Computational Neuroscience: Neuronal Dynamics of Cognition

ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

A: ASSOCIATIVE MEMORY
in a Network of Neurons

Wulfram Gerstner

EPFL, Lausanne, Switzerland

Reading for this week:
NEURONAL DYNAMICS
- Ch. 17.1 - 17.2.4

Cambridge Univ. Press

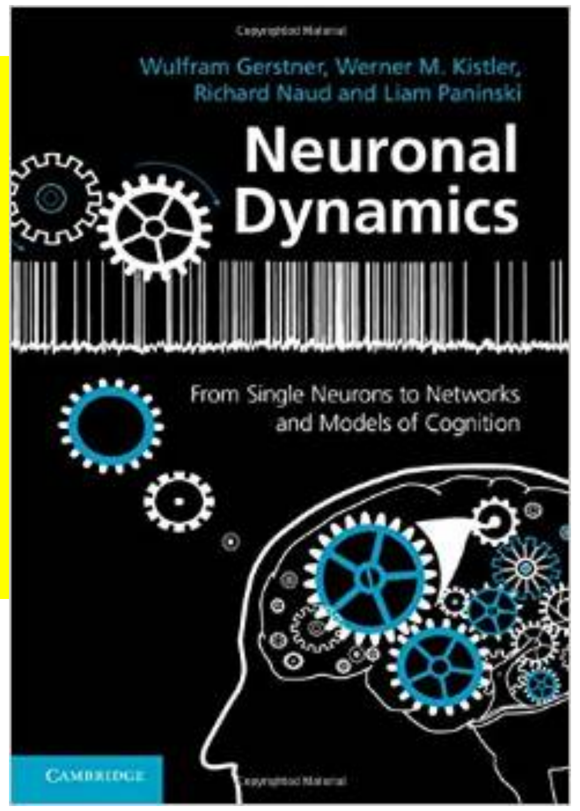

1 Introduction
- networks of neuron
- systems for computing
- associative memory

2 Classification by similarity

3 Detour: Magnetic Materials

4 Hopfield Model

5 Learning of Associations

6 Storage Capacity

1. memory in the brain

-president
-first day of undergraduate
-apple

Our memory has multiple aspects
- recent and far-back
- events, places, facts, concepts

1. memory in the brain

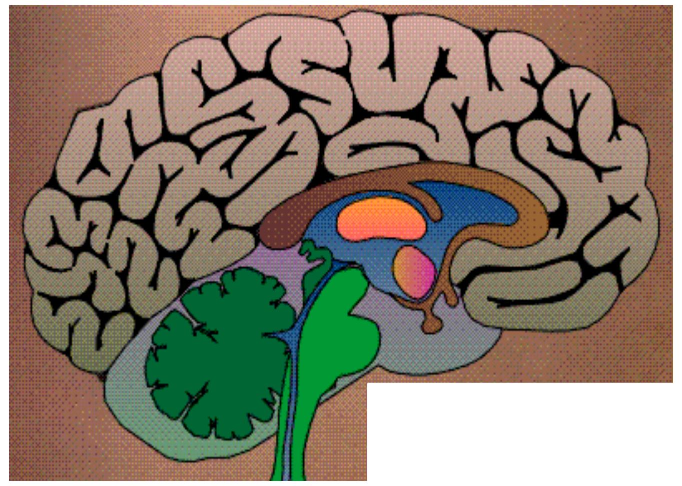

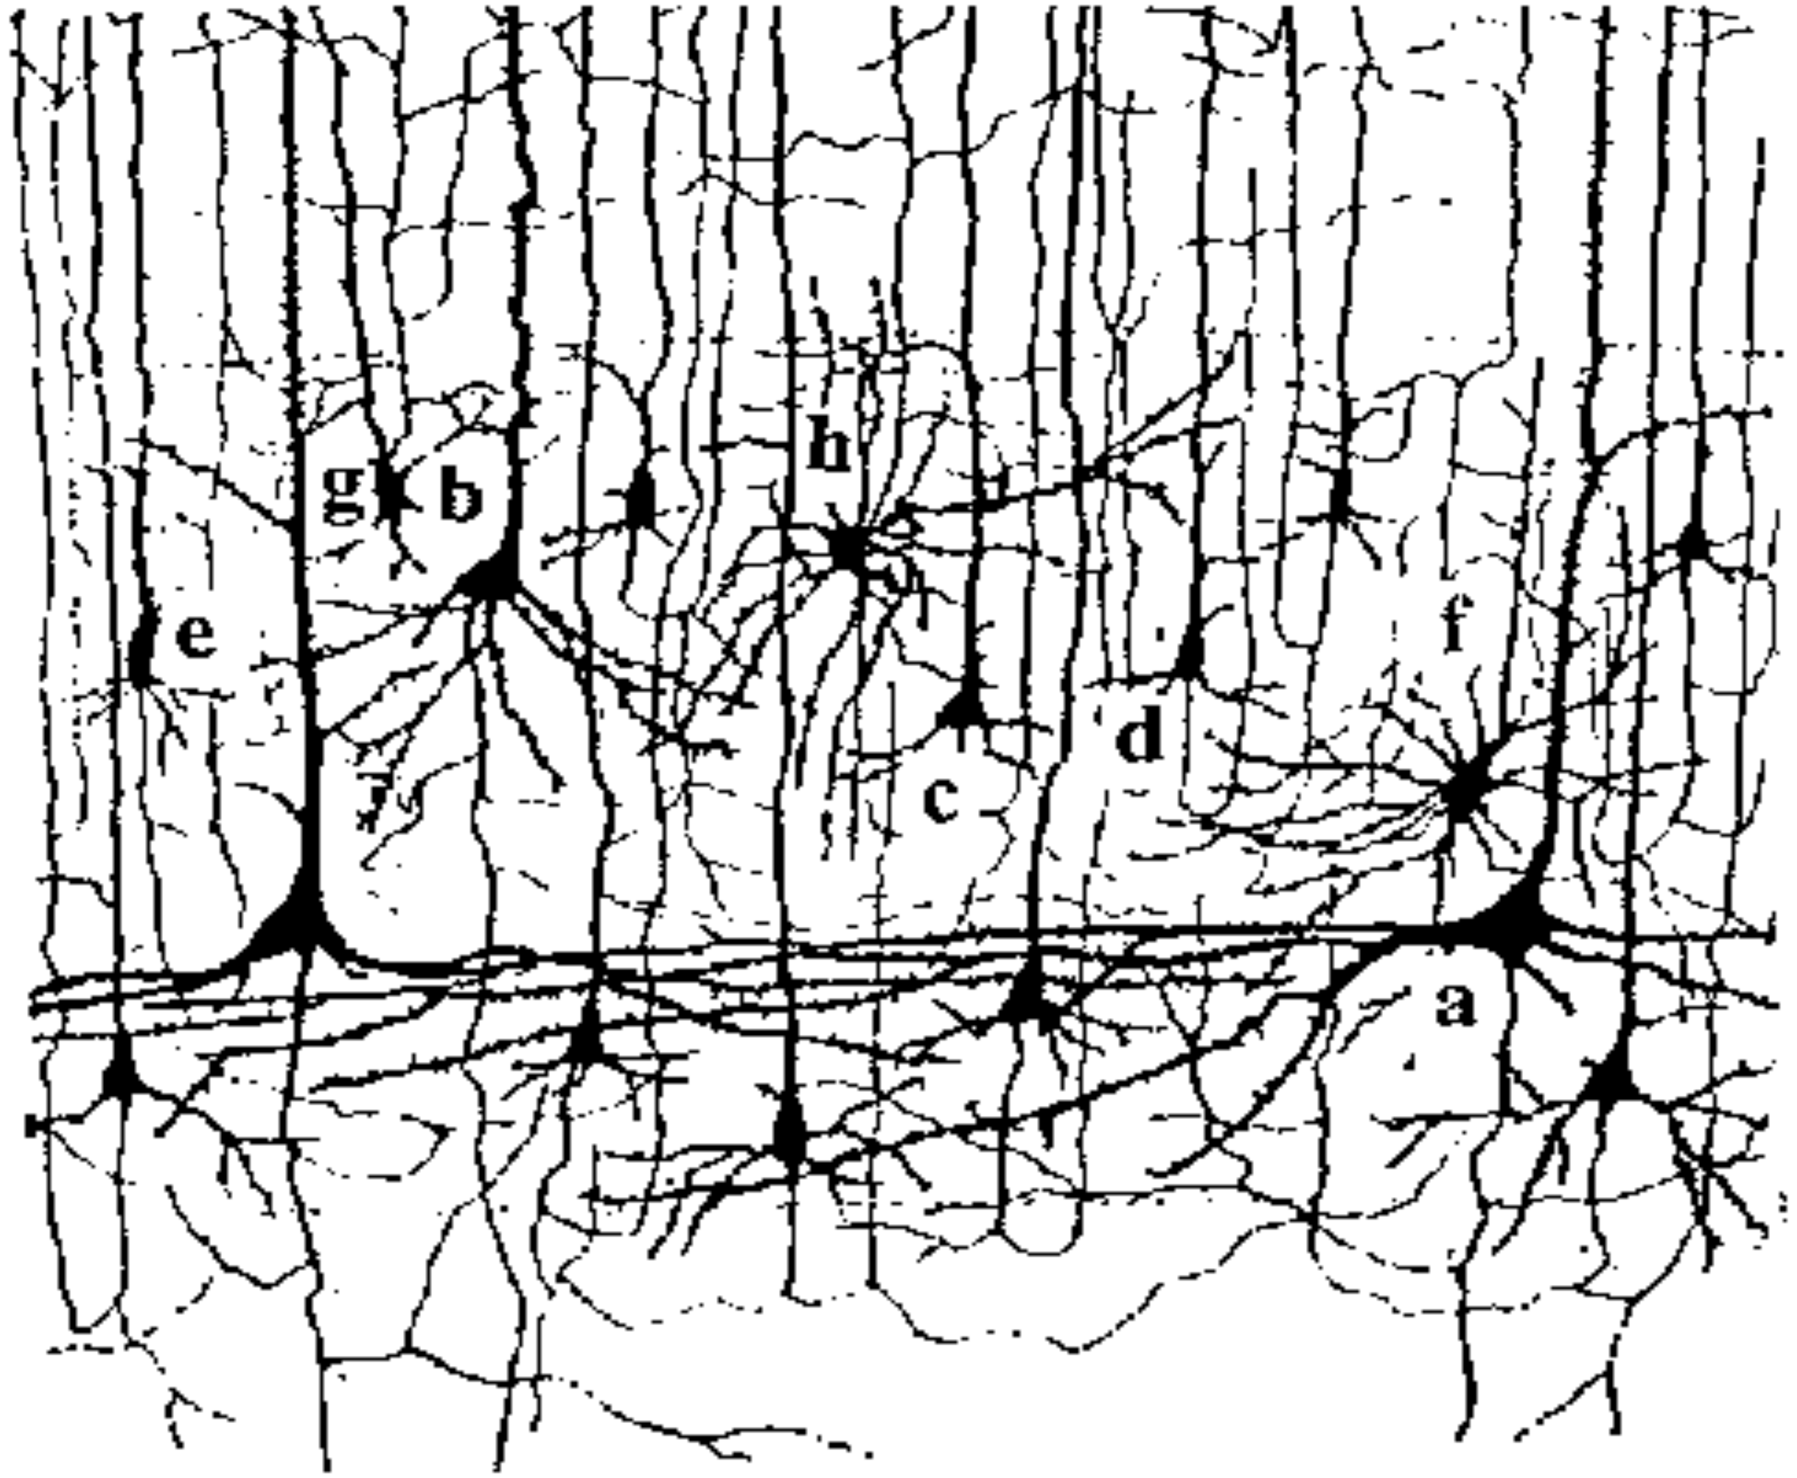

1. Neuronal Networks in the Brain

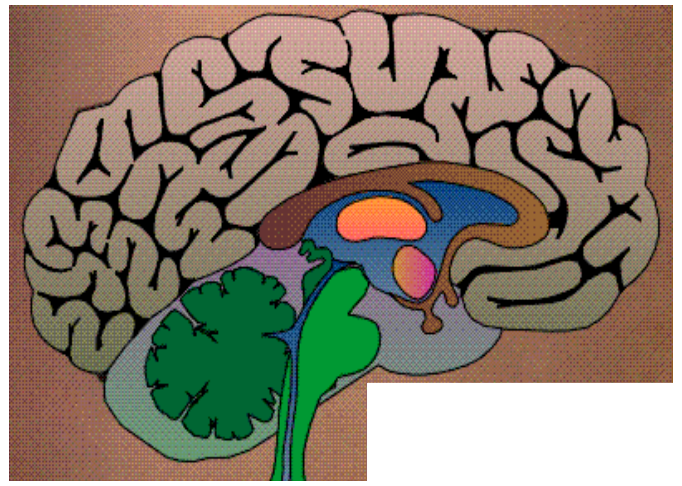

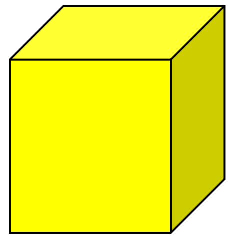

1mm

10 000 neurons
3 km of wire

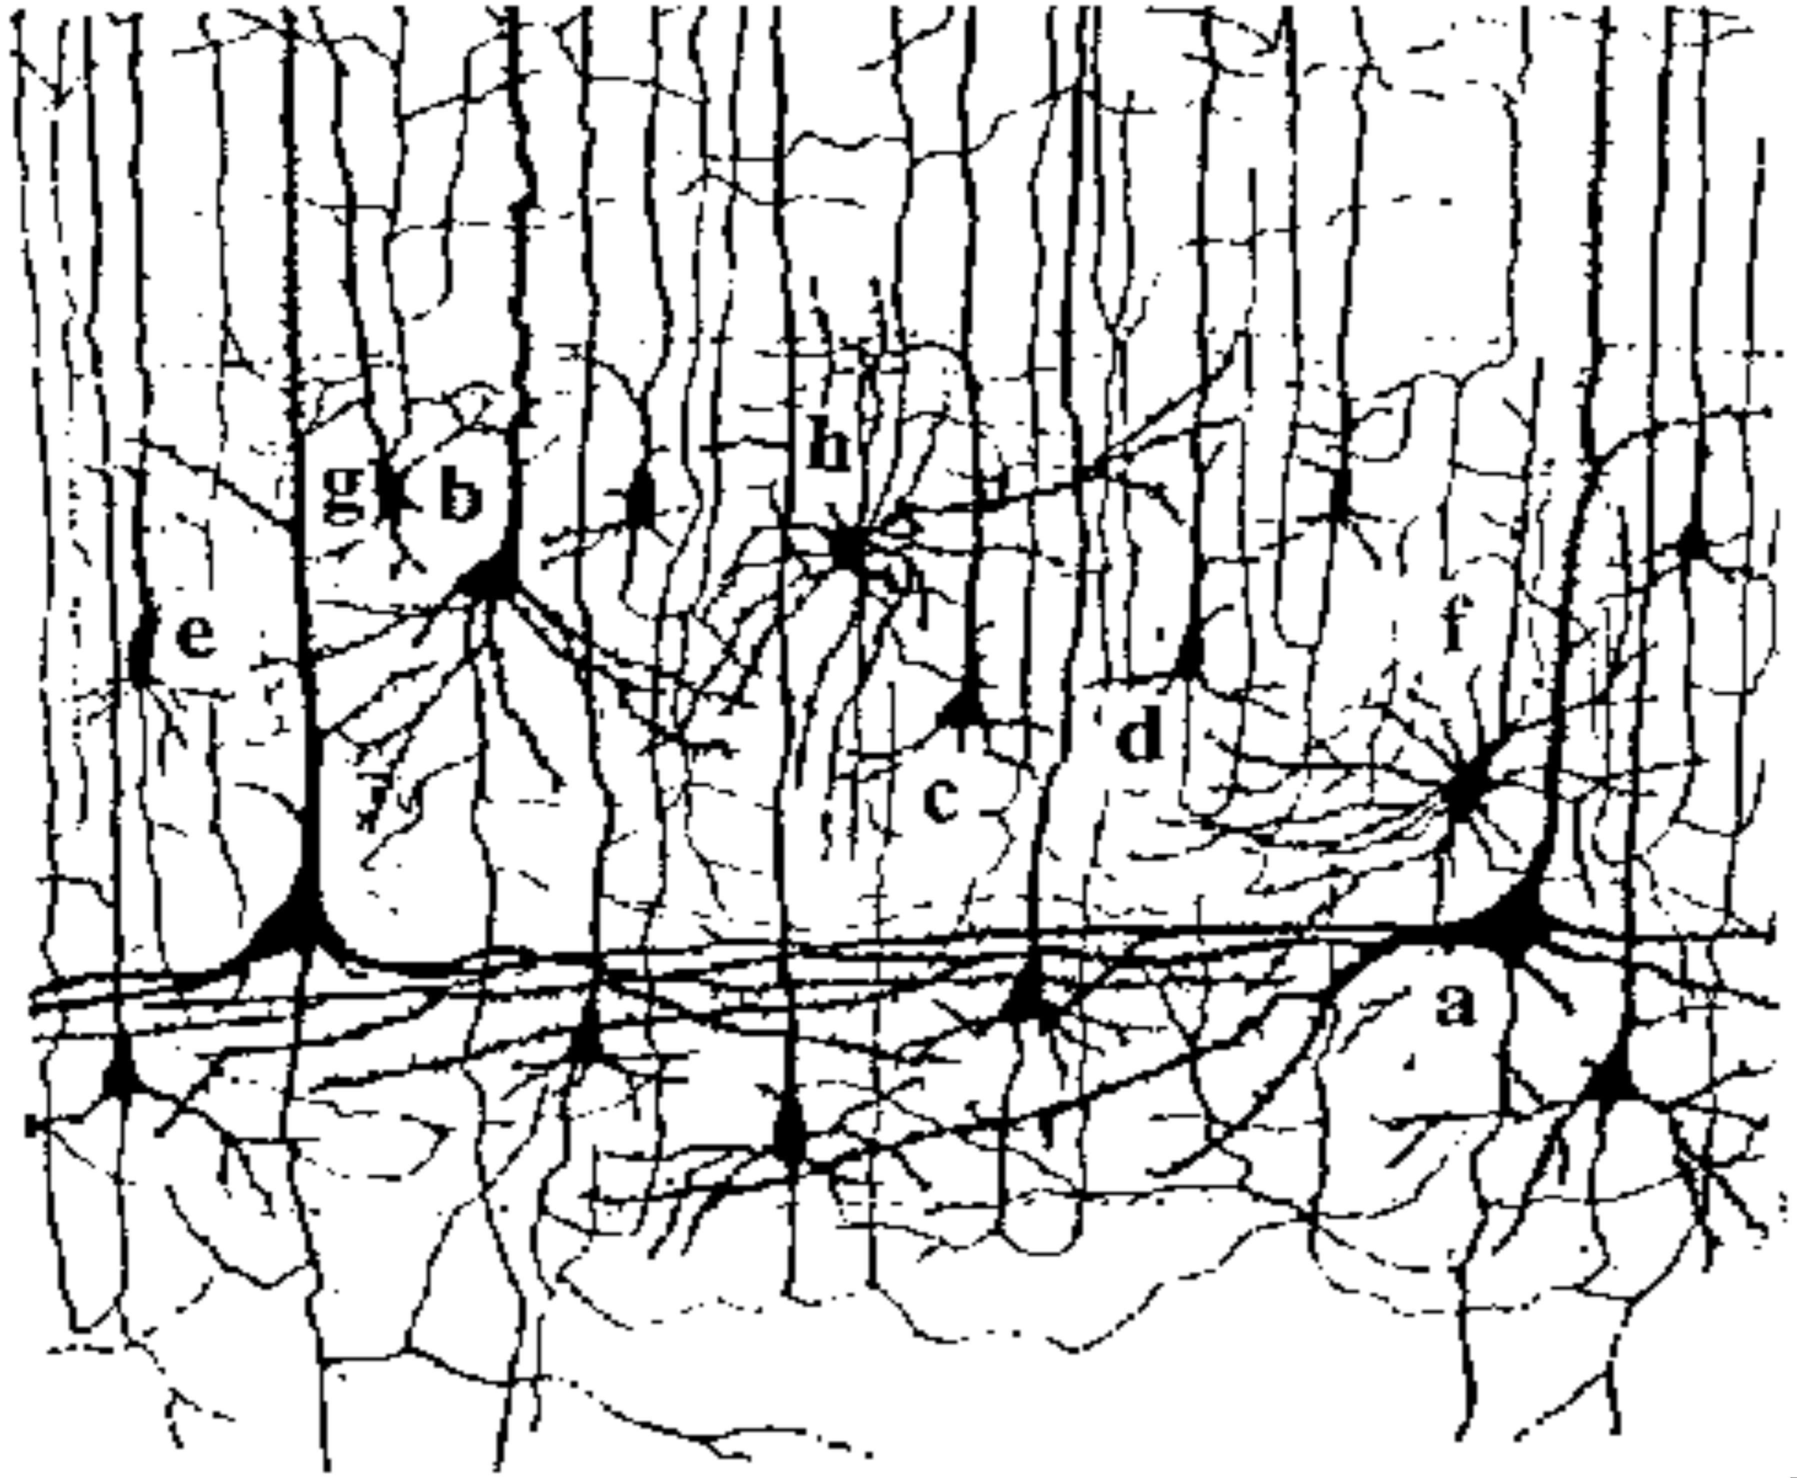

1. Systems for computing and information processing

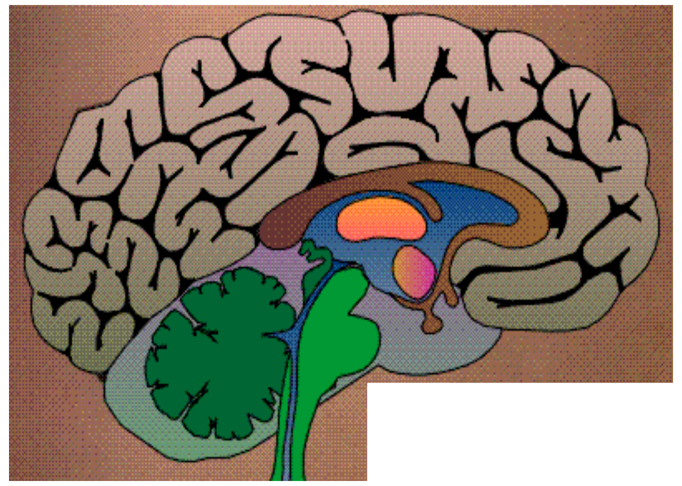

Brain

Computer

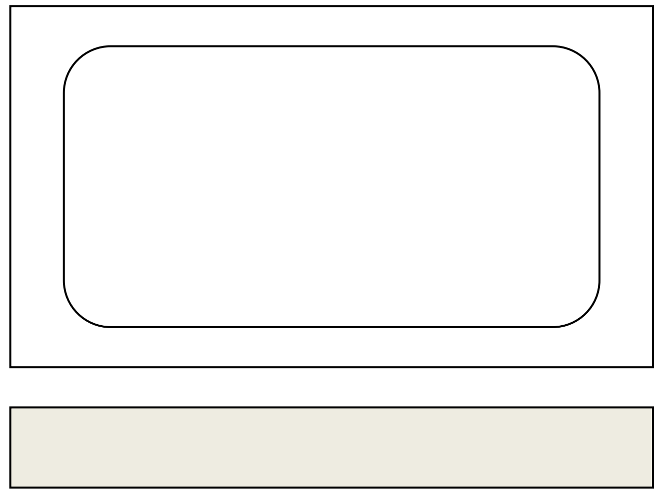

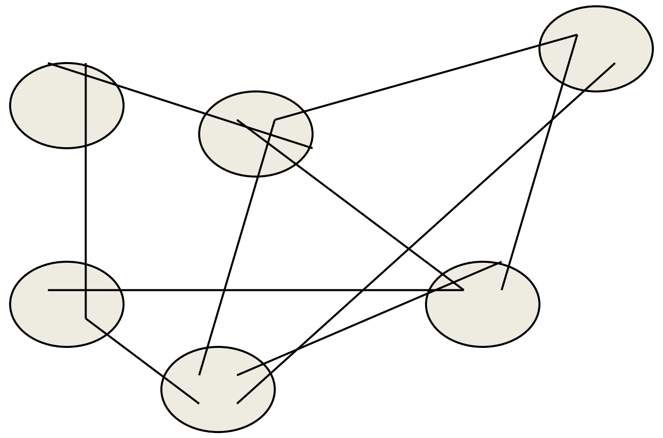

Distributed architecture

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

No separation of processing and memory

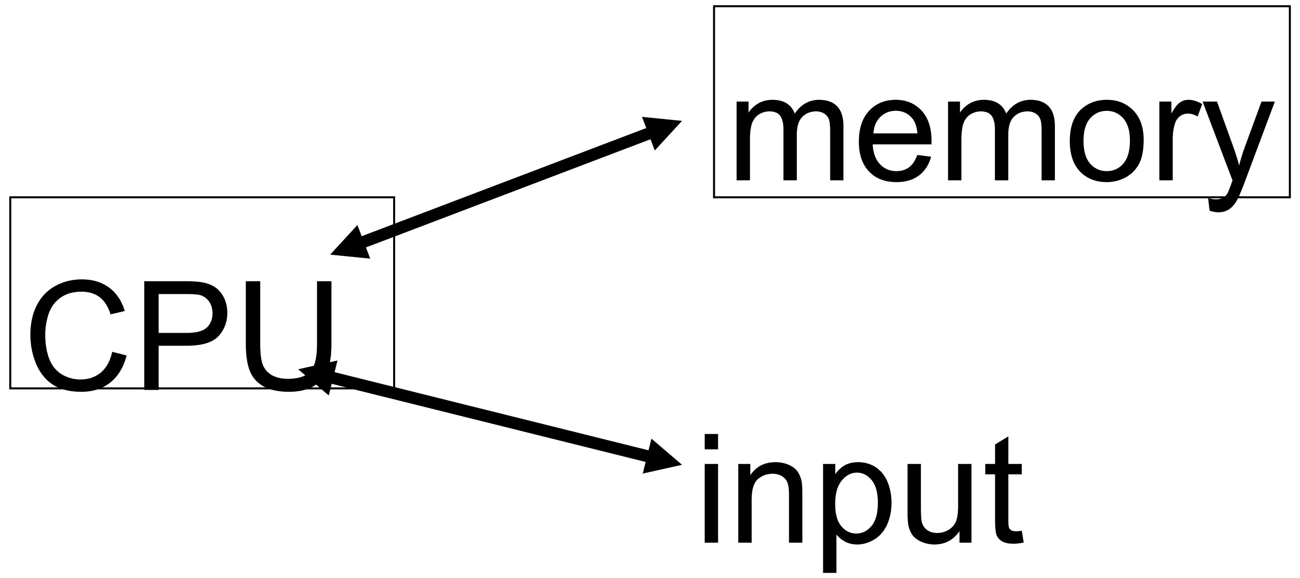

Von Neumann architecture
1 CPU
(10^{10} transistors)

