Stable learning in networks of unreliable, memristive nanodevices

Adaptive Massively Parallel Computing Workshop

March 2-3, 2009
Portland, OR, USA

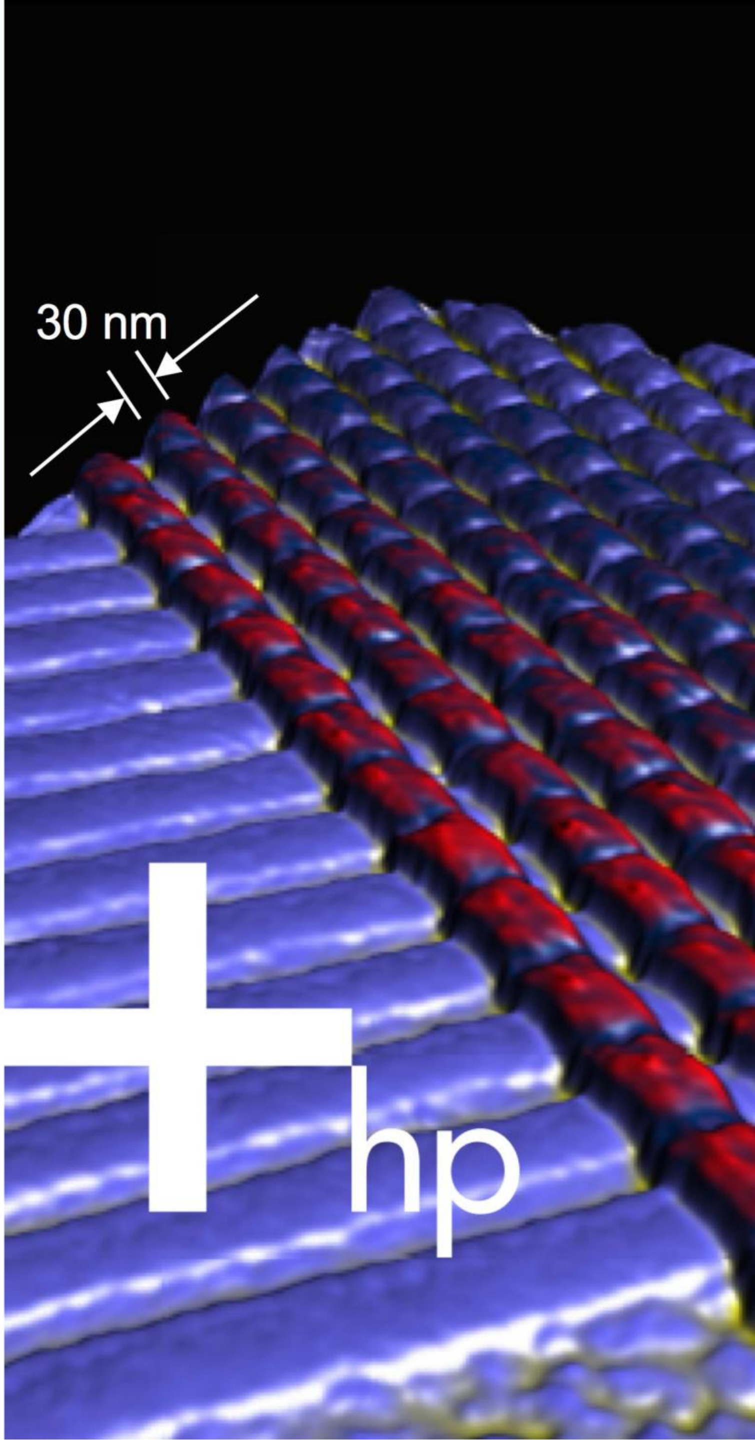

Greg Snider
IQSL
Hewlett-Packard Laboratories

Outline

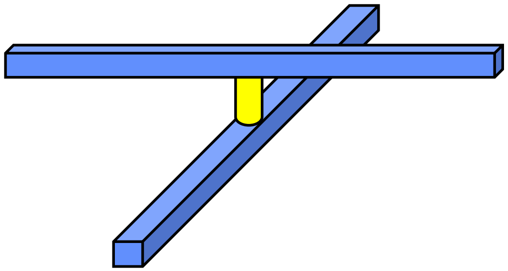

1. **Memristive nanodevices**

   (synapses)

Outline

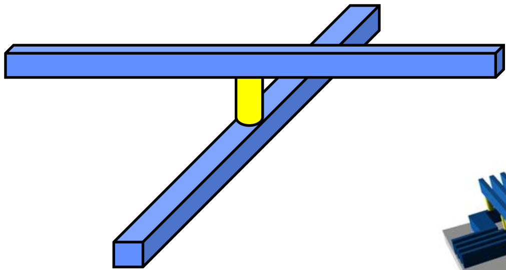

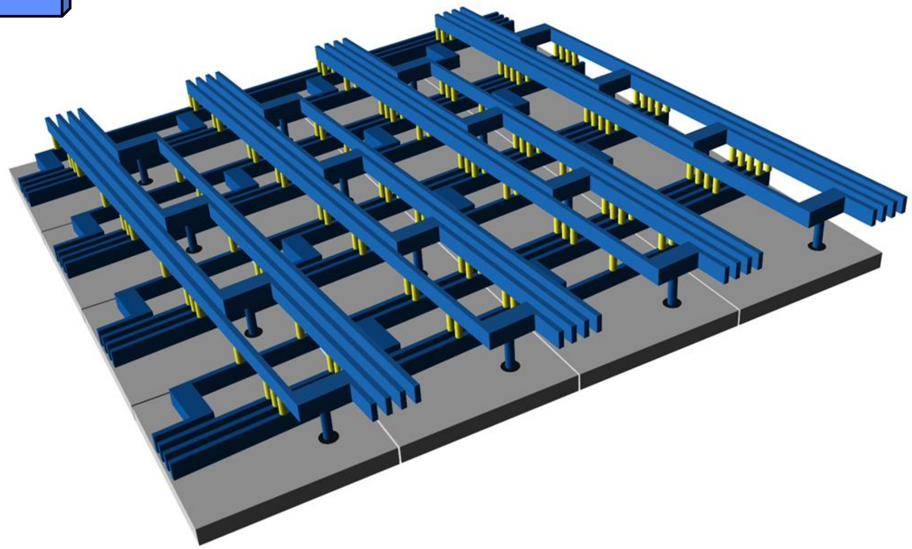

1. Memristive nanodevices

(synapses)

2. Adaptive resonance networks

(nano + CMOS)

Outline

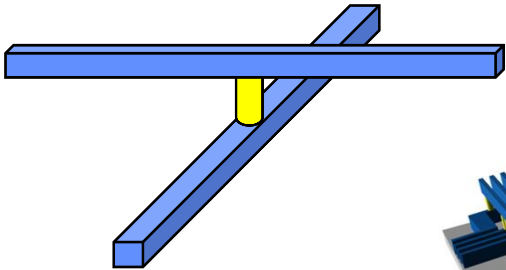

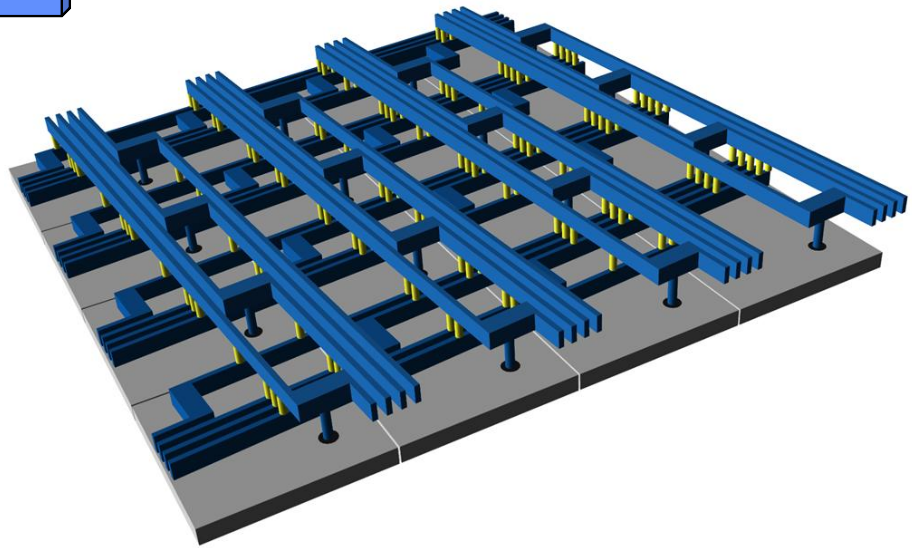

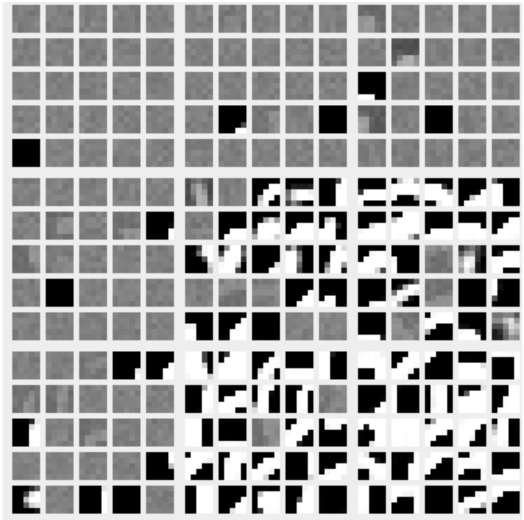

1. Memristive nanodevices

(synapses)

2. Adaptive resonance networks

(nano + CMOS)

3. Simulations

(fast, stable, incremental learning)

Orthogonal nanowire layers separated by memristive material

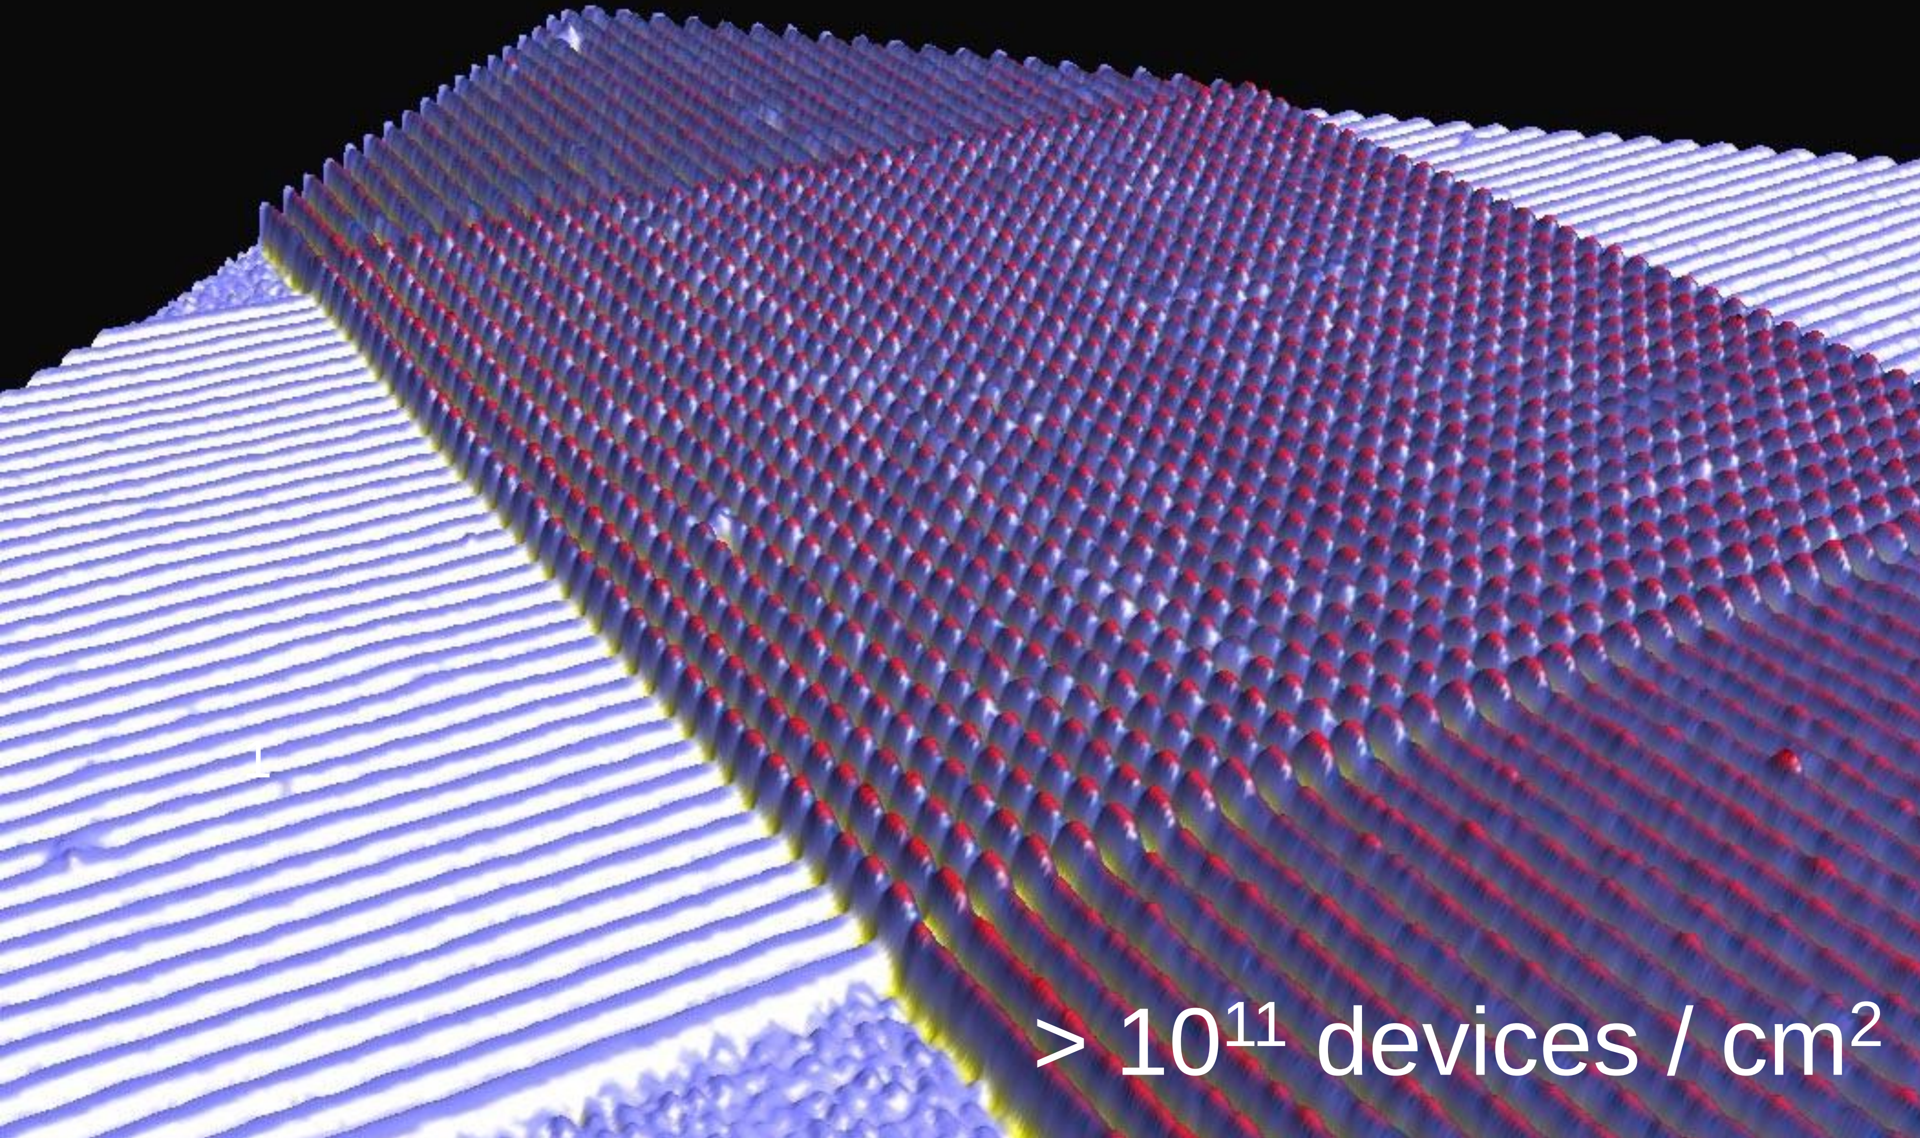

$> 10^{11}$ devices / cm^2

Definition: Memristive Device

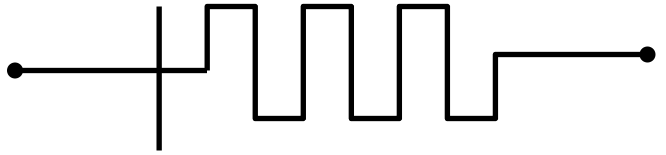

schematic symbol

Definition: Memristive Device

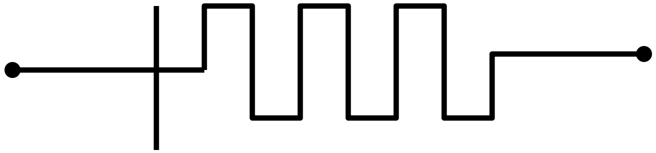

$$\frac{d\mathbf{w}}{dt} = f(\mathbf{w}, v)$$

$$i = g(\mathbf{w}, v)v$$

or

$$\frac{d\mathbf{w}}{dt} = f(\mathbf{w}, i)$$

$$v = r(\mathbf{w}, i)i$$

voltage-controlled

current-controlled

Definition: Memristive Device

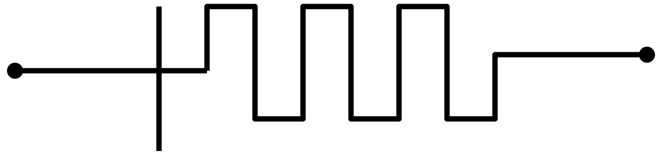

state variables

$$\frac{d\mathbf{w}}{dt} = f(\mathbf{w}, v)$$

$$i = g(\mathbf{w}, v)v$$

current

conductance

voltage

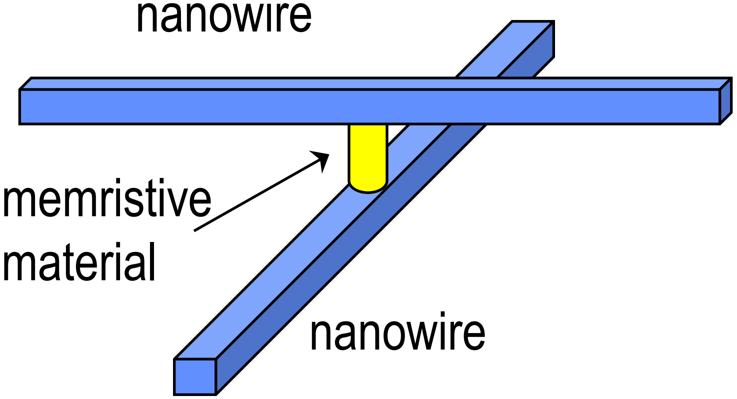

implementation

The exquisite nonlinearity of *dw/dt*

$$\frac{d\mathbf{w}}{dt} = f(\mathbf{w}, v)$$

$$i = g(\mathbf{w}, v)v$$

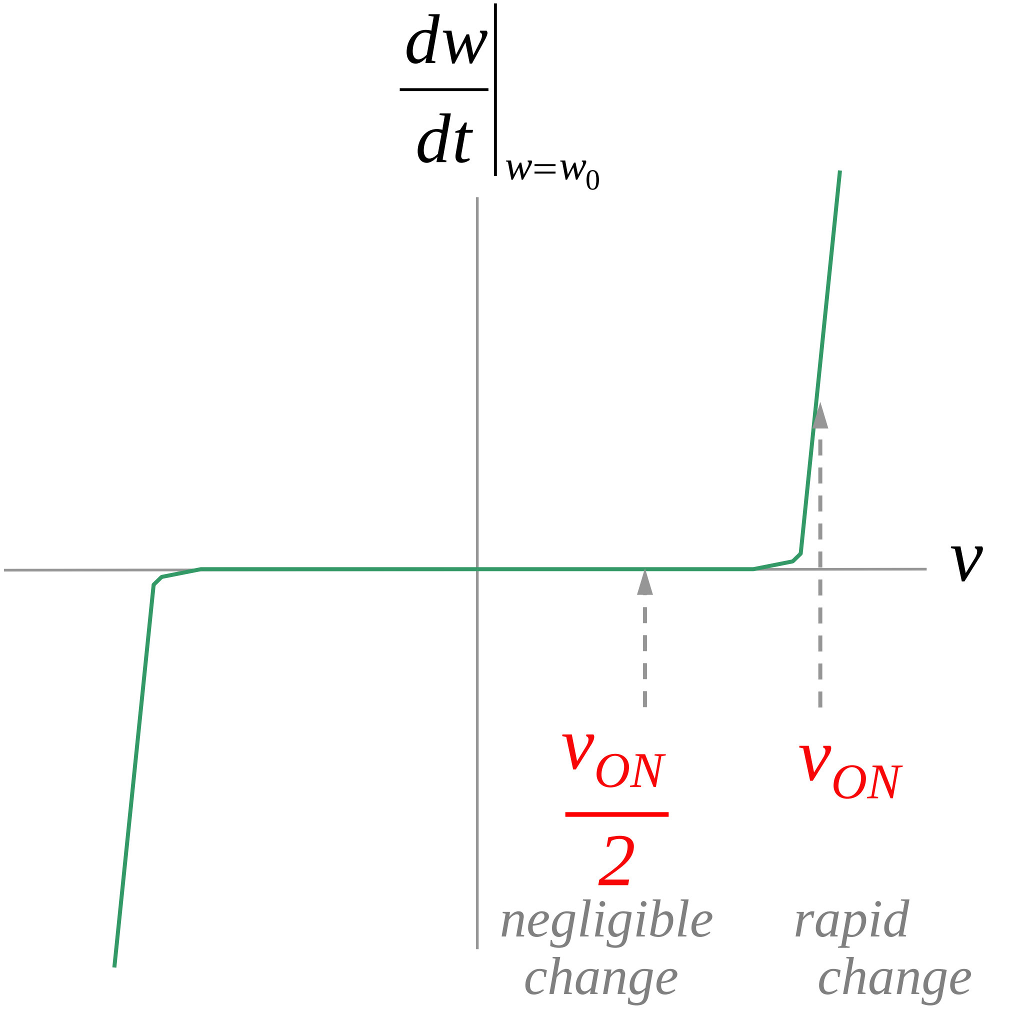

Experimental data

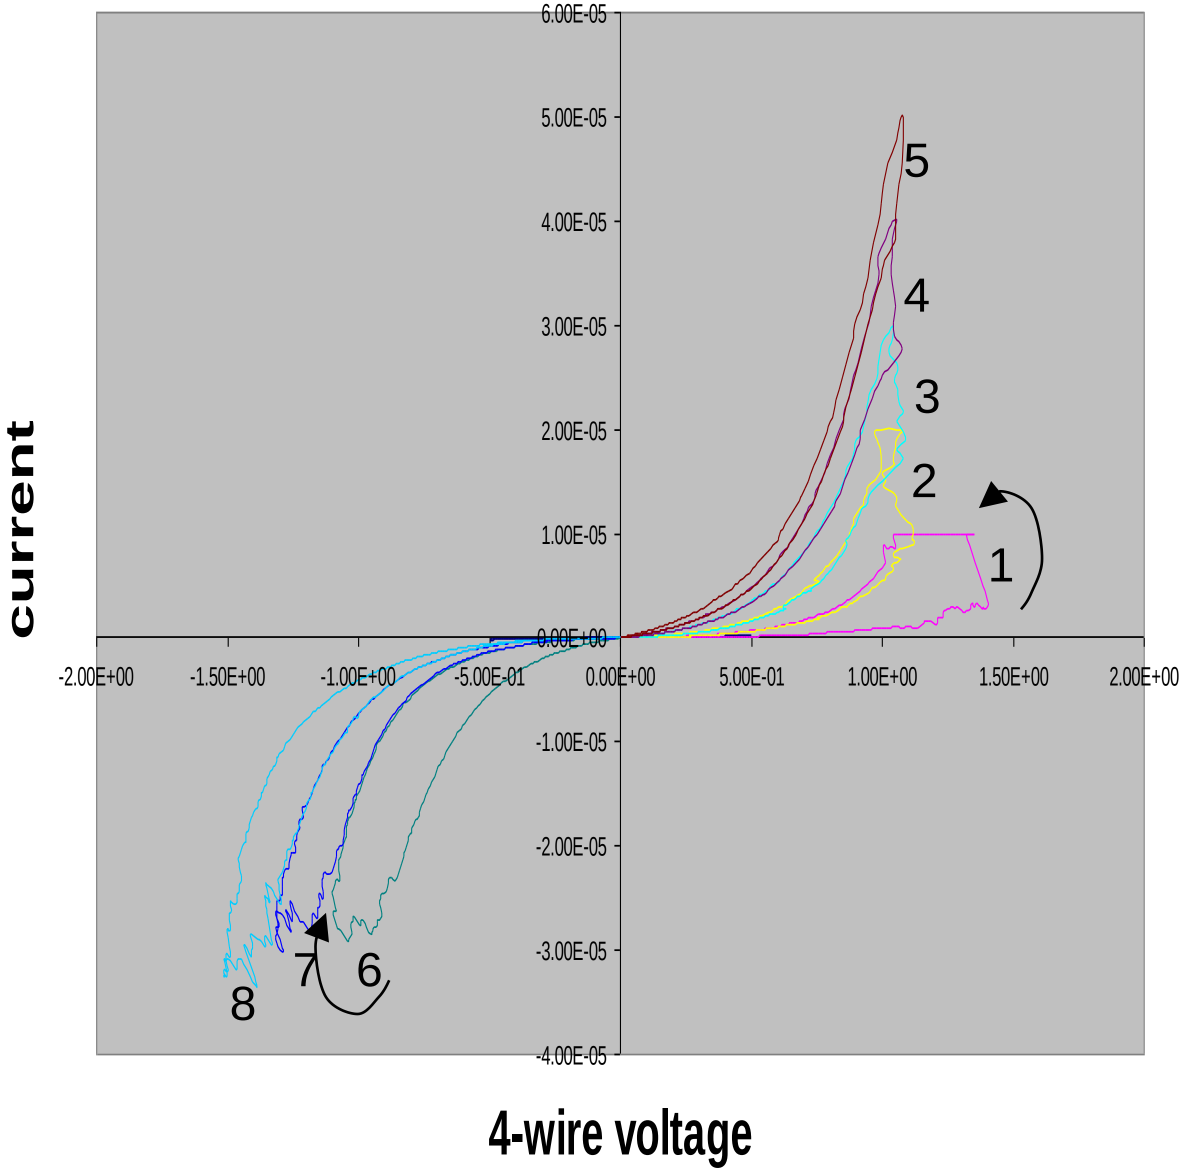

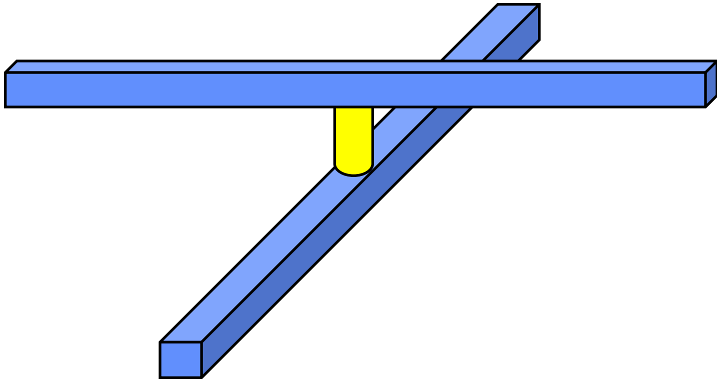

Constructive interference (Hebbian learning)

