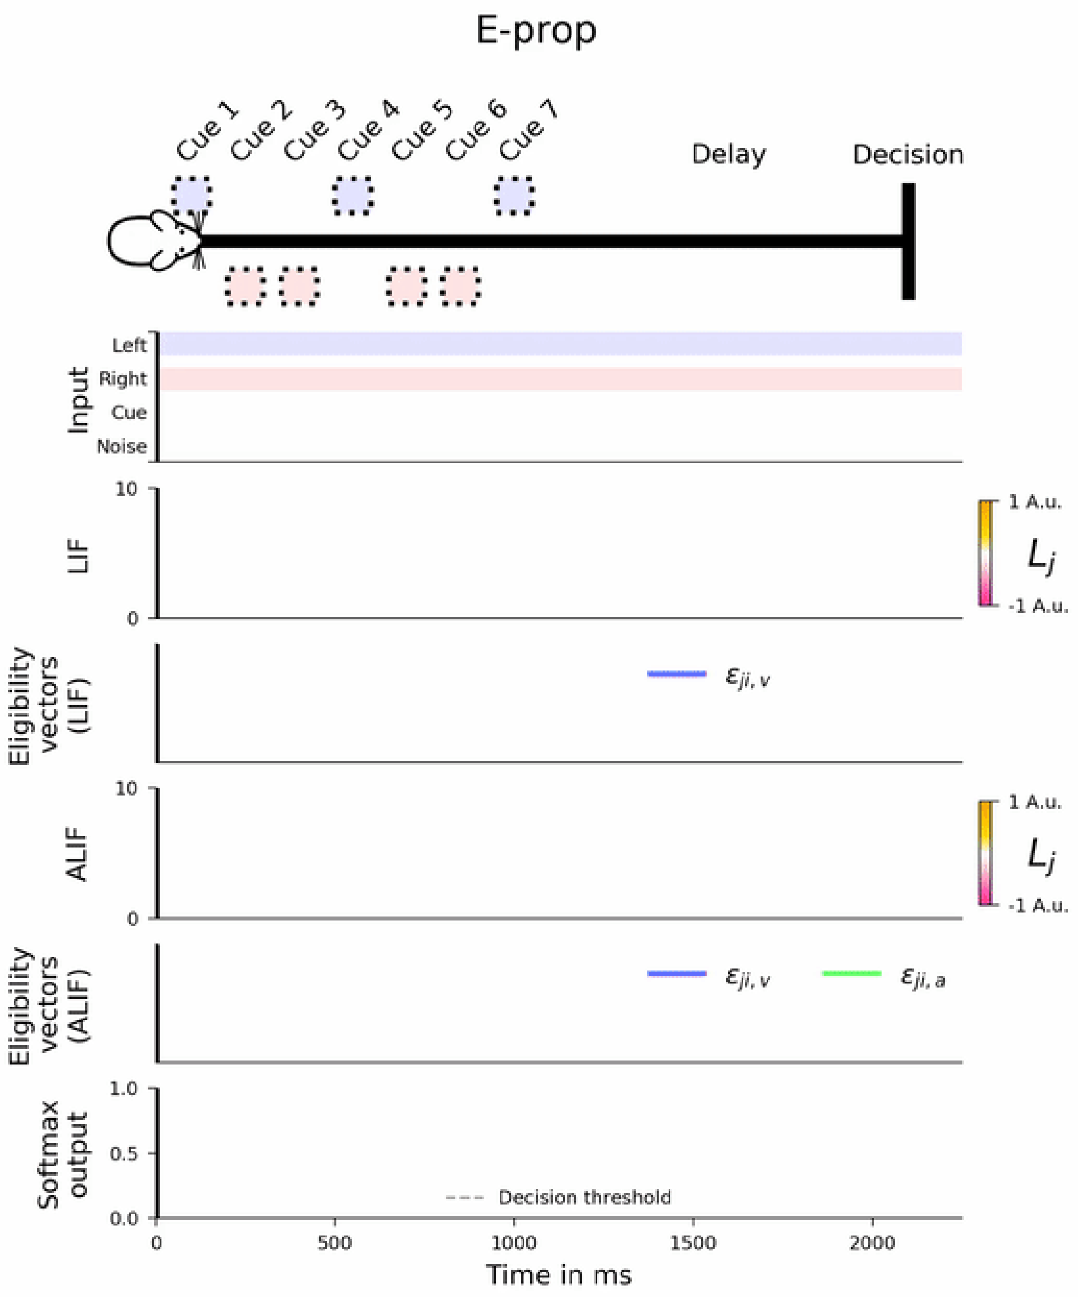

Bellec, G., Scherr, F., Subramoney, A., Hajek, E., Salaj, D., Legenstein, R., & Maass, W. (2020). A solution to the learning dilemma for recurrent networks of spiking neurons. *Nature communications*, *11*(1), 1-15.

