How to build a Cyborg? A brain based architecture for perception, cognition and action

Paul F.M.J. Verschure

Catalan Institute of Advanced Studies (ICREA),

Laboratory of Synthetic Perceptive, Emotive and Cognitive Systems – SPECS

Institute of Audiovisual studies (IUA),

University Pompeu Fabra

Barcelona, Spain

paul.verschure@upf.edu

specs.upf.edu

Robocop

Halo 3:
Master Chief Petty Officer
John-117

Ghost in the Shell (攻殻機動隊, Kōkaku Kidōtai), *Masamune Shirow* (1989).
Motoko Kusanagi: a cyborg employed as the squad leader of Public Security Section 9, of the Japanese National Public Safety Commission.

Cyborgs and Space (1960), by Manfred E. Clynes and Nathan S. Kline, Astronautics.

Manfred Clynes: **Cyborgs** are a "new frontier", "not merely space, but more profoundly the relationship between 'inner space' to 'outer space' – **a bridge...between mind and matter.**"

Mean while back on earth iCub is exercising

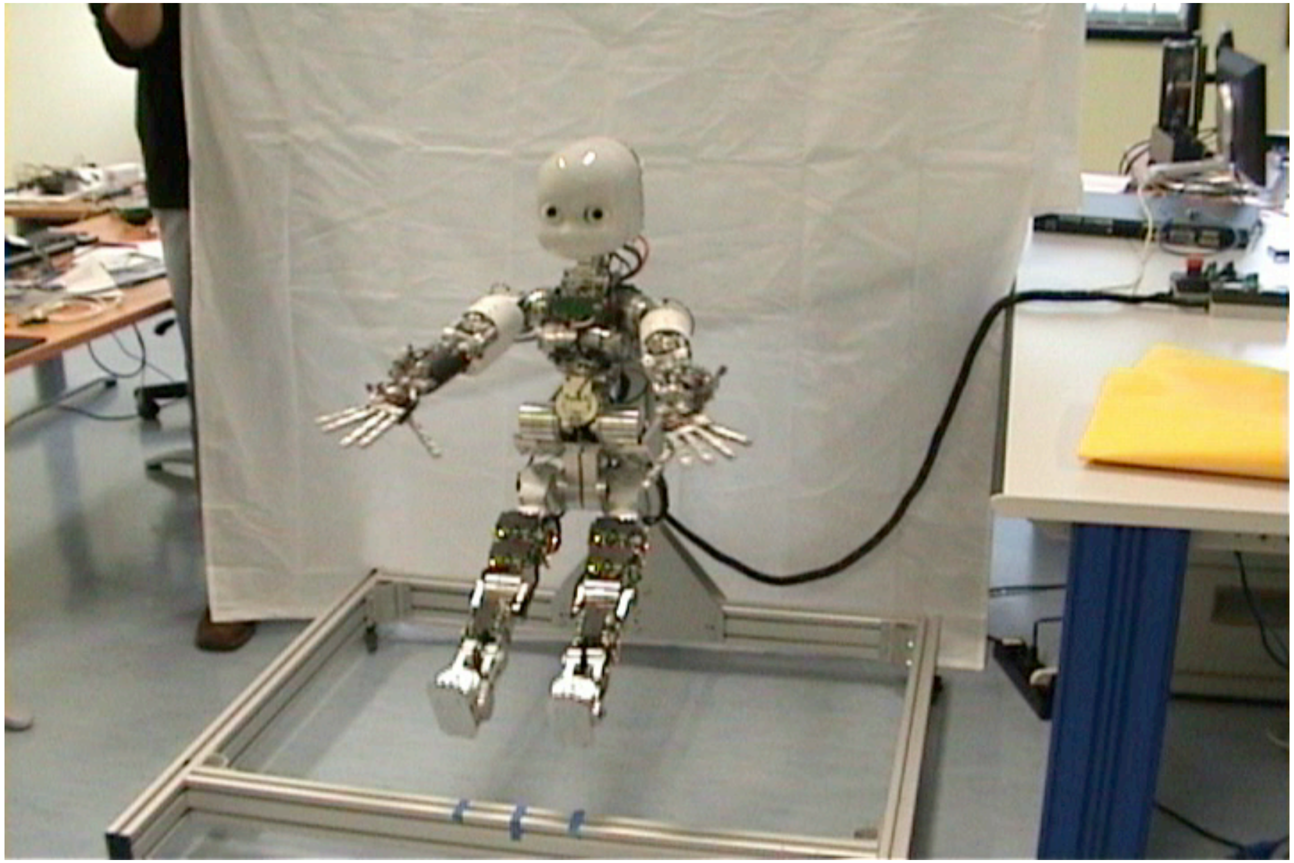

RobotCub.org

How to create a cyborg?

