Interplay of synaptic plasticity and probabilistic brain theories

Wulfram Gerstner
EPFL, Lausanne

Joint work with:
Danilo J. Rezende

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

Funding acknowledged:
ERC, BrainScales, HBP,
Swiss Natl. Sci. Foundation,

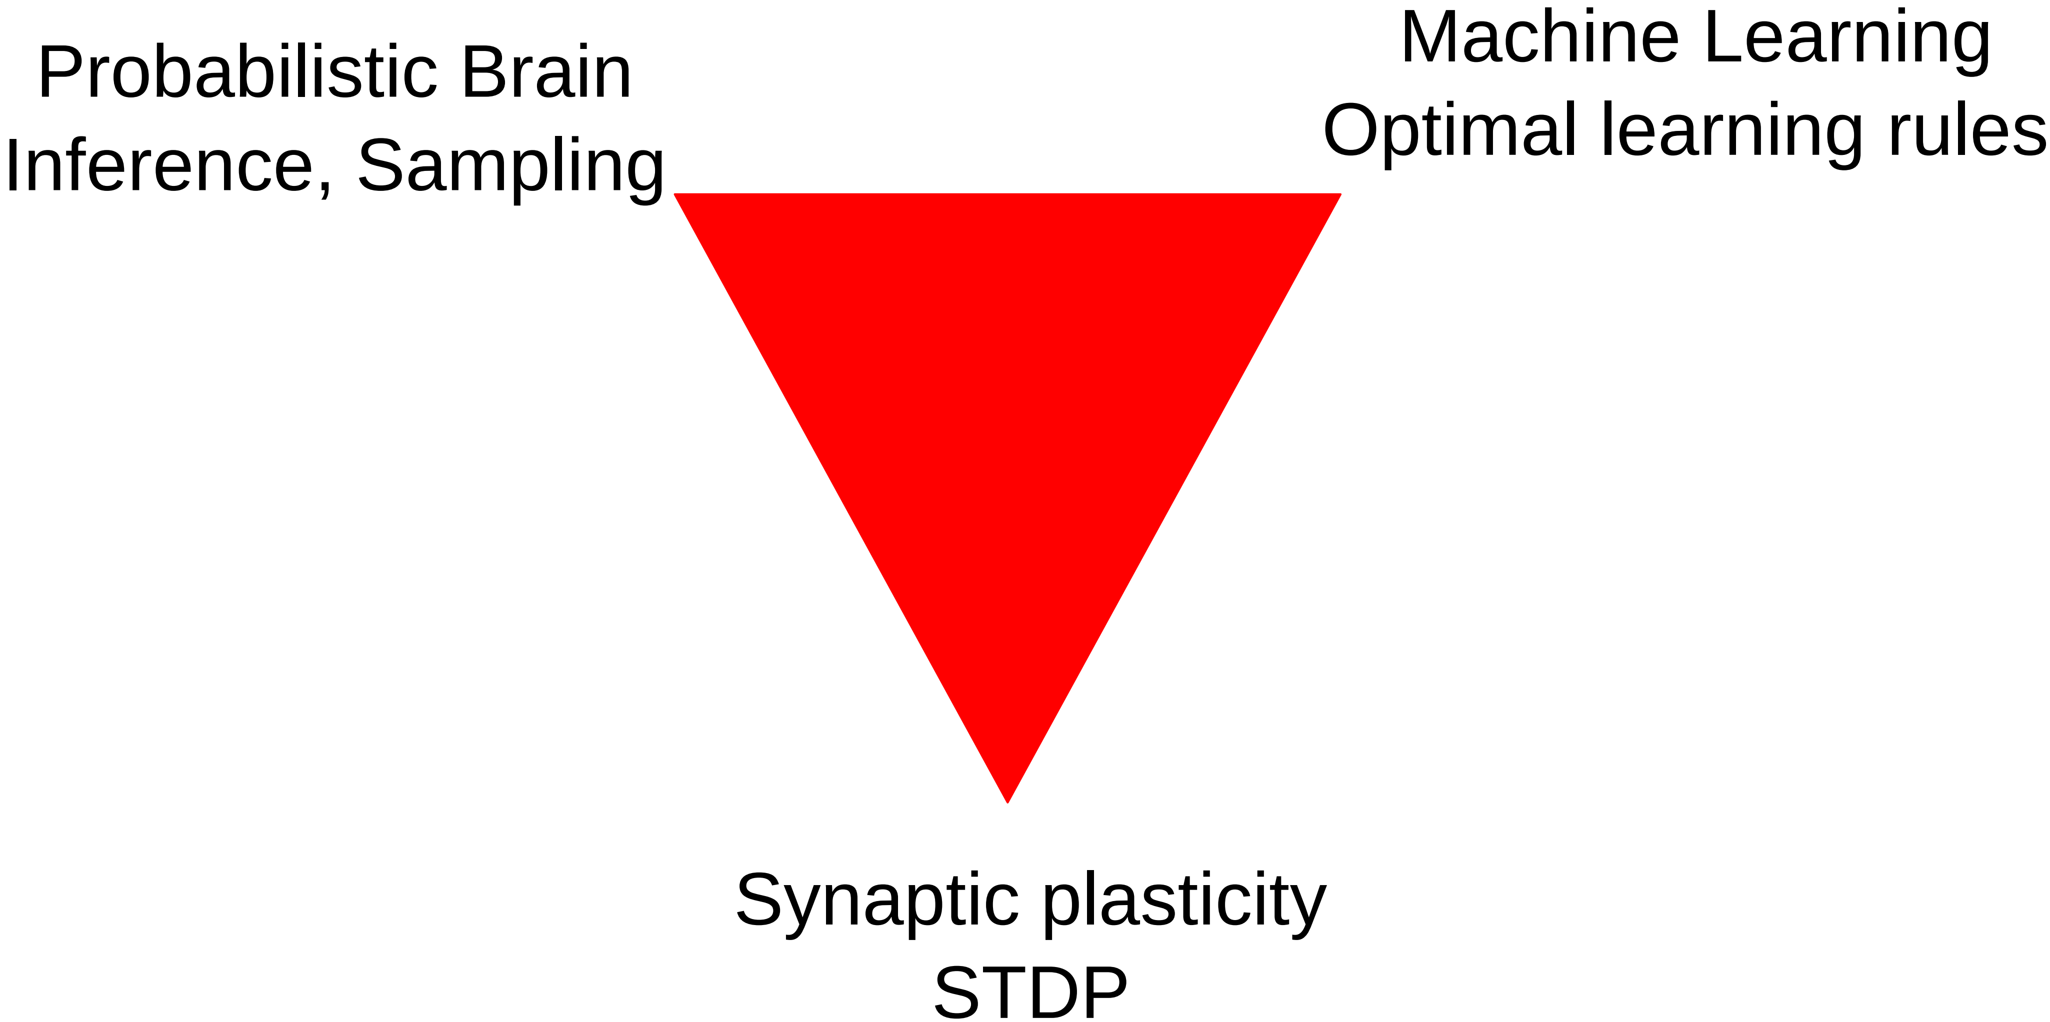
Probabilistic Brain
Inference, Sampling
Machine Learning
Optimal learning rules
Synaptic plasticity
STDP