Temporal credit assignment architecture

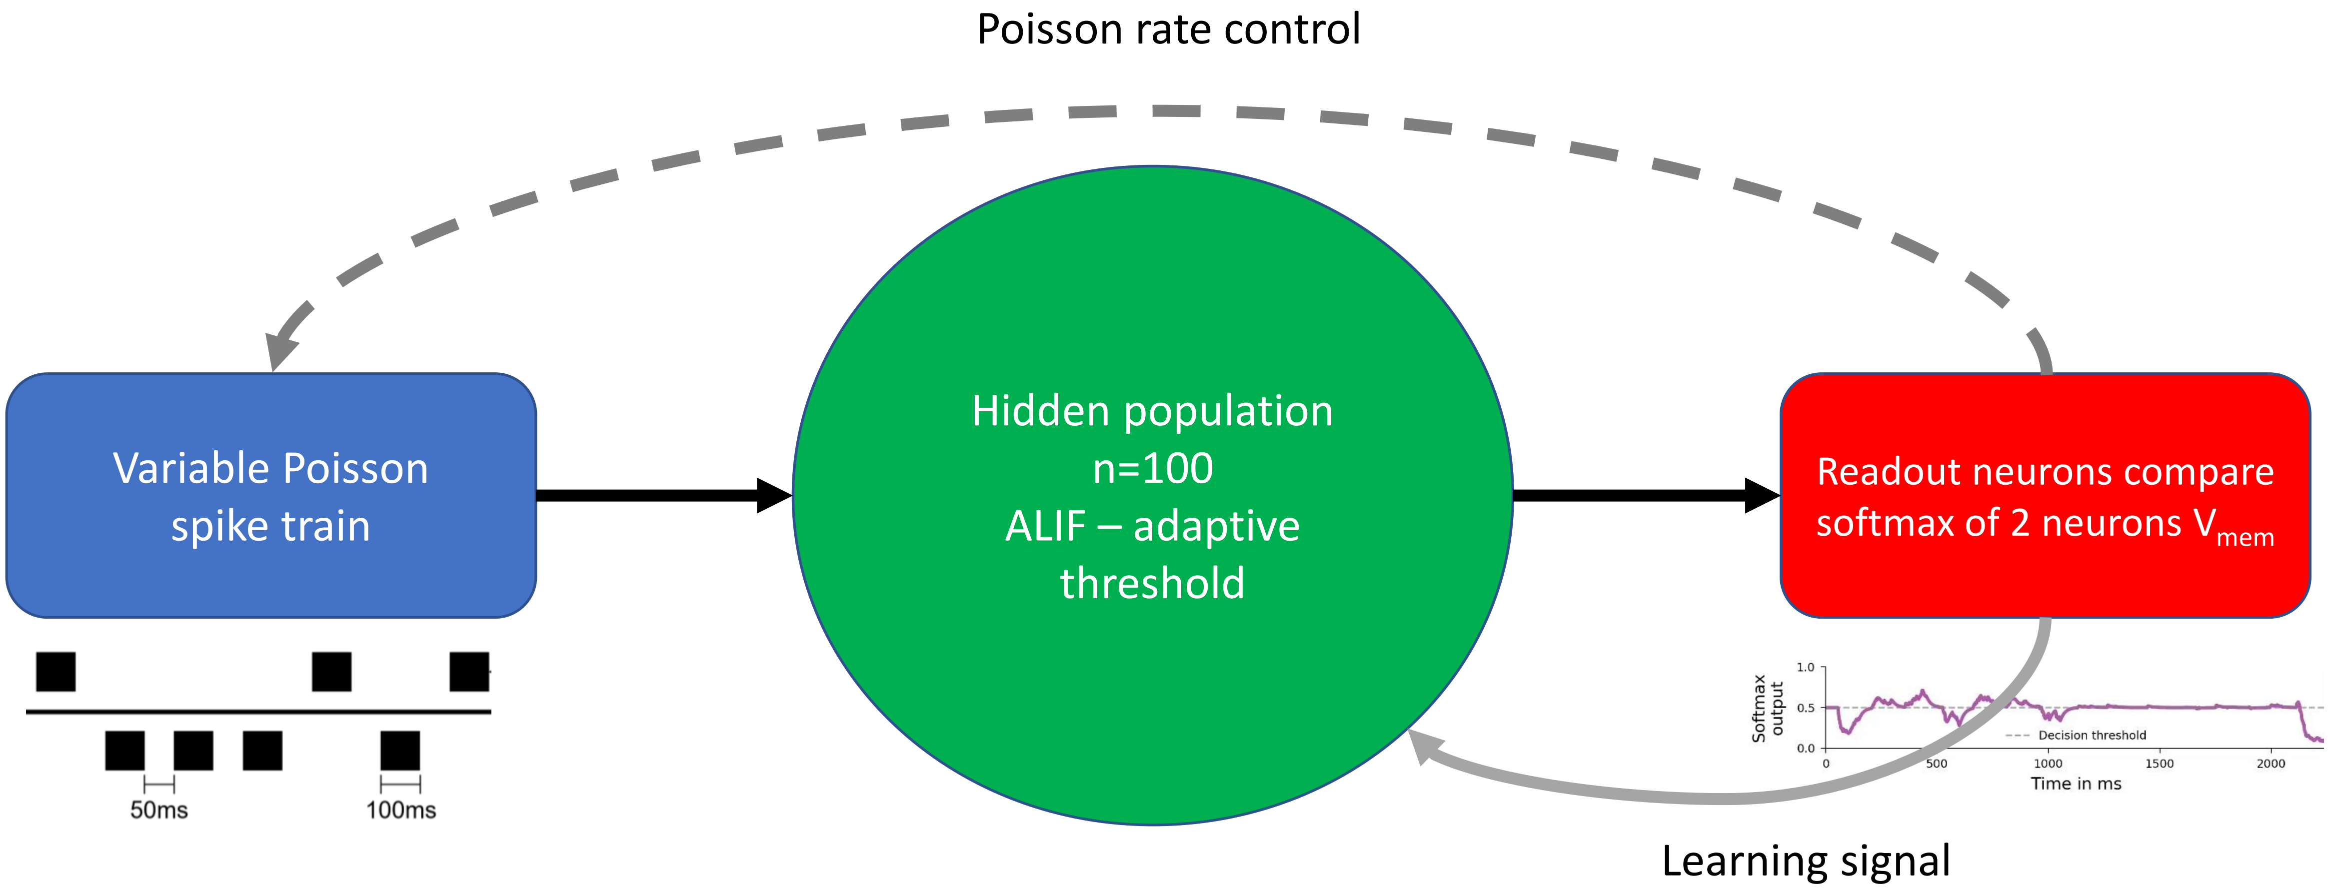

Temporal credit assignment

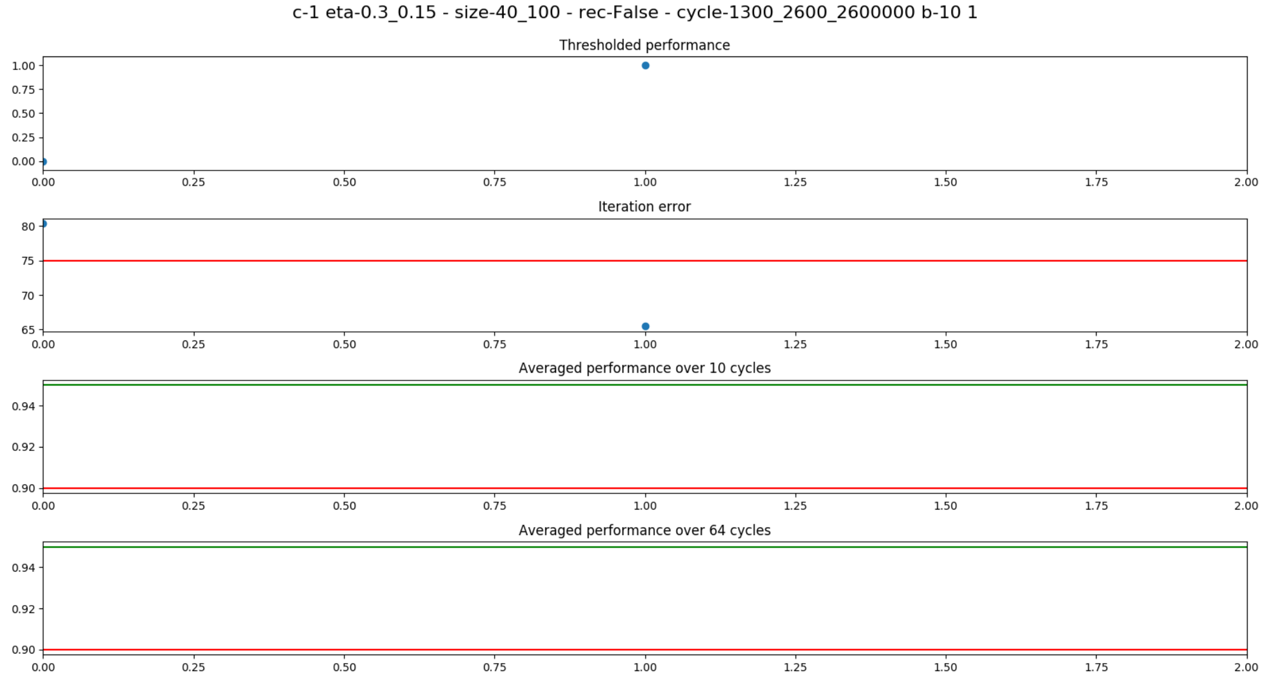