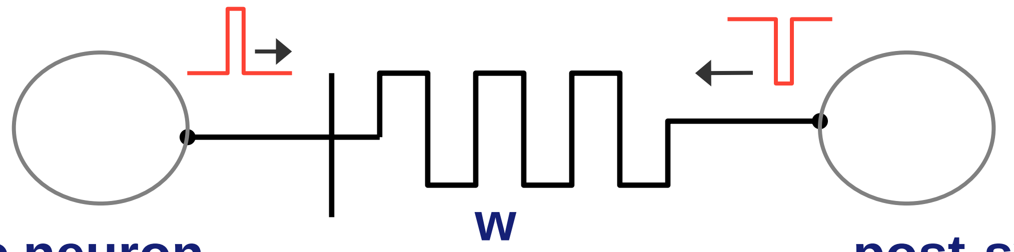

pre-synaptic neuron

post-synaptic neuron

Constructive interference (Hebbian learning)

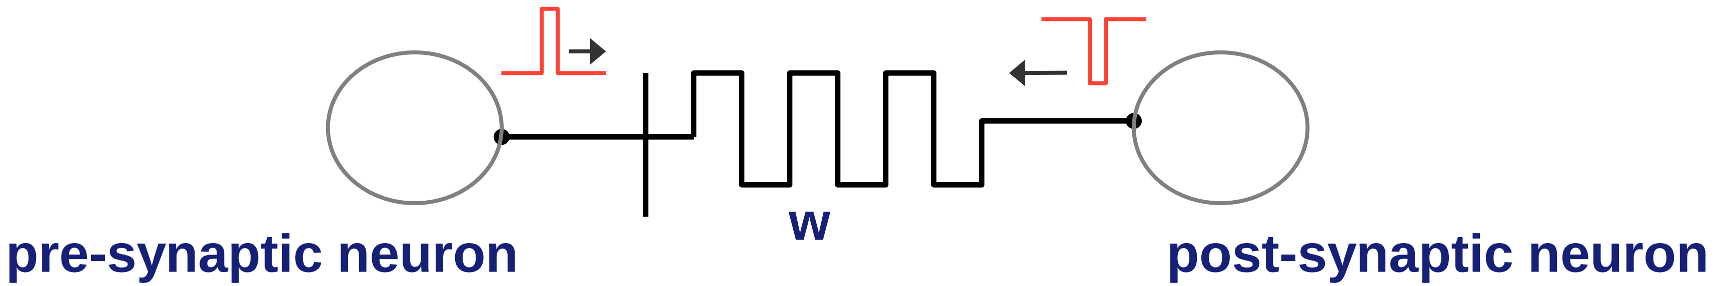

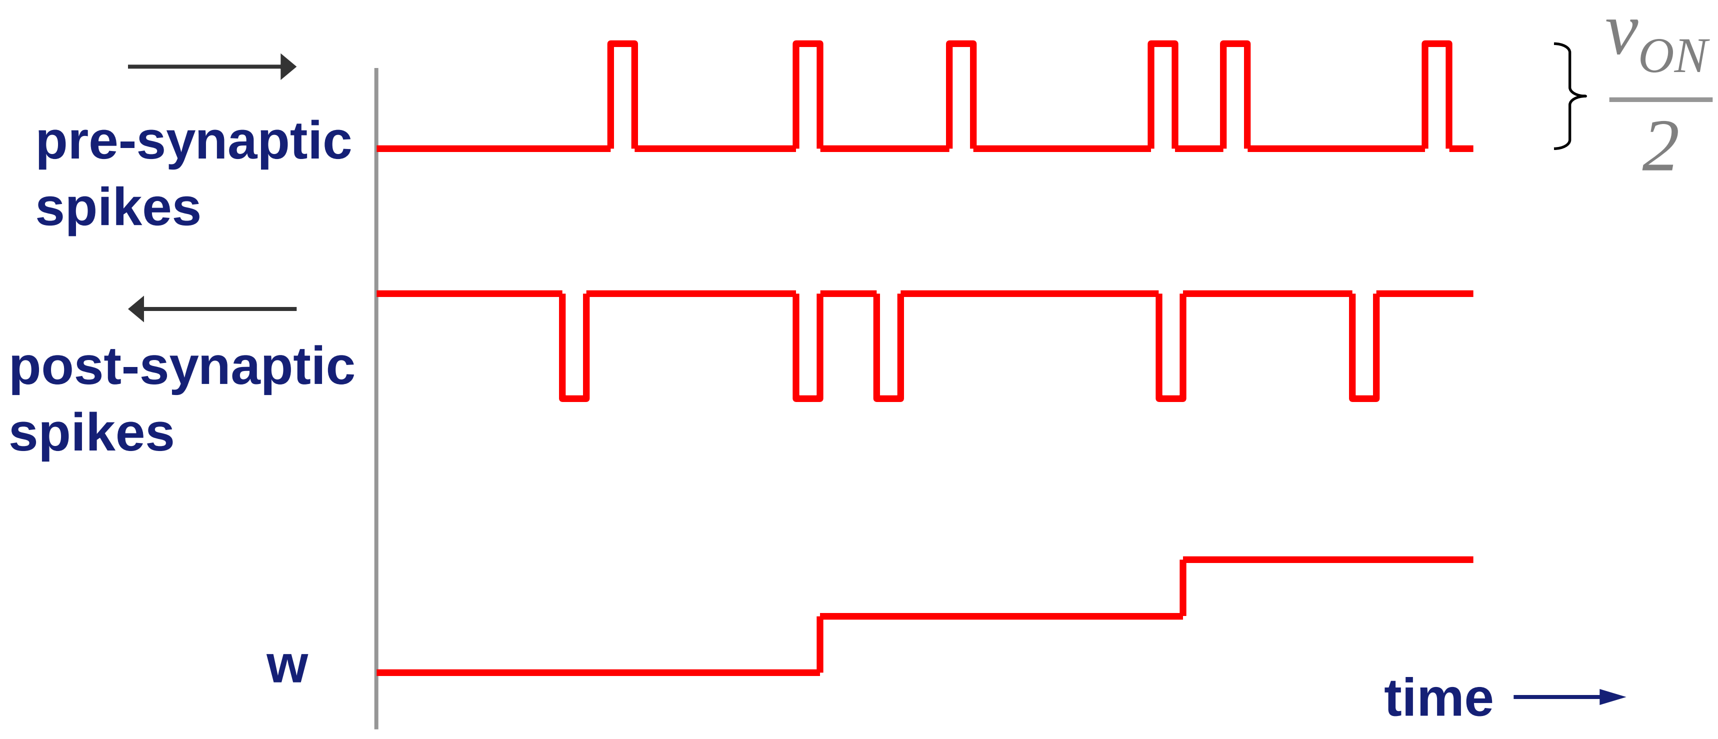

ART networks

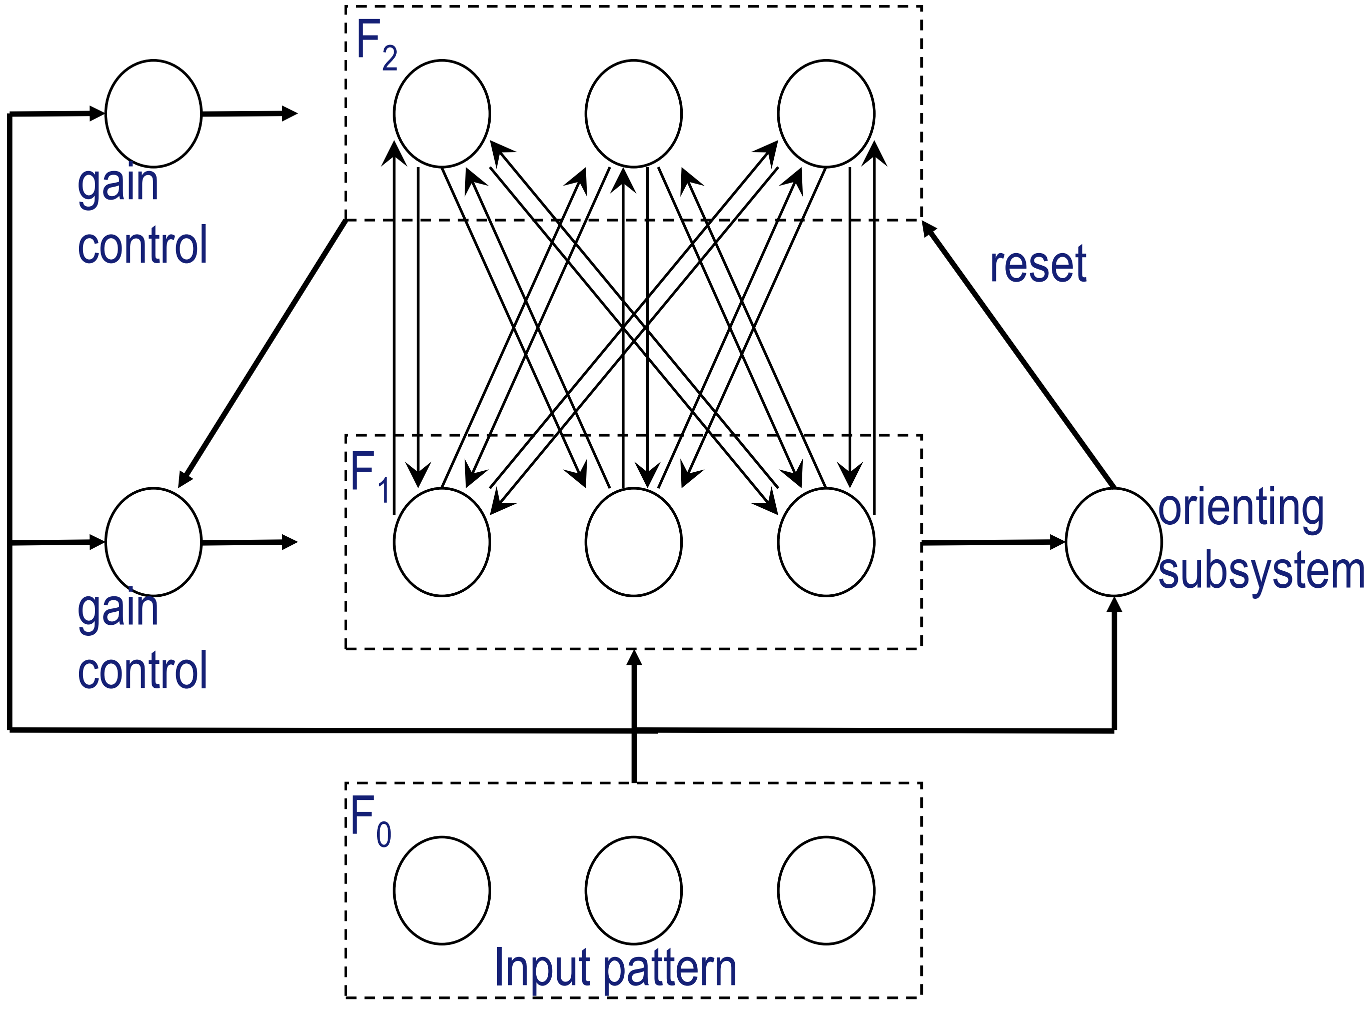

Carpenter and Grossberg, “A Massively Parallel Architecture for a Self-organizing Neural Pattern Recognition Machine,” Computer Vision, Graphics, and Image Processing, 1987

Adaptive Resonance Oversimplified[1]

1. Initially no neurons in upper layer

watchdog
"O.S."

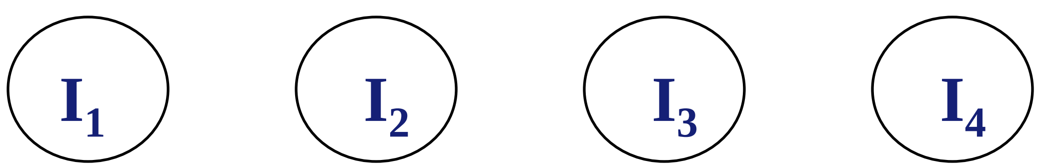

1. with apologies to Gail and Steve

Adaptive Resonance Oversimplified

1. Initially no neurons in upper layer.
2. Input pattern comes in.

watchdog
“O.S.”

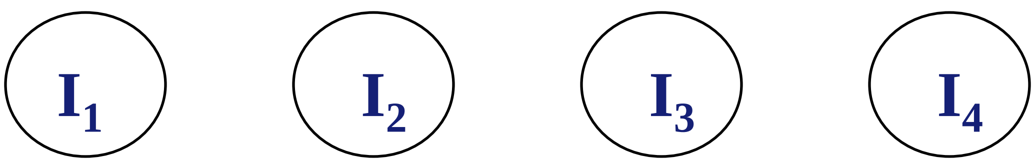

Adaptive Resonance Oversimplified

1. Initially no neurons in upper layer.
2. Input pattern comes in.

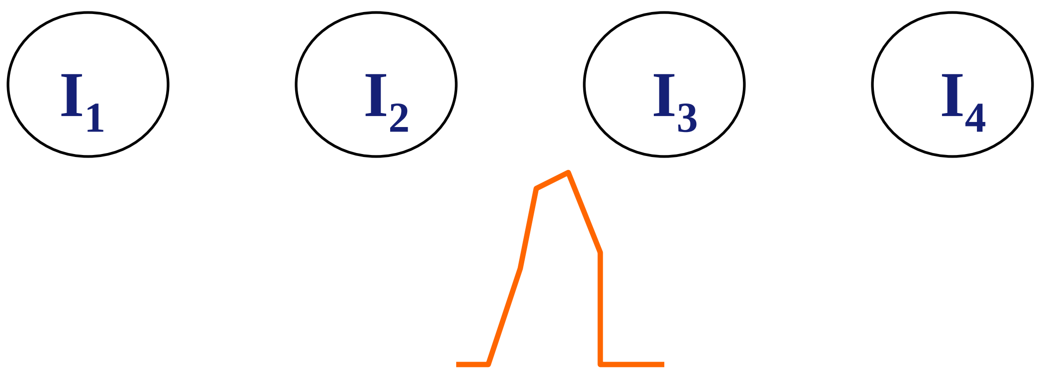

watchdog
"O.S."

Adaptive Resonance Oversimplified

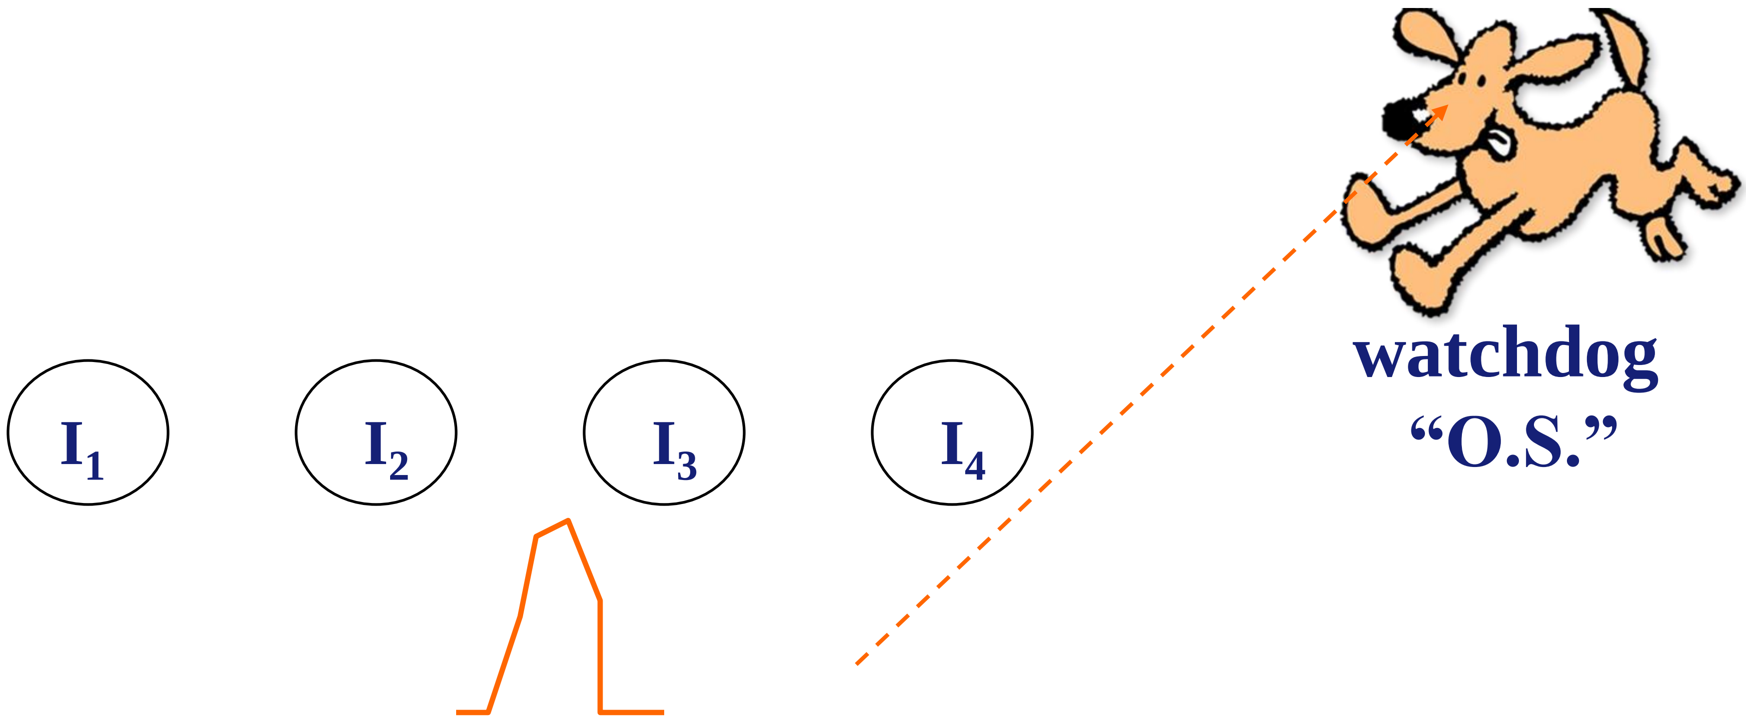

Adaptive Resonance Oversimplified

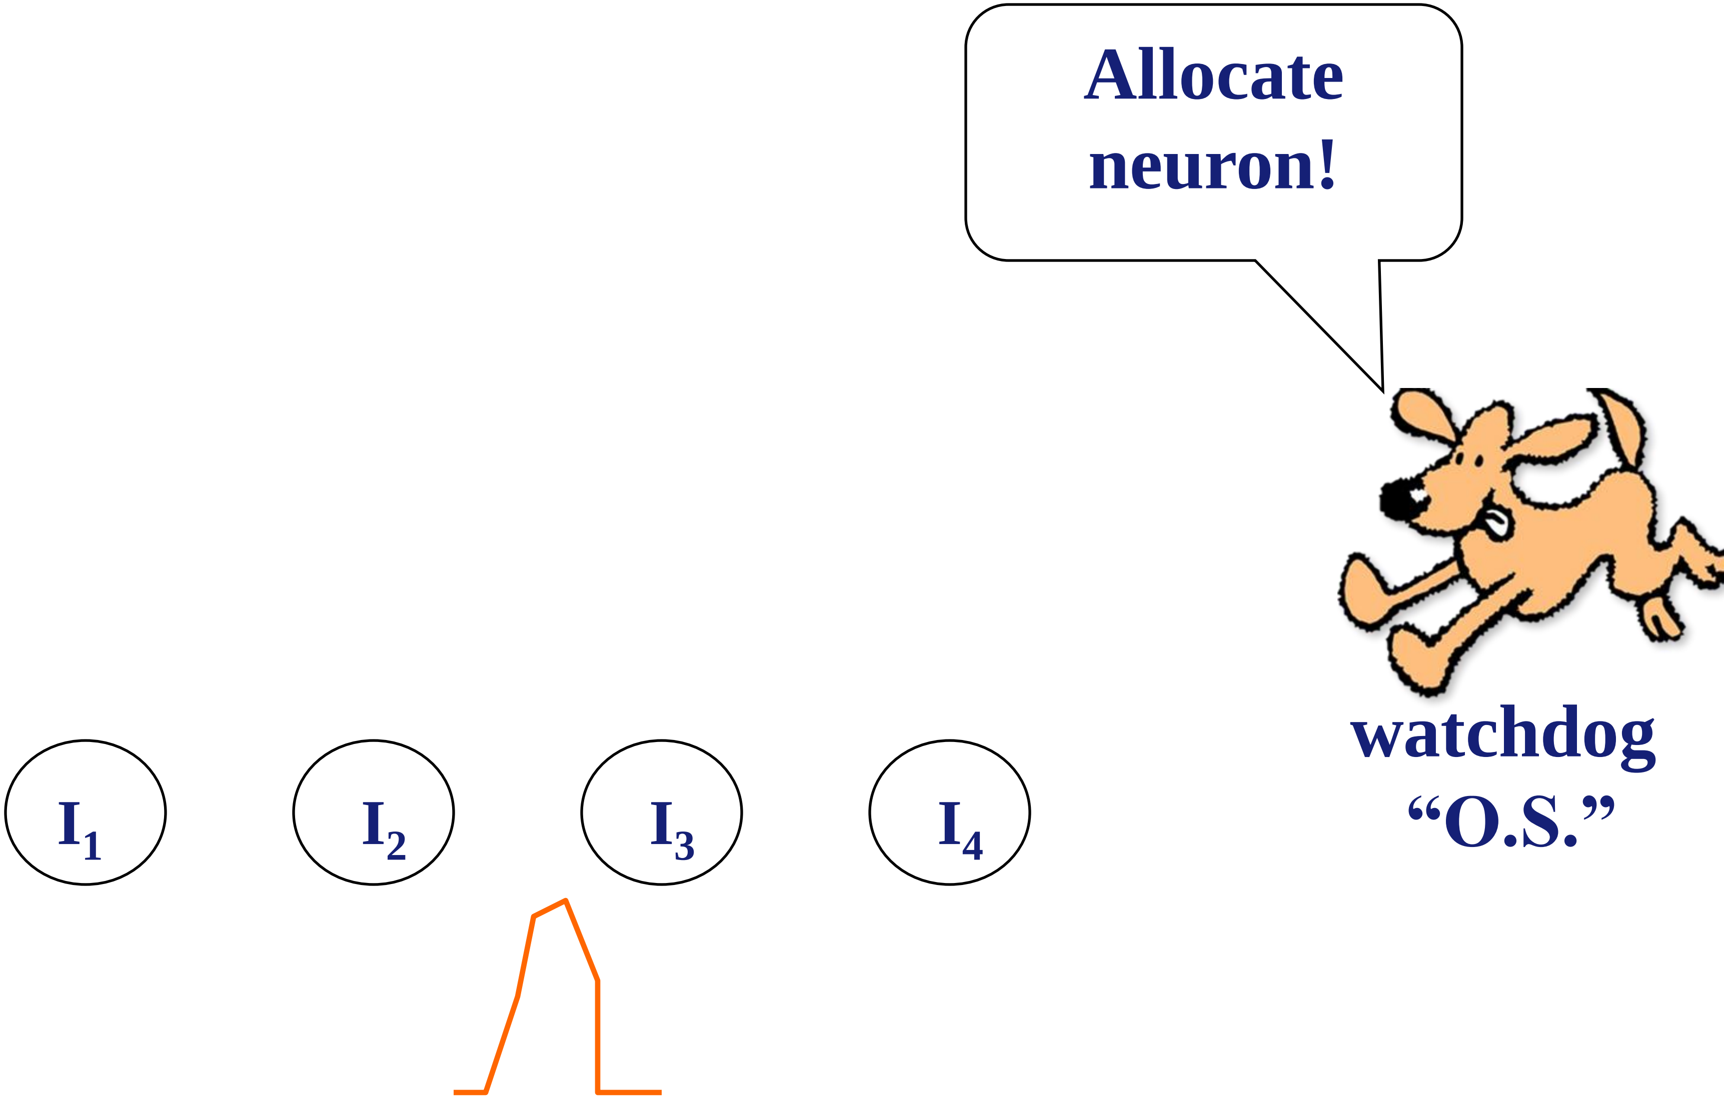

Adaptive Resonance Oversimplified

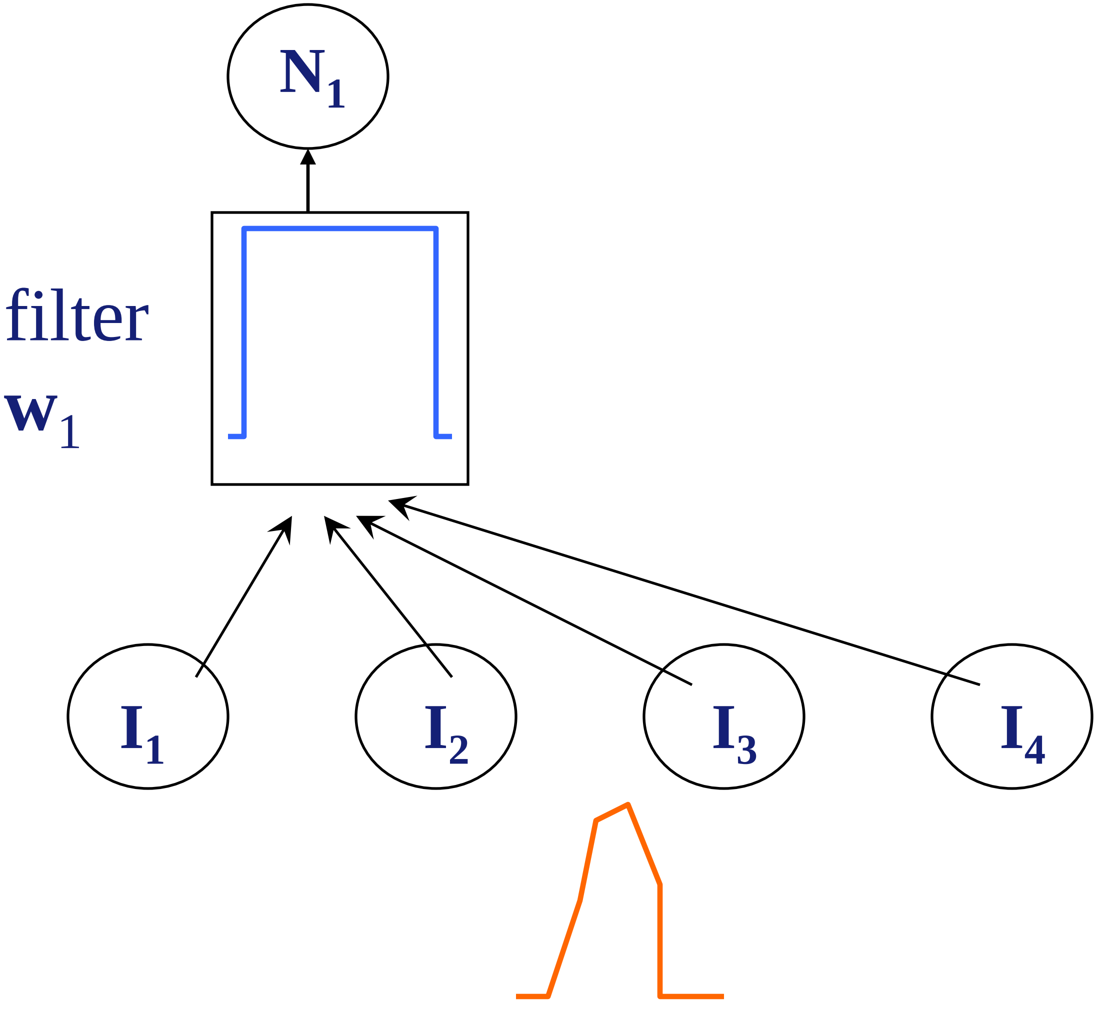

watchdog
"O.S."

