Week 5: Networks of Neurons-Introduction

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

Biological Modeling of Neural Networks

Week 5

NETWORKS of NEURONS and ASSOCIATIVE MEMORY

Wulfram Gerstner

EPFL, Lausanne, Switzerland

5.1 Introduction
- networks of neuron
- systems for computing
- associative memory

5.2 Classification by similarity

5.3 Detour: Magnetic Materials

5.4 Hopfield Model

5.5 Learning of Associations

5.6 Storage Capacity

Systems for computing and information processing

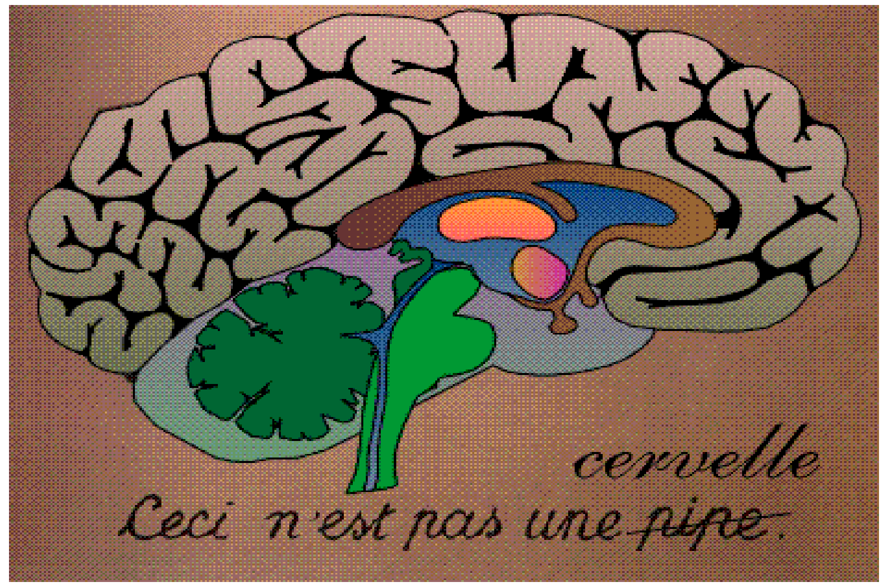

Brain

Computer

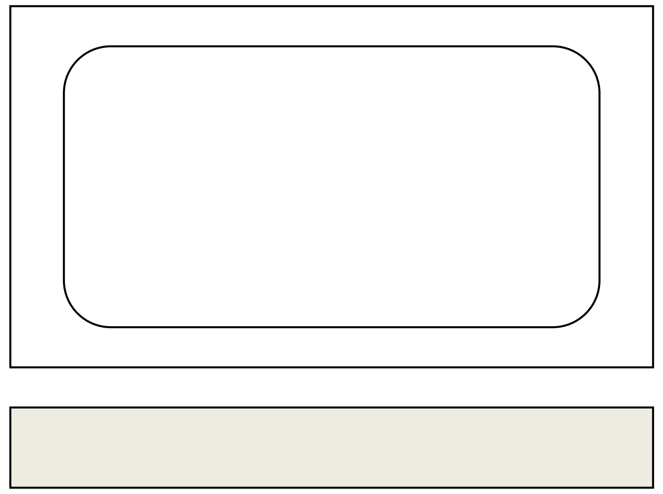

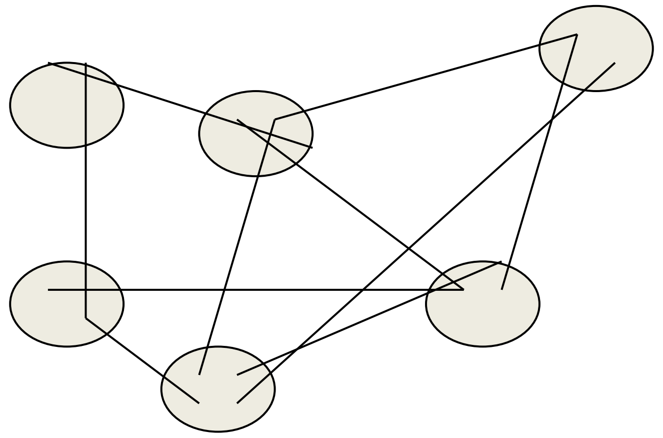

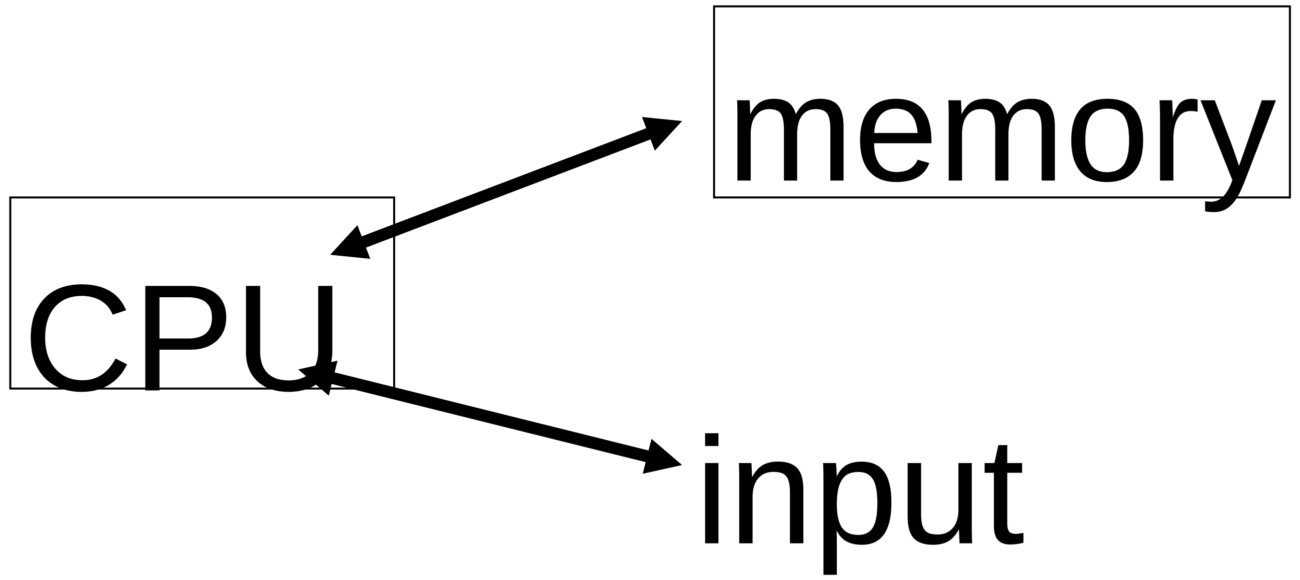

Distributed architecture

(10^{10} proc. Elements/neurons)

No separation of processing and memory

Von Neumann architecture

1 CPU

(10^{10} transistors)

Systems for computing and information processing

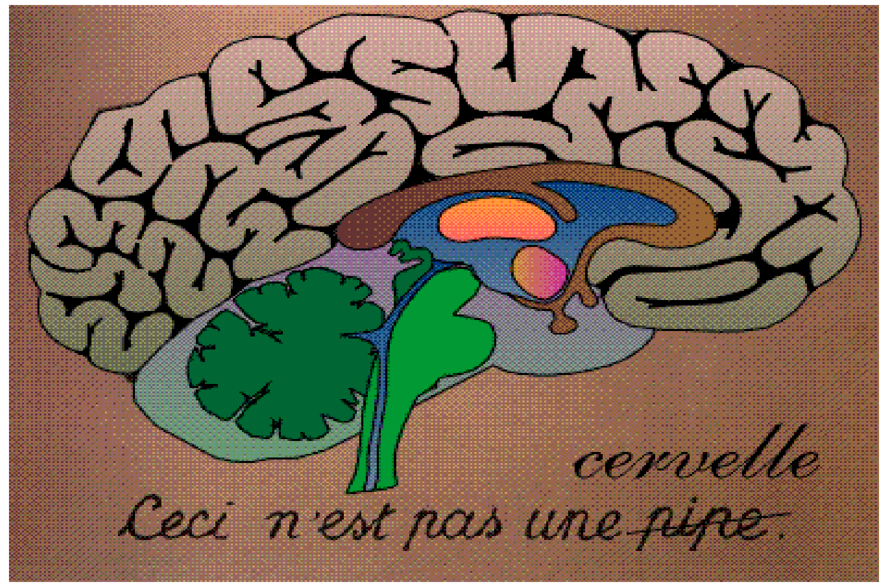

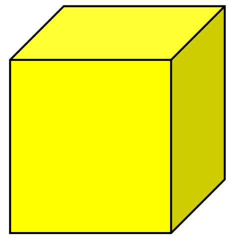

1mm

10 000 neurons
3 km wire

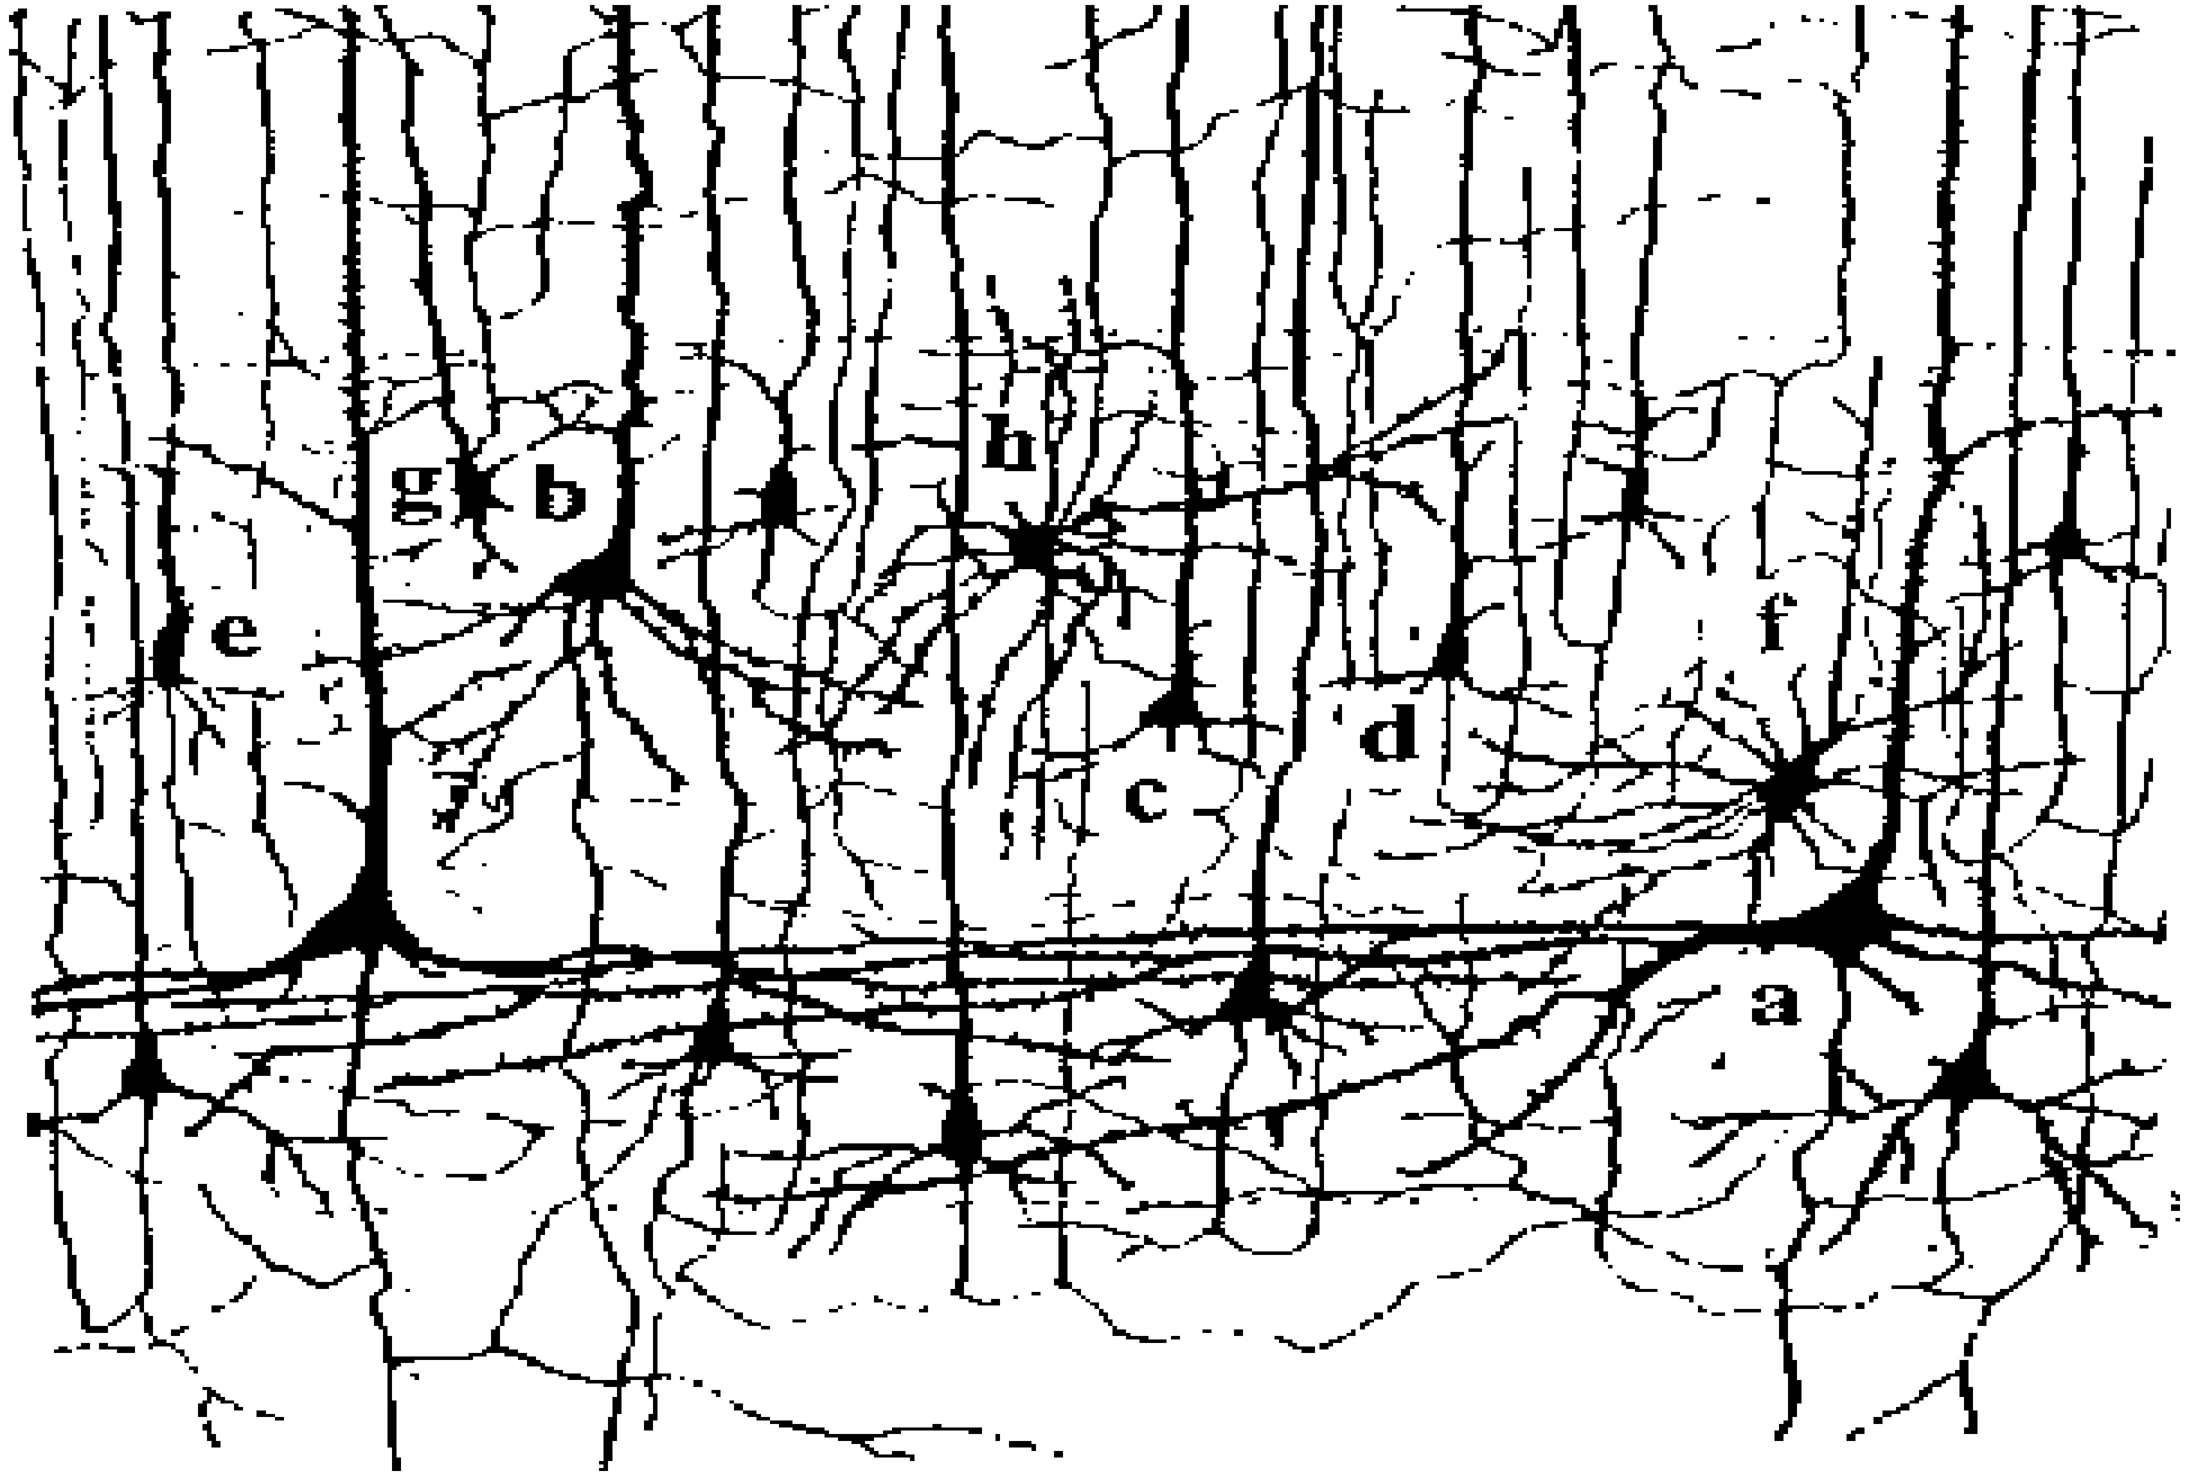

Systems for computing and information processing

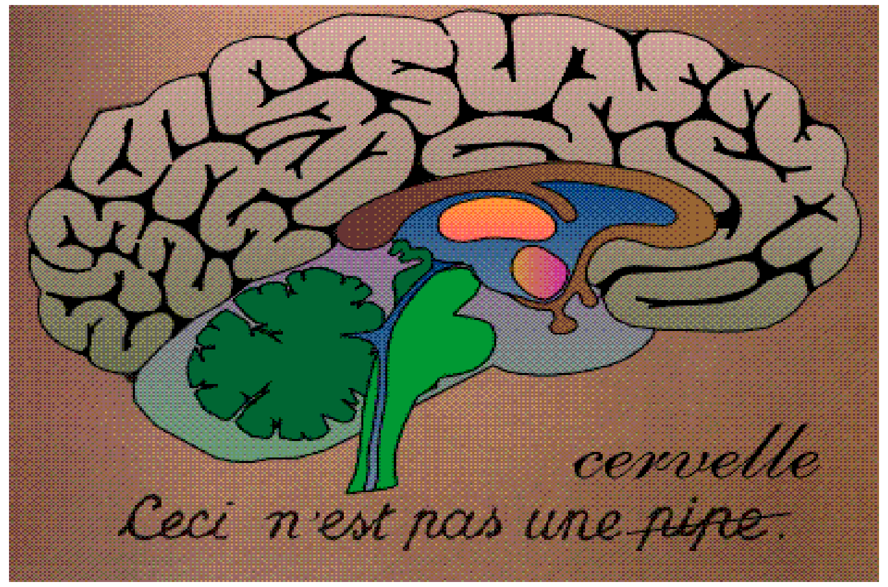

Brain

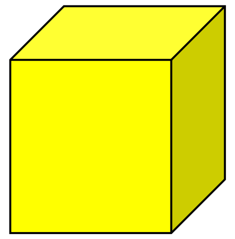

1mm

10 000 neurons
3 km wire

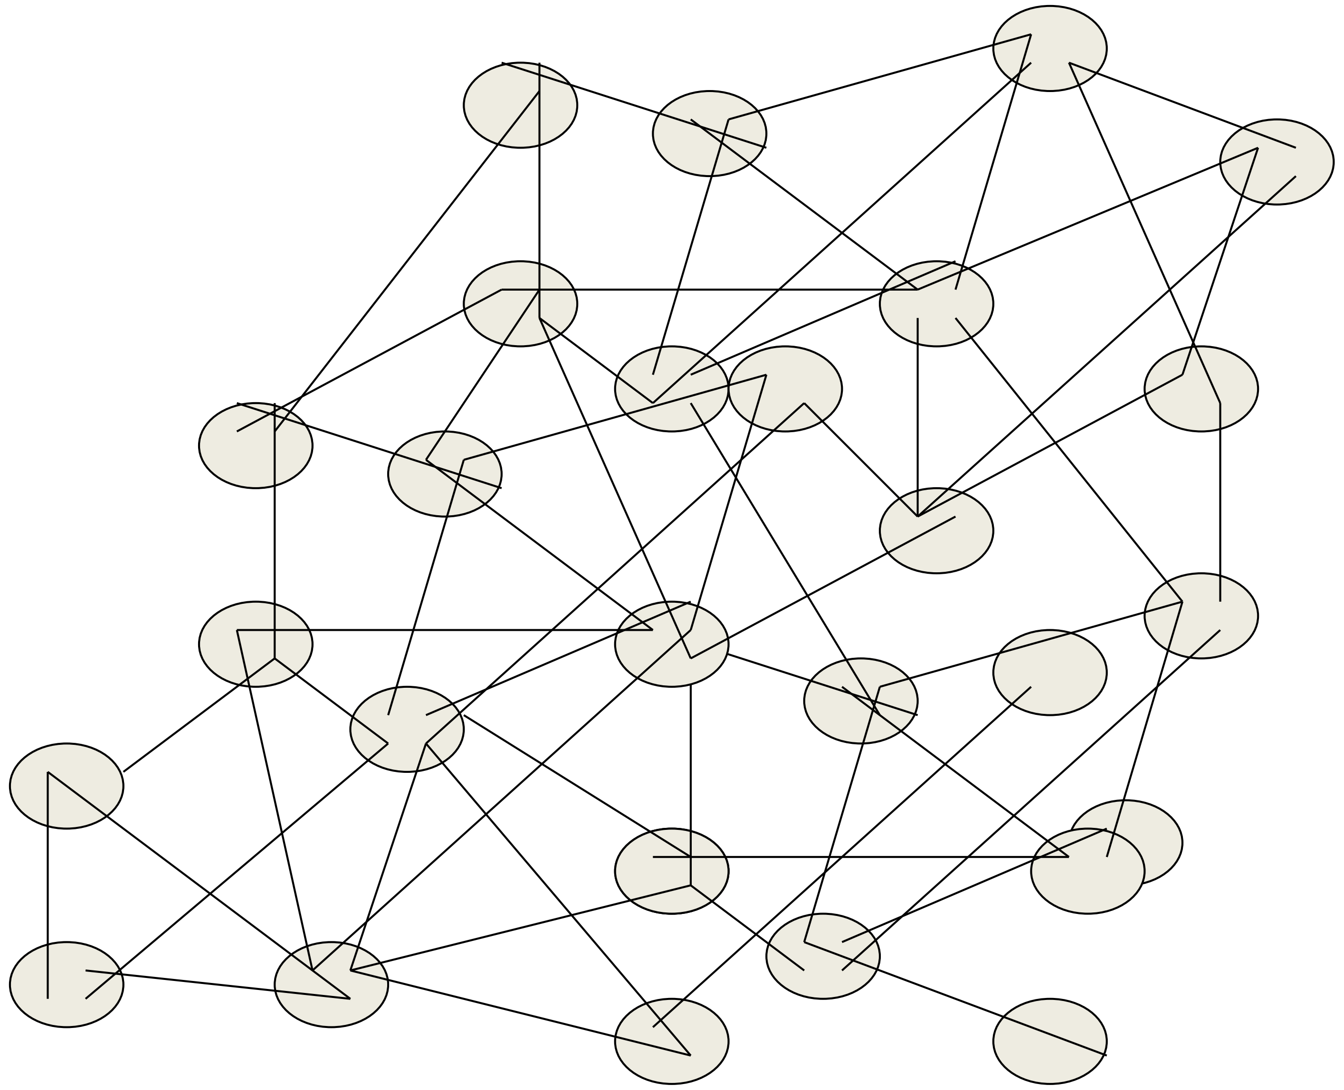

Distributed architecture

10^{10} neurons

10^{4} connections/neurons

No separation of processing and memory

Associations, Associative Memory

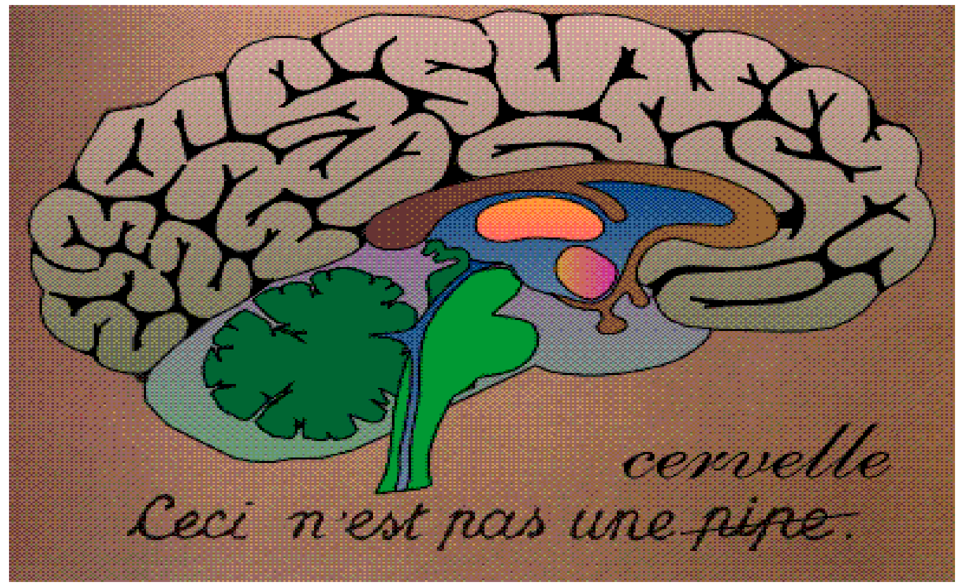

Read this text ***NOW!***

pattern completion/word recognition

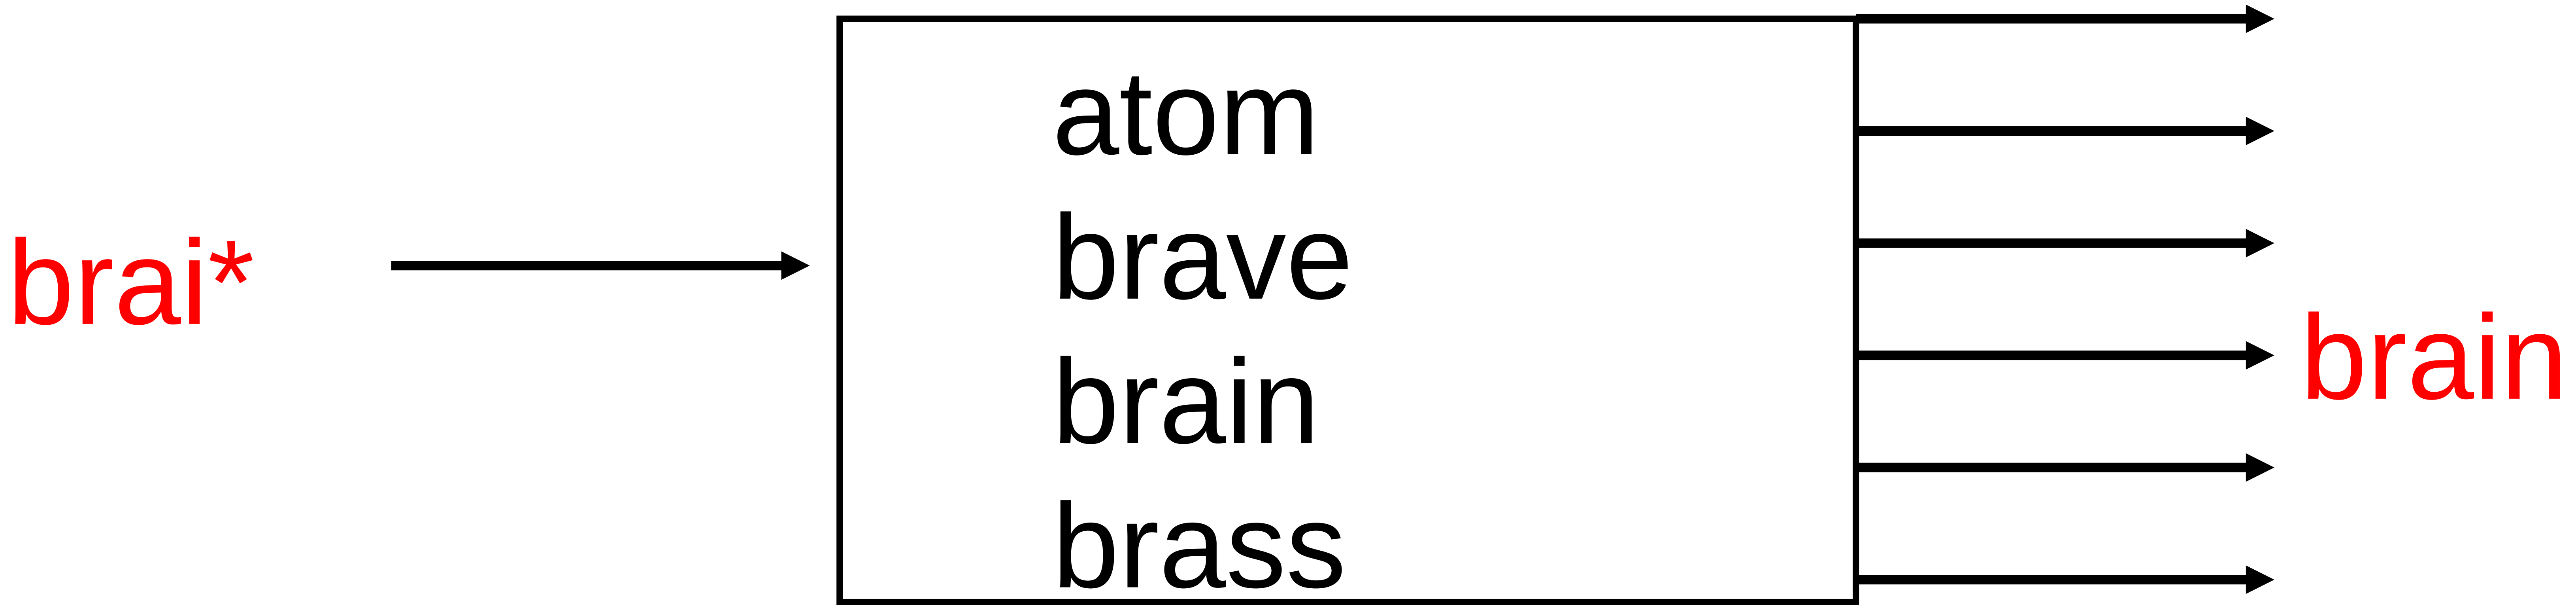

Noisy word

List of words

Output the closest one

Your brain fills in missing information: 'associative memory'

Week 5: Networks of Neurons-Introduction

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

Biological Modeling of Neural Networks

Week 5
NETWORKS of NEURONS and ASSOCIATIVE MEMORY

Wulfram Gerstner
EPFL, Lausanne, Switzerland

5.1 Introduction
- networks of neuron
- systems for computing
- associative memory

5.2 Classification by similarity

5.3 Detour: Magnetic Materials

5.4 Hopfield Model

5.5 Learning of Associations

5.6 Storage Capacity

5.2 Classification by similarity: **pattern recognition**

image

Classification:
comparison
with prototypes

A
B
T
Z

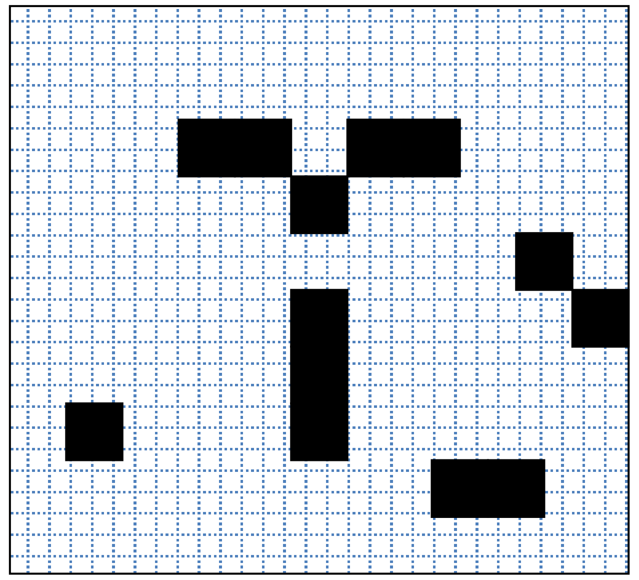

Noisy image

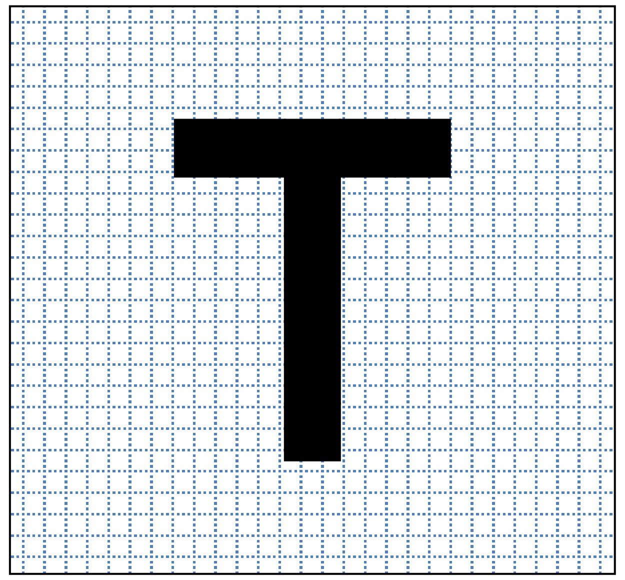

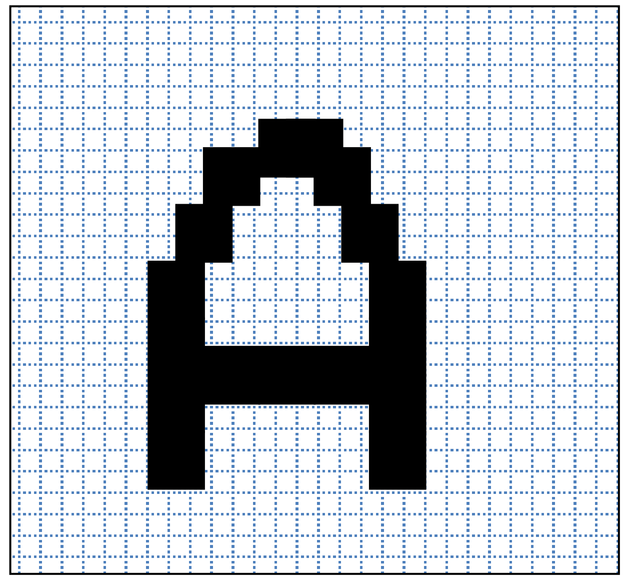

Prototypes

5.2 Classification by similarity: **pattern recognition**

Classification by closest prototype

Blackboard:

$$\left| x - p^{T} \right| \leq \left| x - p^{A} \right|$$

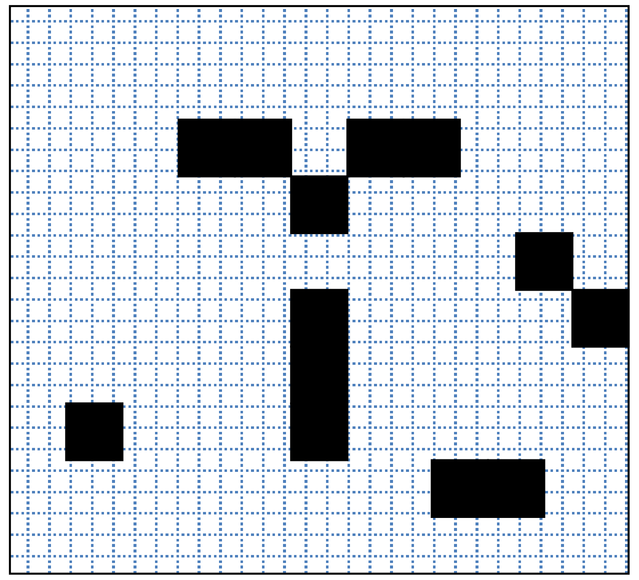

Noisy image

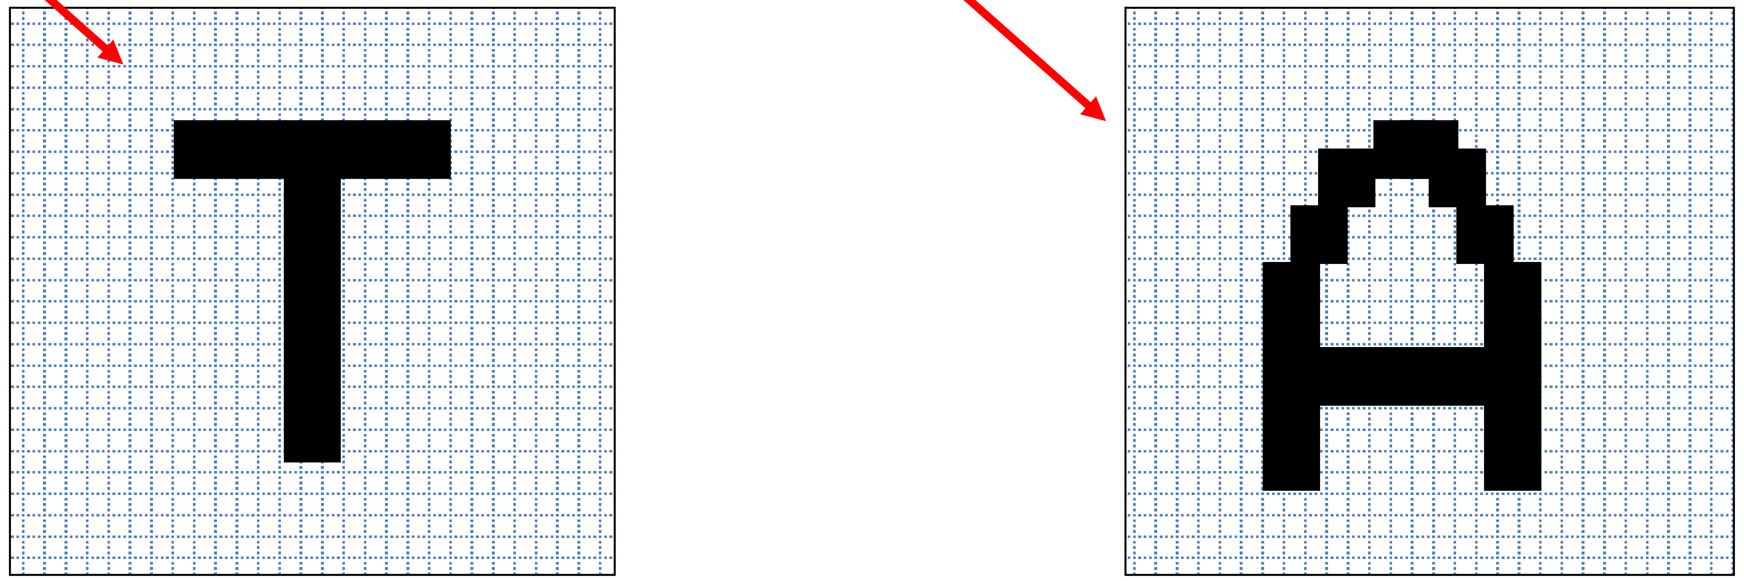

Prototypes

5.2 pattern recognition and Pattern completion

Aim: Understand Associative Memory

Noisy image

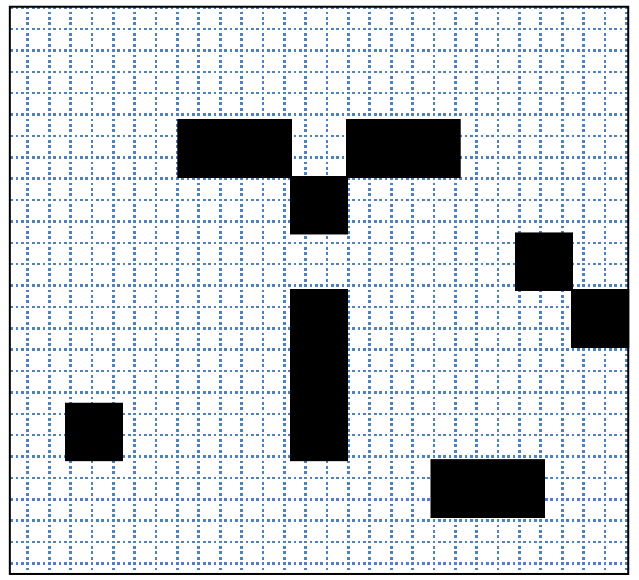

Partial word

Associative memory/ collective computation

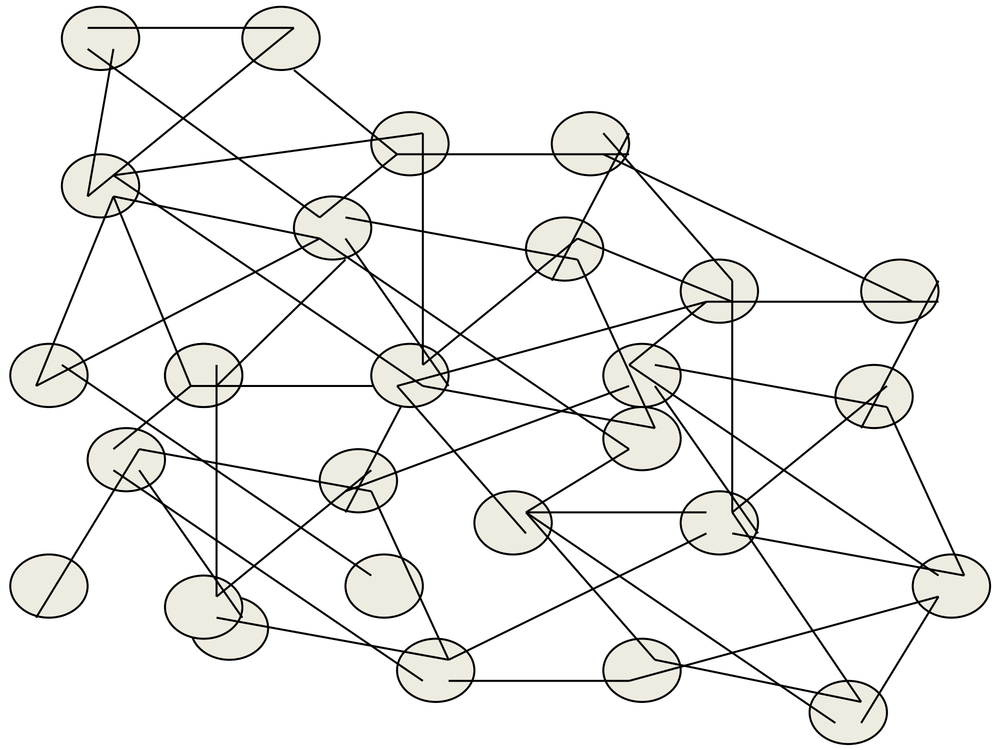

Brain-style computation

Full image

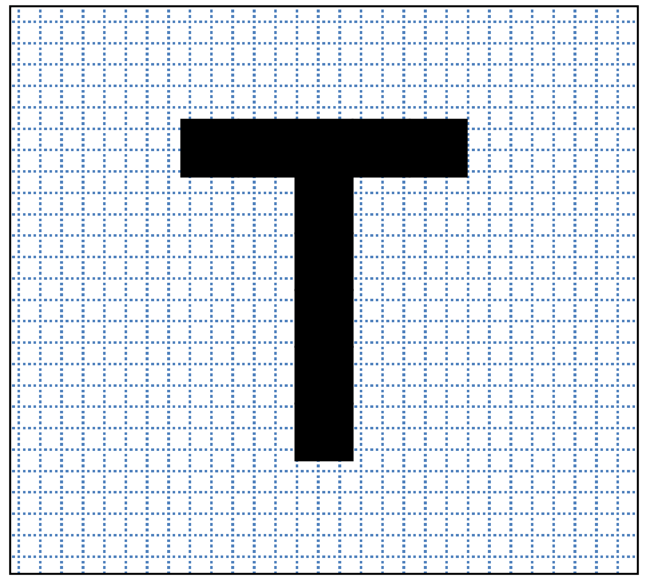

Full word

Quiz 5.1: Connectivity

A typical neuron in the brain makes connections
-To 6-20 neighbors
-To 100-200 neurons nearby
-To more than 1000 neurons nearby
-To more than 1000 neurons nearby or far away.

In a typical cristal in nature, each atom interacts
-with 6-20 neighbors
-with 100-200 neurons nearby
-with more than 1000 neurons nearby
-with more than 1000 neurons nearby or far away.