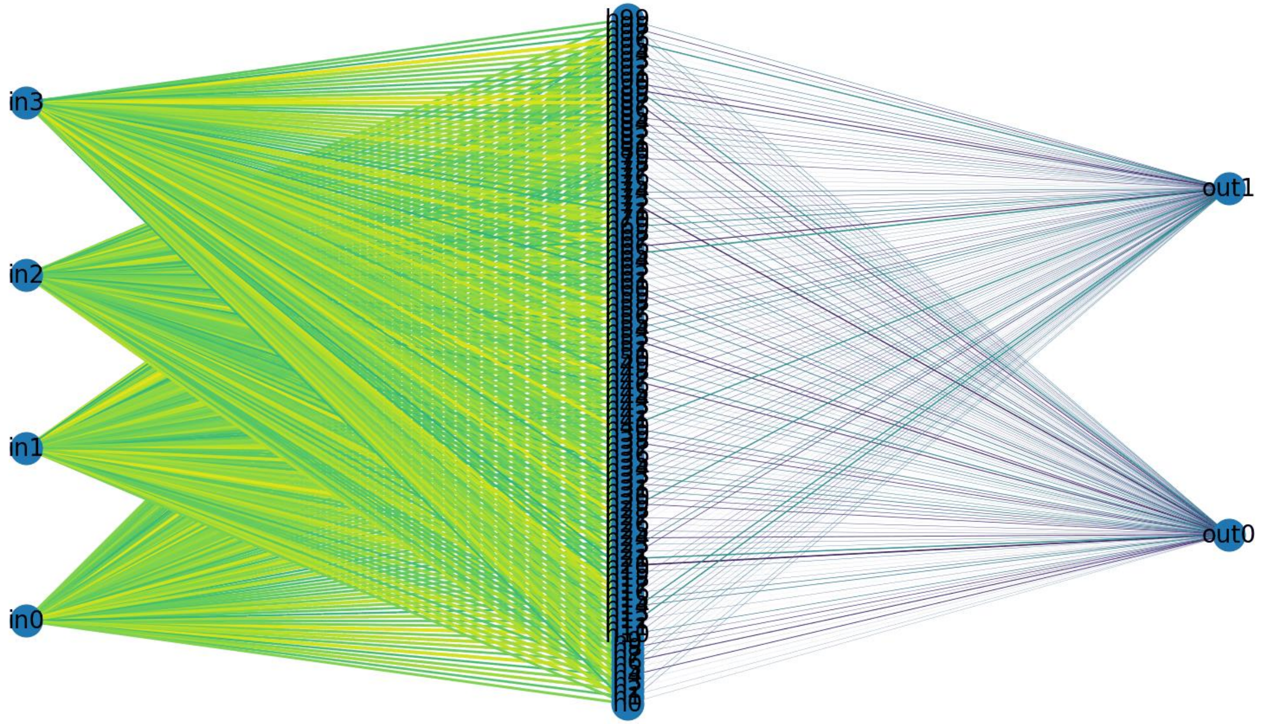

Typical result for temporal credit assignment

- Significant learning takes place during the first stage
