Spike-based Image Processing:

Can we reproduce biology in hardware?

Simon Thorpe

Centre de Recherche Cerveau & Cognition

& SpikeNet Technology SARL,

Toulouse, France

Plan

- Ultra-Rapid Visual Processing
  - Recognition and localisation of complex visual stimuli in 100 to 150 ms
- Spike based processing
  - Using a wave of spikes to process information
  - Selective responses with just one spike per neuron
- Learning mechanisms
  - STDP (Spike-Time Dependent Plasticity) makes neurons selective to input patterns that occur repeatedly
  - A few tens of presentations are enough for selectivity to develop
- Can we build hardware systems using the same principles?
  - Spiking retinas
  - Spiking cochleas
  - Memristor devices

Ultra Rapid Scene Categorisation

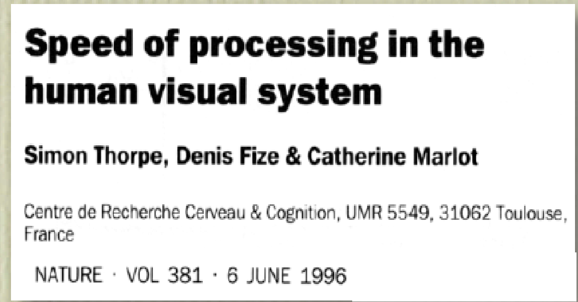
Speed of processing in the human visual system

Simon Thorpe, Denis Fize & Catherine Marlot

Centre de Recherche Cerveau & Cognition, UMR 5549, 31062 Toulouse, France

NATURE · VOL 381 · 6 JUNE 1996

Behavioural Reaction Times

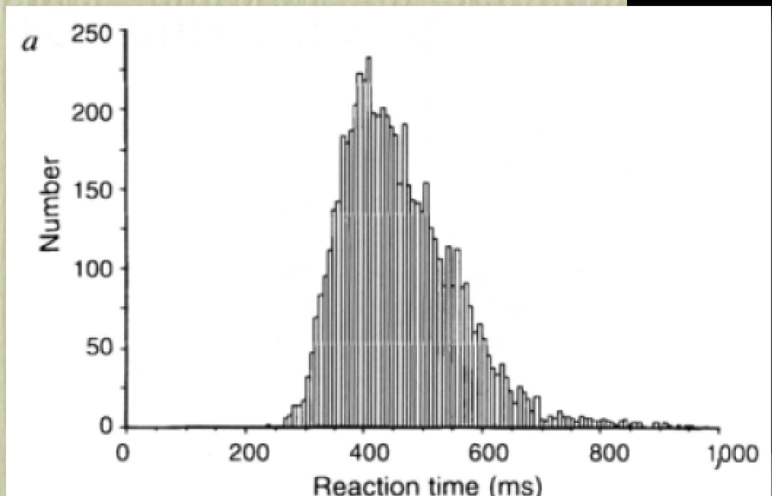

Event Related Potentials

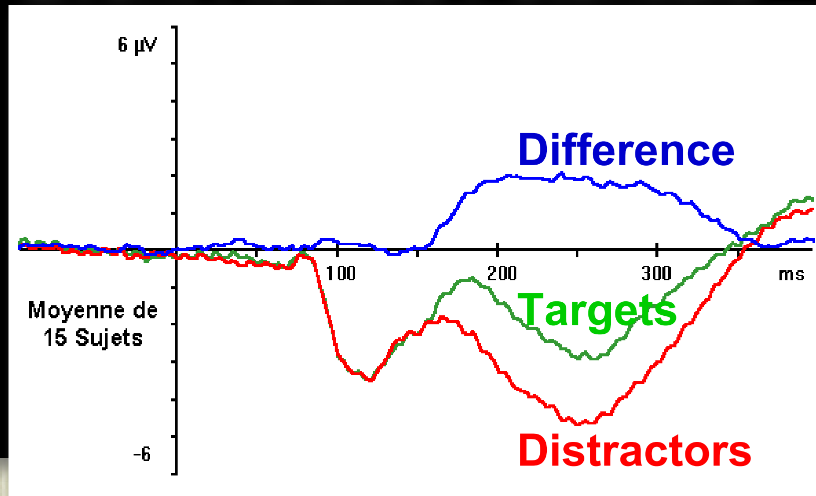

Scene Processing in 150 ms

Temporal Constraints

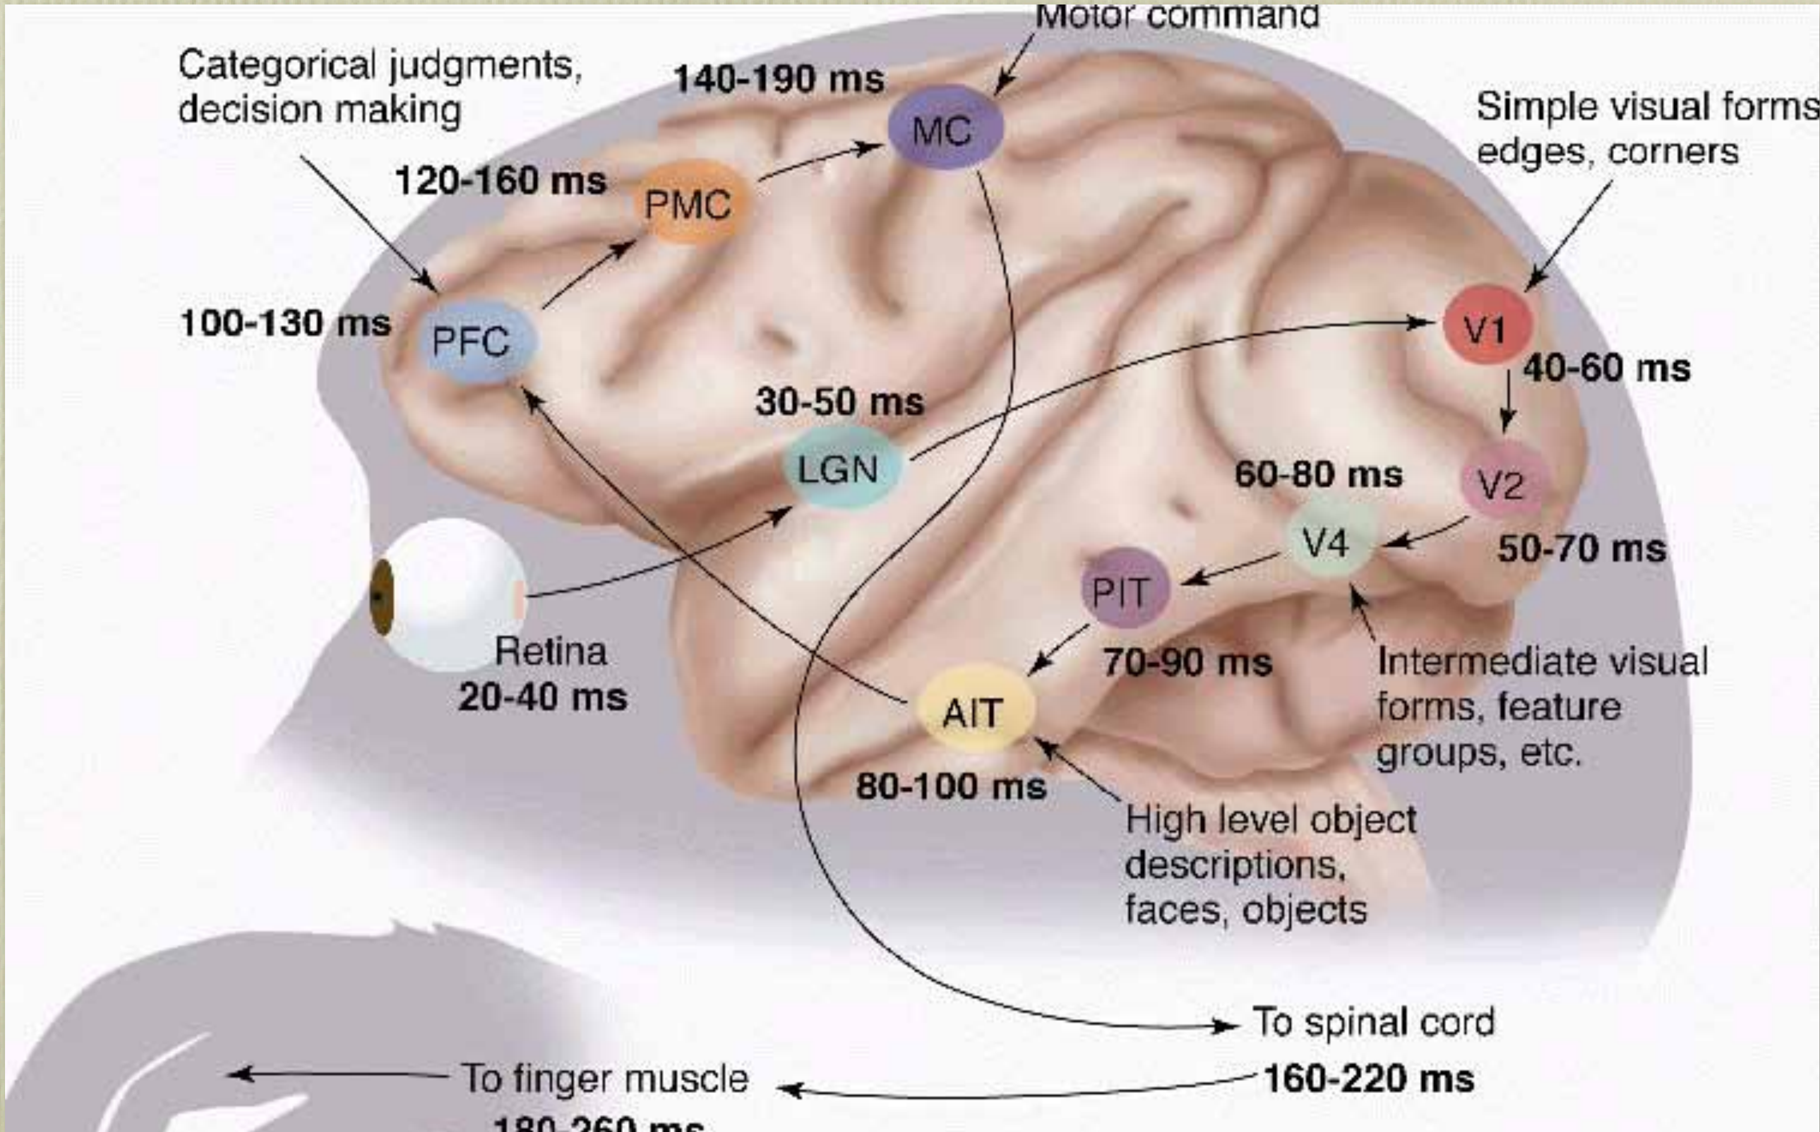

Seeking Categories in the Brain

Simon J. Thorpe and Michèle Fabre-Thorpe

SCIENCE VOL 291 12 JANUARY 2001

Even faster processing

Fast saccades toward faces: Face detection in just 100 ms

Sébastien M. Crouzet Holle Kirchner Simon J. Thorpe

Journal of Vision (2010) 10(4):16, 1–17

Saccadic Choice Task

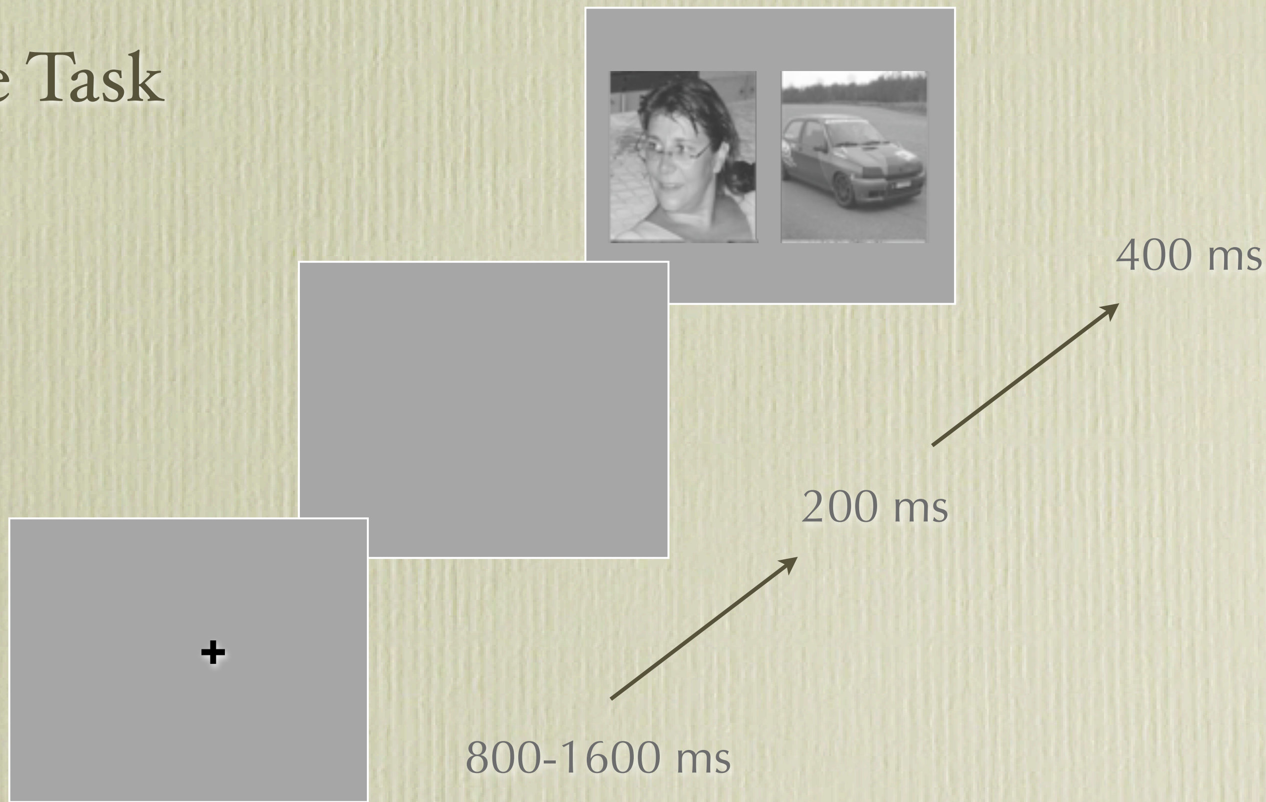

+

+


116

+

+


RAIL MOTOR Nº 34

+

+


FOCUS

+

+

+


116

+


1187

+

Face Saccadic Choice Task

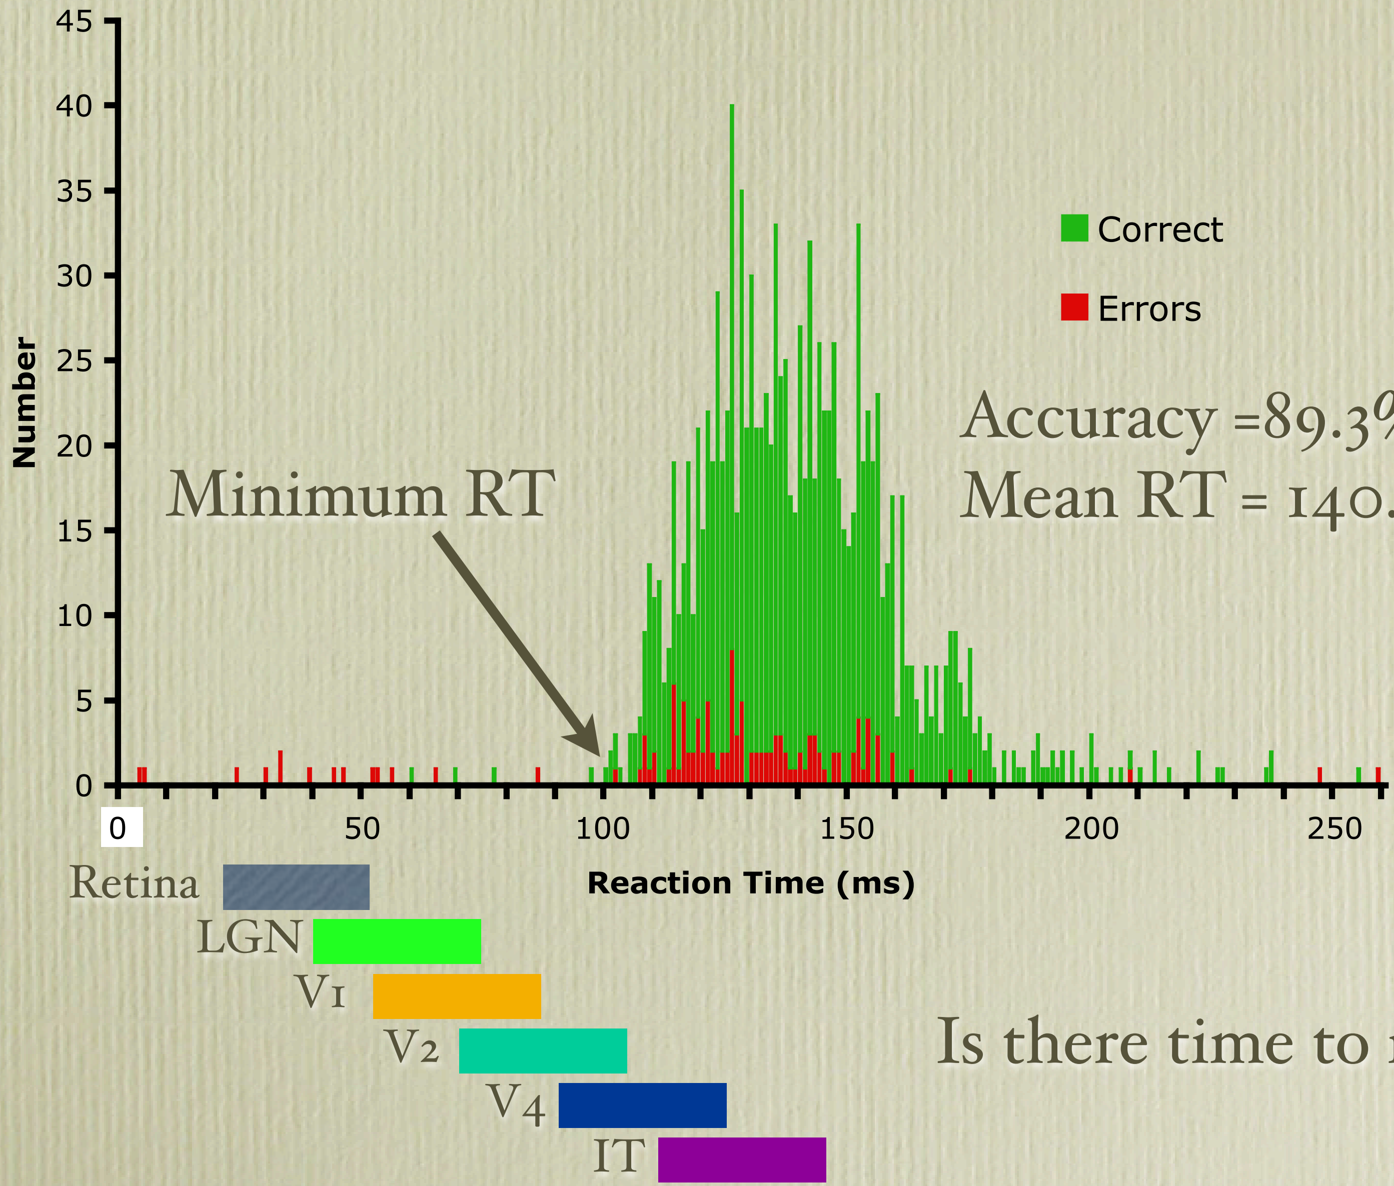

Accuracy =89.3%
Mean RT = 140.7 ms

Is there time to reach IT??

How accurate are saccades to faces?

