Neuromorphic Computing as a new Computing Paradigm

Simon Thorpe

Centre de Recherche Cerveau & Cognition

& SpikeNet Technology SARL,

Toulouse, France

Plan

- Biological and Computer Vision - Can we simulate the Visual System?
  - Classical Artificial Neural Networks
  - Spike based Computing
- STDP-based learning mechanisms
  - Finding the start of repeating patterns
  - Competitive Learning Networks
- Towards Neuromorphic hardware
  - Address Event Representation (AER) Coding
  - Spiking Retinas
  - Spiking Cochlears
  - Memristor based hardware

Plan

- Biological and Computer Vision - Can we simulate the Visual System?
  - Classical Artificial Neural Networks
  - Spike based Computing
- STDP-based learning mechanisms
  - Finding the start of repeating patterns
  - Competitive Learning Networks
- Towards Neuromorphic hardware
  - Address Event Representation (AER) Coding
  - Spiking Retinas
  - Spiking Cochlears
  - Memristor based hardware

Biological and Computer Vision

Common Problems

- Detection, identification and localisation of objects and events in complex dynamically changing natural environments
  - As fast as possible
  - As reliably as possible
  - Using the most energy efficient hardware possible
  - Using the smallest size and weight footprint

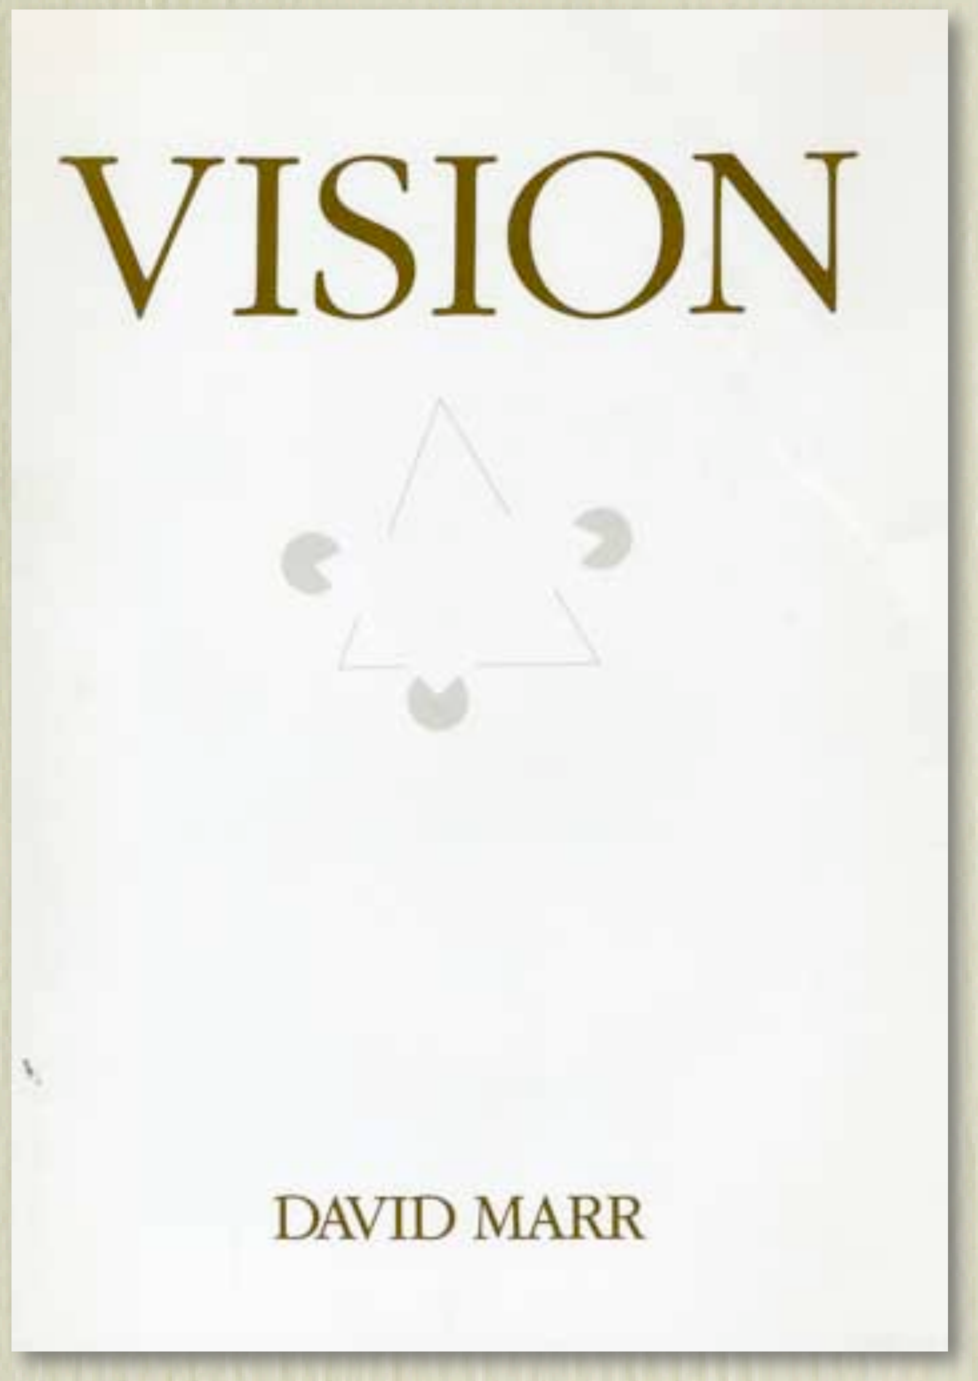

Common Solutions?

- David Marr (1982) "Vision : A computational investigation into the Human Representation and Processing of Visual Information"
- Recent years - is there convergence?

Can we simulate the visual system?

Brain

- 86 billion neurons
  - 16 billion in the neocortex
  - 4 billion for vision?

Questions

Are we going to be able to implement brain style computing with conventional computing?

How many teraflops does the brain need?

How much memory bandwidth?

Response

It depends if we can understand how the brain computes

Classic Neural Computing

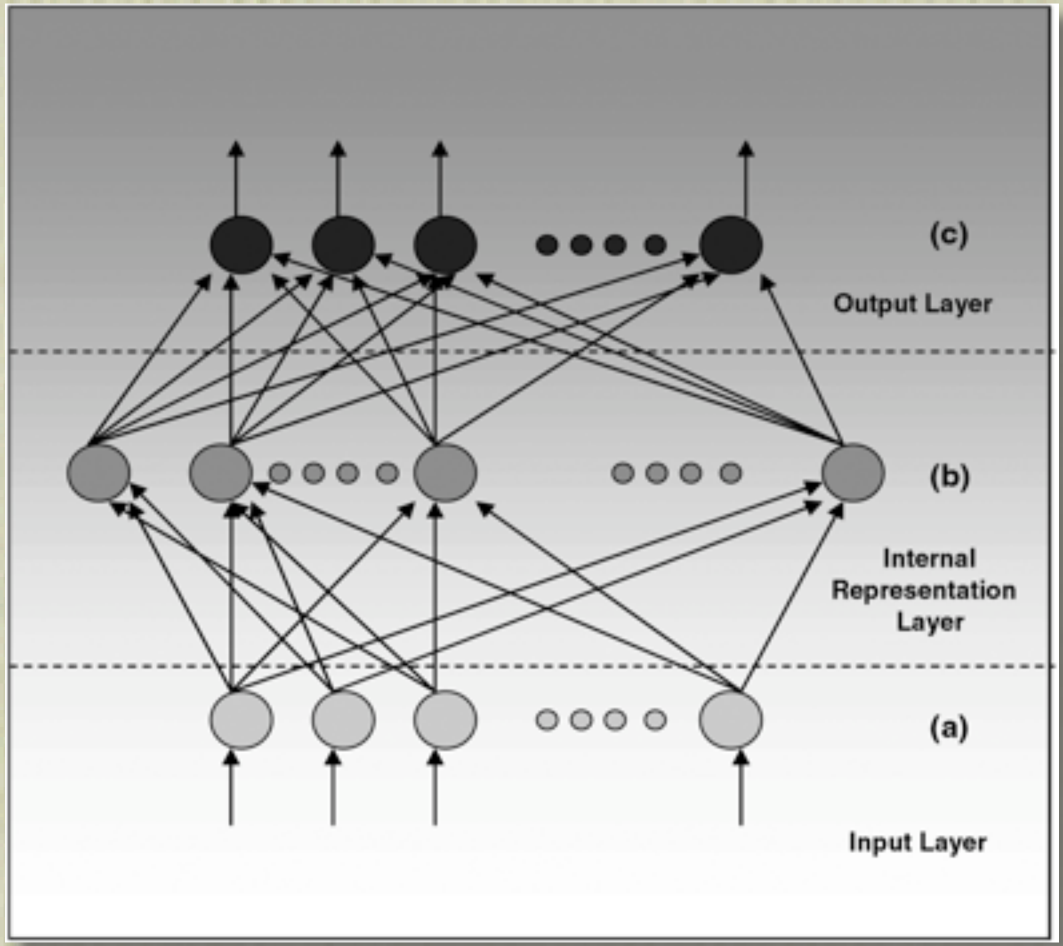

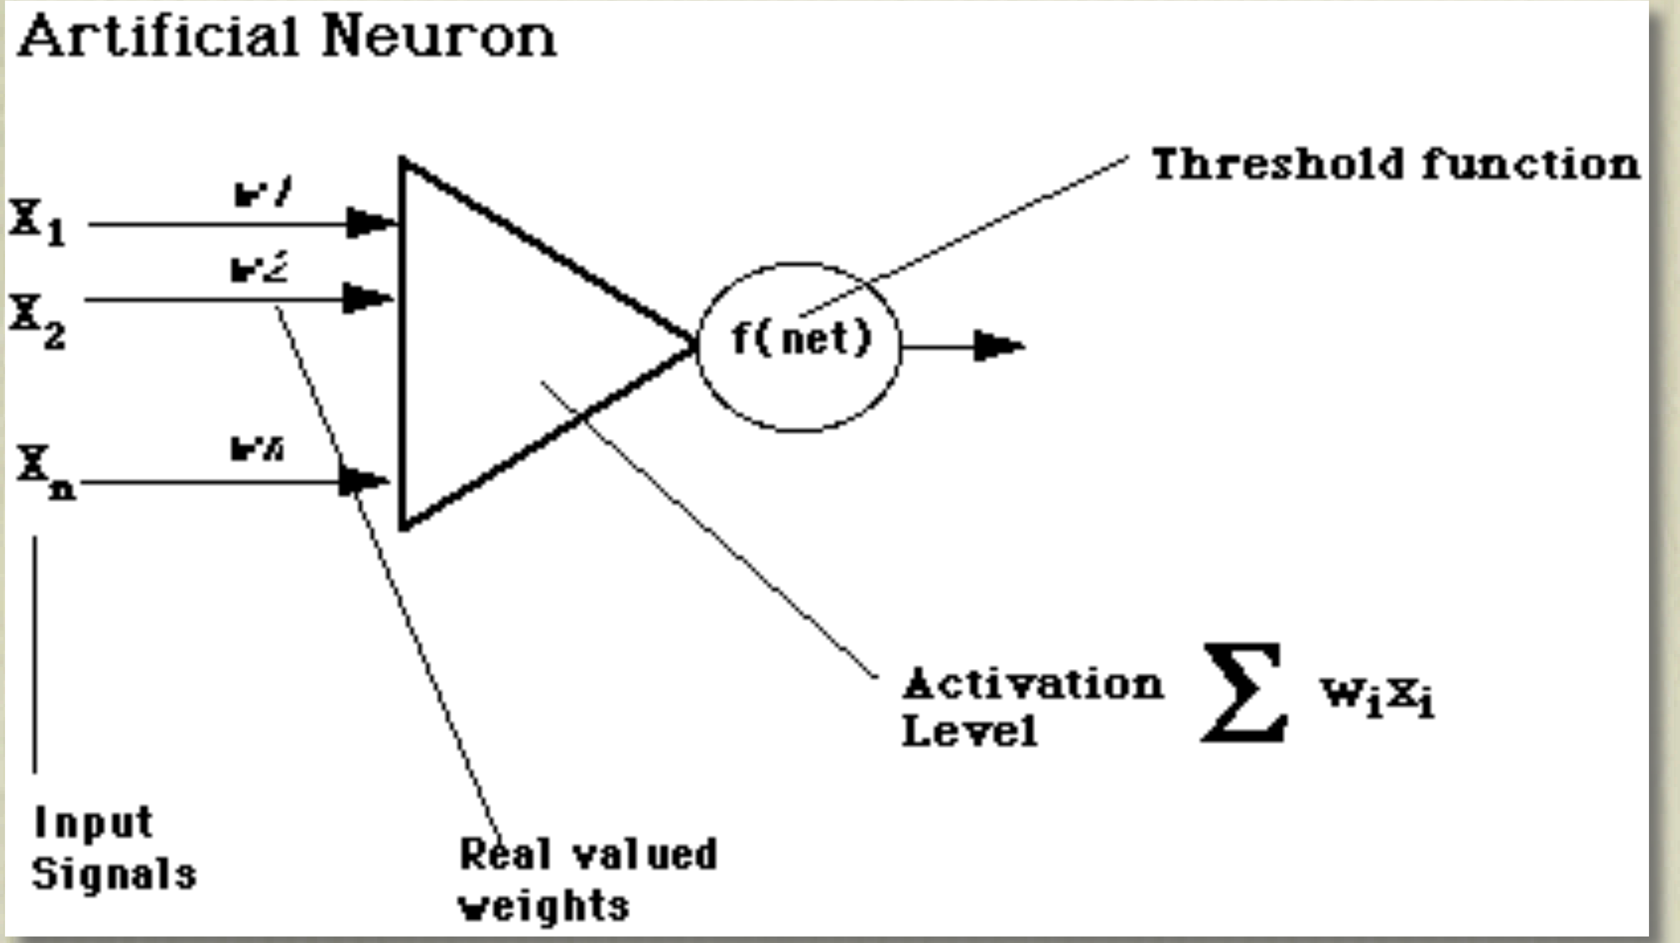

- To simulate the visual system
  - 4 billion neurons
  - 10000 connections each
  - Update at 1 kHz
  - 40 Petaflops

The Blue Brain Project

NATURE REVIEWS | NEUROSCIENCE FEBRUARY 2006

Henry Markram

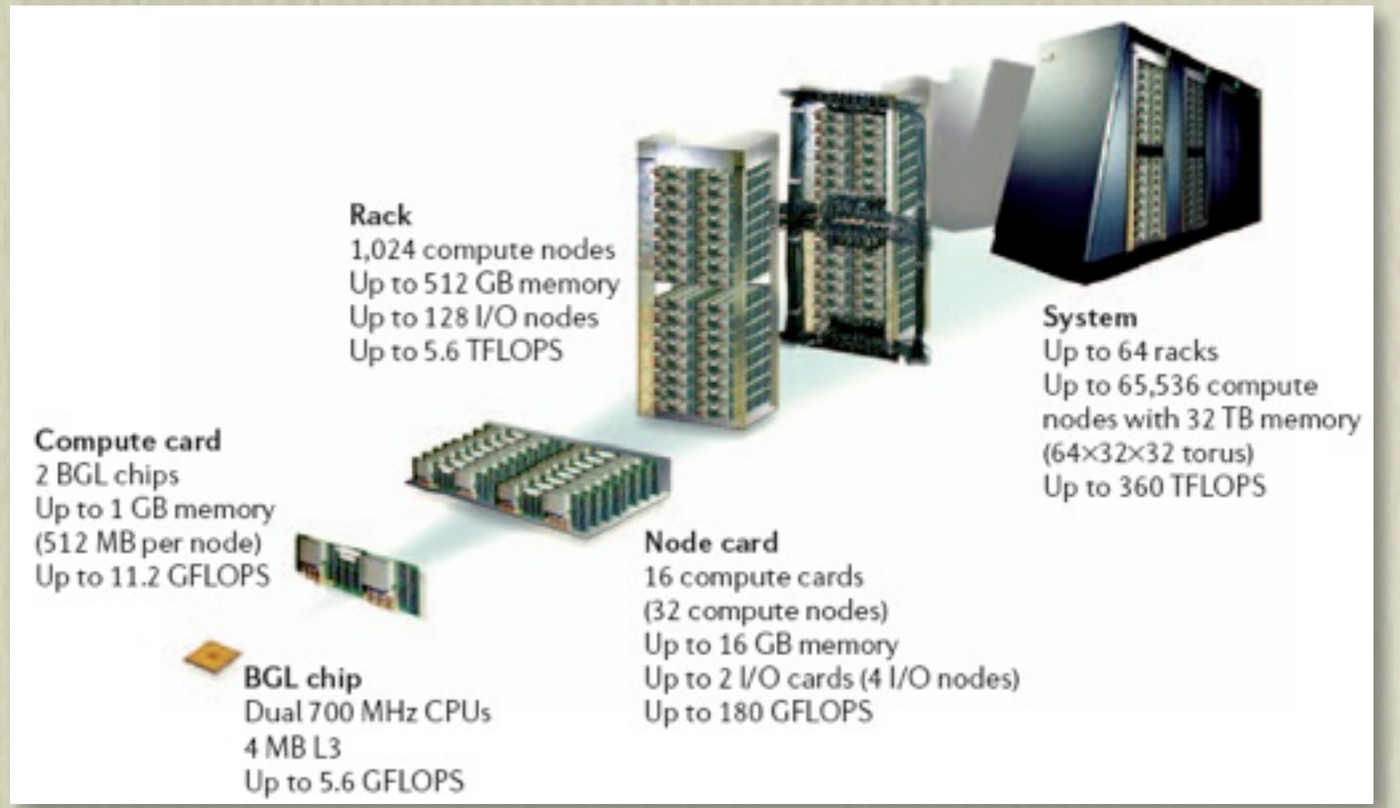

What's missing?

- Real brains use spikes

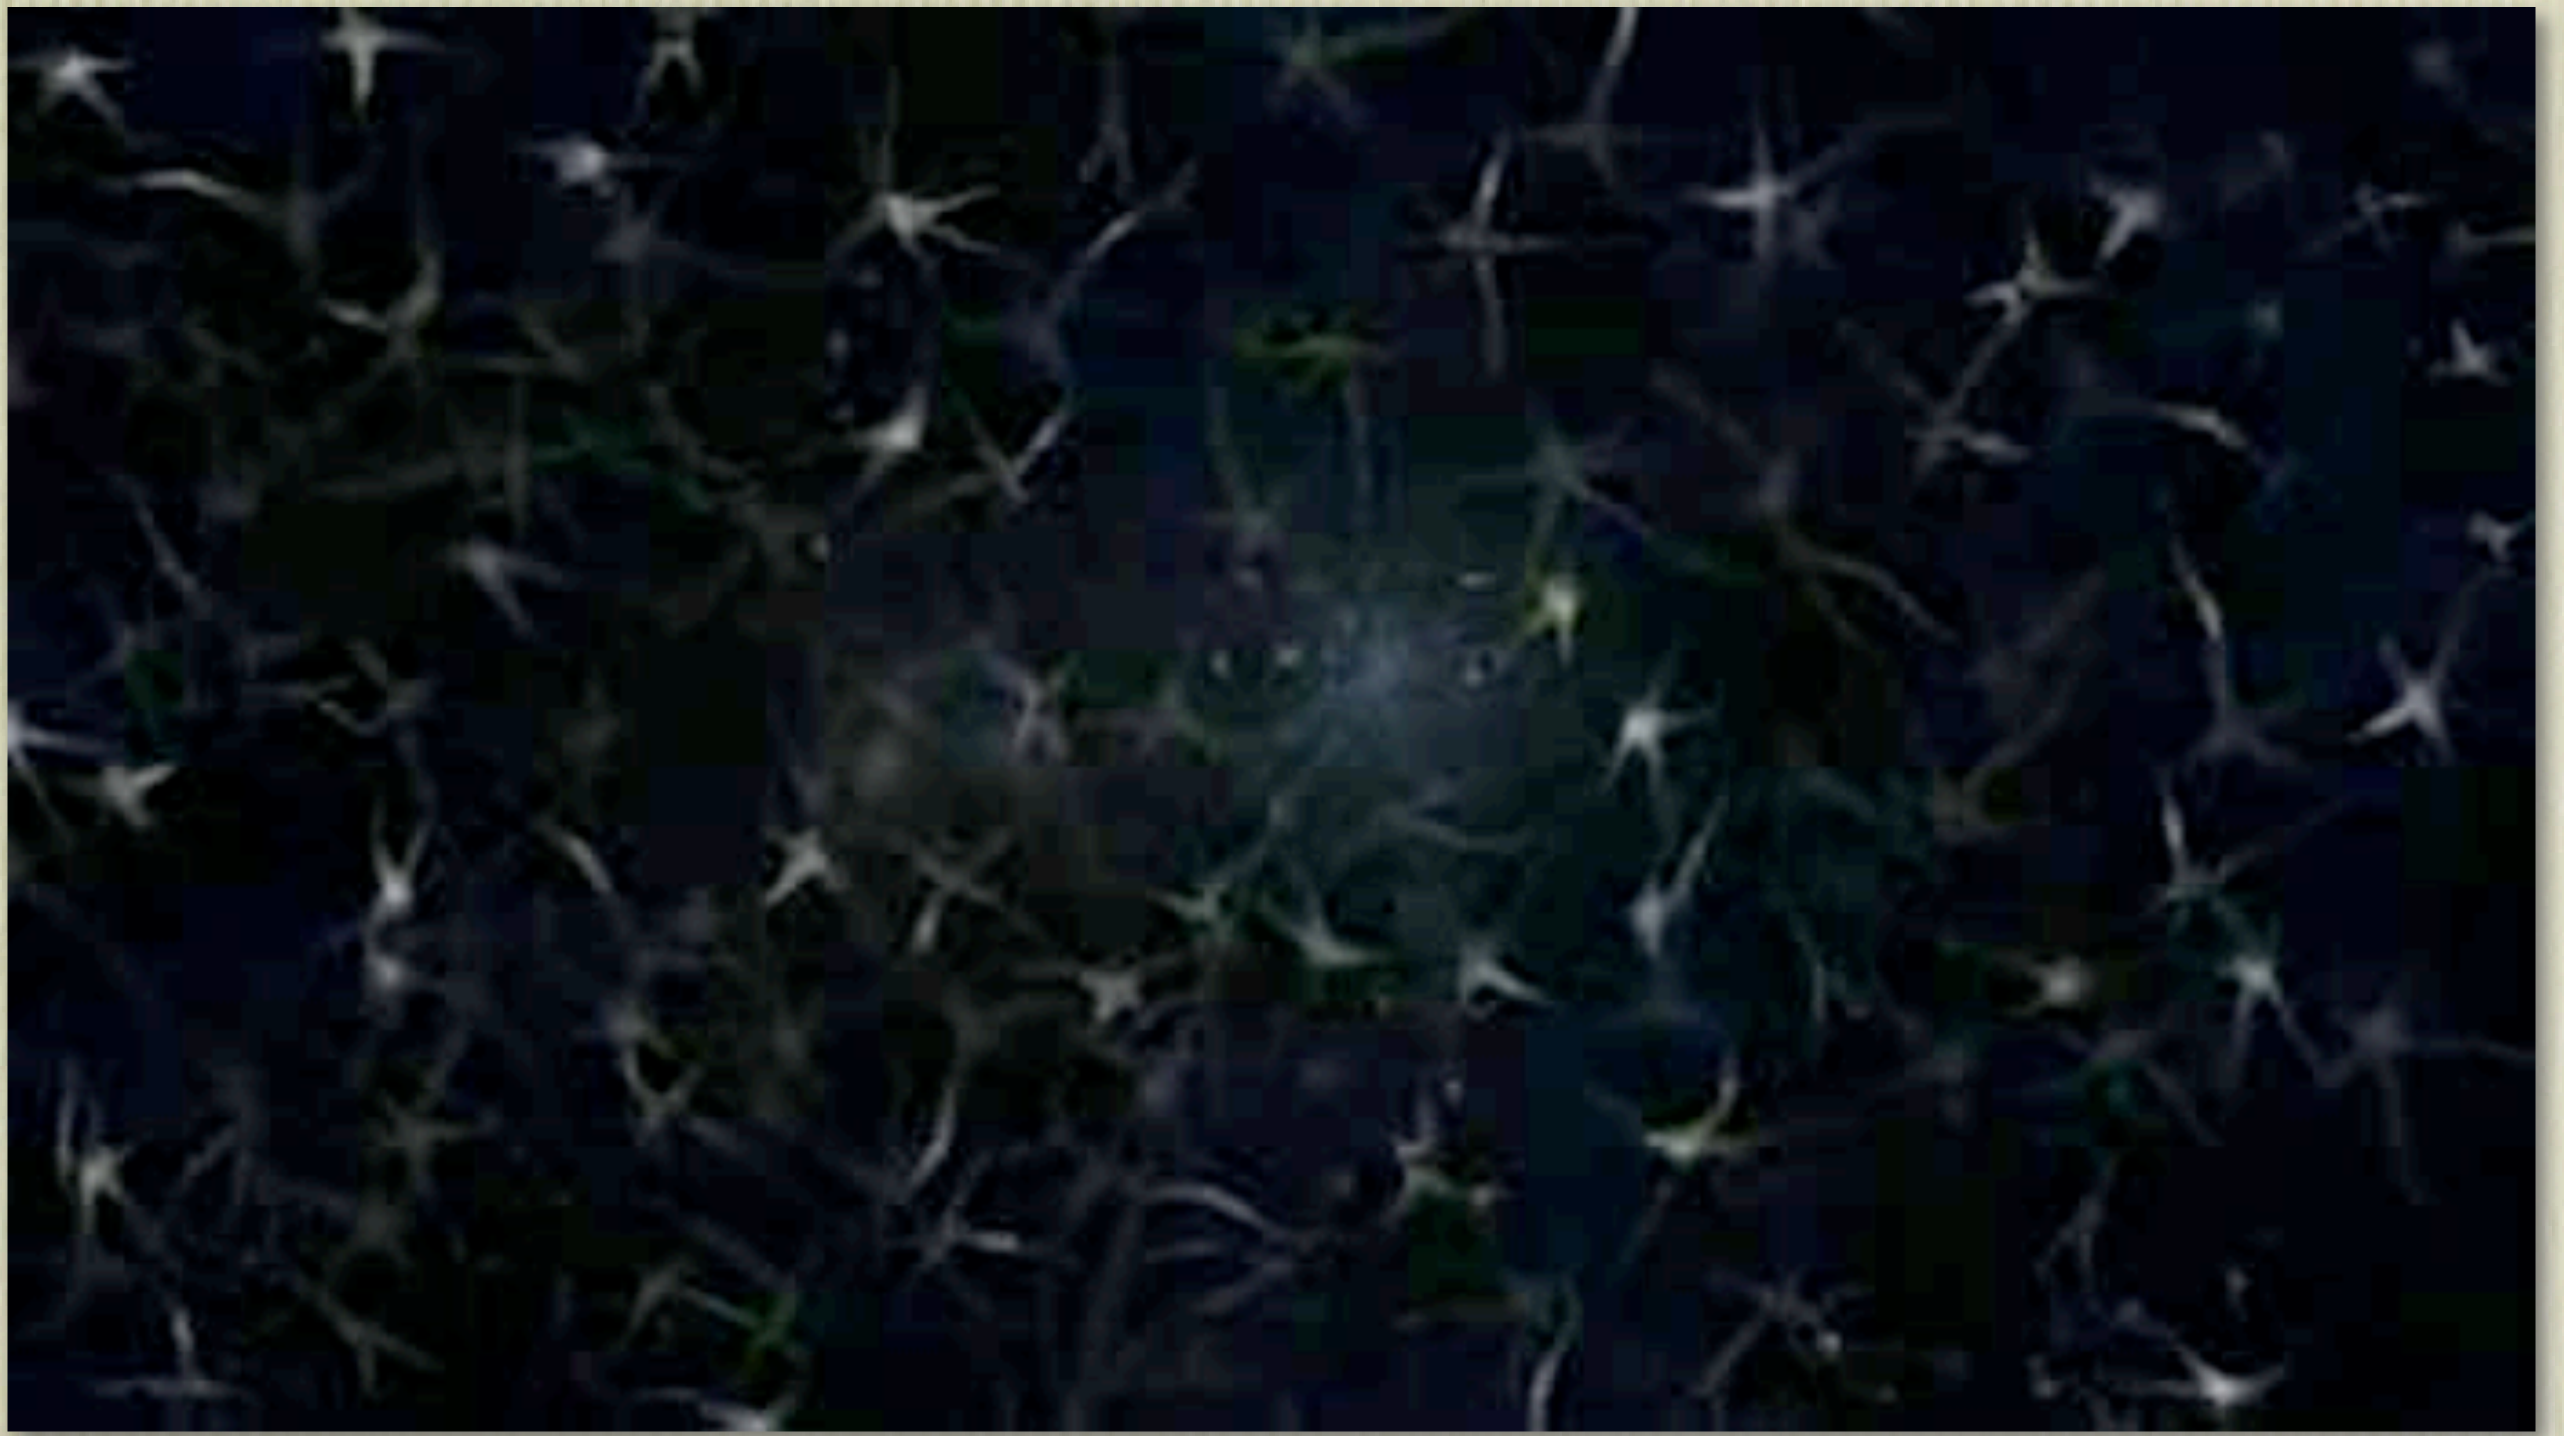

The Classic View

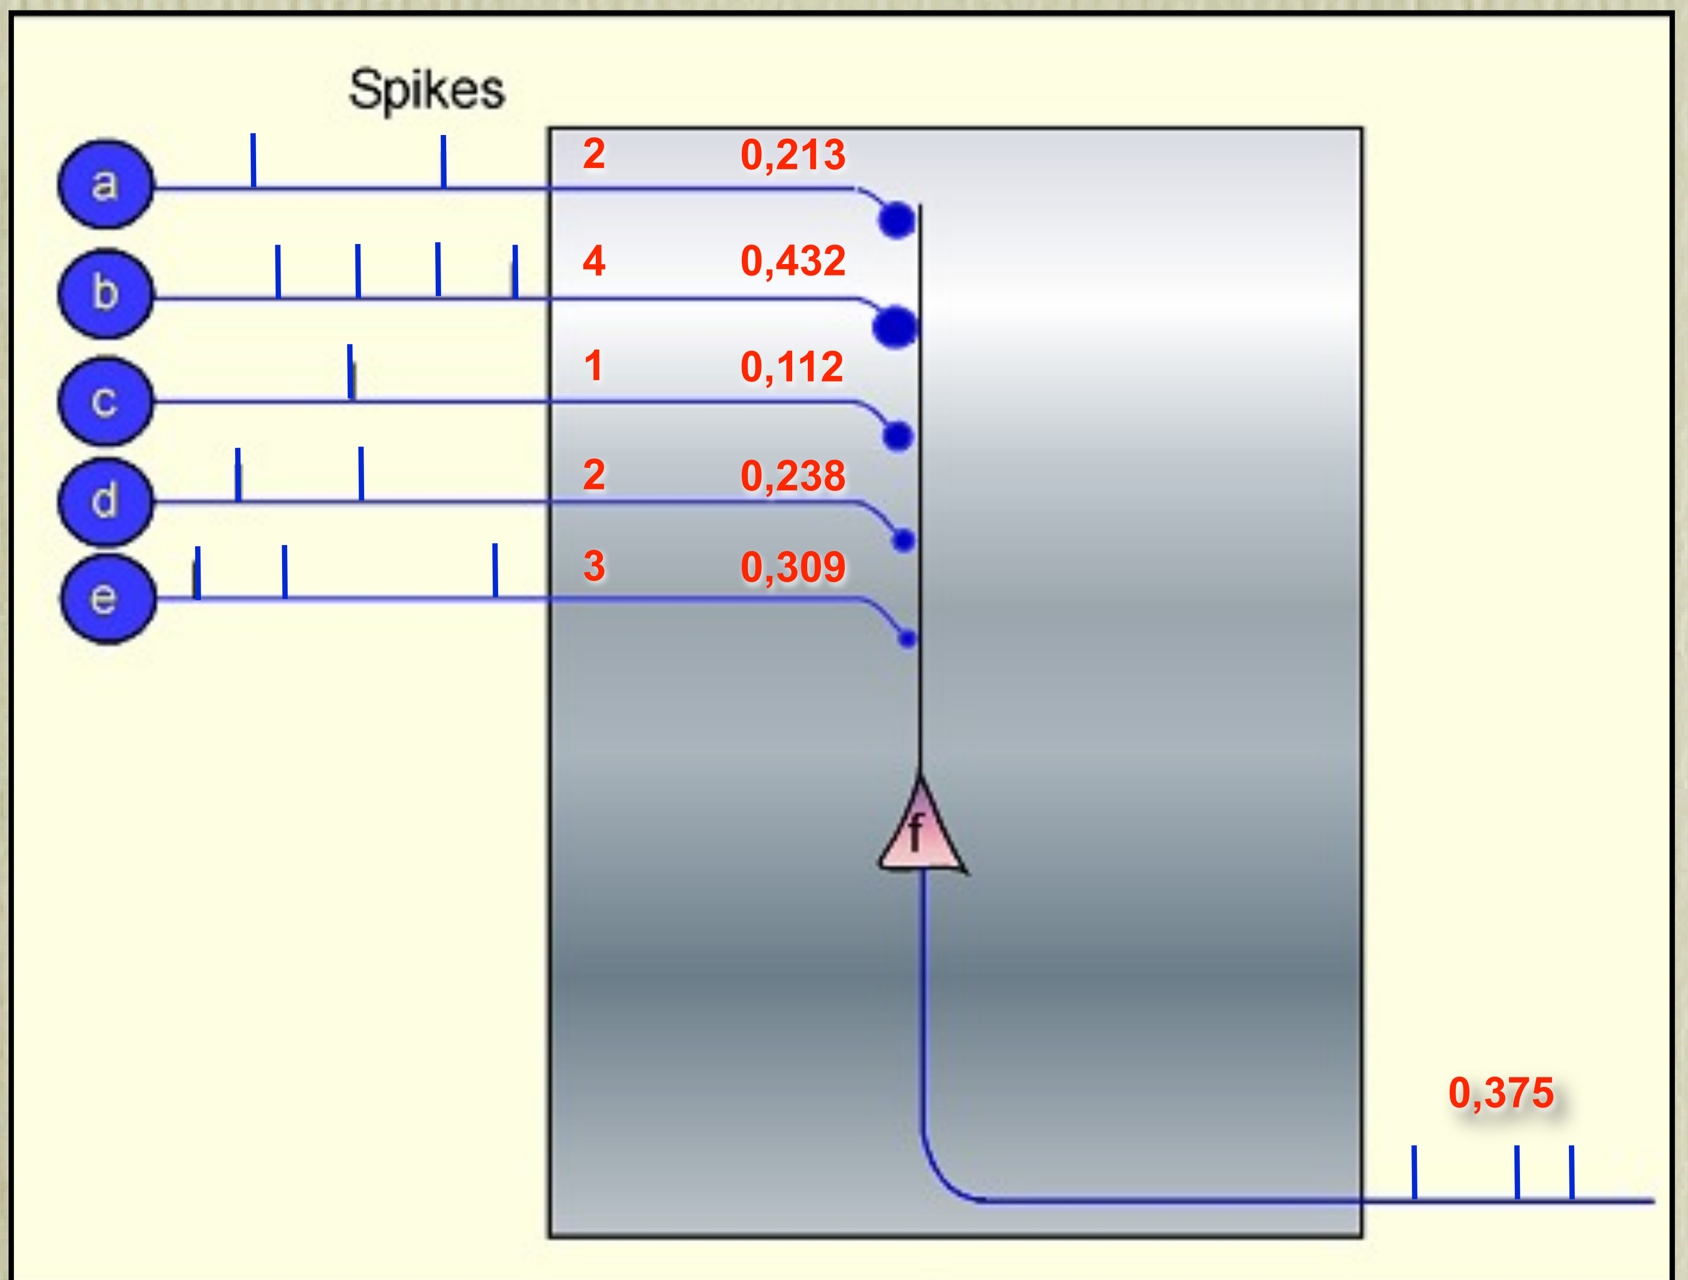

- Spikes don't really matter
- Neurons send floating point numbers
- The floating point numbers are transformed into spikes trains using a Poisson process
- God plays dice with spike generation

Temporal Coding Option

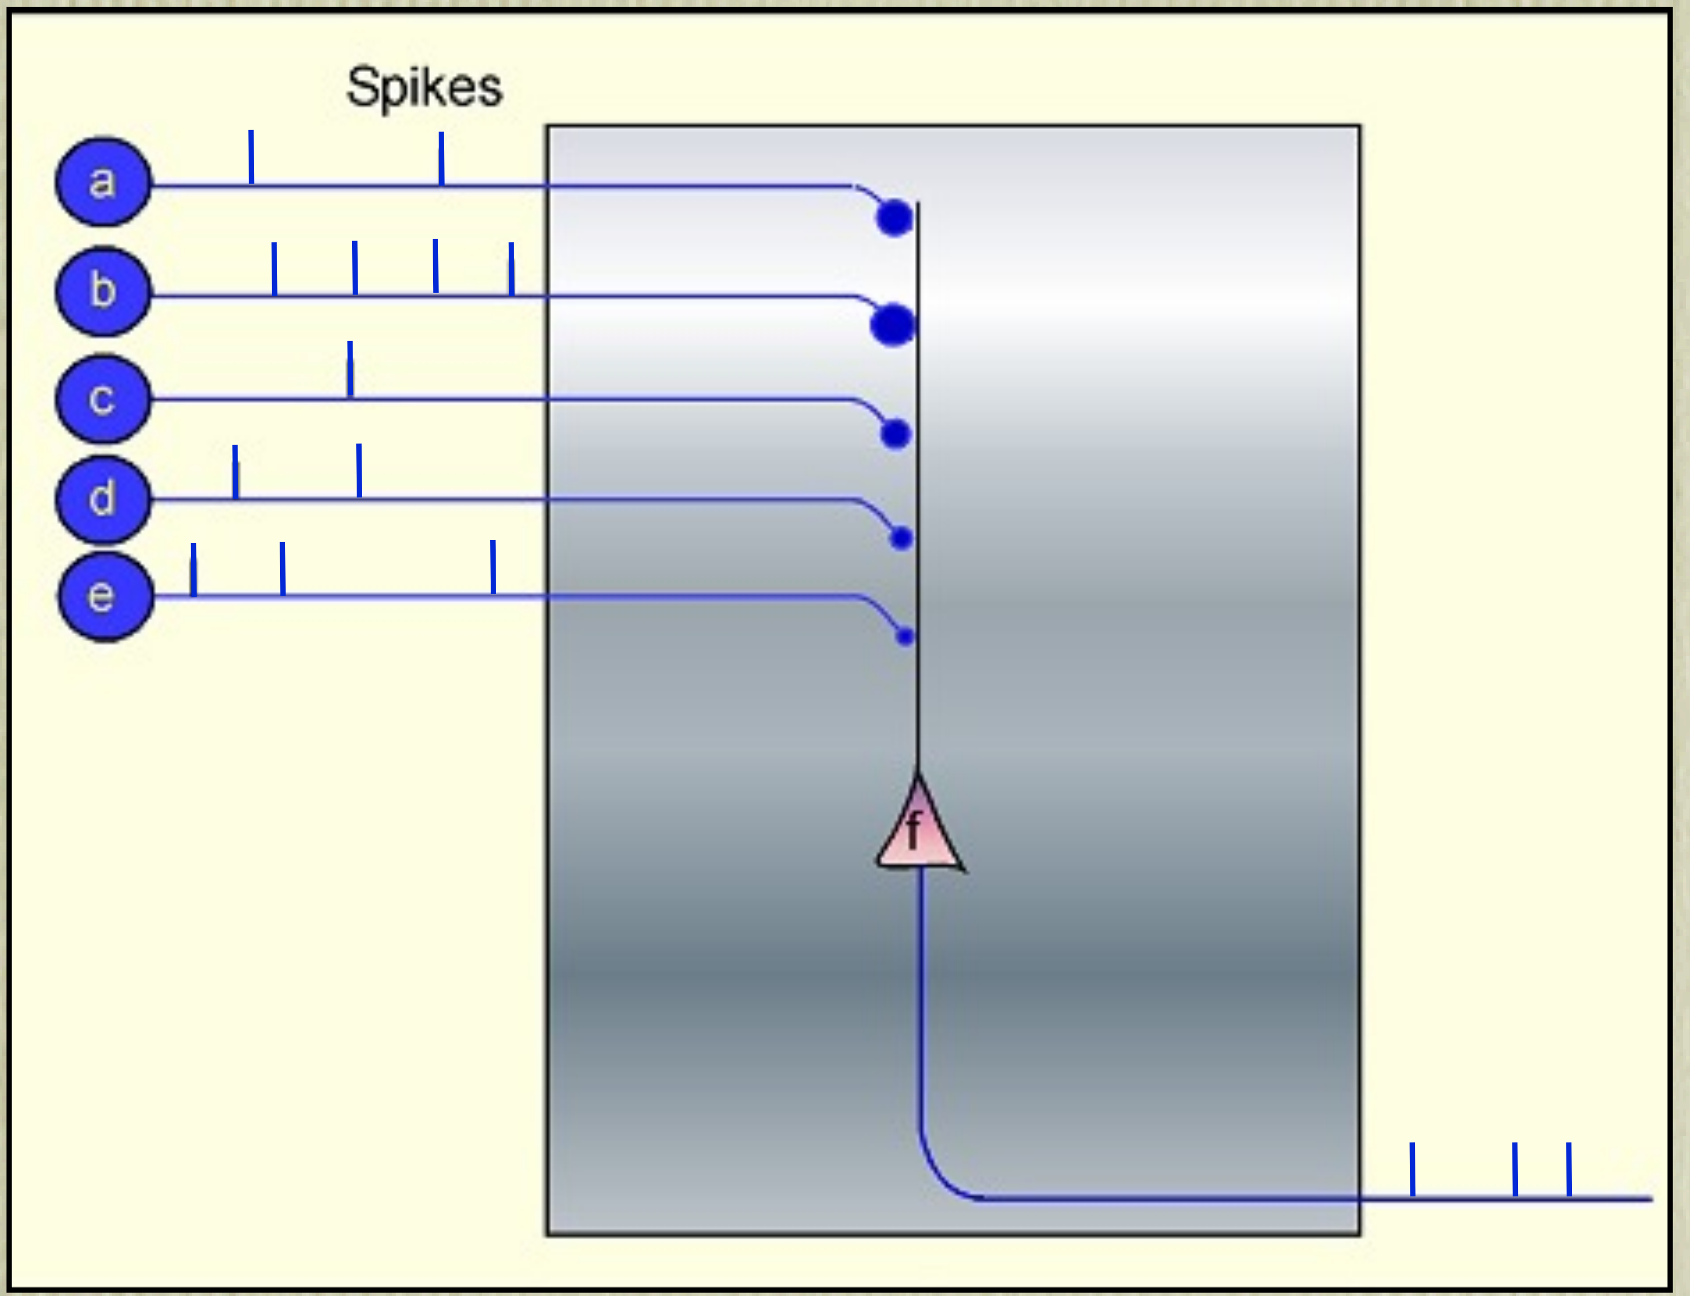

- Spikes do really matter
- The temporal patterning of spikes across neurons is critical for computation
  - Synchrony
  - Repeating patterns
  - etc
- The apparent noise in spiking is unexplained variation