Hardware replacement

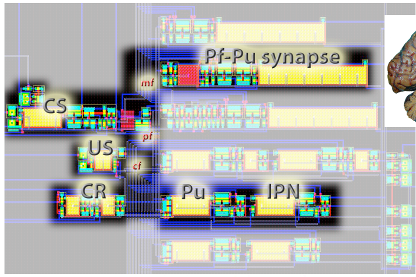

Silicon Cerebellum

Distributed Adaptive Control

Ada at Expo.02

Body

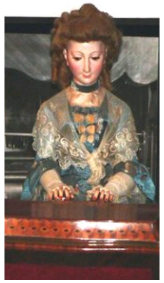

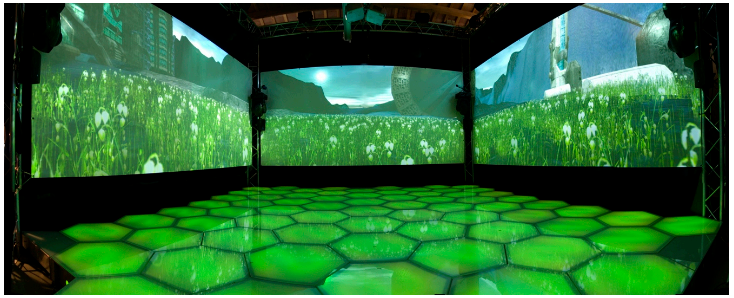

The Experience Induction Machine – XIM

Sensory stimulation

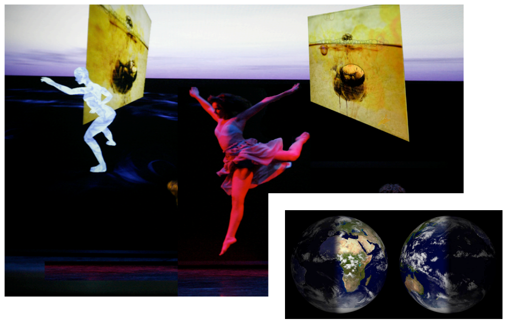

Mixed reality performance re(PER)curso

The Rehabilitation Gaming System - RGS

The brain

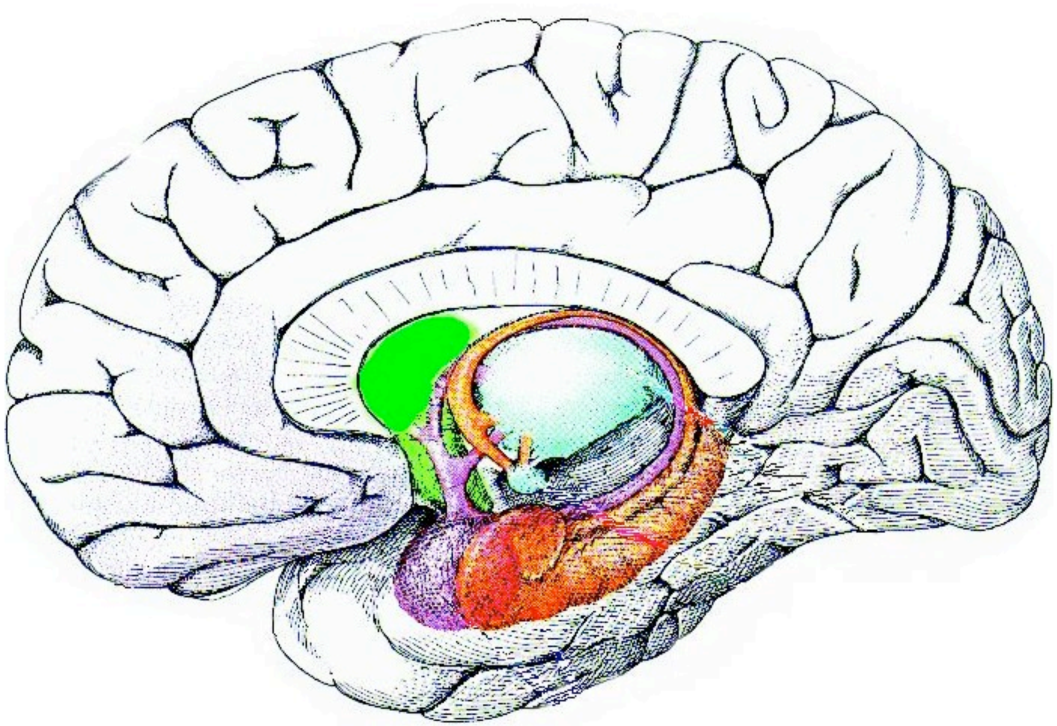

The brain contains 100.000.000.000.000 neurons, 3.200.000.000 km of wires, 1.000.000.000.000.000 connections, weighs 1.5 kg and uses 10 Watt of energy