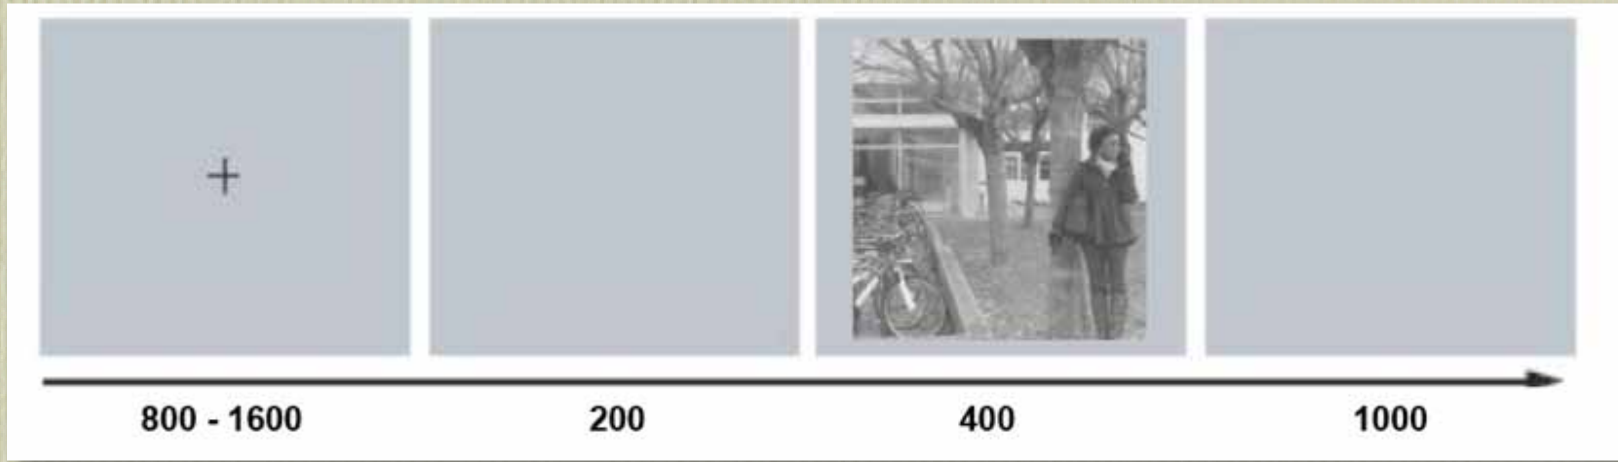

- 18 x 18° image
- 1° face
- 16 locations
  - 8 directions and 2 eccentricities (3.5° and 7°)


WIN
450
520
660
820
980
530
620
640
640
660

Results

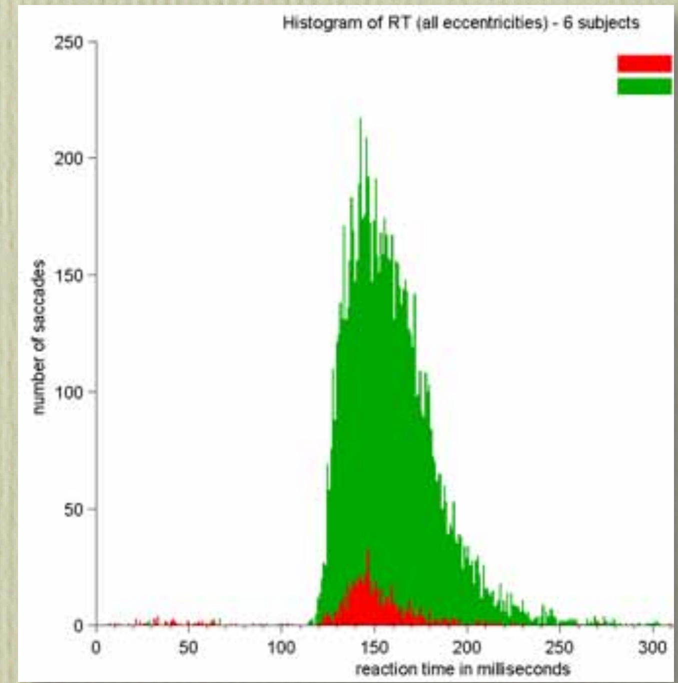

- 92.4% correct
- Mean RT 158 ms

Remarkable overall accuracy

All Saccades

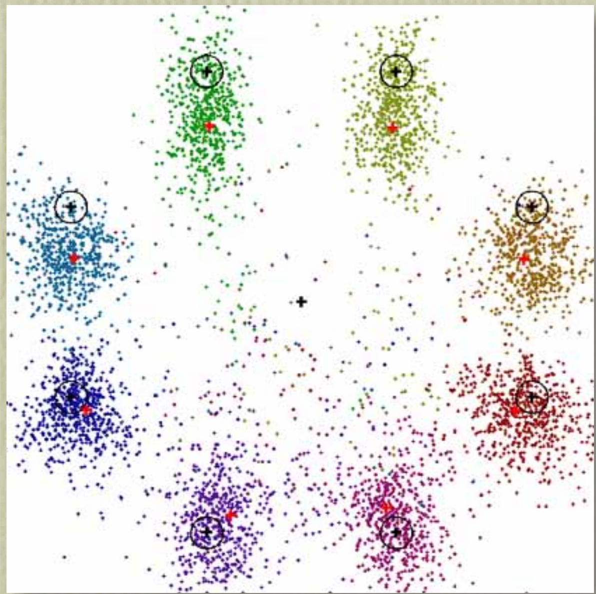

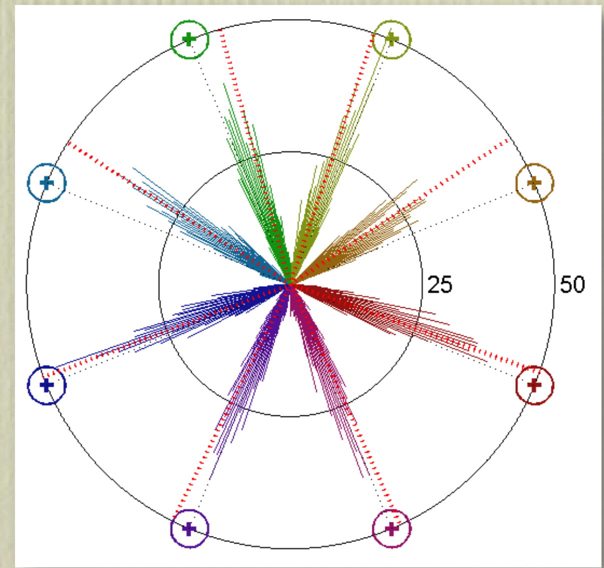

Even for the fastest saccades

Saccades 100-150ms

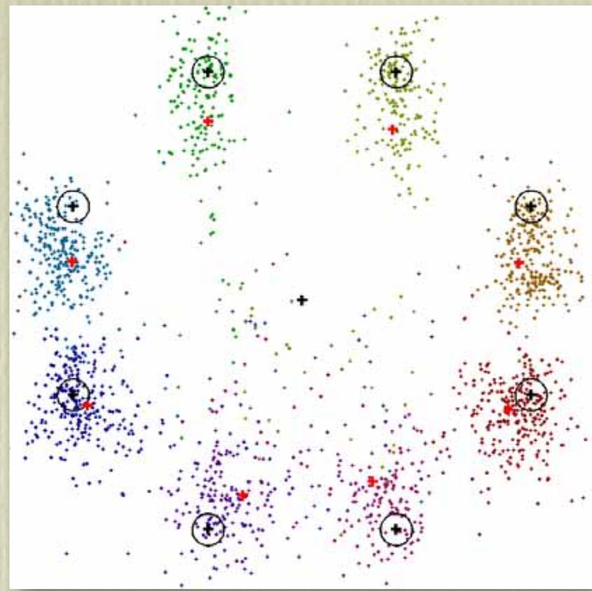

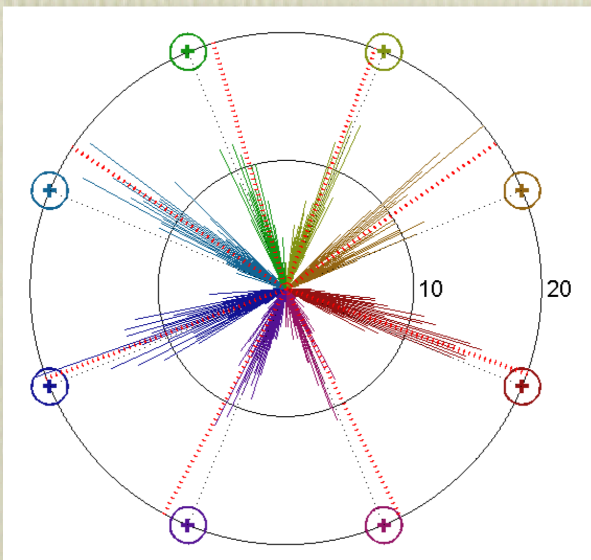

Finding Faces in Cluttered Scenes

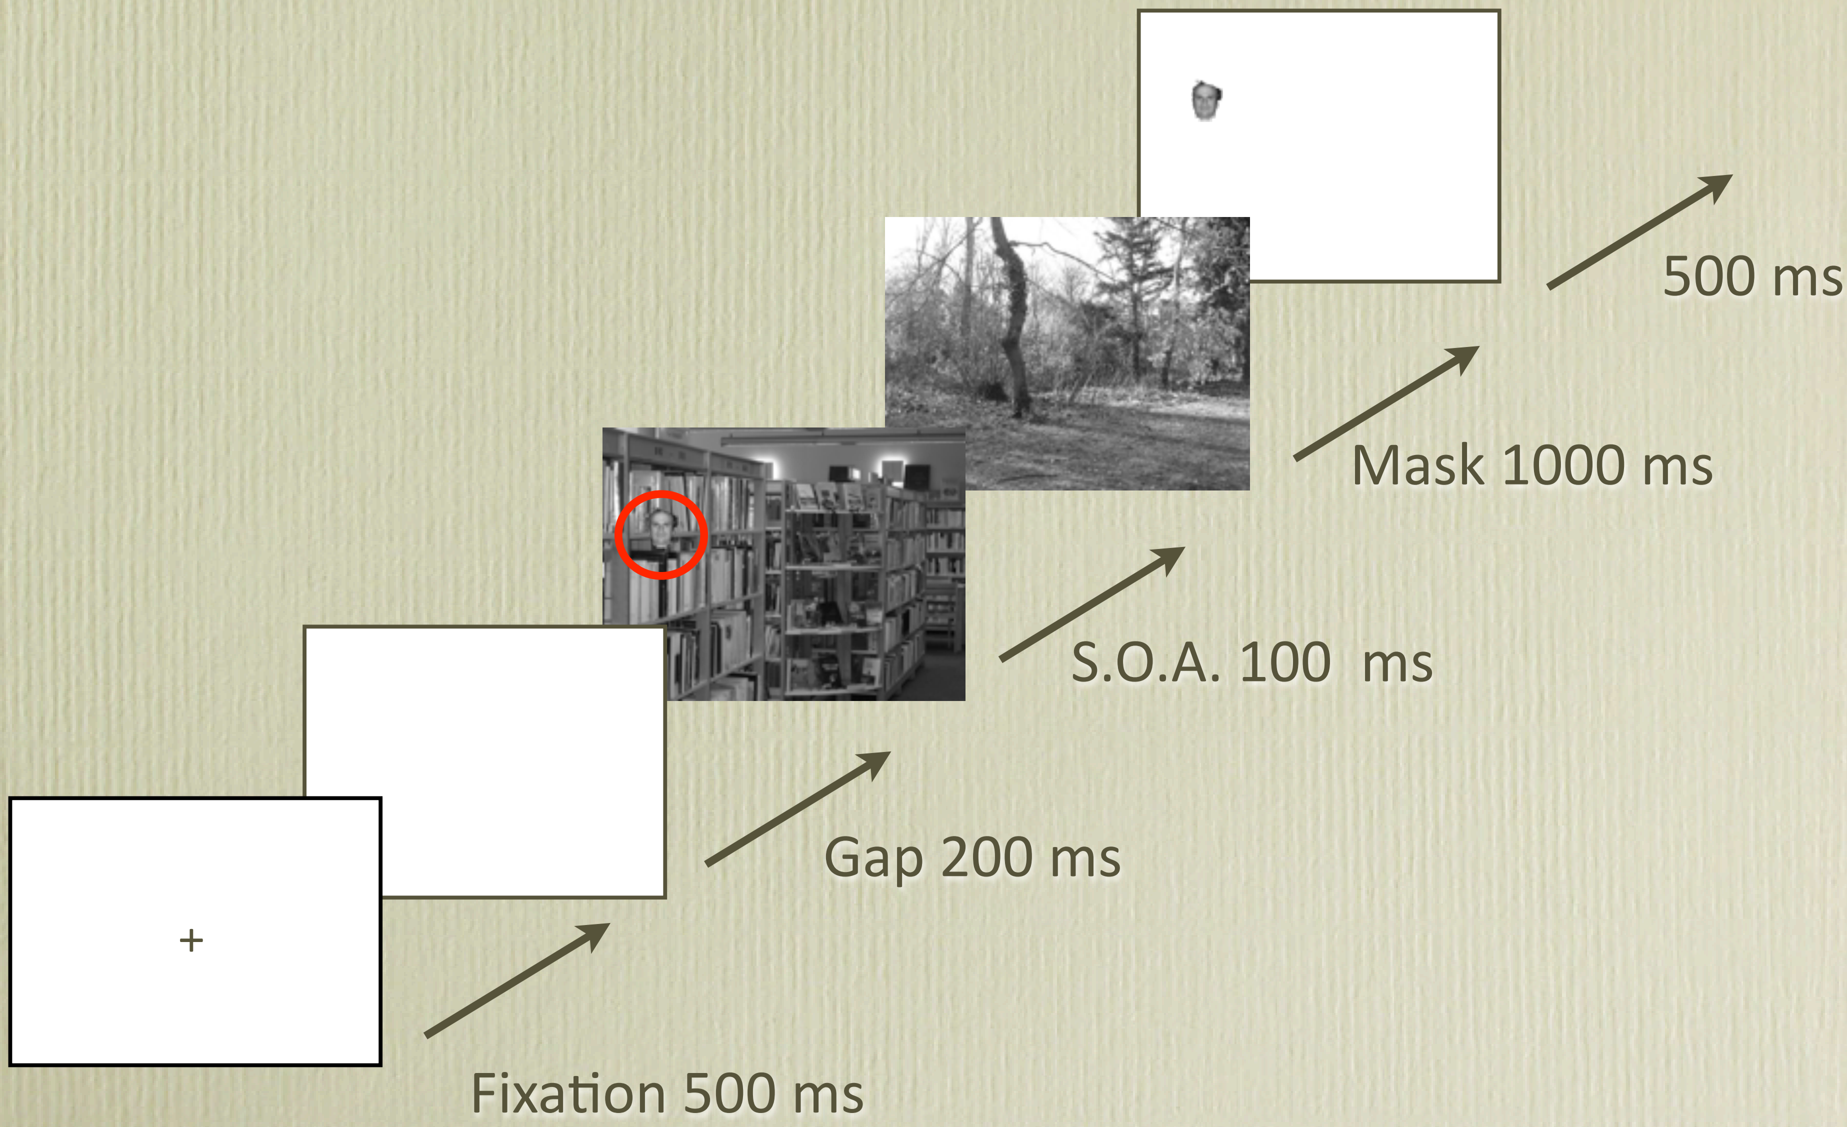


ARMY
NAVY
NIPITS
NIPITS







Biological Image Processing

- Ultra Rapid
  - Go/Nogo manual responses
    - from 270 ms in man
    - from 180 ms in monkeys
  - Saccadic choice task
    - Responses to animals in 120-130 ms
    - Responses to faces in 100-110 ms
  - Accurate face localisation in cluttered scenes
    - 1° face at 7-10° eccentricity
- Biological hardware is slow
    - < 1KHz clock
    - 1-2 $m.s^{-1}$ conduction velocity

An electronic implementation could be orders of magnitude more powerful!

Spike-based Processing

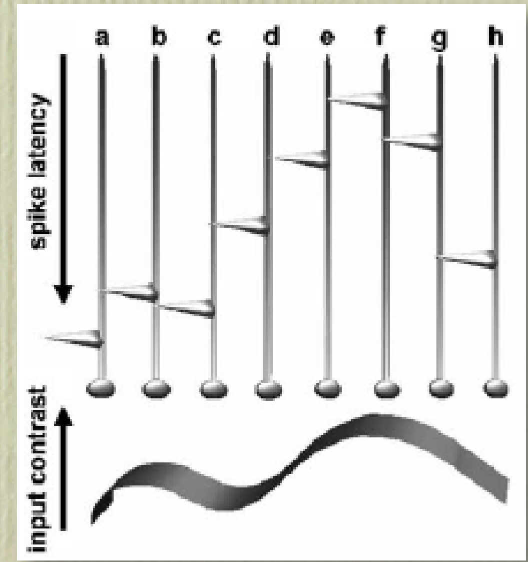

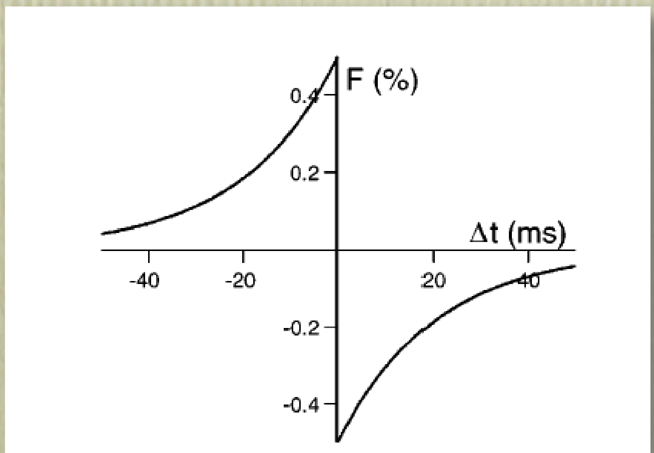

Song, Miller & Abbott, 2000

- Processing with a wave of spikes
  - The most strongly activated cells fire first
  - Information can be encoded in the order of firing
- Spike-Time Dependent Plasticity
  - With repeated presentation, high synaptic weights concentrate on the early firing inputs
  - Allows the development of fast selective responses

