Making Cell Assemblies:
What can we learn about plasticity from spiking neural network models?

Friedemann Zenke

EP Cognitive and Behavioural Neuroscience Seminar

Tuesday, 17 October 2017

https://fzenke.net

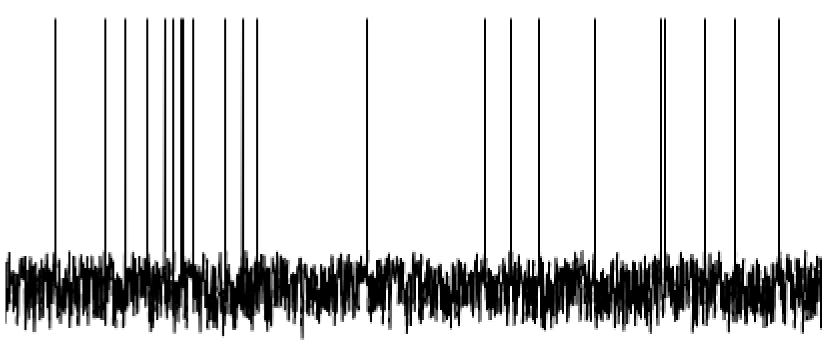

Outline

- Part 1: Cell assemblies and “The temporal paradox of Hebbian and homeostatic plasticity”
- Part 2: Supervised learning in spiking neural networks starting from a cost function

Network activity in an unstructured balanced network

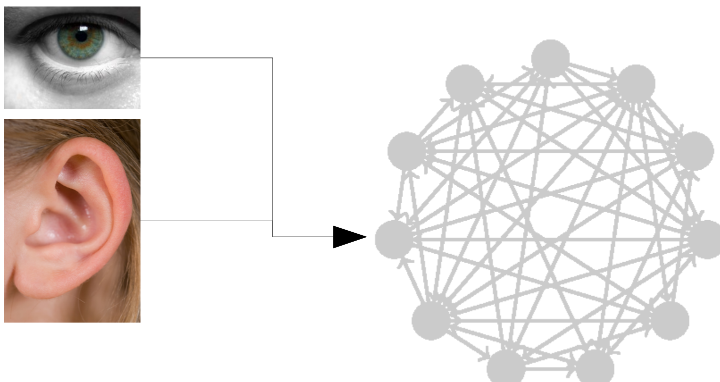

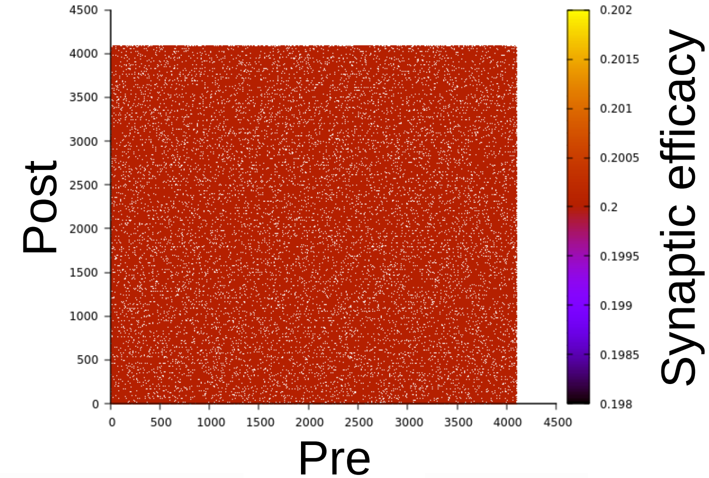

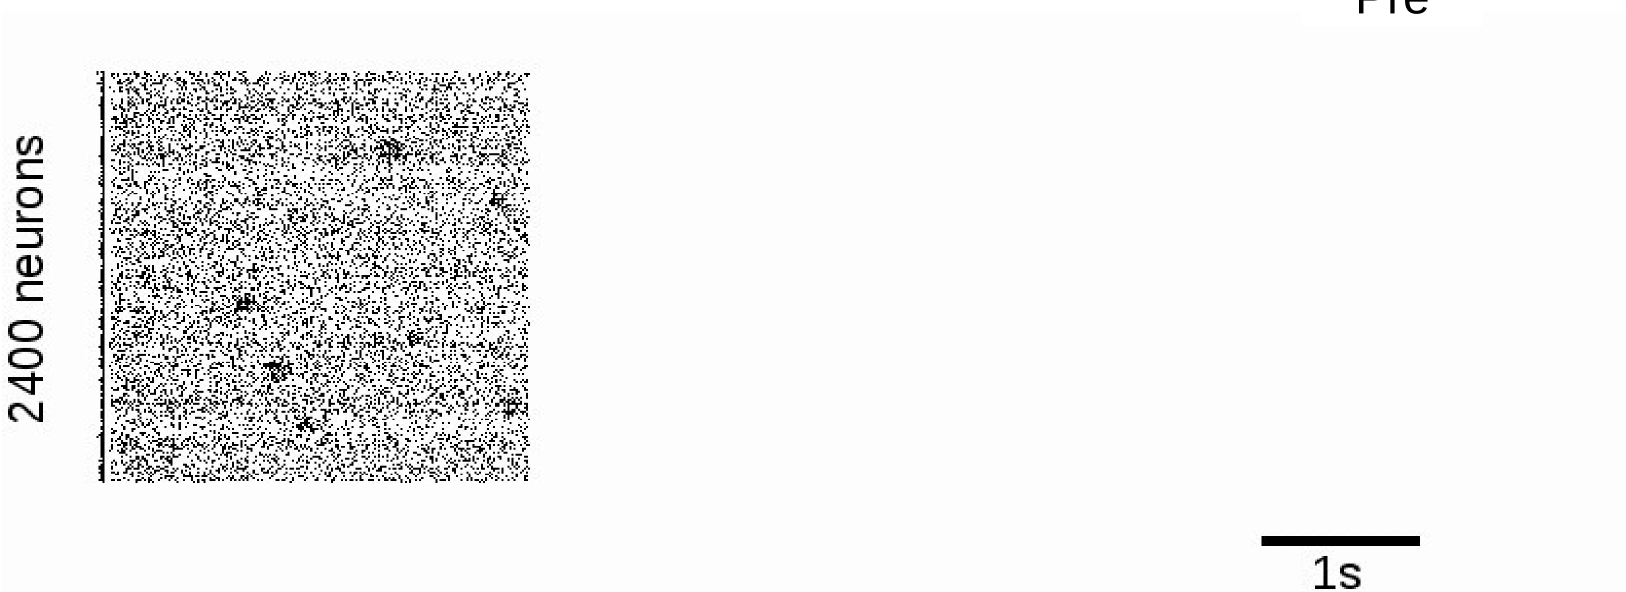

Outline

- Part 1: Cell assemblies and “The temporal paradox of Hebbian and homeostatic plasticity”
- Part 2: Supervised learning in spiking neural networks starting from a cost function

Network activity in a balanced network with cell assemblies

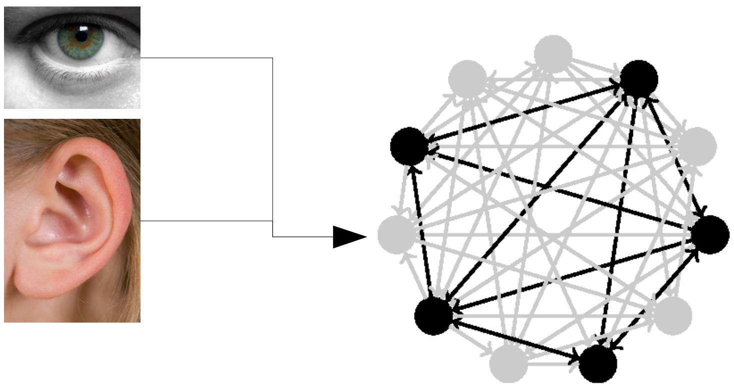

Functions
- Associative
- Amplify
- Selective
- Working memory

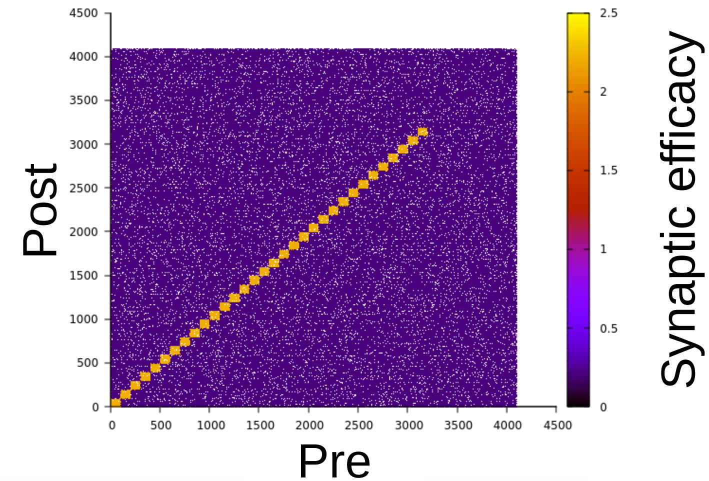

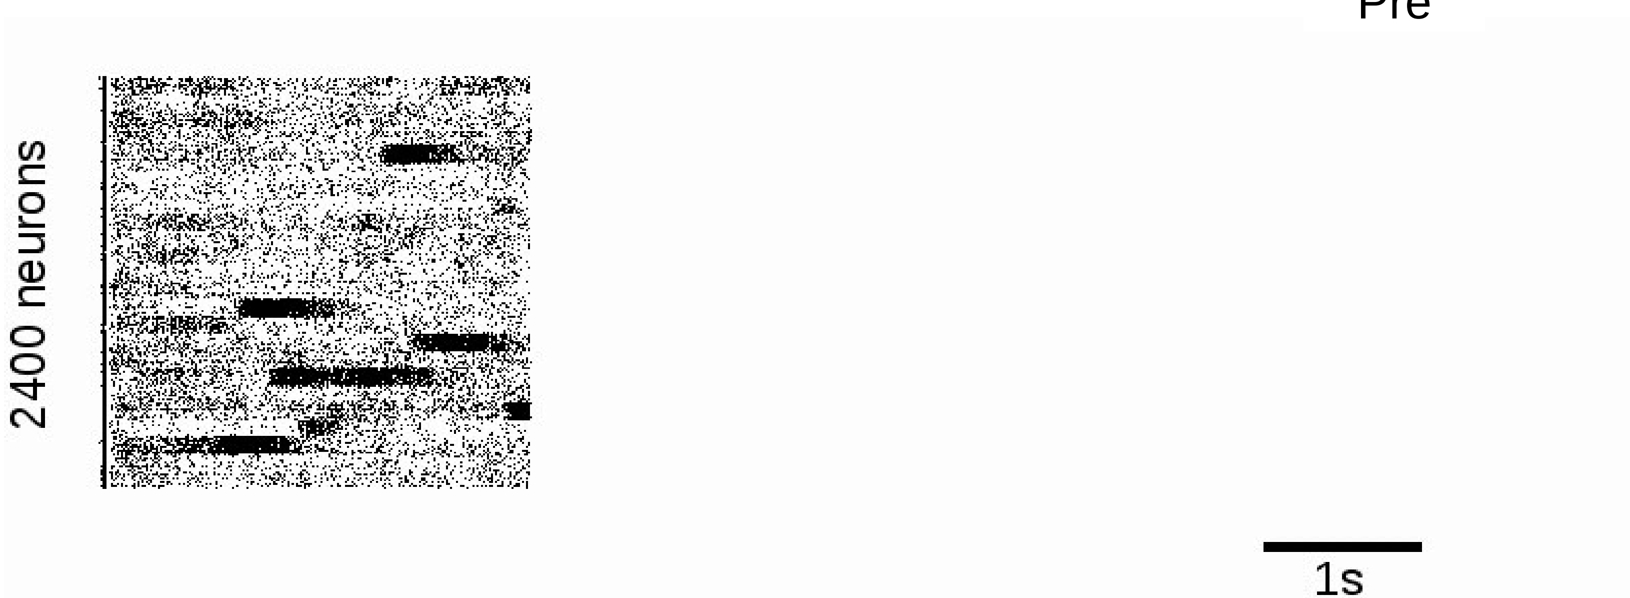

Cell assemblies are formed through Hebbian plasticity?

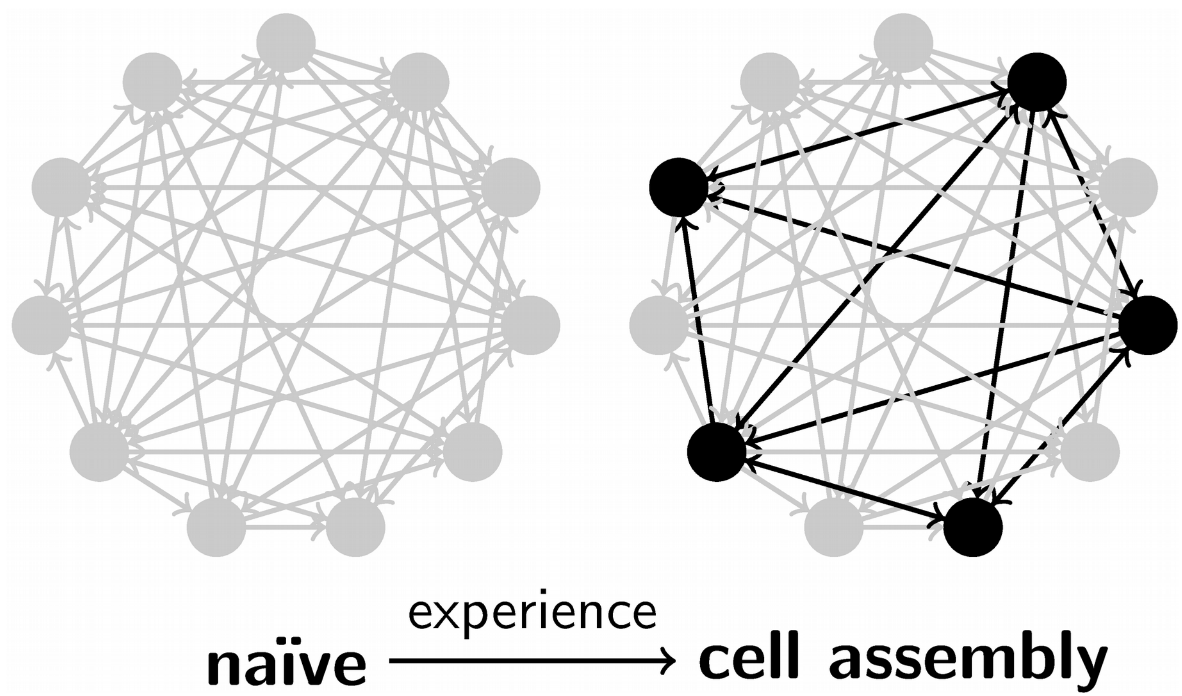

Hebb 1949

Plasticity model: Triplet STDP

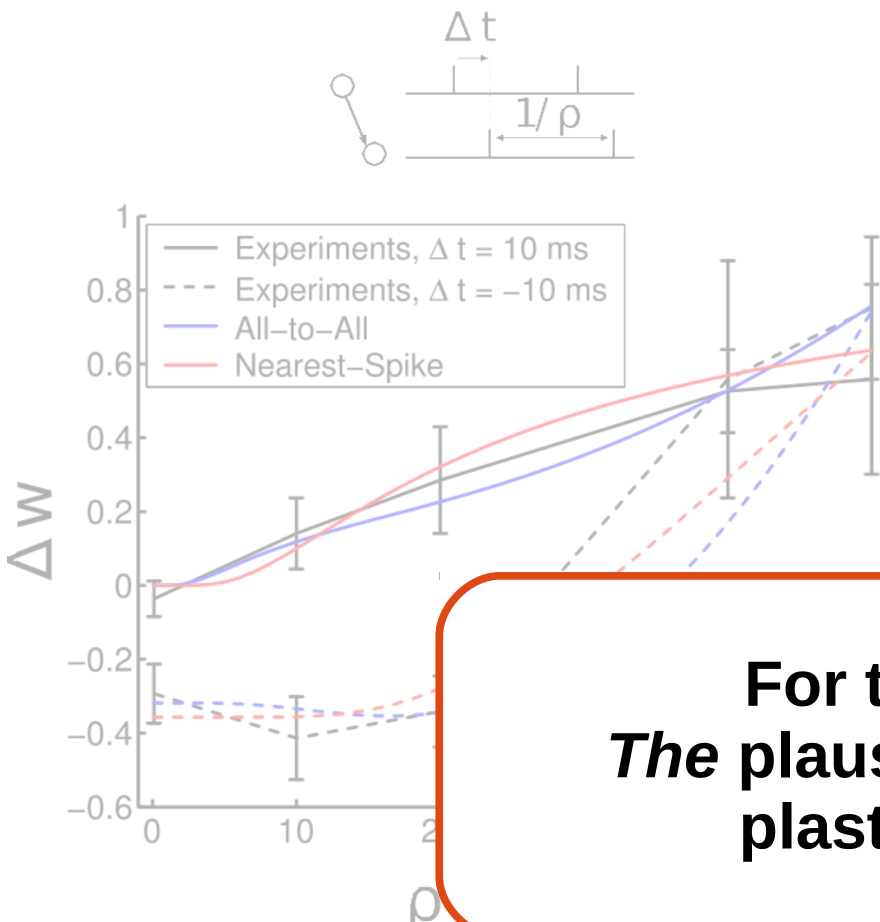

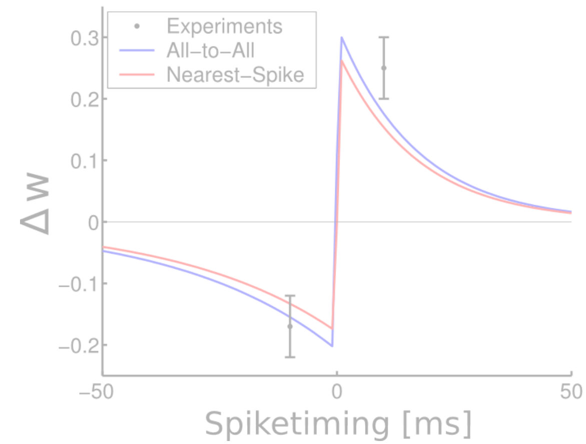

- Reproduces experiments
- riplet protocols
- ence
- er set

For this talk:
***The* plausible Hebbian plasticity rule**

Pfister & Gerstner (2006)
Experiments: Sjöström et al. (2001)

Cell assemblies and Hebbian plasticity → Positive feedback loop

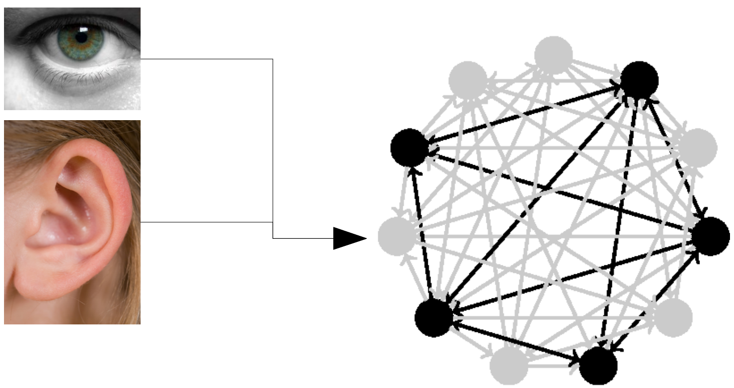

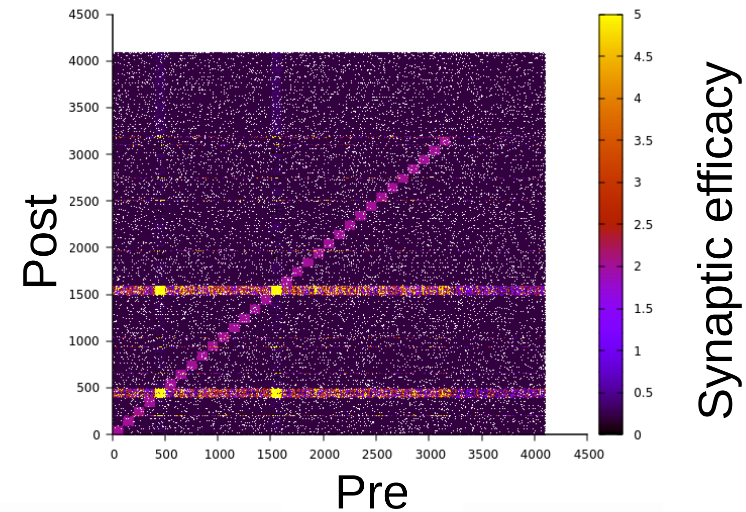

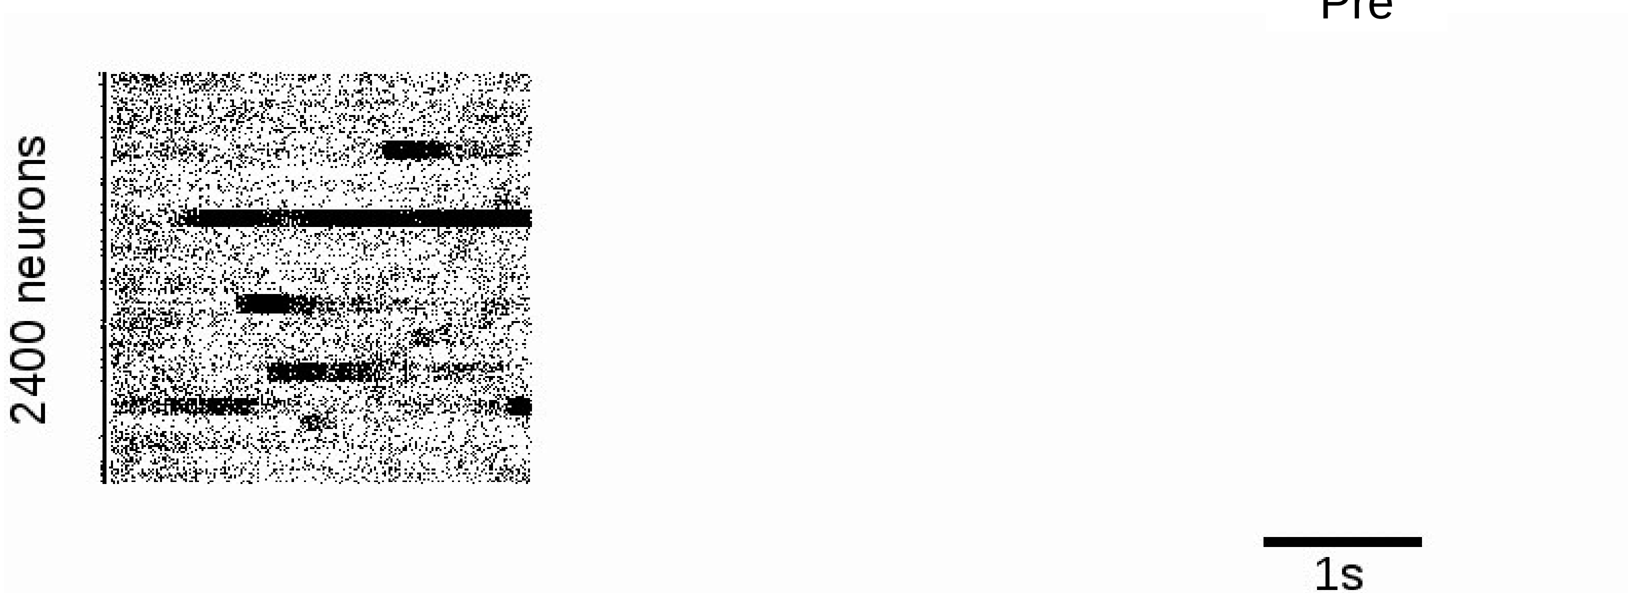

Zenke et al. (2015)

Bistable activity states are unstable in conjunction with triplet STDP

No plasticity

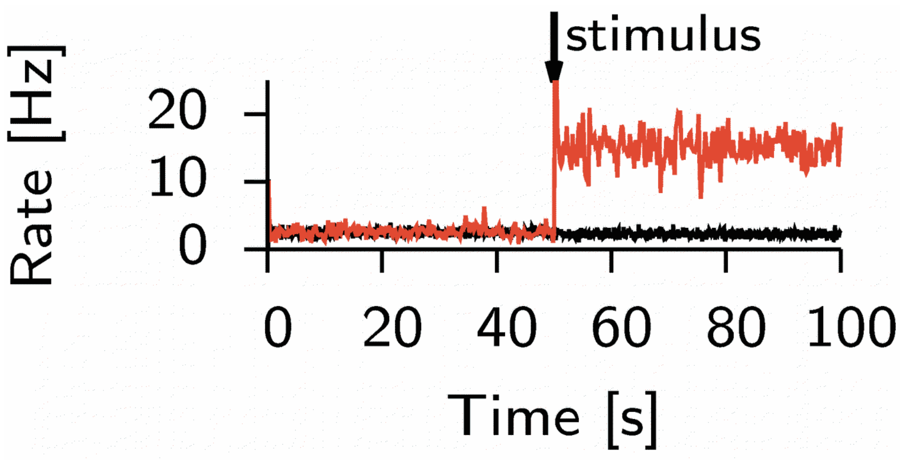

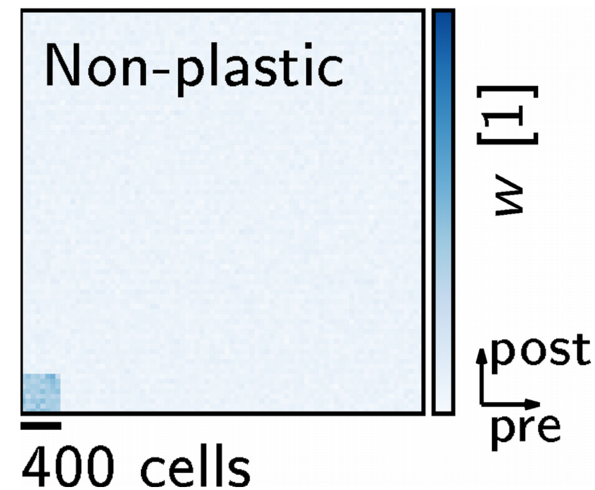

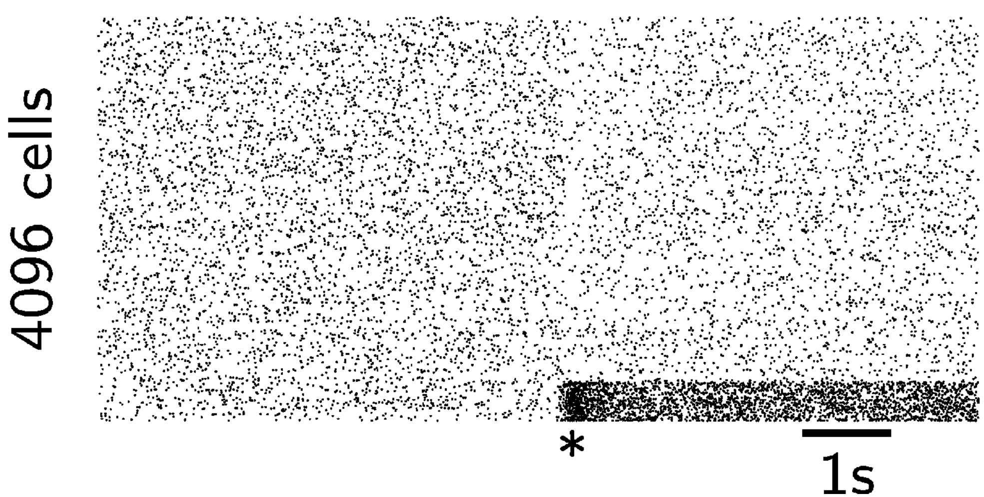

Background neurons

Assembly neurons

Zenke et al. (2015)

Bistable activity states are unstable in conjunction with triplet STDP

No plasticity

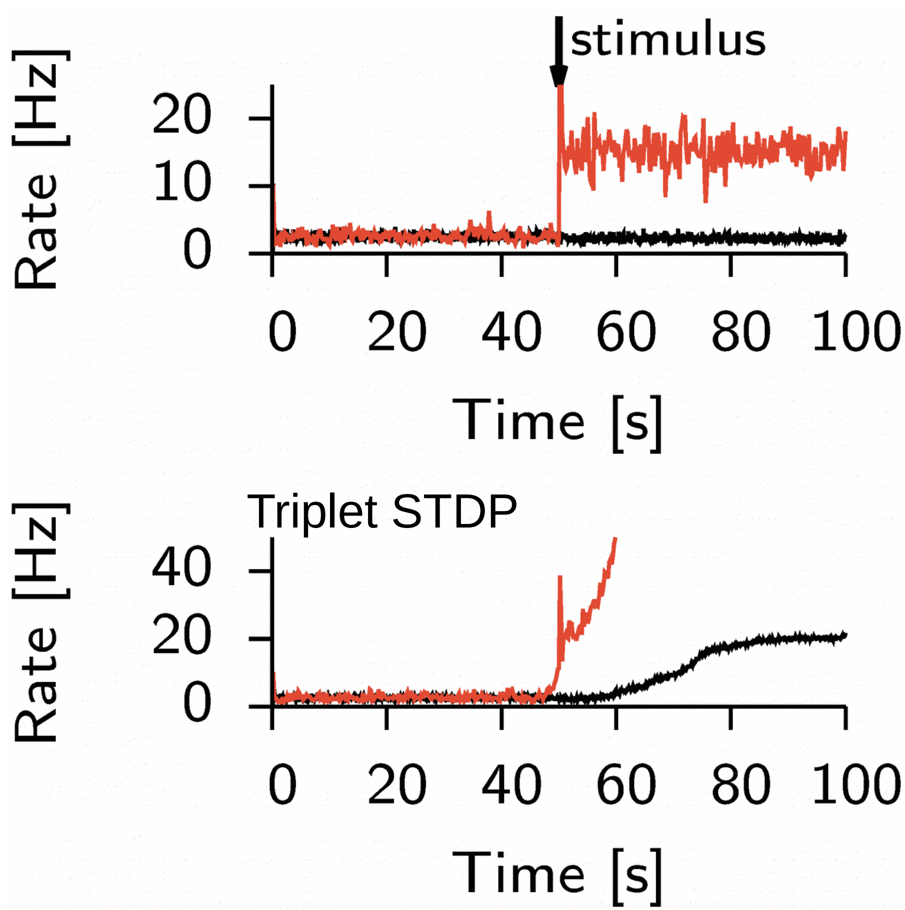

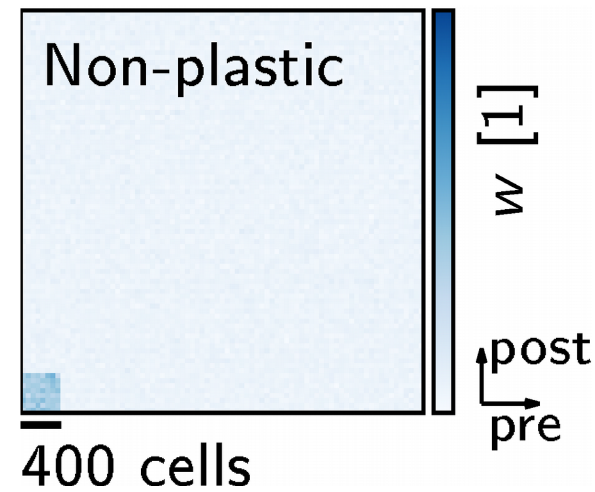

Zenke et al. (2015)

Bistable activity states are unstable in conjunction with triplet STDP

No plasticity

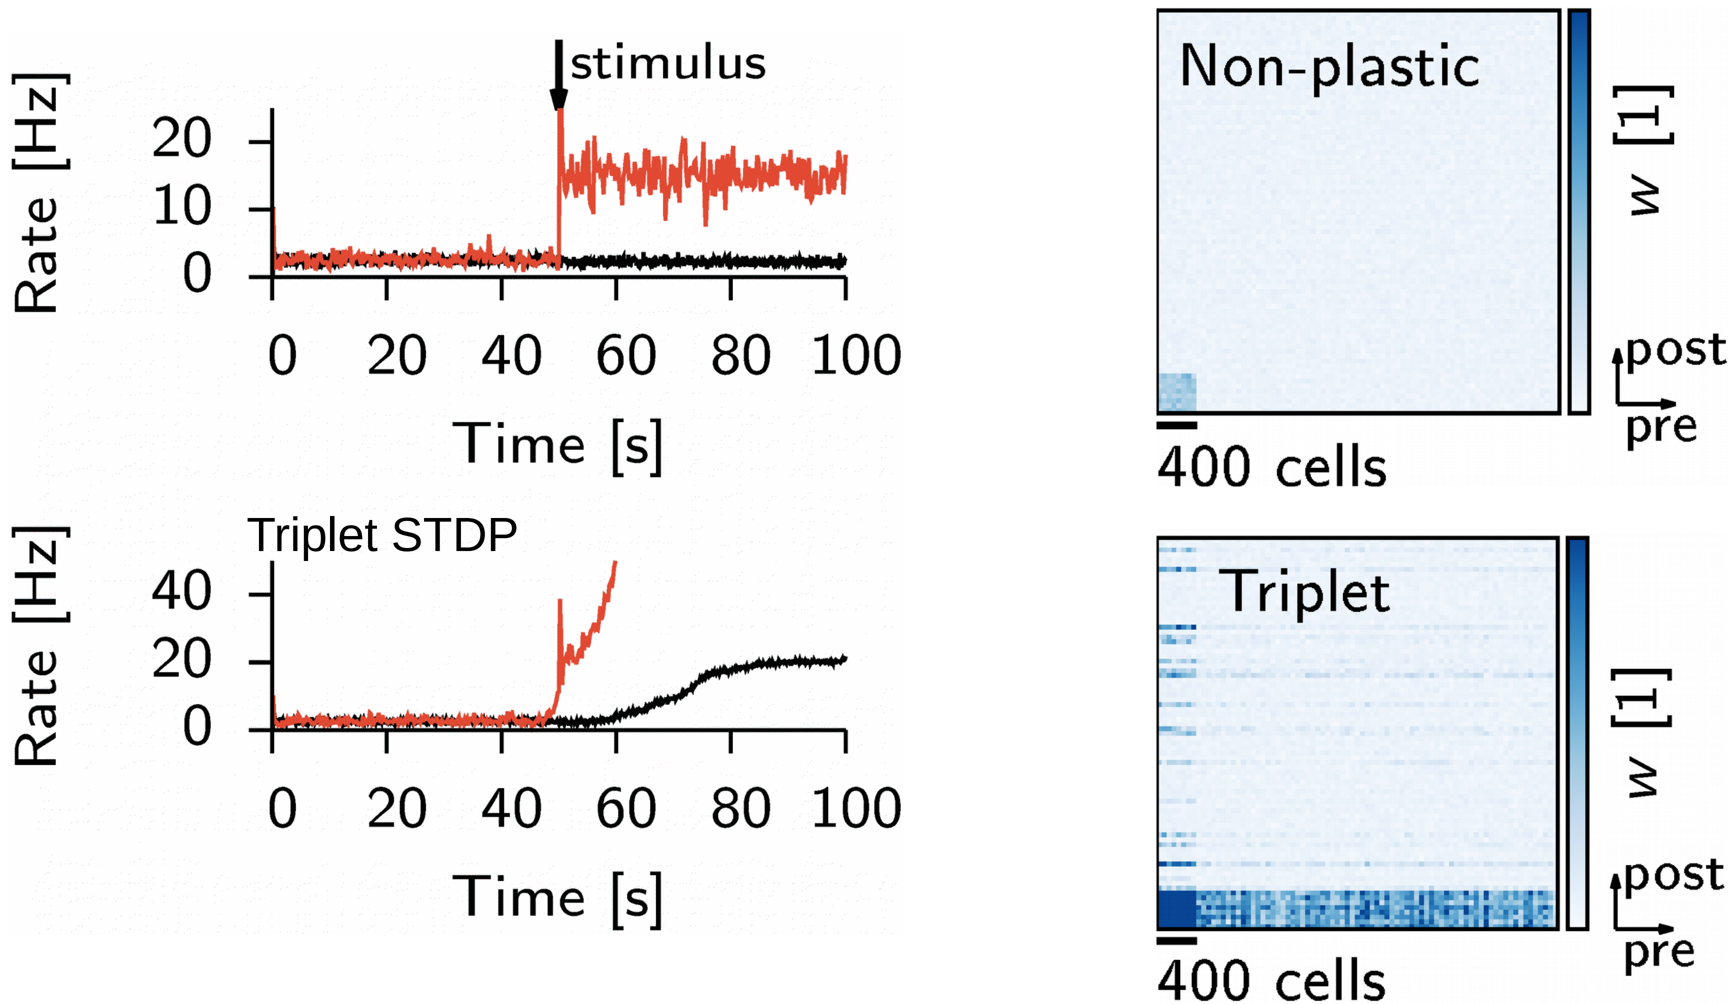

Triplet STDP

Zenke et al. (2015)

Bistable activity states are unstable in conjunction with triplet STDP

No plasticity

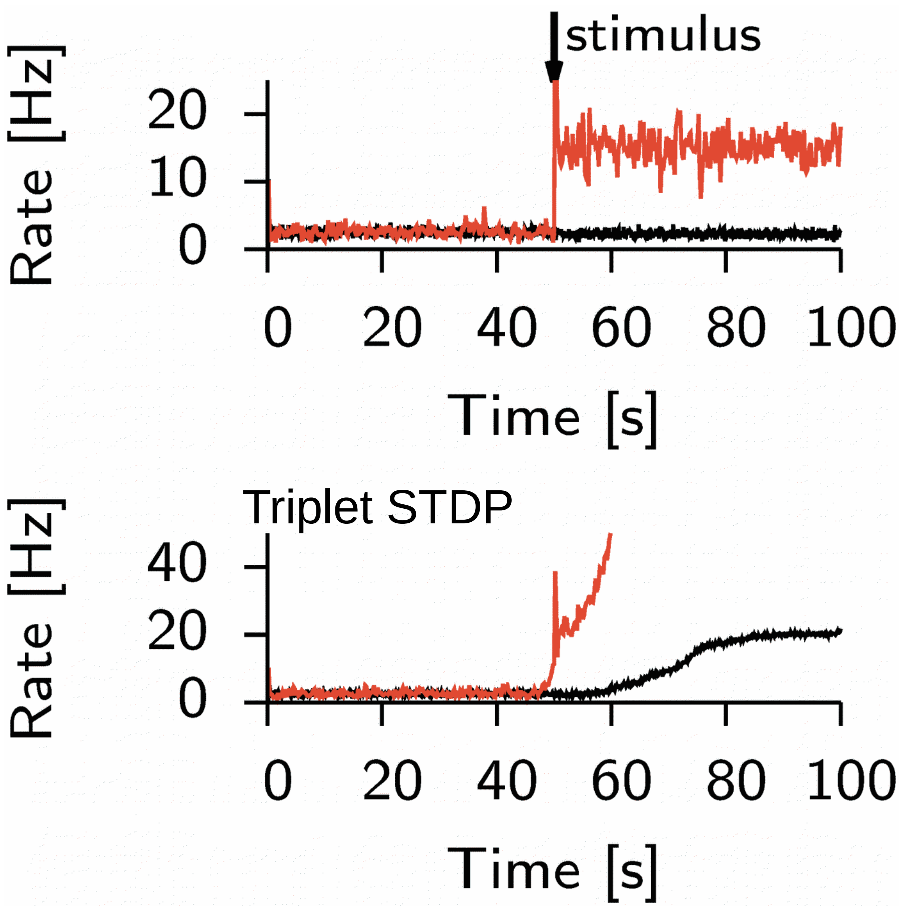

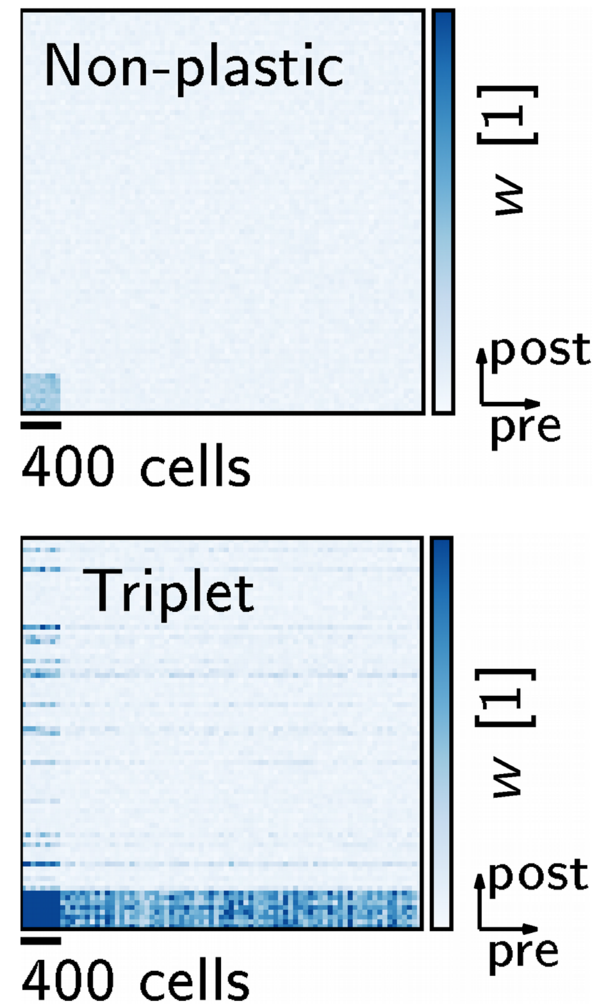

Zenke et al. (2015)

Bistable activity states are unstable in conjunction with triplet STDP

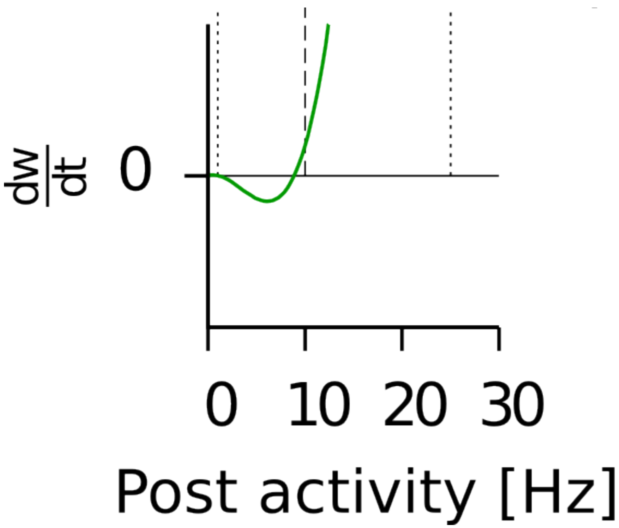

Three important facts

- Run-away potentiation introduces cross-talk and destroys memories
- Timescale of run-away
- Hebbian plasticity is unstable
  → Need "homeostasis"

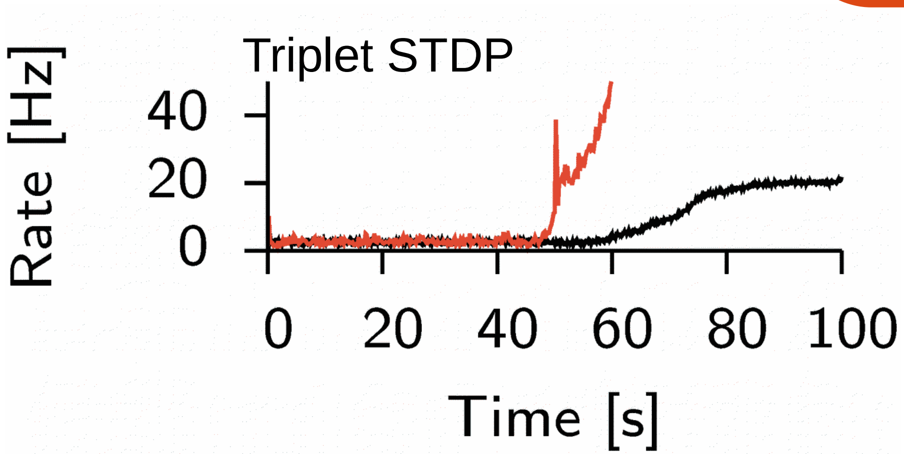

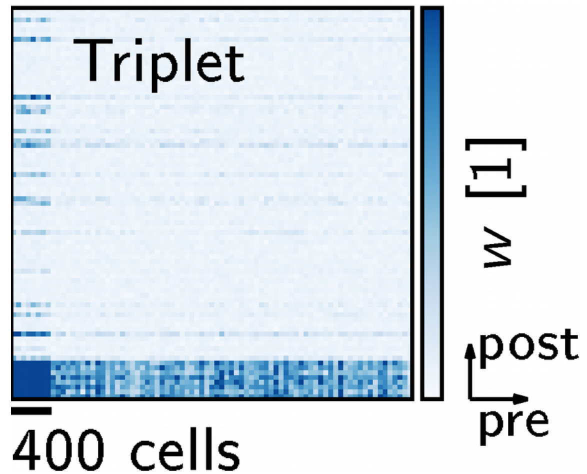

Zenke et al. (2015)

The temporal paradox of Hebbian and homeostatic plasticity

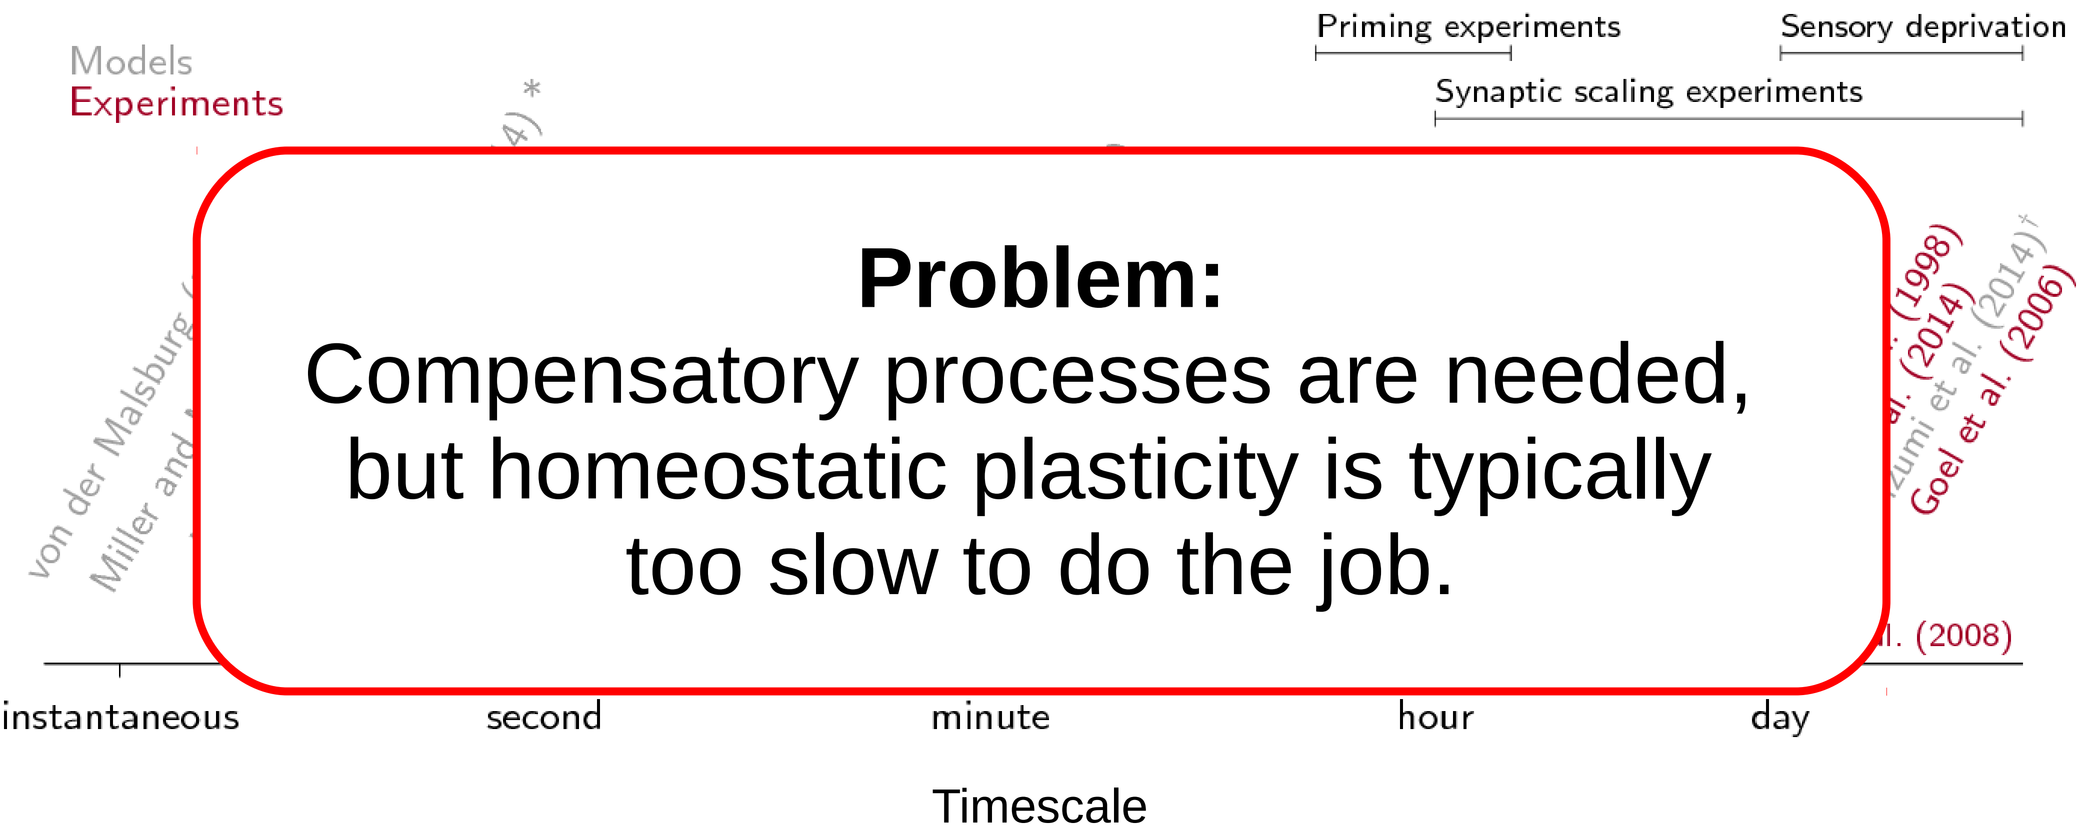

Problem:
Compensatory processes are needed, but homeostatic plasticity is typically too slow to do the job.