1. Systems for computing and information processing

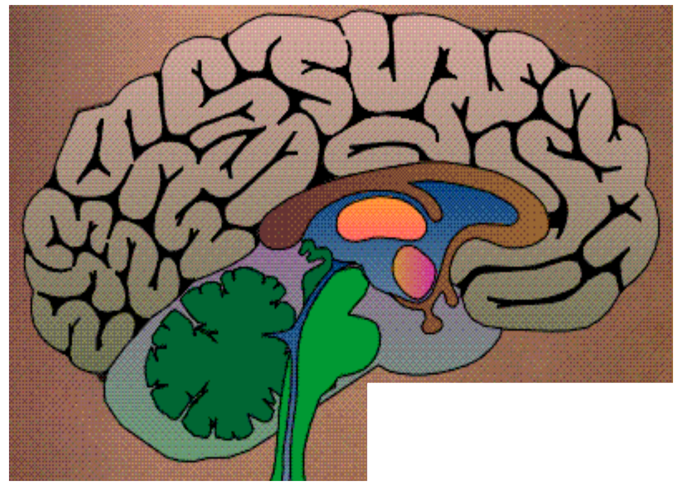

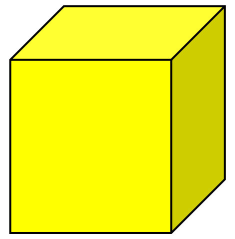

1mm

10 000 neurons
3 km of wire

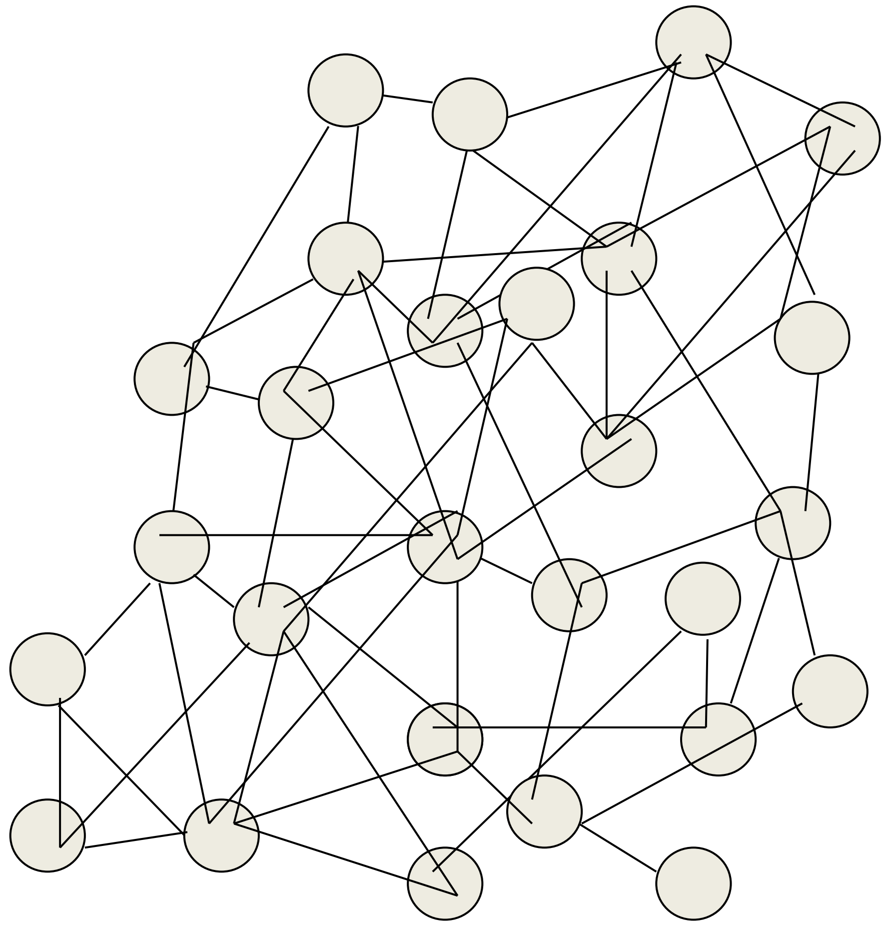

Distributed architecture

10^{10} neurons

10^{4} connections/neurons

No separation of processing and memory

1. Associations, Associative memory

Read this text ***NOW!***

pattern completion/word recognition

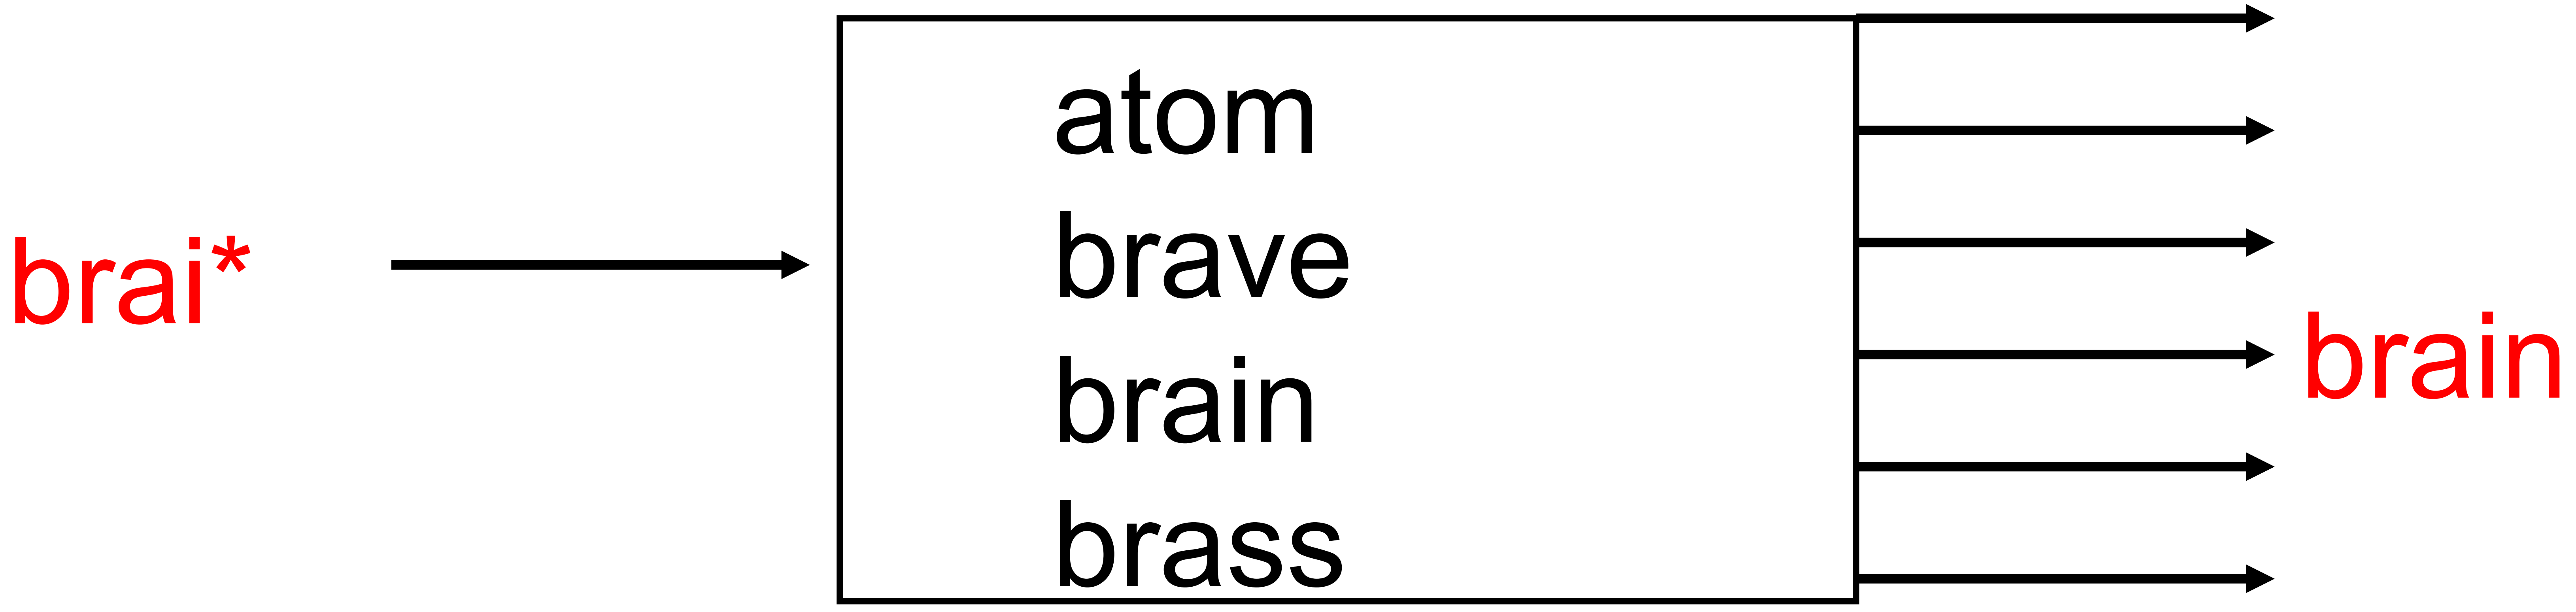

Noisy word

List of words

Output the closest one

Your brain fills in missing information: 'auto-associative memory'

1. Associations, Associative memory

brai* → brain   ***'auto-associative memory'***

bird → swan

vacation → beach   ***'associative memory'***

Quiz 1: Connectivity and Associations

Tick one or several answers

A typical neuron in the brain makes connections

[ ] To 6-30 neighbors

[ ] To 100-500 neurons nearby

[ ] To more than 1000 neurons nearby

[ ] To more than 1000 neurons nearby or far away.