Week 5: Networks of Neurons-Introduction

ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

Biological Modeling of Neural Networks

Week 5
NETWORKS of NEURONS and ASSOCIATIVE MEMORY

Wulfram Gerstner

EPFL, Lausanne, Switzerland

√ 5.1 Introduction
- networks of neuron
- systems for computing
- associative memory

√ 5.2 Classification by similarity

5.3 Detour: Magnetic Materials

5.4 Hopfield Model

5.5 Learning of Associations

5.6 Storage Capacity

5.3 Detour: magnetism

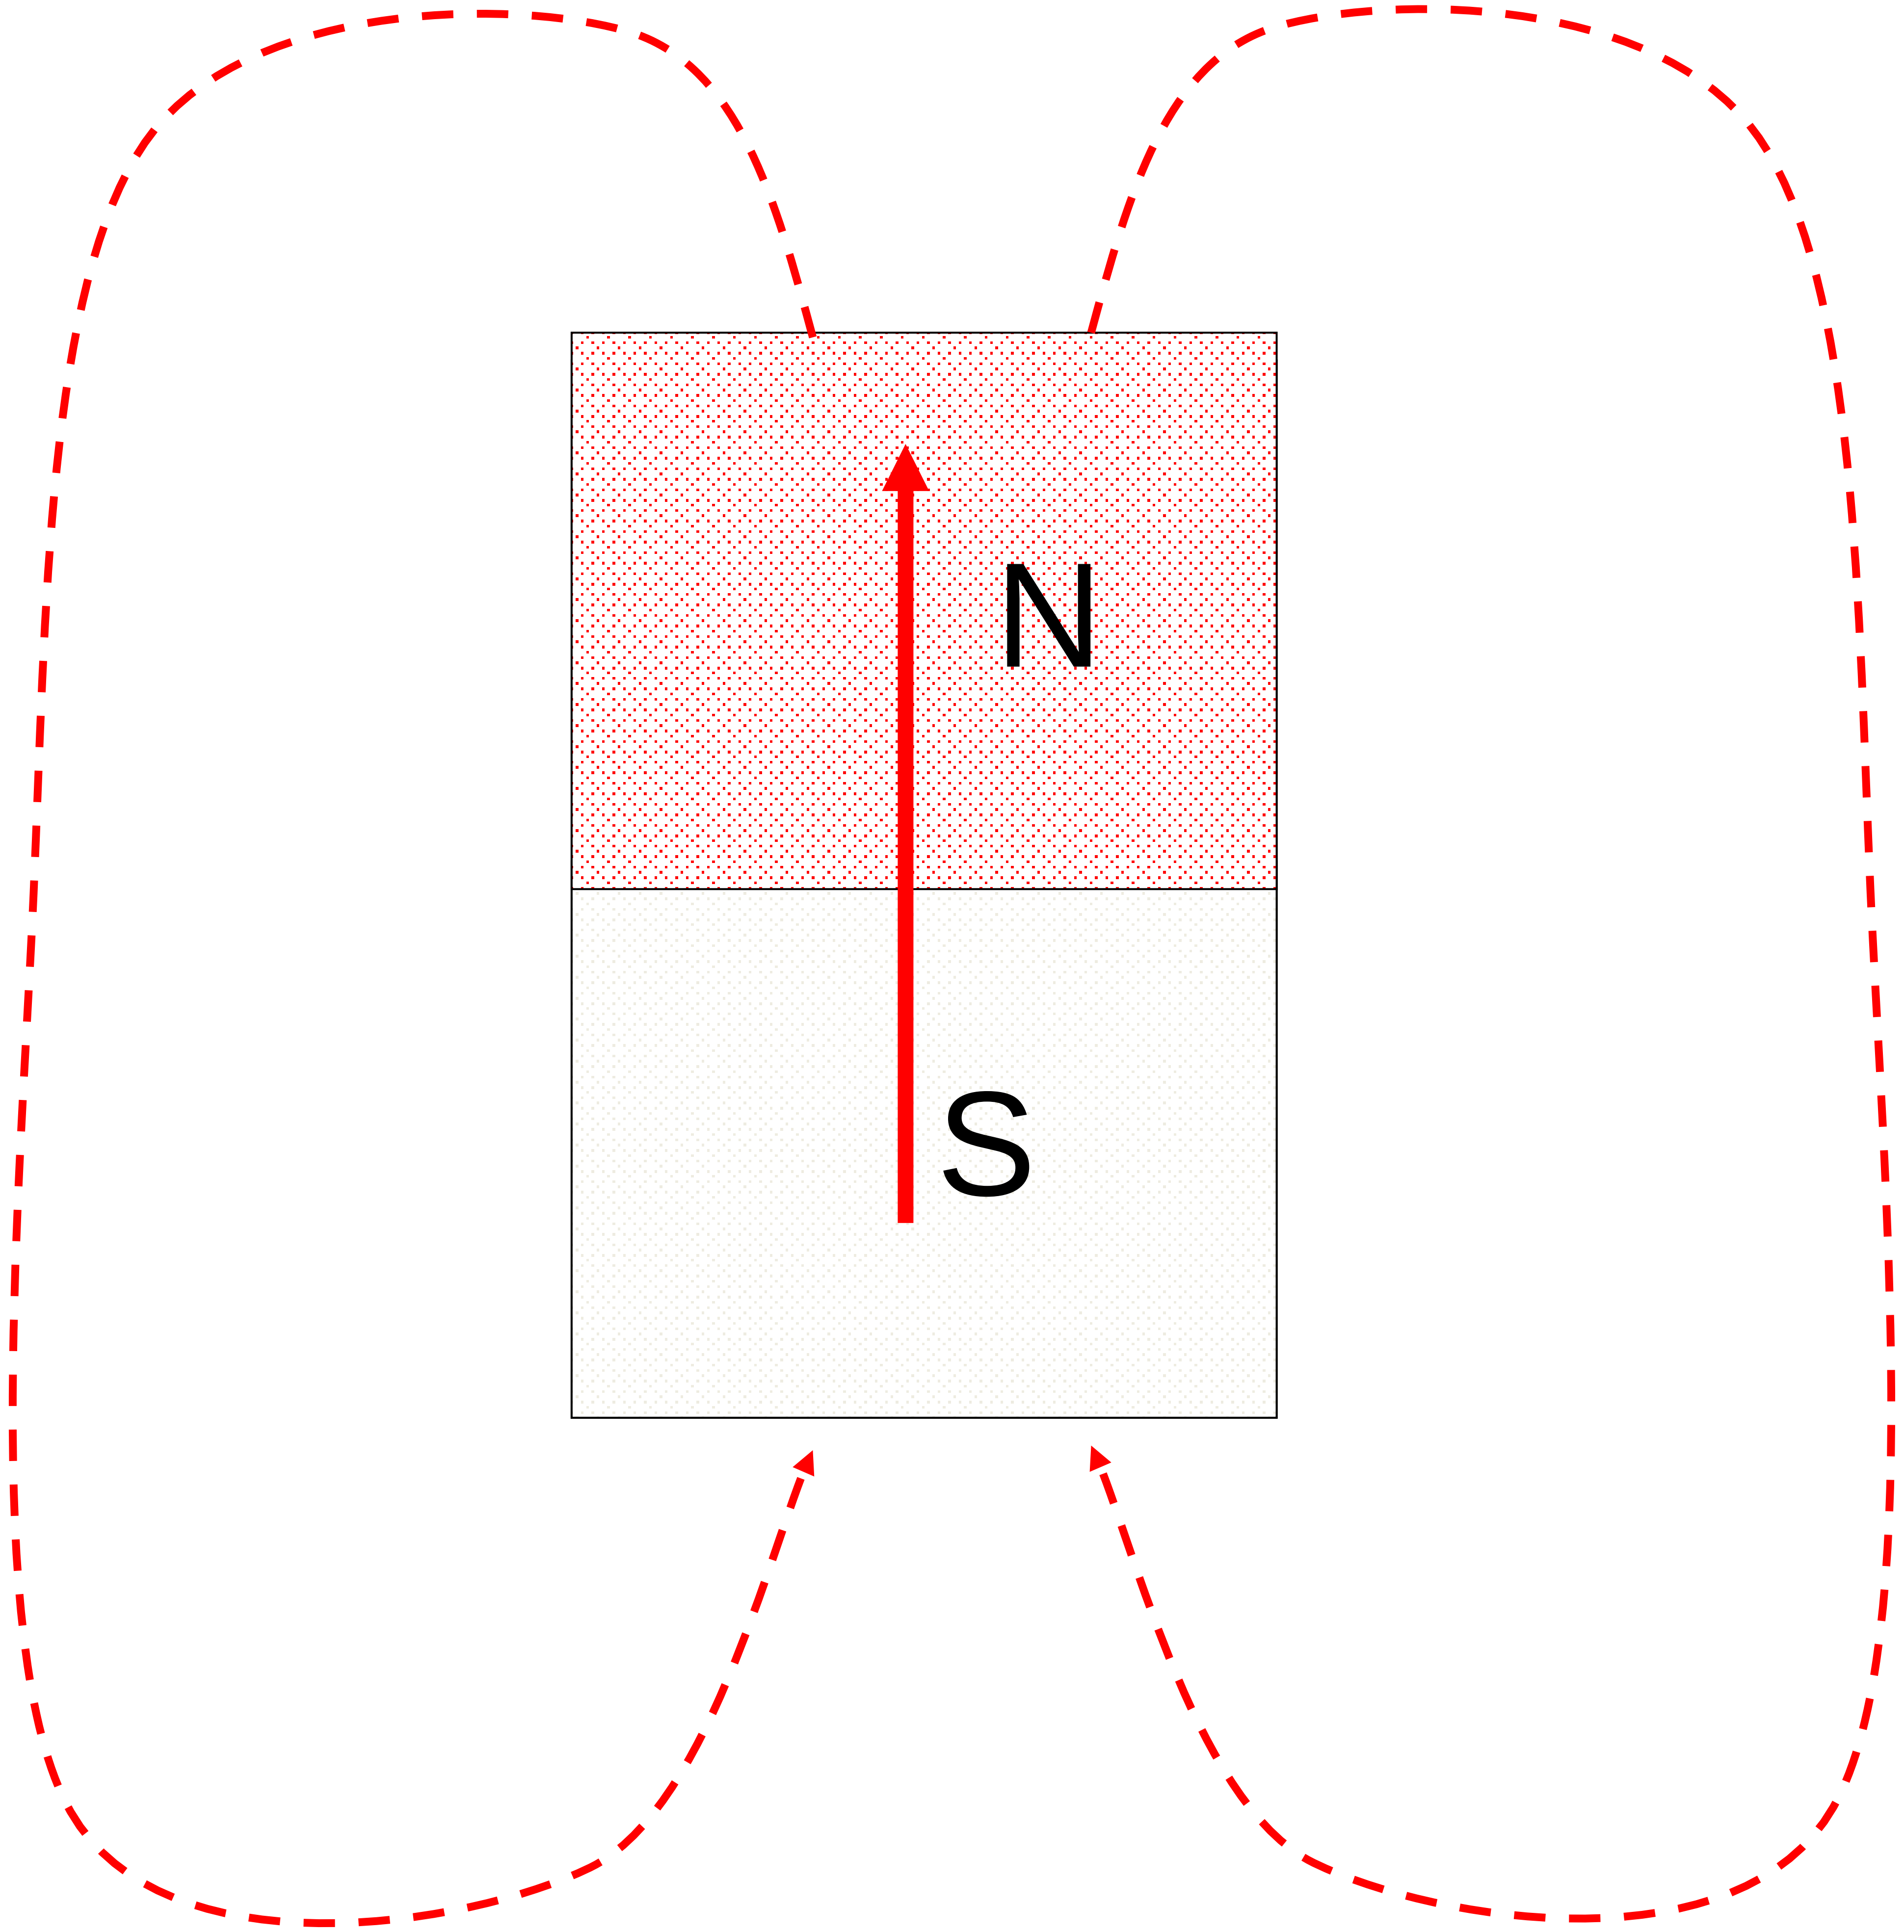

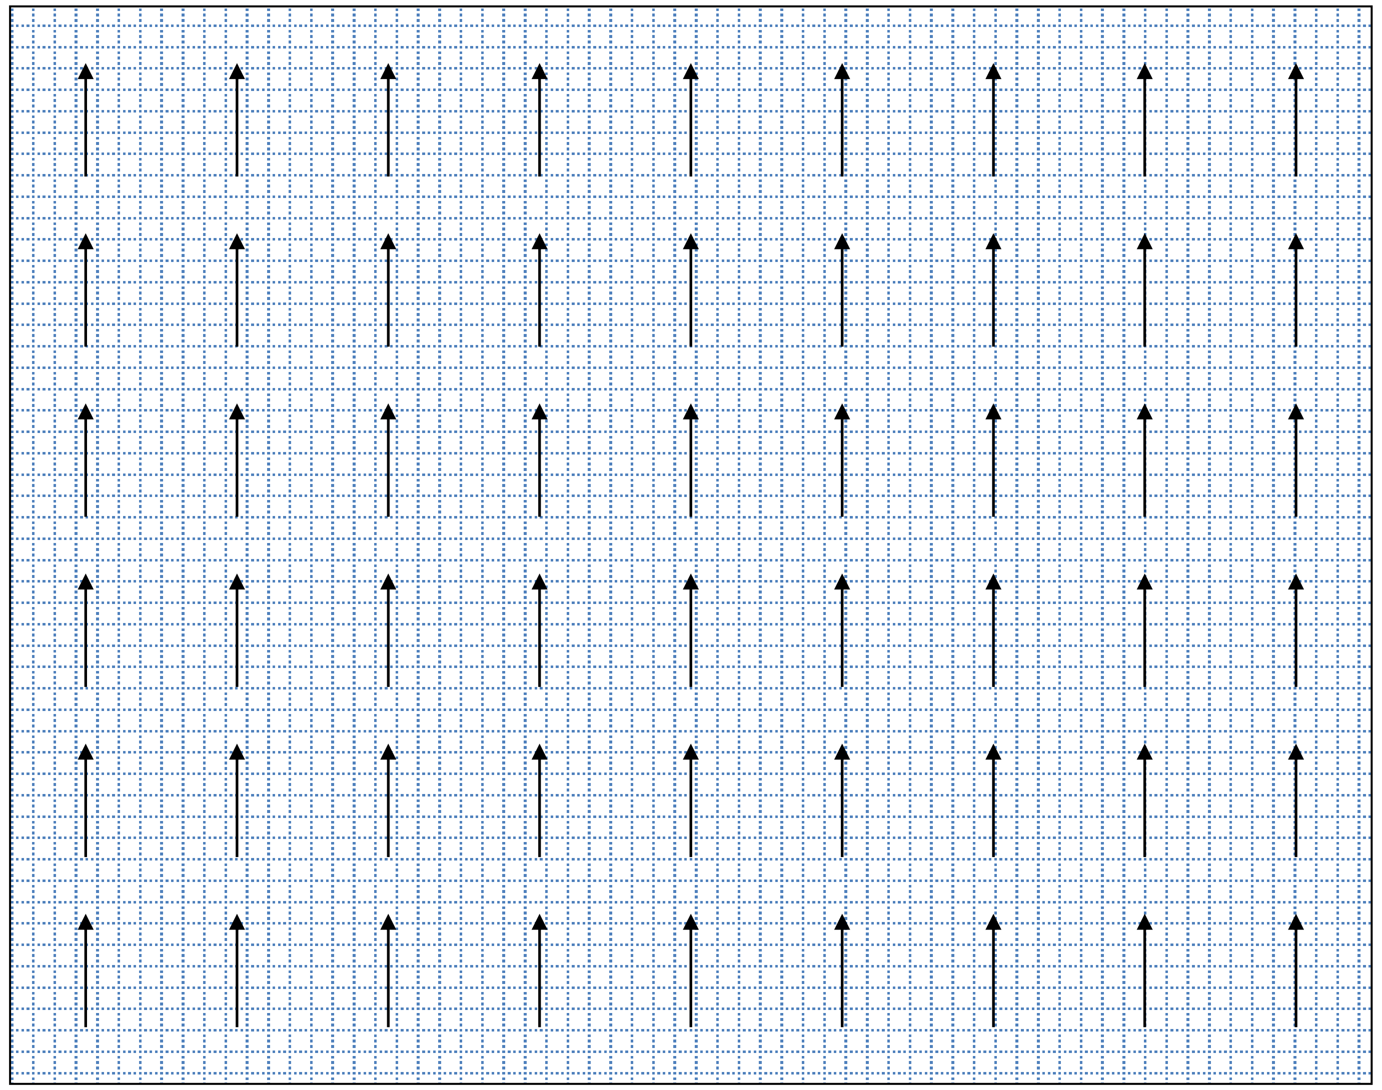

5.3 Detour: magnetism

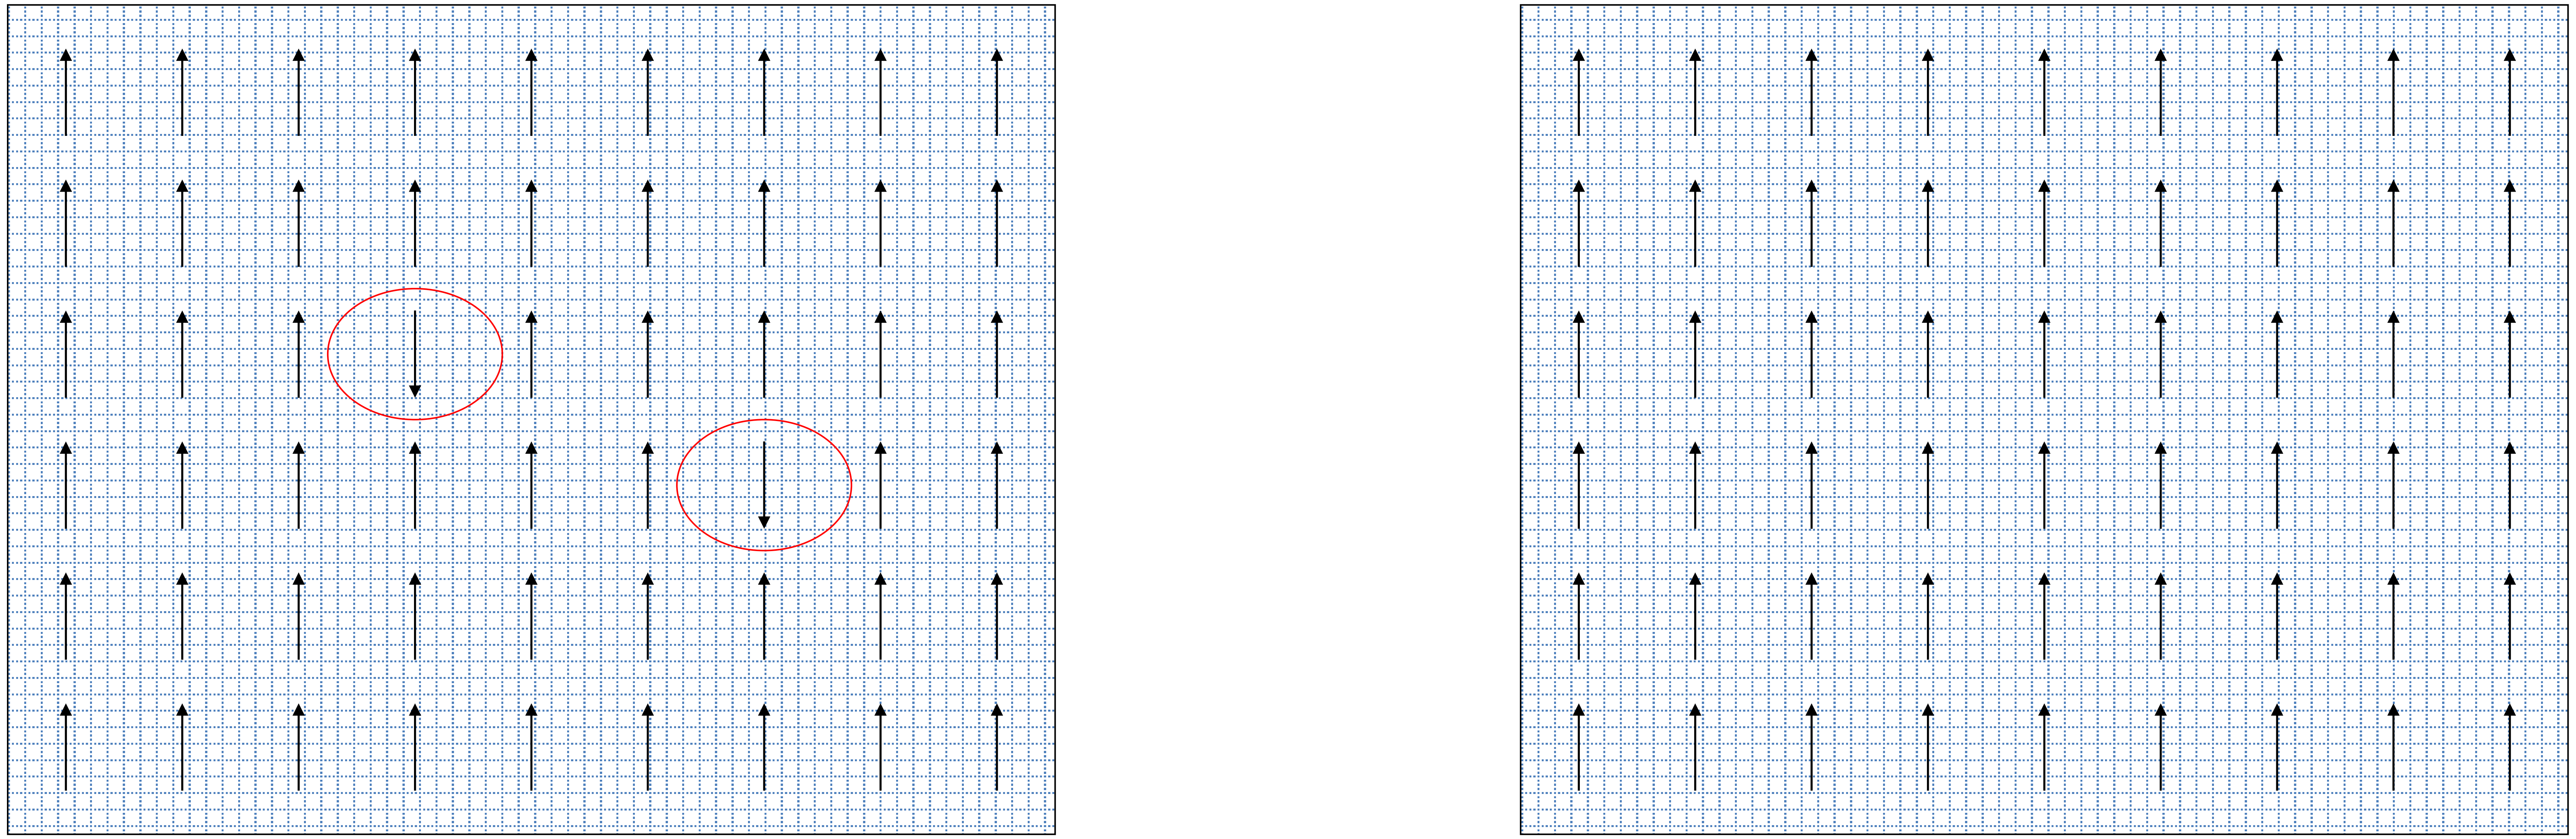

Noisy magnet → pure magnet

5.3 Detour: magnetism

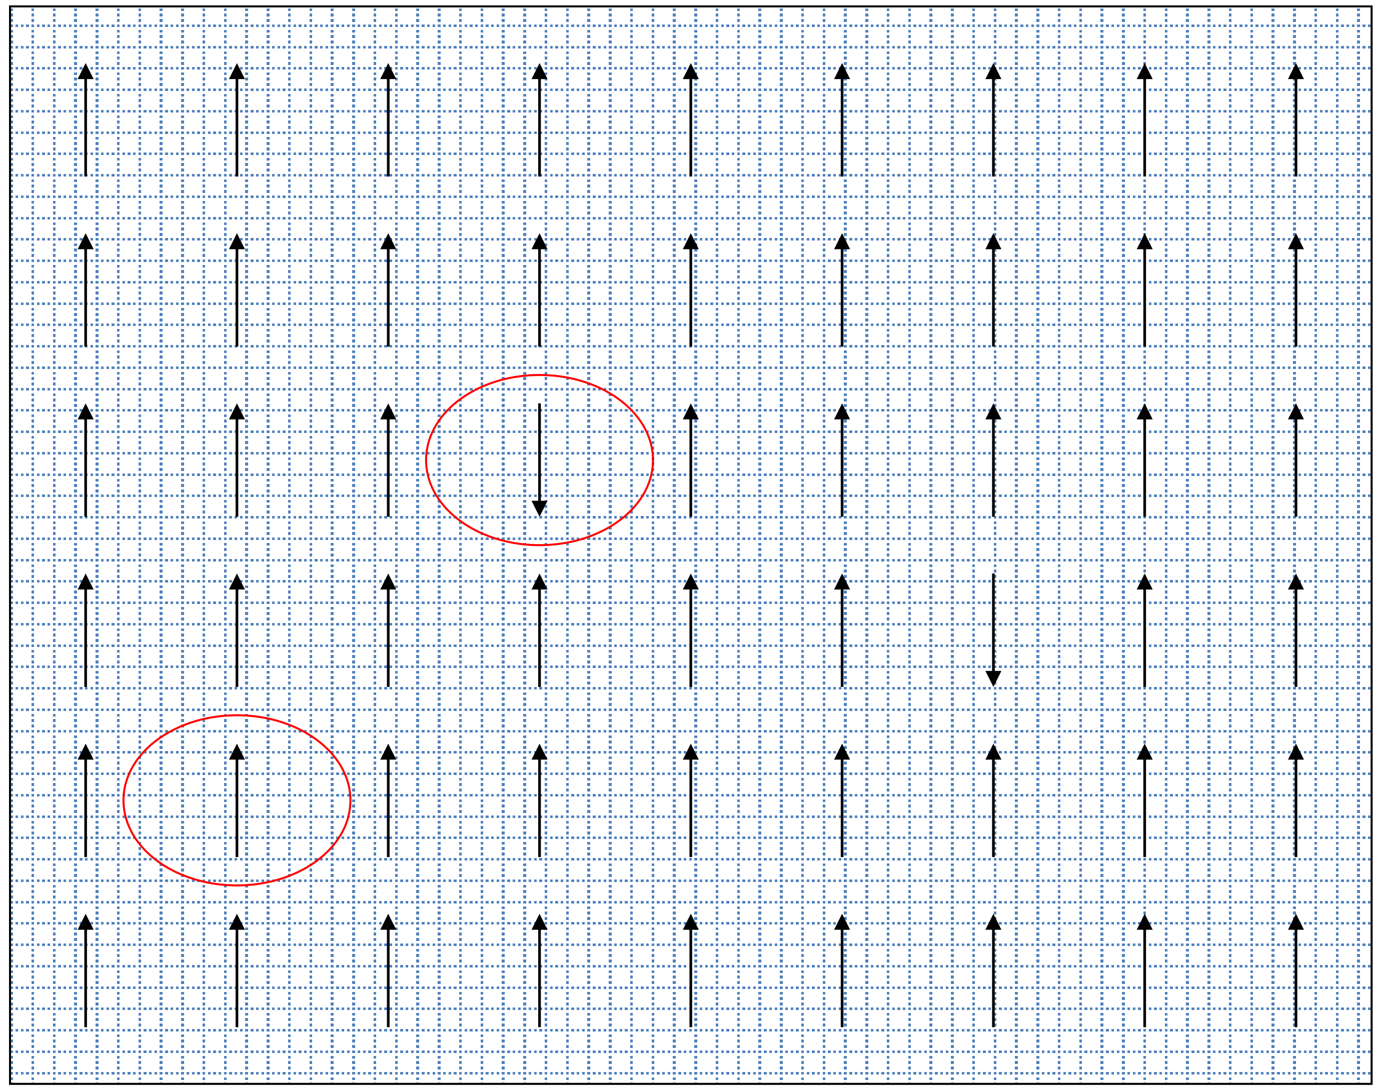

Elementary magnet

$S_i = +1$

$S_i = -1$

Blackboard: example

dynamics

$$S_i(t+1) = \mathrm{sgn}[\sum_j S_j(t)]$$

Sum over all interactions with i

5.3 Detour: magnetism

Anti-ferromagnet

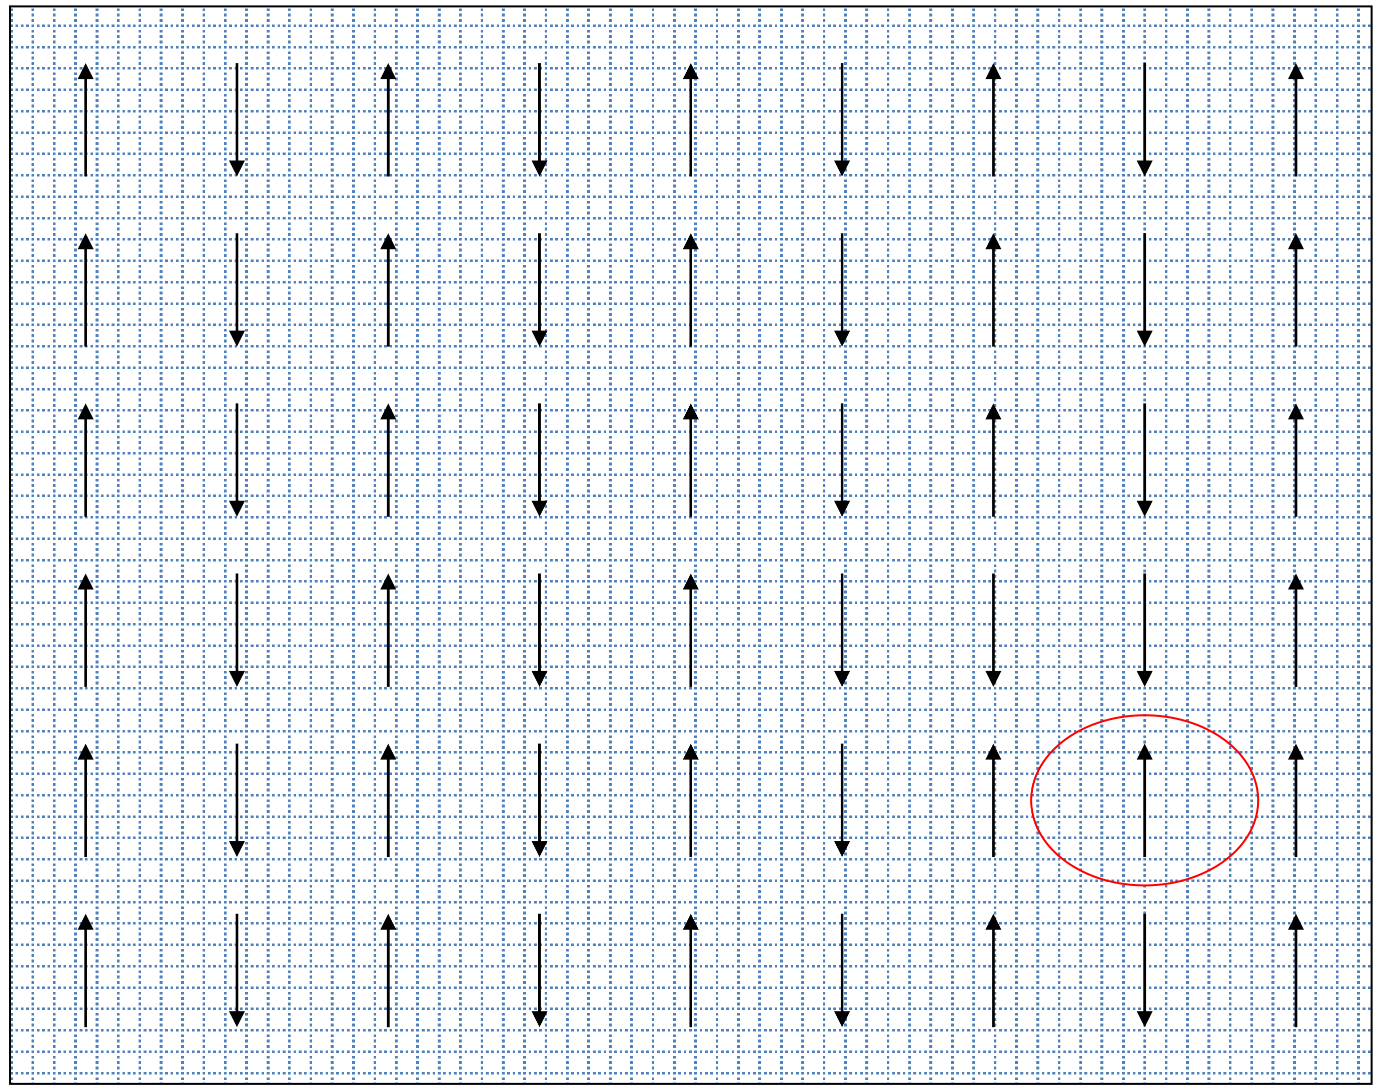

Elementary magnet

$S_i = +1$

$S_i = -1$

$w_{ij} = +1$

$w_{ij} = -1$

dynamics

$$S_i(t+1) = \mathrm{sgn}[\sum_j w_{ij} S_j(t)]$$

Sum over all interactions with i

blackboard

5.3 Magnetism and memory patterns

blackboard

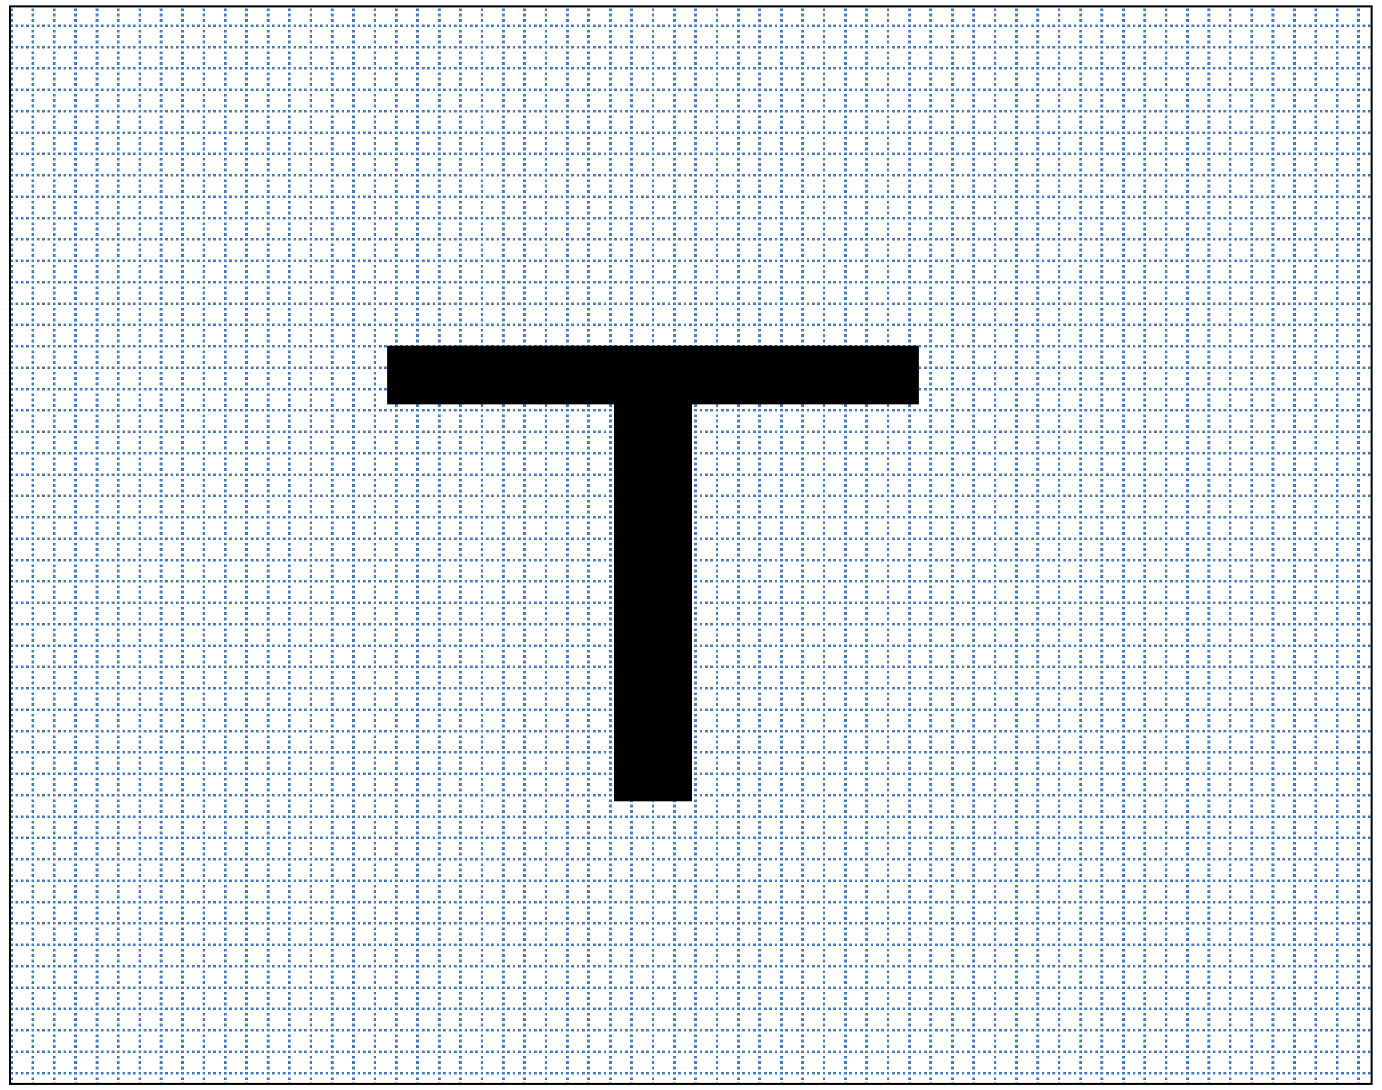

Elementary pixel

■ $S_i = +1$

□ $S_i = -1$

■ ↔ ■ $w_{ij} = +1$

□ ↔ □ $w_{ij} = +1$

□ ↔ ■ $w_{ij} = -1$

dynamics

$$S_i(t+1) = \mathrm{sgn}[\sum_j w_{ij} S_j(t)]$$

Sum over all interactions with i

Hopfield model:
Several patterns→ next section

Exercise 1: Associative memory (1 pattern)

Next lecture at
10h15

Elementary pixel

$S_i = +1$

$S_i = -1$

$w_{ij} = +1$

$w_{ij} = +1$

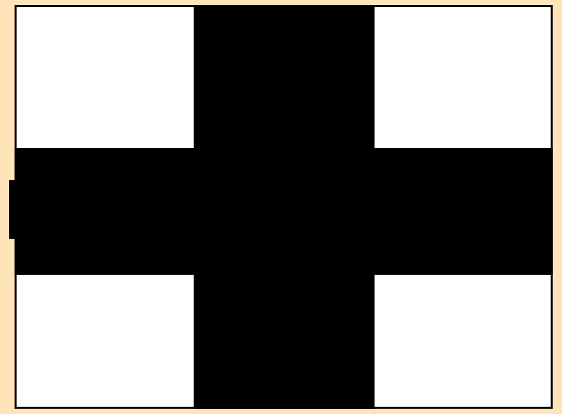

dynamics

$$S_i(t+1) = \mathrm{sgn}[\sum_j w_{ij} S_j(t)]$$

Sum over all
interactions with i

9 neurons

- define appropriate weights
- what happens if one neuron wrong?
- what happens if *n* neurons wrong?