Zenke & Gerstner (2017) Phil. Trans. R. Soc. B

Control theory: Timescales of stabilizing processes needs to be matched to unstable system

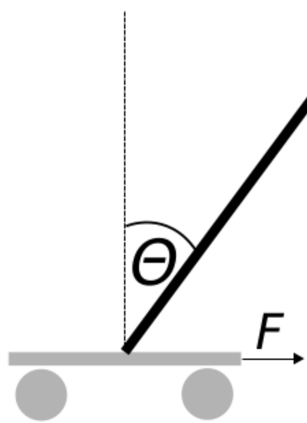

Zenke et al. (2017)

Purely Hebbian learning with slow homeostasis is unstable

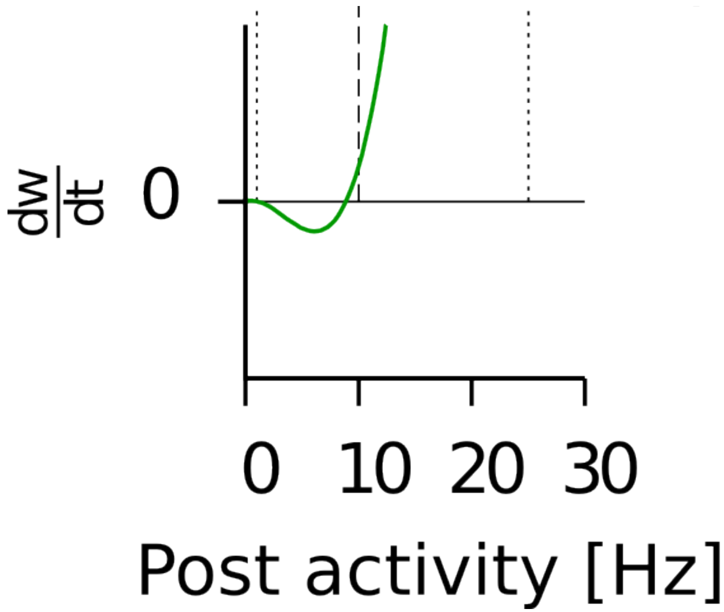

For any presynaptic activity > 0

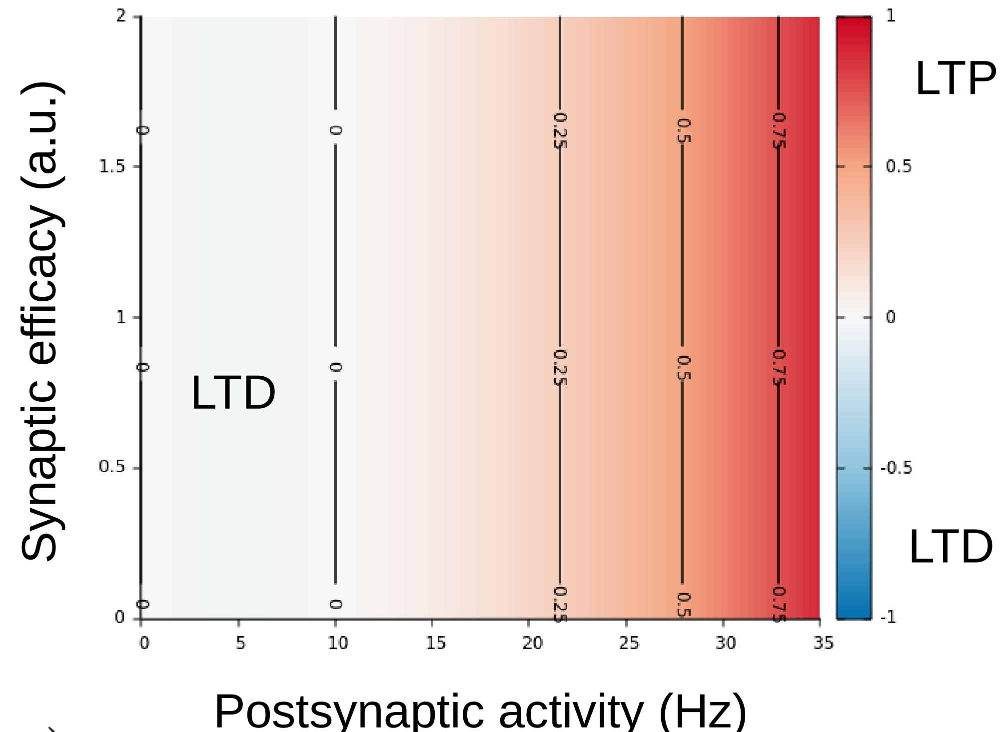

$$\frac{dw}{dt} \propto (\mathrm{pre})\,(\mathrm{post})\,((\mathrm{post}) - \kappa)$$

Purely Hebbian

Postsynaptic activation determines LTP/LTD

Zenke et al. (2015, 2017)

Experimental evidence for heterosynaptic plasticity

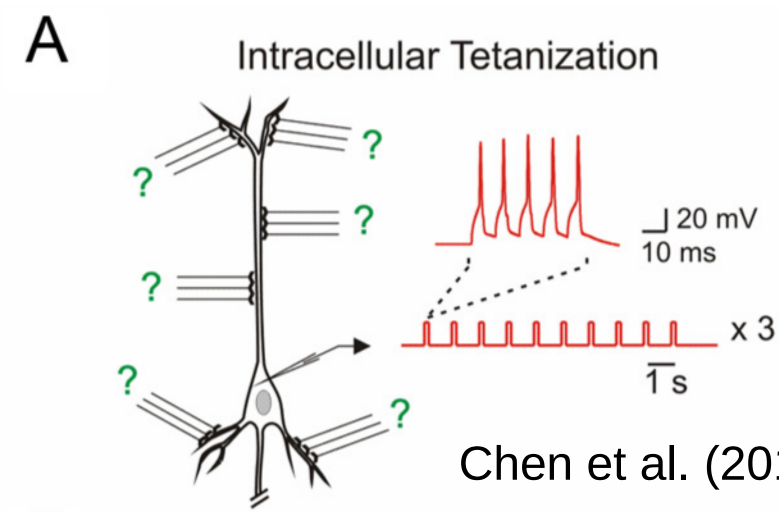

Chen et al. (2013)

$$\frac{dw}{dt} \propto -\beta \left(\mathrm{post}\right)^4 \left(w - \tilde{w}\right)$$

Zenke et al. (2015)

Heterosynaptic plasticity acts on the right timescale

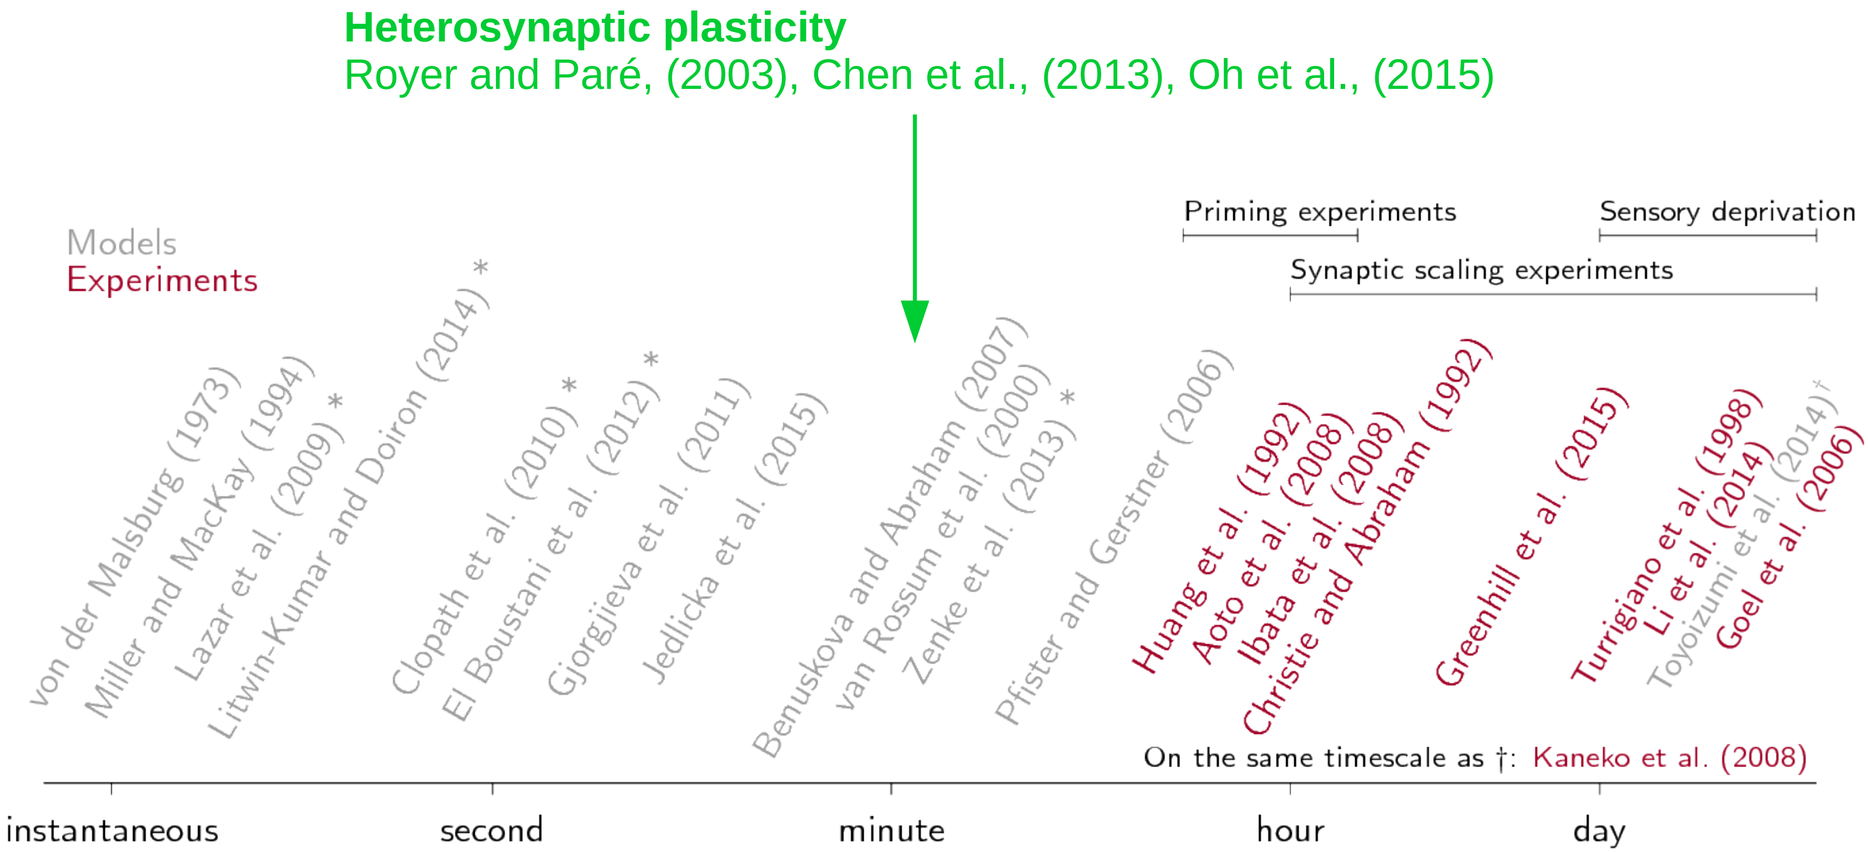

Non-Hebbian plasticity can act as rapid compensatory process

Triplet STDP

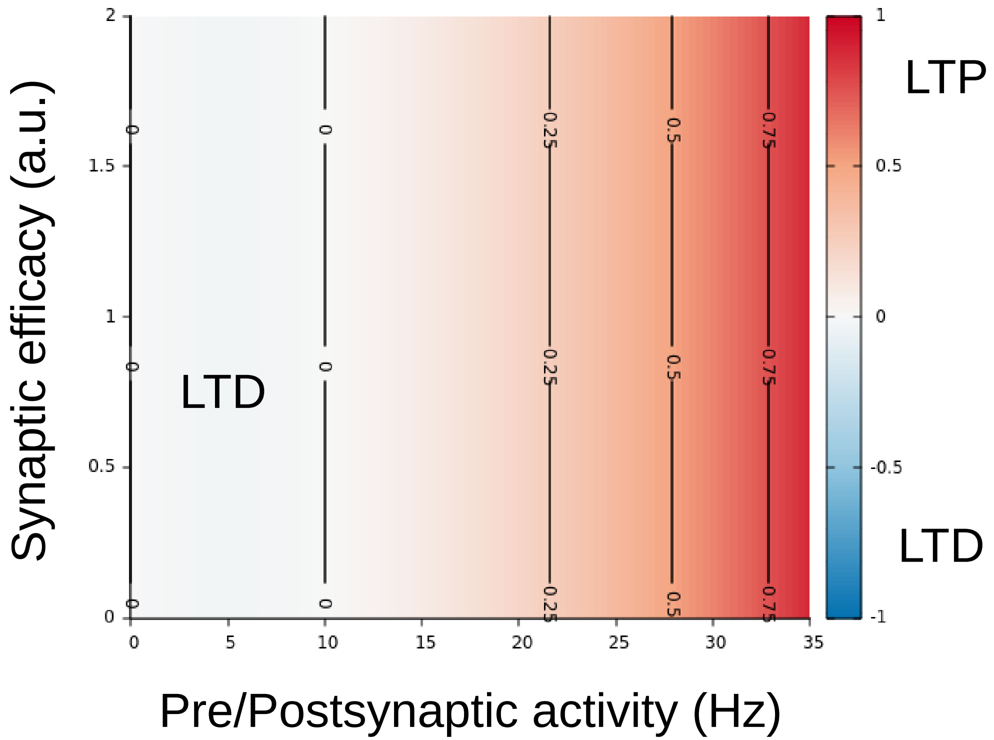

$$\frac{dw}{dt} \propto (\mathrm{pre})\,(\mathrm{post})\,((\mathrm{post}) - \kappa)$$

Zenke et al. (2015, 2017)

Non-Hebbian plasticity can act as rapid compensatory process

Triplet STDP + heterosynaptic plasticity

Low presynaptic activity (1Hz)

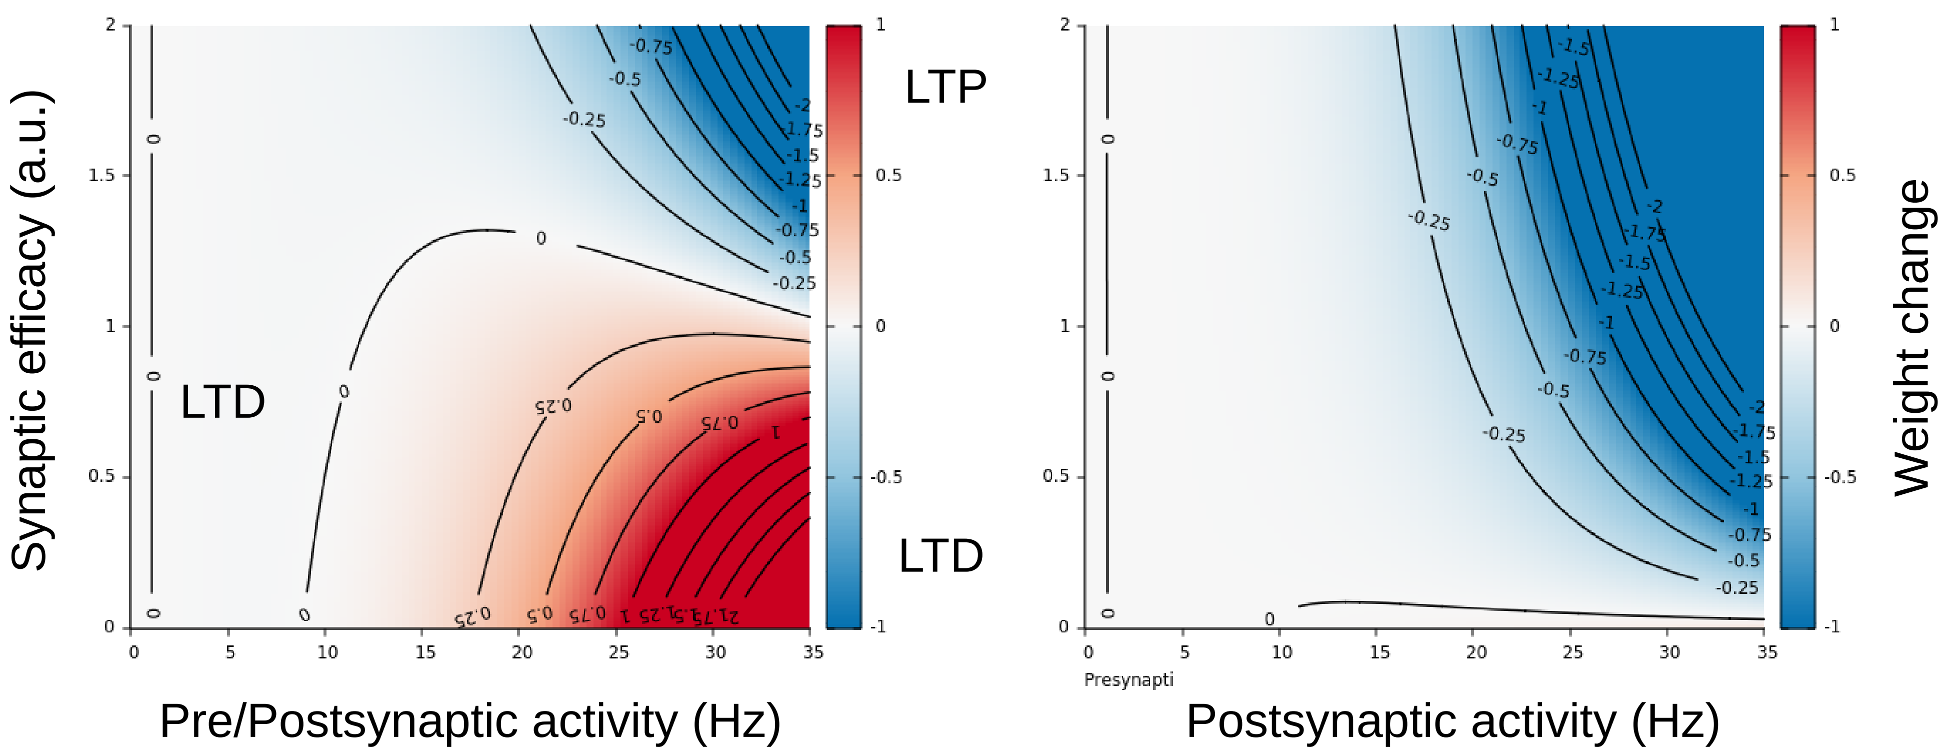

Heterosynaptic plasticity`

$$\frac{dw}{dt} \propto (\text{pre})\,(\text{post})\,((\text{post}) - \kappa) - \beta\,(\text{post})^{k}\,(w - \tilde{w})$$

- Local
- Non-Hebbian
- Stabilizing

Zenke et al. (2015, 2017)

Rapid compensatory processes can stabilize plastic network simulations

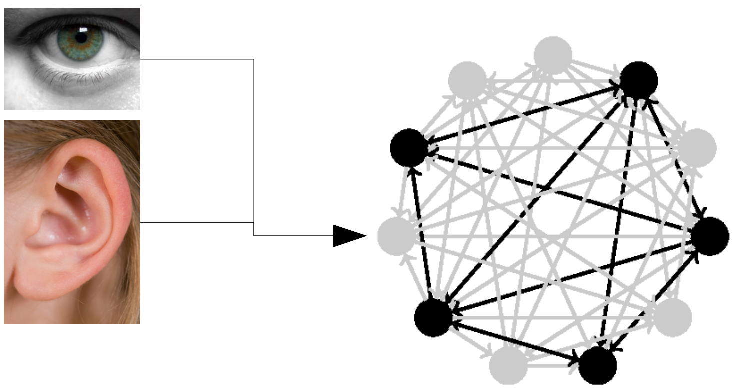

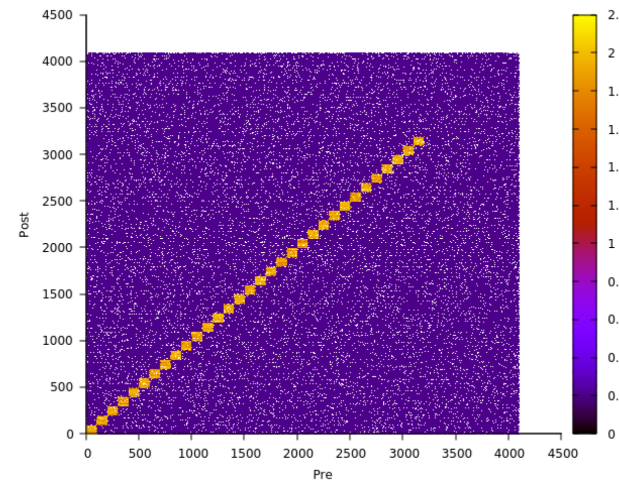

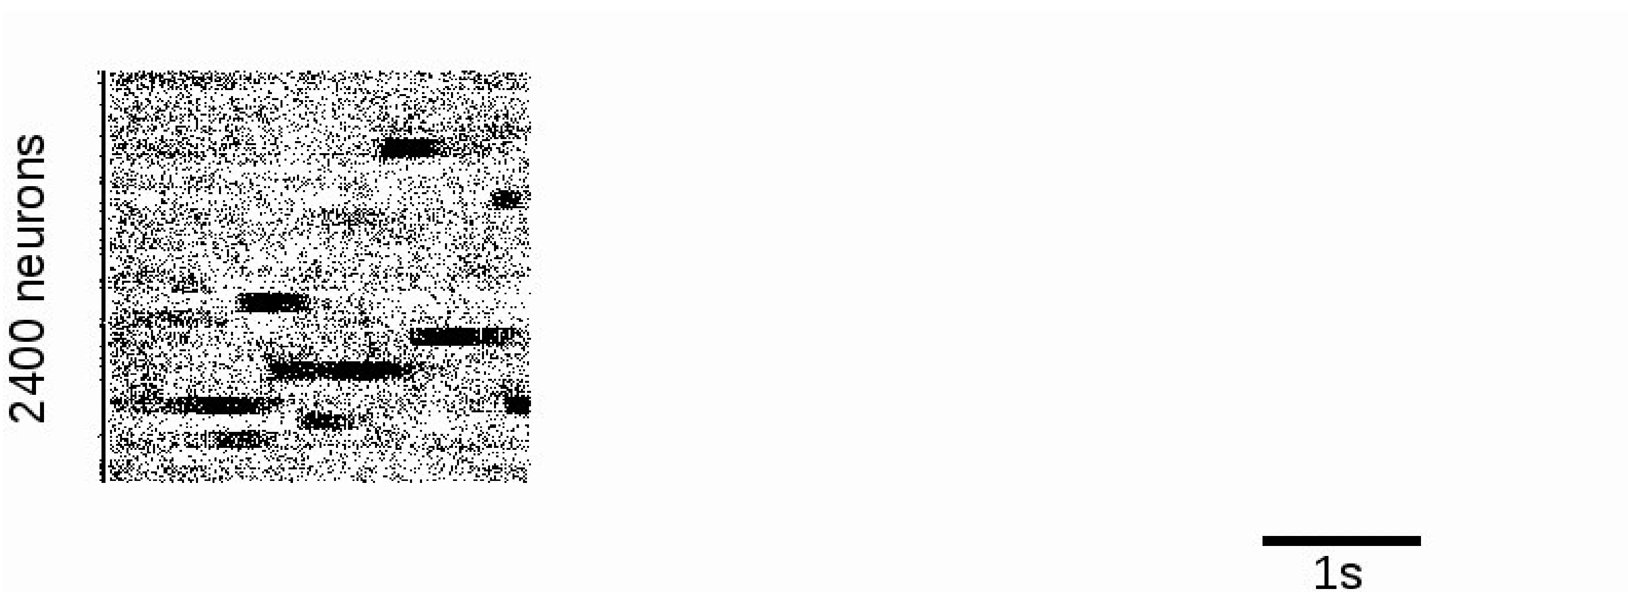

Orchestrated forms of plasticity in a recurrent network model

- Plasticity
  - Triplet STDP
  - Heterosynaptic plasticity
  - Transmitter triggered plasticity
  - Inhibitory plasticity
  - Consolidation: $\tilde{w}(t)$
- Spiking network
  - 4096 excitatory IF
  - 1024 inhibitory IF
  - Conductance based synapses
  - Short term plasticity
  - Spike triggered adaptation
  - Random sparse connectivity

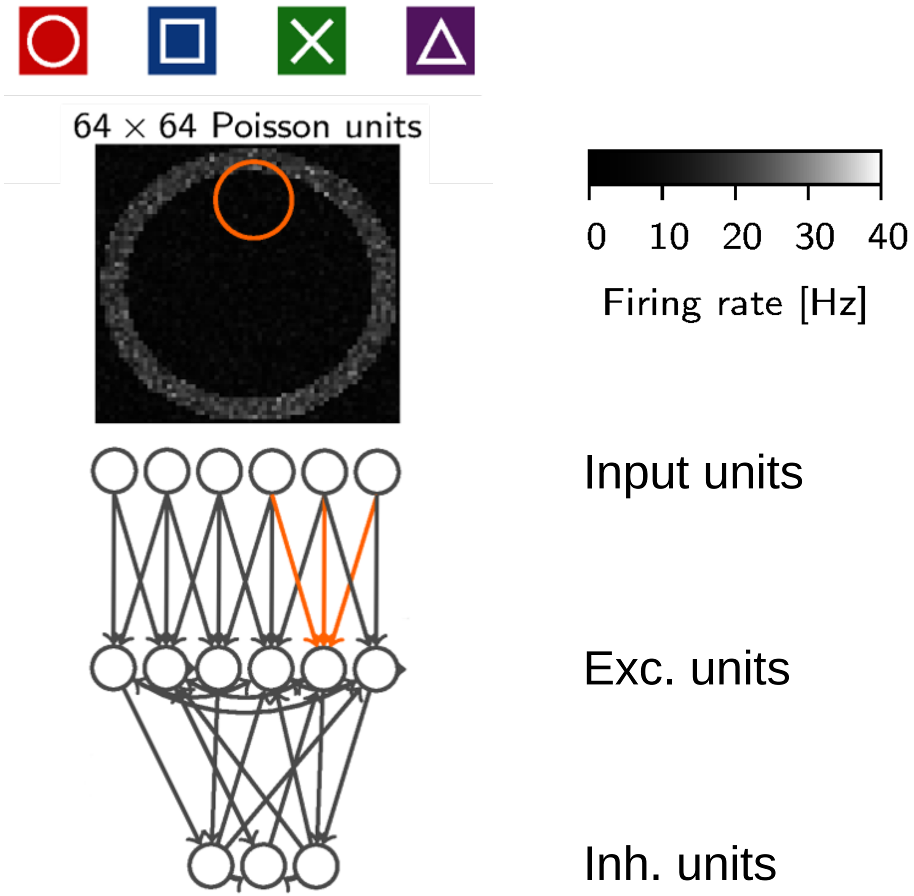

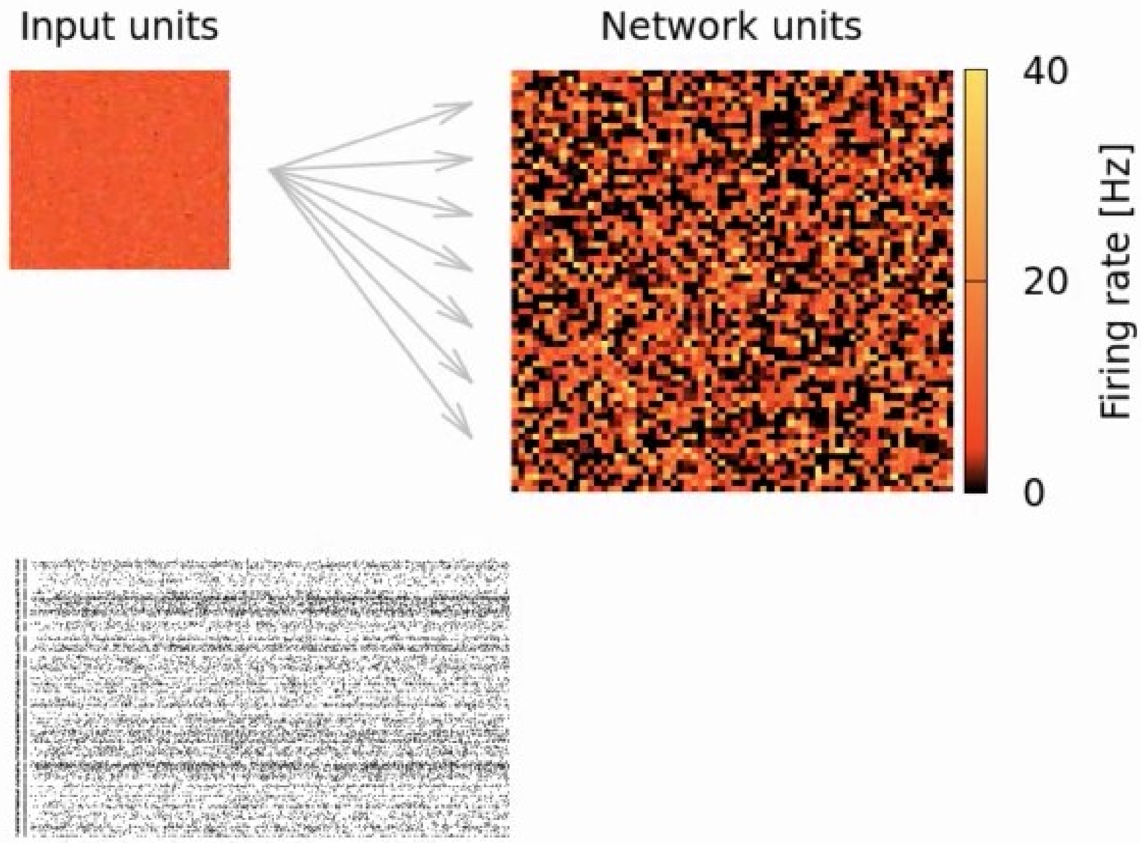

Plasticity inactive

F. Zenke - fzenke.net

Zenke et al. (2015)

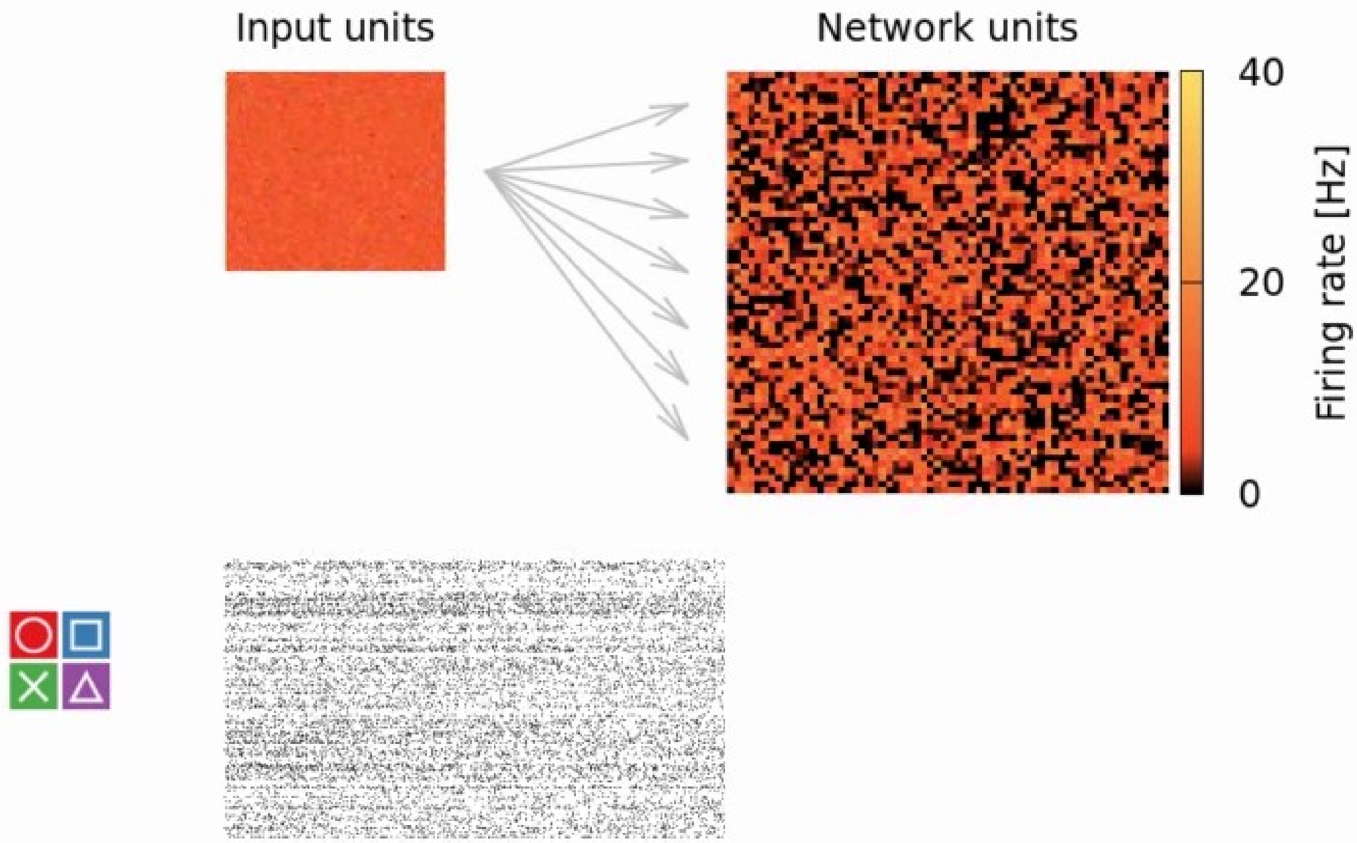

Hebbian & non-Hebbian plasticity active

Zenke et al. (2015)

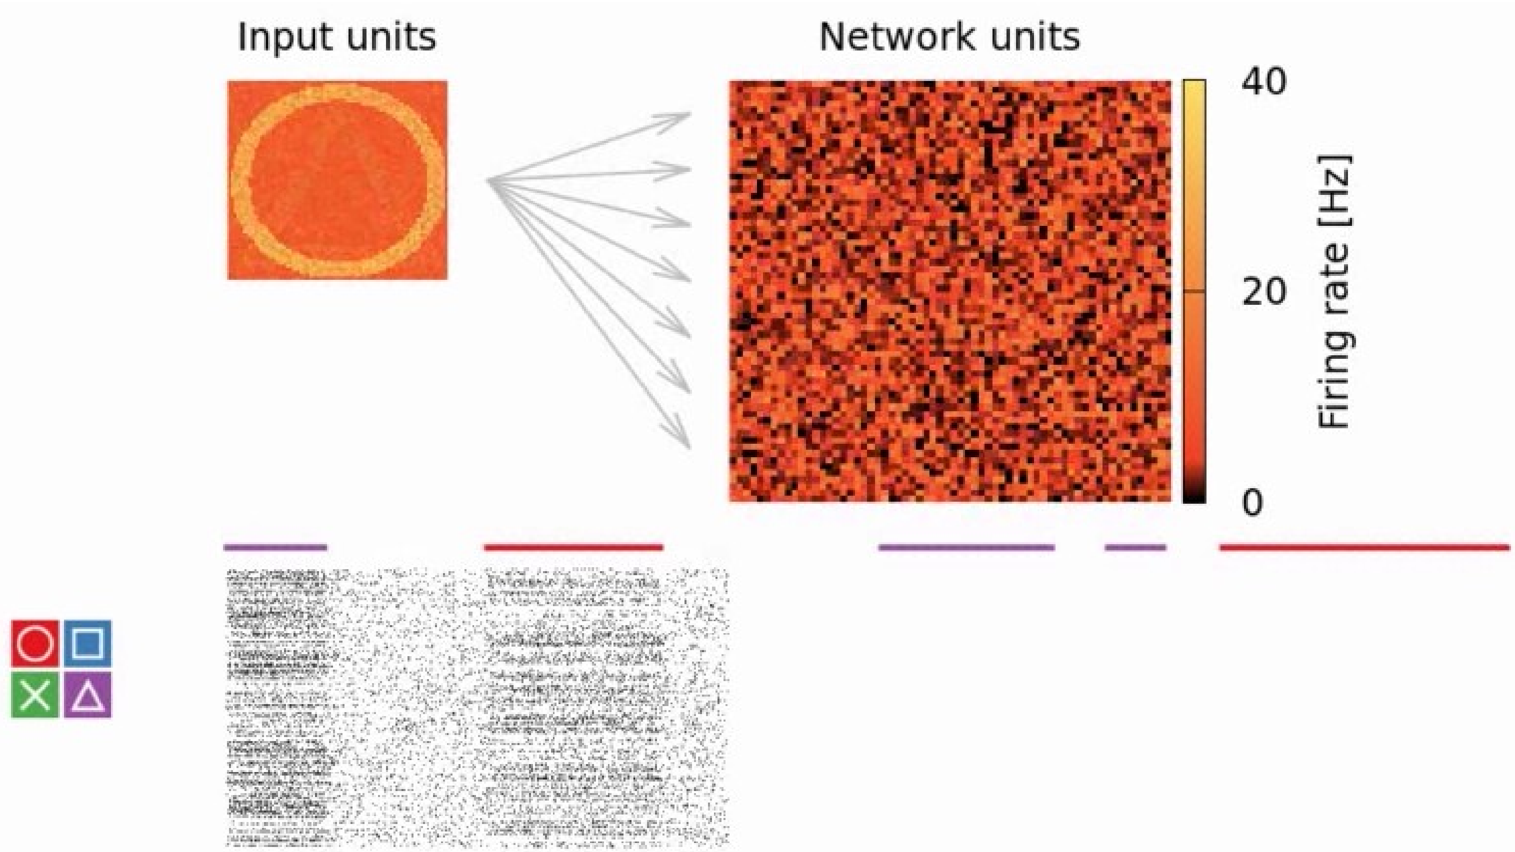

External stimuli

Zenke et al. (2015)

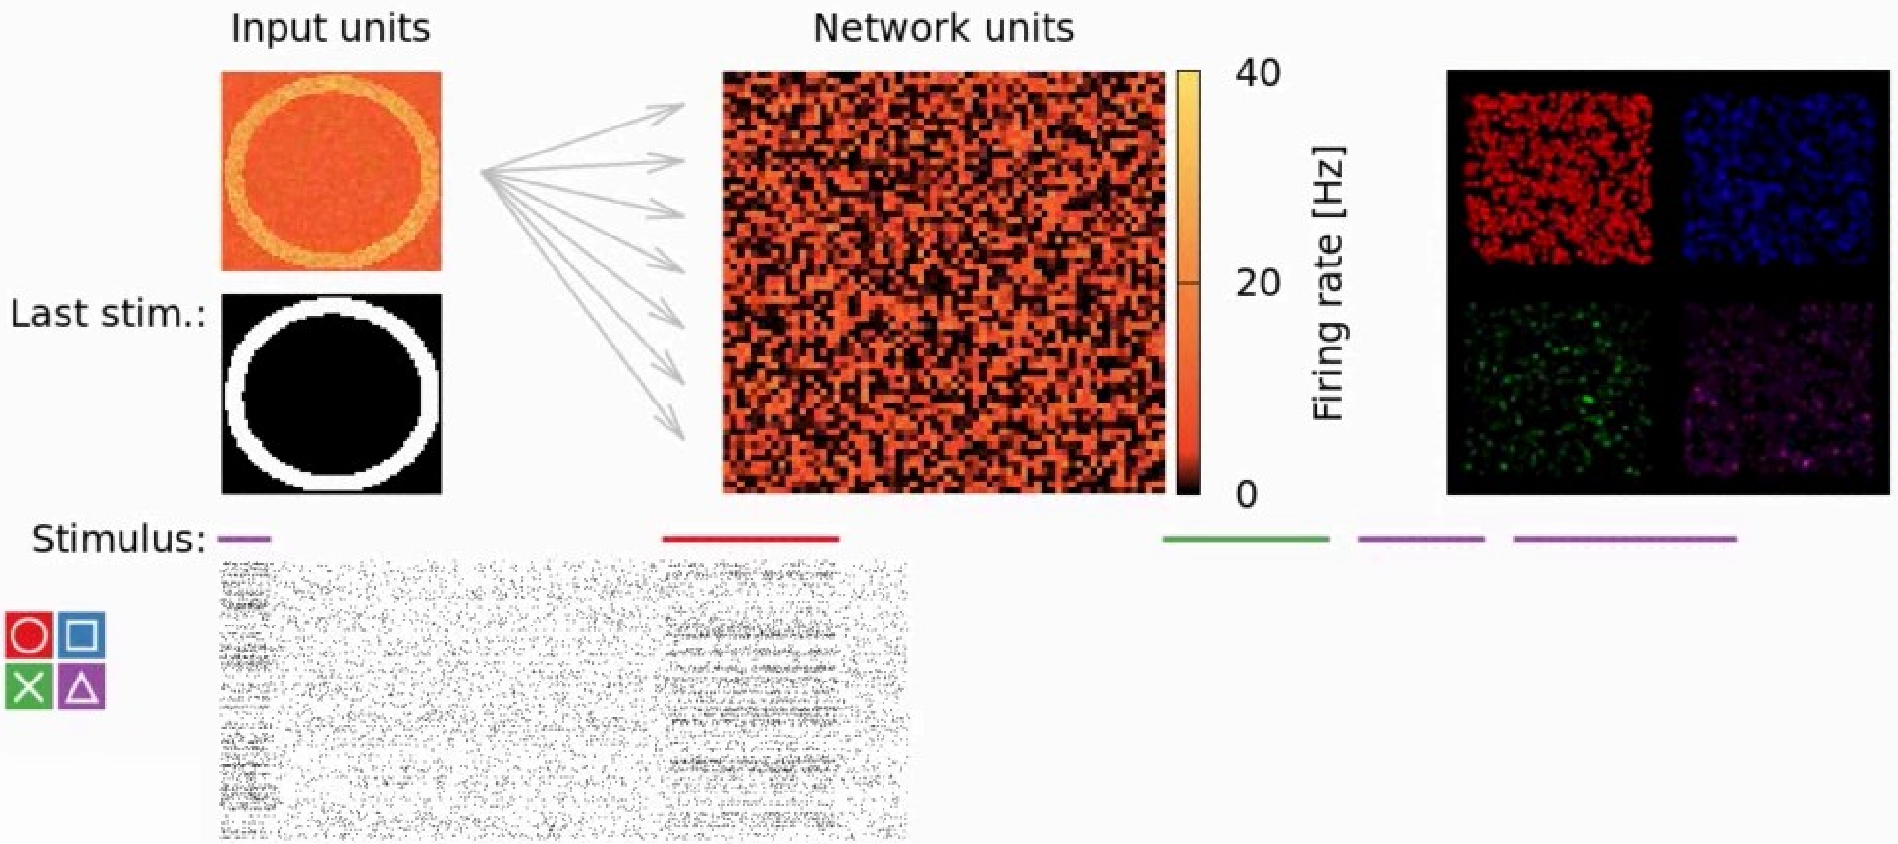

> External stimuli

Zenke et al. (2015)