Associative memory is involved

[ ] If you think of palm trees when you think of a beach

[ ] If partial information helps you to recall a complicated concept

[ ] If a cue helps you to recall a memory

Computational Neuroscience: Neuronal Dynamics of Cognition

ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

A: ASSOCIATIVE MEMORY
in a Network of Neurons

Wulfram Gerstner

EPFL, Lausanne, Switzerland

Reading for this week:
NEURONAL DYNAMICS
- Ch. 17.1 - 17.2.4

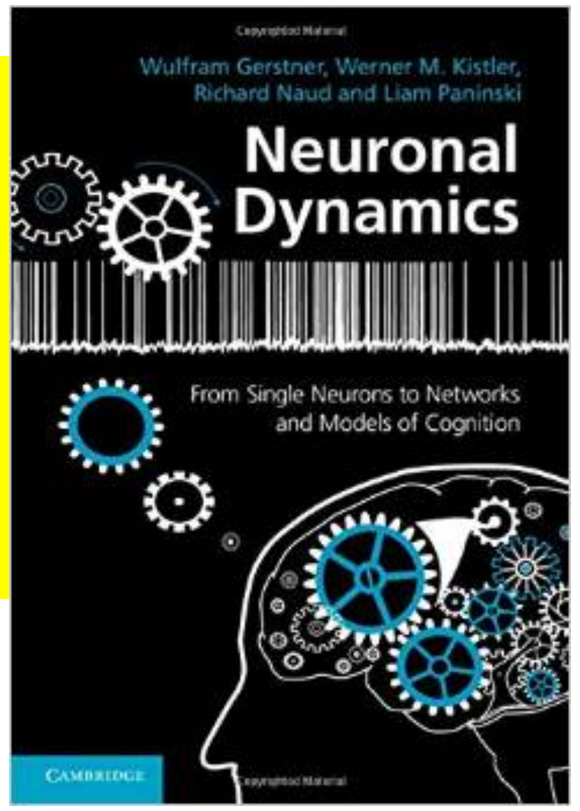

Cambridge Univ. Press

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

2 Classification by similarity

3 Detour: Magnetic Materials

4 Hopfield Model

5 Learning of Associations

6 Storage Capacity

2. Classification by similarity: **pattern recognition**

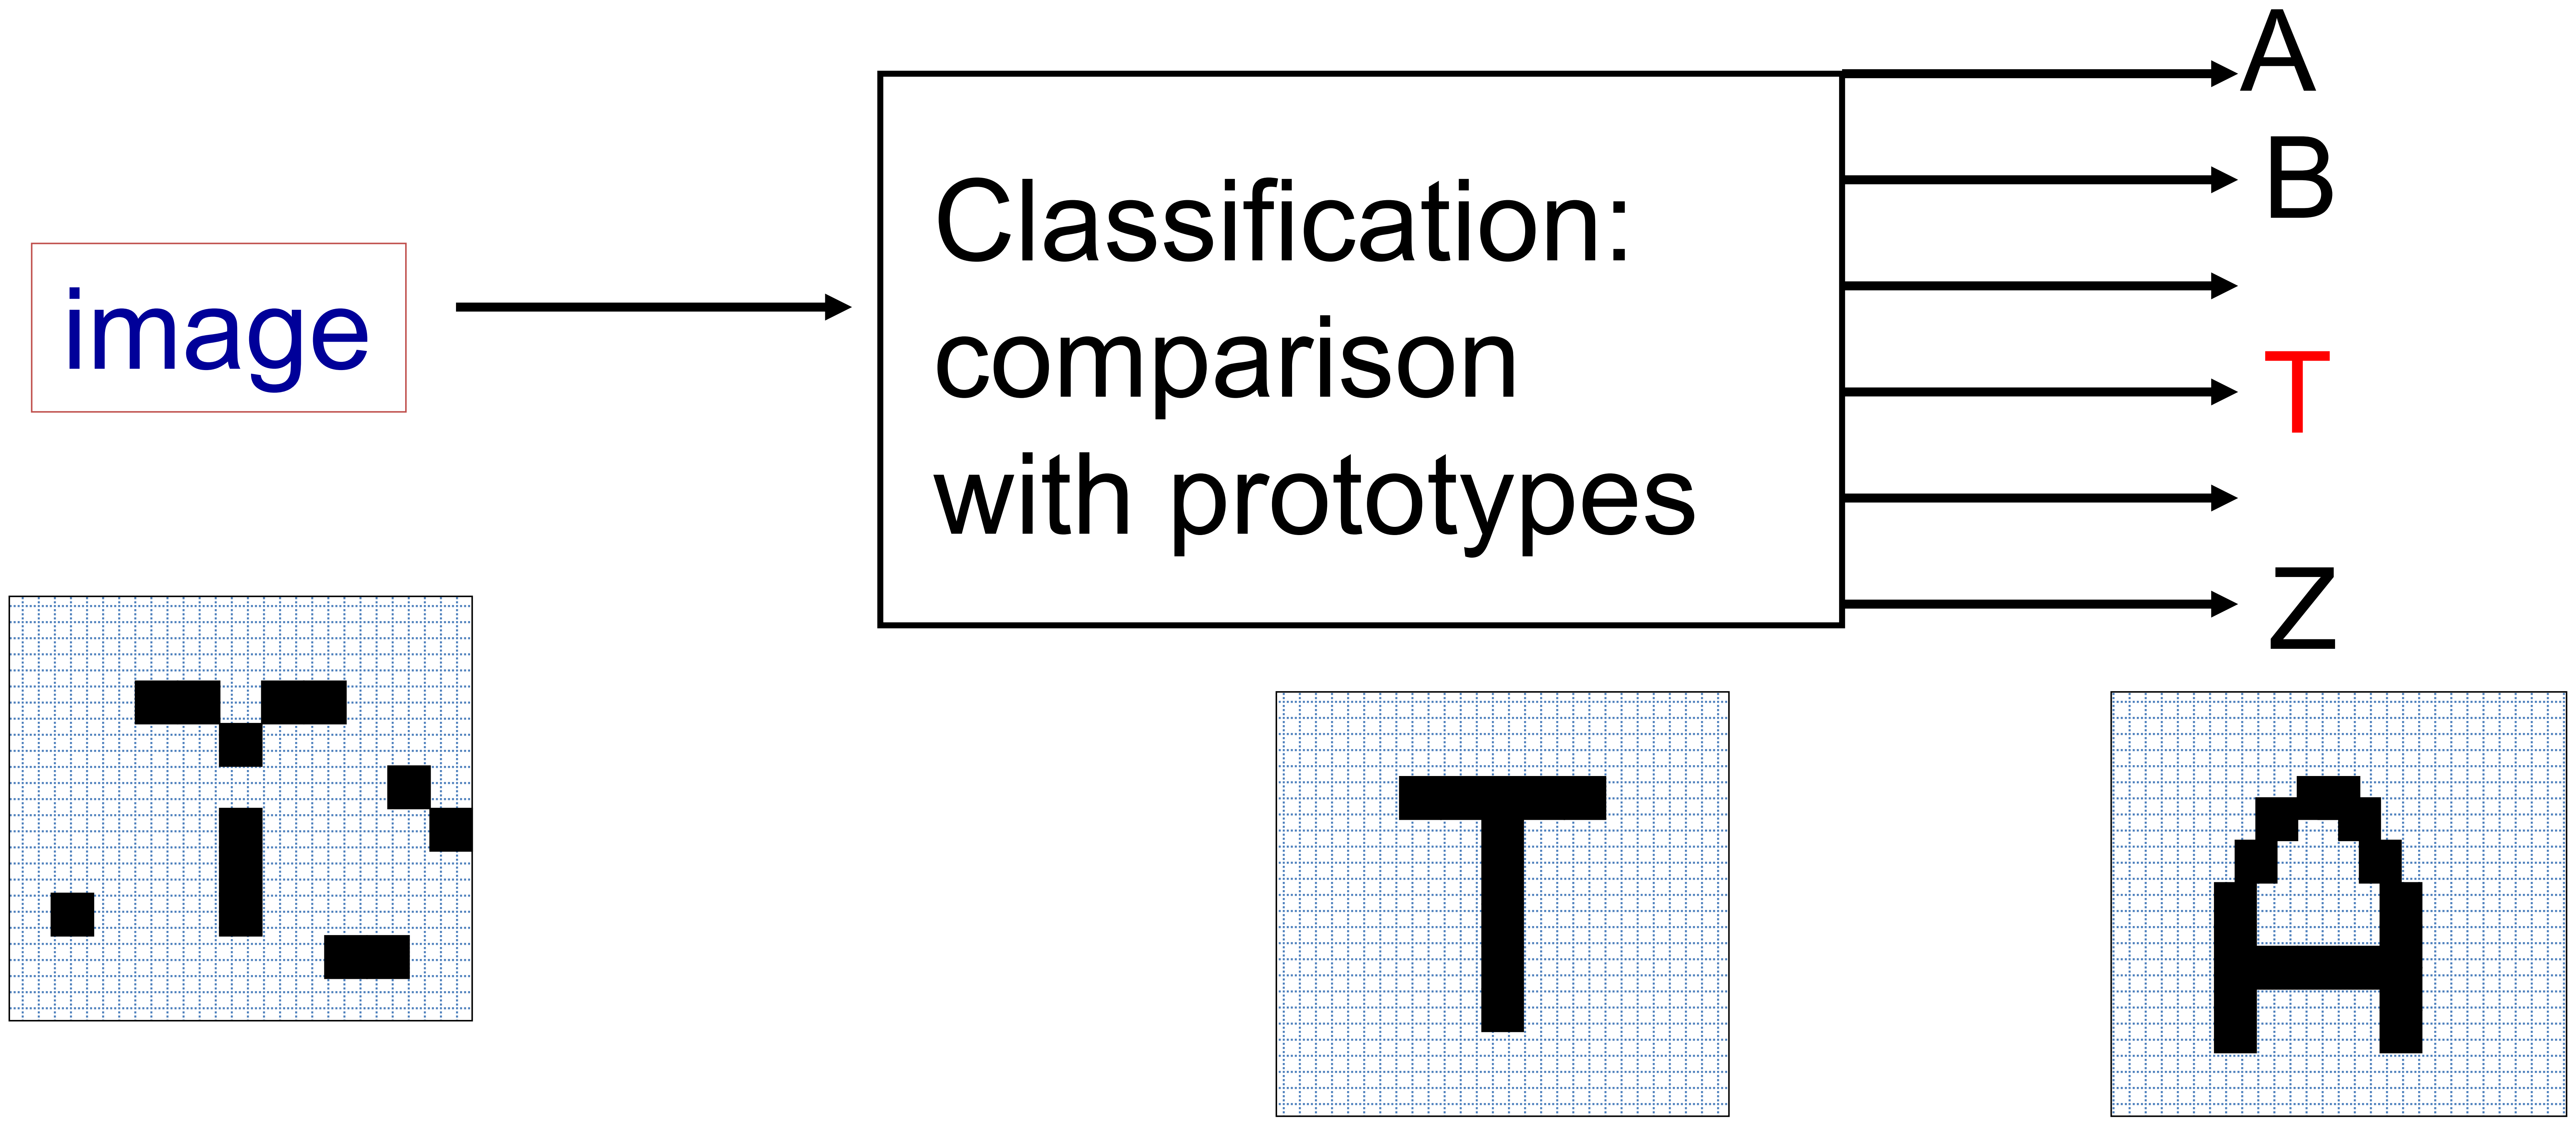

Noisy image

Prototypes

2. Classification by similarity: **pattern recognition**

Classification by closest prototype

x

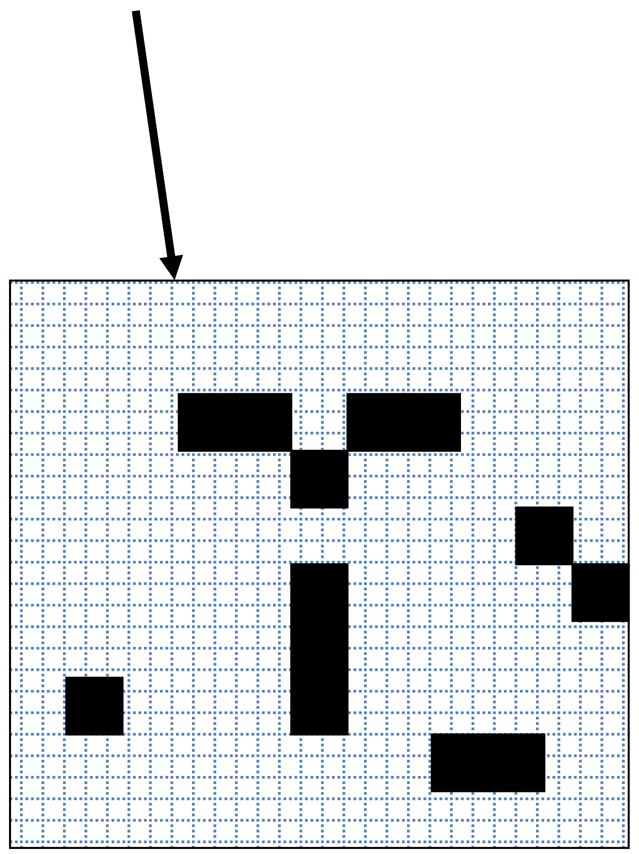

Noisy image

p^T

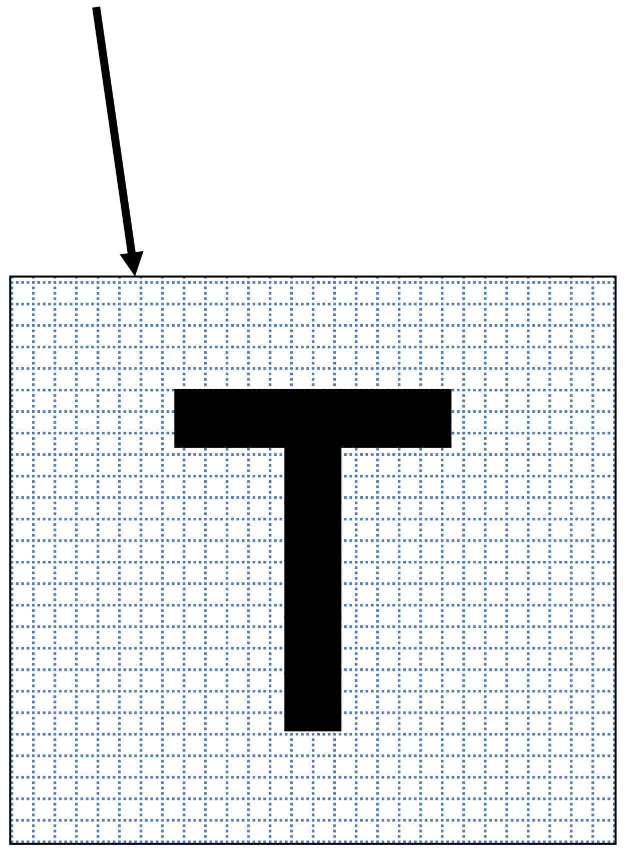

Prototype

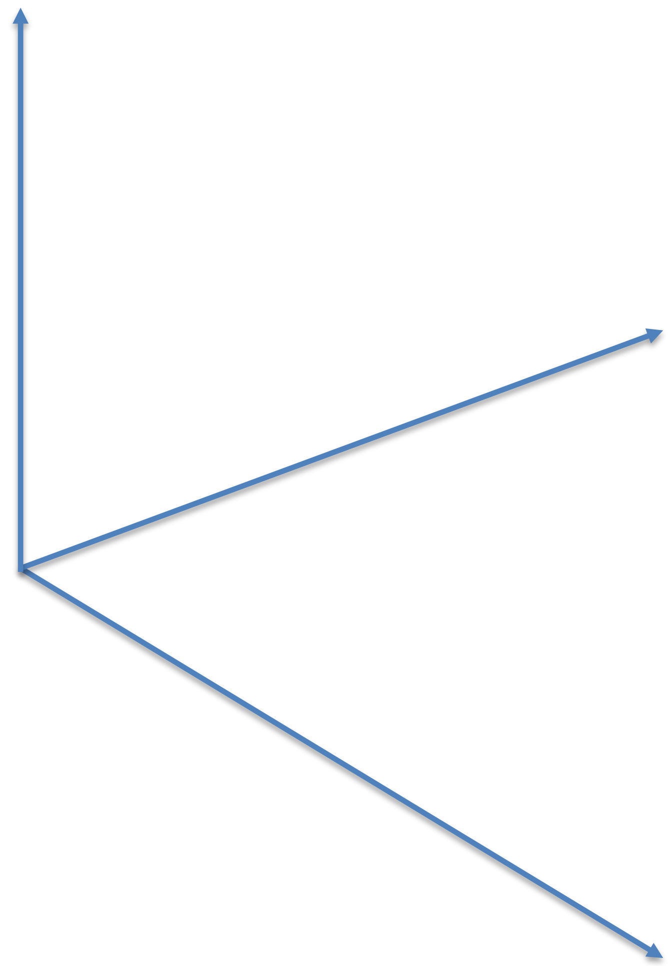

2. Classification by similarity: **pattern recognition**

Classification by closest prototype

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

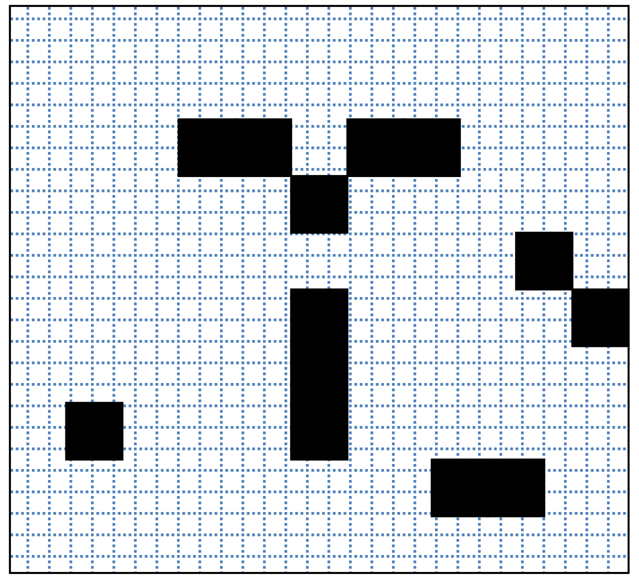

Noisy image

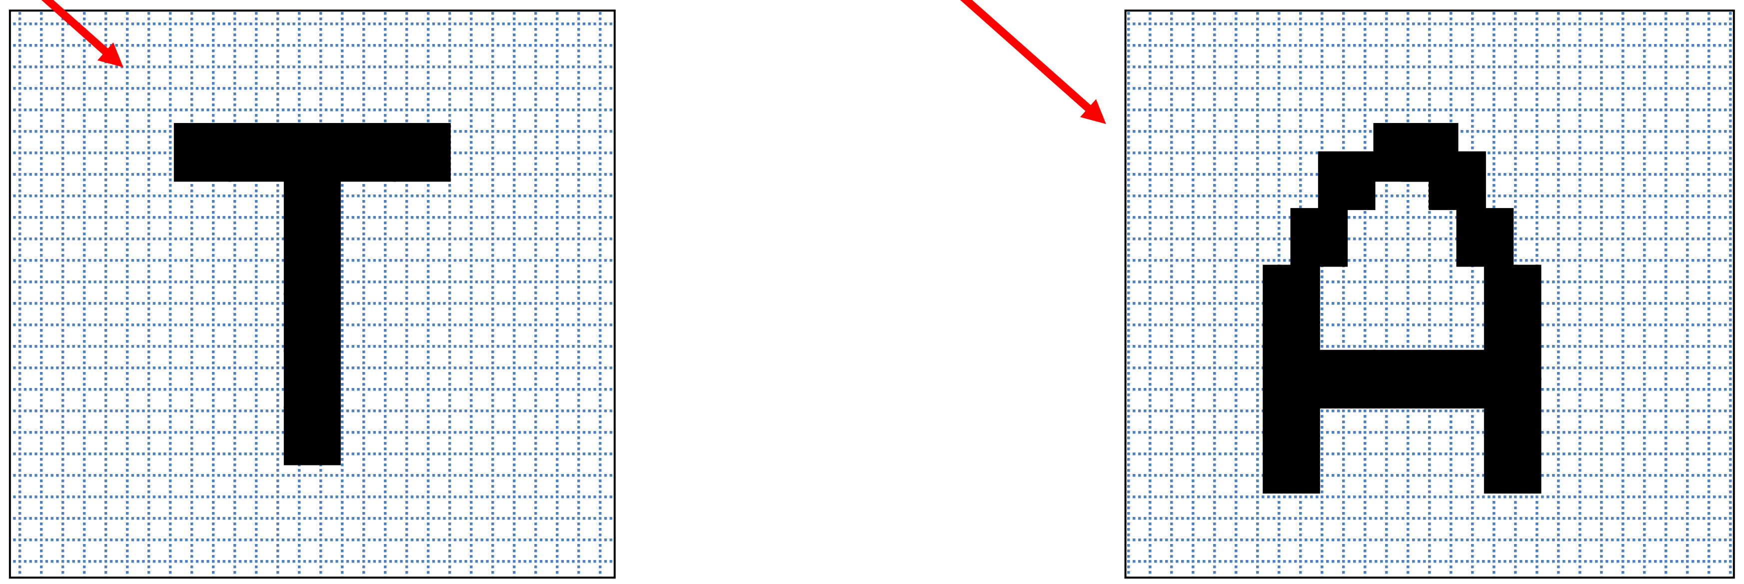

Prototypes

2. pattern recognition and Pattern completion

Aim: Understand Associative Memory

Noisy image

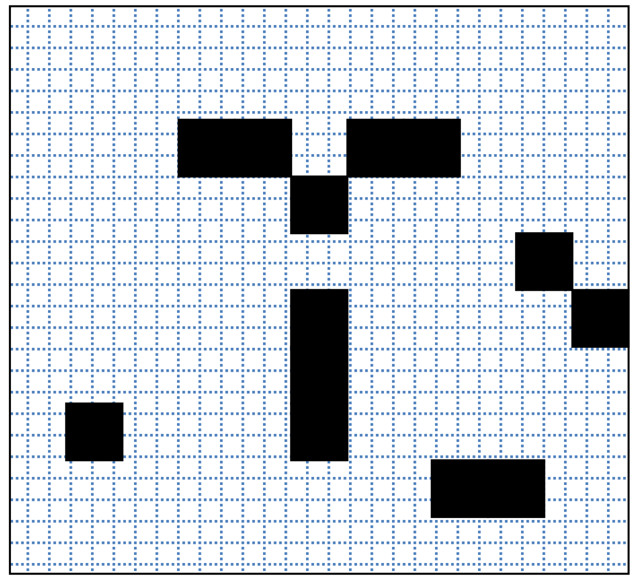

Partial word

Associative memory/ collective computation

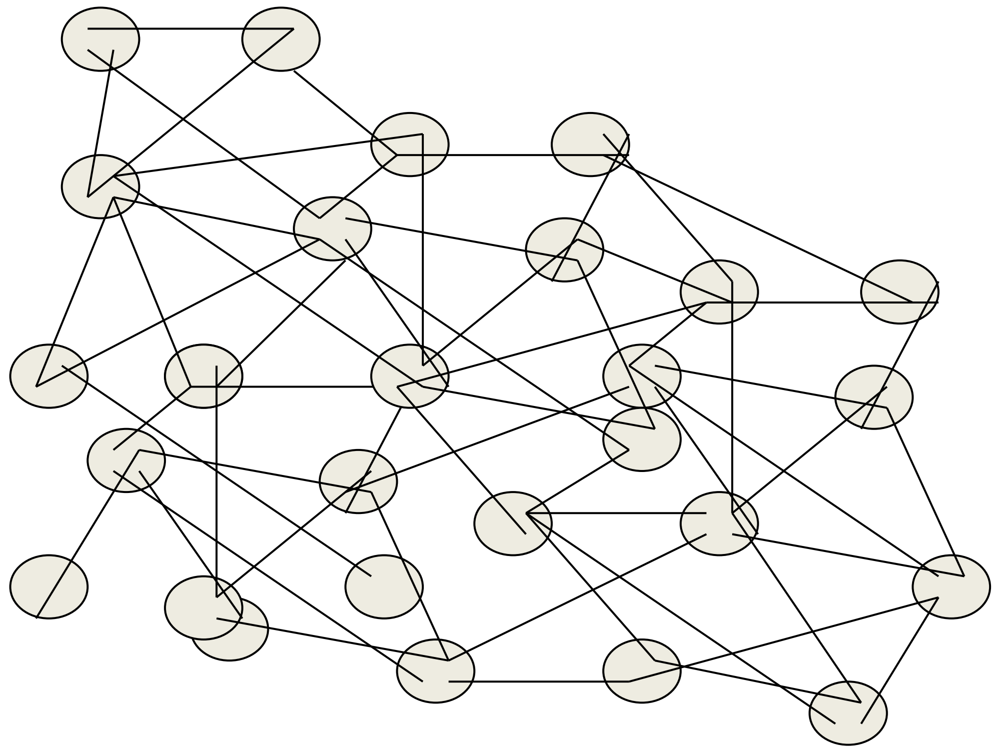

Brain-style computation

Full image

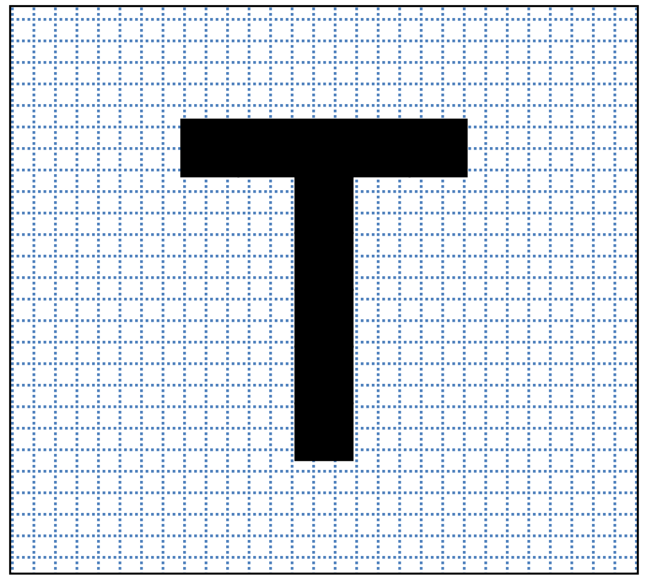

Full word

Quiz 2: Closest prototype

Classification by closest prototype (tick one or several answers)
[ ] Needs a similarity measure
[ ] Needs a distance measure
[ ] Needs a method to find the maximum or minimum