Week 5: Networks of Neurons-Introduction

ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

Biological Modeling of Neural Networks

Week 5
NETWORKS of NEURONS and ASSOCIATIVE MEMORY

Wulfram Gerstner
EPFL, Lausanne, Switzerland

5.4 Hopfield Model of Associative Memory

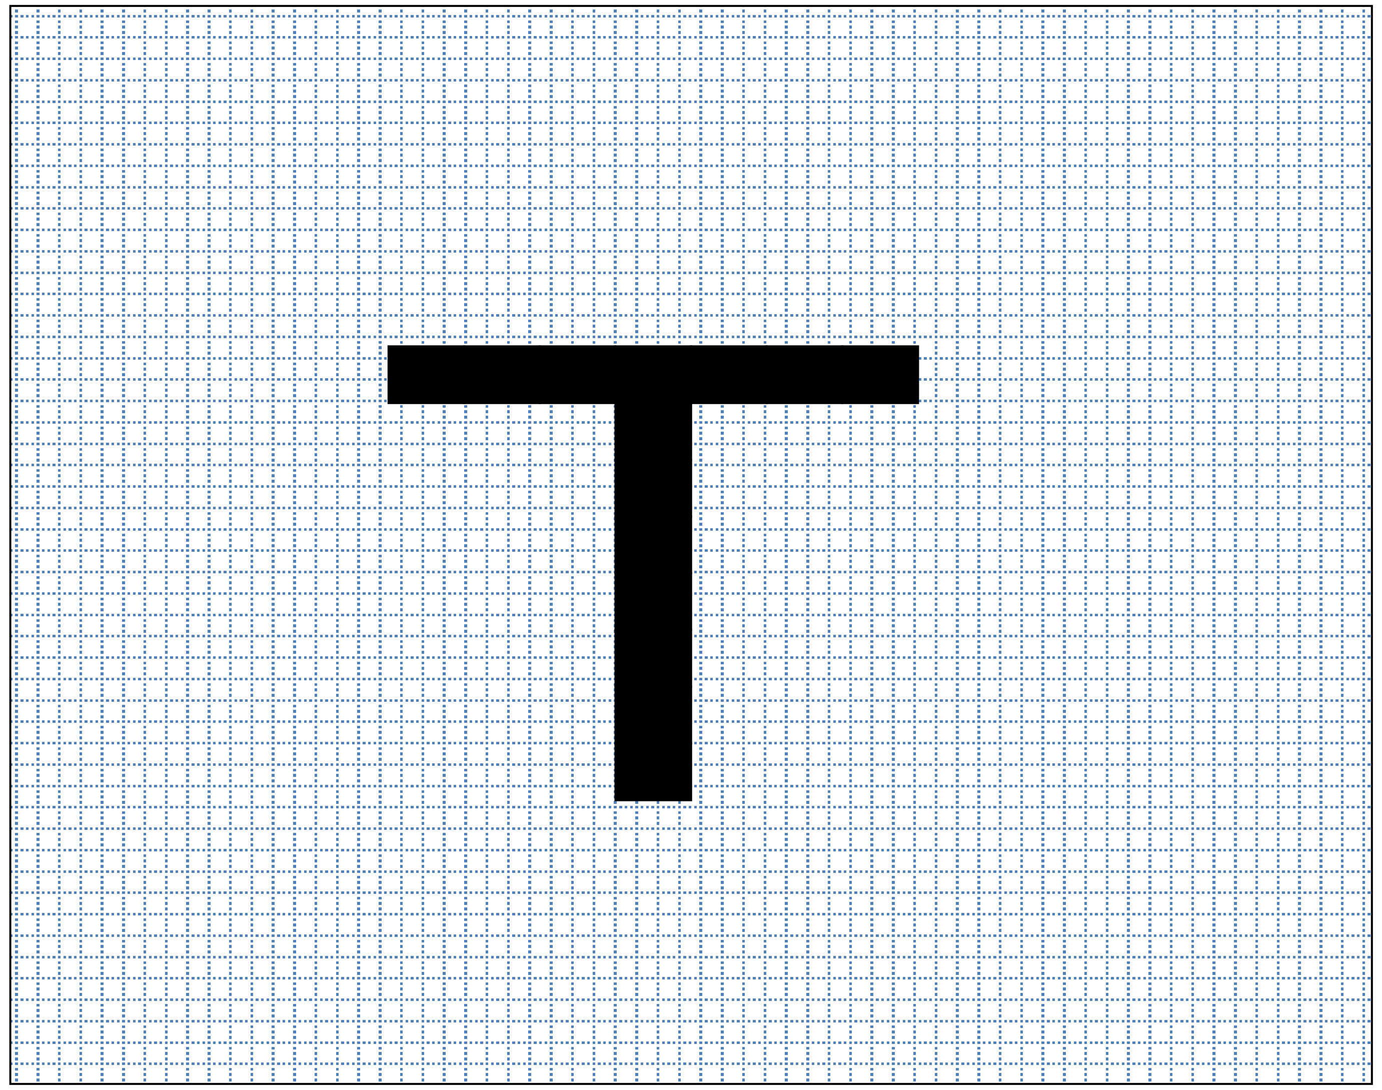

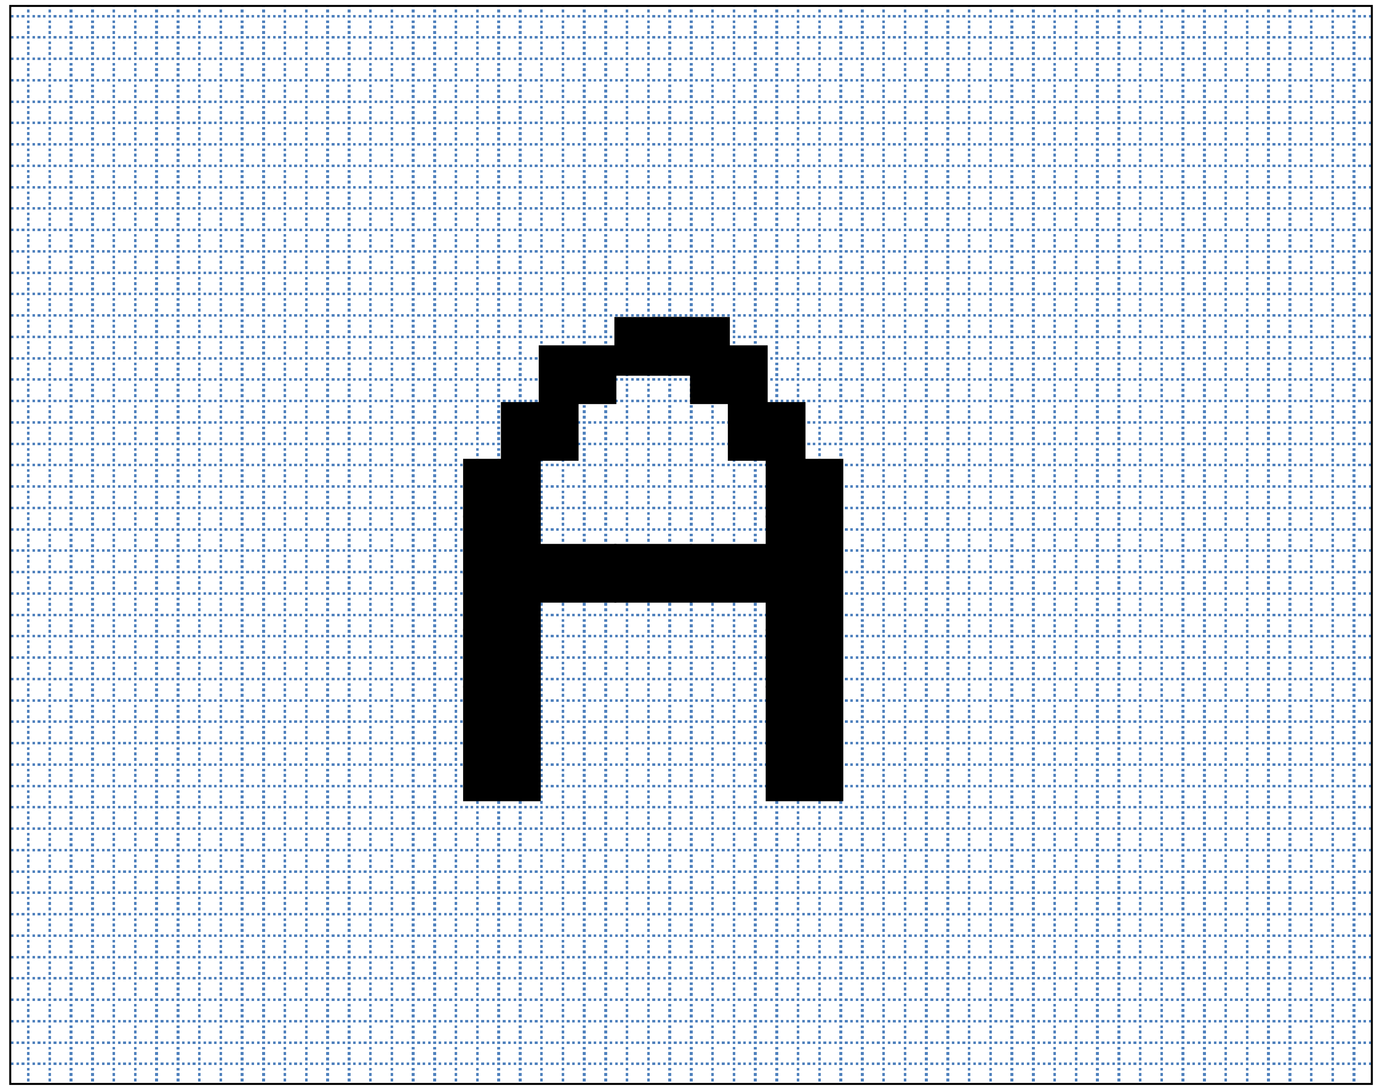

Prototype

$\vec{p}^{1}$

Prototype

$\vec{p}^{2}$

interactions

$$w_{ij} = \sum_{\mu} p_i^{\mu} p_j^{\mu}$$

Sum over all prototypes

dynamics

$$S_i(t+1) = \mathrm{sgn}[\sum_j w_{ij} S_j(t)]$$

Sum over all interactions with i

Hopfield model

5.4 Hopfield Model of Associative Memory

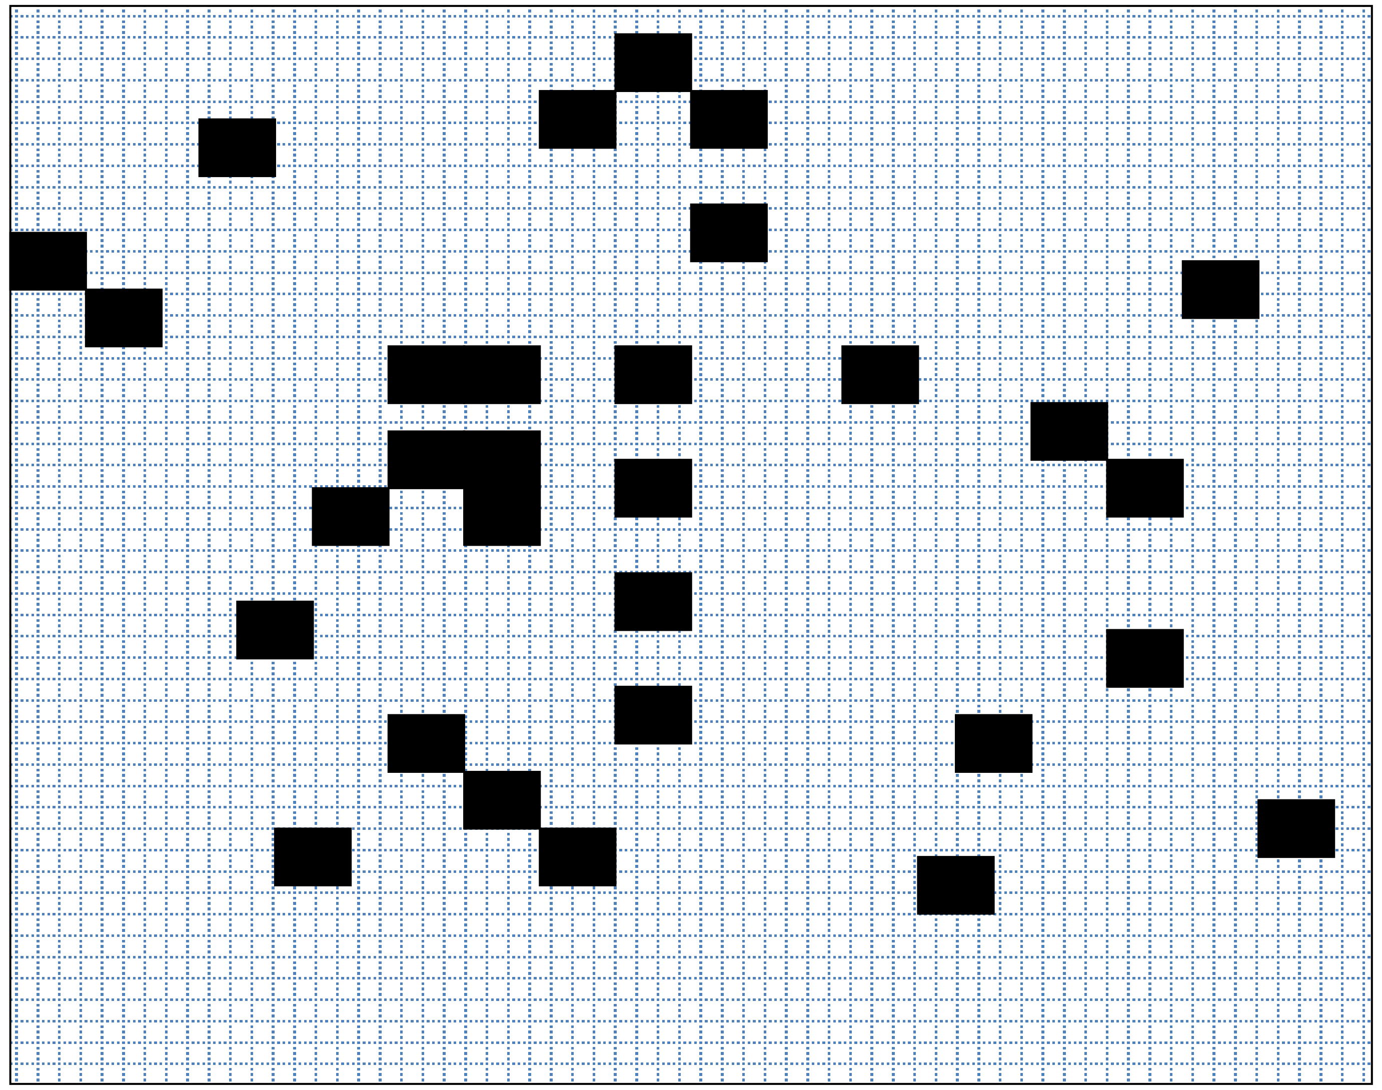

Prototype

$\vec{p}^1$

interactions

$$w_{ij} = \sum_{\mu} p_i^{\mu} p_j^{\mu}$$

Sum over all prototypes

This rule is very good for **random** patterns

It does not work well for correlated patters

DEMO

dynamics

$$S_i(t+1) = \mathrm{sgn}[\sum_j w_{ij} S_j(t)]$$

all interactions with i

Random patterns, fully connected:
Hopfield model

5.4 Hopfield Model of Associative Memory

$$S_i(t+1) = \mathrm{sgn}[\sum_j w_{ij} S_j(t)]$$

$$w_{ij} = \sum_\mu p_i^\mu p_j^\mu$$

Blackboard

overlap $\quad m^\mu(t) = \frac{1}{N}\sum_j \xi^\mu S_j \quad t$

$$m^\mu(t+1) = \frac{1}{N}\sum_j \xi^\mu S_j \quad t+1$$

5.4 Hopfield Model of Associative Memory

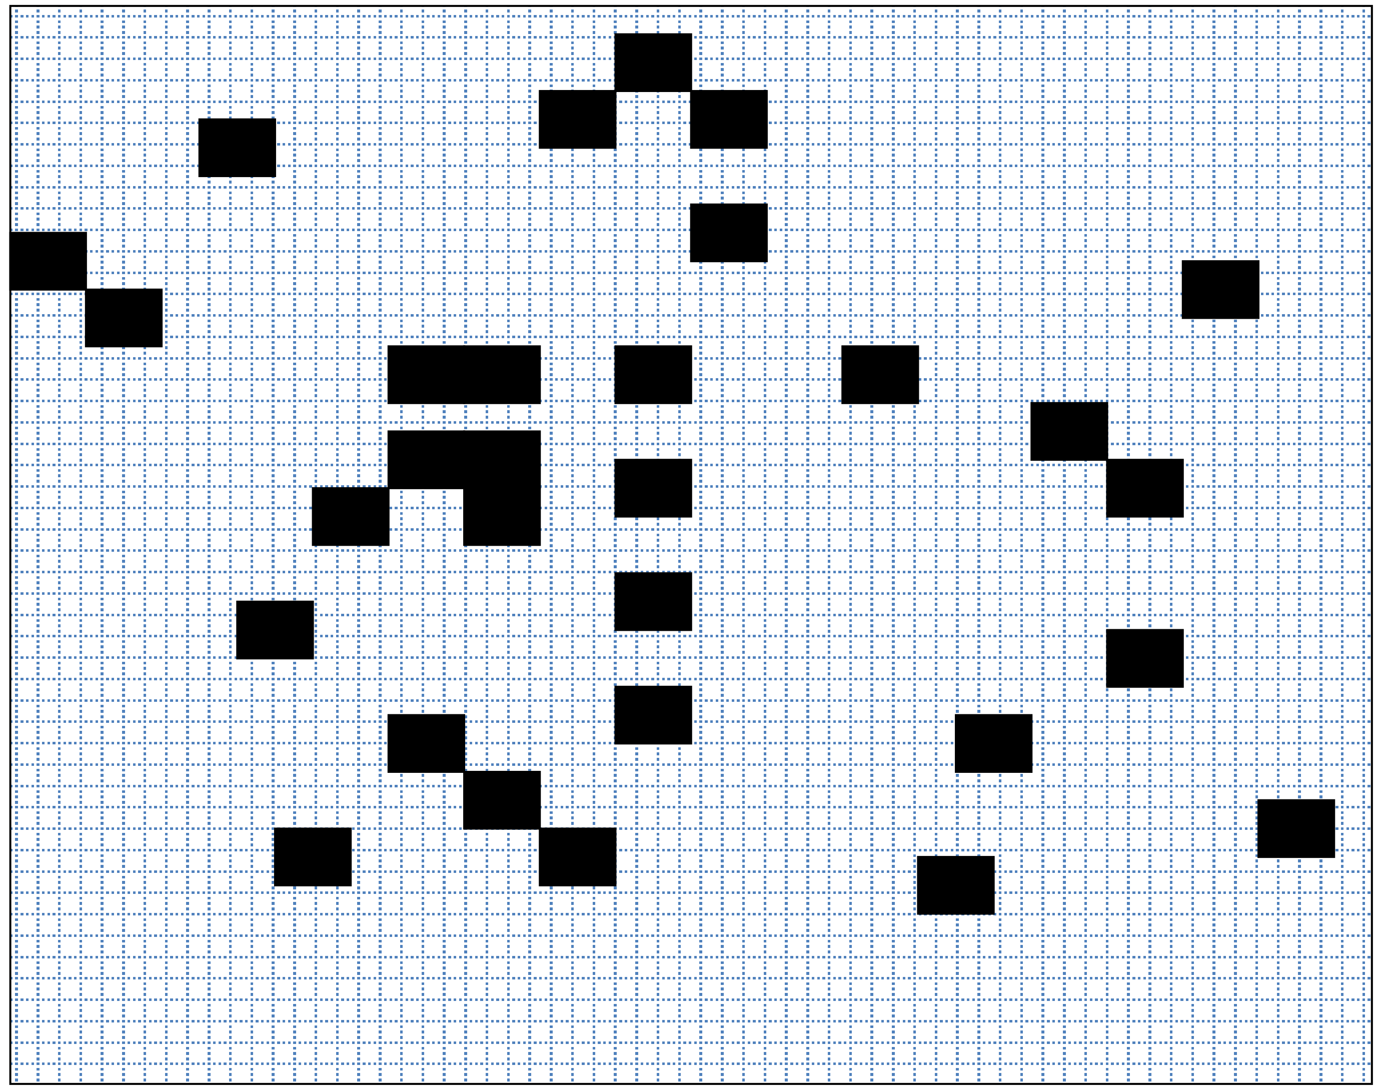

Interacting neurons

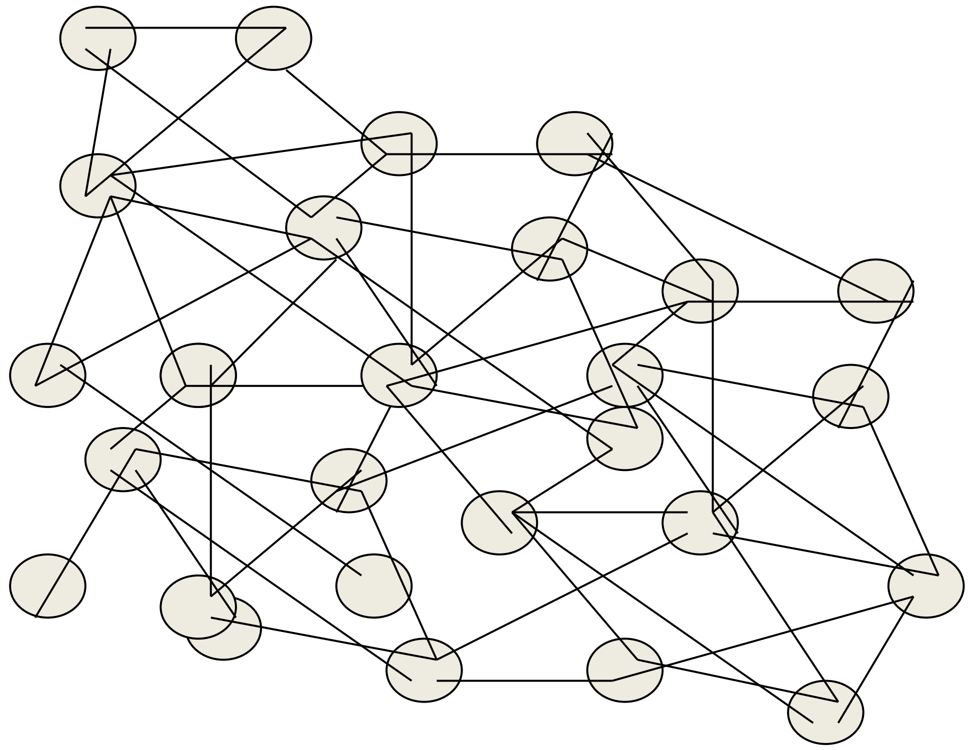

Prototype

$\vec{p}^1$

Finds the closest prototype
i.e. maximal overlap
(similarity) m^μ

Hopfield model

Computation
- without CPU,
- without explicit memory unit

Exercise 3 (now)

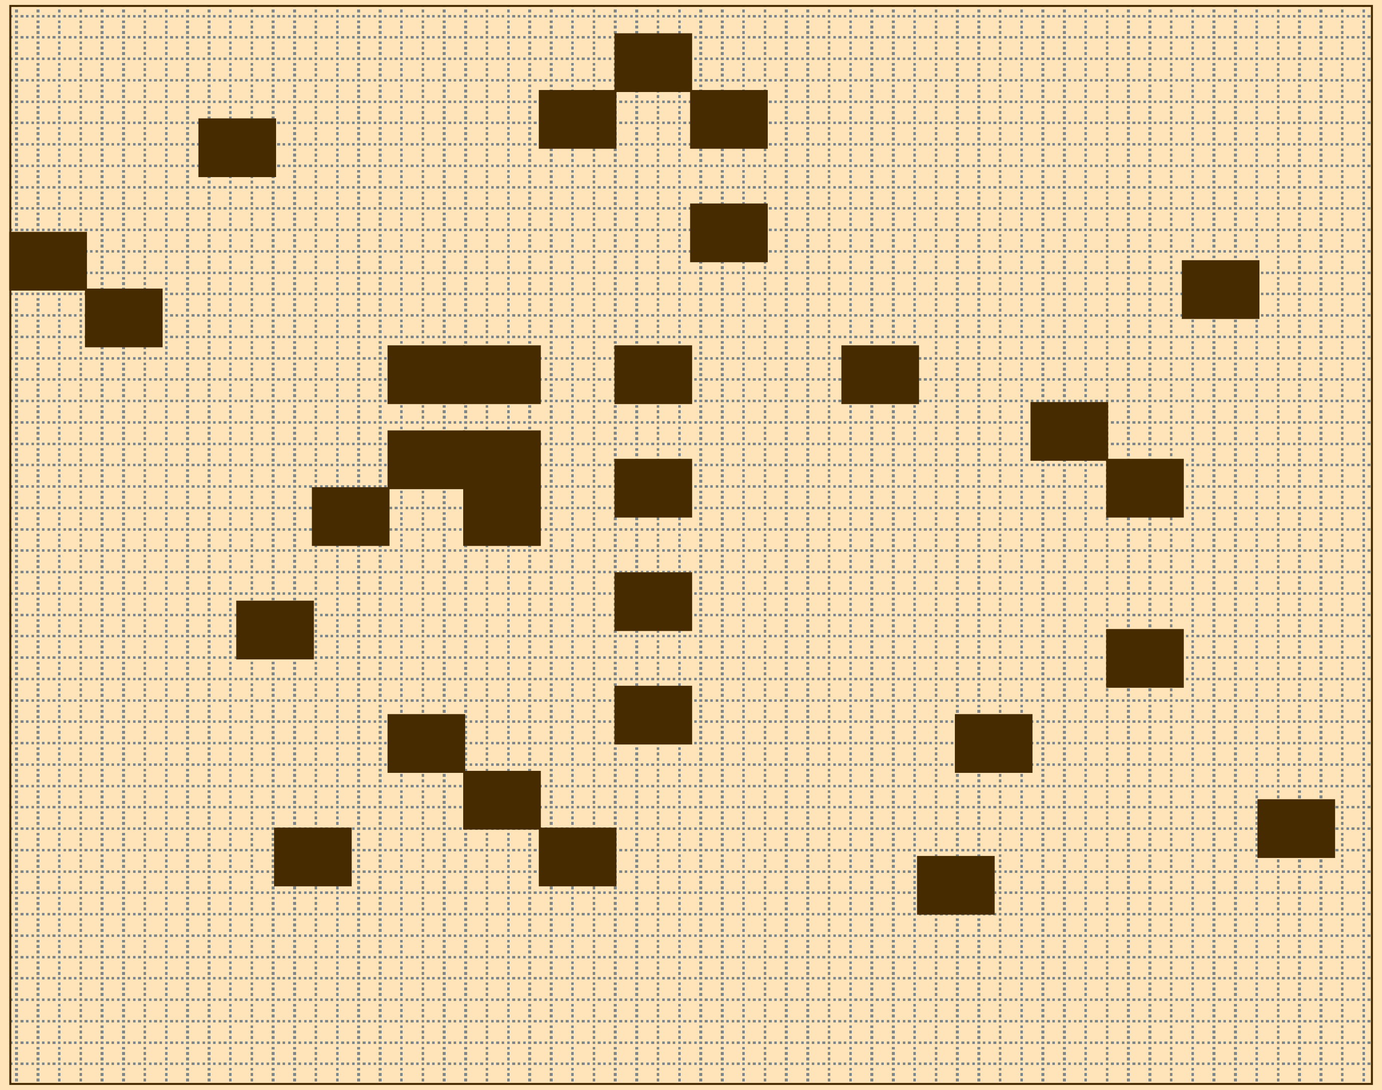

Prototype

$\vec{p}^{1}$

Next lecture at
11h15

$$w_{ij} = \frac{1}{N}\sum_{\mu} p_i^{\mu} p_j^{\mu}$$

$$S_i(t+1) = \mathrm{sgn}[\sum_j w_{ij} S_j(t)]$$

Sum over all interactions with i

Assume 4 patterns. At time t=0, overlap with Pattern 3, no overlap with other patterns.
discuss temporal evolution (of overlaps)
(assume that patterns are orthogonal)

Week 5-5: Learning of Associations

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

Biological Modeling of Neural Networks

Week 5
NETWORKS of NEURONS and ASSOCIATIVE MEMORY

Wulfram Gerstner
EPFL, Lausanne, Switzerland

5.5 Learning of Associations

Where do the connections come from?

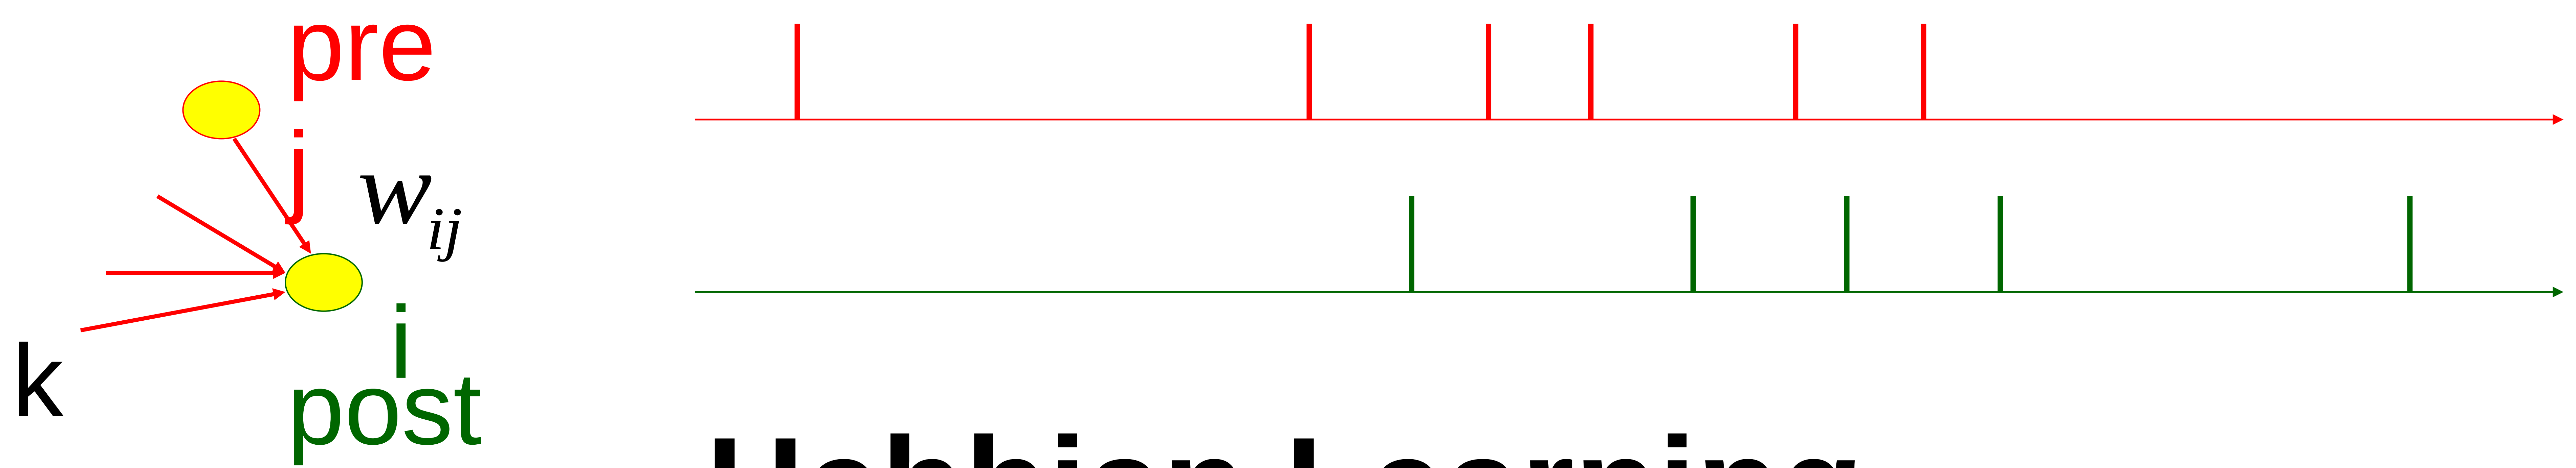

Hebbian Learning

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

- local rule
- simultaneously active (correlations)

5.5 Hebbian Learning of Associations

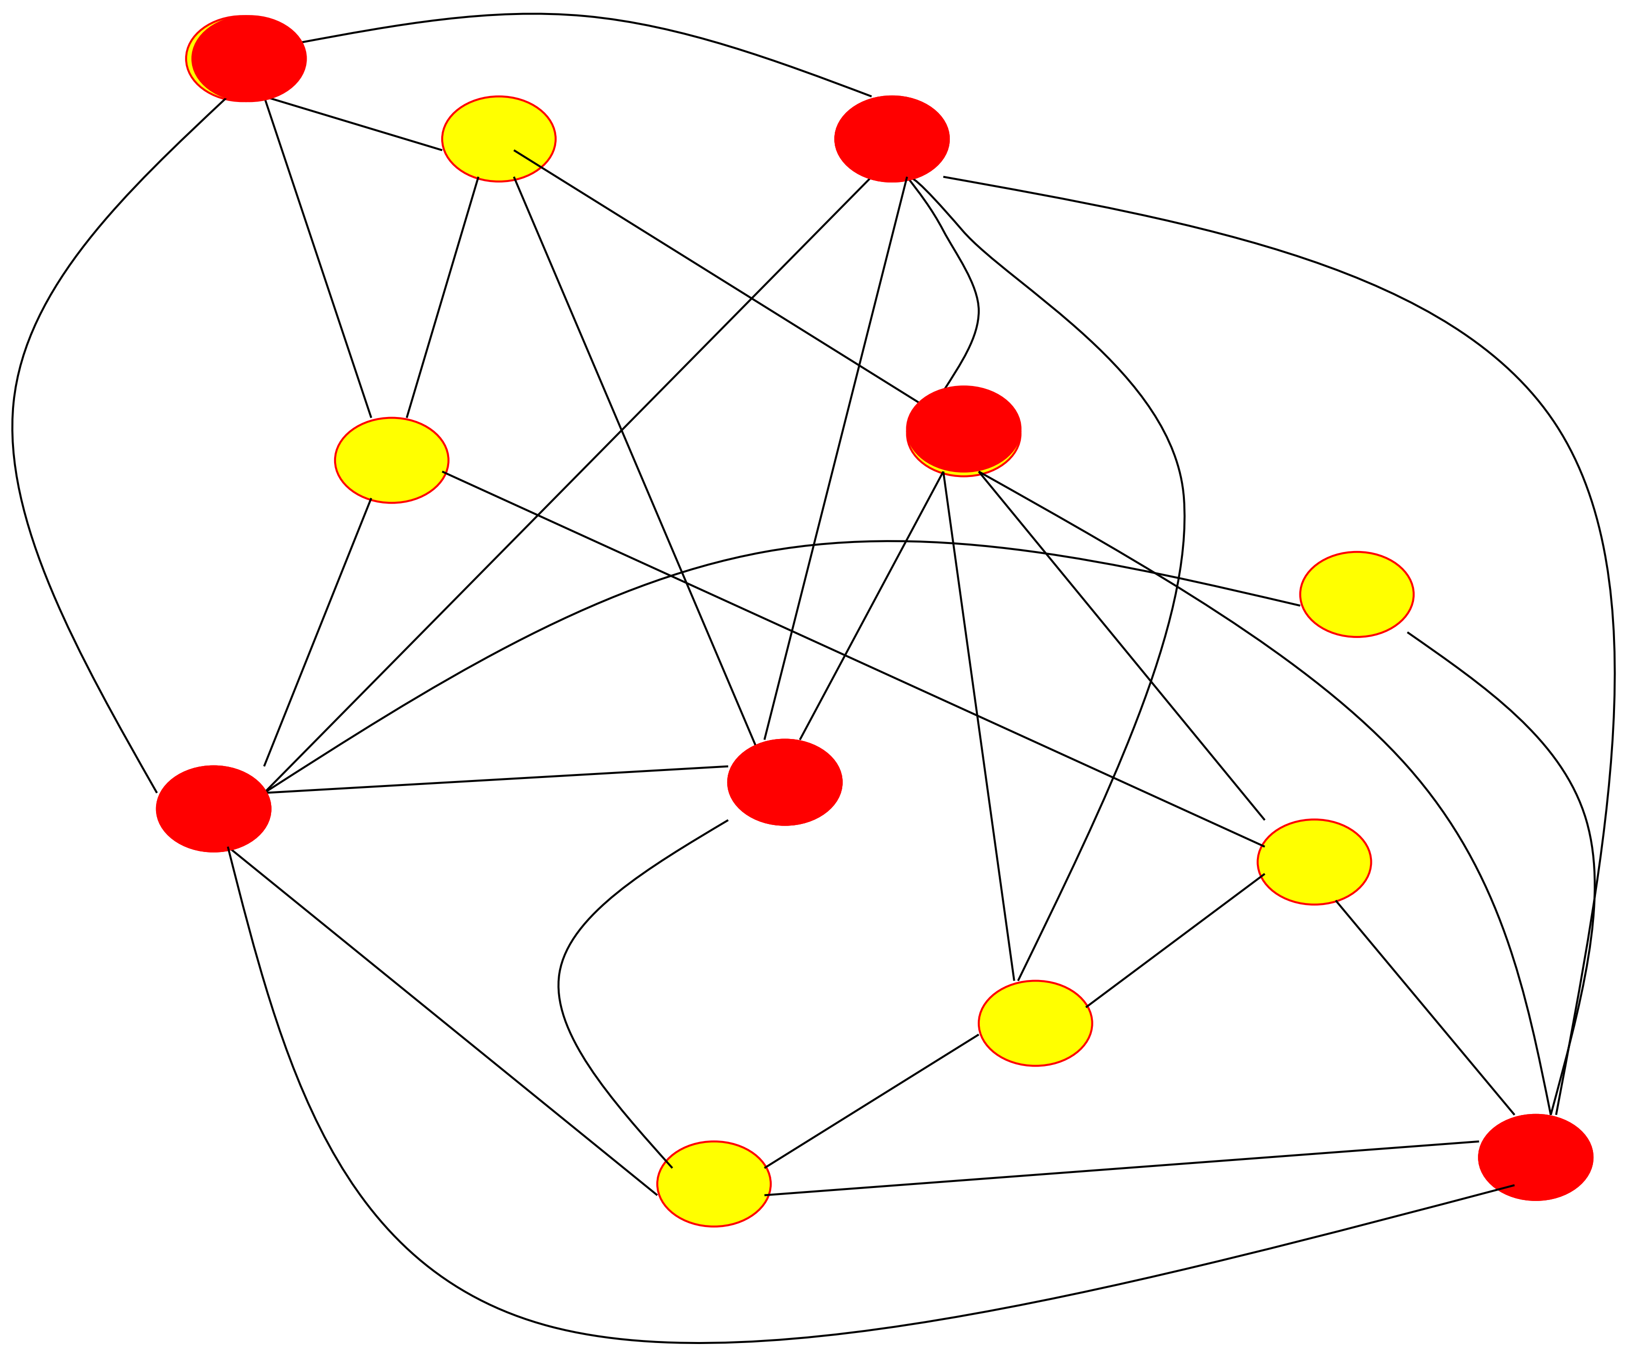

5.5 Hebbian Learning of Associations

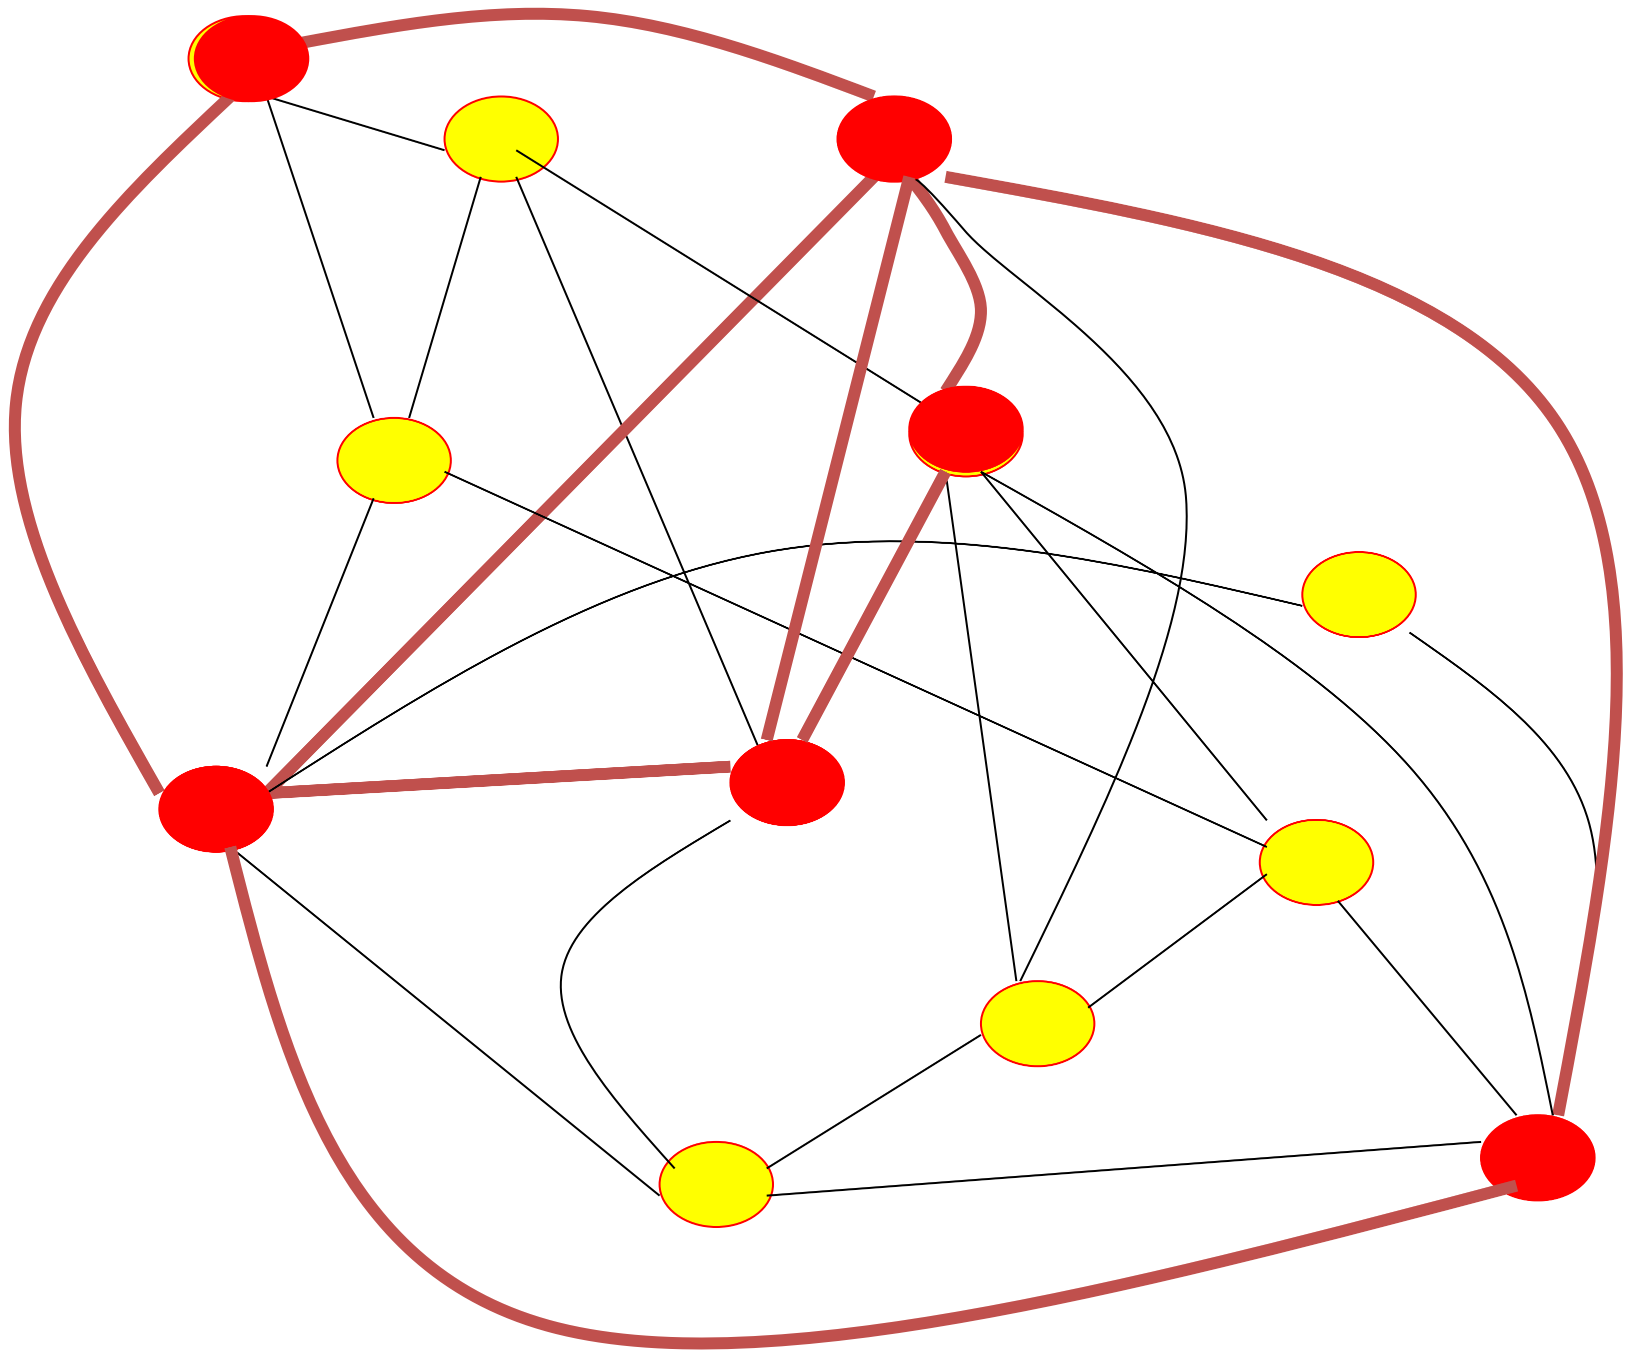

item memorized

5.5 Hebbian Learning: Associative Recall

Recall:
Partial info

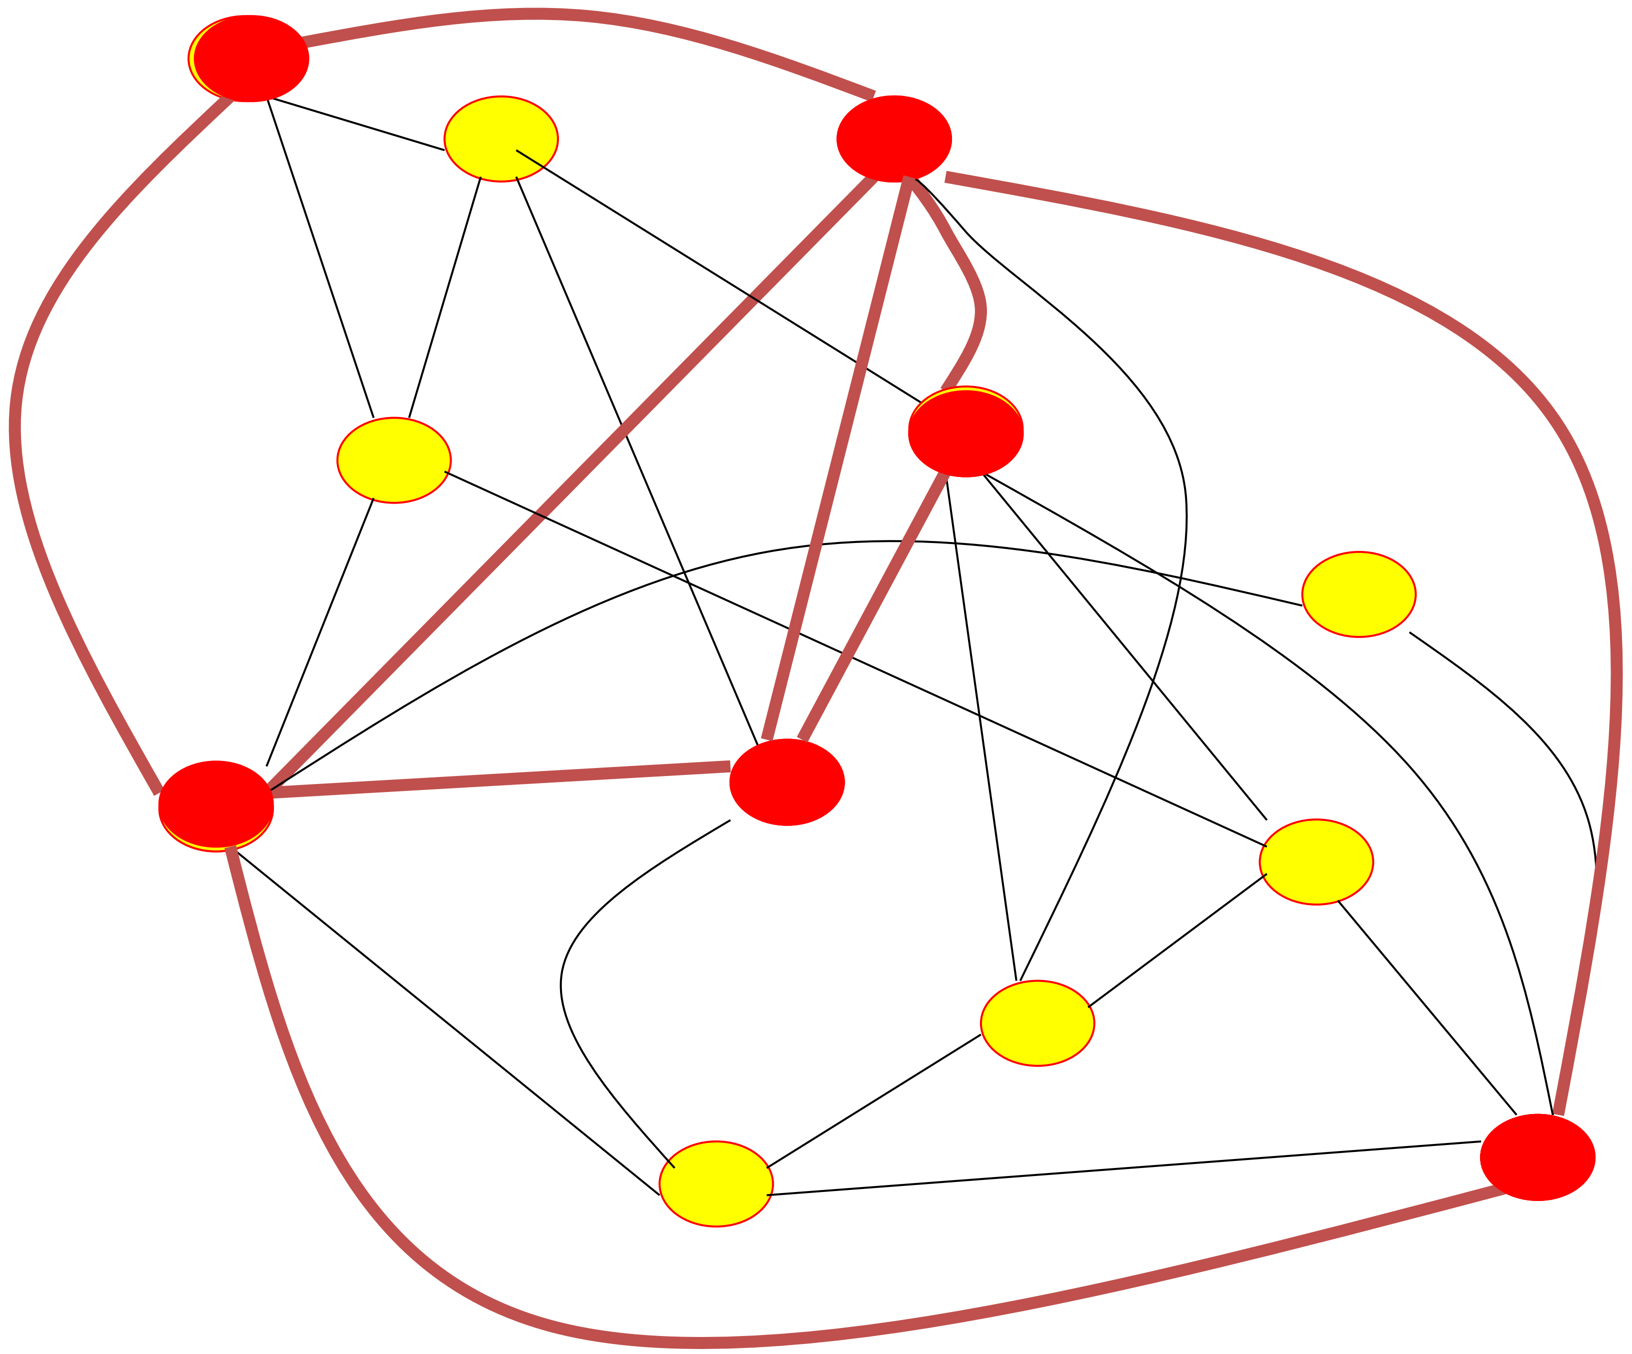

item recalled

5.5 Associative Recall

Tell me the object **shape**

for the following list of 5 items:

Tell me the **color**

for the following list of 5 items:

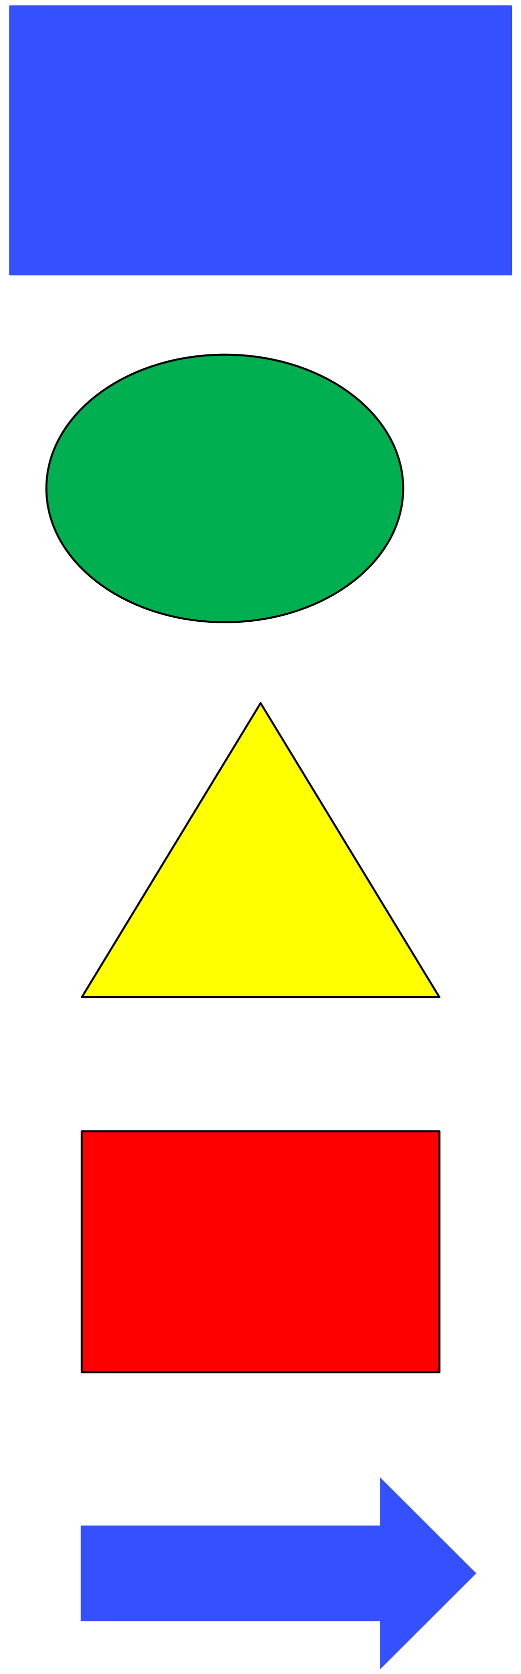

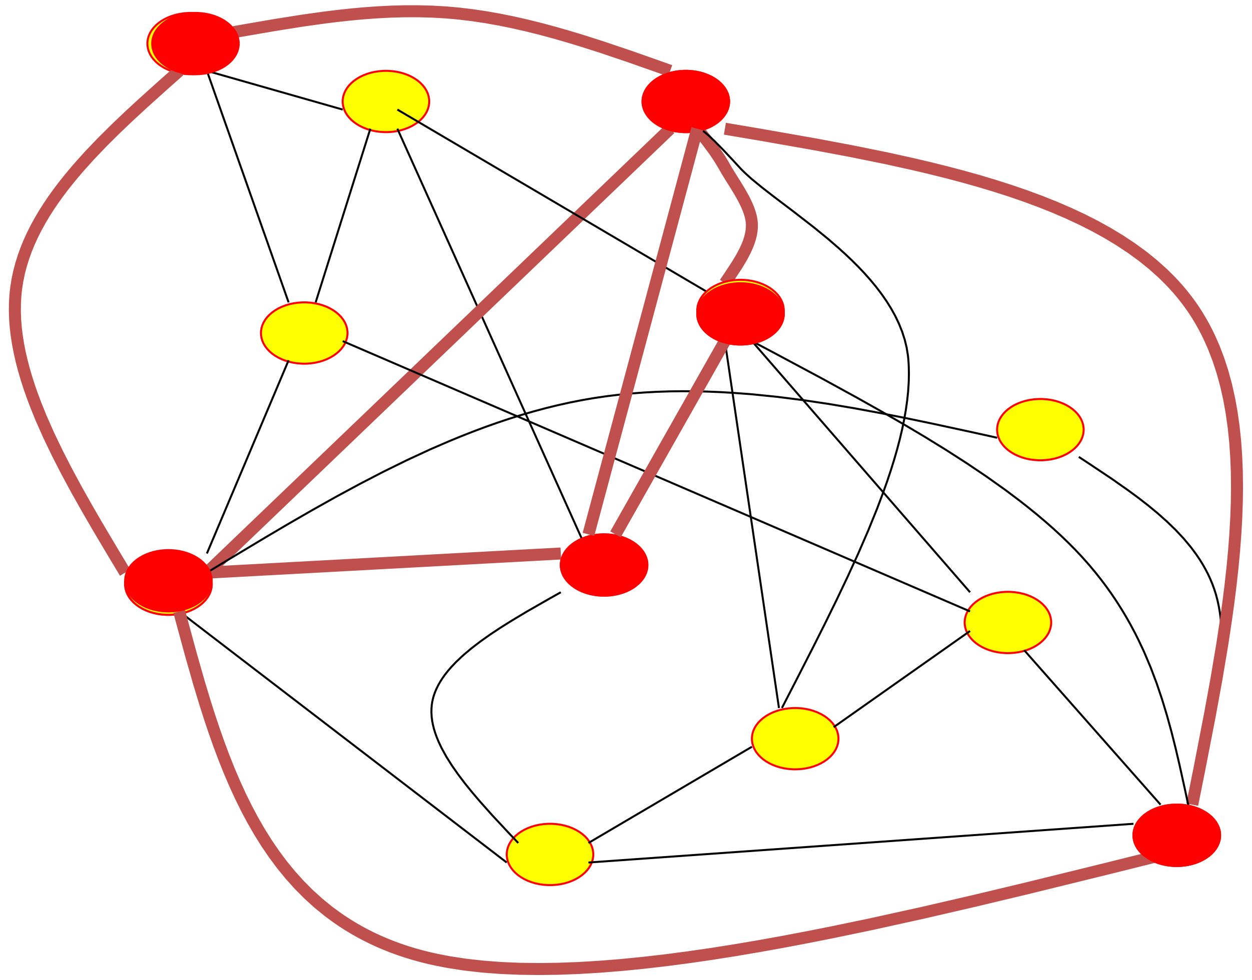

be as fast as possible:

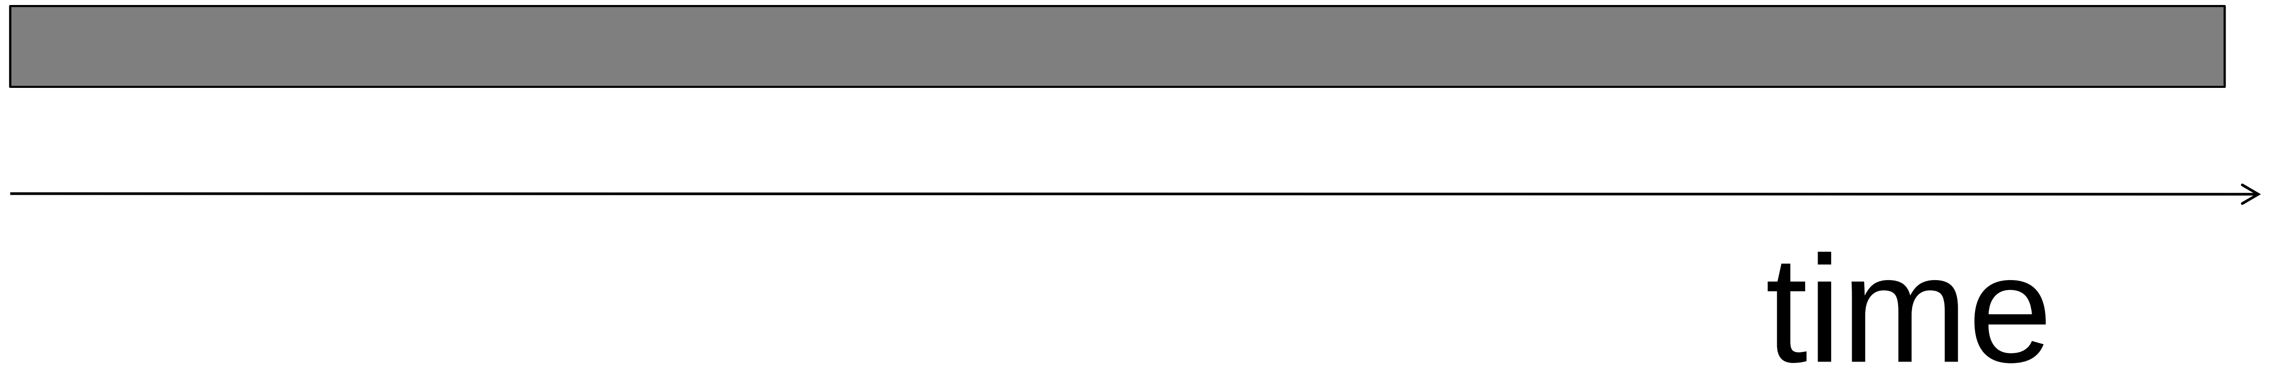

5.5 Associative Recall

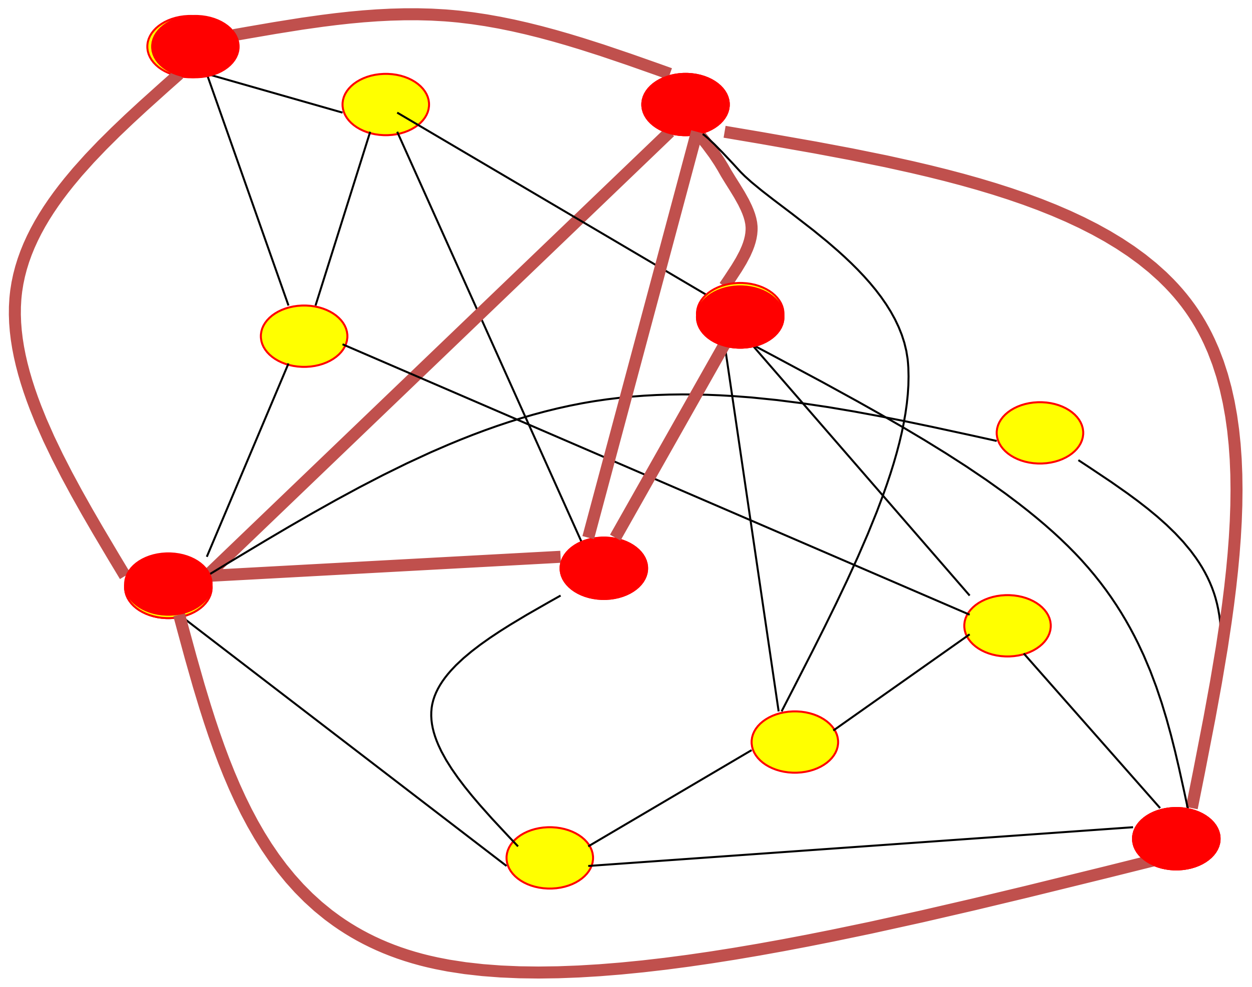

Tell me the **color**

for the following list of 5 items:

be as fast as possible:

time

Stroop effect:

Slow response: hard to work

Against natural associations

5.5 Associative Recall

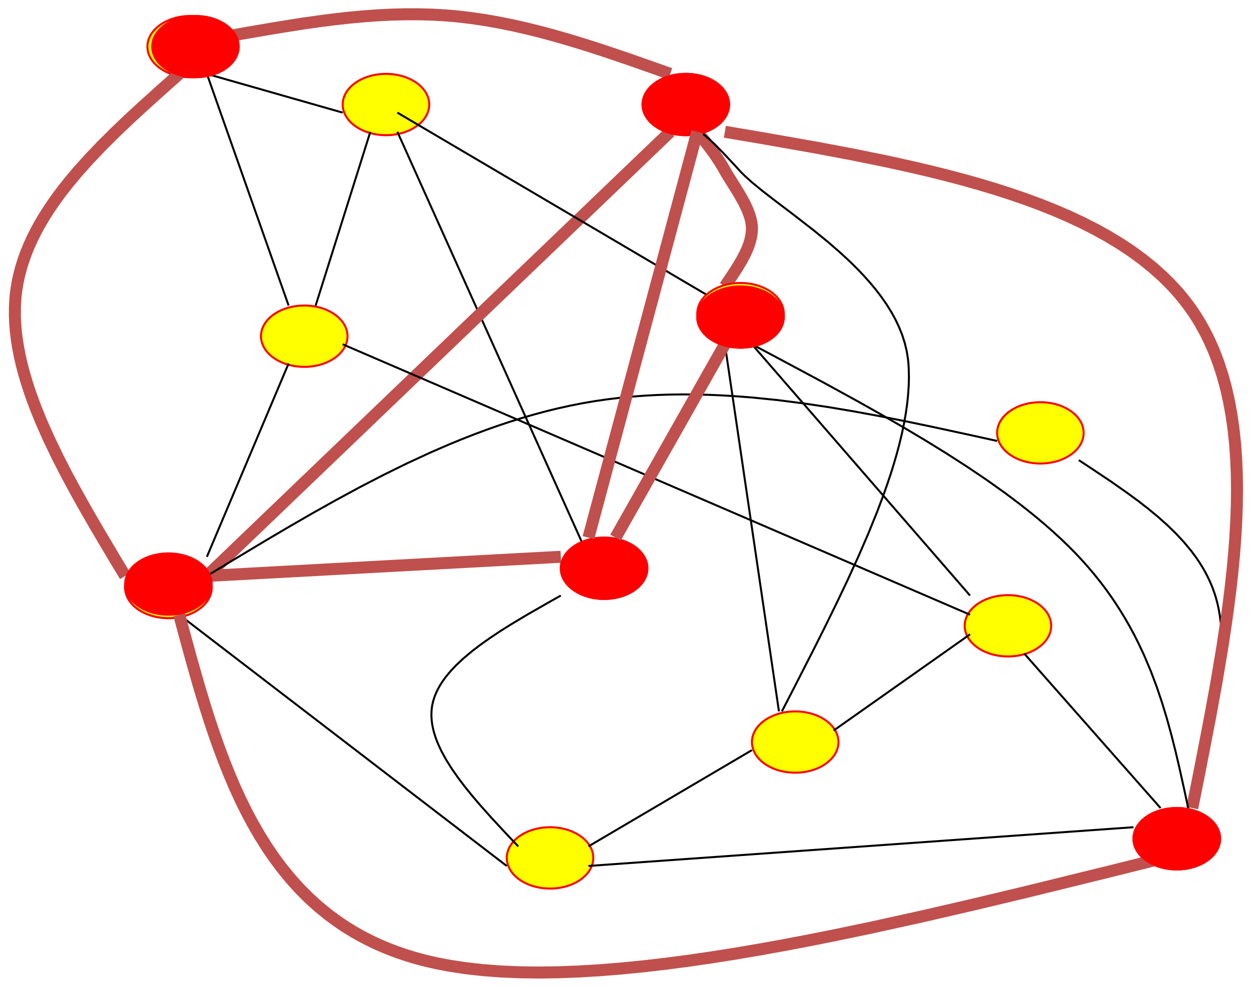

Hierarchical organization of Associative memory

animals

birds fish

Name as fast as possible an example of a bird

swan (or goose or raven or …)

Write down first letter: *s* for *swan* or *r* for *raven* …

5.5 Associative Recall

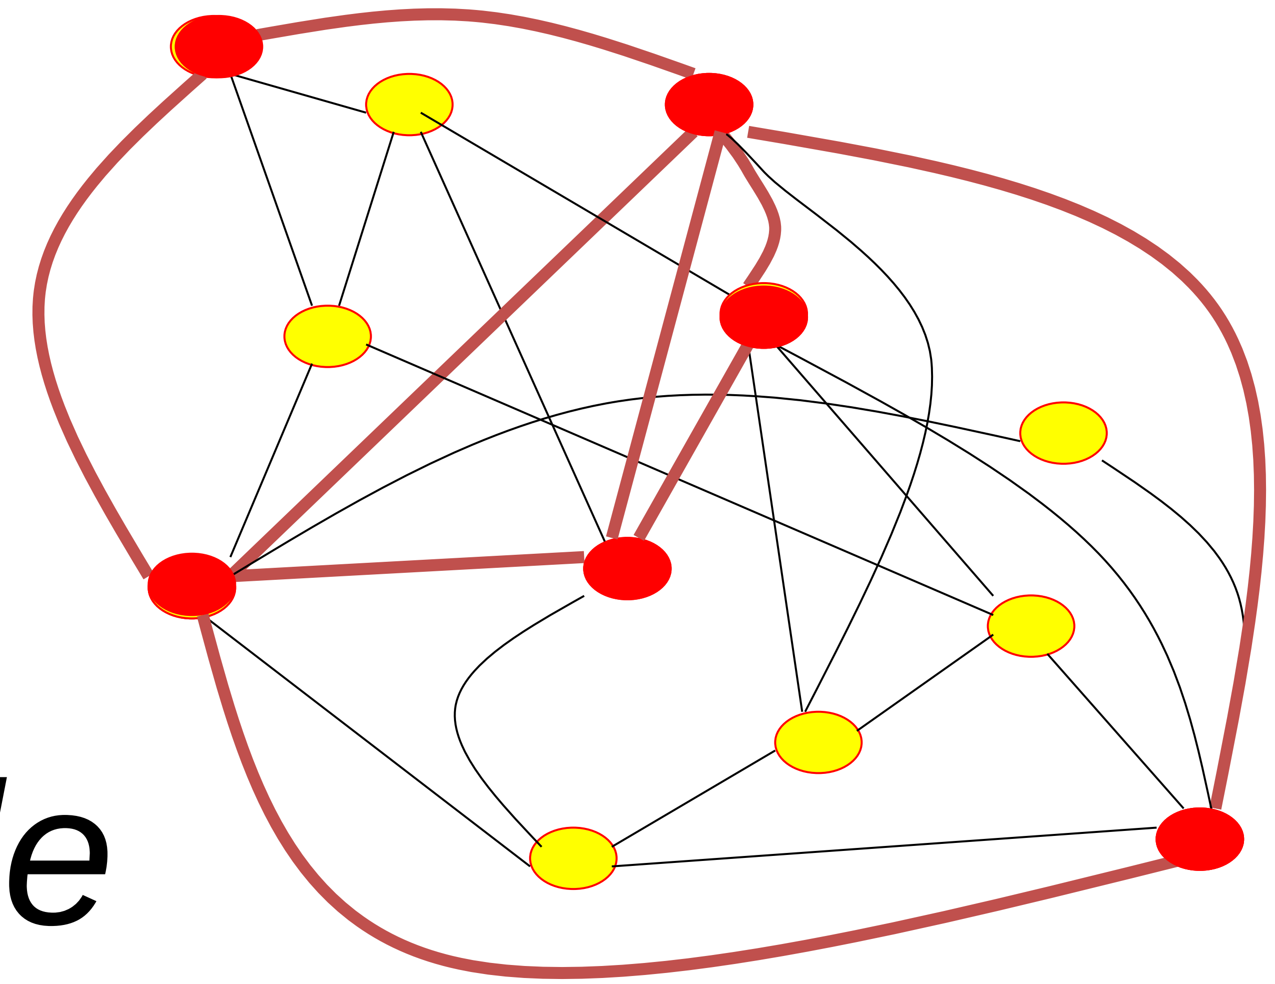

Nommez au plus vite possible un exemple d'un /d'une

name as fast as possible an example of a

outil	**tool**
couleur	**color**
fruit	**fruit**
instrument de musique	**music instrument**

Week 5-5: Learning of Associations

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

Biological Modeling of Neural Networks

Week 5
NETWORKS of NEURONS and ASSOCIATIVE MEMORY

Wulfram Gerstner

EPFL, Lausanne, Switzerland

learning of prototypes

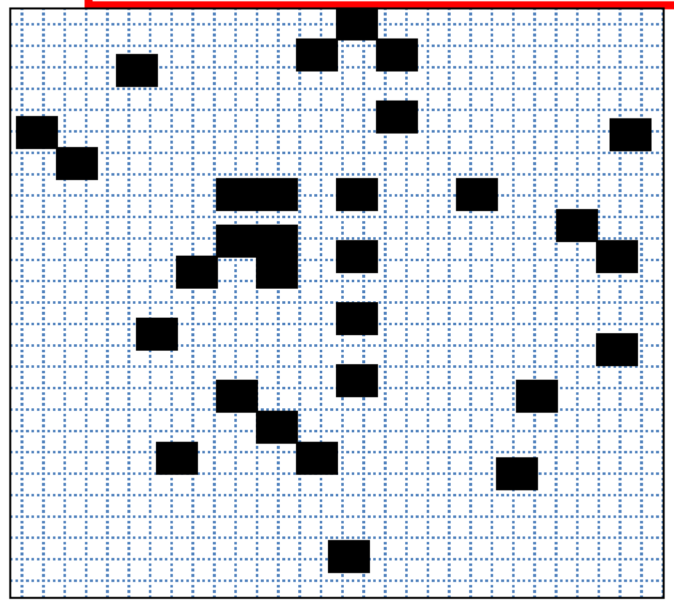

Prototype

$\vec{p}^1$

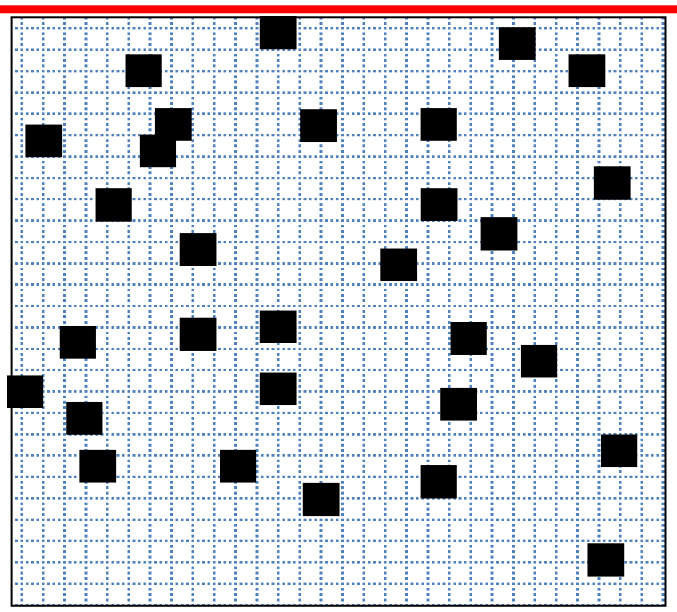

Prototype

$\vec{p}^2$

interactions

$$(1) \quad w_{ij} = \frac{1}{N} \sum_{\mu} p_i^{\mu} p_j^{\mu}$$

Sum over all prototypes

Q; How many prototypes can be stored?

dynamics

$$S_i(t+1) = \mathrm{sgn}\Big[\sum_j w_{ij} S_j(t)\Big]$$

all interactions with i

Q; How many prototypes can be stored?

blackboard

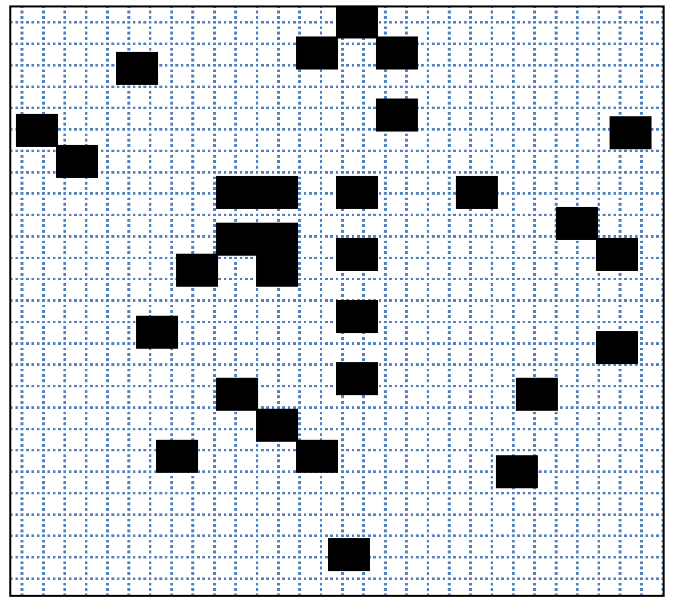

Prototype

$\vec{p}^1$

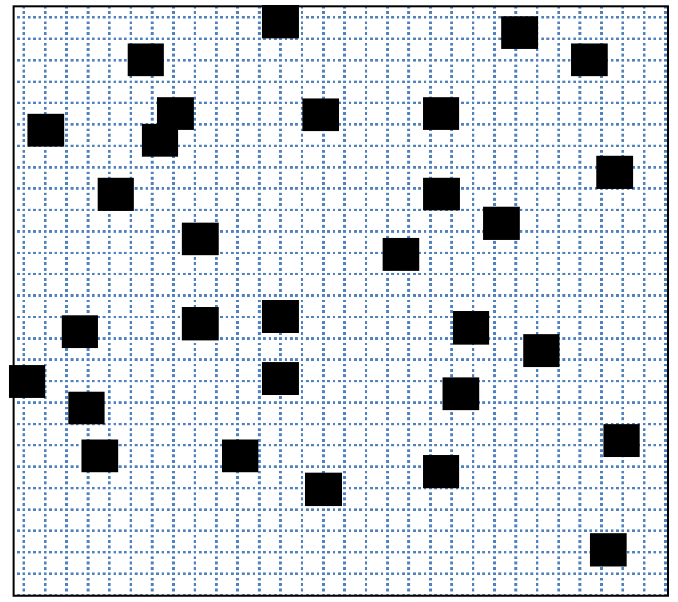

Prototype

$\vec{p}^2$

Random patterns

Interactions (1) $w_{ij} = \sum_{\mu} p_i^{\mu} p_j^{\mu}$

Dynamics (2) $S_i(t+1) = \mathrm{sgn}[\sum_j w_{ij} S_j(t)]$

Minimal condition: pattern is fixed point of dynamics

-Assume we start directly in one pattern

-Pattern stays

Attention: Retrieval requires more (pattern completion)

Exercise 4 now: Associative memory

Q; How many prototypes can be stored?

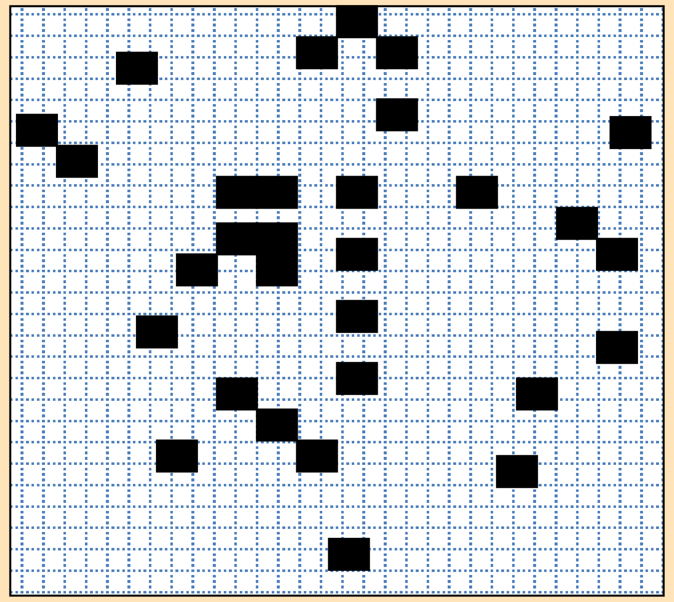

Prototype

$\vec{p}^1$

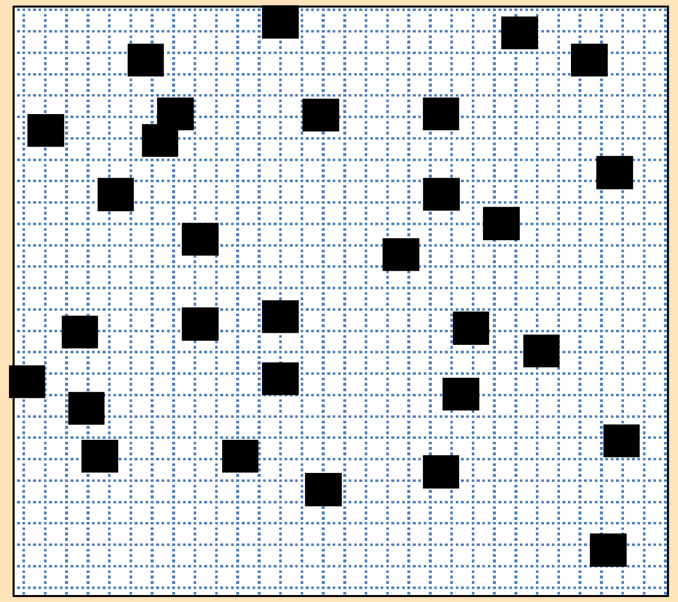

Prototype

$\vec{p}^2$

End of lecture, exercise+
Computer exercise : 12:00

Interactions (1) $w_{ij} = \sum_{\mu} p_i^{\mu} p_j^{\mu}$

Dynamics (2) $S_i(t+1) = \mathrm{sgn}[\sum_j w_{ij} S_j(t)]$

Random patterns → random walk

a) show relation to erf function: importance of p/N

b) network of 1000 neurons – allow at most 1 wrong pixel?

c) network of N neurons – at most 1 promille wrong pixels?

Week 6 Review: storage capacity of Hopfield model

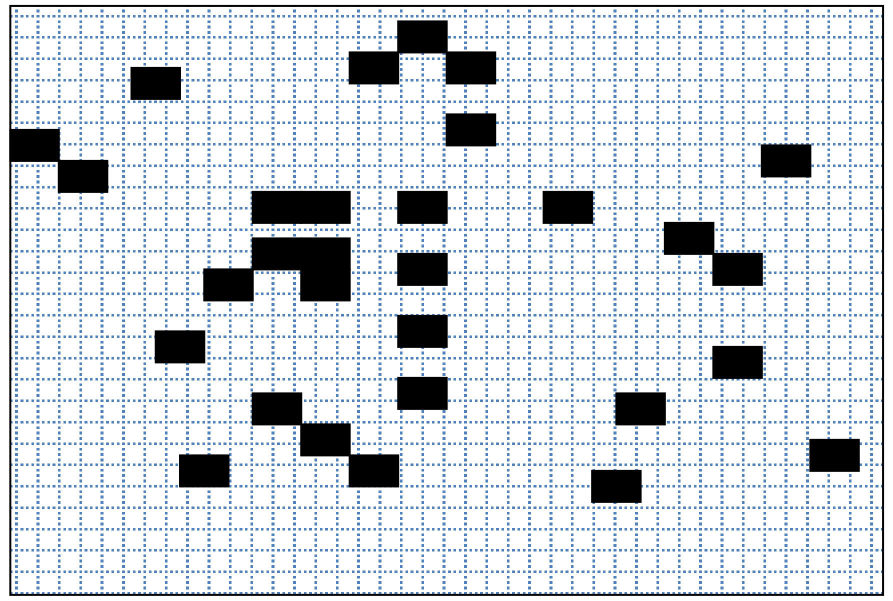

Prototype $\vec{p}^1$

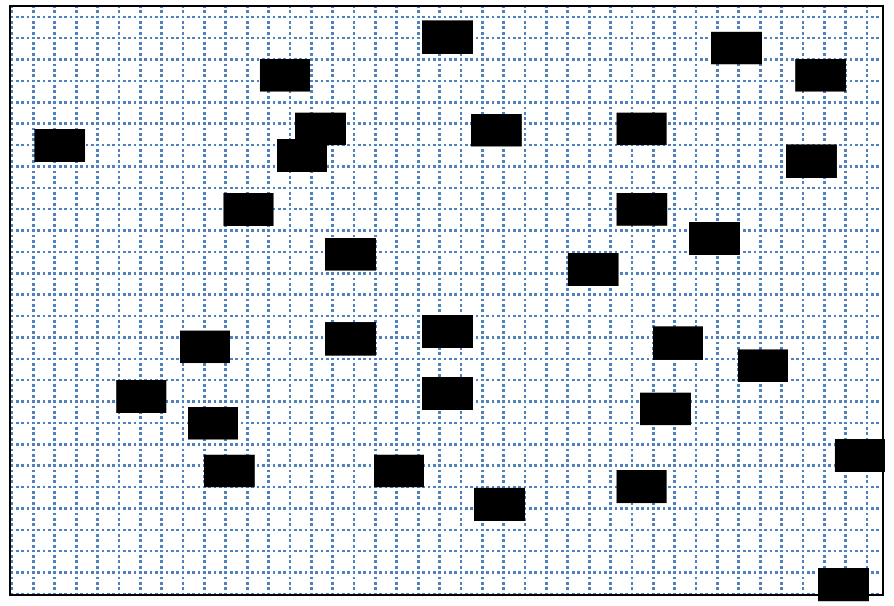

Prototype $\vec{p}^2$

Random patterns

Interactions (1) $w_{ij} = \sum_\mu p_i^\mu p_j^\mu$

Dynamics (2) $S_i(t+1) = \mathrm{sgn}[\sum_j w_{ij} S_j(t)]$

Minimal condition: pattern is fixed point of dynamics

-Assume we start directly in one pattern

-Pattern stays

Attention: Retrieval requires more (pattern completion)

Q; How many prototypes can be stored?

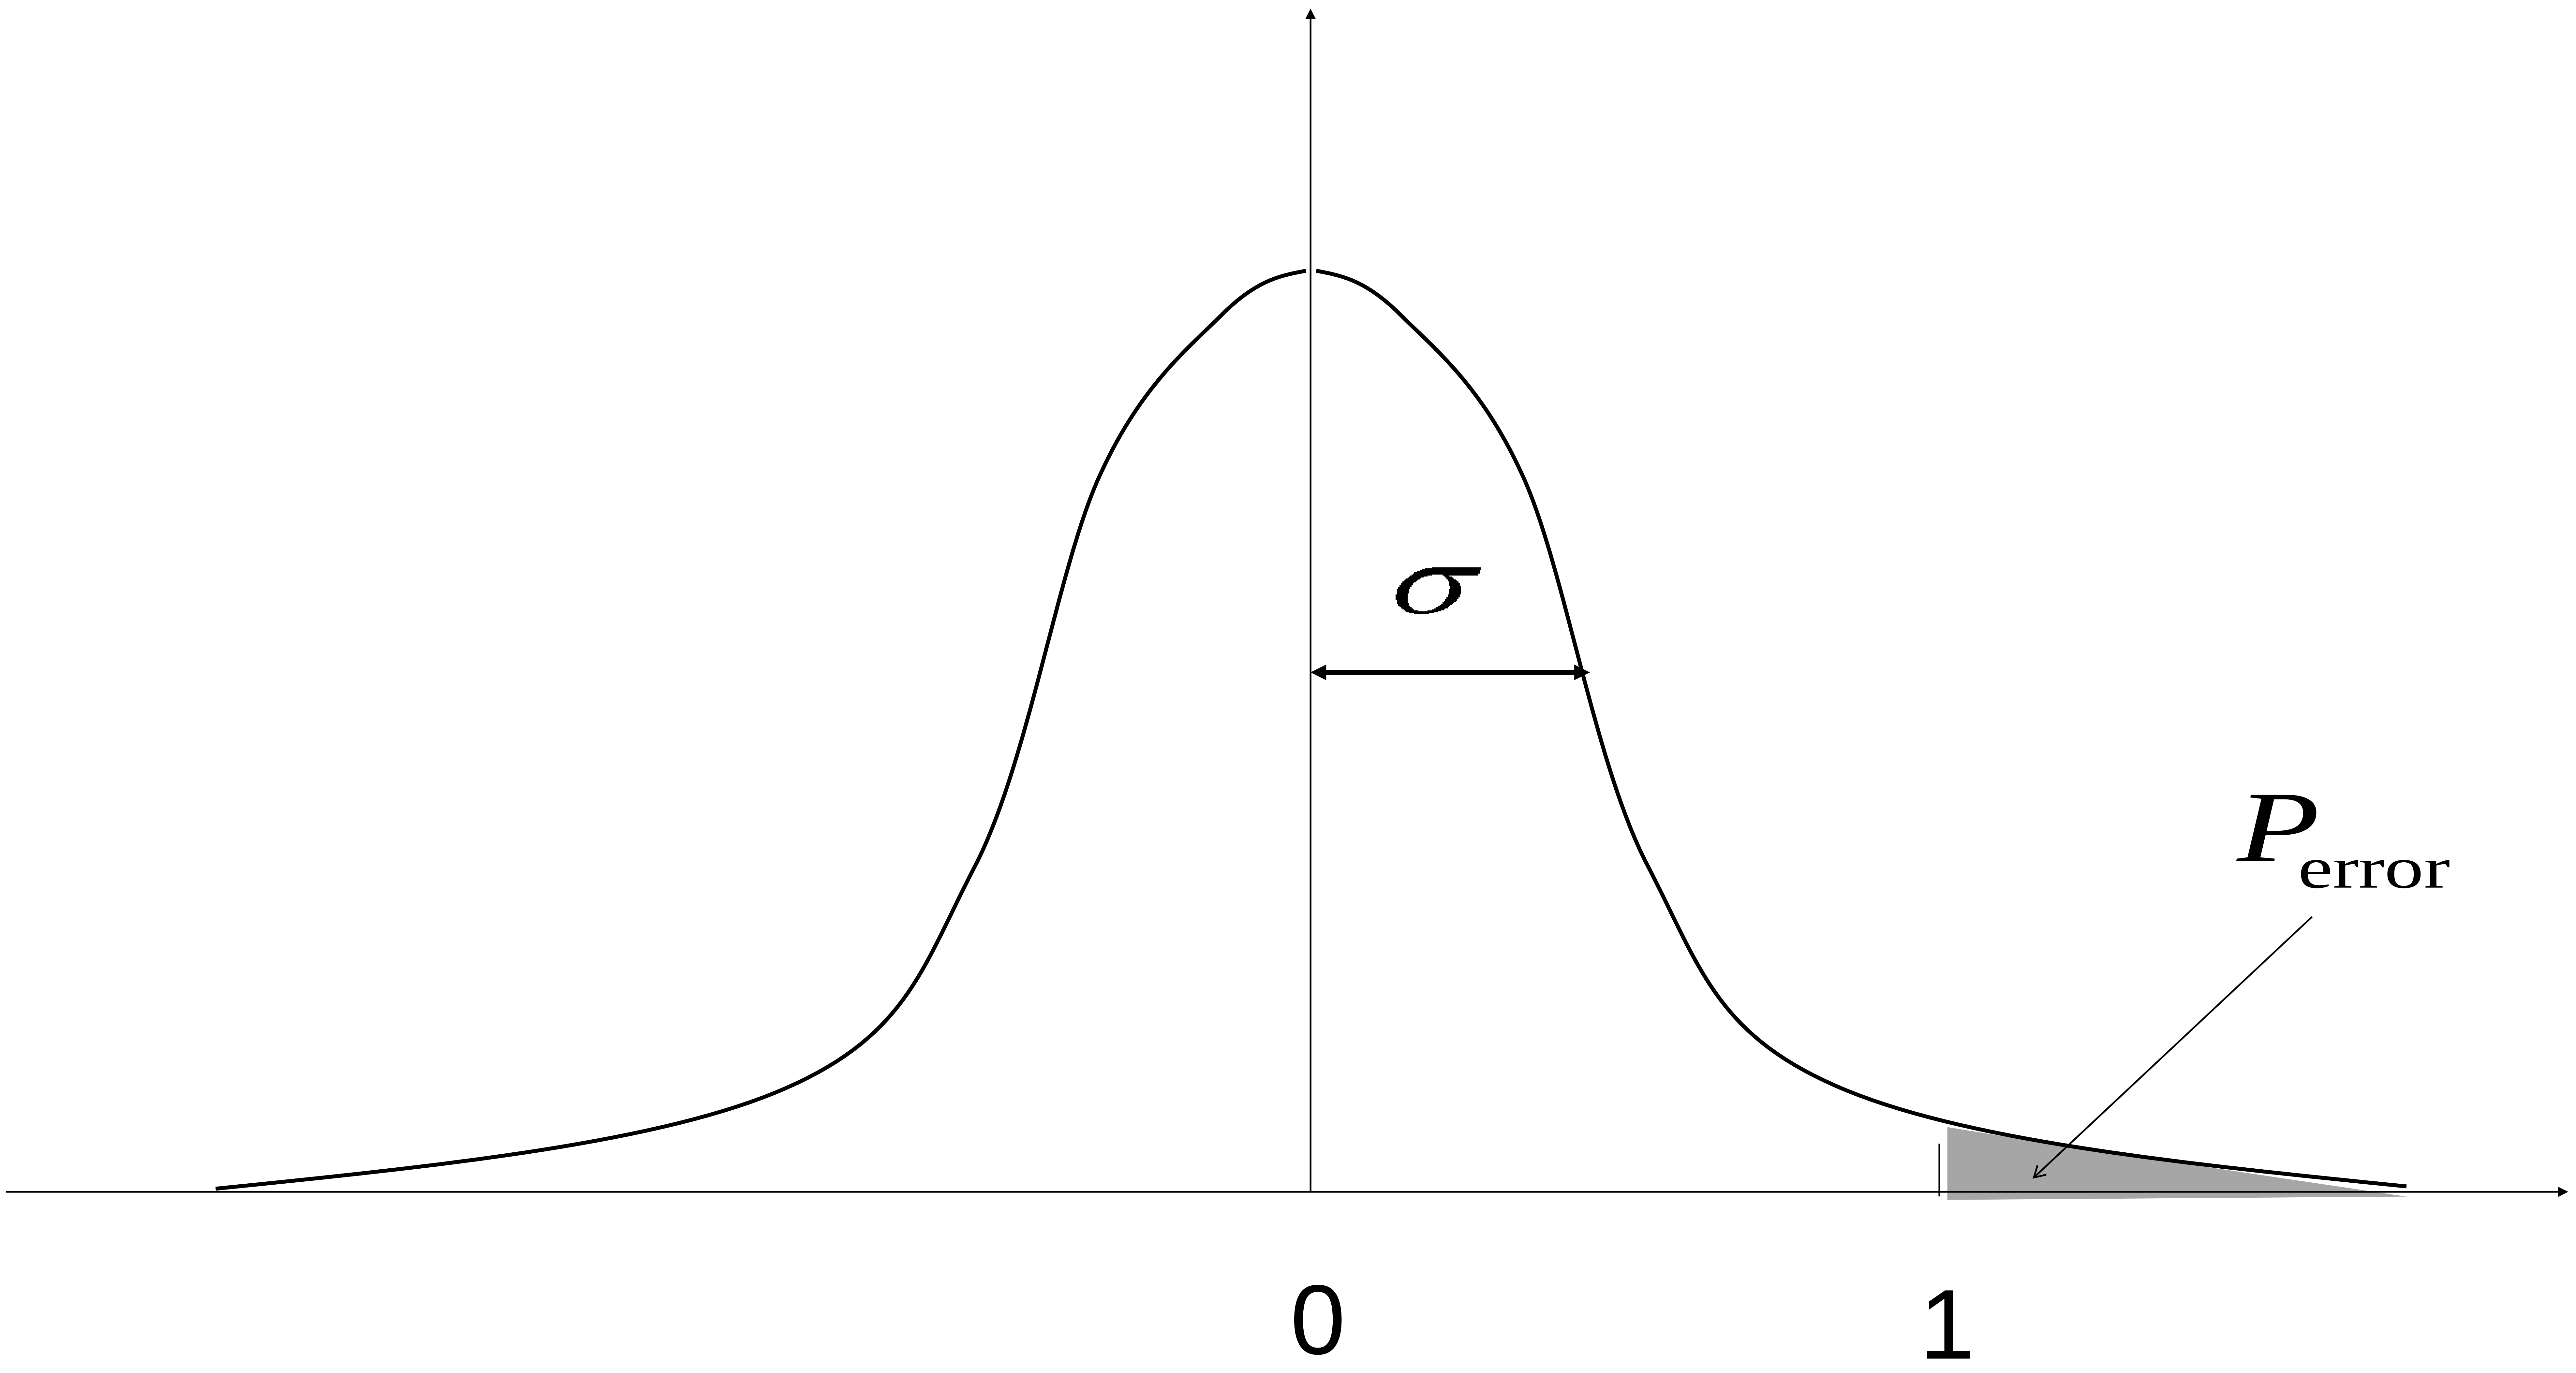

Week 6: Hopfield model continued

ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

Biological Modeling of Neural Networks

Week 6

Hebbian LEARNING and ASSOCIATIVE MEMORY

Wulfram Gerstner

EPFL, Lausanne, Switzerland

6.1 Stochastic Hopfield Model
6.2. Energy landscape
6.3. Low-activity patterns
6.4. Attractor memorie
- spiking neurons
- experimental data