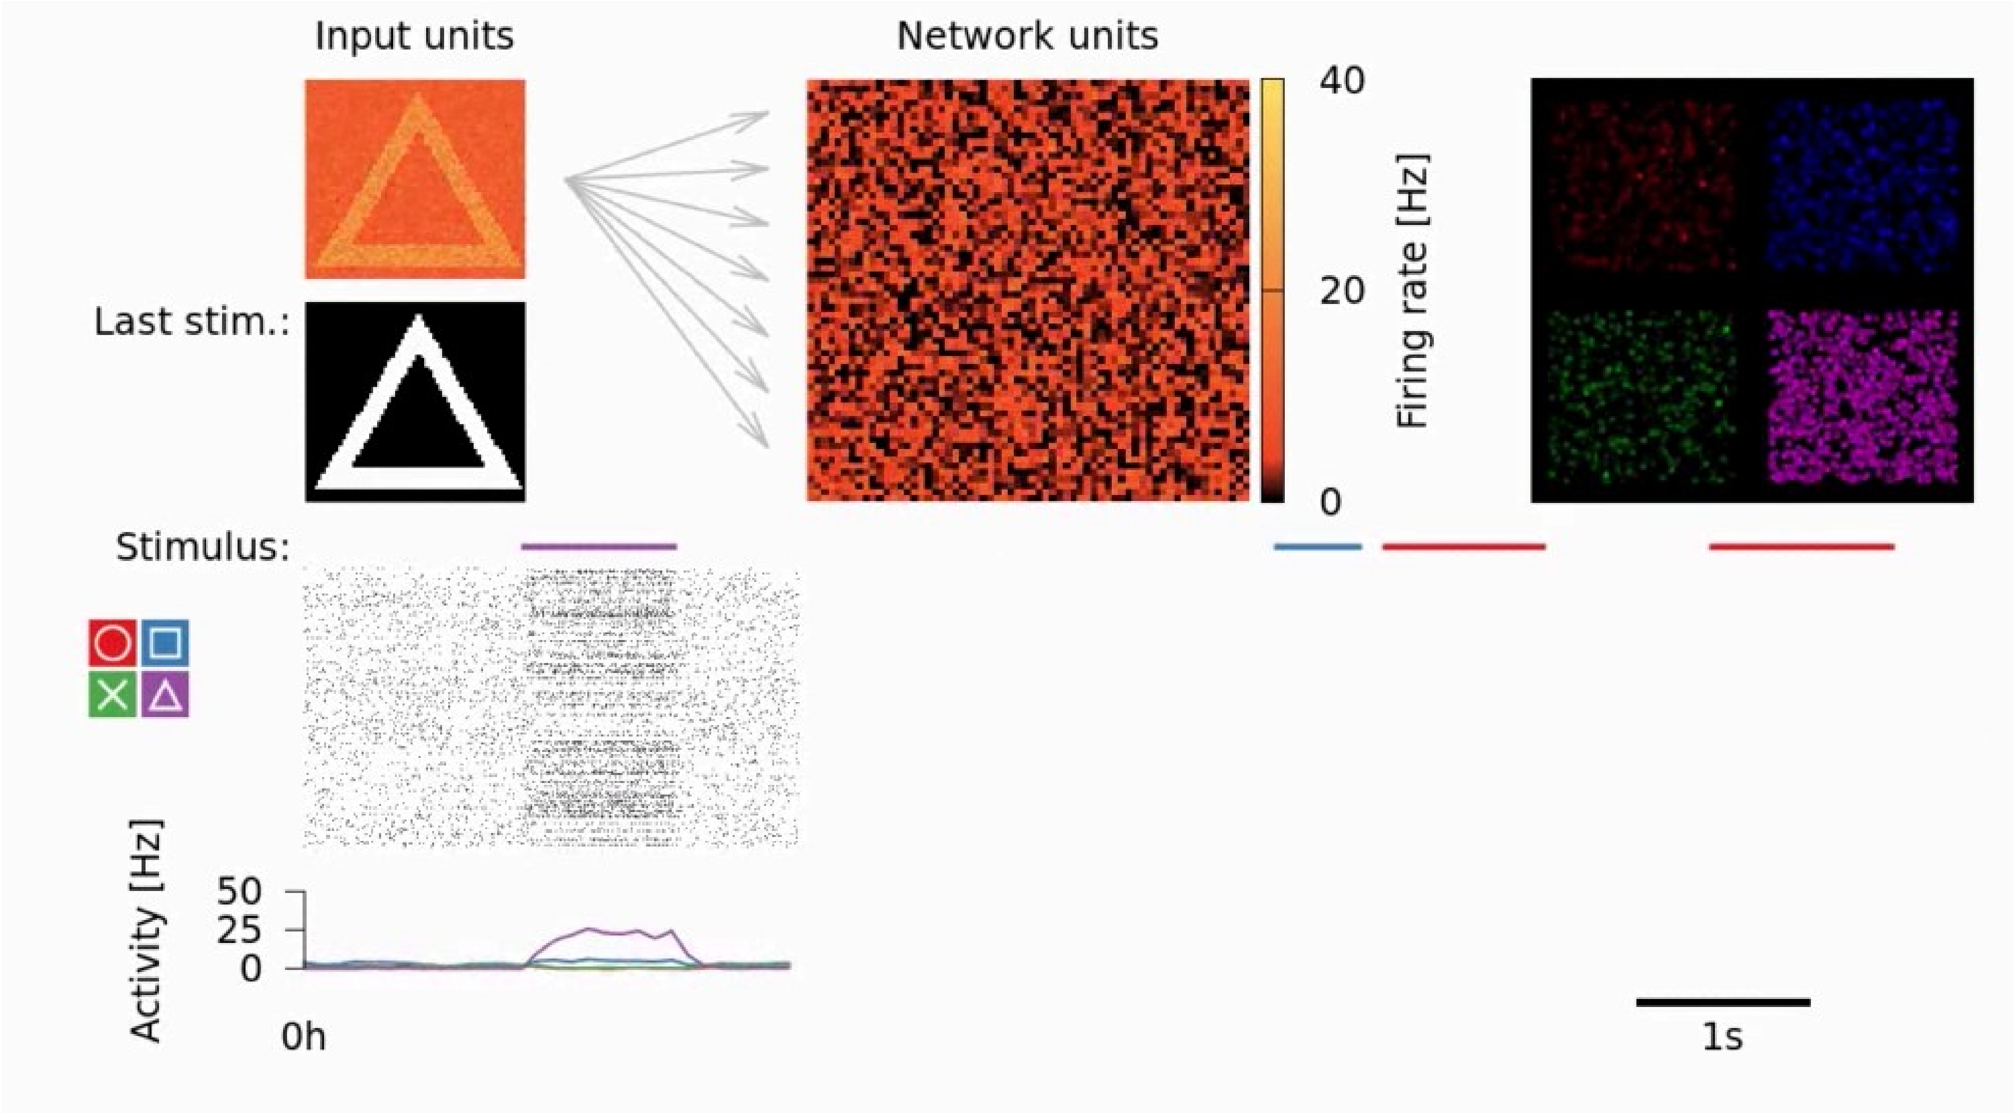

> External stimuli

Zenke et al. (2015)

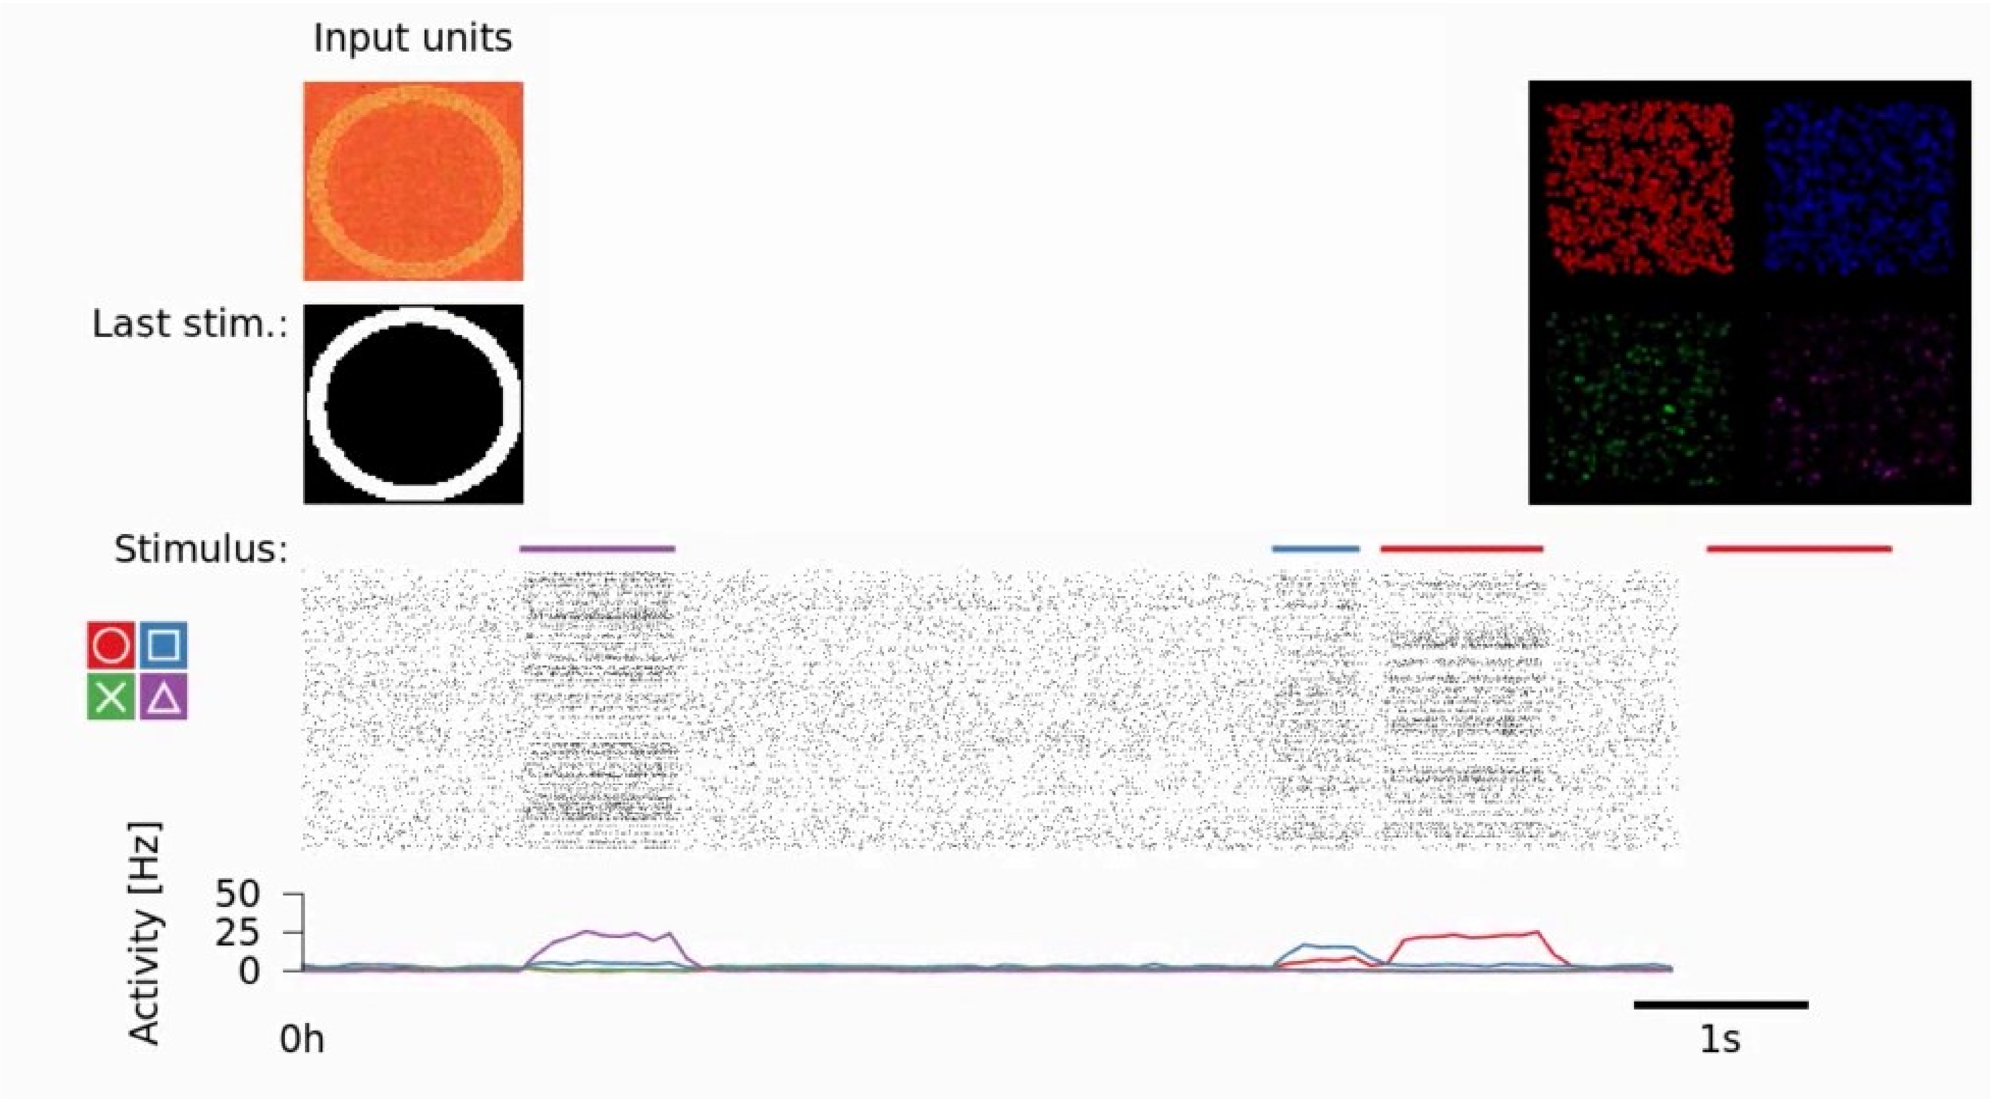

External stimuli

Zenke et al. (2015)

Receptive fields are refined and cell assemblies are formed

Input connections

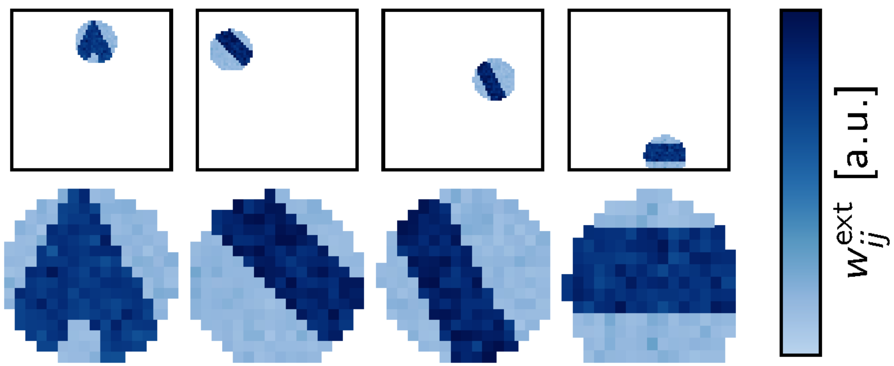

Recurrent connections

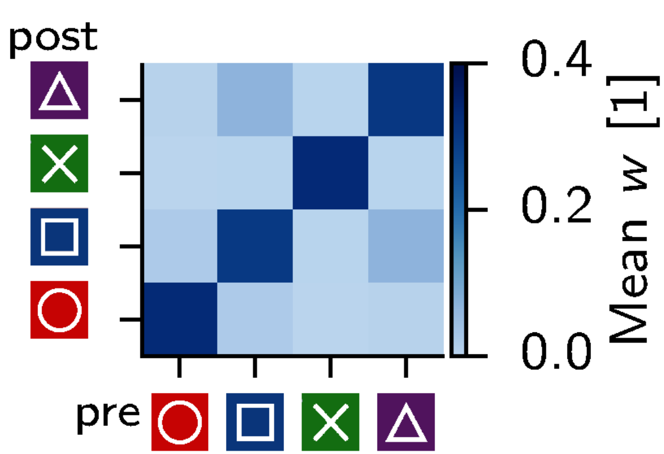

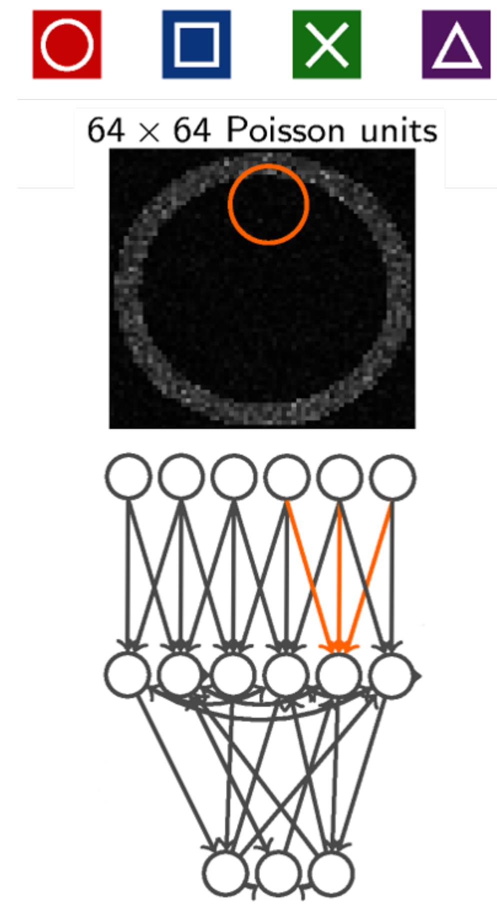

Distorted cues

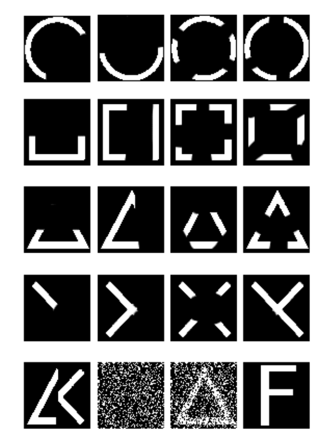

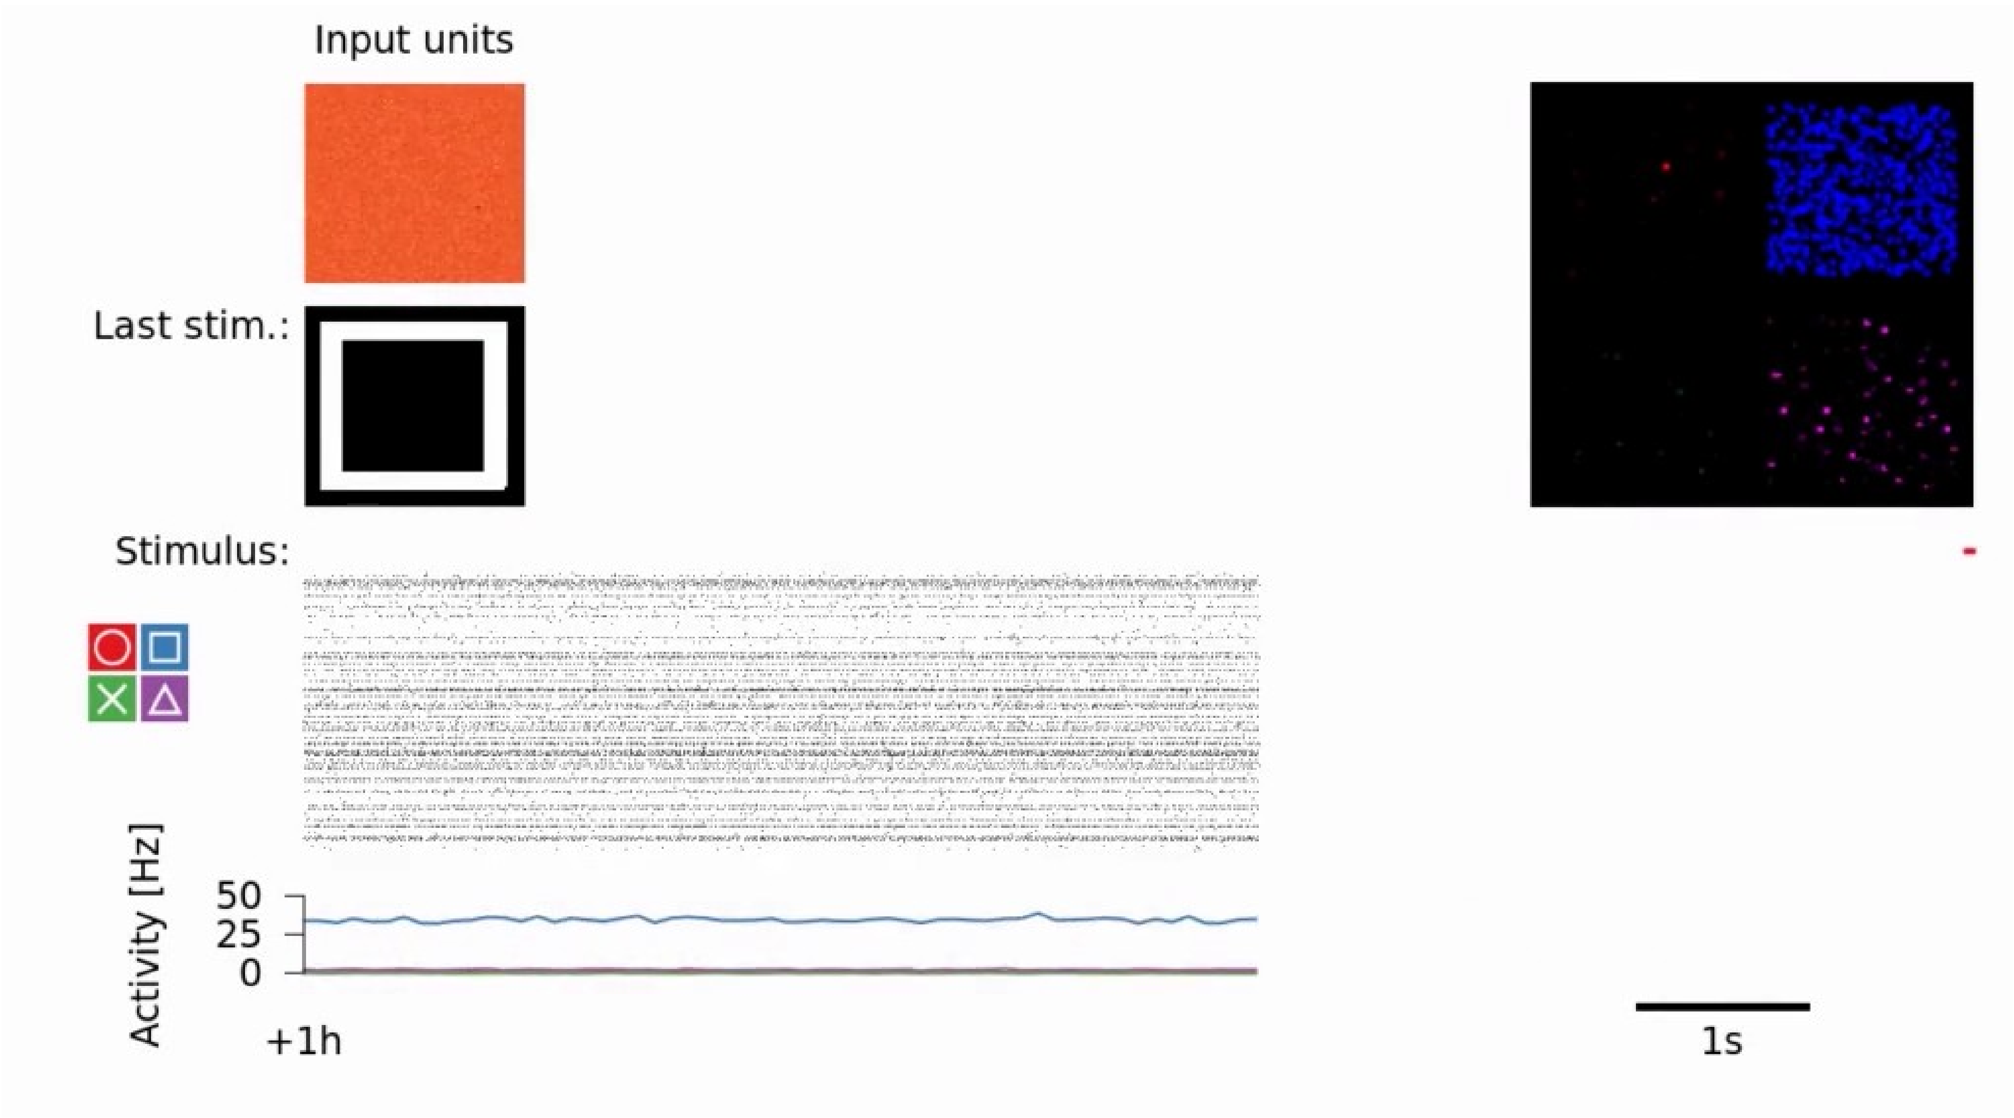

Distorted stimuli & recall (plasticity still active)

Zenke et al. (2015)

Unknown distractors

Zenke et al. (2015)

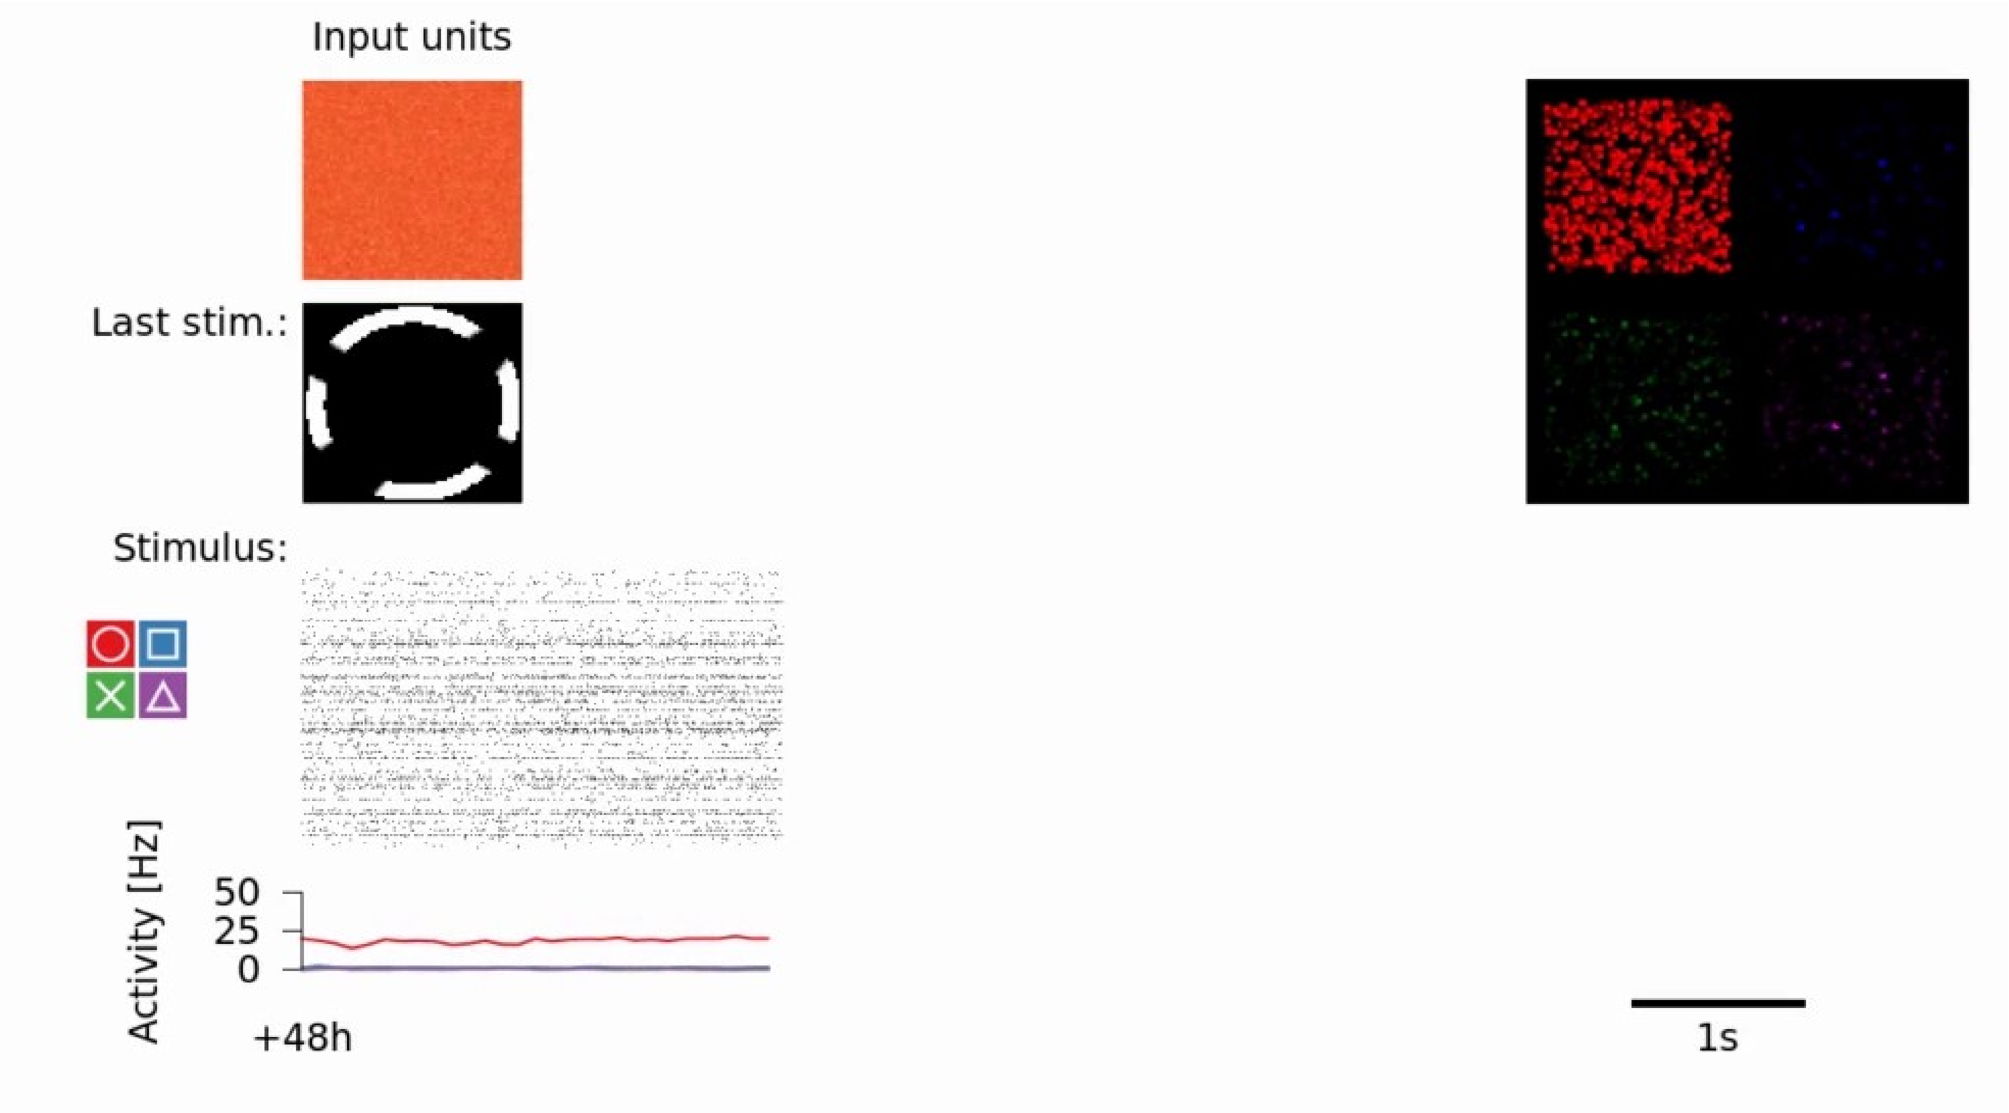

Learned representations stable for days

Zenke et al. (2015)

Part 1: Summary & Conclusion

- Intrinsically stable learning rules necessary to create stable memories
- “Rapid compensatory processes”
  - Have to be on comparable timescale
  - Manifest as non-Hebbian contributions to plasticity

Outline

- Part 1: Cell assemblies and “The temporal paradox of Hebbian and homeostatic plasticity”
- Part 2: Supervised learning in spiking neural networks starting from a cost function

Problem: Understanding neural activity

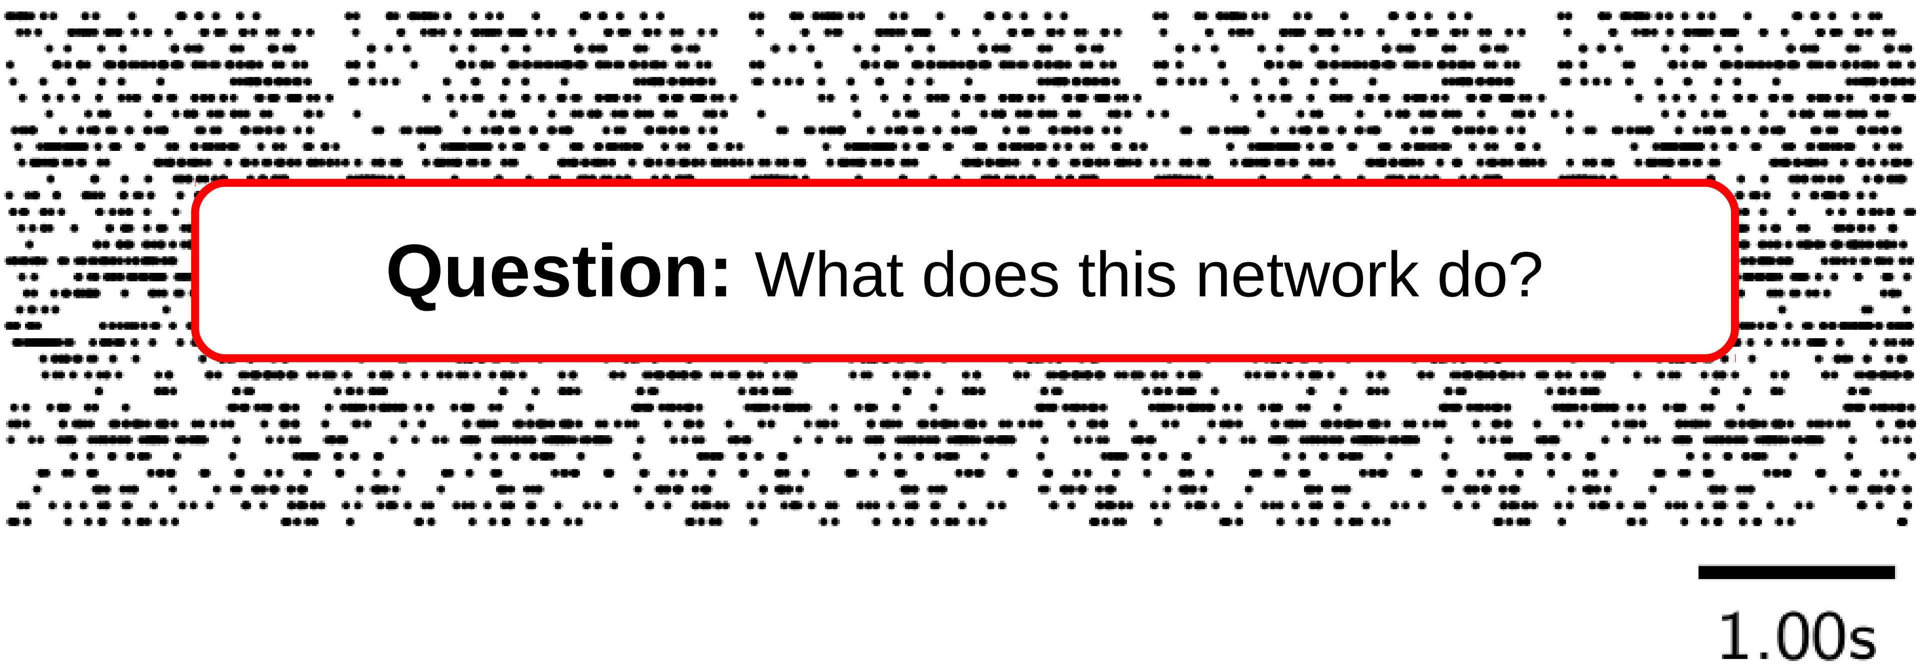

Question: What does this network do?

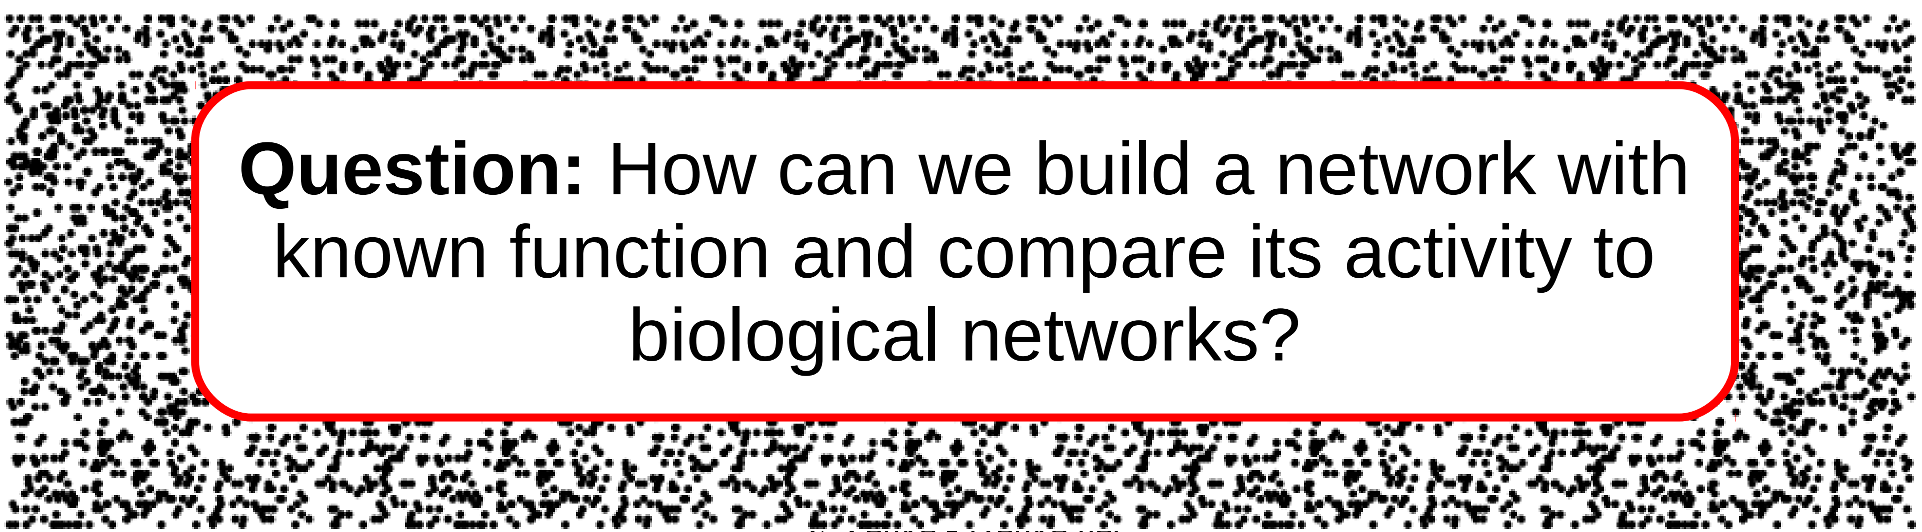

Question: How can we build a network with known function and compare its activity to biological networks?

Our ability to engineer spiking neural networks which perform complex tasks is limited in comparison to rate-based models

Computational neuroscience

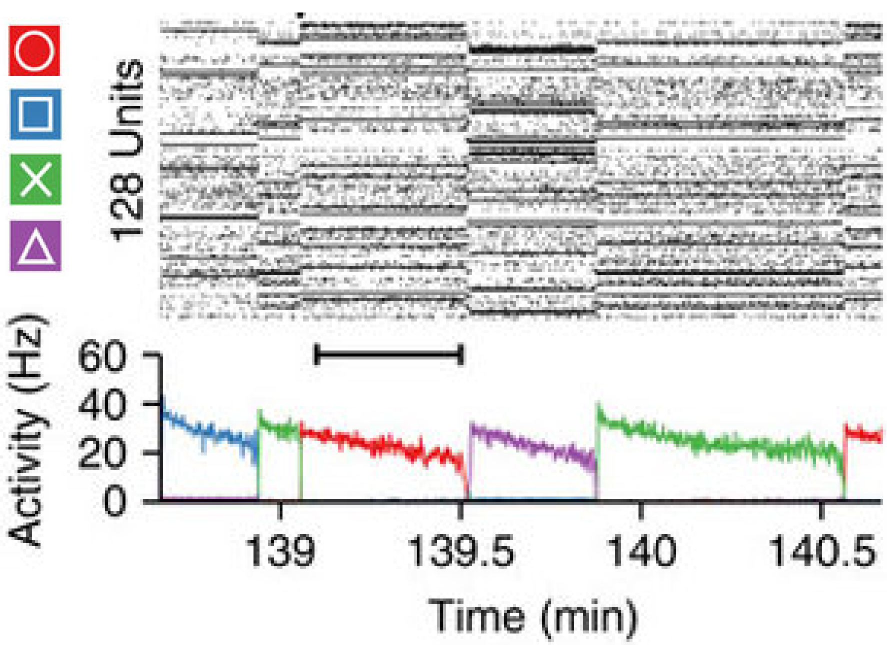

Zenke et al (2015)

Deep learning

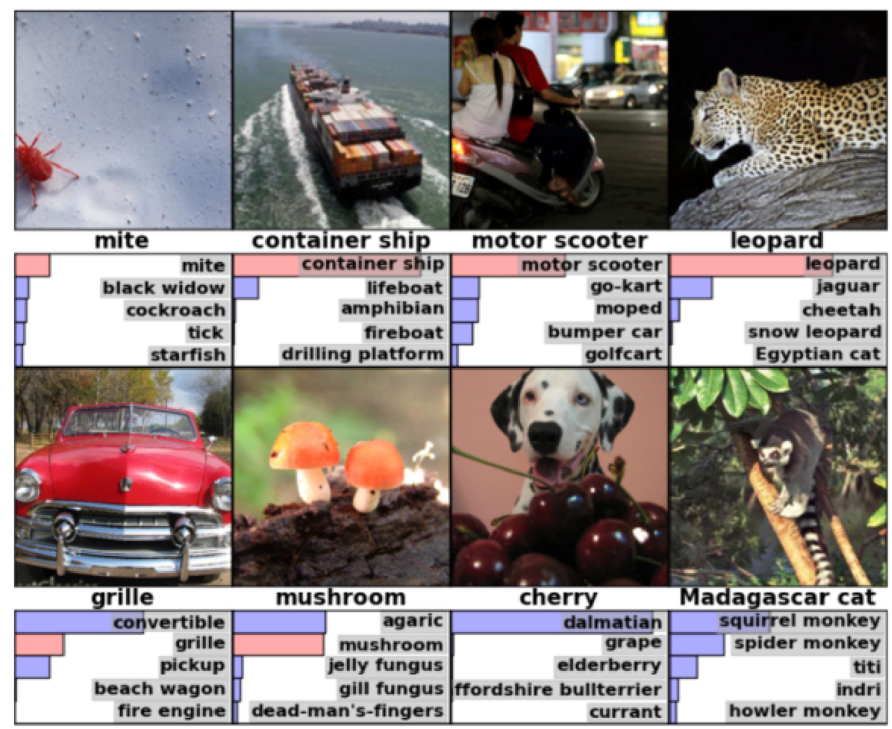

Krizhevsky et al. (2012)

Desiderata

- Spiking network which solves complex task
- Use spike timing
- Algorithm which could be conceivably

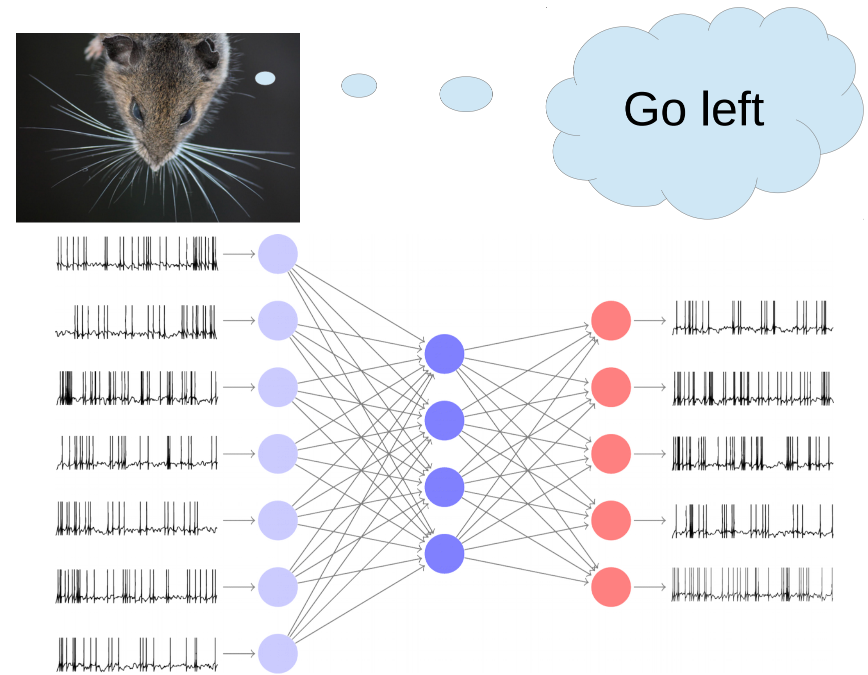

Aim

Get spiking networks to do something interesting, by starting from an objective function approach.