Adaptive Resonance Oversimplified

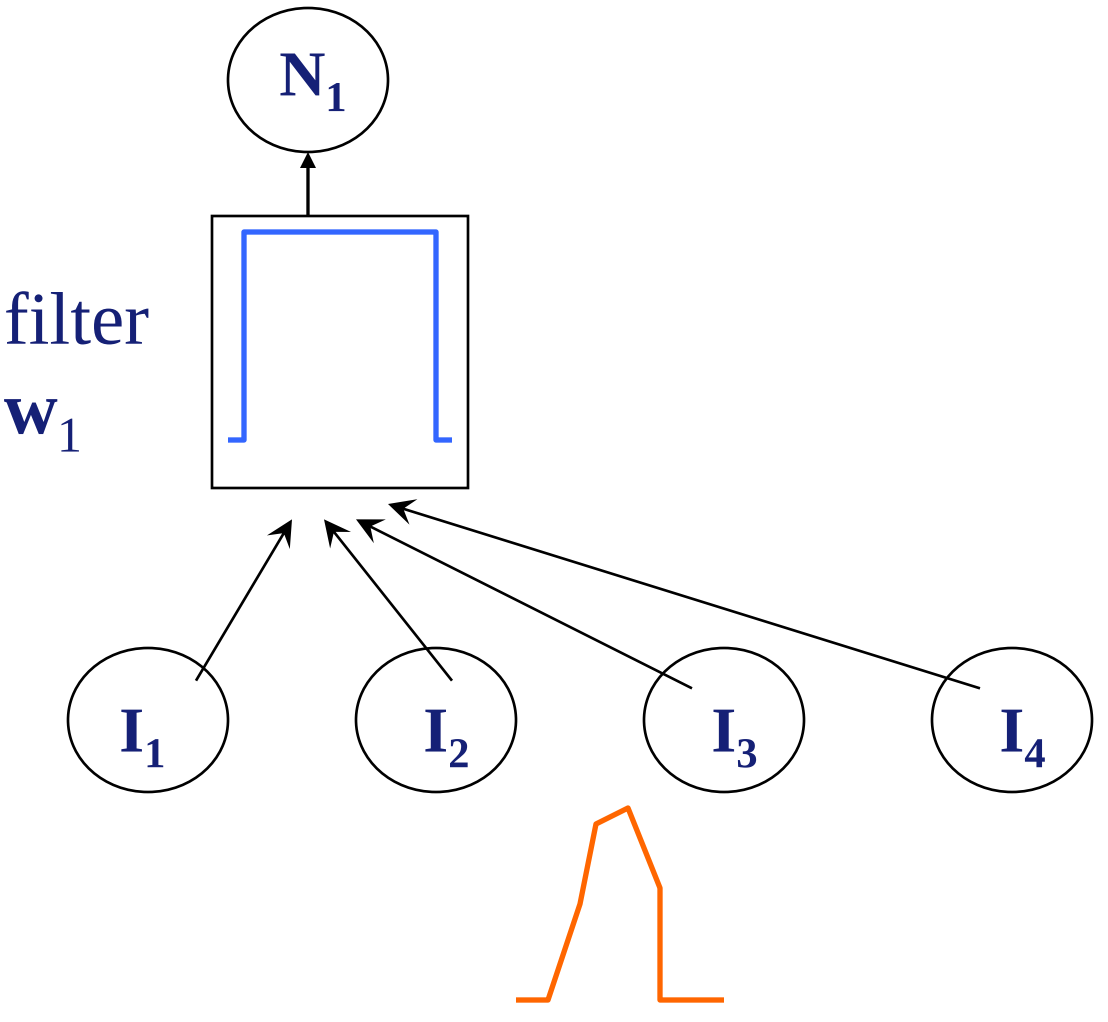

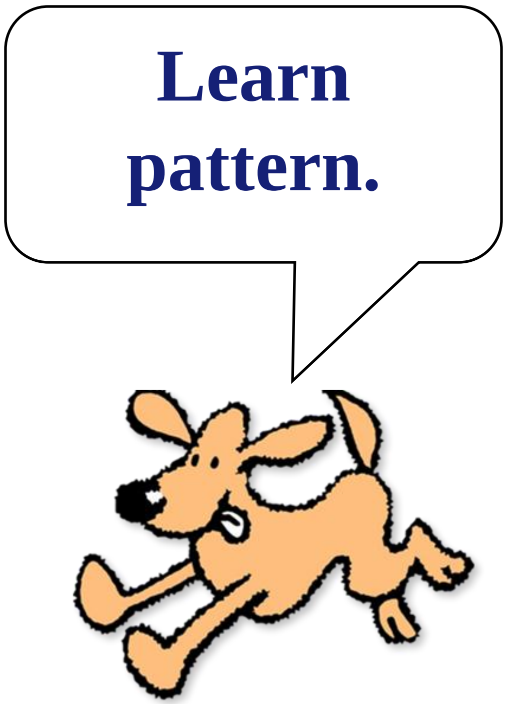

watchdog
“O.S.”

Adaptive Resonance Oversimplified

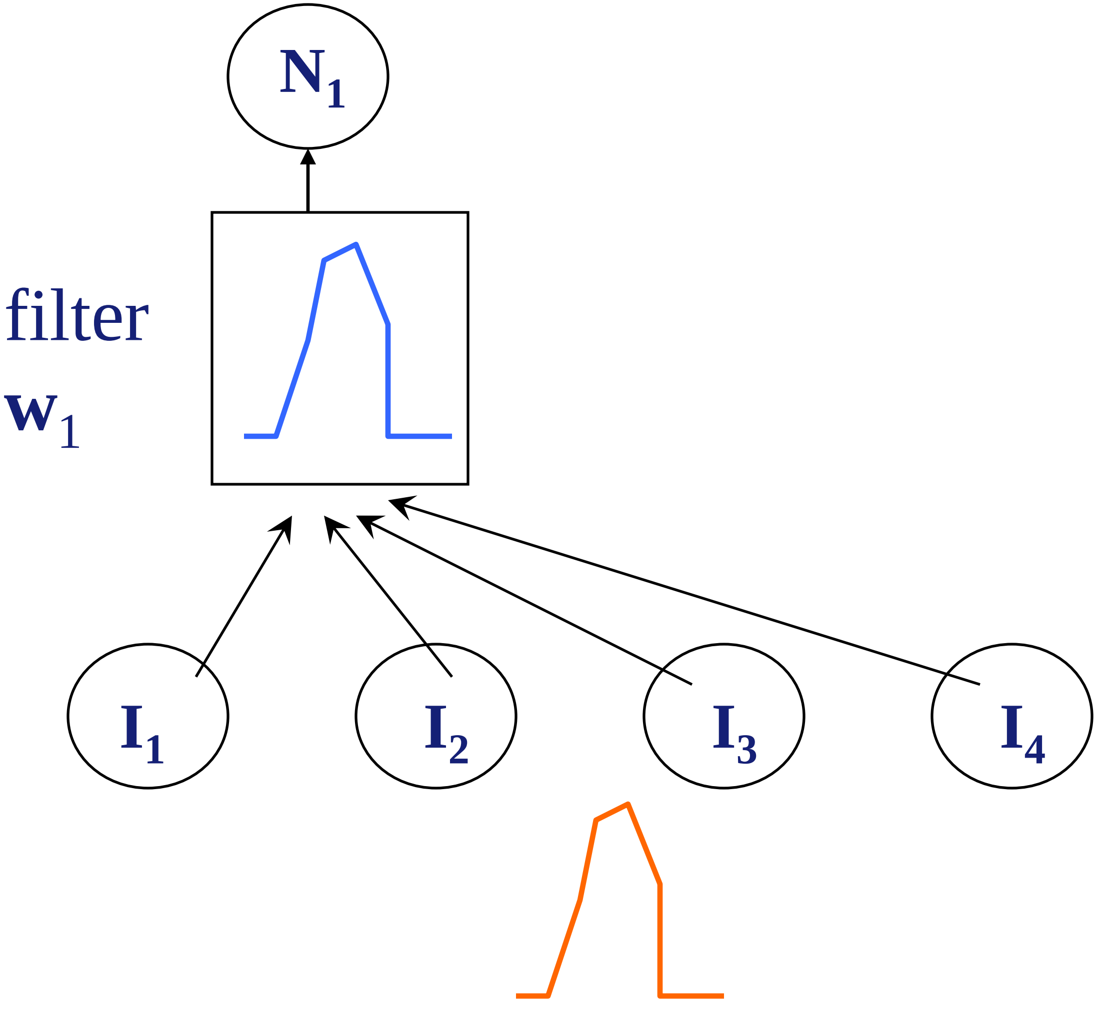

watchdog
"O.S."

Adaptive Resonance Oversimplified

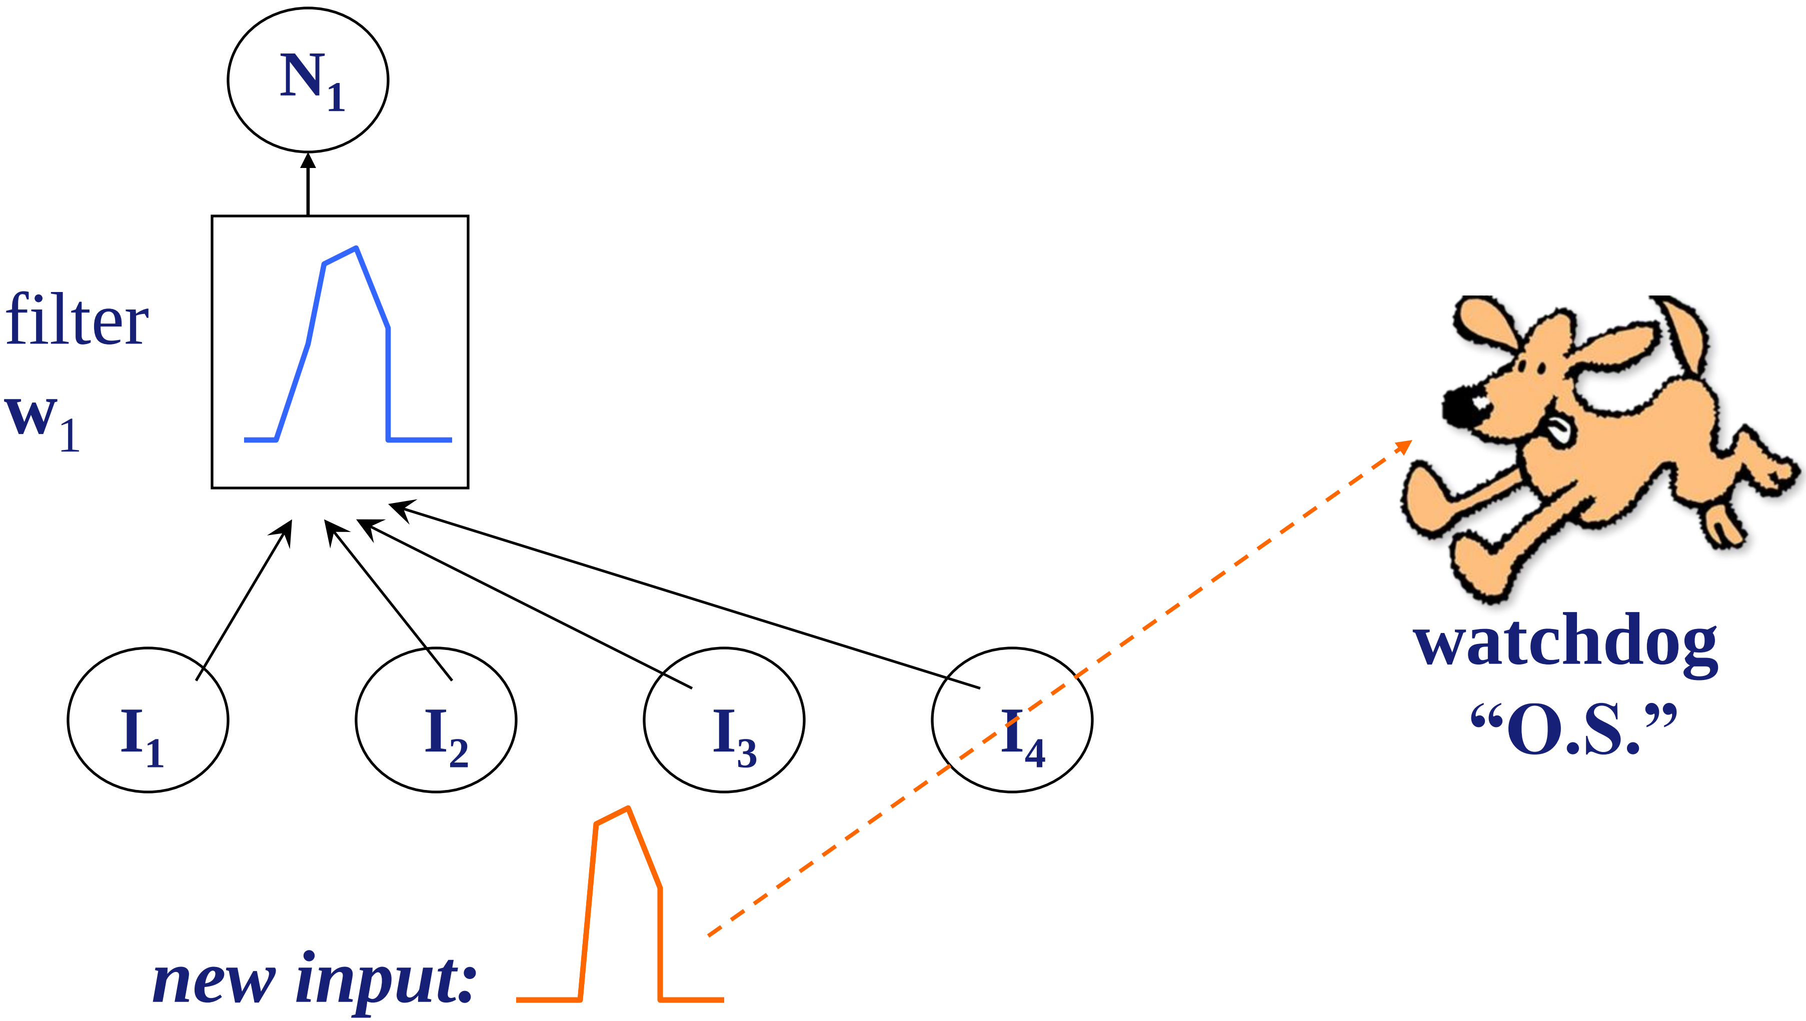

Adaptive Resonance Oversimplified

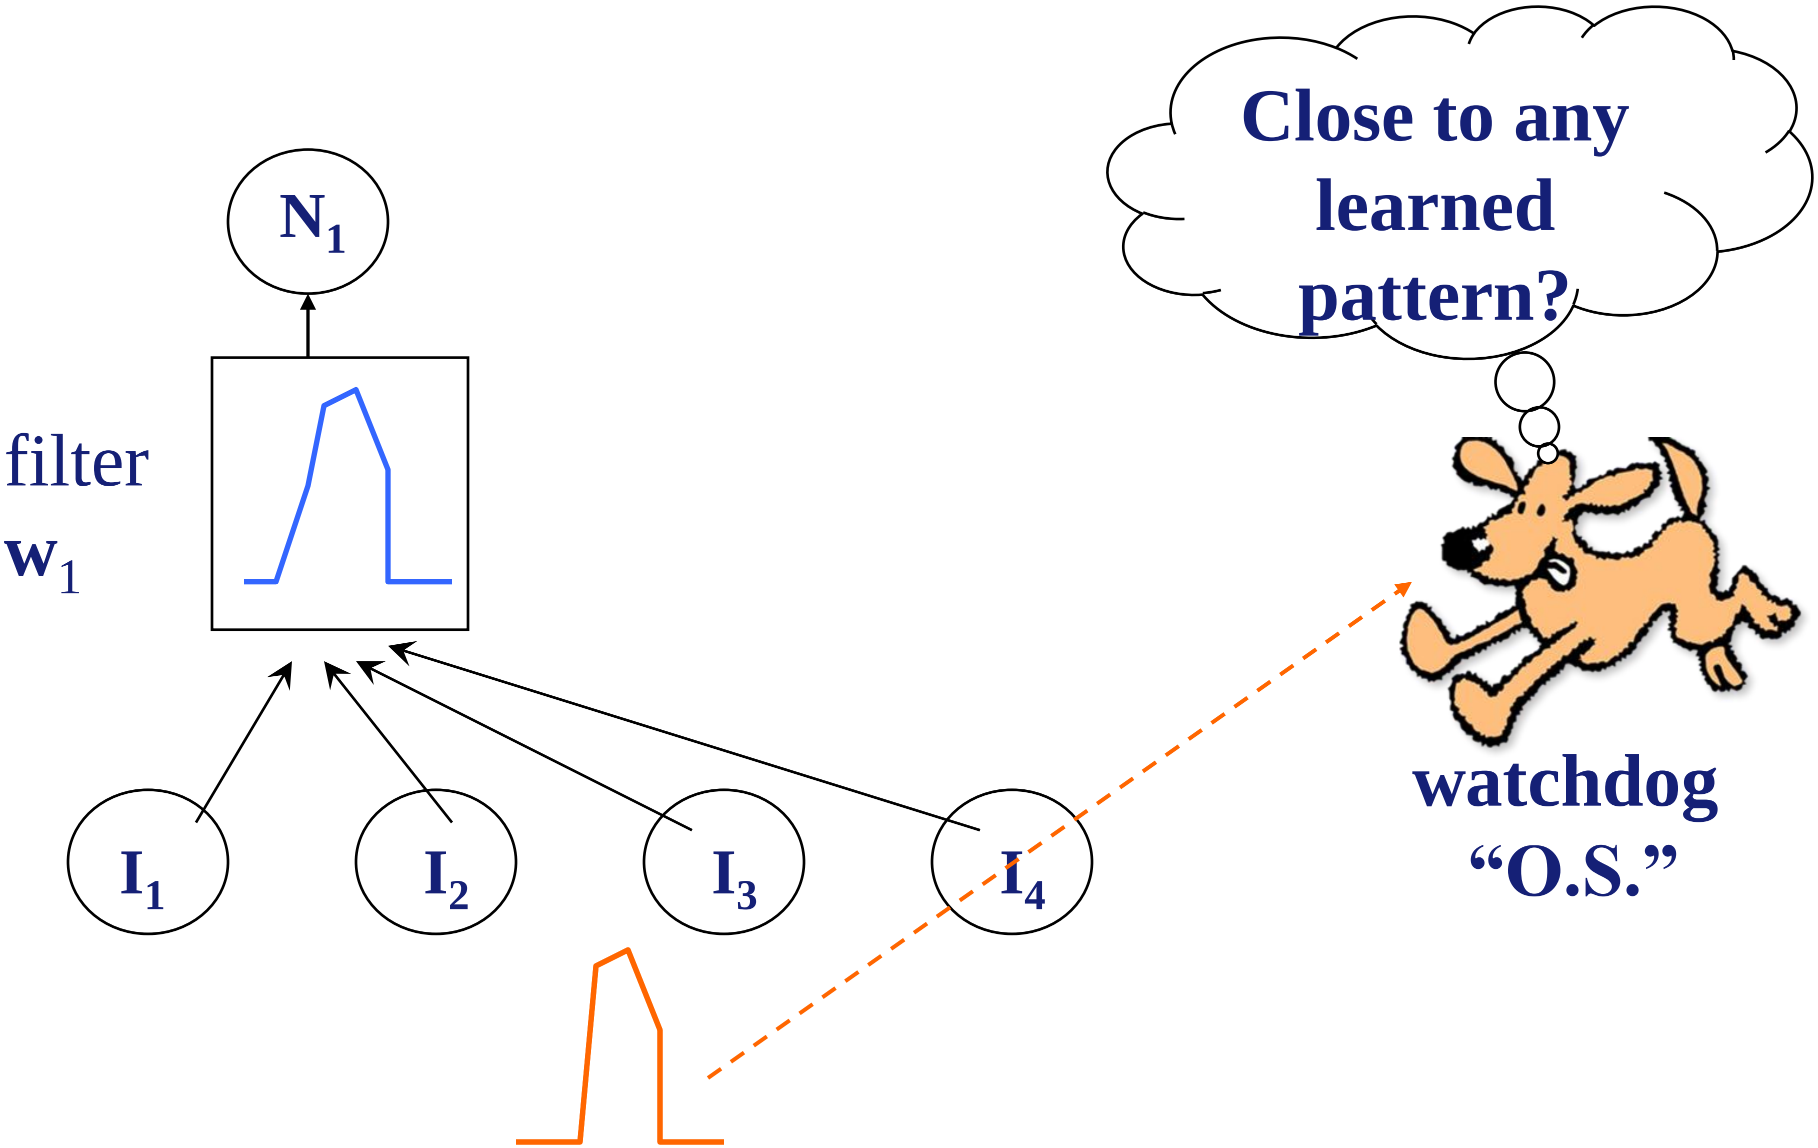

Adaptive Resonance Oversimplified

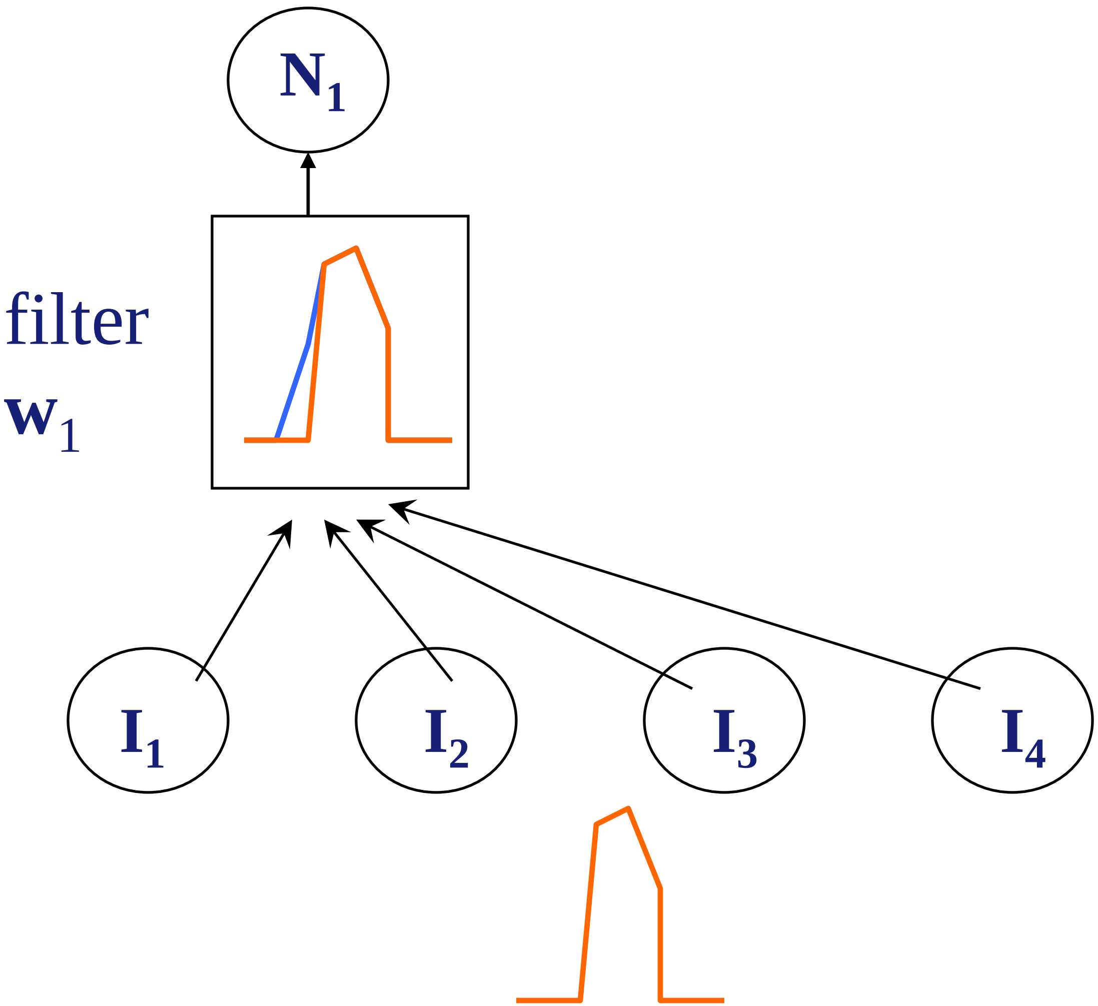

watchdog
"O.S."