6.1 Review: Hopfield model

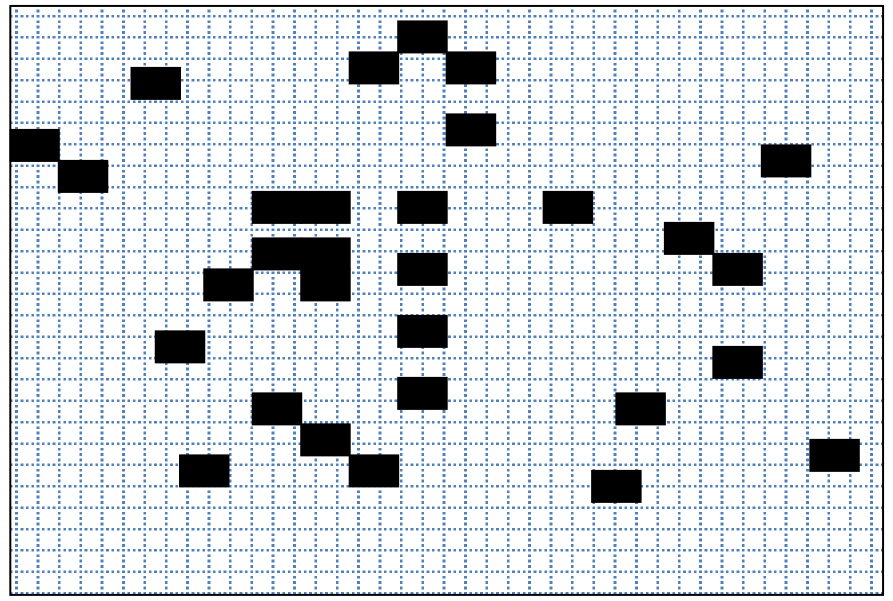

Prototype $\vec{p}^{\,1}$

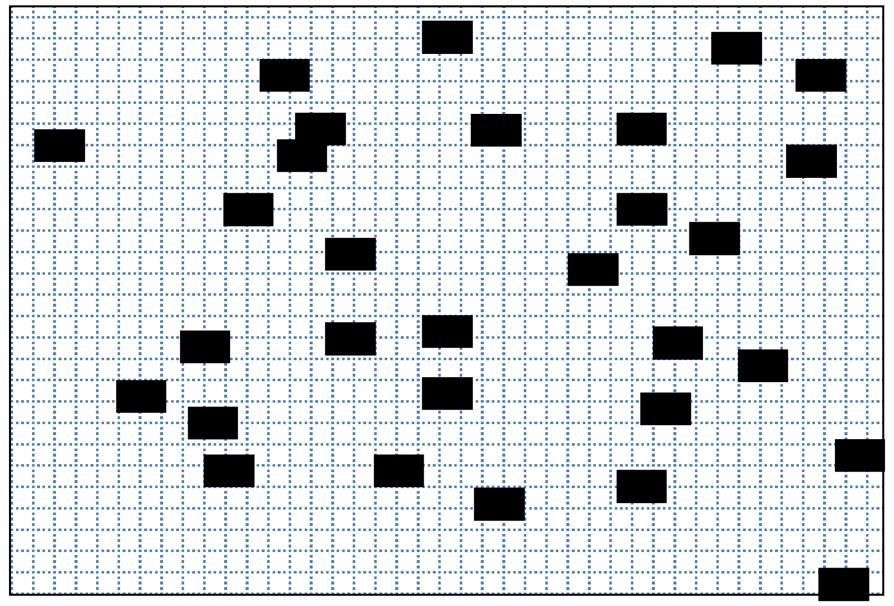

Prototype $\vec{p}^{\,2}$

interactions

$$(1) \qquad w_{ij} = \tfrac{1}{N}\sum_{\mu} p_i^{\mu} p_j^{\mu}$$

Sum over all prototypes

$\boxed{\text{Deterministic dynamics}}$

$$S_i(t+1) = \mathrm{sgn}[\sum_j w_{ij} S_j(t)]$$

6.1 Stochastic Hopfield model

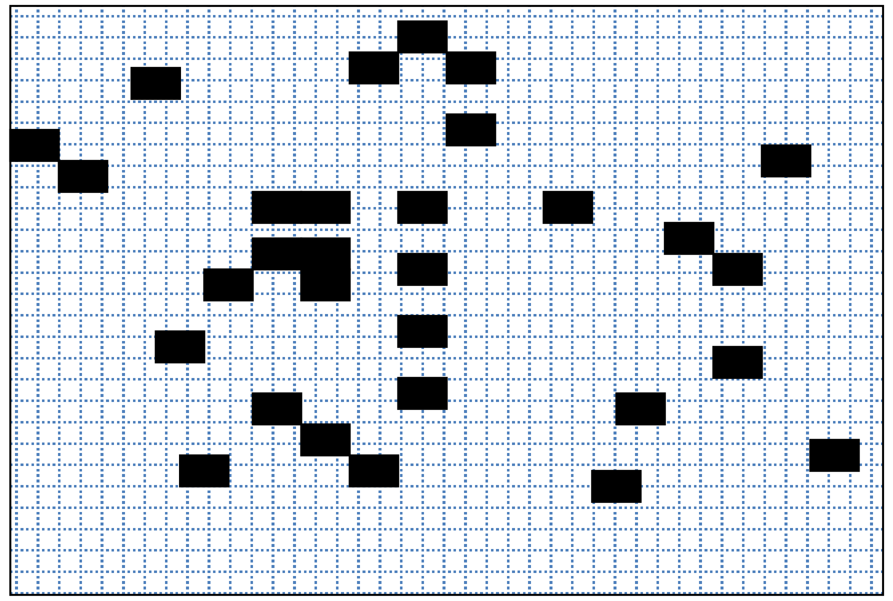

Prototype $\vec{p}^1$

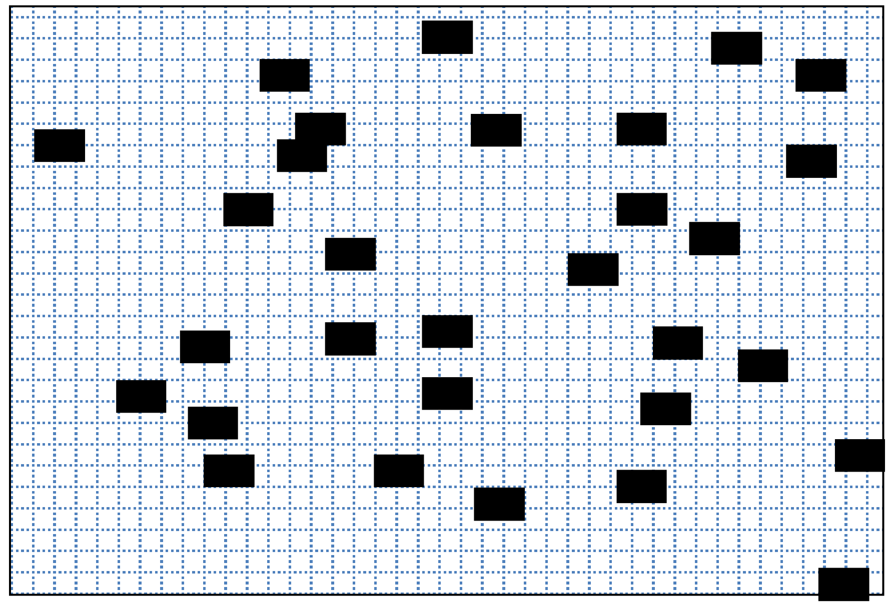

Prototype $\vec{p}^2$

Random patterns

Interactions (1) $w_{ij} = \sum_{\mu} p_i^{\mu} p_j^{\mu}$

Dynamics (2)

$$\Pr\{S_i\,(t+1) = +1 \,|\, h_i\} = g[h_i] = g\left[\sum_j w_{ij} S_j \quad t\right]$$

$$\Pr\{S_i\,(t+1) = +1 \,|\, h_i\} = g\left[\sum_{\mu} p_i^{\mu} m^{\mu} \quad t\right]$$

blackboard

6.1 Stochastic Hopfield model: memory retrieval

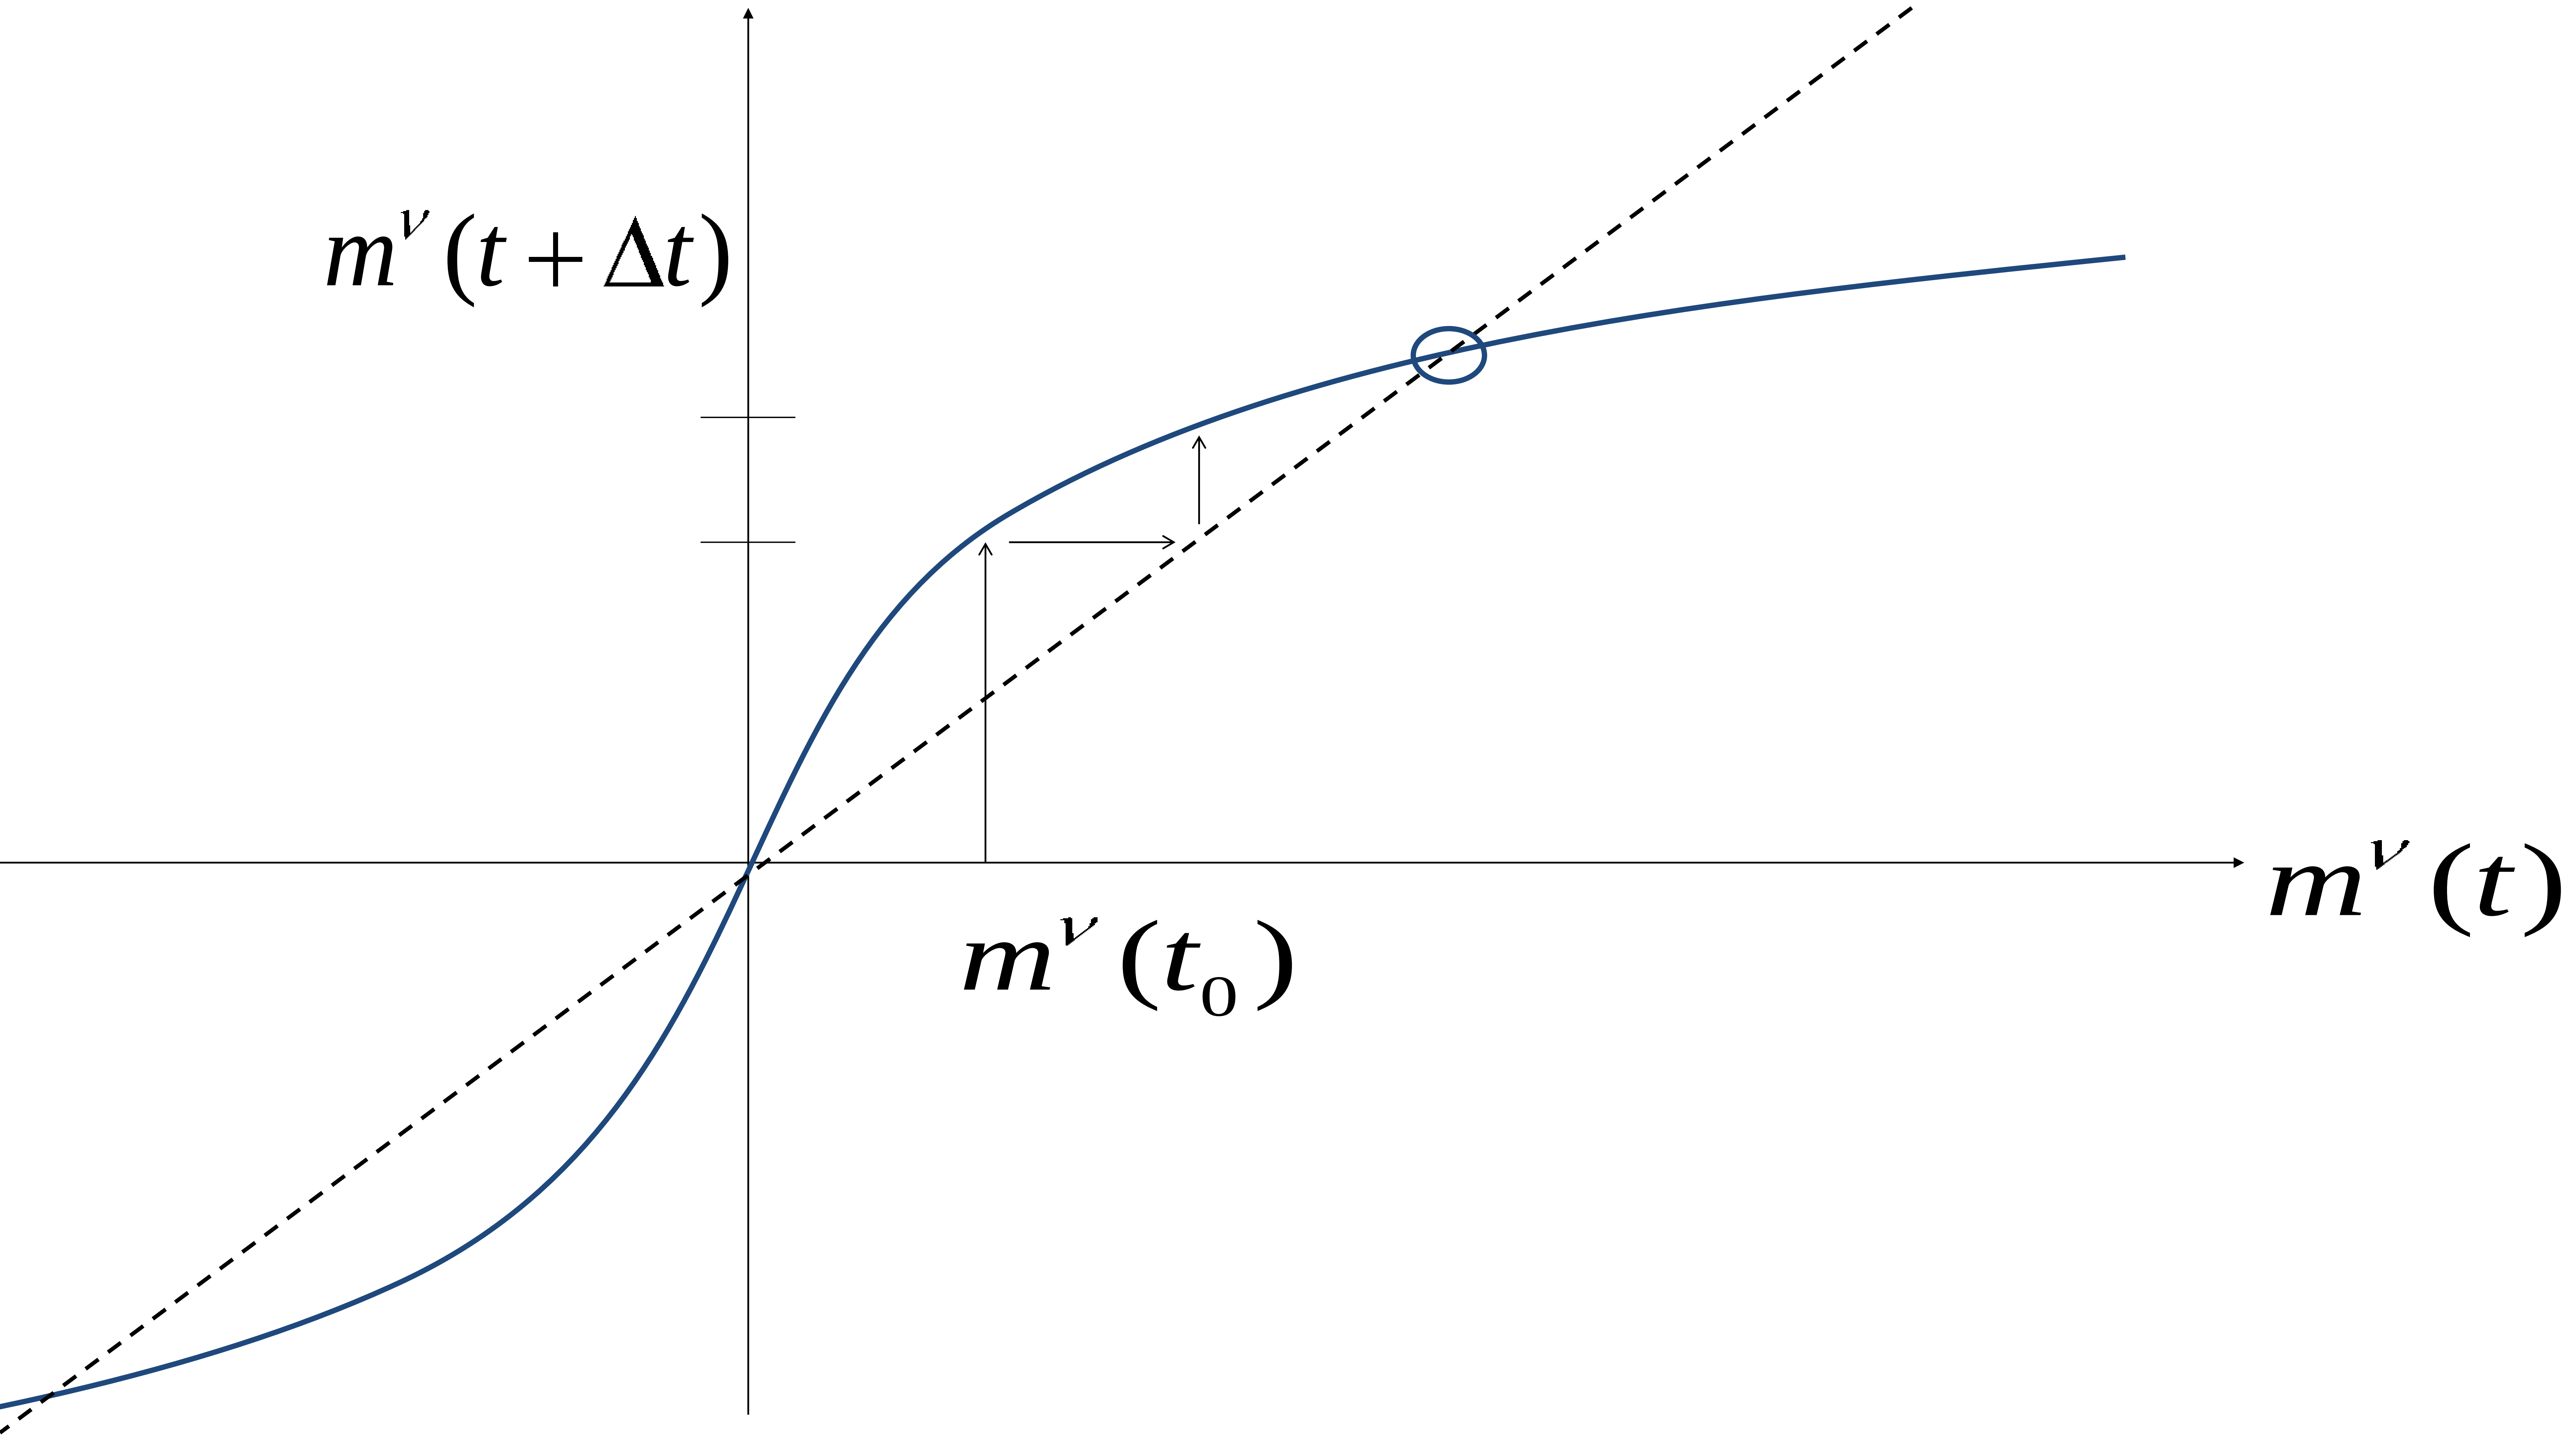

6.1 Stochastic Hopfield model: memory retrieval

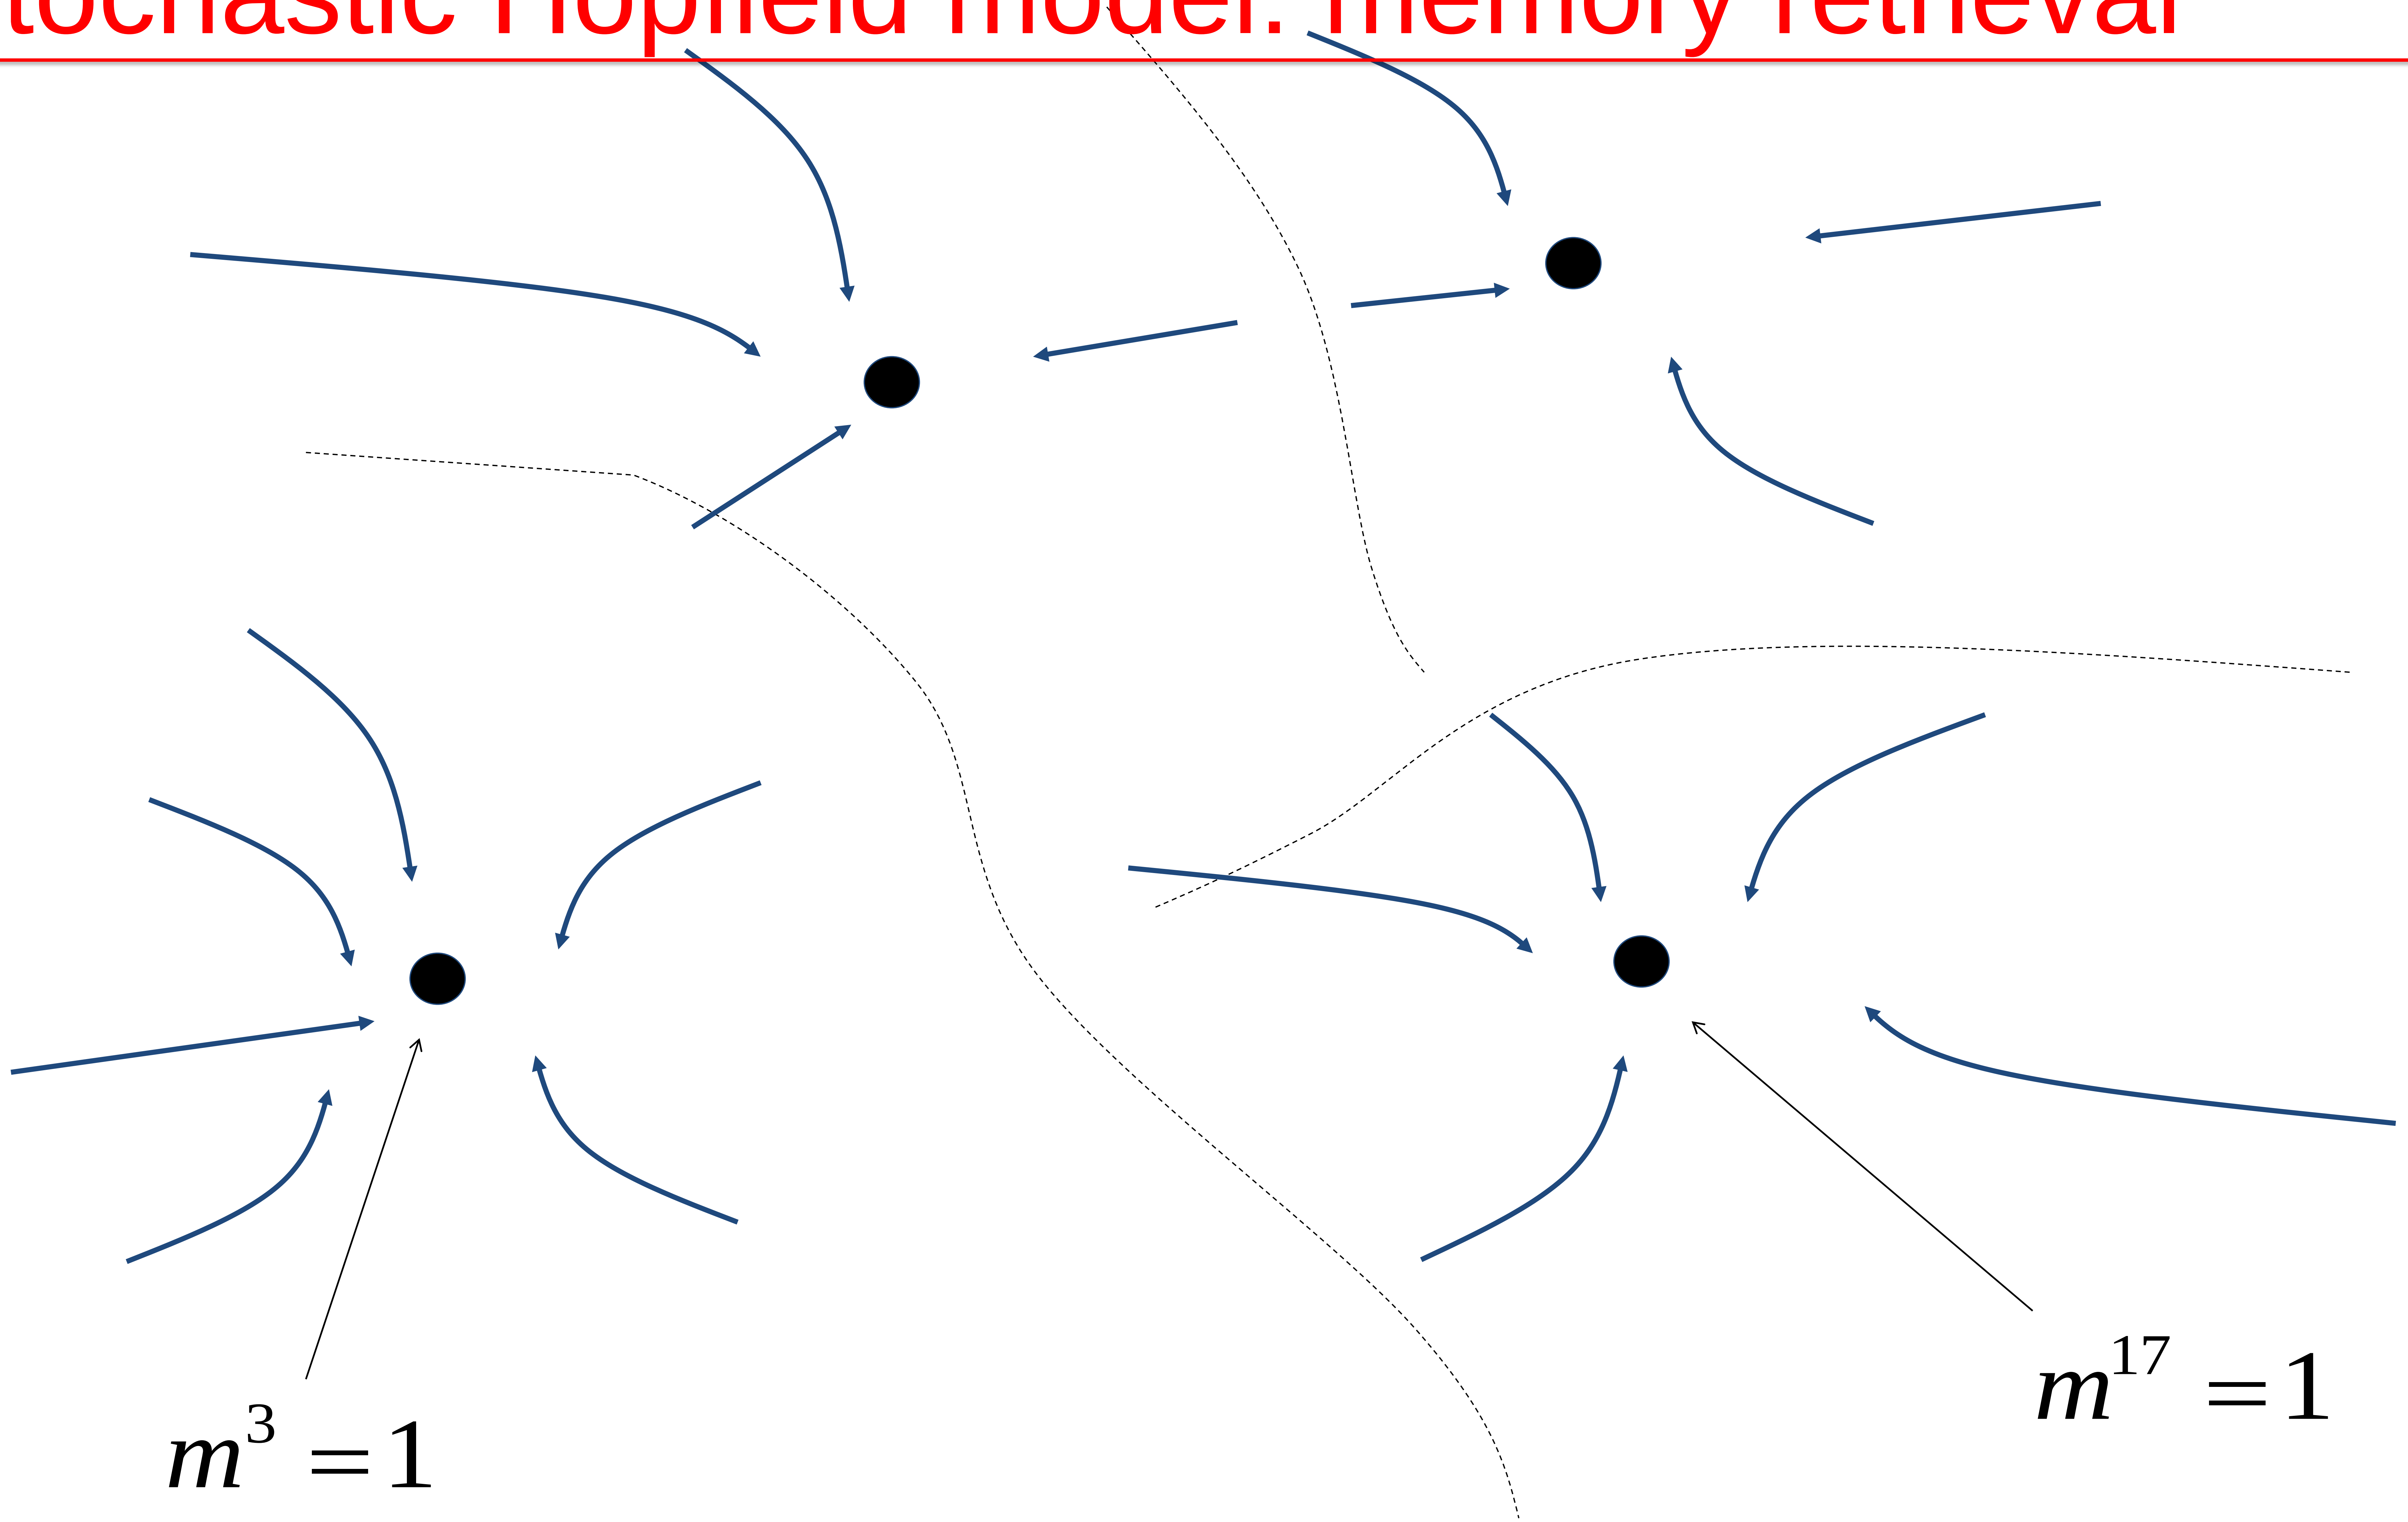

6.1 Stochastic Hopfield model

Dynamics (2)

$$\Pr\{S_i\,(t+1) = +1 \,|\, h_i\} = g[h_i] = g\left[\sum\nolimits_j w_{ij} S_j \;\; t \;\right]$$

$$\Pr\{S_i\,(t+1) = +1 \,|\, h_i\} = g\left[\sum\nolimits_\mu p_i^\mu m^\mu \;\; t \;\right]$$

blackboard

Assume that there is only overlap with pattern 17:

two groups of neurons: those that should be ‘on’ and ‘off’

$$\Pr\{S_i\,(t+1) = +1 \,|\, h_i = h^+\} = g\left[m^{17} \;\; t \;\right]$$

$$\Pr\{S_i\,(t+1) = +1 \,|\, h_i = h^-\} = g\left[-m^{17} \;\; t \;\right]$$

$$2m^{17}(t+1) = g\left[m^{17} \;\; t \;\right] + \{1 - g\left[-m^{17} \;\; t \;\right]\} - g\left[m^{17} \;\; t \;\right] - \{1 - g\left[-m^{17} \;\; t \;\right]\}$$

6.1 Stochastic Hopfield model: memory retrieval

$$2m^{17}(t+1) = g\left[m^{17} \quad t \quad\right] + \{1 - g\left[-m^{17} \quad t \quad\right]\} - g\left[m^{17} \quad t \quad\right] - \{1 - g\left[-m^{17} \quad t \quad\right]\}$$