Simon Thorpe's Version

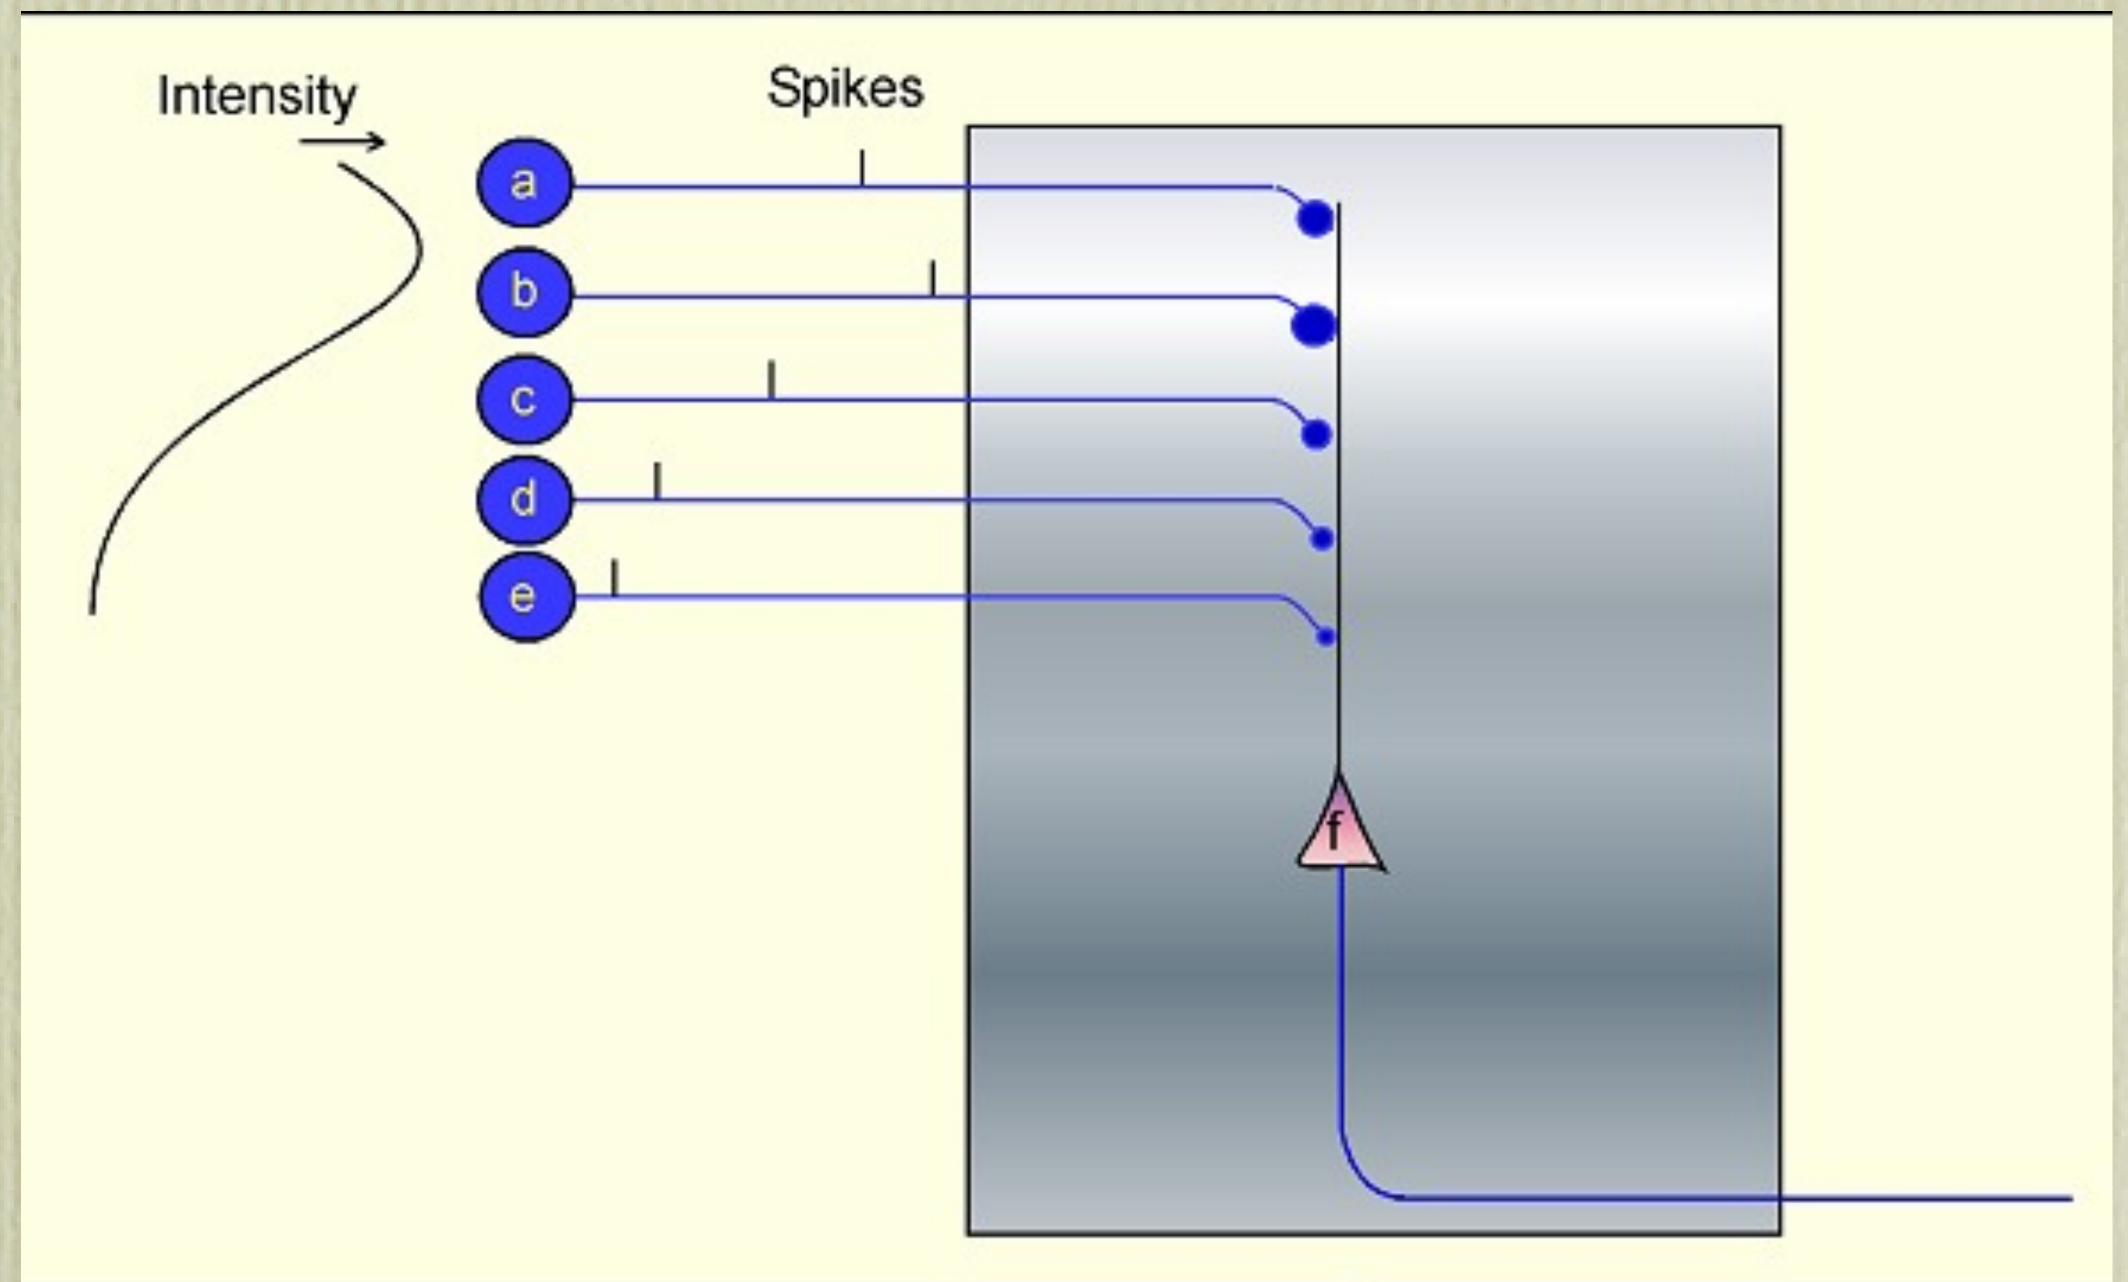

- Ordering of spikes is critical
- The most activated neurons fire first
- Temporal coding is used even for stimuli that are not temporally structured
- Computation theoretically possible even when each neuron emits one spike

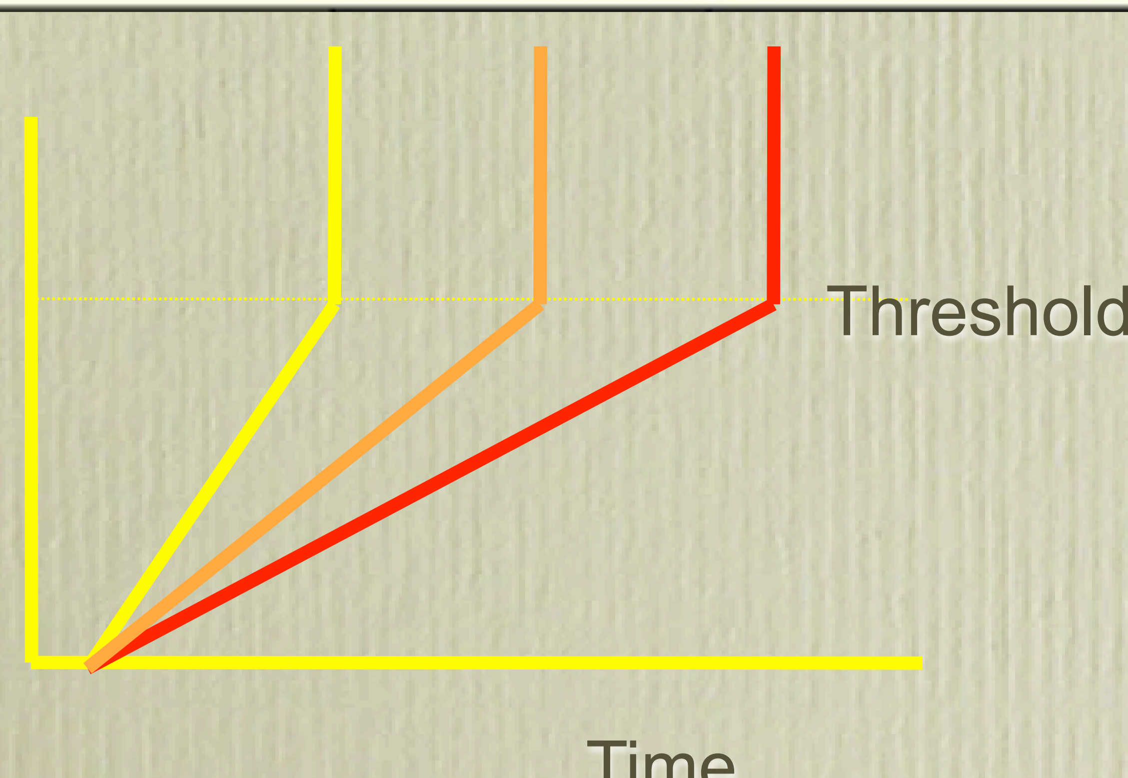

Strong Stimulus

Why compute with spikes?

- Speed constraints
  - We can initiate saccades to animals in complex scenes in 120-130 ms
  - We can initiate saccades to faces from 100-110 ms
  - These saccades are remarkably accurate
  - Neurons can only have time to fire one spike
  - Spike order based coding

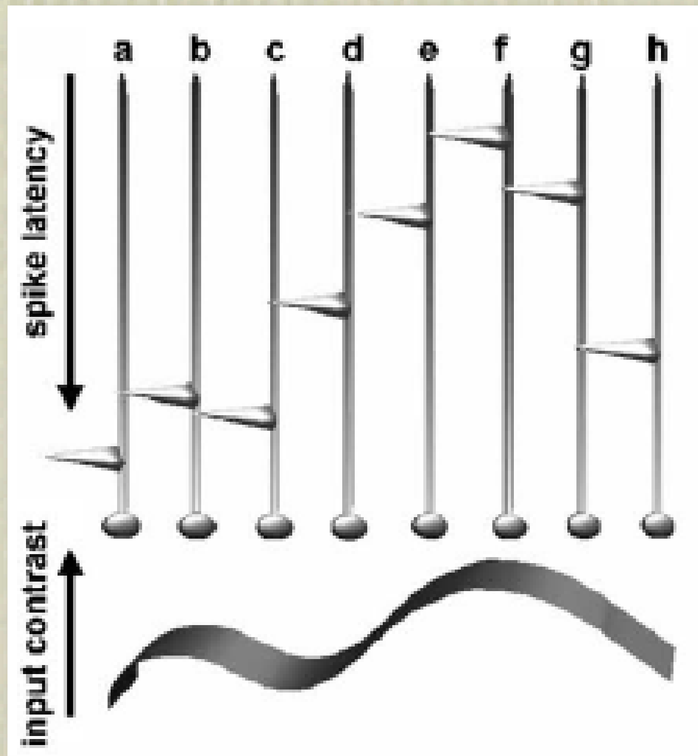

- Spike-based pattern learning
  - STDP finds recurring patterns of spikes
  - Neurons can find the start of the pattern

Plan

- Biological and Computer Vision - Can we simulate the Visual System?
  - Classical Artificial Neural Networks
  - Spike based Computing
- STDP-based learning mechanisms
  - Finding the start of repeating patterns
  - Competitive Learning Networks
- Towards Neuromorphic hardware
  - Address Event Representation (AER) Coding
  - Spiking Retinas
  - Spiking Cochlears
  - Memristor based hardware

A key mechanism

- Spike-Time Dependent Plasticity (STDP)
  - Synapses that fire before the target neuron get strenthened
  - Synapses that fire after the target neuron get weakened

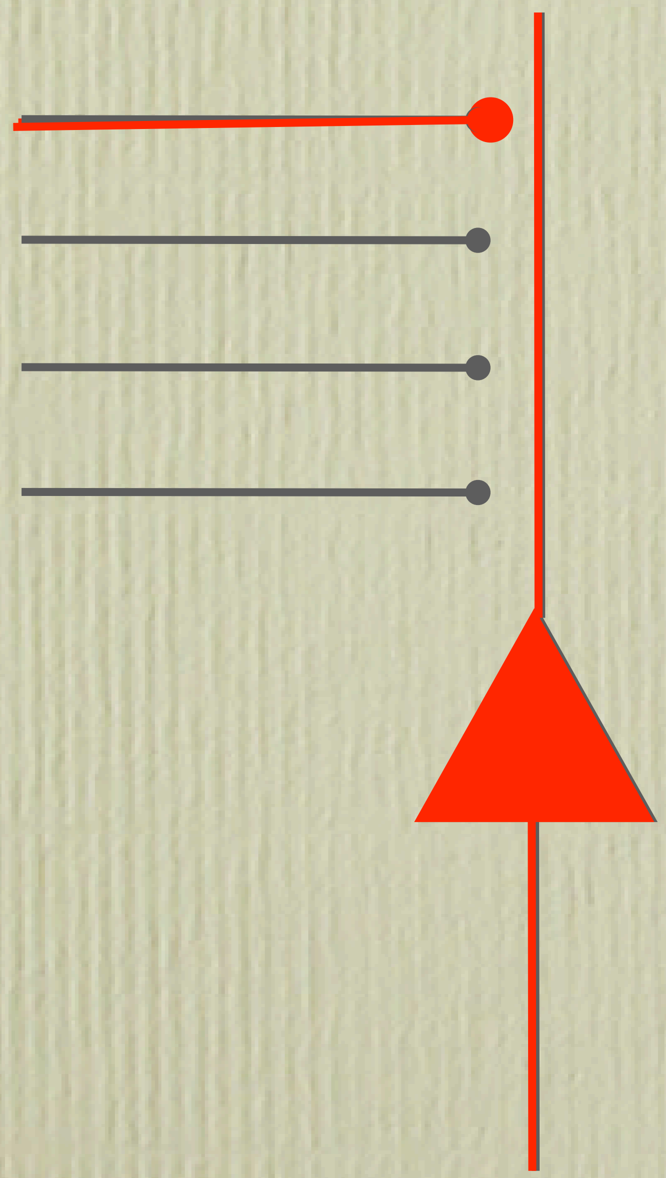

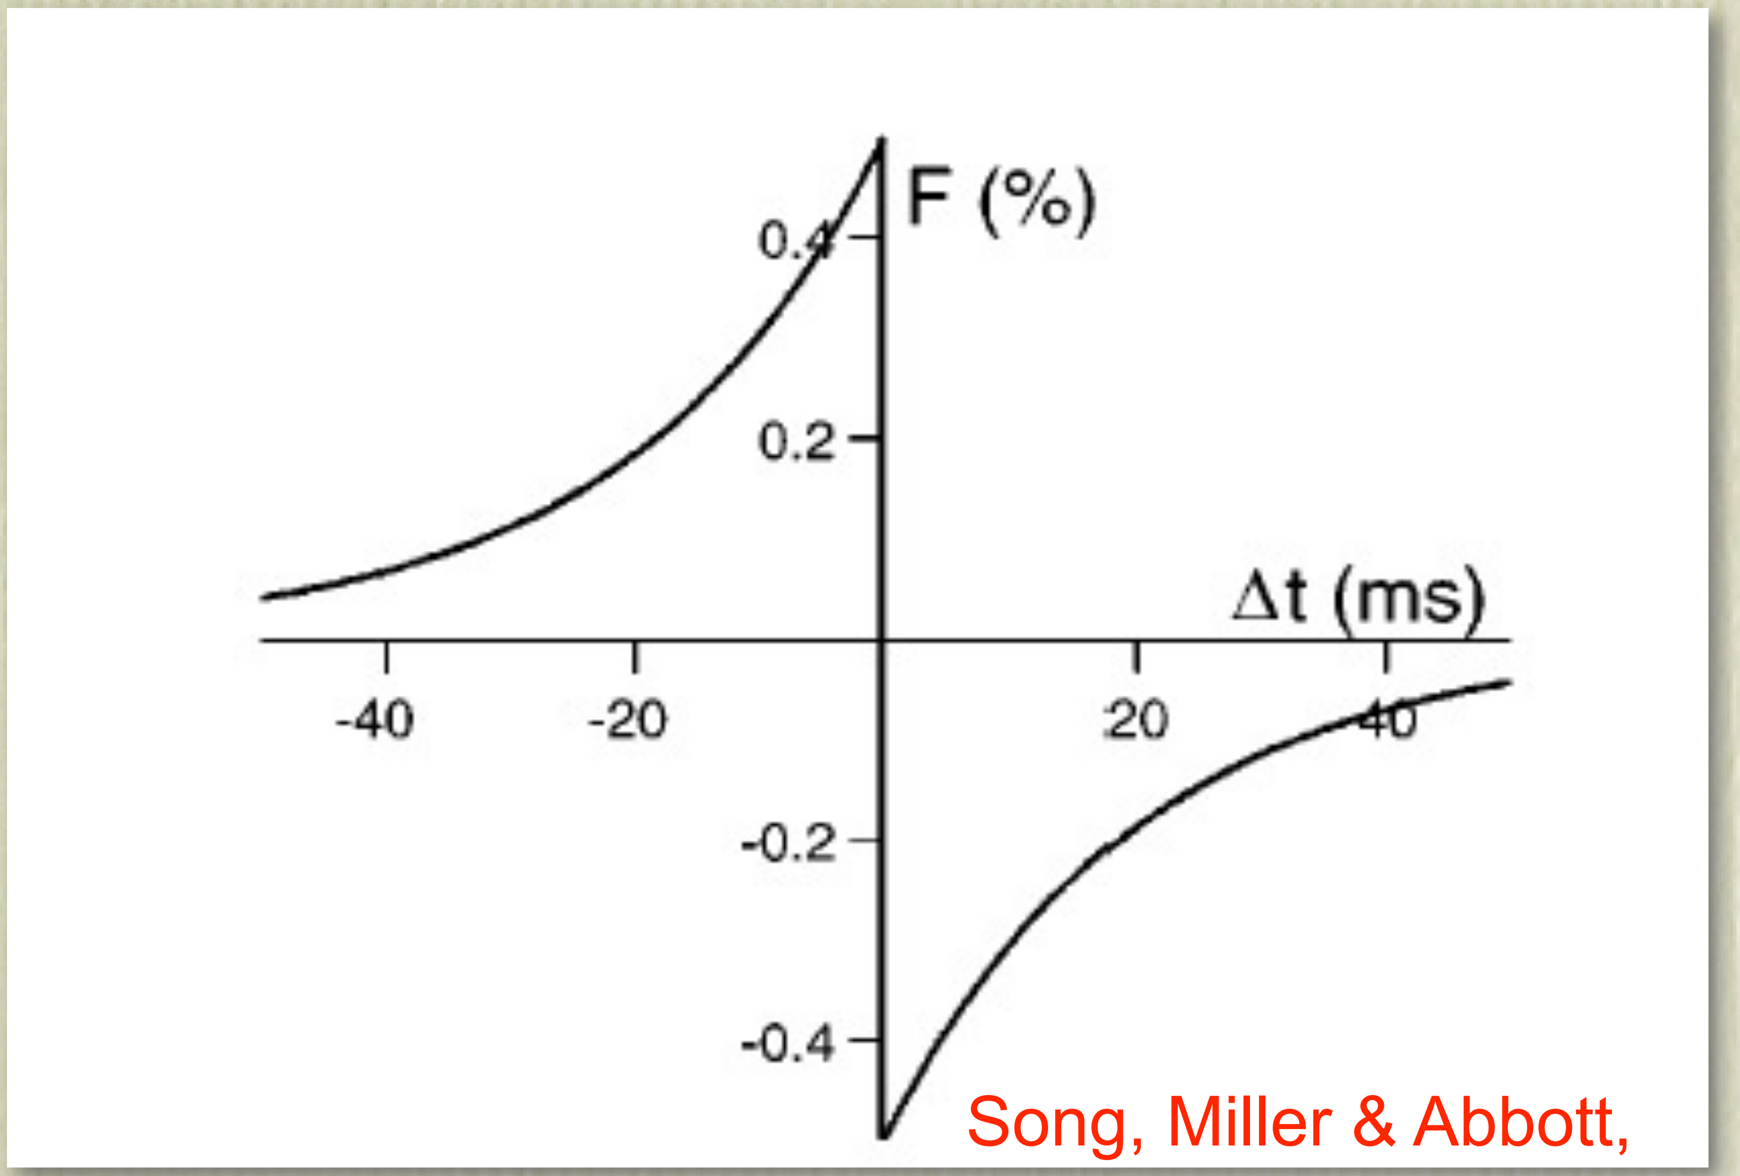

- A natural consequence
  - High synaptic weights will concentrate on early firing inputs