Adaptive Resonance Oversimplified

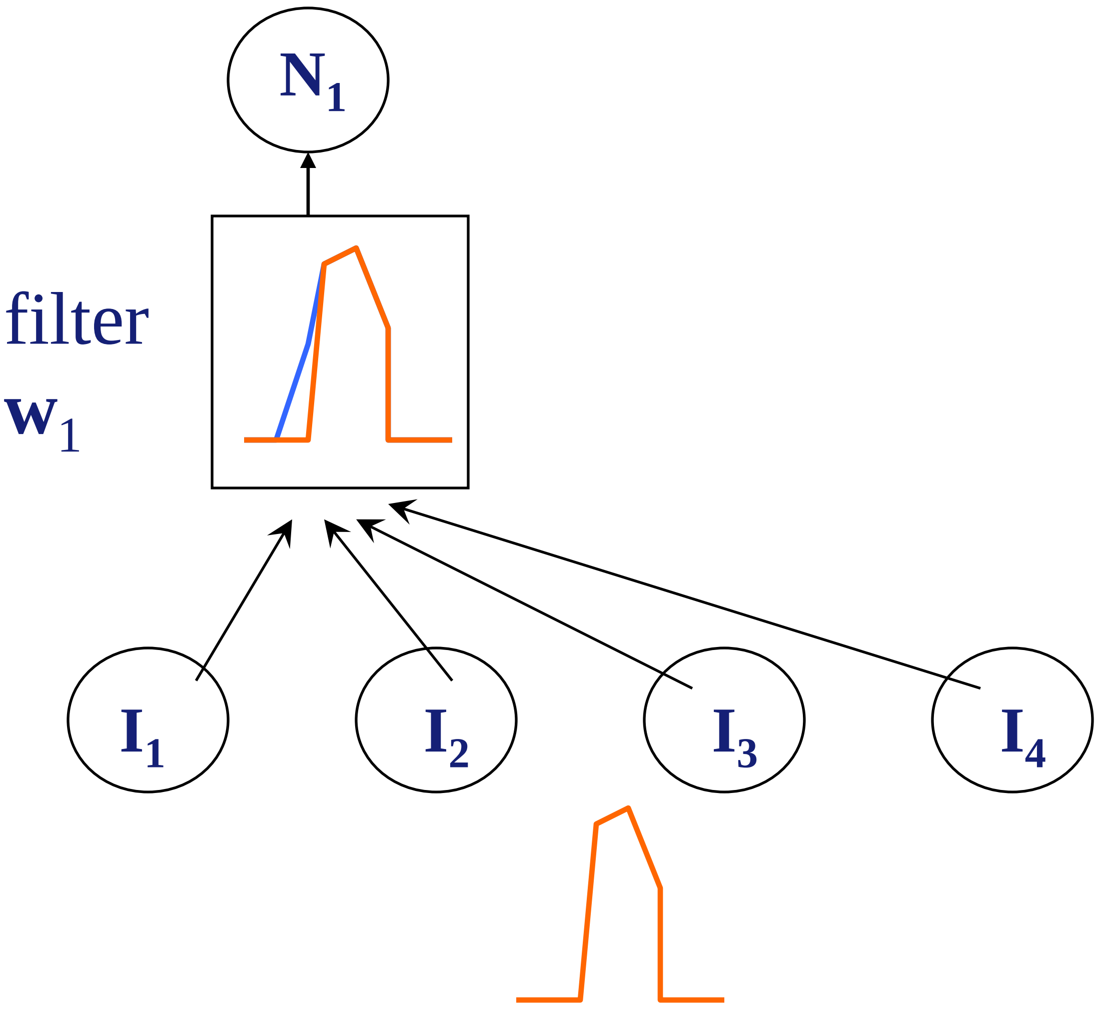

Adaptive Resonance Oversimplified

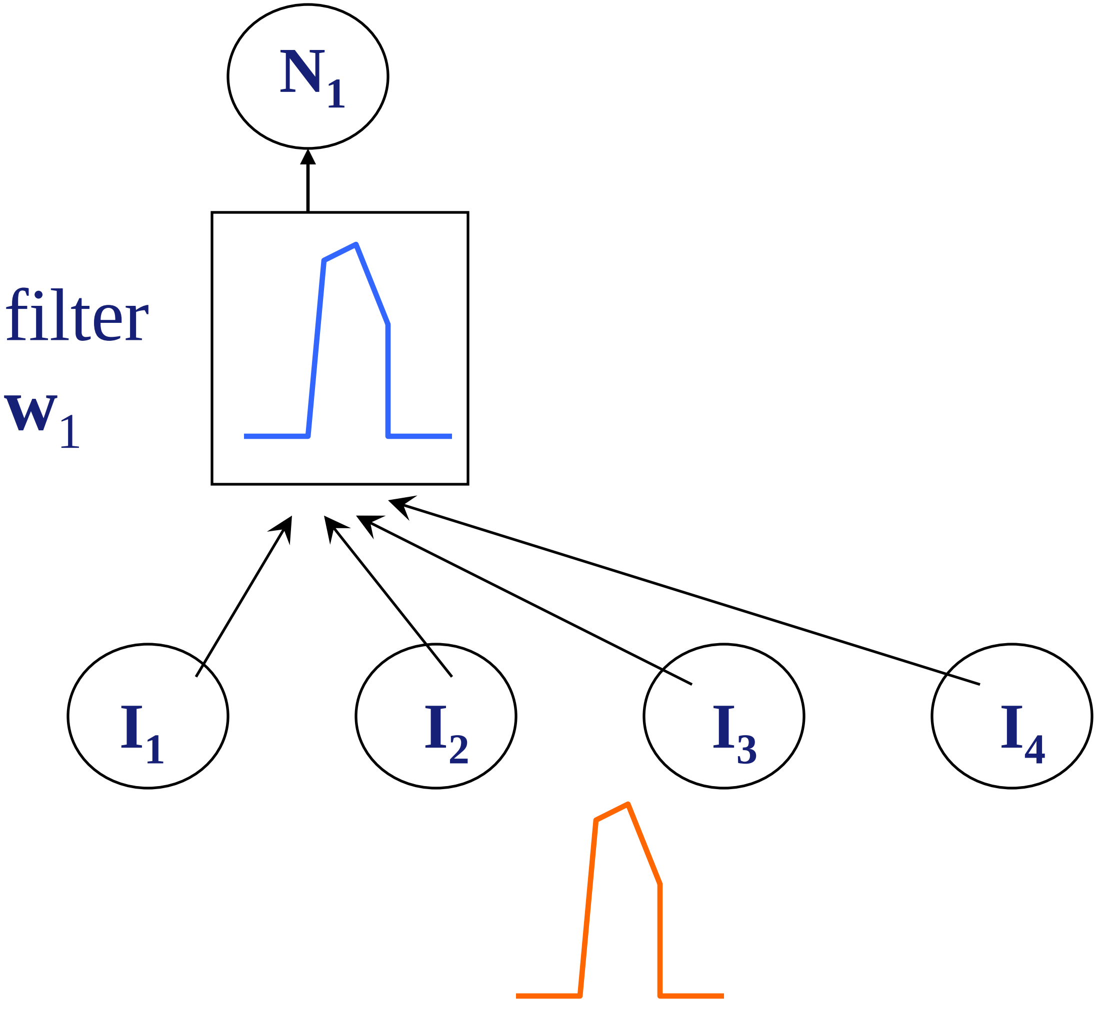

watchdog
"O.S."