$$m^{17}(t+1) = F\left[m^{17} \quad t \quad\right]$$

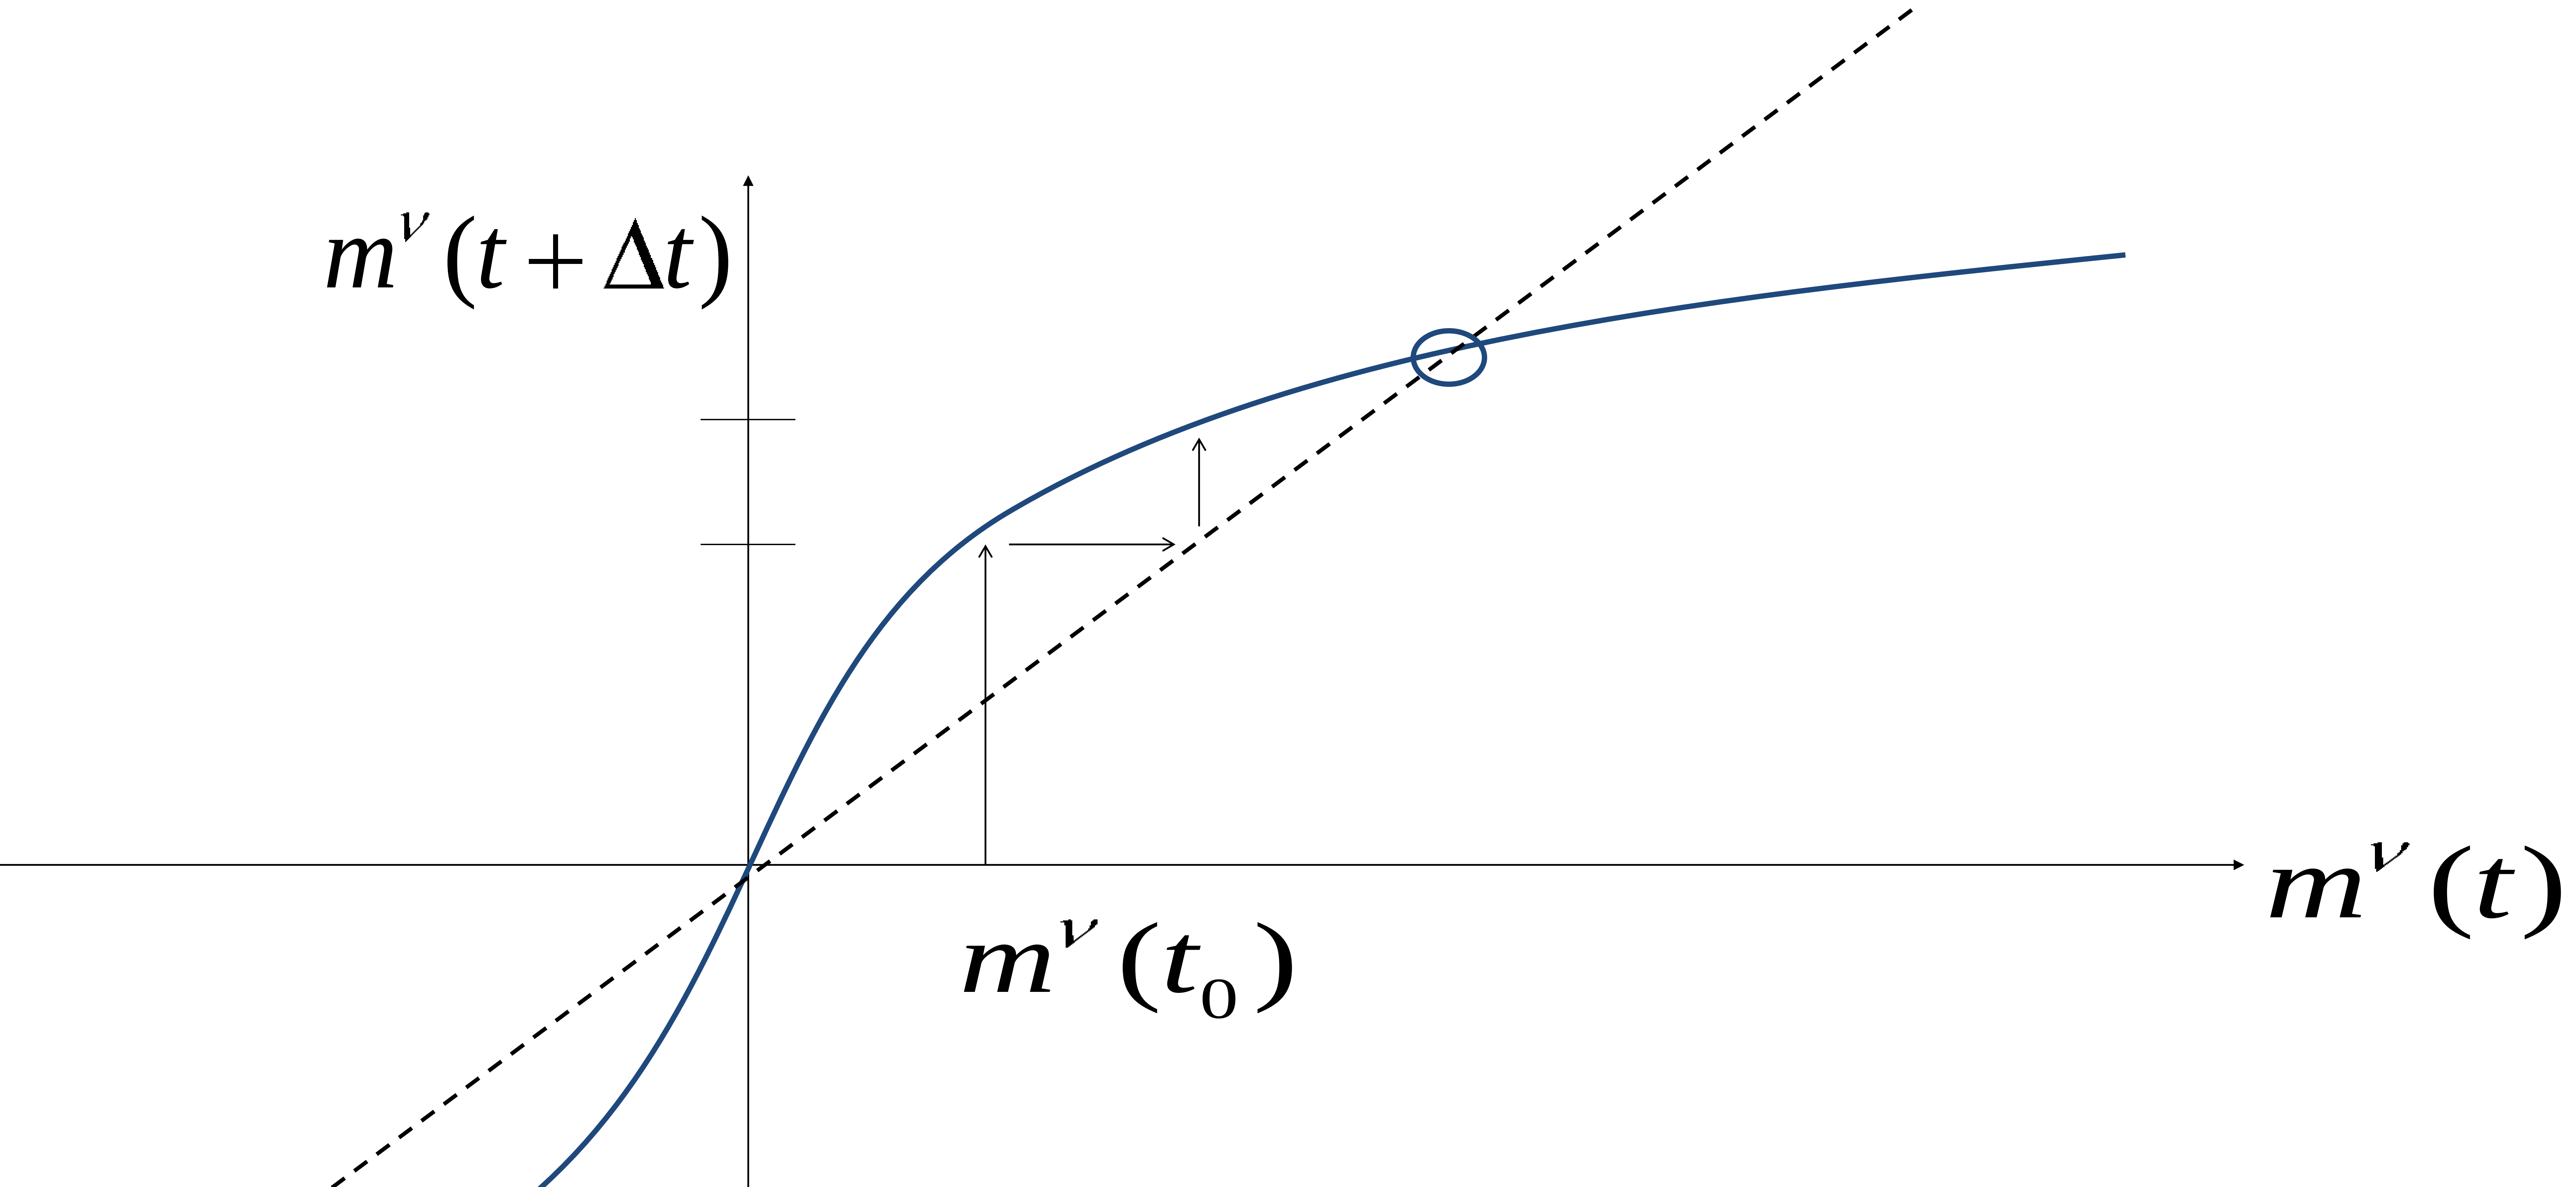

Week 6: Hopfield model continued

ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

Biological Modeling of Neural Networks

Week 6

Hebbian LEARNING and ASSOCIATIVE MEMORY

Wulfram Gerstner

EPFL, Lausanne, Switzerland

6.1 Stochastic Hopfield Model

6.2. Energy landscape

6.3. Low-activity patterns

6.4. Attractor memorie

- spiking neurons
- experimental data

6.2 memory retrieval

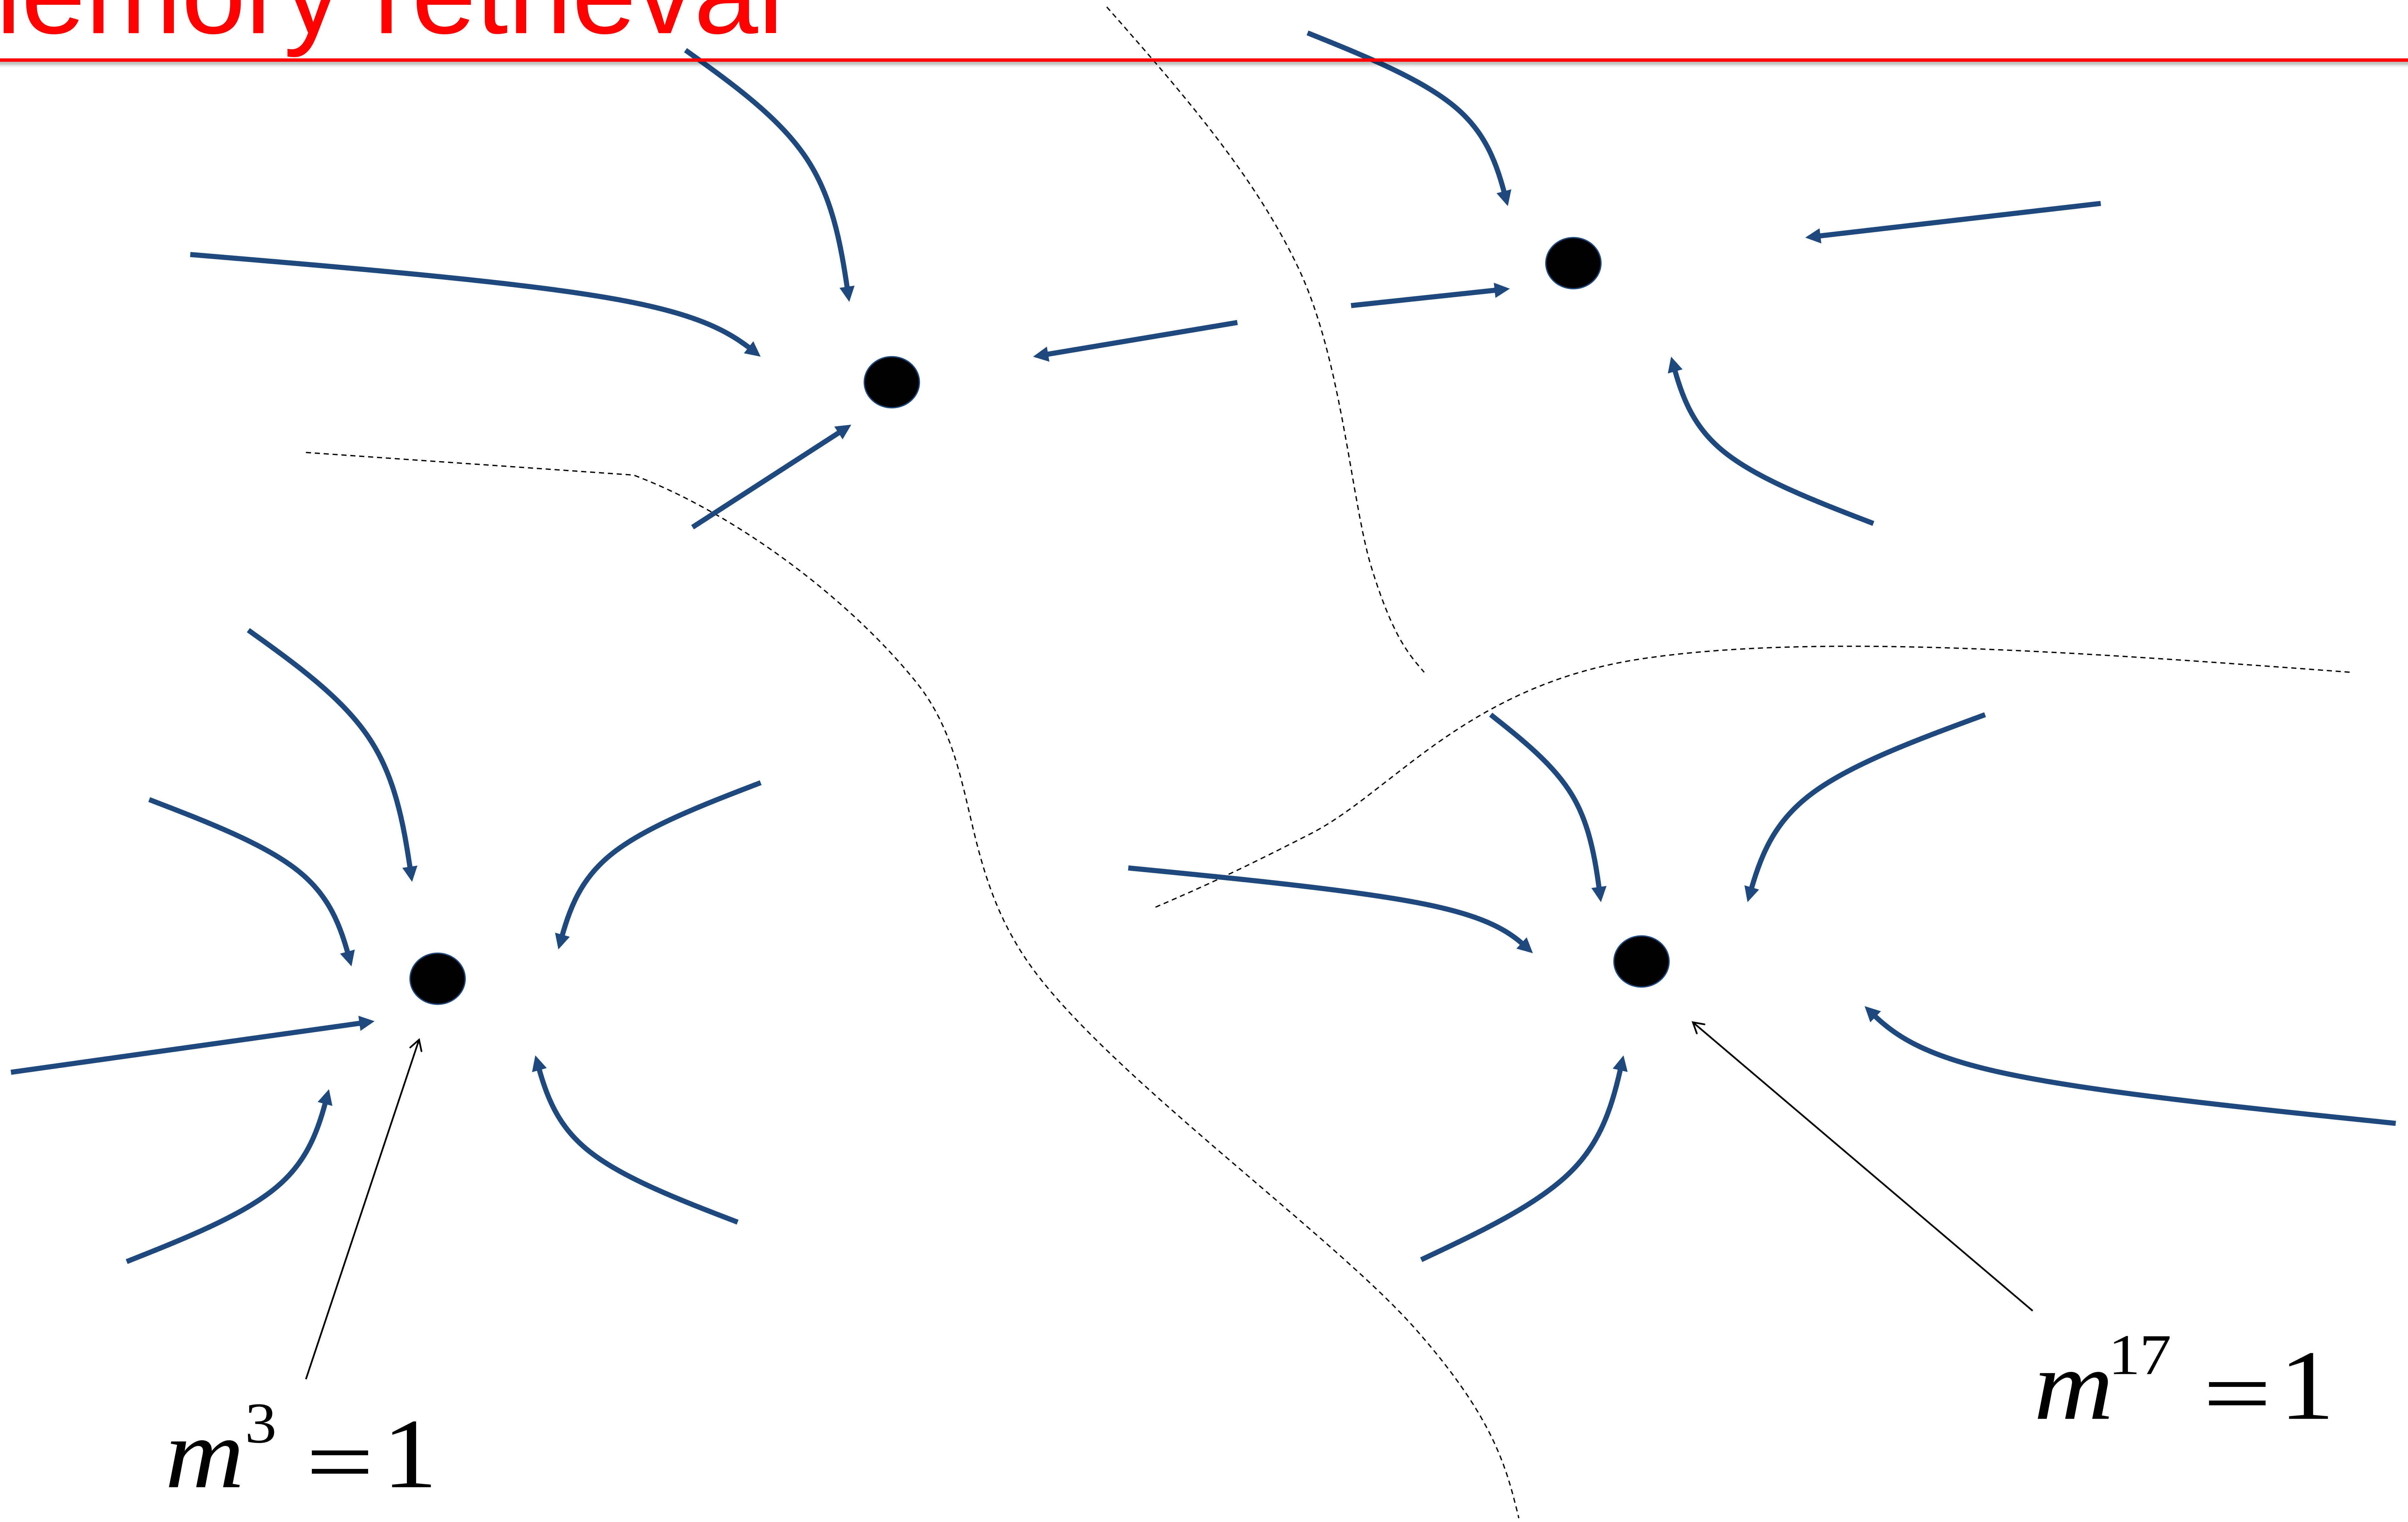

6.2 Symmetric interactions: Energy picture

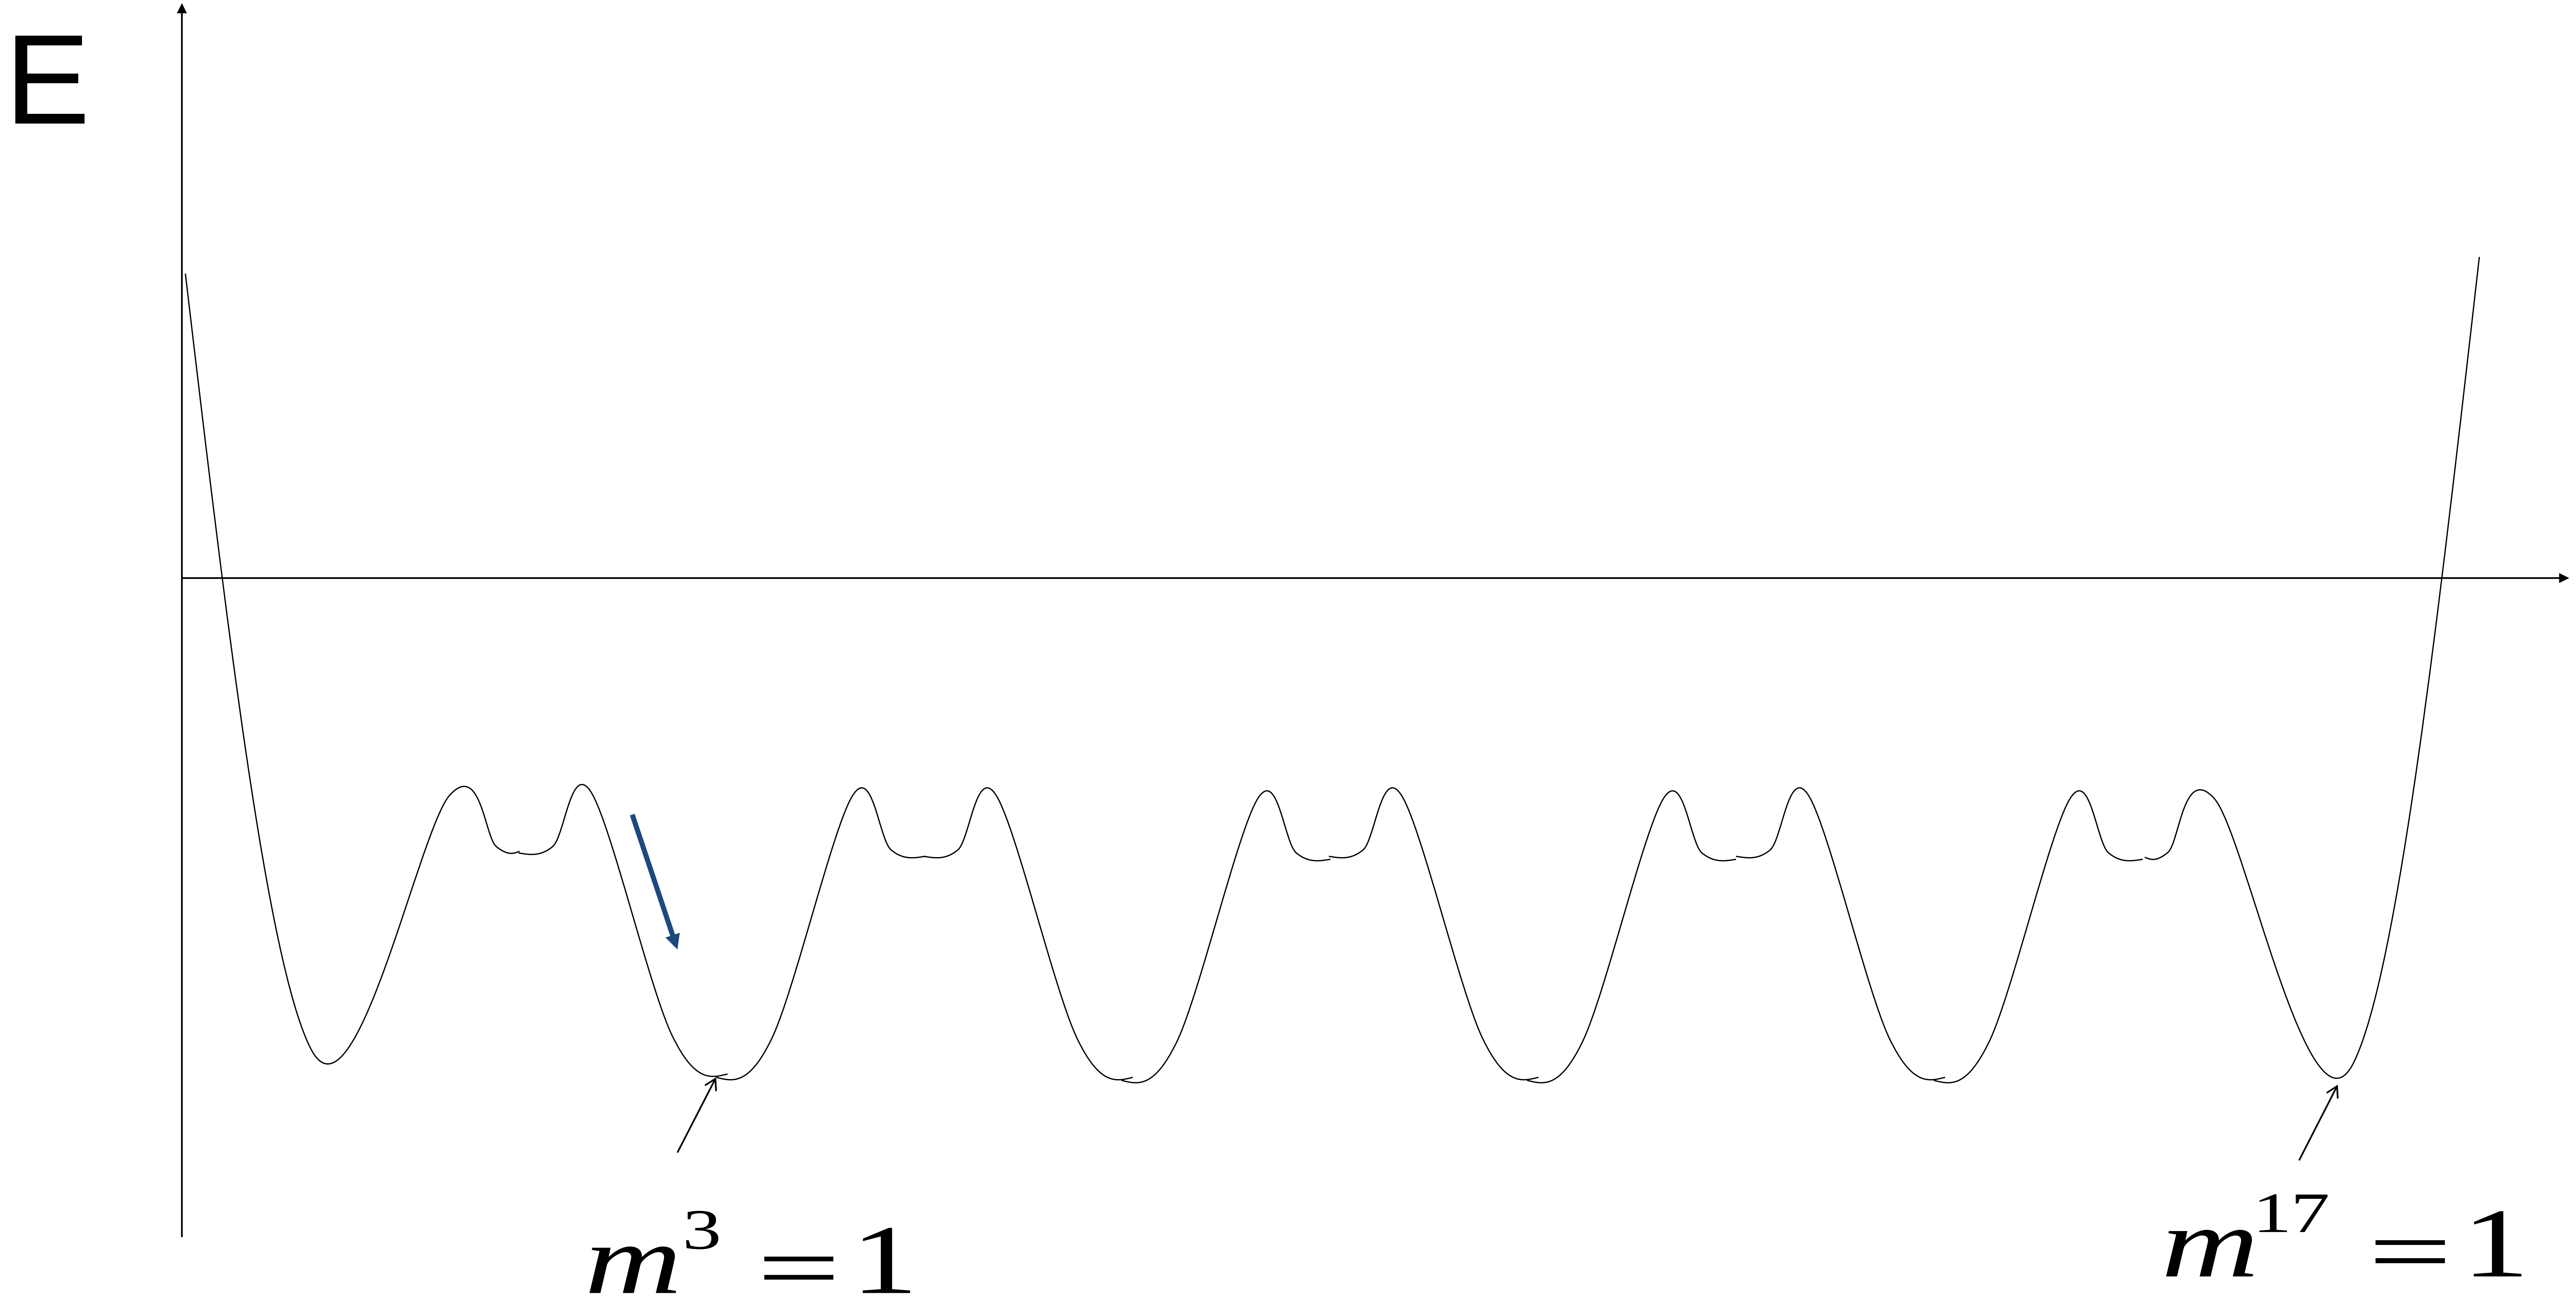

Exercise 2 now: Energy picture

$$S_i(t+1) = \mathrm{sgn}[\sum_j w_{ij} S_j(t)]$$

$$E = -\sum_i \sum_j w_{ij} S_i S_j$$

Next lecture 11:25

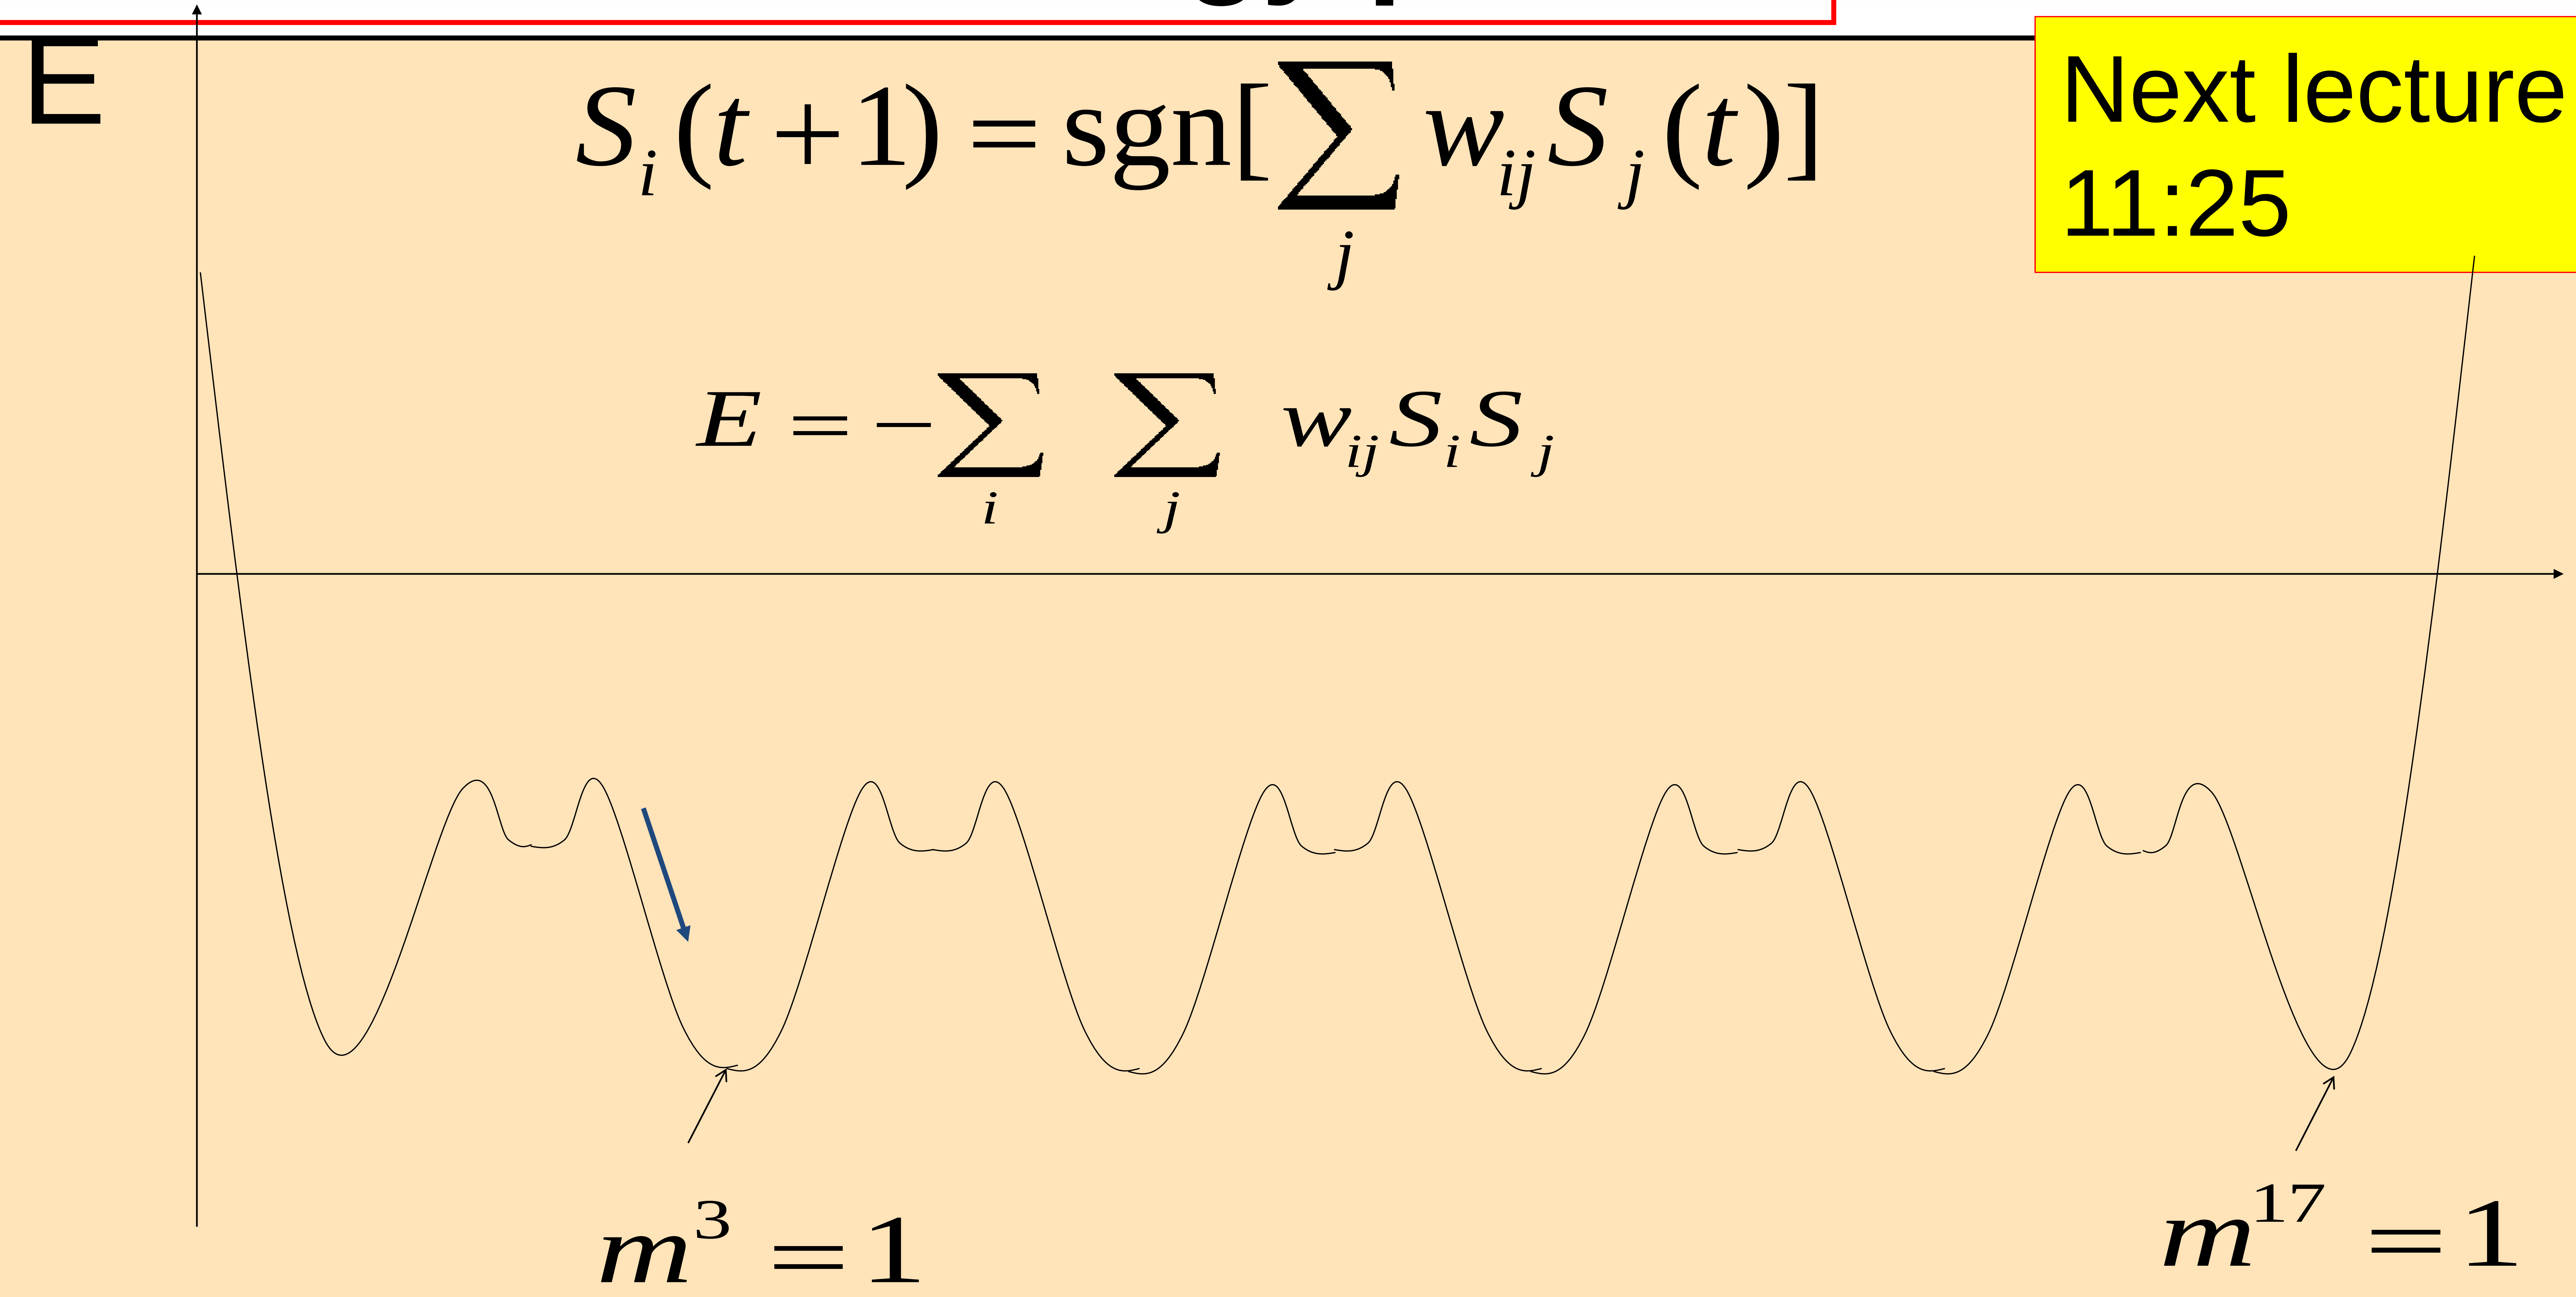

Week 6: Hopfield model continued

ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

Biological Modeling of Neural Networks

Week 6

Hebbian LEARNING and ASSOCIATIVE MEMORY

Wulfram Gerstner

EPFL, Lausanne, Switzerland

6.1 Stochastic Hopfield Model

6.2. Energy landscape

6.3. Low-activity patterns

6.4. Attractor memories

- spiking neurons
- experimental data

6.3 Attractor memory

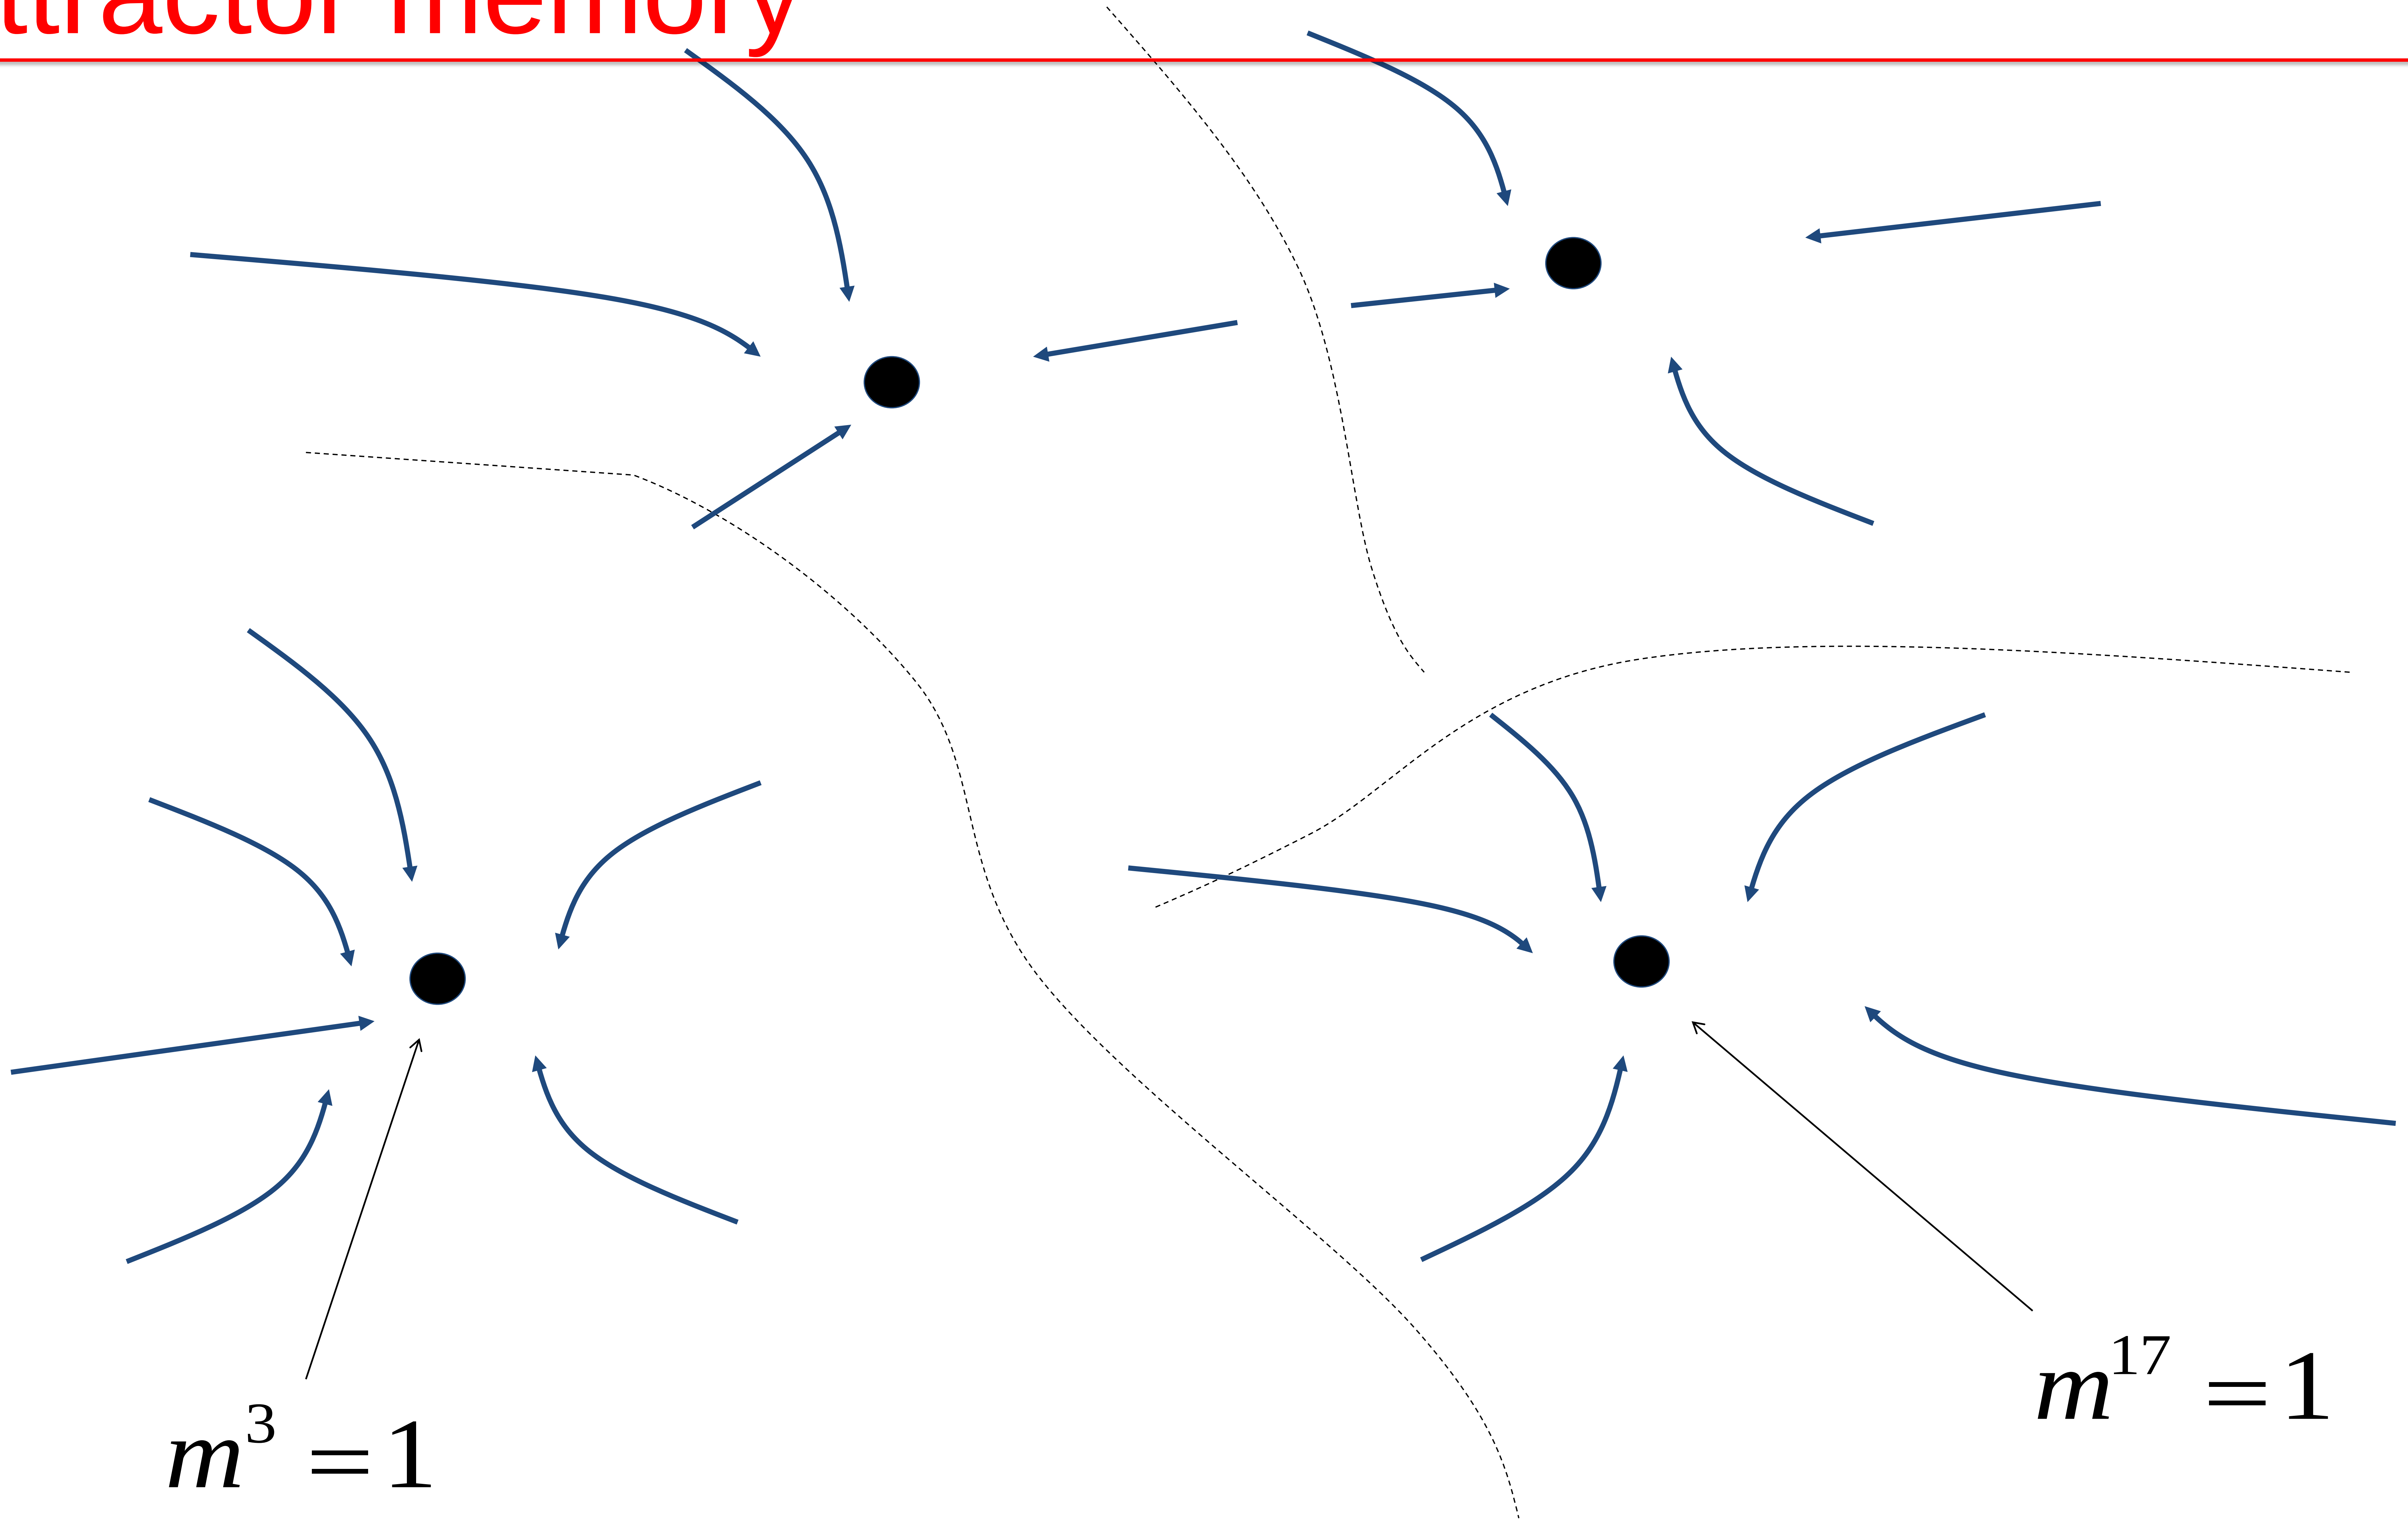

Memory with spiking neurons

-Mean activity of patterns?

-Separation of excitation and inhibition?

-Modeling?

-Neural data?