Deep learning approach: Gradient descent on “smooth” cost function

- Start with suitable cost function L
- Take the derivative w.r.t. to synaptic weights
- Do gradient descent $\frac{\partial w_k}{\partial t} \propto -\frac{\partial L}{\partial w_k}$

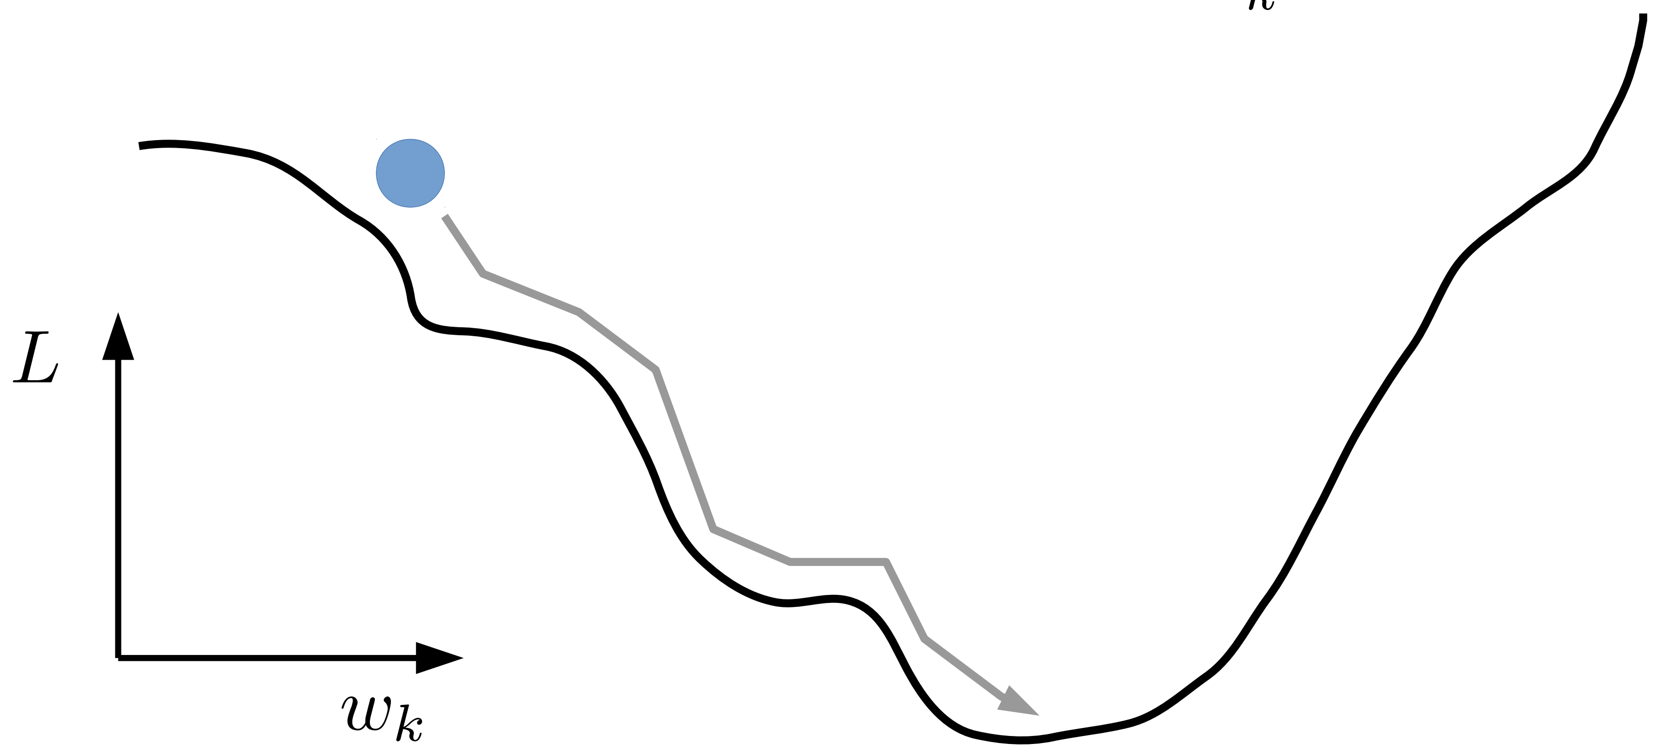

Problems with gradient descent in spiking neural networks

- Suitable cost function for spike trains
- Spikes are inherently non-differentiable
- Spiking neurons have history dependence due to spike reset
- Credit assignment in hidden layers

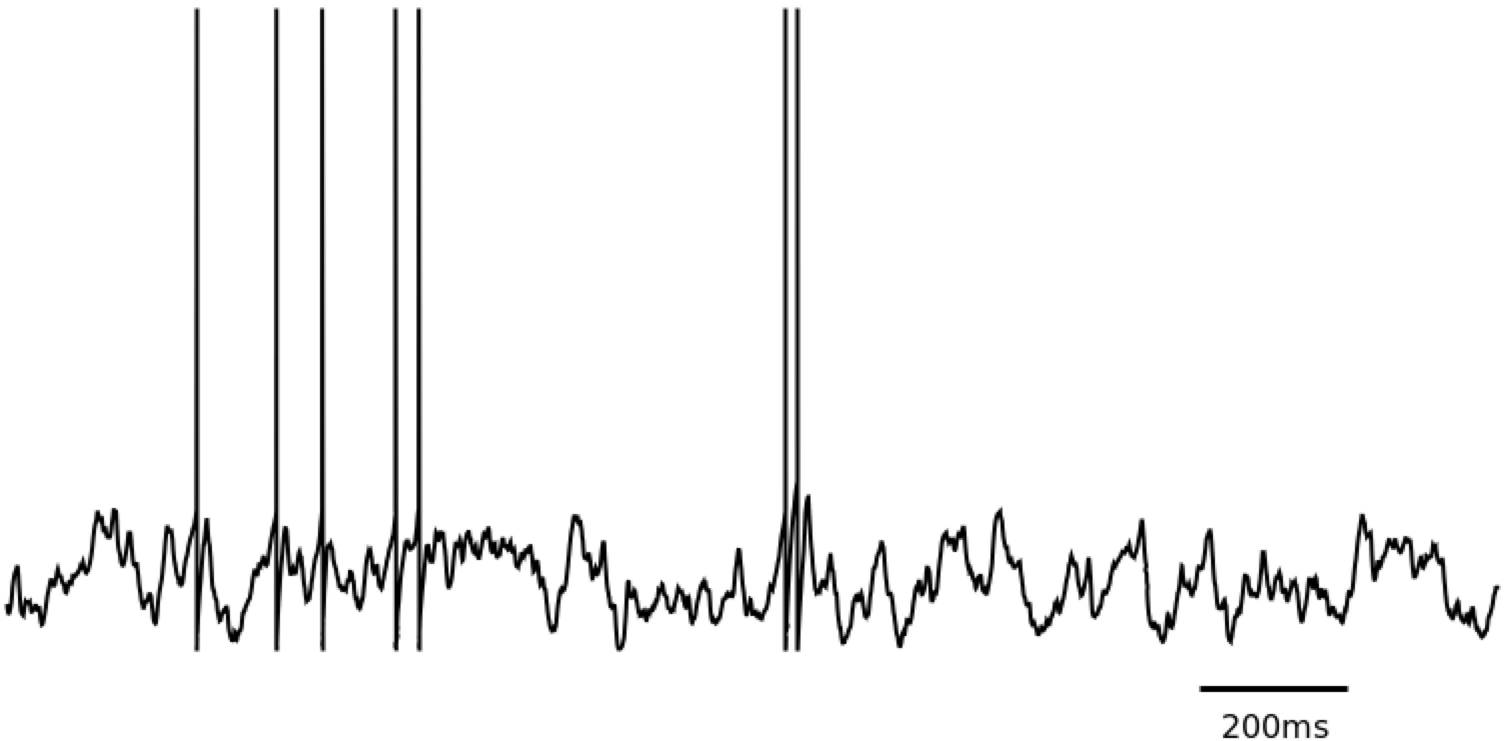
200ms

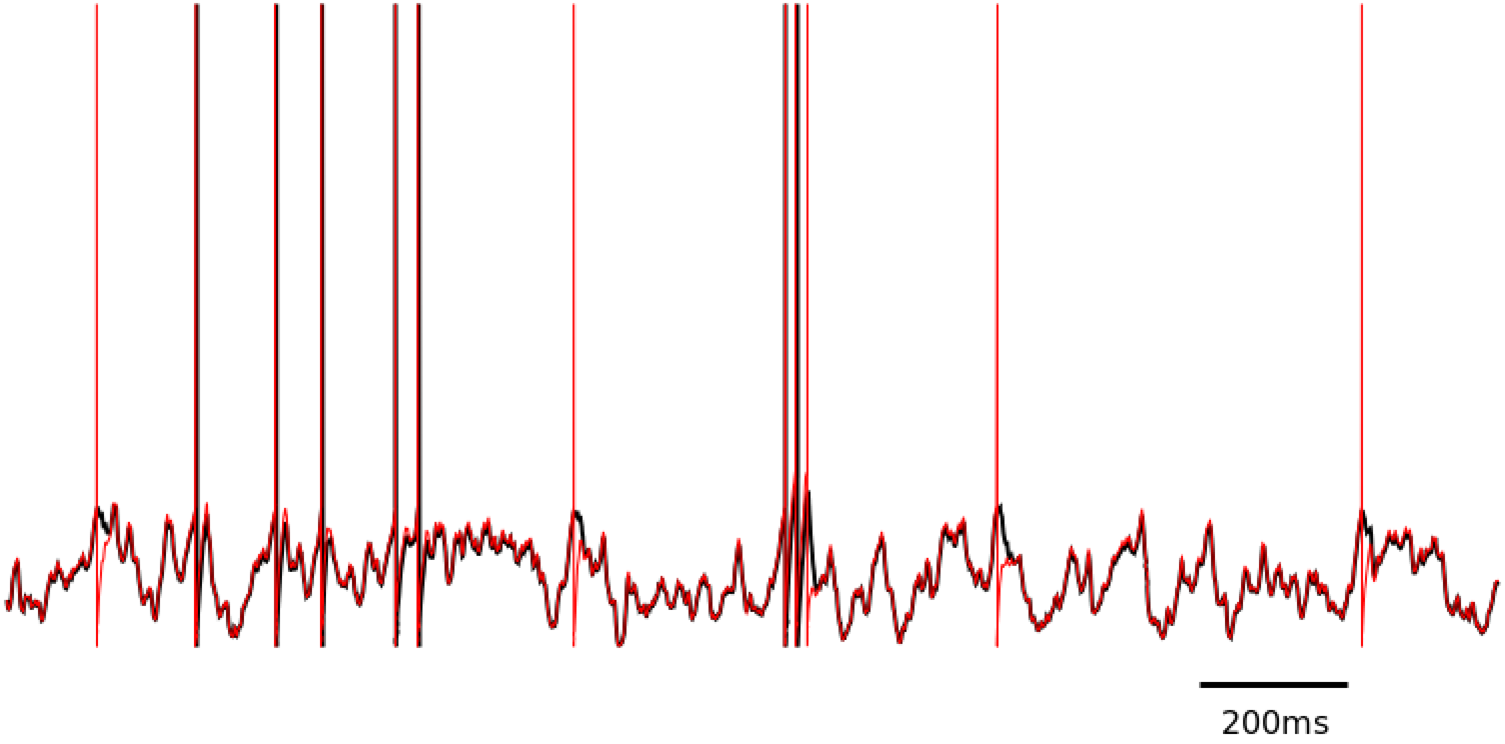
200ms

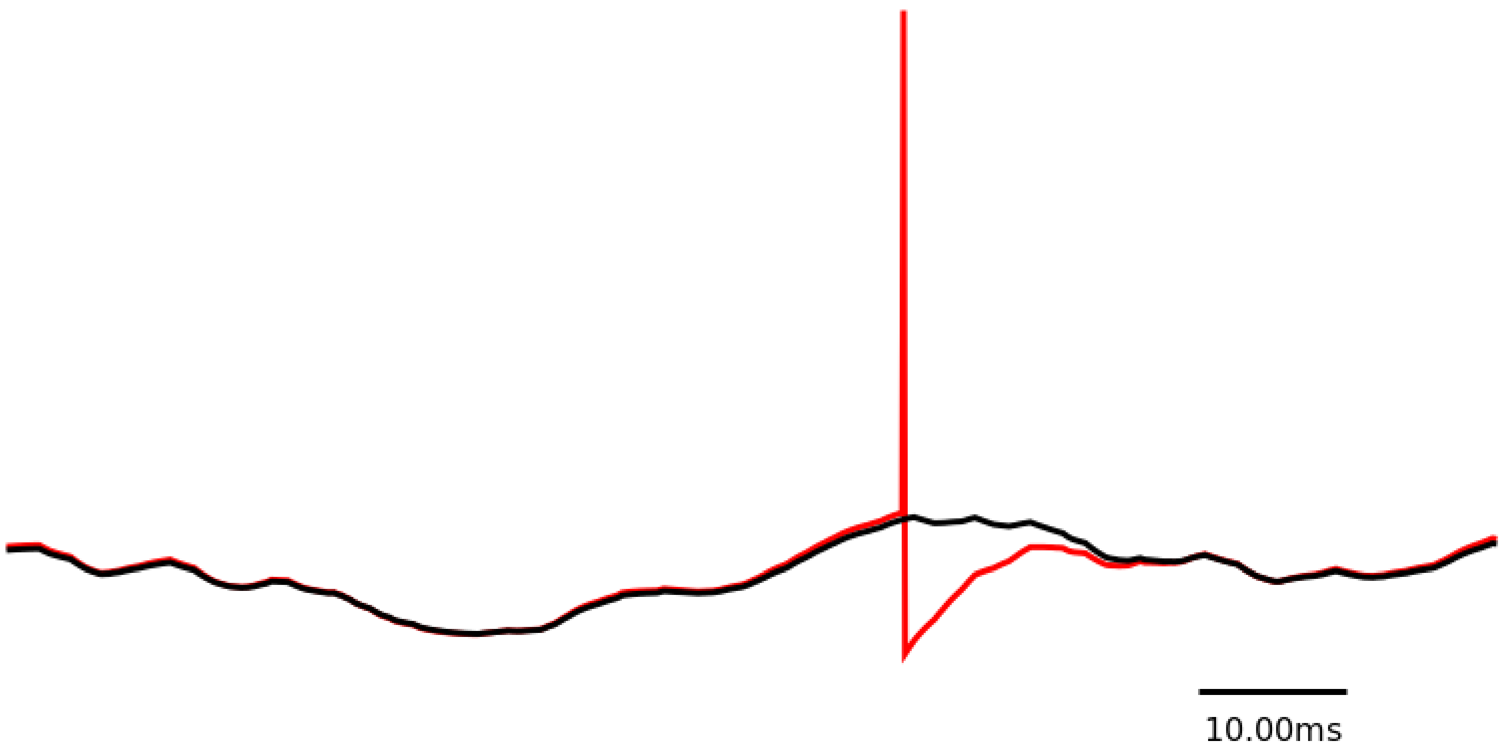

$$S(U_i(t)) \propto \Theta(U_i(t))$$

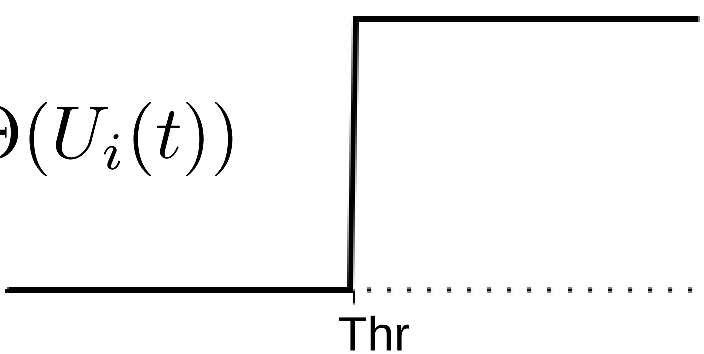

Optimization landscape in spiking neural networks flat everywhere

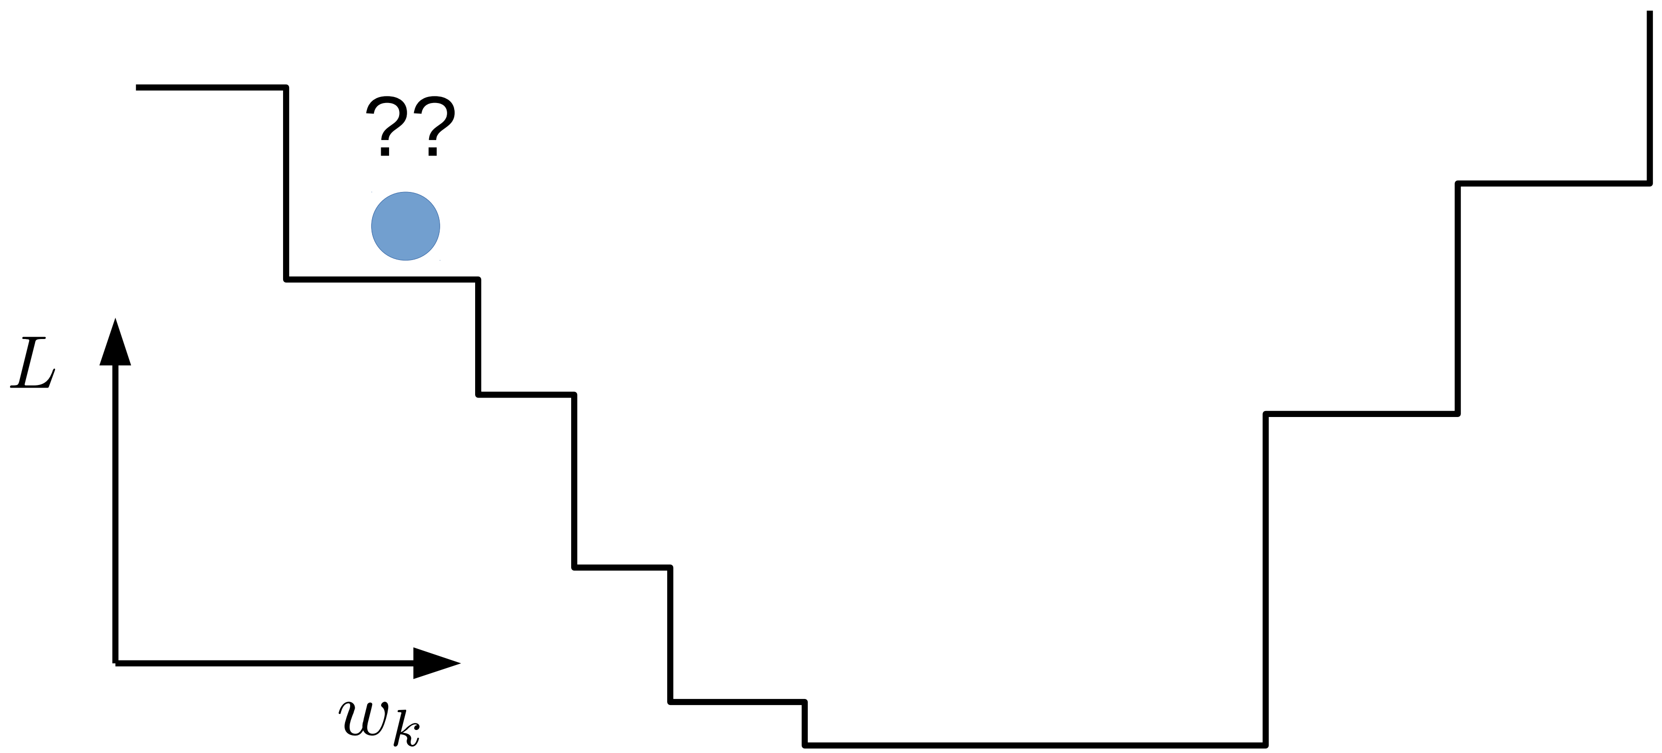

Idea: Smooth out the cost function

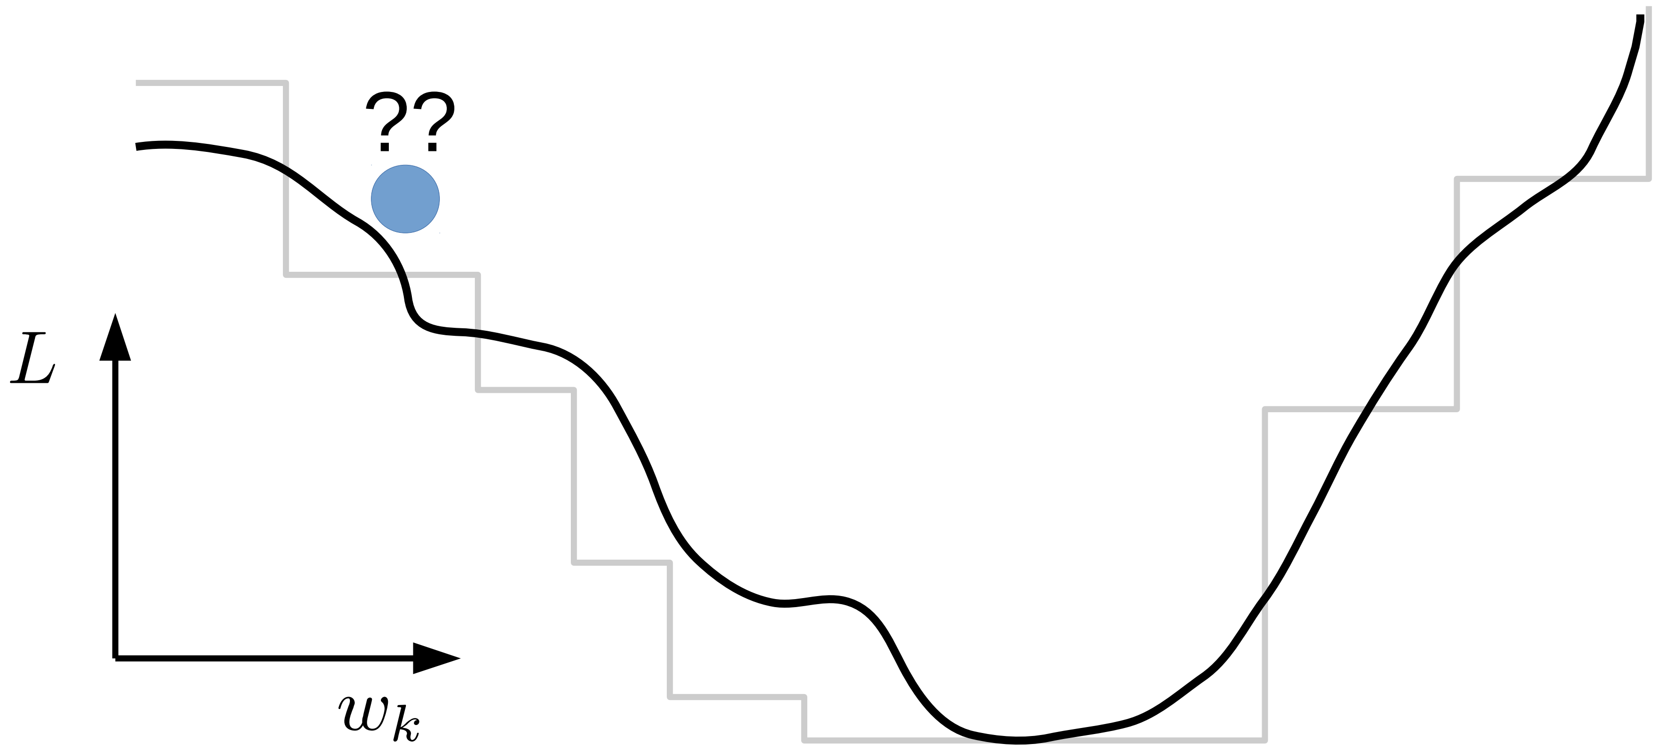

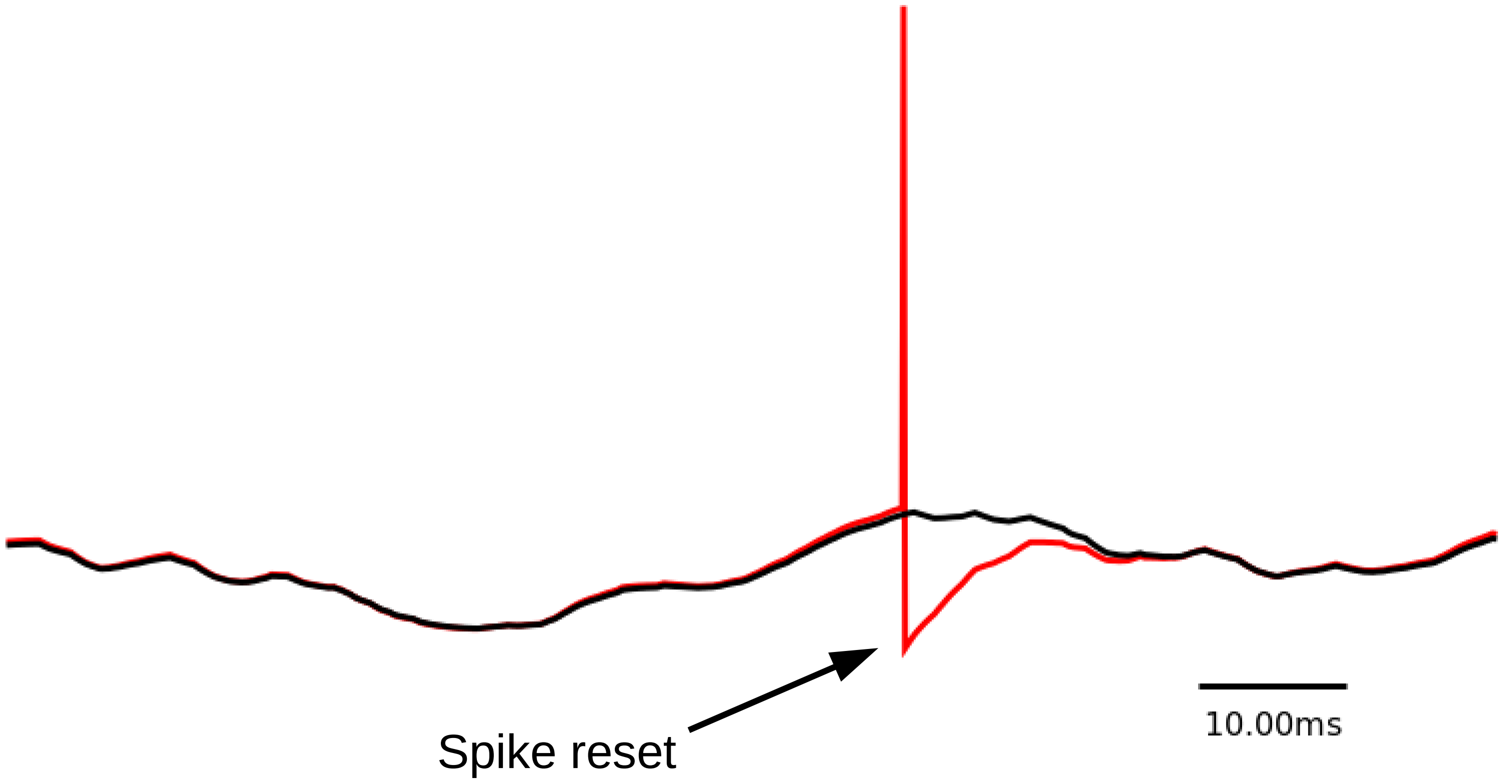
Spike reset
10.00ms

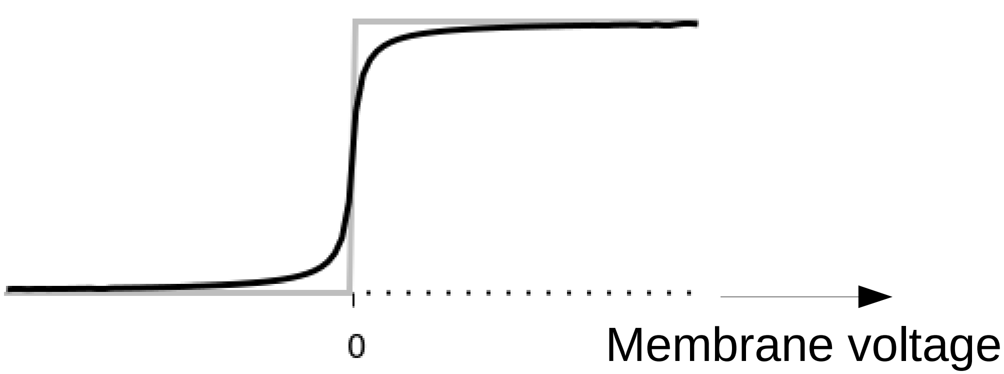
0
Membrane voltage

The learning rule can be interpreted as a three factor rule

$$\frac{\partial w_{ij}}{\partial t} \equiv \underbrace{e_i(t)}_{\text{error signal}} \quad \epsilon * (\underbrace{\epsilon * S_j(t)}_{\text{pre}} \underbrace{\sigma'(U_i)}_{\text{post}})$$

- Three factor rule
- Non-linear Hebbian
- Voltage-based
- Think of outer convolution as eligibility trace

Sequential inputs → single spike

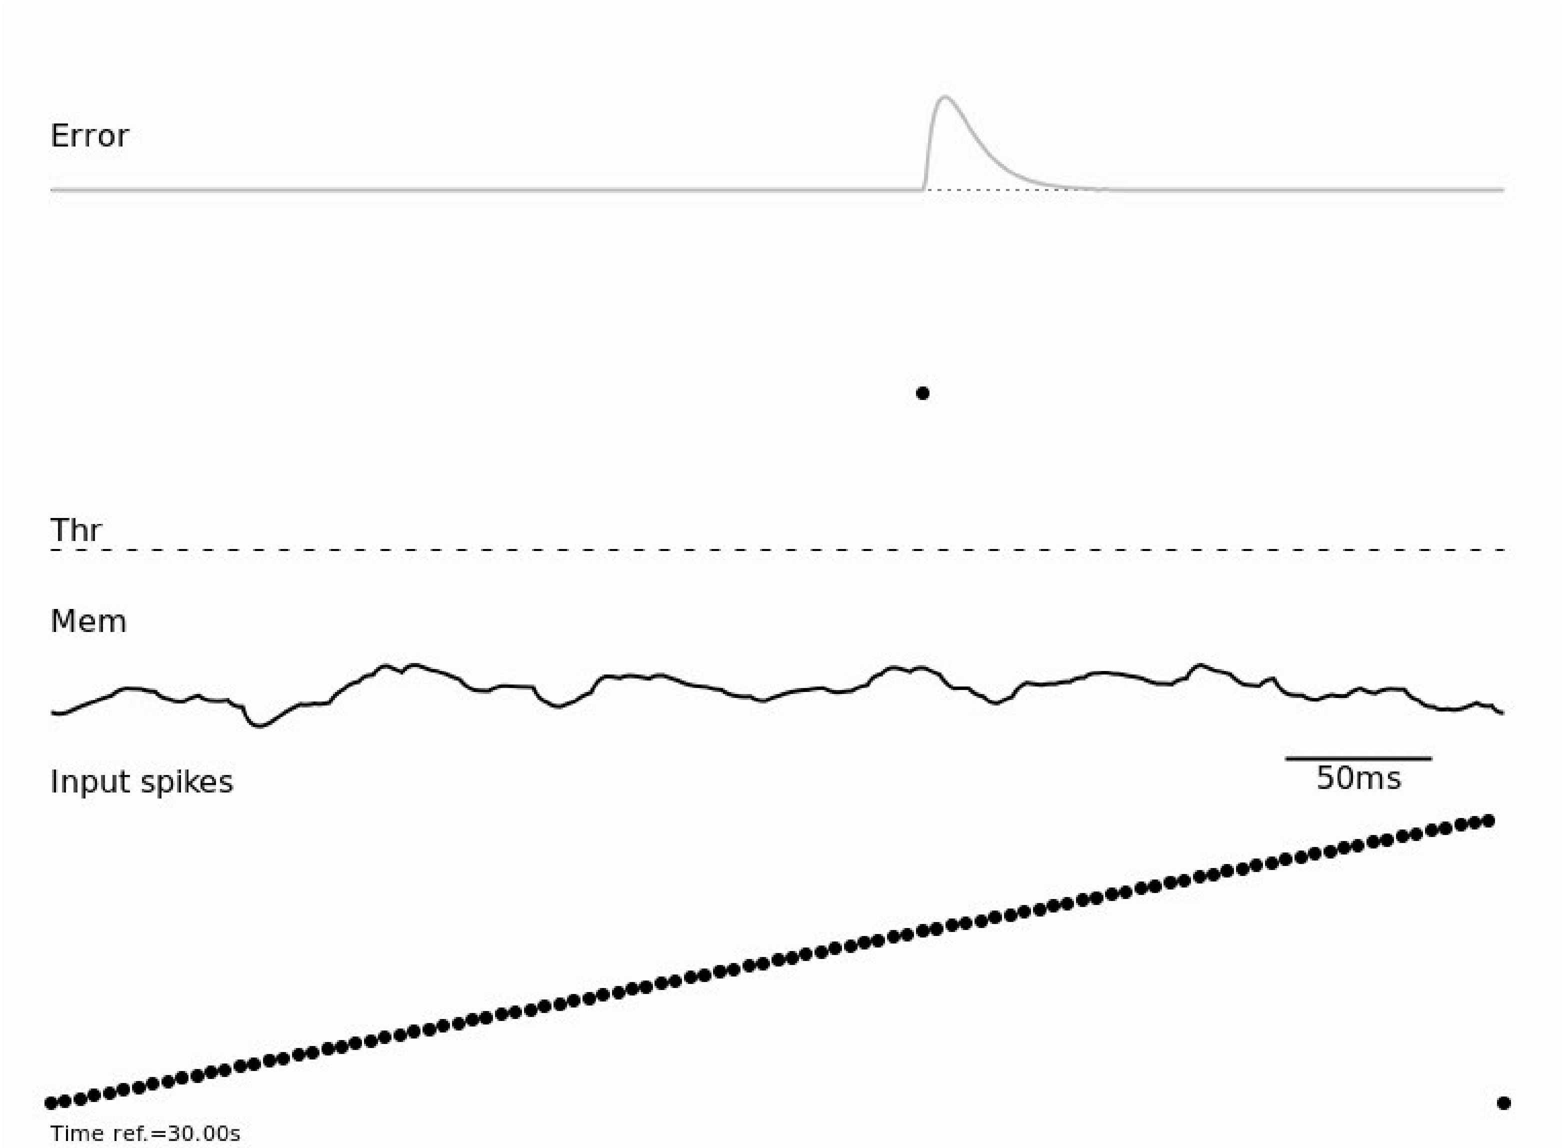

Learning of output spike train

Output:

Input:

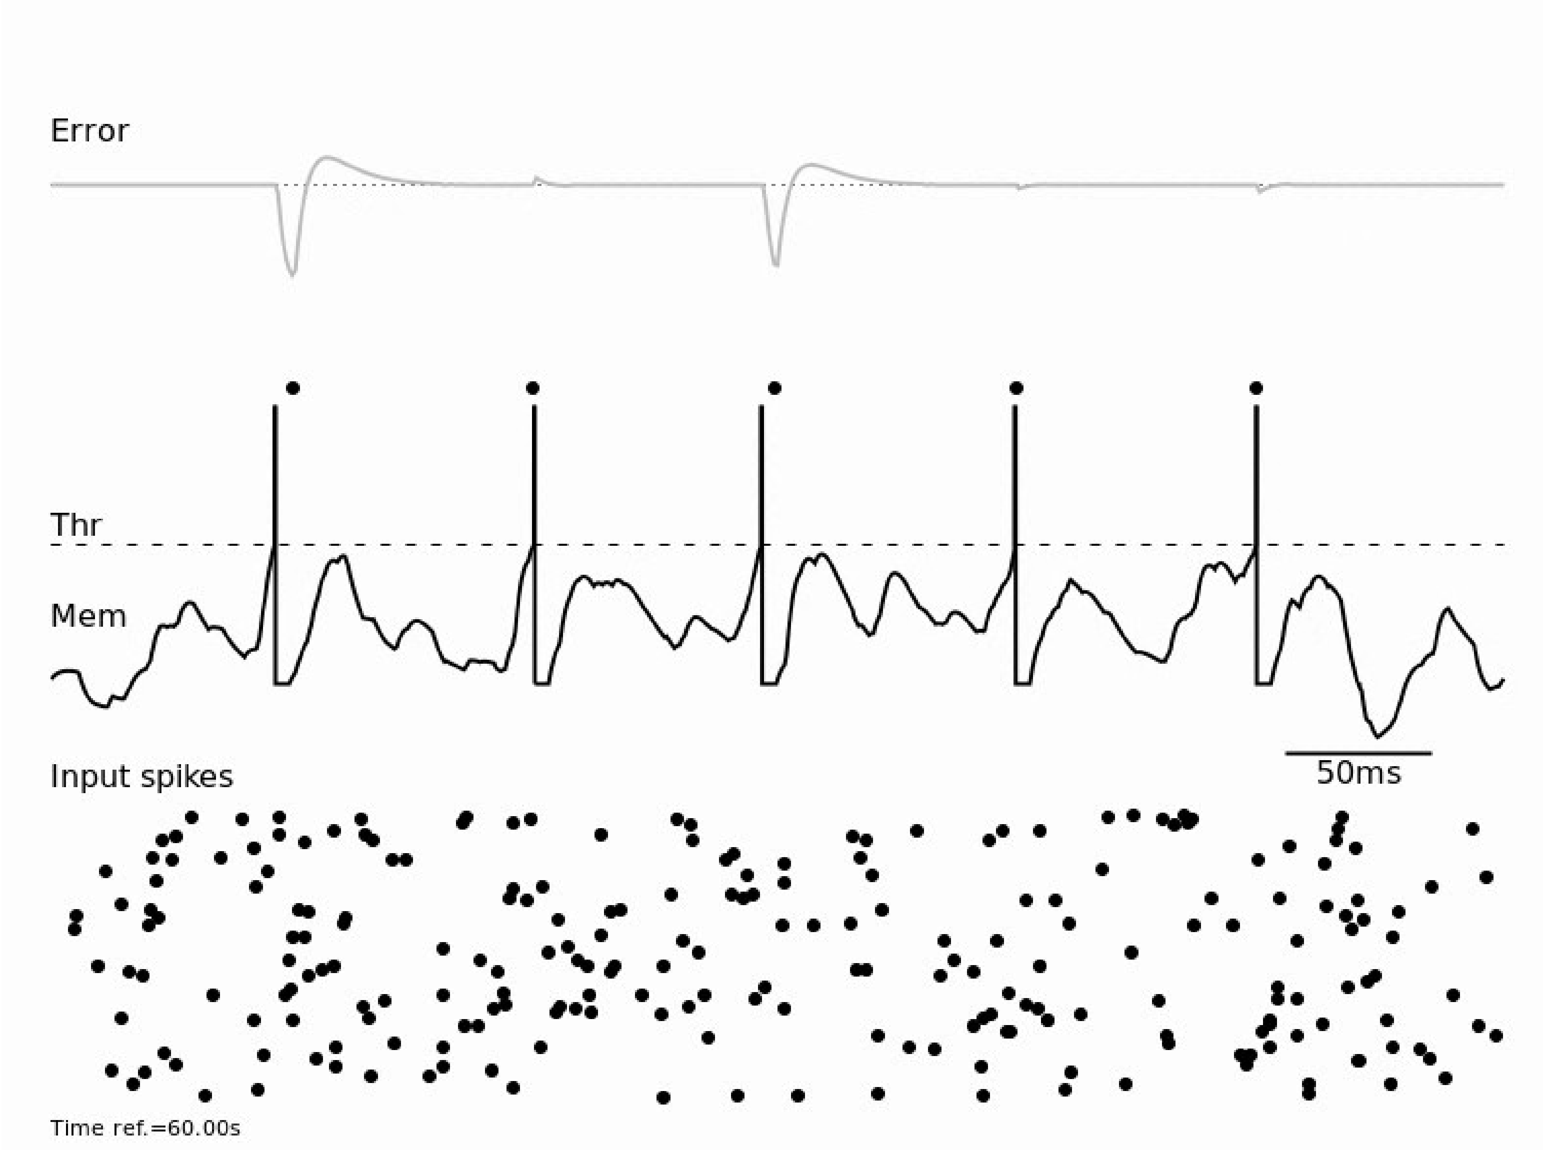

Training many outputs in parallel

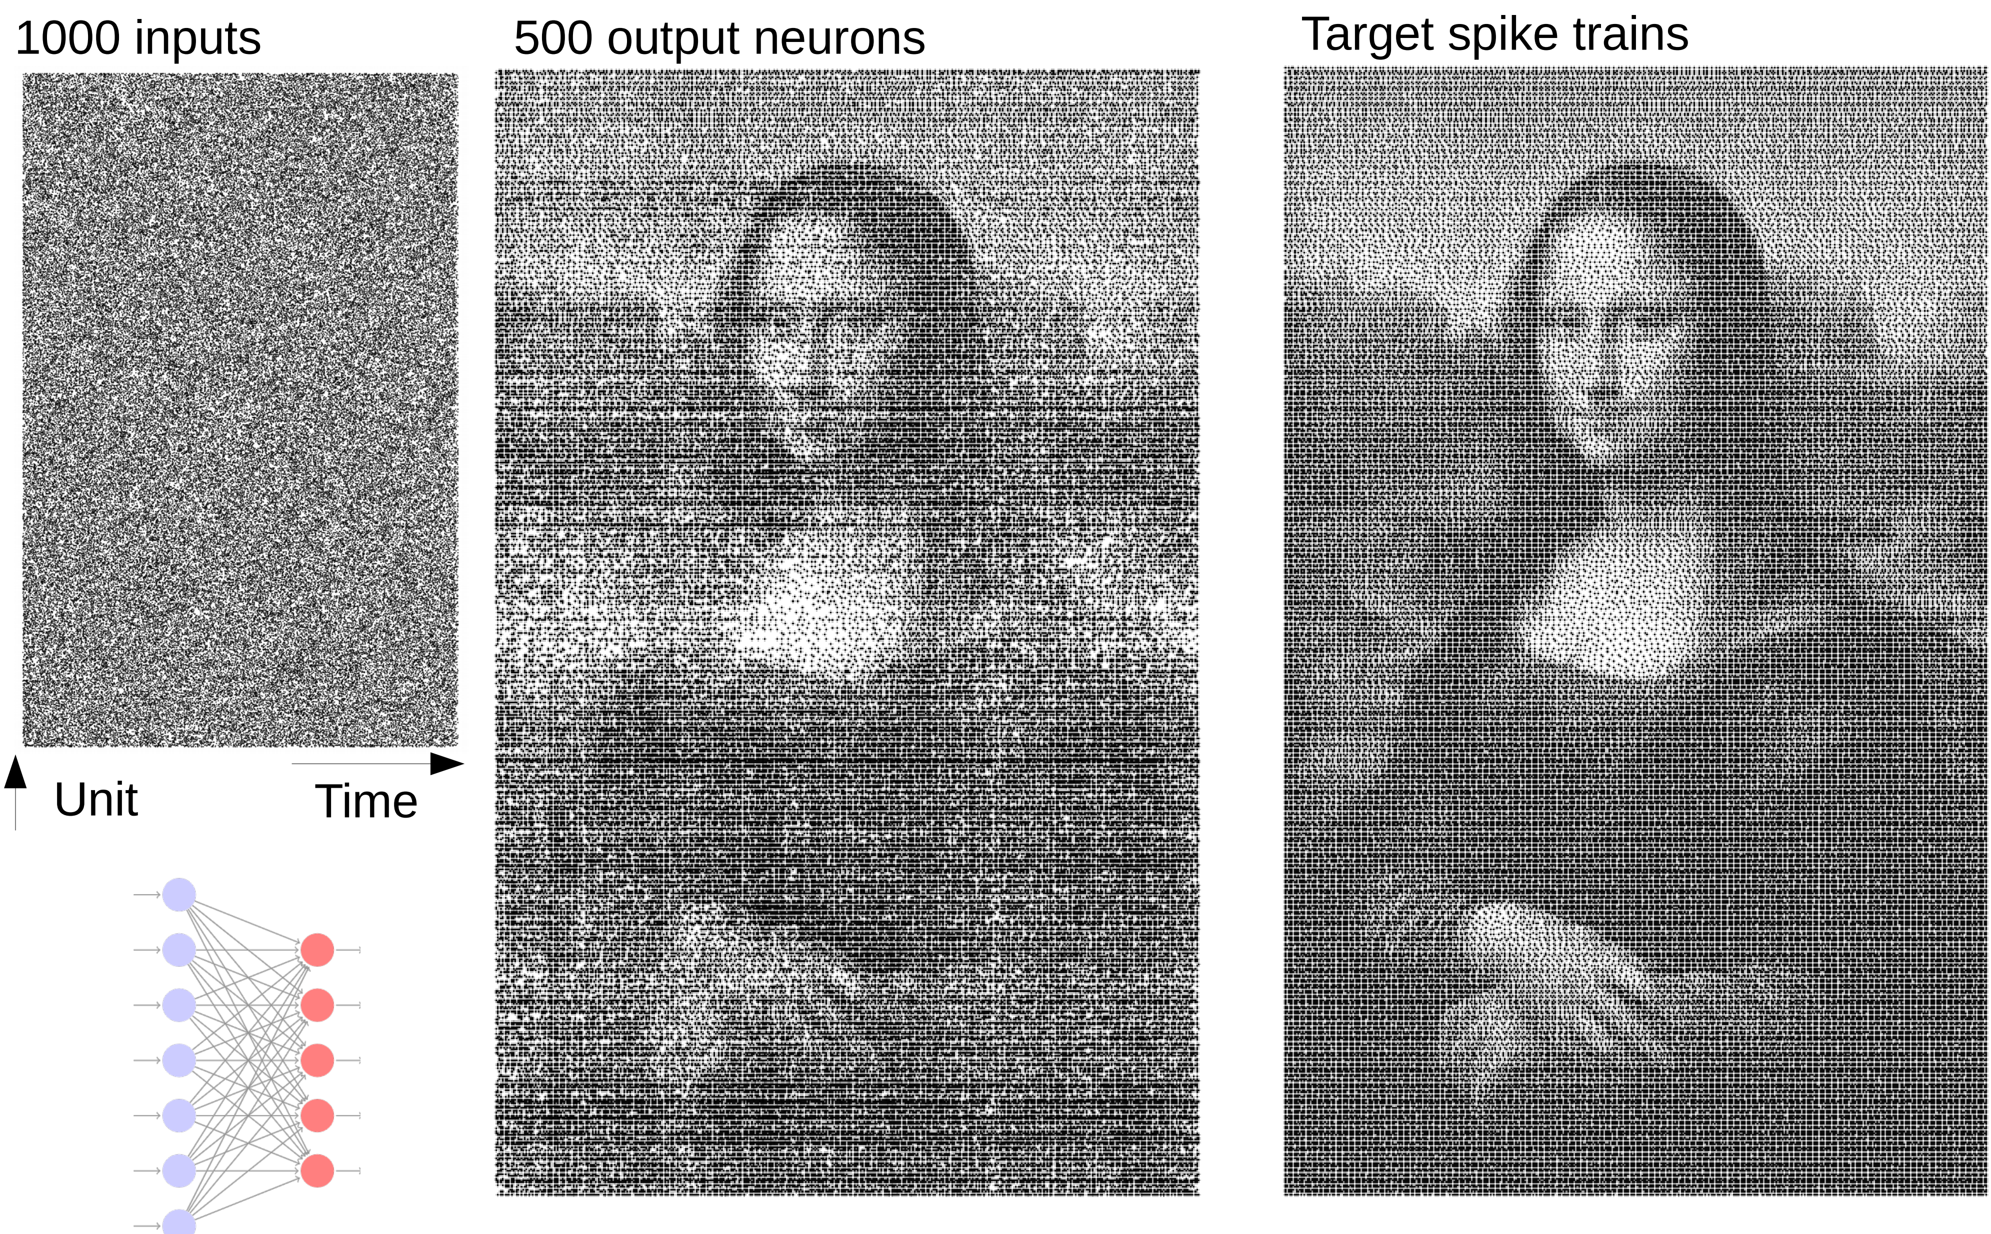

Can we take this to multiple layers?