Adaptive Resonance Oversimplified

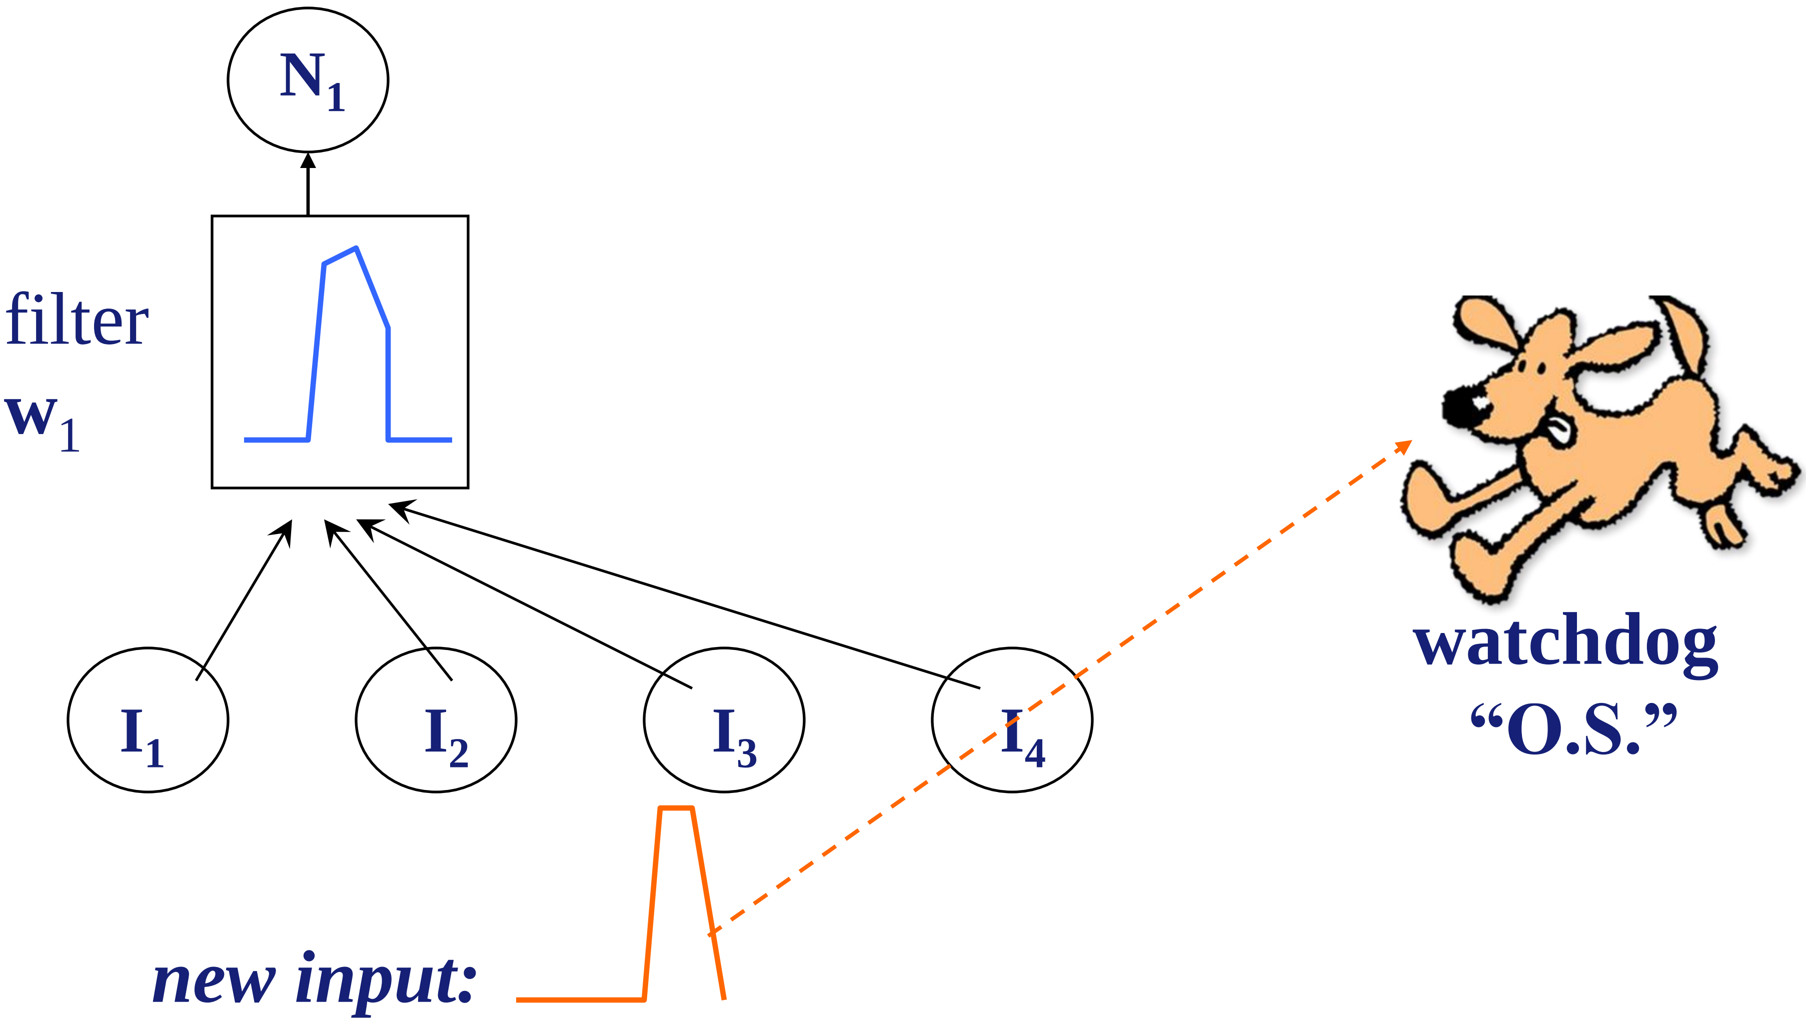

Adaptive Resonance Oversimplified

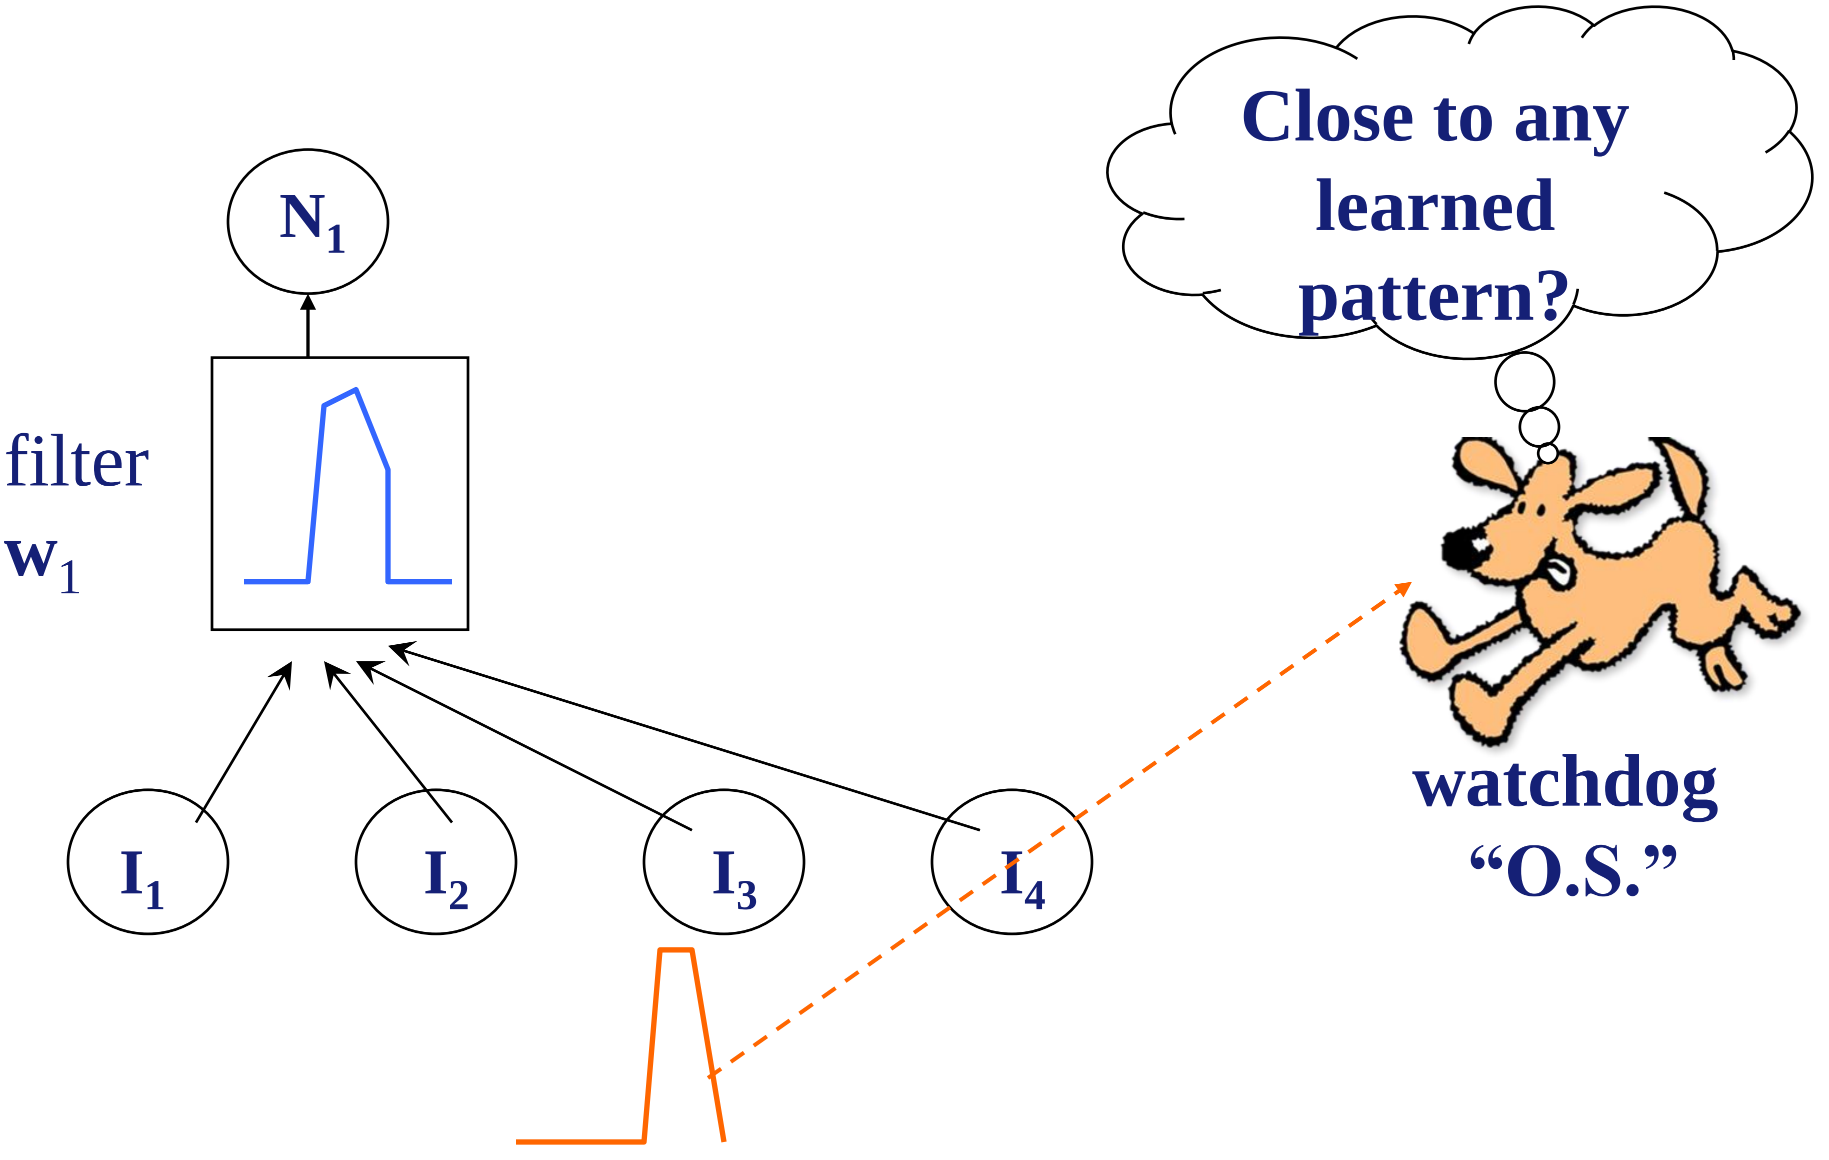

Adaptive Resonance Oversimplified

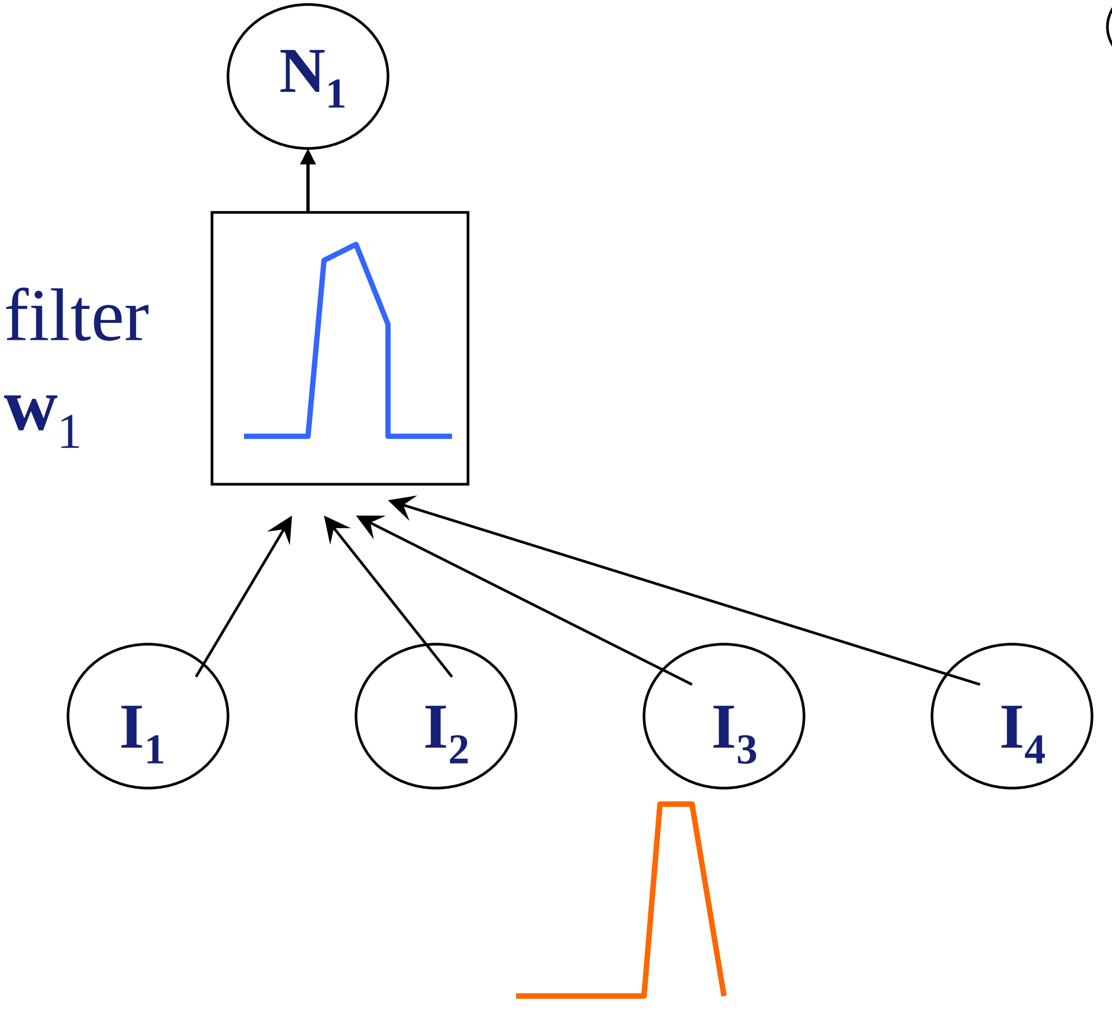

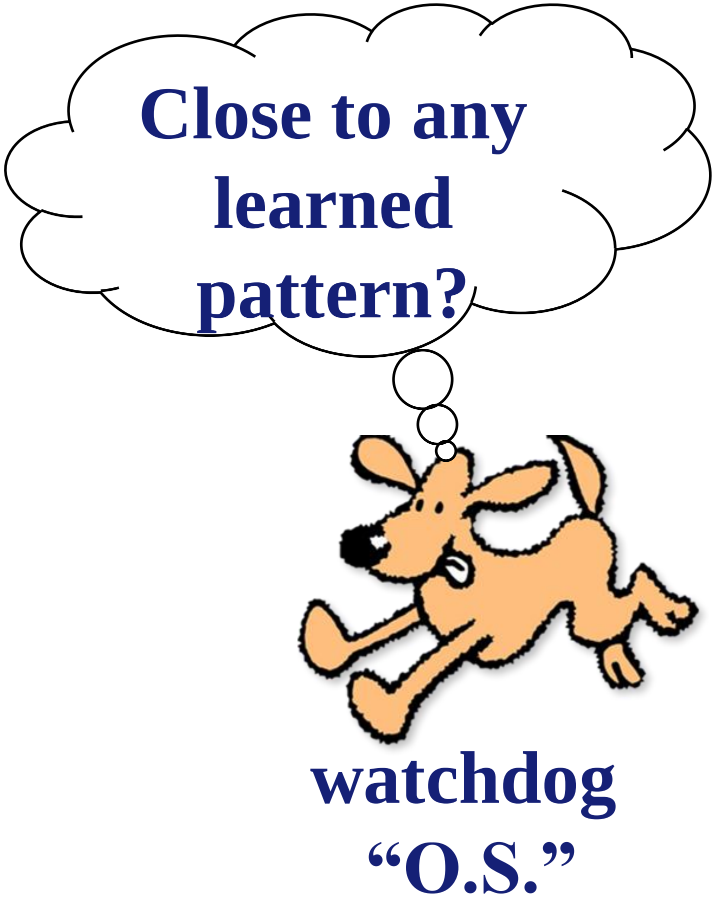

Adaptive Resonance Oversimplified

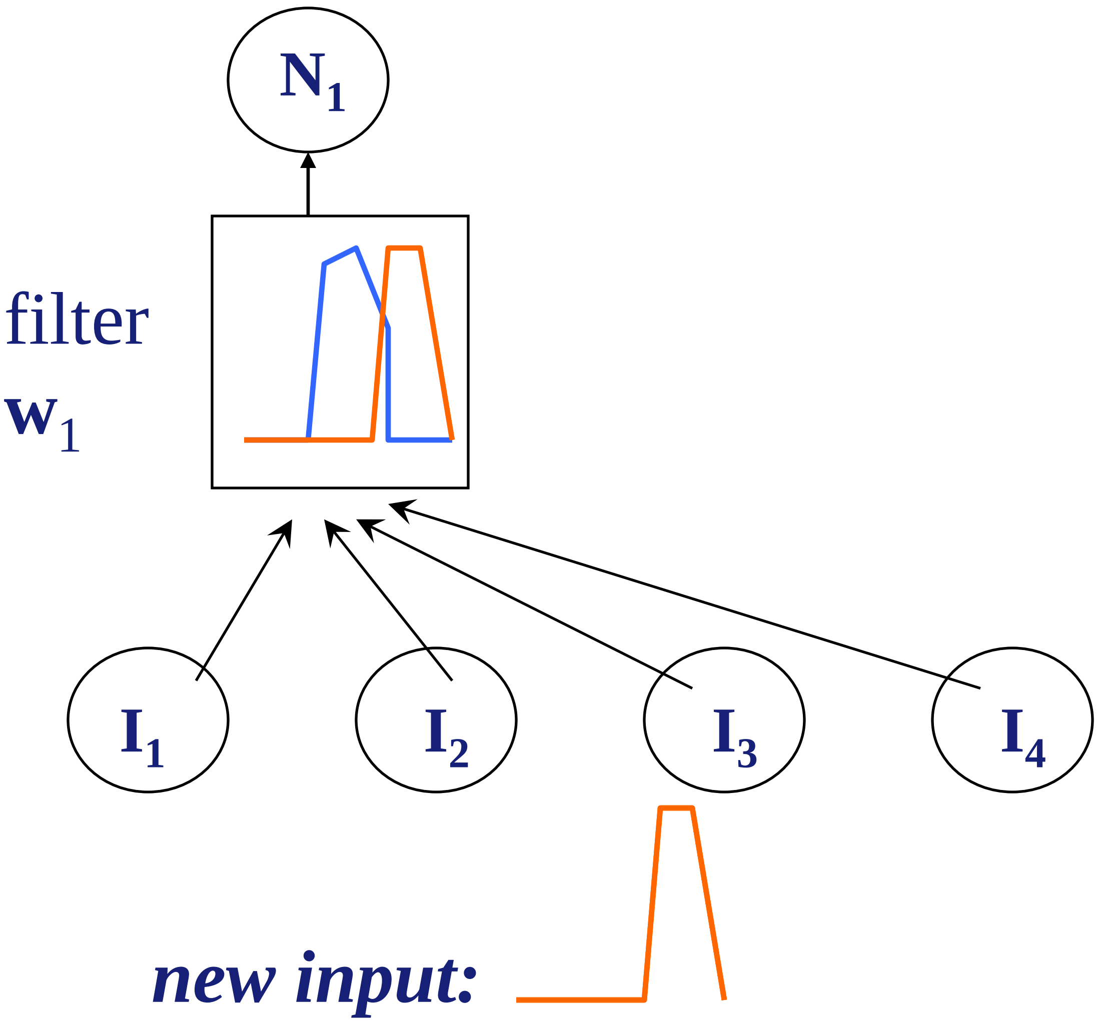

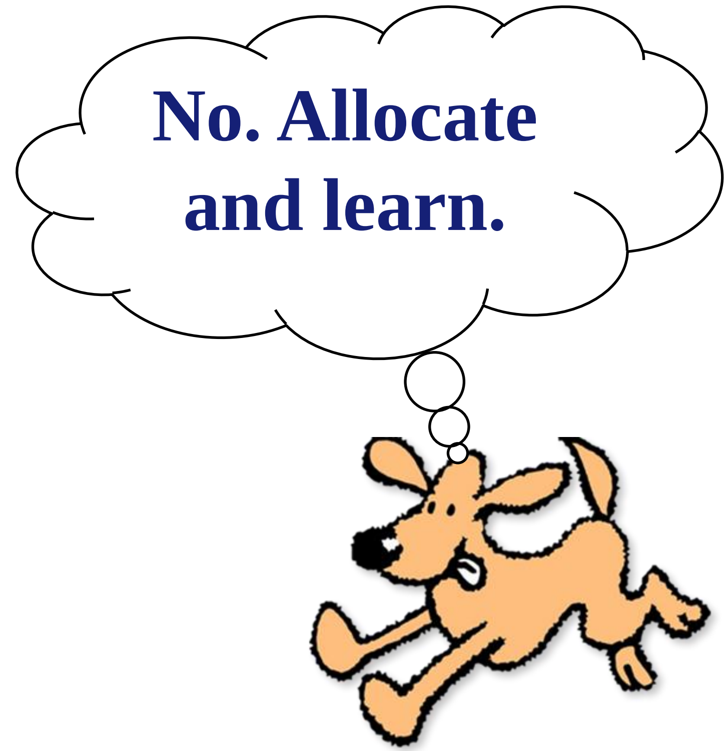

watchdog
“O.S.”

Adaptive Resonance Oversimplified

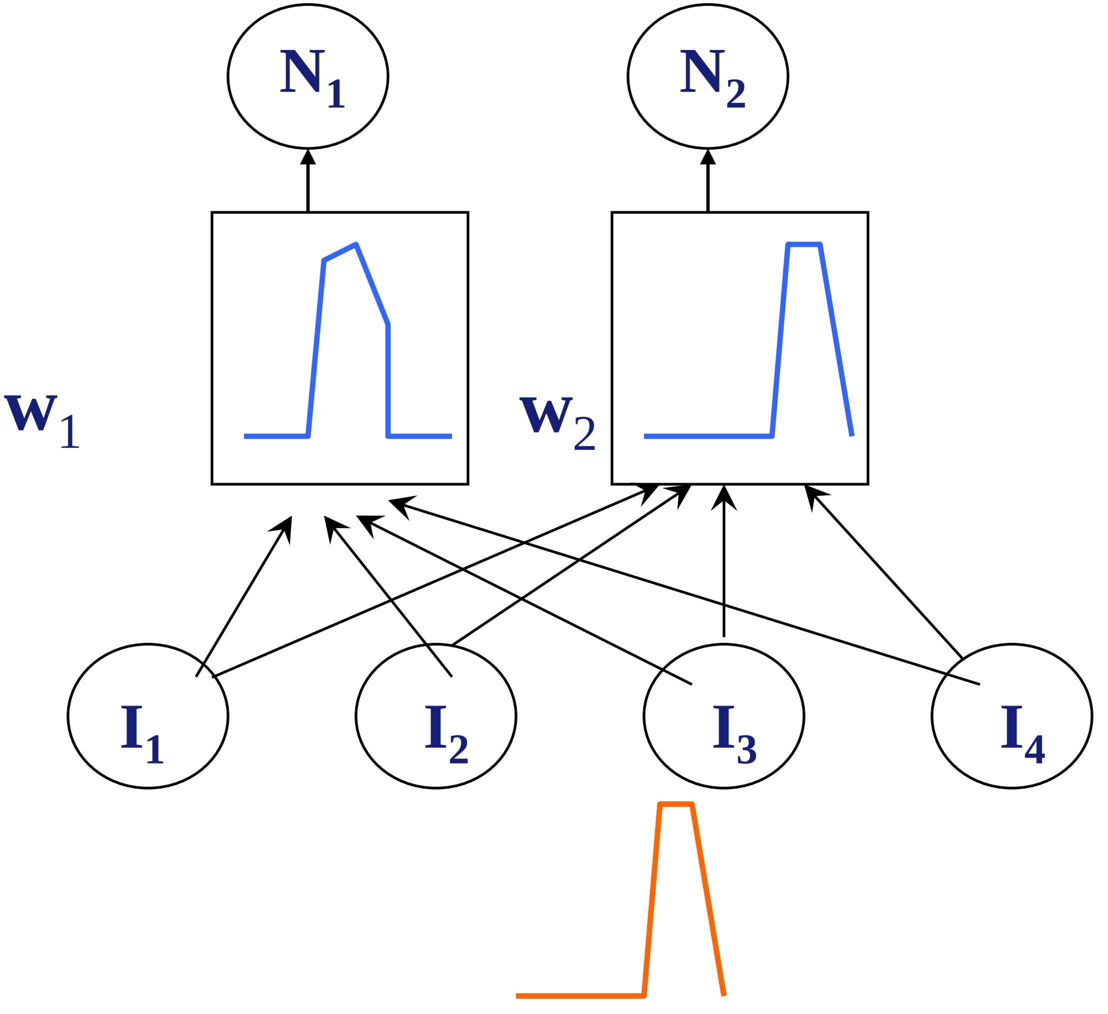

watchdog
“O.S.”

Adaptive Resonance Oversimplified

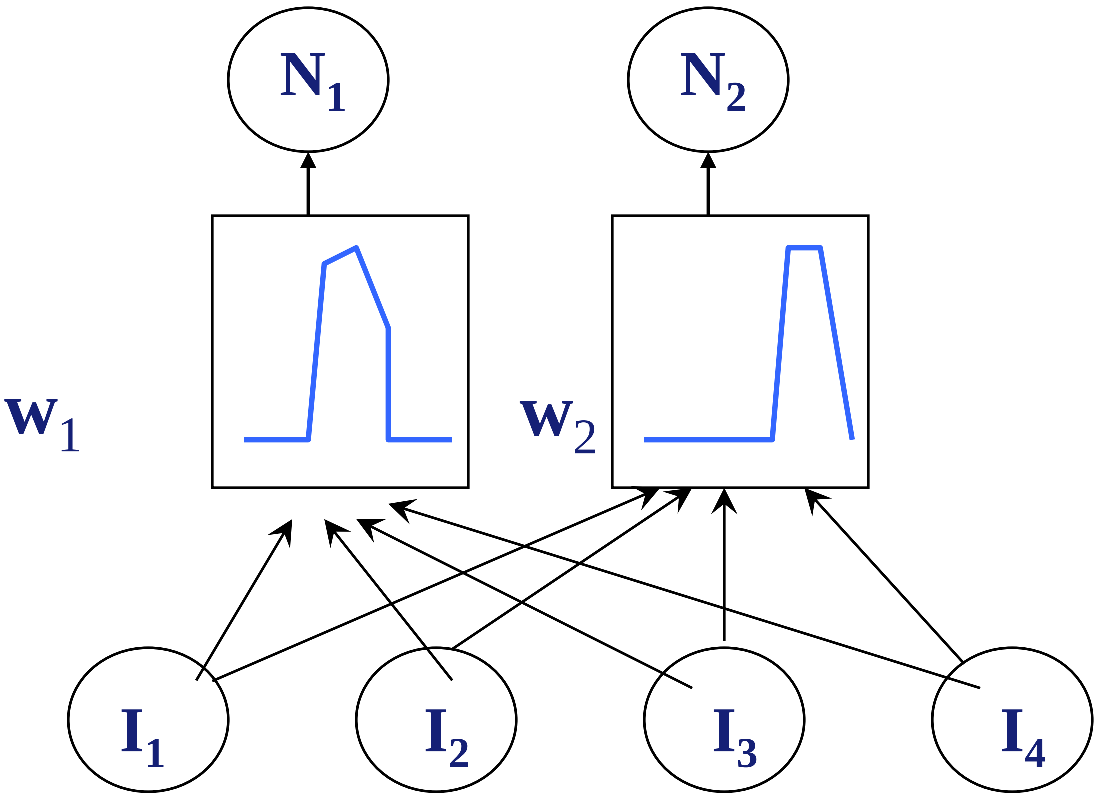

Winner =

$$\arg\max_j \left(\frac{\mathbf{I} \cdot \mathbf{w}_j}{\alpha + |\mathbf{w}_j|} \right)$$

where

$$\frac{\mathbf{I} \cdot \mathbf{w}_j}{|\mathbf{I}|} > \rho$$

Adaptive Resonance Oversimplified

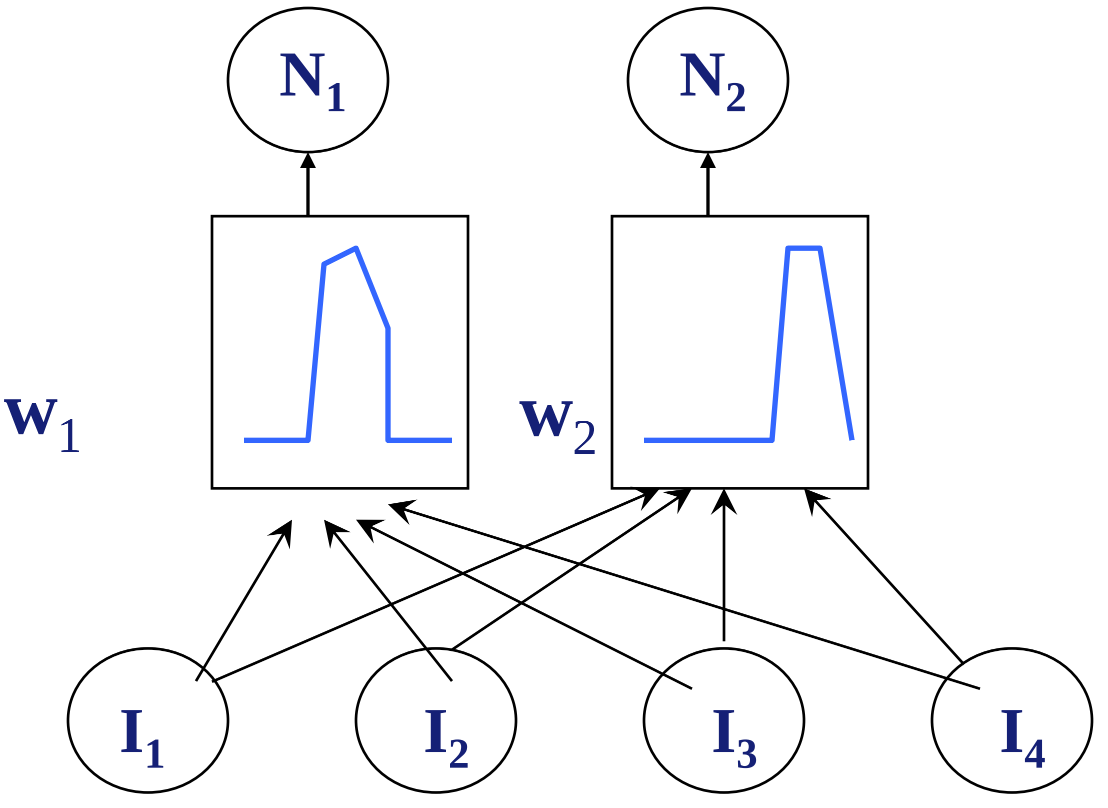

Winner =

$$\arg\max_j \left(\frac{\mathbf{I} \cdot \mathbf{w}_j}{\alpha + |\mathbf{w}_j|} \right)$$

where

$$\frac{\mathbf{I} \cdot \mathbf{w}_j}{|\mathbf{I}|} > \rho$$

Vigilance:
1.0 → perfect fit
0.0 → anything goes

Adaptive Resonance Oversimplified

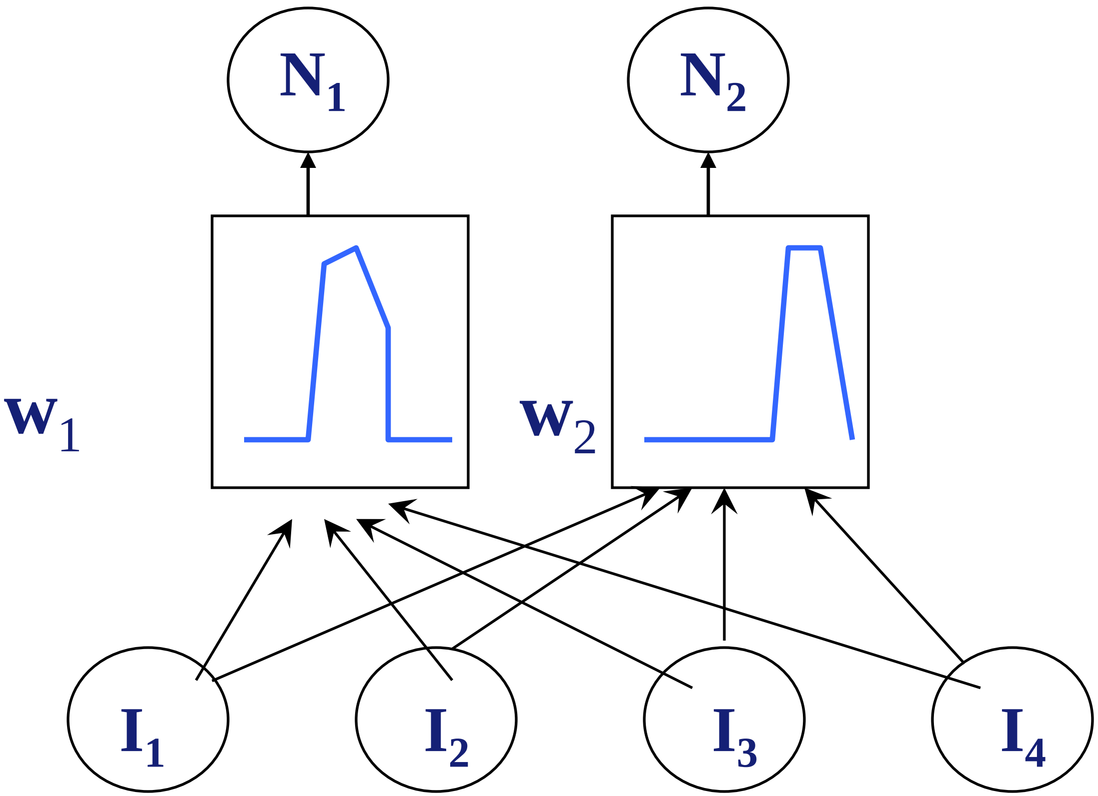

Winner =

$$\arg\max j \left(\frac{\mathbf{I} \cdot \mathbf{w}_j}{\alpha + |\mathbf{w}_j|} \right)$$