The Neuron as an Intensity-Delay Convertor

- Onset latency varies with activation

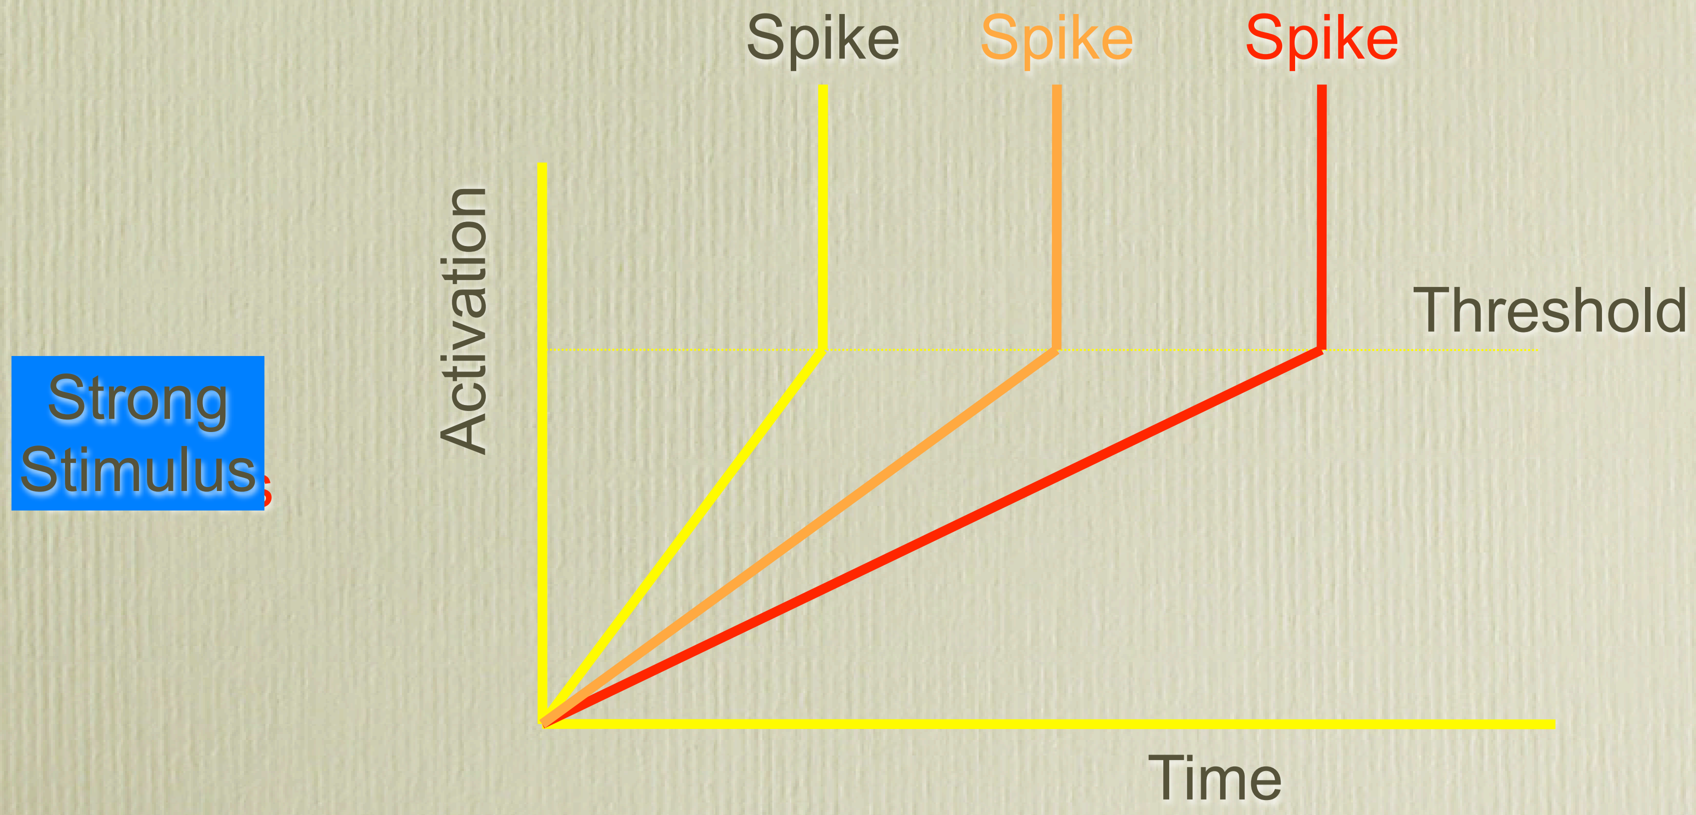

Generating Selectivity

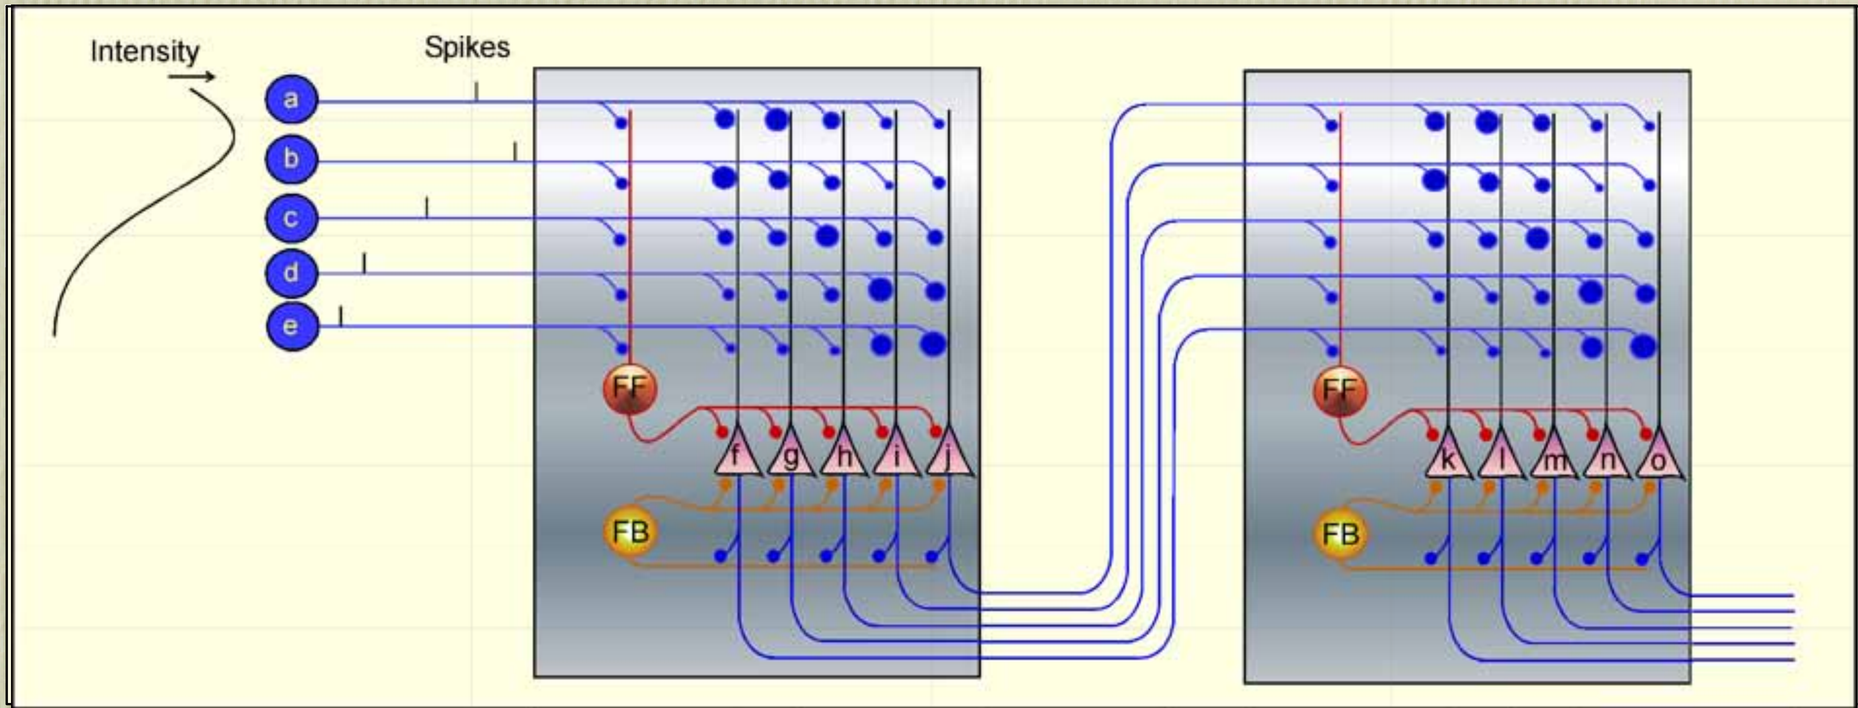

- Feed-forward inhibition
  - desensitisation
  - gives maximum importance to the first spikes
- Feedback inhibition
  - k-Winner take all
  - Controls the number of cells that are allowed to fire

Early Studies

Face identification using one spike per neuron: resistance to image degradations

A. Delorme*, S.J. Thorpe   Neural Networks 14 (2001) 795–803

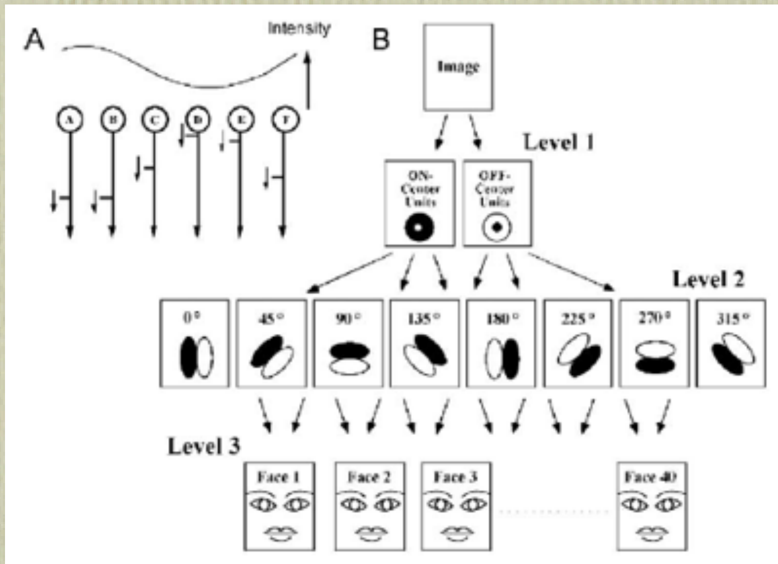

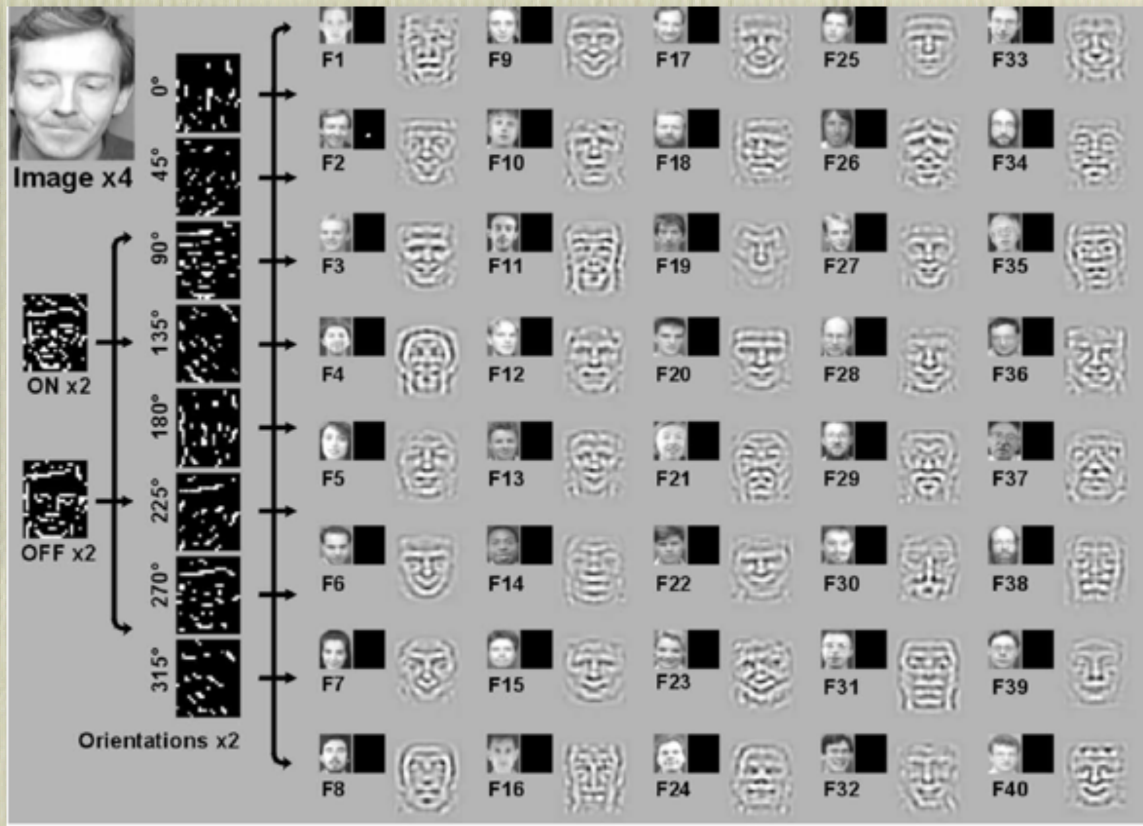

- Face identification directly from the output of oriented filters

Early Studies

Face identification using one spike per neuron: resistance to image degradations

A. Delorme*, S.J. Thorpe

Neural Networks 14 (2001) 795–803

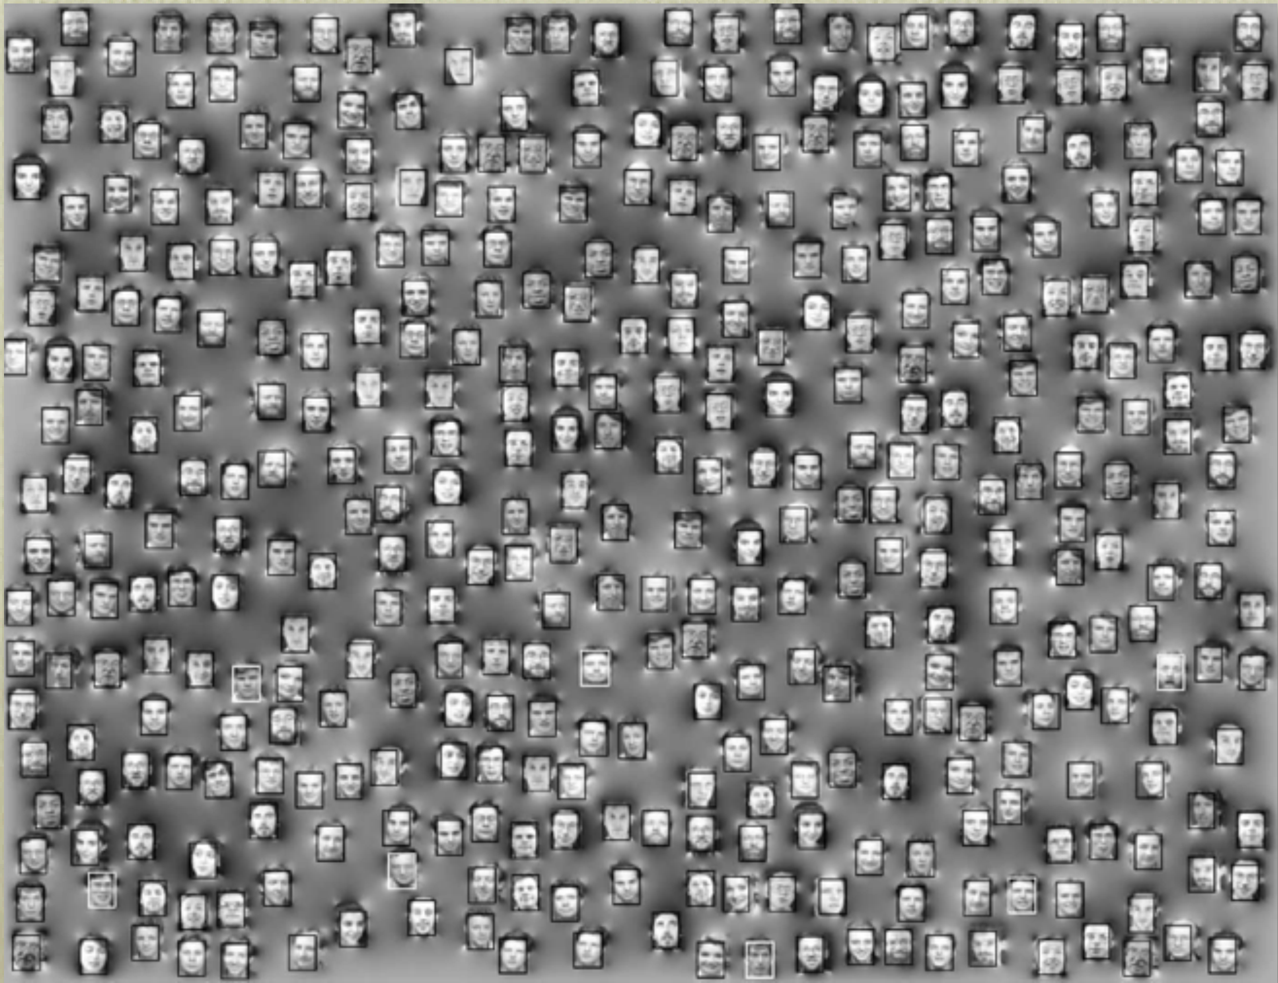

- Virtually all the faces correctly identified

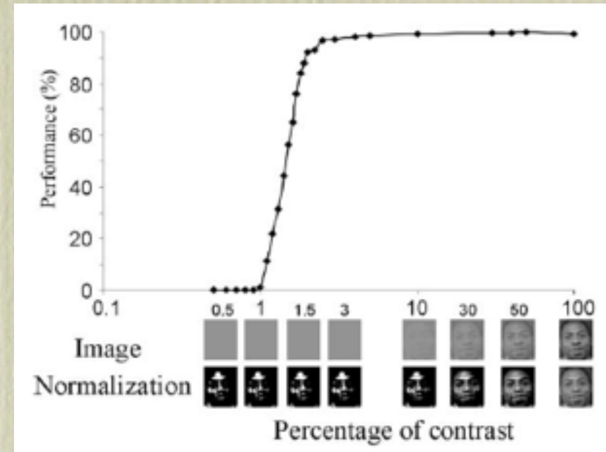

- Very robust to low contrast

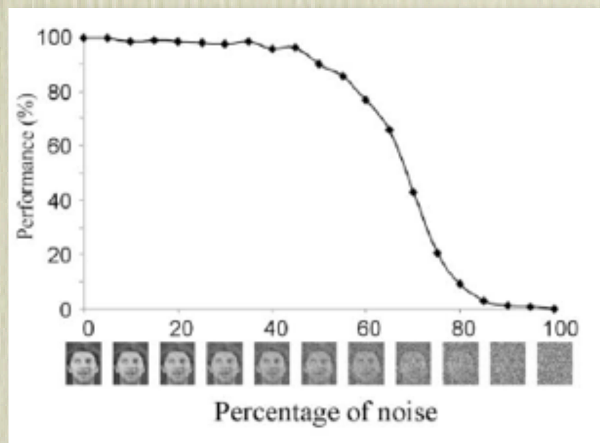

- Very robust to noise

SpikeNet Technology

- Created in 1999
- Currently 11 employees
- Key mechanisms
  - Control the percentage of cells that fire in the input layer (1-2%)
    - This can be done using inhibitory circuits
  - Put high weights on the earliest firing units
  - Put low (or zero) weights on later firing units
  - Set the threshold of the recognition layer units so that only inputs similar to the training stimulus can fire the unit
- Even complex visual forms like faces can be detected and localised using simple 3-layer networks
  - One spike per neuron
  - No feedback

What about learning?

Spike Time Dependent Plasticity (STDP)

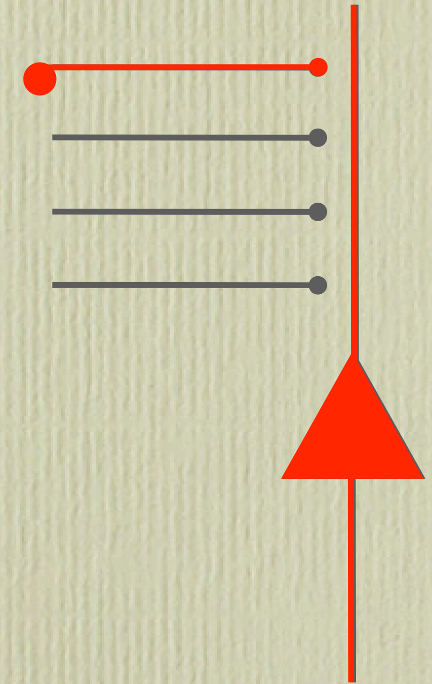

- Synapses that fire before the target neuron get strengthened
- Synapses that fire after the target neuron get weakened

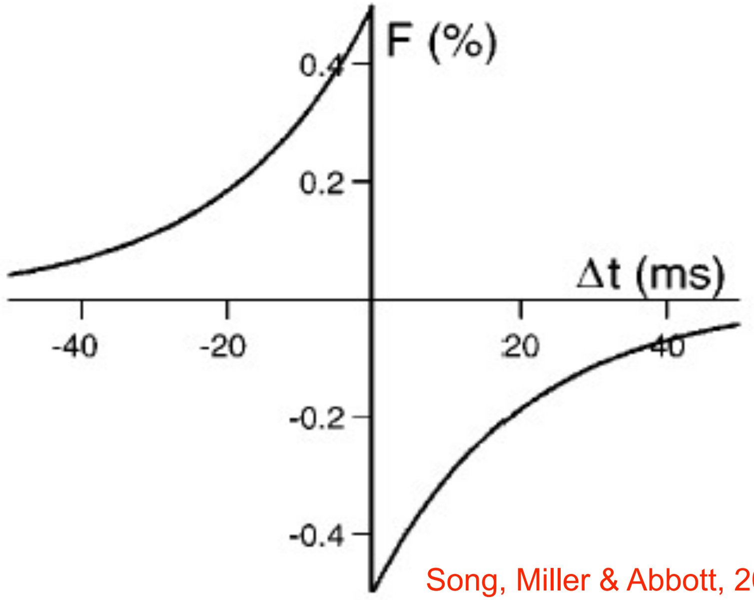

Song, Miller & Abbott, 2000

A natural consequence

- When an input pattern occurs repeatedly, high synaptic weights will concentrate on early firing inputs