Probabilistic Brain,
Sampling, Bayesian Brain

Kording and Wolpert 2004;
Berkes et al. 2011;
Knill and Pouget 2004

Machine Learning
Optimal learning rules

Dayan 2000
Friston and Stephan, 2007
Beal and Gahrahmani, 2006

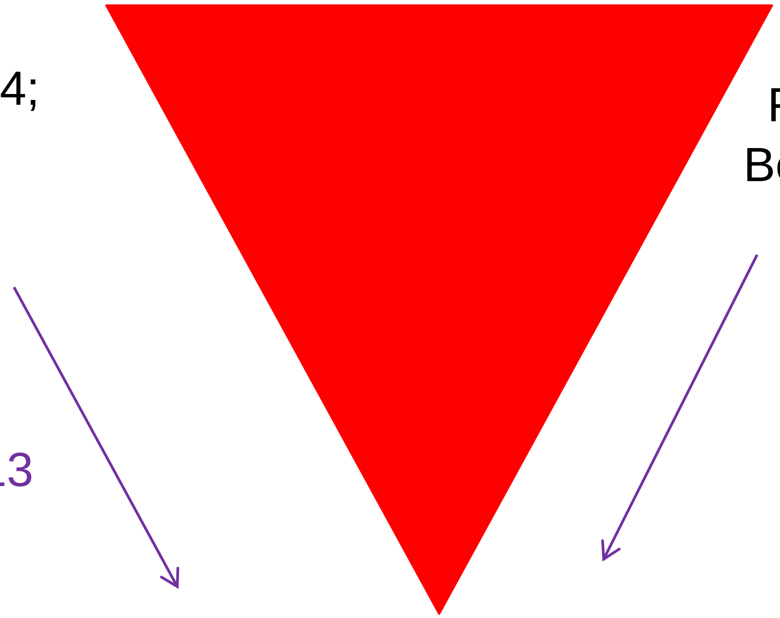

Nessler et al. 2013

Pfister et al. 2006
Brea et al. 2013

Synaptic plasticity
STDP

Markram et al. 1997; Bi and Poo 2001
Artola and Singer 1993

Behavioral Learning – and synaptic plasticity

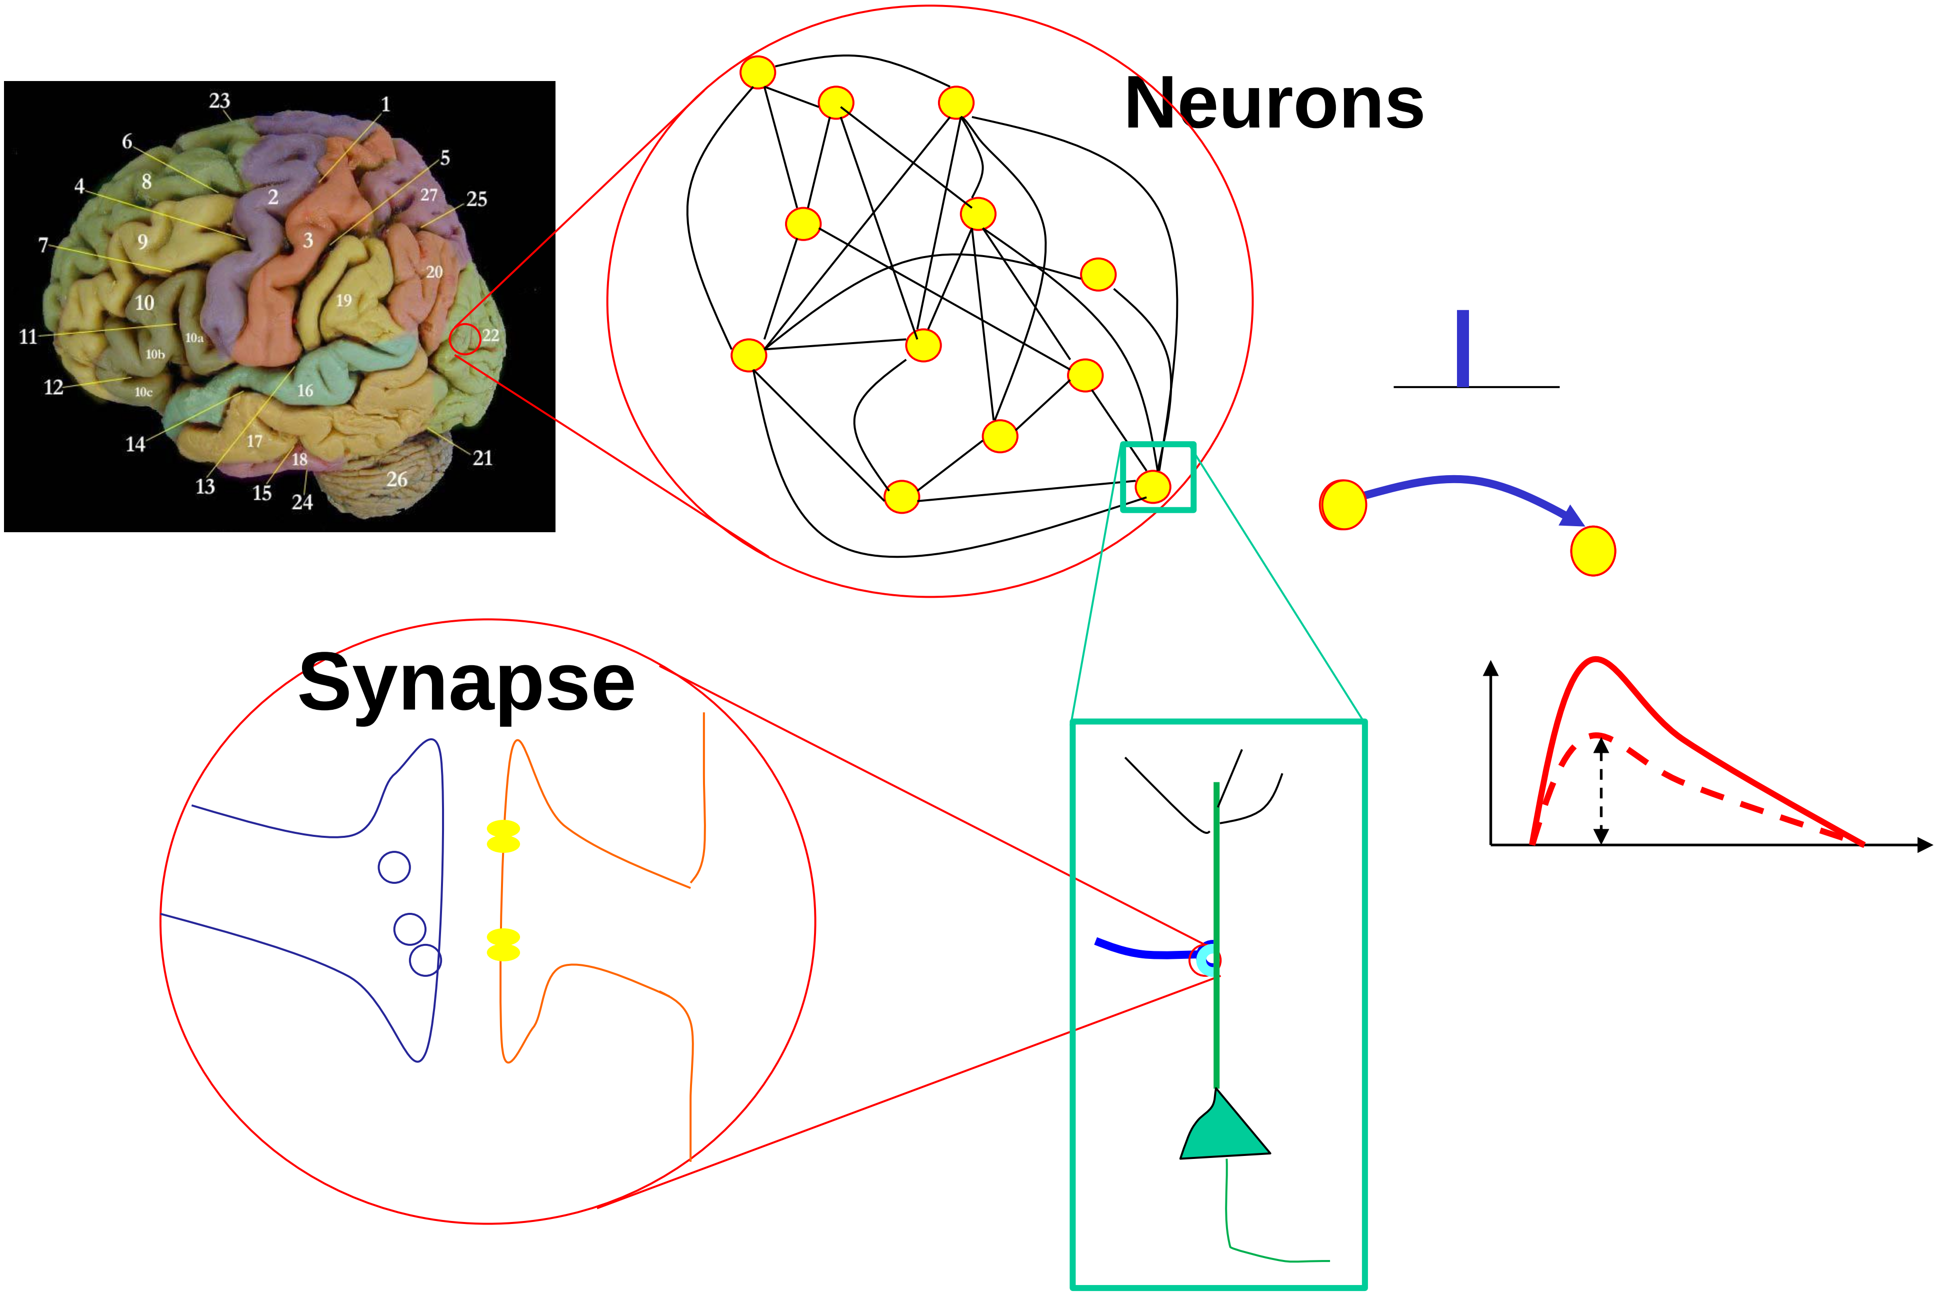

Synaptic Plasticity =Change in Connection Strength

Hebbian Learning

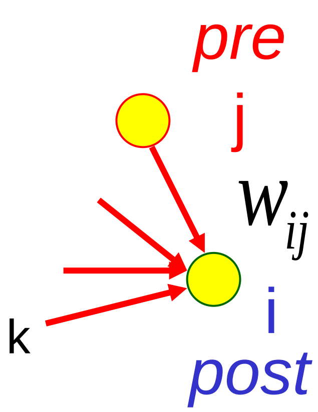

$$\Delta w_{ij} \propto F(pre, post)$$

Spike arrival,
EPSP

Postsyn. Spikes
voltage

When an axon of cell **j** repeatedly or persistently takes part in firing cell **i**, then j's efficiency as one of the cells firing i is increased

Hebb, 1949

Is Hebbian learning useful? –
Developmental learning/rec. field development

Is Hebbian learning linked to experiments?