where

$$\frac{\mathbf{I} \cdot \mathbf{w}_j}{|\mathbf{I}|} > \rho$$

Vigilance:
1.0 → perfect fit
0.0 → anything goes

If winner
 resonate and refine learned pattern
otherwise
 allocate new neuron

Generalization with vigilance, ρ

1 → perfect match
0 → anything goes

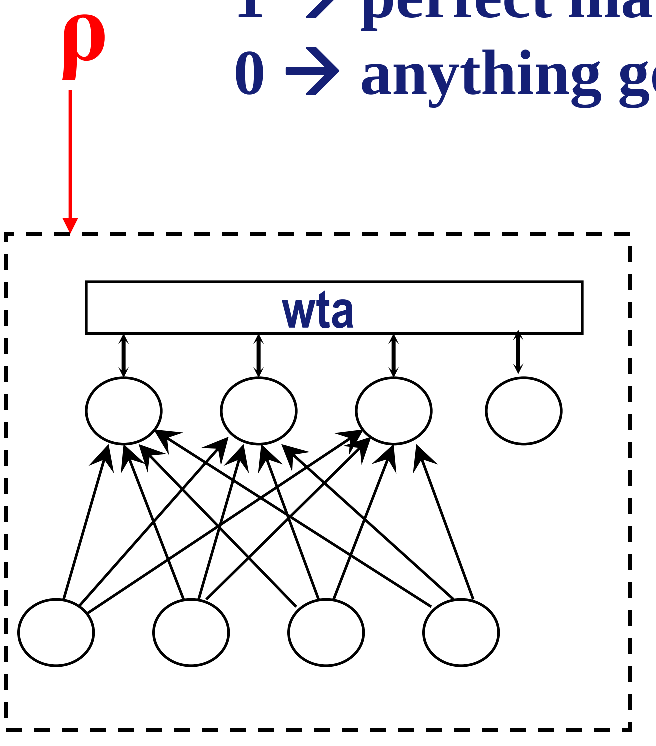

Generalization with vigilance, ρ

1 → perfect match
0 → anything goes

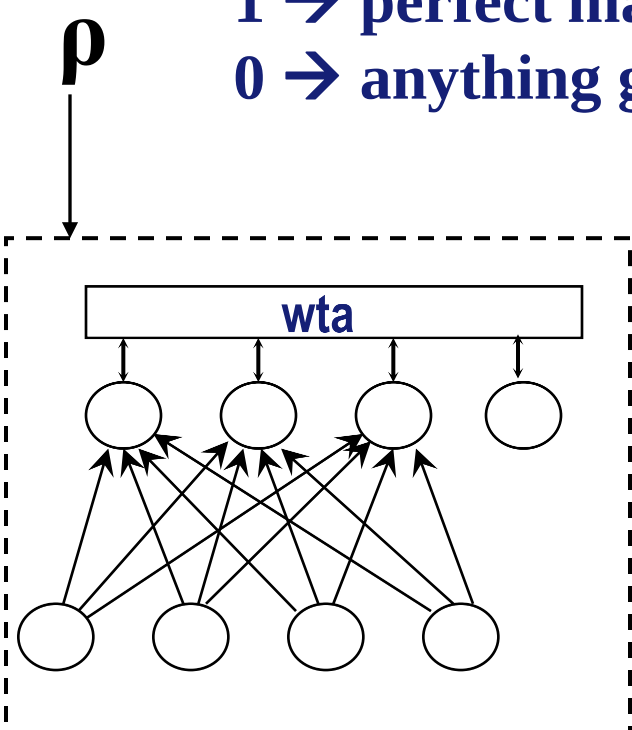

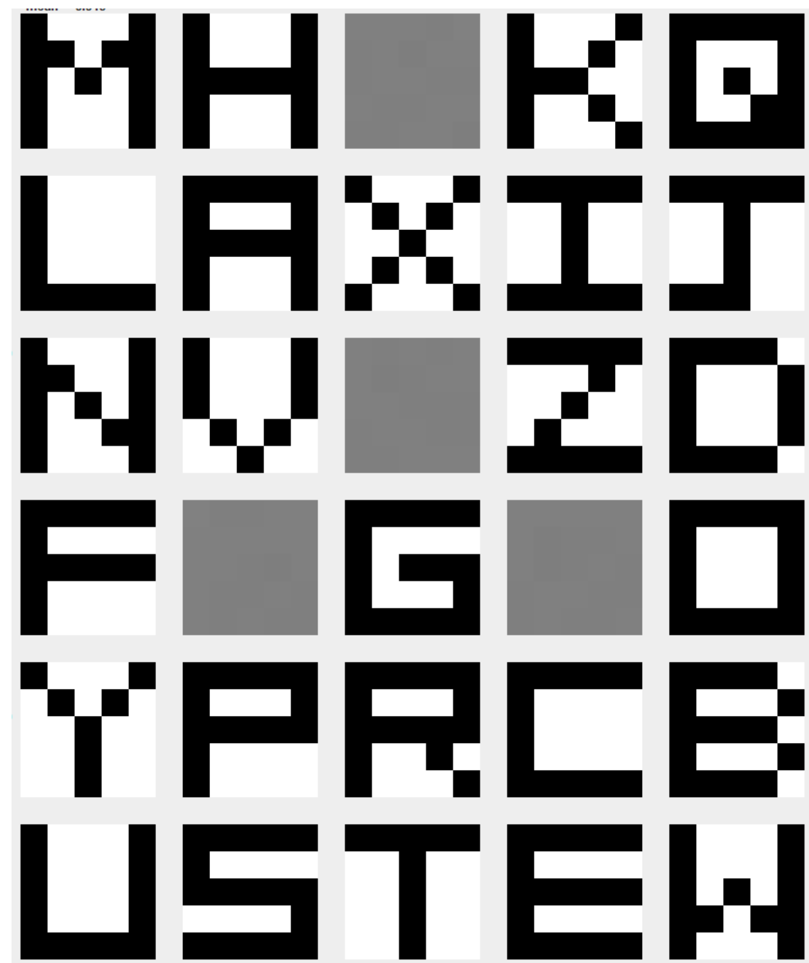

ρ = 0.99

Generalization with vigilance, ρ

1 → perfect match
0 → anything goes

ρ

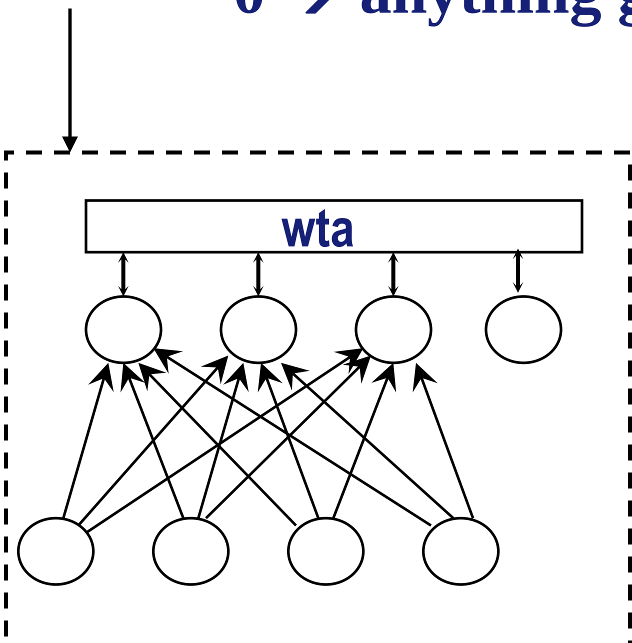

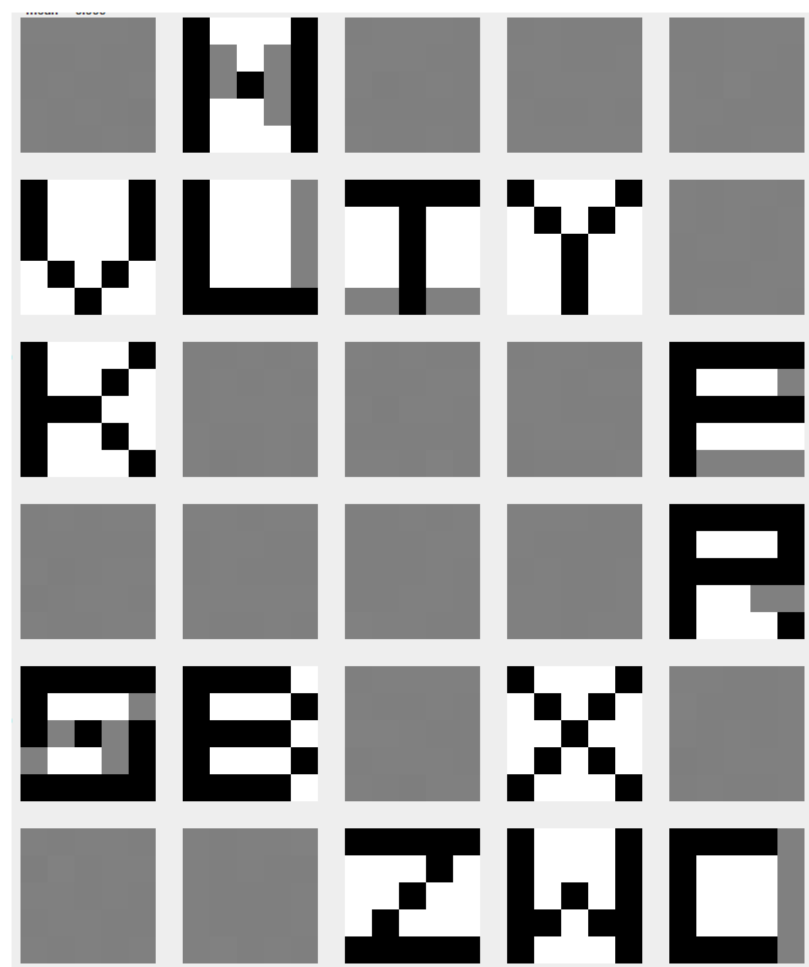

ρ = 0.8

Generalization with vigilance, ρ

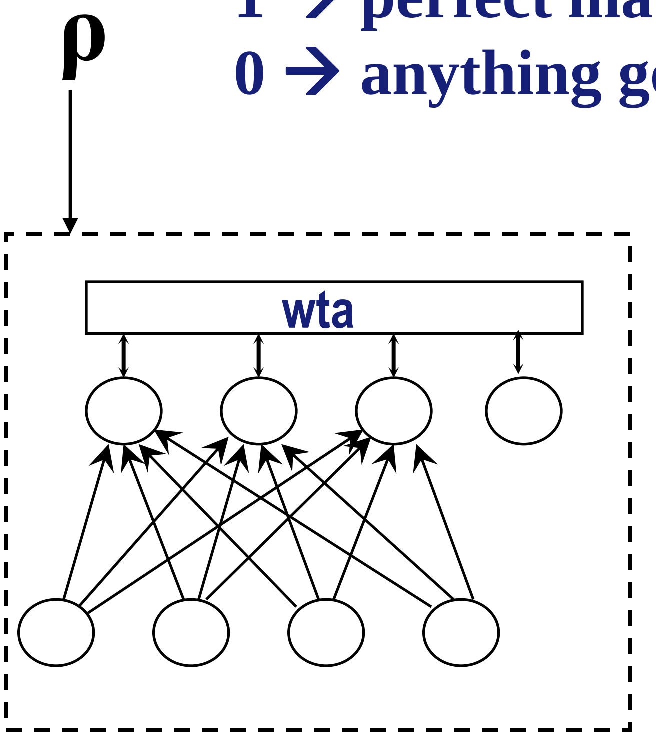

1 → perfect match
0 → anything goes

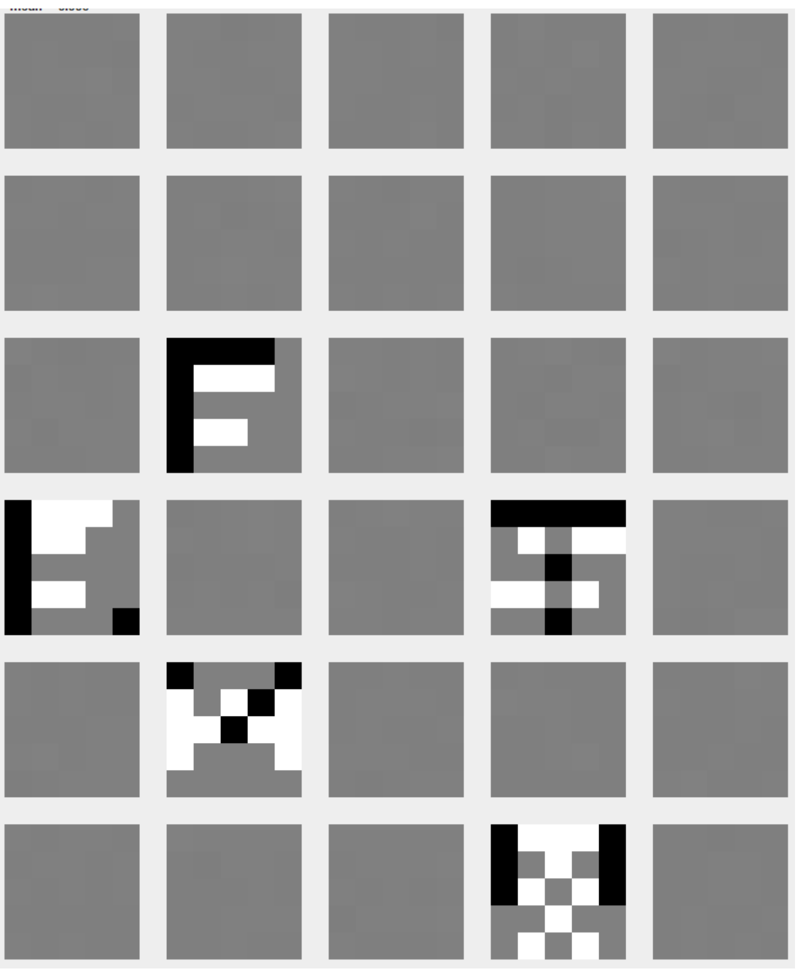

$\rho = 0.5$

Generalization with vigilance, ρ

1 → perfect match
0 → anything goes

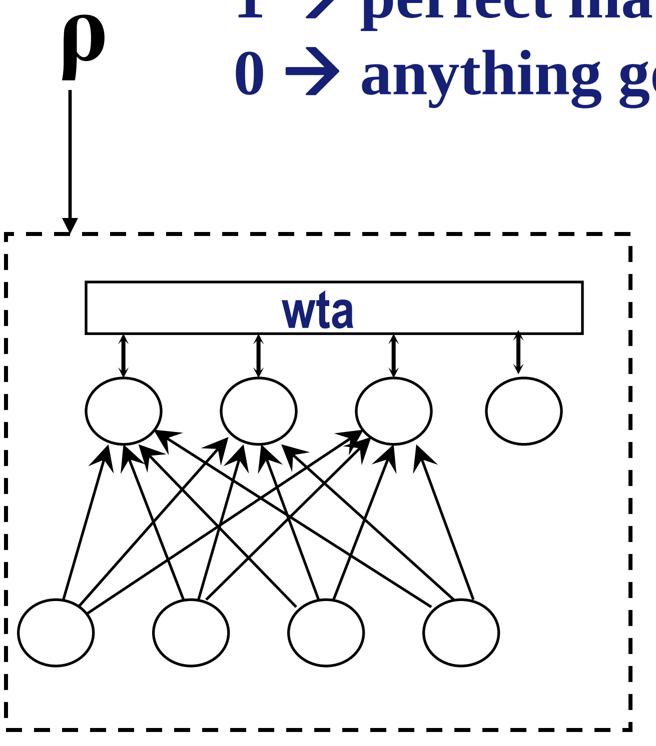

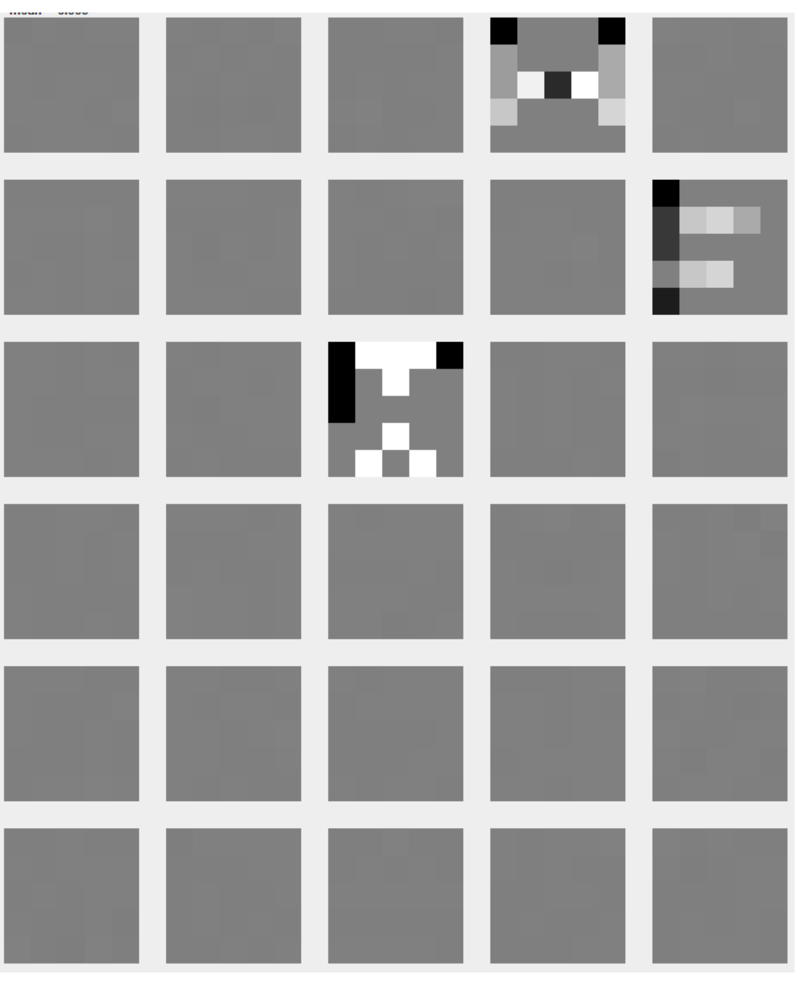

ρ = 0.2

Generalization with vigilance, ρ

ρ

1 → perfect match
0 → anything goes

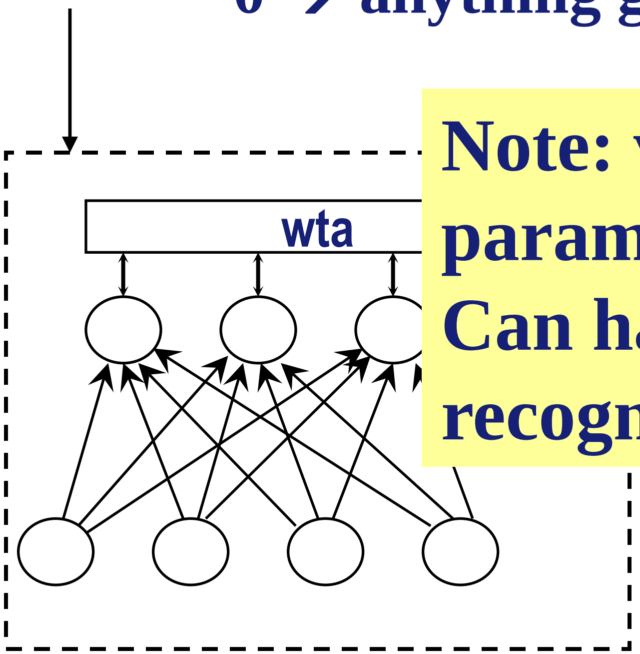

Note: vigilance is a dynamic parameter.
Can have coarse and fine pattern recognition in same network.

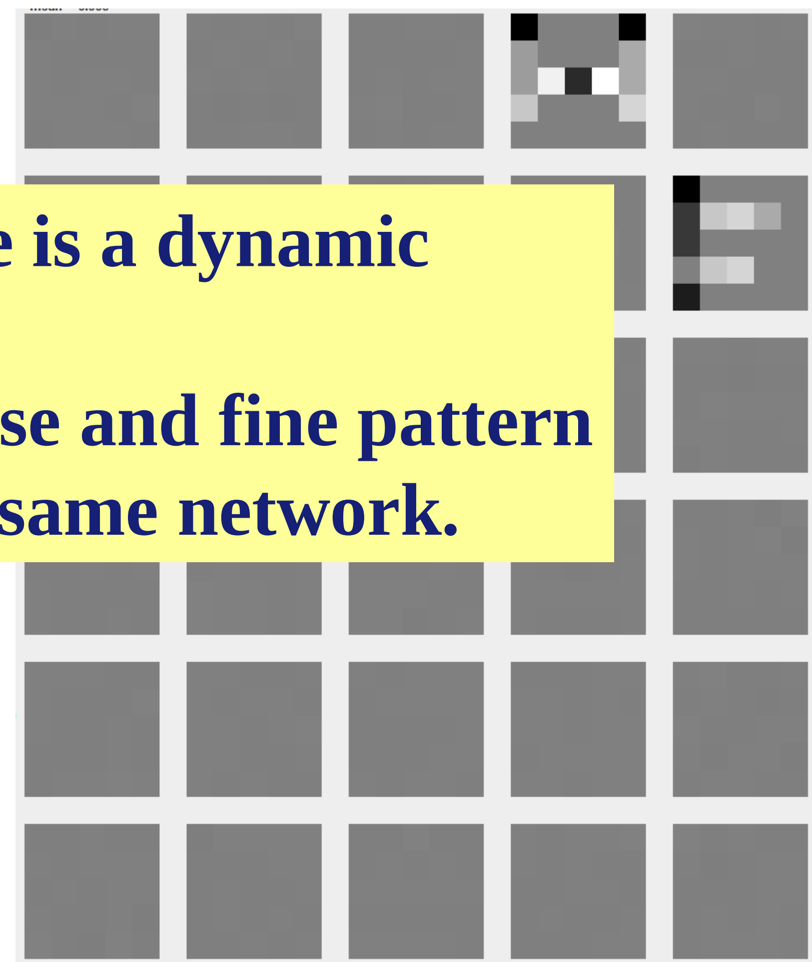

$\rho = 0.2$

Generalization with vigilance, ρ

1 → perfect match
0 → anything goes

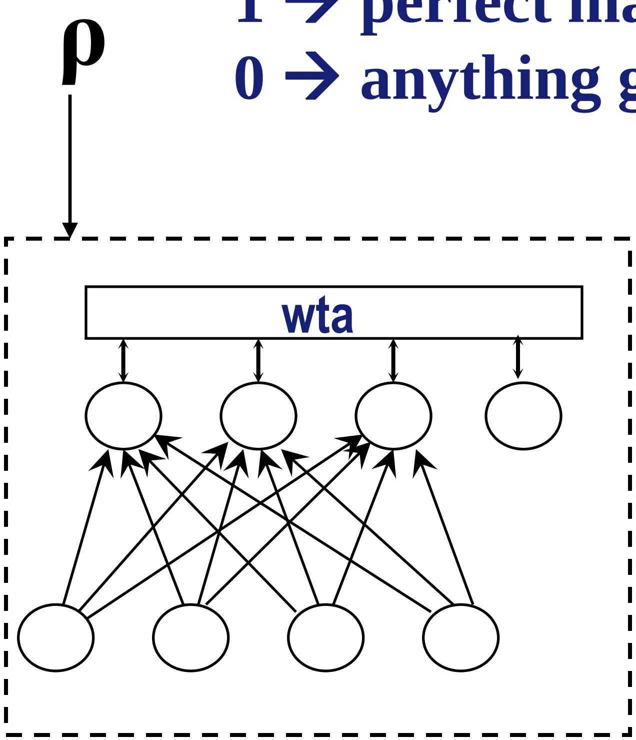

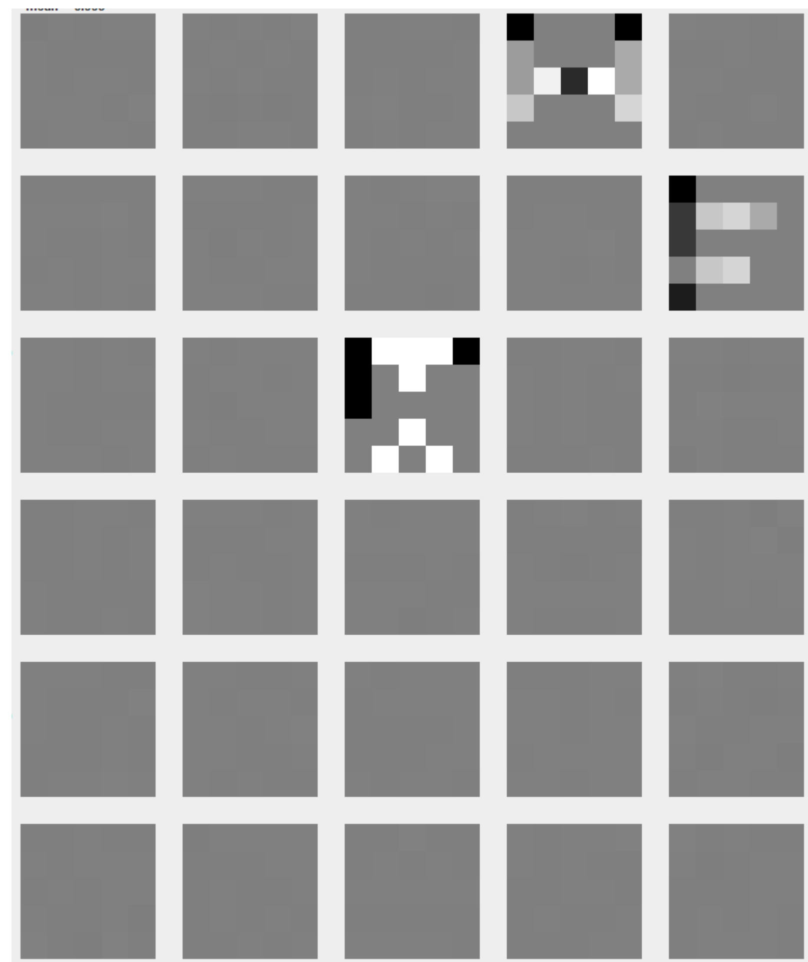

ρ = 0.2: A B C …. X Y Z

ρ = 0.99: A ?

ρ = 0.2

Generalization with vigilance, ρ

1 → perfect match
0 → anything goes

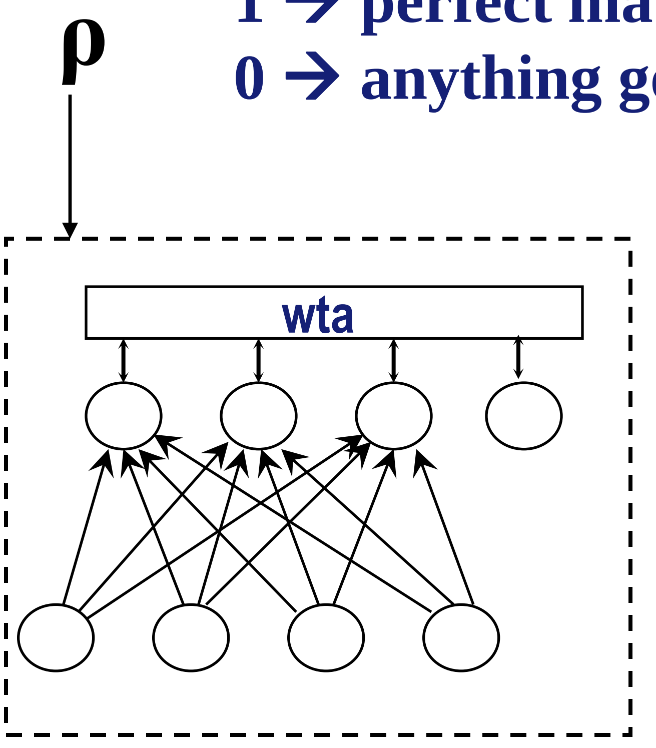

$\rho = 0.2$: A B C …. X Y Z

$\rho = 0.99$: A

Learns it!