$$\frac{\partial w_{ij}}{\partial t} \equiv \sum_k e_k(t)\epsilon * \left[w_{ki}\epsilon * \left(\epsilon * S_j(t)\sigma'(U_i)\right)\sigma'(U_k)\right]$$

Problems:
- Symmetric weights
- Downstream activities

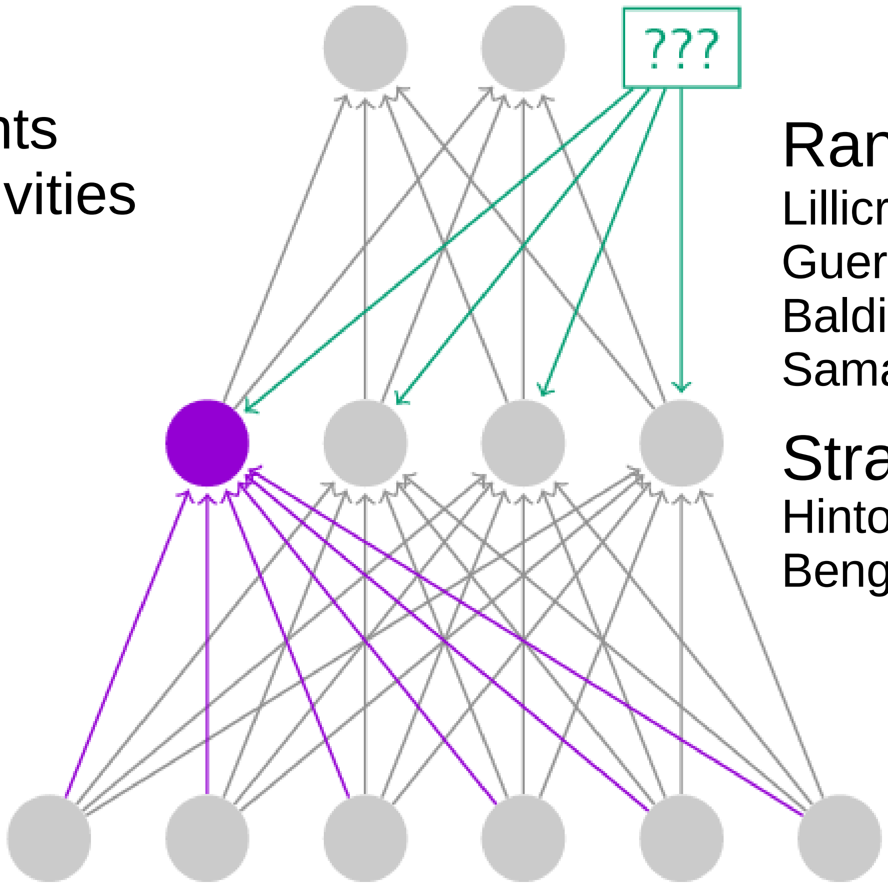

Random feedback

Lillicrap et al. (2014, 2016)
Guergiuev et al. (2016)
Baldi et al. (2016)
Samadi et al. (2017)

Straight-through estim.

Hinton (2012)
Bengio et al. (2013)

Feedback in the hidden layers

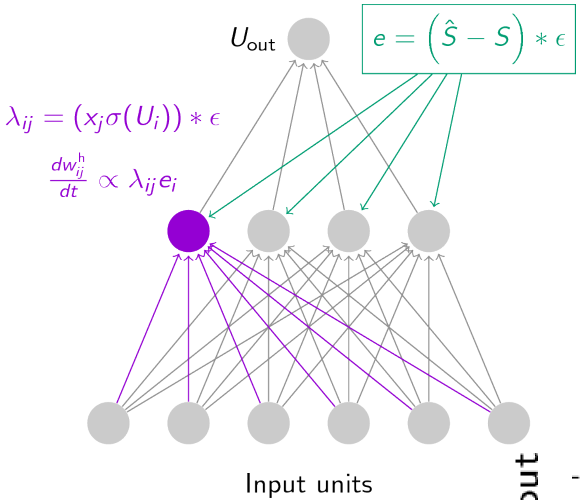

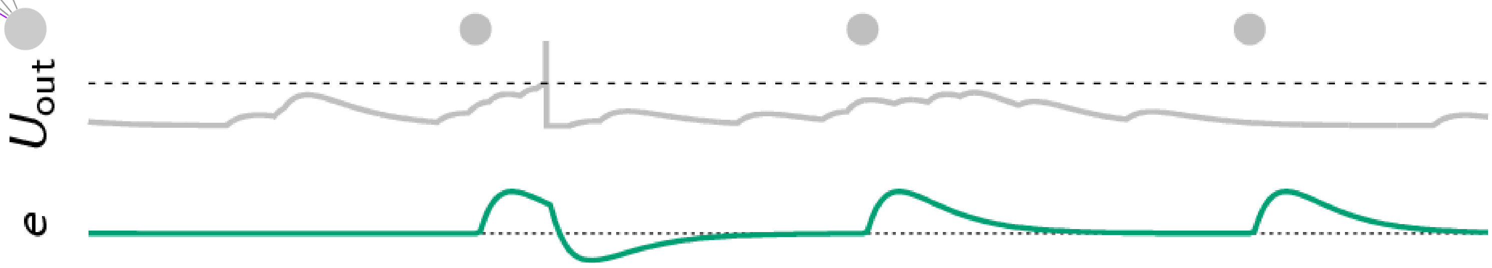

Random feedback, one hidden layer

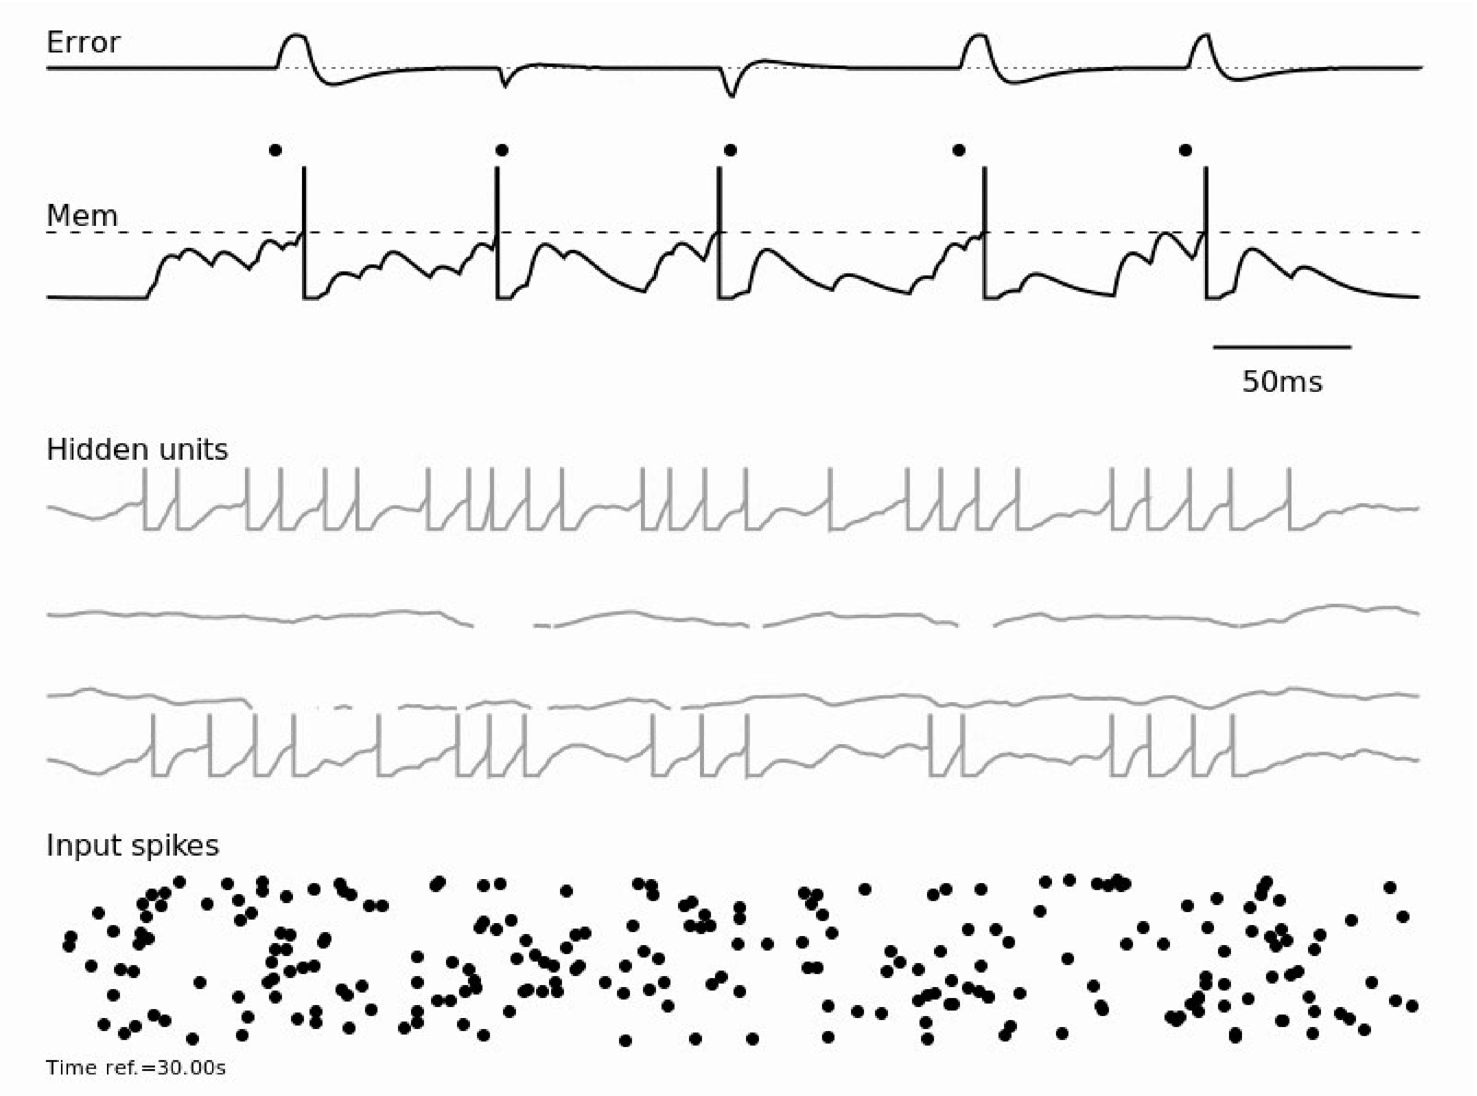

Non-linearly separable problems

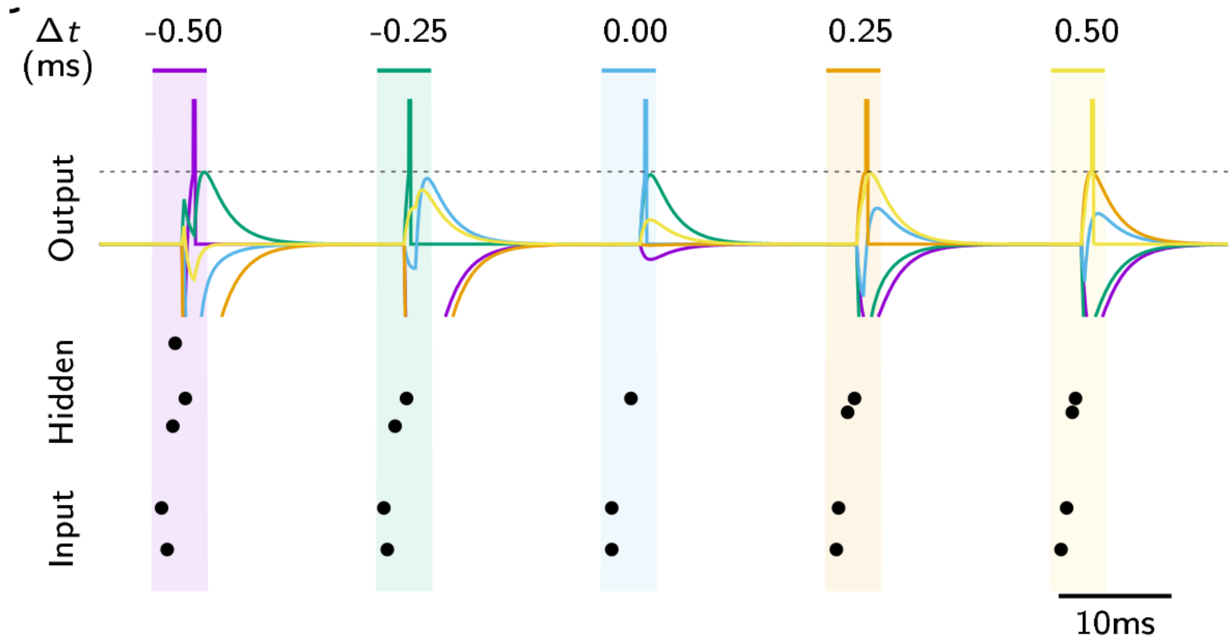

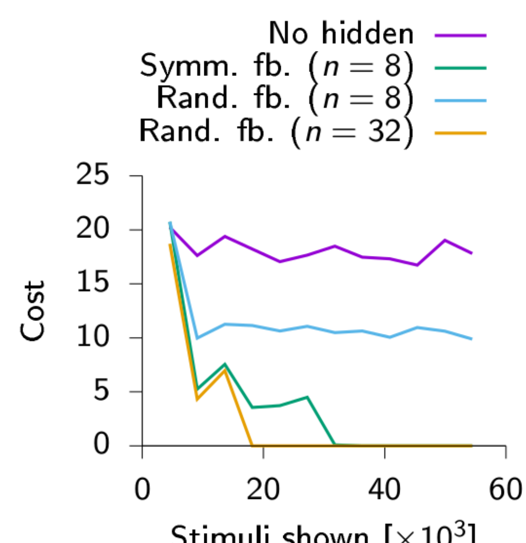

Learning compex spatiotemporal sequences

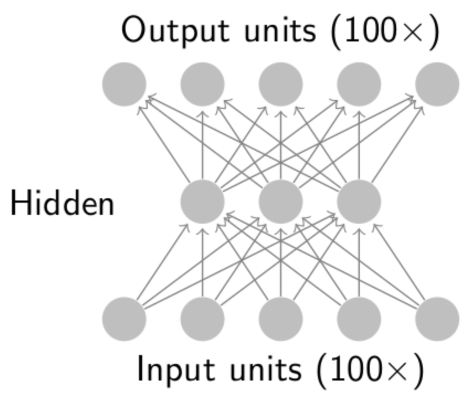

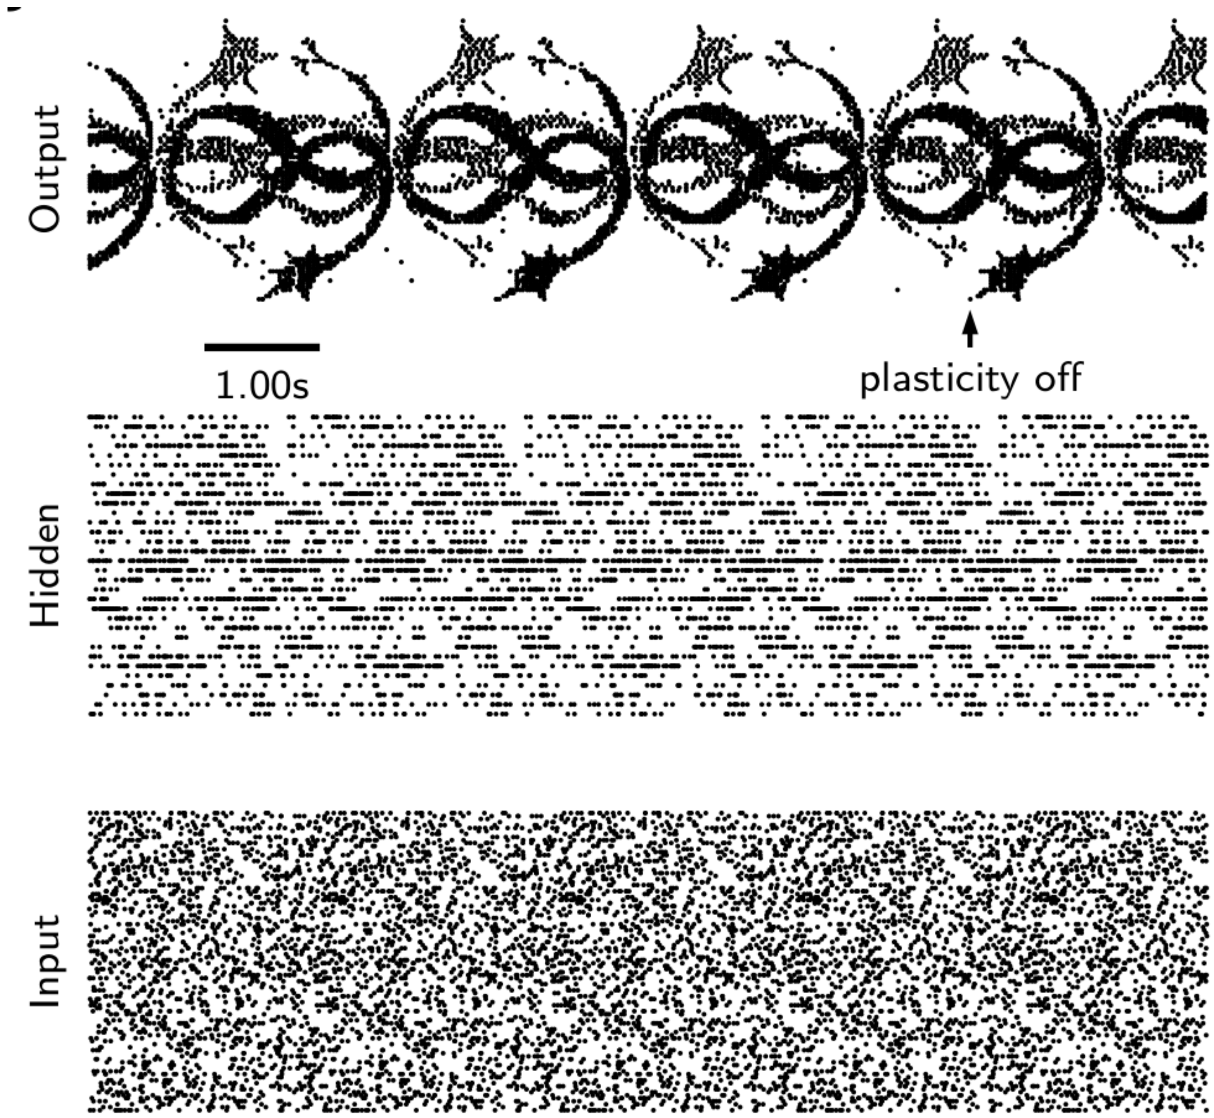

F. Zenke - fzenke.net

Classifying temporally coded input patterns

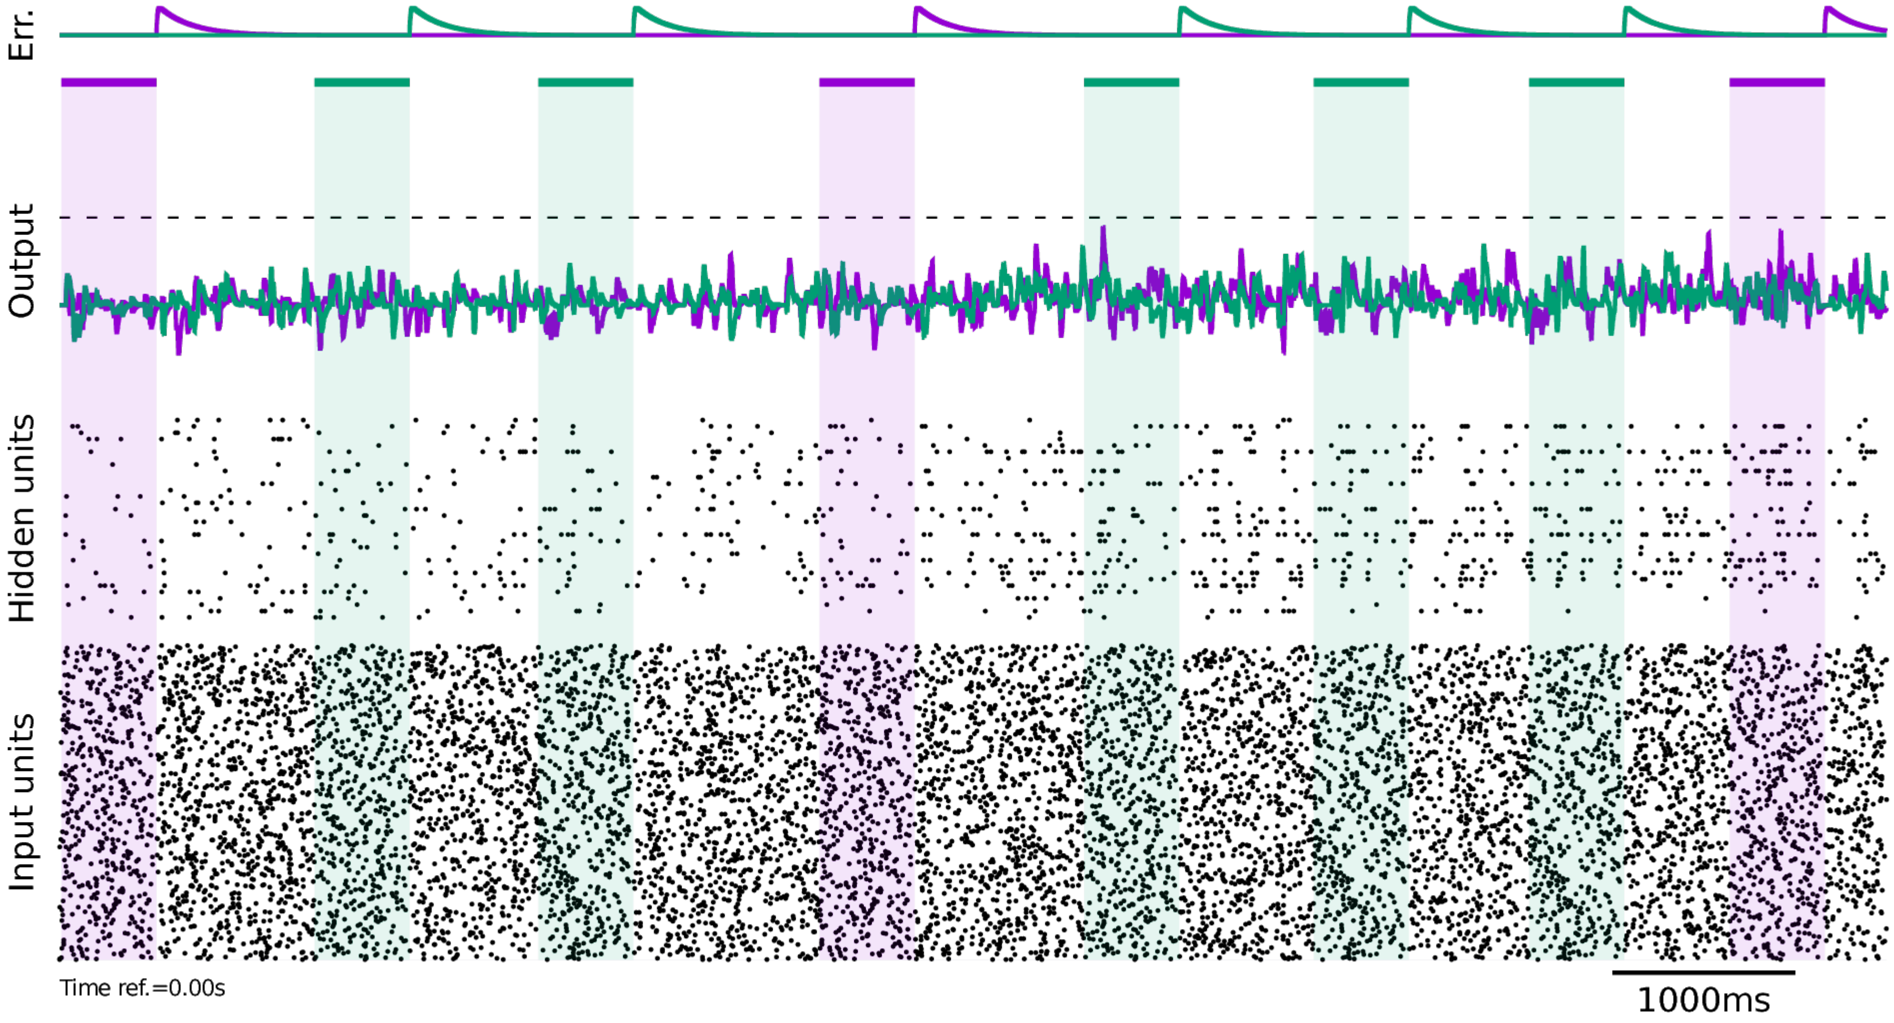

Classifying temporally coded input patterns

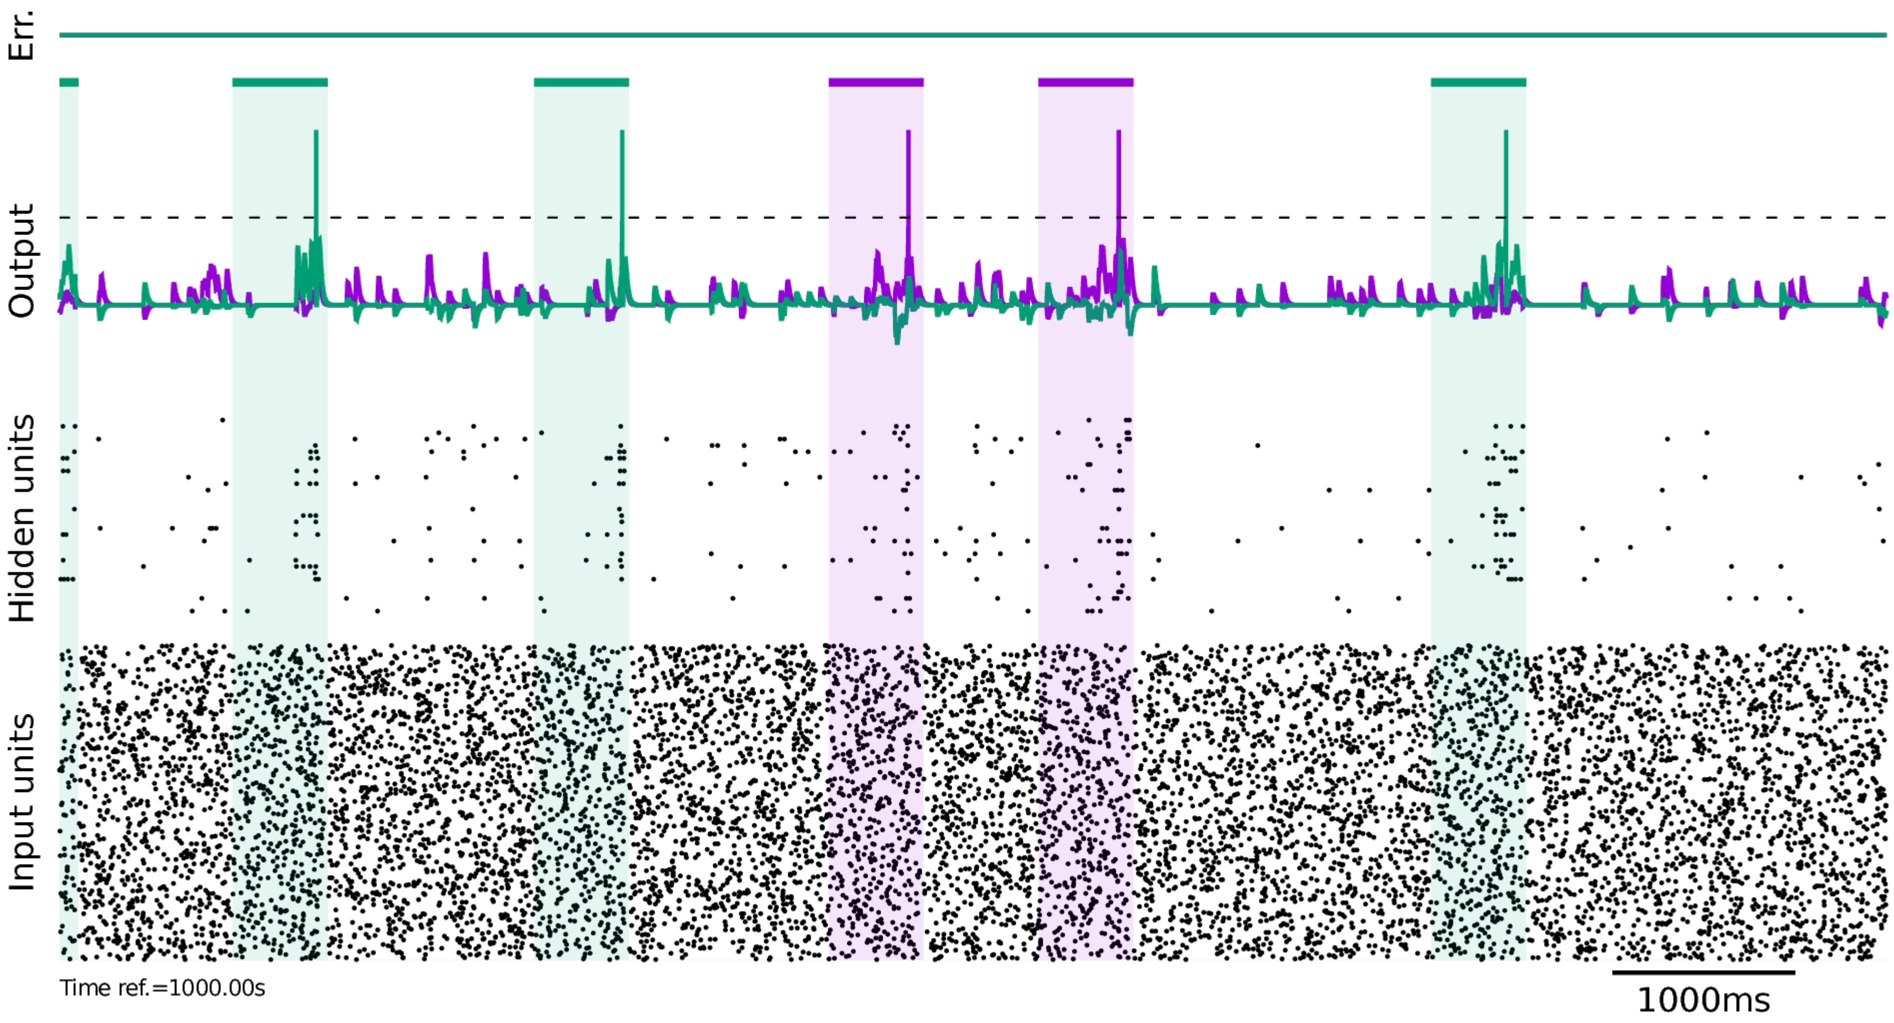

Heterogeneous feedback is important

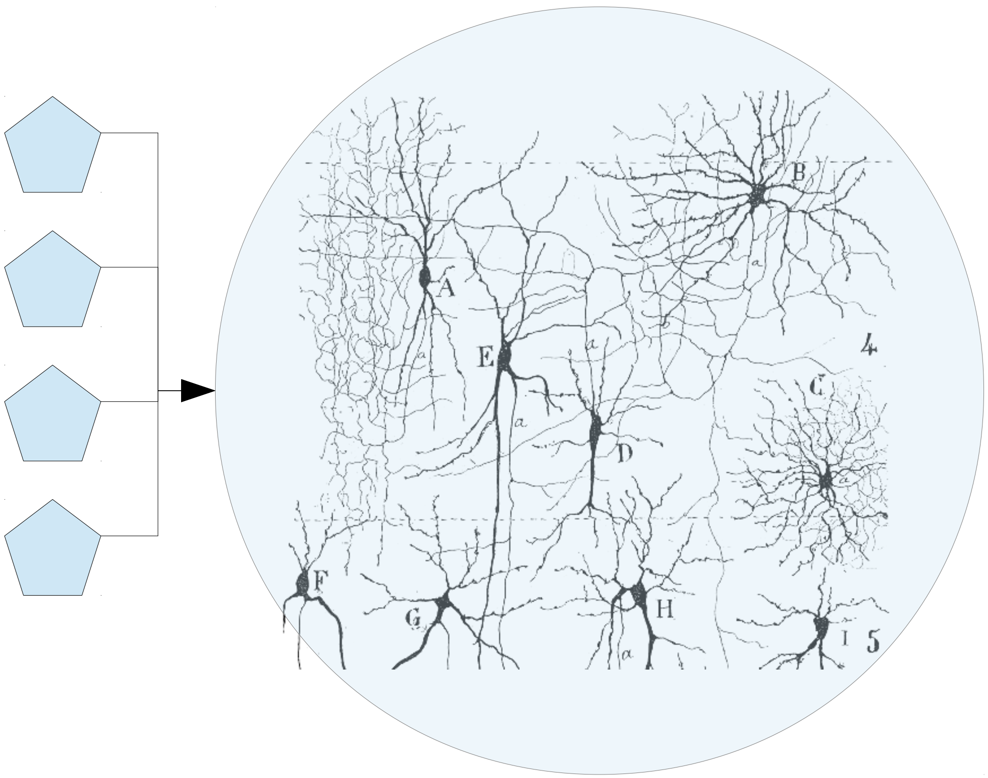

Global third factor

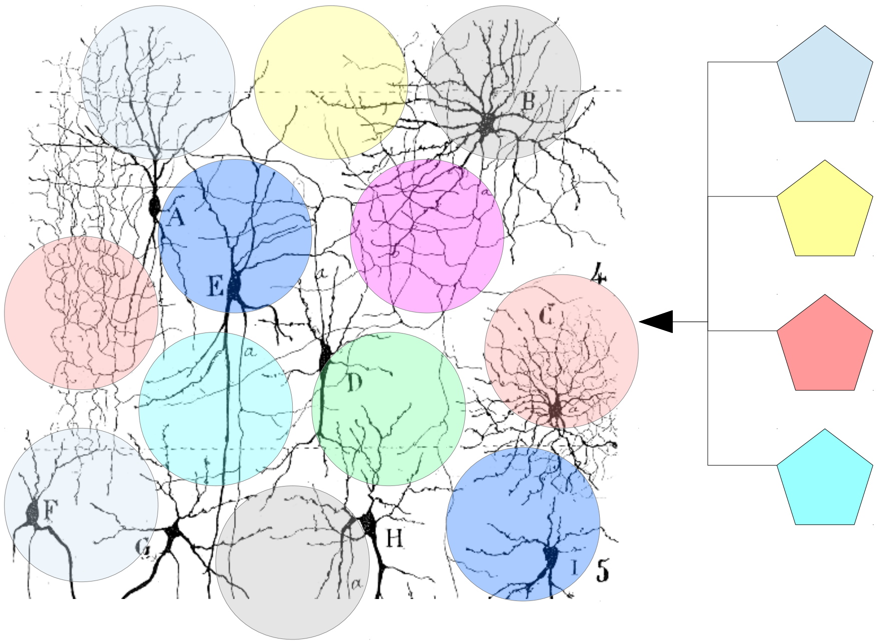

Heterogeneous third factor

Neurons by Ramon y Cajal

Summary

- Started from cost function approach
- Method to teach spiking nets to solve non-trivial temporally coded problems
- Learning rule has a simple interpretation in a biological context

Acknowledgements

- Part 1: Everton Agnes & Wulfram Gerstner
- Part 2: Surya Ganguli & Gang

MARIE CURIE ACTIONS

FNSNF

SWISS NATIONAL SCIENCE FOUNDATION

wellcome

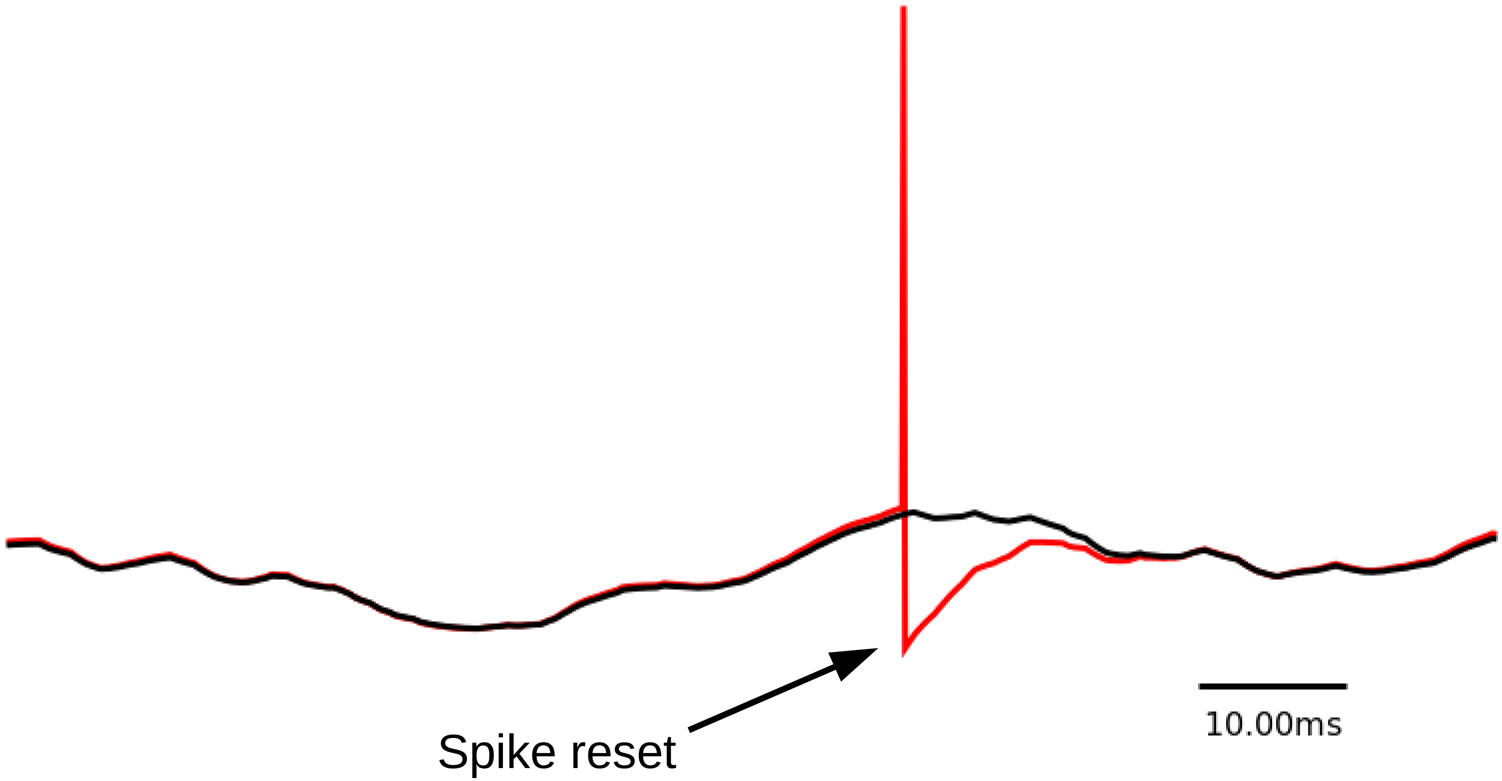

$$S(U_i(t)) \propto \Theta(U_i(t)) \approx \sigma(U_i(t))$$

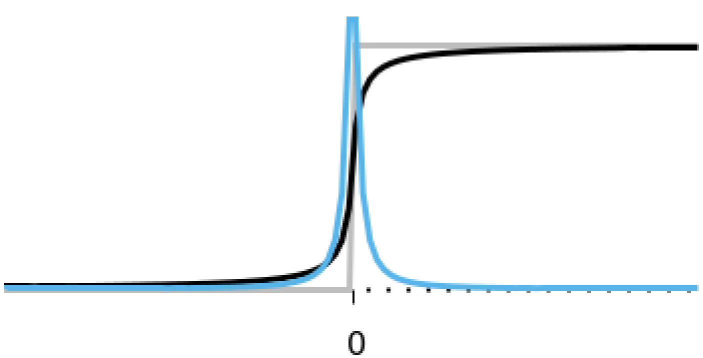

$$\frac{\partial S_i}{\partial w_{ij}} \approx \sigma'(U_i) \frac{\partial U_i}{\partial w_{ij}}$$

need $\frac{\partial U_i}{\partial w_{ij}}$

- Leaky integrate-and-fire neuron model
- Dynamics defined by ODE
- Equivalent to the spike-response-model (SRM0)
- Can be written with temporal convolutions