Computational Neuroscience: Neuronal Dynamics of Cognition

ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

A: ASSOCIATIVE MEMORY
in a Network of Neurons

Wulfram Gerstner

EPFL, Lausanne, Switzerland

Reading for this week:
NEURONAL DYNAMICS
- Ch. 17.1 - 17.2.4

Cambridge Univ. Press

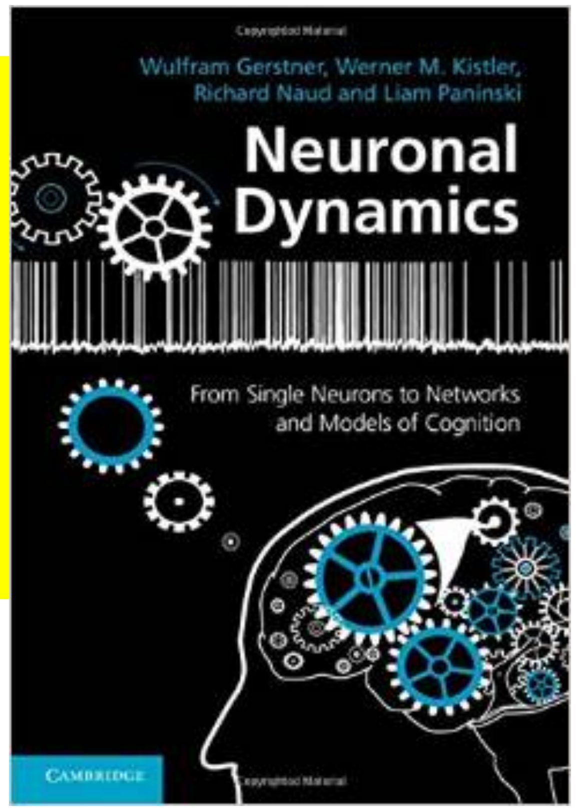

1 Introduction
- networks of neuron
- systems for computing
- associative memory

2 Classification by similarity

3 Detour: Magnetic Materials

4 Hopfield Model

5 Learning of Associations

6 Storage Capacity

3. Detour: magnetism

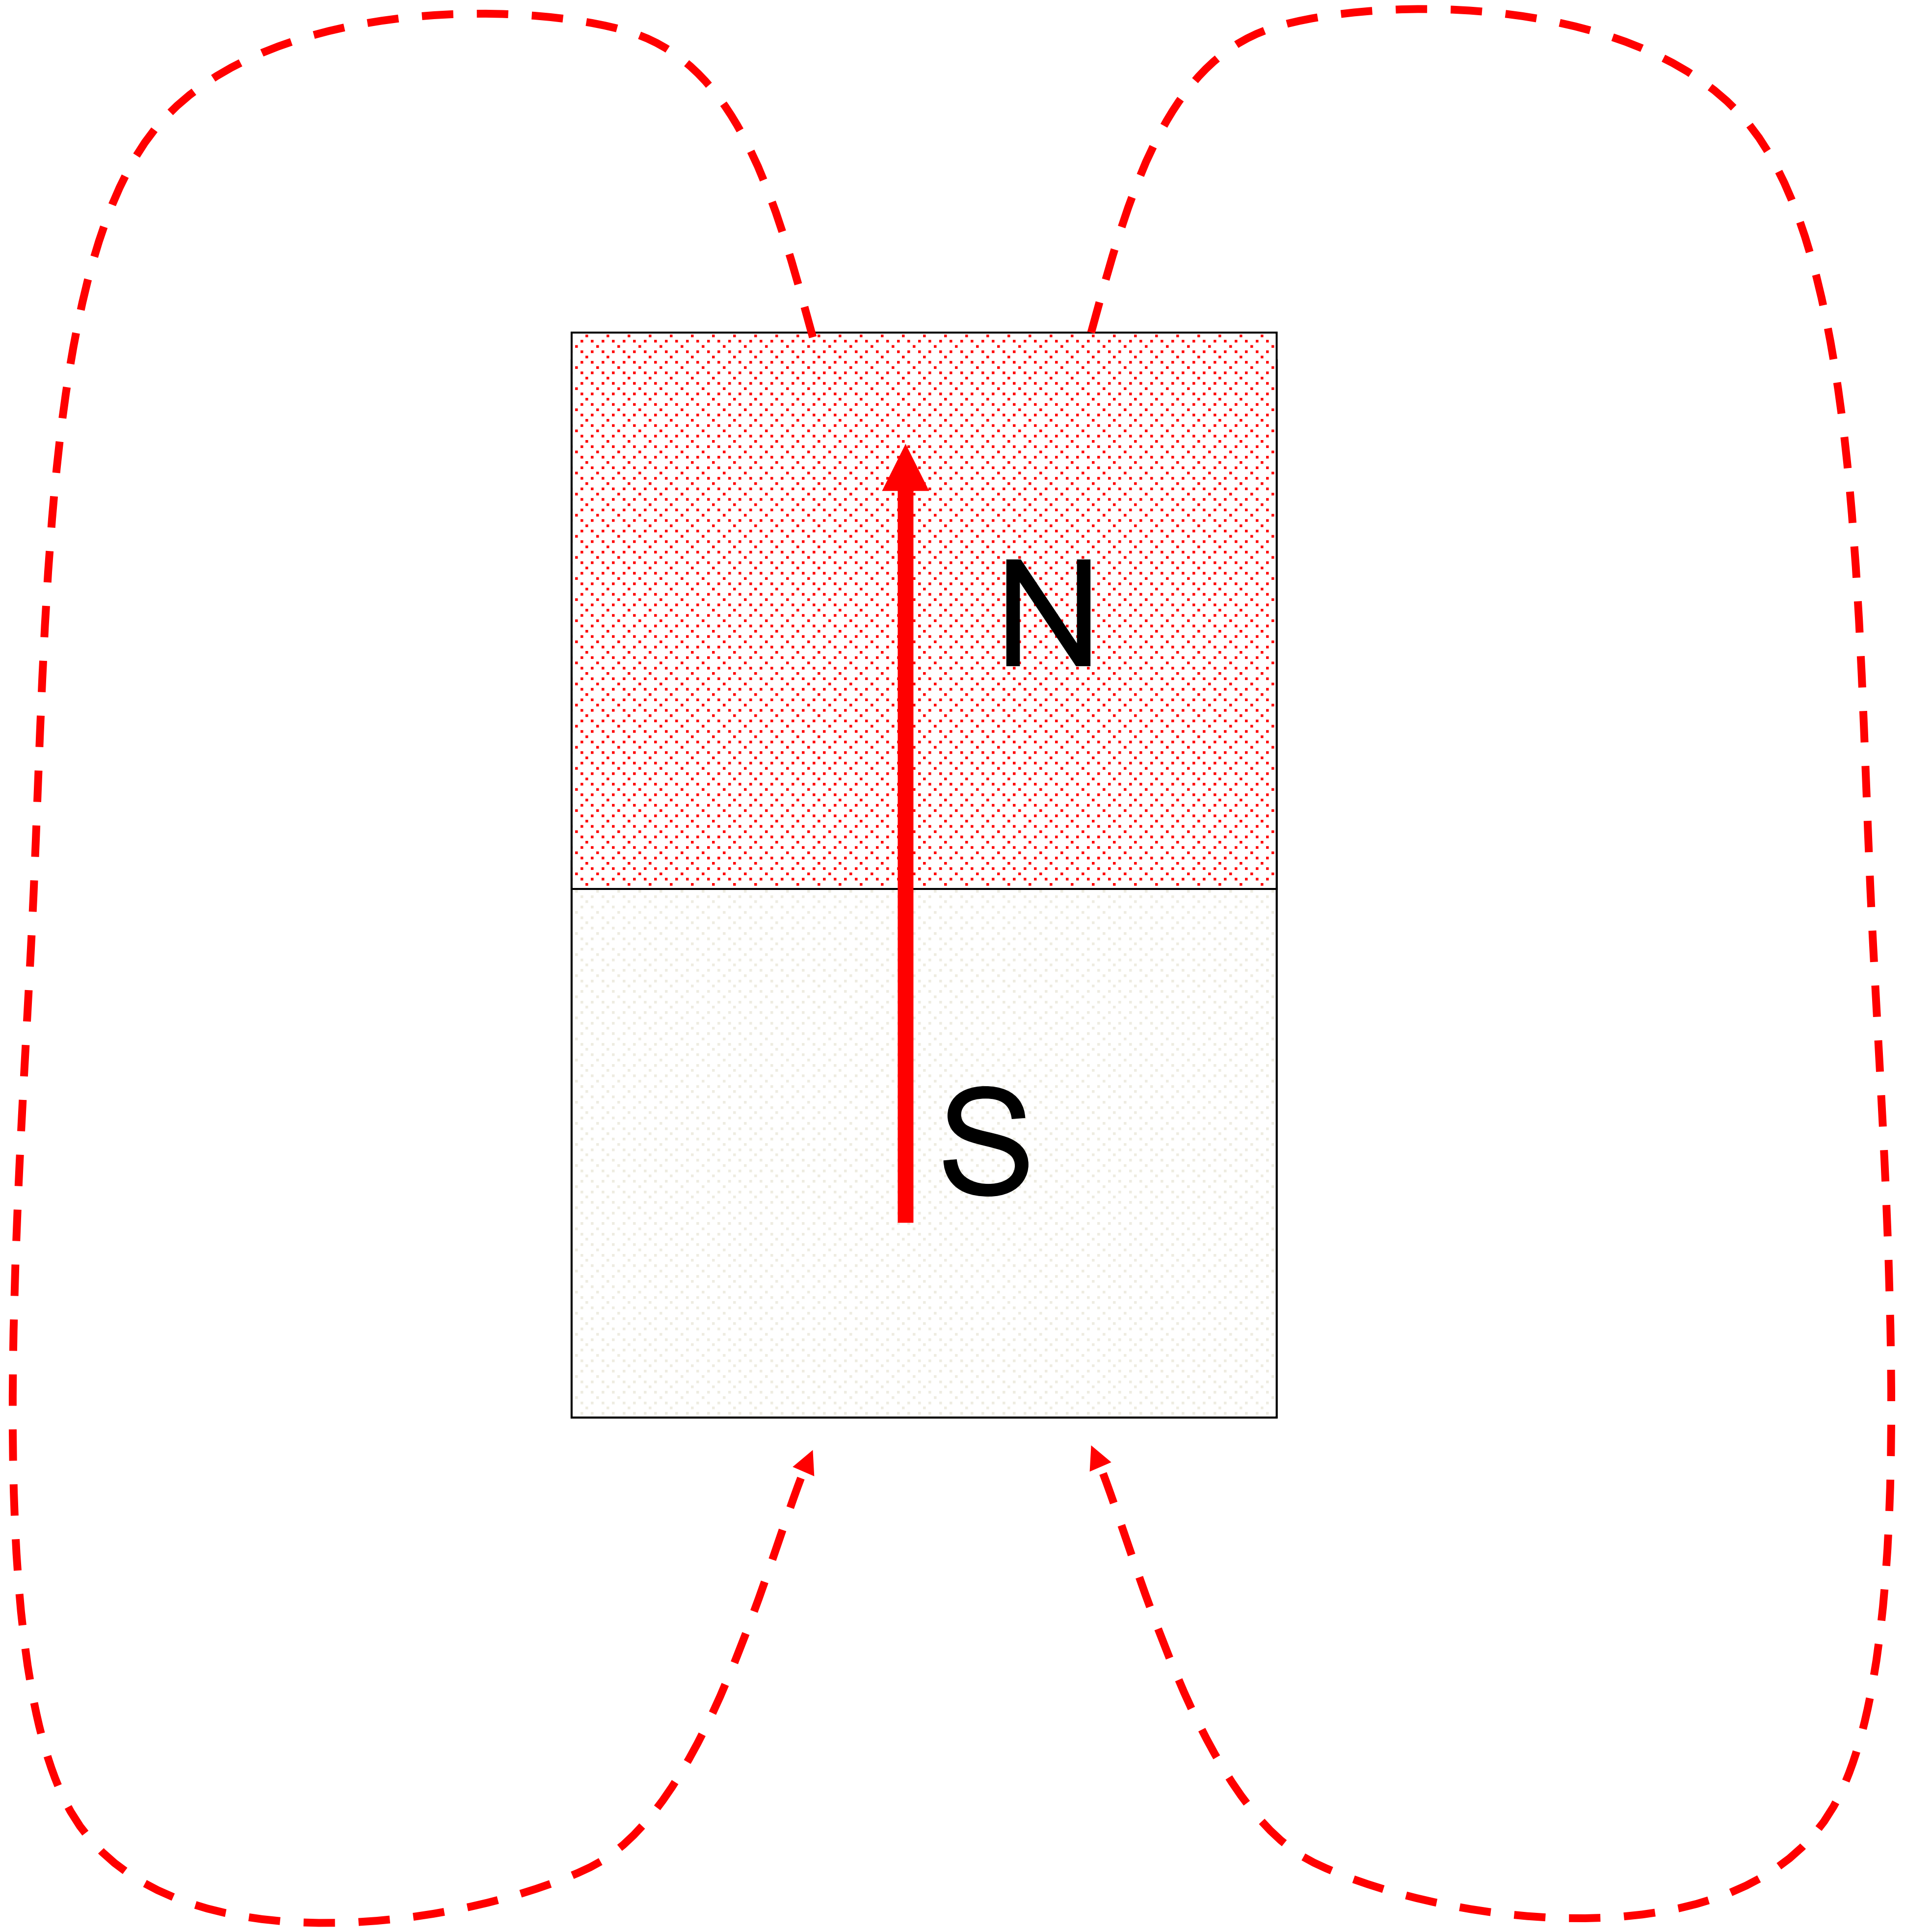

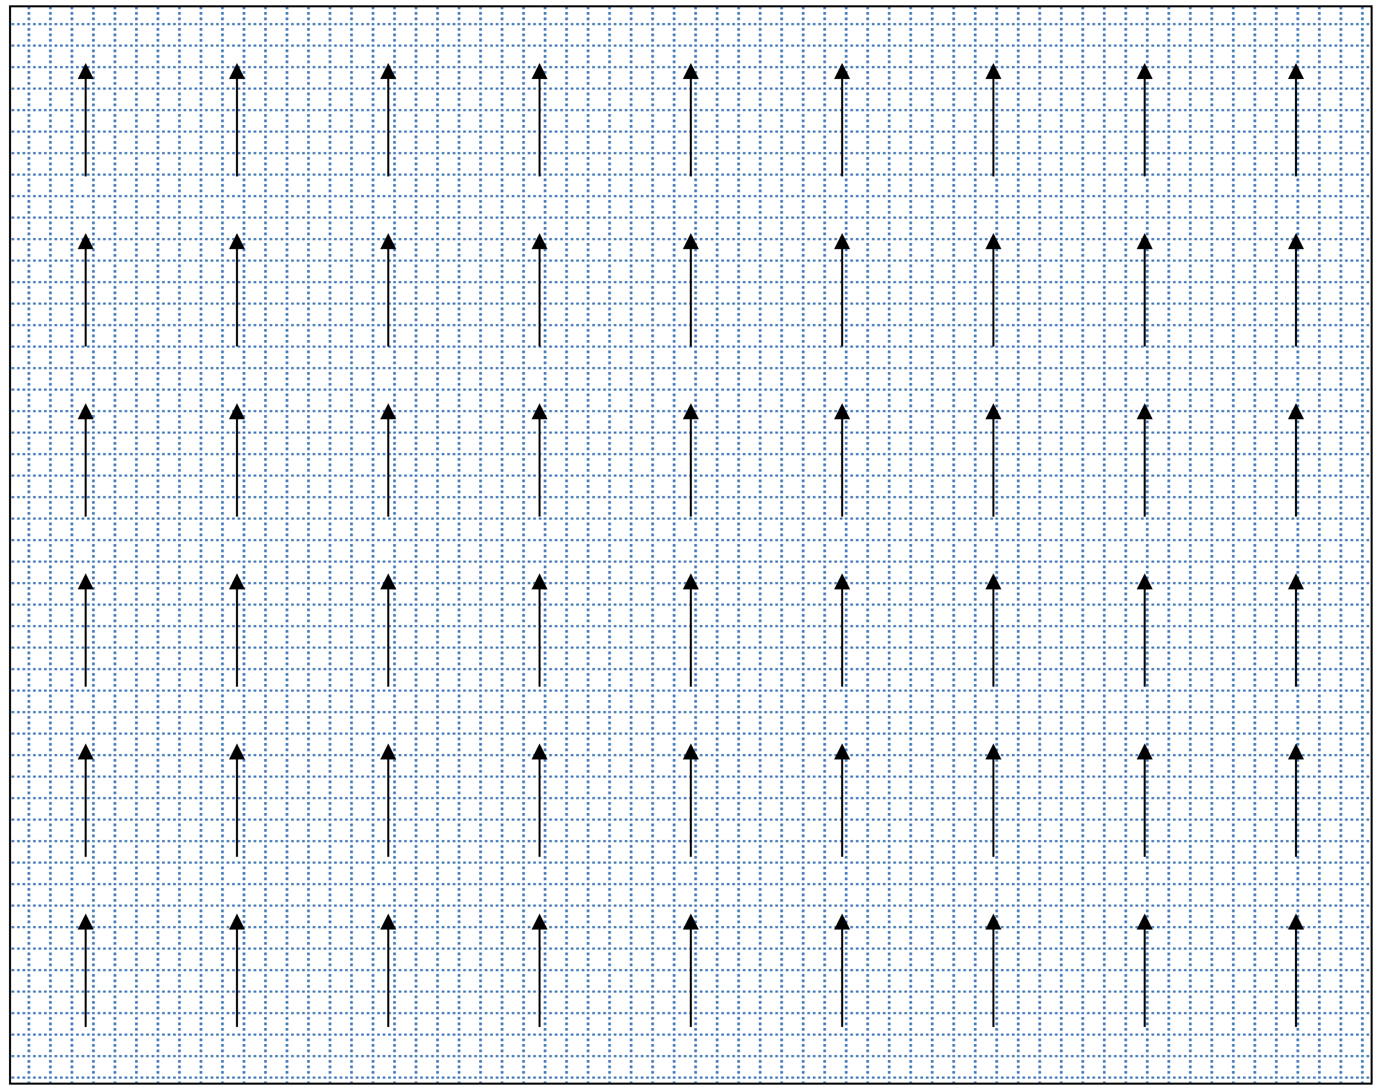

3. Detour: magnetism

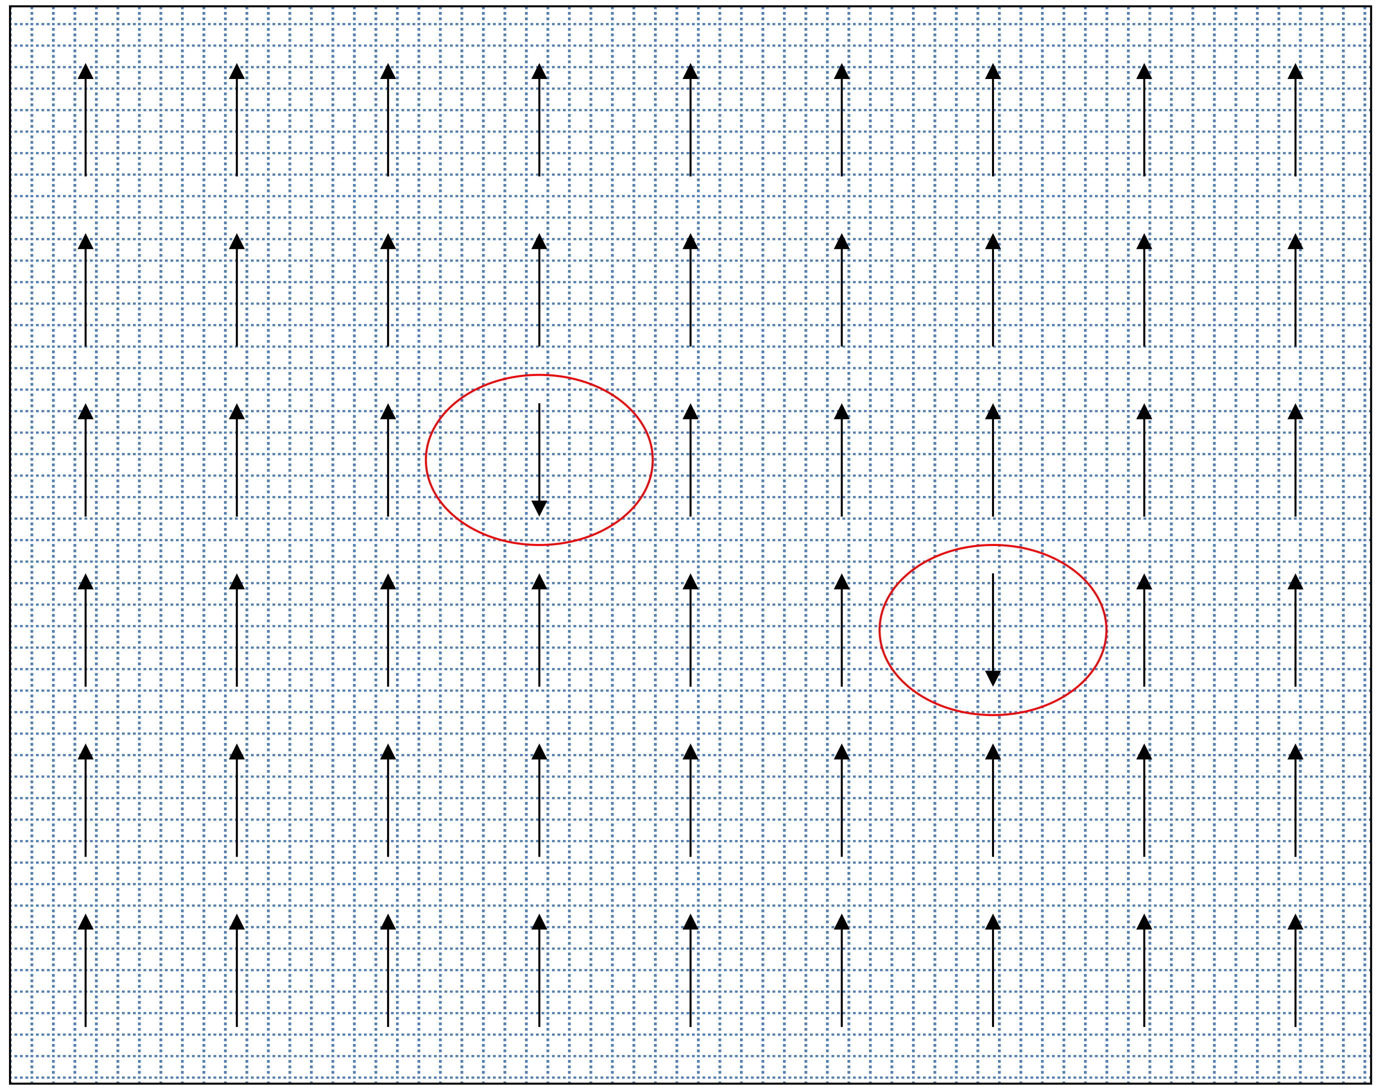

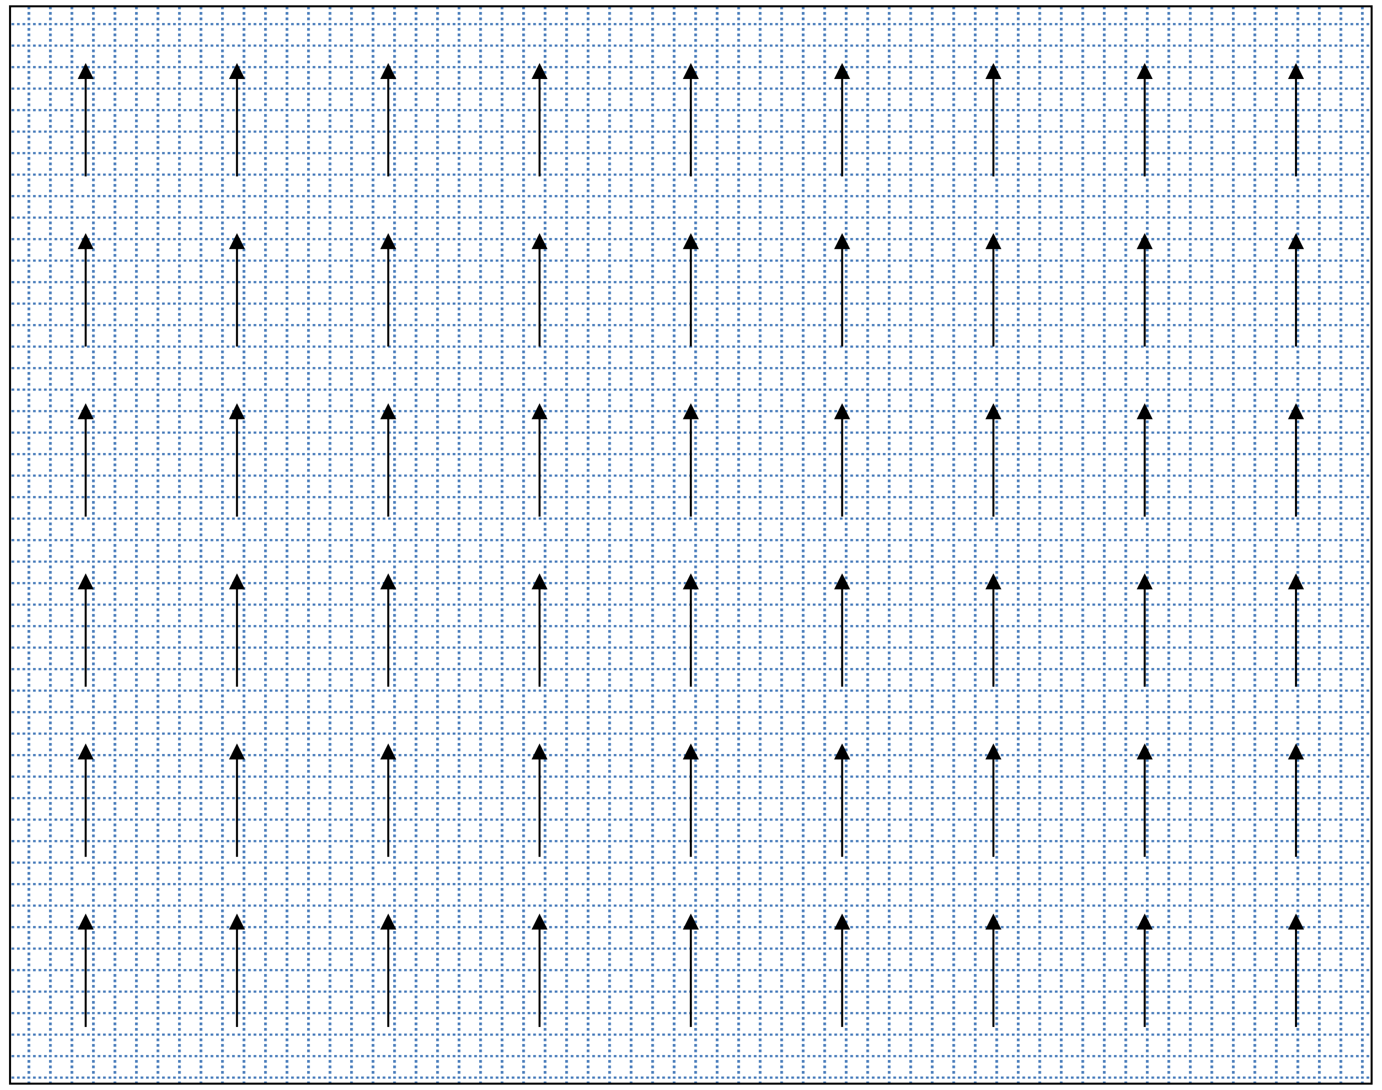

Noisy magnet → pure magnet

3. Detour: magnetism

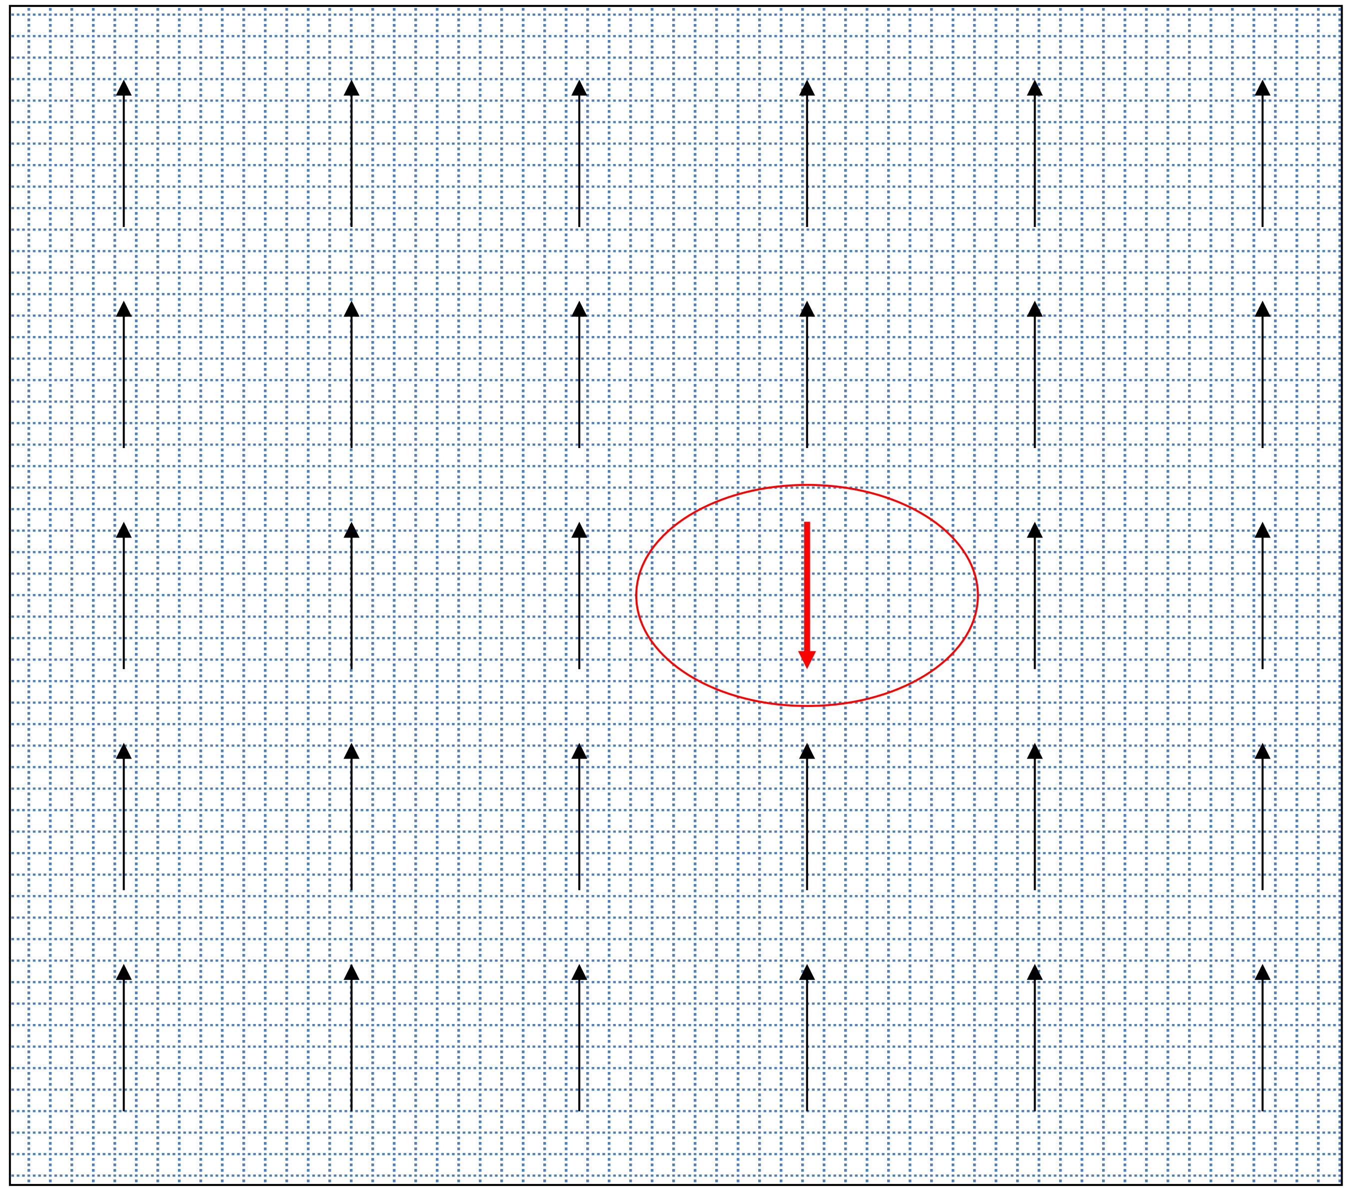

Elementary magnet

↑ $S_i = +1$

↓ $S_i = -1$

dynamics

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

Sum over all interactions with i

3. Detour: magnetism

Anti-ferromagnet

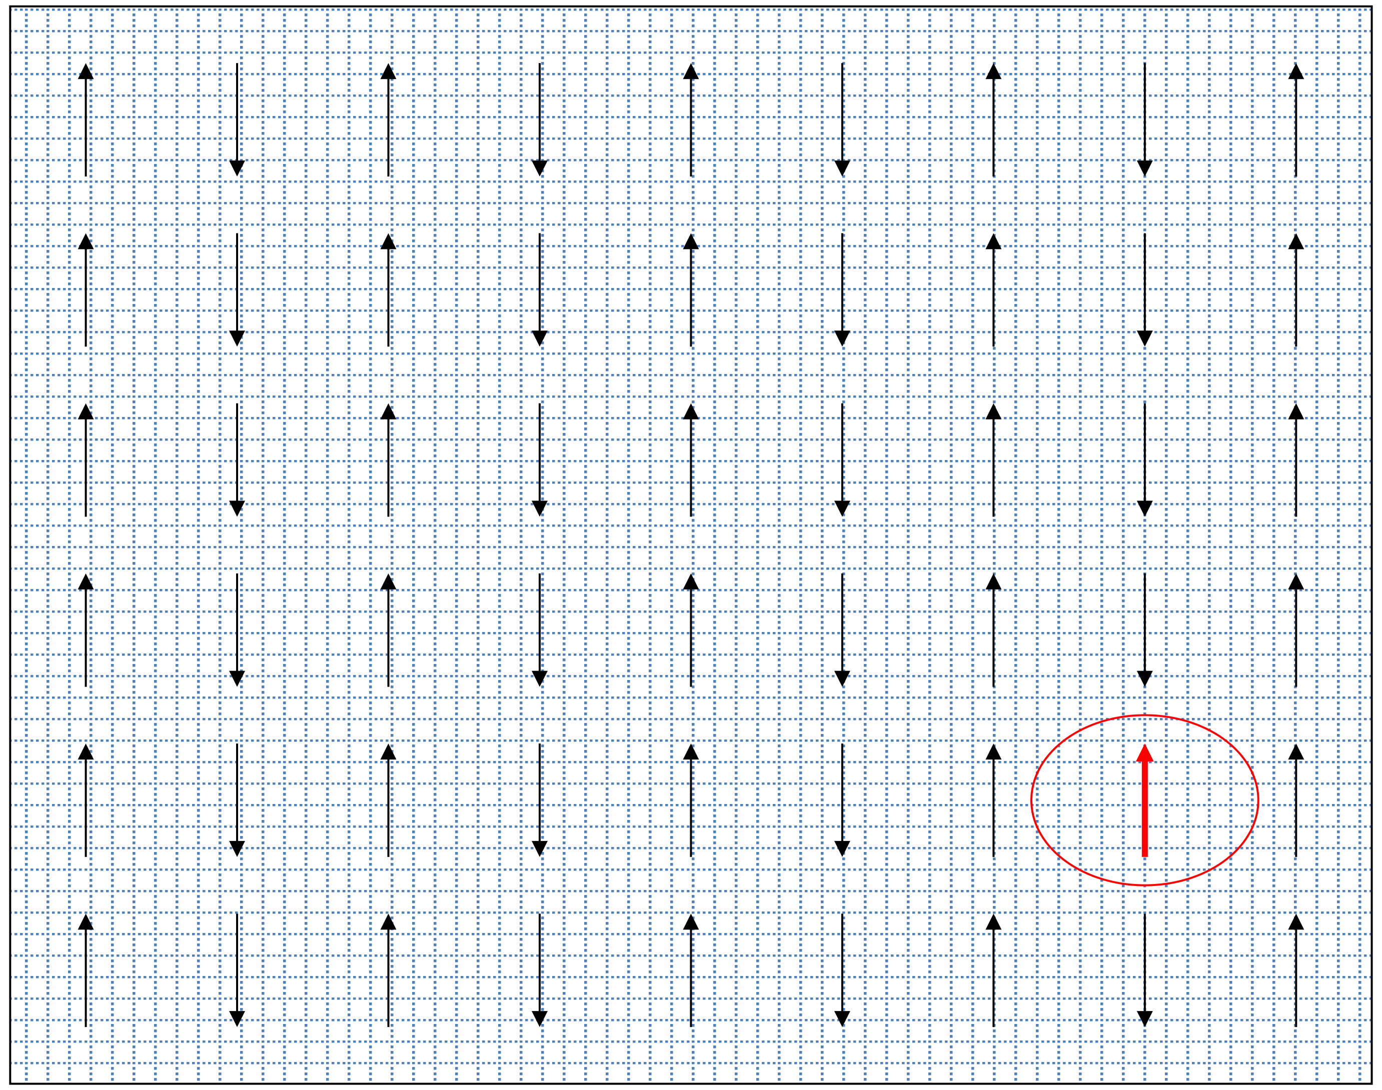

Elementary magnet

$S_i = +1$

$S_i = -1$

$w_{ij} = +1$

$w_{ij} = -1$

dynamics

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

Sum over all interactions with i

3. Detour: magnetism

Anti-ferromagnet

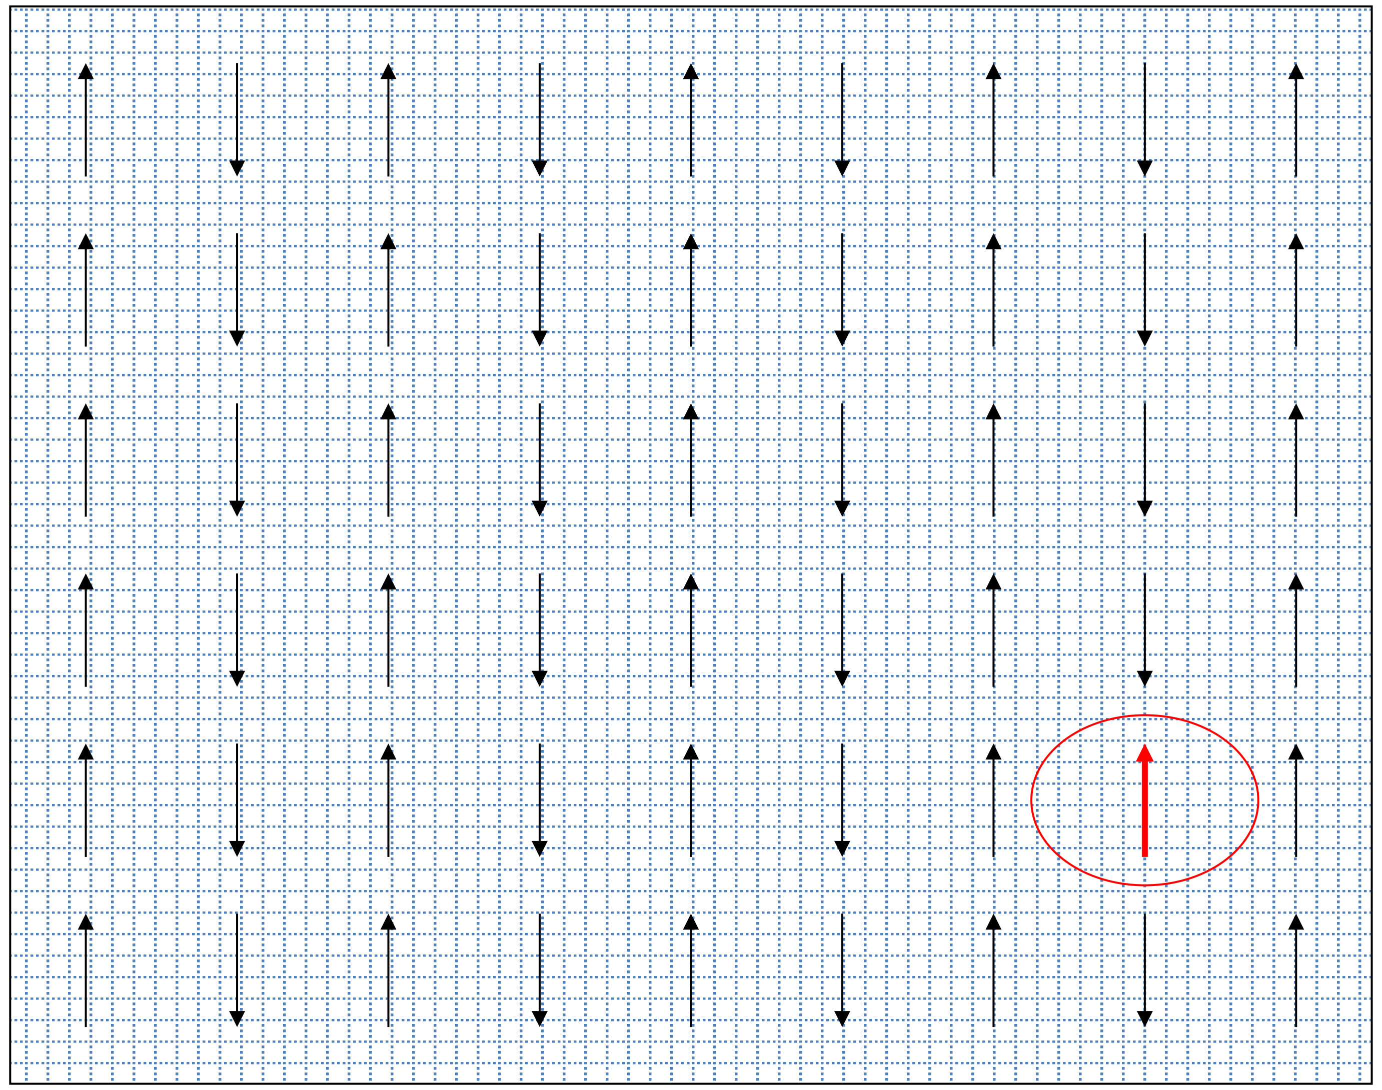

Elementary magnet

↑ $S_i = +1$

↓ $S_i = -1$

↑↑ $w_{ij} = +1$

↑↓ $w_{ij} = -1$

dynamics

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

Sum over all interactions with i

3. Magnetism and memory patterns

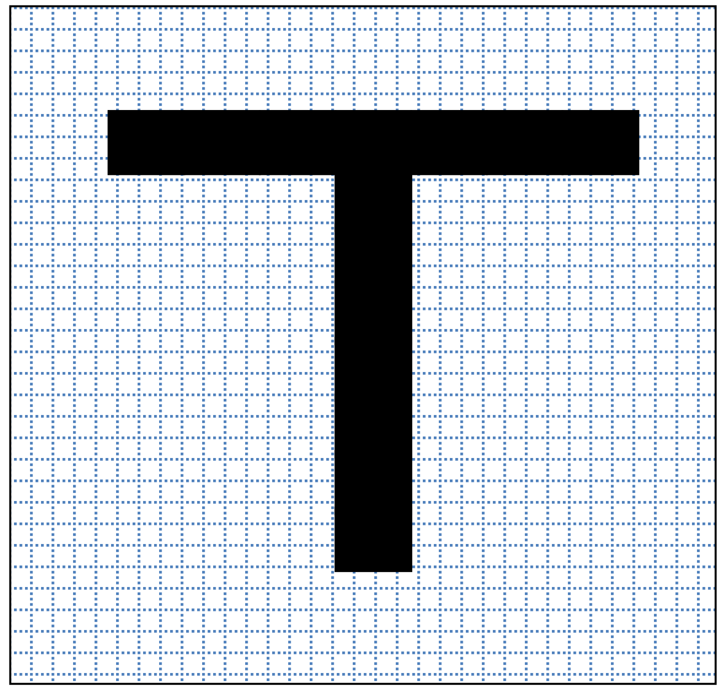

Elementary pixel

■ $S_i = +1$

□ $S_i = -1$

■ ↔ ■ $w_{ij} = +1$

□ ↔ □ $w_{ij} = +1$

□ ↔ ■ $w_{ij} = -1$

dynamics

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

Sum over all interactions with i

Hopfield model:
Several patterns → next section

Exercise 1: Associative memory (1 pattern)