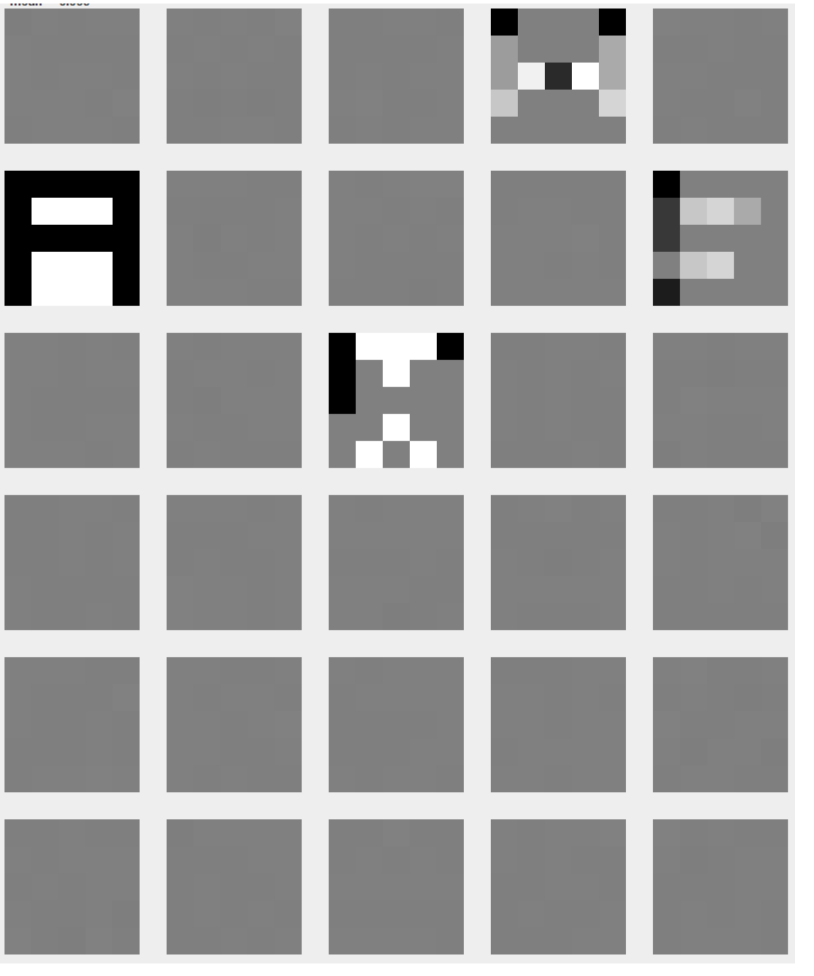

$\rho = 0.2$

ART in hardware? *4 challenges*

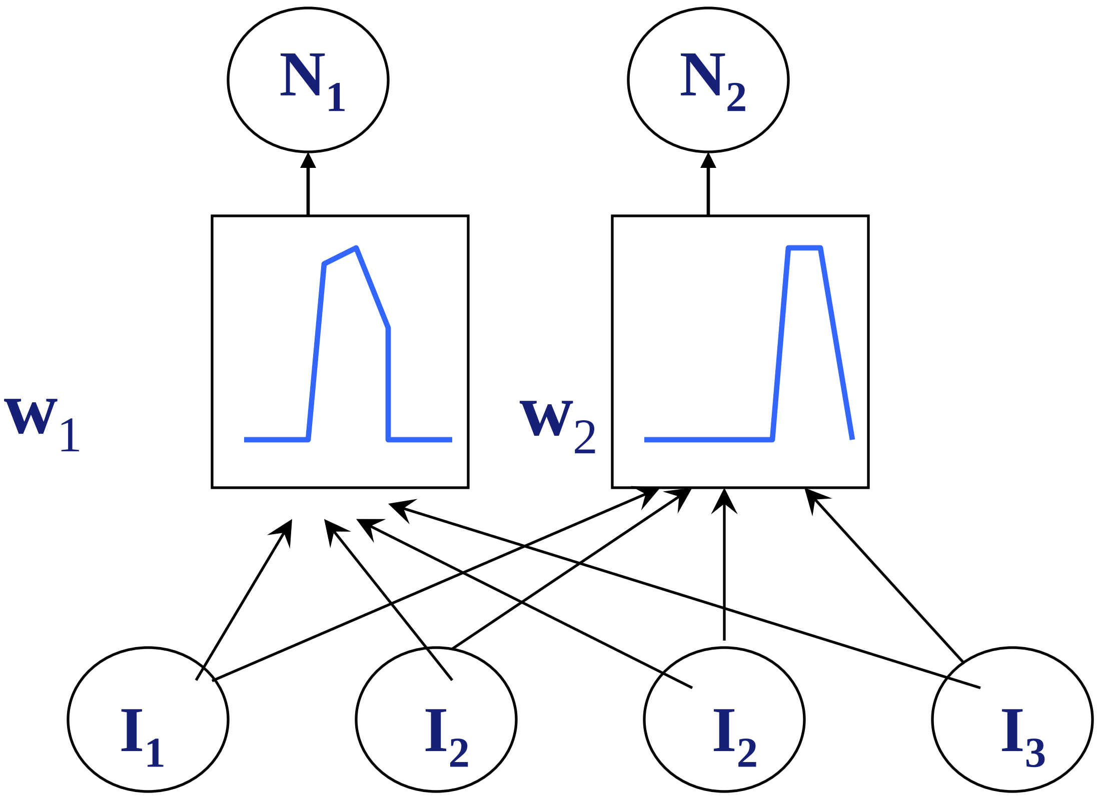

Winner =

$$\max j\left(\frac{\mathbf{I}\cdot\mathbf{w}_j}{\alpha+\left|\mathbf{w}_j\right|}\right)$$

where

$$\frac{\mathbf{I}\cdot\mathbf{w}_j}{\left|\mathbf{I}\right|}>\rho$$

ART in hardware? *4 challenges*

1. Nonregular

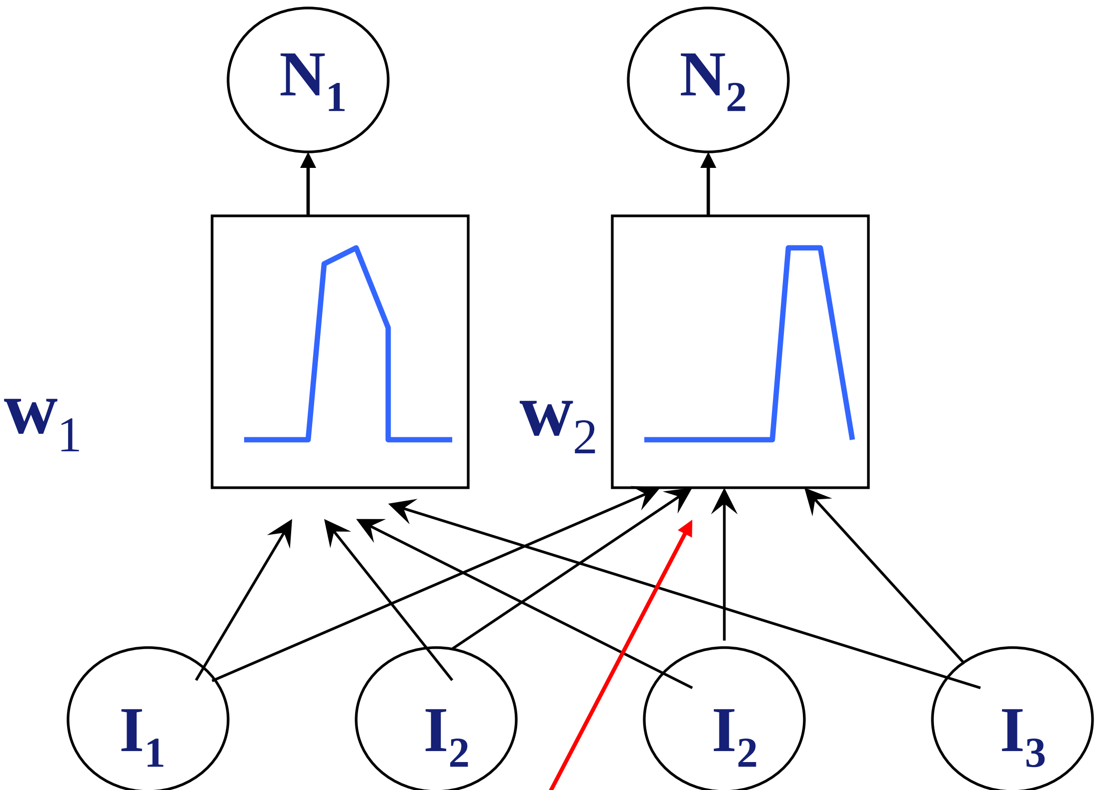

Winner =

$$\max j\left(\frac{\mathbf{I}\cdot\mathbf{w}_j}{\alpha+|\mathbf{w}_j|}\right)$$

where

$$\frac{\mathbf{I}\cdot\mathbf{w}_j}{|\mathbf{I}|} > \rho$$

4. Nonlocal

2. Dynamic allocation

3. Bounded resource

ART in hardware? *4 challenges*

1. Nonregular

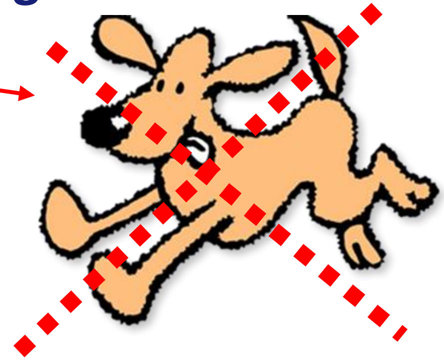

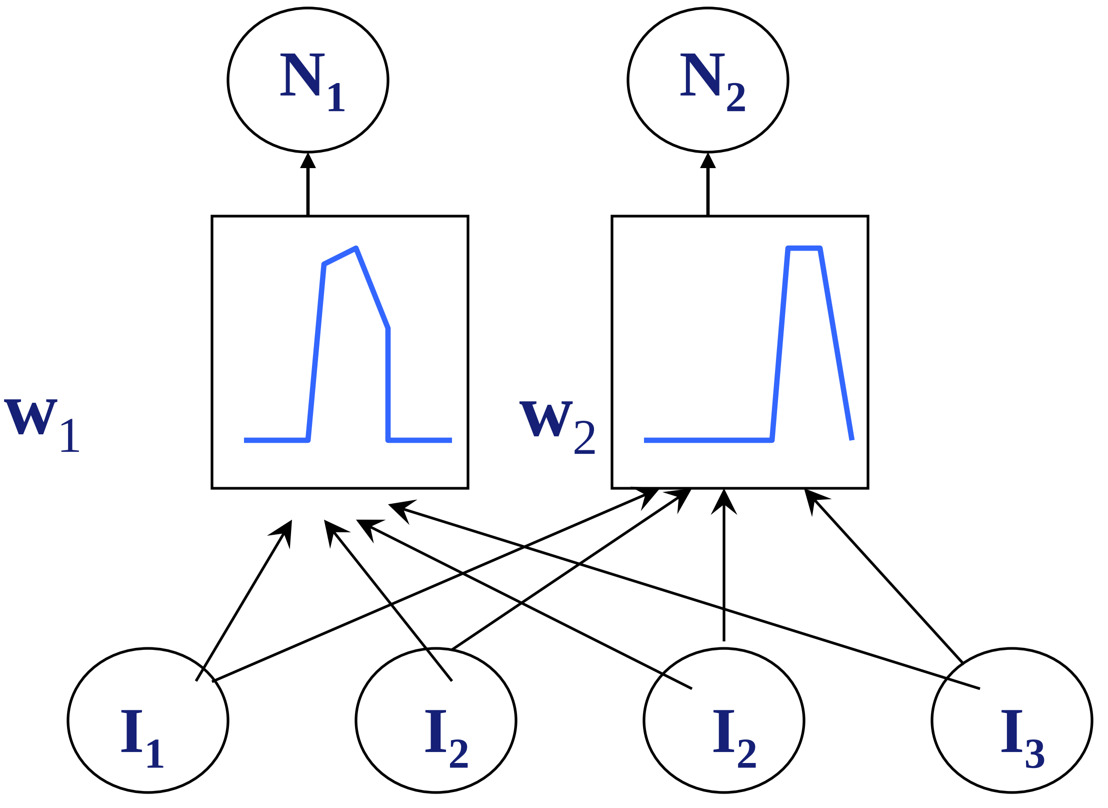

Winner =

$$\max j\left(\frac{\mathbf{I}\cdot\mathbf{w}_j}{\alpha+|\mathbf{w}_j|}\right)$$

where

$$\frac{\mathbf{I}\cdot\mathbf{w}_j}{|\mathbf{I}|} > \rho$$

ART in hardware? *4 challenges*

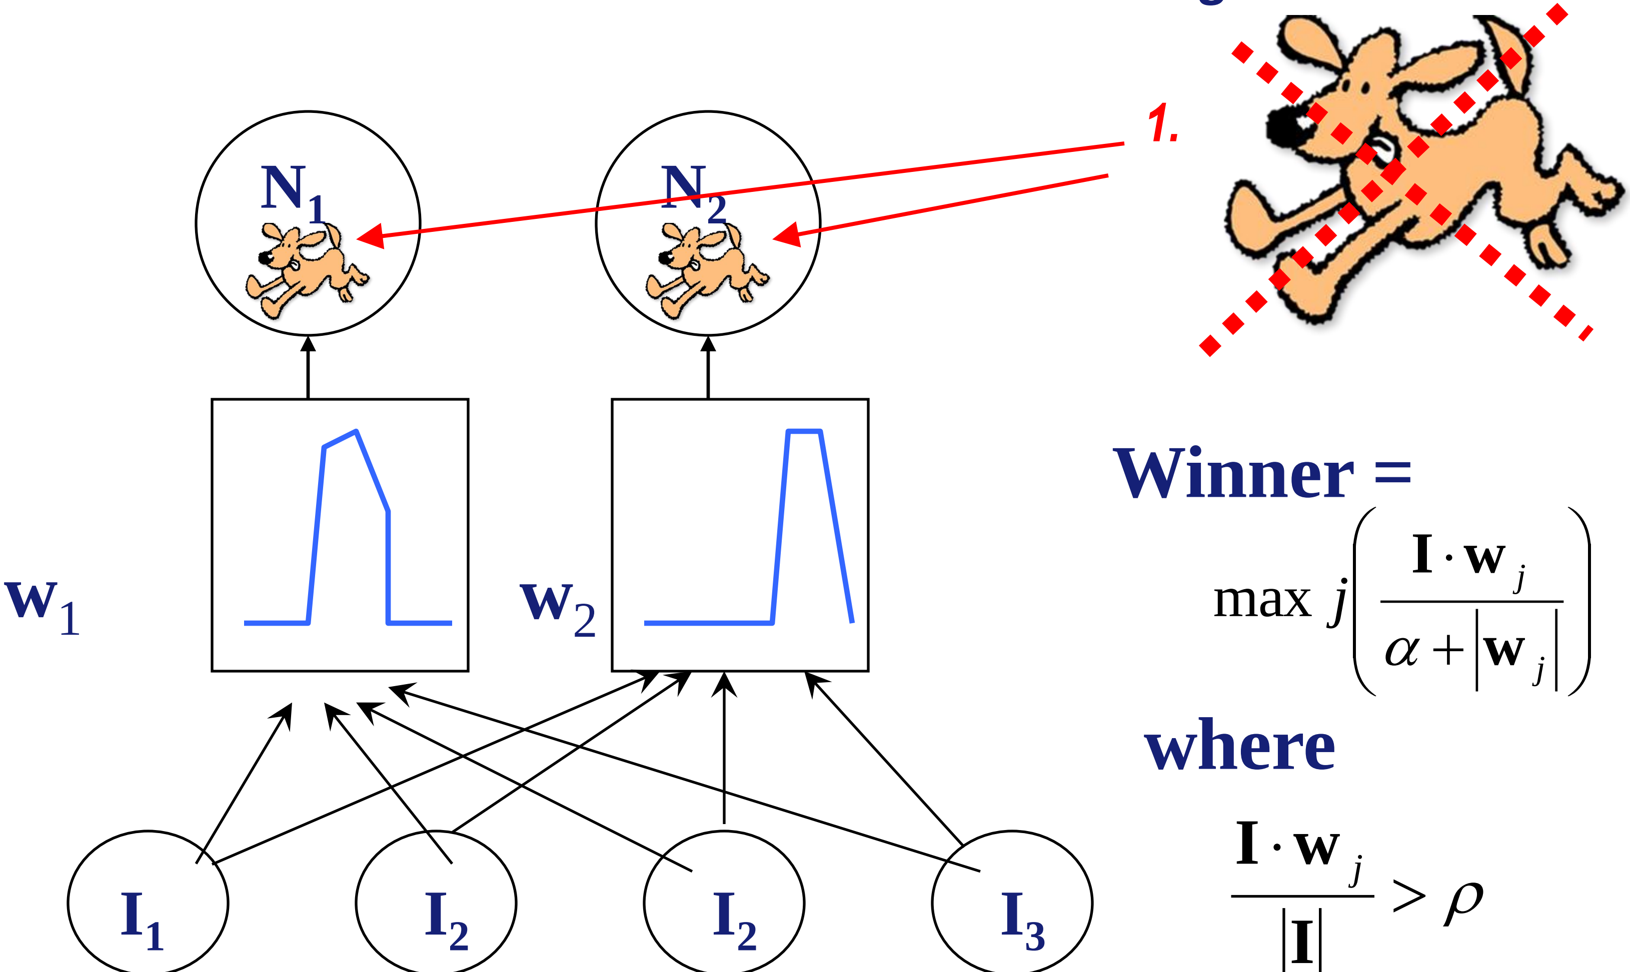

Winner =

$$\max_j \left(\frac{\mathbf{I} \cdot \mathbf{w}_j}{\alpha + |\mathbf{w}_j|} \right)$$

where

$$\frac{\mathbf{I} \cdot \mathbf{w}_j}{|\mathbf{I}|} > \rho$$

ART in hardware? *4 challenges*

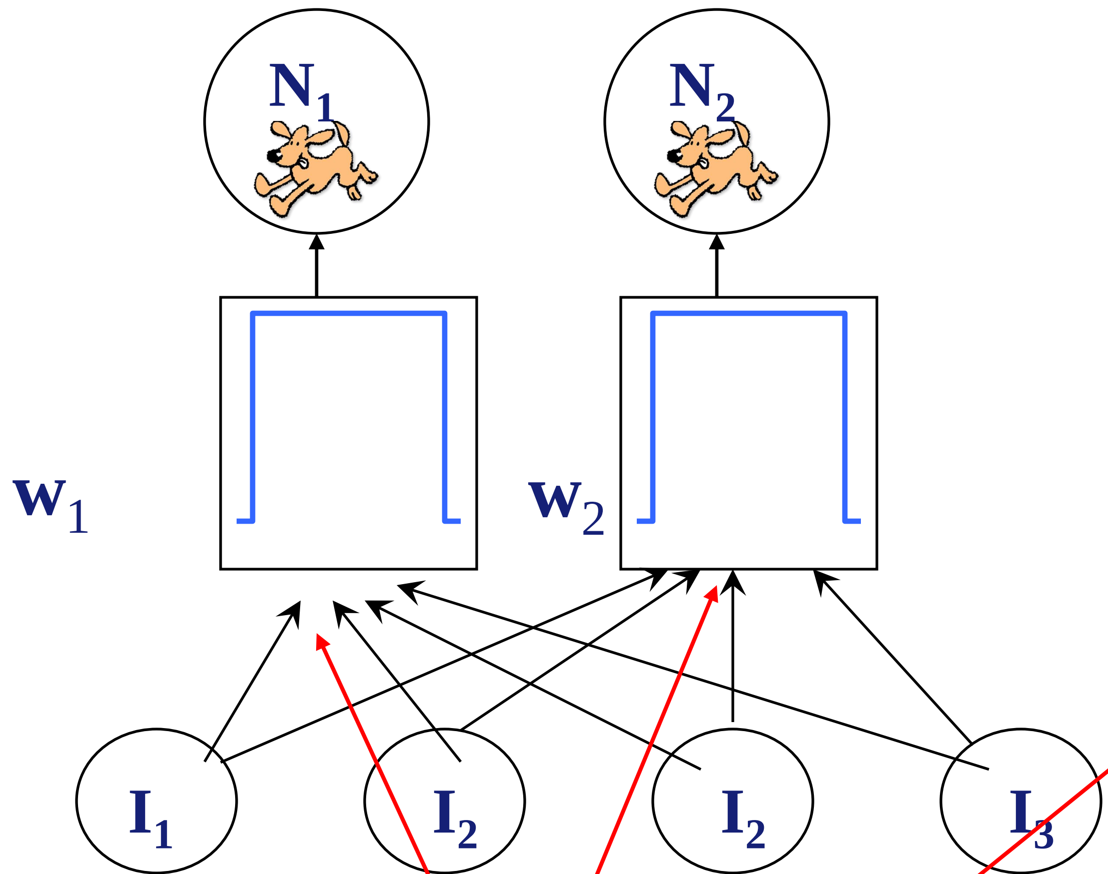

Winner =

$$\max j\left(\frac{\mathbf{I}\cdot\mathbf{w}_j}{\alpha+|\mathbf{w}_j|}\right)$$

where

$$\frac{\mathbf{I}\cdot\mathbf{w}_j}{|\mathbf{I}|} > \rho$$

2. (Dynamic allocation)

Preallocated filters—poor fit to everything, but will resonate if necessary

ART in hardware? *4 challenges*

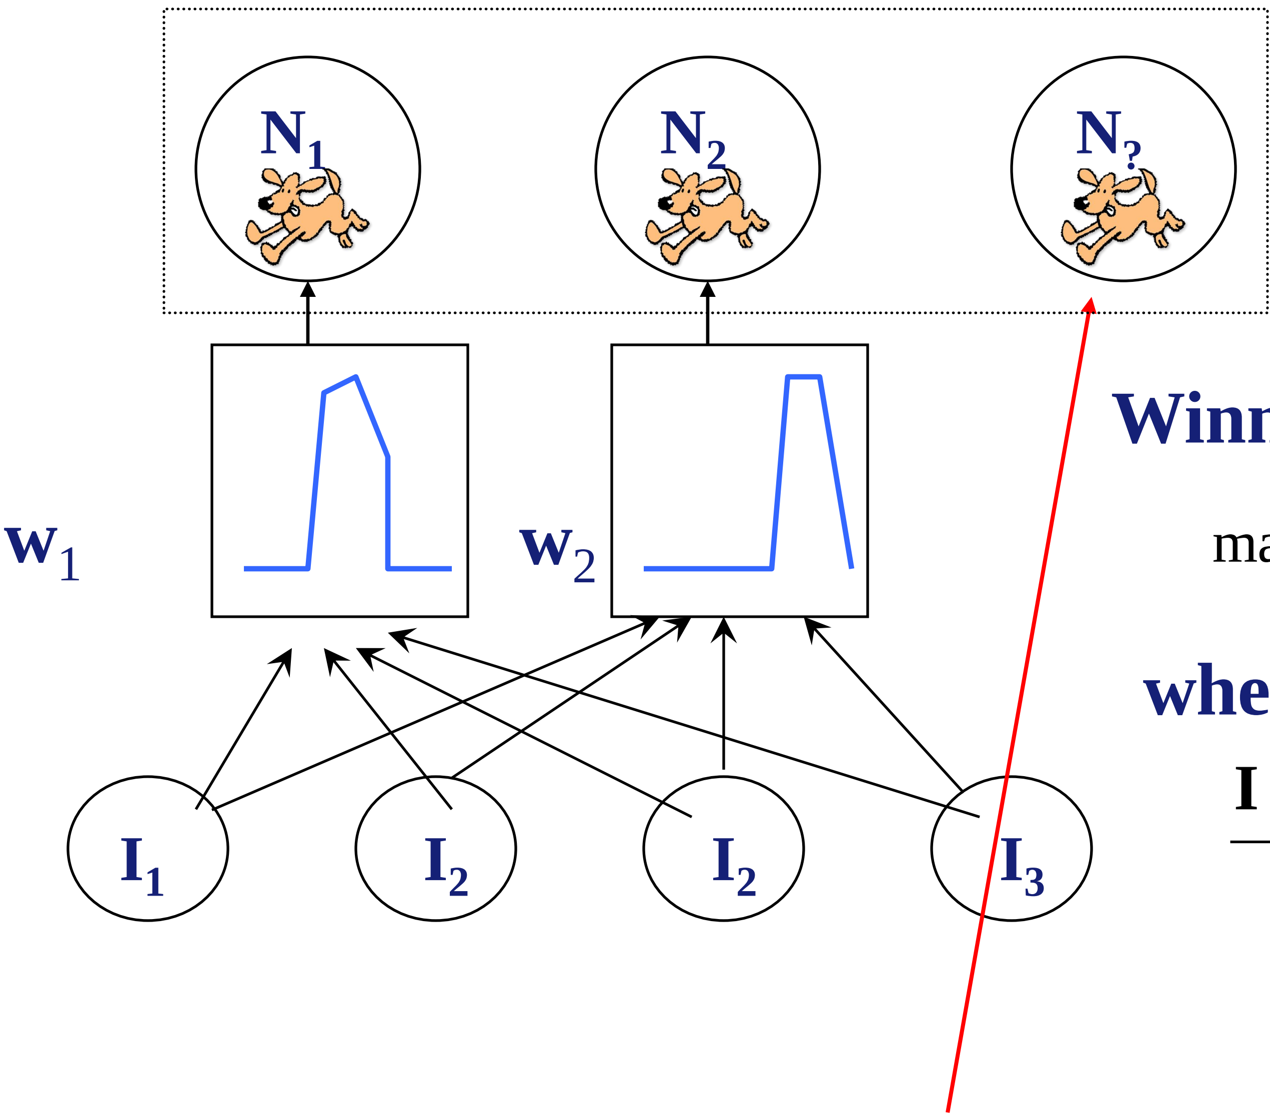

Winner =

$$\max j\left(\frac{\mathbf{I}\cdot\mathbf{w}_j}{\alpha+\left|\mathbf{w}_j\right|}\right)$$

where

$$\frac{\mathbf{I}\cdot\mathbf{w}_j}{\left|\mathbf{I}\right|}>\rho$$

3. (Bounded resource) "I don't know" neuron. If all neurons have learned and novel pattern comes in, the "I don't know" neuron wins.