Spike Time Dependent Plasticity

Presentation 1

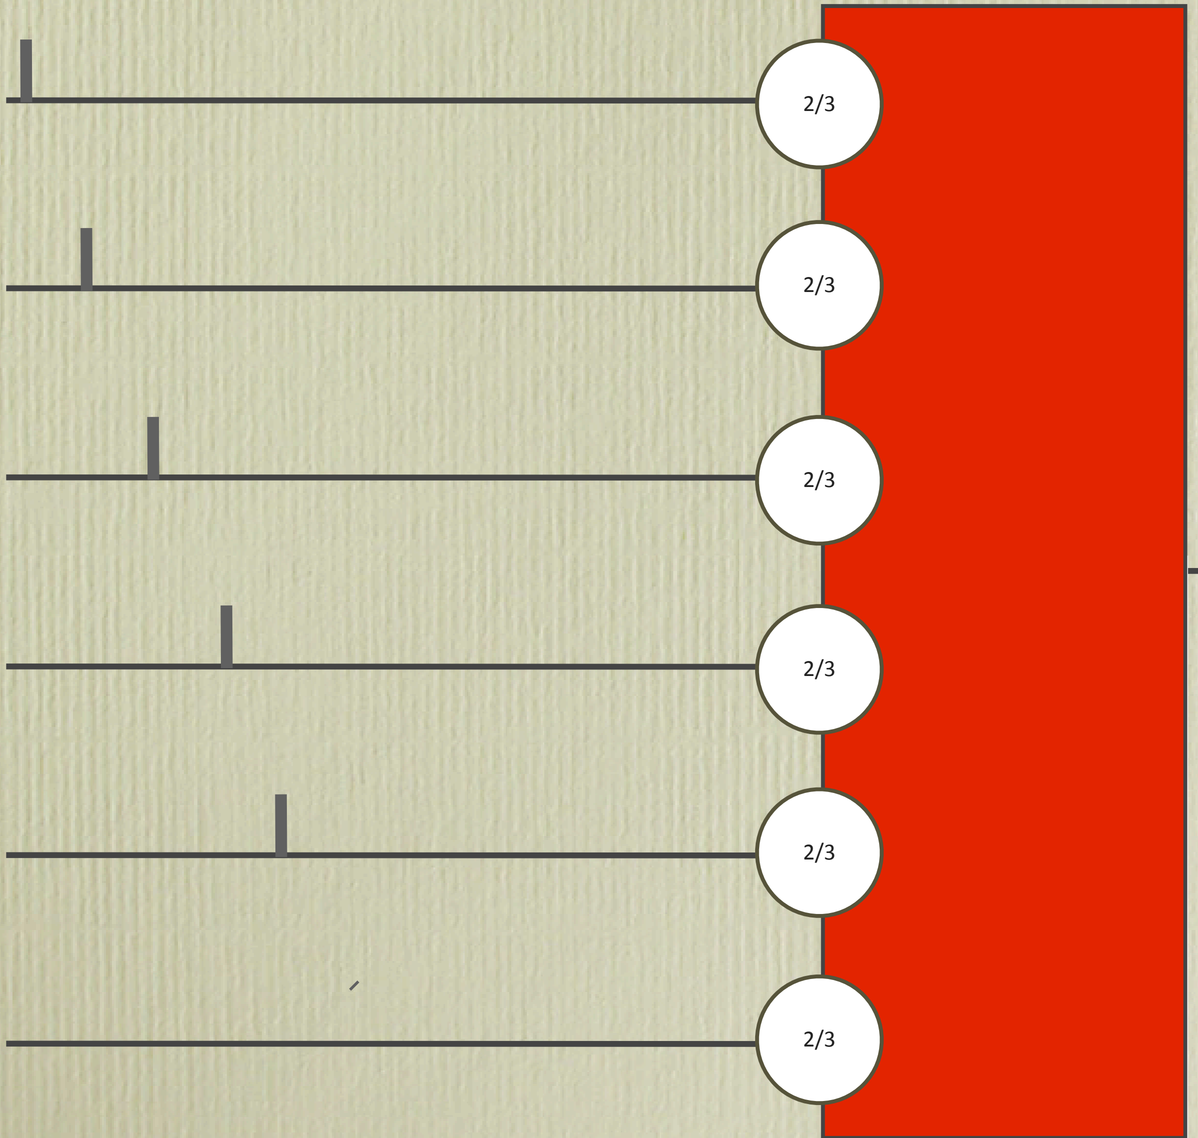

Threshold = 2

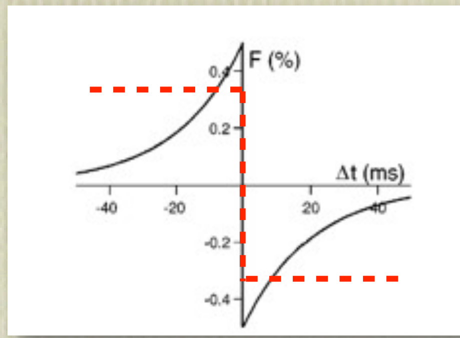

Spike Time Dependent Plasticity

Presentation 2

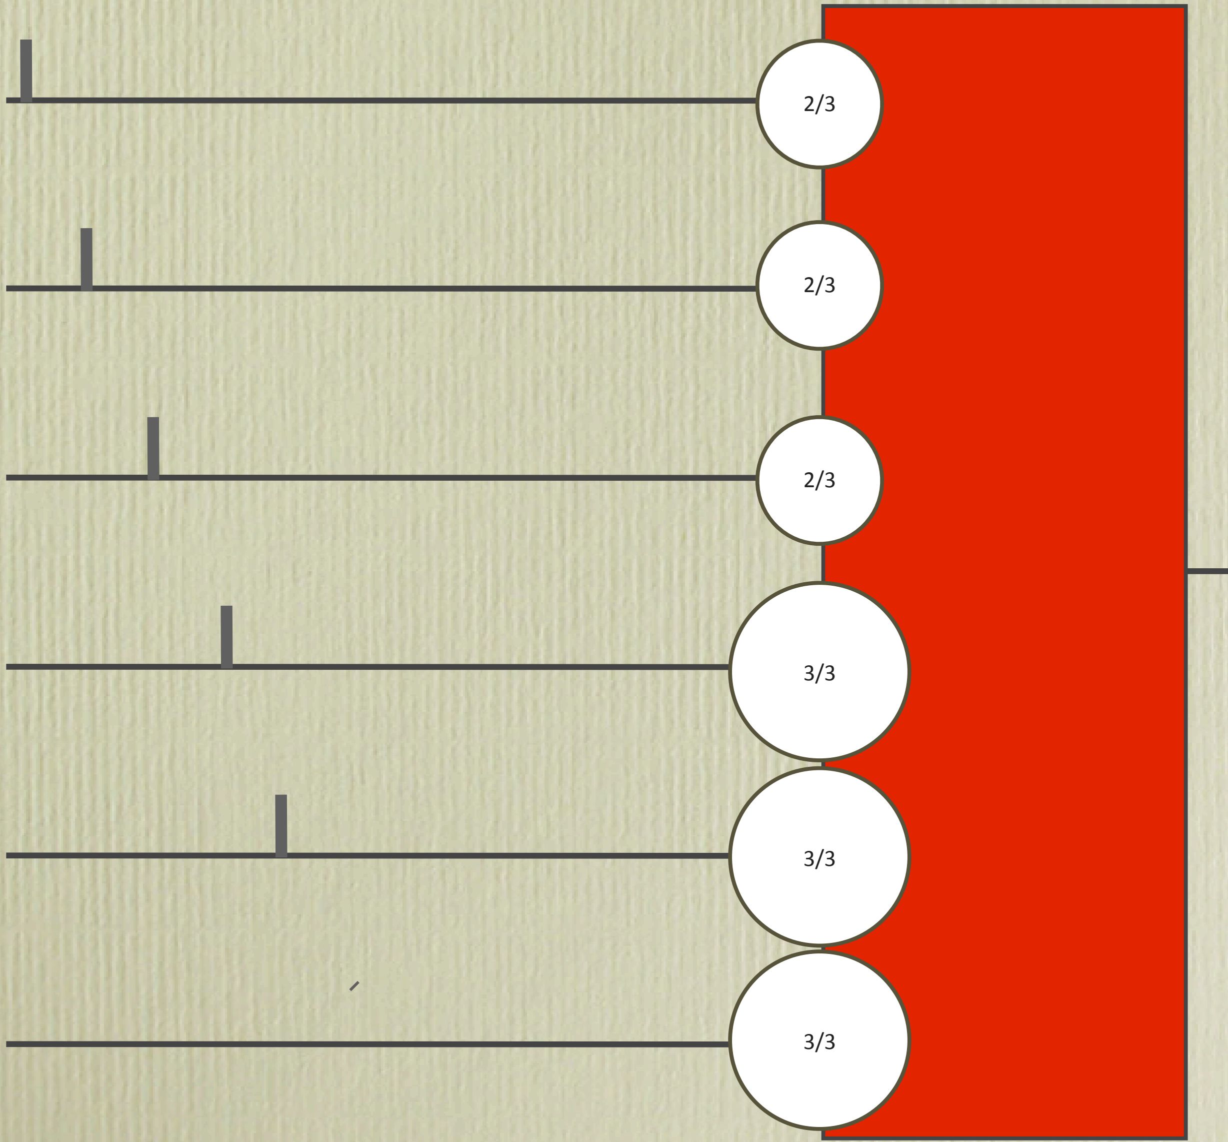

Threshold = 2

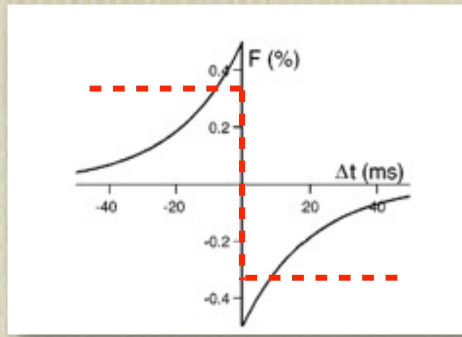

Spike Time Dependent Plasticity

Presentation 3

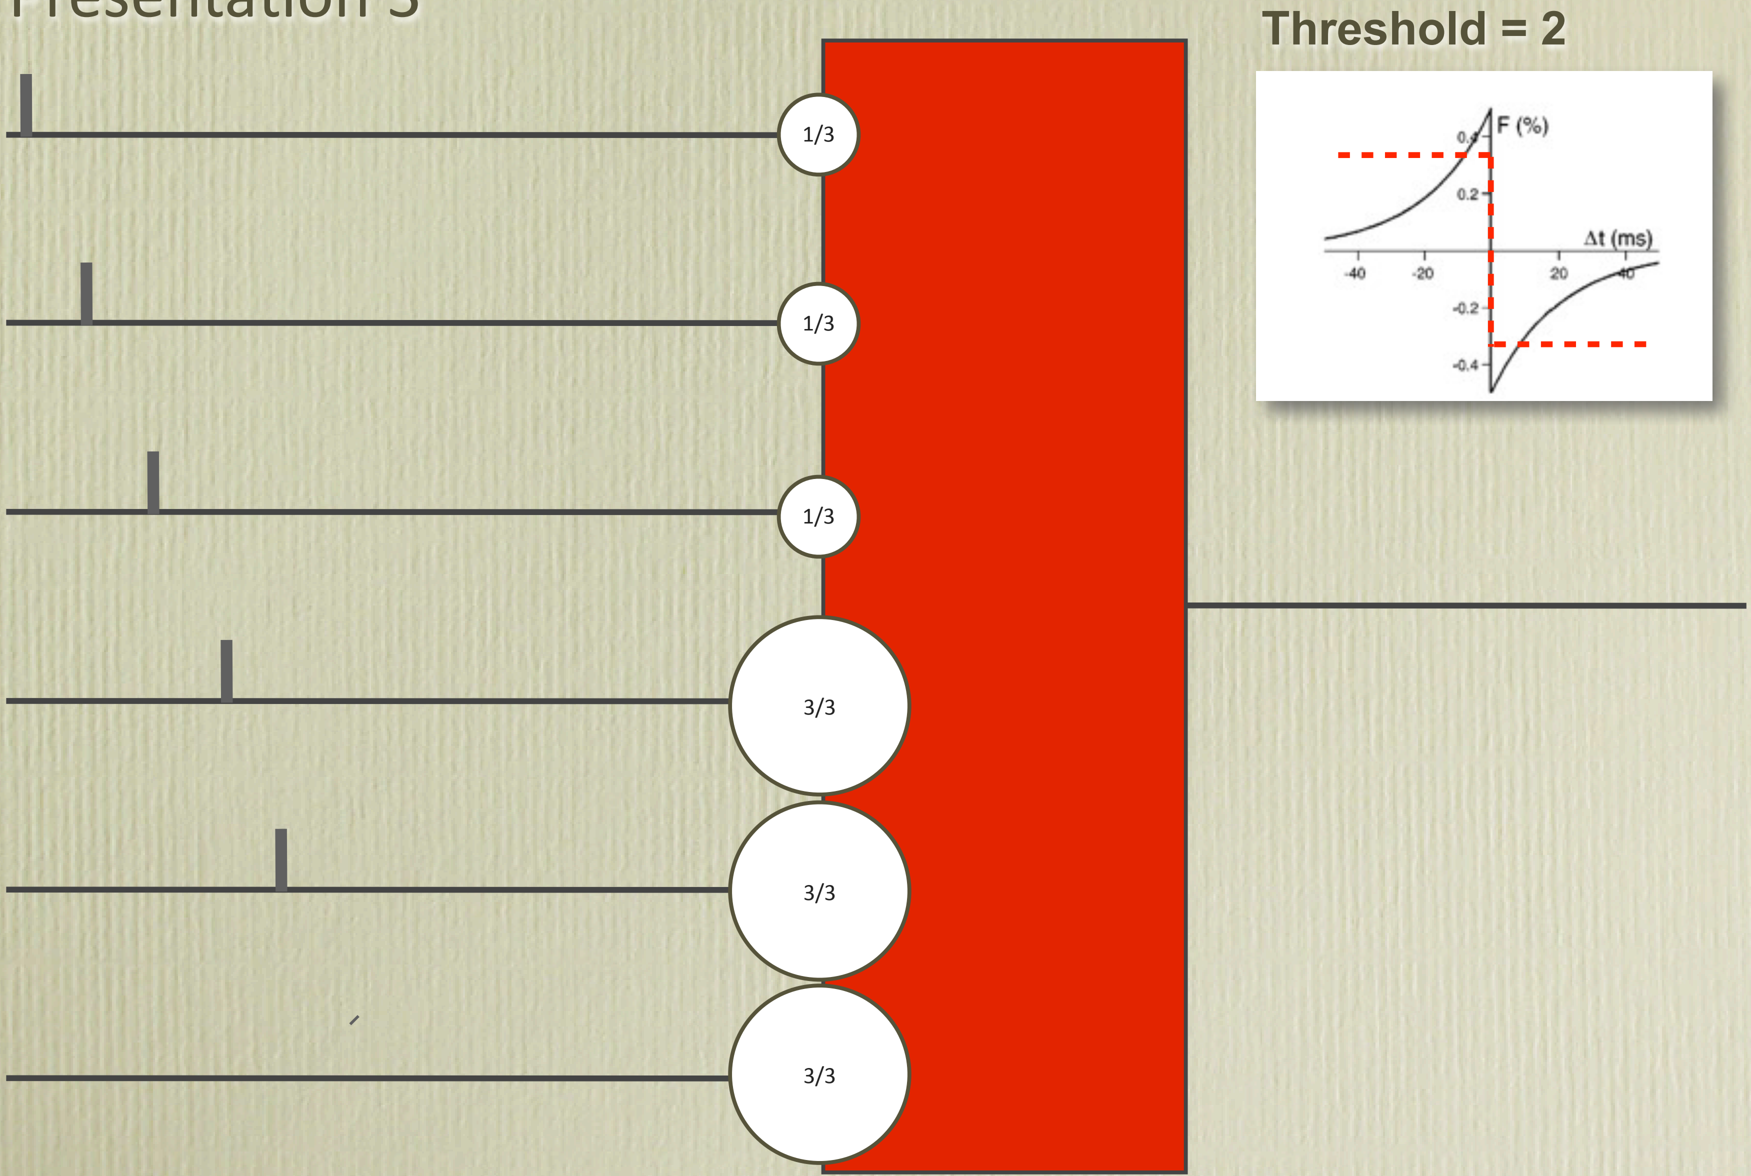

Spike Time Dependent Plasticity

Presentation 4

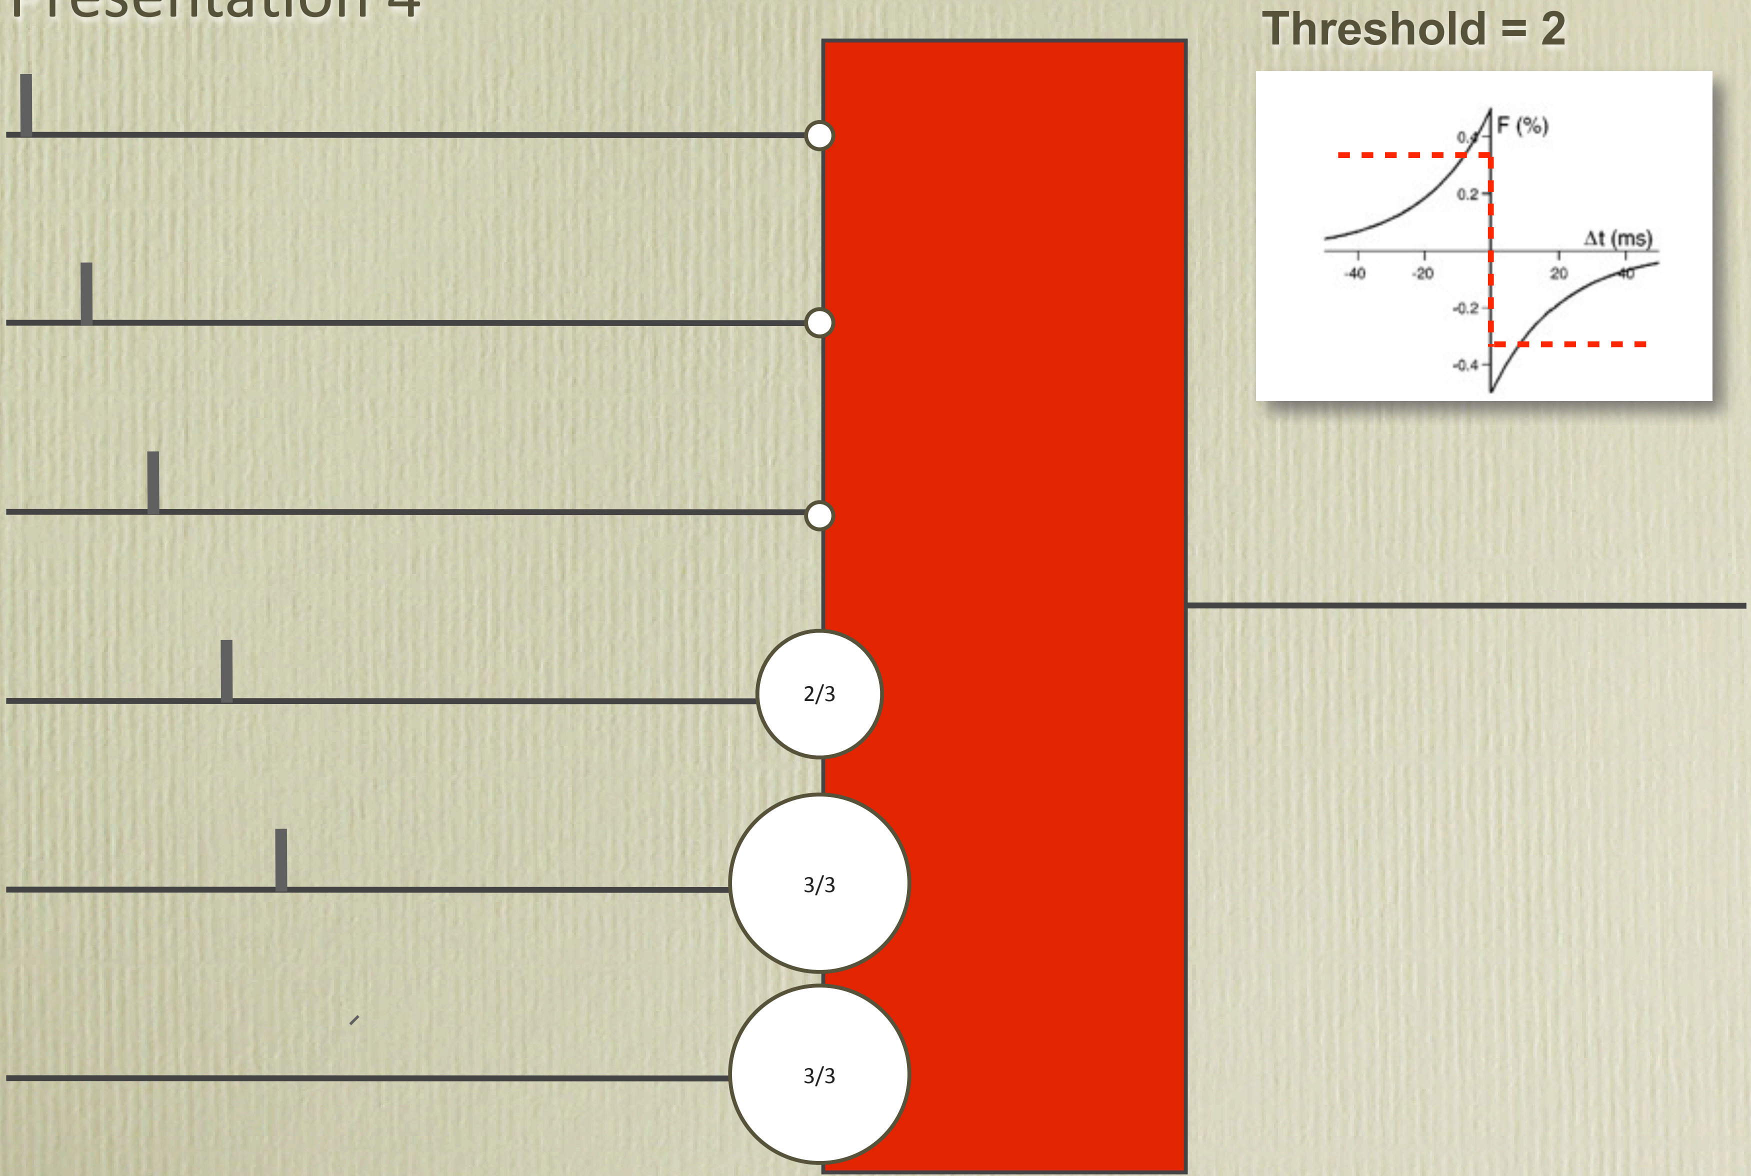

Spike Time Dependent Plasticity

Presentation 5

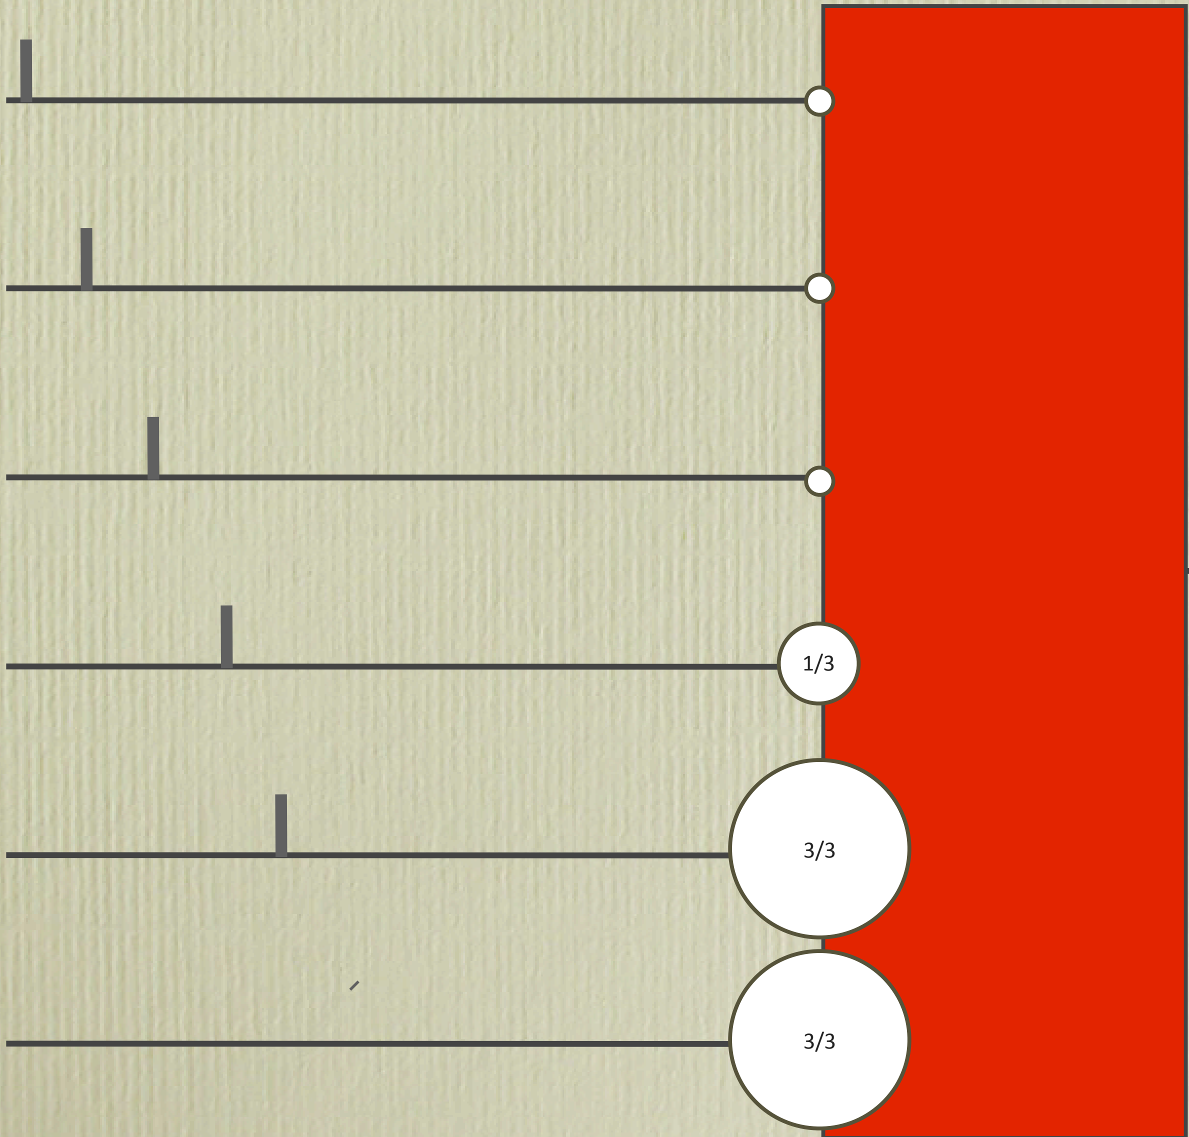

Threshold = 2

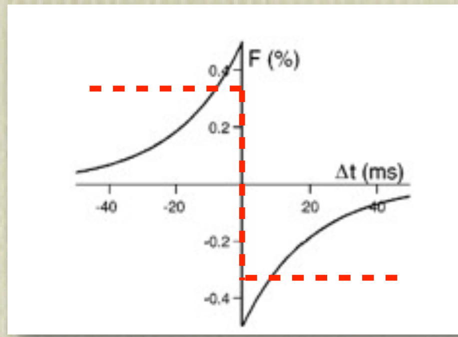

STDP concentrates high weights on the early firing inputs

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)

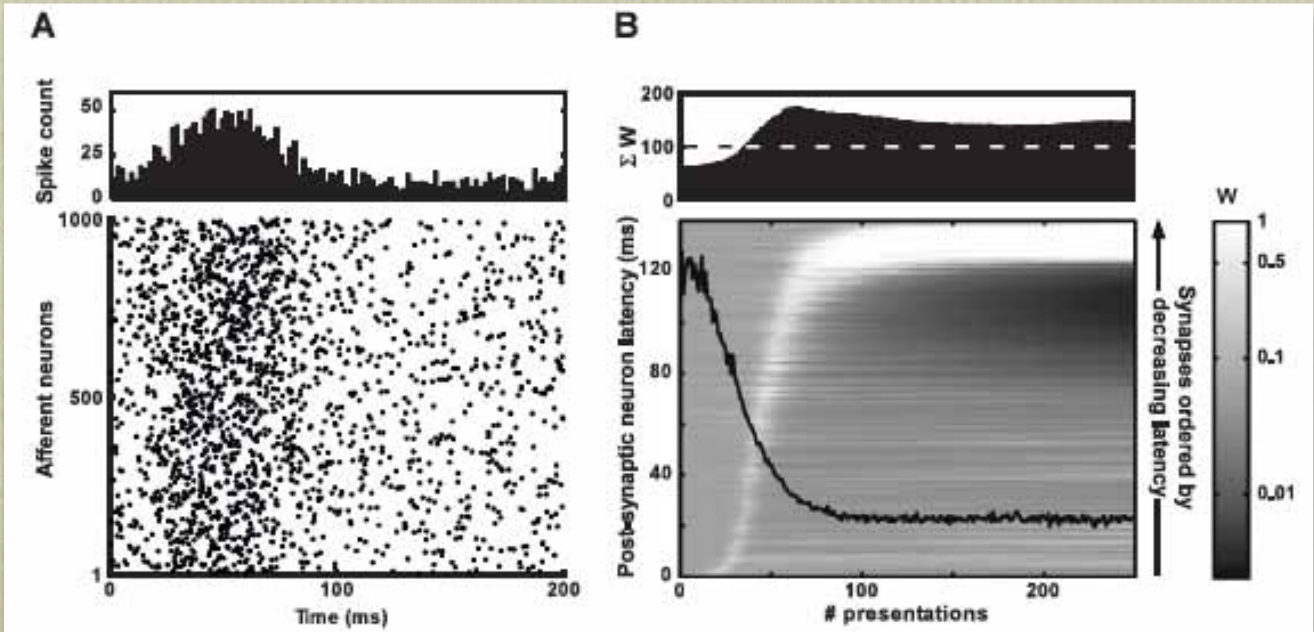