PSP kernel ϵ

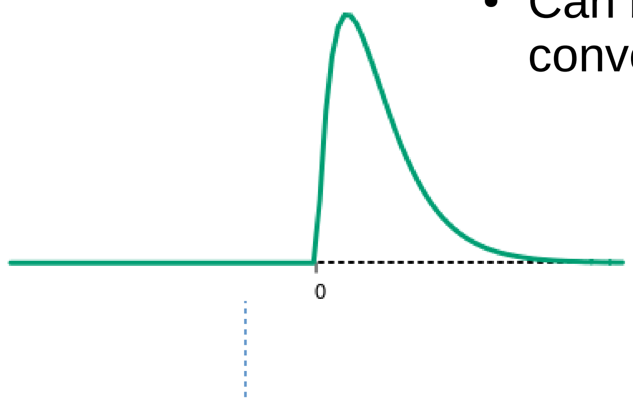

Membrane voltage

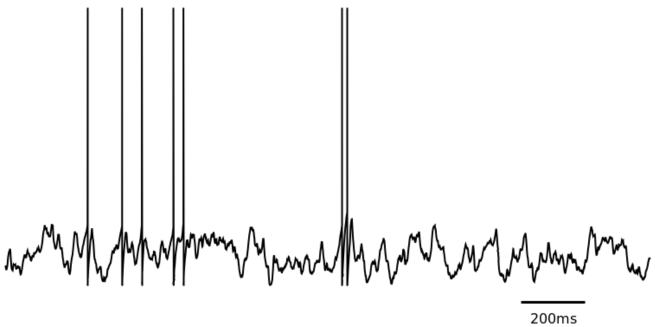

Reset kernel η

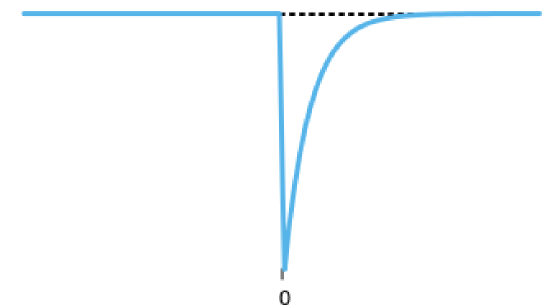

$$U_i(t) = \sum_j w_{ij} \epsilon * S_j(t) + \eta * S_i(t)$$

$$\frac{\partial S_i}{\partial w_{ij}} \approx \underbrace{\sigma'(U_i)}_{\text{post}} \underbrace{\epsilon * S_j(t)}_{\text{pre}}$$

The learning rule can be interpreted as a three factor rule

$$-\frac{\partial L}{\partial w_k} = \int_{-\infty}^{t} \underbrace{\epsilon * \left(\hat{S}_i(s) - S_i(s)\right)}_{\equiv \mathrm{e_i(t)}} \epsilon * \frac{\partial S_i(t)}{\partial w_k} ds$$

$$\frac{\partial w_{ij}}{\partial t} \equiv \underbrace{e_i(t)}_{\text{error signal}} \quad \epsilon * (\underbrace{\epsilon * S_j(t)}_{\text{pre}} \underbrace{\sigma'(U_i)}_{\text{post}})$$

- Three factor rule
- Non-linear Hebbian
- Voltage-based
- Think of outer convolution as eligibility trace

Which spike train metric to use

Spike train:

$$S_i(t) = \sum_k \delta(t - t_i^k)$$

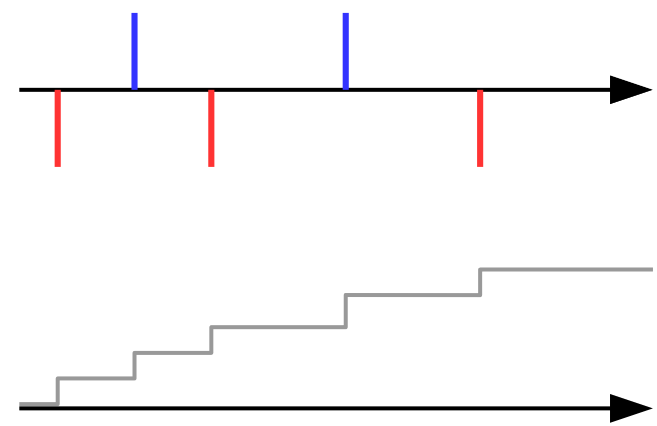

$$L_{\text{naive}} = \int_{-\infty}^{t} \left| \hat{S}_i(s) - S_i(s) \right| ds$$

van Rossum distance:

$$L = \int_{-\infty}^{t} \Big(\epsilon * \hat{S}_i(s) - \epsilon * S_i(s) \Big)^2 ds$$

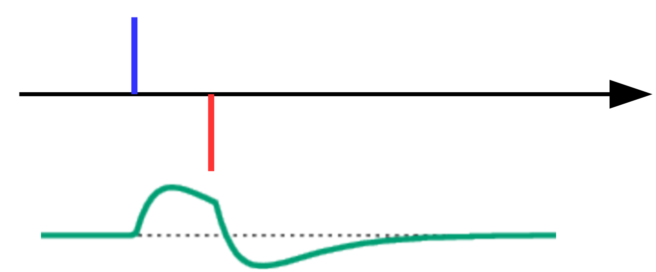

Gardner & Grunig 2015

Derivative of a spike train is zero almost everywhere

$$L = \frac{1}{2} \int_{-\infty}^{t} \left(\epsilon * \left(\hat{S}_i(s) - S_i(s) \right) \right)^2 ds$$

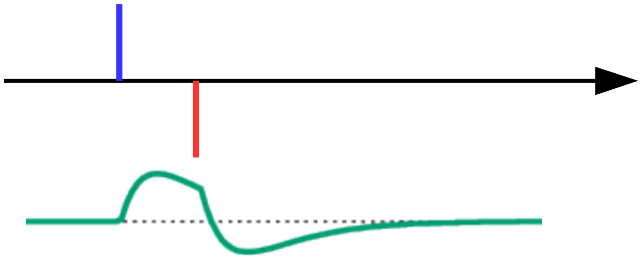

D'Oh!

$$-\frac{\partial L}{\partial w_k} = \int_{-\infty}^{t} \epsilon * \left(\hat{S}_i(s) - S_i(s) \right) \epsilon * \frac{\partial S_i(t)}{\partial w_k} ds$$

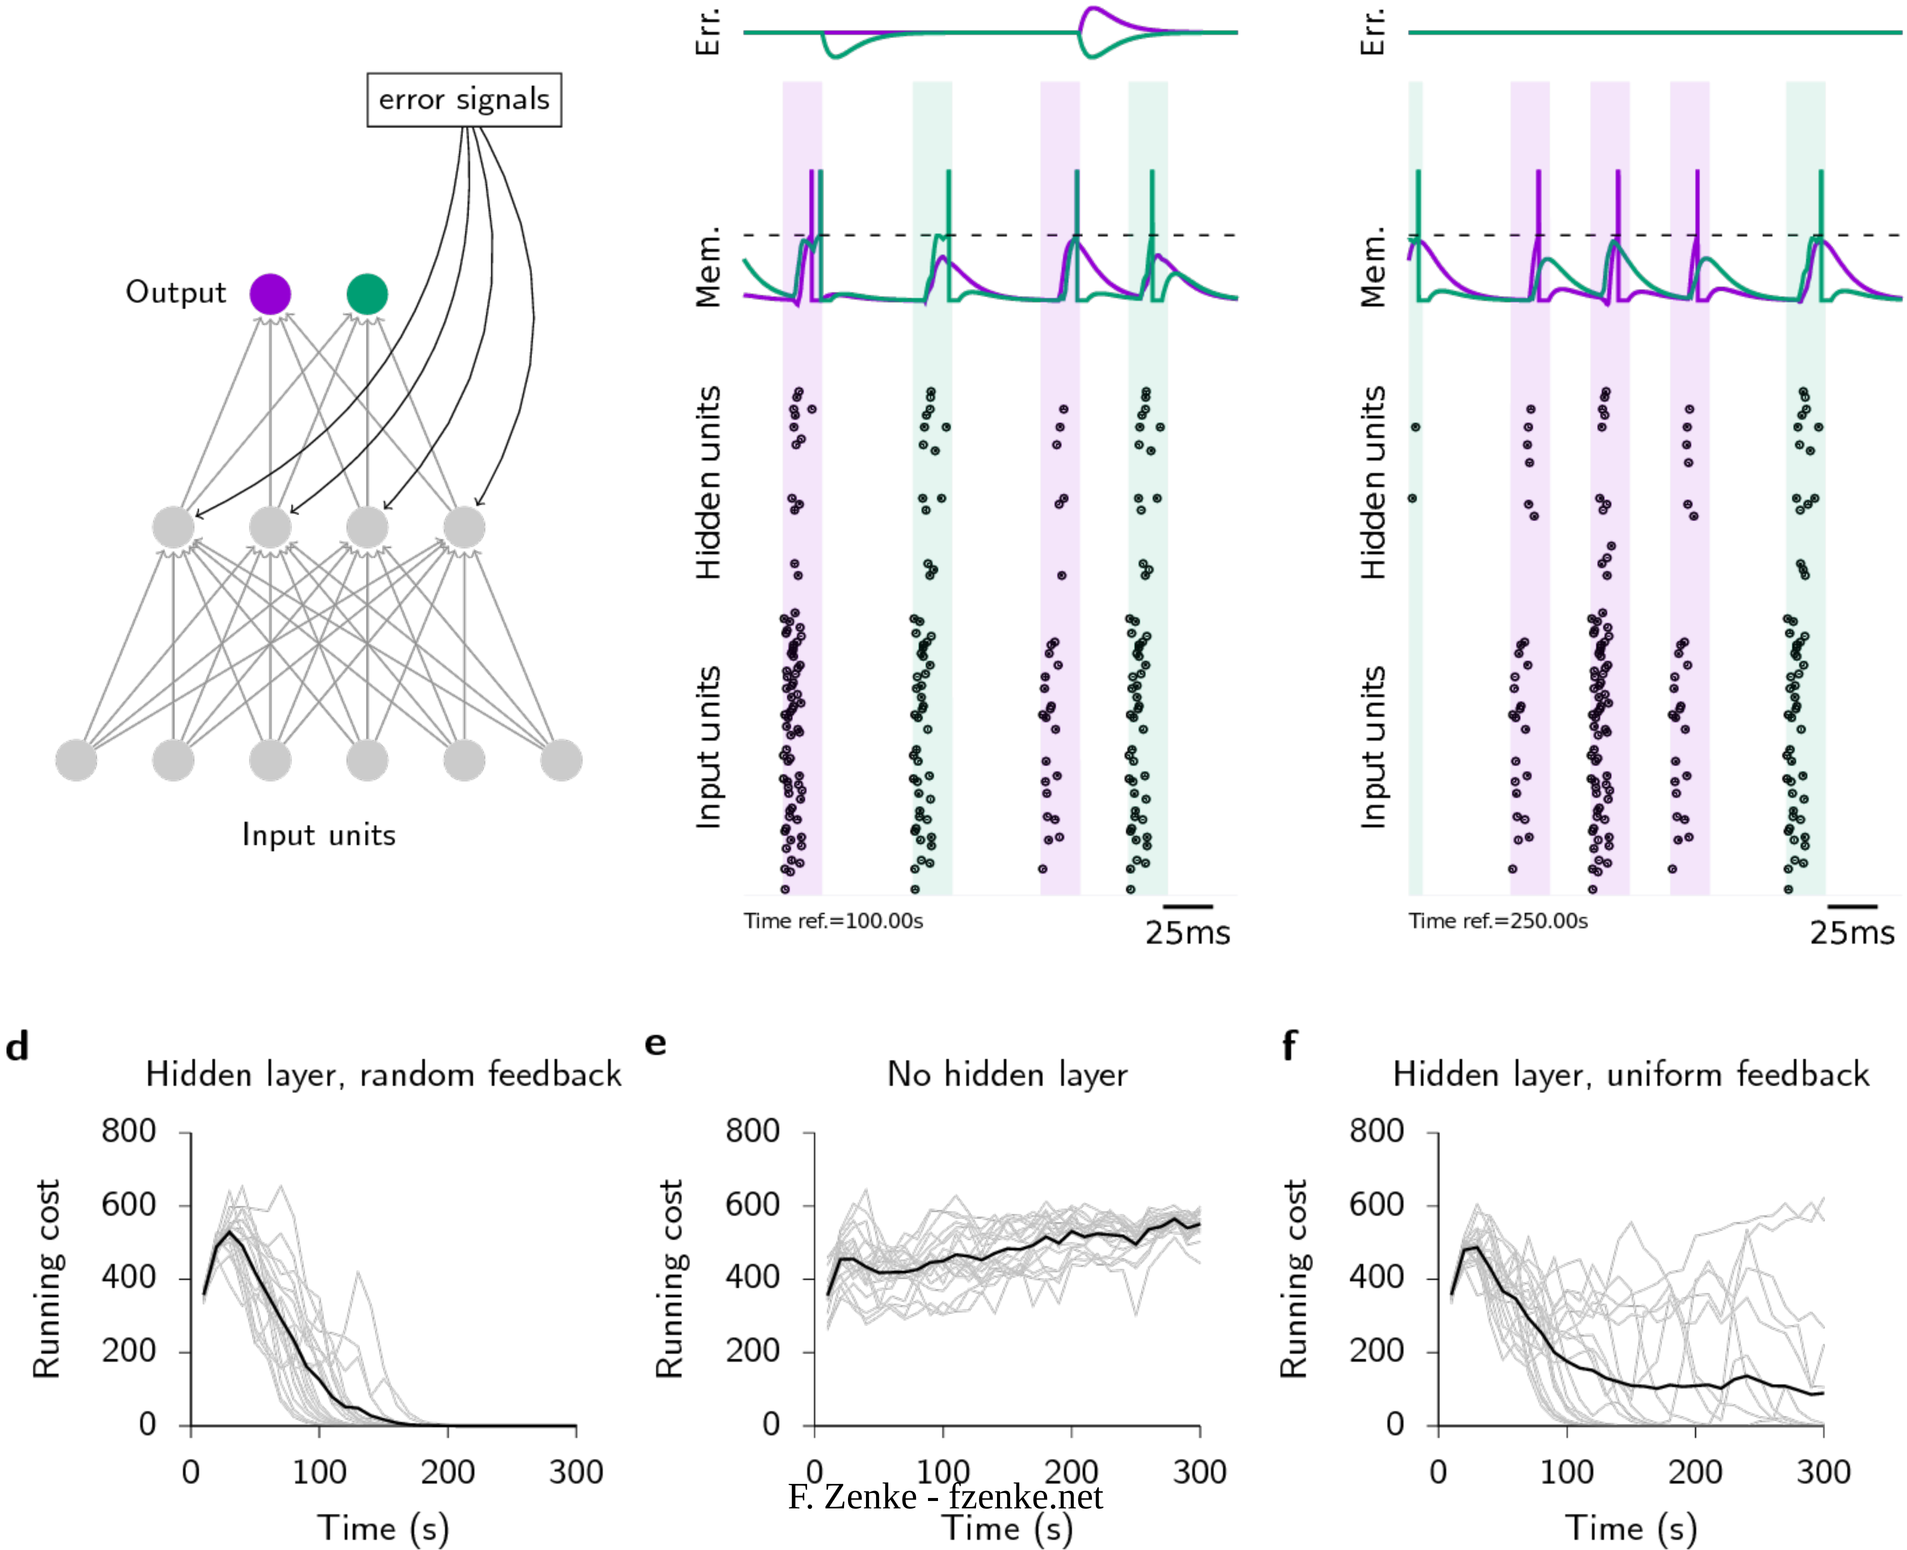
error signals
Output
Input units
Err.
Mem.
Hidden units
Input units
Time ref.=100.00s
25ms
Time ref.=250.00s
25ms
d
Hidden layer, random feedback
e
No hidden layer
f
Hidden layer, uniform feedback
Running cost
800
600
400
200
0
0
100
200
300
Time (s)

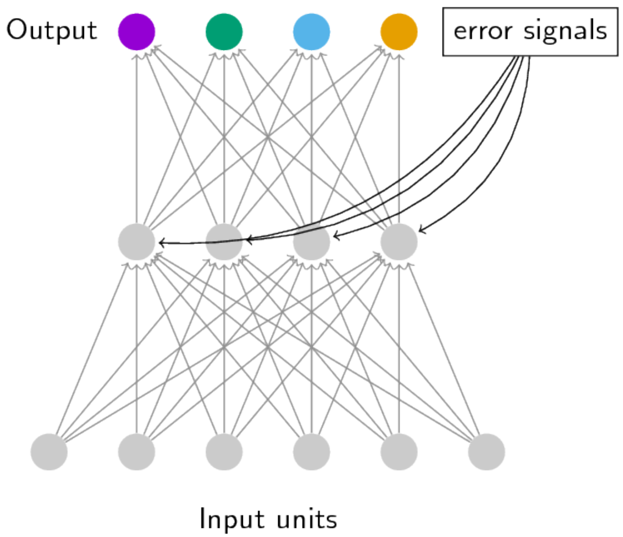
Output
error signals
Input units

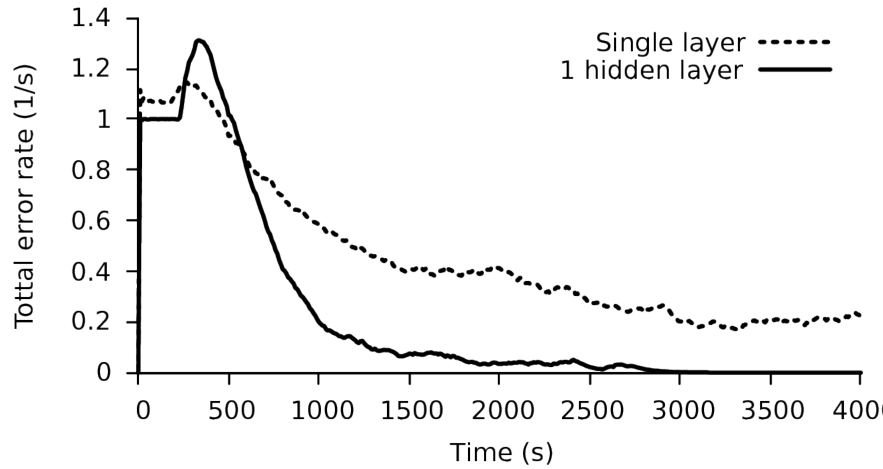
Single layer
1 hidden layer
Tottal error rate (1/s)
Time (s)
1.4
1.2
1
0.8
0.6
0.4
0.2
0
0
500
1000
1500
2000
2500
3000
3500
4000

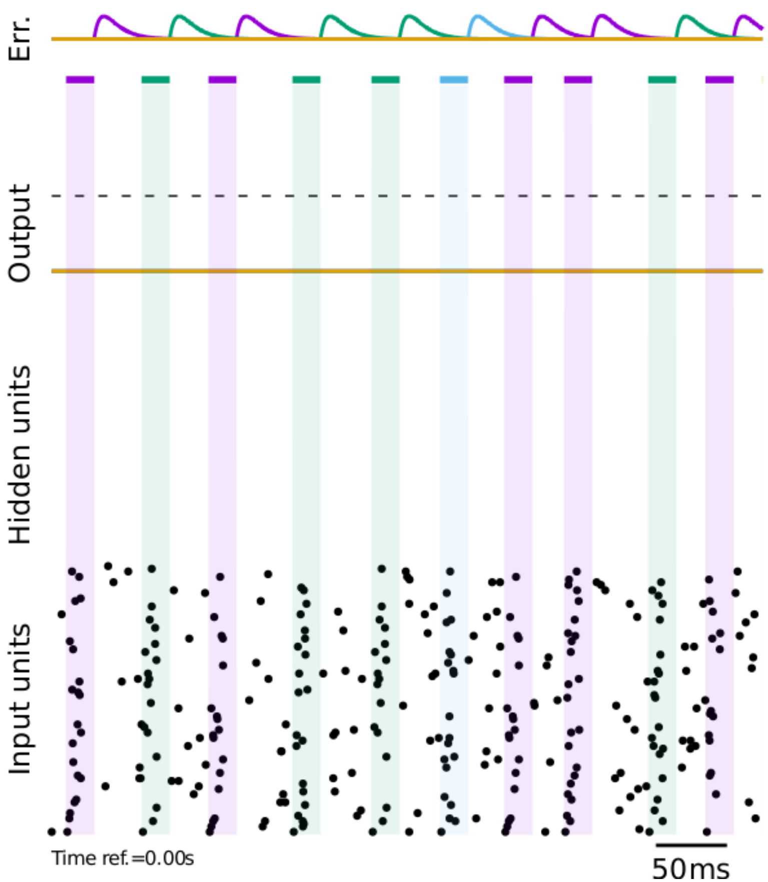
Err.
Output
Hidden units
Input units
Time ref.=0.00s
50ms

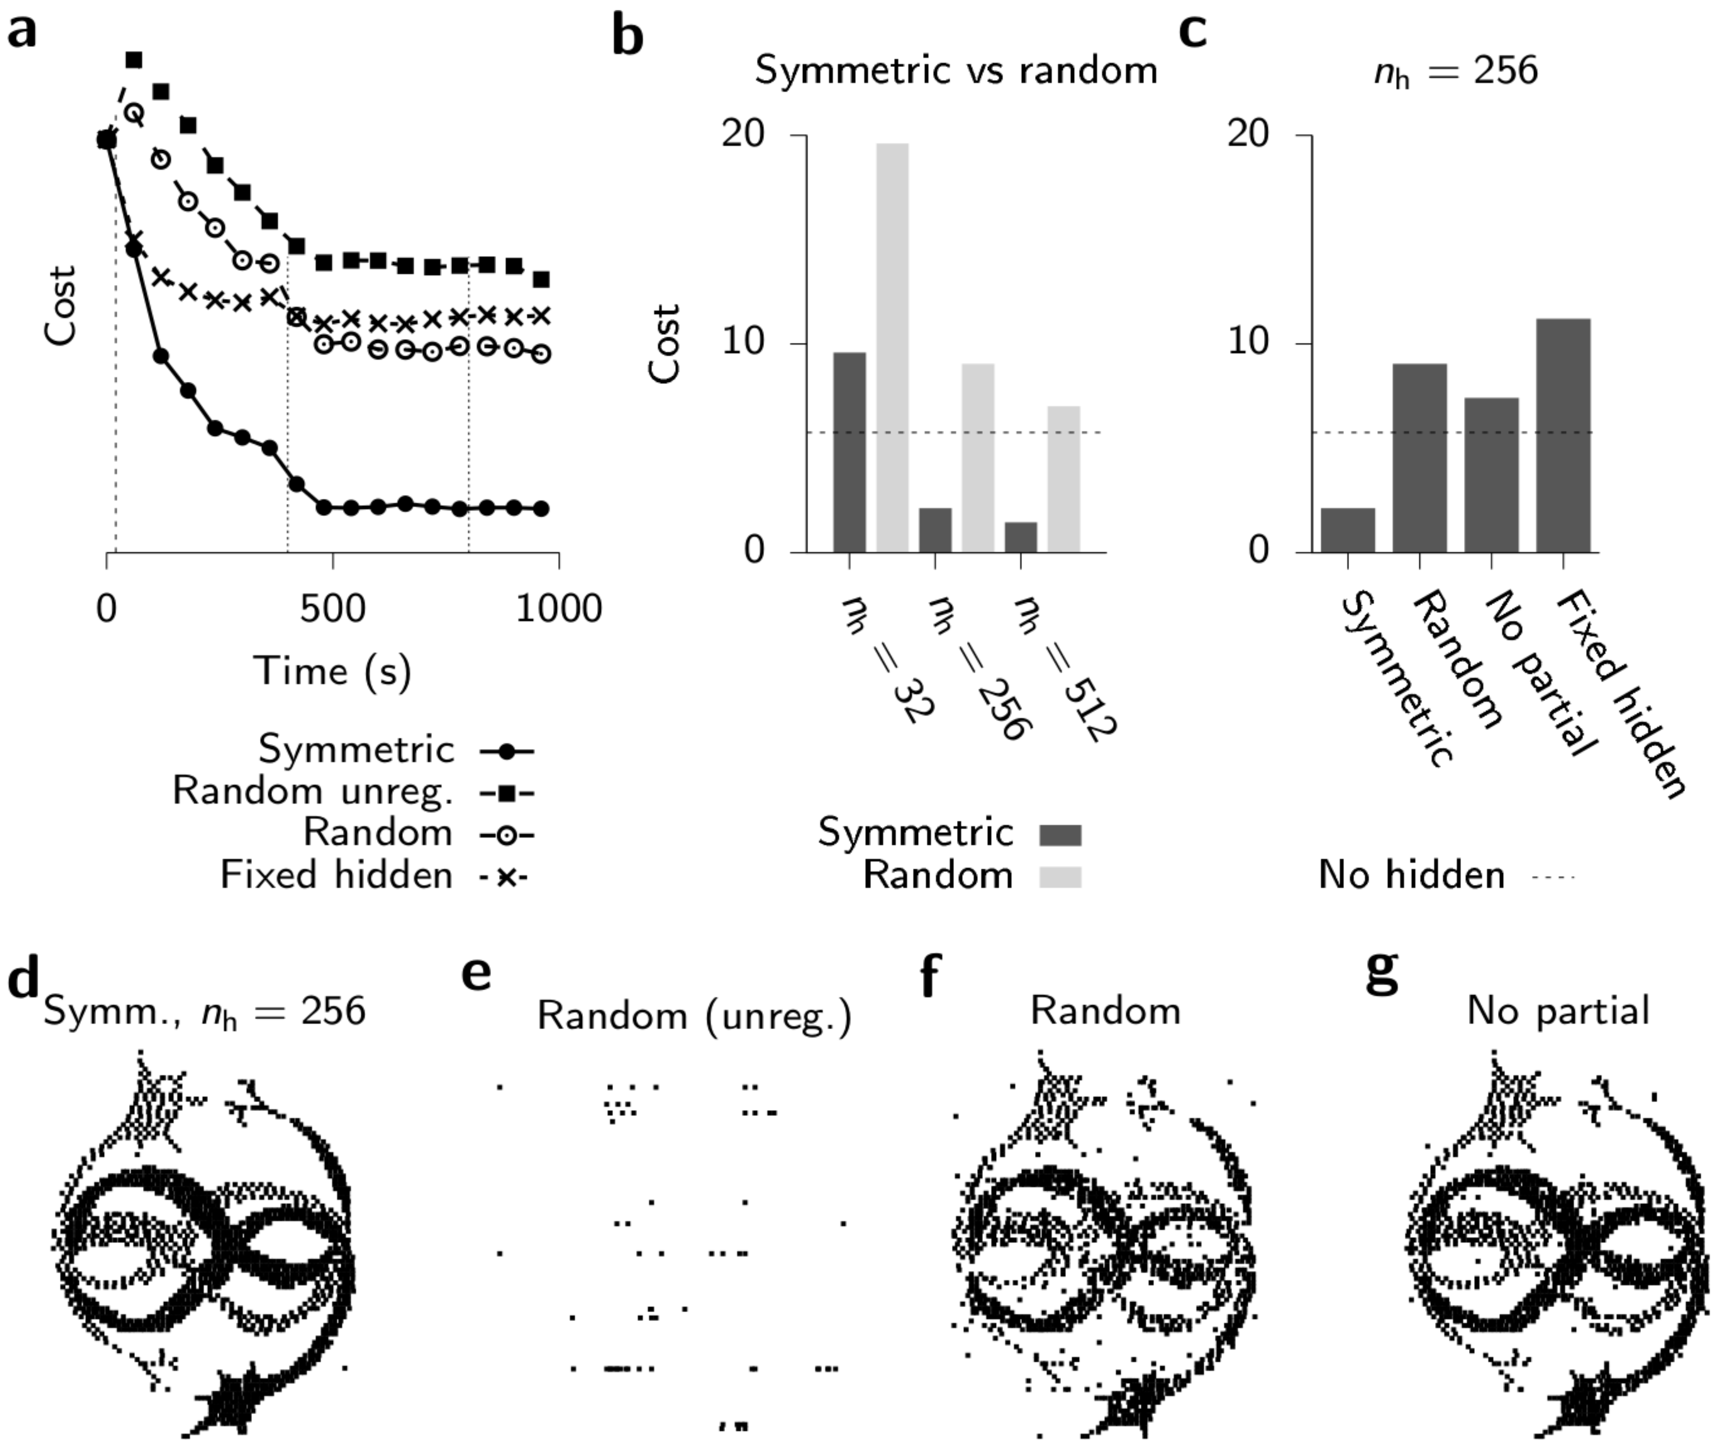
a
Cost
0
500
1000
Time (s)
Symmetric
Random unreg.
Random
Fixed hidden
b
Symmetric vs random
20
10
0
Cost
$n_h = 32$
$n_h = 256$
$n_h = 512$
Symmetric
Random
c
$n_h = 256$
20
10
0
Symmetric
Random
No partial
Fixed hidden
No hidden
d
Symm., $n_h = 256$
e
Random (unreg.)
f
Random
g
No partial

Subspace learning

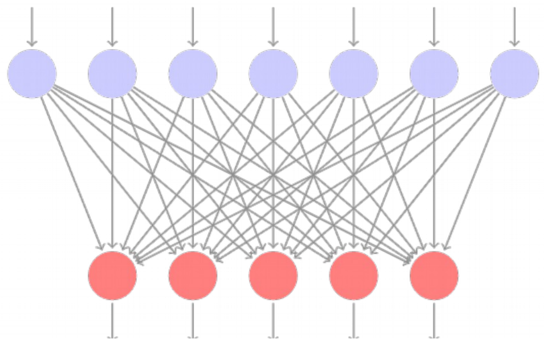

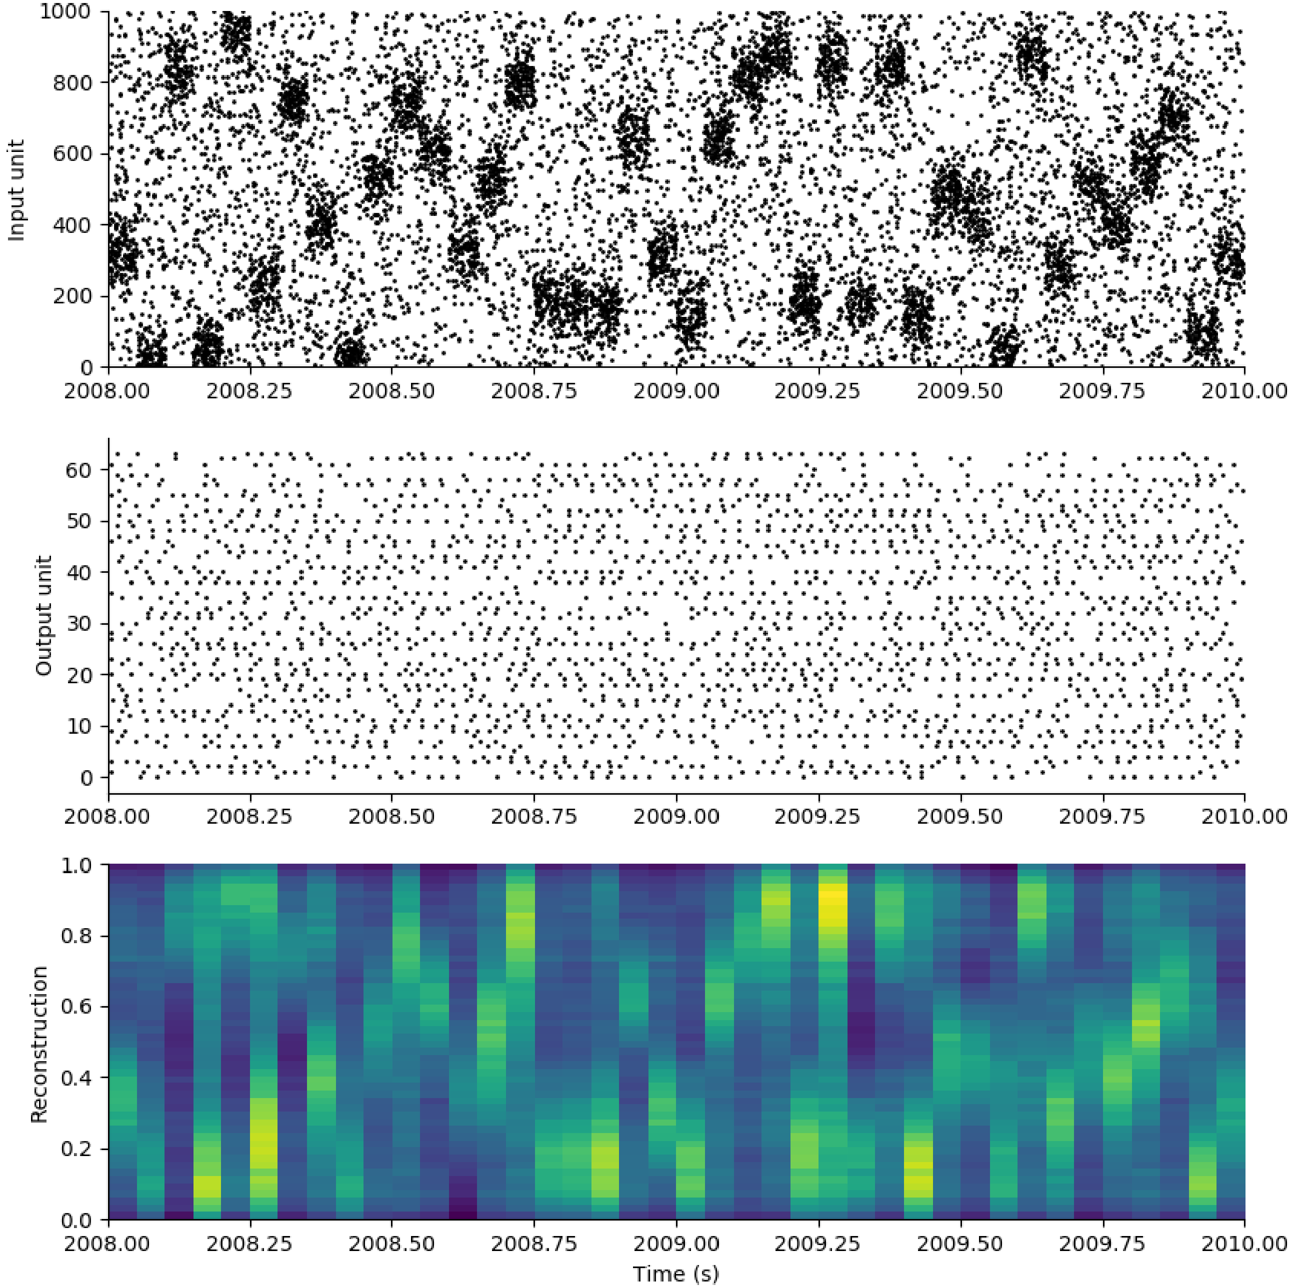

Random feedback, one hidden layer

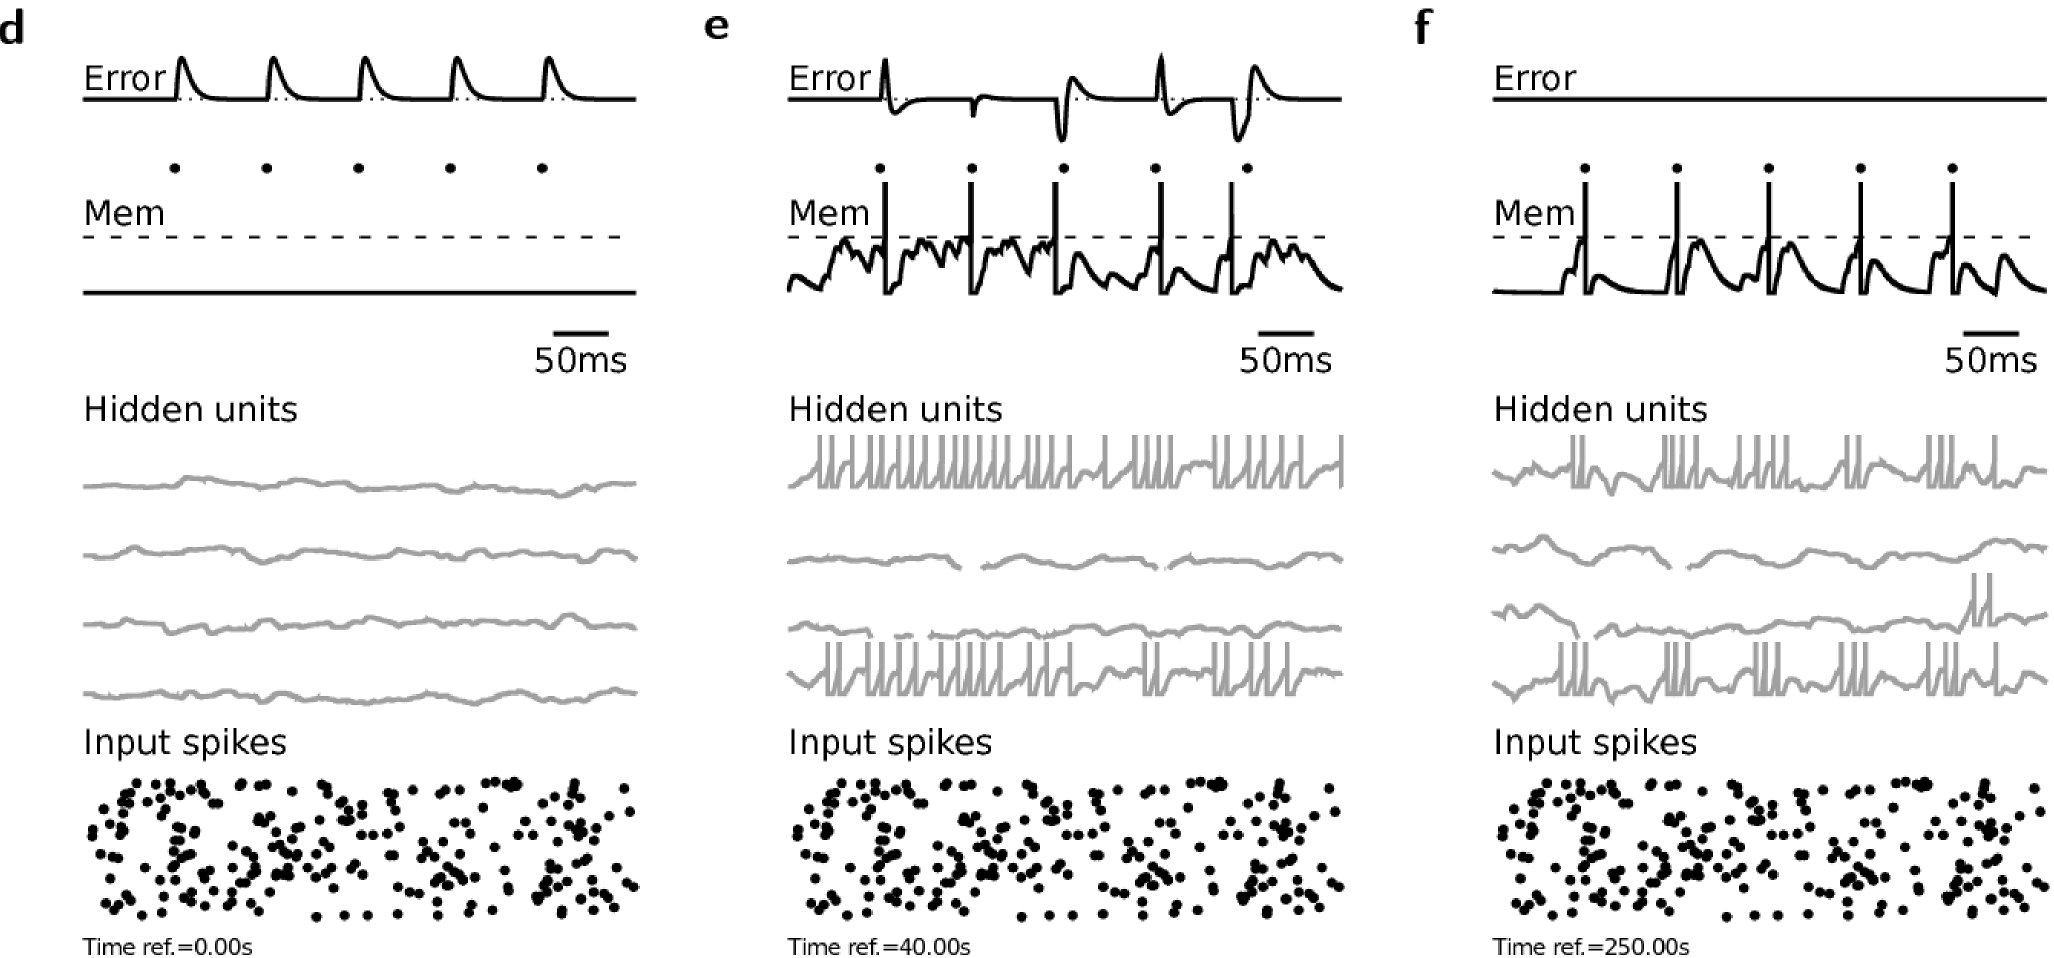

Why Inhibitory Plasticity cannot solve all the problems

- LTP/LTD are not reversed by LTDi/LTPi
  →edge effects
- ISP would have to be fast
  → Experimental evidence suggests that ISP follows plasticity at excitatory synapses
  → If it was fast: anti memories
- Presumably important for de-correlation (and balance)