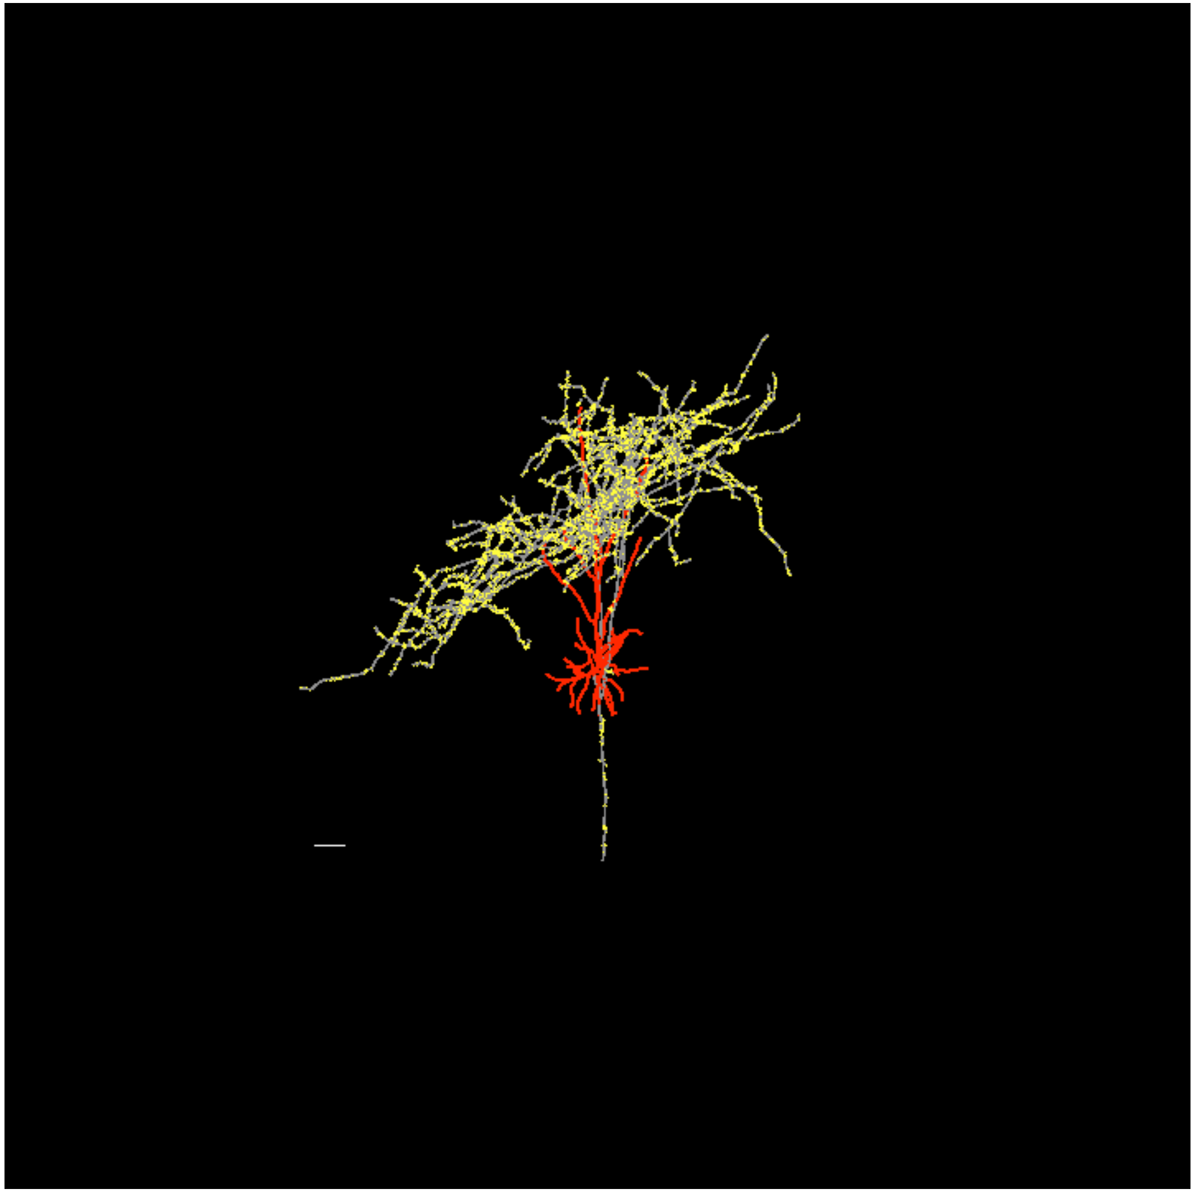

Reconstructed layer Pyramidal neuron of cat visual cortex (courtesy Binziger, Anderson & Martin - INI, Zurich)

some things you can do with neurons

A typical problem...

- Sensory information is local and noisy
- „landmarks“ must be identified
- Reward is intermittent
- How to organize perception and behavior in the face of uncertainty?

Distributed Adaptive Control A multi-layer architecture

Contextual layer

Planning
Operant conditioning

Adaptive layer

Stimulus/Action shaping
Classical conditioning

Reactive layer

Reflex
Autonomic control

DAC is based on the behavioral paradigms of classical and operant conditioning

Verschure et al (2003) Nature (425) 620
Verschure et al (2003) Cogn. Sci. (27) 561
Verschure & Voegtlin (1998) Neural Netw
Verschure et al (1991) Rob. Aut. Sys.

The scientific study of learning & memory starts here

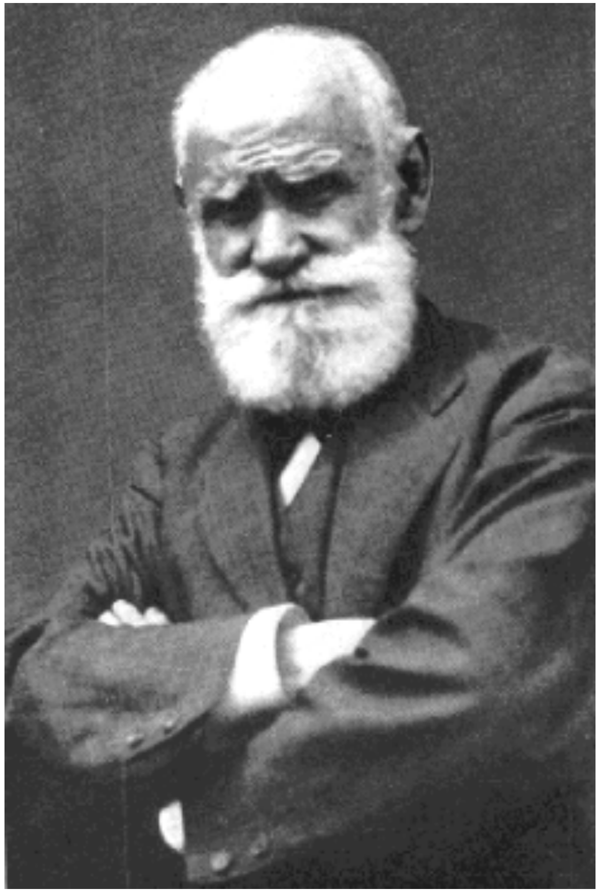

I. Pavlov (1849-1936)

The study of the psychic reflex

**``The dog sees, hears and sniffs all these things, directs his attention to them, tries to obtain them if they are eatable or agreeable, but turns away from them and evades their introduction into the mouth if they are undesired or disagreeable. Every one would say that this is a psychical reaction of the animal, a psychical excitation of the salivary glands. How should the physiologist treat such facts? How can he state them, how analyze them? What are their common and what their individual
characteristics? To understand these phenomena, are we obliged to enter into the inner state of the animal, and to fancy his feelings and wishes as based on our own? For the investigator, I believe there is only one possible answer to the last question - an absolute ``No".``**