Spike Time Dependent Plasticity

Threshold = 2

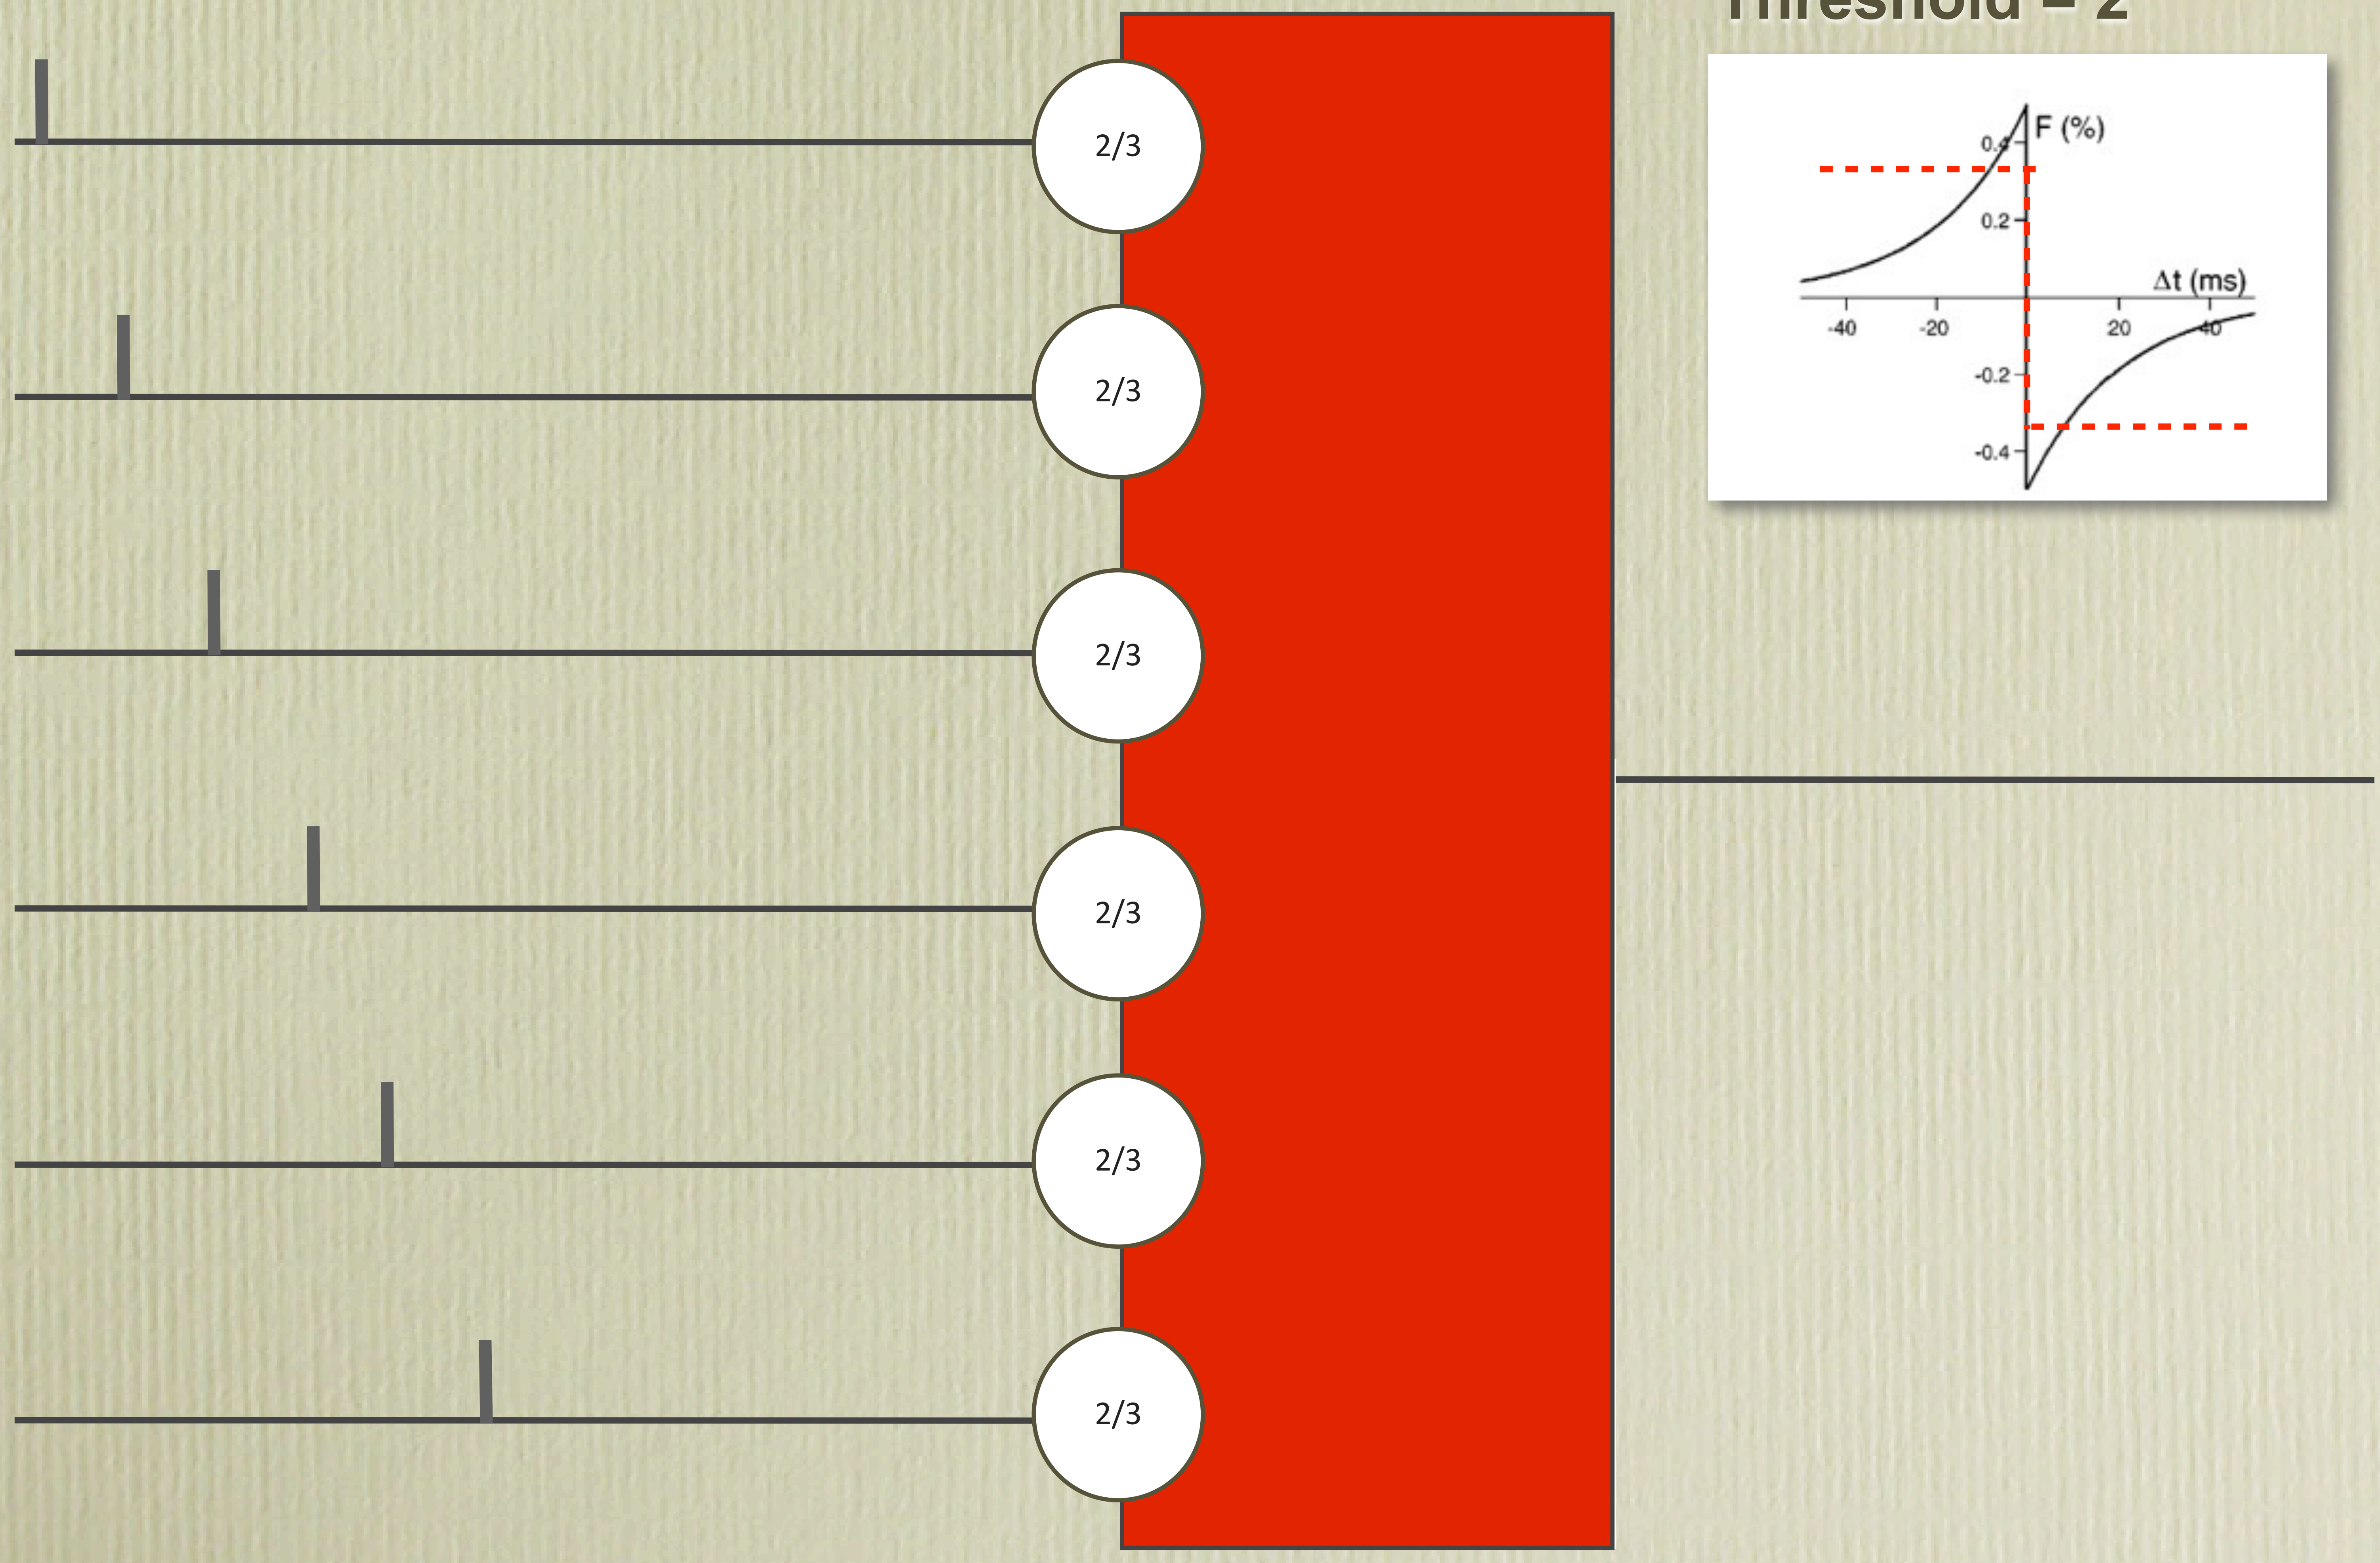

Spike Time Dependent Plasticity

Threshold = 2

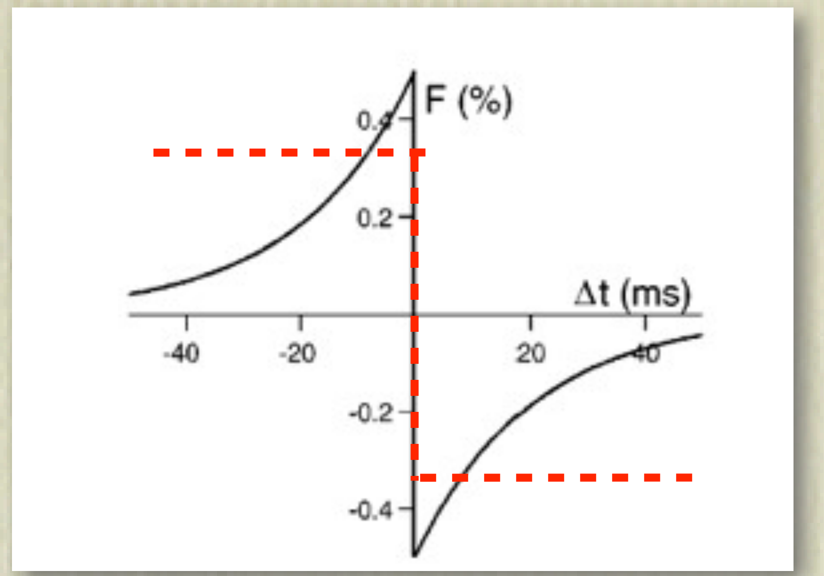

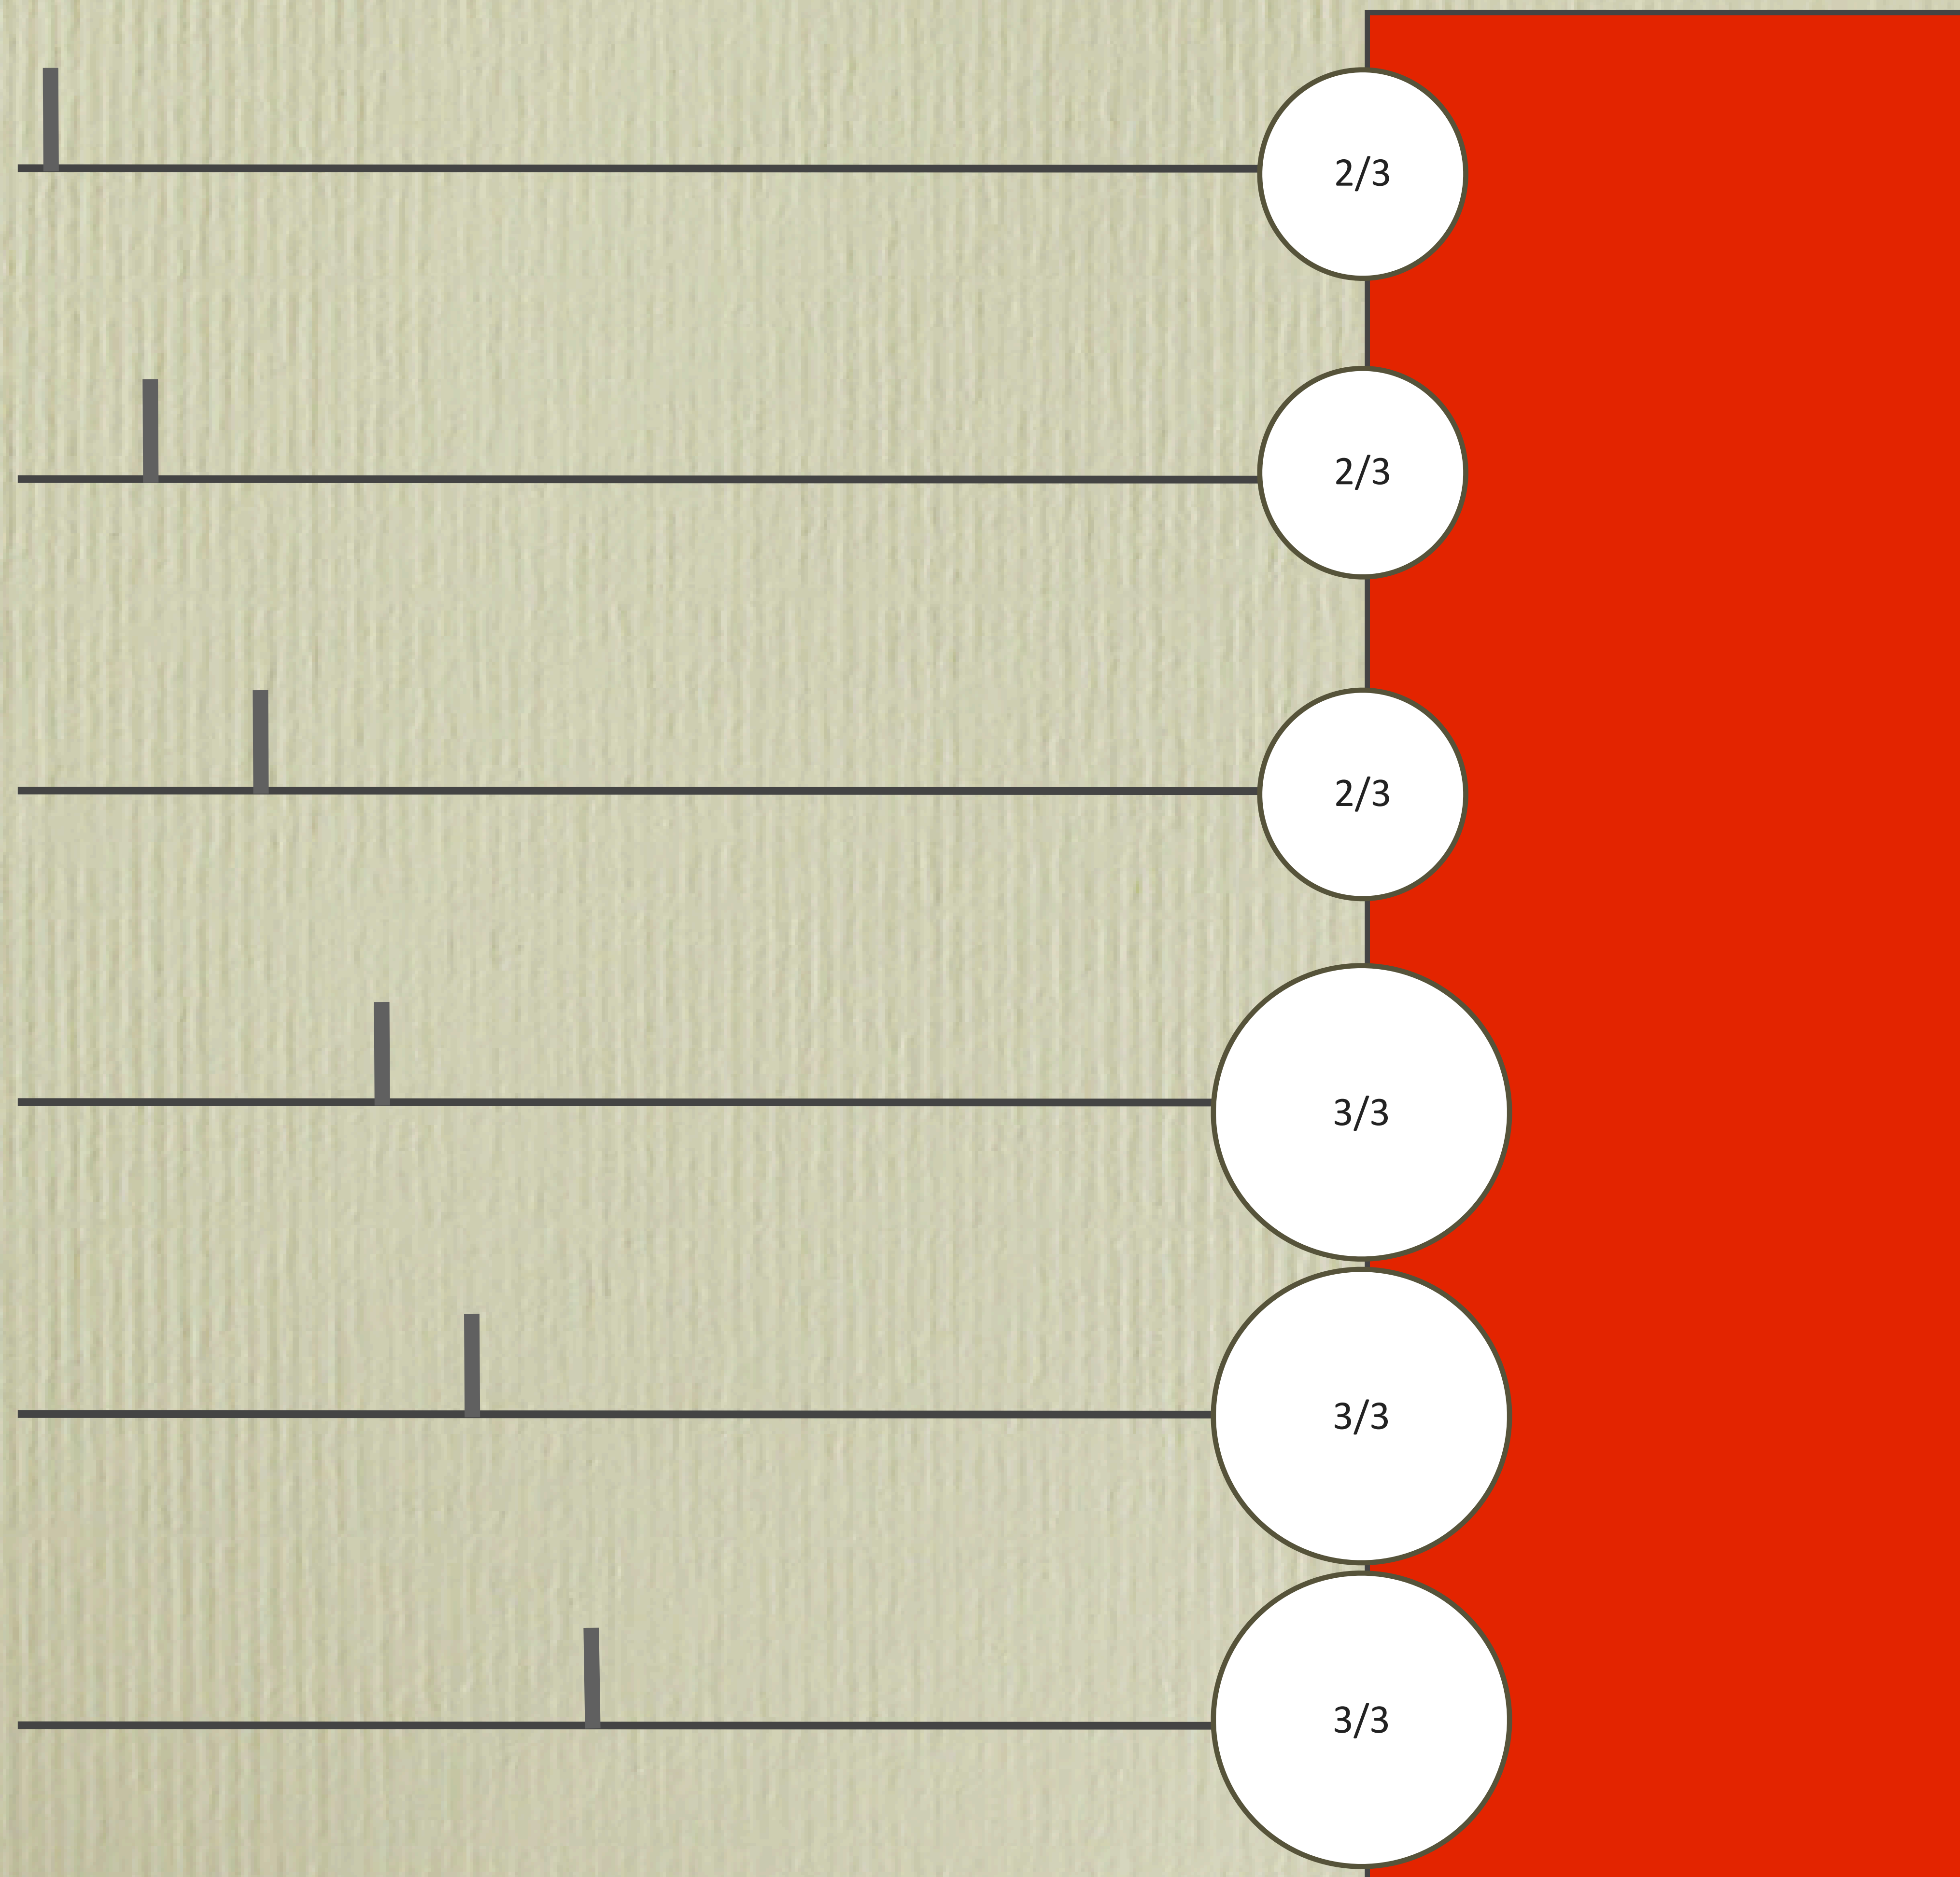

Spike Time Dependent Plasticity

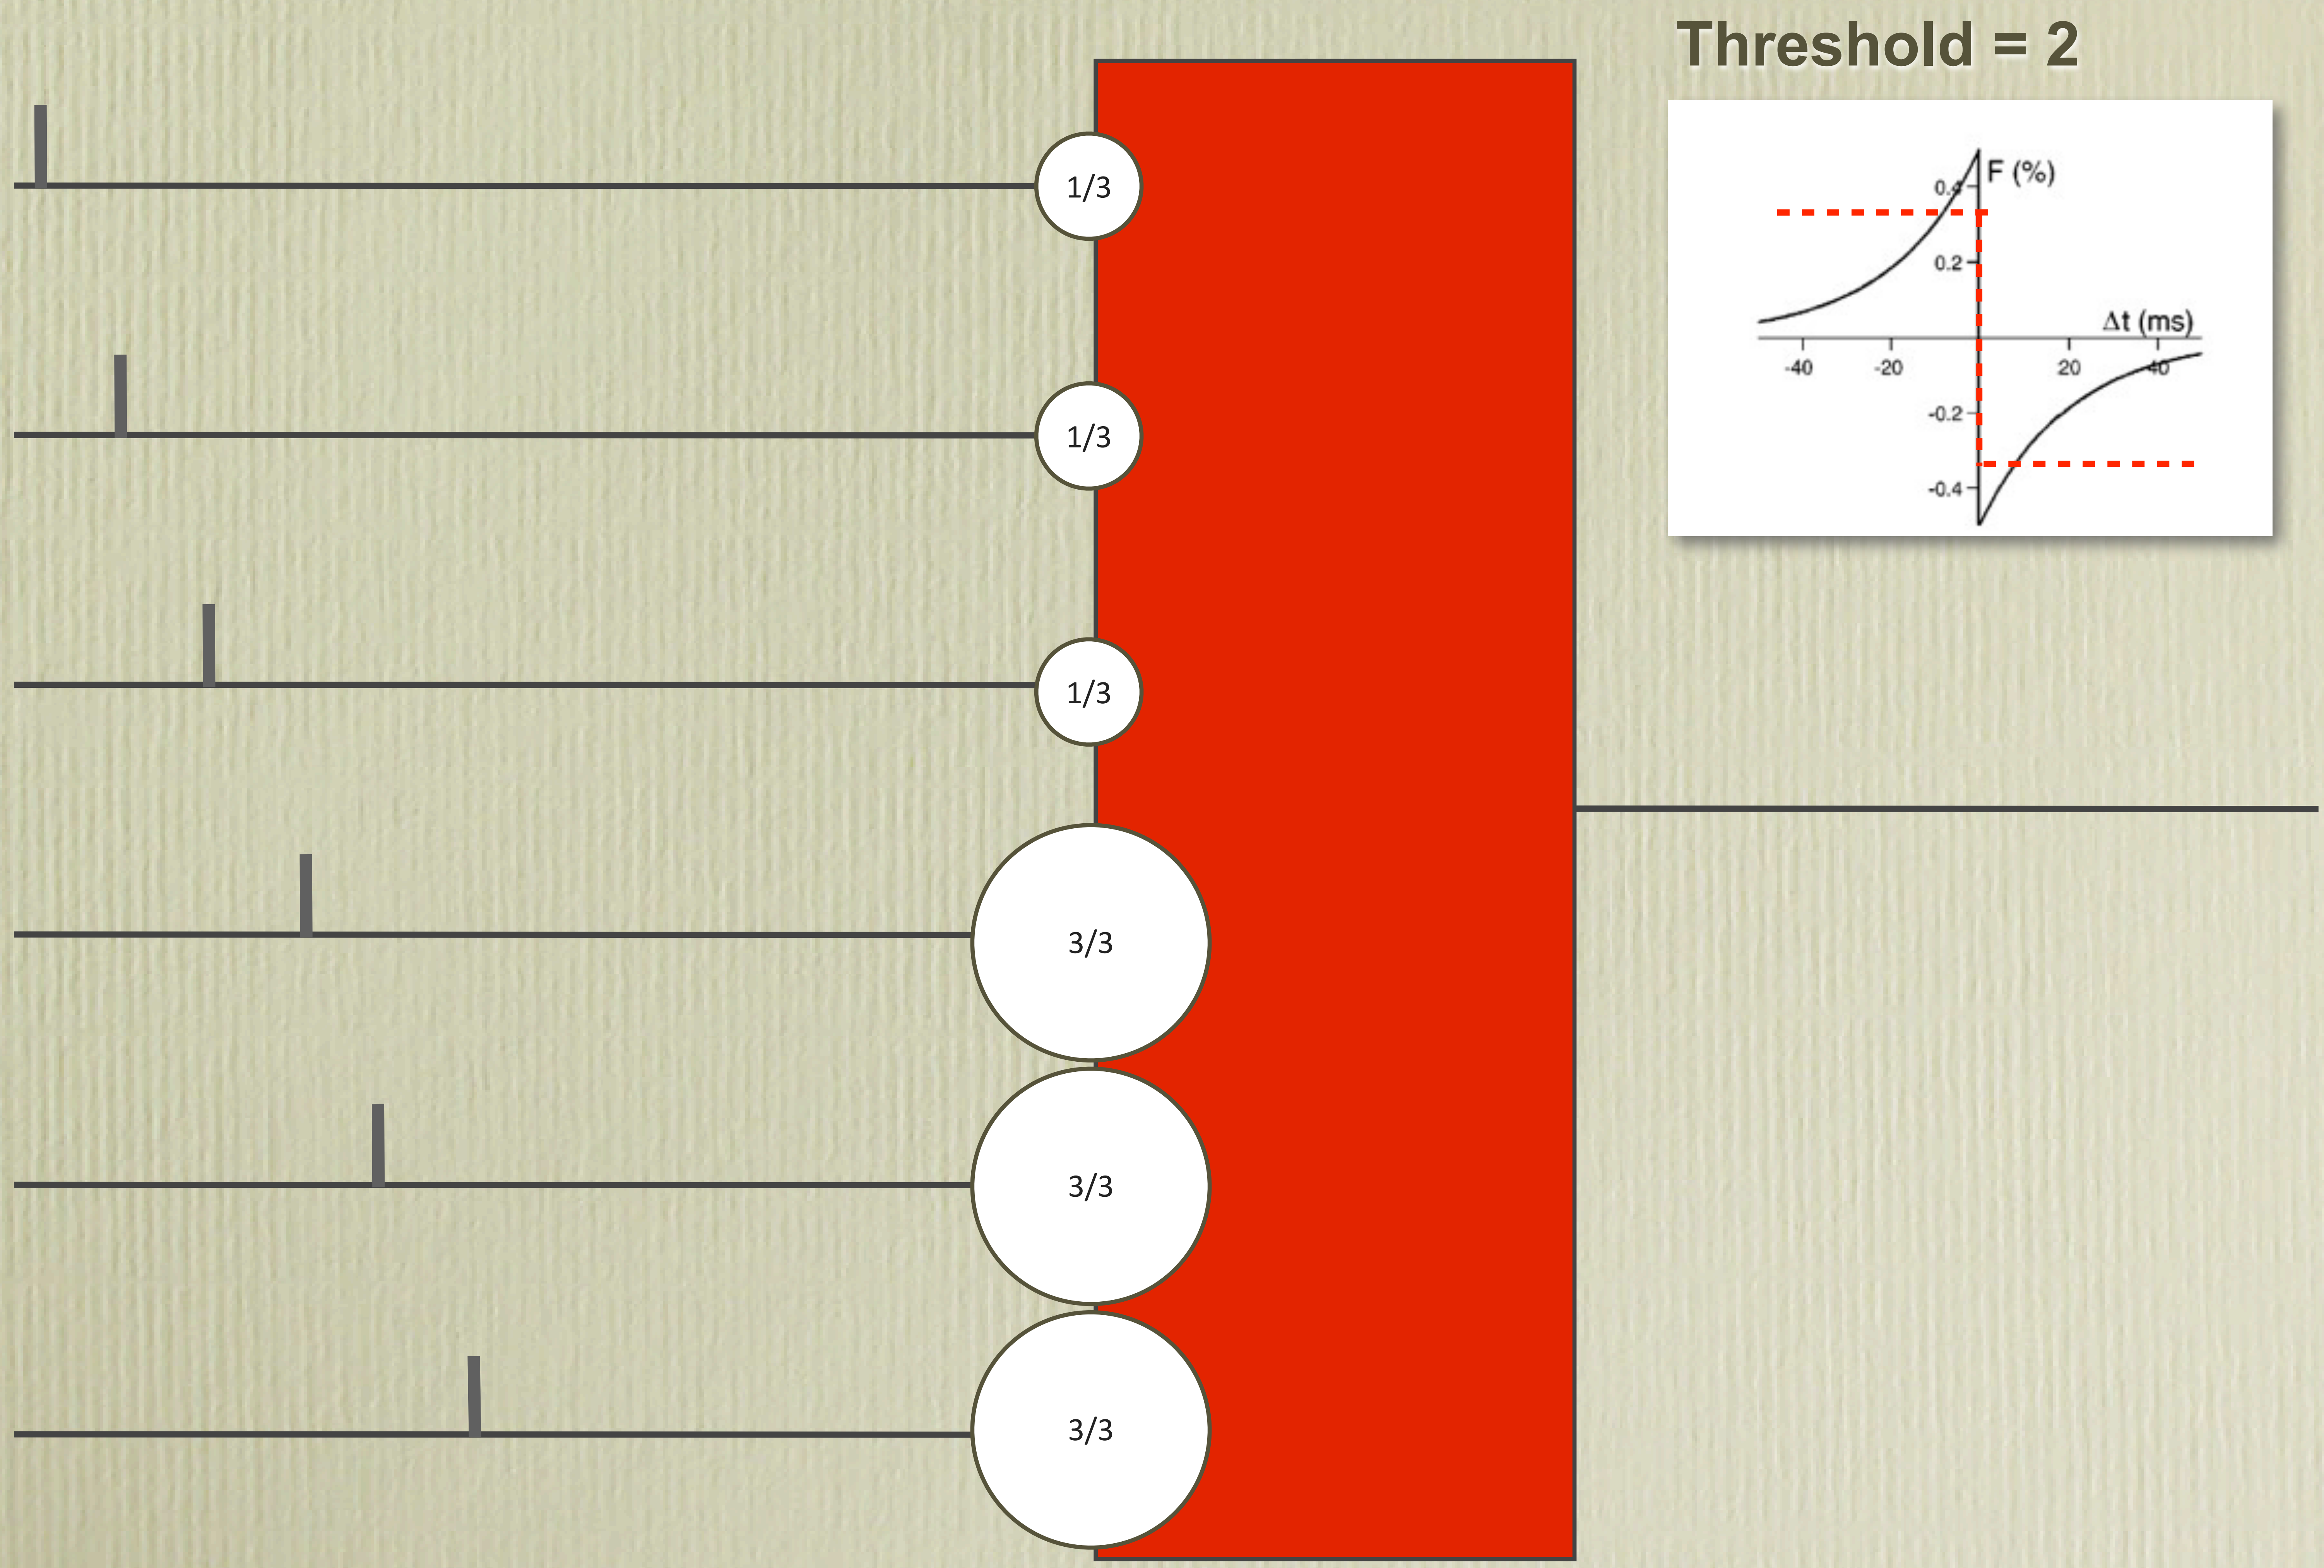

Spike Time Dependent Plasticity

Threshold = 2

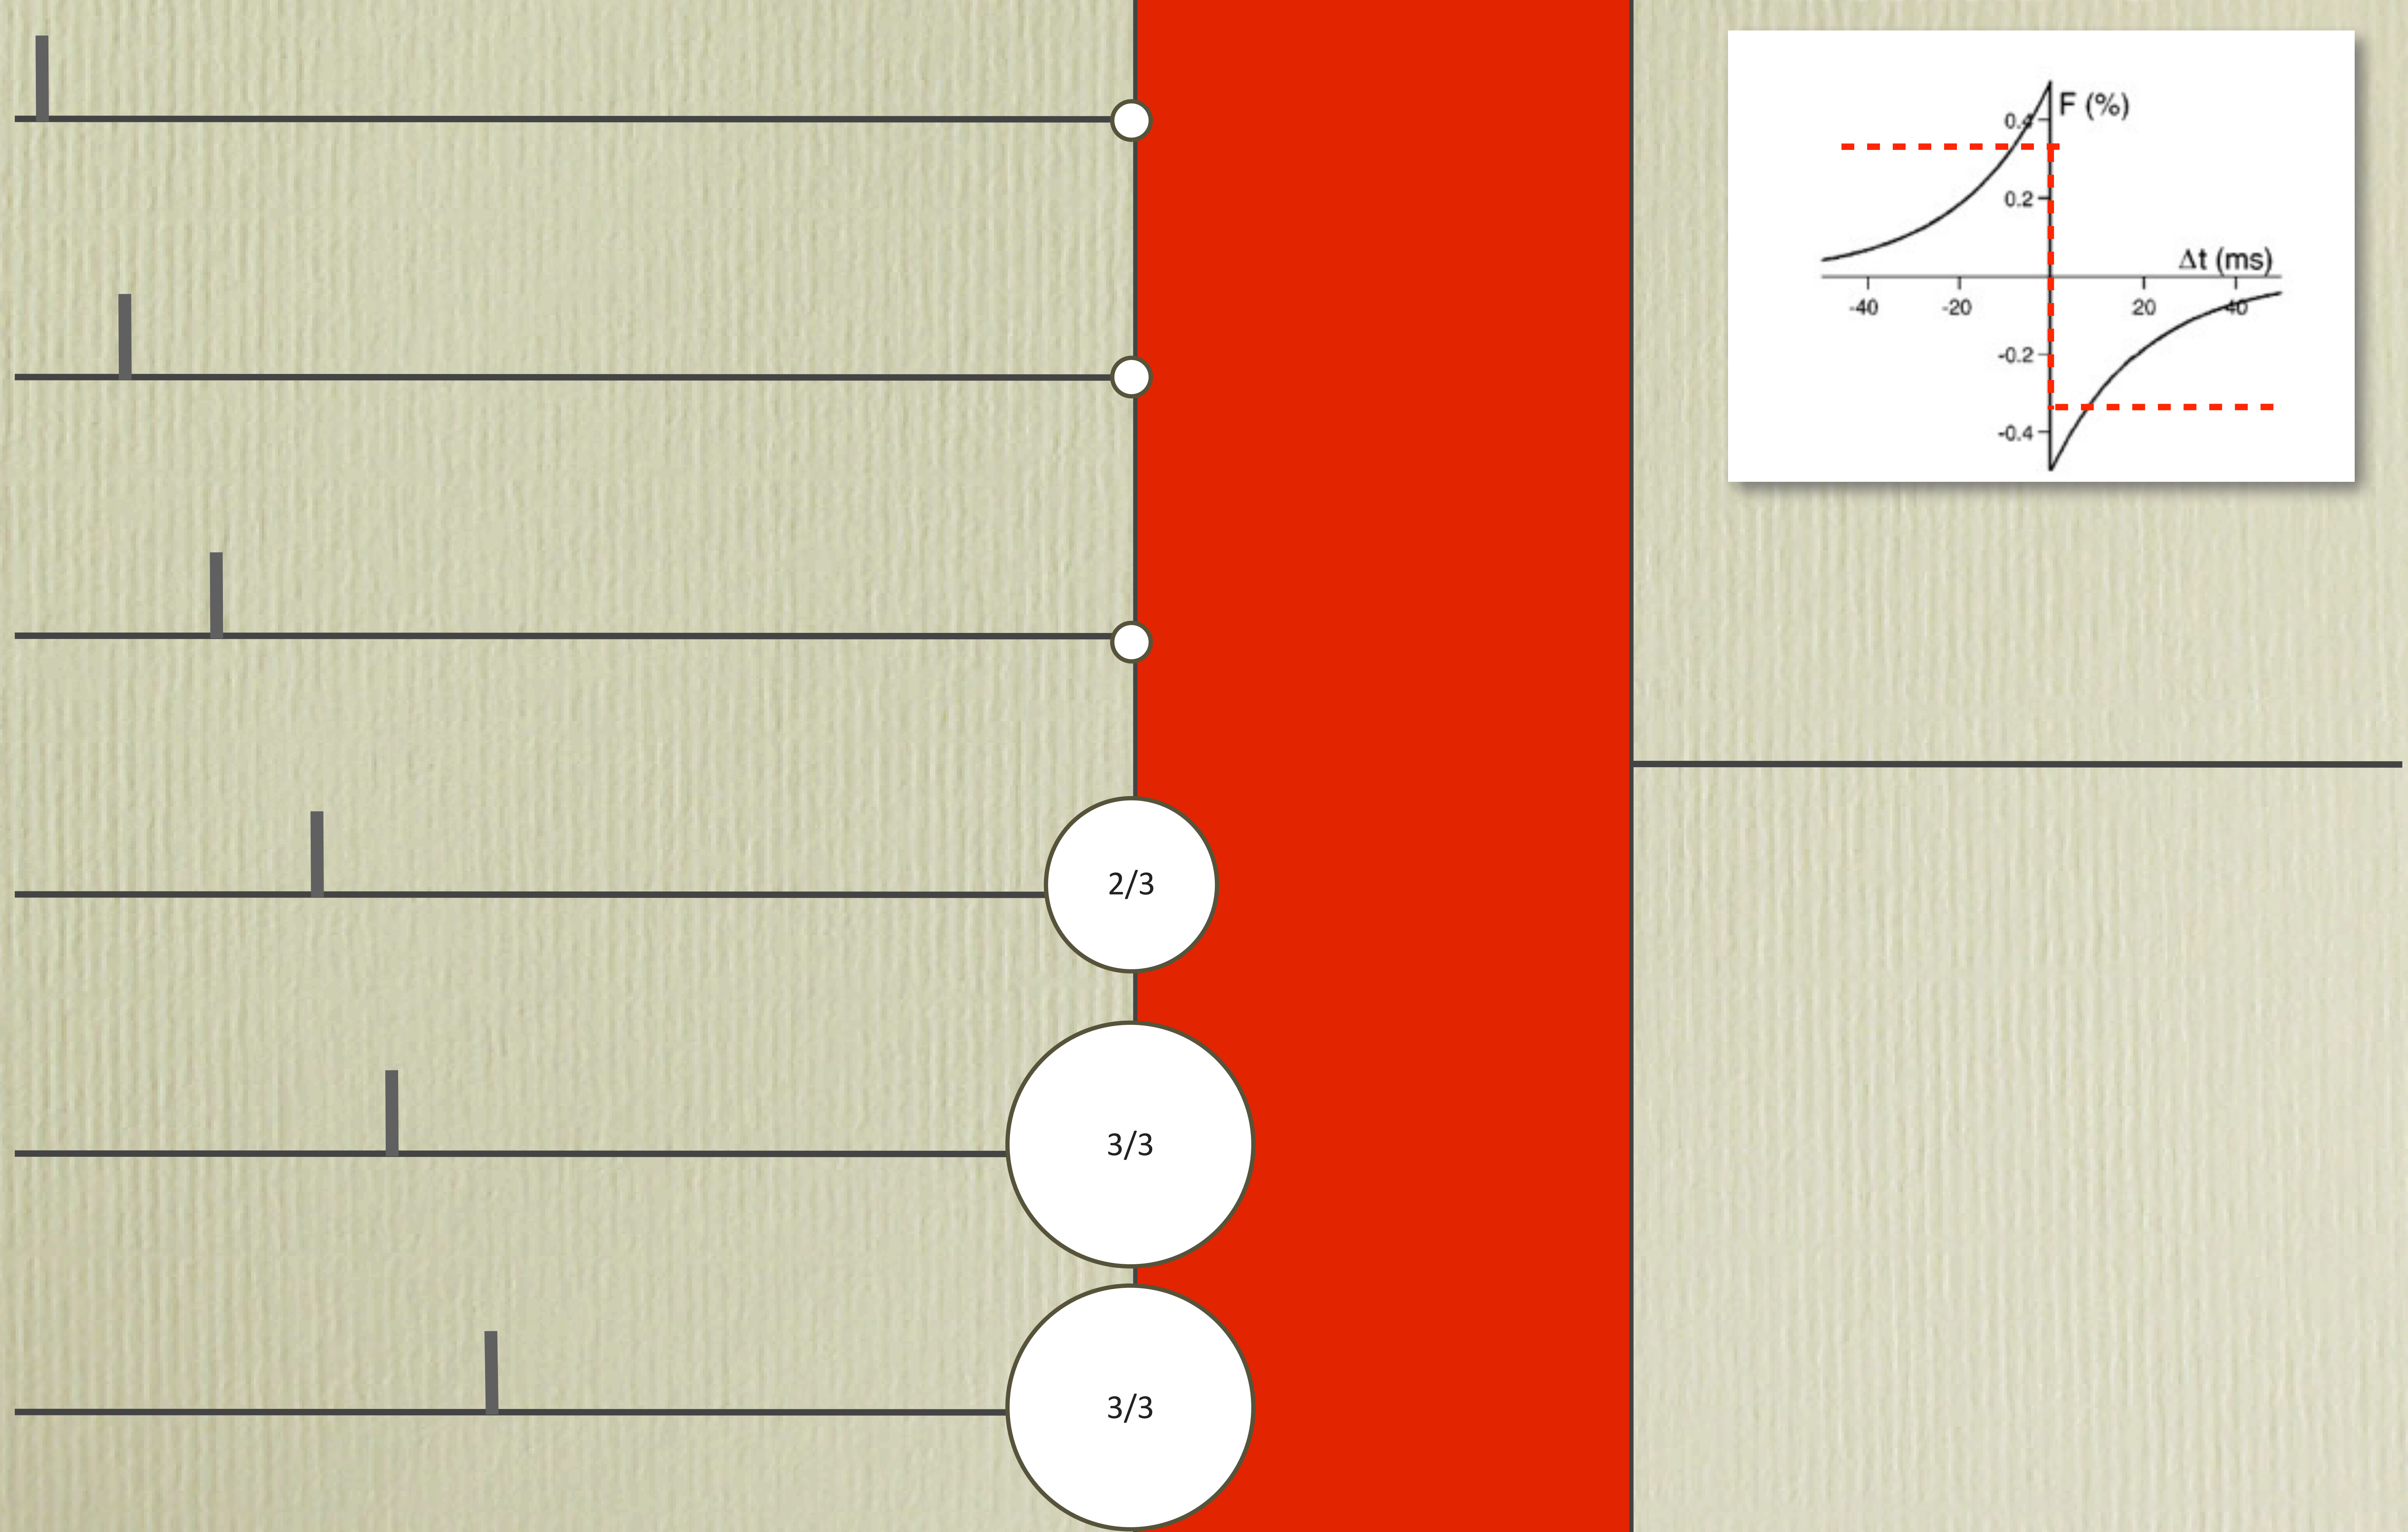

Spike Time Dependent Plasticity

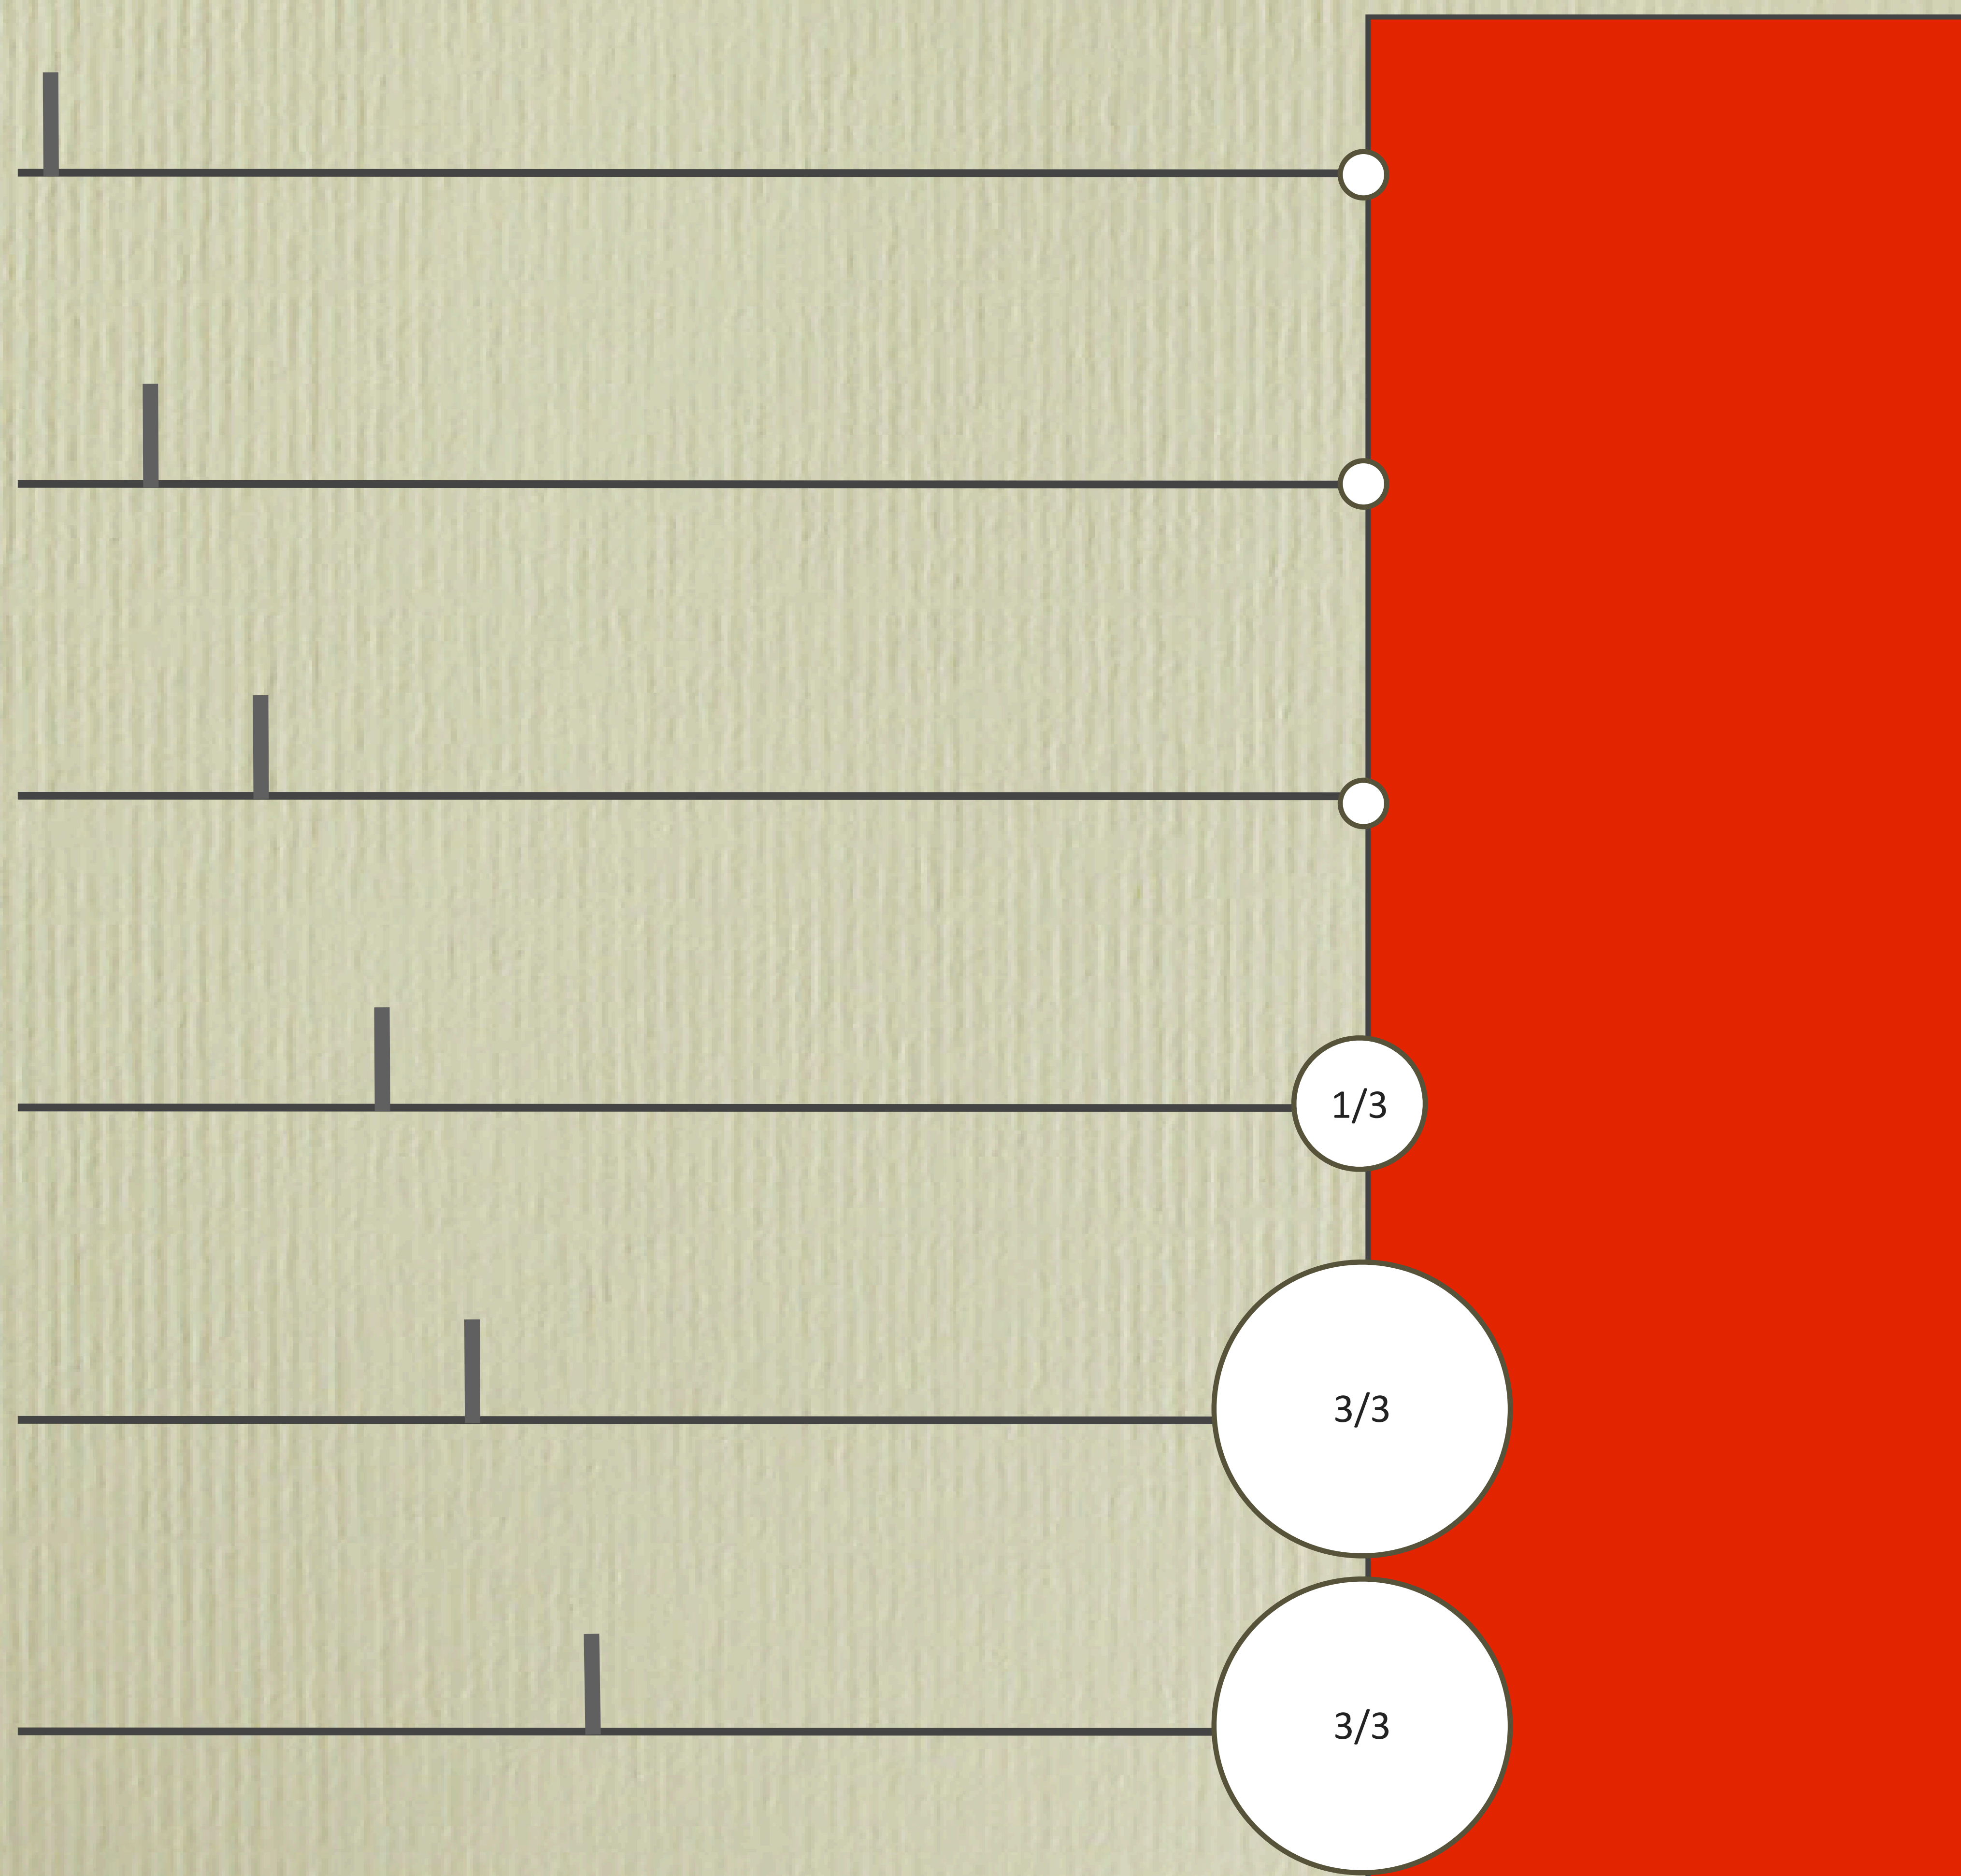

Threshold = 2

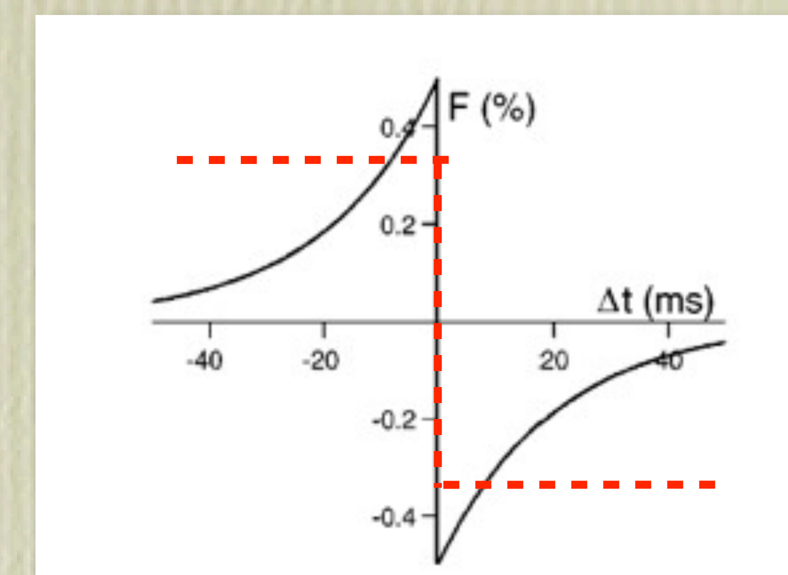

STDP concentrates high weights on the early firing inputs

Finding the earliest spikes

Neurons Tune to the Earliest Spikes Through STDP

Rudy Guyonneau Rufin VanRullen Simon J. Thorpe

Neural Computation **17, 859–879** (2005)