- Eventually the task becomes difficult enough to degrade performance requiring substantial retraining

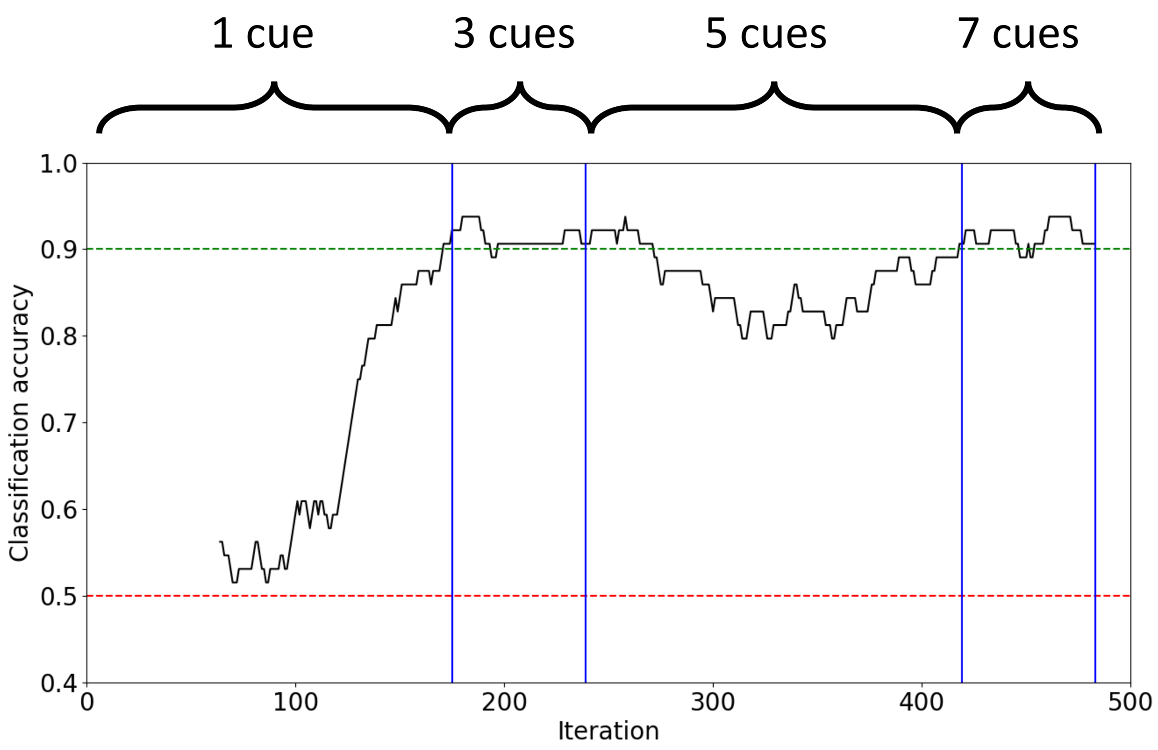

Overview

- Motivation
- Memory in neural networks
- E-prop
- SpiNNaker
- Wave form matching
- Temporal credit assignment
- Conclusions