With a few tens of presentations, high weights concentrate on the earliest firing inputs

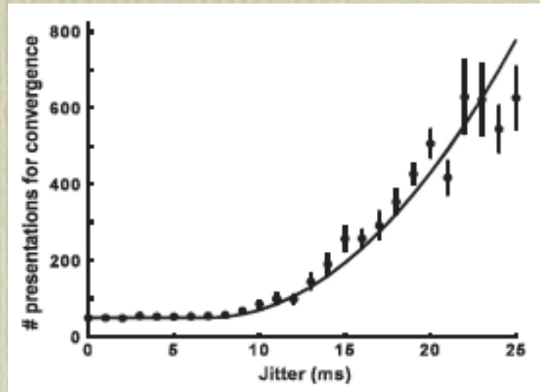

Even with jitter

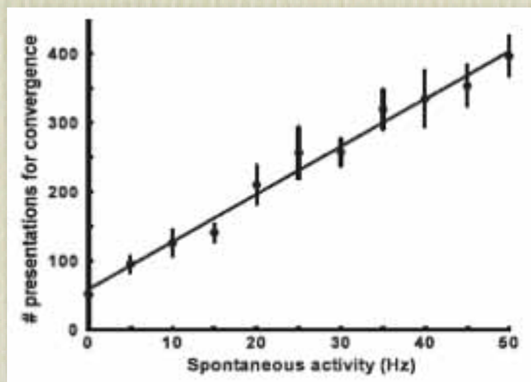

Even with spontaneous background activity

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]

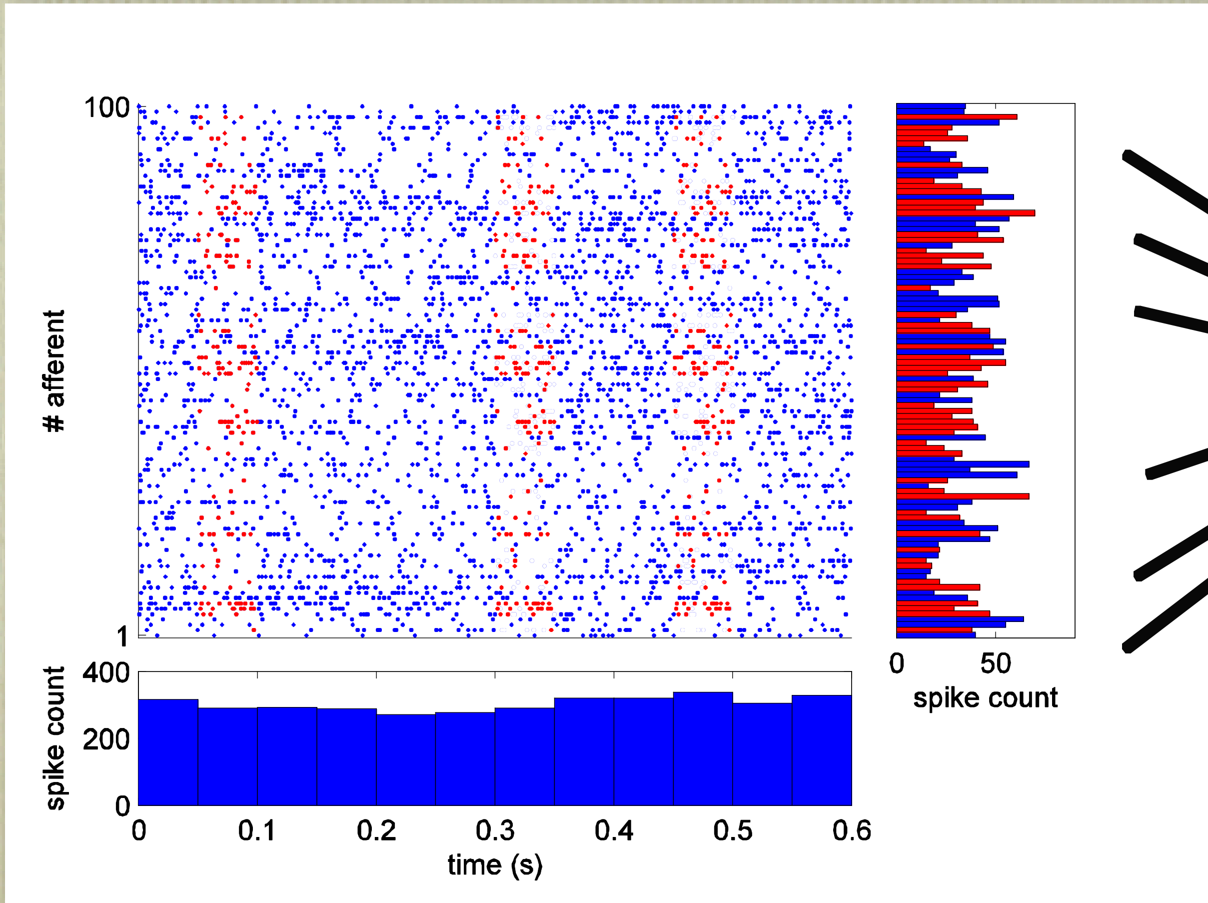

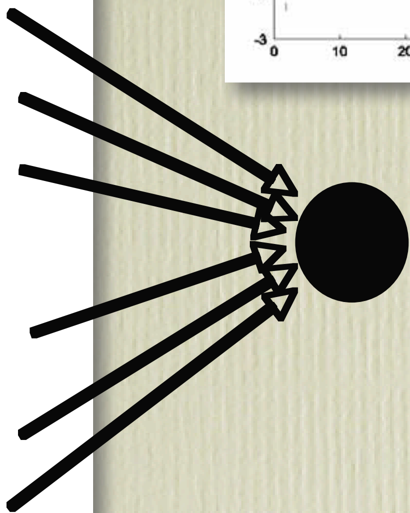

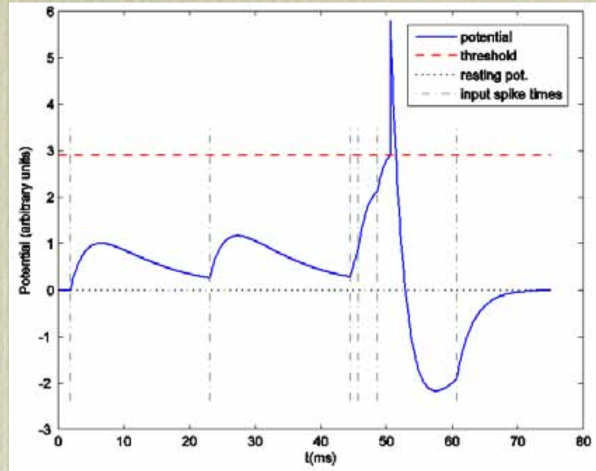

Learning Spike Sequences with STDP

- Initial State

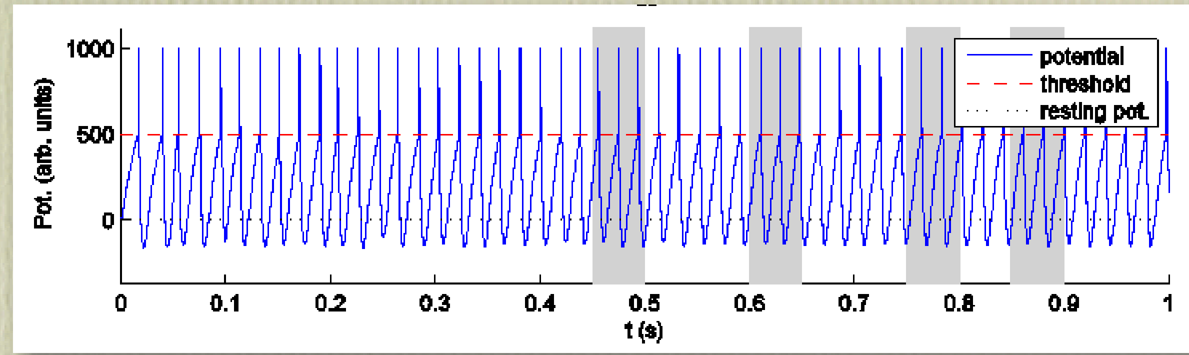

- During Learning

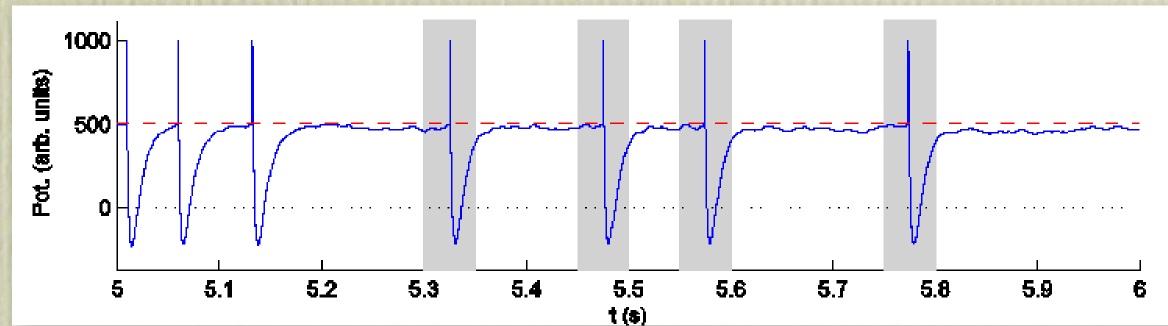

- After Learning

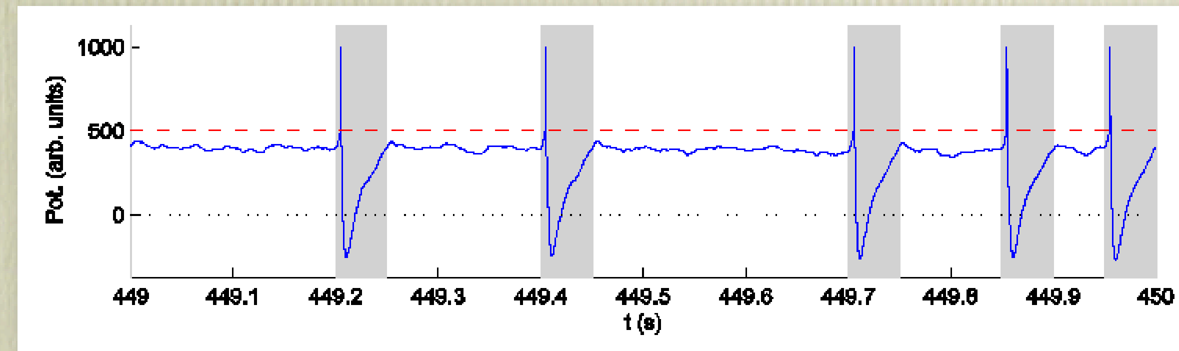

Learning Spike Sequences with STDP

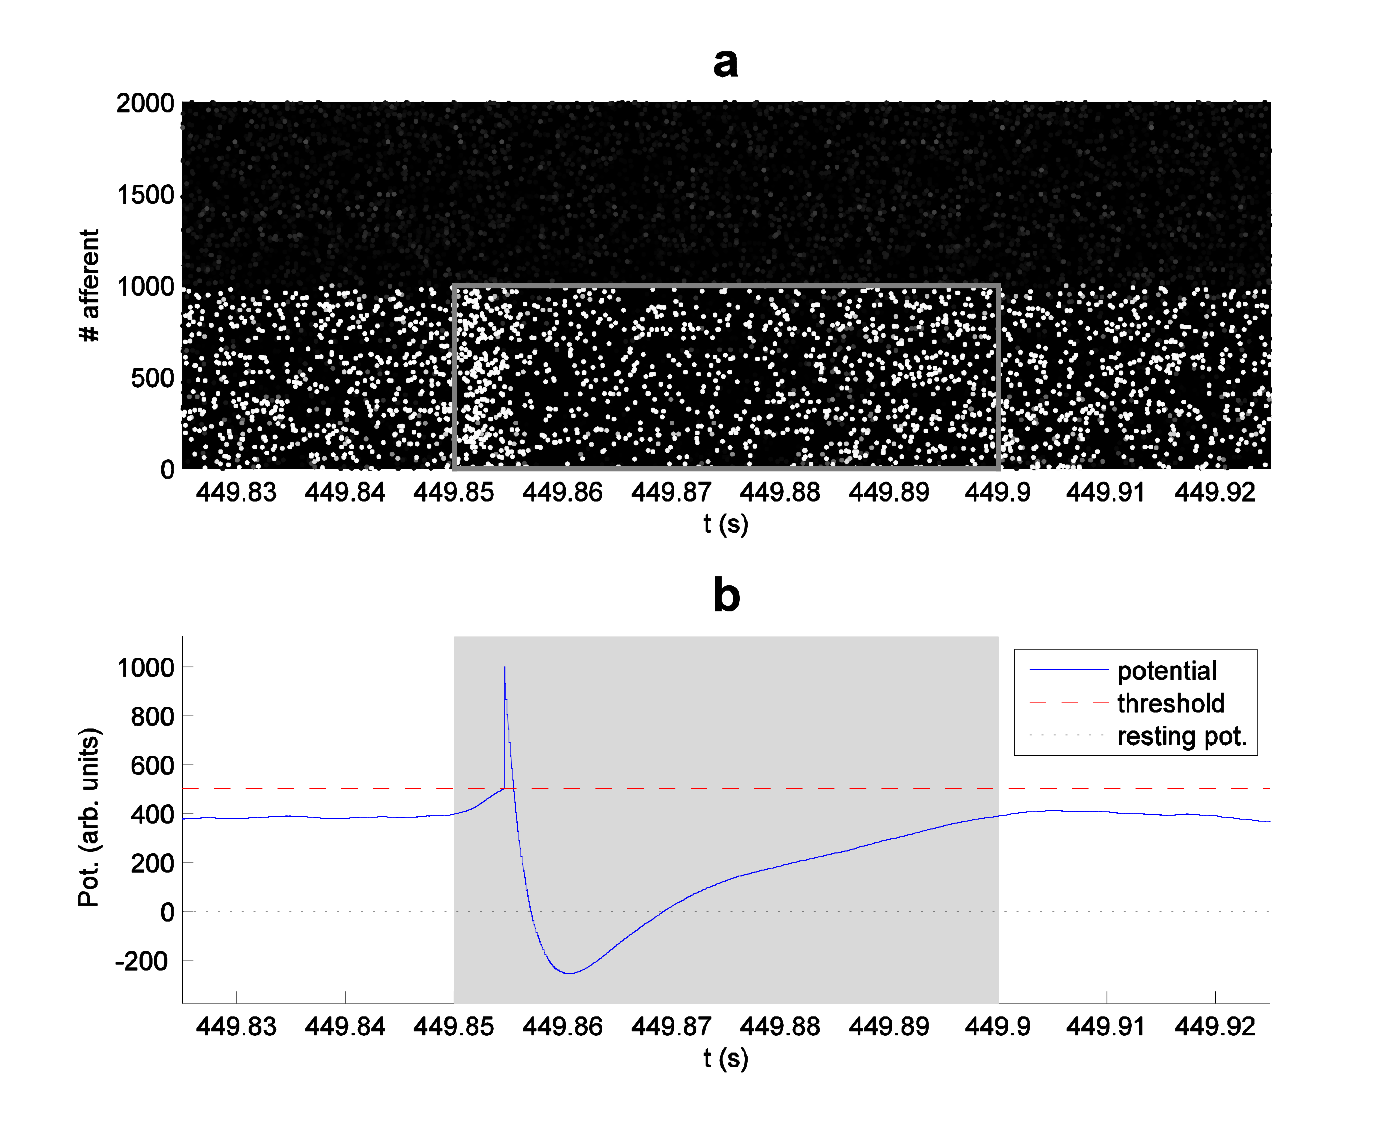

The neuron responds to the near synchronous firing in the afferents with high synaptic weights at the start of the pattern