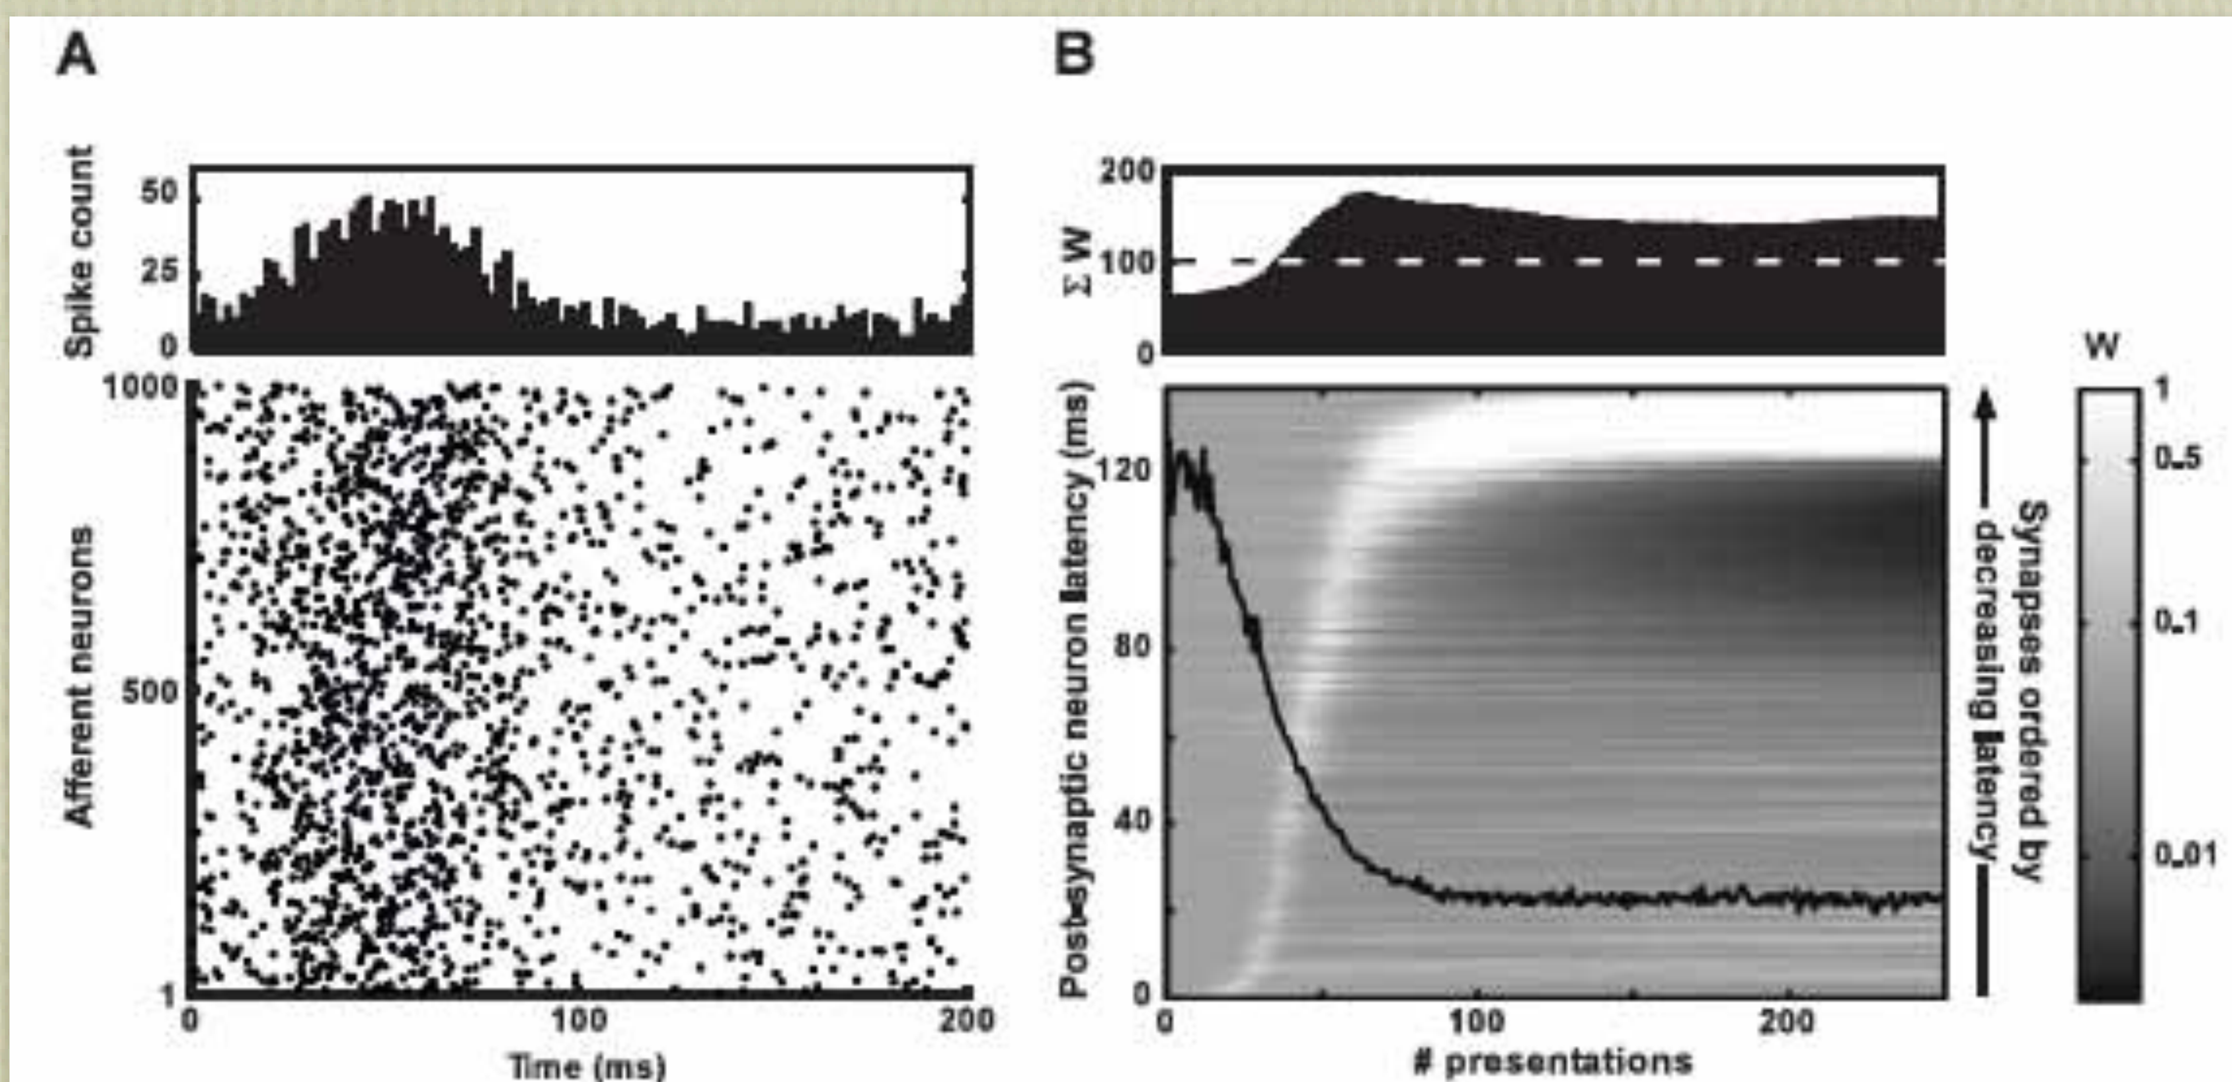

With a few tens of presentations, high weights concentrate on the earliest firing inputs

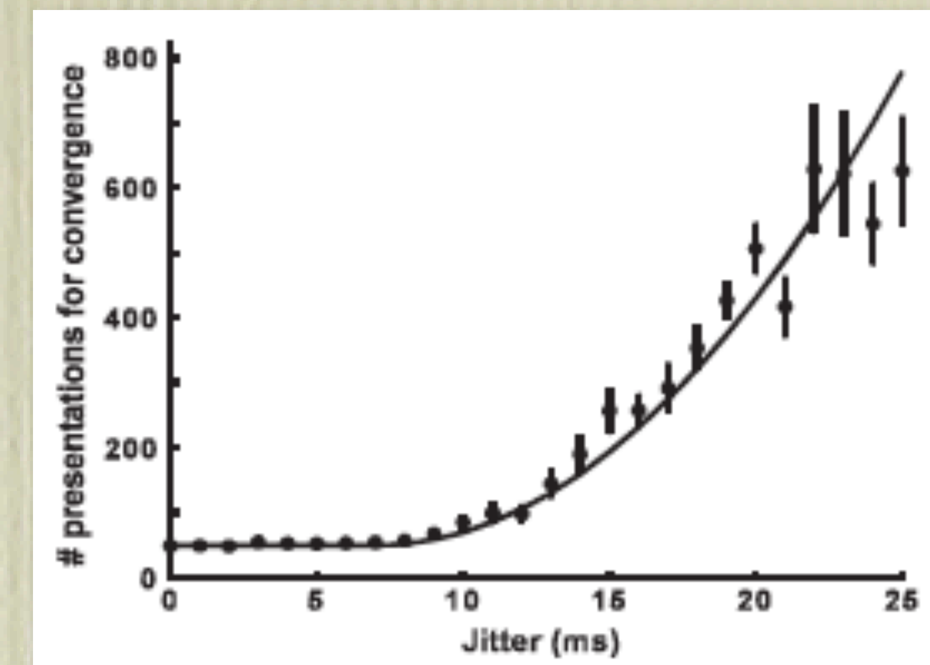

Even with jitter

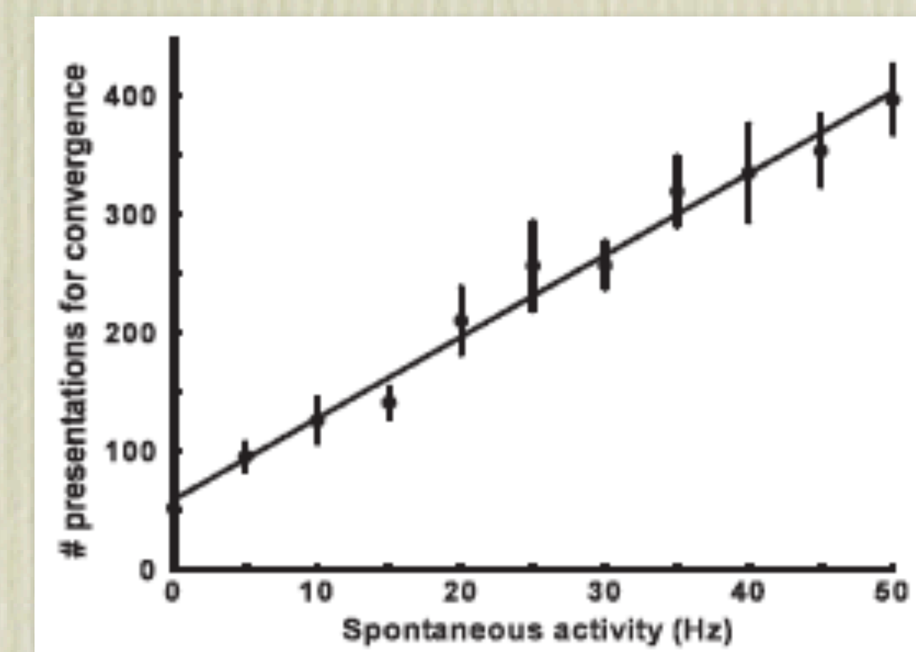

Even with spontaneous background activity

Learning Spike Sequences with STDP

OPEN ACCESS Freely available online PLoS one

Spike Timing Dependent Plasticity Finds the Start of Repeating Patterns in Continuous Spike Trains

Timothée Masquelier[1,2*], Rudy Guyonneau[1,2], Simon J. Thorpe[1,2]

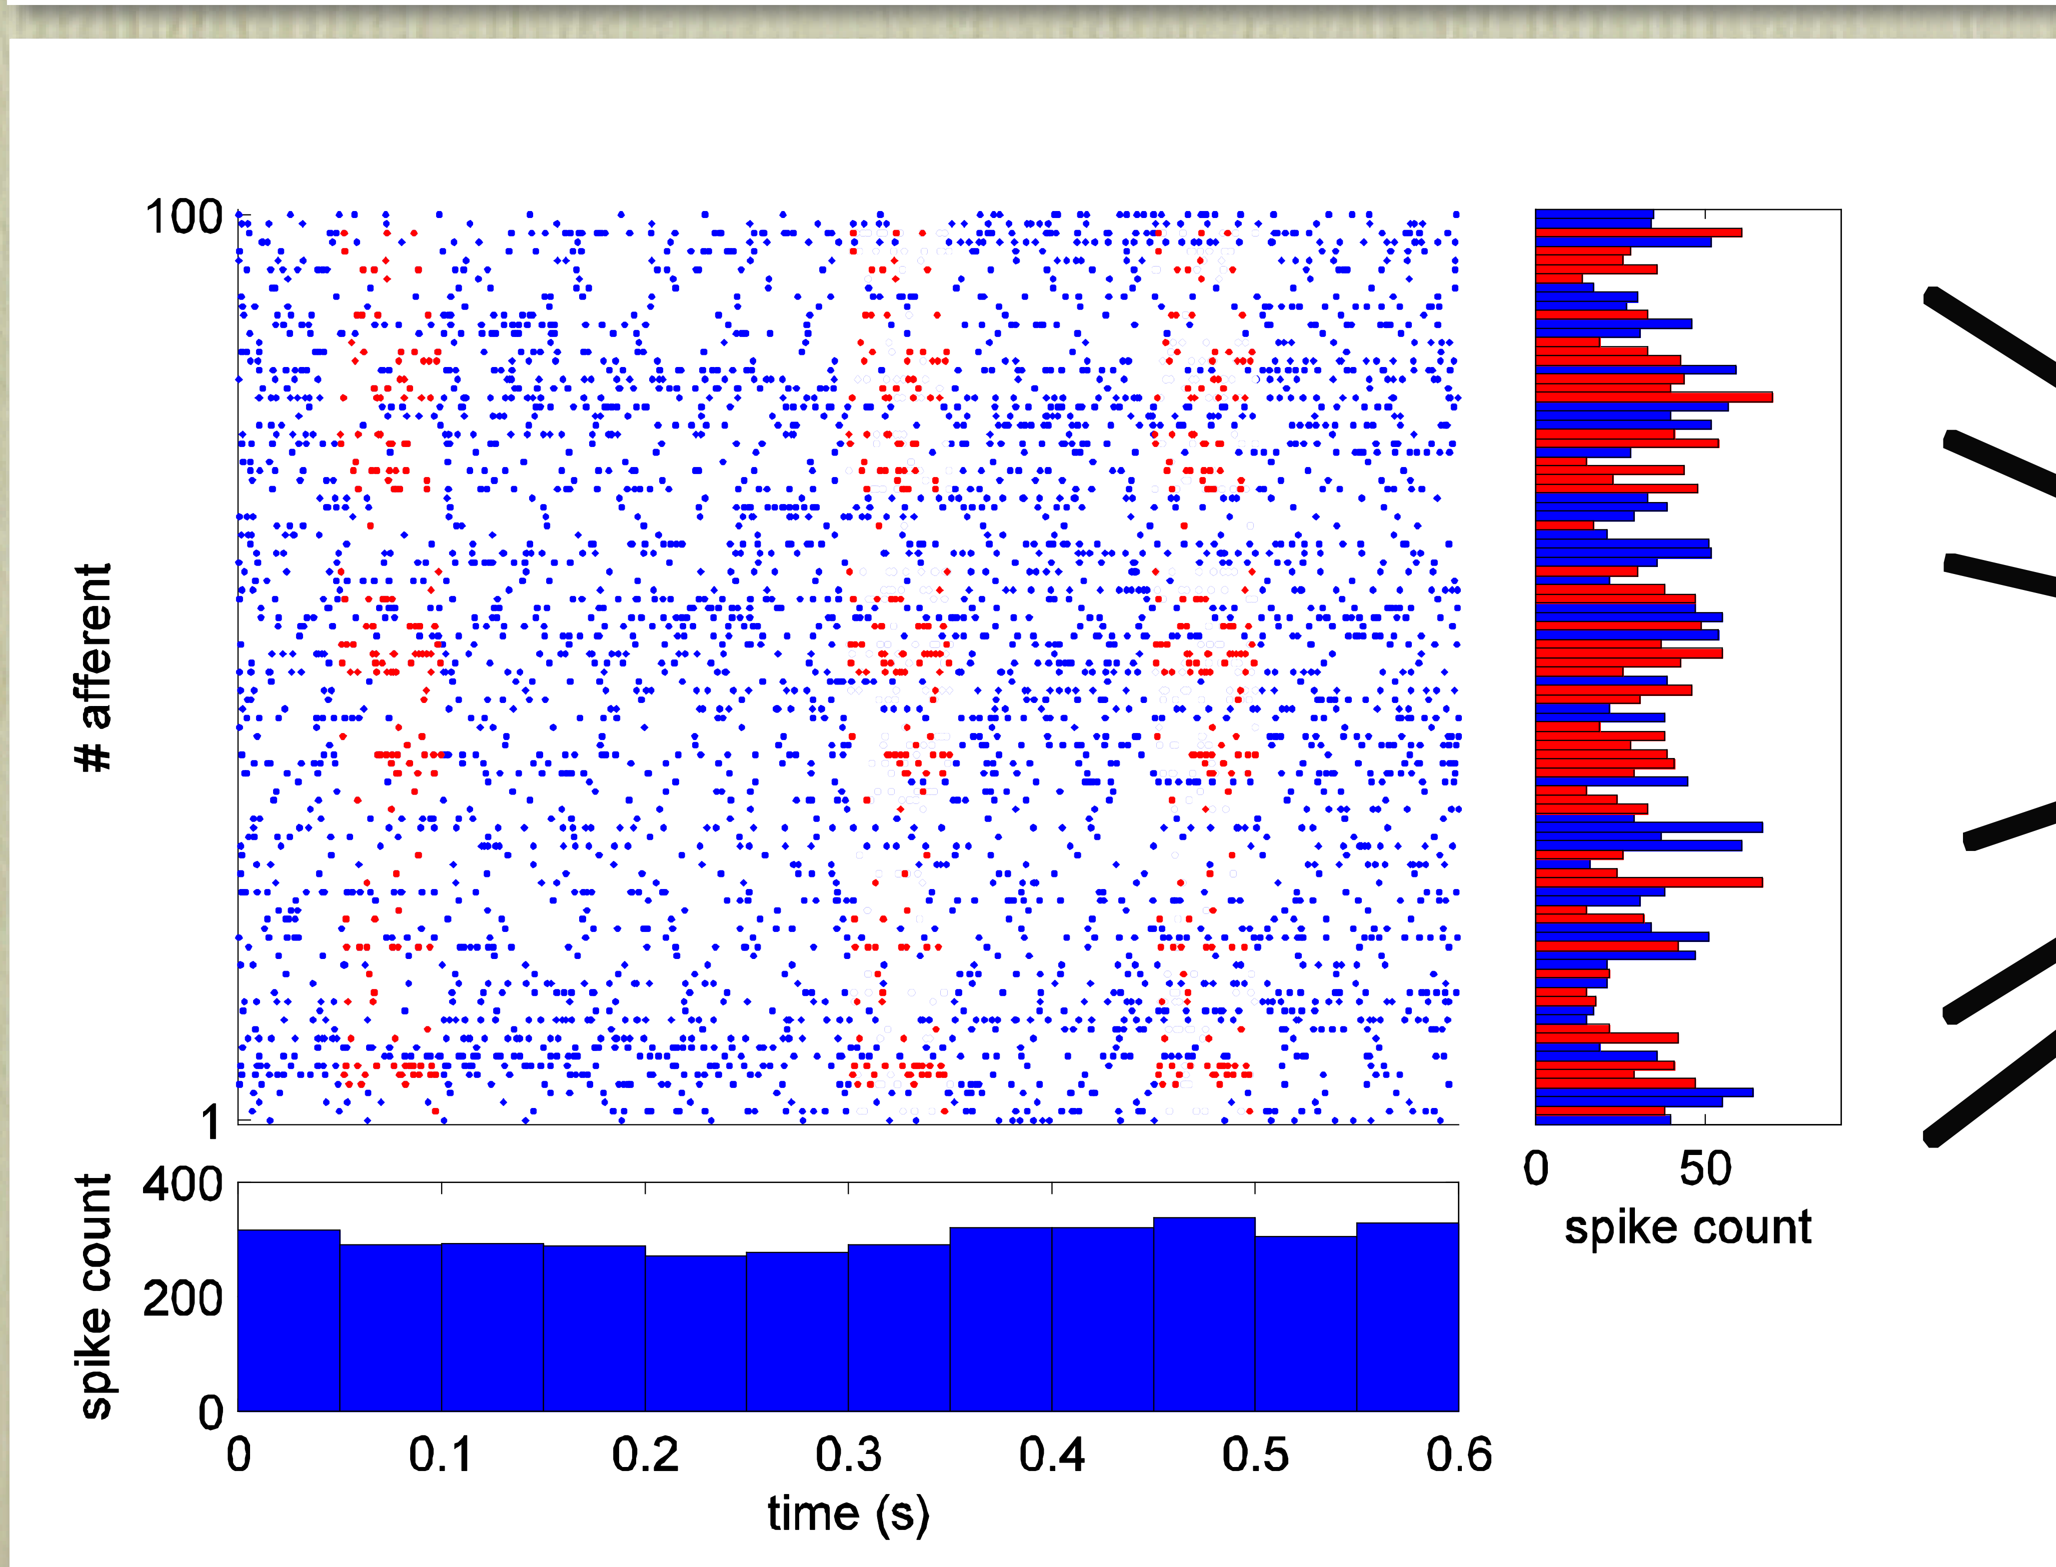

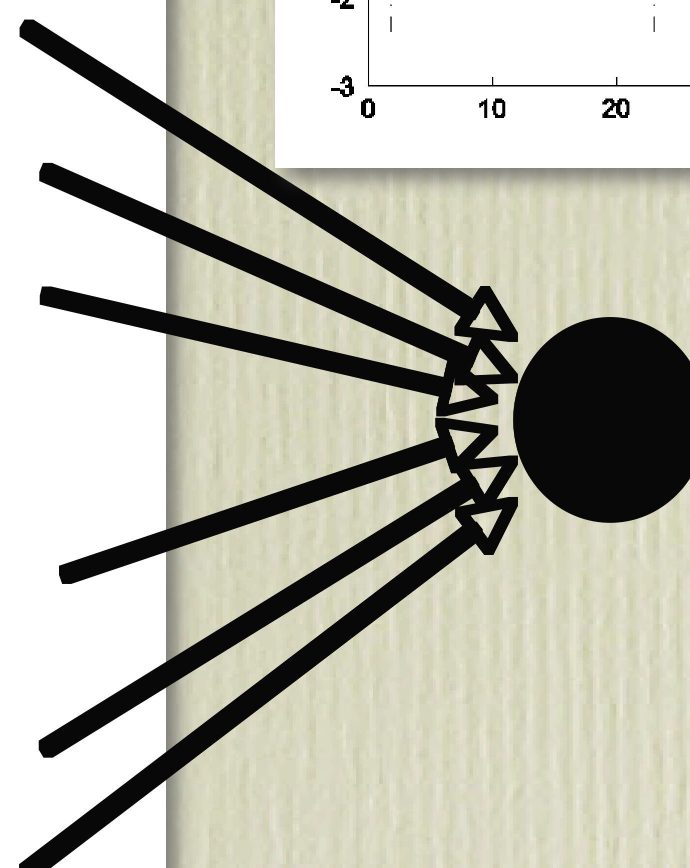

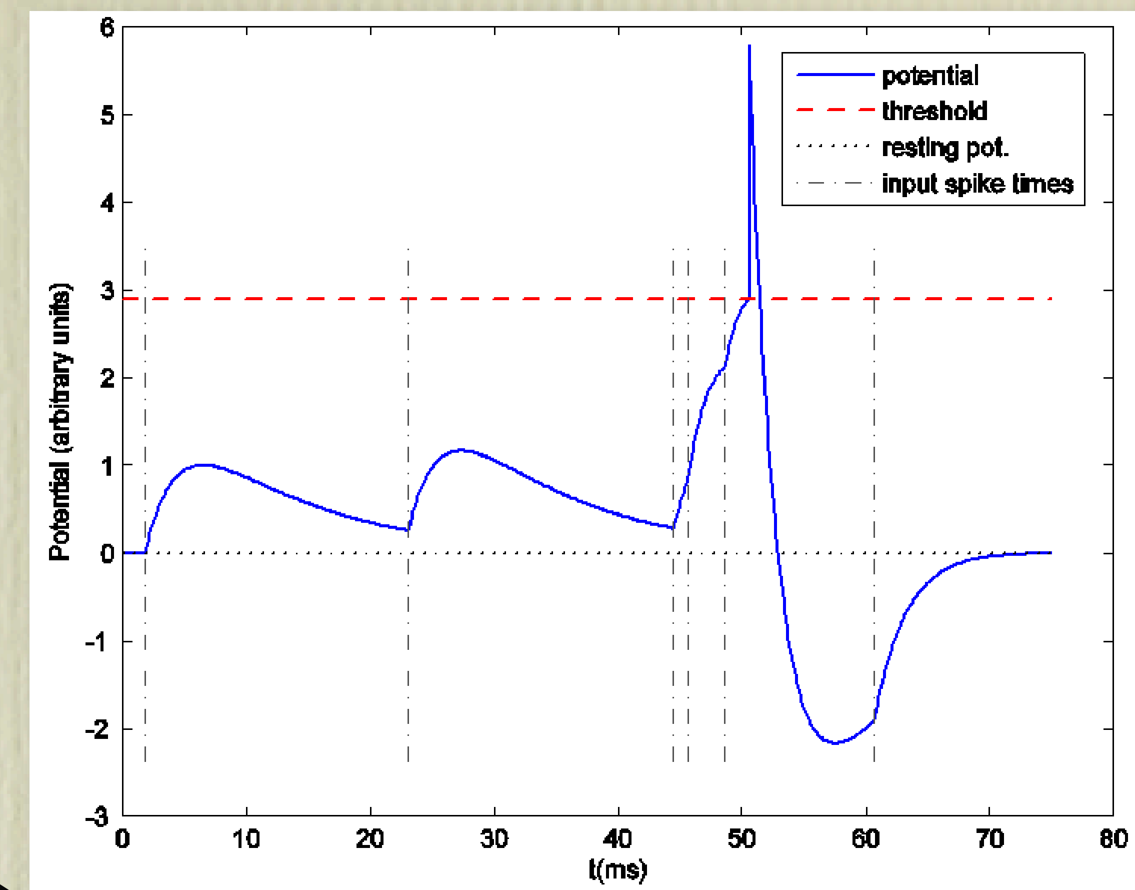

Learning Spike Sequences with STDP

- Initial State

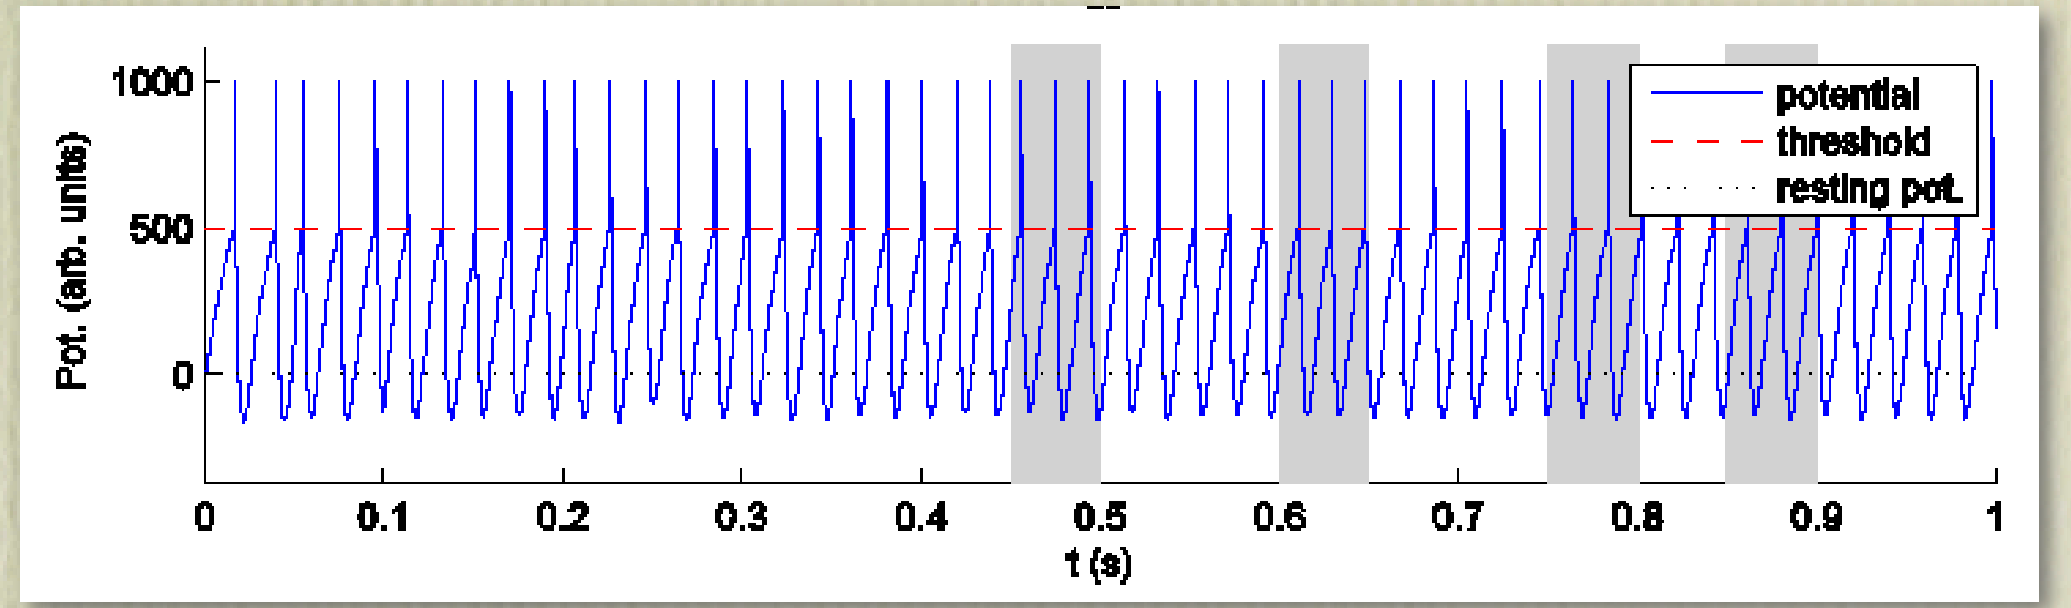

- During Learning

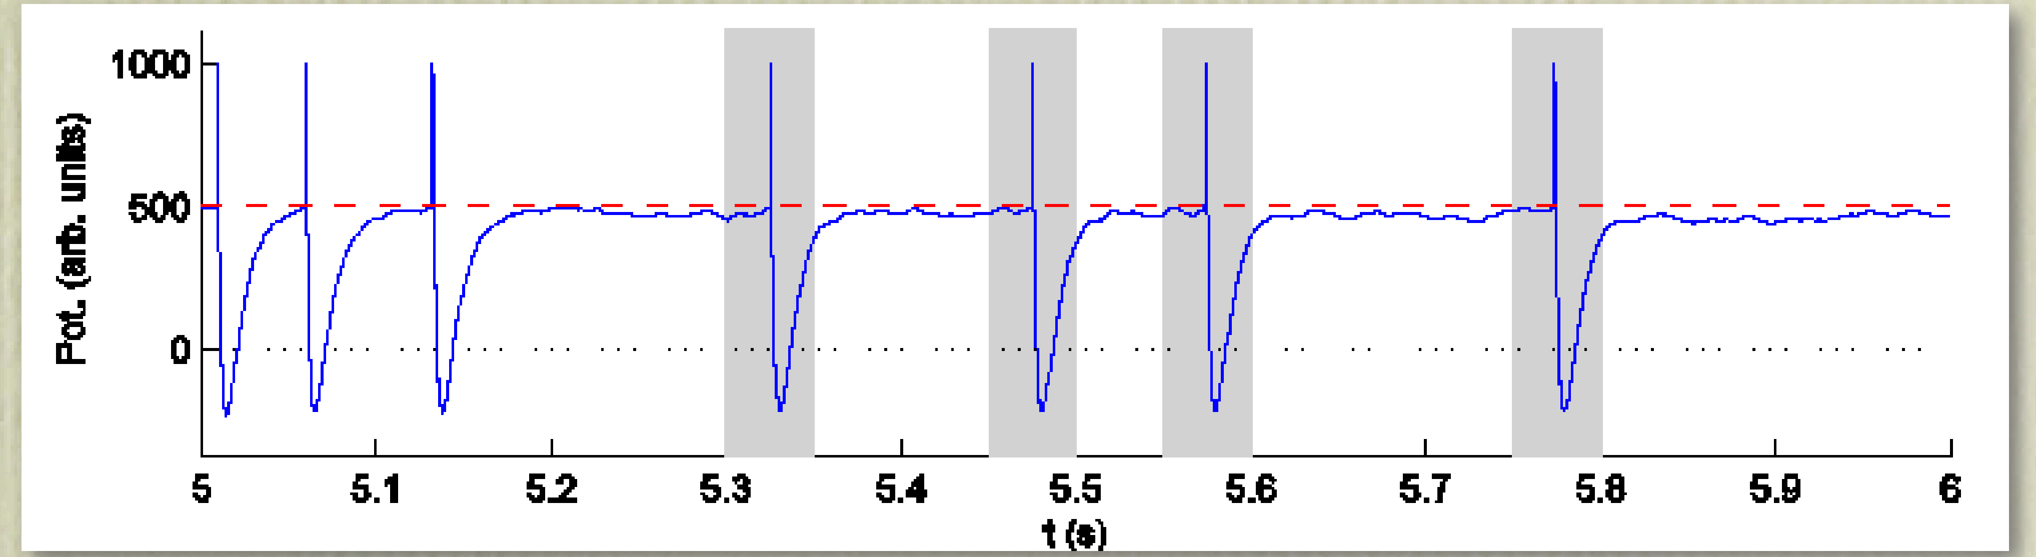

- After Learning

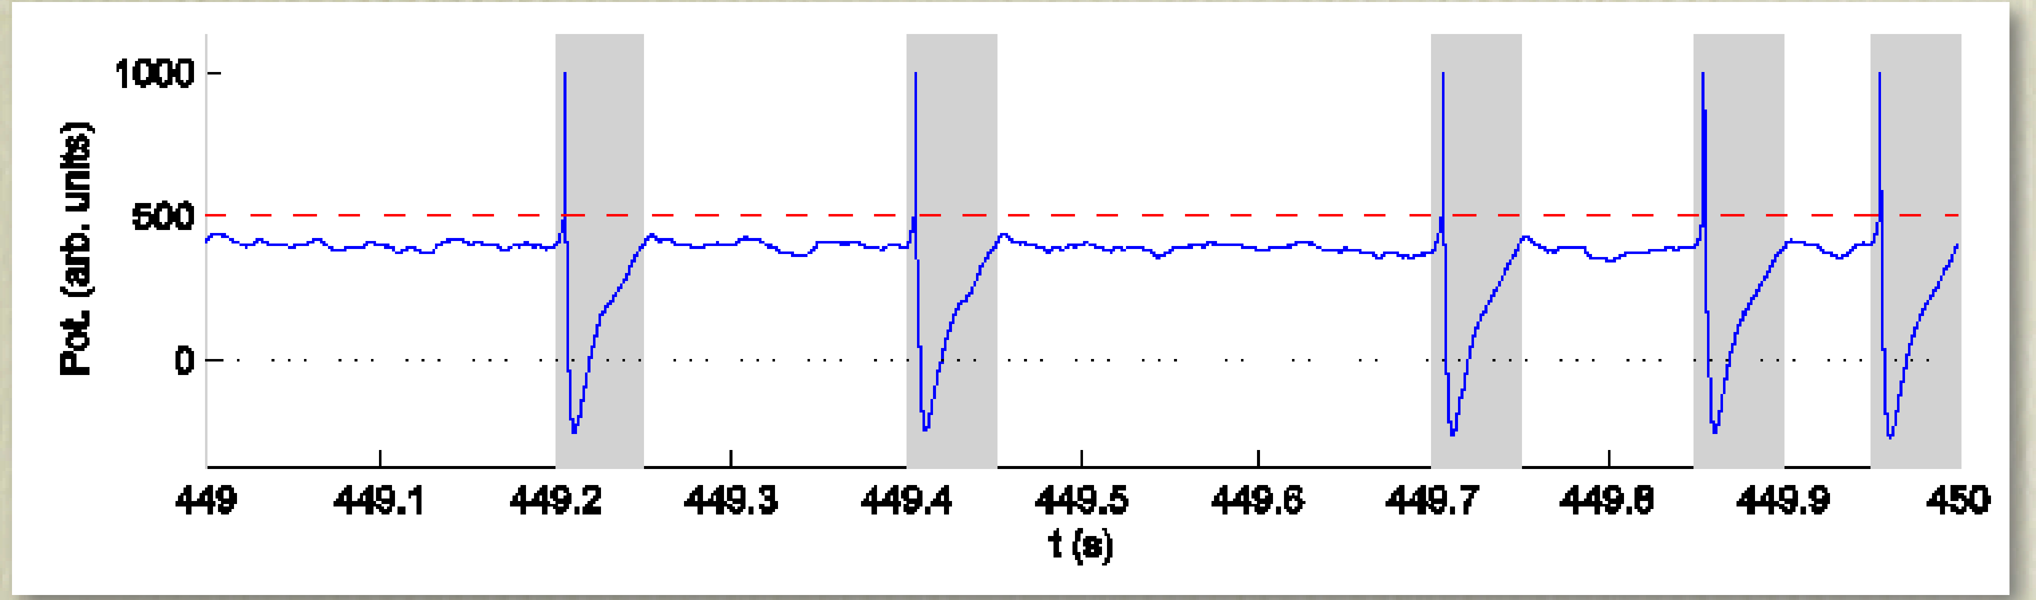

Competitive Networks

Competitive STDP-Based Spike Pattern Learning

Timothée Masquelier Rudy Guyonneau Simon J. Thorpe

Neural Computation **21**, 1259–1276 (2009)