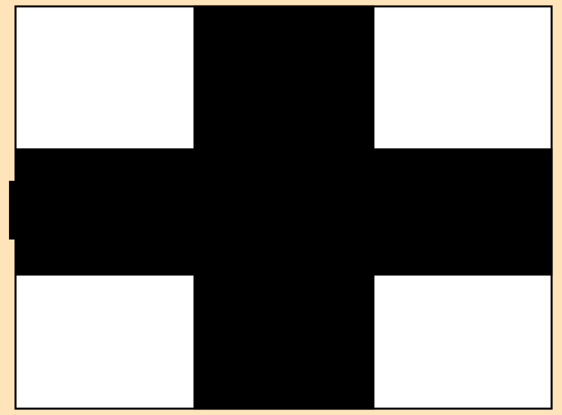

Elementary pixel

$S_i = +1$

$S_i = -1$

$w_{ij} = +1$

$w_{ij} = +1$

9 neurons, connected all-to-all

- define appropriate weights:
  what is the weight
  $w_{79} = ?$
- what happens if neuron 7 is +1?
- what happens if *3* neurons wrong?

dynamics

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

Sum over all
interactions with i

Computational Neuroscience: Neuronal Dynamics of Cognition

ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

A: ASSOCIATIVE MEMORY
in a Network of Neurons

Wulfram Gerstner

EPFL, Lausanne, Switzerland

Reading for this week:
NEURONAL DYNAMICS
- Ch. 17.1 - 17.2.4

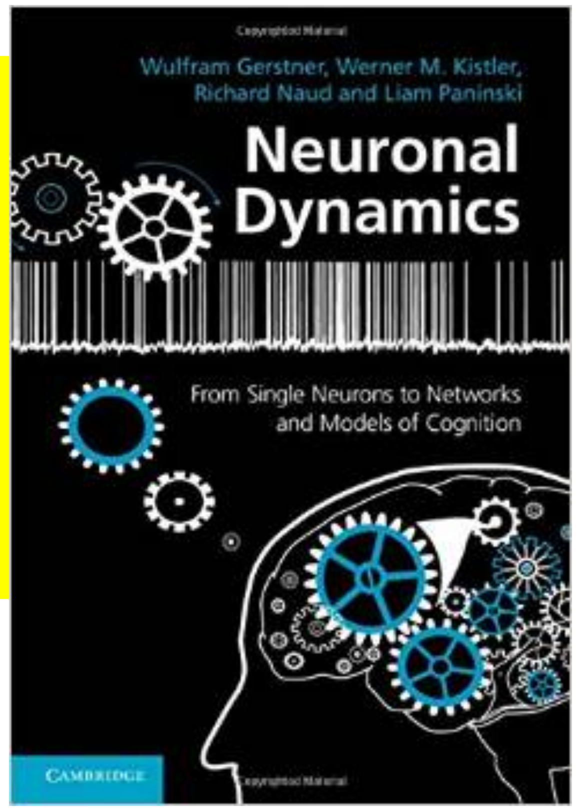

Cambridge Univ. Press

4. Single pattern

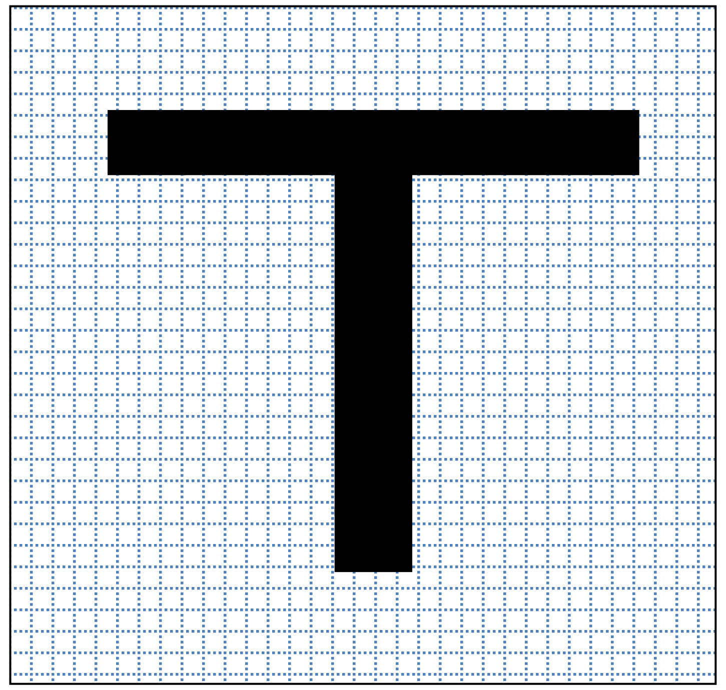

Elementary pixel
(target pattern)

■ p_i = +1

□ p_i = -1

■ ↔ ■ w_{ij} = +1

□ ↔ □ w_{ij} = +1

□ ↔ ■ w_{ij} = -1

$$w_{ij} =$$

dynamics

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

Sum over all
interactions with i

4. Hopfield Model of Associative Memory

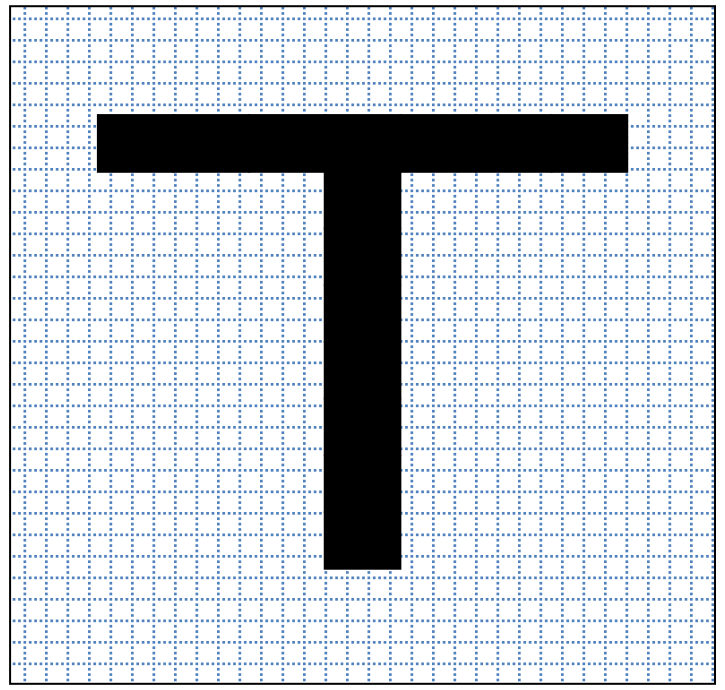

Prototype $\vec{p}^{\,1}$

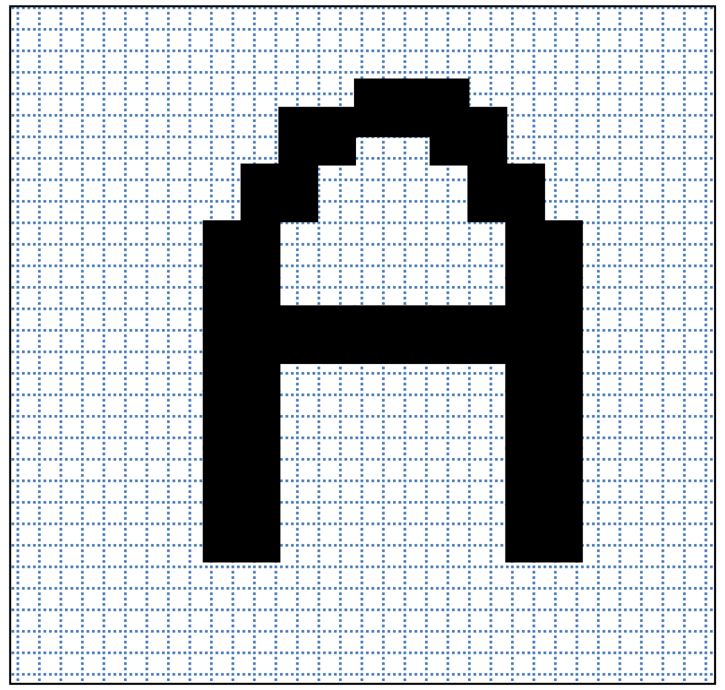

Prototype $\vec{p}^{\,2}$

several patterns

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

4. Hopfield Model of Associative Memory

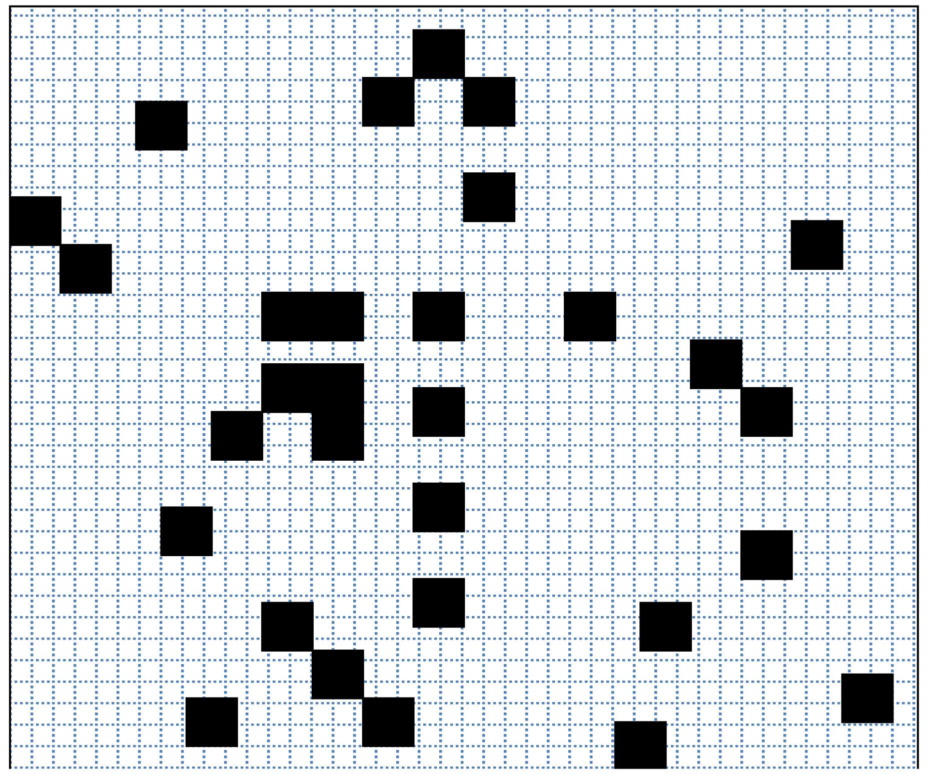

Pattern

$\vec{p}^{1}$

interactions

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

Sum over all prototypes

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

It does not work well for correlated patters

Hopfield model (1982)
- **several random patterns**
- **fully connected network**
- **binary neurons**
- **weights (1); dynamics (2)**

dynamics

$$S_i(t+1) = \mathrm{sgn}[\sum_{j} w_{ij} S_j(t)] \quad \textbf{(2)}$$

all interactions with i

J. Hopfield, 1982

4. Overlap: a measure of similarity

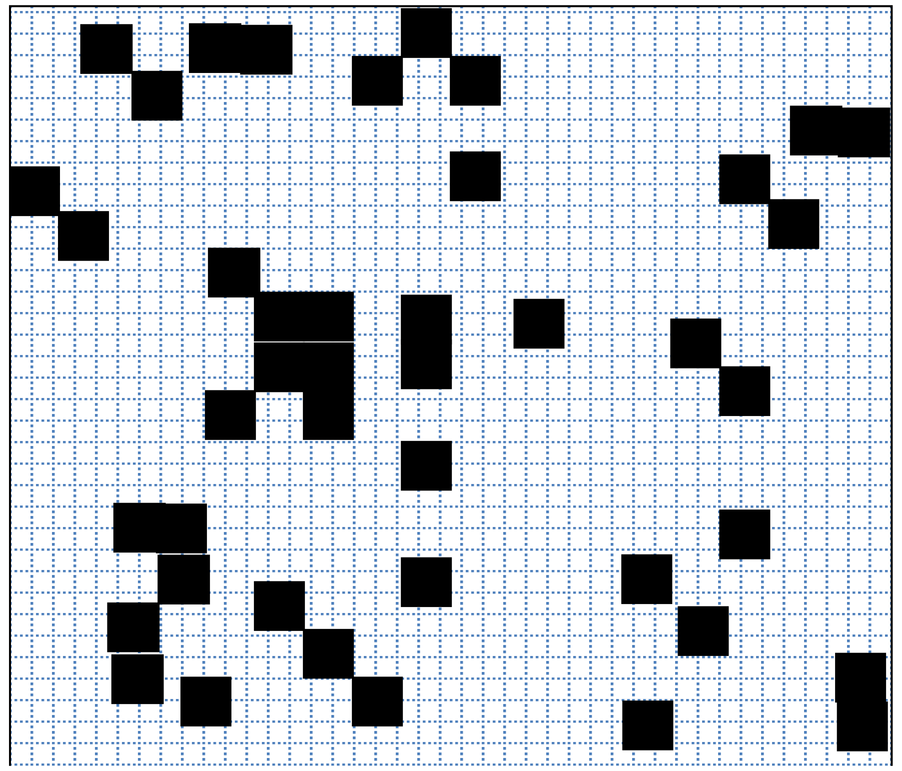

current state: (+1,-1,-1,+1,-1,+1,+1,-1)

target pattern, (+1,+1,-1,+1,-1,-1,-1,-1)
prototype

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

4. Hopfield Model of Associative Memory

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

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

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

4. Hopfield Model of Associative Memory

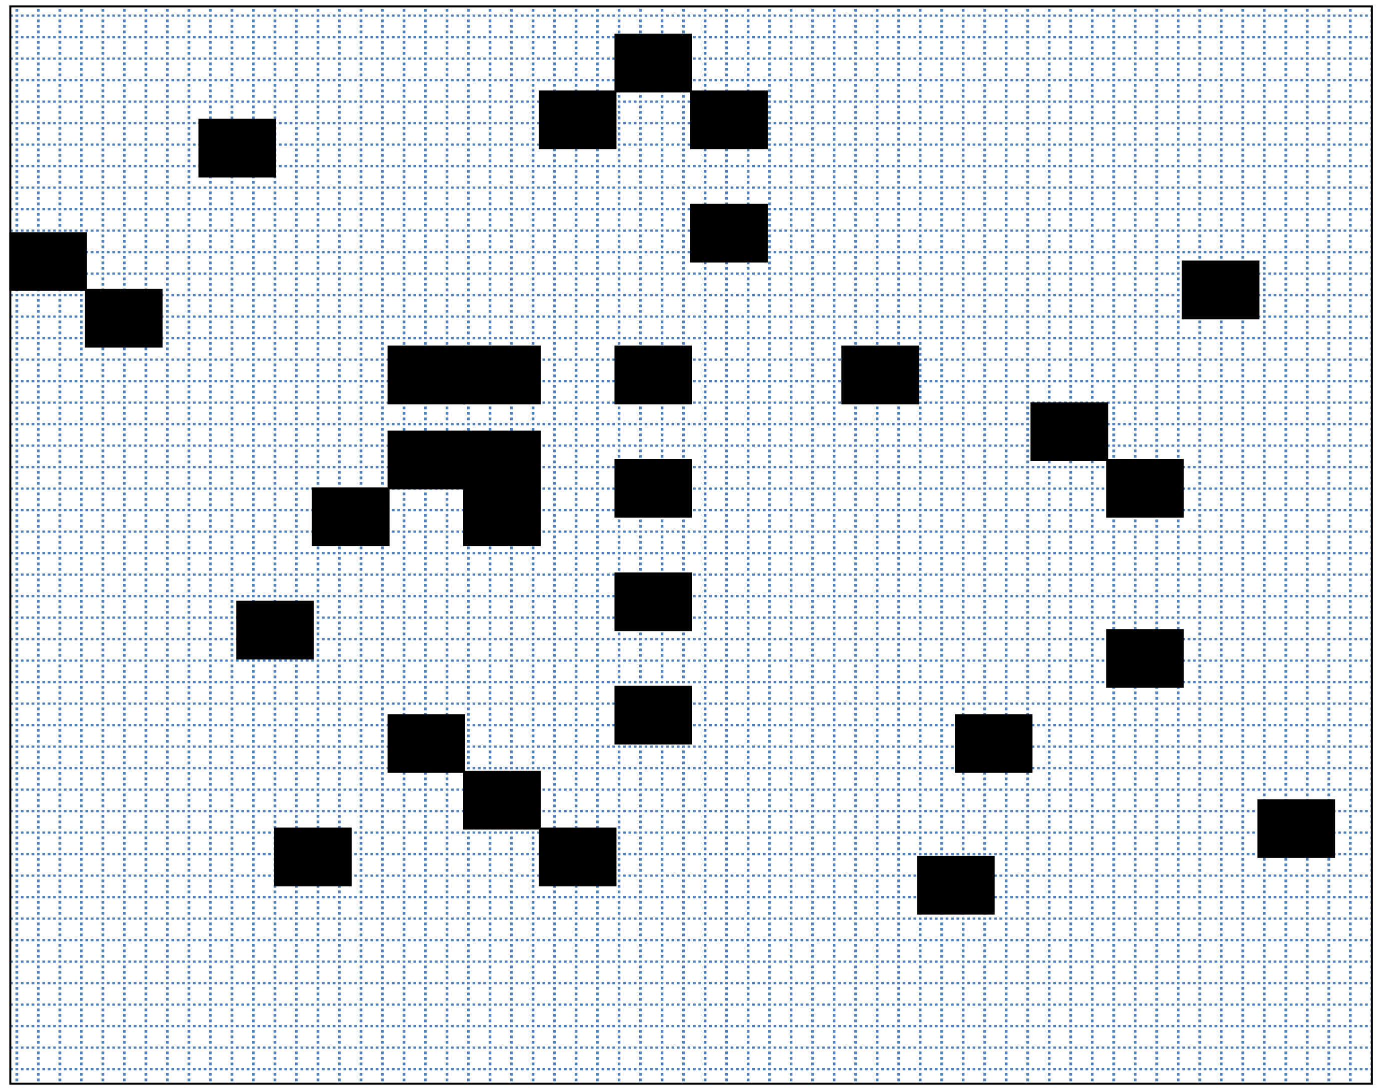

Interacting neurons

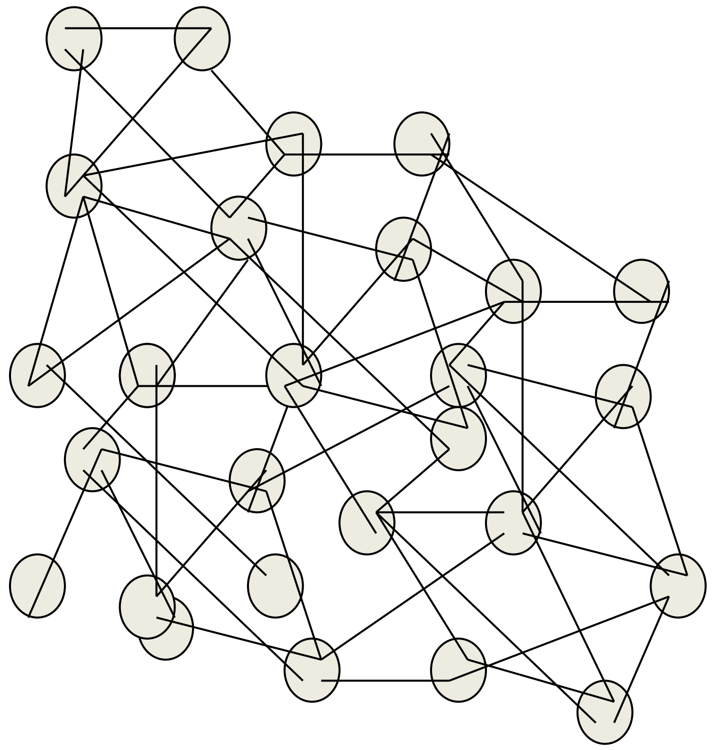

Prototype

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

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

Hopfield model

Computation
- without CPU,
- without explicit memory unit

4. Correlated patterns, orthogonal patterns

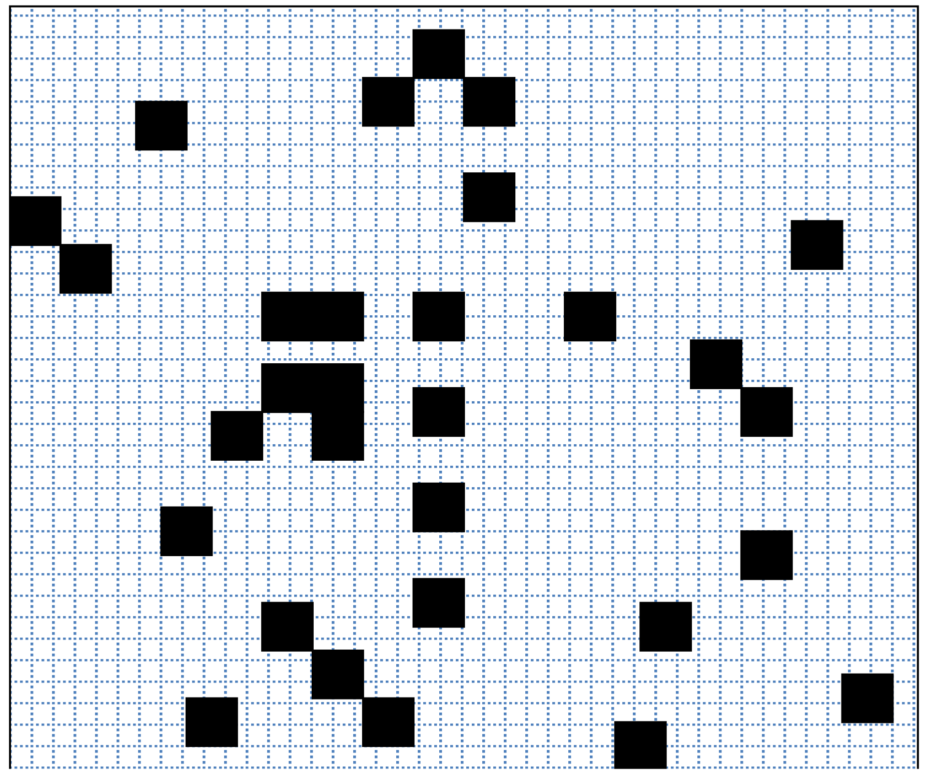

target pattern, prototype 3 (+1,-1,+1,+1,-1,+1,+1,-1)

target pattern, prototype 7 (+1,+1,-1,+1,-1,-1,-1,-1)

Similarity of two patterns:

Orthogonal patterns:

Random patterns

overlap $m^{\mu}(t) = \frac{1}{N}\sum_{j} p_j^{\mu} S_j(t)$

Exercise 2 (now)

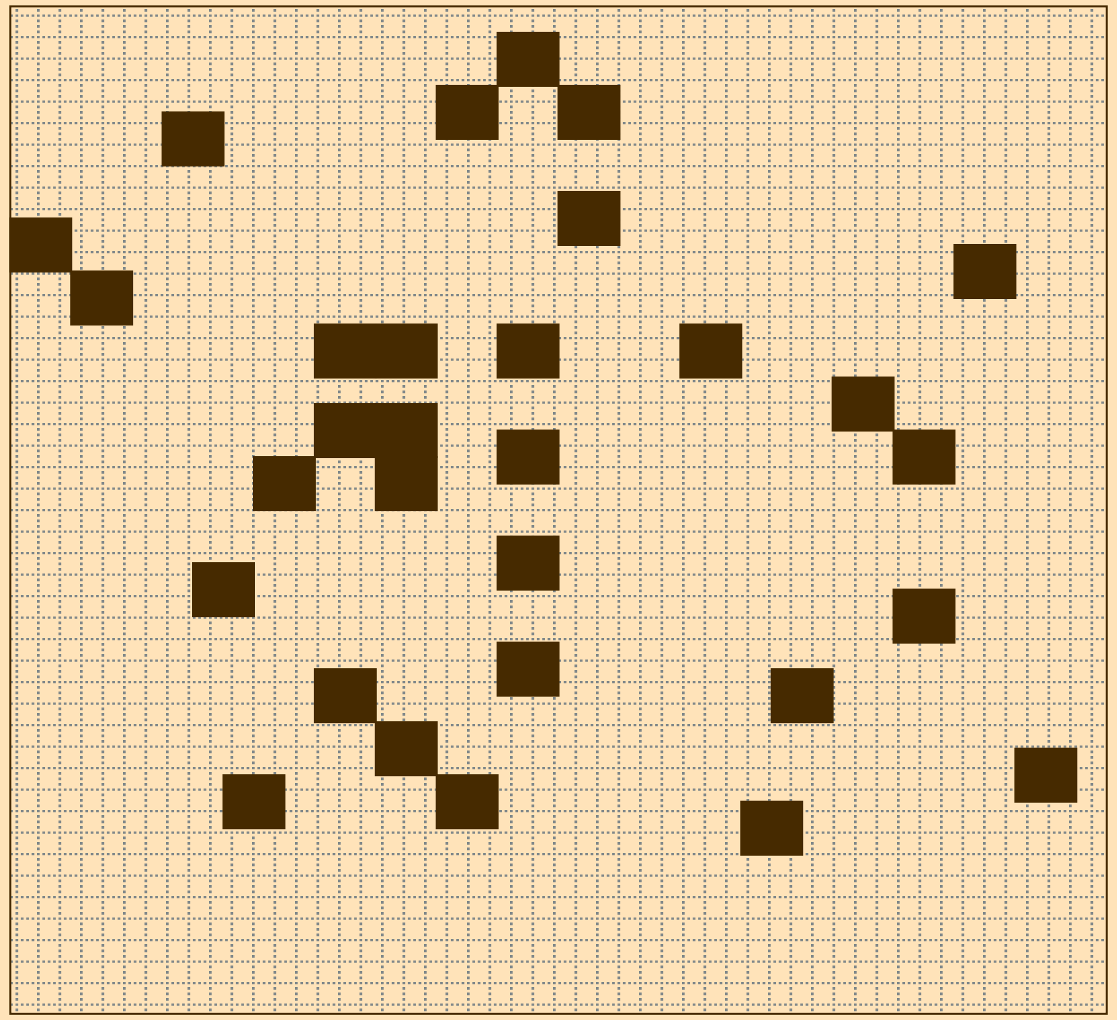

$$w_{ij} = \tfrac{1}{N}\sum_{\mu} p_i^{\mu} p_j^{\mu}$$

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

Sum over all interactions with i

Assume 4 orthogonal patterns.
At time t=0, overlap with pattern 3, no overlap with other patterns.

Calculate the overlap at t=1!