Note
It would also respond to the end of the pattern if it was reversed

Competitive Networks

Competitive STDP-Based Spike Pattern Learning

Timothée Masquelier   Rudy Guyonneau   Simon J. Thorpe

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

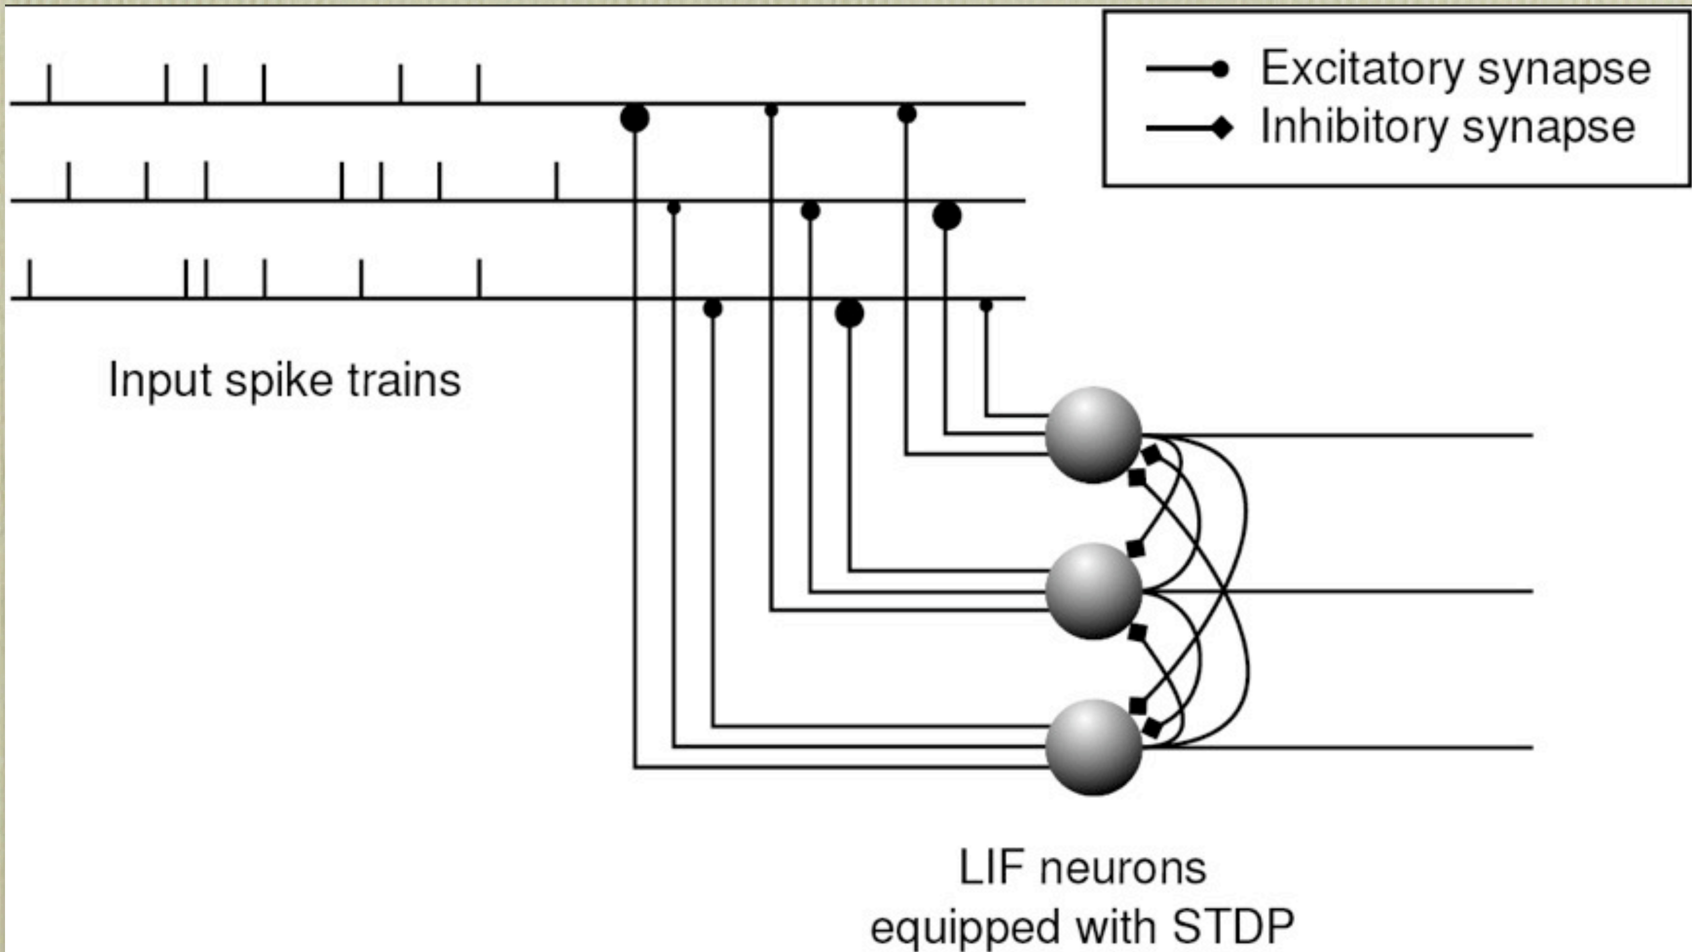

Learning with multiple neurons

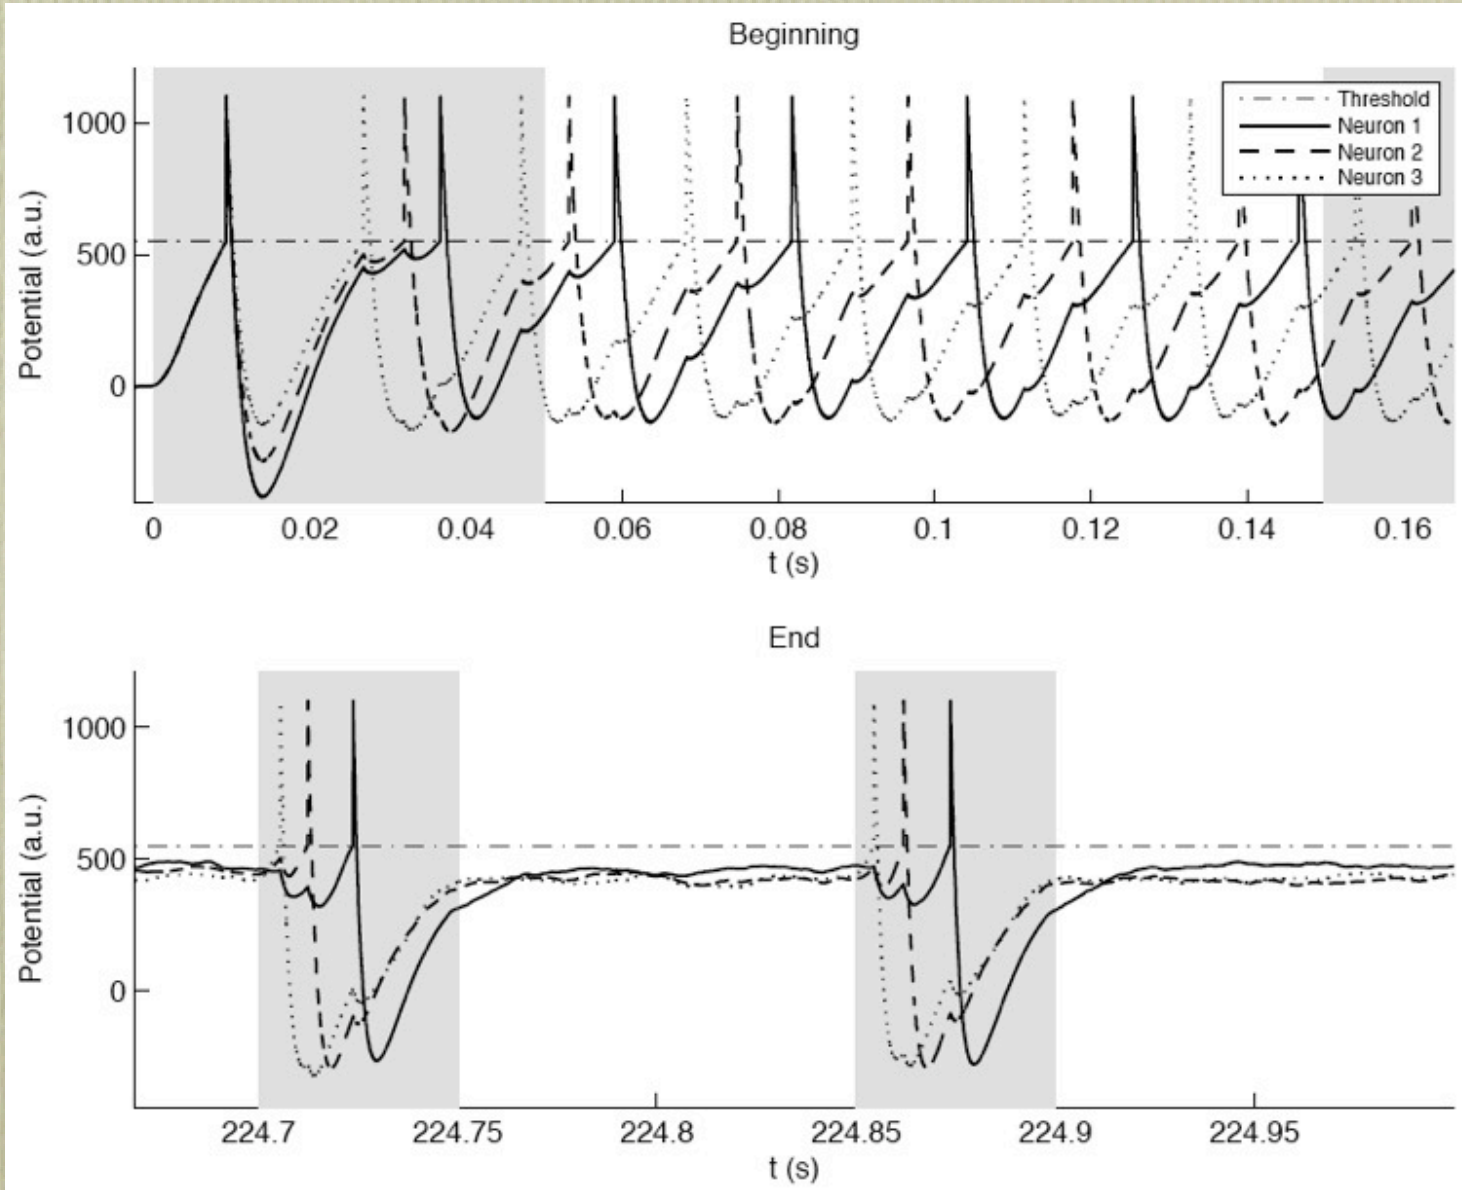

Multiple Patterns

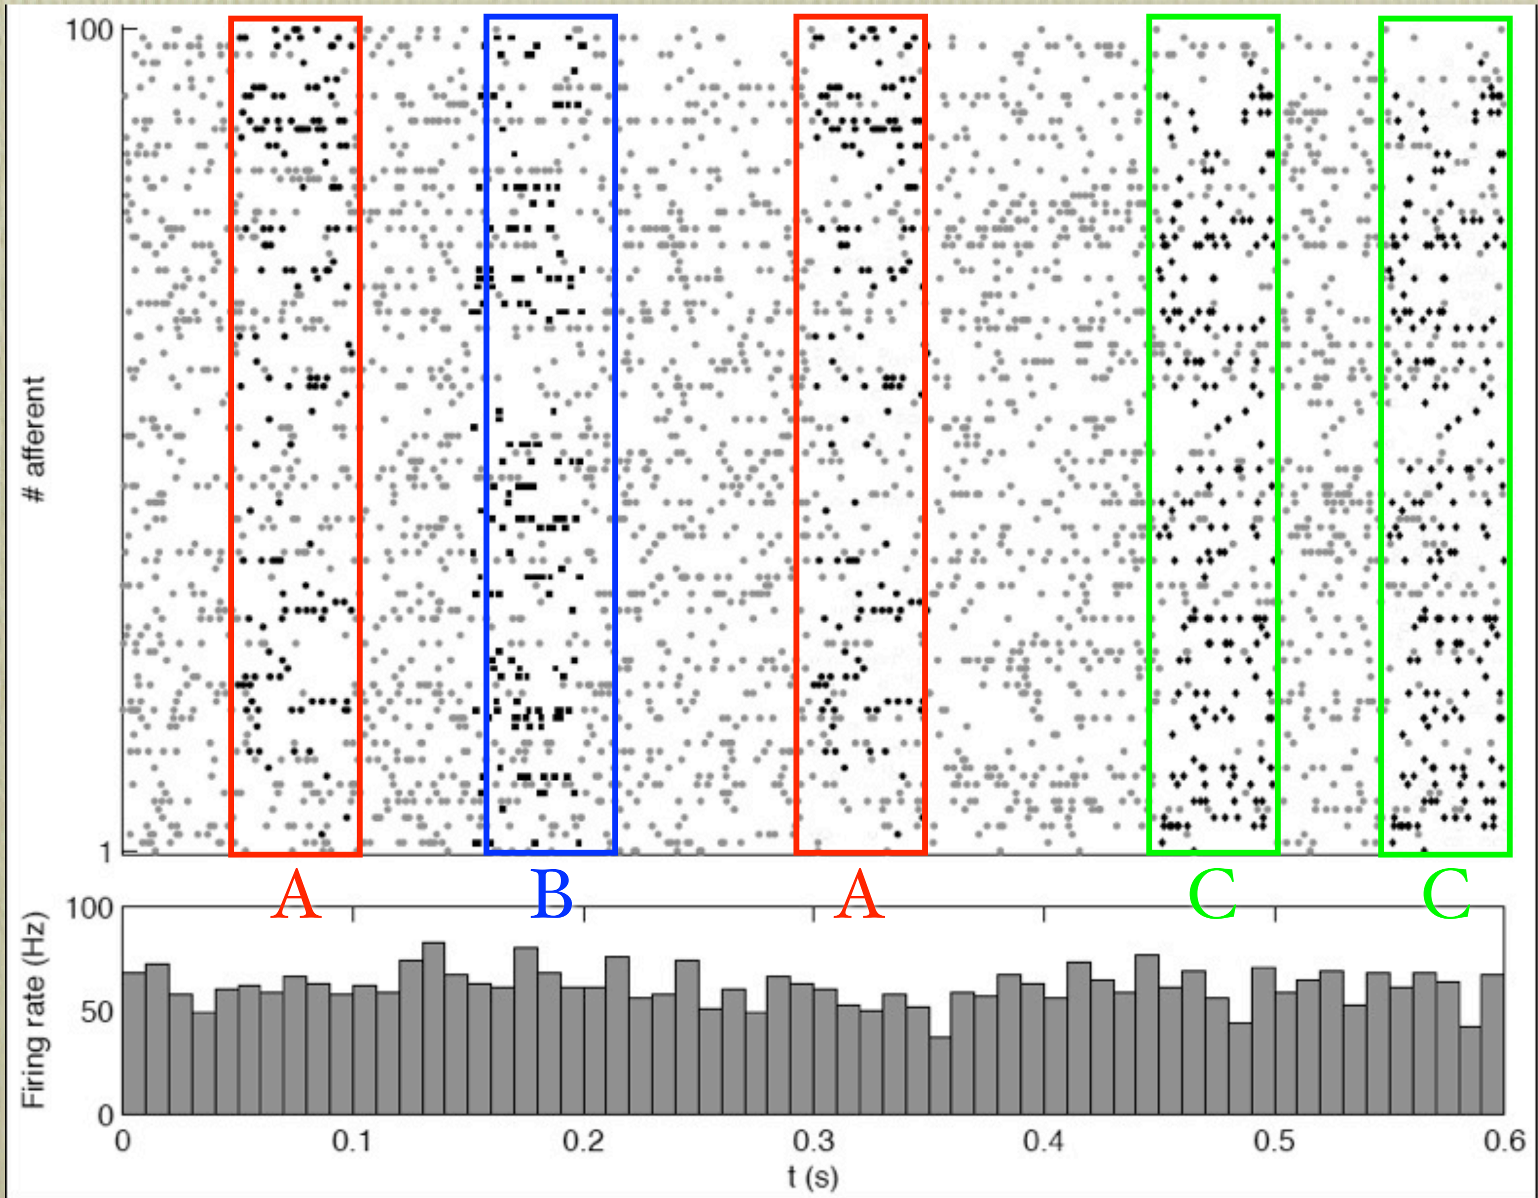

Does this happen in biological systems?