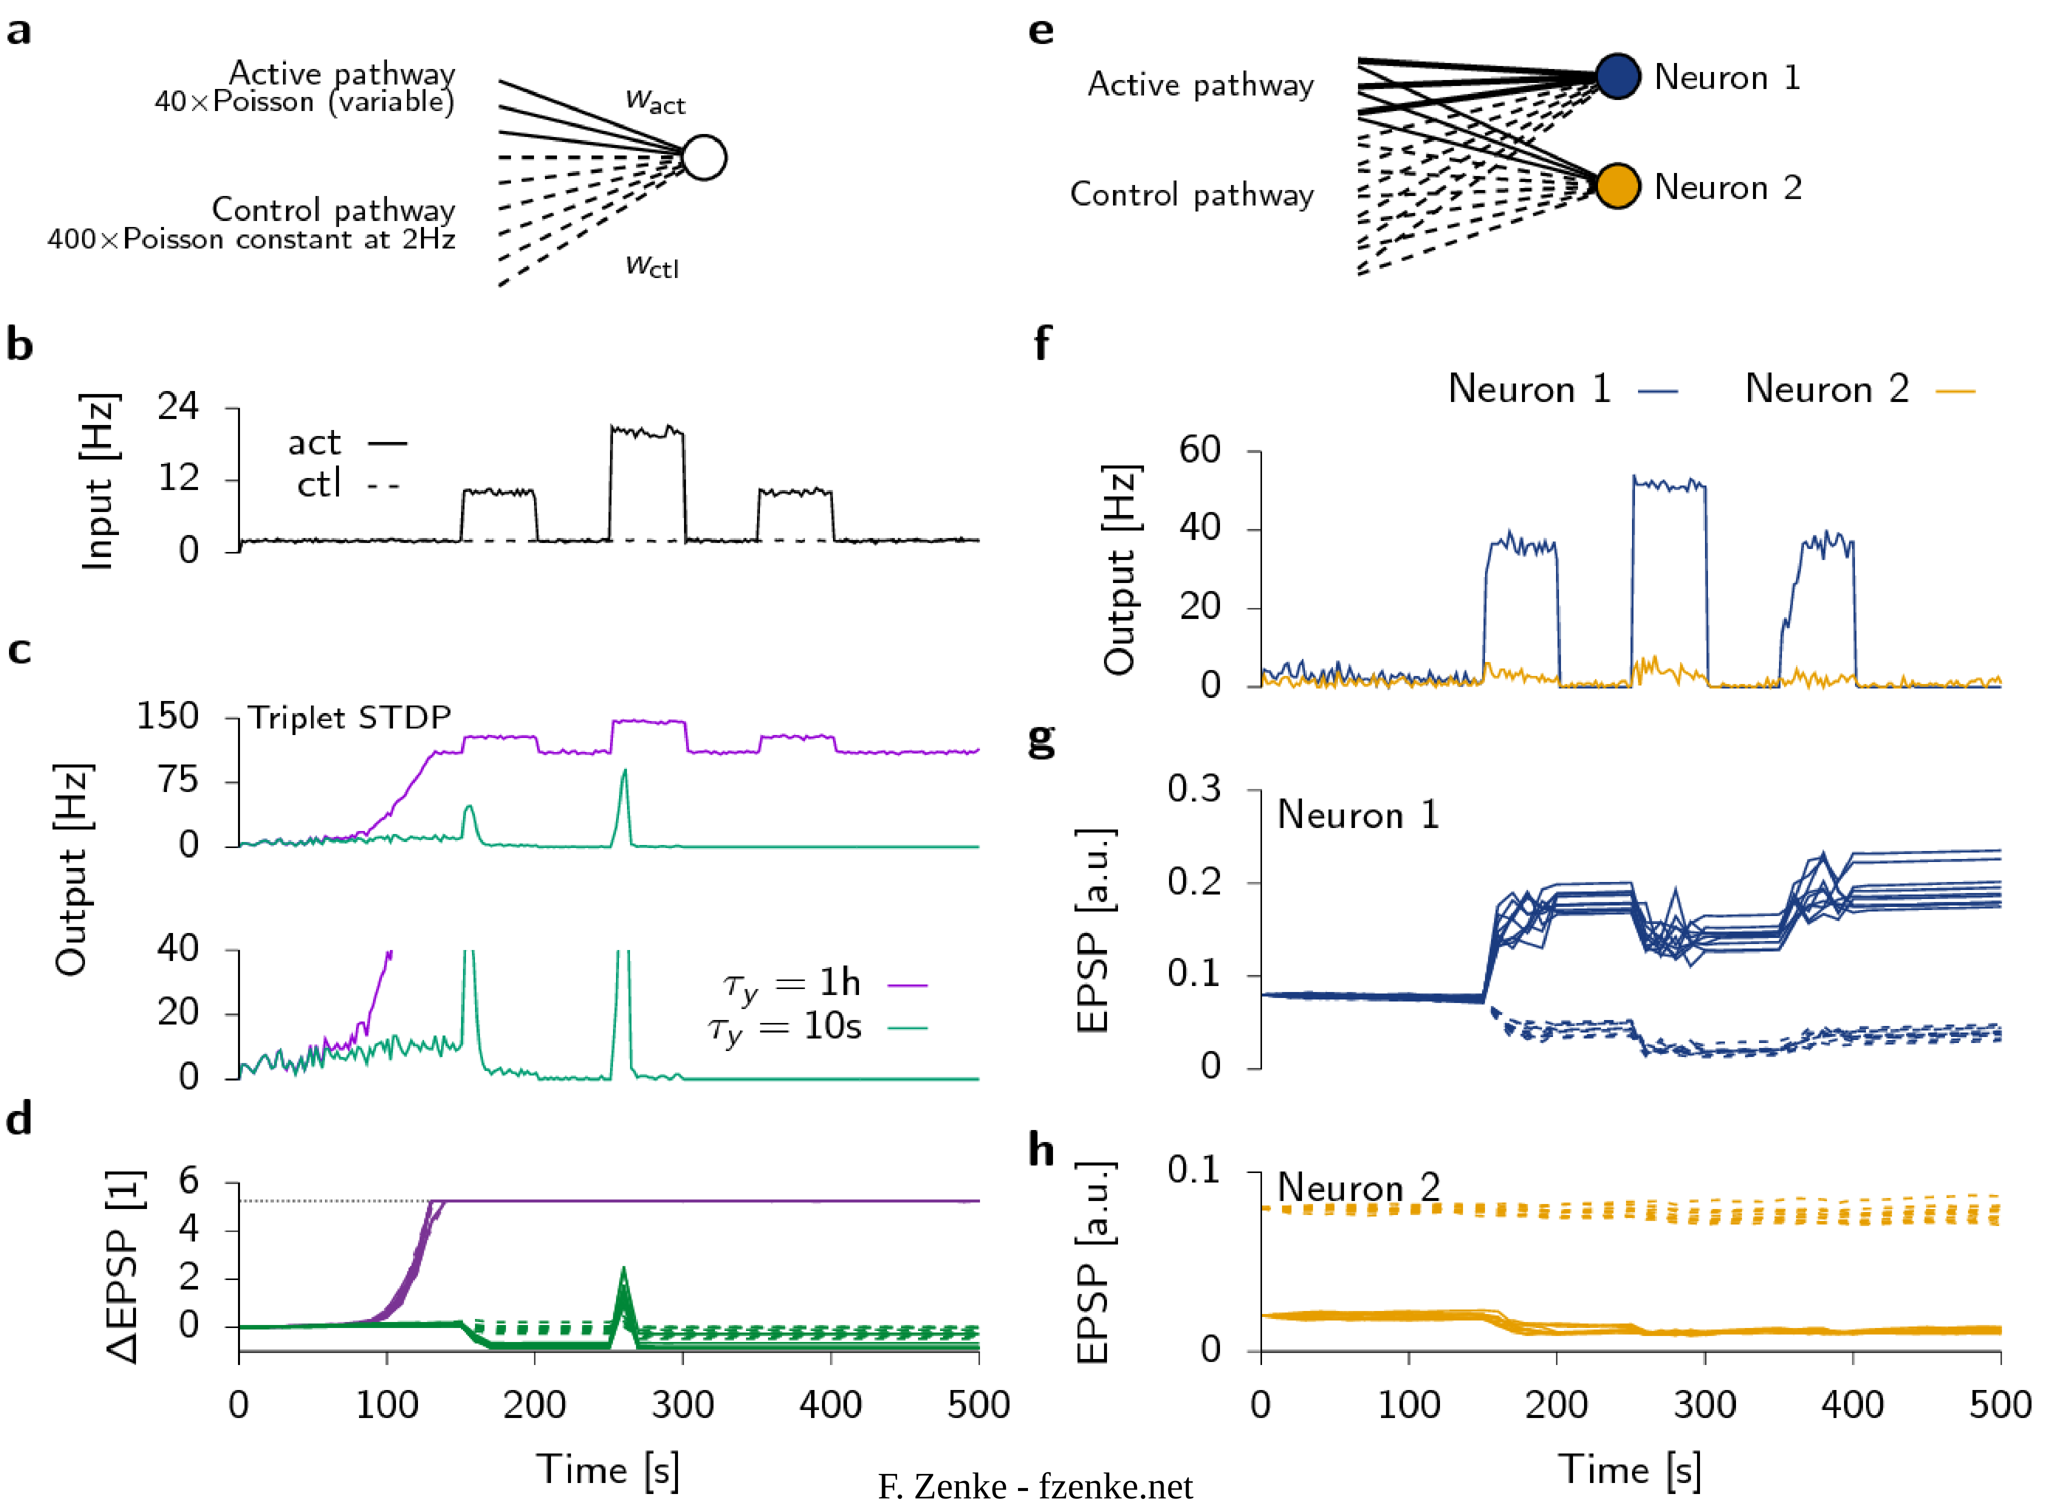

F. Zenke - fzenke.net

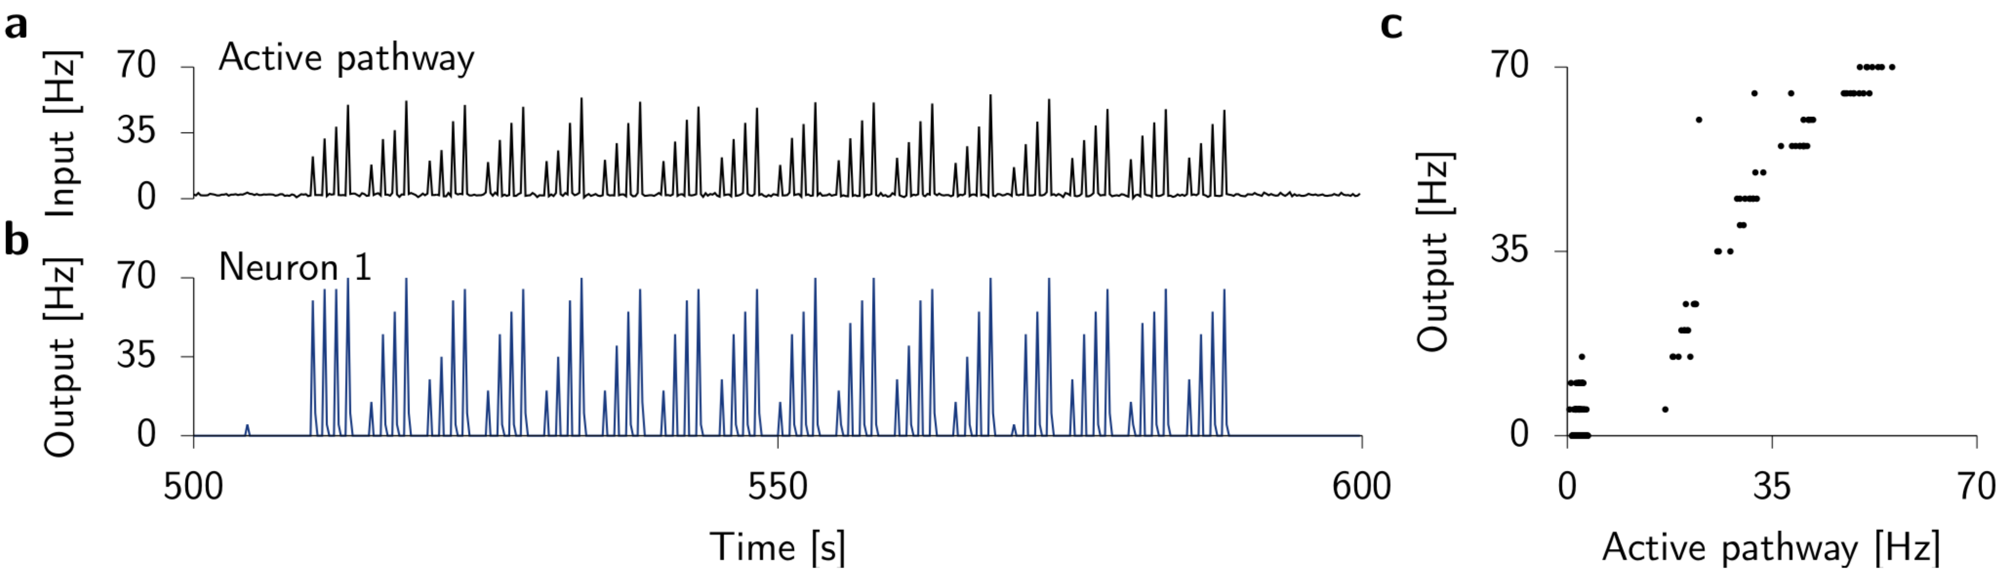
a
Active pathway
Input [Hz]
70
35
0
b
Neuron 1
Output [Hz]
70
35
0
500
550
600
Time [s]
c
Output [Hz]
70
35
0
0
35
70
Active pathway [Hz]

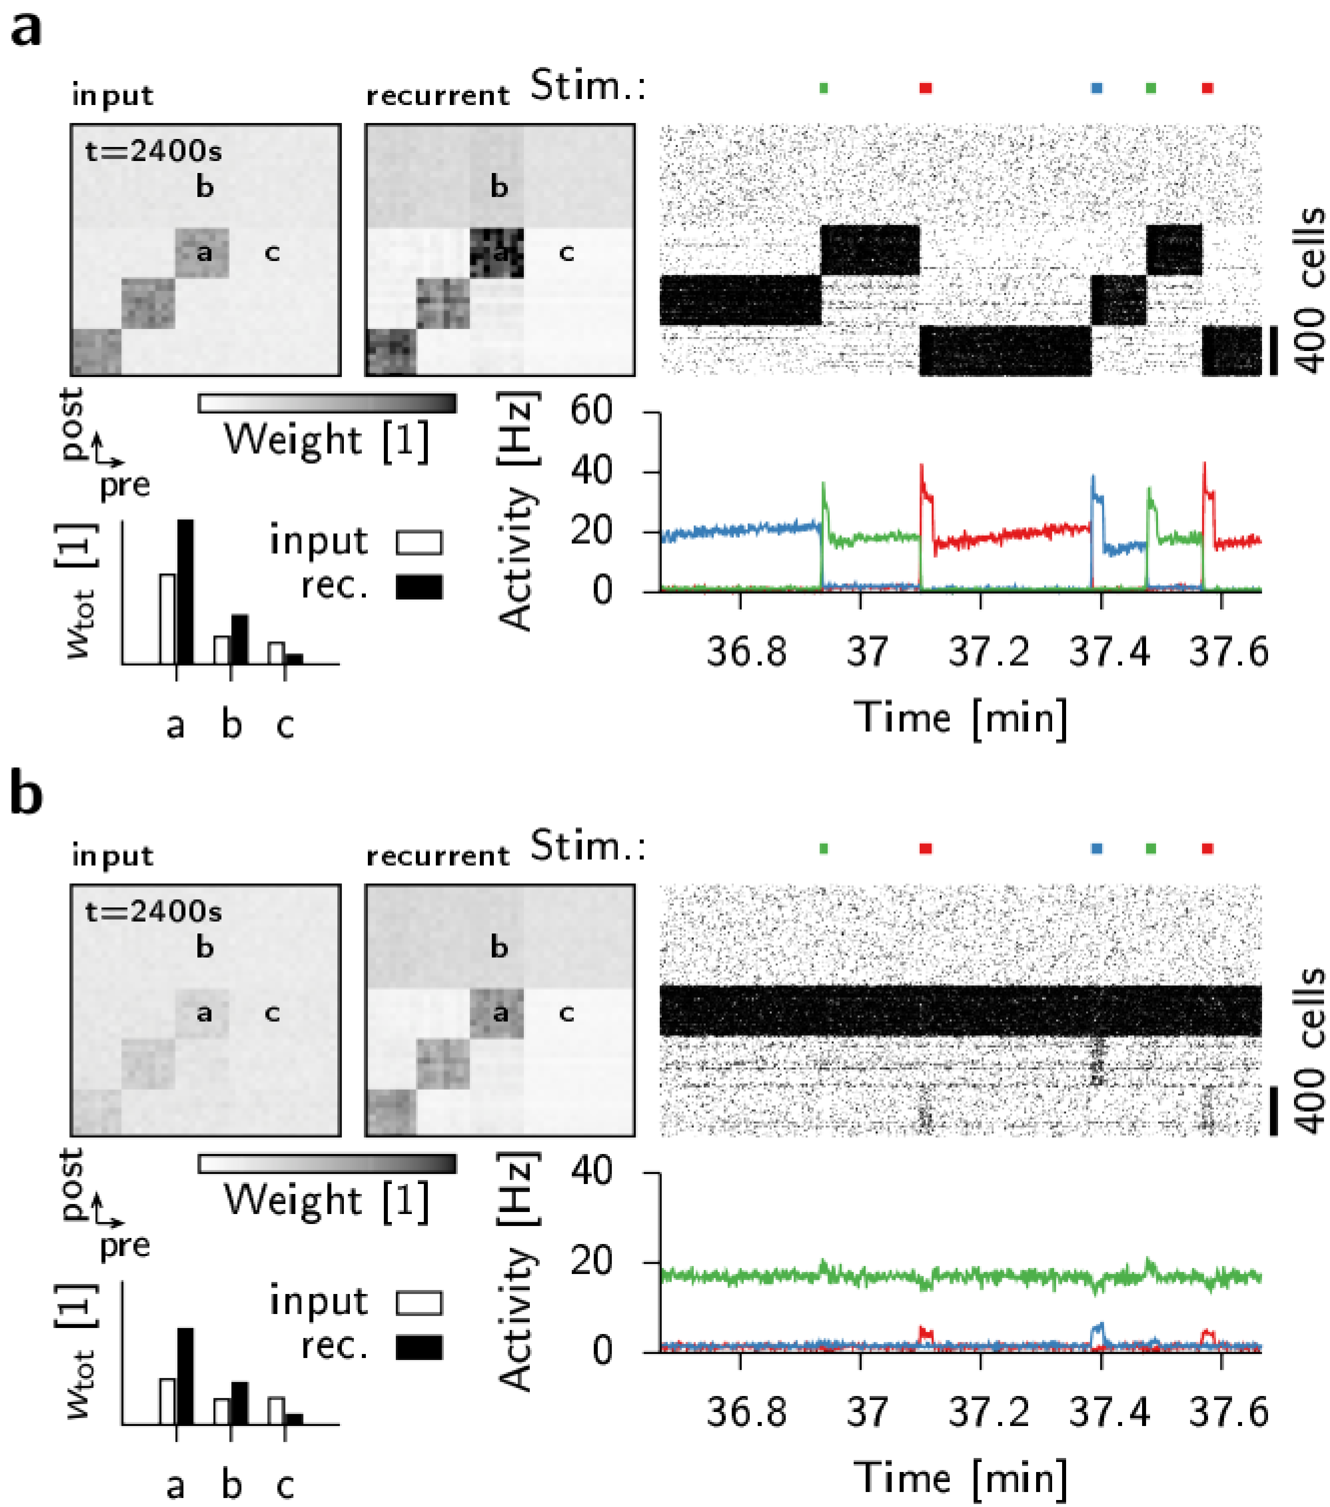
a
input
recurrent
Stim.:
t=2400s
b
a
c
400 cells
post
pre
Weight [1]
Activity [Hz]
60
40
20
0
36.8
37
37.2
37.4
37.6
Time [min]
w_{tot} [1]
input
rec.
a
b
c
b
input
recurrent
Stim.:
t=2400s
b
a
c
400 cells
post
pre
Weight [1]
Activity [Hz]
40
20
0
36.8
37
37.2
37.4
37.6
Time [min]
w_{tot} [1]
input
rec.
a
b
c

How does this form of plasticity reveal itself in experiments?

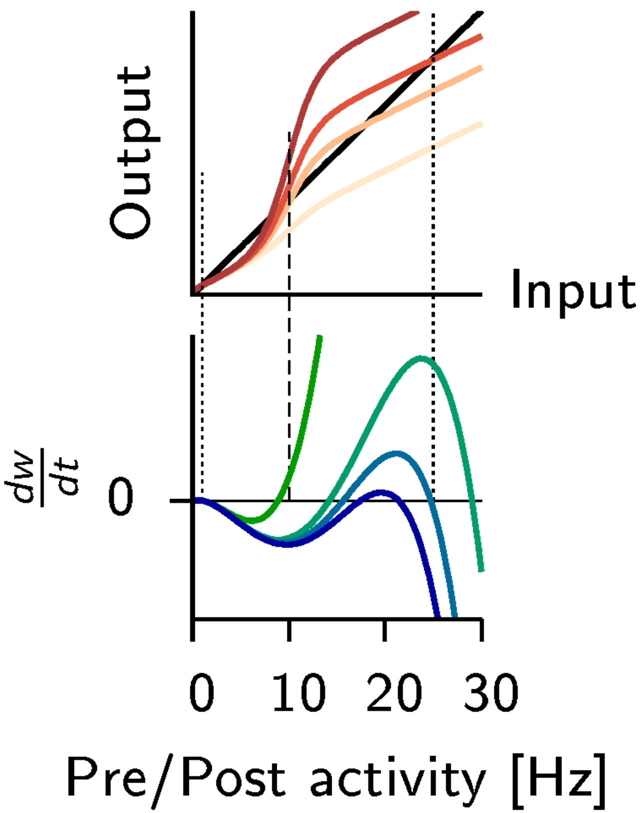

- Visible in experiments in which (pre)=0
- Burst detector (only at high rates)
- Bi-directional

$$\frac{dw}{dt} \propto (\text{pre})\,(\text{post})\,((\text{post}) - \kappa)$$ Triplet

$$-\beta\,(\text{post})^k\,(w - \tilde{w})$$ +Non-Hebbian

Heterosynaptic plasticity

$$k \geq 4$$

Zenke et al. (2015)

Cell assembly fixed points self-tune to fixed points of network activity

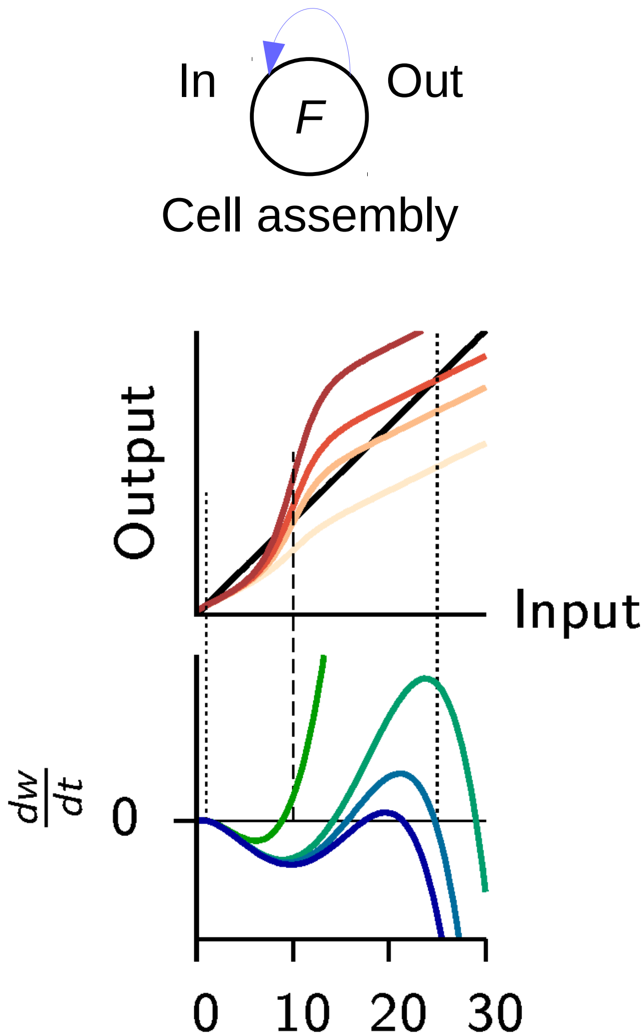

Zenke et al. (2015)

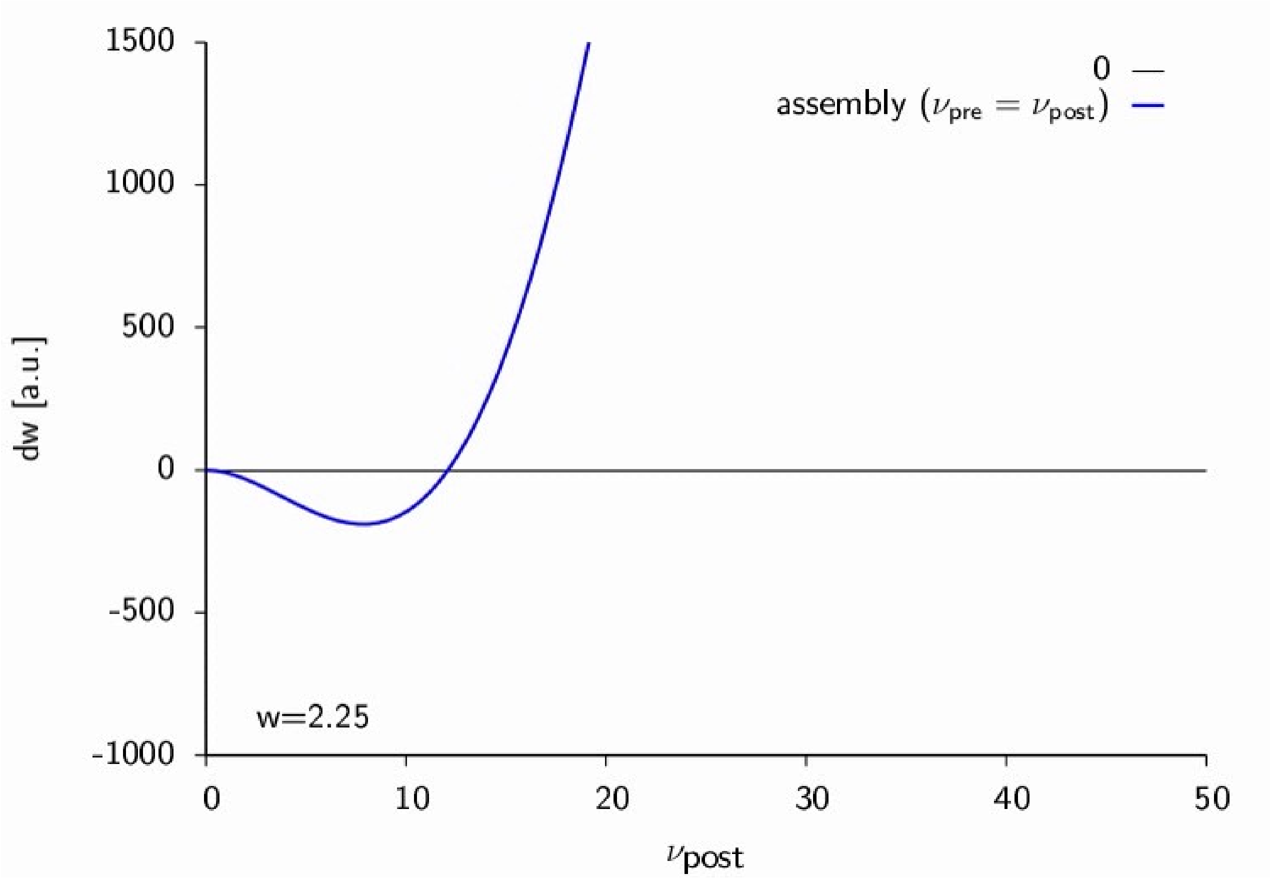

$$\frac{dw}{dt} \propto (\mathrm{pre})\,(\mathrm{post})\,((\mathrm{post}) - \kappa) \\ -\beta\,(\mathrm{post})^k\,(w-1)$$

Afferent connections increase little during assembly activation

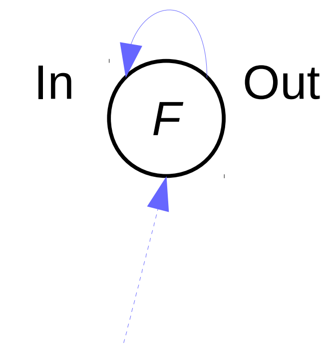

Background neurons

$$\nu_{\text{pre}} = 1\text{Hz}$$

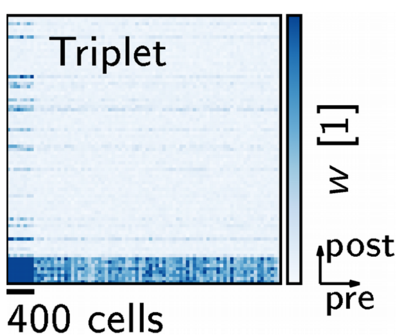

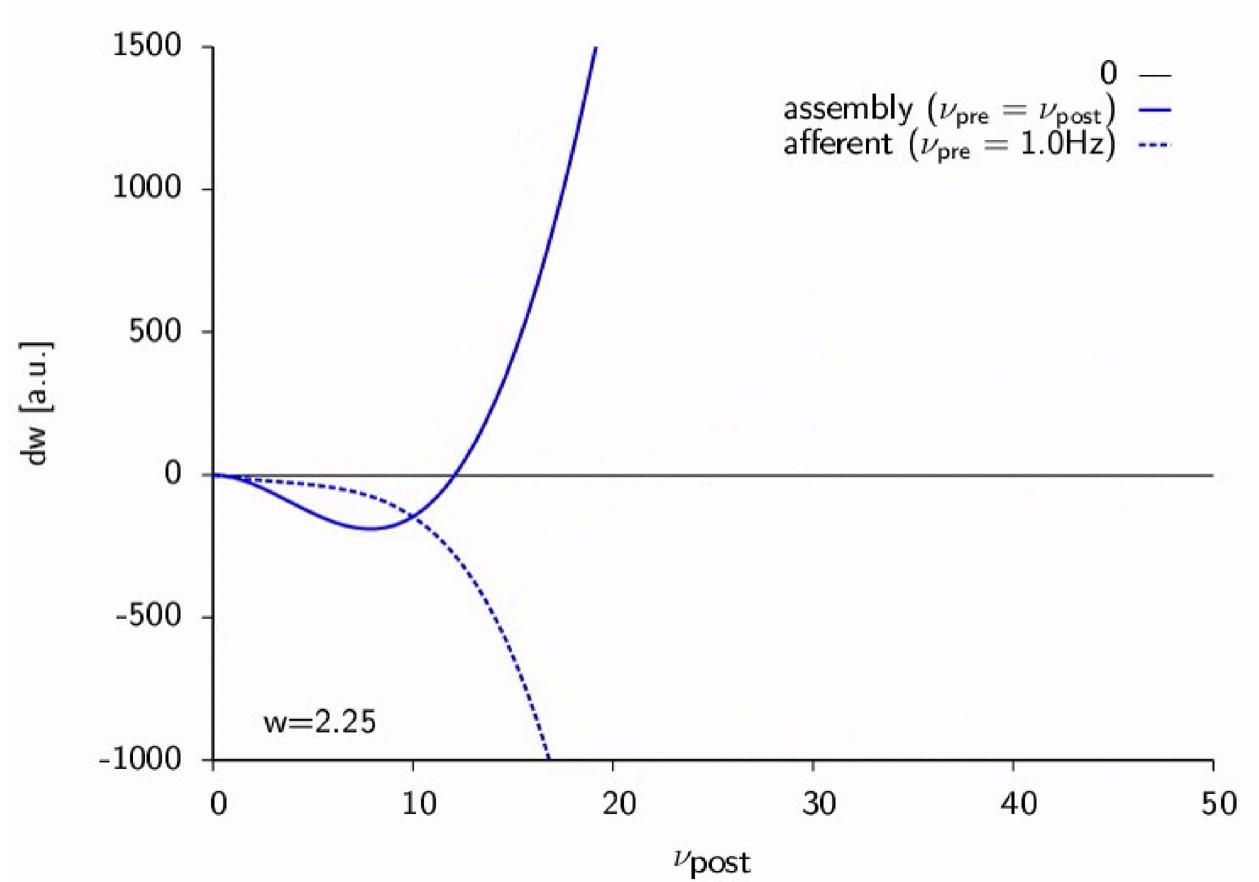

$$\frac{dw}{dt} \propto (\text{pre})\,(\text{post})\,((\text{post}) - \kappa) - \beta\,(\text{post})^k\,(w-1)$$

Zenke et al. (2015)

Nonlinearity of heterosynaptic plasticity is crucial for bistability

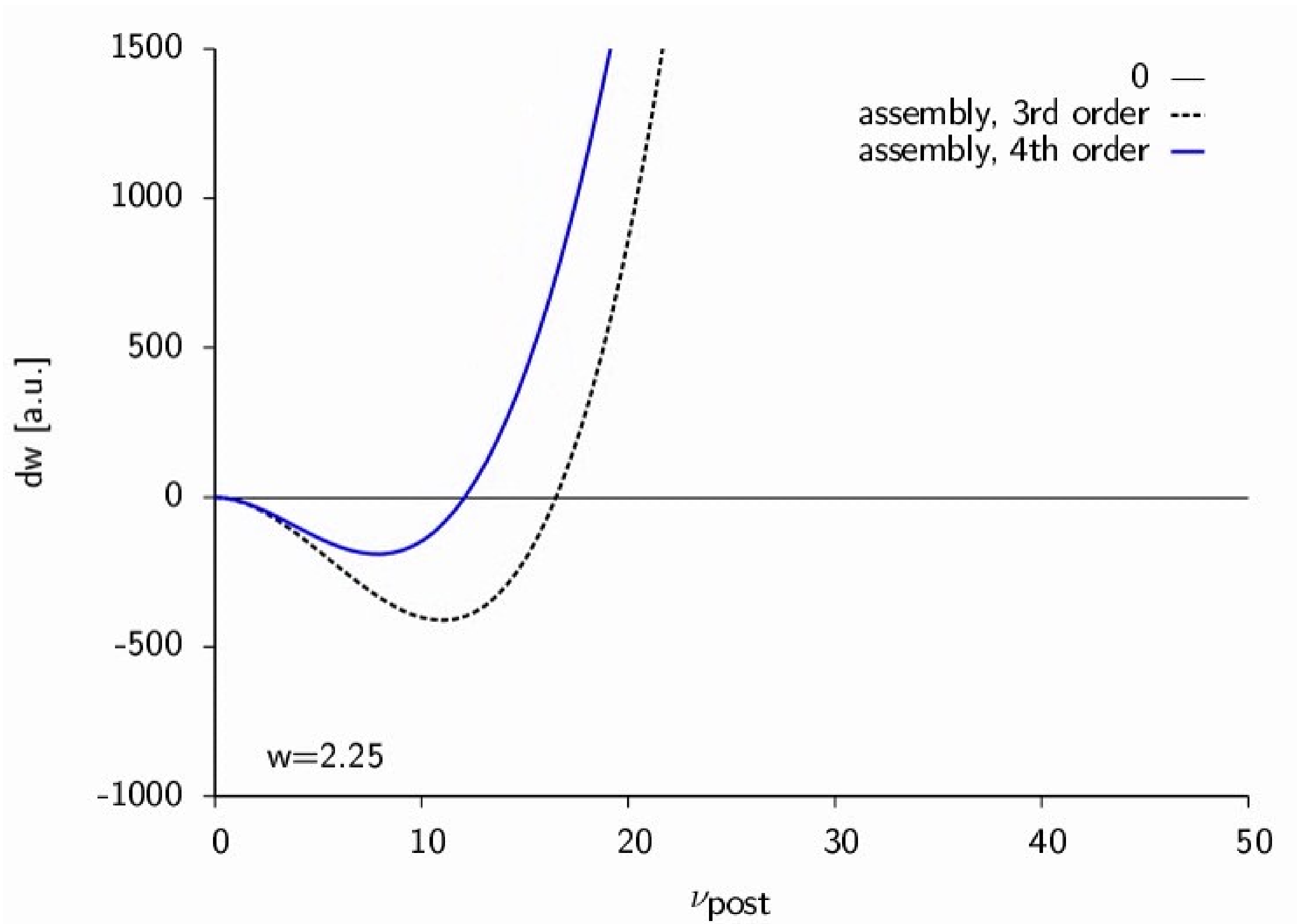

$$\frac{dw}{dt} \propto (\text{pre})\,(\text{post})\,((\text{post}) - \kappa) - \beta\,(\text{post})^k\,(w-1)$$

Zenke et al. (2015)

Silent memories are stable

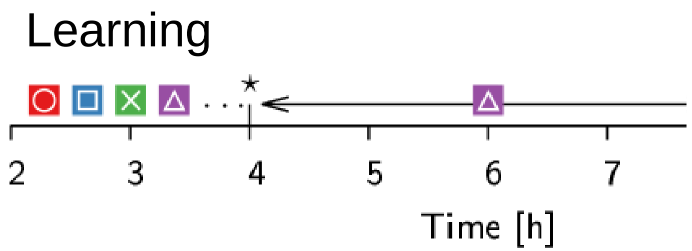

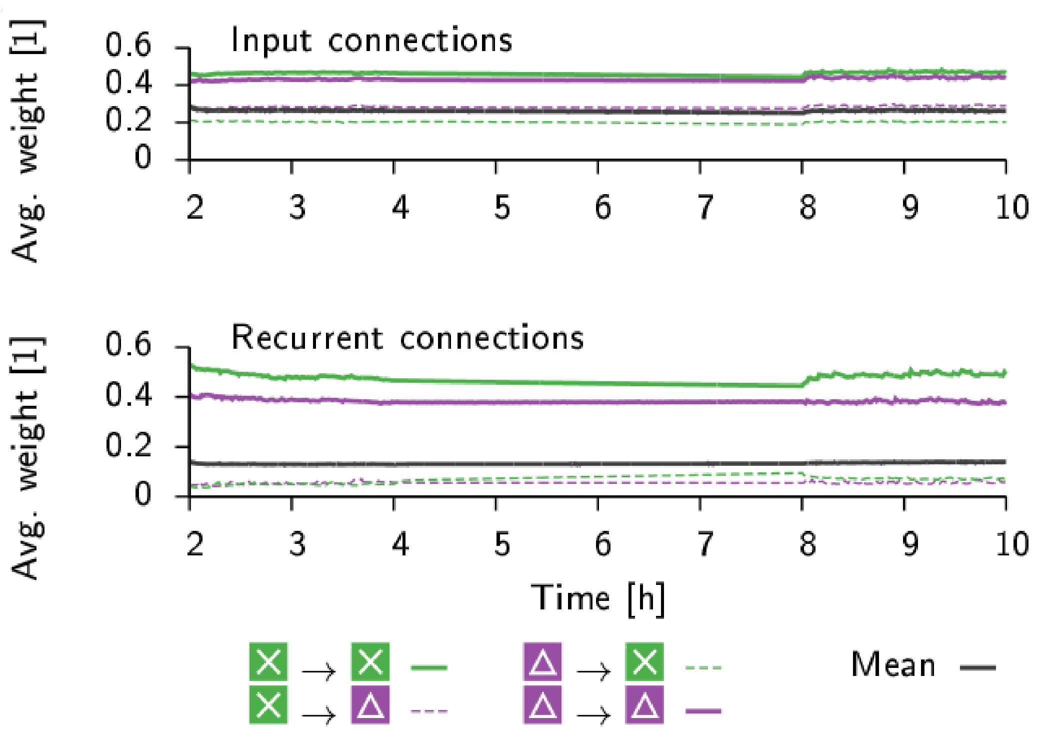

Learning of new assemblies

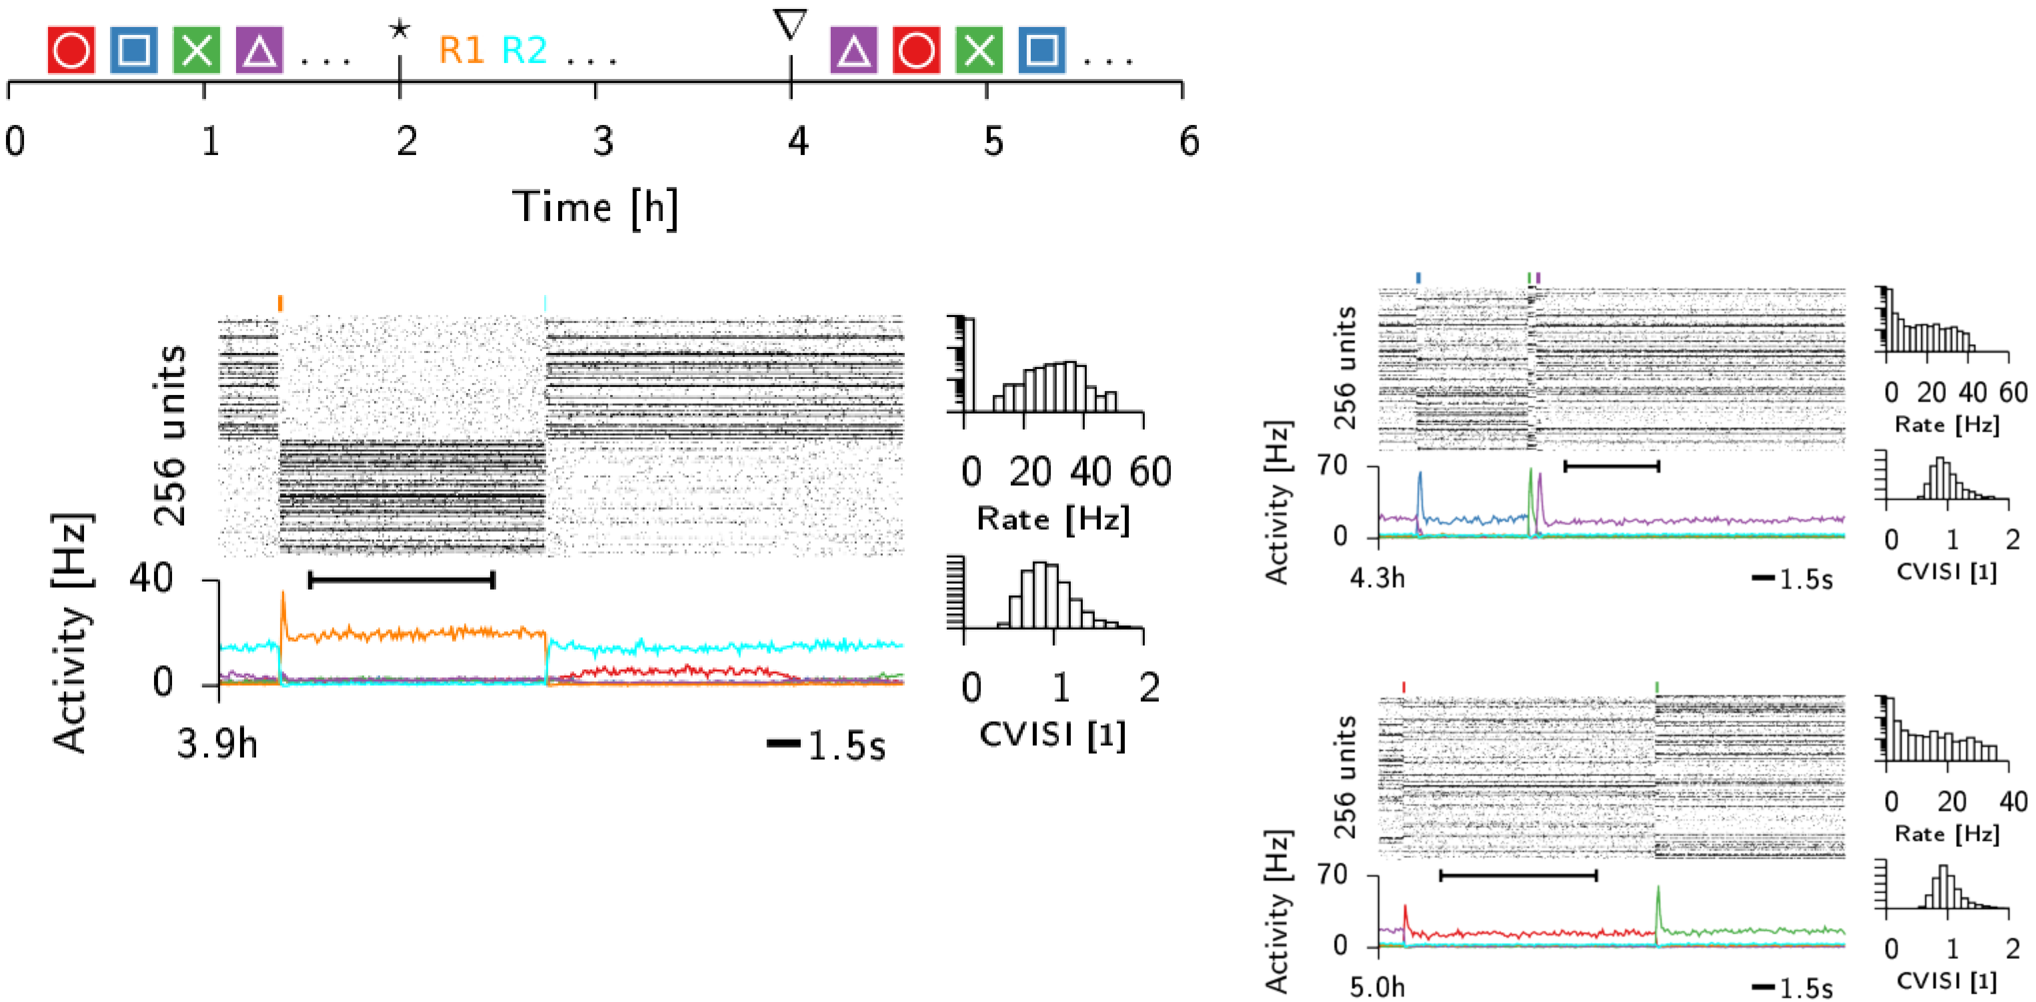

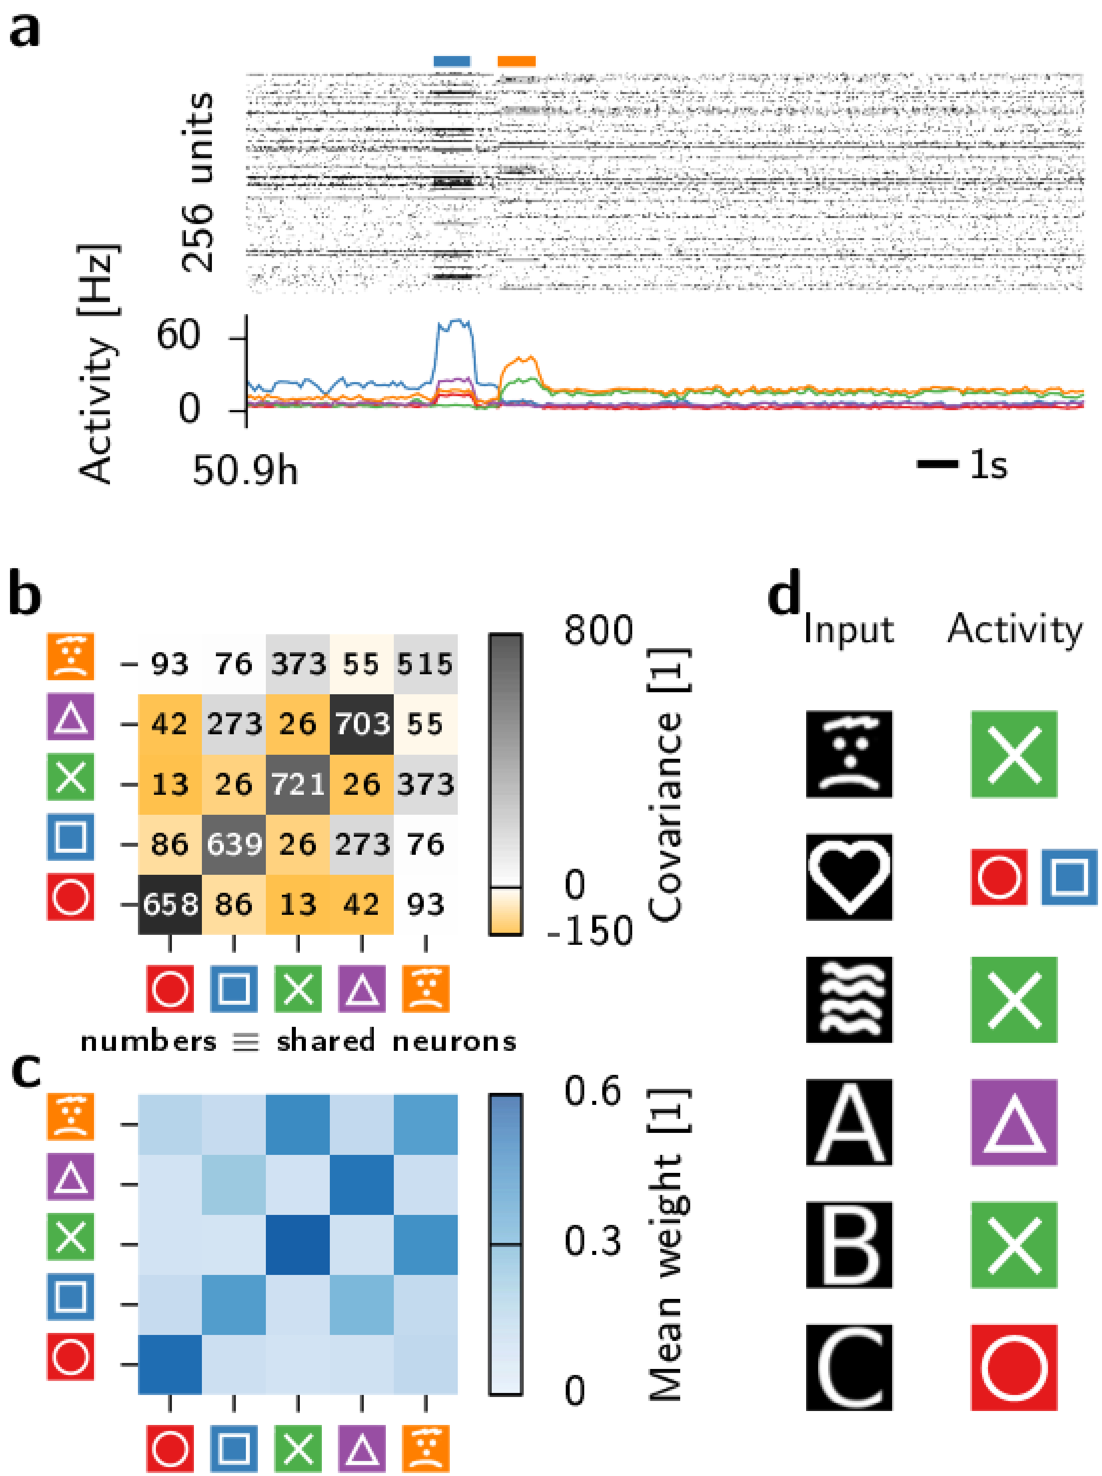
a
256 units
Activity [Hz]
60
0
50.9h
1s
b
93 76 373 55 515
42 273 26 703 55
13 26 721 26 373
86 639 26 273 76
658 86 13 42 93
800
0
-150
Covariance [1]
numbers ≡ shared neurons
c
0.6
0.3
0
Mean weight [1]
d
Input
Activity
A
B
C