Pavlov (1928) Lectures on Conditioned Reflexes, pp. 50

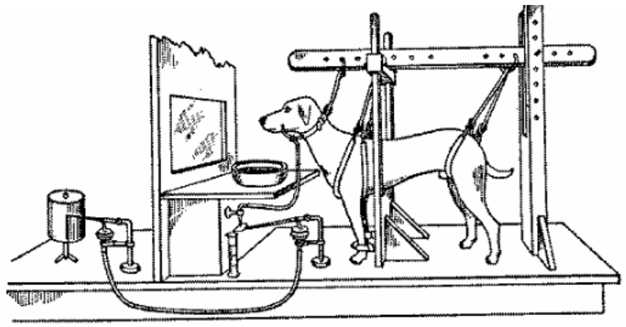

Pavlov's Classical Conditioning

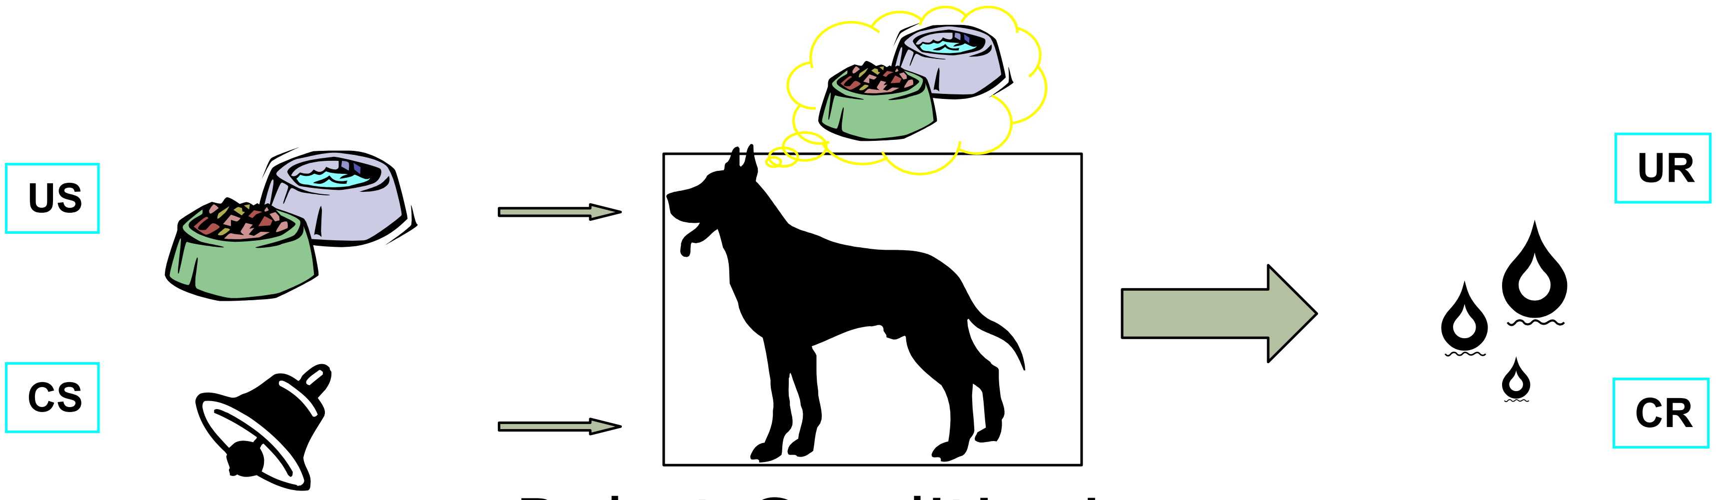

Robot Conditioning

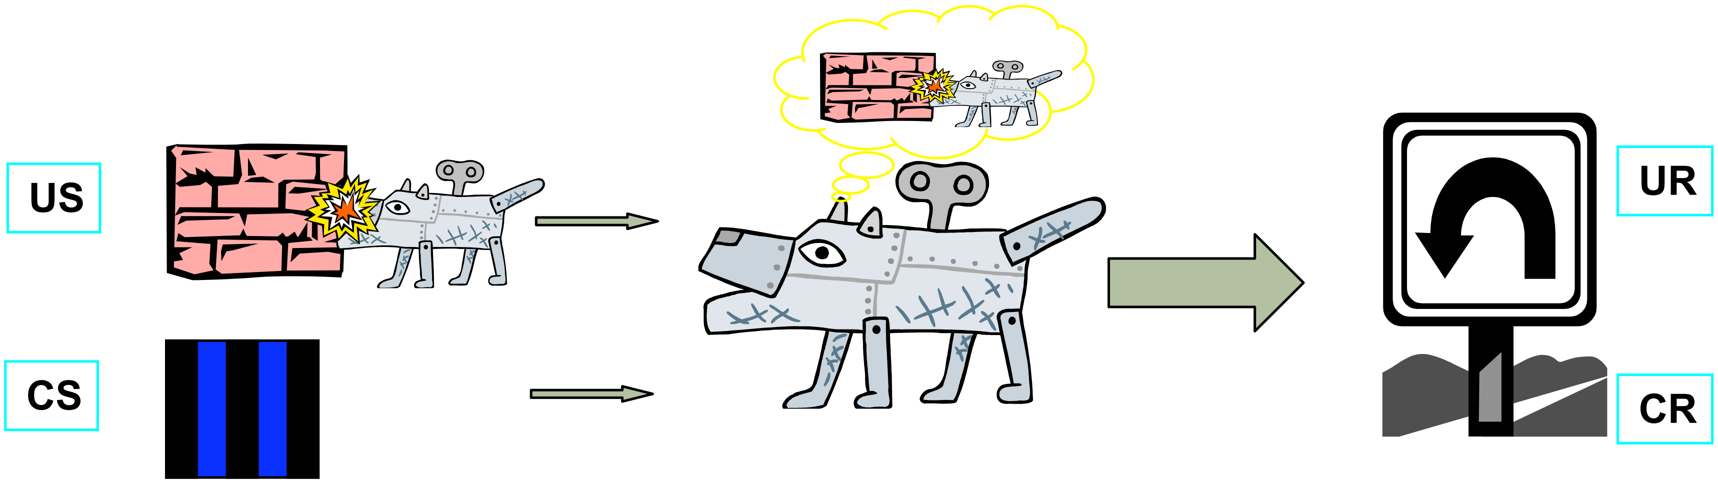

CS: Conditioned stimulus
US: Unconditioned stimulus
CR: Conditioned response
UR: Unconditioned response

Classical conditioning CS substitutes the US in invoking the CR
In the DAC framework it provides for:

- adaptive stimulus identification
- action shaping

Kamin's blocking

Training	Test:
CS_A + US	$CS_A \rightarrow CR$
CS_B + US	$CS_B \rightarrow CR$
CS_A + US	$CS_A \rightarrow CR$
CS_A + CS_B + US	$CS_B \not\rightarrow CR$

Learning to respond to CS_A "blocks" learning to CS_B

Overexpectation

Training

CS_A + US
CS_B + US

Test:

$CS_A \longrightarrow CR$
$CS_B \longrightarrow CR$

Initially:
$CS_A + CS_B$ + US $\quad CS_A + CS_B \longrightarrow CR$

After additional presentations:
$CS_A + CS_B$ + US $\quad CS_A + CS_B \not\longrightarrow CR$

Two effective CSs presented together causes extinction

Rescorla-Wagner laws of associative competition (1972)

$$\Delta V = \alpha \beta(\lambda - V)$$

α = salience of CS
β = strength of the US
V = the current associative value of all CSs paired with this US
λ = maximum associative value given the US (100%)

As learning proceeds, associative value V increases and new associative value decreases
With two CSs A and B the combined association is: $V_{A+B} = V_A + V_B$