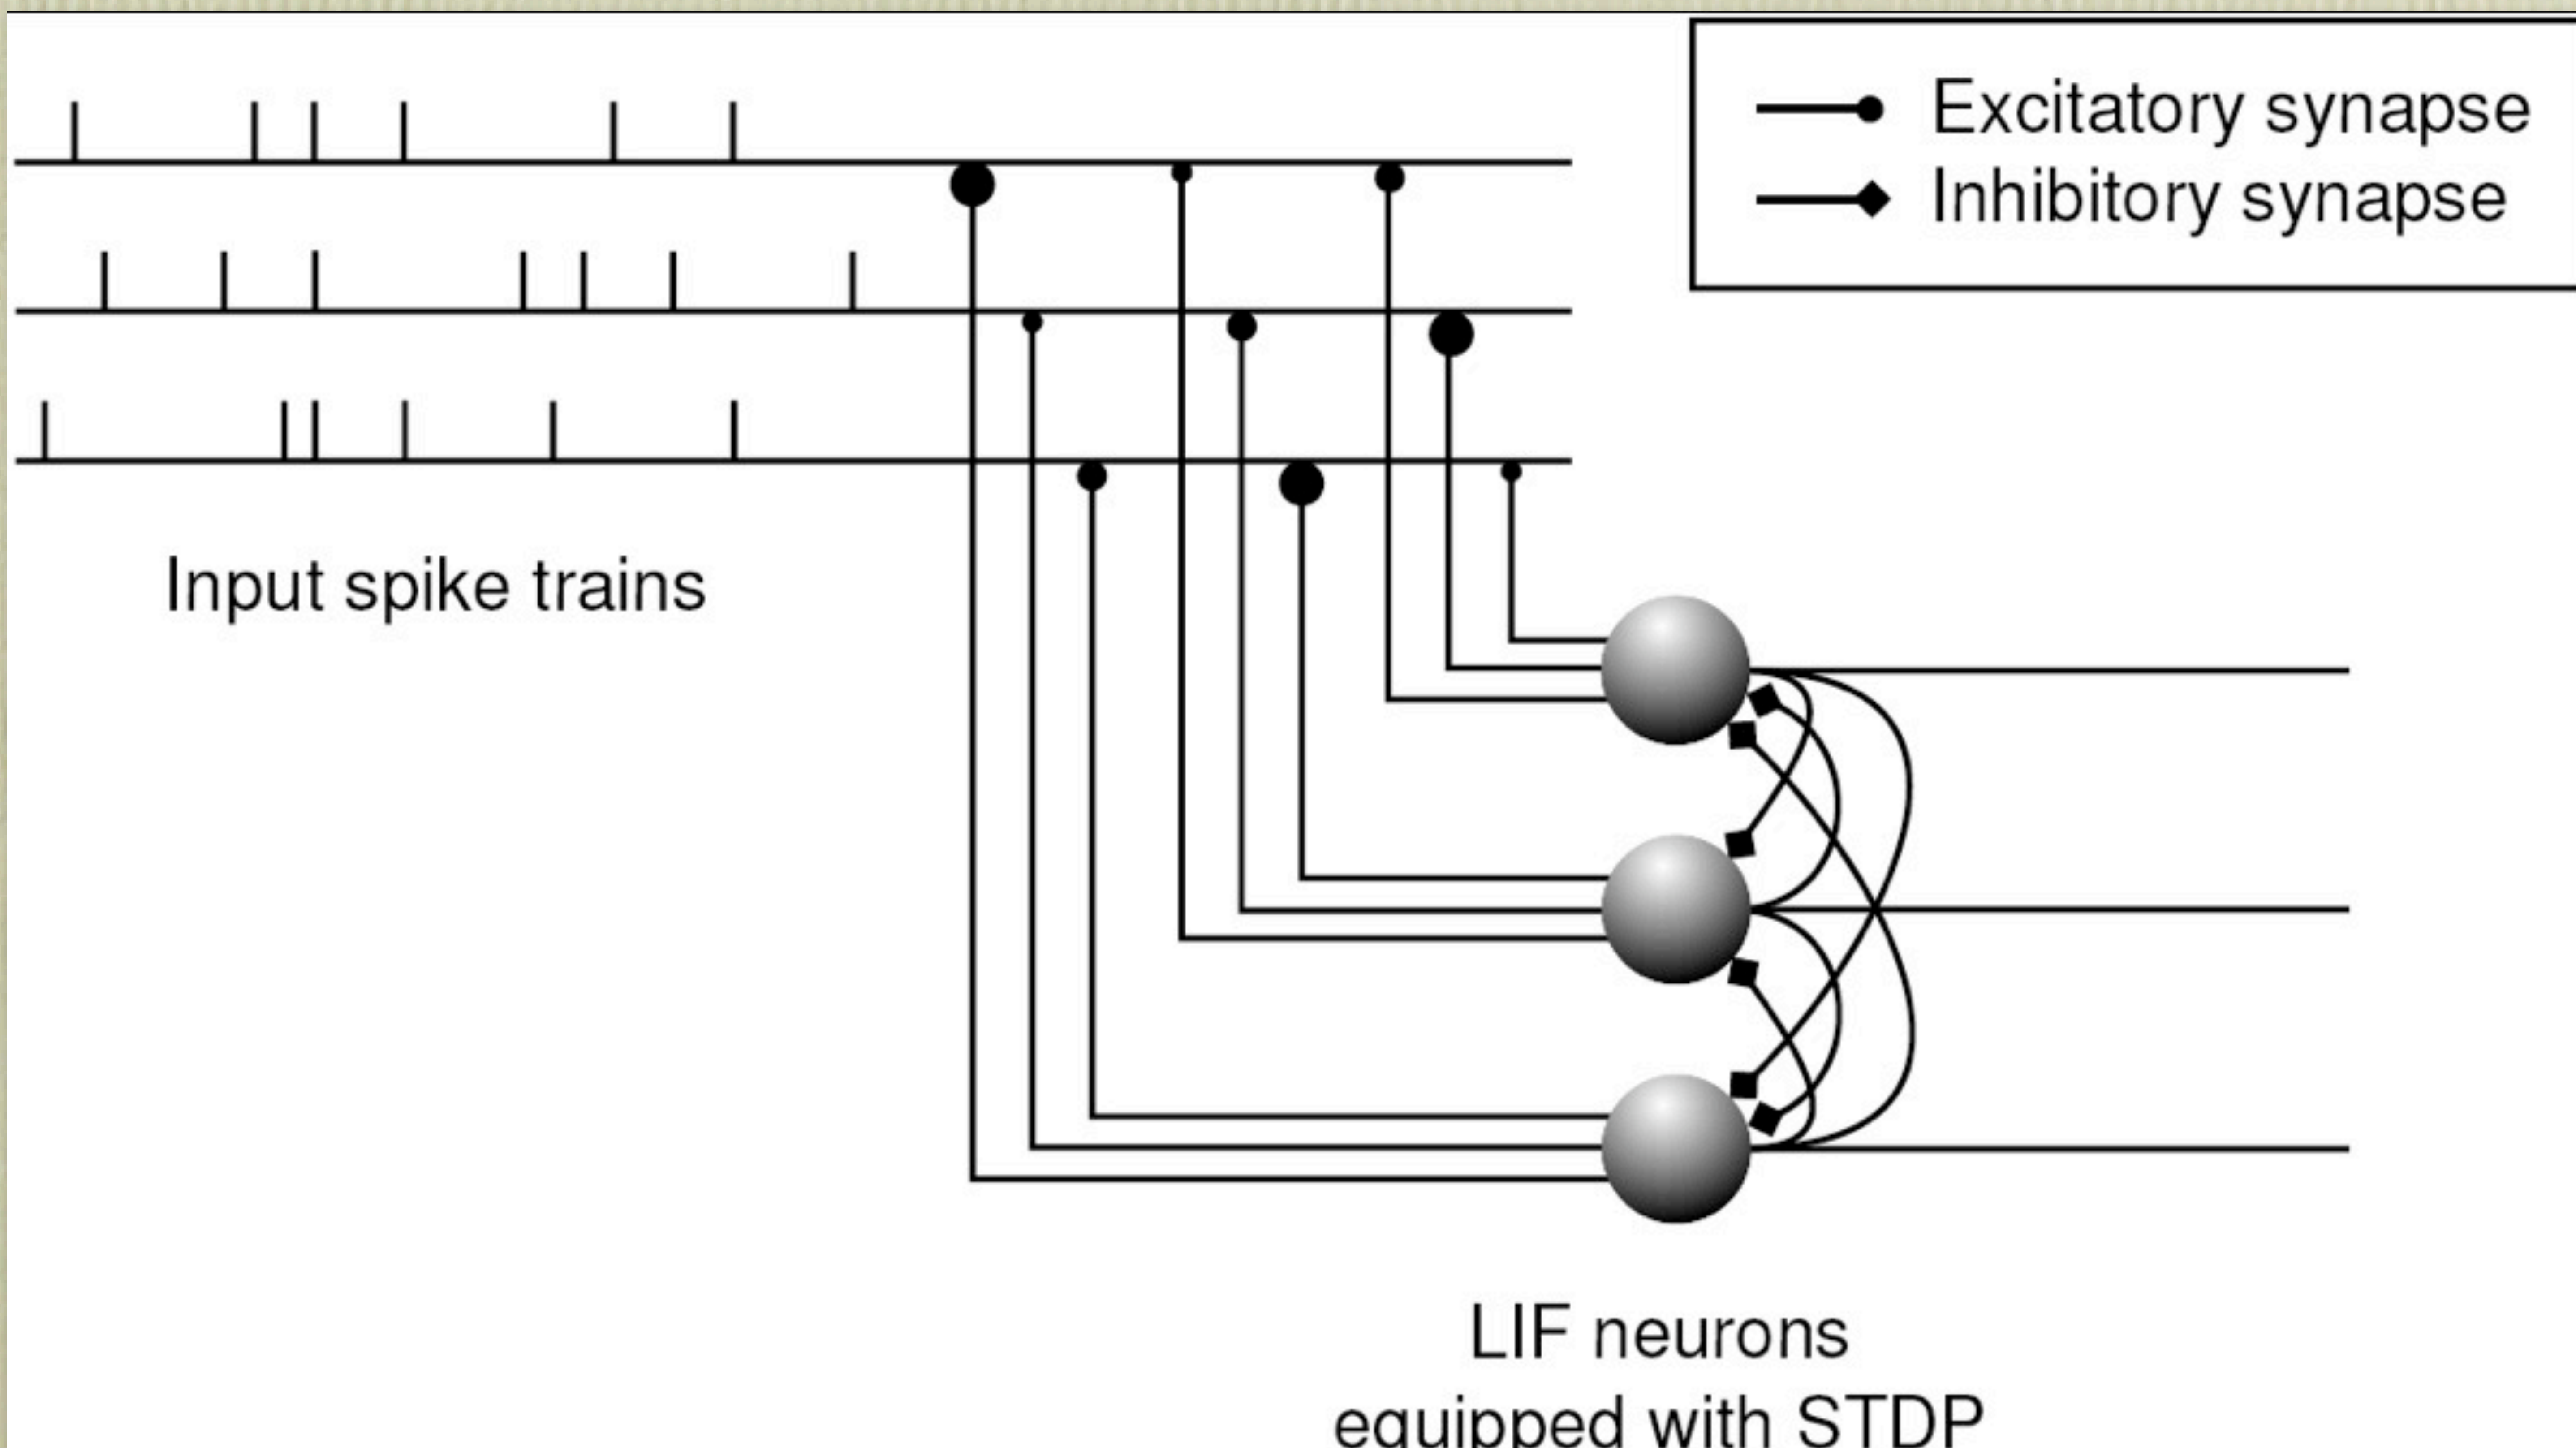

Learning with multiple neurons

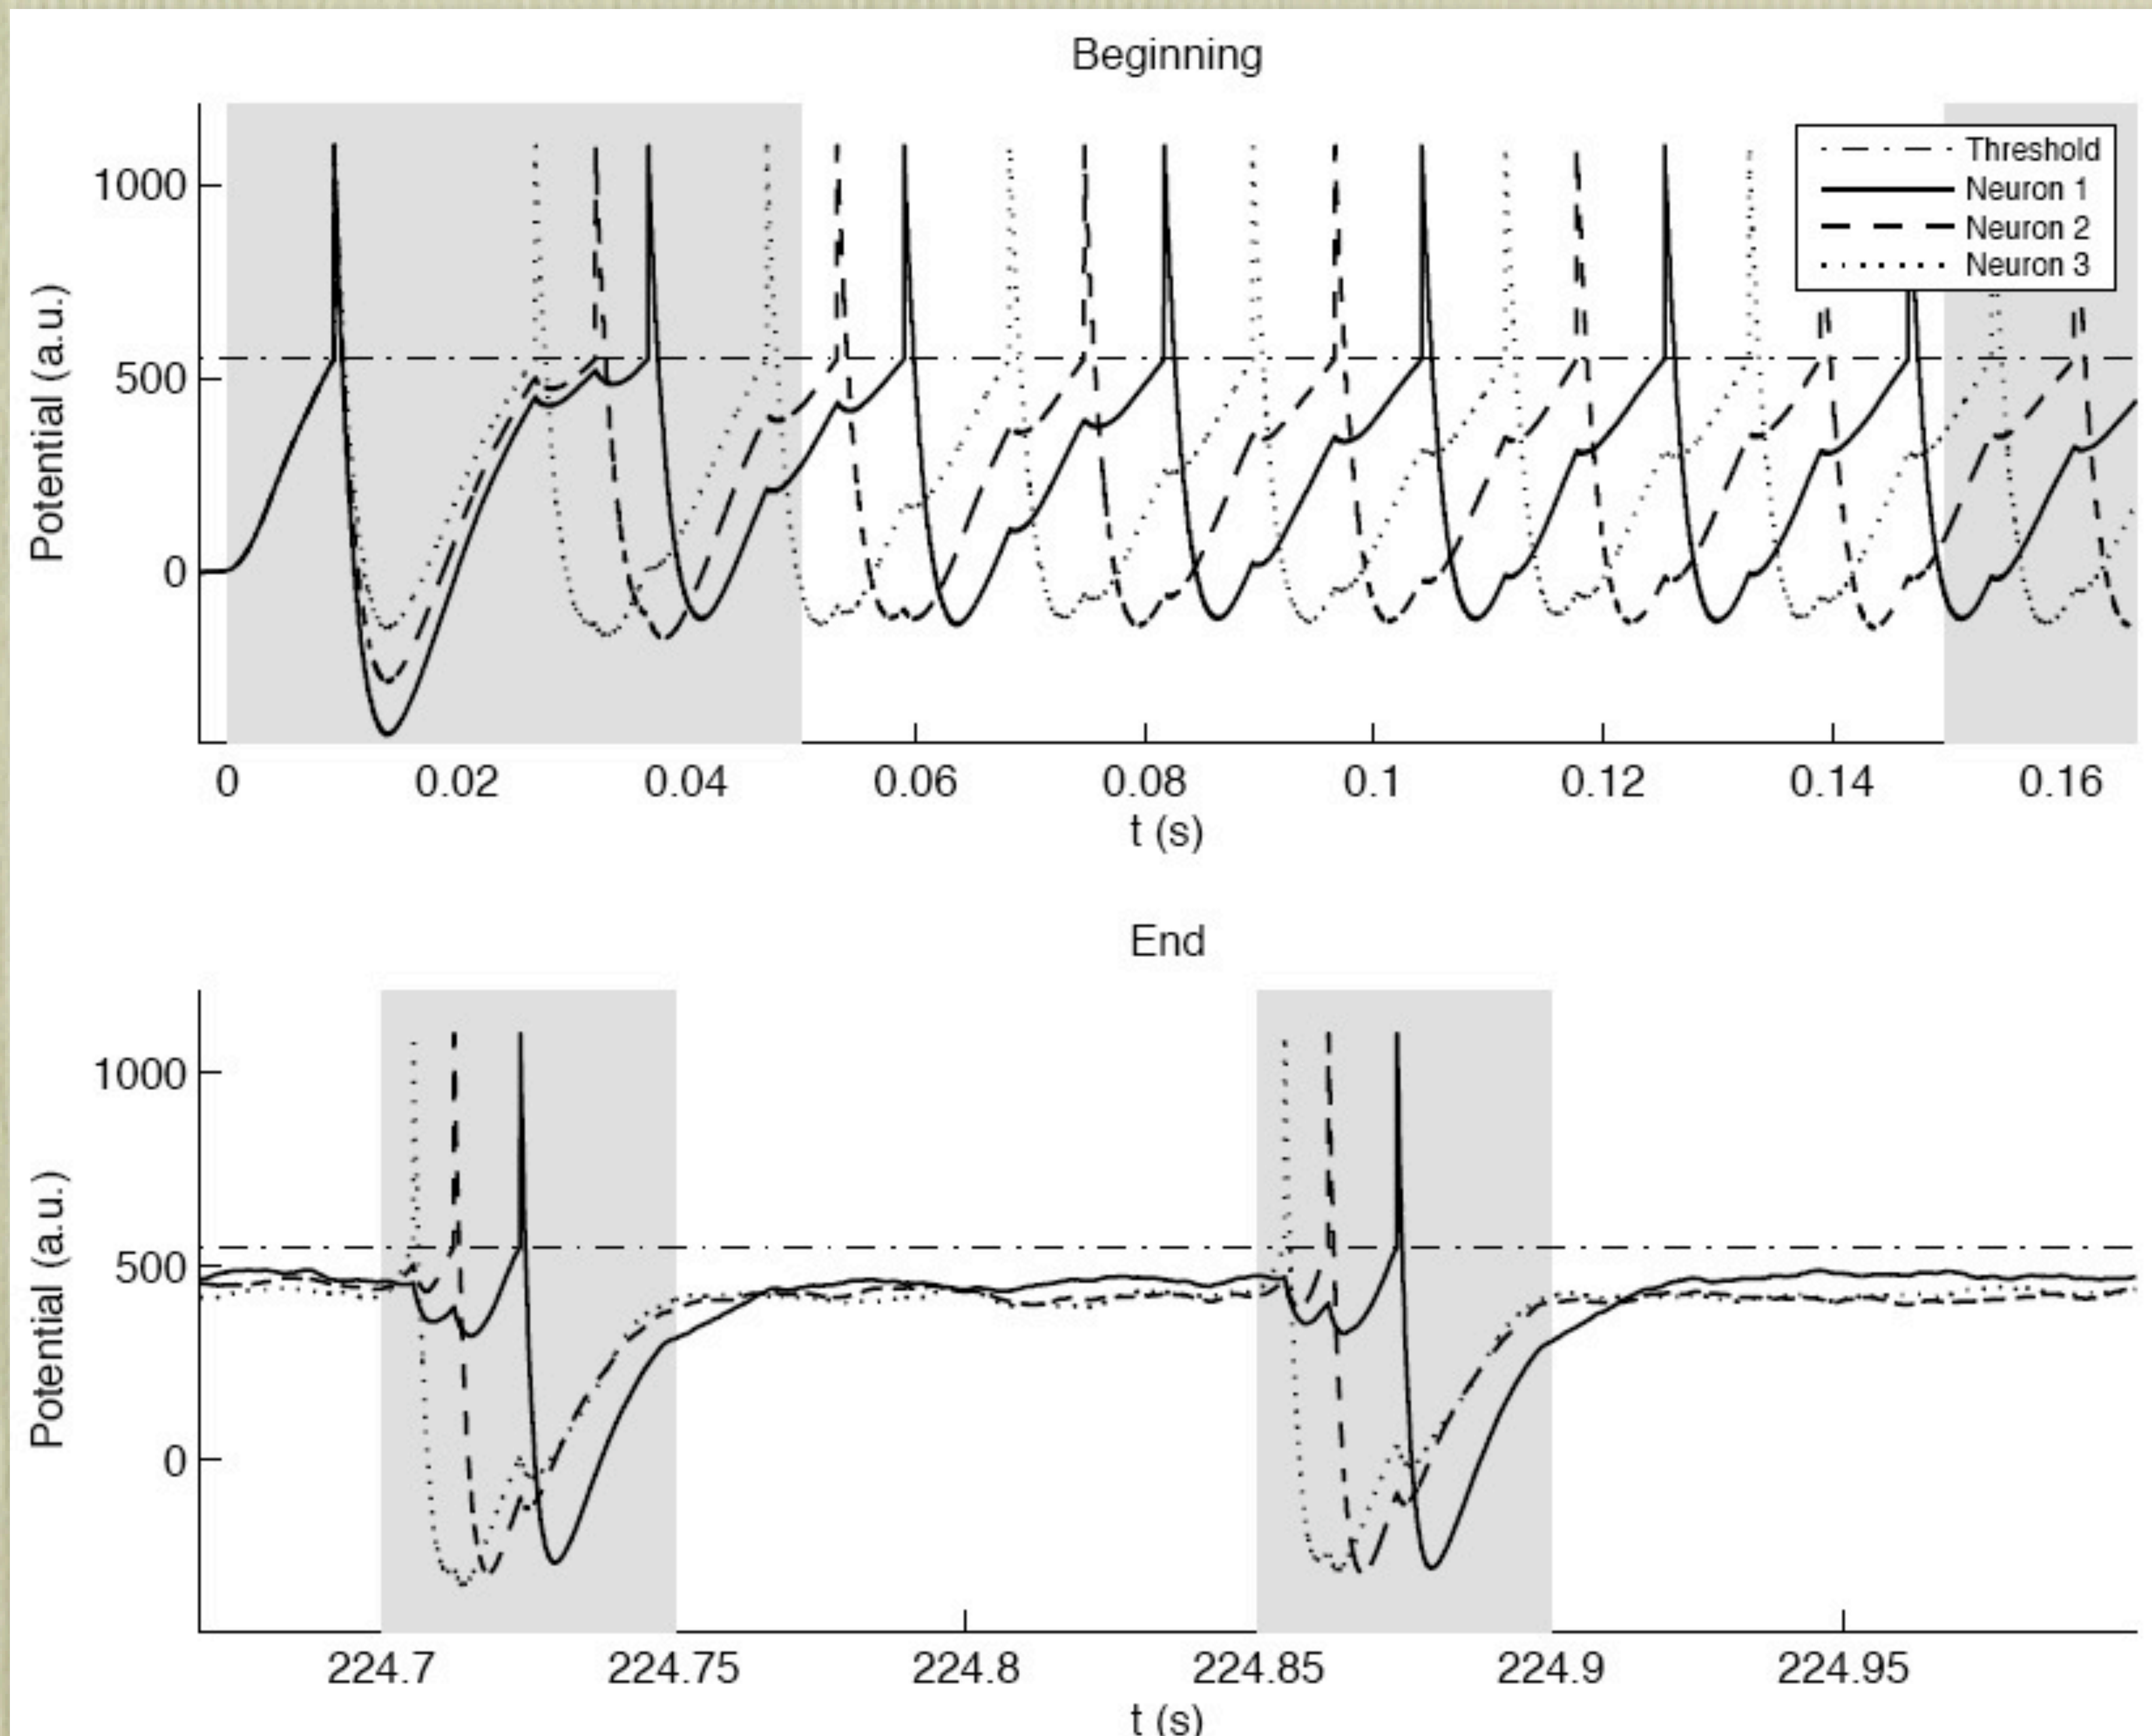

Multiple Patterns

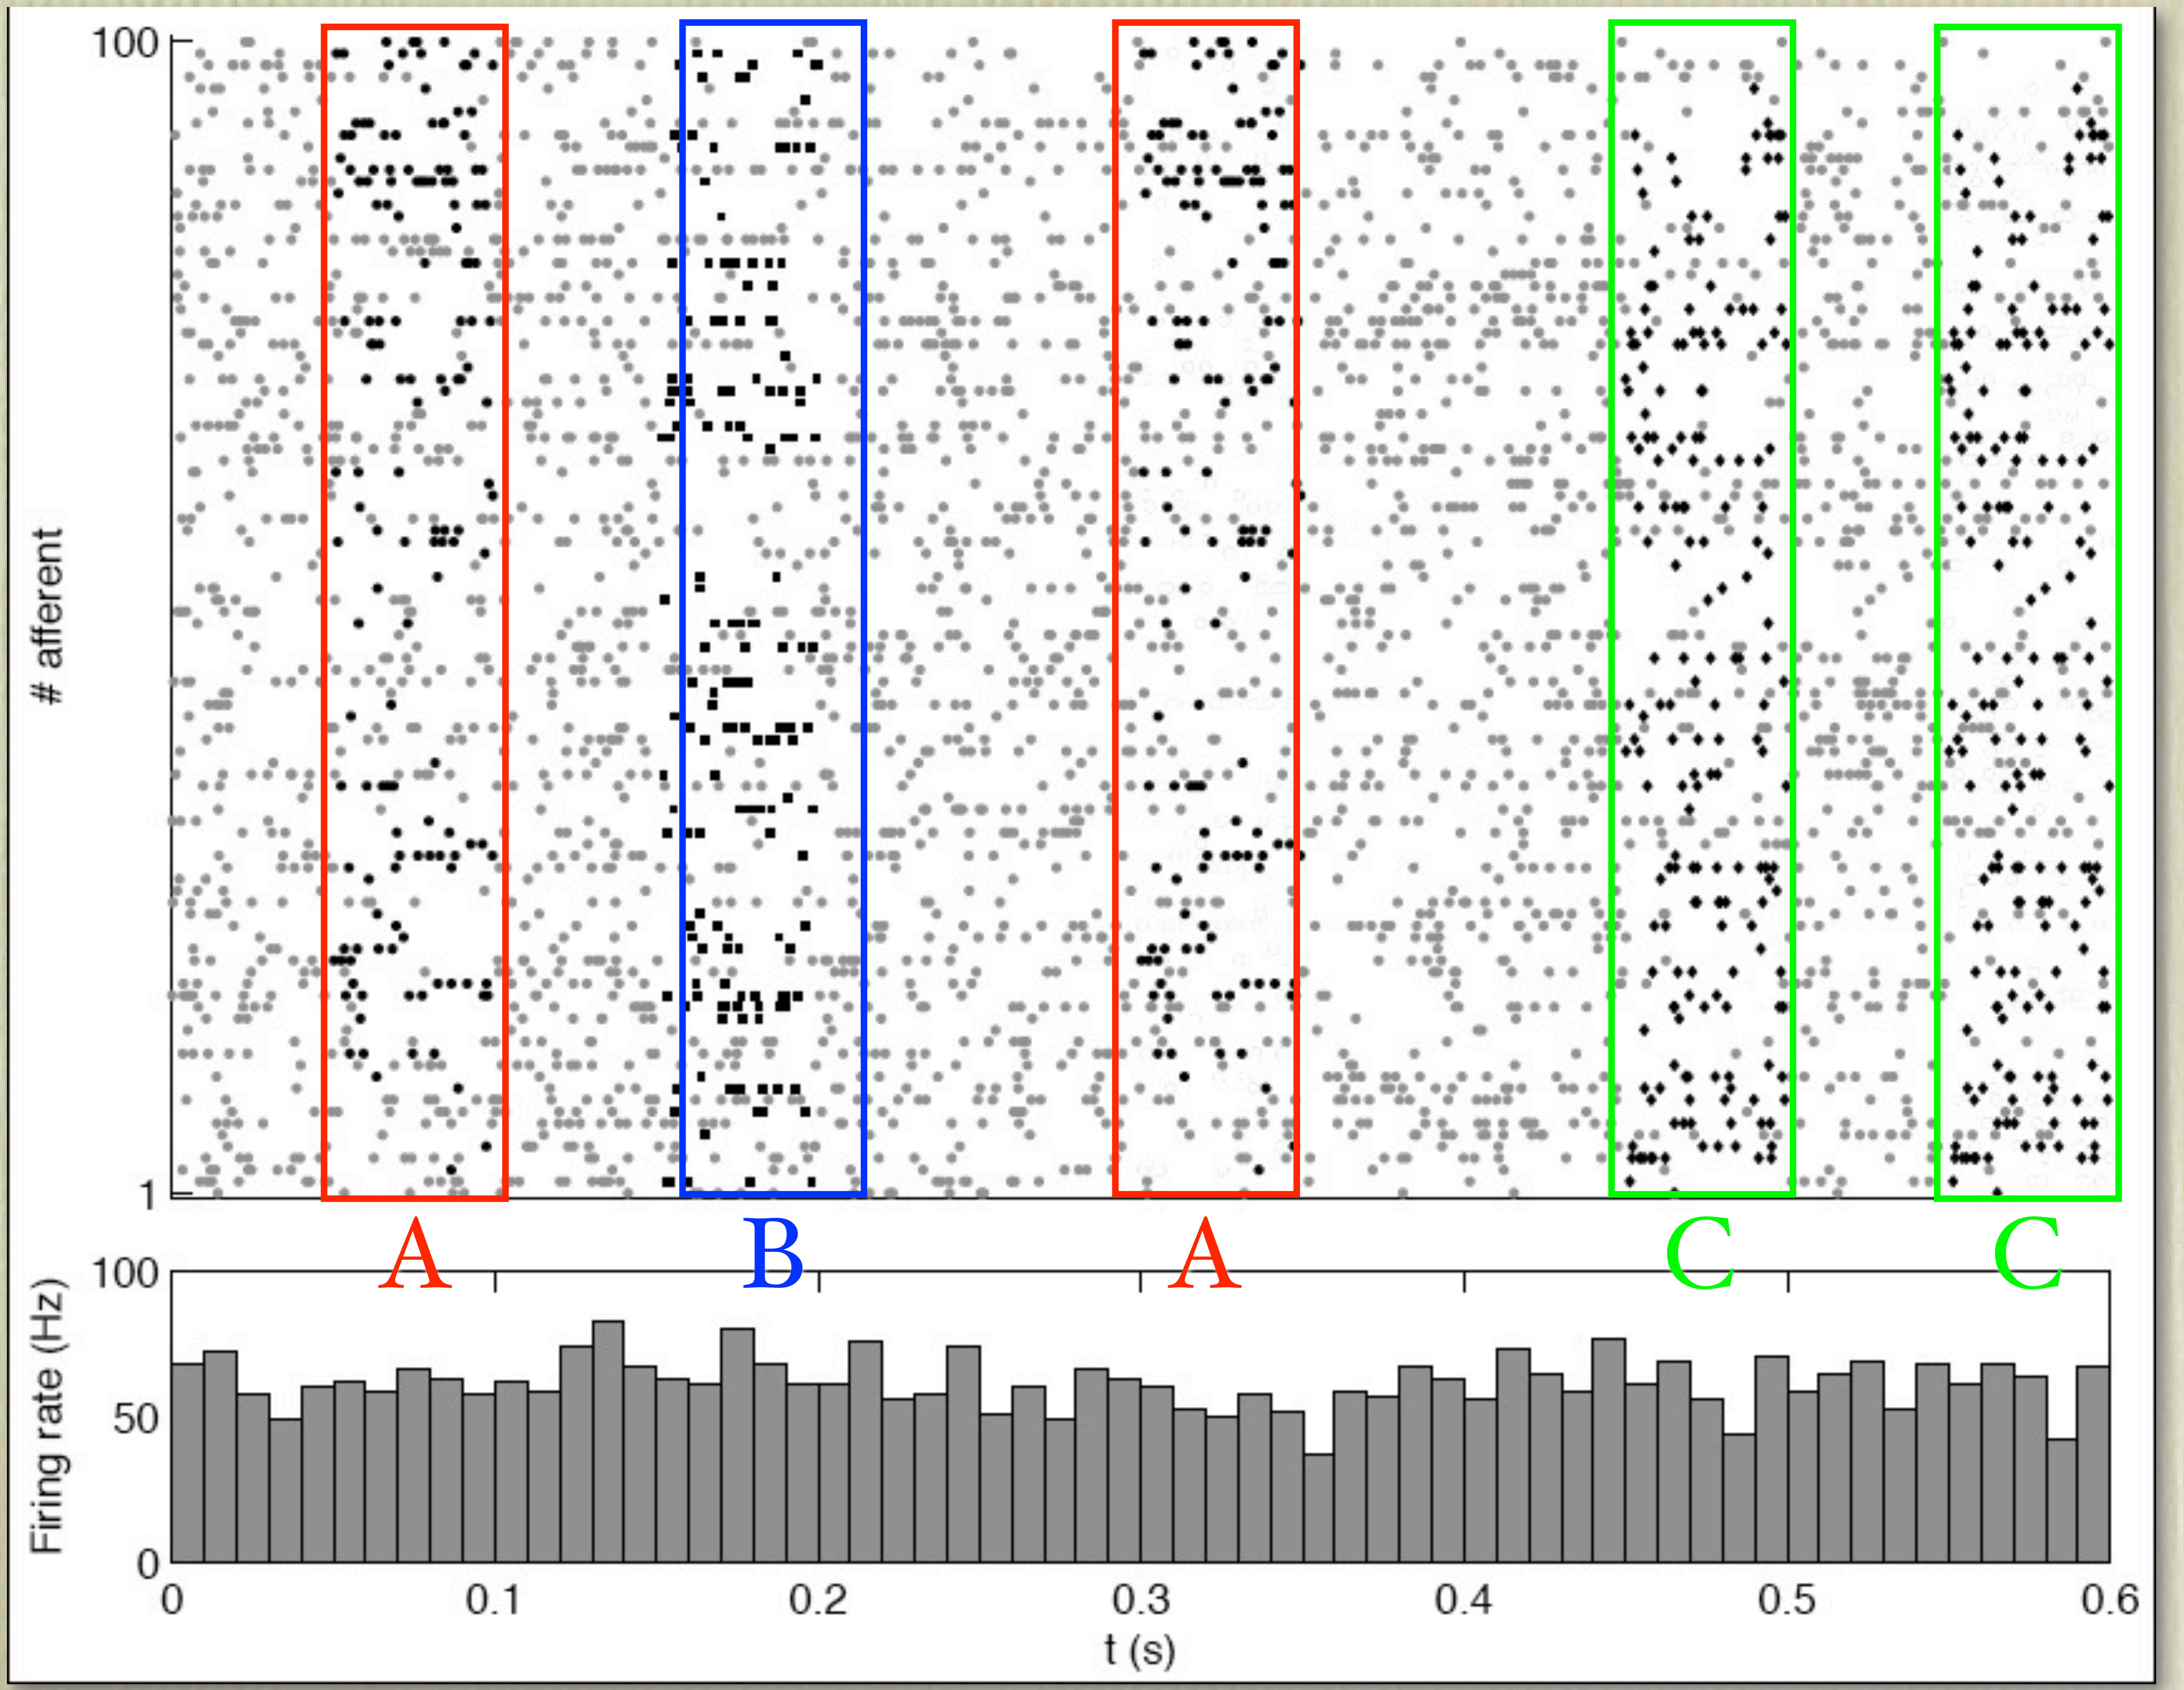

Learning to Detect Faces with STDP

Unsupervised Learning of Visual Features through Spike Timing Dependent Plasticity

Timothée Masquelier[1,2*], Simon J. Thorpe[1,2] PLoS COMPUTATIONAL BIOLOGY February 2007