Consolidation

$$\frac{dw}{dt} \propto (\mathrm{pre})\,(\mathrm{post})\,((\mathrm{post}) - \kappa)$$
$$-\beta\,(\mathrm{post})^k\,(w - \tilde{w})$$

a With consolidation

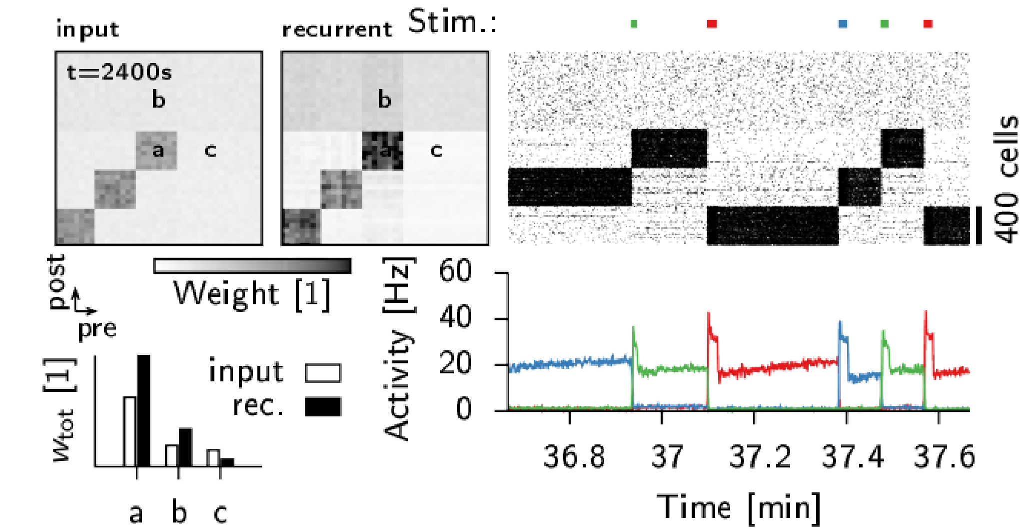

Slow consolidation dynamics

$$\tau_{\mathrm{cons}} \frac{d\tilde{w}}{dt} = -\tilde{w} + w + f(\tilde{w})$$

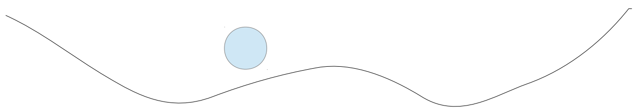

Ziegler, Zenke et al. 2015

Network stays plastic: Learning of new assemblies

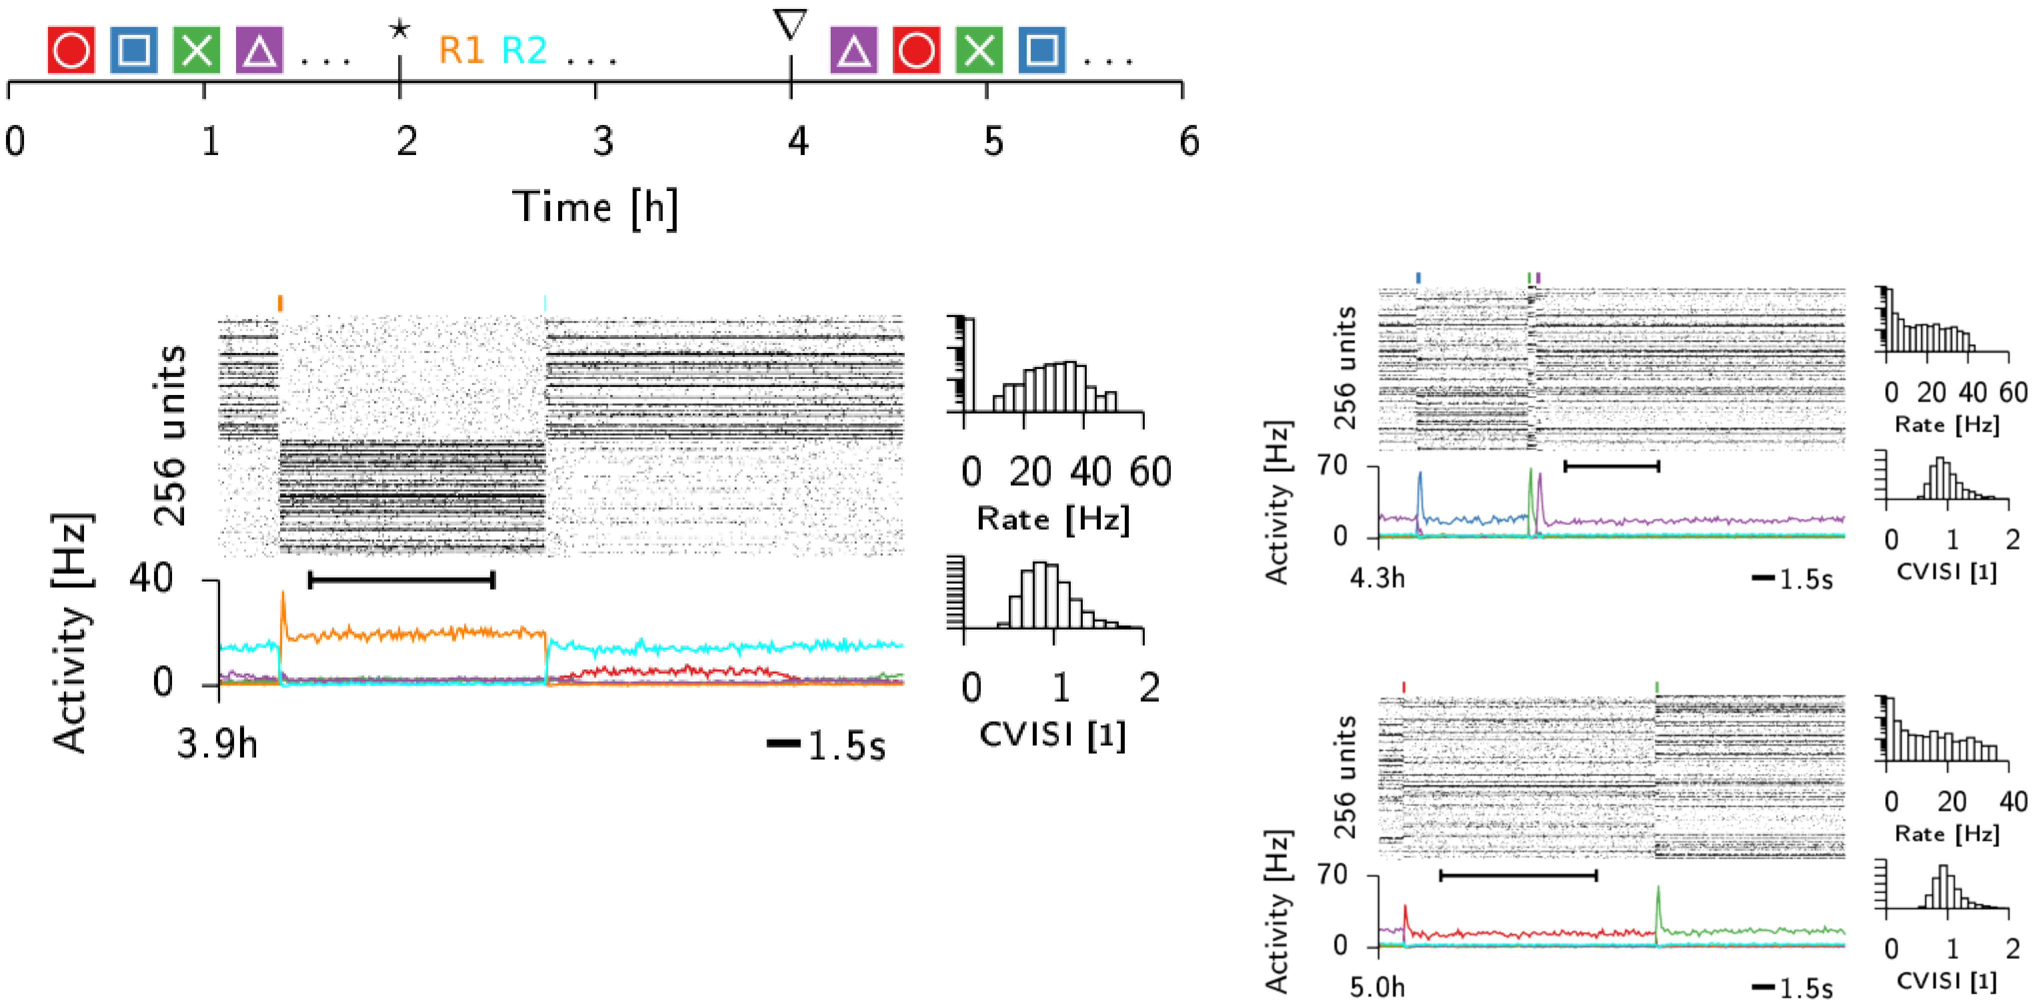

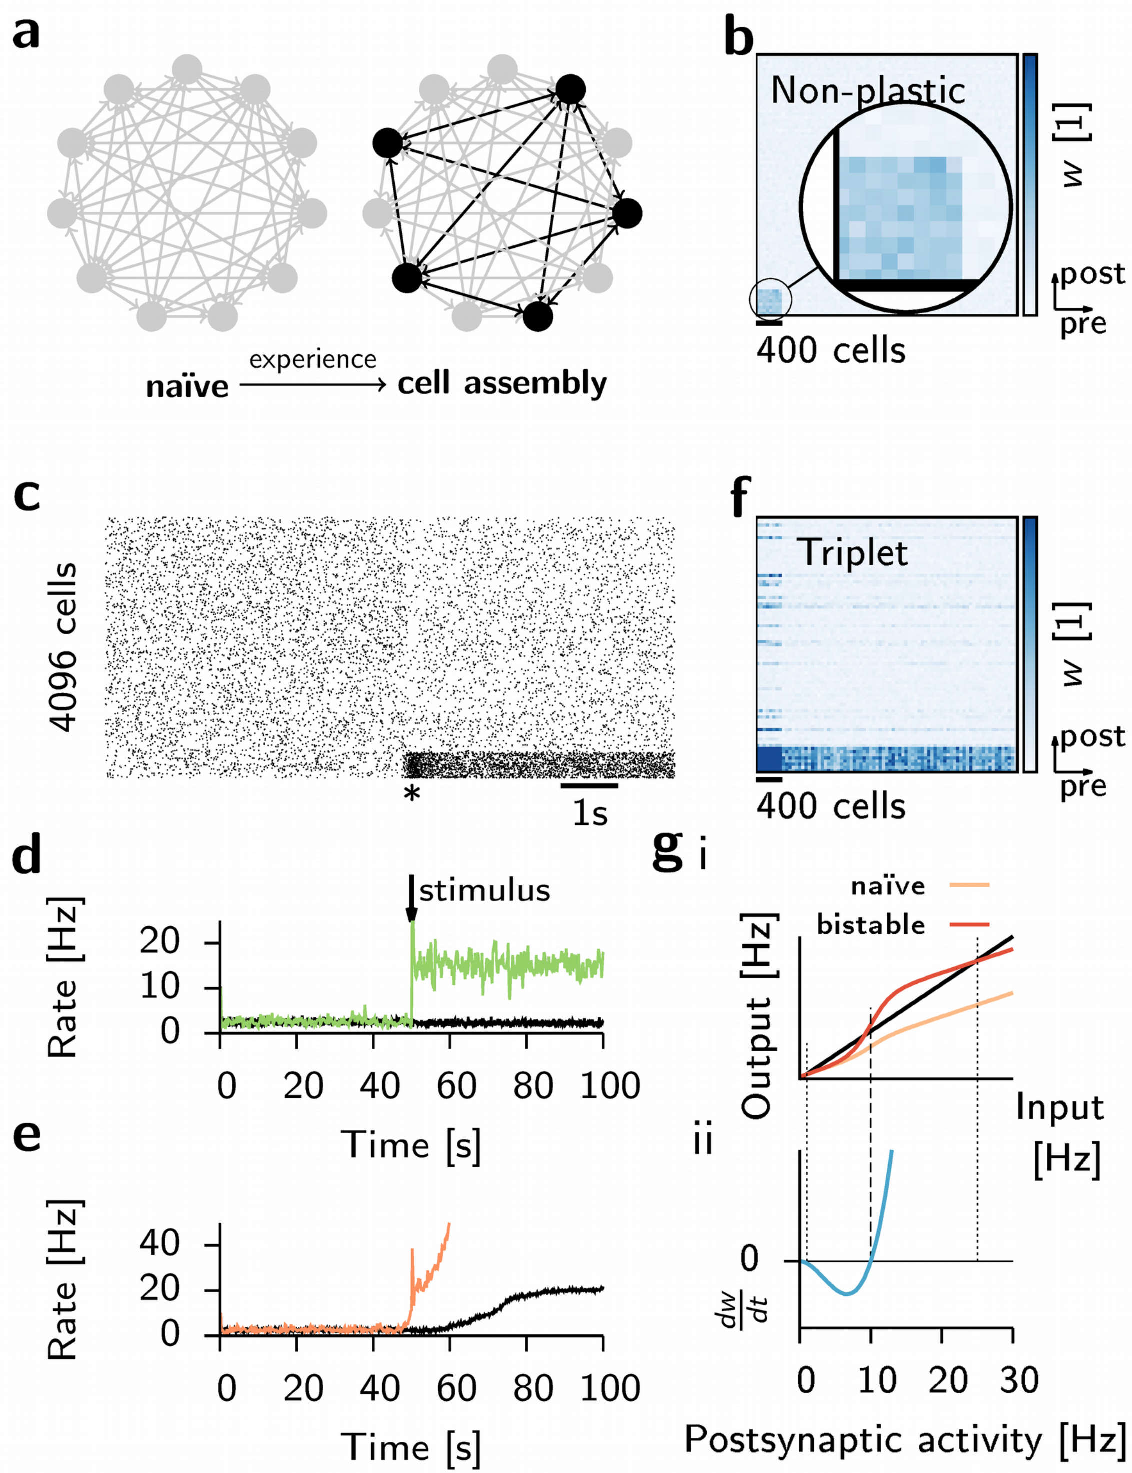
a
experience
naïve
cell assembly
b
Non-plastic
w [1]
post
pre
400 cells
c
4096 cells
*
1s
f
Triplet
w [1]
post
pre
400 cells
d
stimulus
Rate [Hz]
20
10
0
0
20
40
60
80
100
Time [s]
e
Rate [Hz]
40
20
0
0
20
40
60
80
100
Time [s]
g i
naïve
bistable
Output [Hz]
Input
[Hz]
ii
0
dw/dt
0
10
20
30
Postsynaptic activity [Hz]

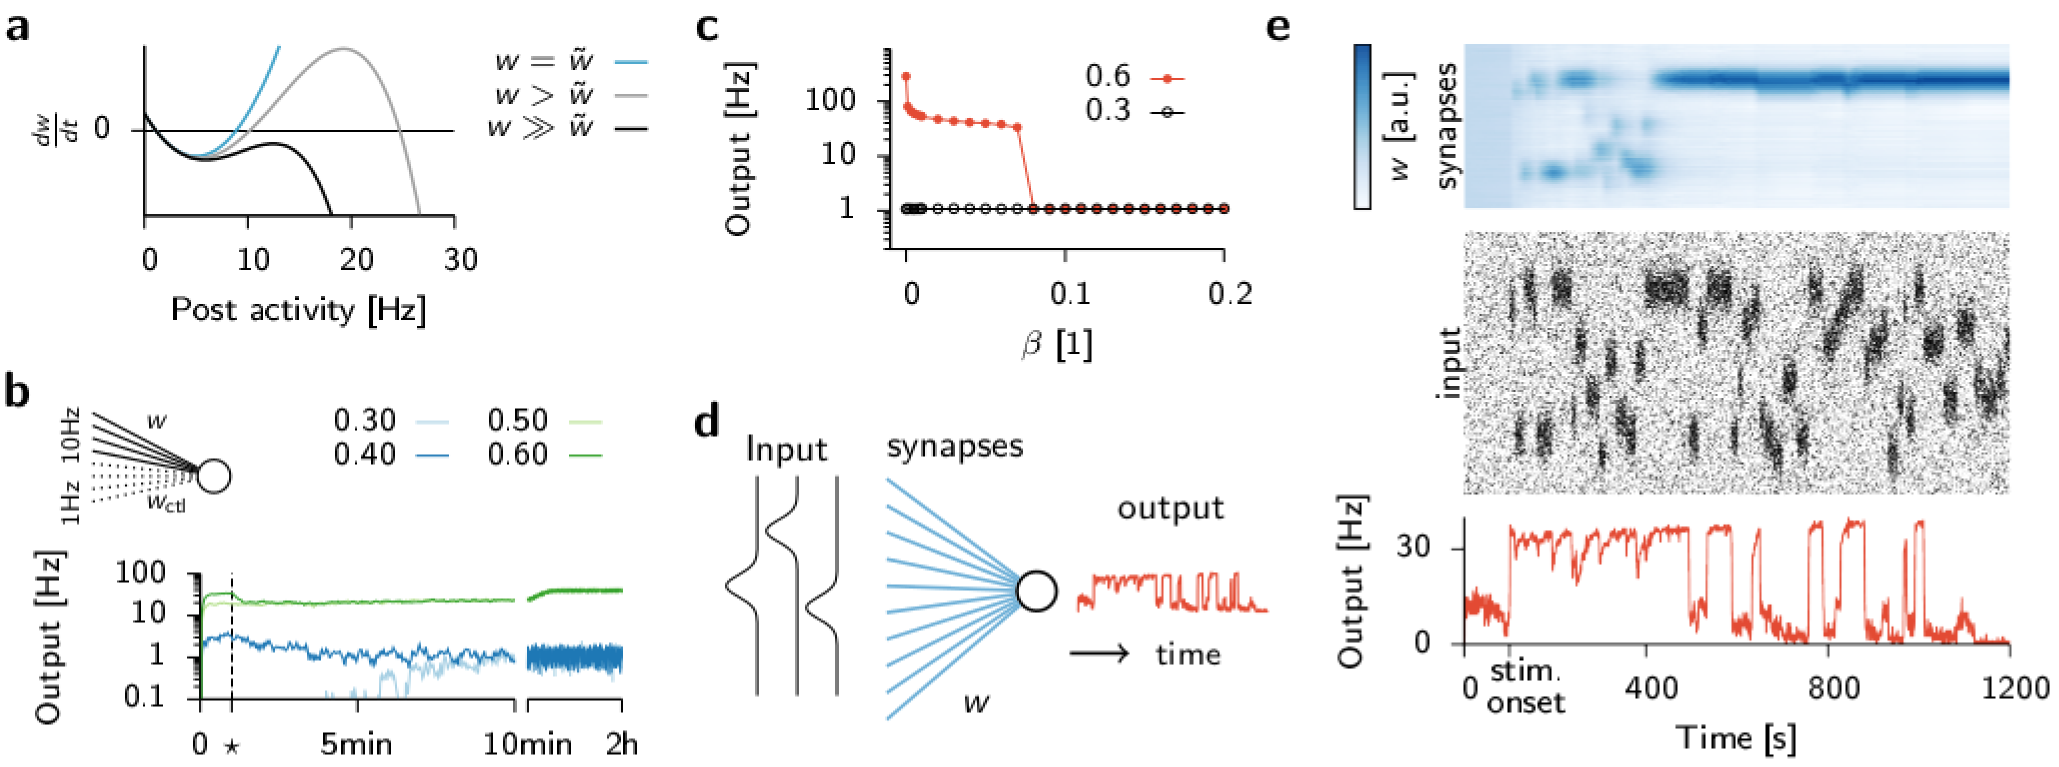
a
dw/dt
0
w = w̃
w > w̃
w ≫ w̃
0 10 20 30
Post activity [Hz]
b
w
10Hz
1Hz
w_ctl
0.30
0.40
0.50
0.60
Output [Hz]
100
10
1
0.1
0 ⋆ 5min 10min 2h
c
Output [Hz]
100
10
1
0.6
0.3
0 0.1 0.2
β [1]
d
Input
synapses
output
w
time
e
w [a.u.]
synapses
input
Output [Hz]
30
0
0 stim. onset 400 800 1200
Time [s]

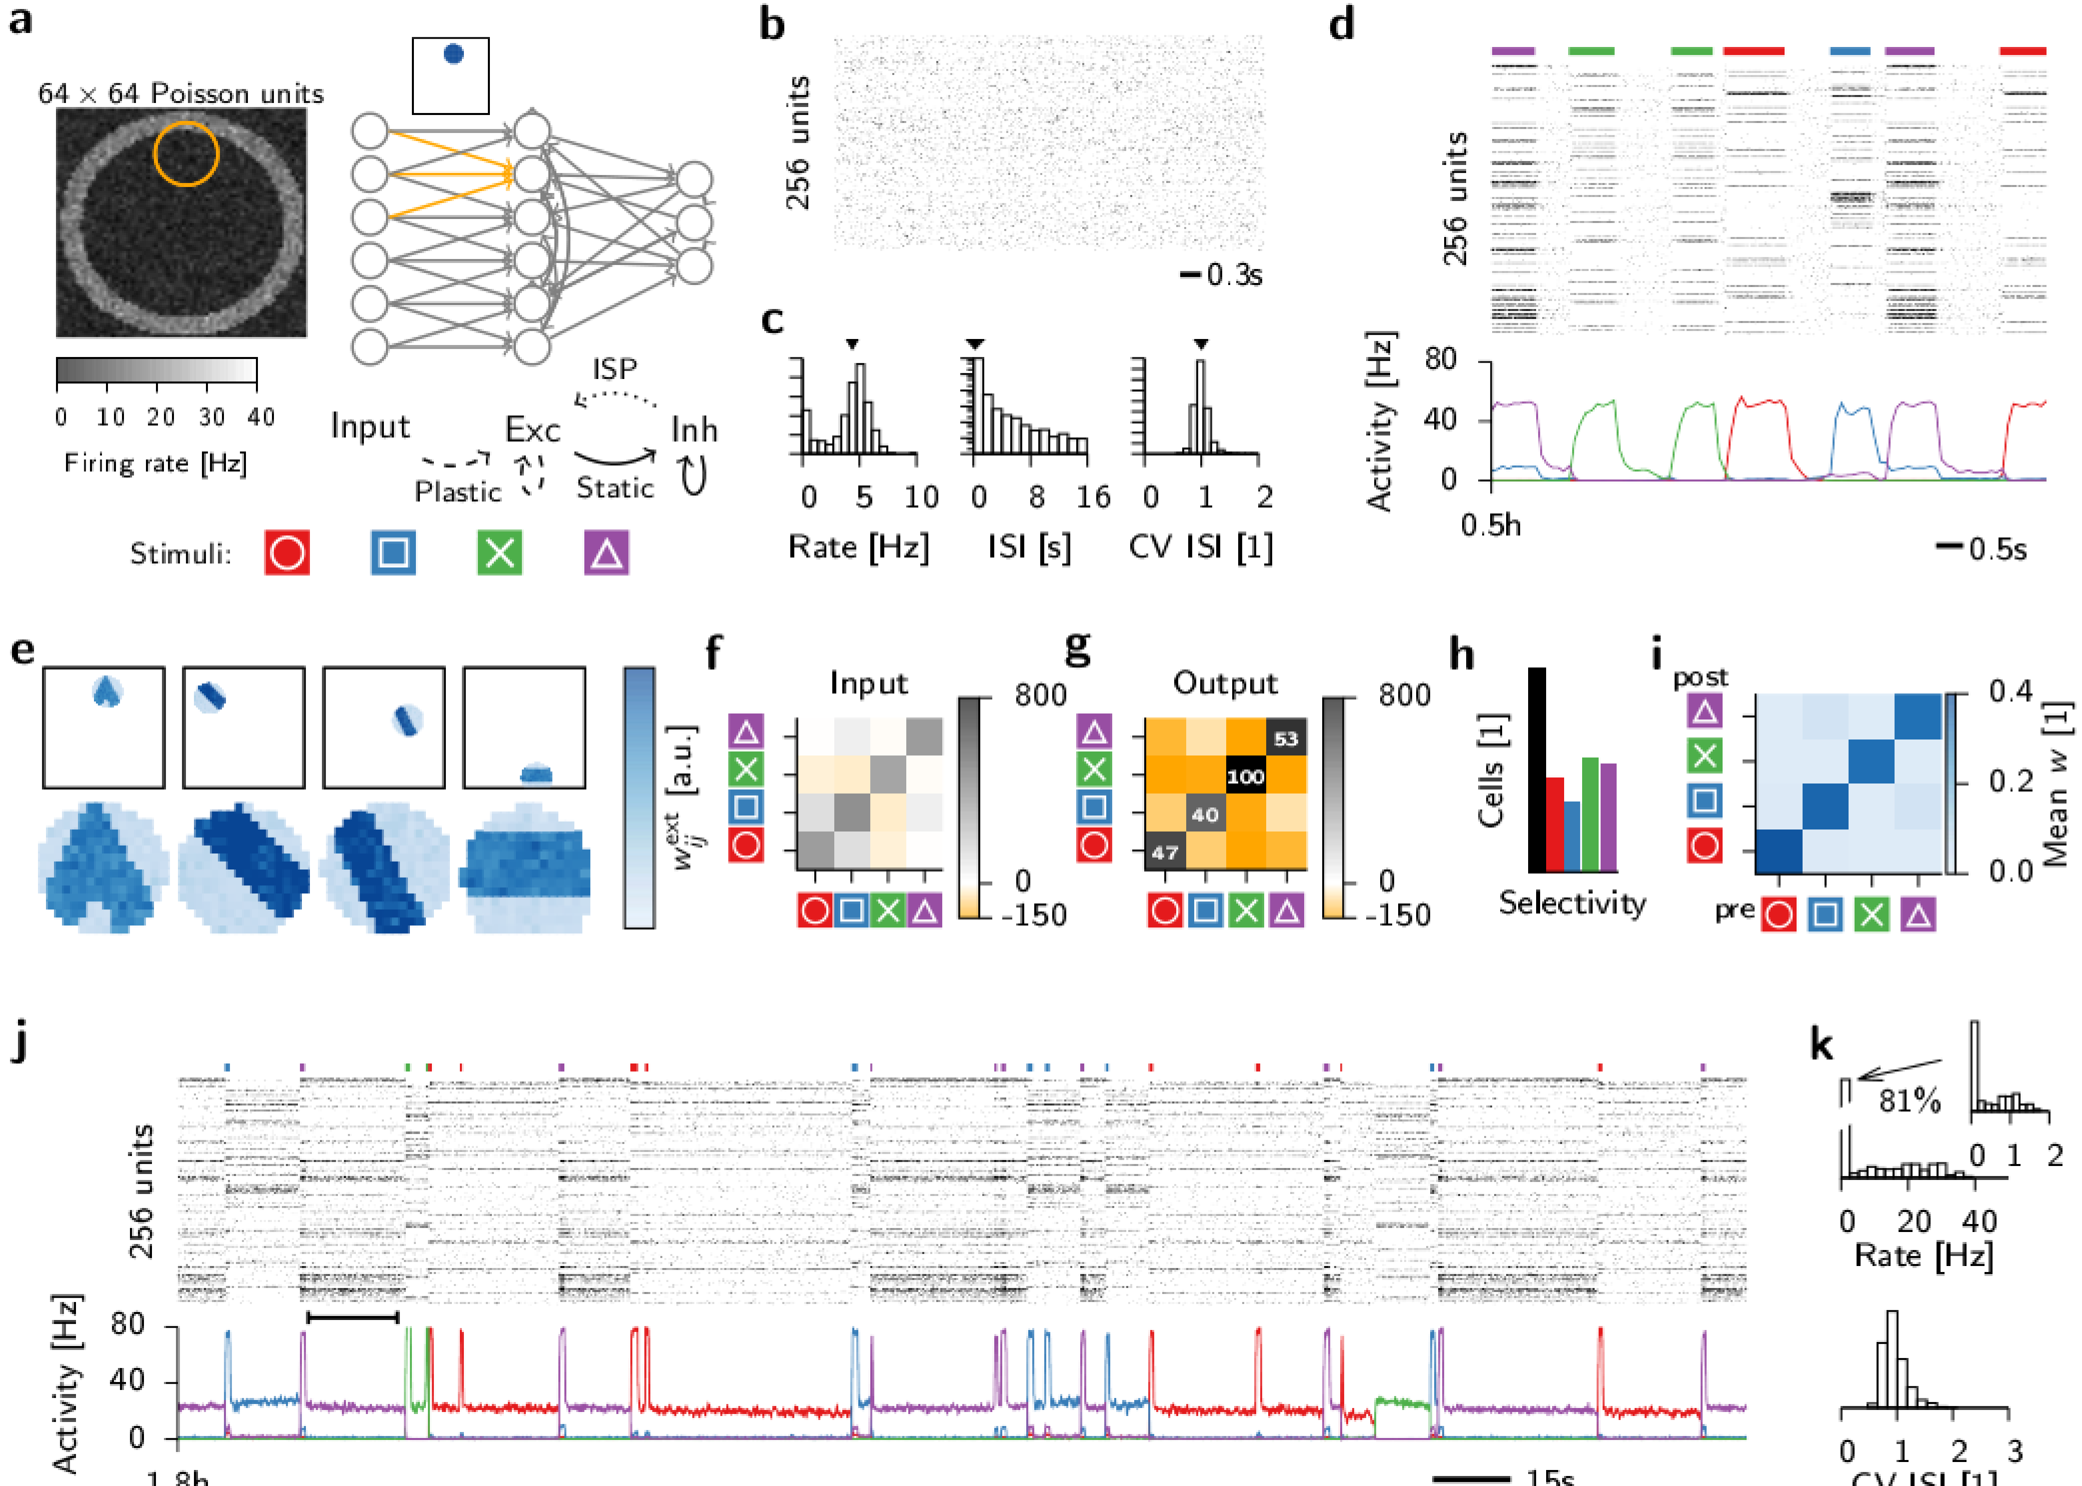

a
64 × 64 Poisson units
0 10 20 30 40
Firing rate [Hz]
Input
Exc
Inh
ISP
Plastic
Static
Stimuli:
b
256 units
0.3s
c
Rate [Hz]
ISI [s]
CV ISI [1]
d
256 units
Activity [Hz]
0.5h
0.5s
e
w_{ij}^{ext} [a.u.]
f
Input
800
0
-150
g
Output
53
100
40
47
h
Cells [1]
Selectivity
i
post
pre
Mean w [1]
j
256 units
Activity [Hz]
1.8h
15s
k
81%
Rate [Hz]
CV ISI [1]

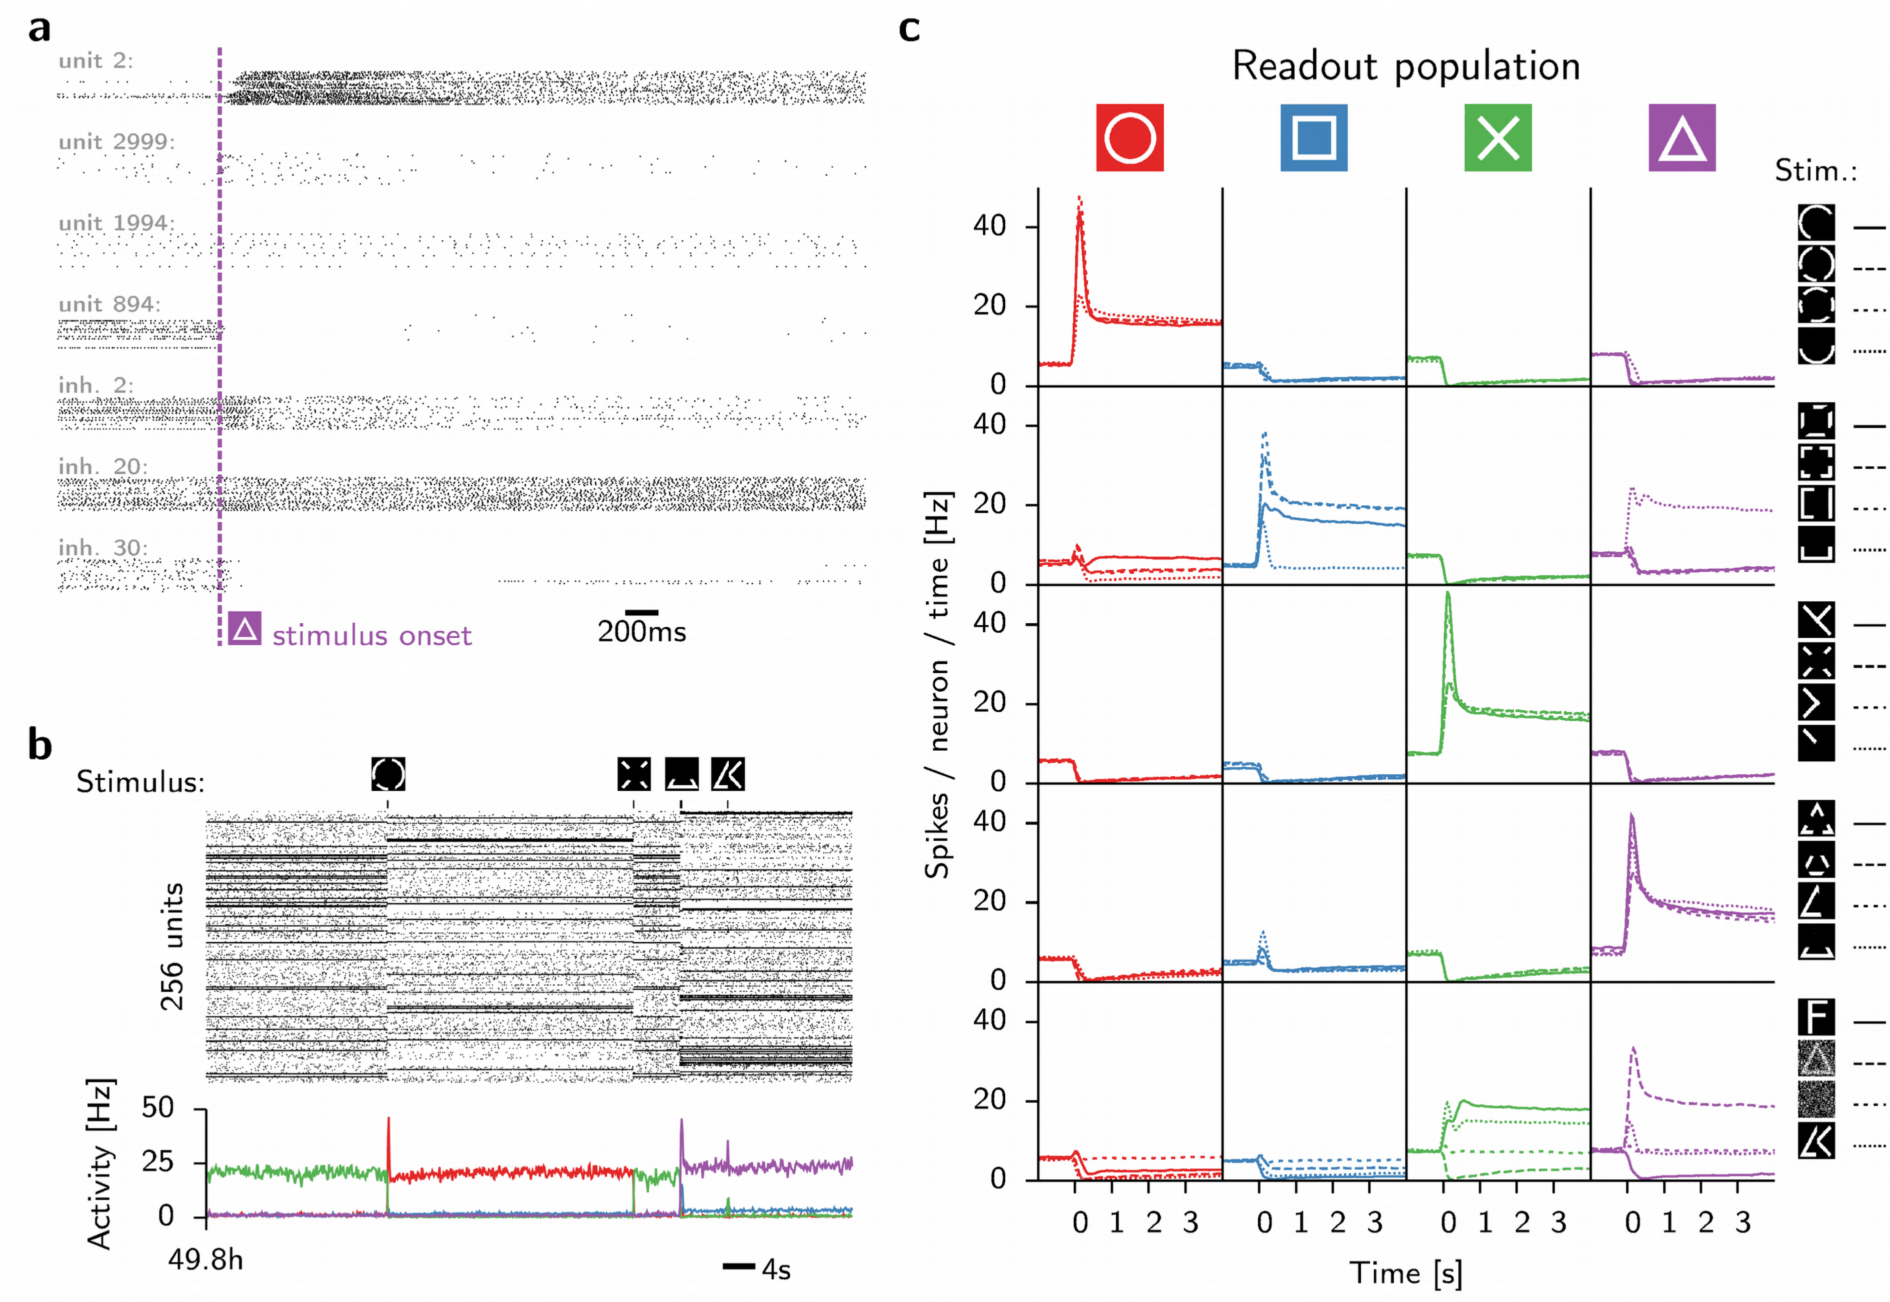

F. Zenke - fzenke.net

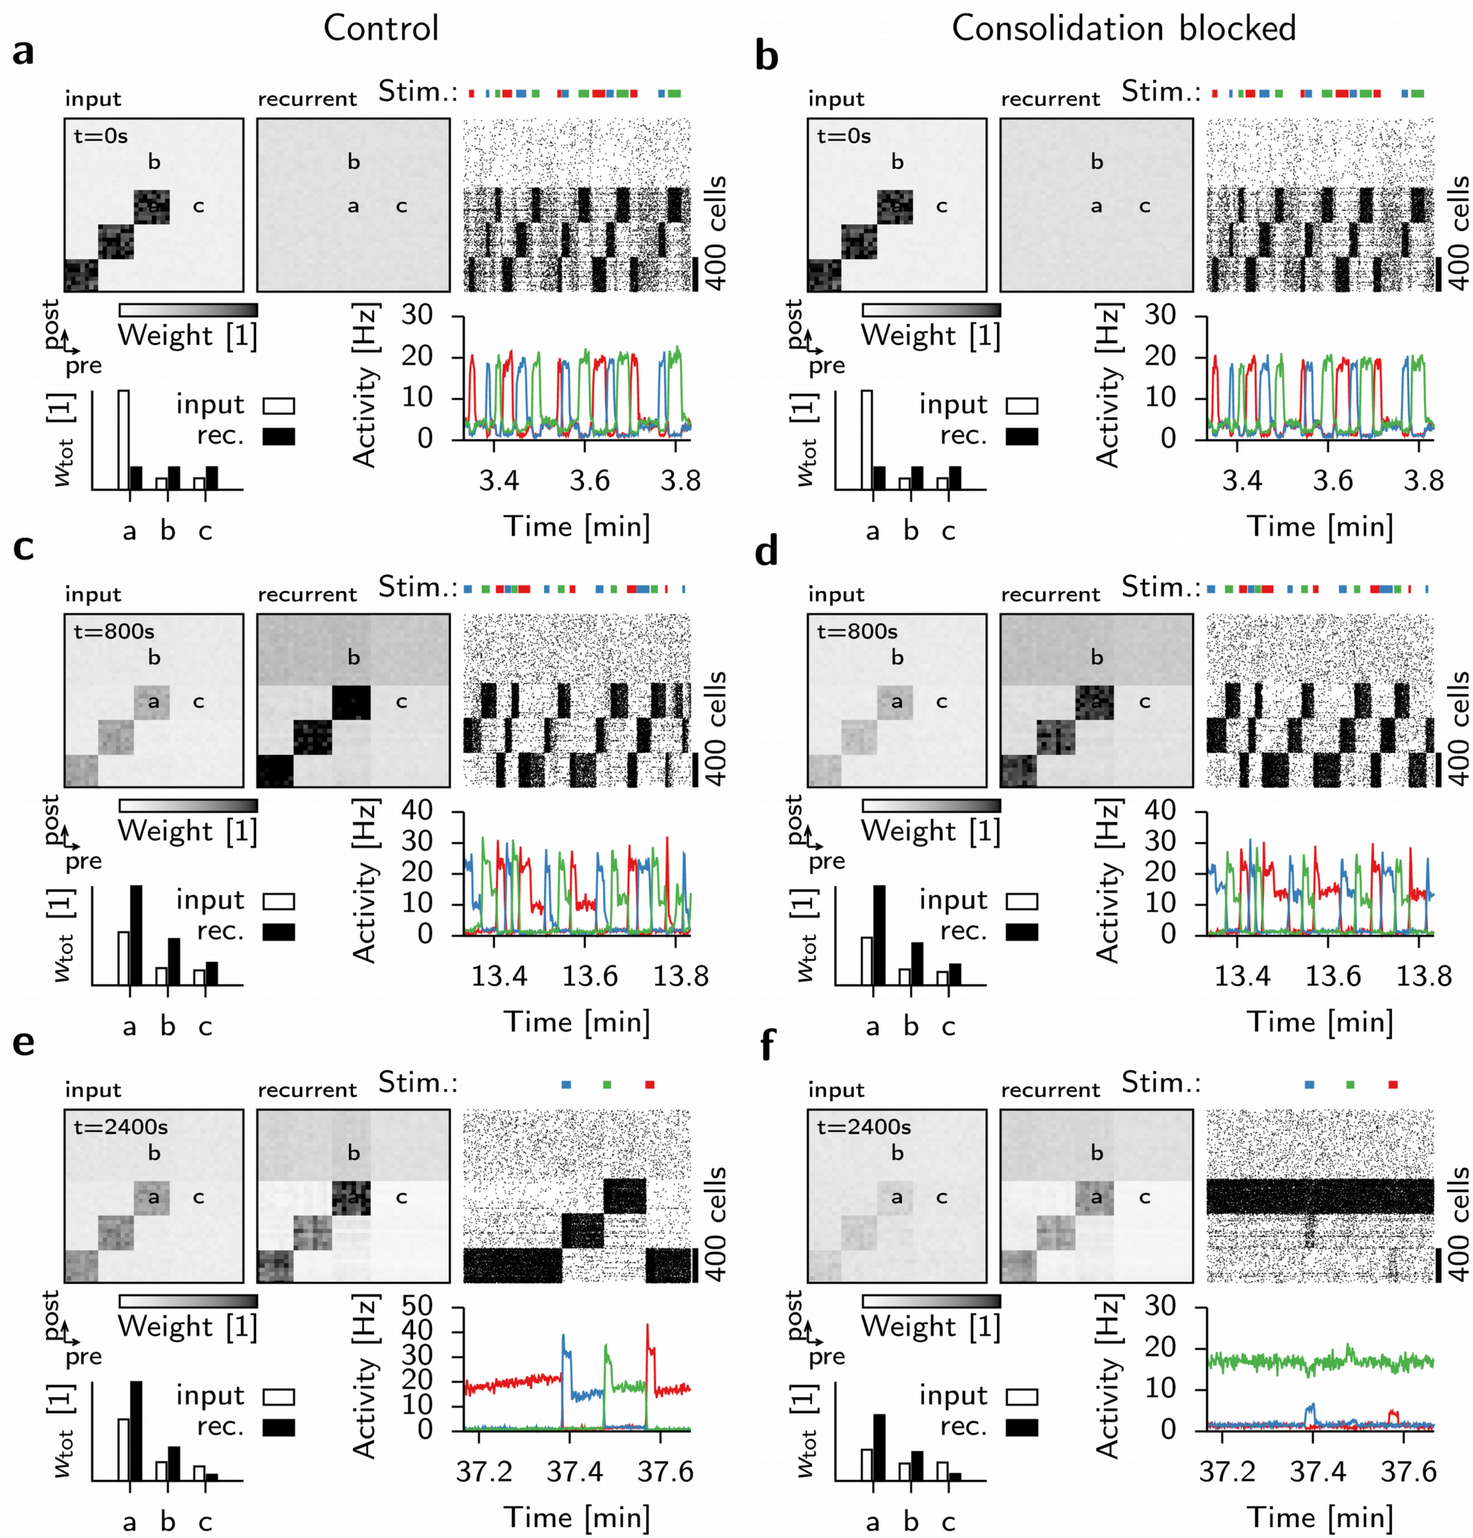
Control
Consolidation blocked
a
b
c
d
e
f
input
recurrent
Stim.:
t=0s
t=800s
t=2400s
a
b
c
400 cells
post
pre
Weight [1]
w_{tot} [1]
input
rec.
Activity [Hz]
Time [min]

Weak initial state

No transmitter triggered

No heterosynaptic plasticity

No inhibitory plasticity