6.3 attractor memory with low activity patterns

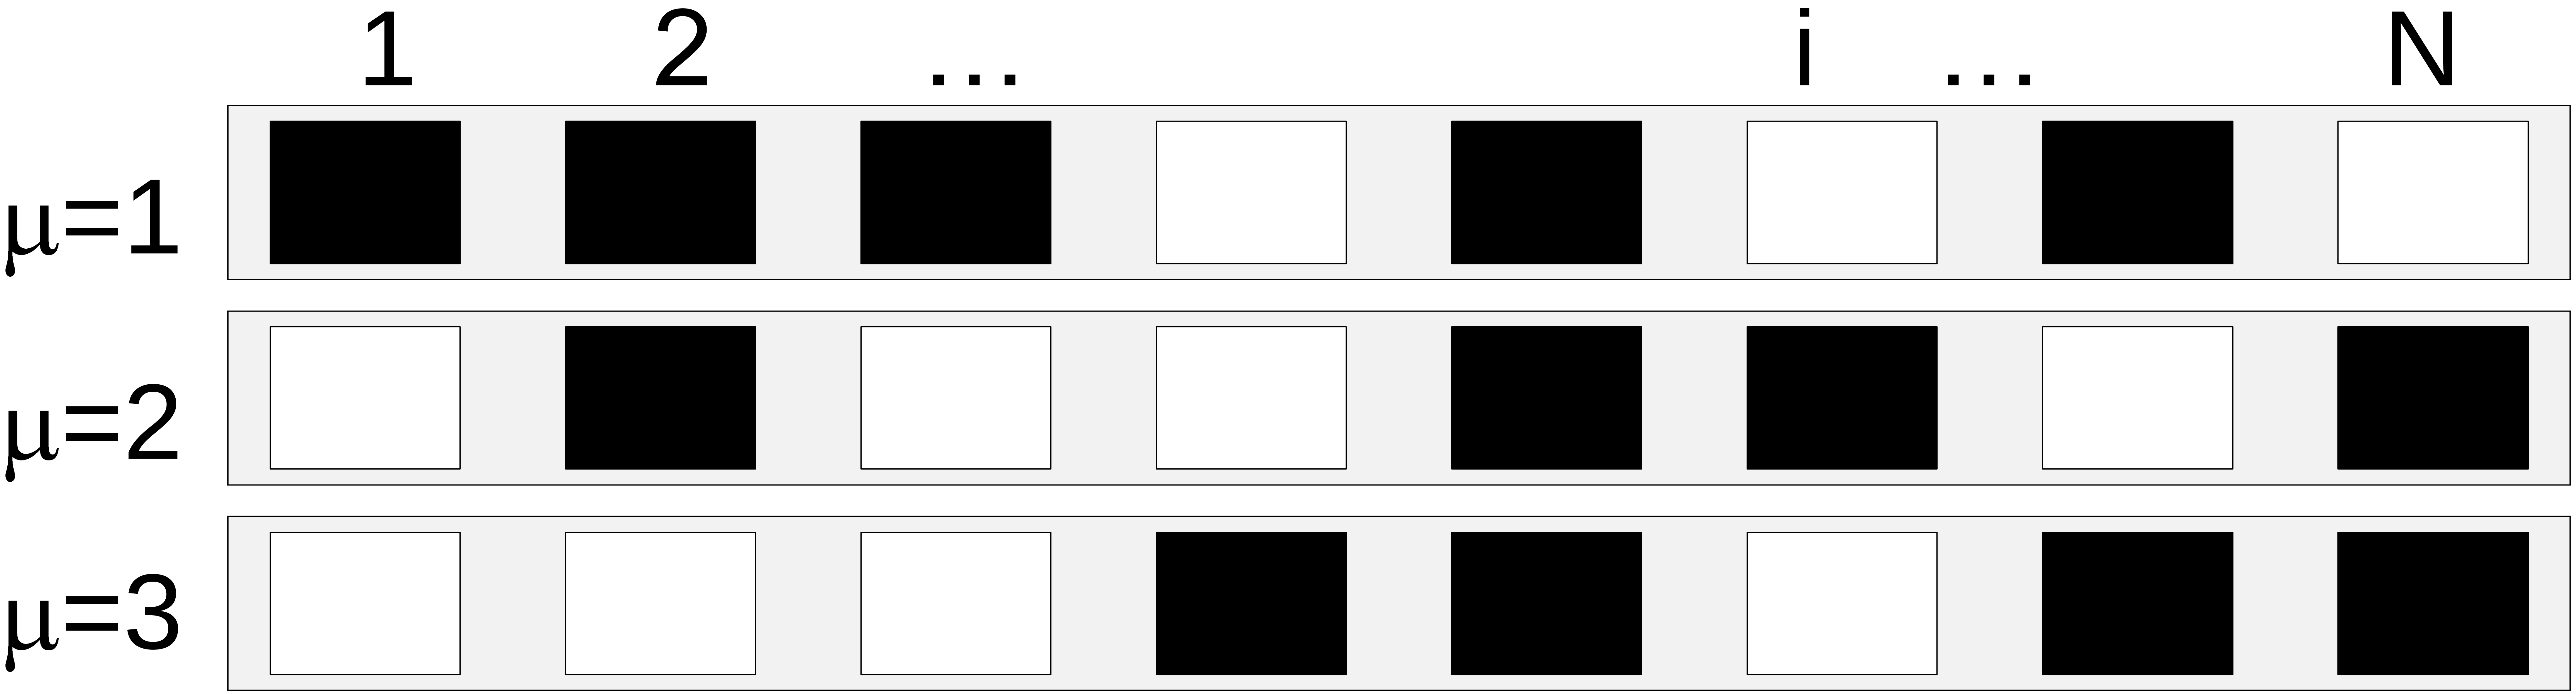

Random patterns +/-1 with zero mean →
50 percent of neurons should be active in each pattern

$$w_{ij} = \frac{1}{N} \sum_{\mu} p_i^{\mu} p_j^{\mu}$$

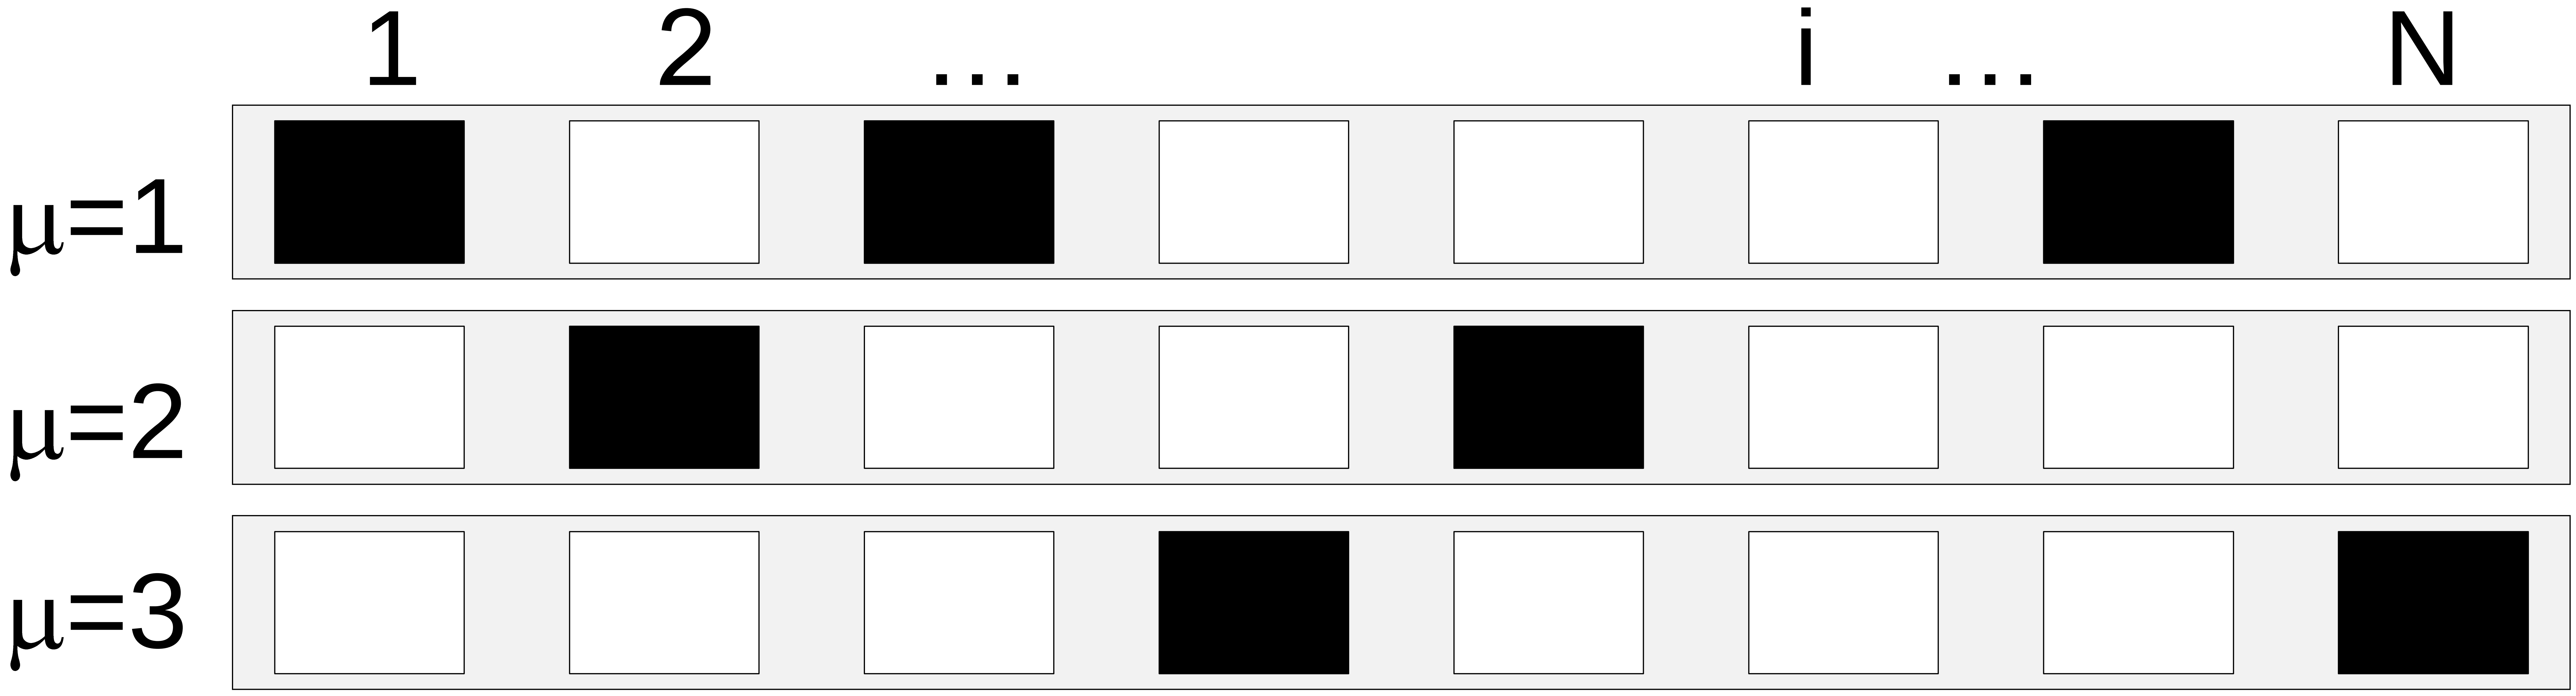

Random patterns +/-1 with **low activity (mean =a<0)** →
20 percent of neurons should be active in each pattern

$$w_{ij} = \frac{1}{N}\sum_{\mu}(p_i^{\mu} - b)(p_j^{\mu} - a)$$

Some constant

activity

Week 6: Hopfield model continued

ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

Biological Modeling of Neural Networks

Week 6

Hebbian LEARNING and ASSOCIATIVE MEMORY

Wulfram Gerstner

EPFL, Lausanne, Switzerland

6.1 Stochastic Hopfield Model
6.2. Energy landscape
6.3. Low-activity patterns
6.4. Attractor memories
- spiking neurons
- experimental data

6.4 attractor memory with spiking neurons

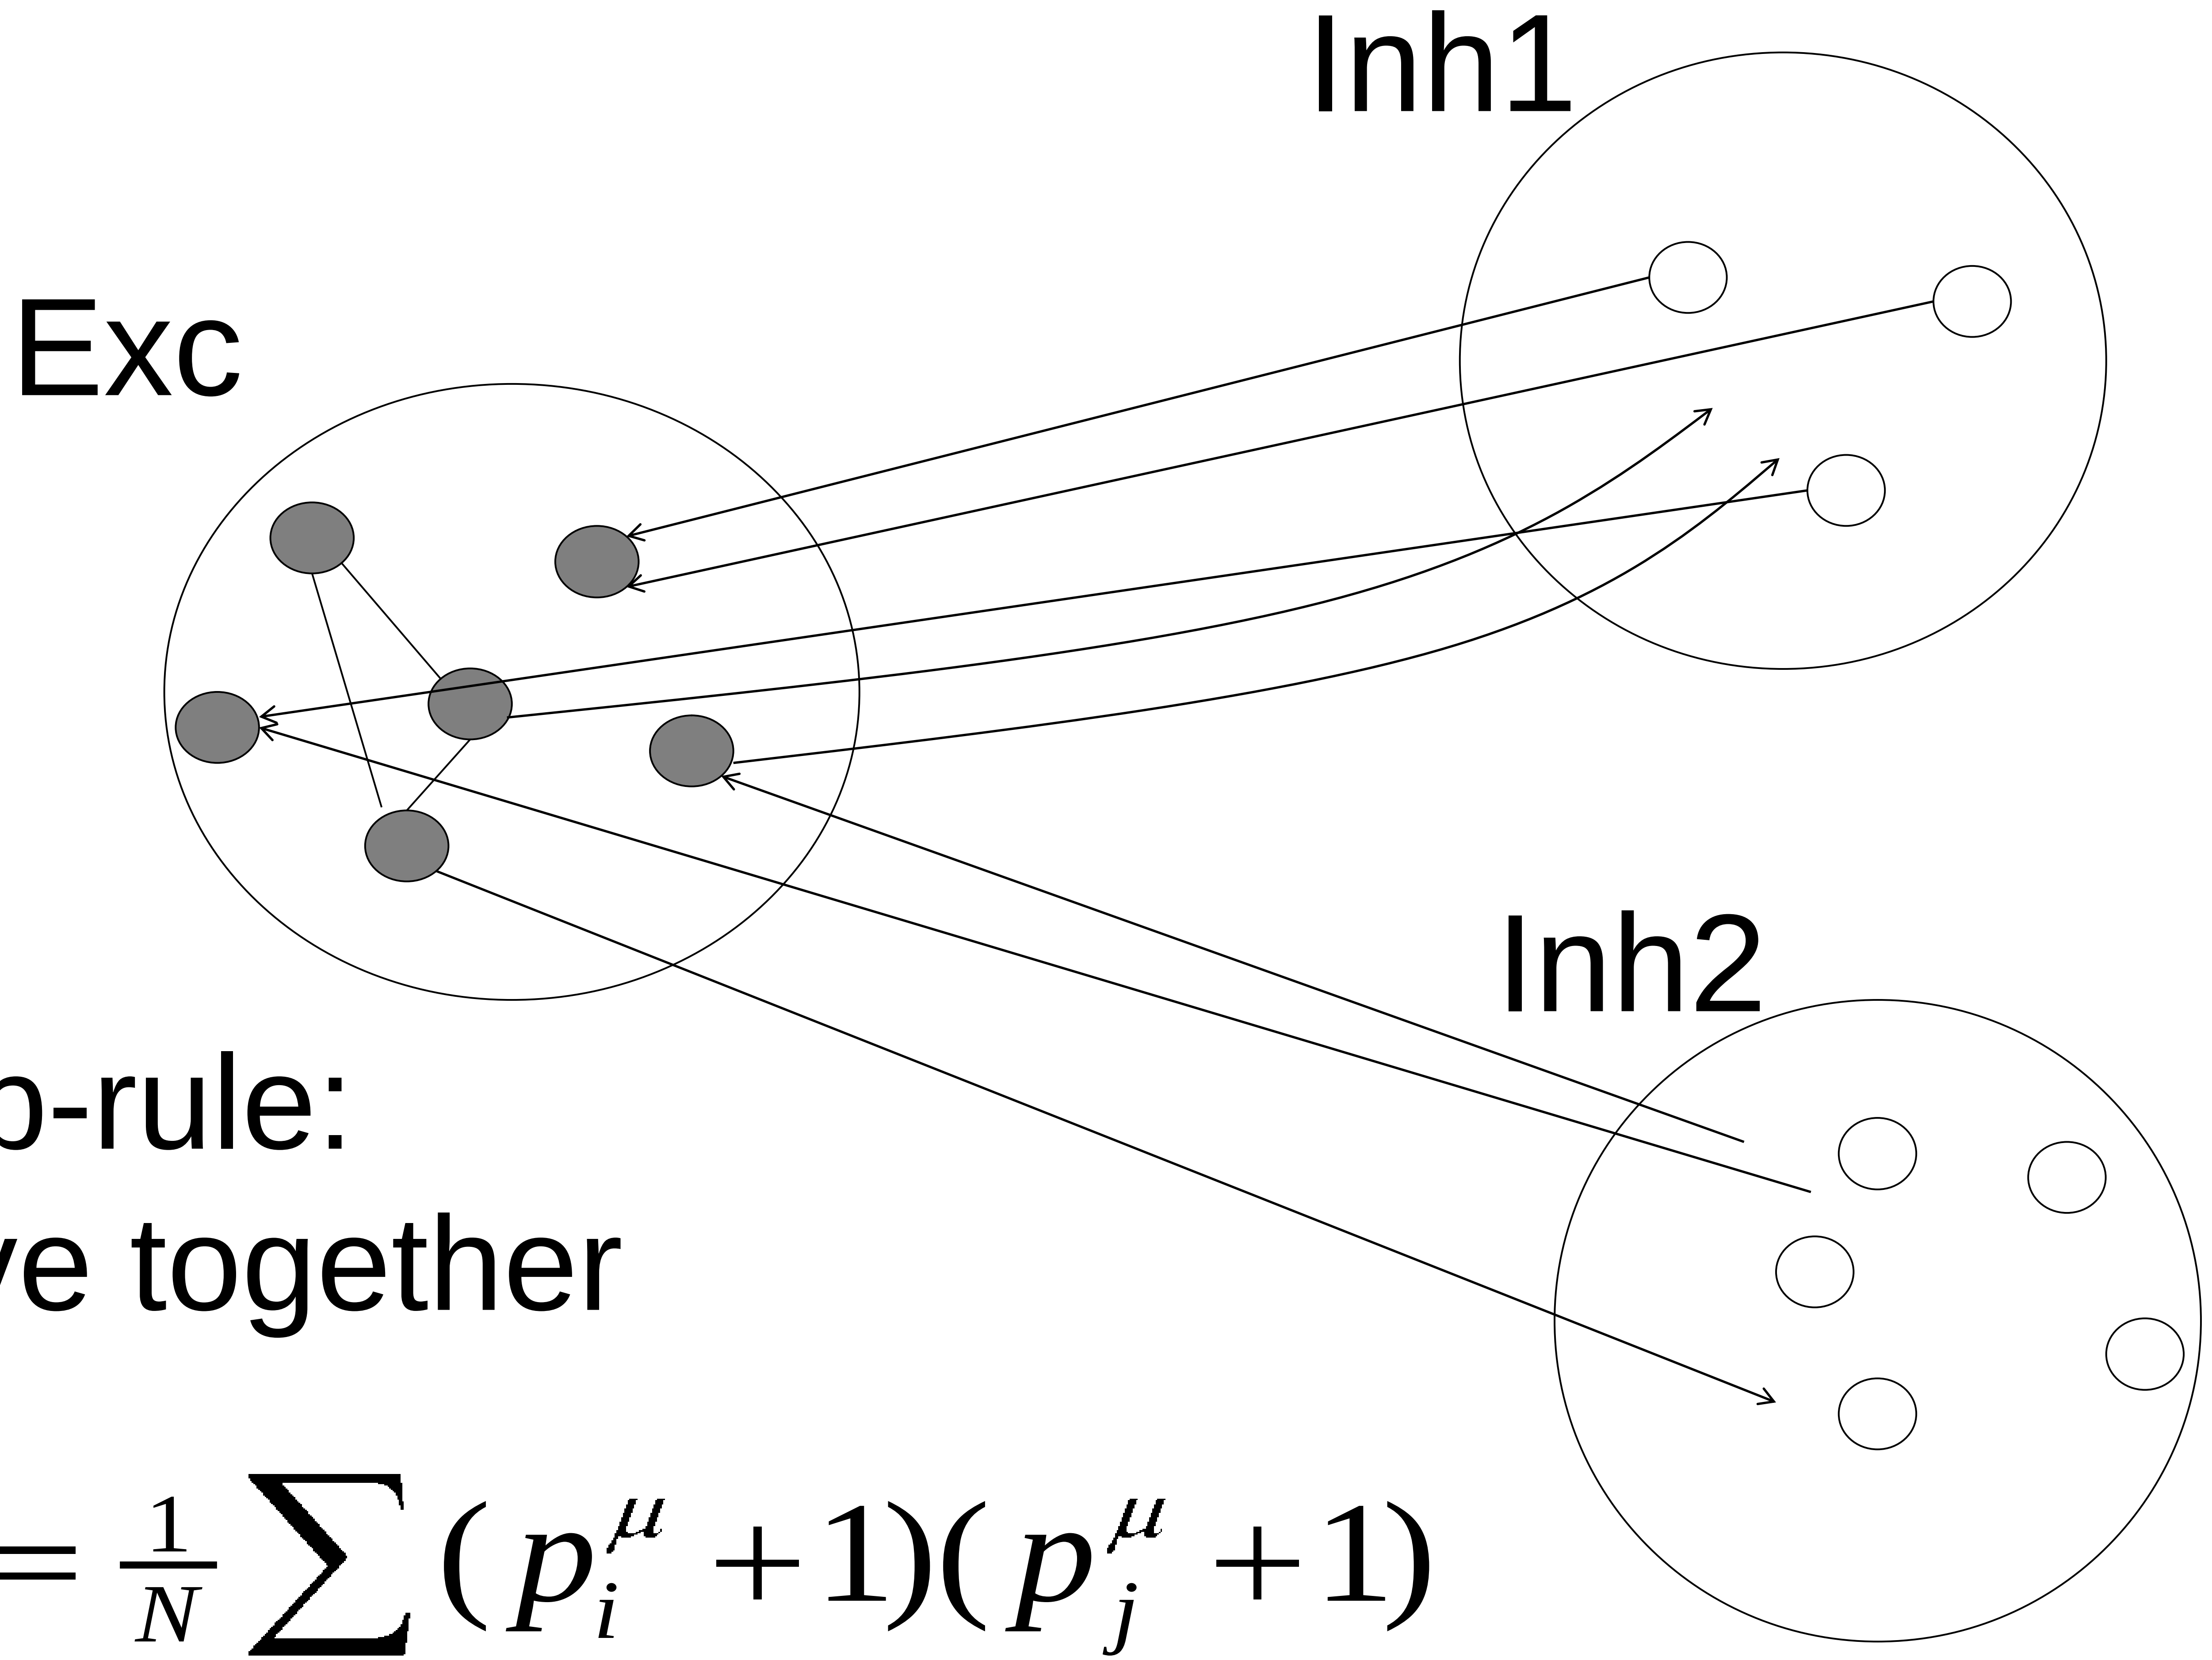

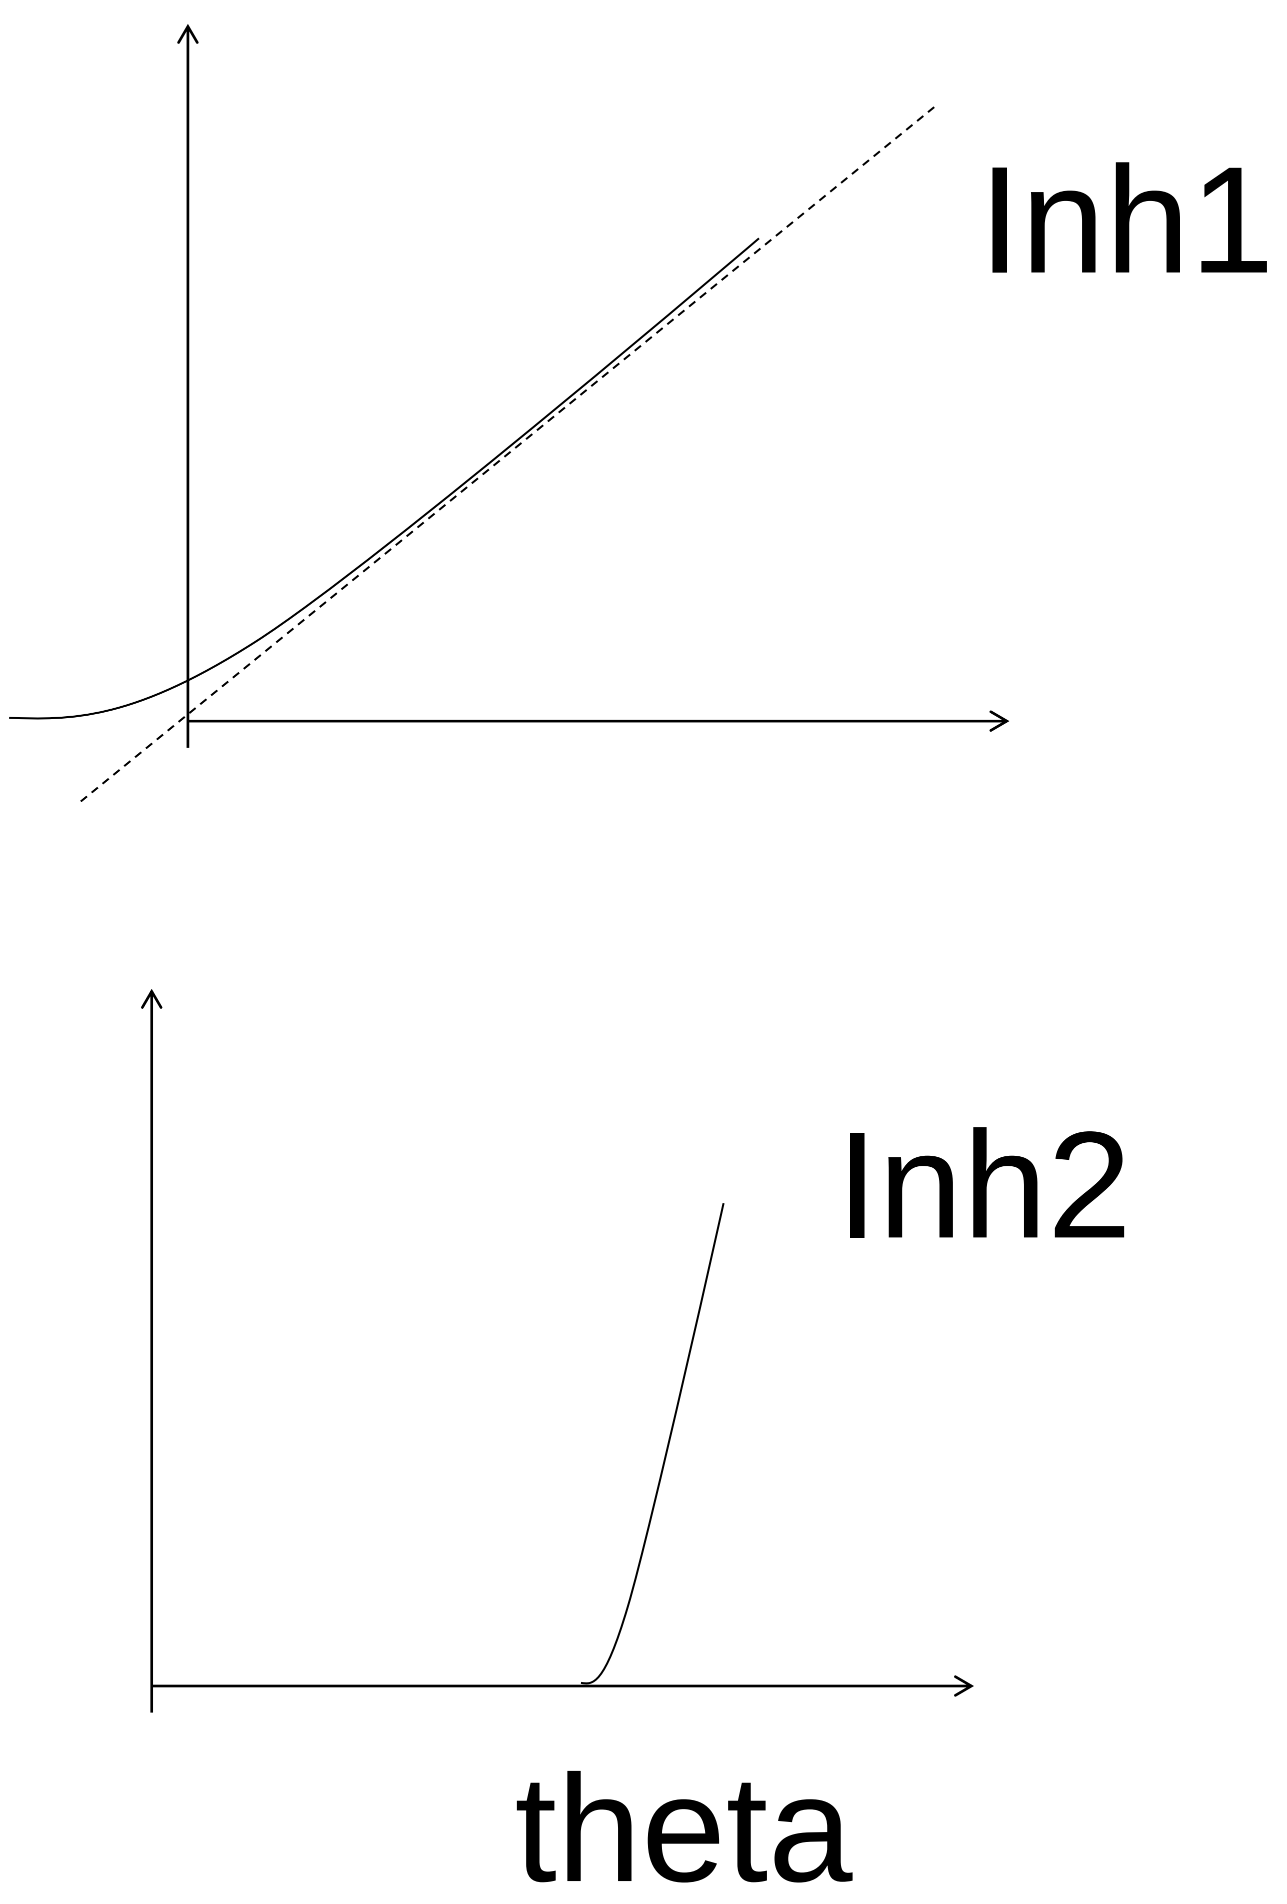

Hebb-rule:
Active together

$$w_{ij} = \frac{1}{N} \sum_{\mu} (p_i^{\mu} + 1)(p_j^{\mu} + 1)$$

Overlap with patterns 1 … 3

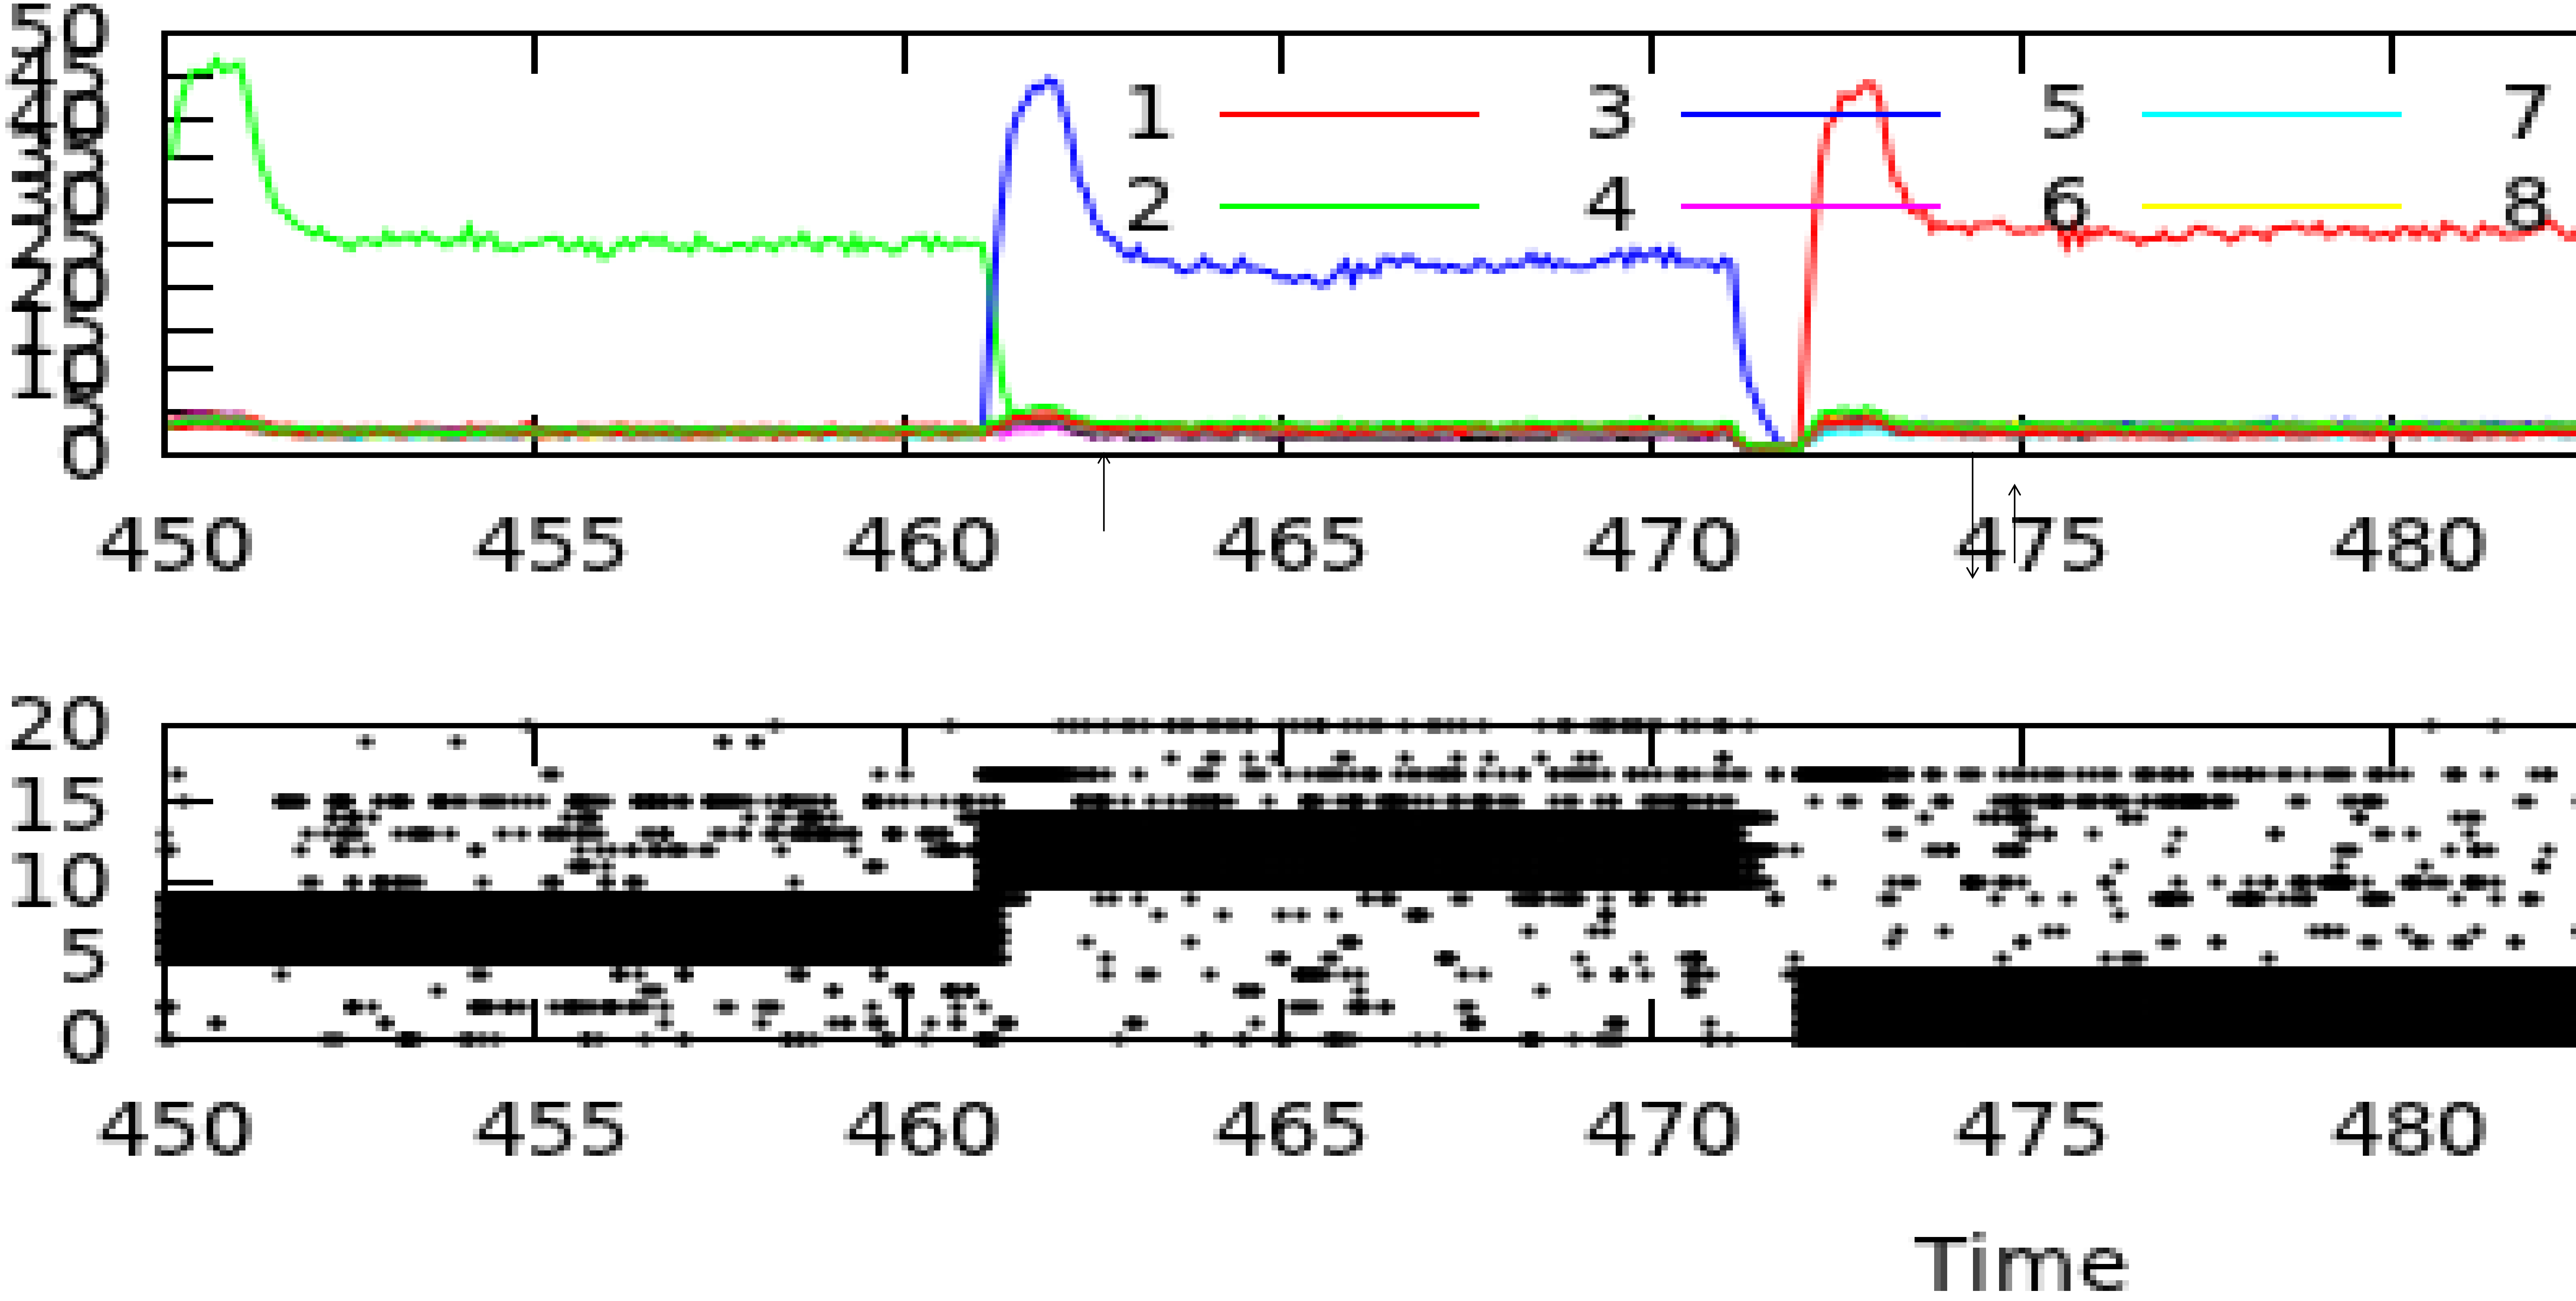

Overlap with patterns 1 … 11 (80 patterns stored!)

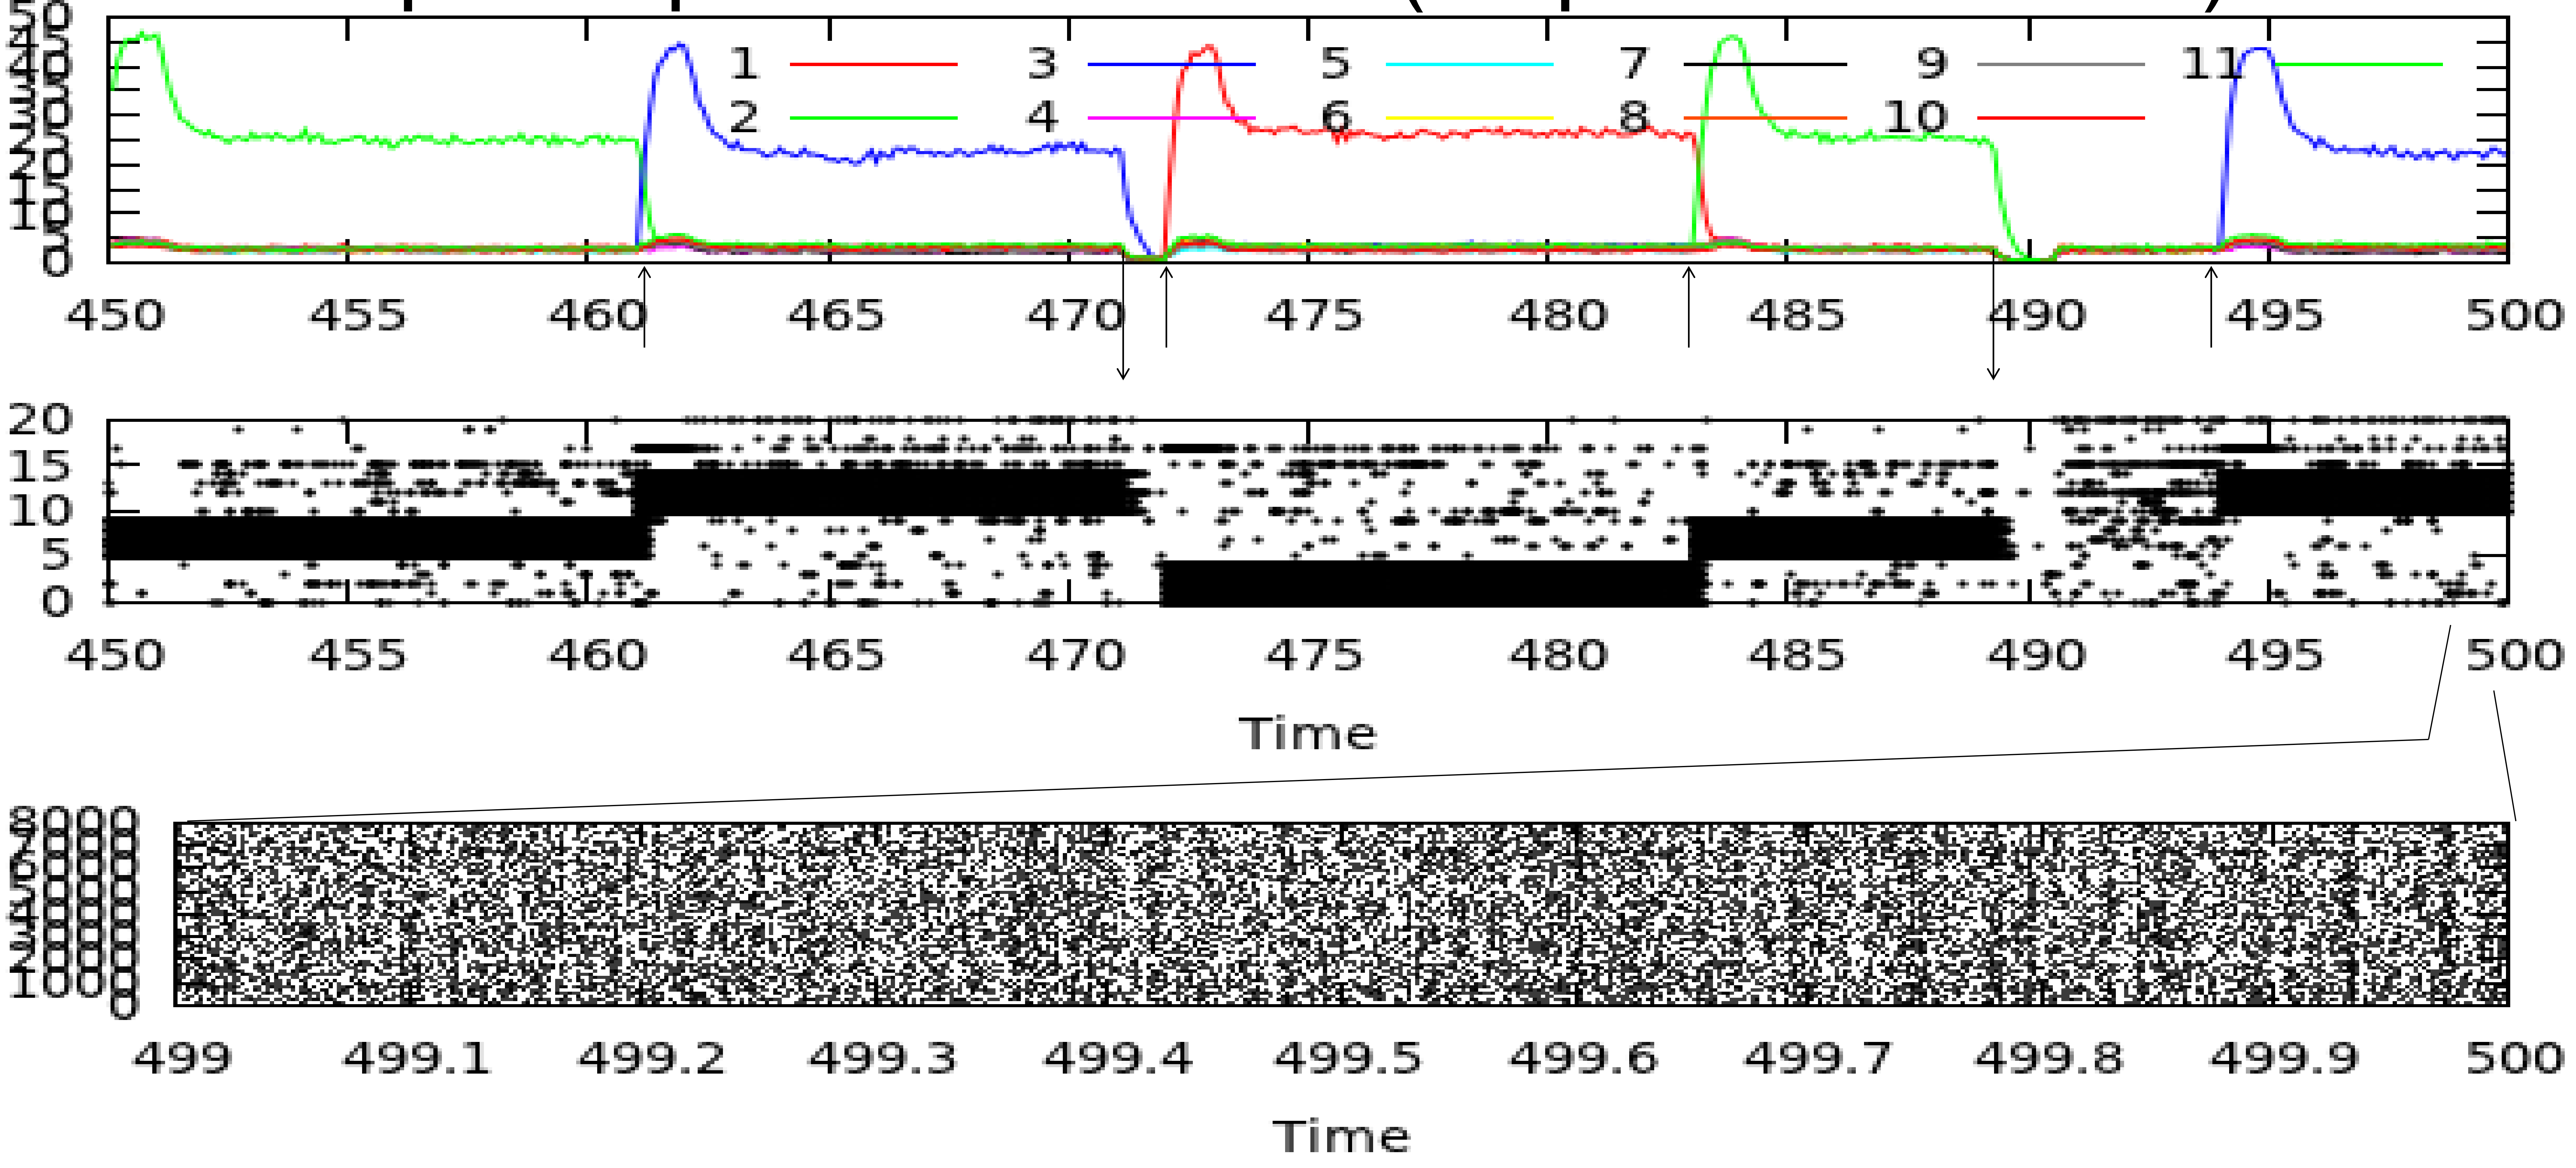

Memory with spiking neurons

-Low activity of patterns?