Computational Neuroscience: Neuronal Dynamics of Cognition

ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

A: ASSOCIATIVE MEMORY
in a Network of Neurons

Wulfram Gerstner

EPFL, Lausanne, Switzerland

Reading for this week:
NEURONAL DYNAMICS
- Ch. 17.1 - 17.2.4

Cambridge Univ. Press

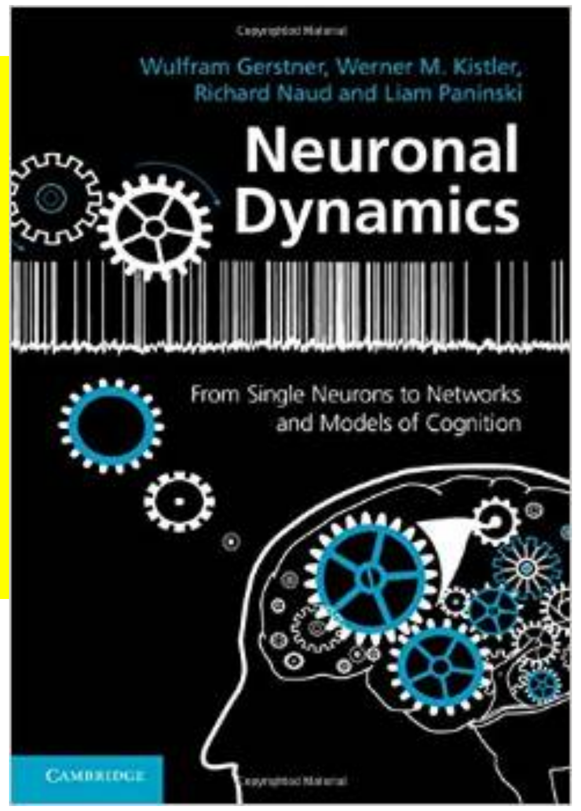

1 Introduction
- networks of neuron
- systems for computing
- associative memory

2 Classification by similarity

3 Detour: Magnetic Materials

4 Hopfield Model

5 Learning of Associations

6 Storage Capacity

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. Hebbian Learning of Associations