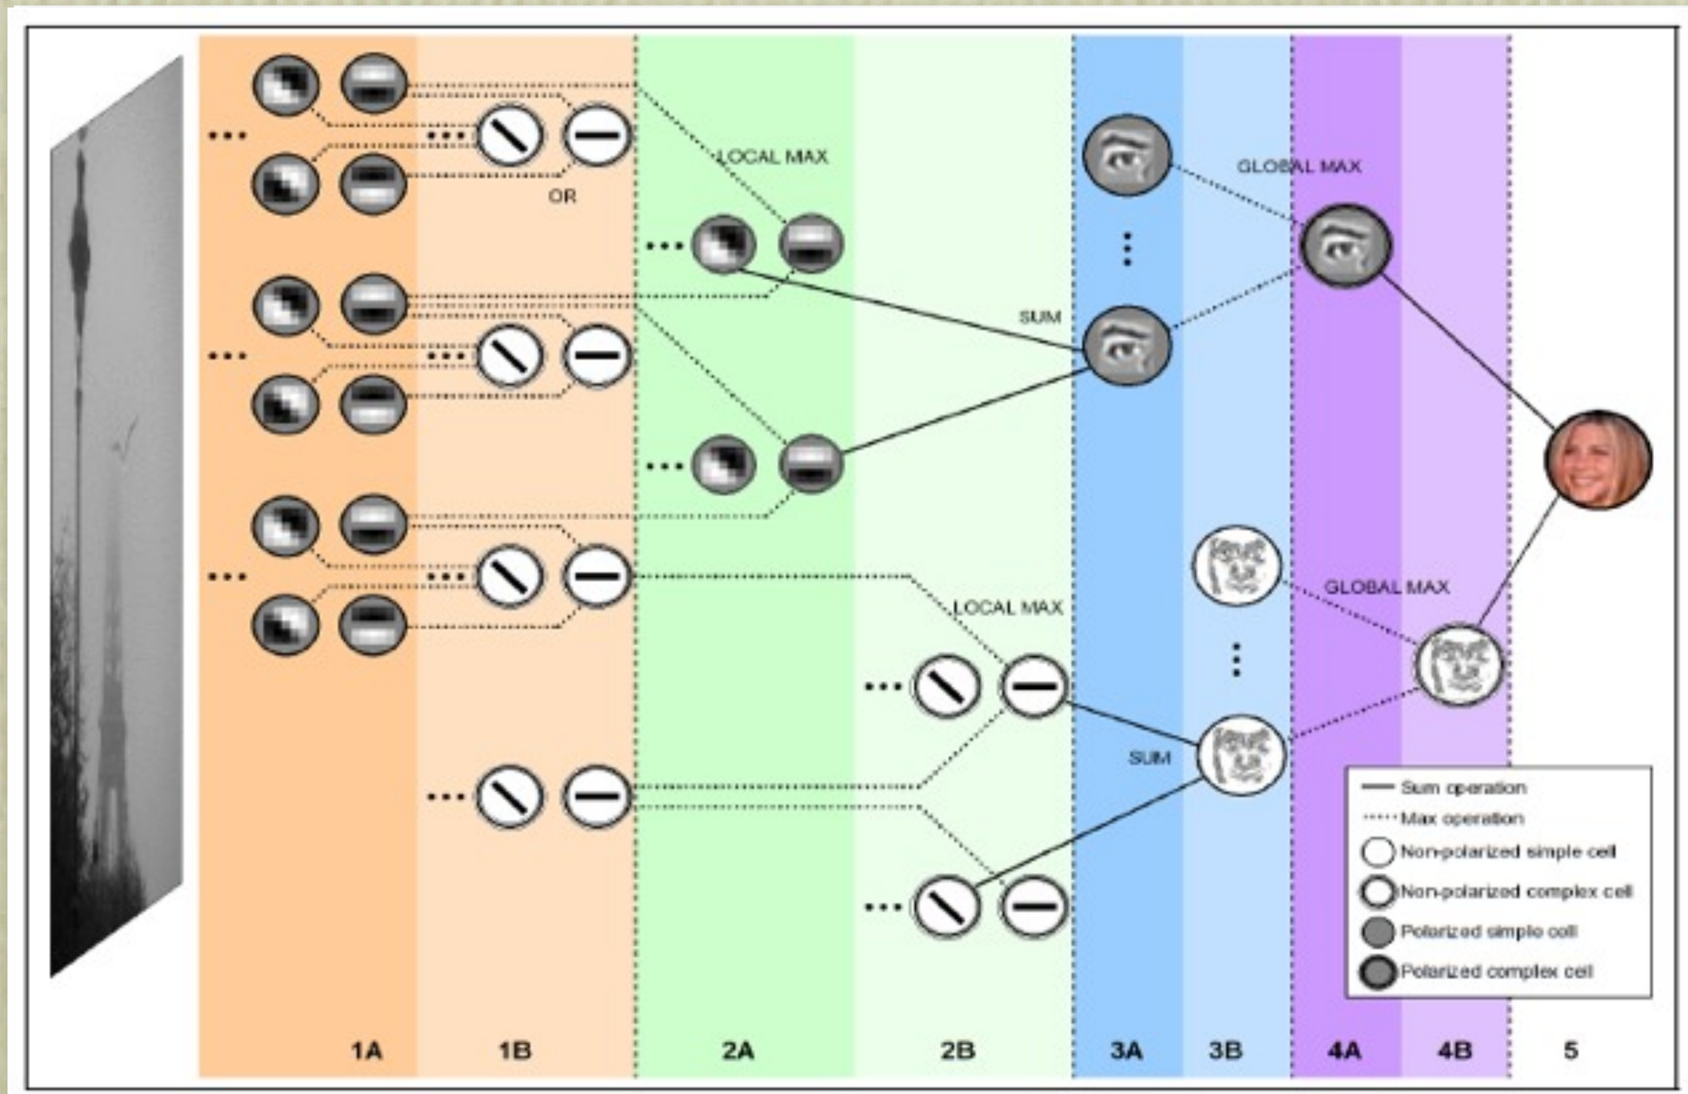

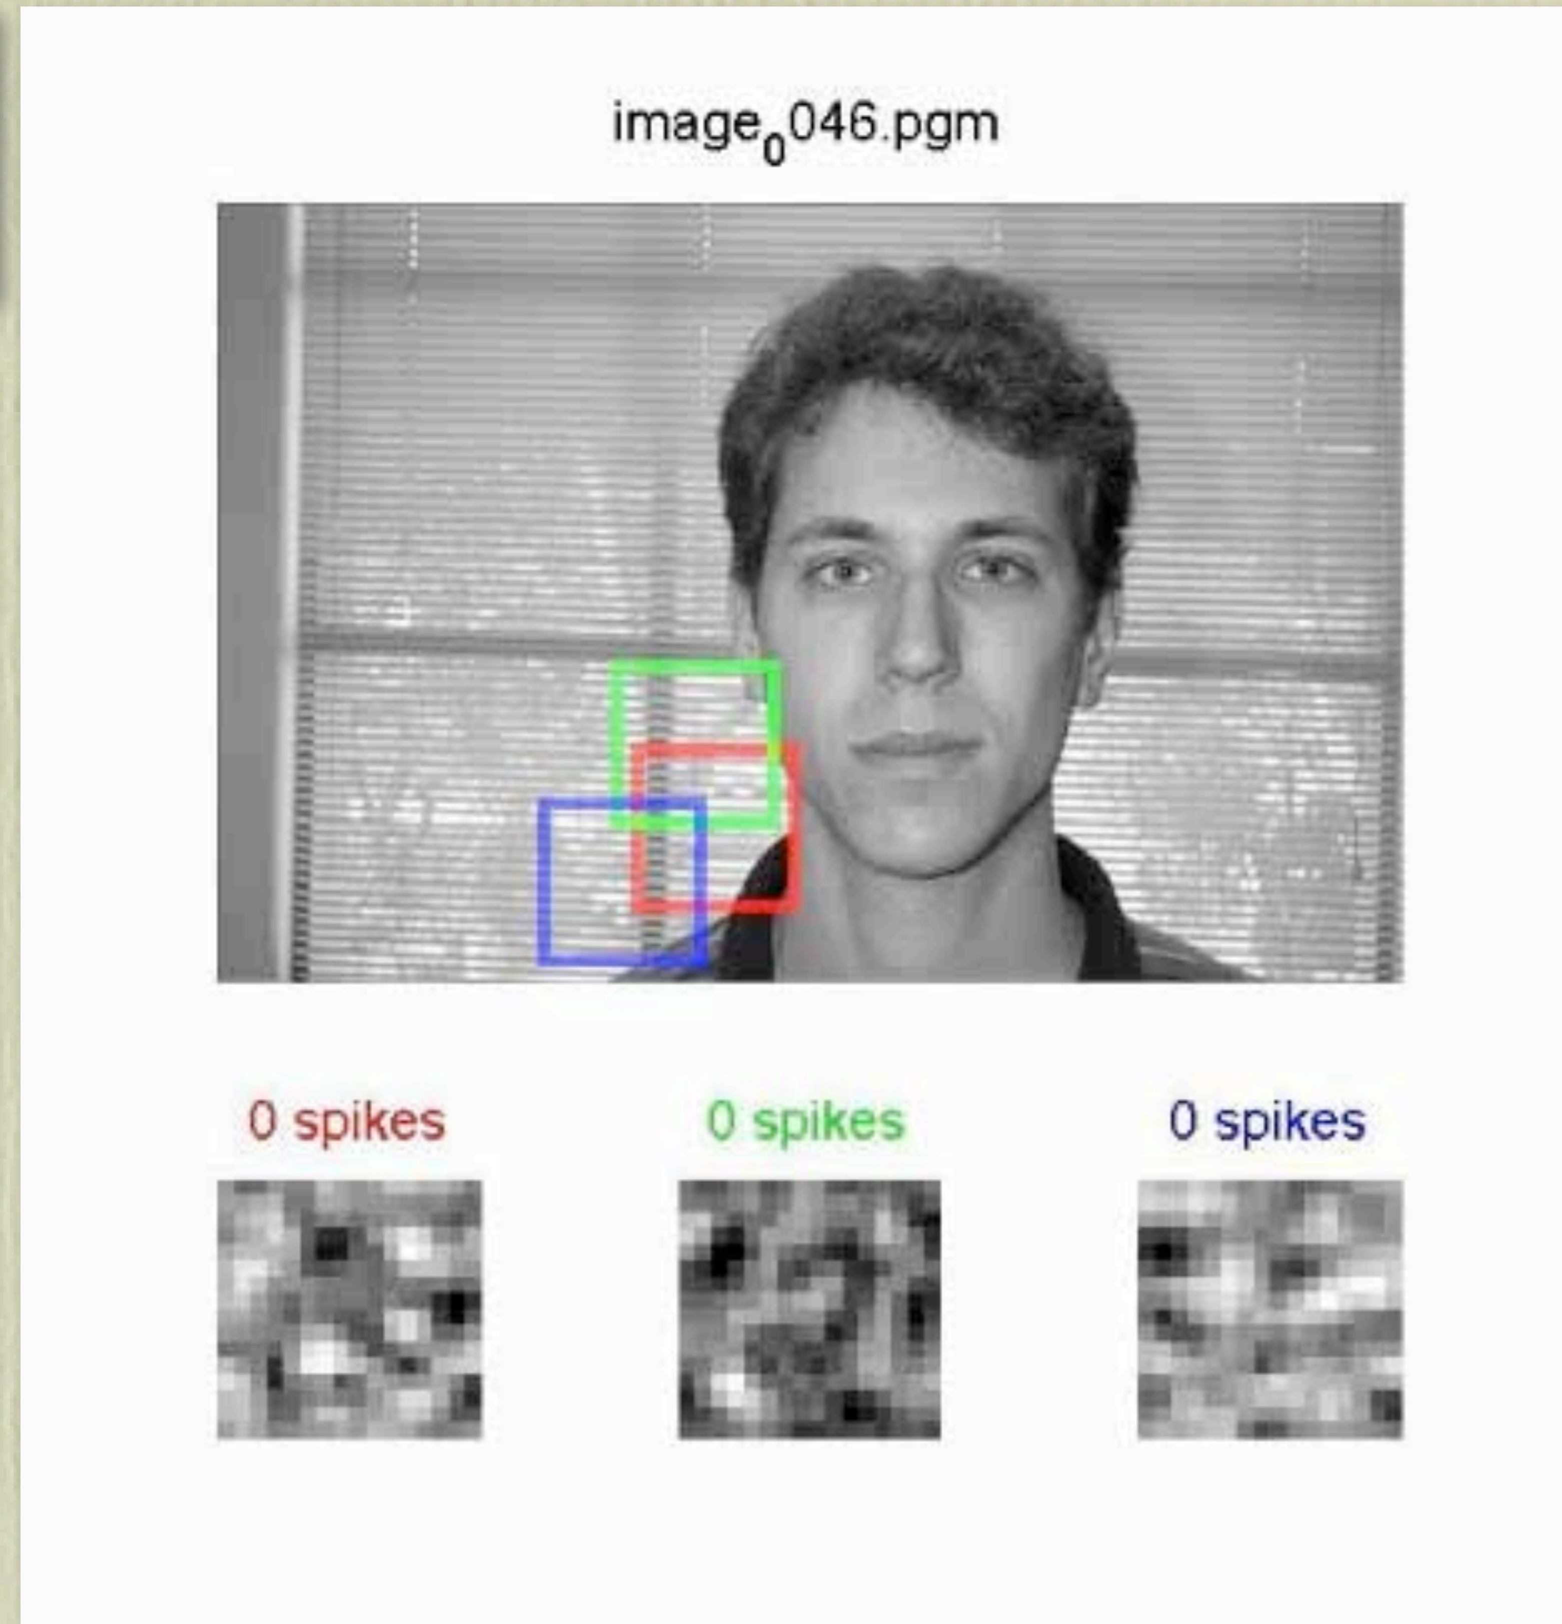

- More complete hierarchical architecture
- Similar to Serre, Wolf & Poggio
- Modified STDP Rule

STDP based learning

- STDP concentrates high synaptic weights on early firing inputs
- Inhibitory connections between neurons allows them to function as a competitive learning system in which different neurons will tend to learn different stimuli
- Different neurons will learn to respond to different parts of the same pattern
- Only a small number of presentations may be needed for changes to occur

Hypothesis

- 16 billion neurons in the neocortex
- Each tries to find repeating spike patterns that have not already been found by other neurons
- Patterns can be found
  - In sensory inputs from the retina, cochlear, somatosensory system, olfactory system
  - In feedback from later processing stages
  - In the motor system
- A key to intelligence?
- Can this be simulated?
- Can it be implemented in hardware?

Plan

- Biological and Computer Vision - Can we simulate the Visual System?
  - Classical Artificial Neural Networks
  - Spike based Computing
- STDP-based learning mechanisms
  - Finding the start of repeating patterns
  - Competitive Learning Networks
- Towards Neuromorphic hardware
  - Address Event Representation (AER) Coding
  - Spiking Retinas
  - Spiking Cochlears
  - Memristor based hardware

Address Event Representation (AER) Coding

- Code the identity of the neuron that spikes with a number

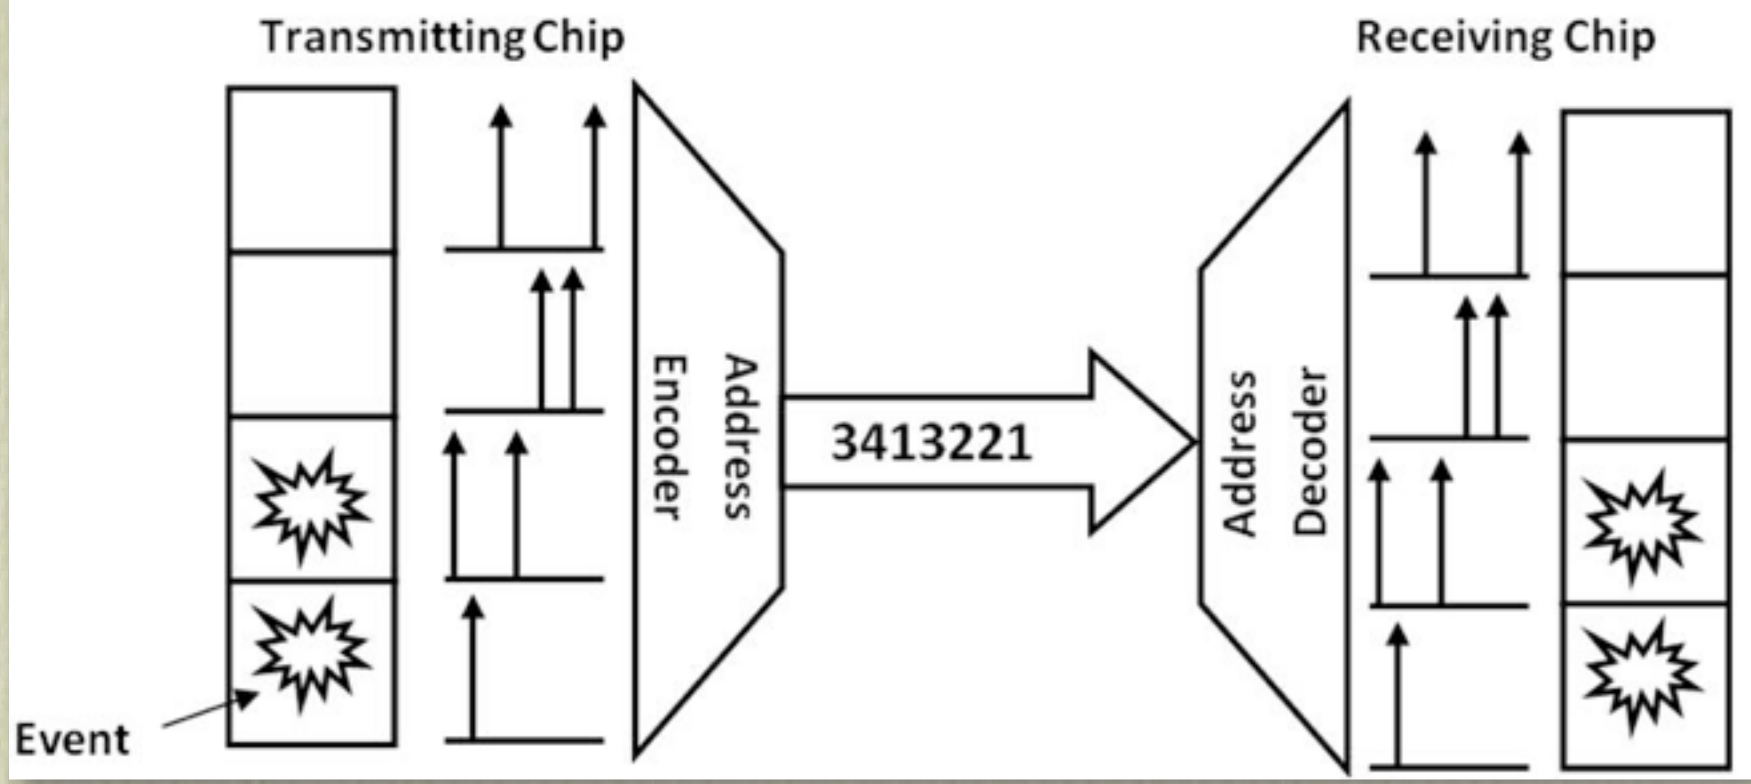

EU funded Caviar project

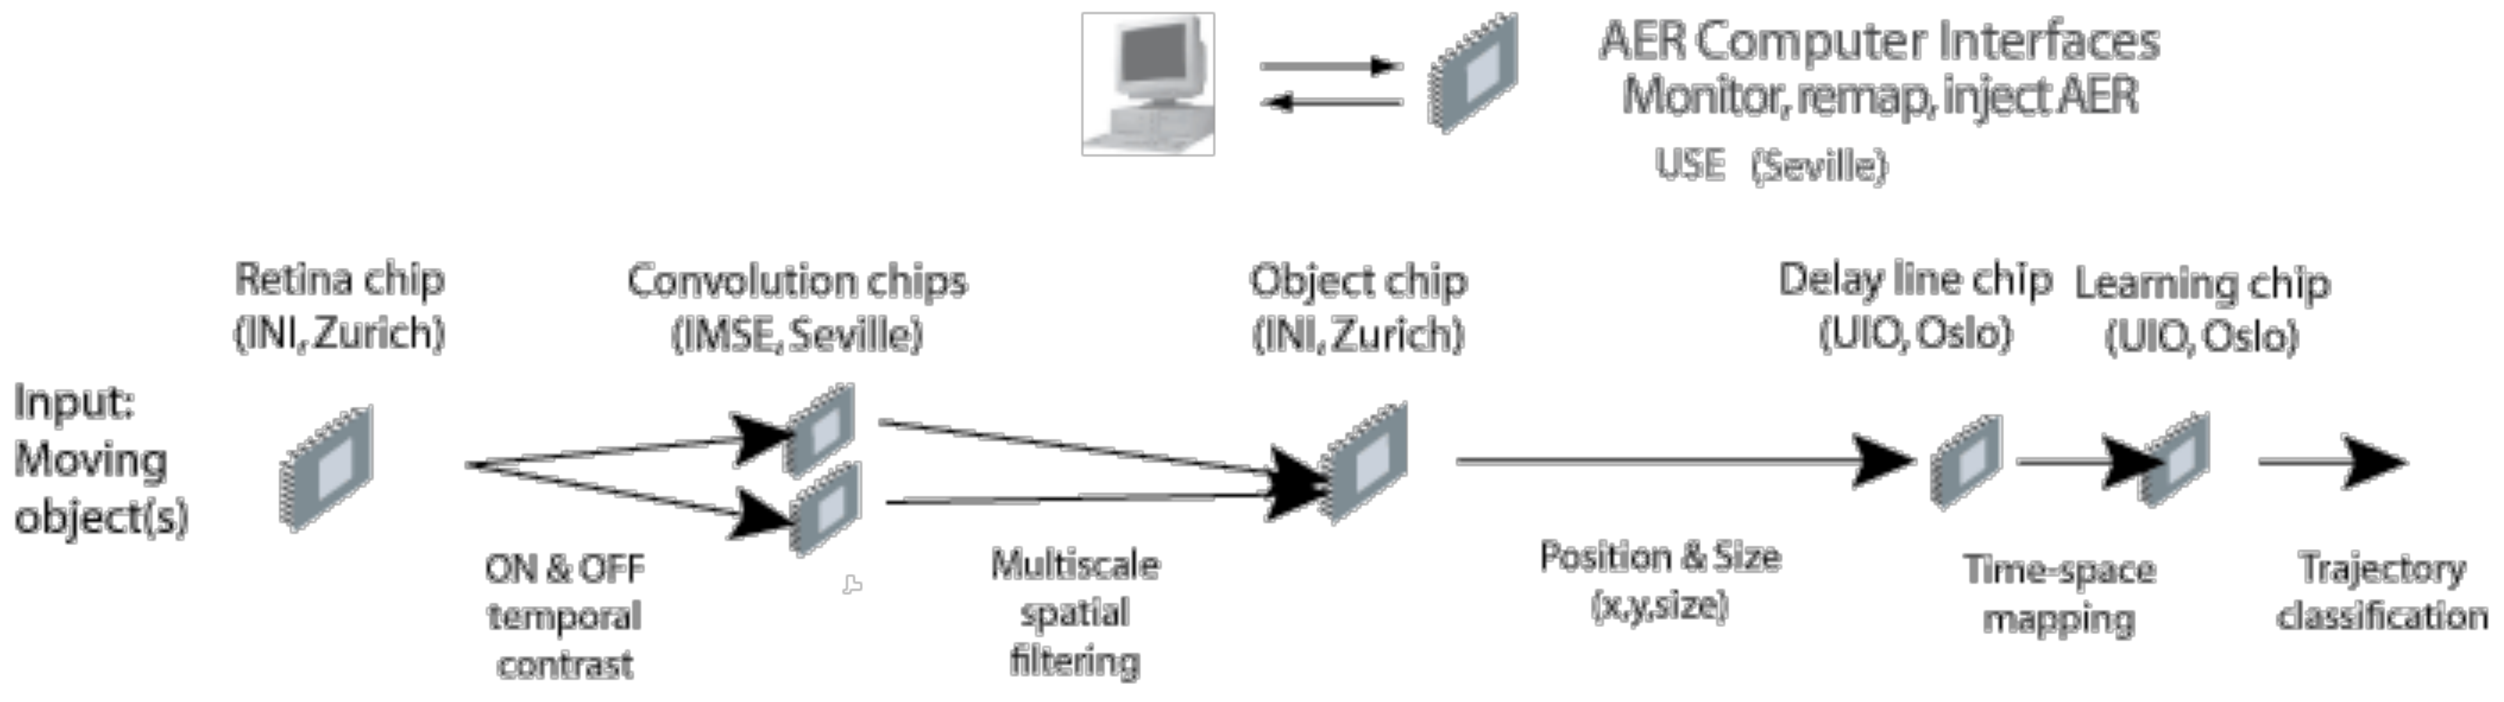

Applications in Vision

Contents lists available at SciVerse ScienceDirect

Neural Networks

journal homepage: www.elsevier.com/locate/neunet

ELSEVIER

2012 Special Issue

Extraction of temporally correlated features from dynamic vision sensors with spike-timing-dependent plasticity

Olivier Bichler[a,*], Damien Querlioz[b], Simon J. Thorpe[c], Jean-Philippe Bourgoin[d], Christian Gamrat[a]

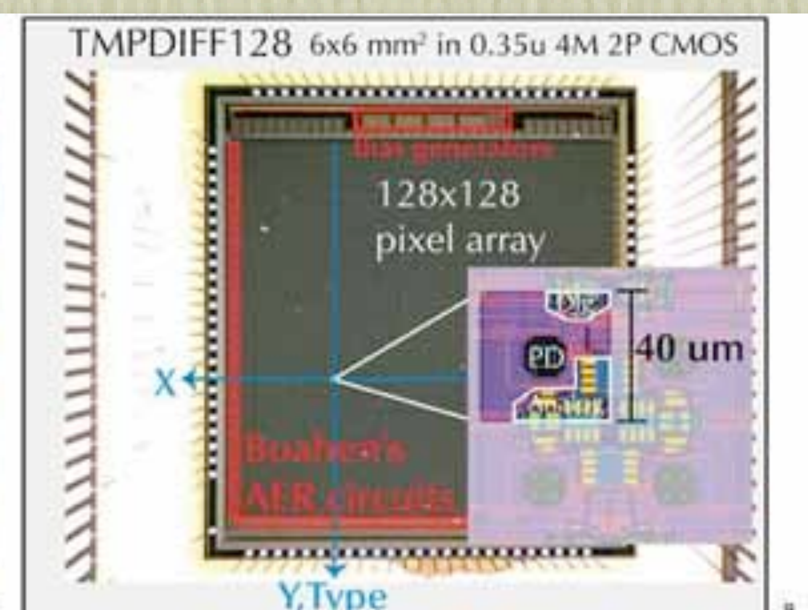

566 IEEE JOURNAL OF SOLID-STATE CIRCUITS, VOL. 43, NO. 2, FEBRUARY 2008

A 128×128 120 dB 15 μs Latency Asynchronous Temporal Contrast Vision Sensor

Patrick Lichtsteiner, *Member, IEEE*, Christoph Posch, *Member, IEEE*, and Tobi Delbruck, *Senior Member, IEEE*