Experimental Induction Protocols

Markram et al. 1997, Bi an Poo 1998, Sjostrom et al. 2001

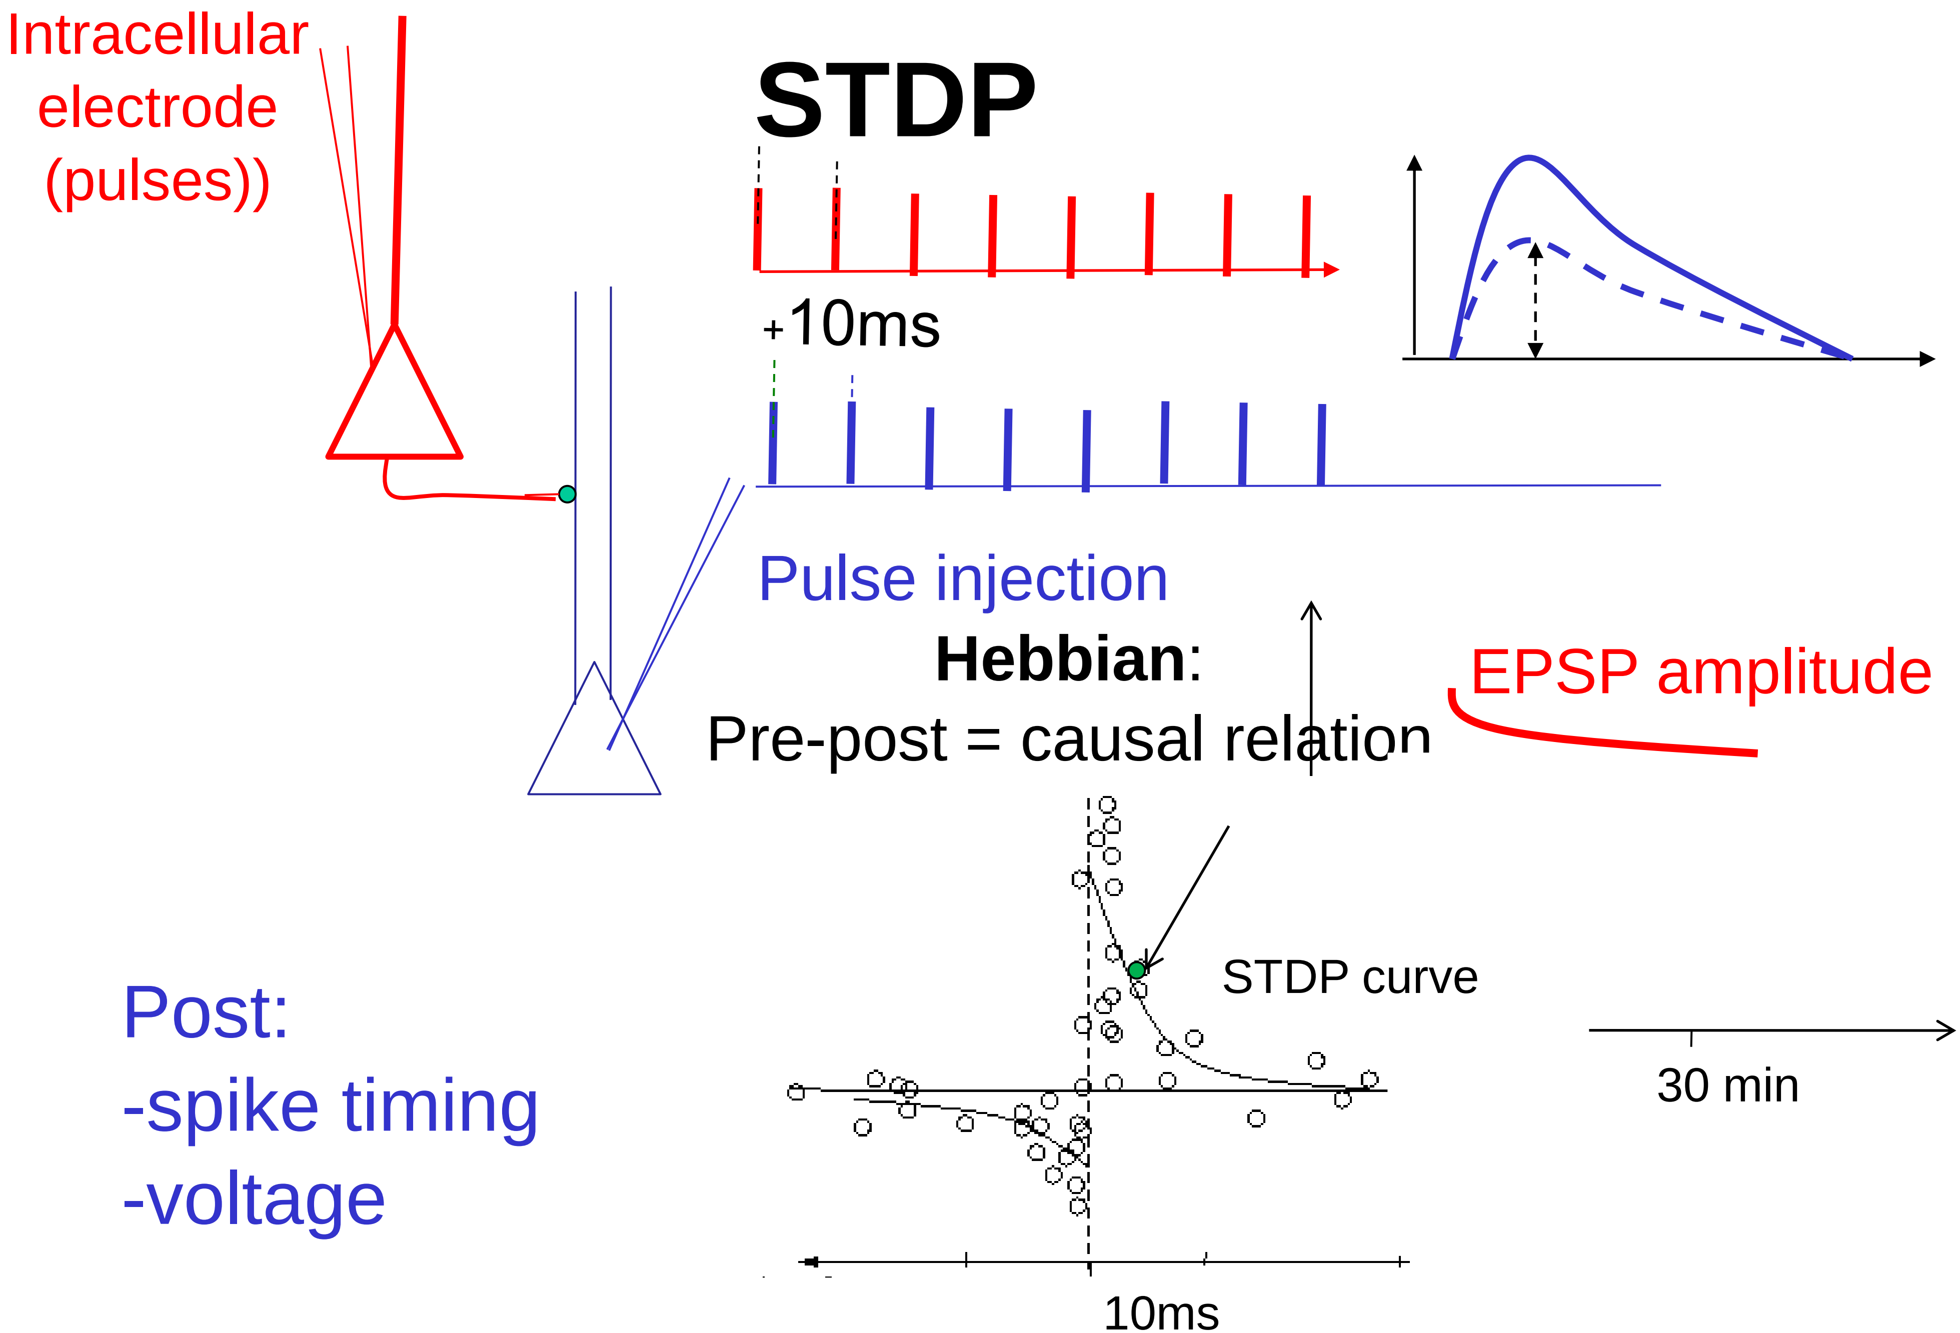

Post:
-spike timing
-voltage

Three factor rules (schematic)

Eligibility trace:
Synapse keeps memory
Of pre-post Hebbian events

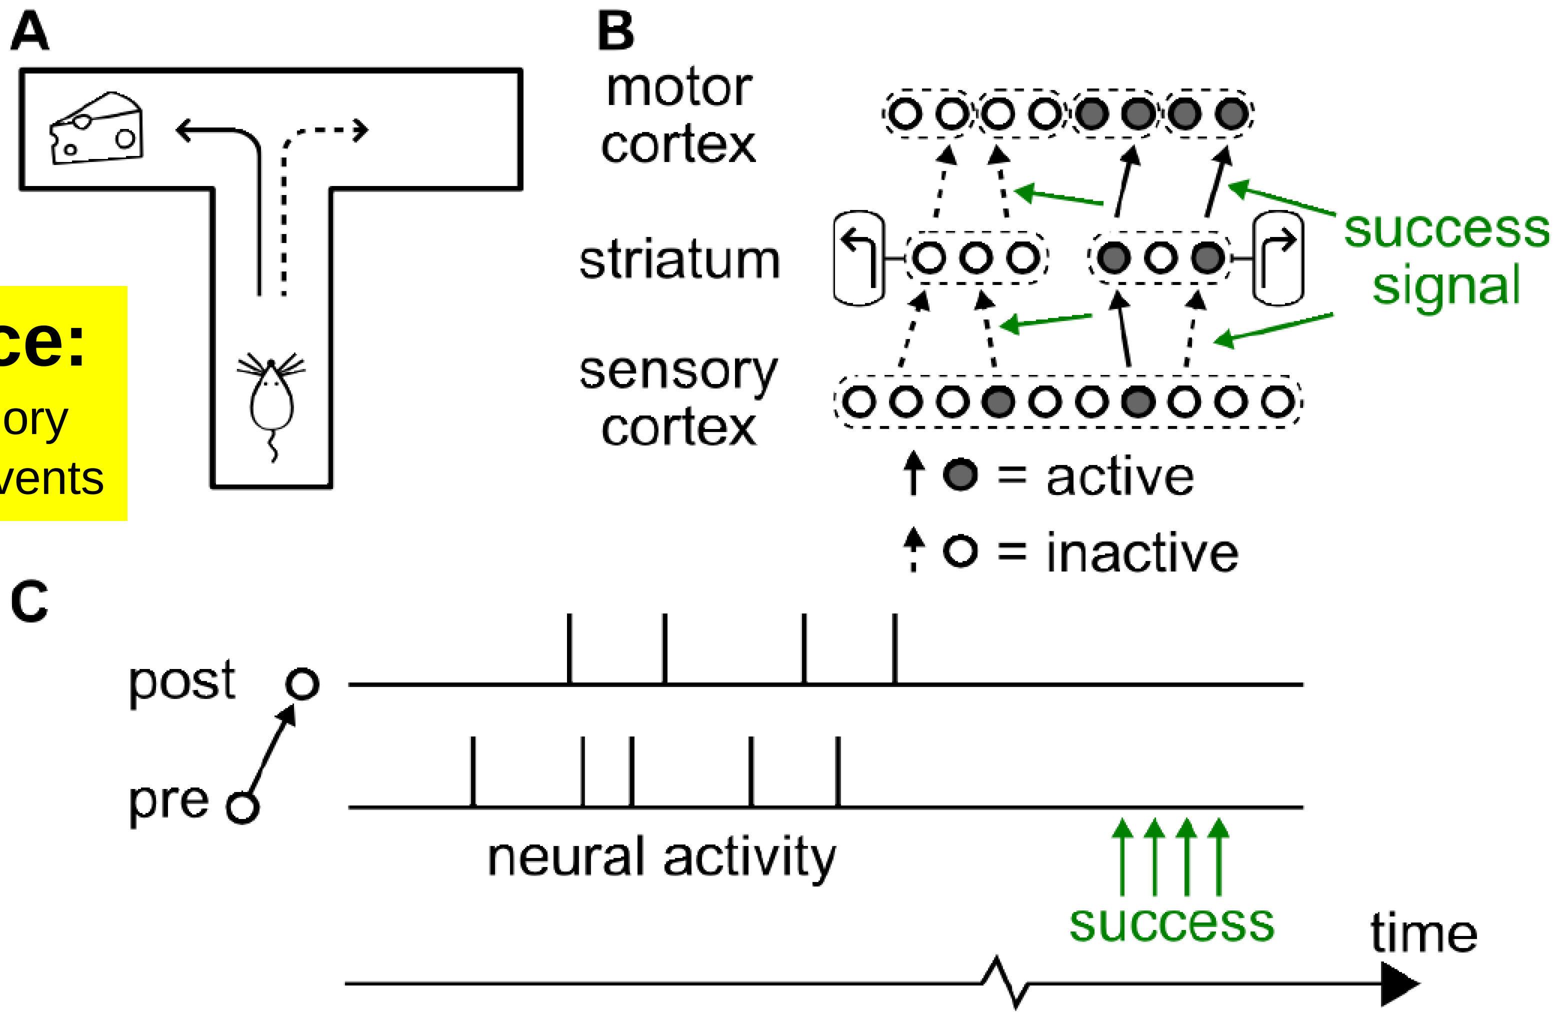

→ Reinforcement learning: success = reward – (expected reward)

Three factor rules (schematic)

Neuromodulator:
Novelty/surprise
Reward/success

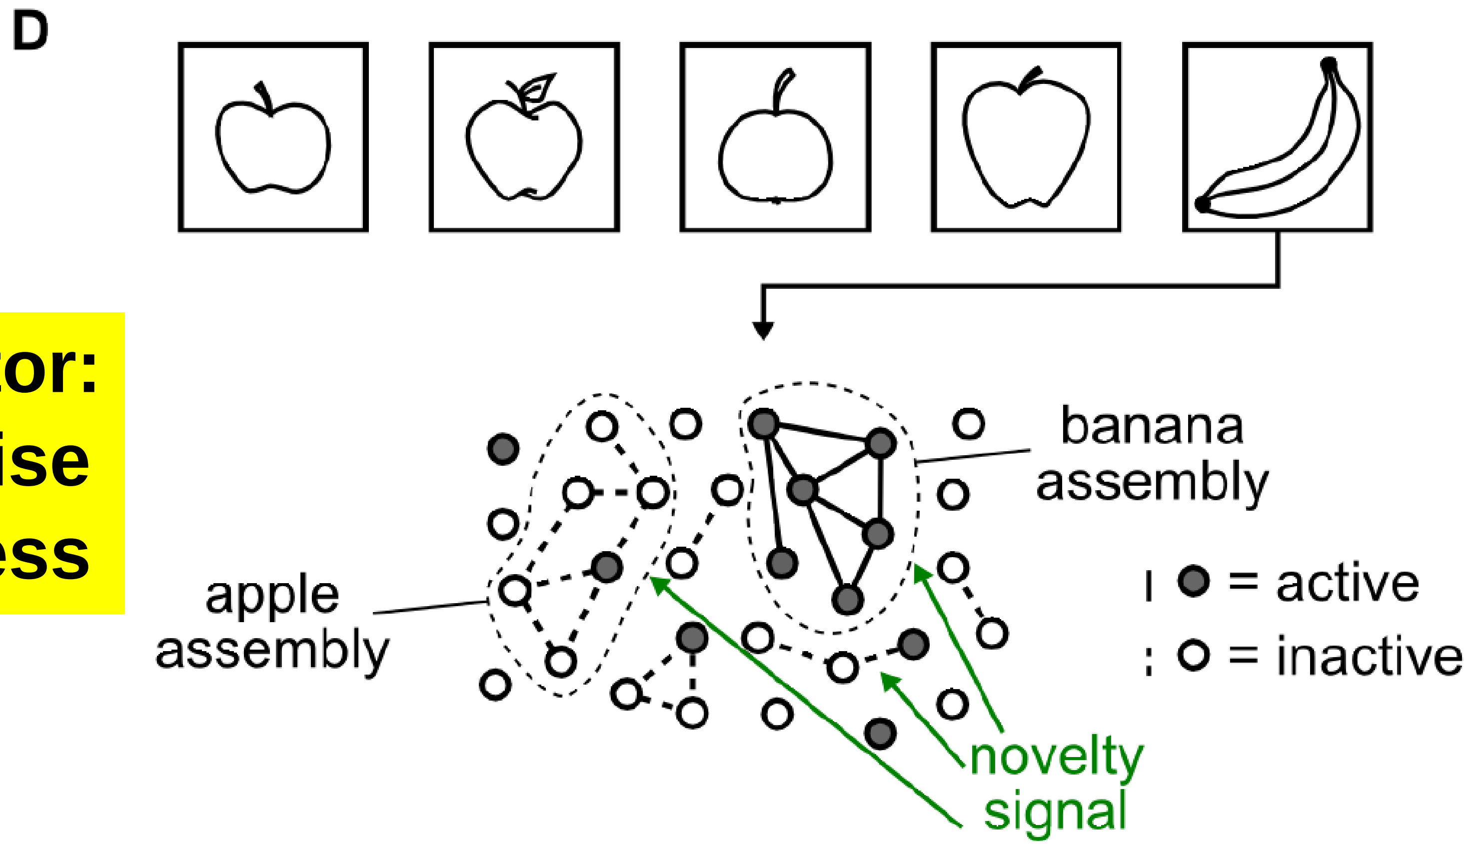

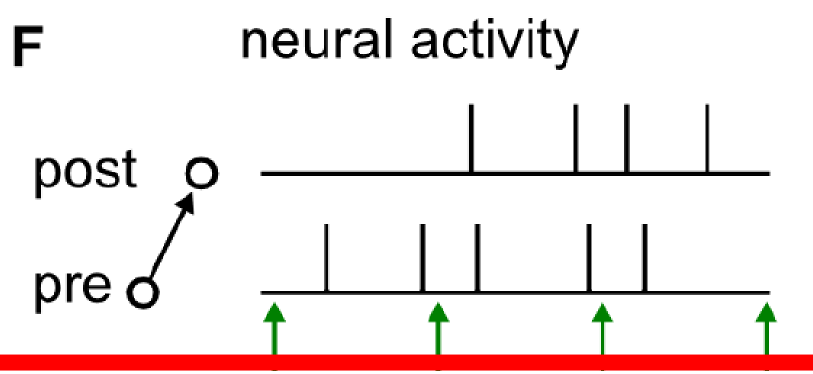

→ Novelty based learning: novelty = ?