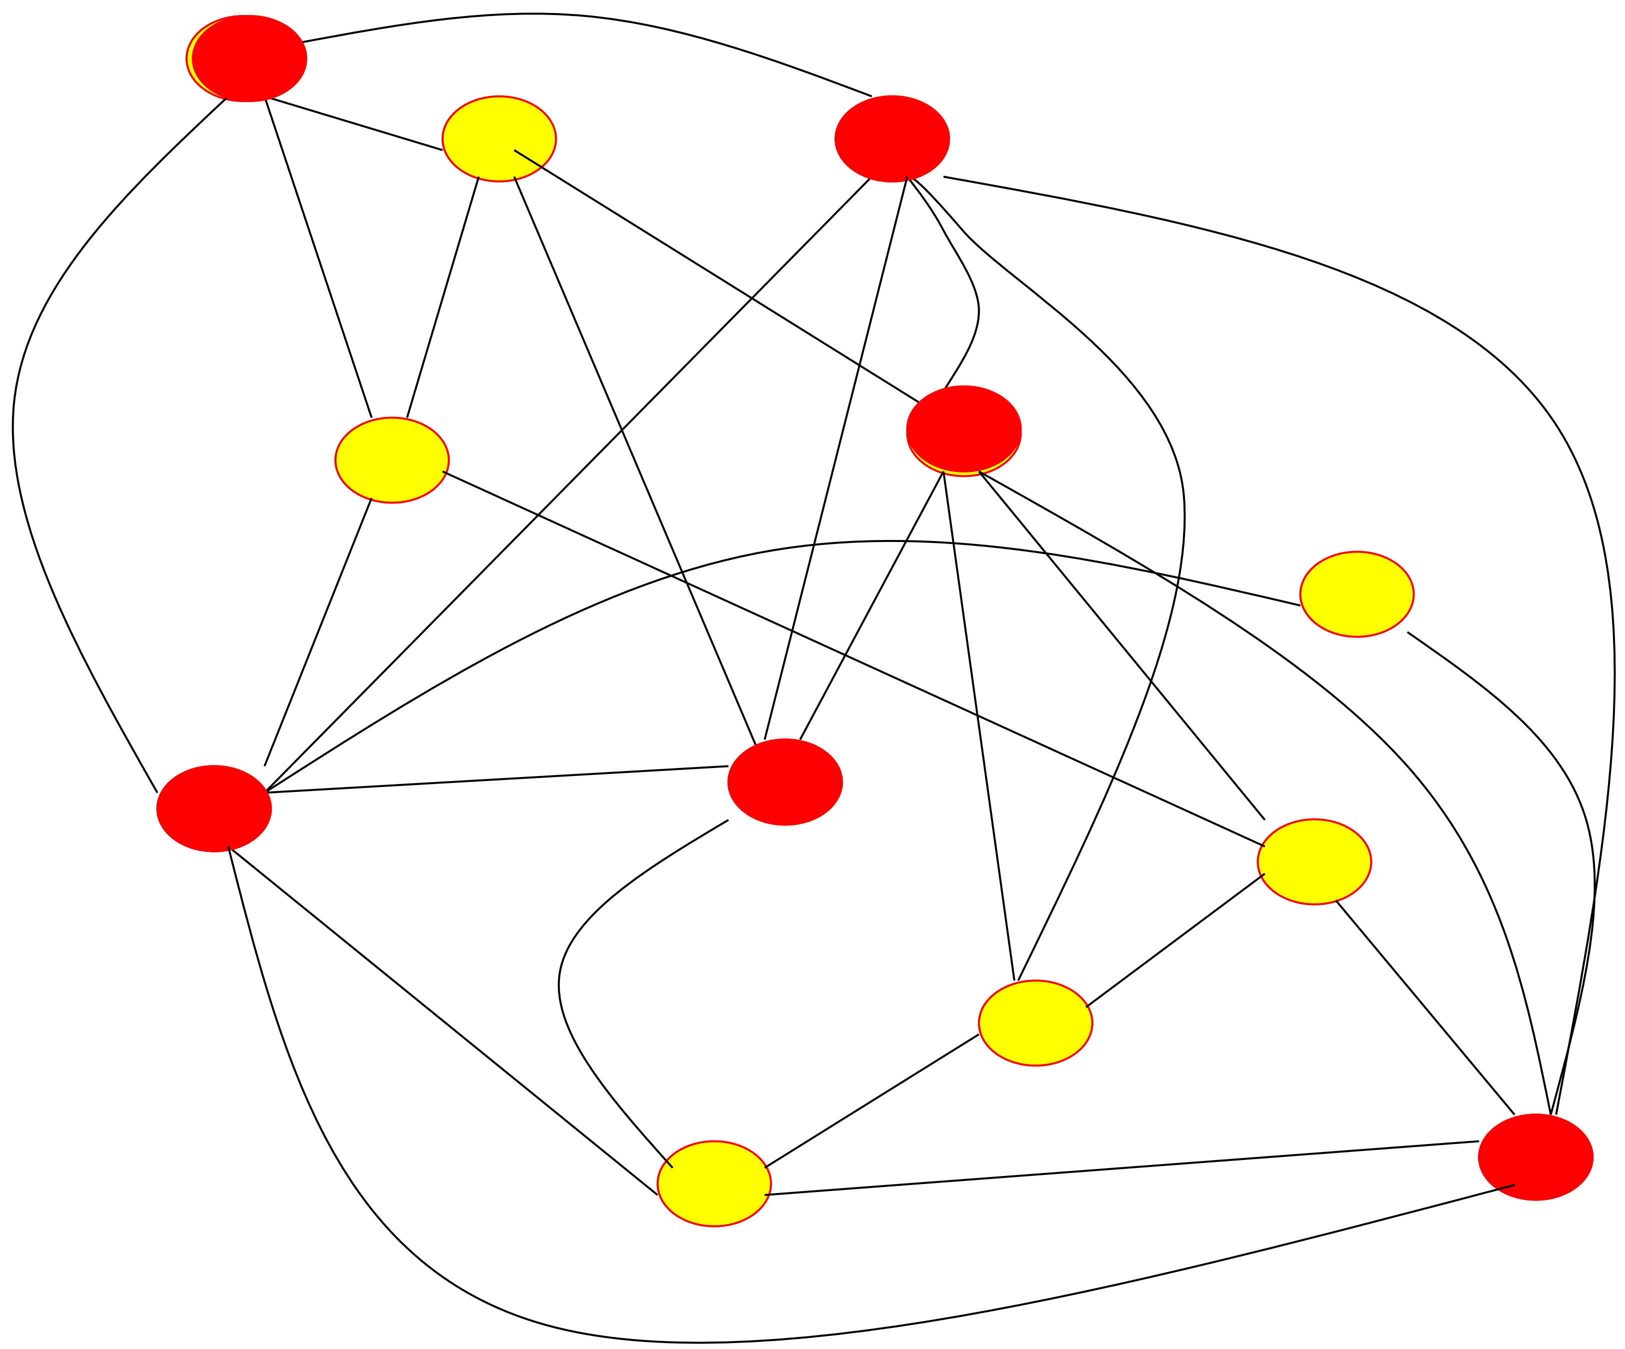

5. Hebbian Learning of Associations

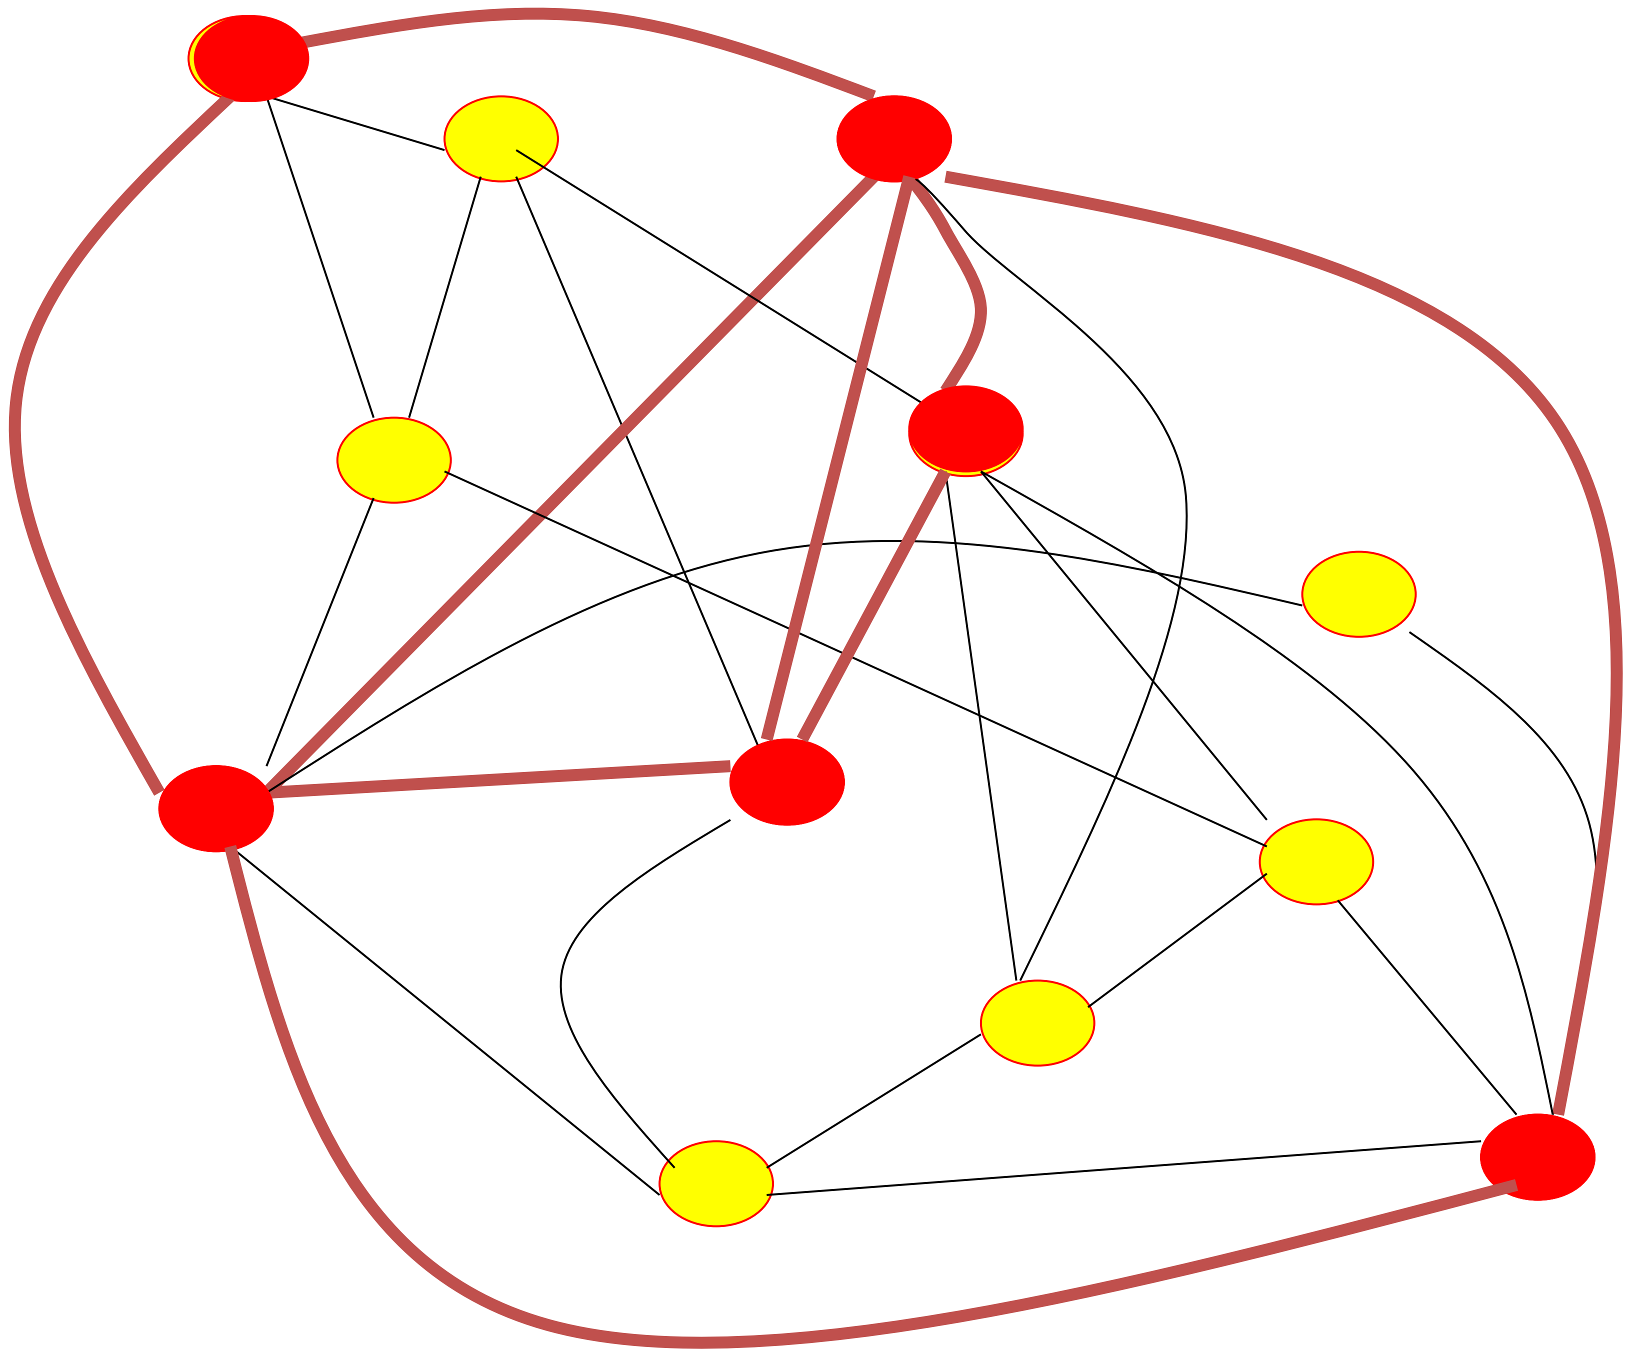

item memorized

5. Hebbian Learning: Associative Recall

Recall:
Partial info

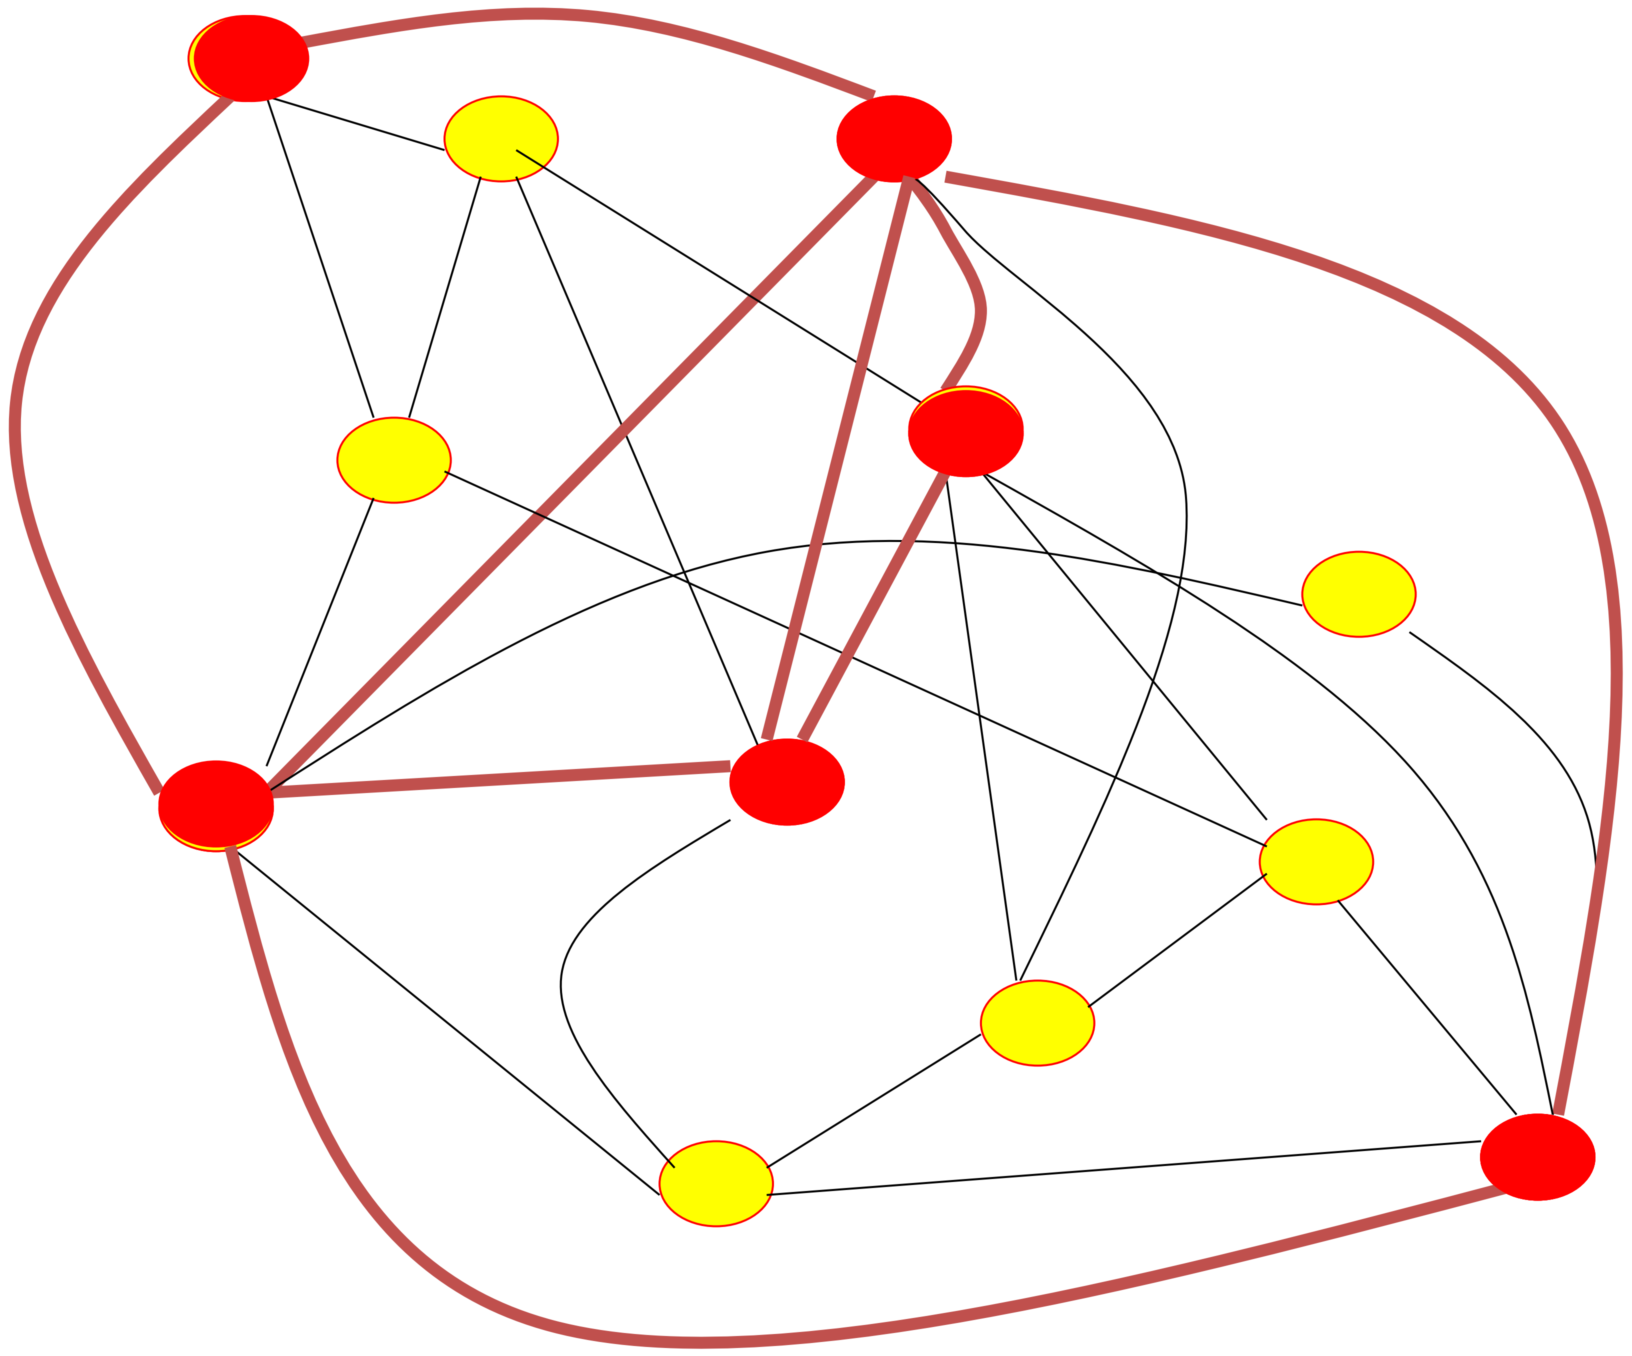

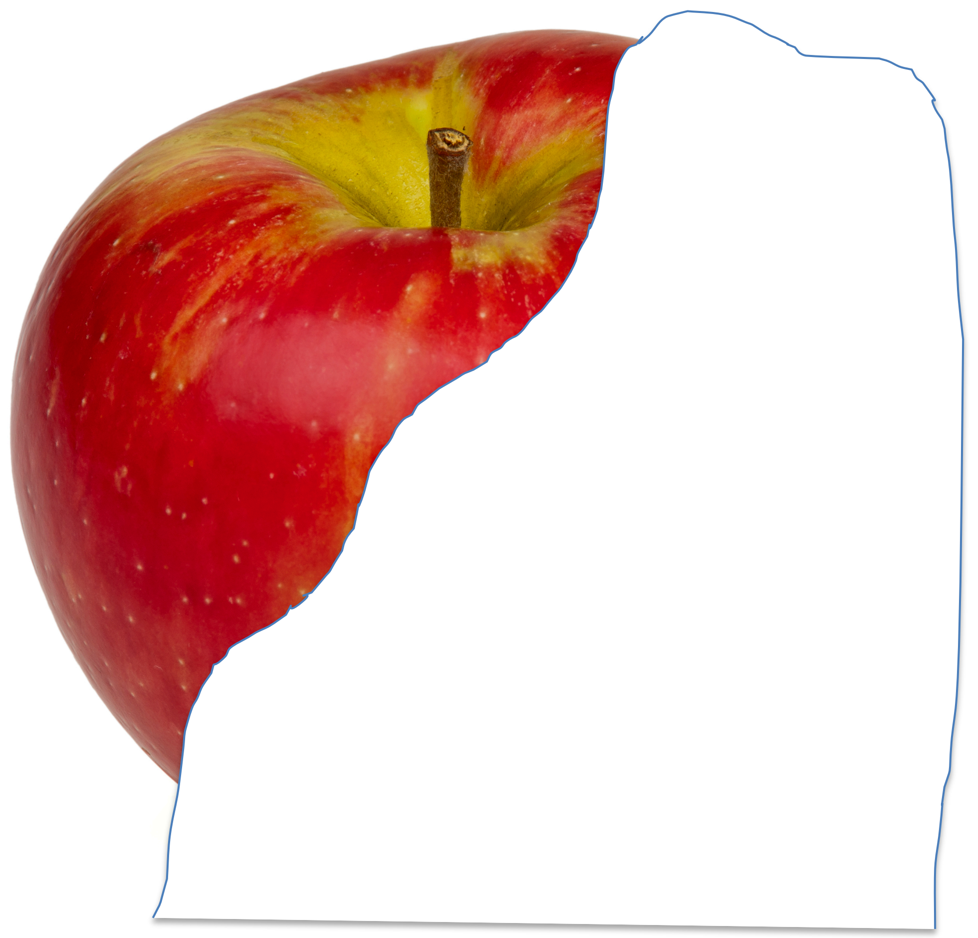

item recalled

5. Learned concepts

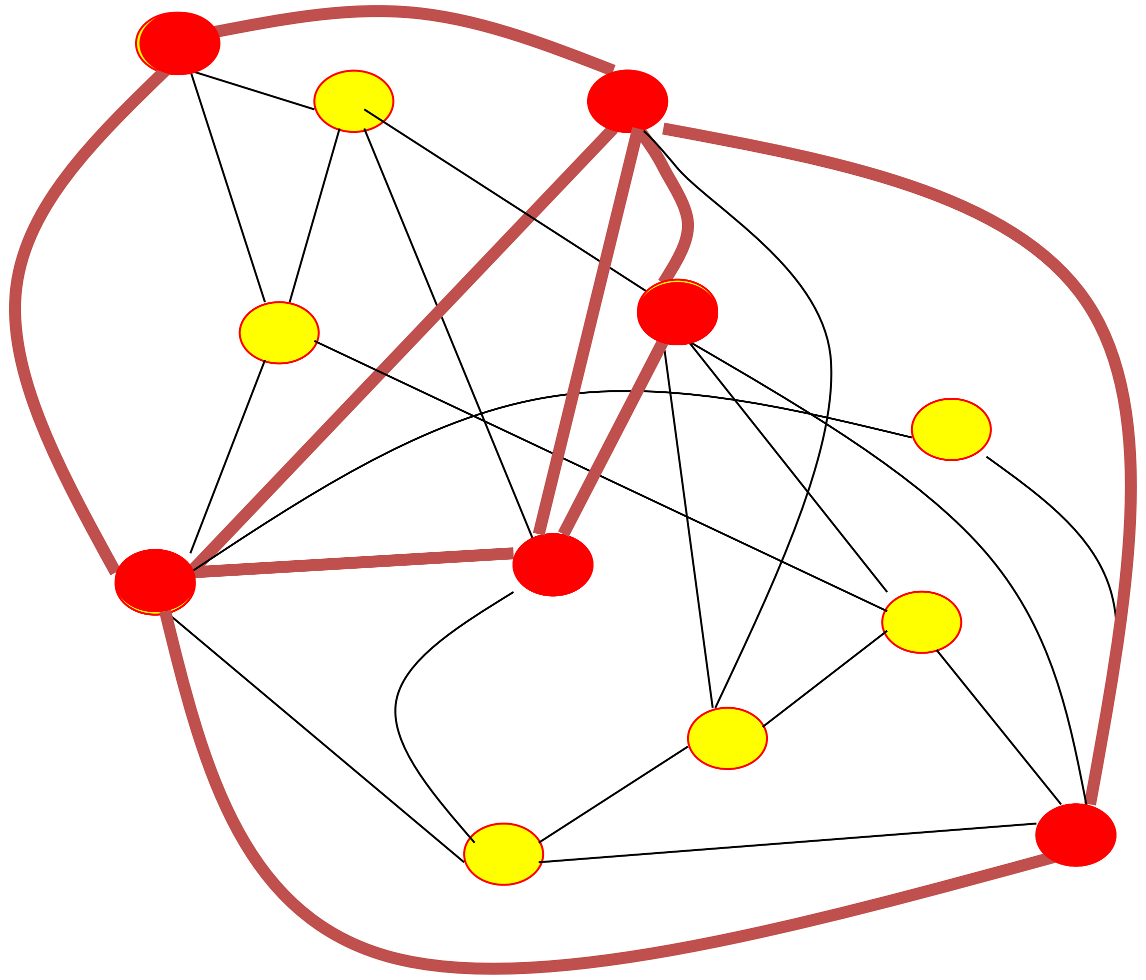

Activity of neurons in human brain

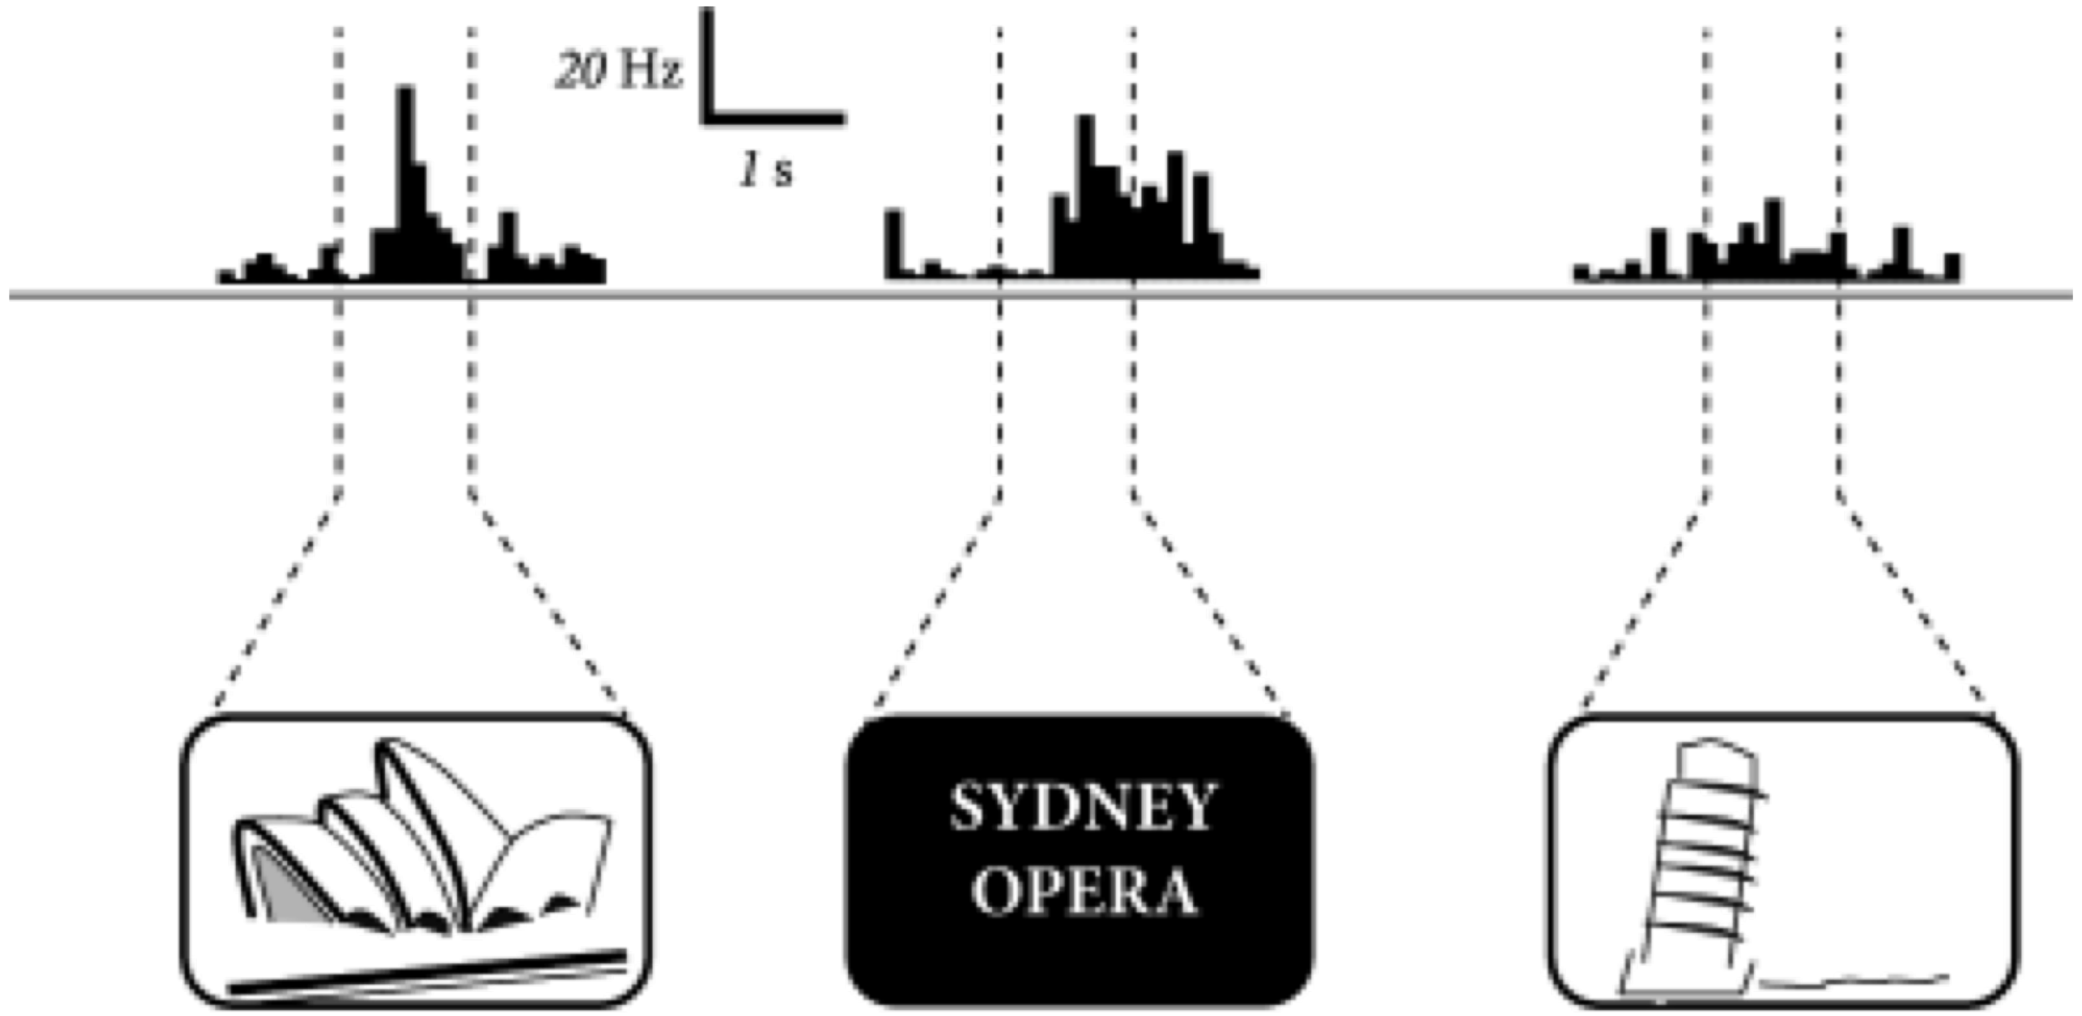

assembly of neurons

Image: Neuronal Dynamics, Gerstner et al., Cambridge Univ. Press (2014), Adapted from ***Quiroga et al. (2005),*** *Nature 435:1102-1107*

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:

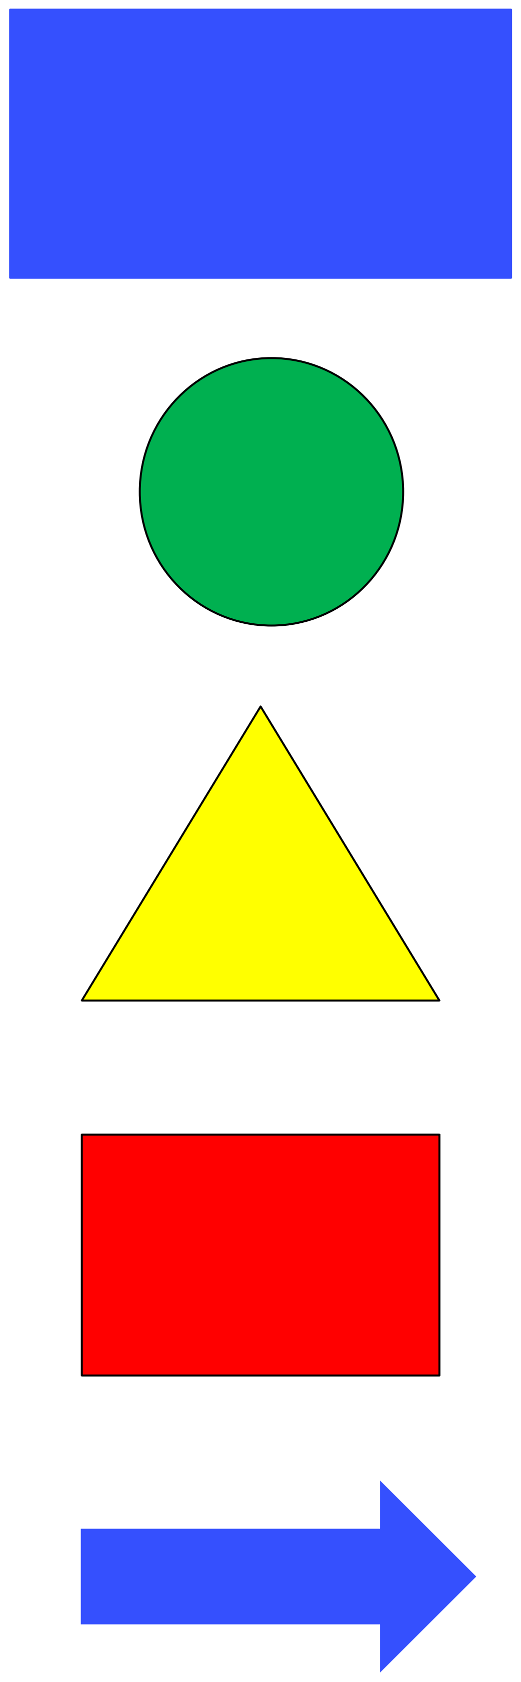

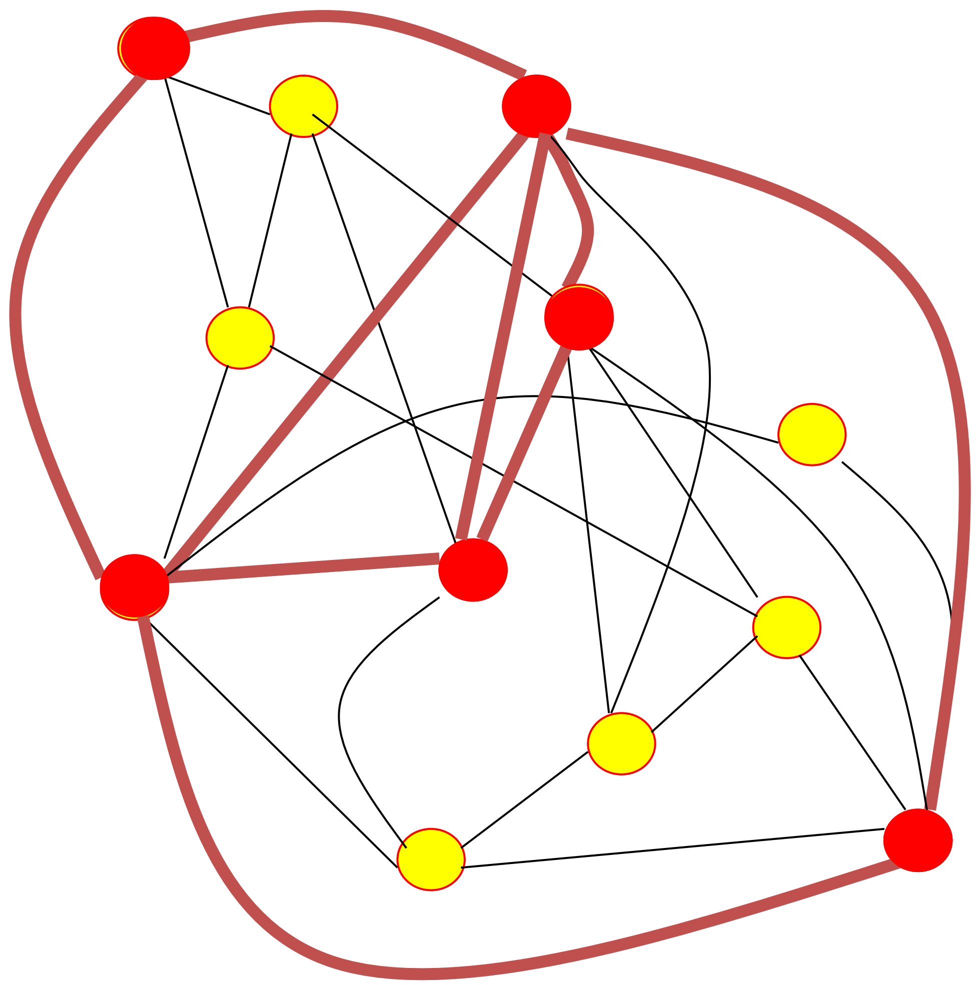

be as fast as possible:

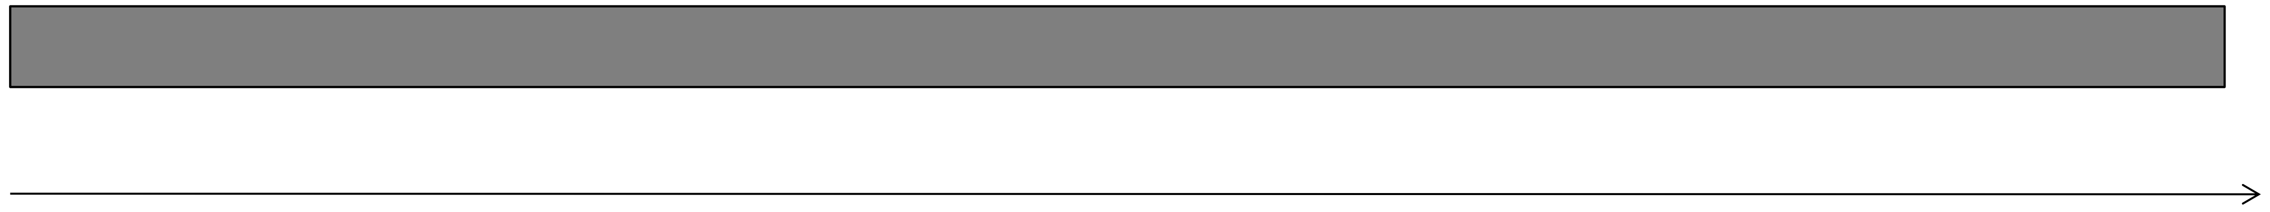

time

5. Associative Recall

Tell me the **color**

for the following list of 5 items:

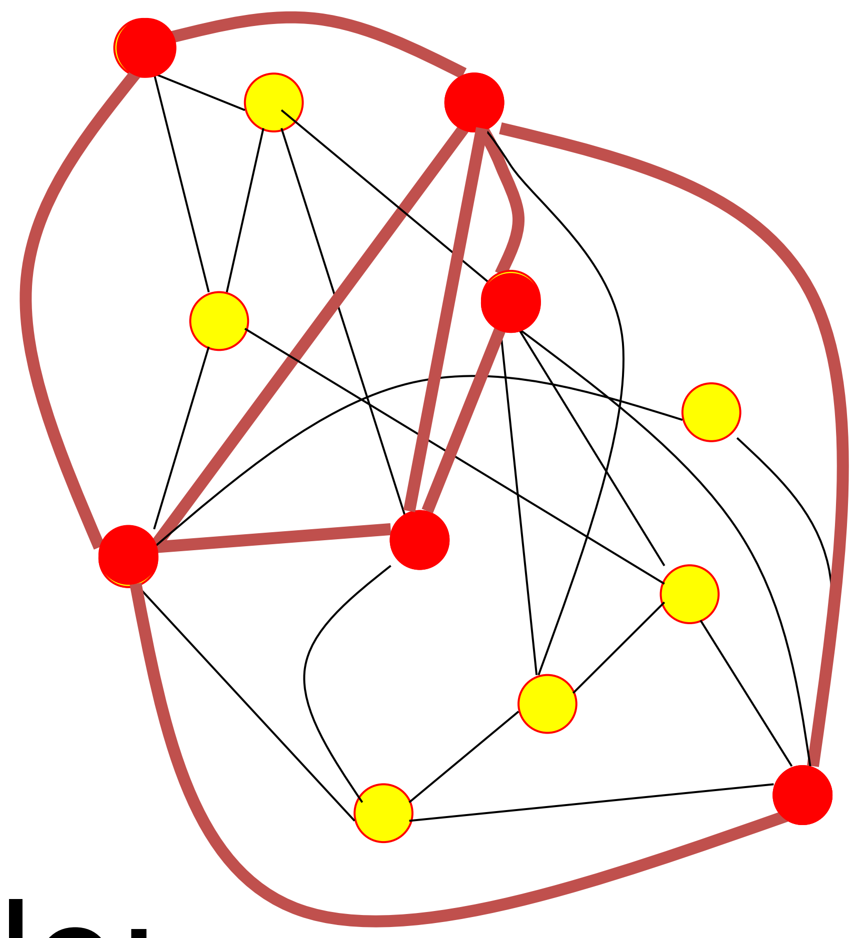

be as fast as possible:

time

Stroop effect:

Slow response: hard to work

Against natural associations

5. Associative Recall

Hierarchical organization of
Associative memory

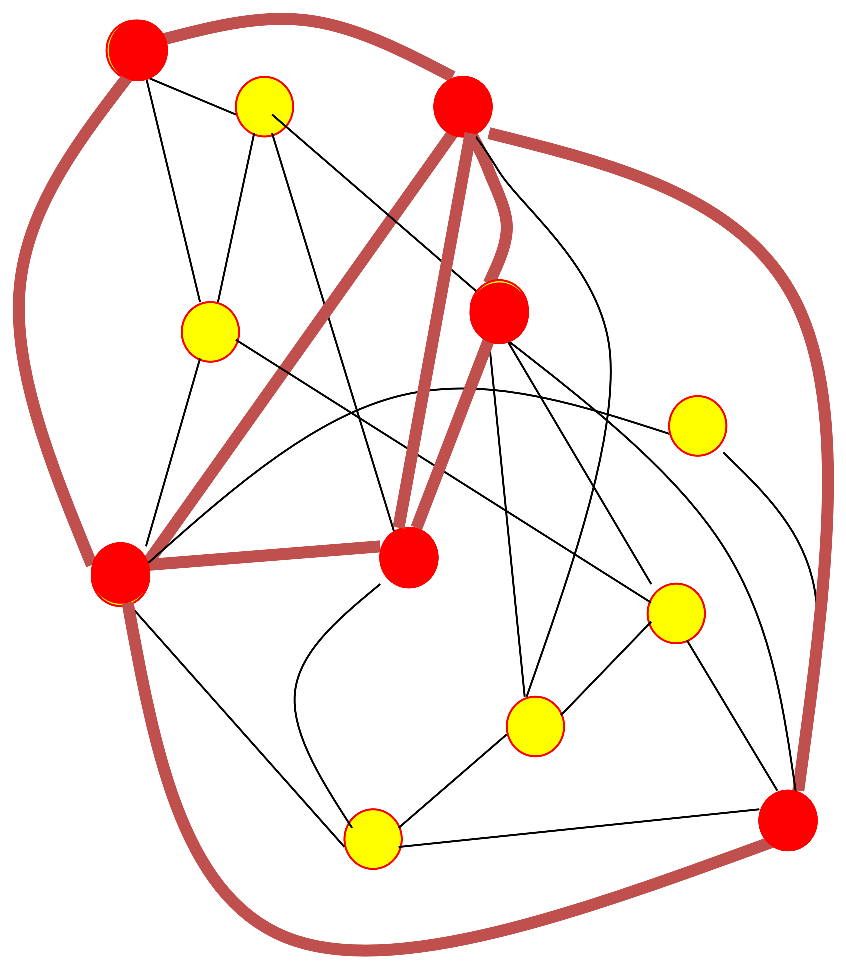

animals

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. Associative Recall

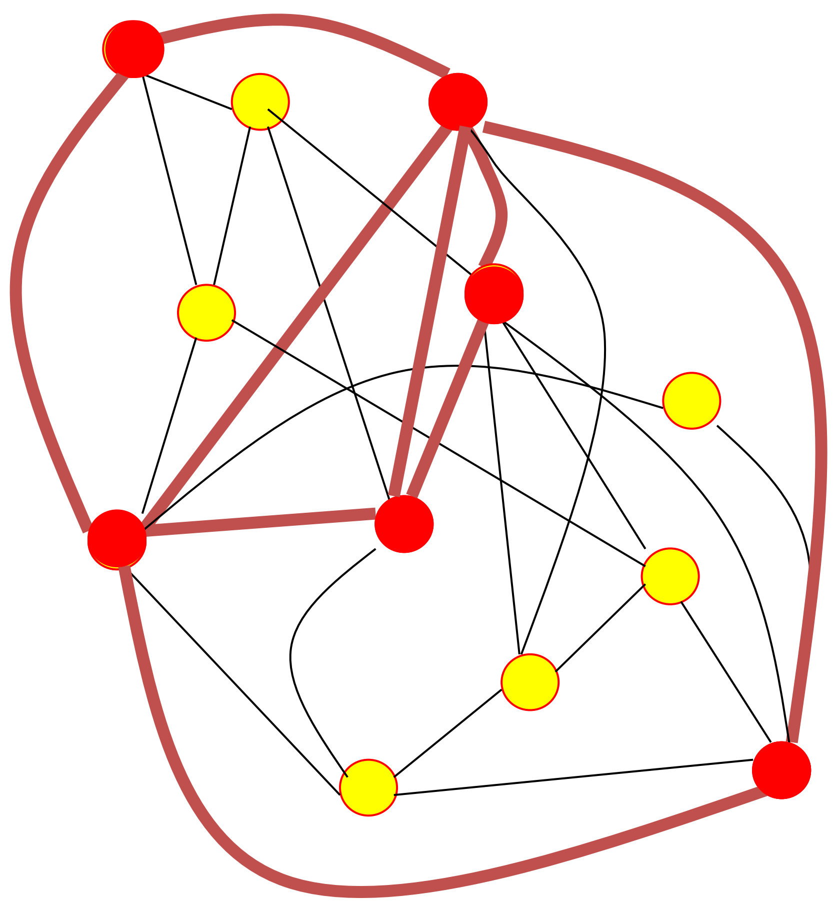

name as fast as possible
an example of a

5. Associative Recall

Associative memory

animals

birds fish

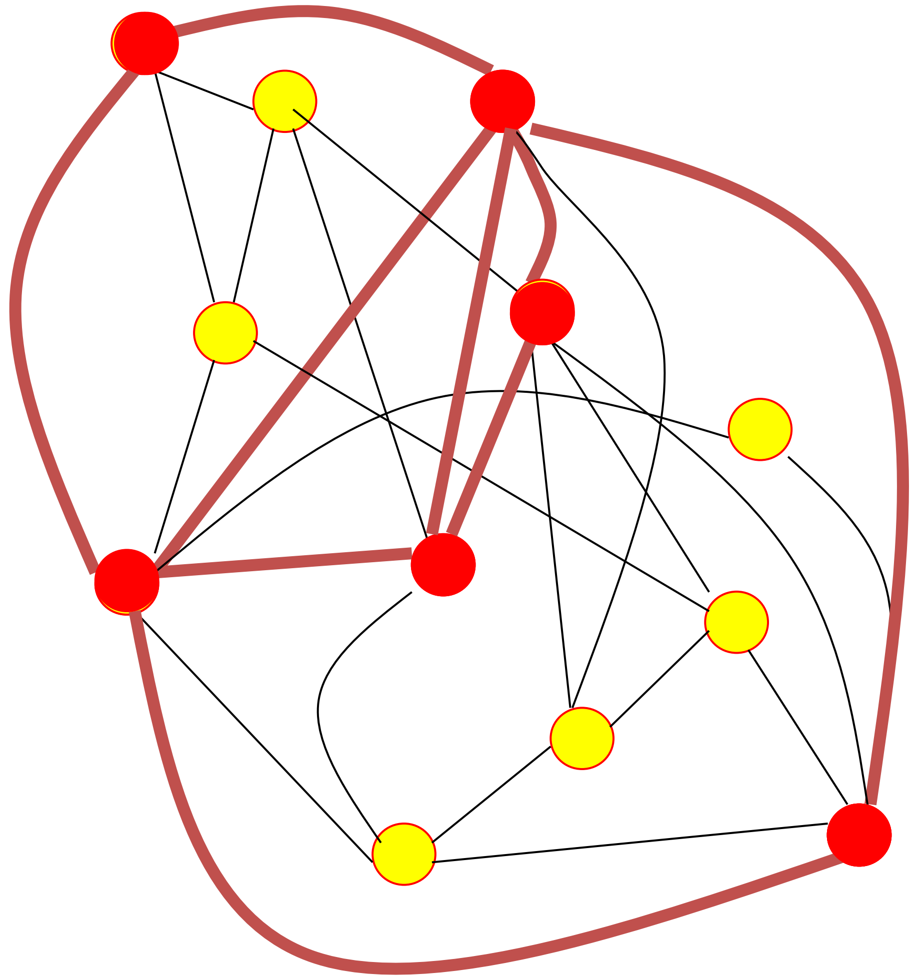

- Associations can be very strong!
- It is hard to go against natural associations!
- Different aspects of a ‘concept’ are bound together!
- Assocations have been learned!