Formalizes Kamin's blocking

Animals only learn when events violate their expectations

Micro robot foraging

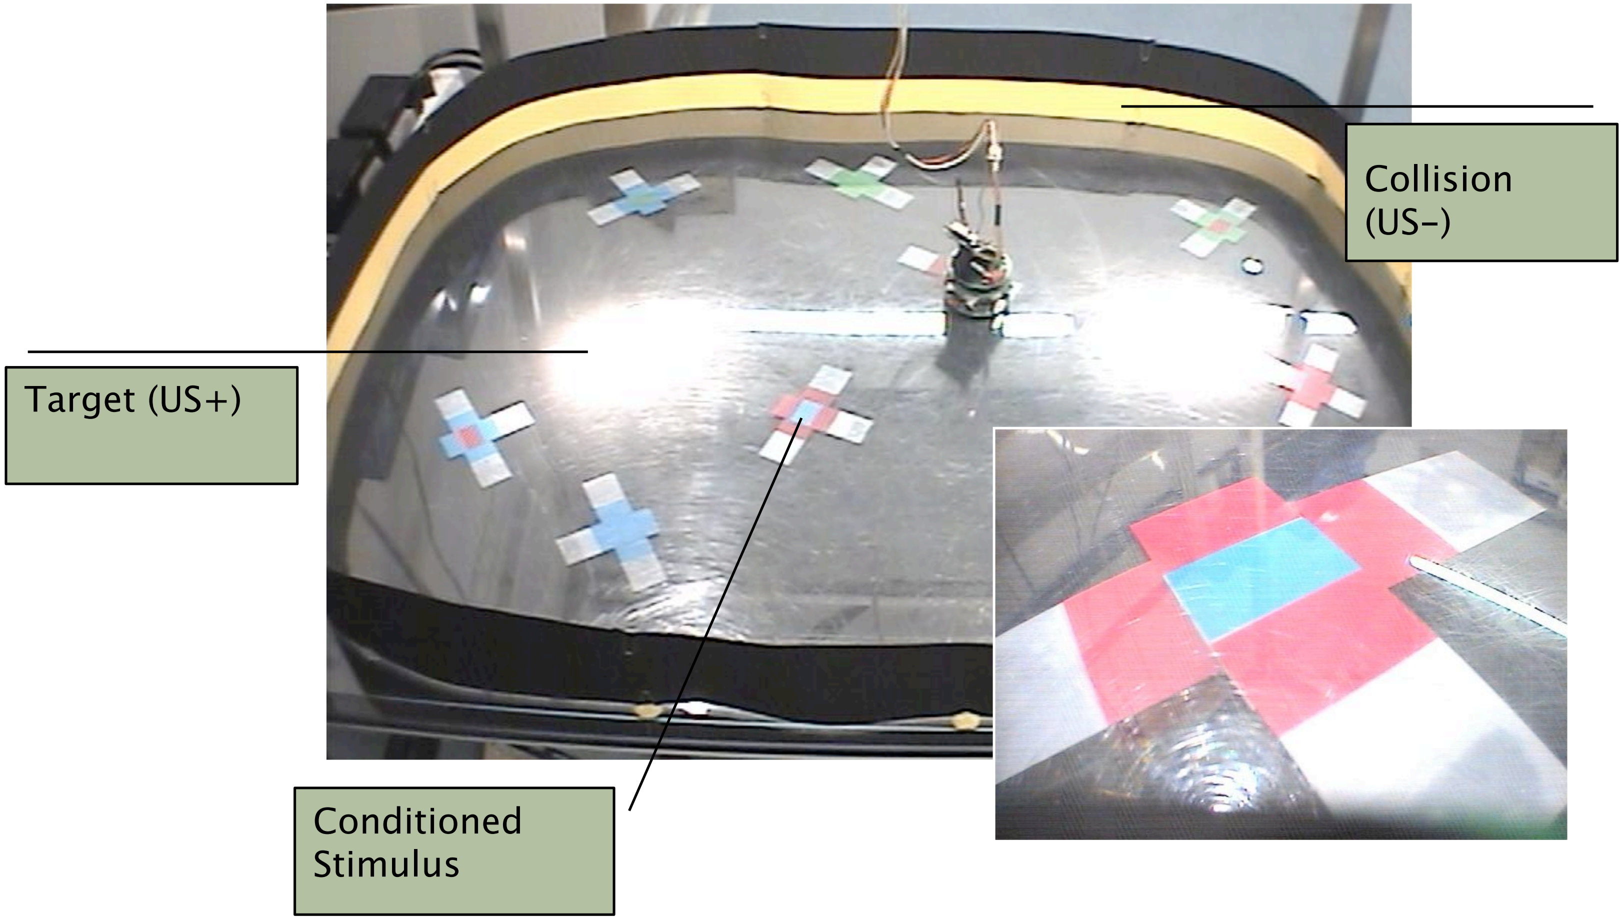

DAC OVERVIEW

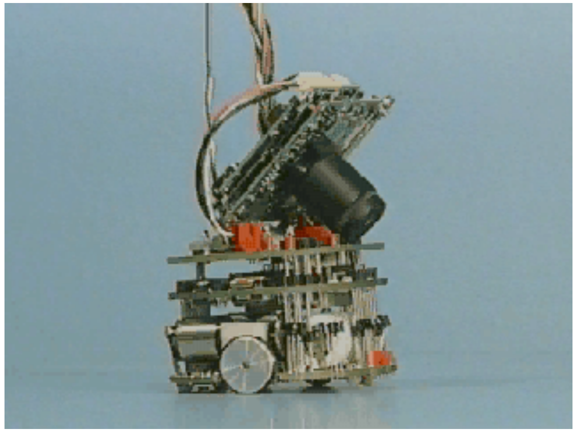

ROBOT

camera

sensors

CS: color

US: Light (US+)

Collision (US-)