Three-factor rules in theory

$$\Delta w_{ij} \propto F(pre, post, 3rd\ \ factor)$$

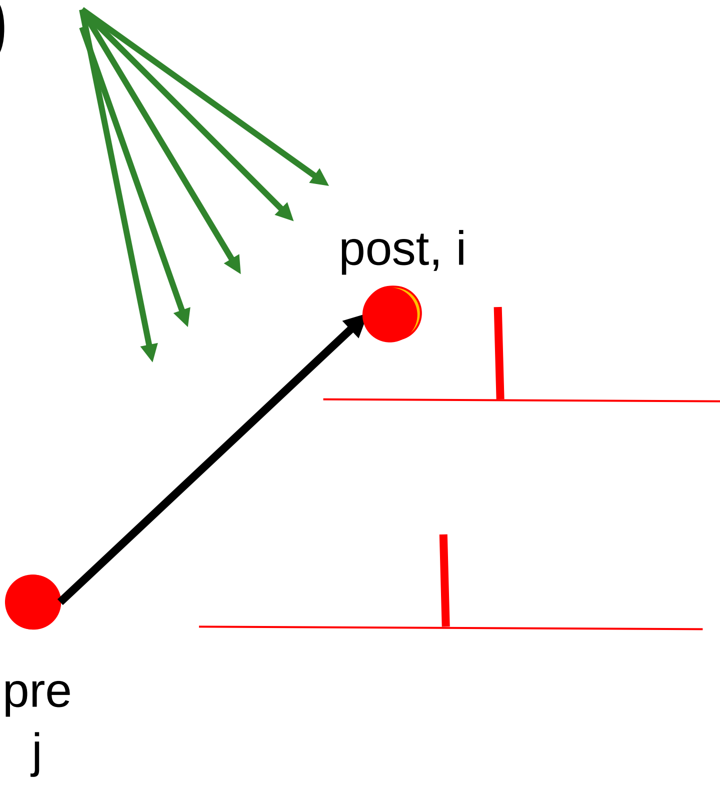

R-STDP

(Izhikevich, 2007;
Florian, 2007;
Legenstein et al., 2008)

Success/gating/novelty

$$\Delta w_{ij} \propto F(pre, post, 3rd\ factor)$$

$$\Delta w_{ij} \propto STDP \times D$$

$$\Delta w_{ij} = H_{ij} \times D$$

Hebb rule

Success signal

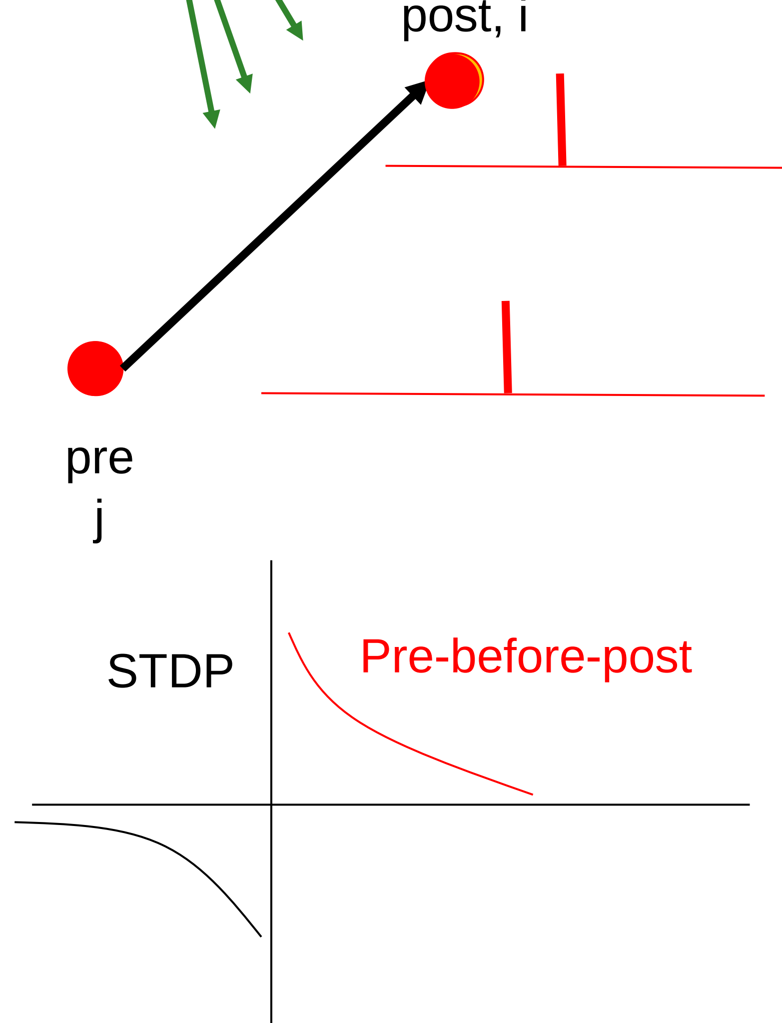

Neuromodulated STDP experiments

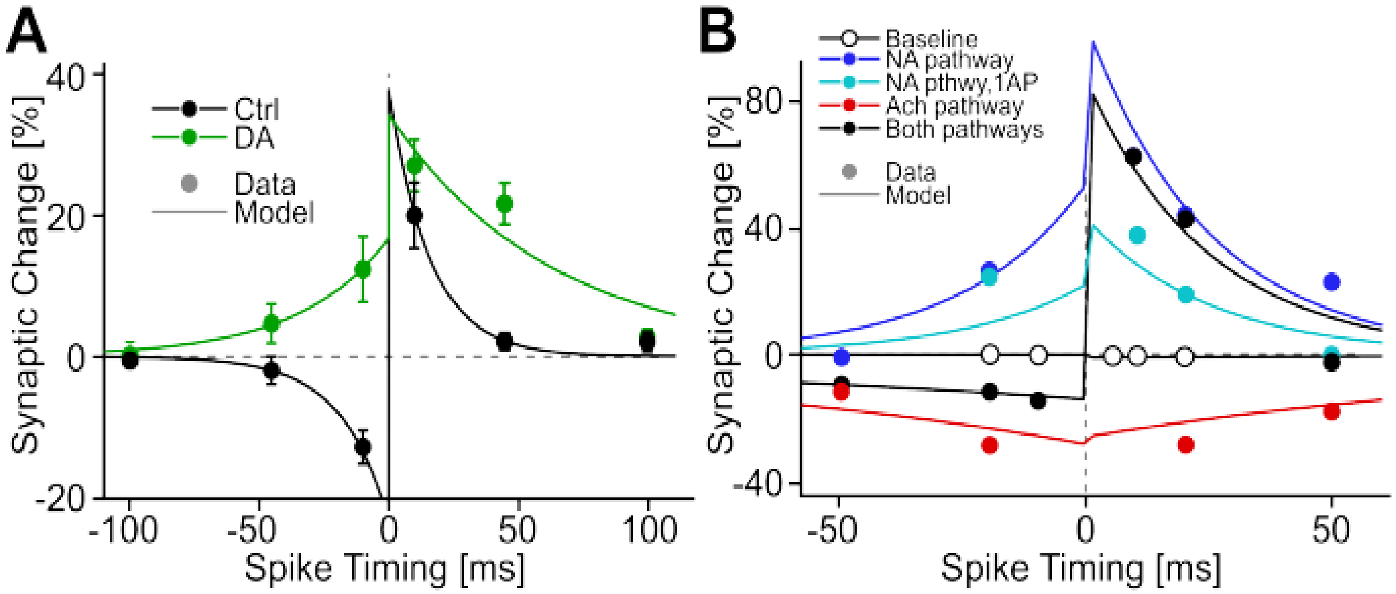

Figure 2: **Fit of experimental neuromodulated STDP data with a simple phenomenological model.** A: Fit of the data shown in Figure 3 in [Zhang et al., 2009]. Dots show the data (same as in Figure 1C), lines show the model fit, with parameters as in Table 2. B: Fit of the data in Figures 2B, 3C and 4B in [Seol et al., 2007]. Dots show the data, lines show the model fit, with parameters as in Table 2. Cyan points refer to a protocol using single postsynaptic action potentials for STDP induction, all the other points correspond to postsynaptic burst induction.