Quiz 3: Assocations

The Stroop effect implies that you are faster,
if the color does not match the meaning of the color-word

[ ] Yes

[ ] No

Hebbian learning strengthens links between neurons that

[ ] are simultaneously active

[ ] belong to the same ‘concept’ (assembly)

Computational Neuroscience: Neuronal Dynamics of Cognition

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

A: ASSOCIATIVE MEMORY in a Network of Neurons

Wulfram Gerstner

EPFL, Lausanne, Switzerland

Reading for this week:
NEURONAL DYNAMICS
- Ch. 17.1 - 17.2.4

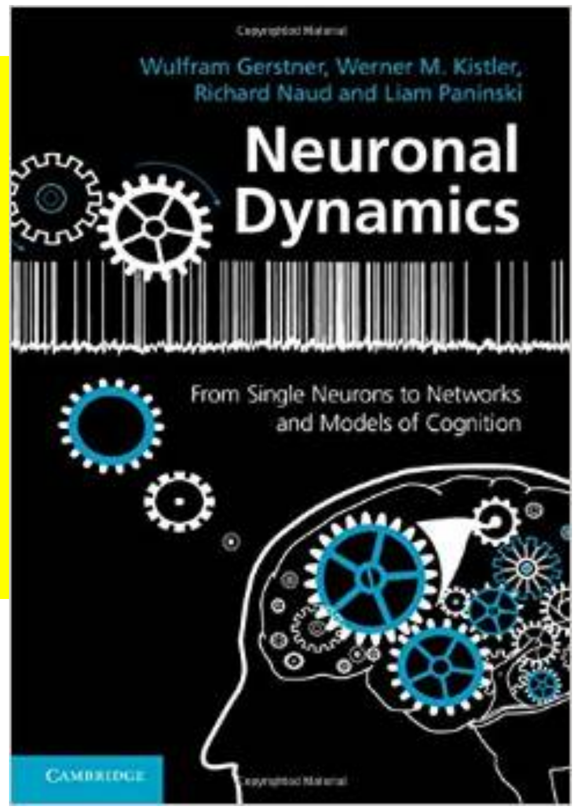

Cambridge Univ. Press

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

2 Classification by similarity

3 Detour: Magnetic Materials

4 Hopfield Model

5 Learning of Associations

6 Storage Capacity

6. learning of several prototypes

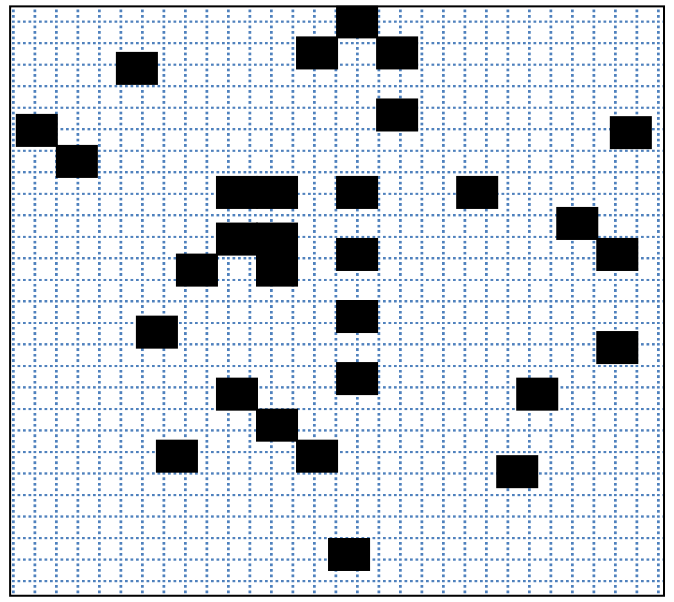

Prototype $\vec{p}^1$

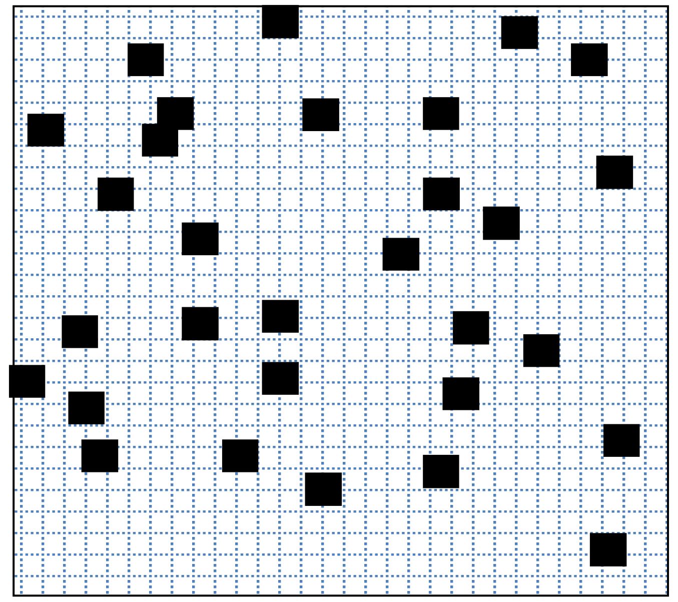

Prototype $\vec{p}^2$

interactions

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

Sum over all prototypes

Question: How many prototypes can be stored?

dynamics

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

all interactions with i

6. Storage capacity: How many prototypes can be stored?

-Assume we start directly in one pattern (say pattern 7)
-Pattern must stay

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

Interactions (1)

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

6. Storage capacity: How many prototypes can be stored?

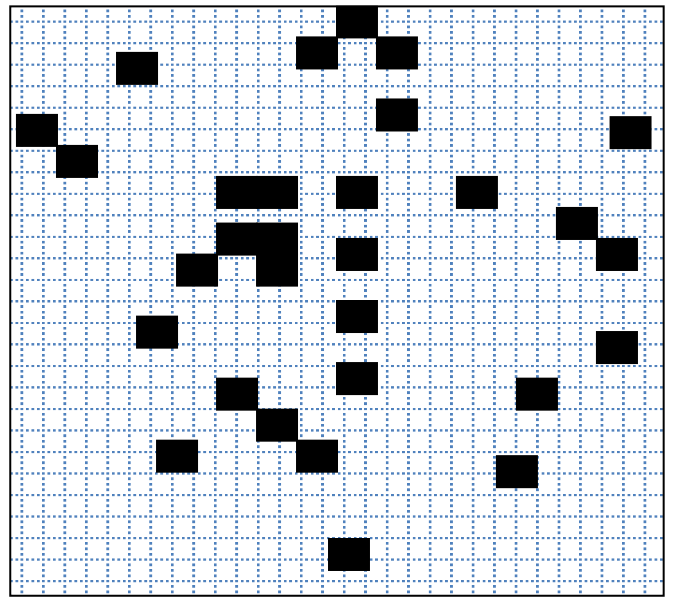

Prototype

$\vec{p}^1$

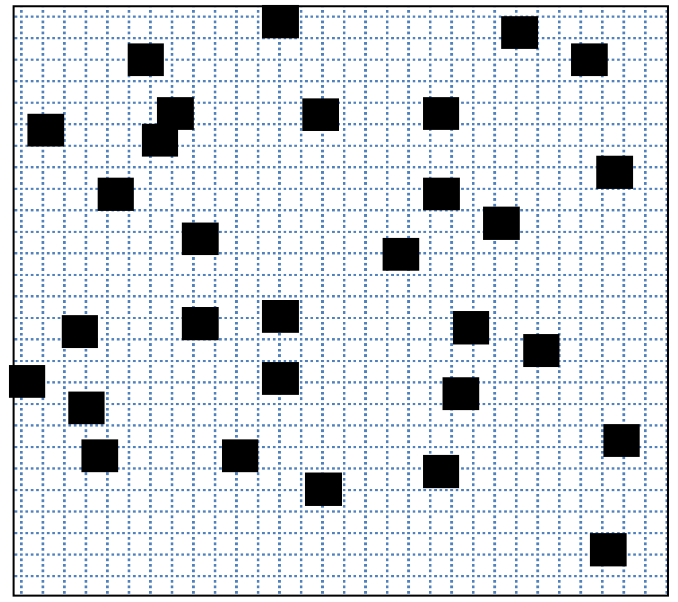

Prototype

$\vec{p}^2$

Random patterns

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

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 (say pattern ν)

-Pattern must stay

Attention: Retrieval requires more (pattern completion)

Q: How many prototypes can be stored?

A: If too many prototypes, errors (wrong pixels) show up.
The number of prototypes M that can be stored
is proportional to number of neurons N;
memory load = M/N

$$S_i(t+1) = p_i^\nu \operatorname{sgn}[1 + \tfrac{1}{N} \sum_{\mu=1,\mu\neq\nu}^{M} \sum_{j=1}^{N} p_i^\mu p_i^\nu p_j^\mu p_j^\nu]$$

$$= p_i^\nu \operatorname{sgn}[1 - a_i^\nu]$$

Error-free if

$$S_i(t+1) = p_i^\nu$$

Gaussian

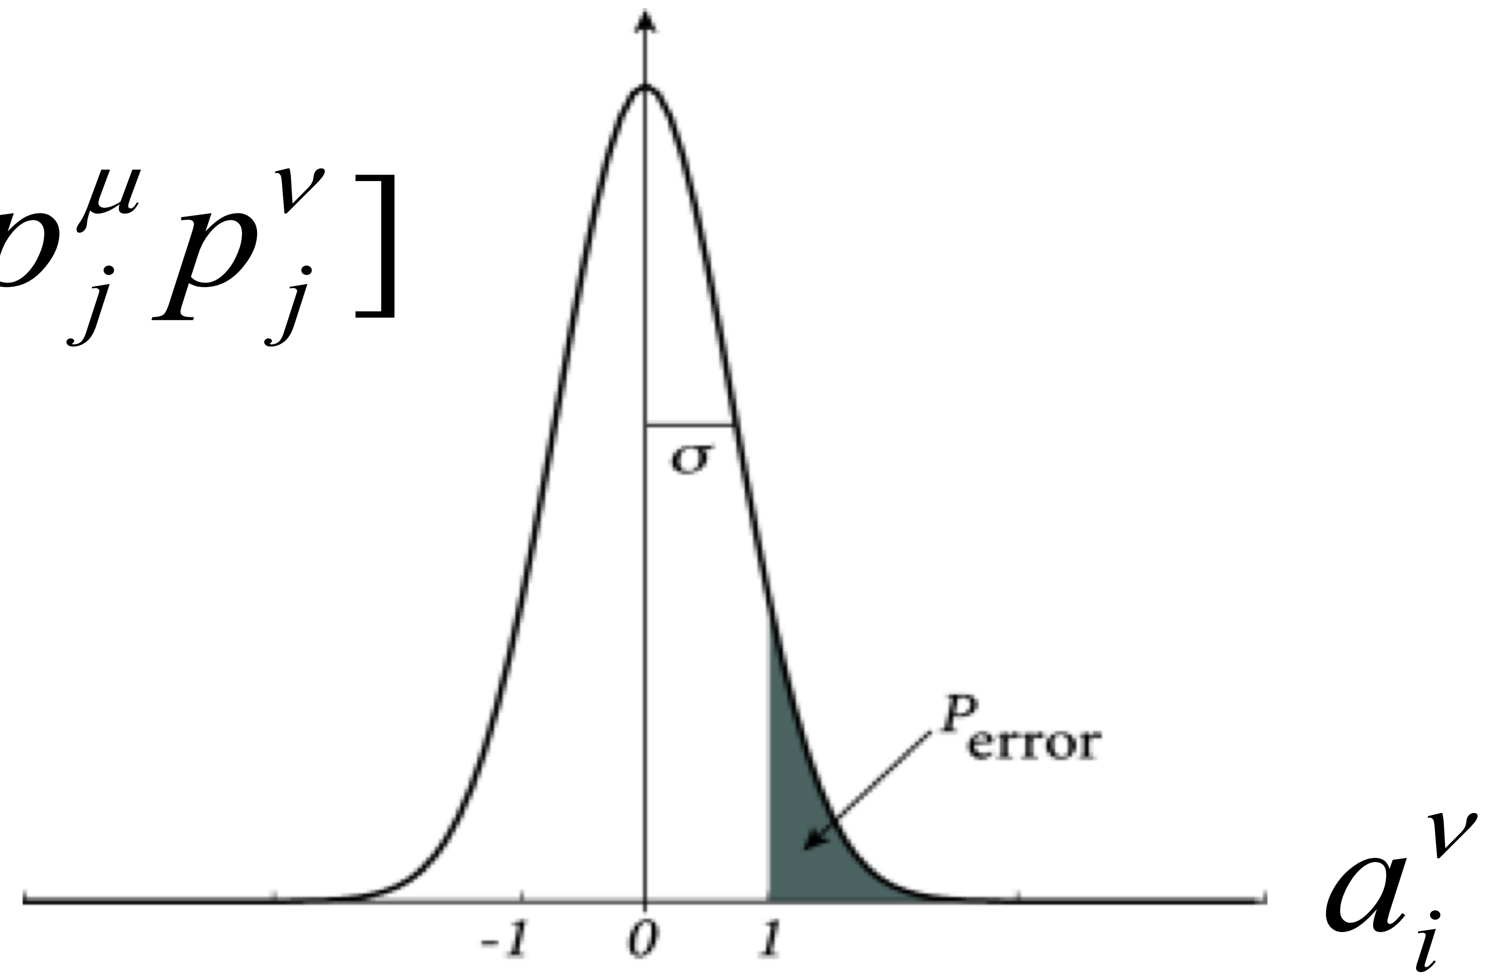

Image: Neuronal Dynamics, Gerstner et al., Cambridge Univ. Press (2014),

6. Storage capacity: How many prototypes can be stored?

Random walk with steps

Standard deviation

$$S_i(t+1) = p_i^\nu \,\mathrm{sgn}[1 + \tfrac{1}{N} \sum_{\mu=1,\mu\neq\nu}^{M} \sum_{j=1}^{N} p_i^\mu p_i^\nu p_j^\mu p_j^\nu]$$

$$= p_i^\nu \,\mathrm{sgn}[1 - a_i^\nu]$$

Error-free if

$$S_i(t+1) = p_i^\nu$$

Gaussian

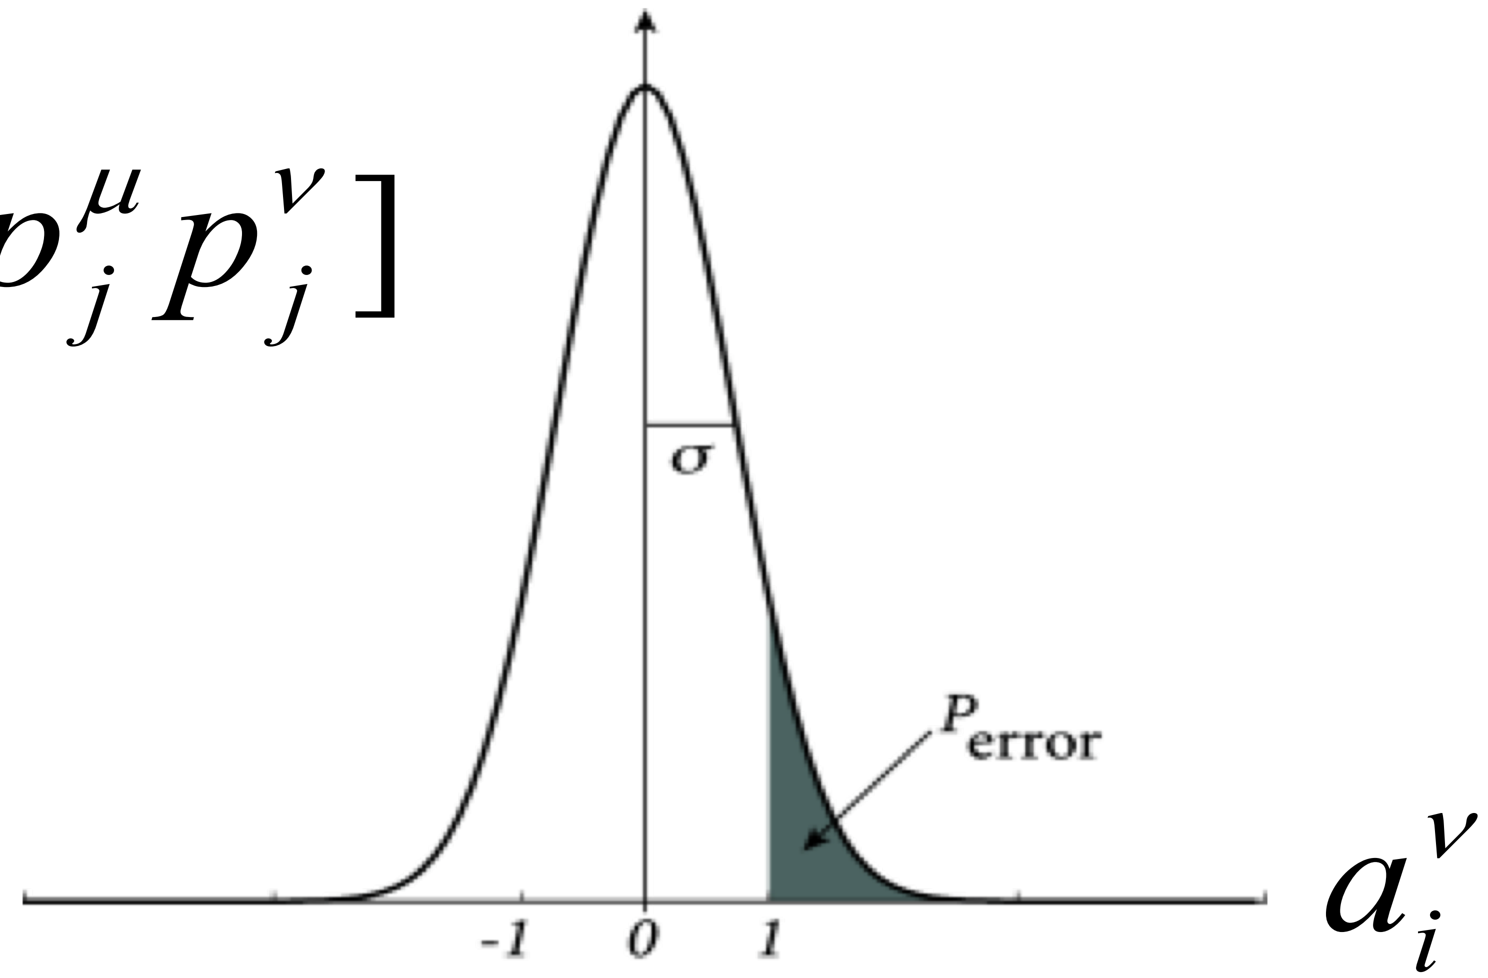

Image: Neuronal Dynamics, Gerstner et al., Cambridge Univ. Press (2014),

This week: Understand Associative Memory

Partial information → Associative memory/ collective computation → Full concept

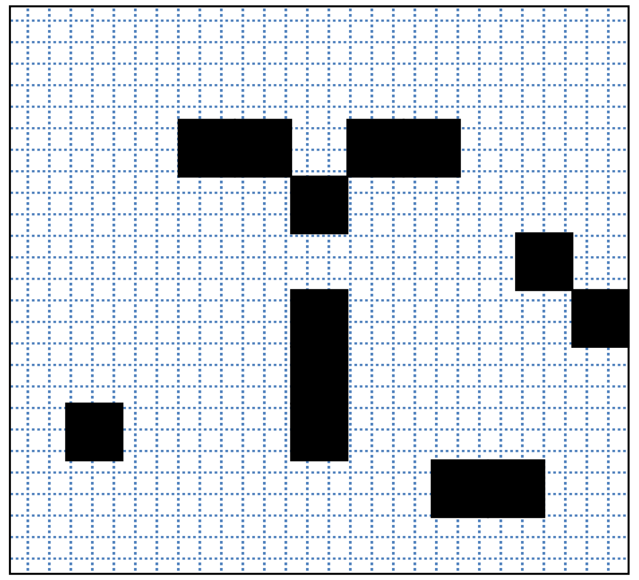

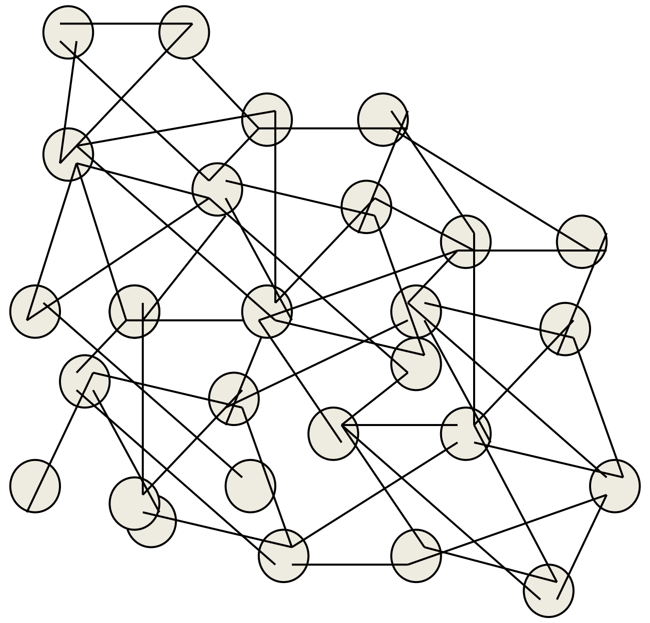

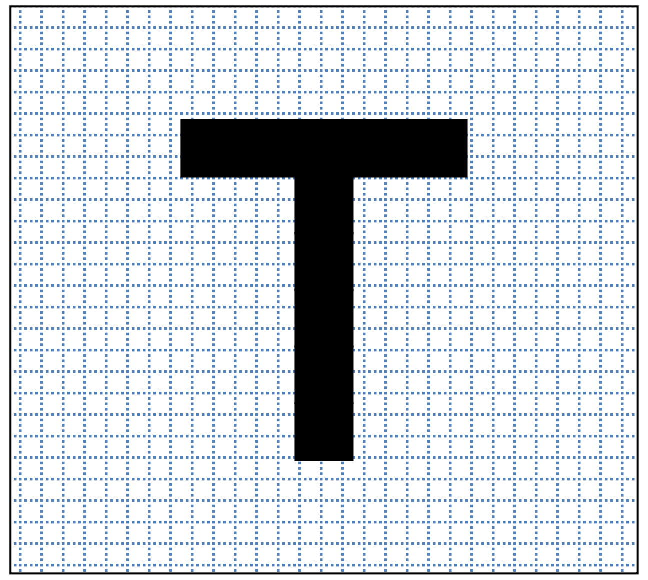

Brain-style computation

- Memory stored in connections
- Many memories can be stored in same network
- Retrieval of memories without centralized controller
- Interactions of neurons makes network converge to most similar pattern

References: Associative Memory Models

D. J. Willshaw, O. P. Bunemann and H. C. Longuet-Higgins (1969)
Non-holographic associative memory. Nature 222, pp. 960–962

J. A. Anderson (1972)
A simple neural network generating an interactive memory.
Math. Biosc. 14, pp. 197–220

T. Kohonen (1972)
Correlation matrix memories. IEEE trans. comp. C-21,
pp. 353–359.

W. A. Little (1974)
The existence of persistent states in the brain.
Math. Biosc. 19, pp. 101–120.

J.J. Hopfield (1982) Neural networks and physical
systems with emergent collective computational abilities.
Proc. Natl. Acad. Sci. USA 79, pp. 2554–2558

The end

Documentation:
http://neuronaldynamics.epfl.ch/

Online html version available

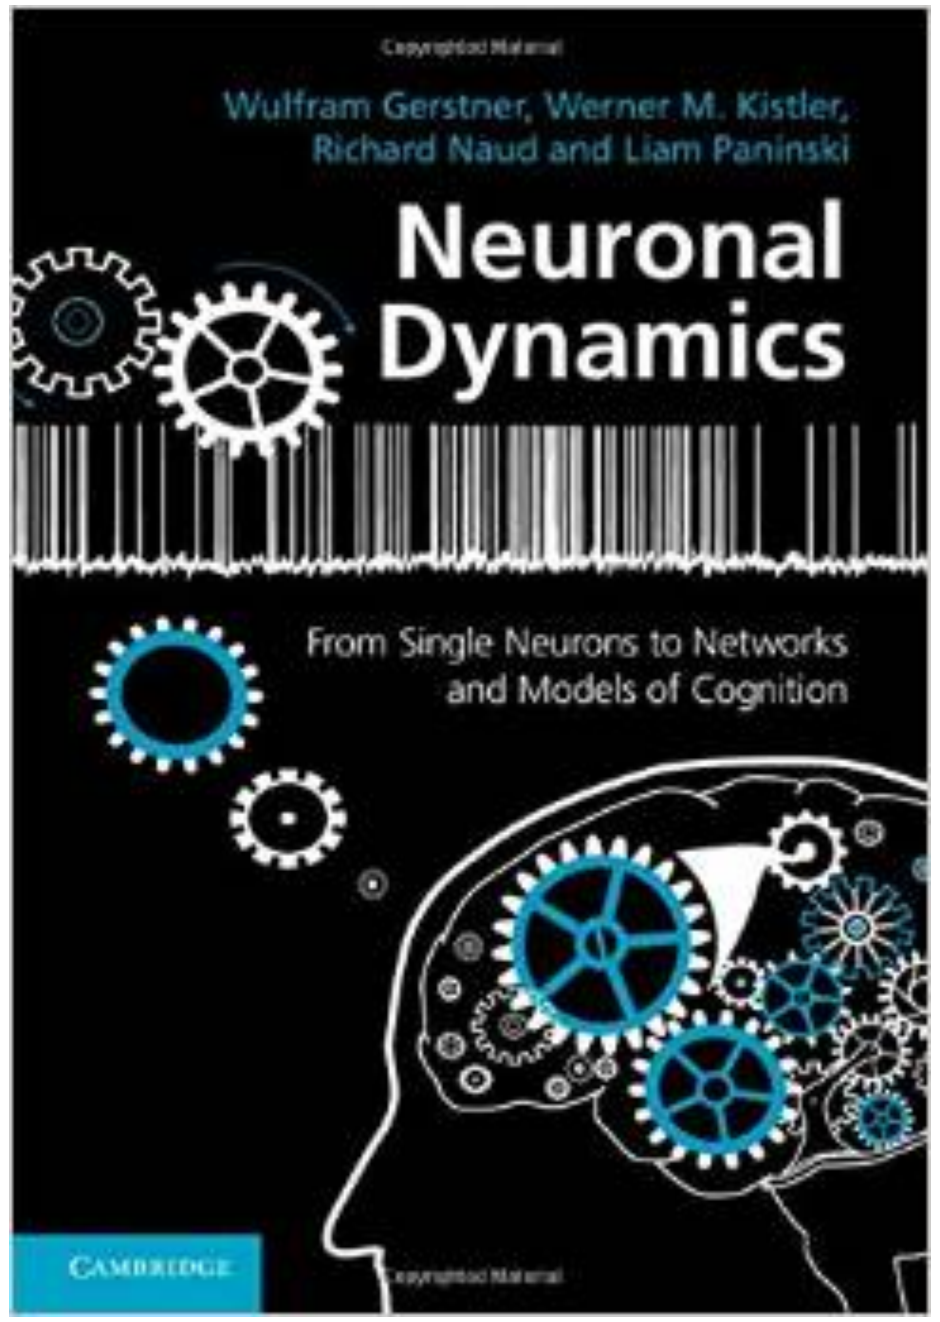

Reading for this week:
NEURONAL DYNAMICS
- Ch. 17.1 - 17.2.4

Cambridge Univ. Press