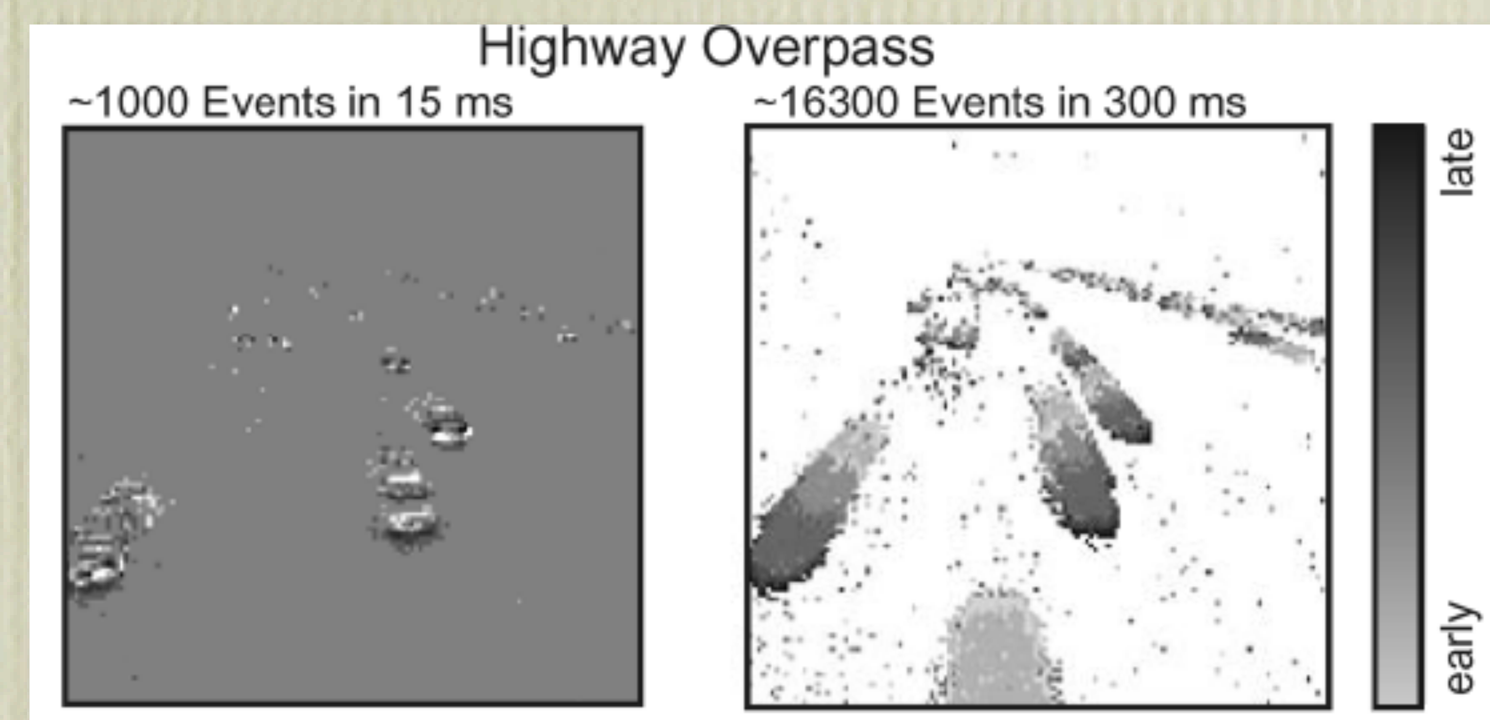

Simulation Studies

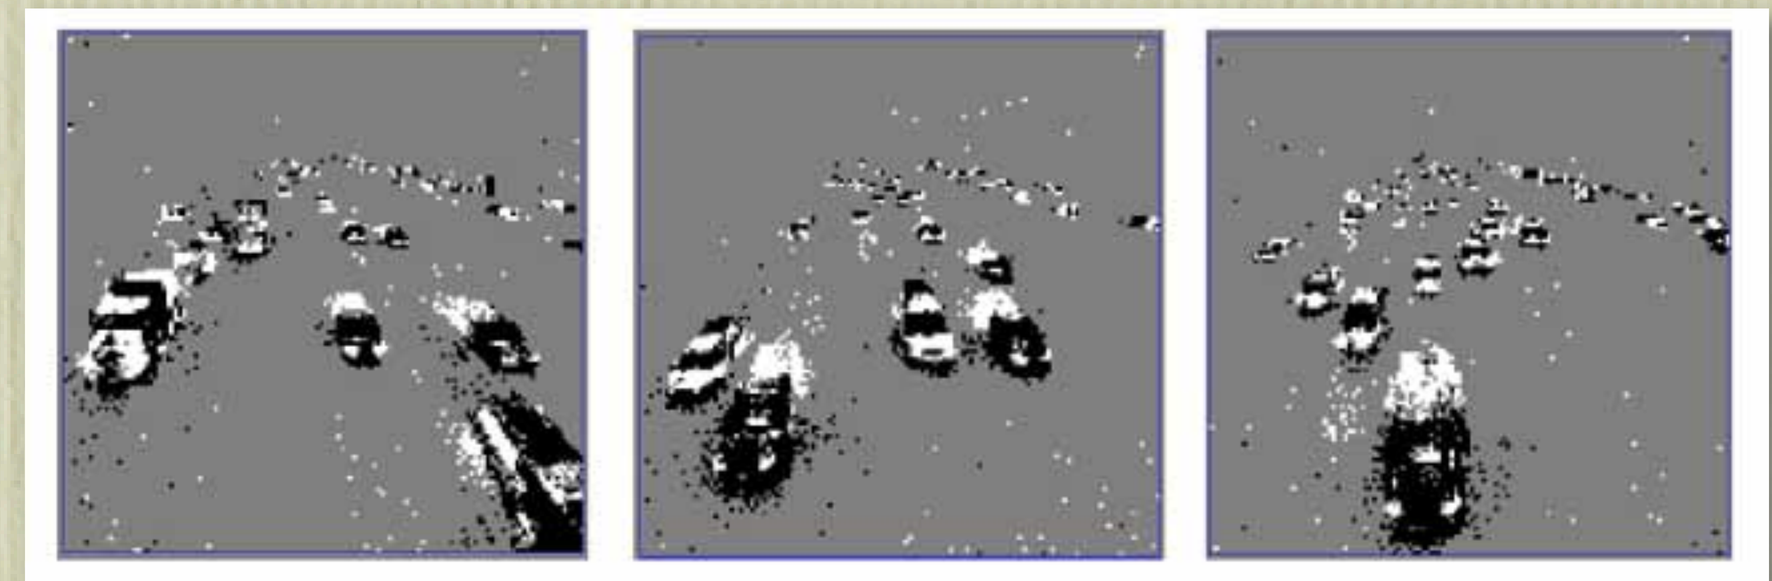

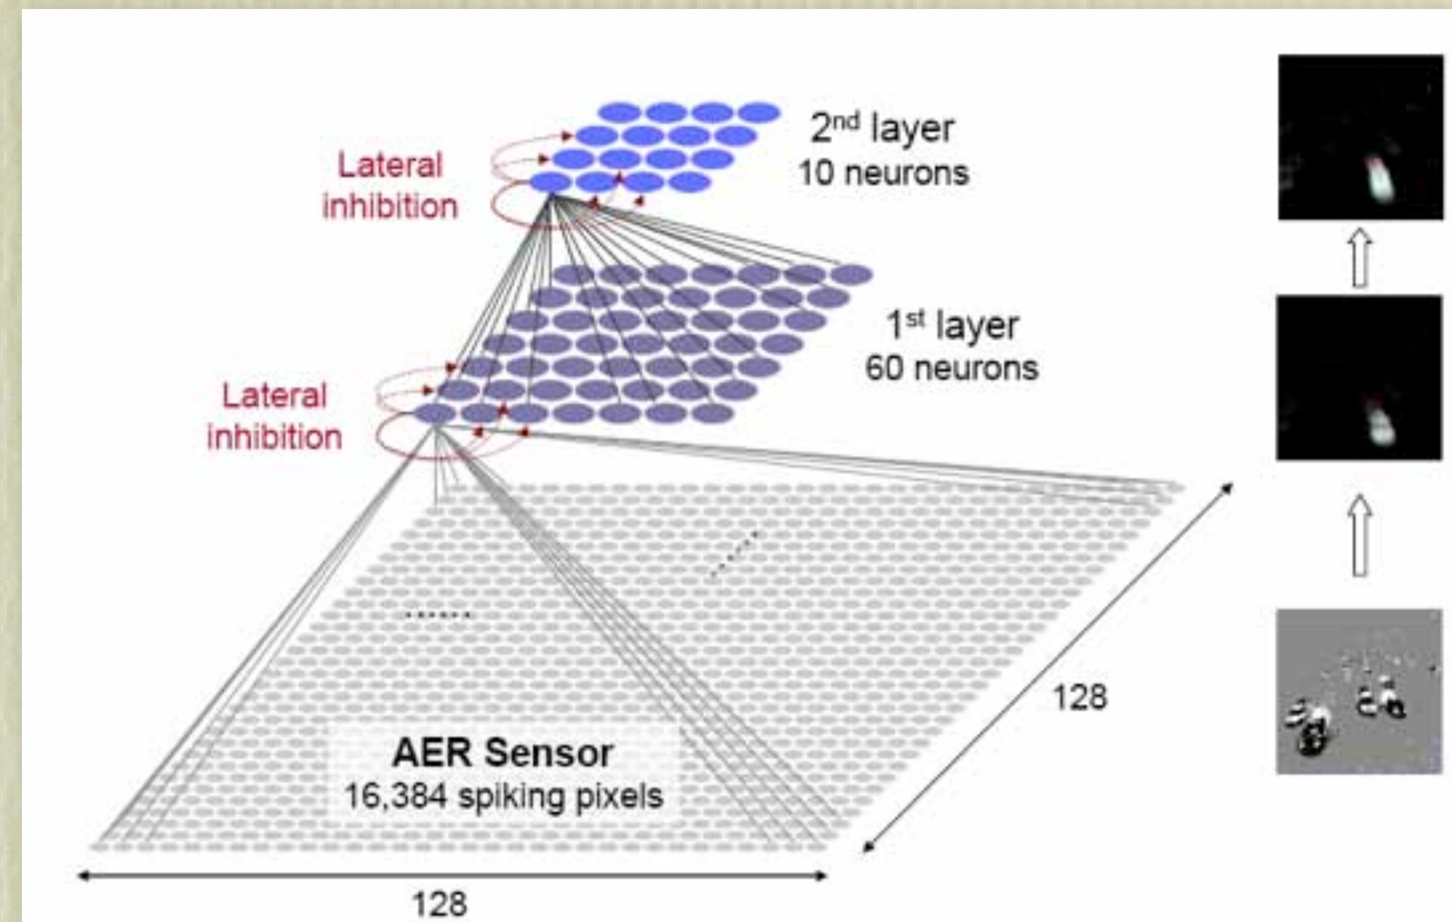

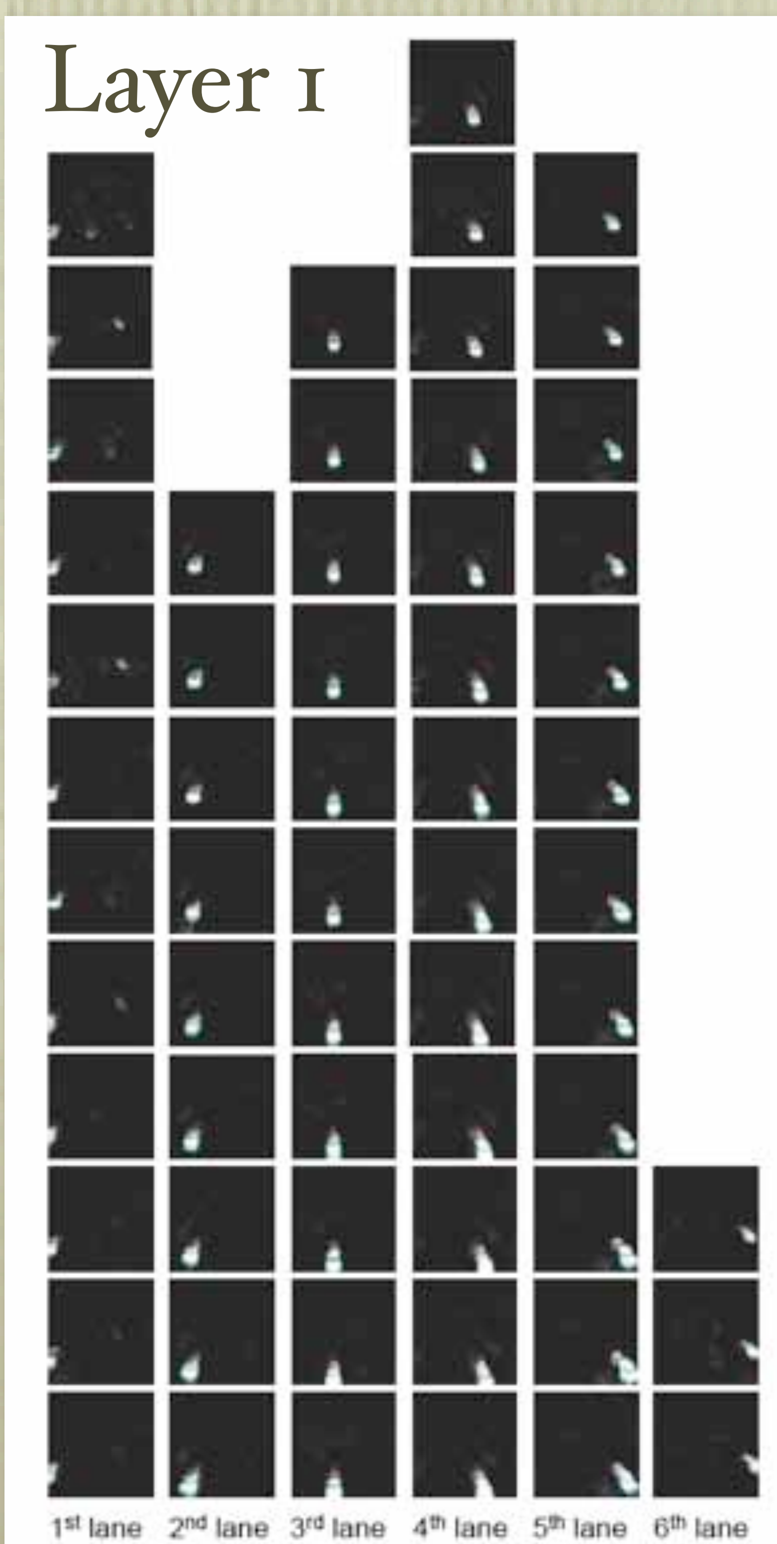

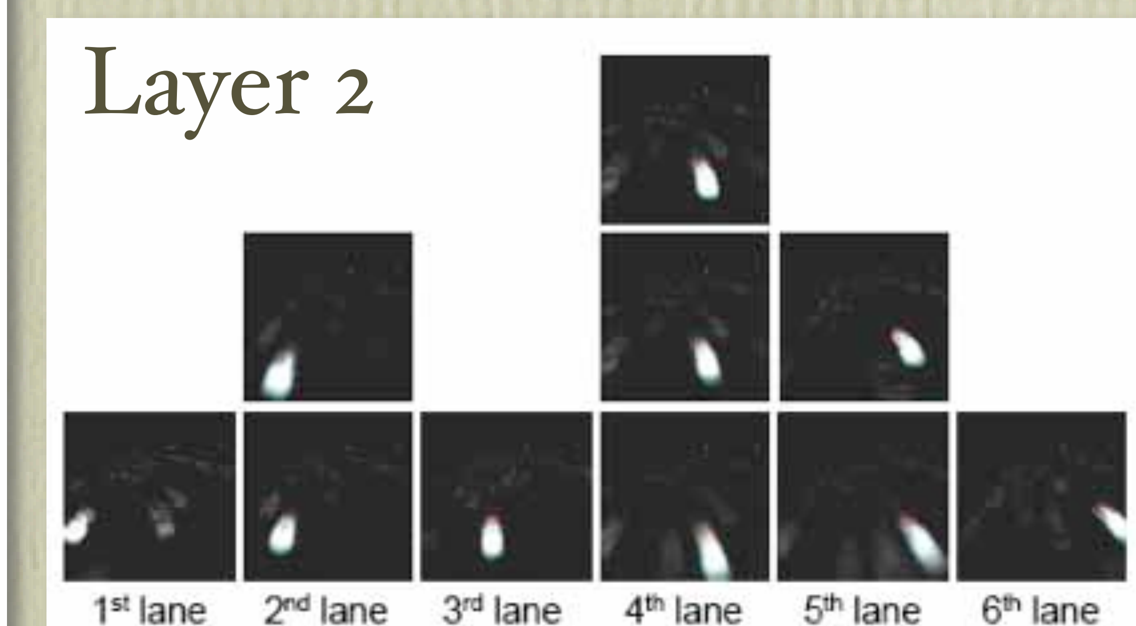

Neural Events
Reference Events

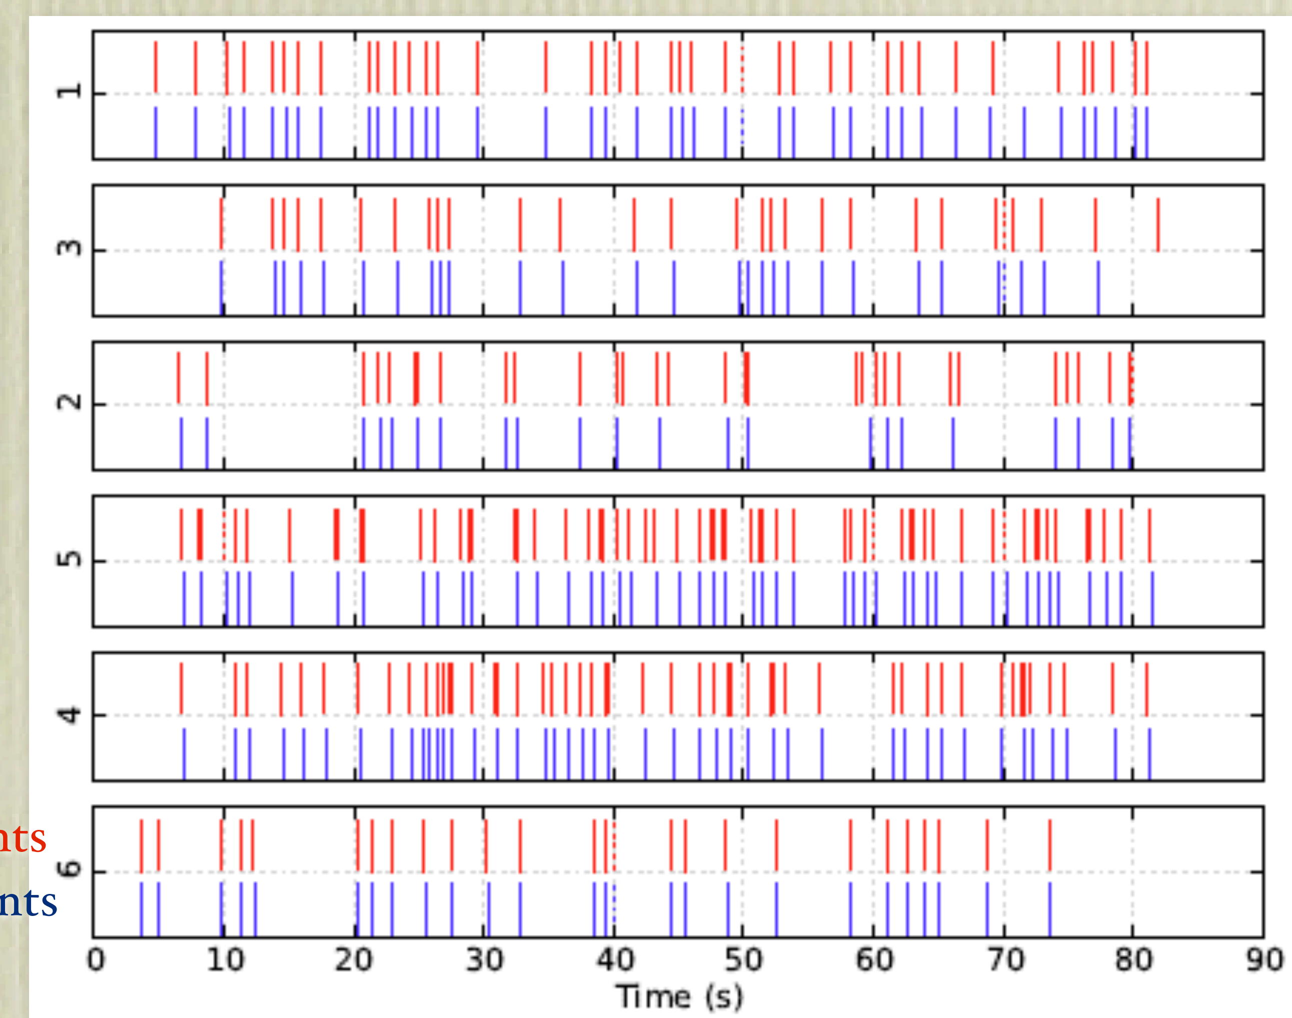

Performance 98% overall

Auditory Noise Learning in Humans

Rapid Formation of Robust Auditory Memories: Insights from Noise

Trevor R. Agus,[1,2,*] Simon J. Thorpe,[3,4] and Daniel Pressnitzer[1,2]

Neuron *66*, 610–618, May 27, 2010

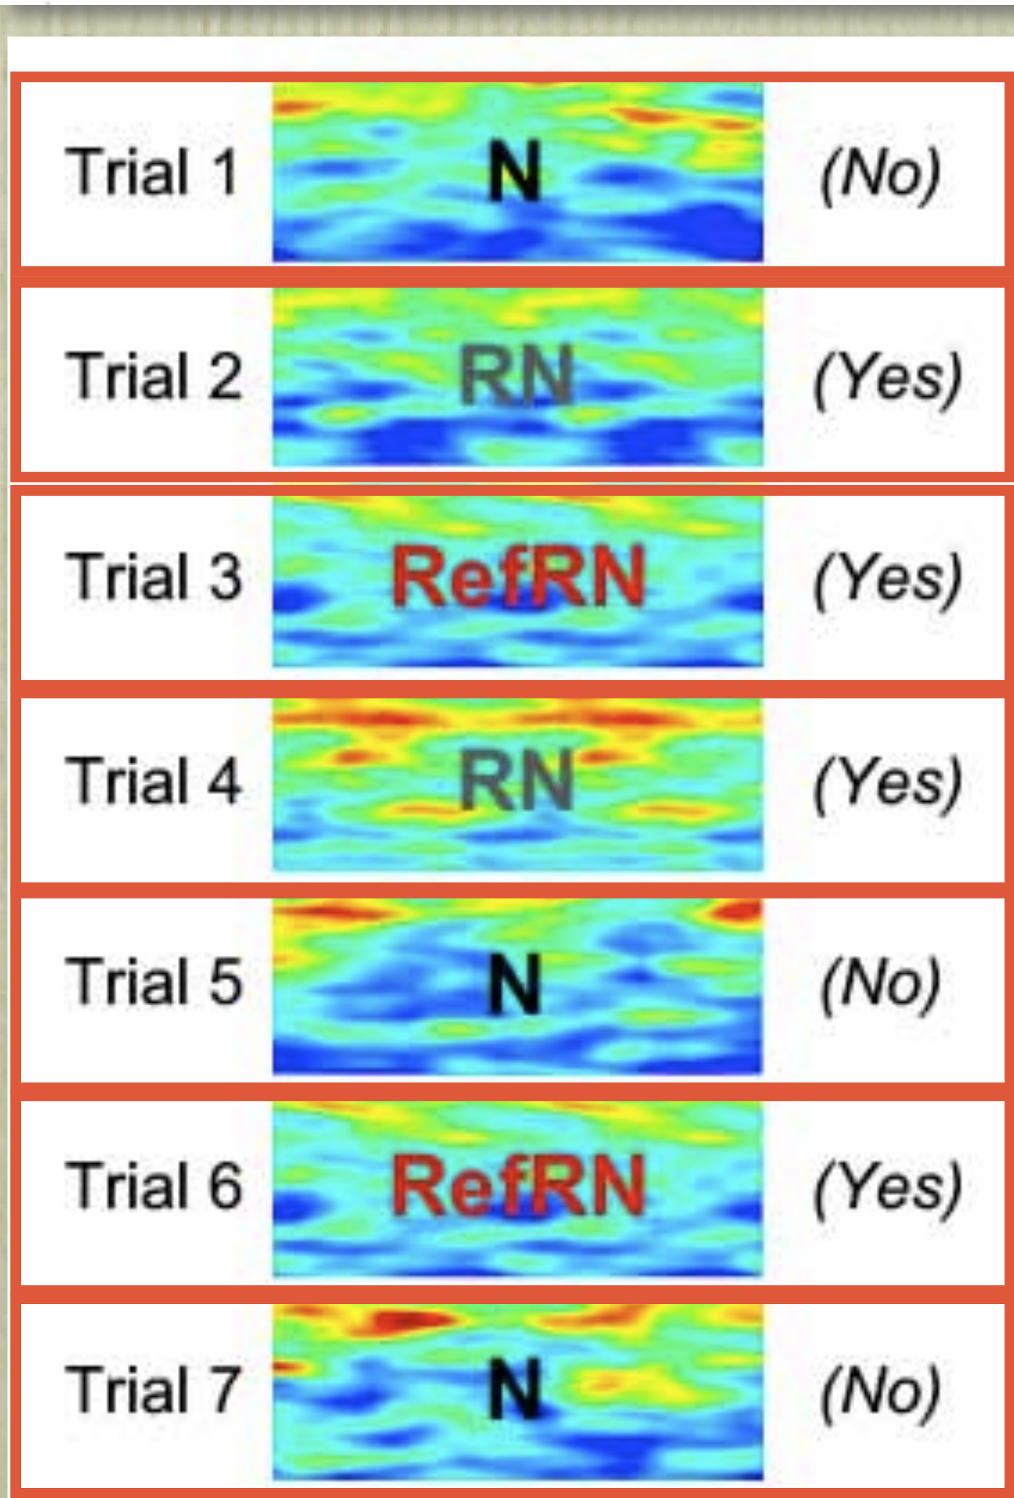

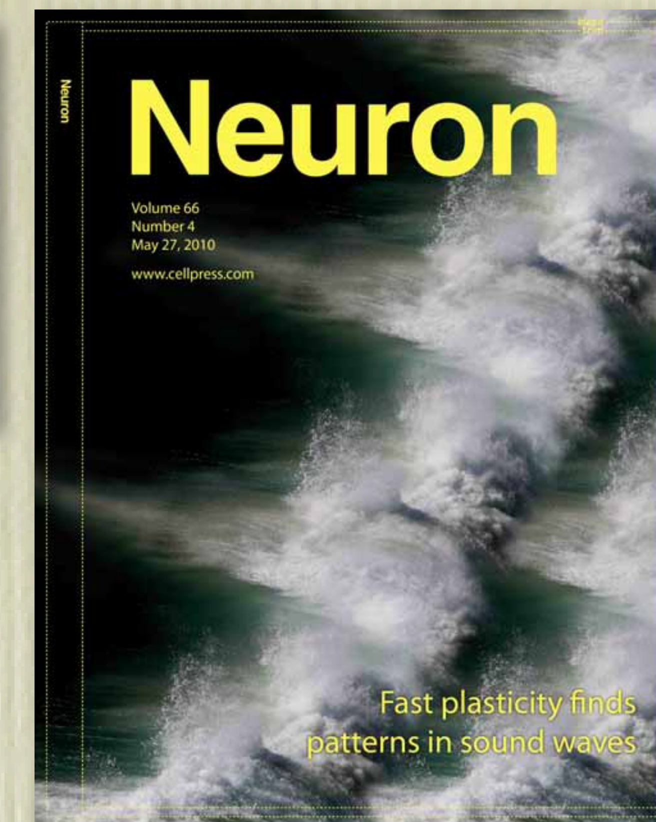

- Learning of random noise patterns
- Roughly 10 repetitions are sufficient!
- Learning appears to be all-or-none
- Lasts for weeks at least

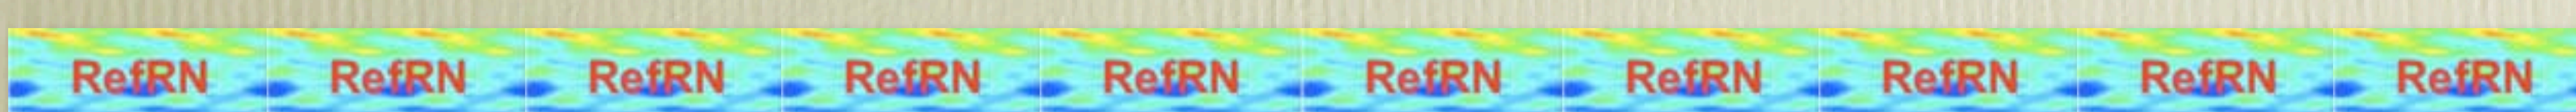

Auditory Noise Learning with STDP

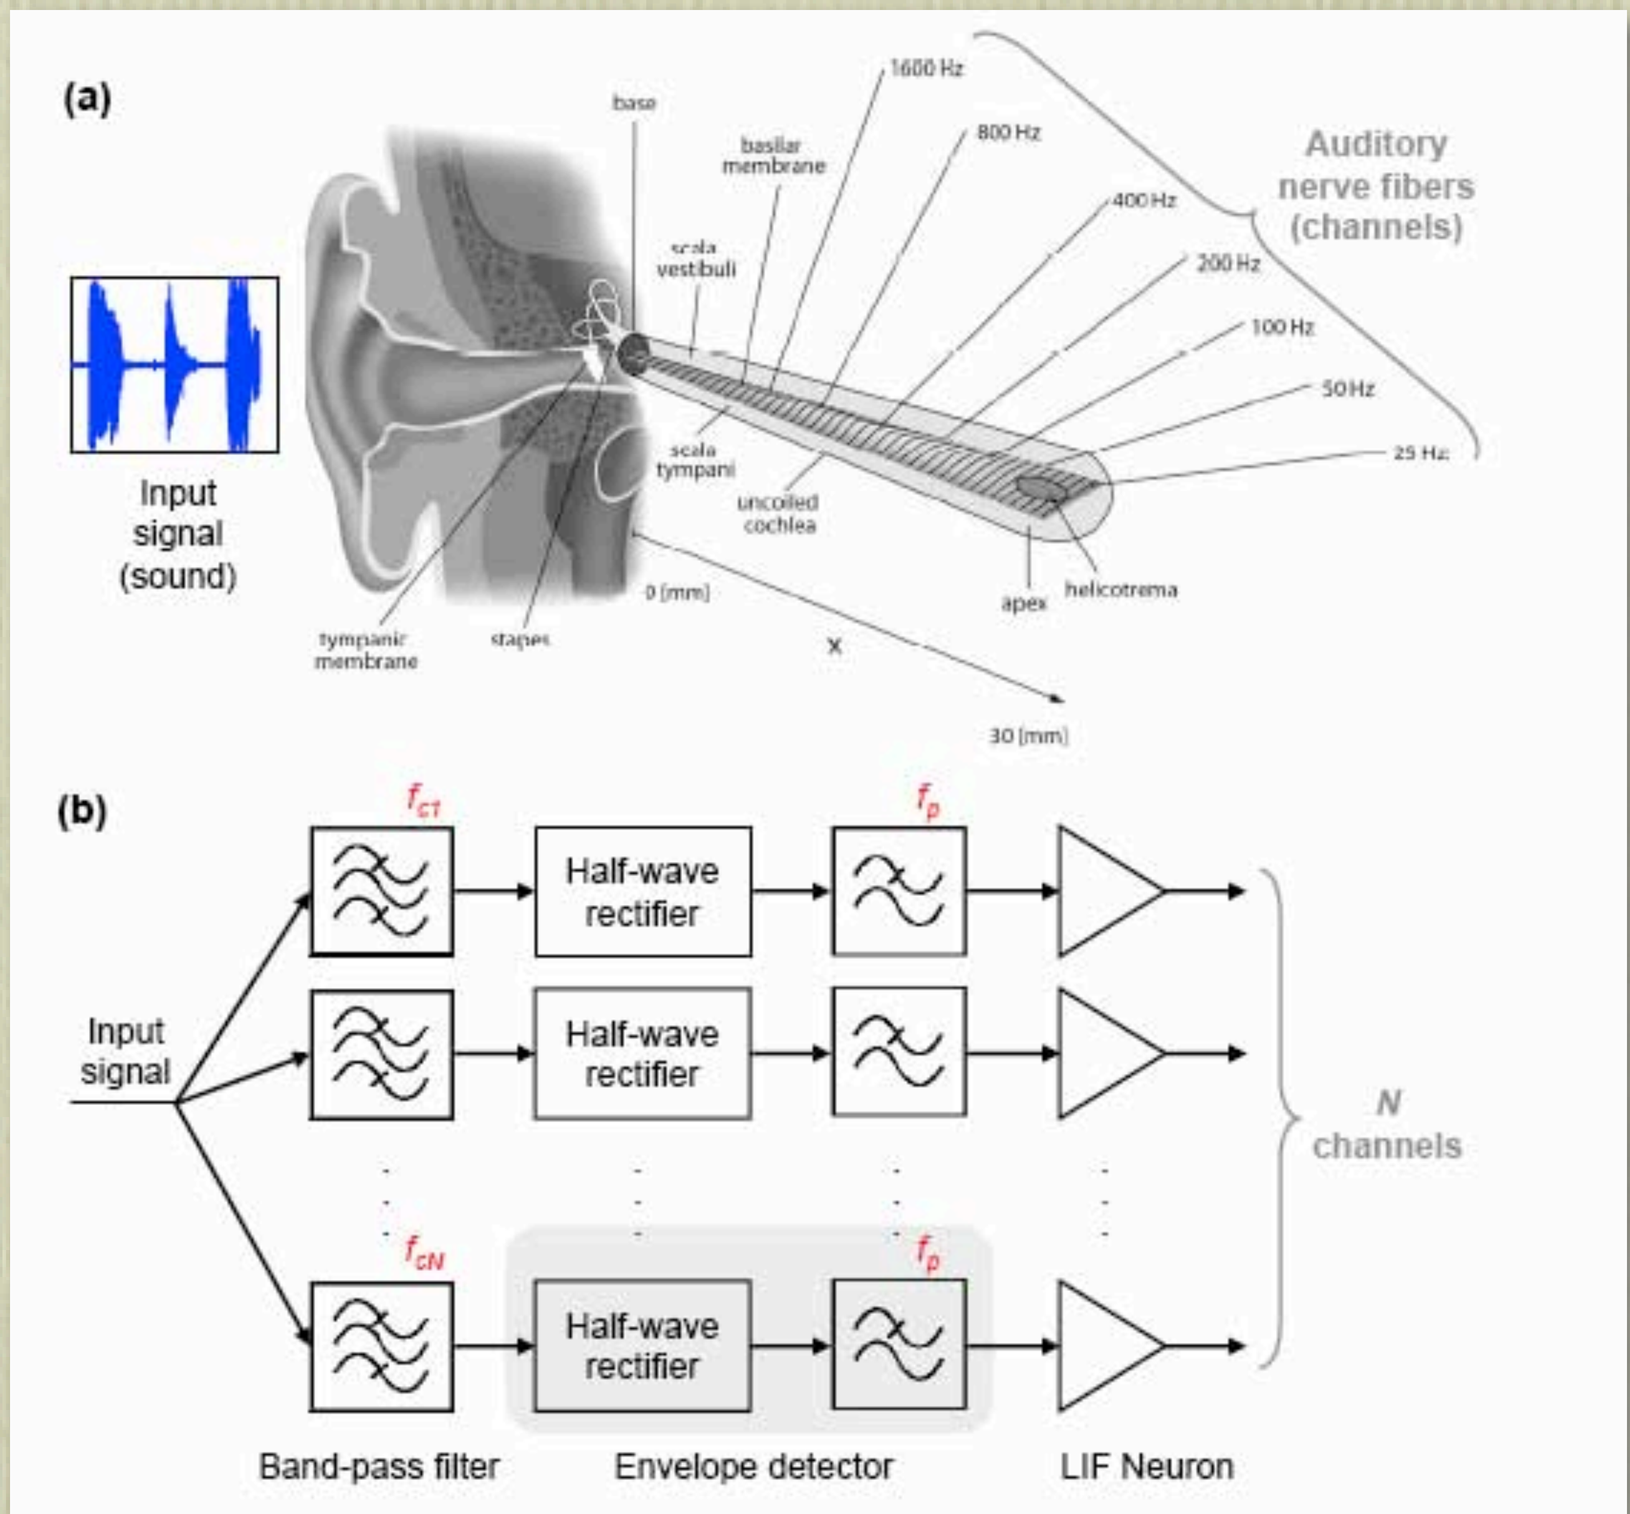

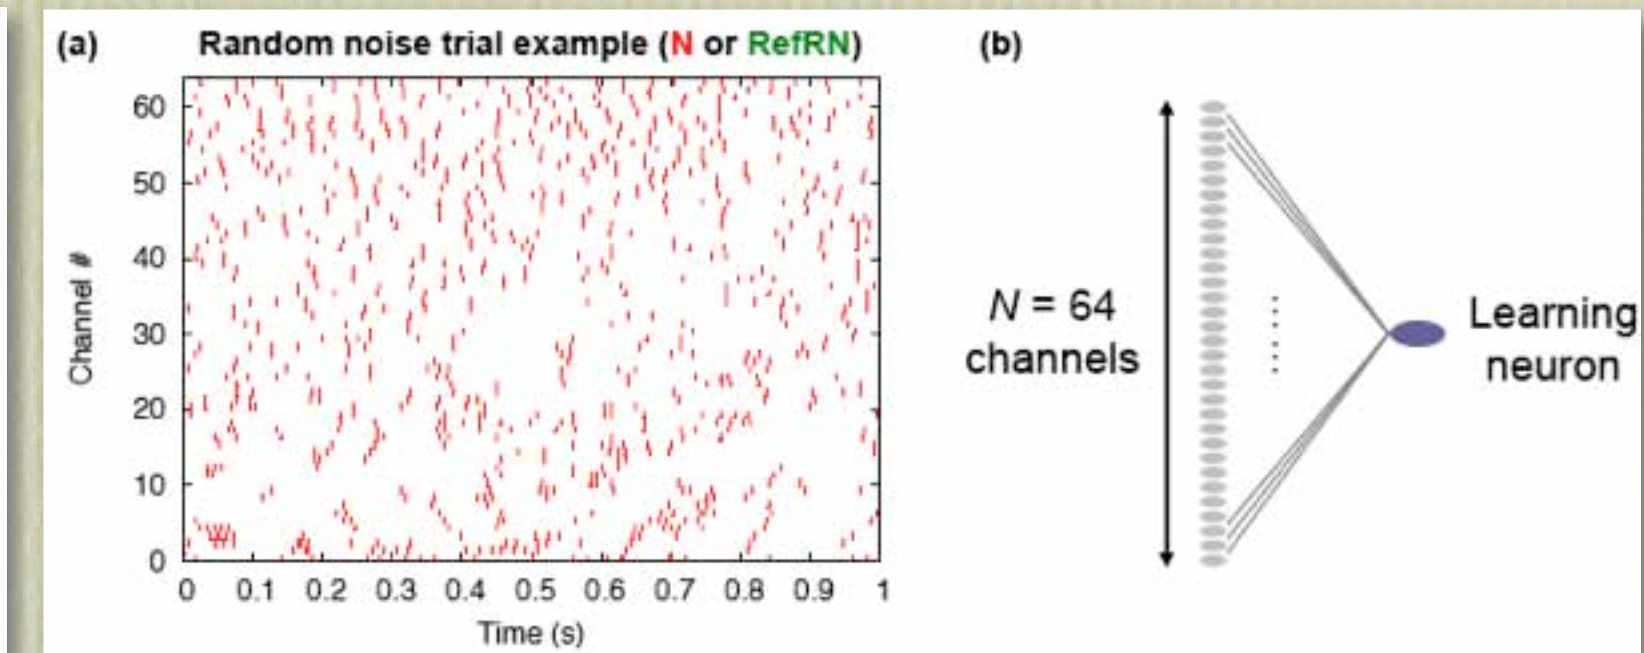

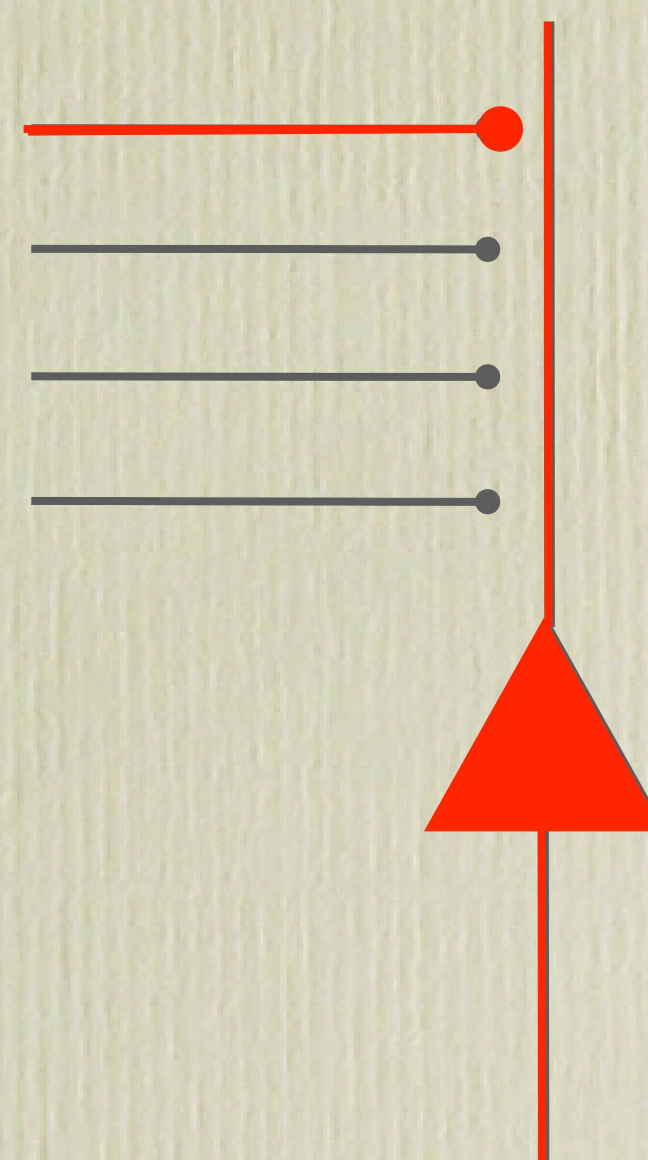

Olivier Bichler, Thesis

Modified STDP rule

- Post synaptic spike - depresses all synapses except those activated recently

Auditory Noise Learning with STDP

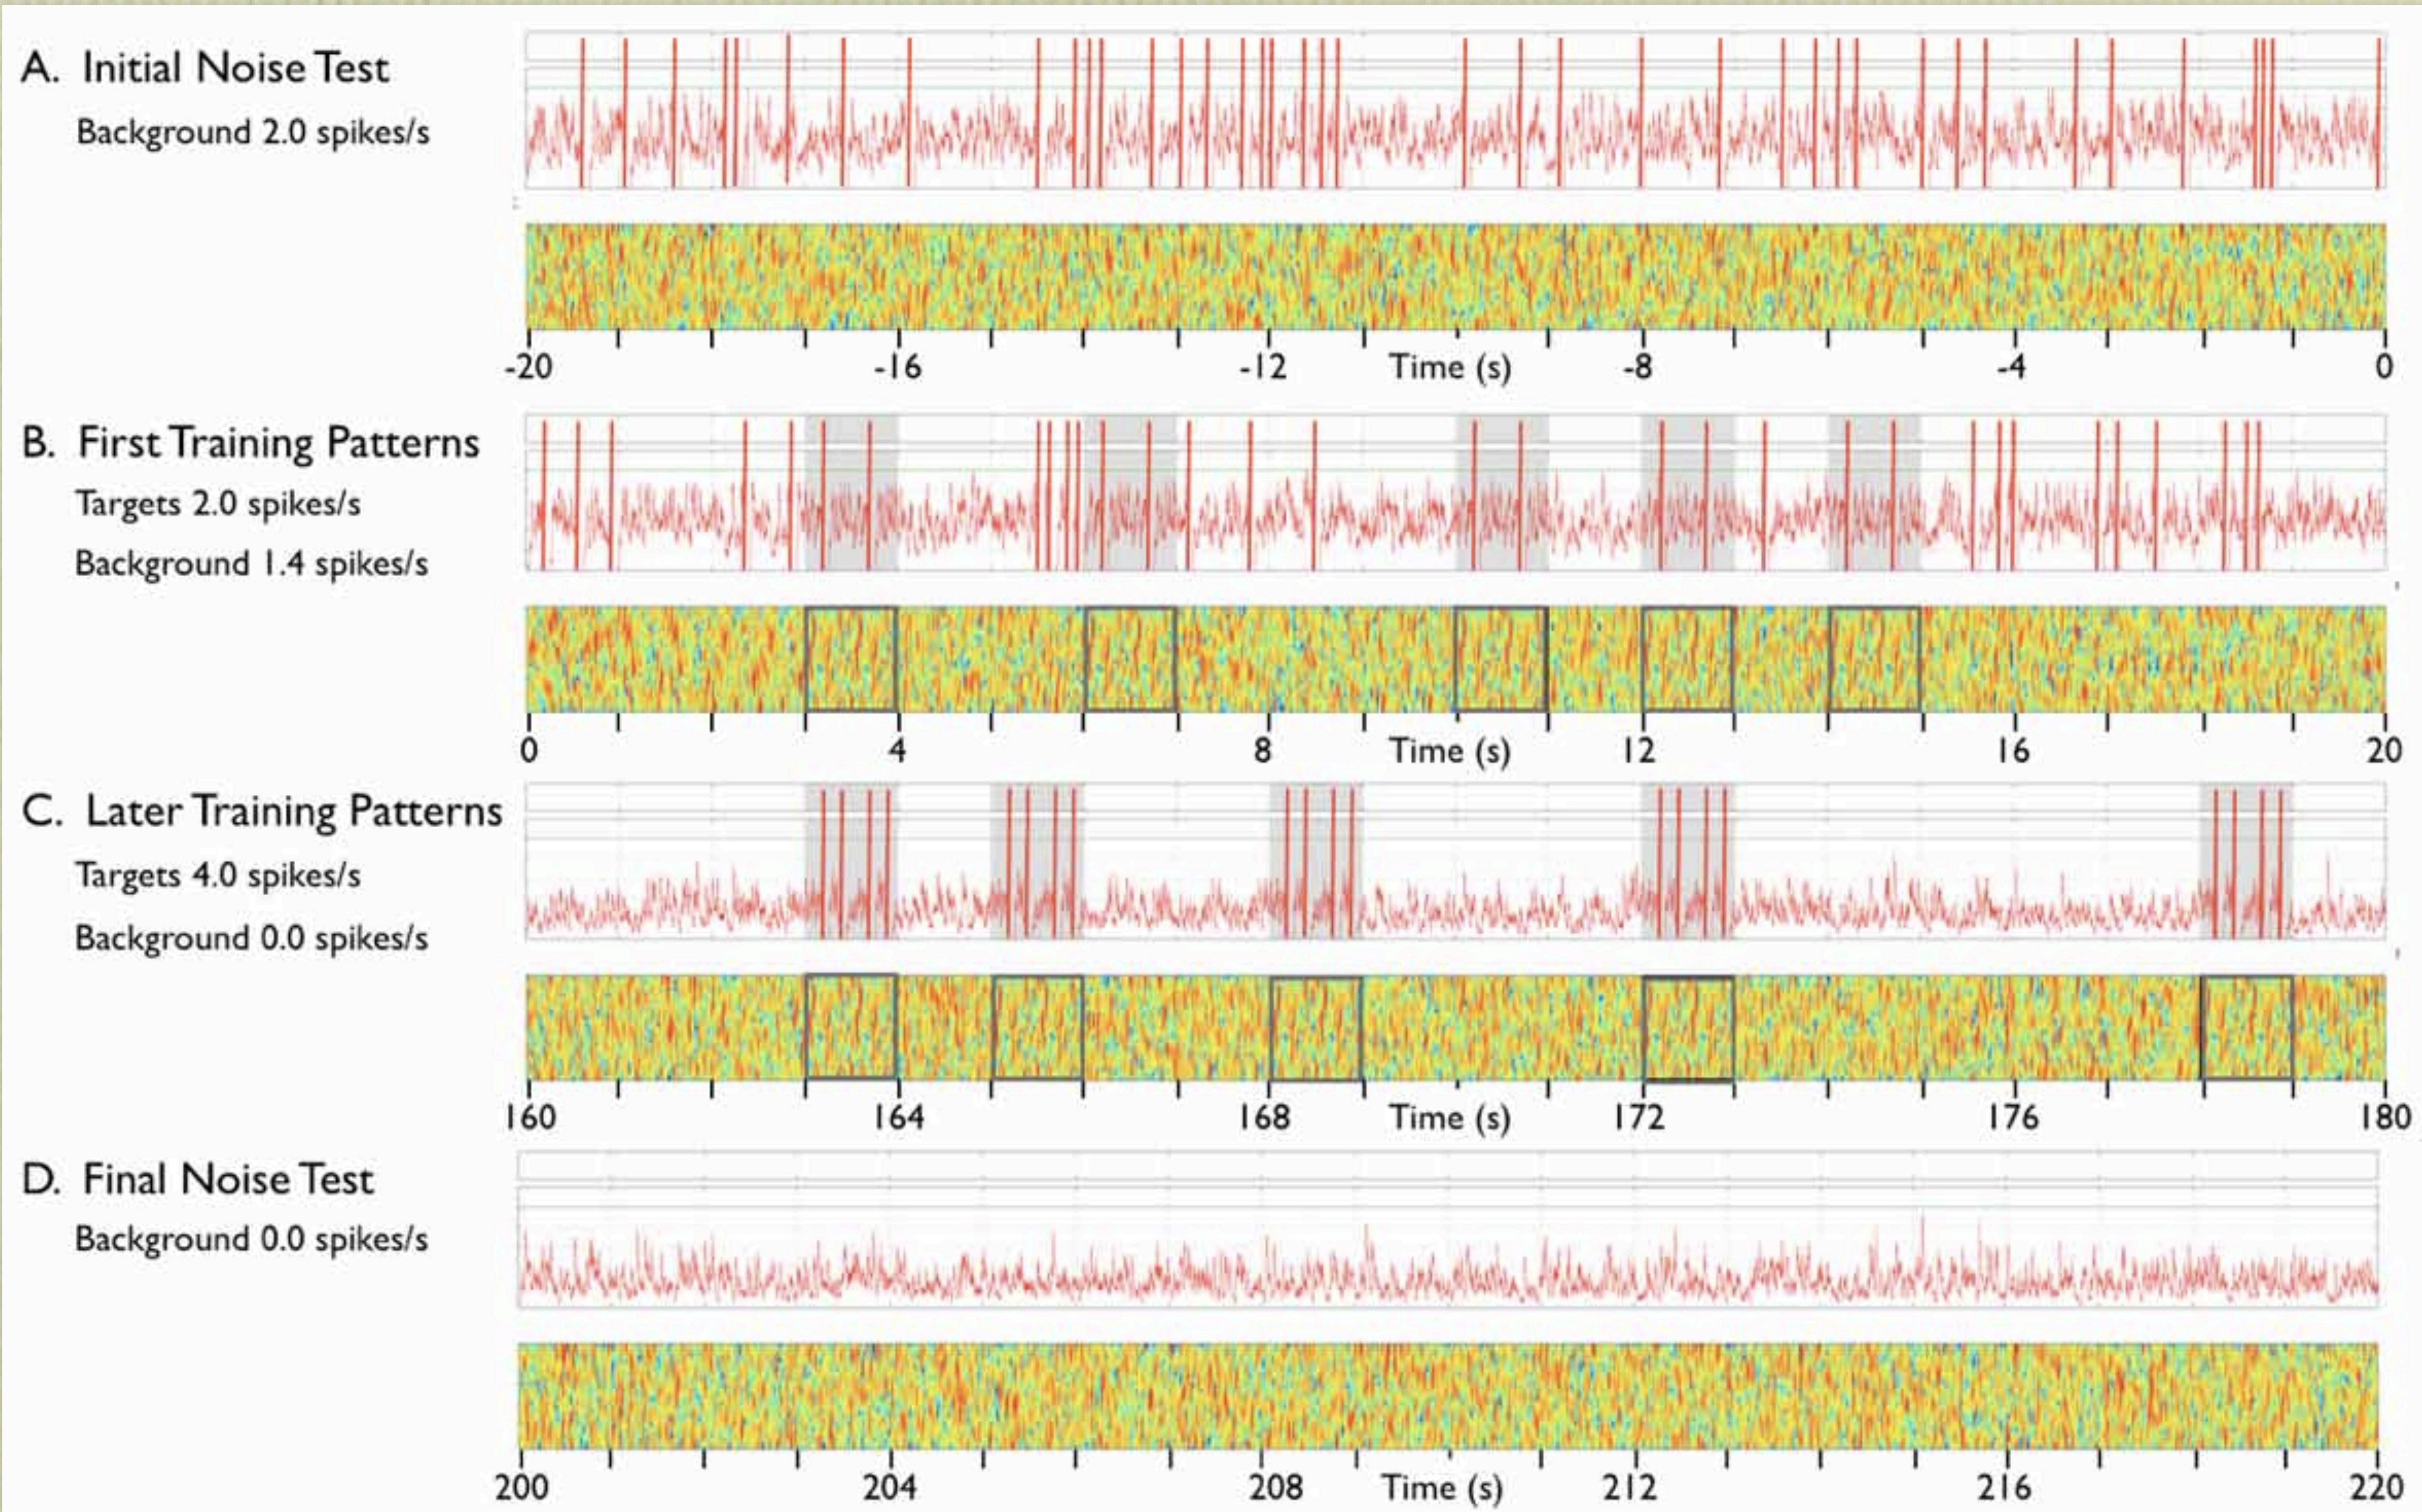

STDP and Spiking Neurons

- "Intelligent" sensory processing with small numbers of neurons connected to spiking sensory devices
  - 2 layers of neurons connected to a spiking retina
    - Unsupervised learning of car counting
  - A single neuron connected to a spiking cochlear
    - Unsupervised learning of auditory noise patterns
- What would happen with 16 billion neurons and multiple sensory inputs?
- Can we implement the learning in hardware?