Experimental Support

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

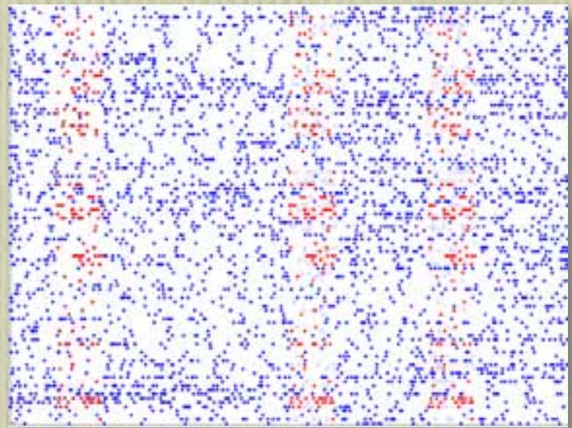

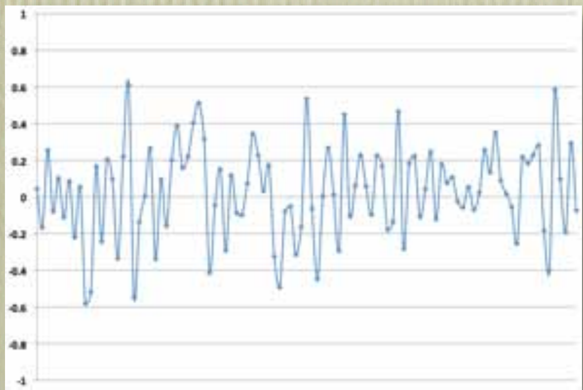

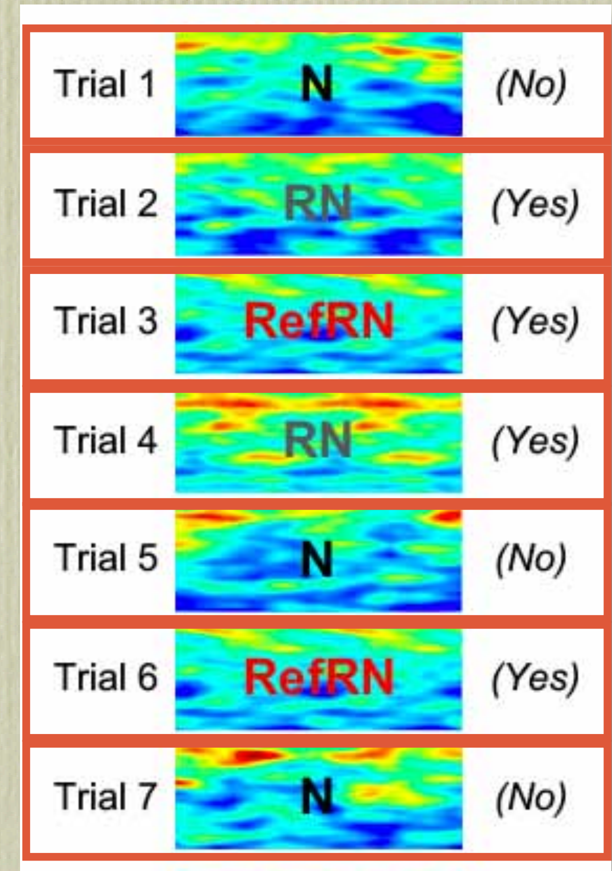

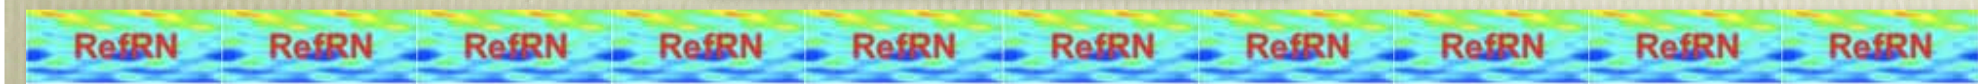

Experimental Support

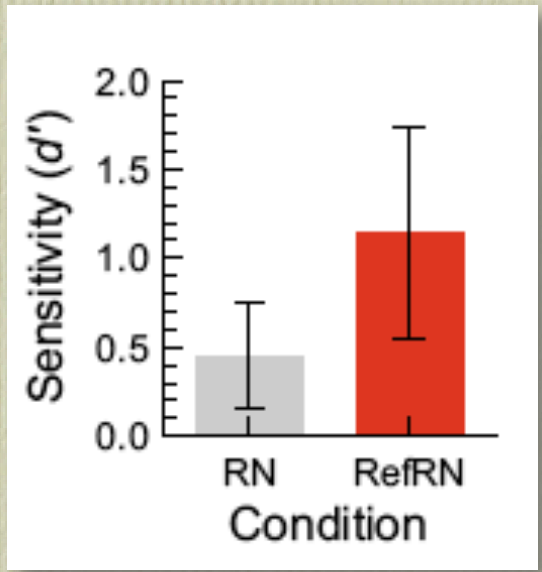

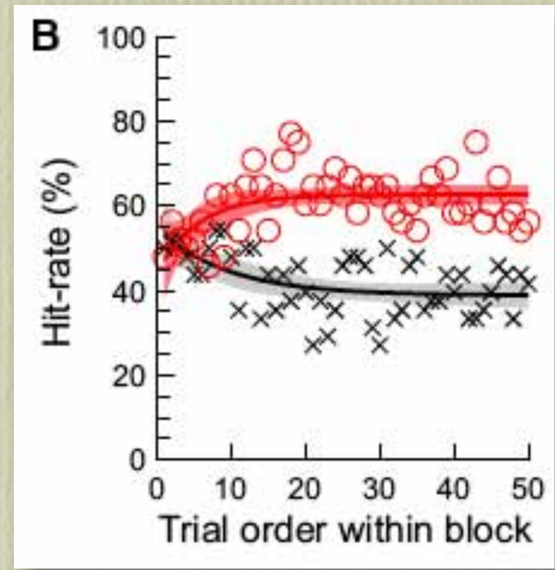

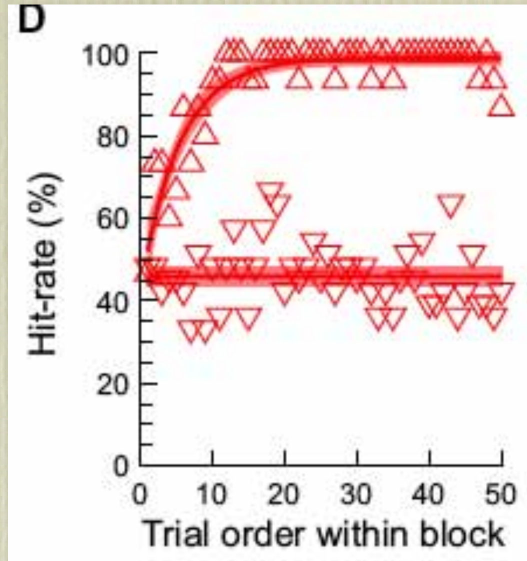

- Learning of random noise patterns
- Roughly 10 repetitions are sufficient!
- Learning appears to be all-or-none

How long does it last?

First Block

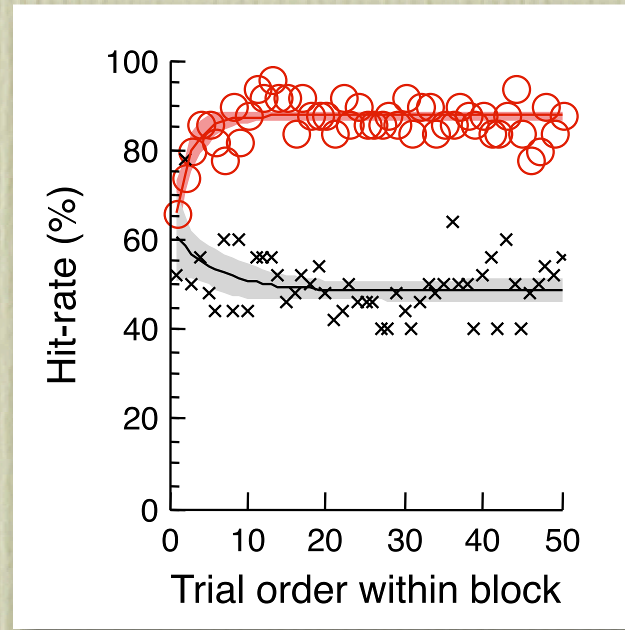

Second Block

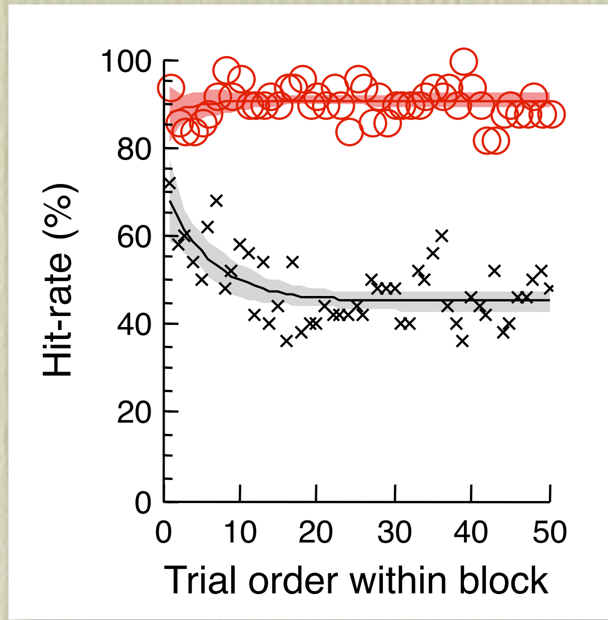

- Listeners are more sensitive from the first trial of the second block
- Median interval 17 days
- The memories last for weeks!

How invariant is the learning?

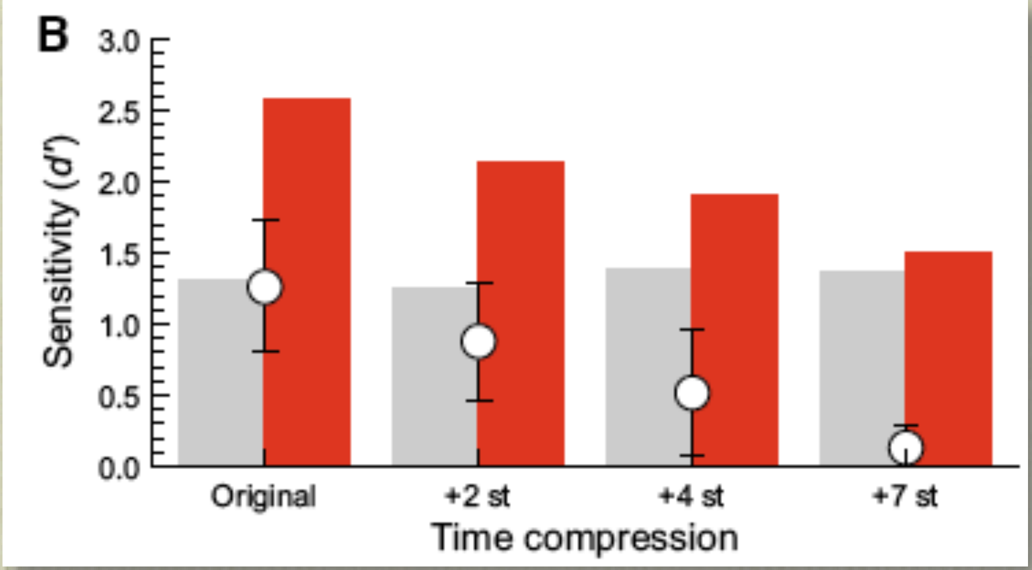

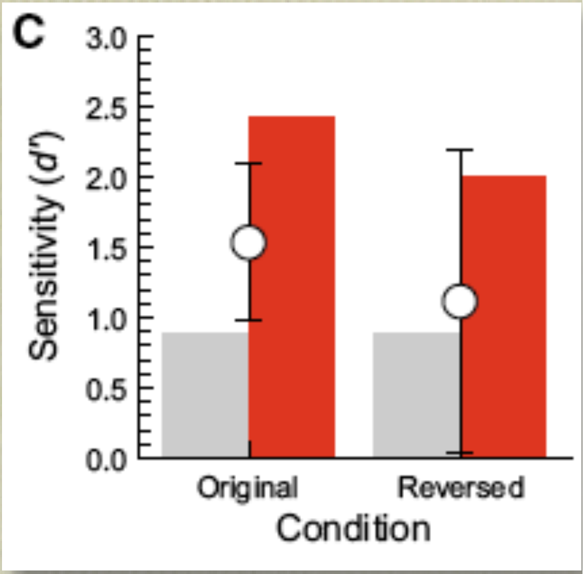

- Still works with speeded up patterns
  - Not learning of the precise timing
- Reversed patterns also work
  - Fits with STDP learning data

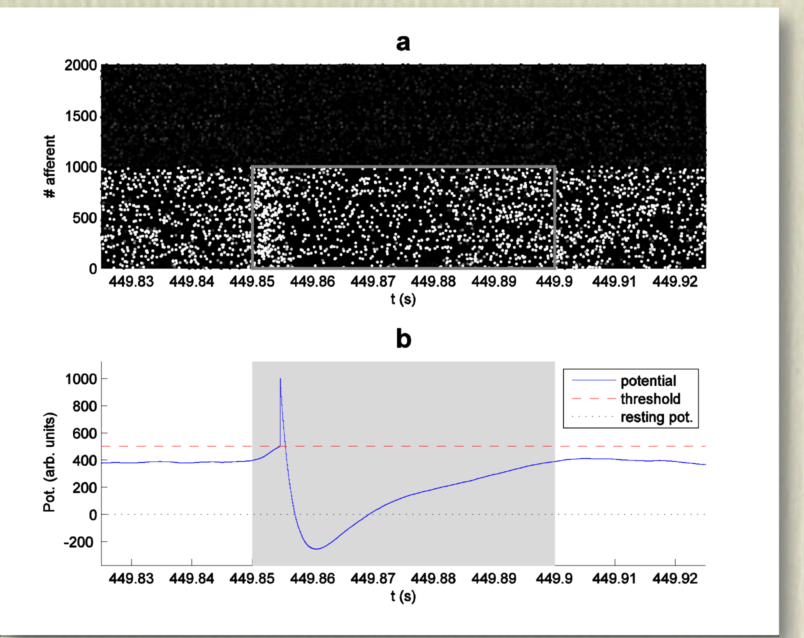

Can STDP learn Auditory Noise?

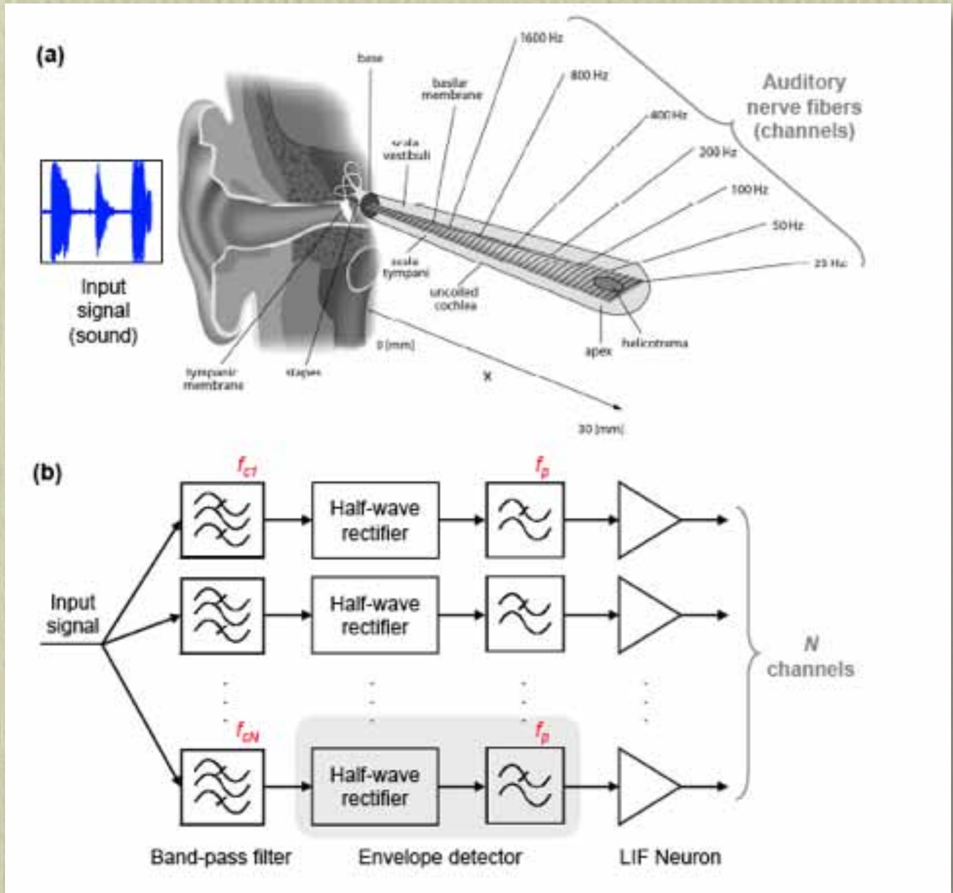

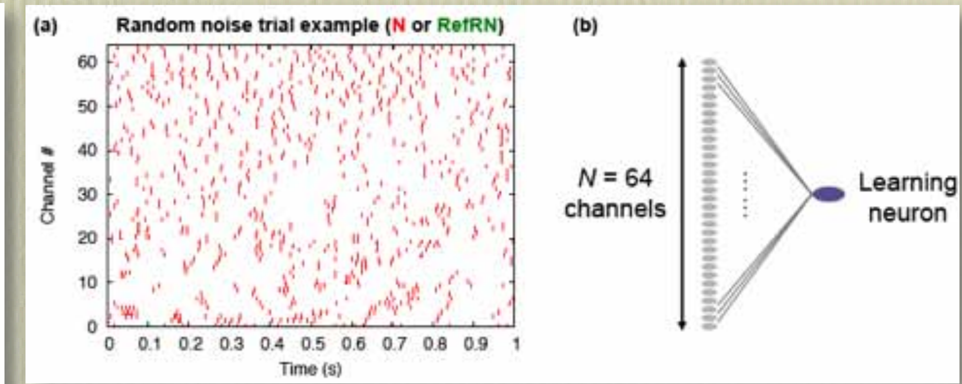

Olivier Bichler, Thesis

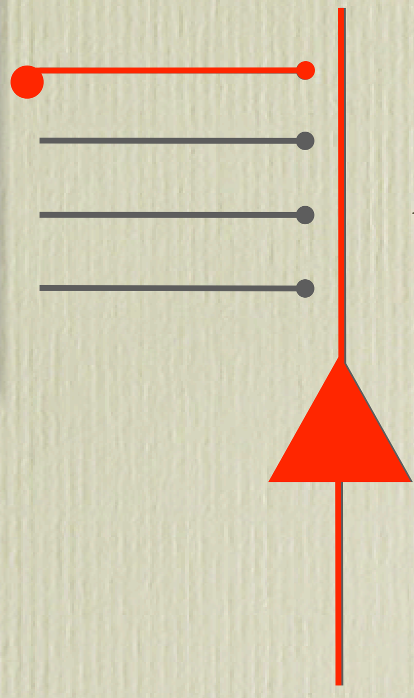

Modified STDP rule

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

Can STDP learn Auditory Noise?

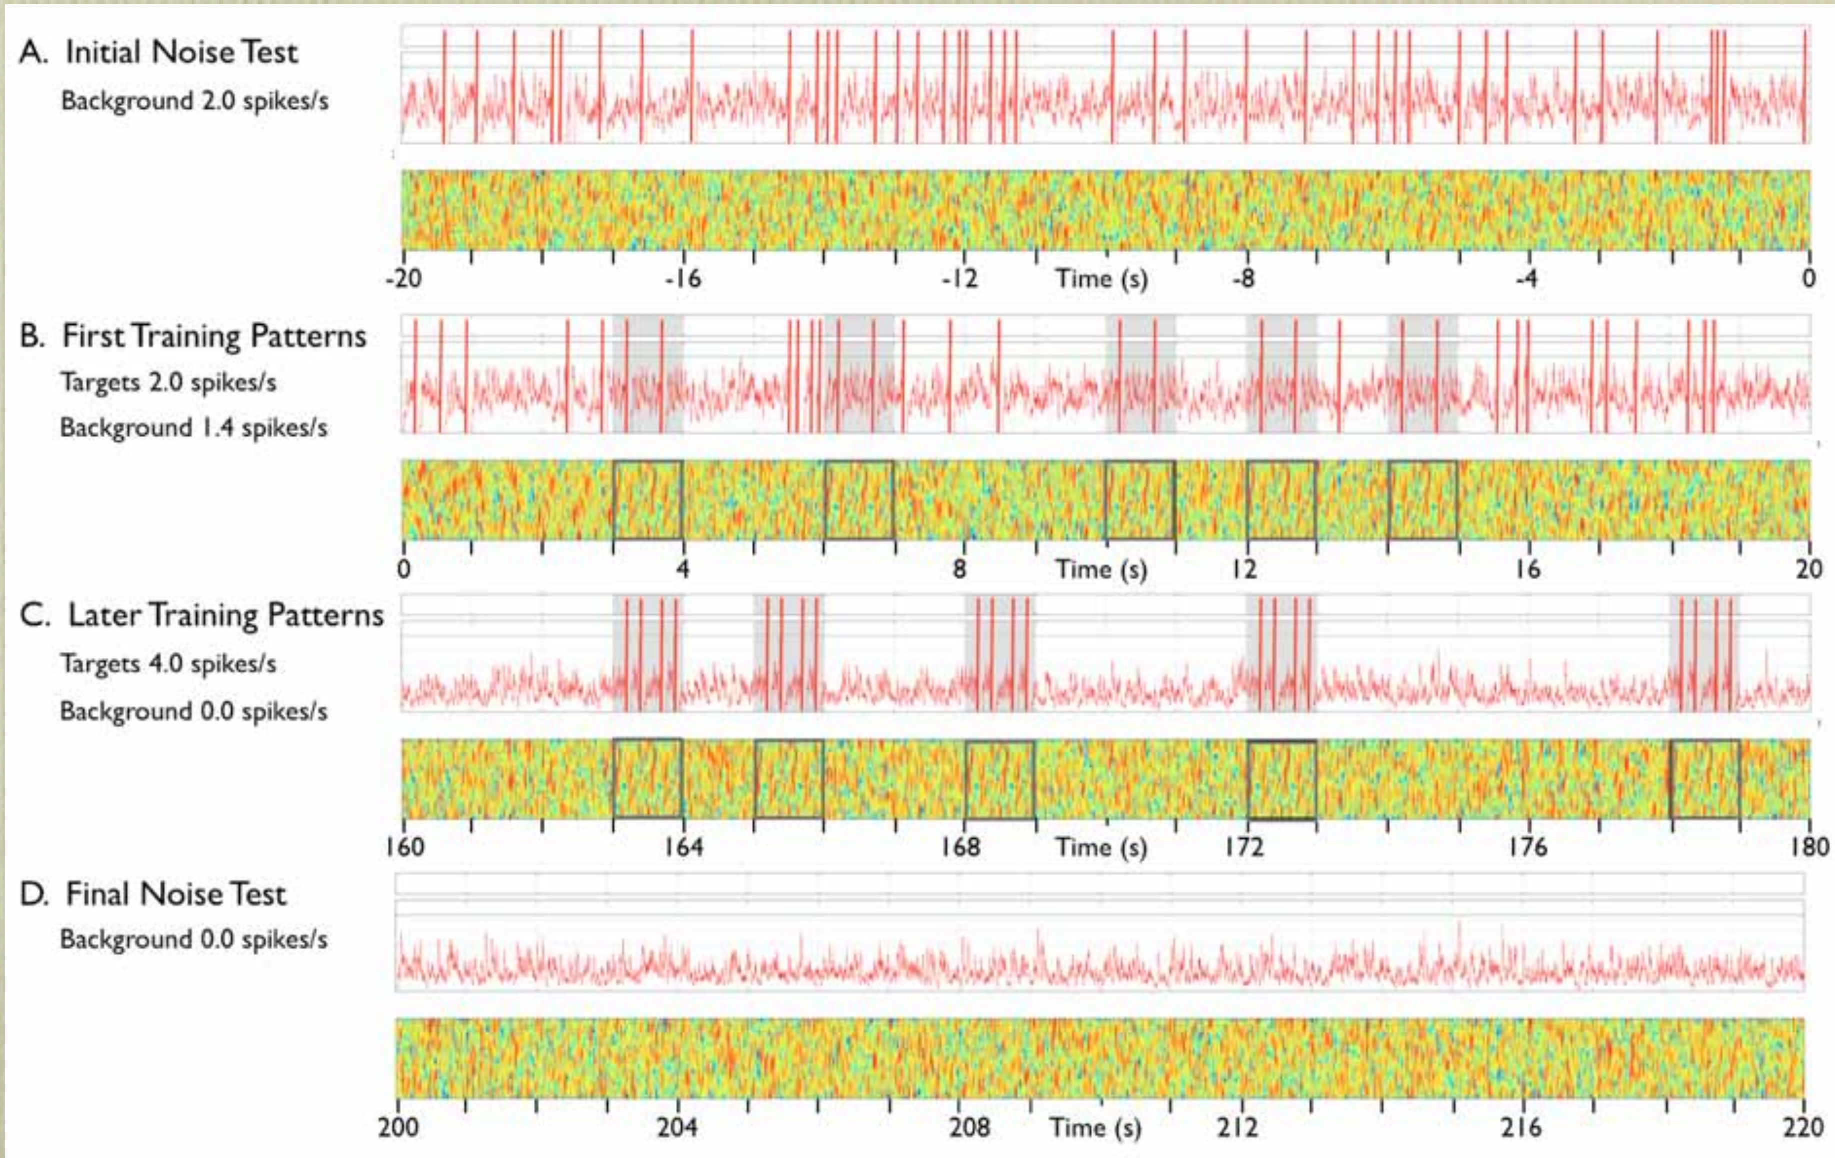

Applications in Vision

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]