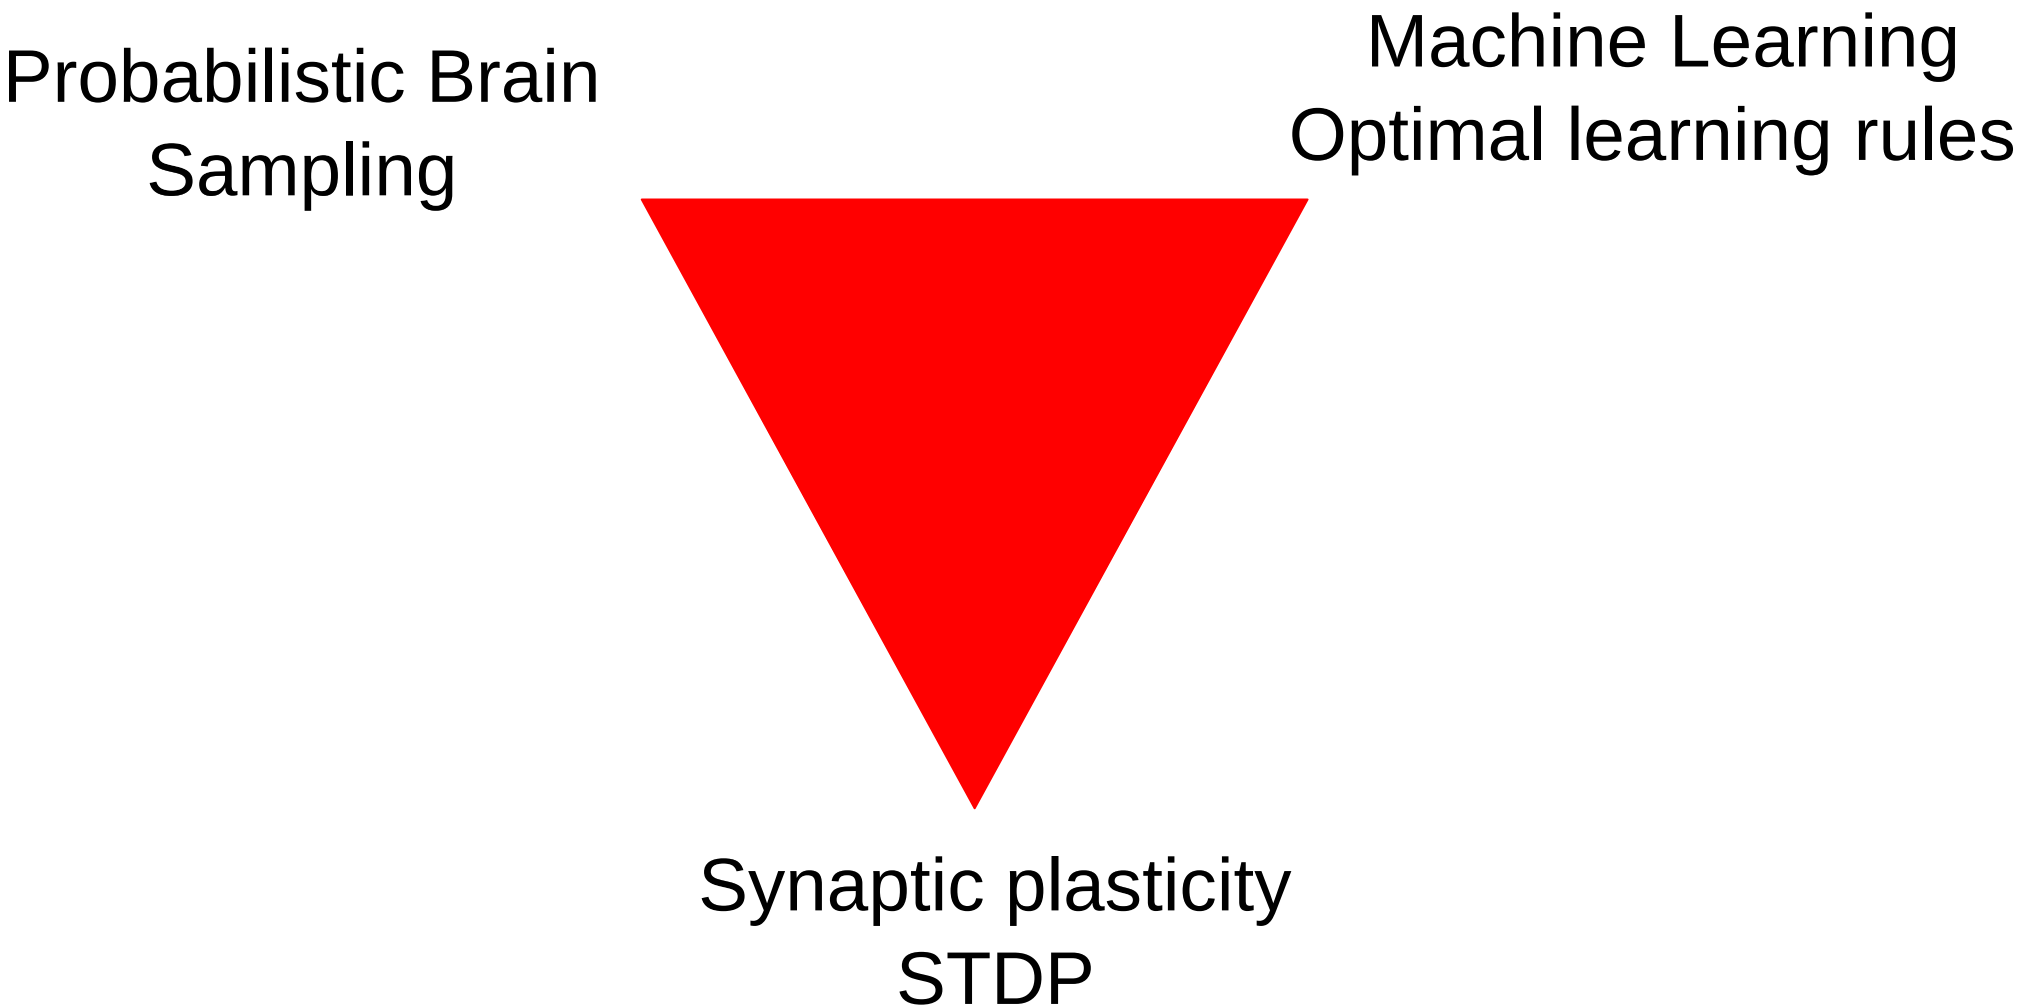

Conclusion:

STDP experiments are unconclusive

Lot's of things possible

Interplay of synaptic plasticity and probabilistic brain theories

- Introduction

- Network and Task

- Math1: Likelihood in Spiking Neurons

- Math 2: Learning rule for network with hidden units

Spiking Neural Network

$\mathcal{V}$

all neurons are
'visible'

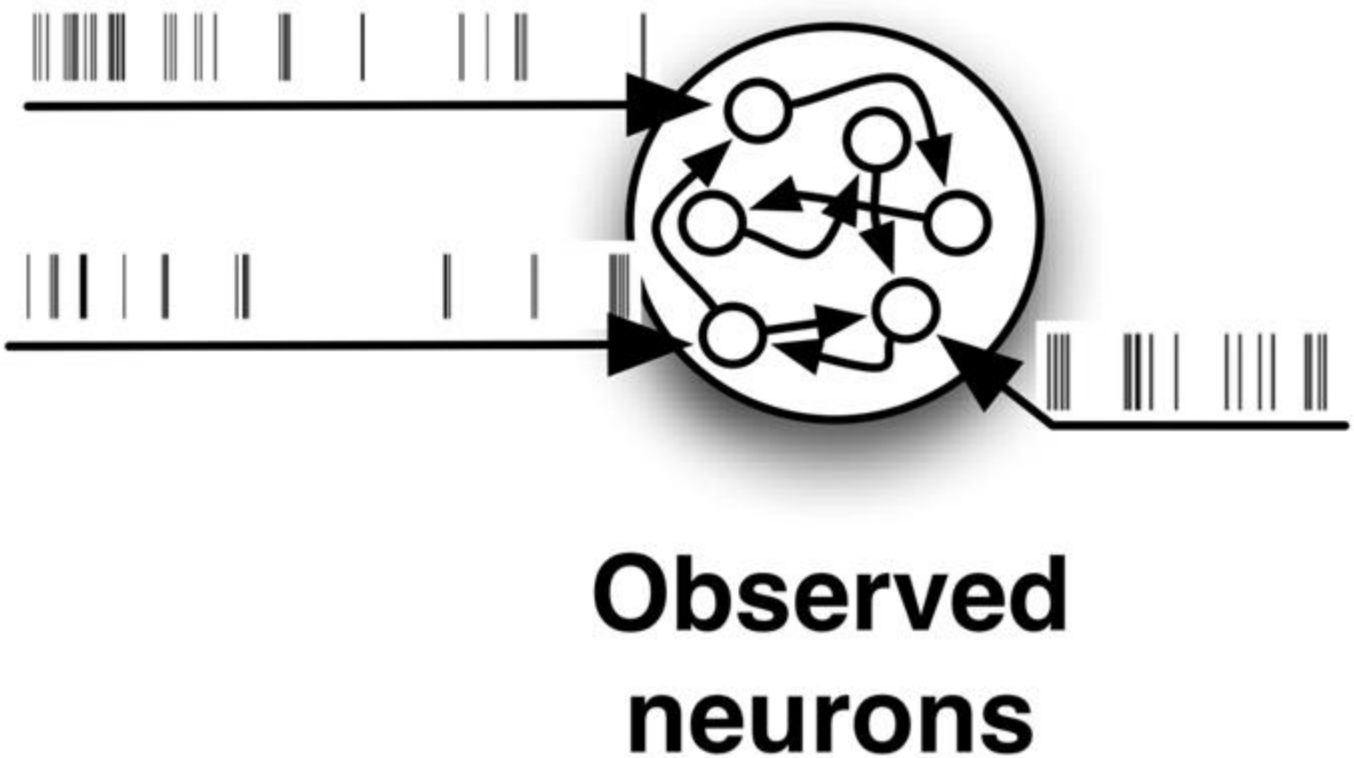

The task:

- ‘slow’ sequence

‘lullaby, children song’

frere Jacques, ….

Task: Sequence learning and Sequence Generation

Target pattern (sample)

30 neurons

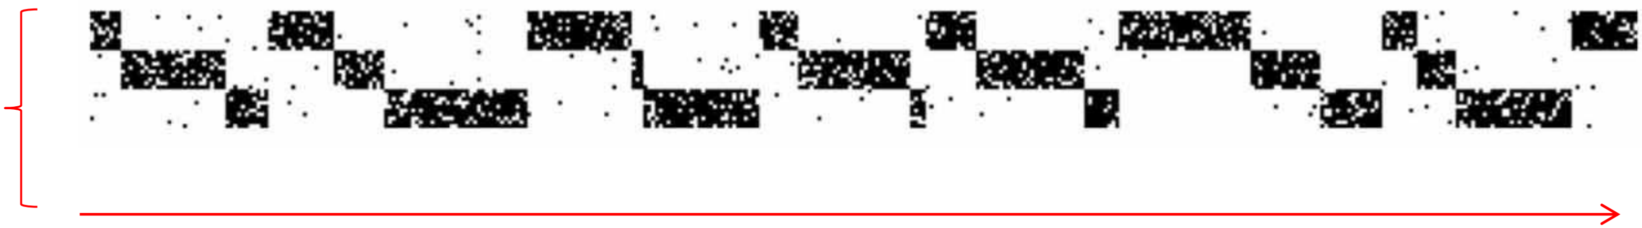

400ms

After learning, sequence generation, network of 30 neurons (all visible)

Integrate-and-fire neurons

Neuronal time constant: 10ms

Duration of each step in sequence: 30ms +/- 20ms

→ Network must keep 'memory'

Rezende and Gerstner,
Frontiers Comp. Neurosci. 2014

Task: Sequence learning and Sequence Generation

Target pattern (sample)

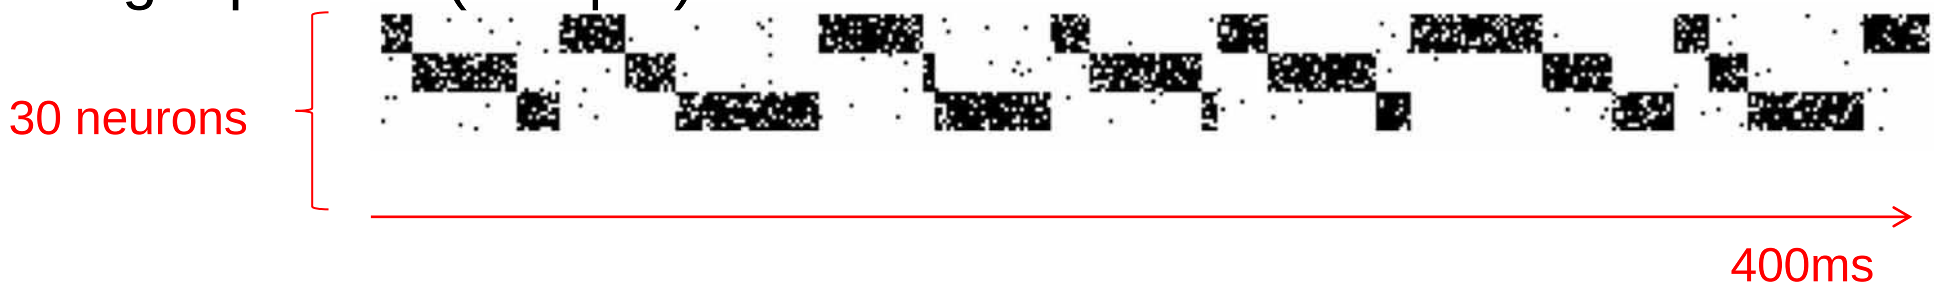

After learning, sequence generation, network with
30 visible neurons
50 hidden neurons

Rezende and Gerstner,
Frontiers Comp. Neurosci. 2014

Hidden neurons

- Memory
- Hidden causes
- Compressed explanation

… all well known in machine learning, Bayes theory, artificial neural networks, deep learning etc

Big question:

-How can we learn the hidden representation?

-Biologically plausible?

Interplay of synaptic plasticity and probabilistic brain theories

- Introduction
- Network and Task
- Math1: Likelihood in Spiking Neurons
- Math 2: Learning rule for network with hidden units
- Surprise

Neuron model: Spike response model with stochastic firing.