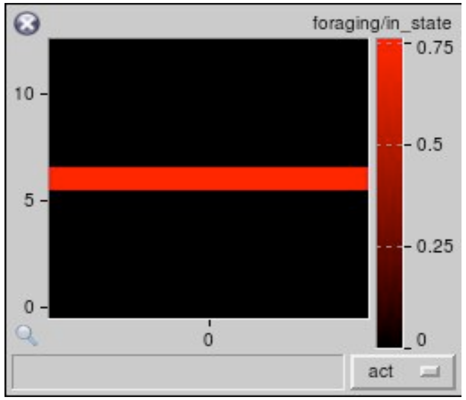

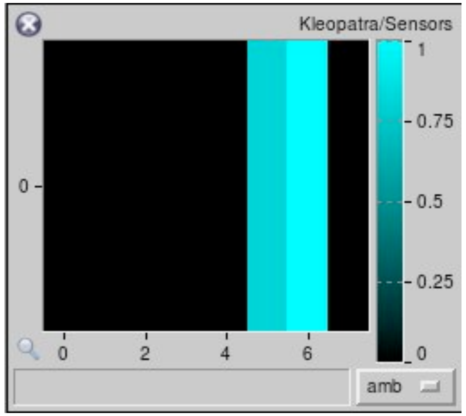

ambient

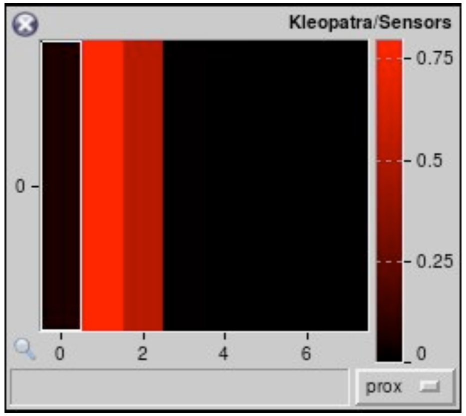

proximity

DAC: Reactive layer

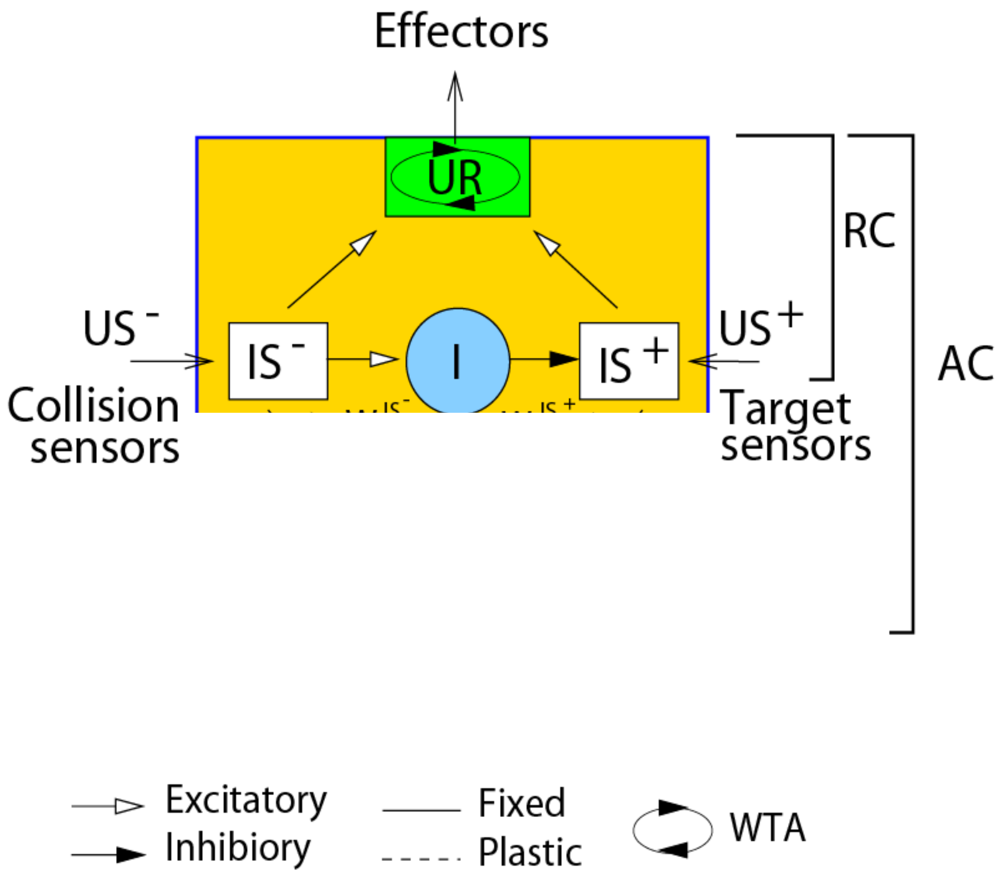

Reactive layer performance