ART in hardware? *4 challenges*

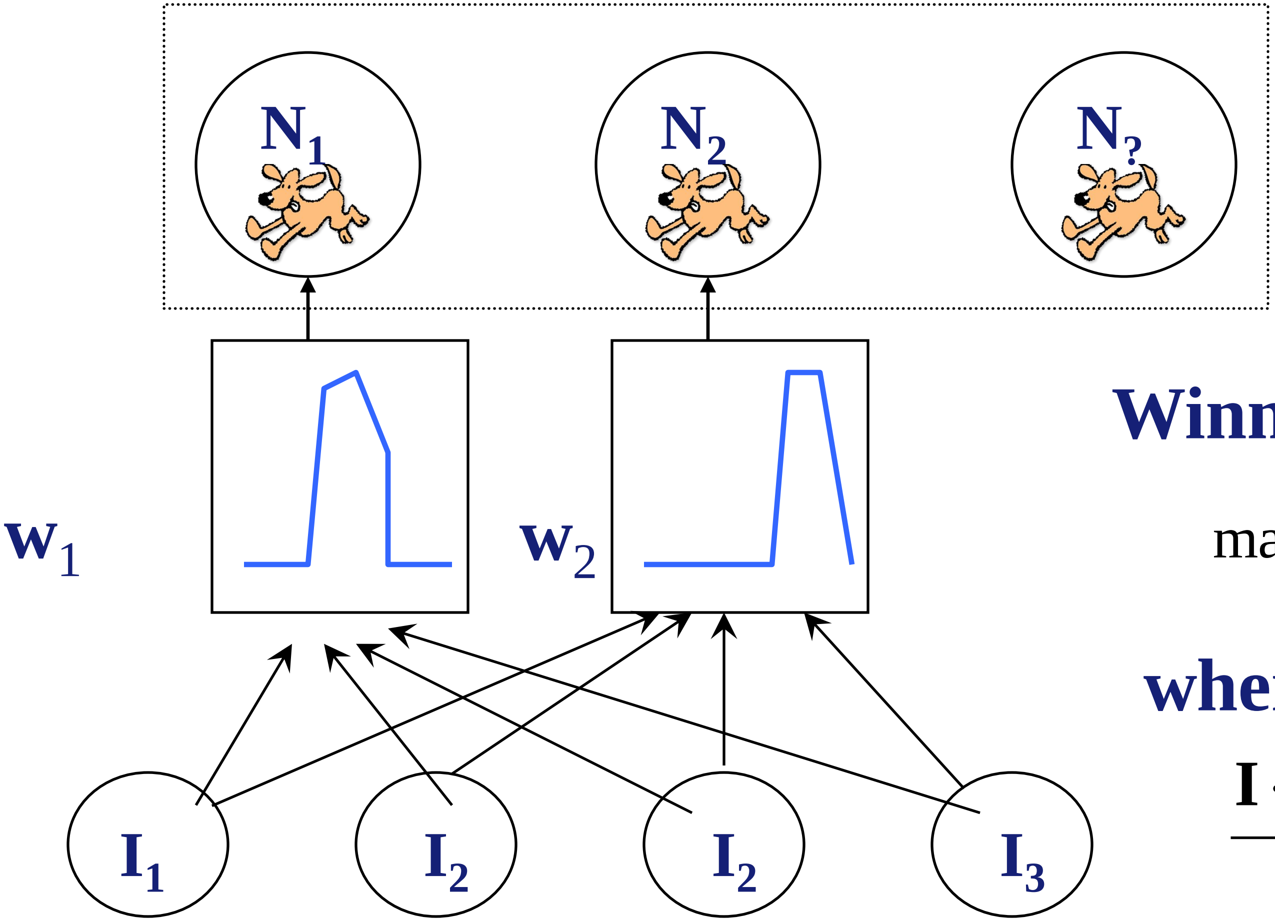

Winner =

$$\max j\left(\frac{\mathbf{I}\cdot\mathbf{w}_j}{\alpha+|\mathbf{w}_j|}\right)$$

where

$$\frac{\mathbf{I}\cdot\mathbf{w}_j}{|\mathbf{I}|}>\rho$$

4. (Nonlocal) Adaptive neurons (synaptic scaling)

Nano/CMOS layout, mART-1

Adaptive filters (synapses) in nano

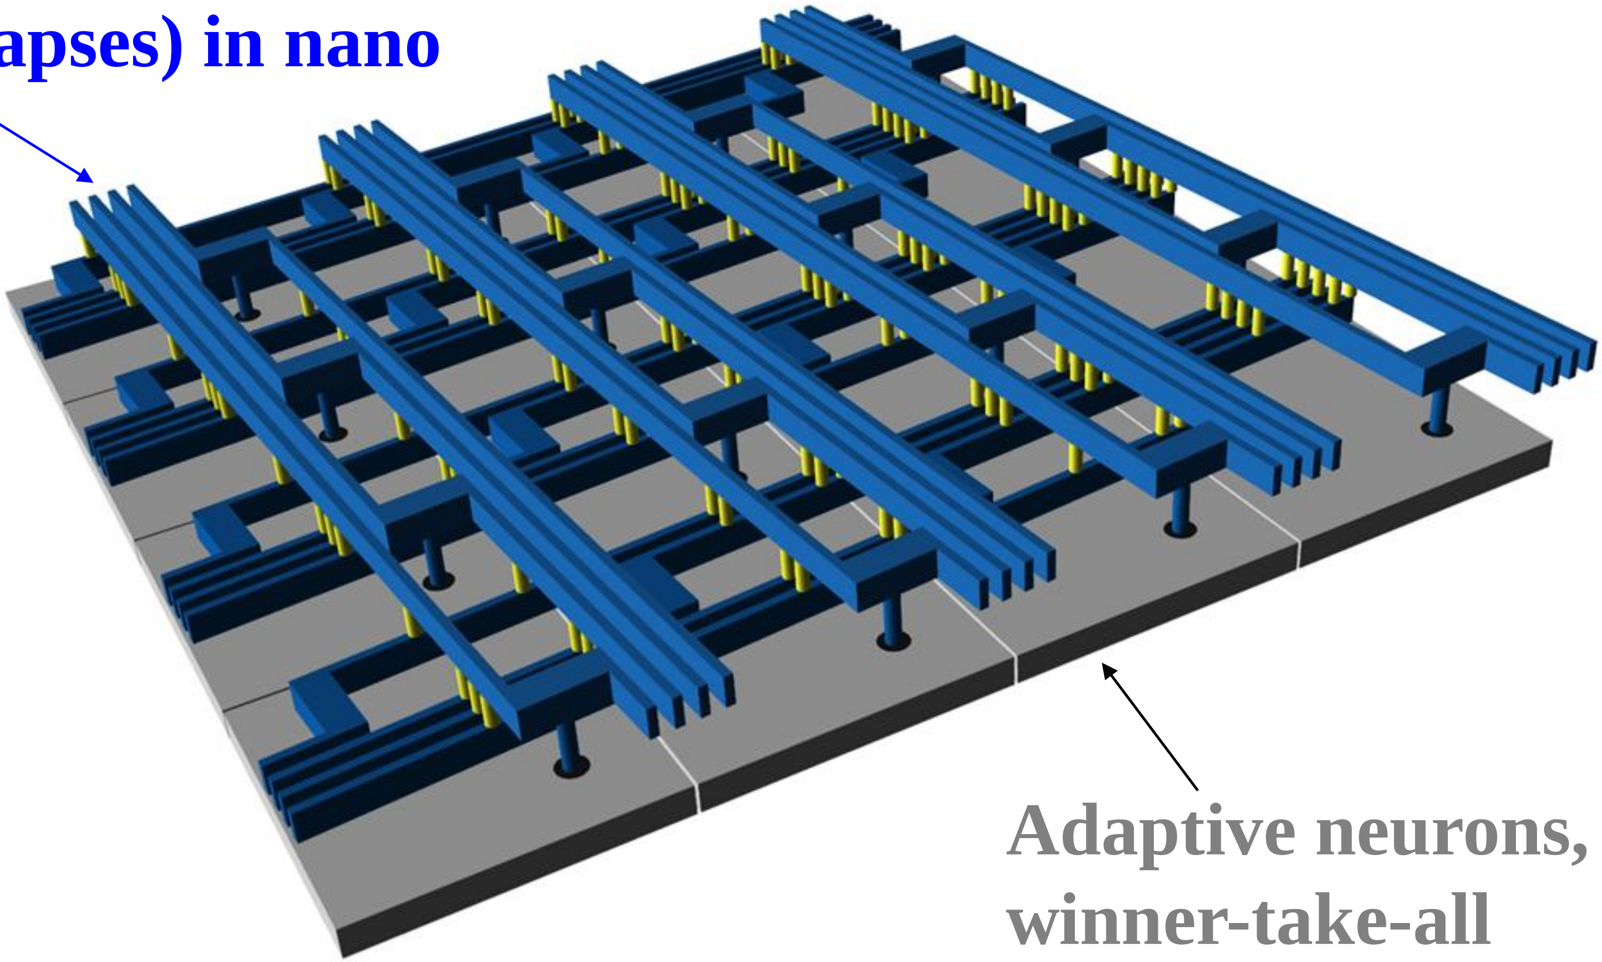

Adaptive neurons, winner-take-all circuits in CMOS

Fast, incremental learning

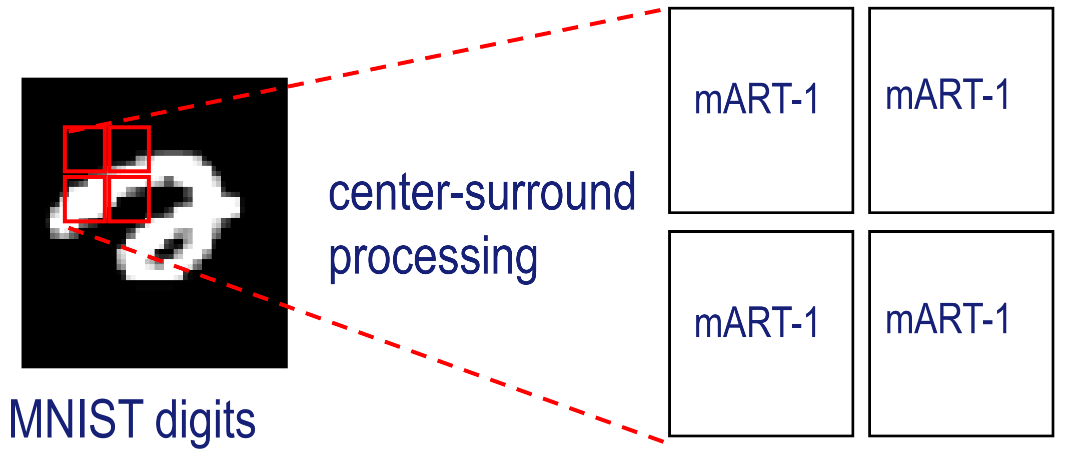

Low vigilance

Fast, incremental learning

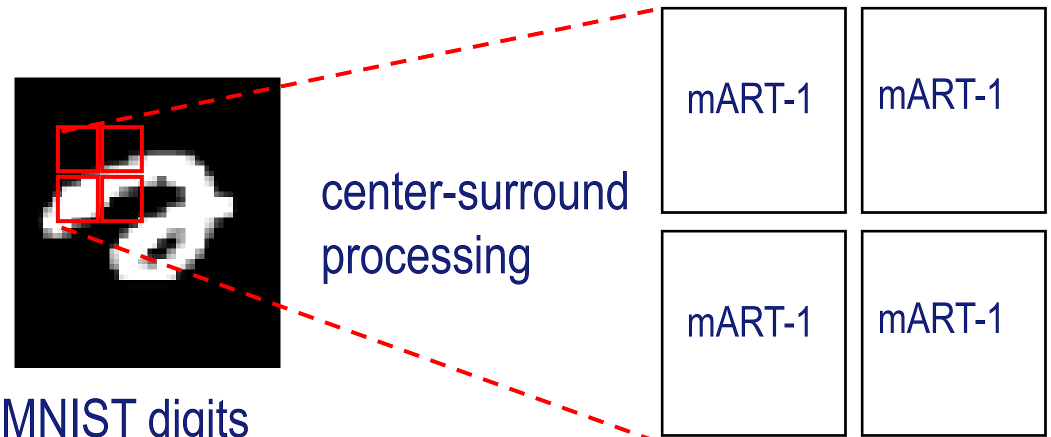

MNIST digits

Low vigilance

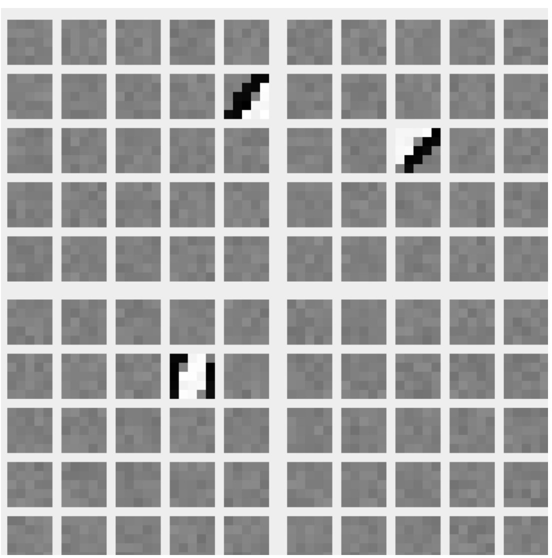

After 1 digit

Fast, incremental learning

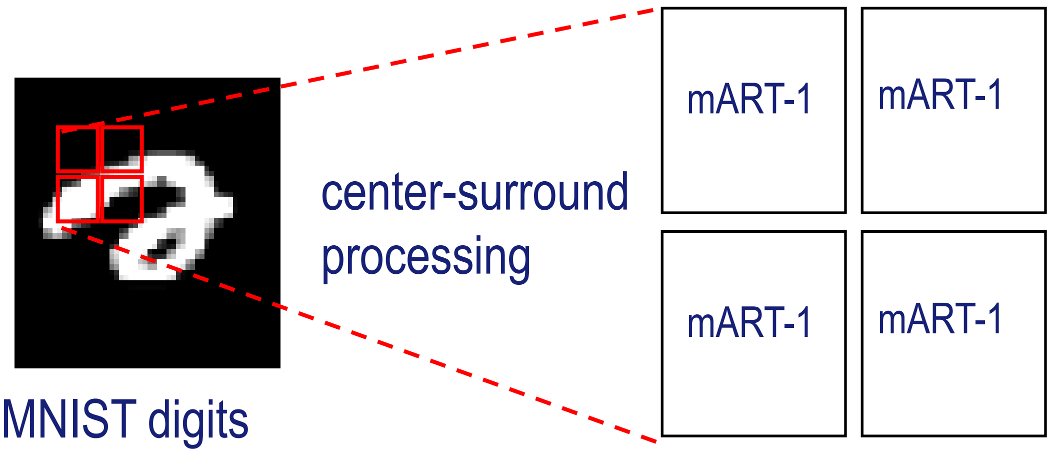

MNIST digits

Low vigilance

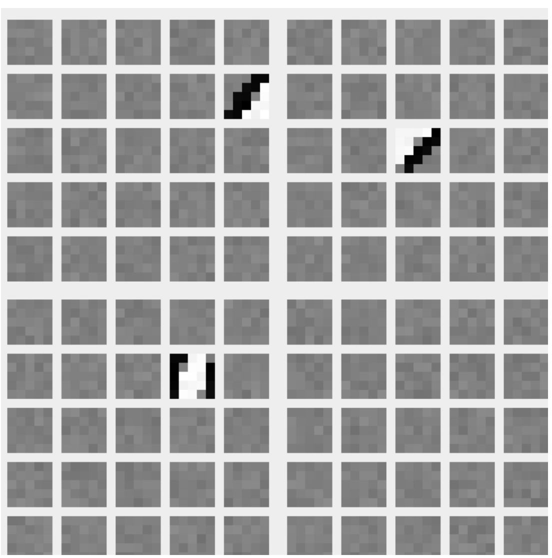

After 1 digit

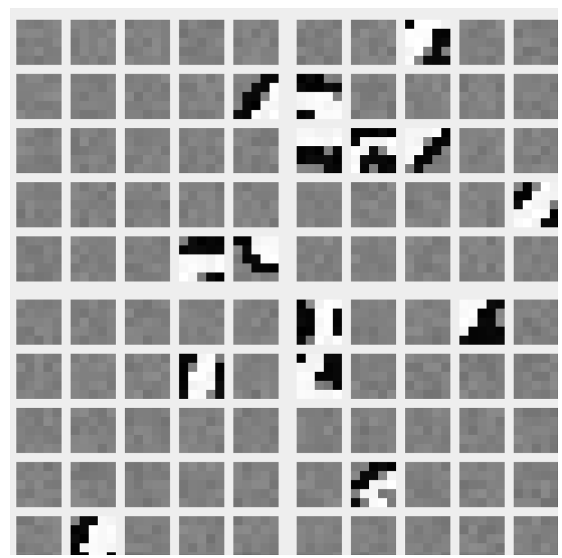

After 30 digits

Fast, incremental learning

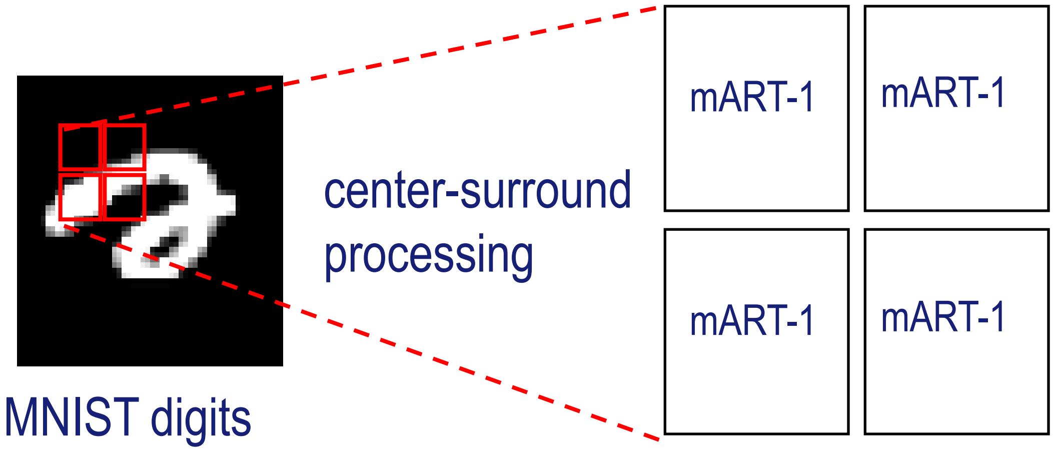

Low vigilance

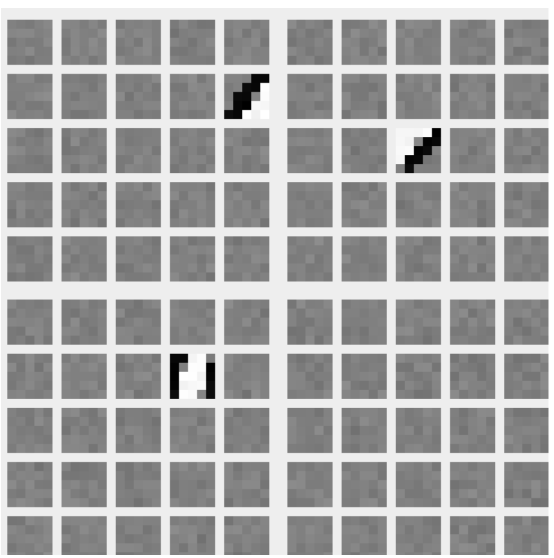

After 1 digit

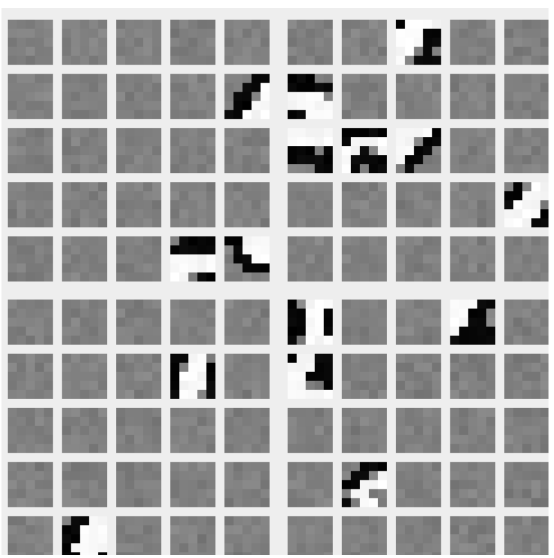

After 30 digits

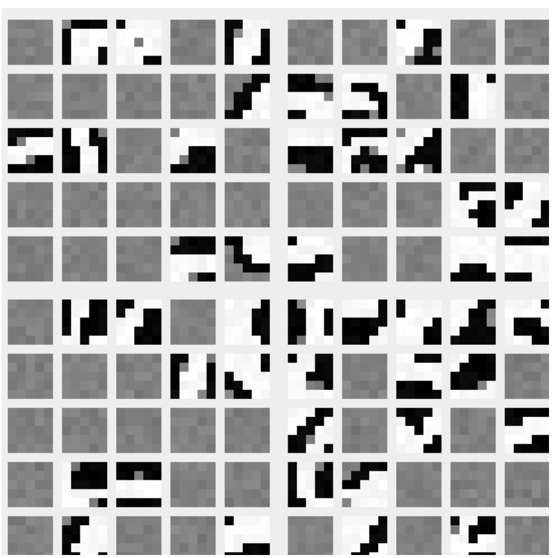

After 60,000 digits

Defect, device-variation strategy: *Ignore them*

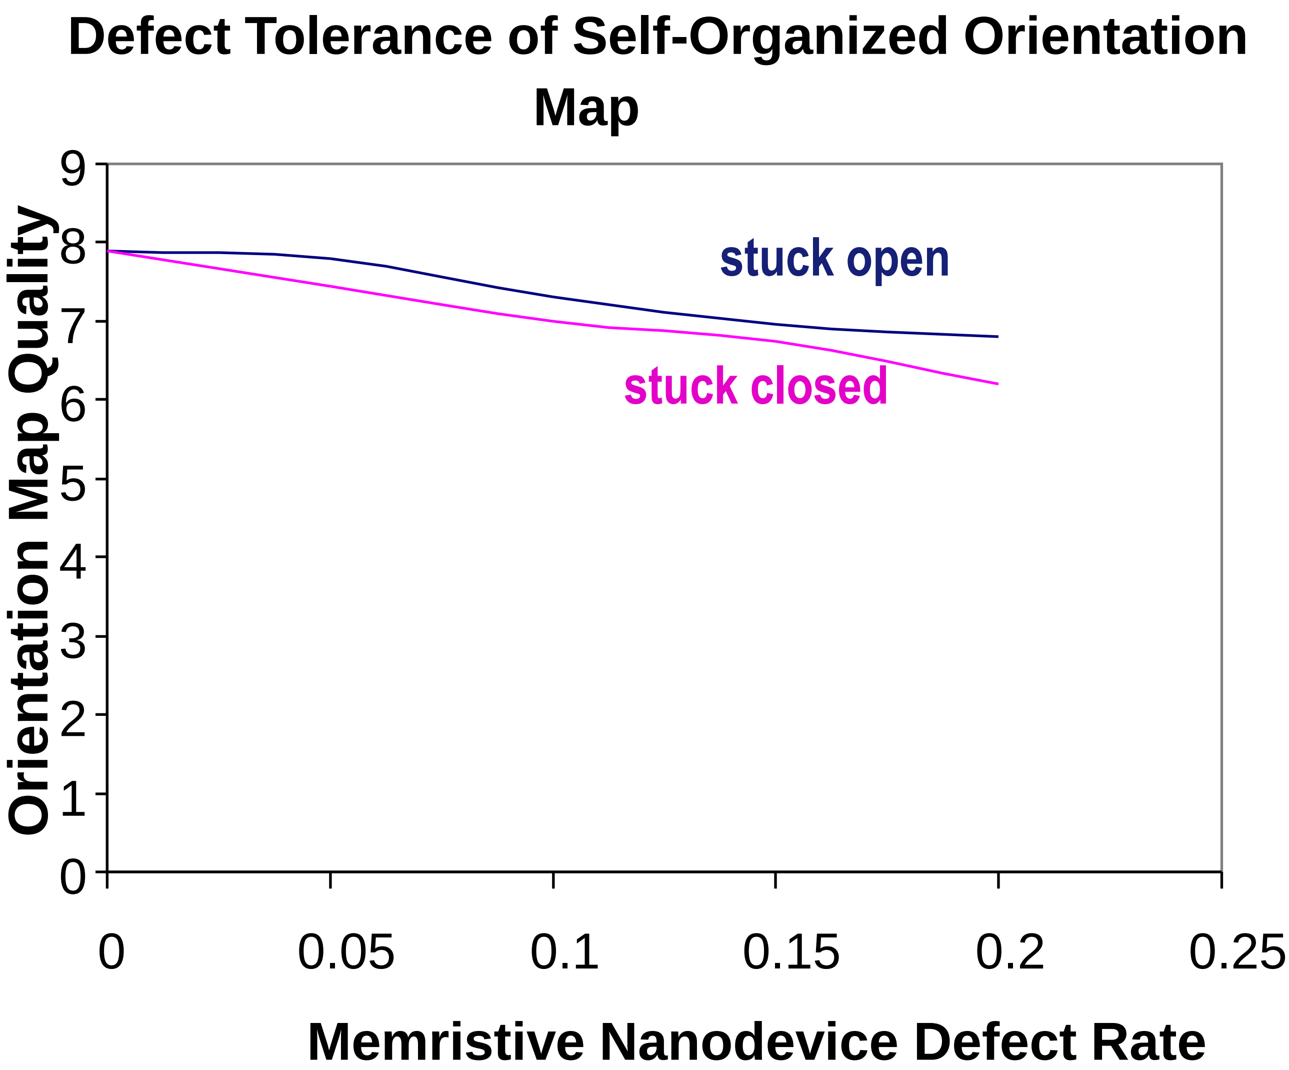

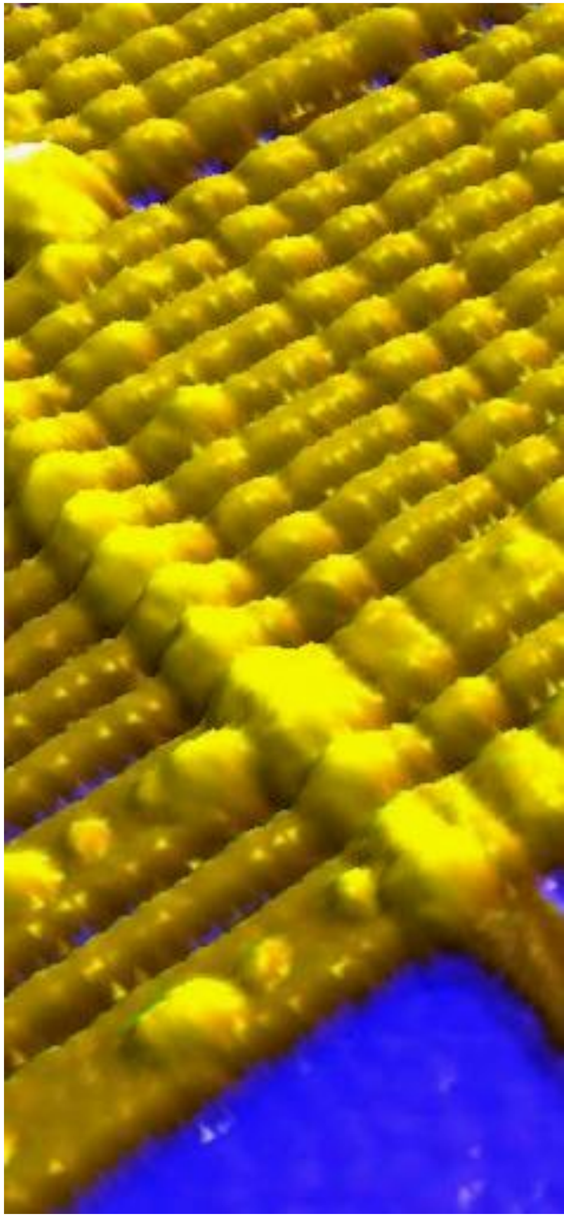

Summary

Memristive nanodevices → synapses

ART networks → regular structure (CMOS + nano)

Fast or slow learning

Tolerant of defects and device variation