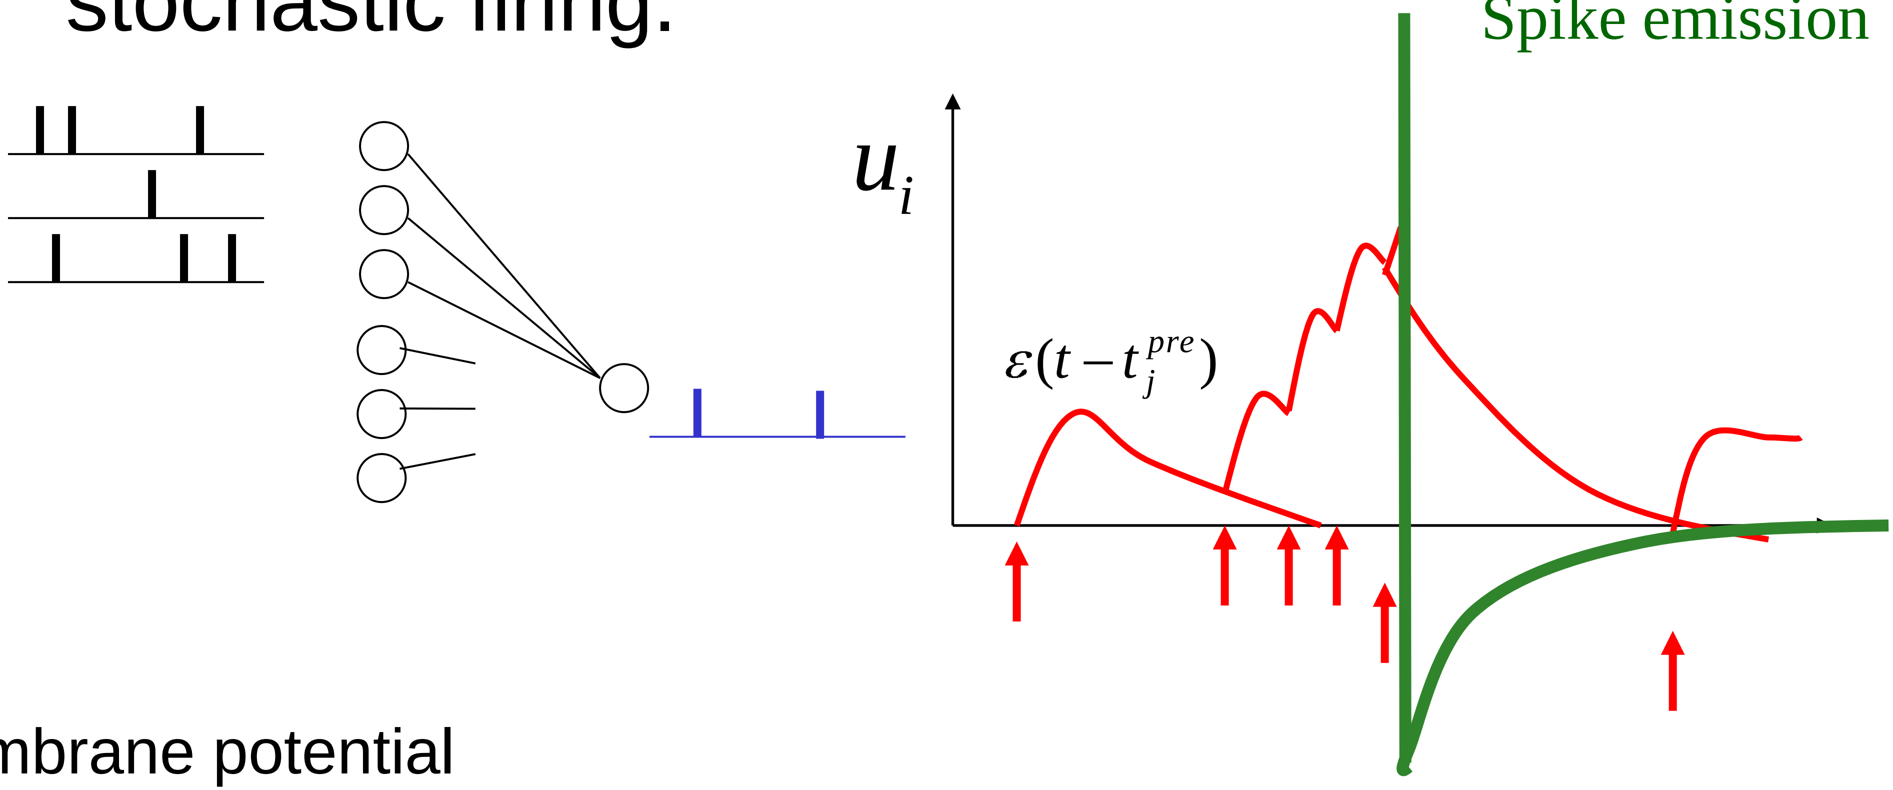

Membrane potential

$$u_i(t\,|\,input) = u_{rest} + \sum_j w_{ij} \sum_{input\ spikes} \varepsilon(t-t_j^f) + \sum_{output\ spikes} \underbrace{\eta(t-t_i^f)}_{reset}$$

Spike generation (probabilistic)

$$\rho(t\,|\,u_i) = \rho(u(t)) \propto \exp(\beta\, u(t))$$

The higher the potential, the more likely the neuron is to fire

Likelihood of a spike train

Spike Response Model with escape noise

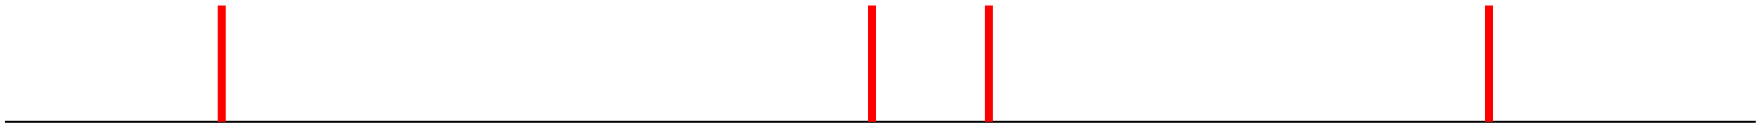

$$L = P(Spiketrain|input) = \Pi_n \rho\left(\hat{t}^n \middle| u(\hat{t}^n)\right) \exp\left[-\int_{\hat{t}^{n-1}}^{\hat{t}^n} \rho\left(t' \middle| u(t')\right) dt'\right]$$

Past spike times, input spikes

known

$$u_i(t) = \sum_{spikes\ of\ i} \eta(t - t_i^f) + \sum_j \sum_{\substack{input \\ spikes}} w_{ij} \varepsilon(t - t_j^f)$$

Maximize likelihood to generate observed spike train:

→ Adjust weights

Pillow et al. 2008, Paninski 2004, Pfister et al. 2006

Derivation of learning rule

- Maximize likelihood that (observed) spikes could have been generated by model

$$\boxed{\frac{d}{dt} w_{ij} = \eta \frac{d}{dw_{ij}} \log \langle L \rangle}$$

Pfister, Barber, Gerster. 2006

spiking neurons

$$\frac{d}{dt} w_{ij}(t) \propto \quad \varepsilon(t - t_j^{pre})[\delta(t - t_i^f) - \rho(u(t))]$$

EPSP

Post

Optimization is Convex

Work of Paninski

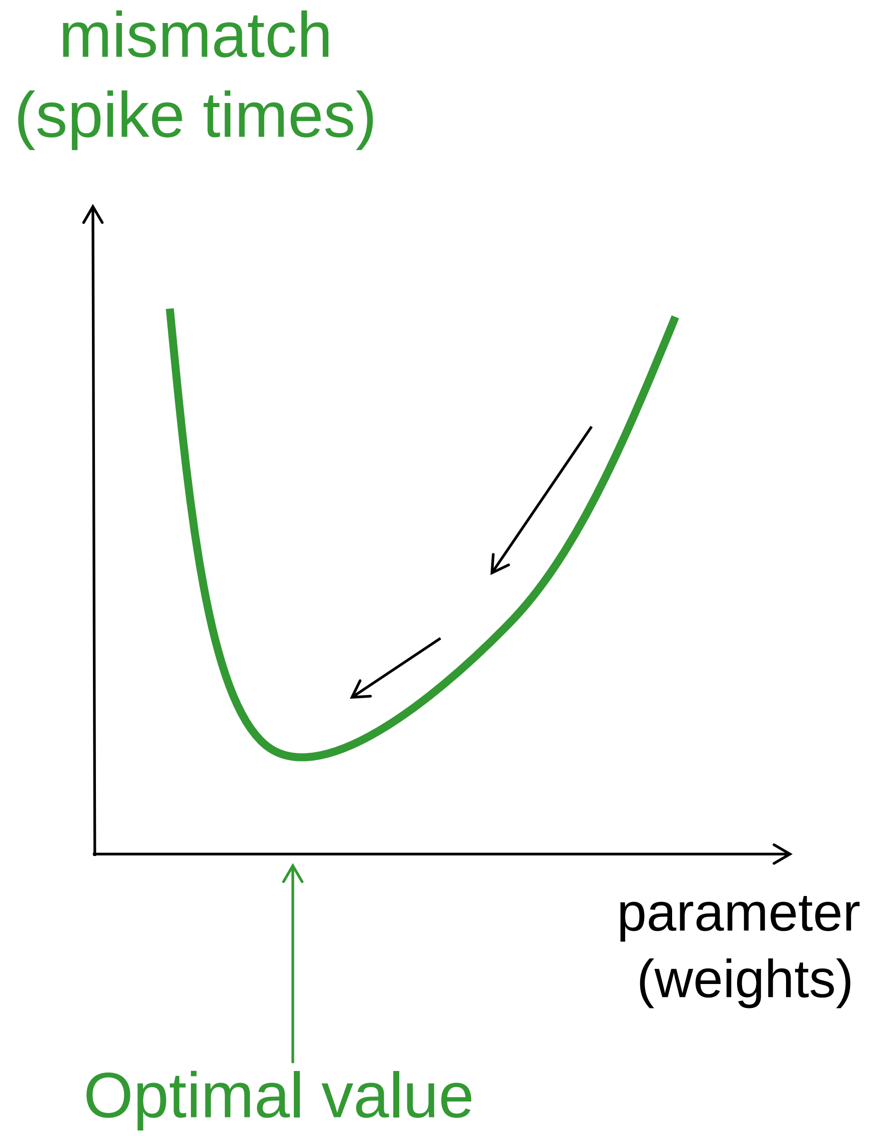

Learning rule for fully observable network

$$\boxed{\frac{d}{dt} w_{ij} = \eta \frac{d}{dw_{ij}} \log \langle L \rangle}$$

$$\frac{d}{dt} w_{ij}(t) \propto \quad \varepsilon(t - t_j^{pre})[\delta(t - t_i^{f}) - \rho(u_i(t))]$$

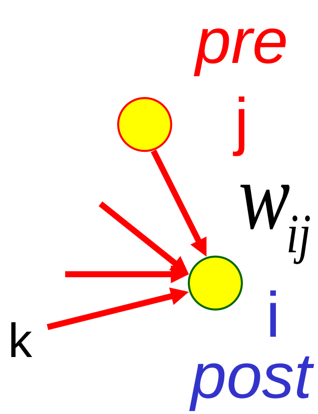

EPSP

spike arrival

Post

Spike time, membrane potential

$$\Delta w_{ij} \propto F(pre, post)$$

Pfister et al. 2006

Interplay of synaptic plasticity and probabilistic brain theories

- Introduction
- Network and Task
- Math1: Likelihood in Spiking Neurons
- Math 2: Learning rule for network with hidden units
- Surprise/novelty