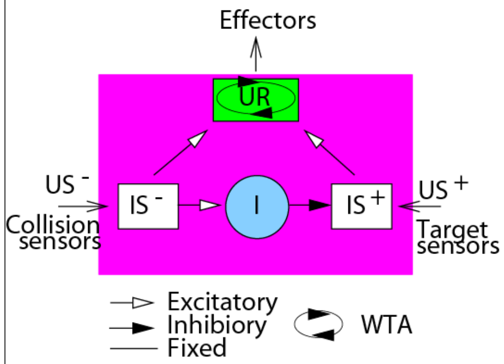

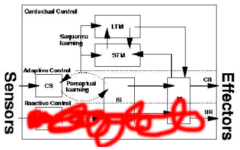

DAC: Adaptive layer

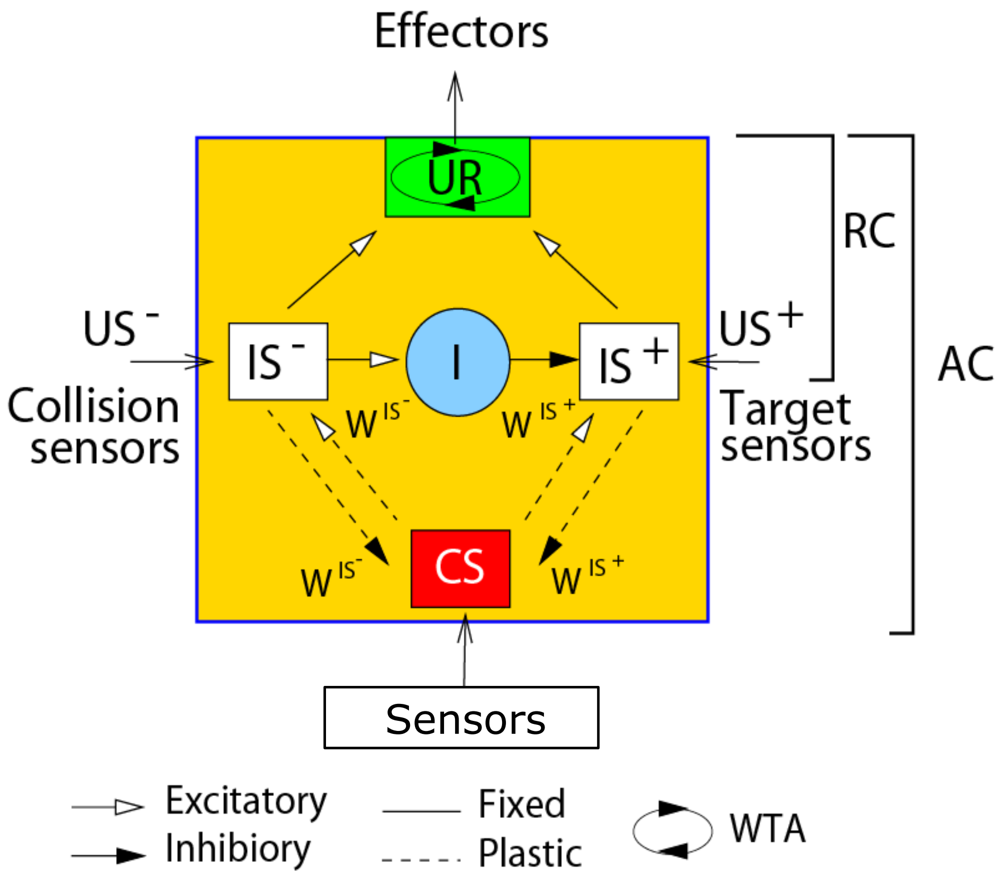

DAC: Reactive/Adaptive layer

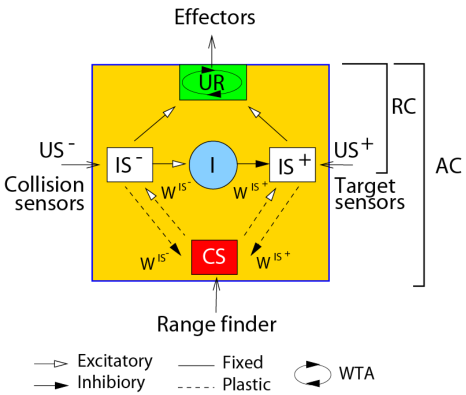