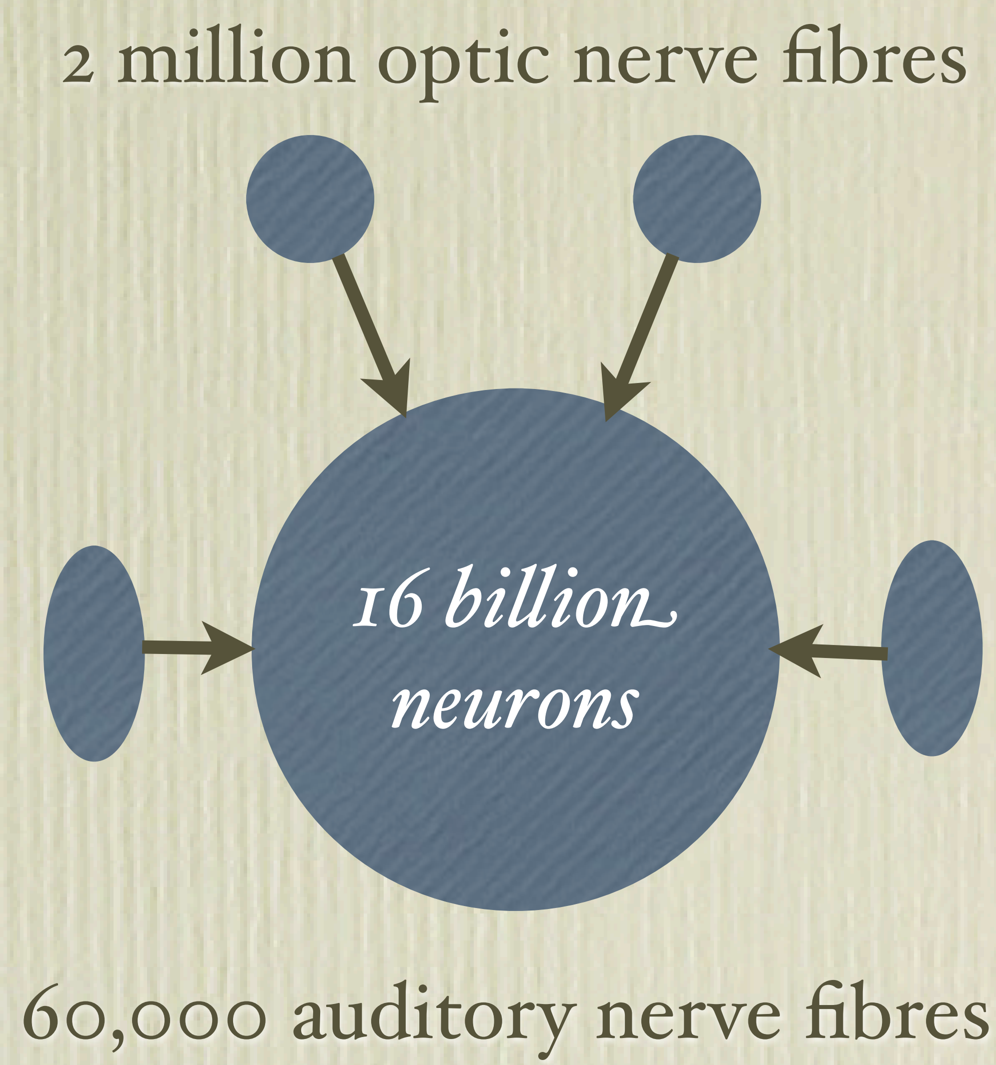

STDP based hardware?

- Use memristor devices to implement STDP in electronics

Circuit Elements With Memory: Memristors, Memcapacitors, and Meminductors

Nanoscale devices, that store information without need for a power source, can be used for non-volatile memory, and promise to allow simulation of learning, adaptation and spontaneous behavior.

By Massimiliano Di Ventra, Yuriy V. Pershin, and Leon O. Chua, *Fellow IEEE*

Memristance can explain Spike-Time-Dependent-Plasticity in Neural Synapses

Bernabé Linares-Barranco and Teresa Serrano-Gotarredona

Nature Precedings 2009

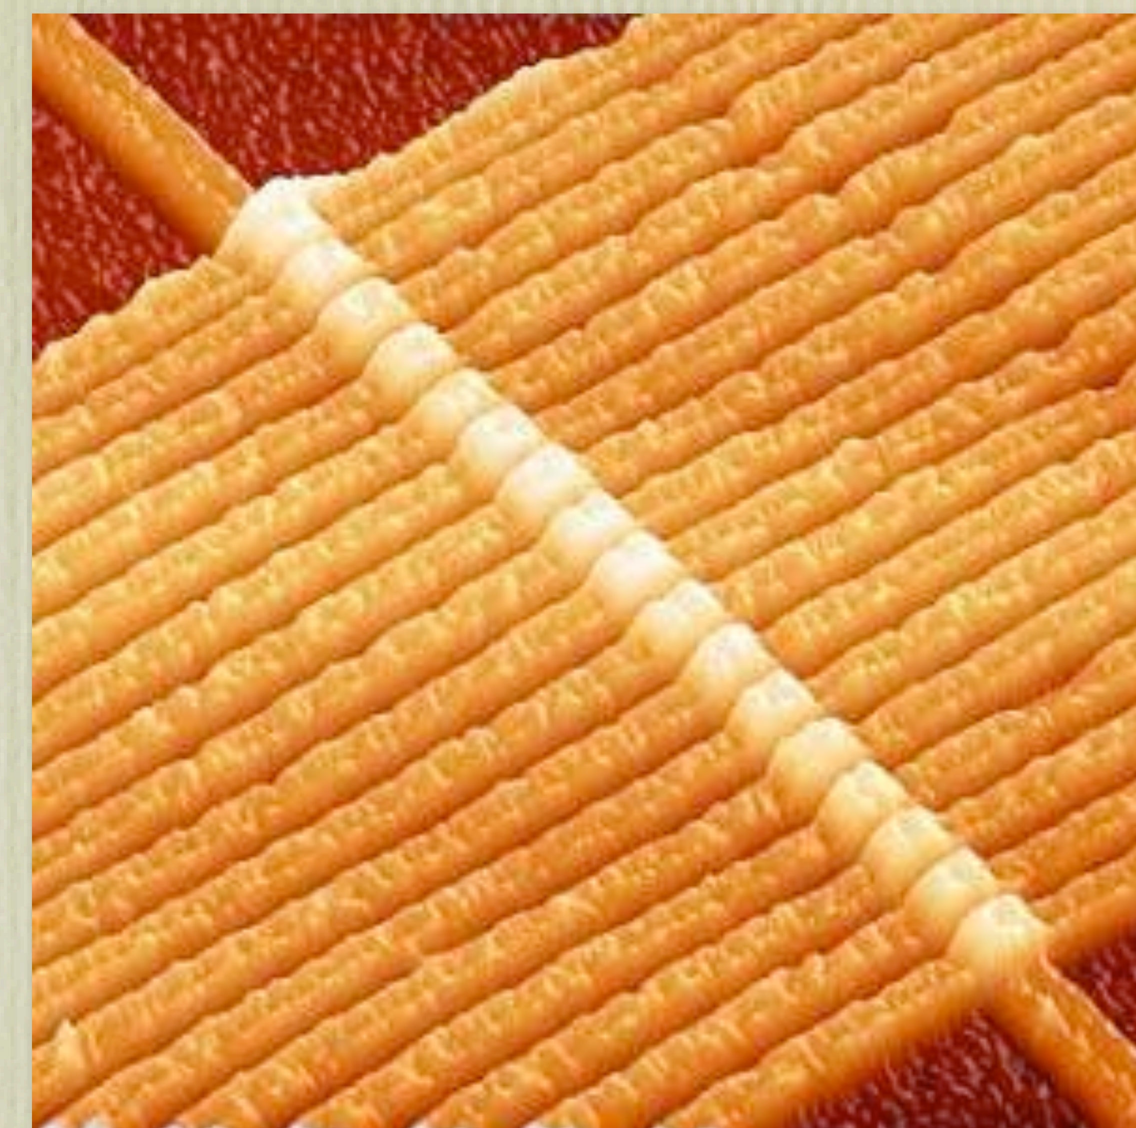

- Hewlett Packard claim to be able to put 100 Gigabits of memristor memory per cm^2

- 10 million synapses on a chip!

Memristor Based STDP

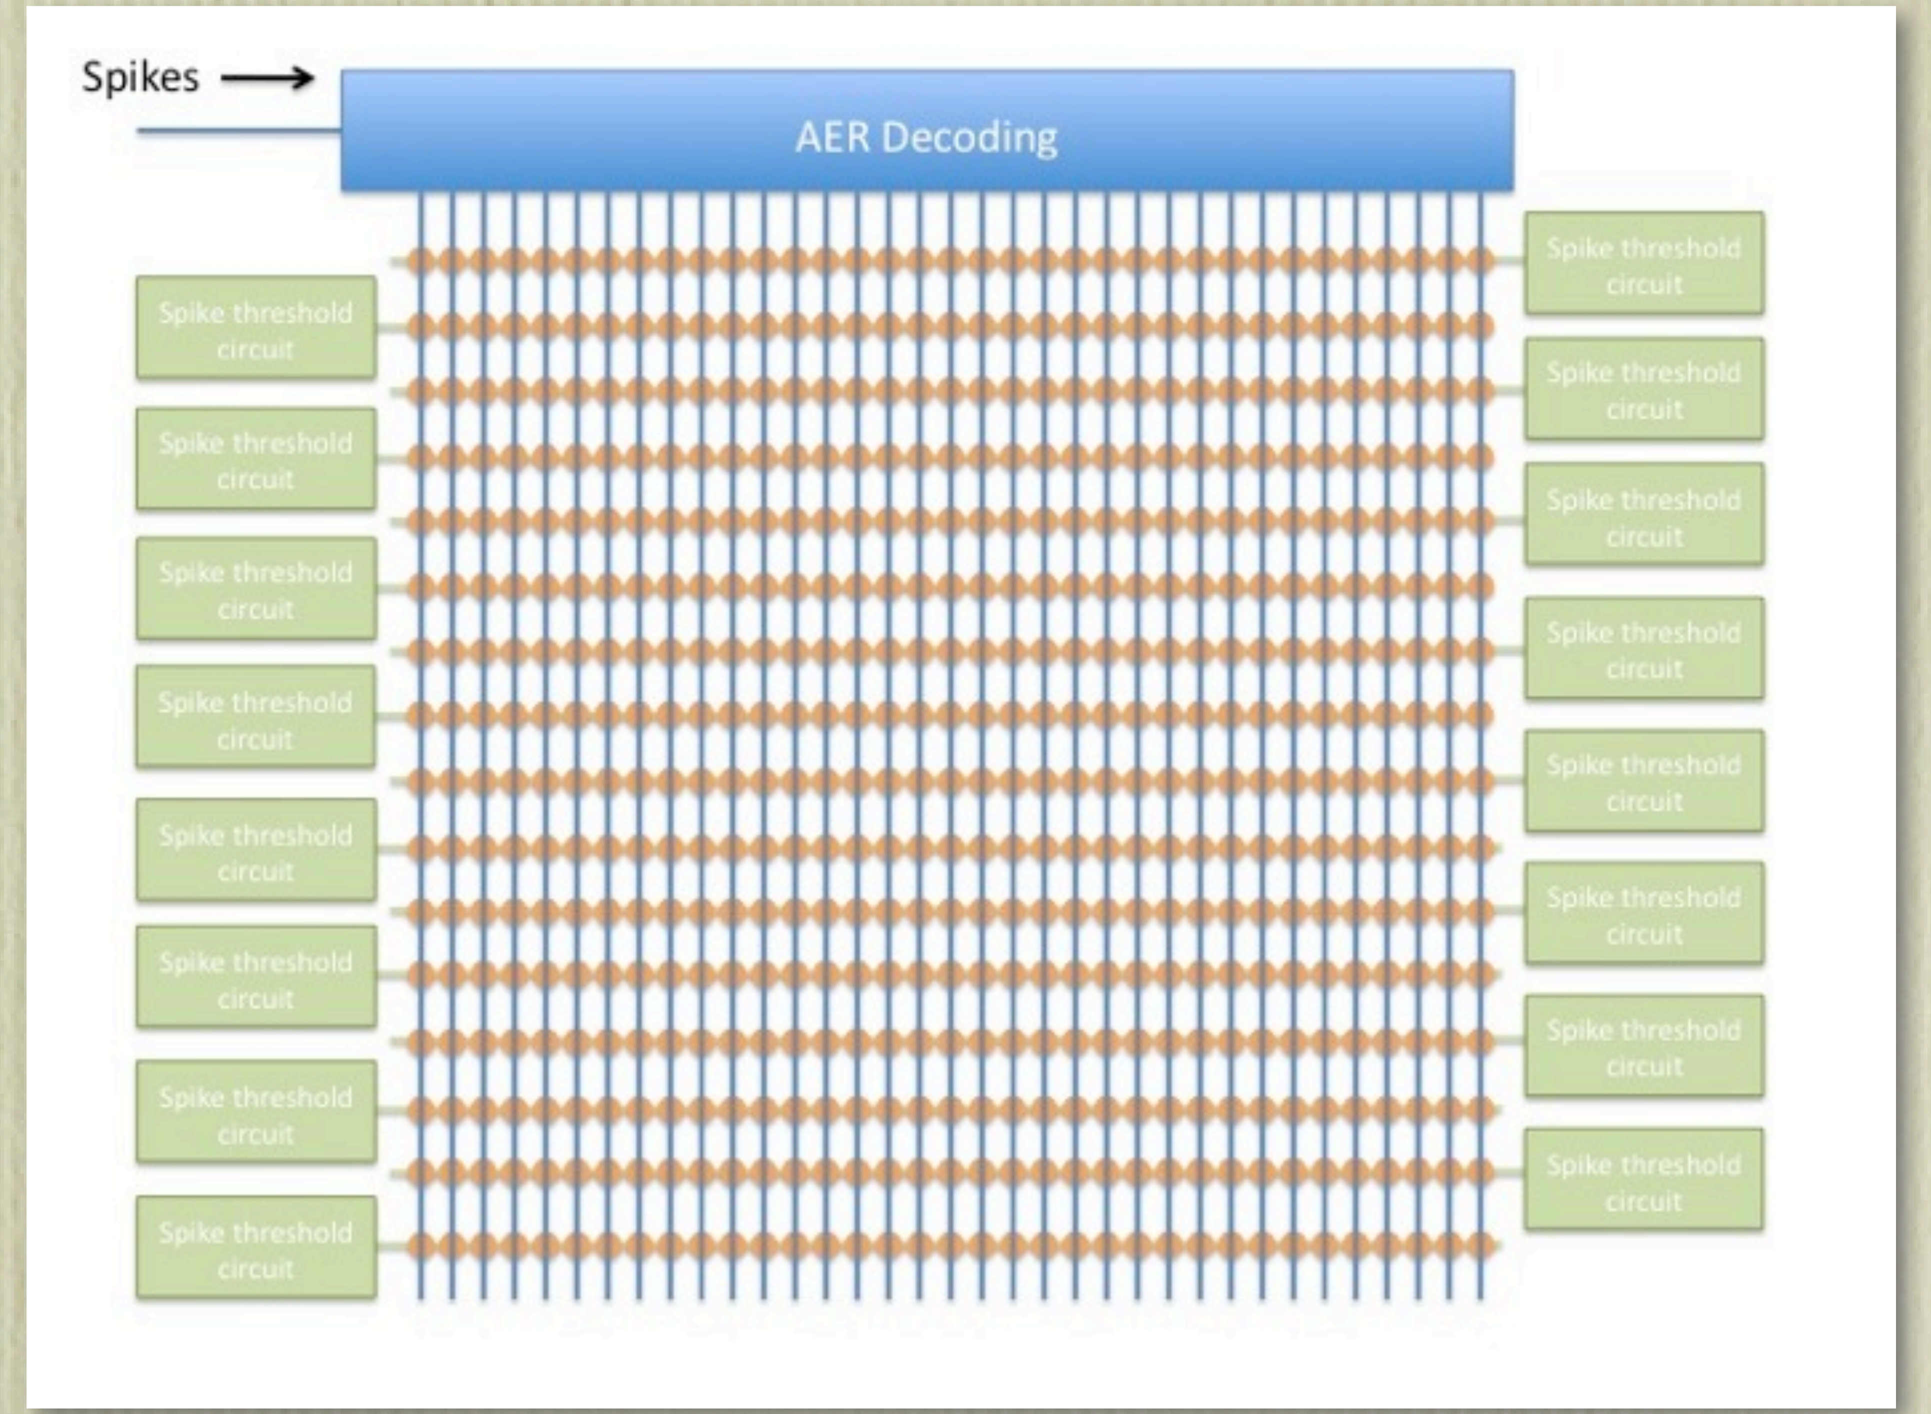

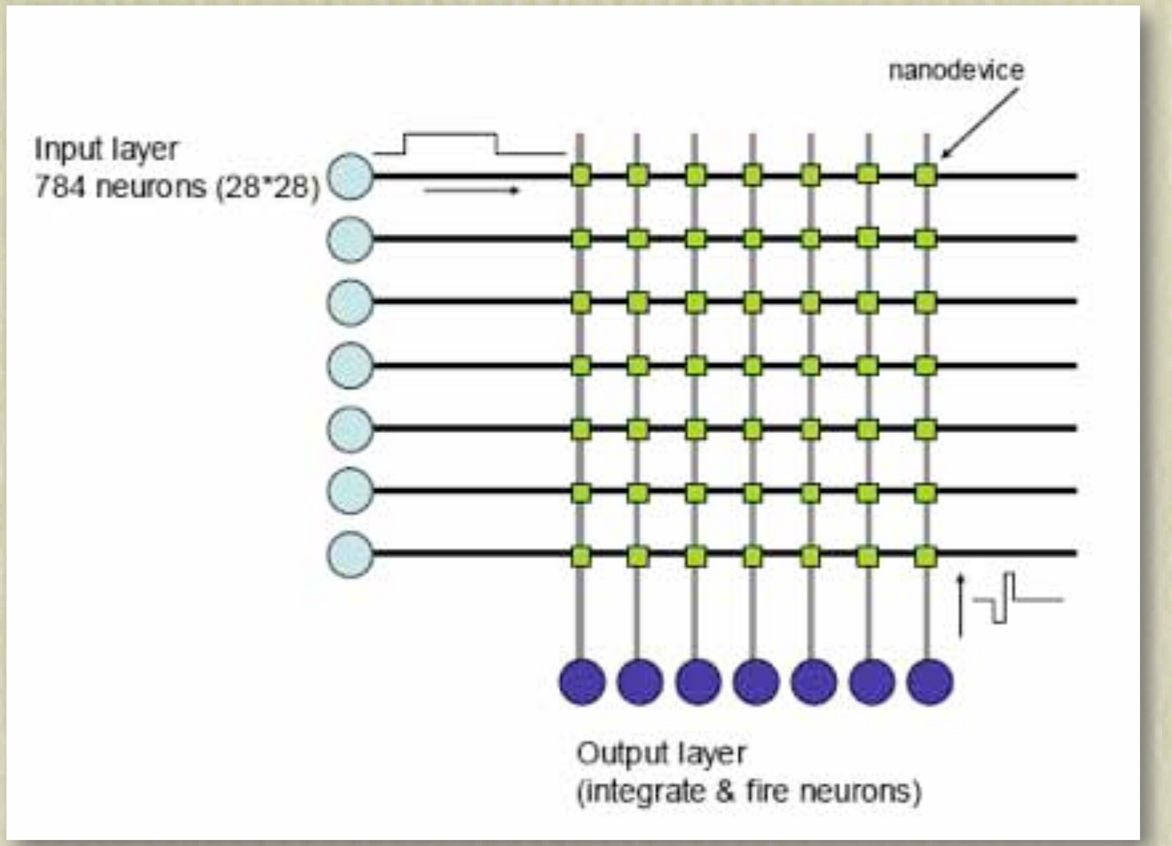

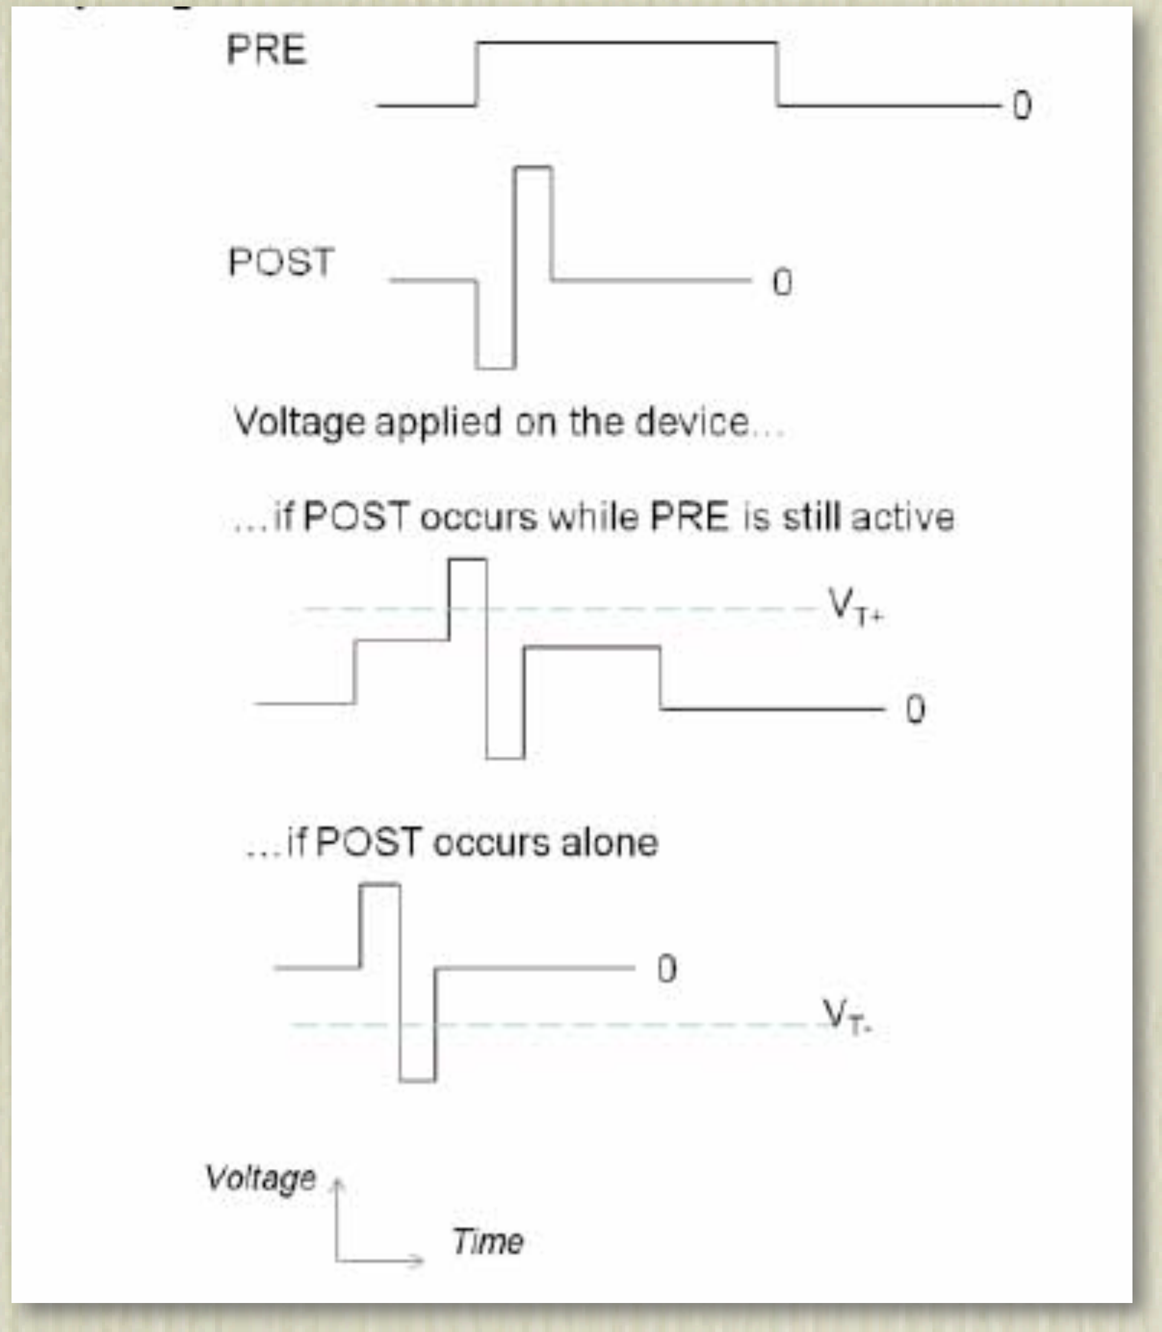

- When the neuron fires a spike
  - All synapses are depressed
  - Except those active just before

Memristive Technologies

- Nanoparticle-Organic Memory Transistor (NOMFET)
- Phase Change Memories (PCM, PRAM or PCRAM)
- Conductive Bridging RAM (CBRAM)
- Resistive RAM (RRAM or ReRAM)

Pavlov Circuit using a NOMFET

Bichler et al (2012) “Pavlov’s dog associative learning demonstrated on synaptic-like organic transistors”

Neural Computation (in press)

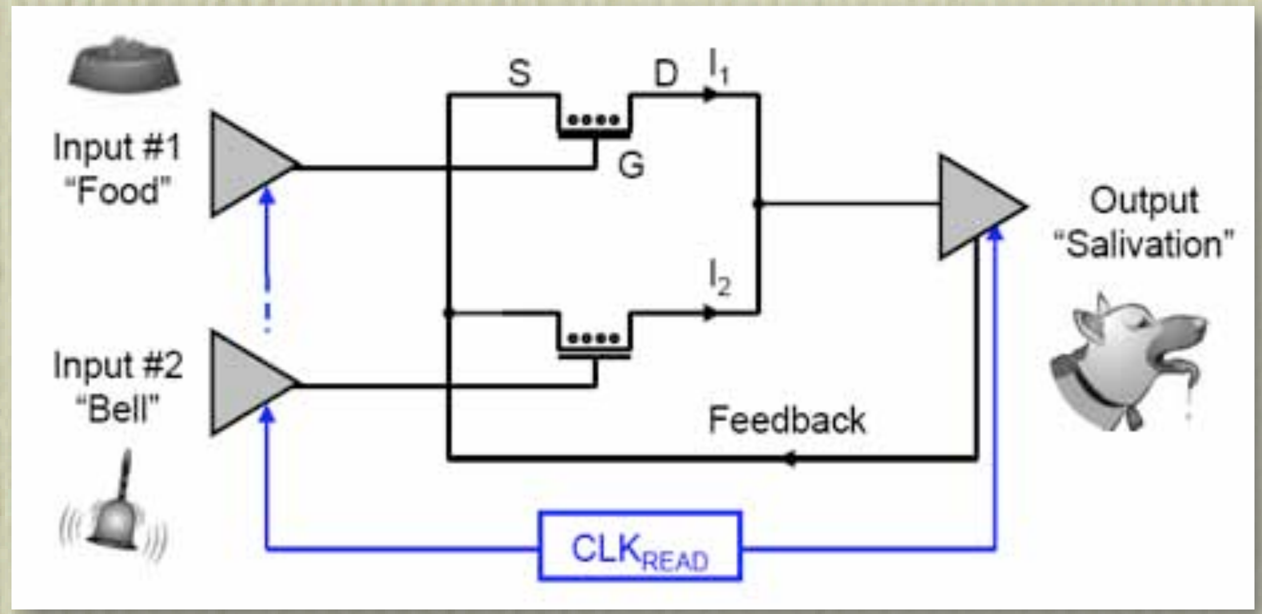

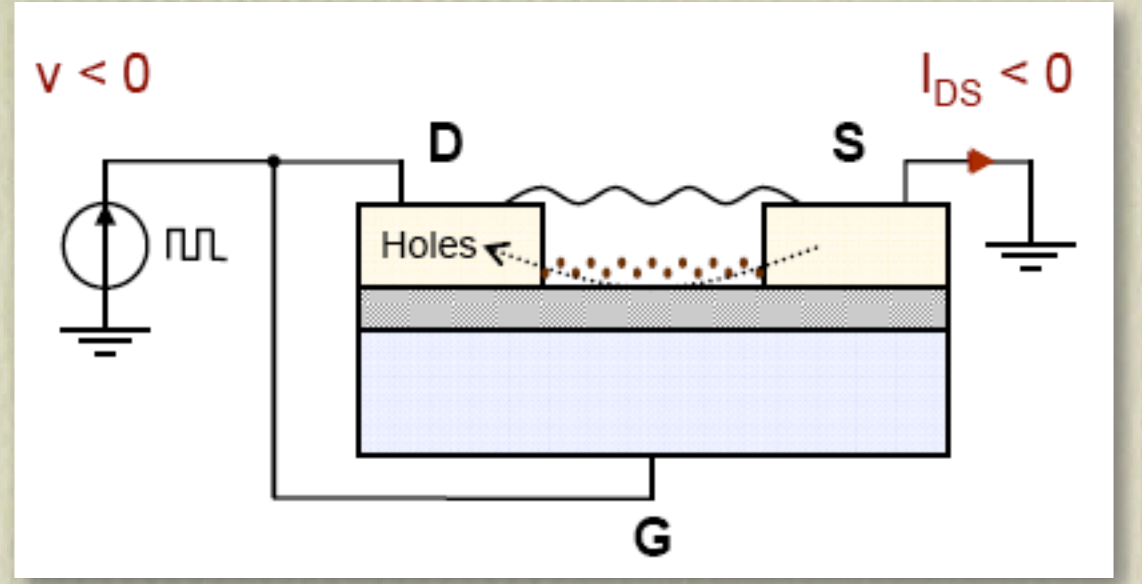

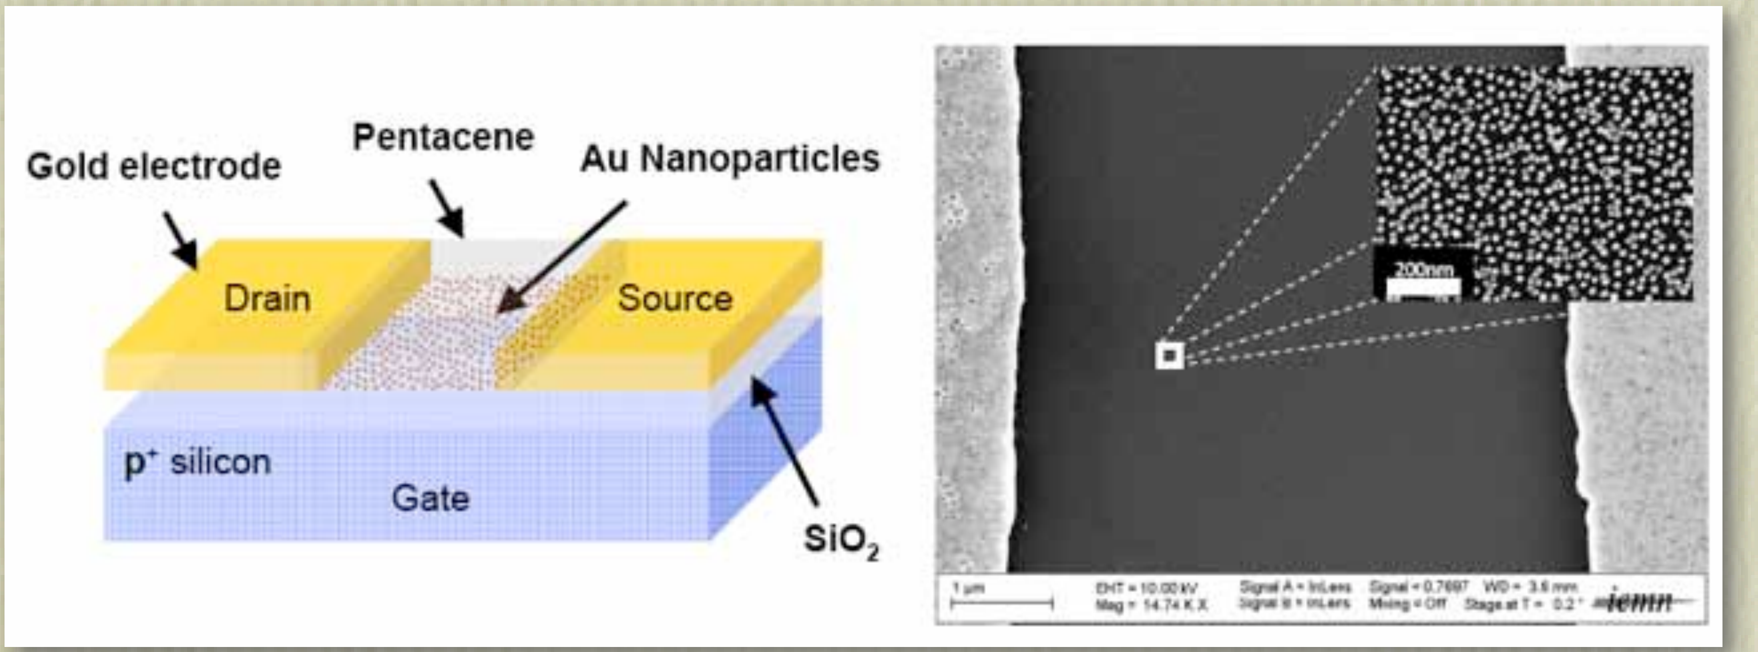

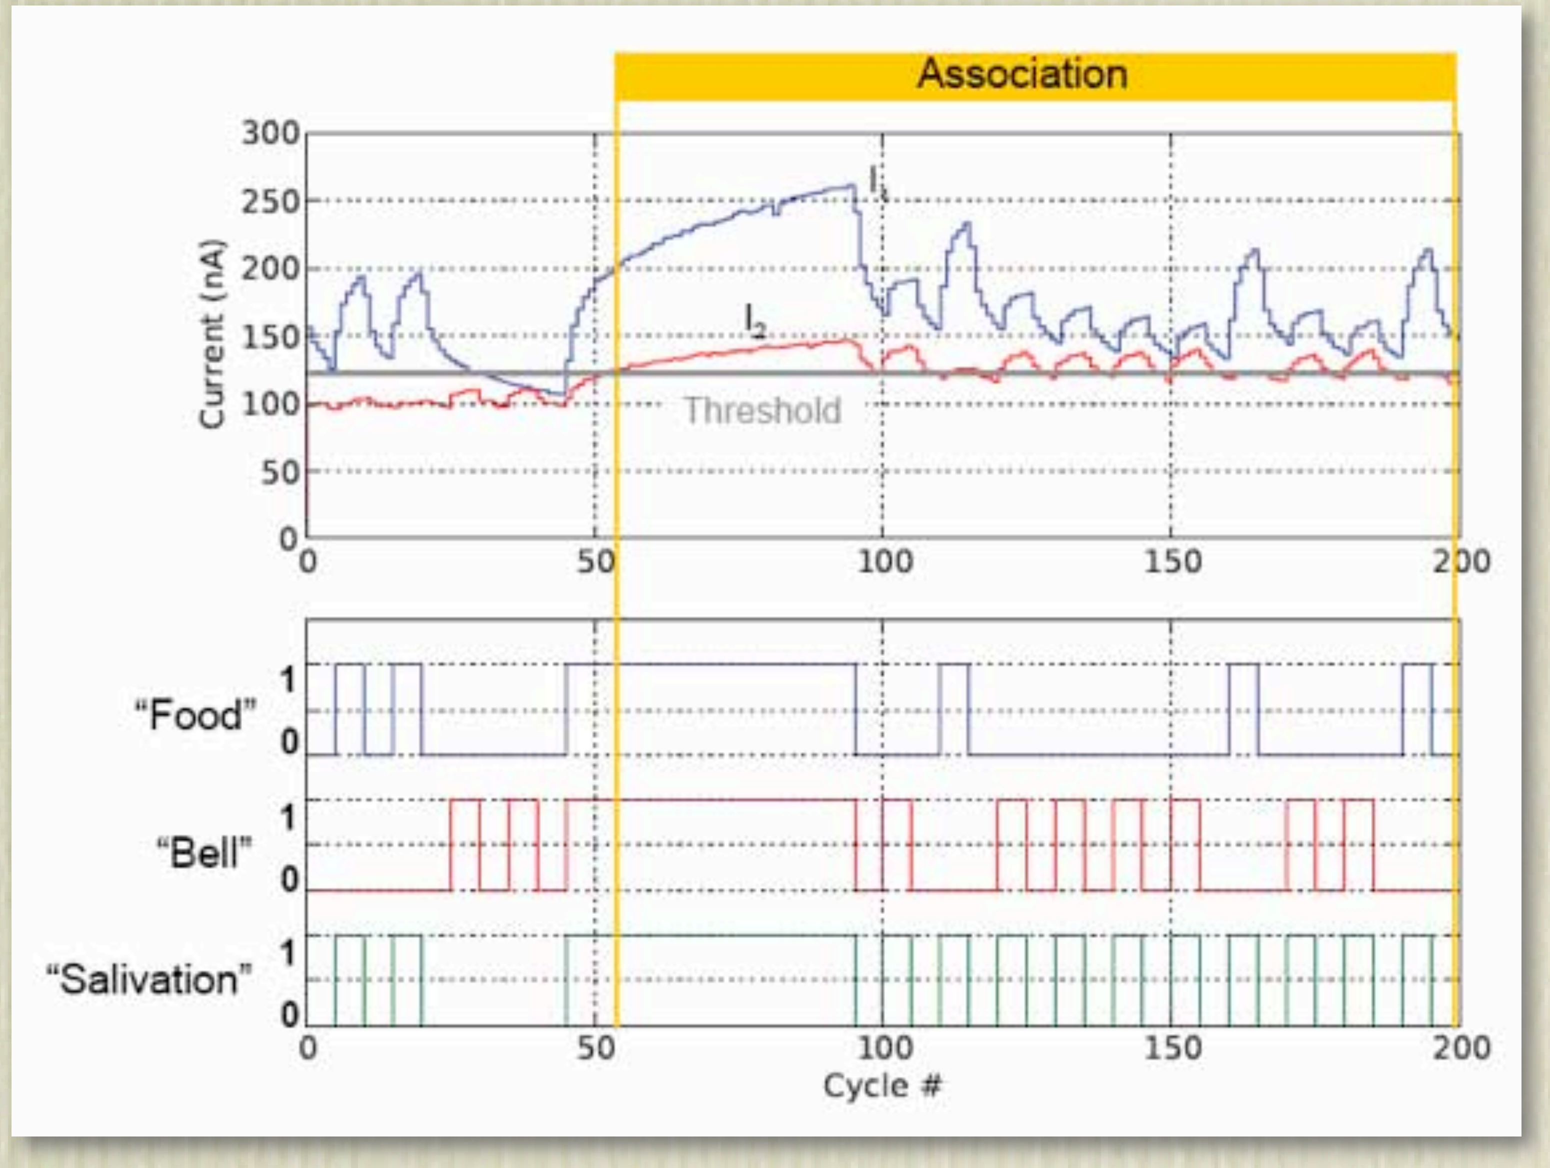

Memristor Based Hardware

- STDP based hardware
  - Remarkable ability to learn to detect repeating patterns within the input
  - Automatically finds the start - shortest possible latency
  - Applications
    - vision (asynchronous or frame-based...)
    - audition (speech, music....)
    - somatosensory processing
- Beyond biology?
  - Synapse sizes - 10 nm - density 10^{11} per cm^2
    - 3D stacking?
  - Conduction velocity - $1 ms^{-1}$ for biology, unlimited for electronics
  - Number of synapses
    - Biology - 10^3-10^4
    - Electronics - ?

Final Comments

- STDP and Spikes
  - A very powerful way to process information
  - Naturally intelligent
- Implementation
  - Software
  - FPGA
  - ASIC
  - Memristors?
- Why use Nanotechnology?
  - Size
  - Power
  - Fault resistance
  - Complete elimination of the Von Neumann bottleneck

Credits

Auditory Noise Learning

Trevor
Agus

Daniel
Pressnitzer

STDP Modelling

Rufin
VanRullen

Rudy
Guyonneau

Tim
Masquelier

Memristor Learning Architectures

Olivier
Bichler

Damien
Querlioz

Christian
Gamrat

Jean-Philippe
Bourgoin