Math 2: Learning rule for network with hidden units

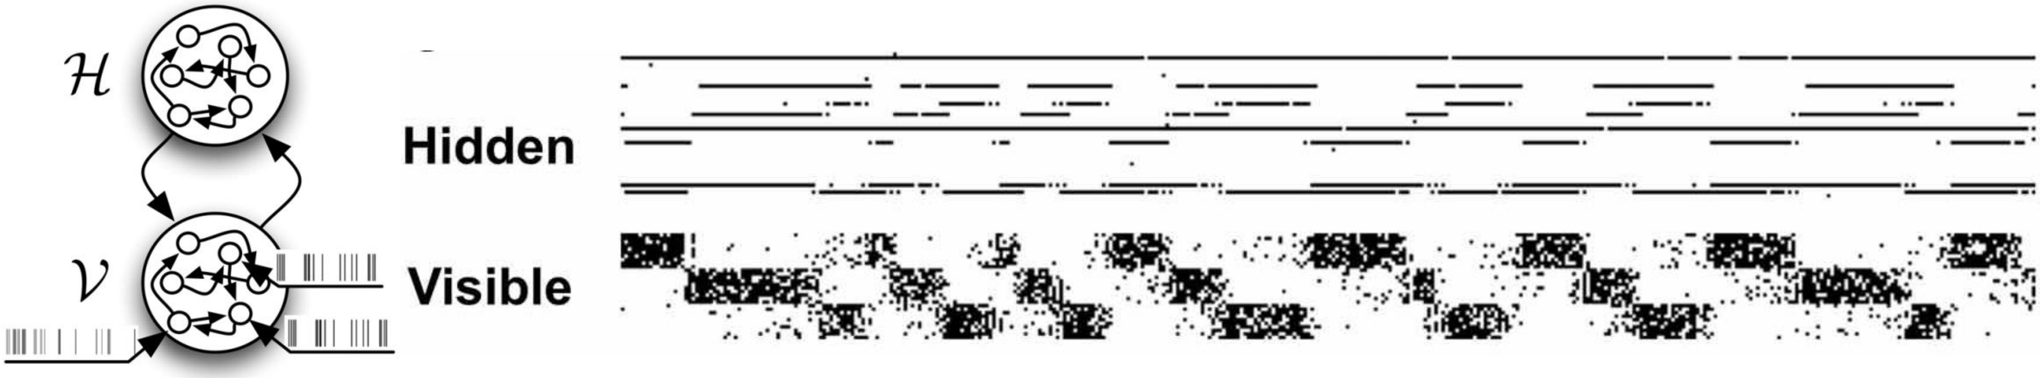

Hidden units

+ memory
+ hidden causes
+ compact representations

- hard to train

Aim: **visible** units

$$L = P(observed\ Spiketrain)$$

$$= P(X_v)$$

$$= \int P(X_v, X_H)\, dX_H$$

all hidden states

Trick: Variational Learning (a.k.a: Free Energy)

e.g. Friston 2005

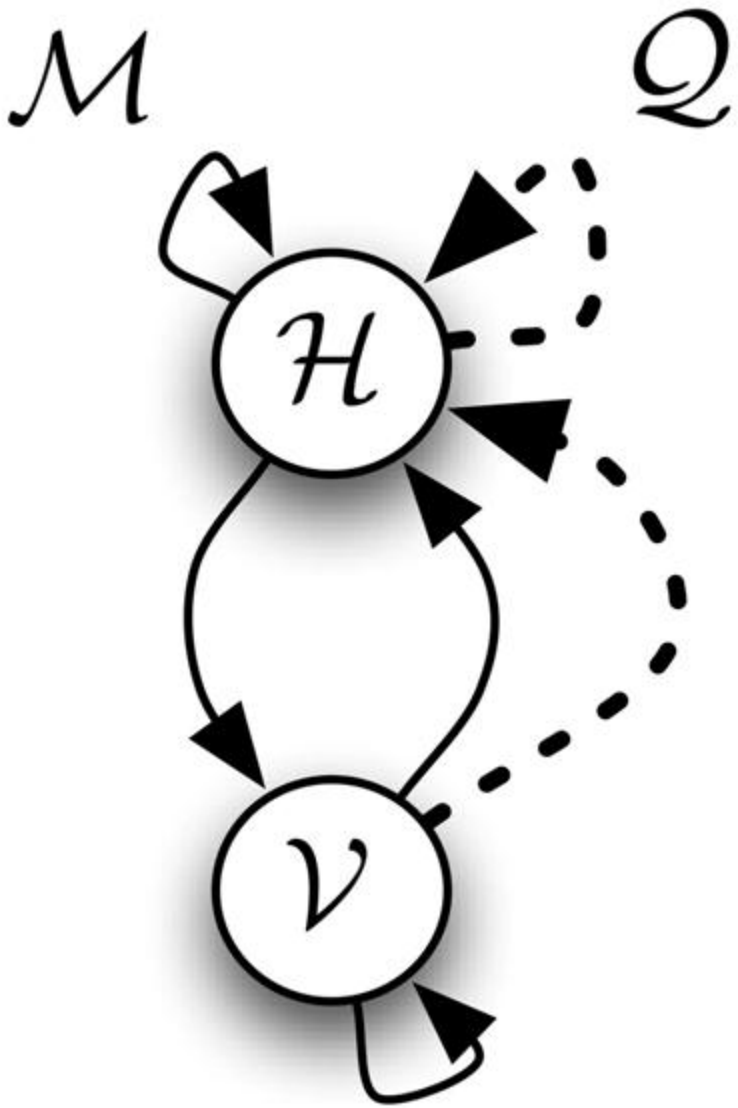

Approximate complicated network *M*, by simpler network *Q*

Minimize KL – divergence

$$KL(q;p) = \int q(X_H|X_v)\log\frac{q(X_H|X_v)}{p(X_H|X_v)}dX_H$$

$$KL(q;p) = F + \log p(X_v)$$

$$F = \left\langle \log q(X_H|X_v) - \log p(X_H|X_v)\right\rangle_{q(X_H|X_v)}$$

average over samples from simple network

Trick: Variational Learning (aka: Free Energy)

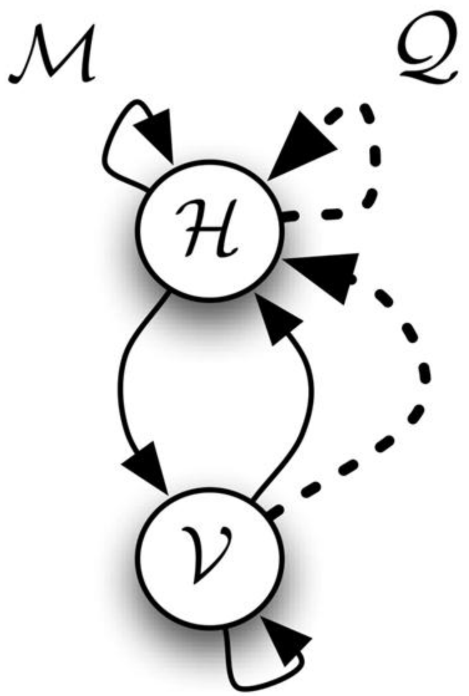

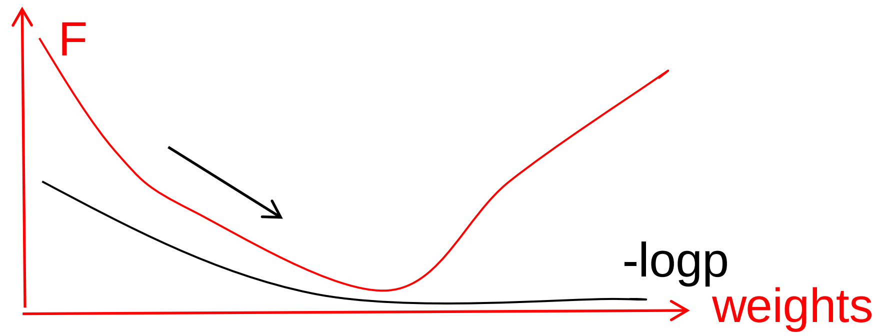

Minimize F→ maximize (upper bound of) log p

$$-\log p(X_v) \leq F$$

$$KL(q;p) = F + \log p(X_v) \quad \geq 0$$

$$F = \left\langle \log q(X_H | X_v) - \log p(X_H | X_v) \right\rangle_{q(X_H | X_v)}$$

average over samples from simple network

e.g., Dayan 2000
Friston, 2005
Friston and Stephan, 2007
Beal and Gahrahmani, 2006

Math 2: Learning rule for network with hidden units

$$\boxed{\frac{d}{dt} w_{ij} = -\eta \frac{d}{dw_{ij}} \log \langle F \rangle}$$

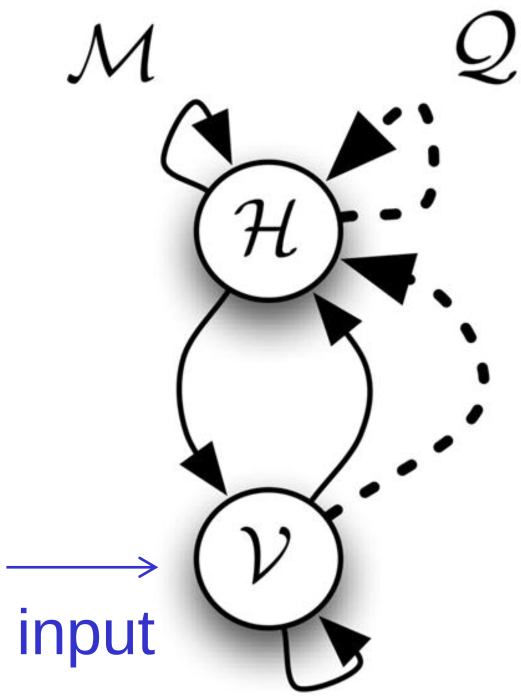

For *Q* network

$$\frac{d}{dt} w_{ij}(t) \propto F\, \varepsilon(t - t_j^{pre})[\delta(t - t_i^{f}) - \rho(u_i(t))]$$

generated by Q-network

For *M* network

$$\frac{d}{dt} w_{ij}(t) \propto \quad \varepsilon(t - t_j^{pre})[\delta(t - t_i^{f}) - \rho(u_i(t))]$$

→ M network takes Q network as teacher

Biological implementation of learning rule

novelty/surprise

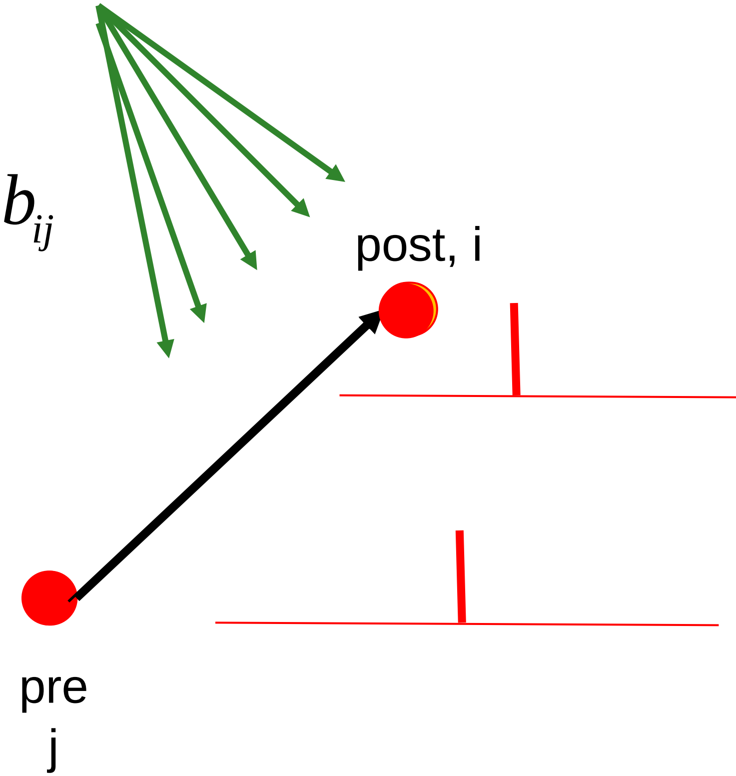

Hebbian learning leaves trace:

$$\frac{d}{dt} Hebb_{ij}(t) = \varepsilon(t - t_j^{pre})[\delta(t - t_i^{f}) - \rho(u_i(t))] - Hebb_{ij}$$

$$\frac{d}{dt} w_{ij}(t) = \eta \cdot Nov \cdot Hebb_{ij}$$

novelty/surprise

$$Nov = \hat{F} - \bar{F}$$

online estimate - running average
of
free energy

weights updated when we are more surprised then normally

Interplay of synaptic plasticity and probabilistic brain theories

- Introduction
- Network and Task
- Math1: Likelihood in Spiking Neurons
- Math 2: Learning rule for network with hidden units
- Surprise/novelty

Surprise/novelty

Neuromodulator Ach

See: Gu 2002, Ranganath and Rainer 2003, Yu and Dayan, 2005

Neuromodulators act on:
- Activity of neurons
- Synaptic plasticity

Our proposition: neuromodulator signal

Novelty = surprise – expected surprise

compare: reinforcement learning

Success = reward – expected reward

Surprise/novelty

Maze of 16
MNIST pixel patterns
28x28 visible neurons
30 hidden neurons,
Actions are random,
every second