Conclusions

Challenges:

- Neuromorphic restrictions
- Firing rate regularisation

Future work:

- Move to multi-core model
- Recurrent connections
- Reinforcement learning tasks (eg Pong)

8 neurons constant firing

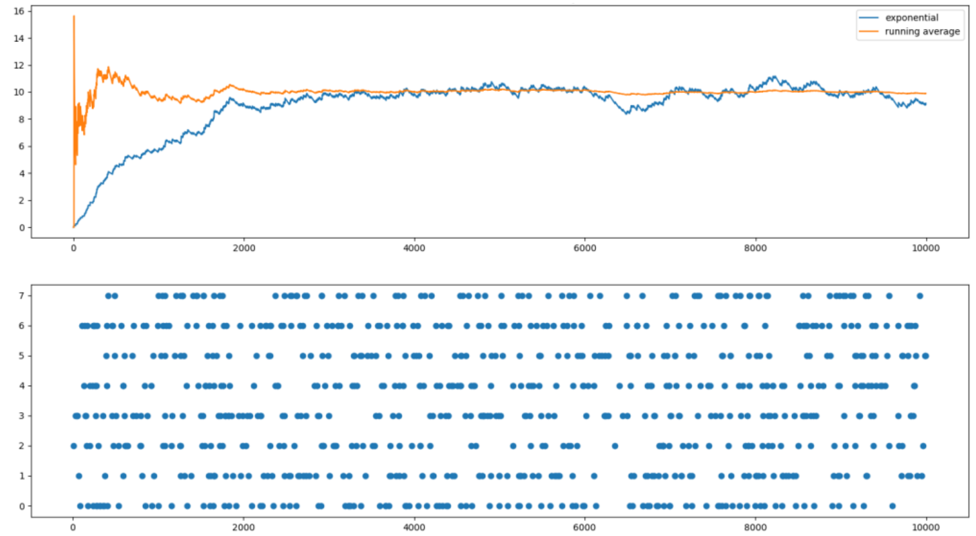

8 neurons burst firing

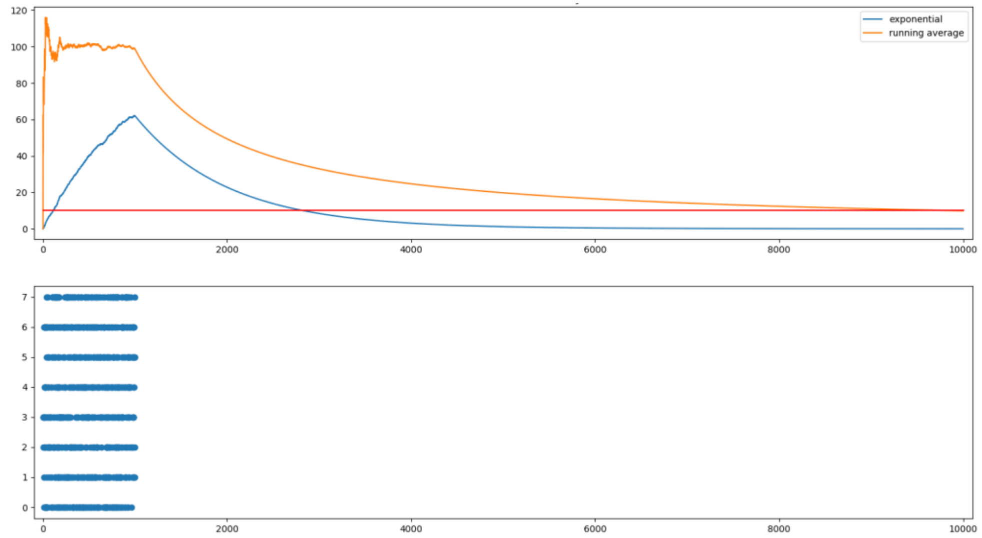

Questions?