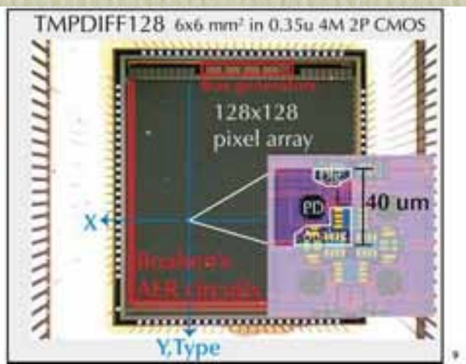

566 IEEE JOURNAL OF SOLID-STATE CIRCUITS, VOL. 43, NO. 2, FEBRUARY 2008

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*

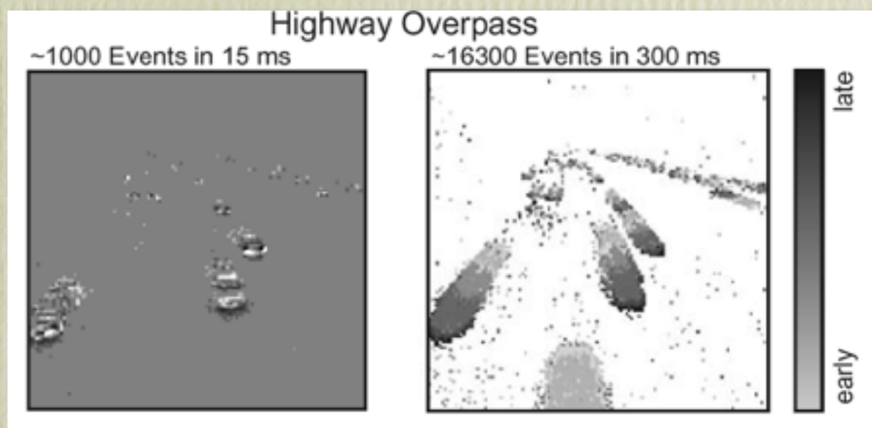

Simulation Studies

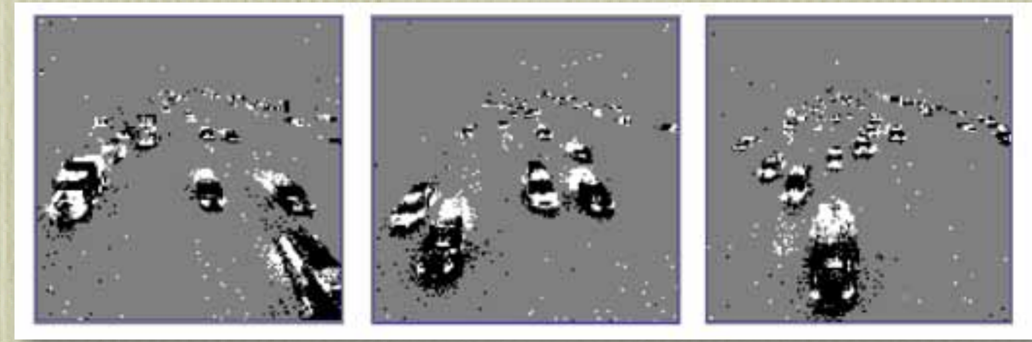

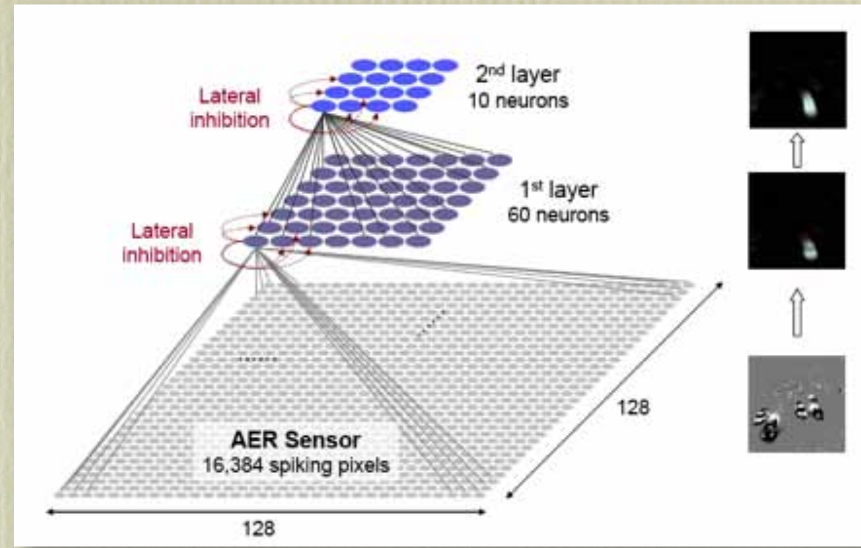

Layer 1

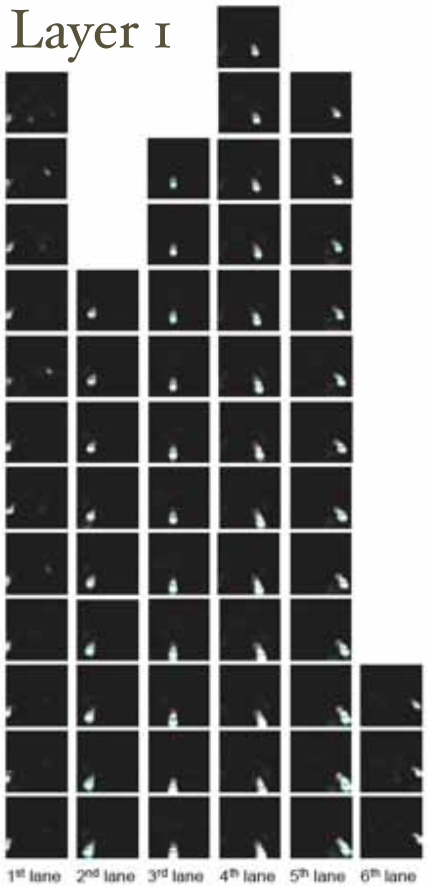

Layer 2

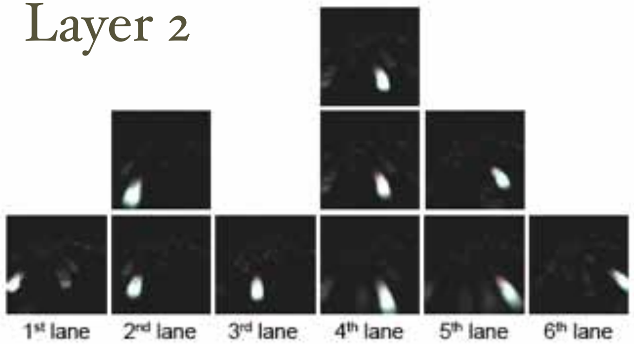

Neural Events
Reference Events

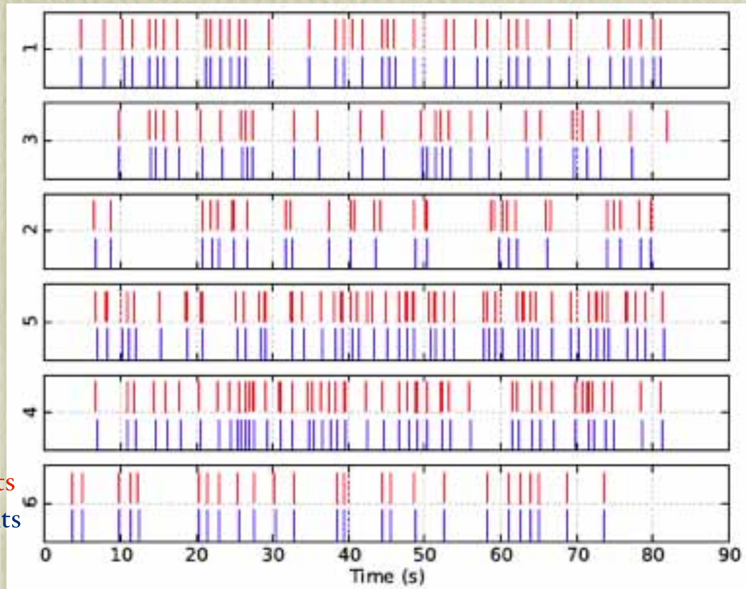

Performance 98% overall

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
- Potential for implementation in hardware?

Memristor Based STDP

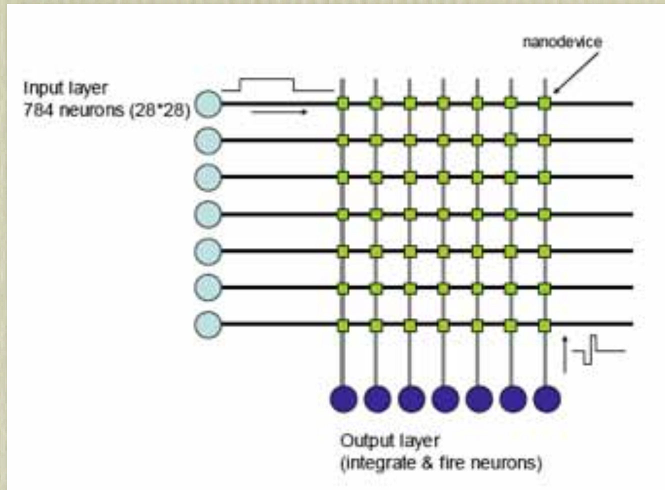

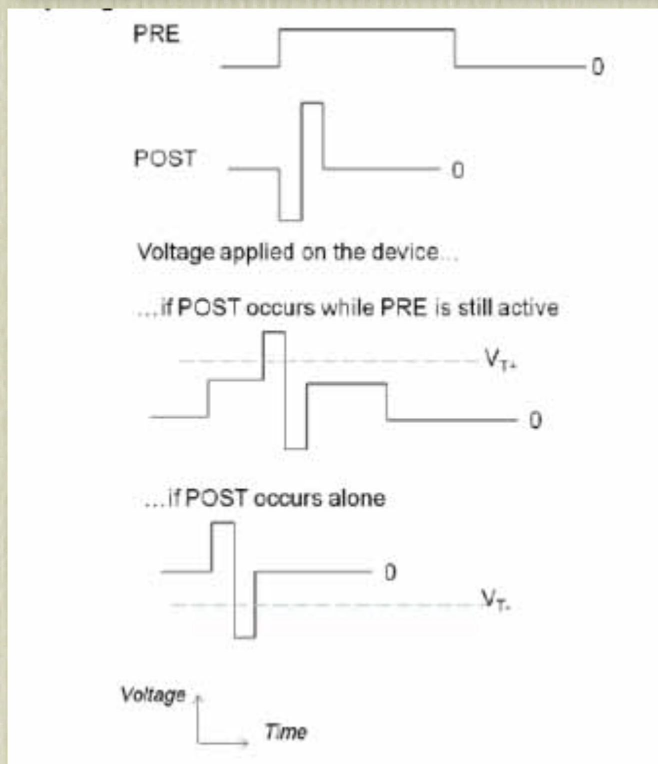

- When the neuron fires a spike
  - All synapses are depressed
  - Except those active just before

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)

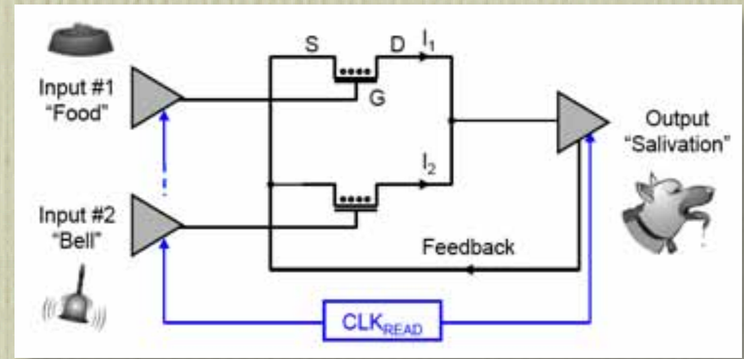

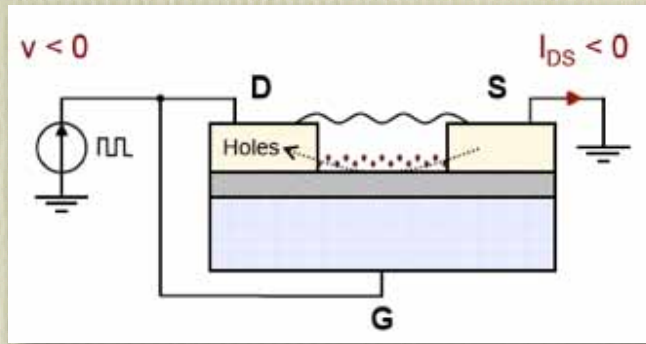

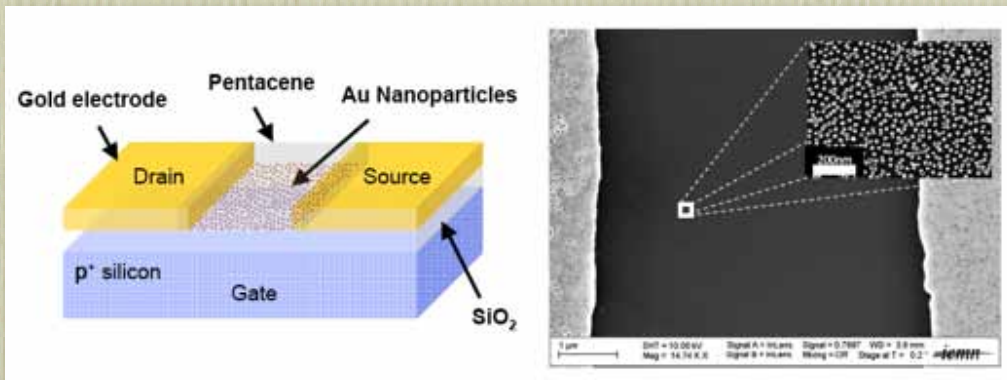

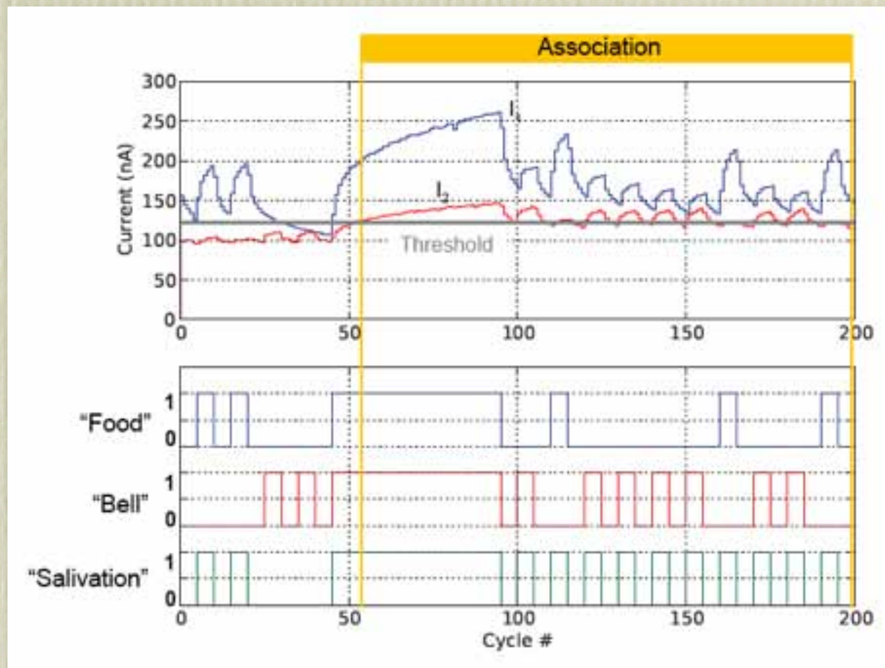

Final Conclusions

- Spikes make sense
- Sophisticated processing with only one spike per neurone
  - No conventional rate coding
  - No feedback
- STDP based learning
  - Neurones become selective to repeating patterns
  - A few tens of repetitions are enough
  - “Grandmother Cell” selectivity
  - Allows memories to be maintained for decades
- Towards hardware implementation
  - Memristor crossbar architectures
  - Implementation using (for example) NOMFETS
- A completely new way to compute

Credits

Psychophysics

Sébastien
Crouzet

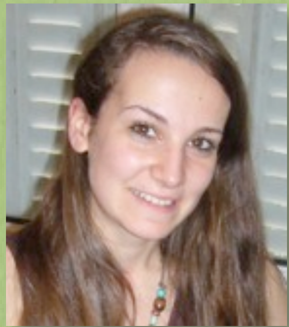

Marie
Mathey

Adrien
Brilhault

STDP Modelling

Rudy
Guyonneau

Tim
Masquelier

Olivier
Bichler

Auditory Noise Learning

Trevor
Agus

Daniel
Pressnitzer

Memristor Learning Architectures

Damien
Querlioz

Christian
Gamrat

Jean-Philippe
Bourgoin

Resistance to device variability

- **Querlioz, D., Bichler, O., and Gamrat, C. (2011).** Simulation of a memristor-based spiking neural network immune to device variations: IJCNN 2011

- MNIST character data set (60,000 handwritten numerals)
- 28x28 pixel image
- Unsupervised learning
- Each neuron labelled based on its best response

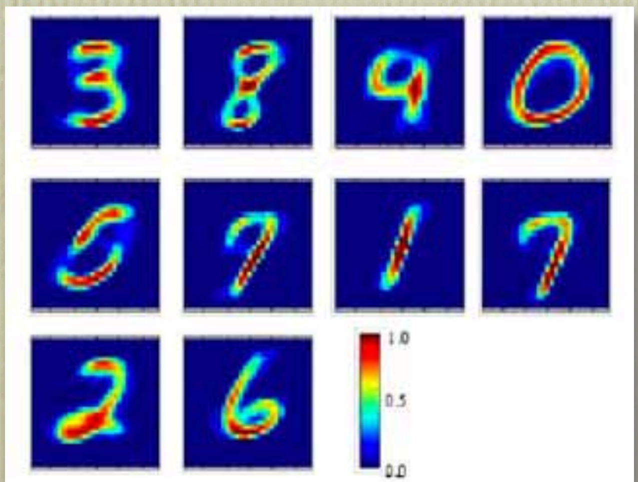

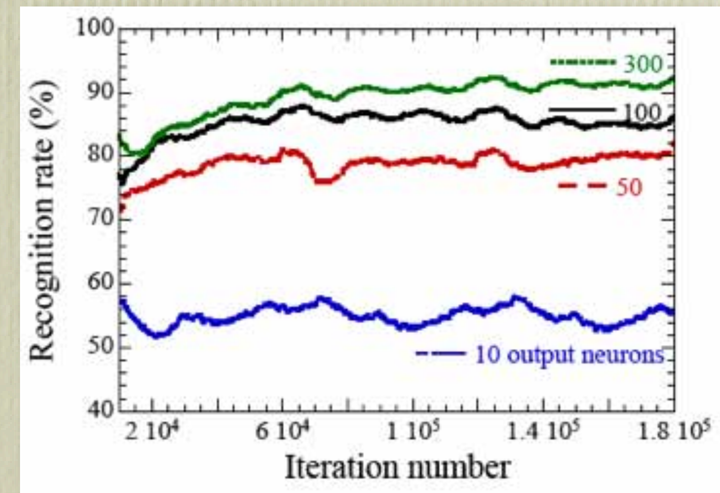

- Remarkable resistance to variability
- 25% no effect