Expected surprise (free energy)

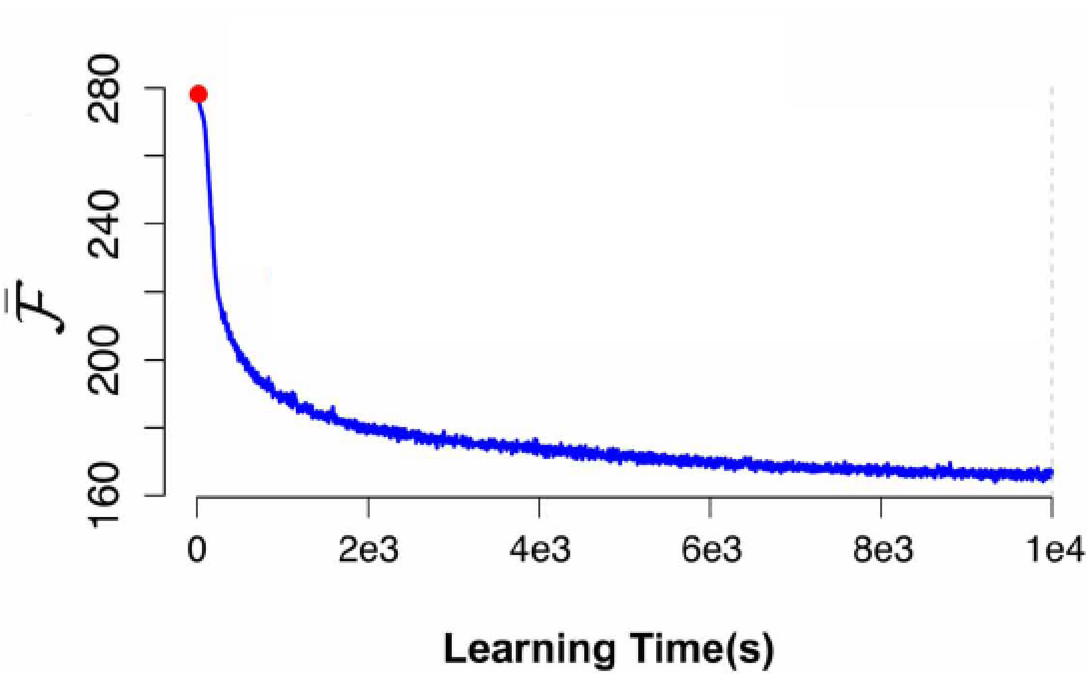

novelty $Nov = \hat{F} - \bar{F}$

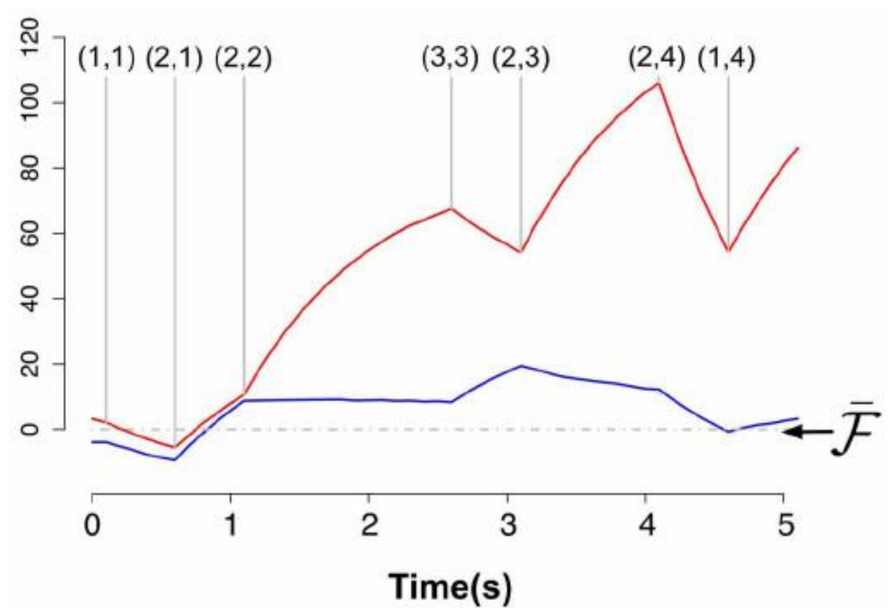

Free Energy 'measures' mismatch of model to data

$$F \geq -\log p(X_v)$$

$\hat{F} =$ Online-single-sample estimate of *F*

$\hat{F}$ 'measures' surprise of present input

Our proposition: neuromodulator signal

Novelty = surprise – expected surprise

$$Nov = \hat{F} - \bar{F}$$

weights updated
when we are more
surprised then normally

Conclusions

$$\Delta w_{ij} \propto F(pre, post, 3rd\ \ factor)$$

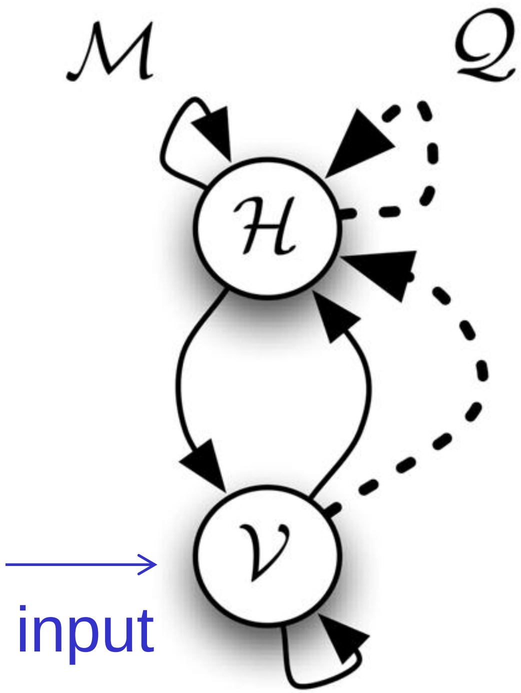

$$Nov = \hat{F} - \bar{F}$$

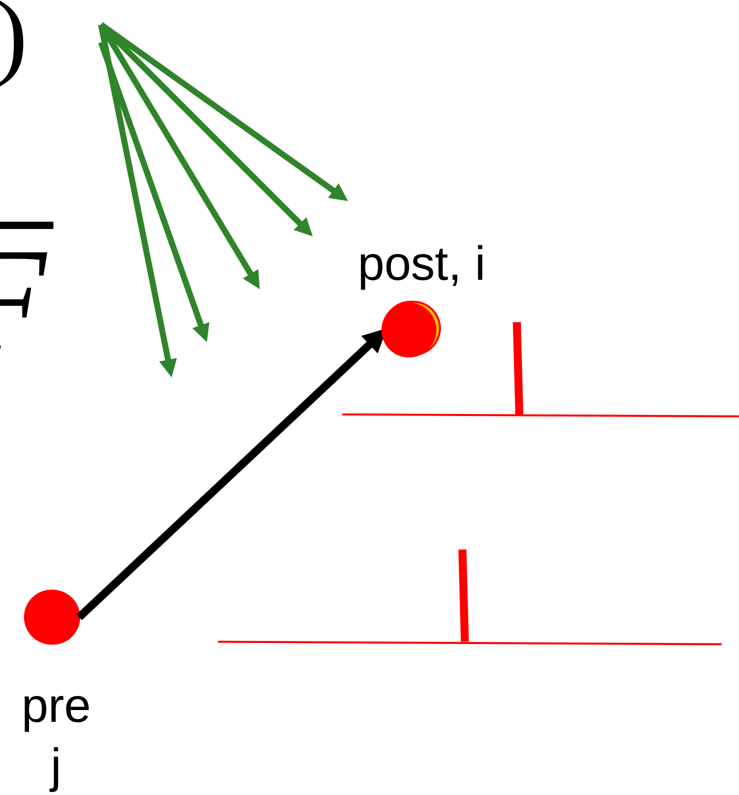

-Hidden neurons/network structure to form memories

-Hebb-rule/STPD for feedback connections

-3-factor rule with ‘novelty’ for feedforward connections

The end

Thanks to
Danilo REZENDE
Daan WIERSTRA
Johanni BREA

Jimenez-Rezende, Wierstra, Gerstner, NIPS 2011
Brea, Senn, Pfister, J. Neuroscience 2013
Rezende and Gerstner, Frontiers Comp. Neurosci. 2014

Classification

R-max Xie&Seung 2004, Pfister et al. 2006, Florian 2007, …

$$\dot{w} = R \times \left(H(\mathrm{pre}, \mathrm{post}) - \langle H(\mathrm{pre}, \mathrm{post}) | \mathrm{pre} \rangle \right) \; ,$$

$$\longrightarrow \qquad \langle \dot{w} \rangle = \mathrm{Cov}(R, H(\mathrm{pre}, \mathrm{post})) \, .$$

R-STDP

Florian 2007, Farries&Fairhall 2008, Legenstein 2008, …

$$\dot{w} = (R - \langle R | \mathrm{pre} \rangle) \times H(\mathrm{pre}, \mathrm{post}) \, .$$

$$\longrightarrow \qquad \langle \dot{w} \rangle = \mathrm{Cov}(R, H(\mathrm{pre}, \mathrm{post})) \, .$$

R-STDP with gating effect Izhikevich 2007

$$\langle \dot{w} \rangle = \mathrm{Cov}(R, H(\mathrm{pre}, \mathrm{post})) + \langle R \rangle \langle H(\mathrm{pre}, \mathrm{post}) \rangle \, .$$

TD-STDP Fremaux et al. 2013, to appear, PLOS Comput. Biol.

$$\dot{w} = \delta \times H(\mathrm{pre}, \mathrm{post}) \, .$$

Classification

$$\dot{w} = M \times H(\text{pre}, \text{post})$$

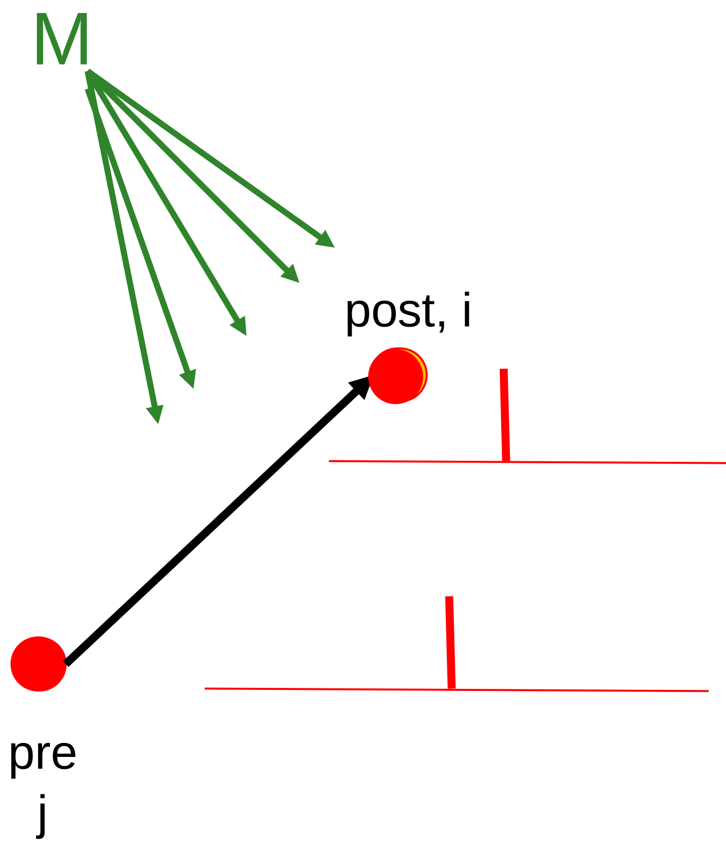

$$M = \begin{cases} R - \langle R \rangle & \longrightarrow \quad \text{covariance-rule} \\ \delta^{\text{TD}} & \longrightarrow \quad \text{TD learning} \\ R & \longrightarrow \quad \text{gated Hebbian learning} \\ S & \longrightarrow \quad \text{surprise/novelty-modulated STDP} \\ const & \longrightarrow \quad \text{non-modulated STDP}\,, \end{cases}$$