-Separation of excitation and inhibitio ?

-Modeling?

All possible

-Neural data?

Human Hippocampus

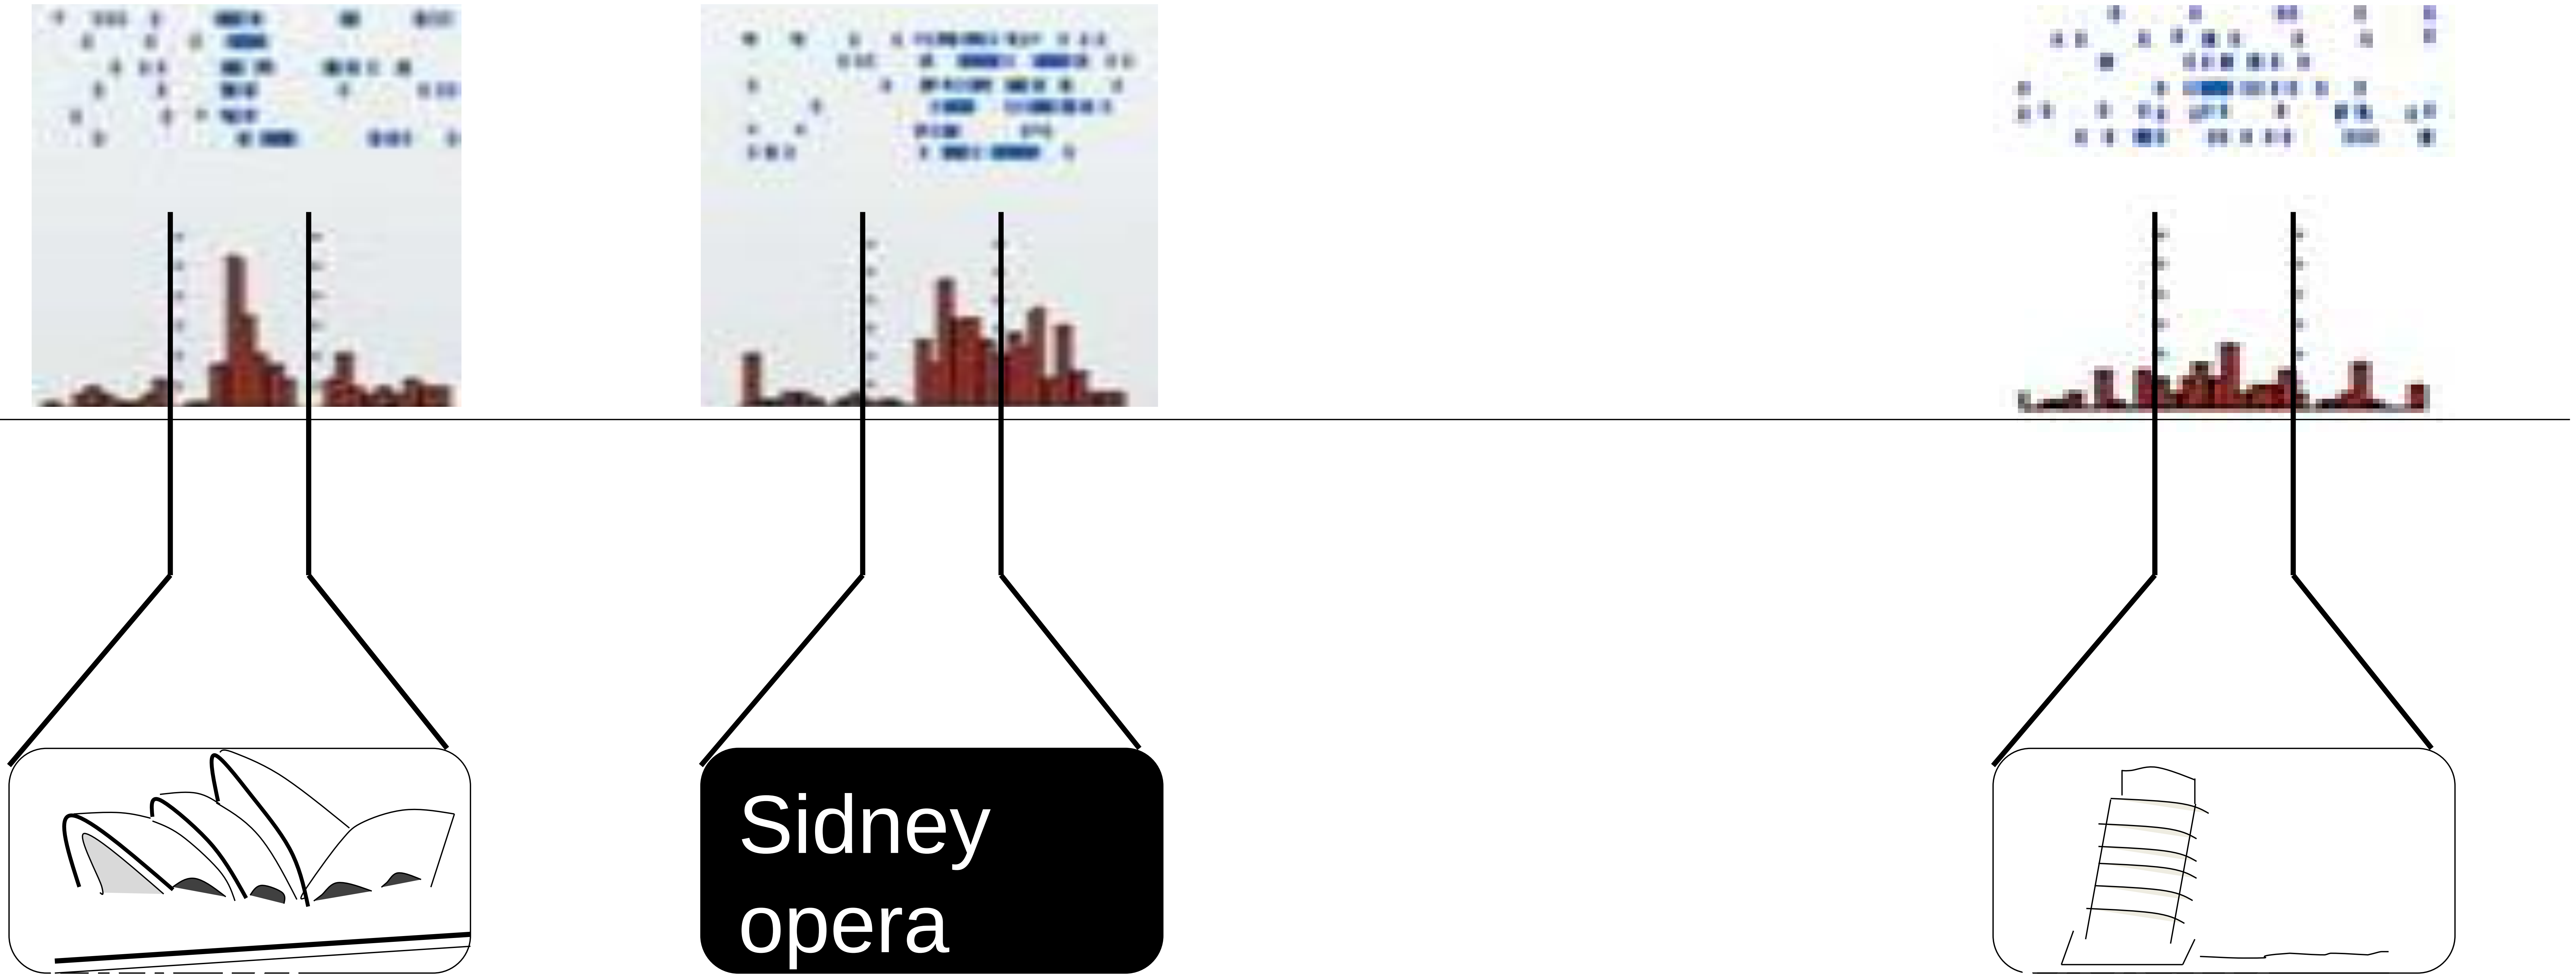

Quiroga, R. Q., Reddy, L., Kreiman, G., Koch, C., and Fried, I. (2005).
Invariant visual representation by single neurons in the human brain.
Nature, 435:1102-1107.

6.4 memory data

Delayed Matching to Sample Task

Animal experiments

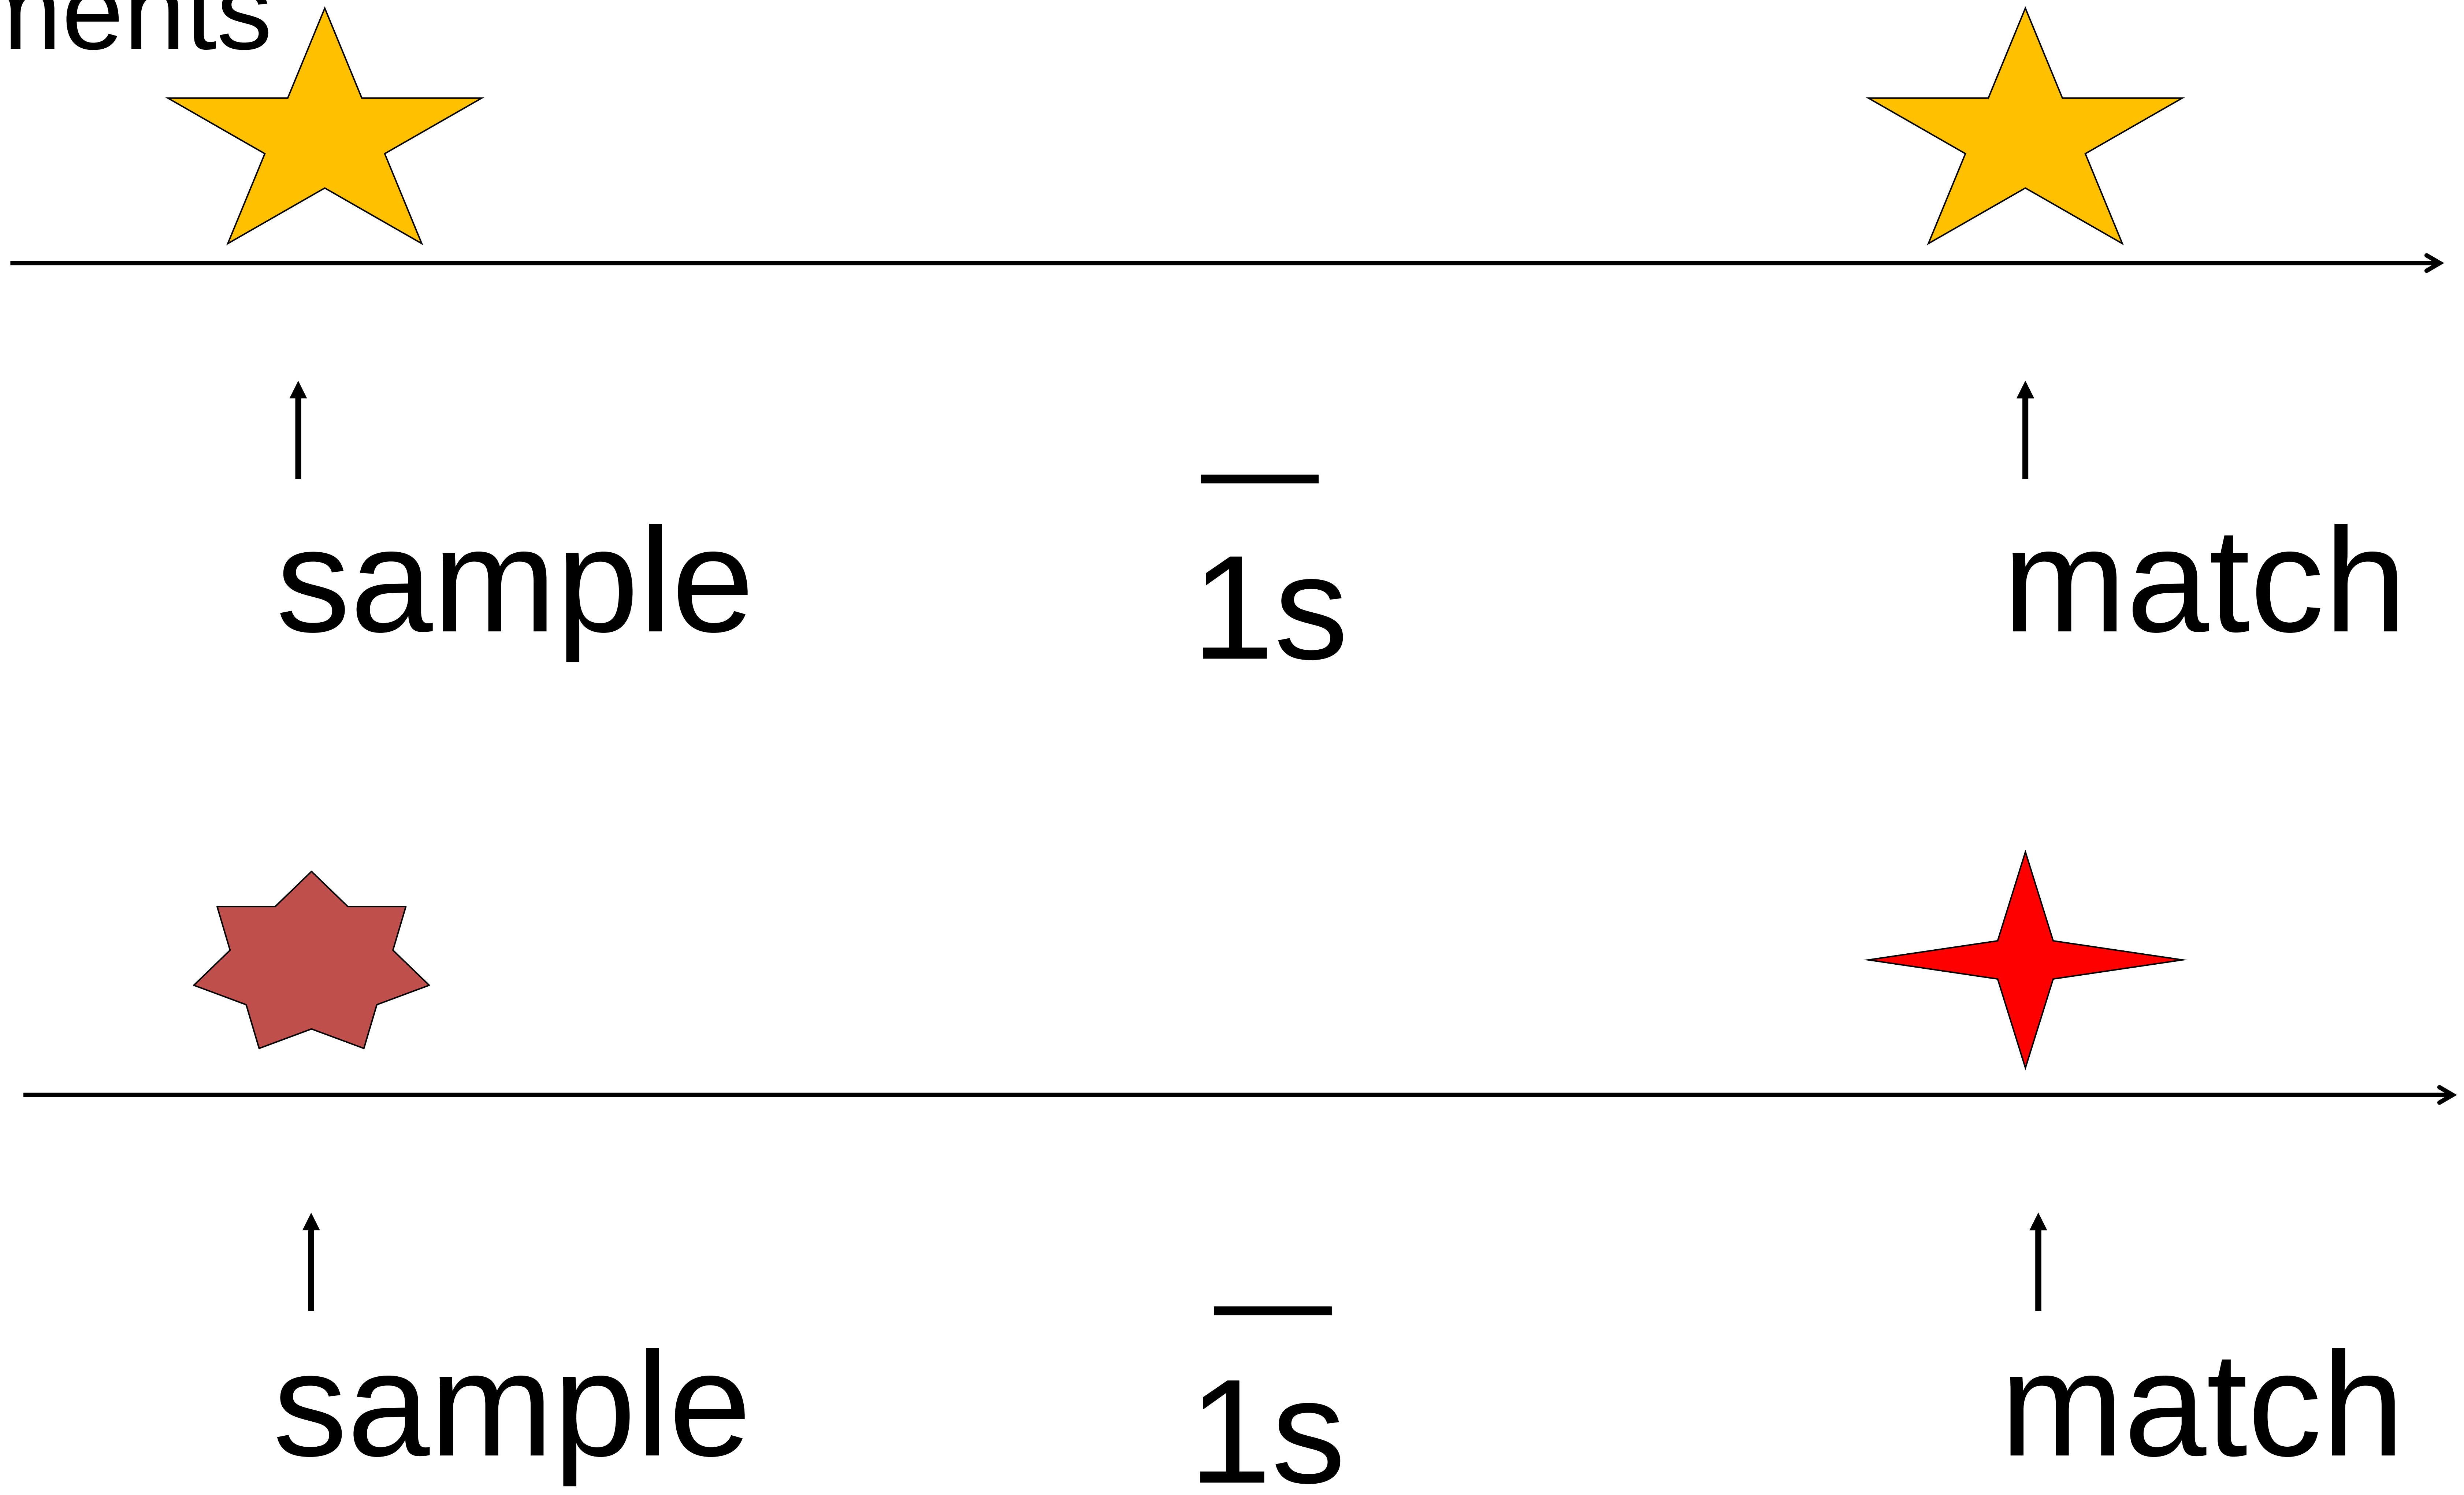

6.4 memory data

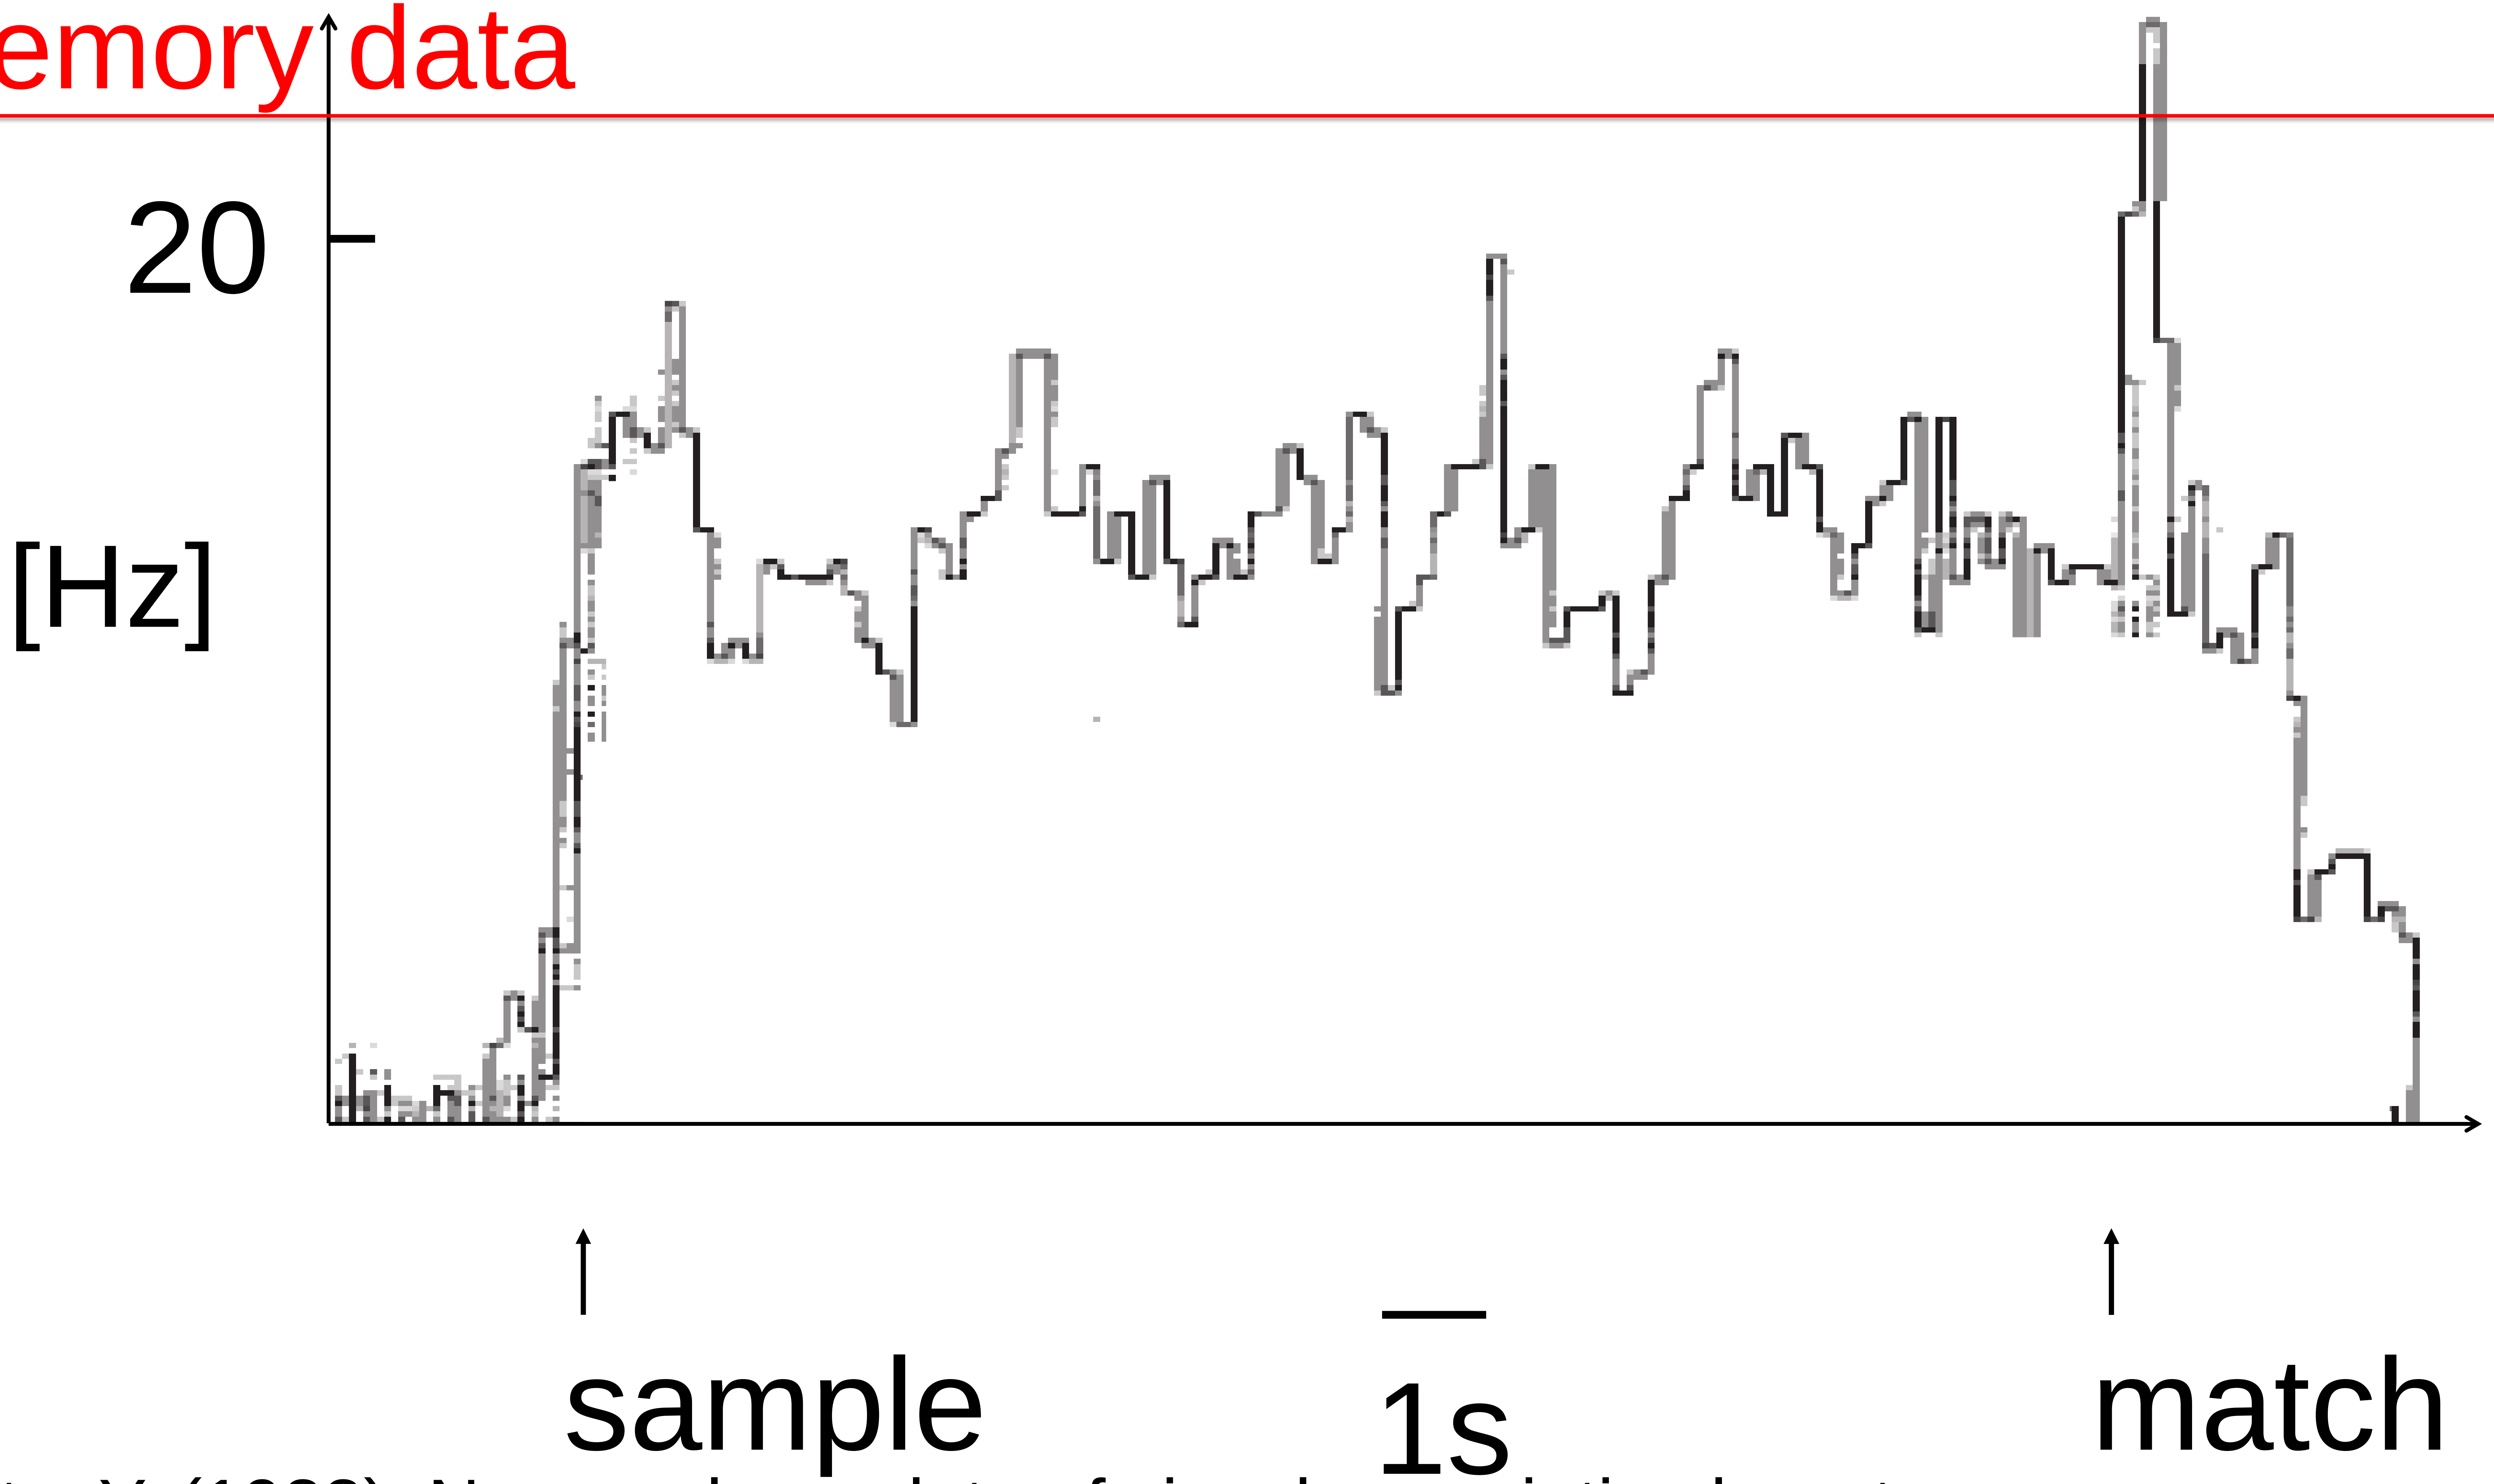

Miyashita, Y. (1988). Neuronal correlate of visual associative long-term memory in the primate temporal cortex. Nature, 335:817-820.

6.4 memory data

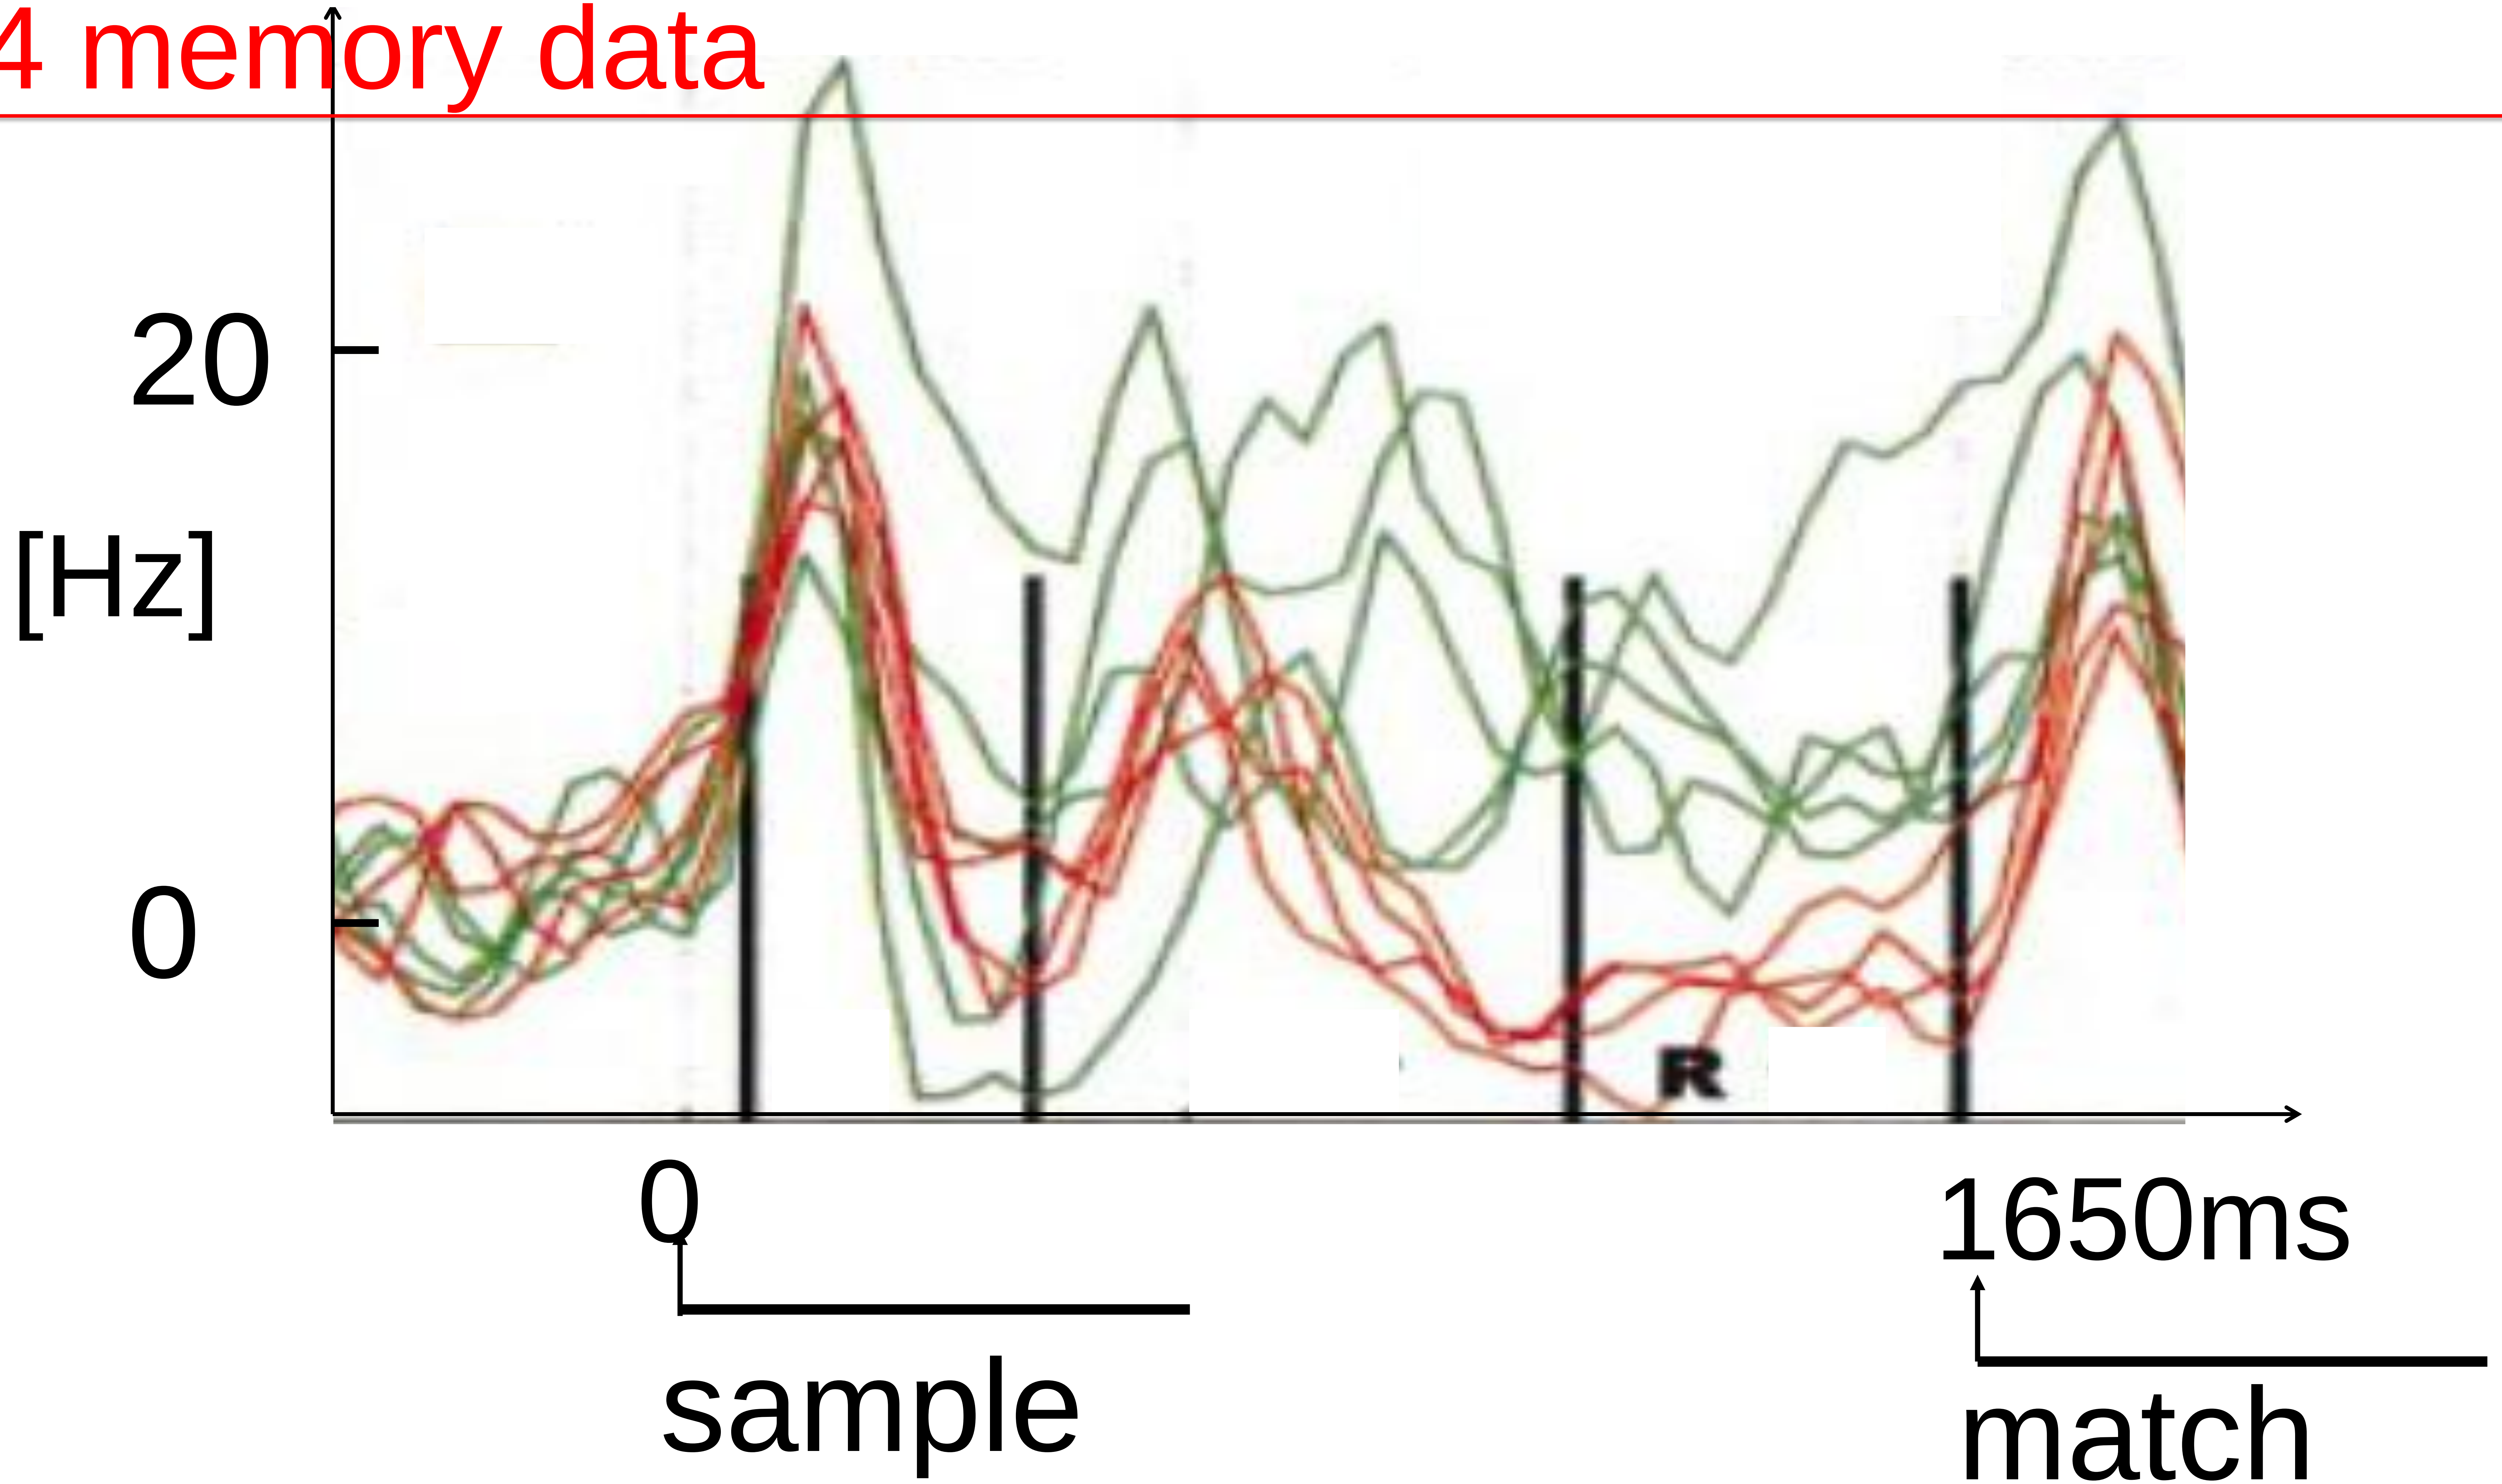

Rainer and Miller (2002). Timecourse of object-related neural activity in the primate prefrontal cortex during a short-term memory task. Europ. J. Neurosci., 15:1244-1254.

Exercise 3 NOW- from Hopfield to spikes

In the Hopfield model, neurons are characterized by a binary variable $S_i = +/-1$. For an interpretation in terms of spikes it is, however, more appealing to work with a binary variable x_i which is zero or 1.

(i) Write $S_i = 2x_i - 1$ and rewrite the Hopfield model in terms of x_i. What are the conditions so that the input potential is

$$h_i = \sum_j w_{ij} x_j$$

(ii) Repeat the same calculation for low-activity patterns and and weights

$$w_{ij} = \frac{1}{N} \sum_{\mu} (p_i^{\mu} - b)(p_j^{\mu} - a)$$

with some constants a and b

The end