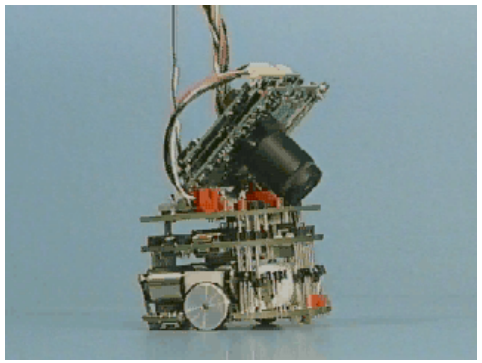

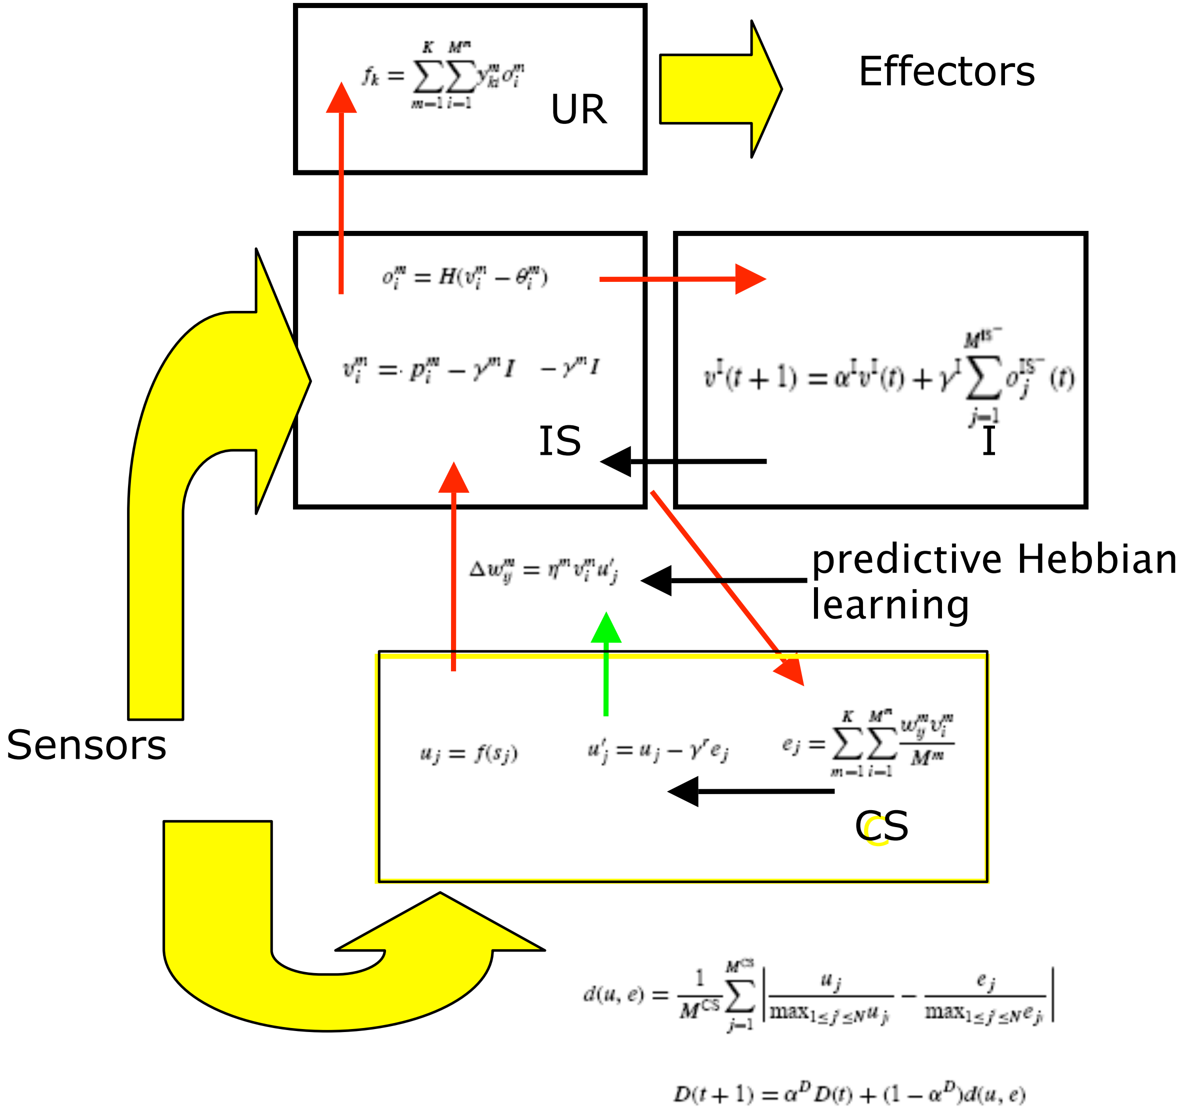

Adaptive layer performance

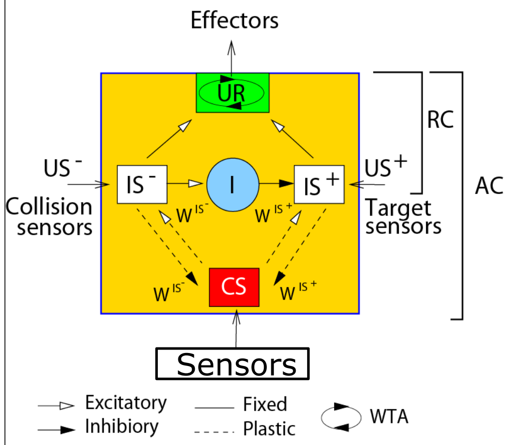

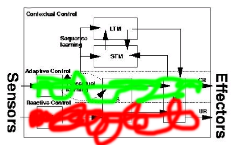

DAC: Contextual control

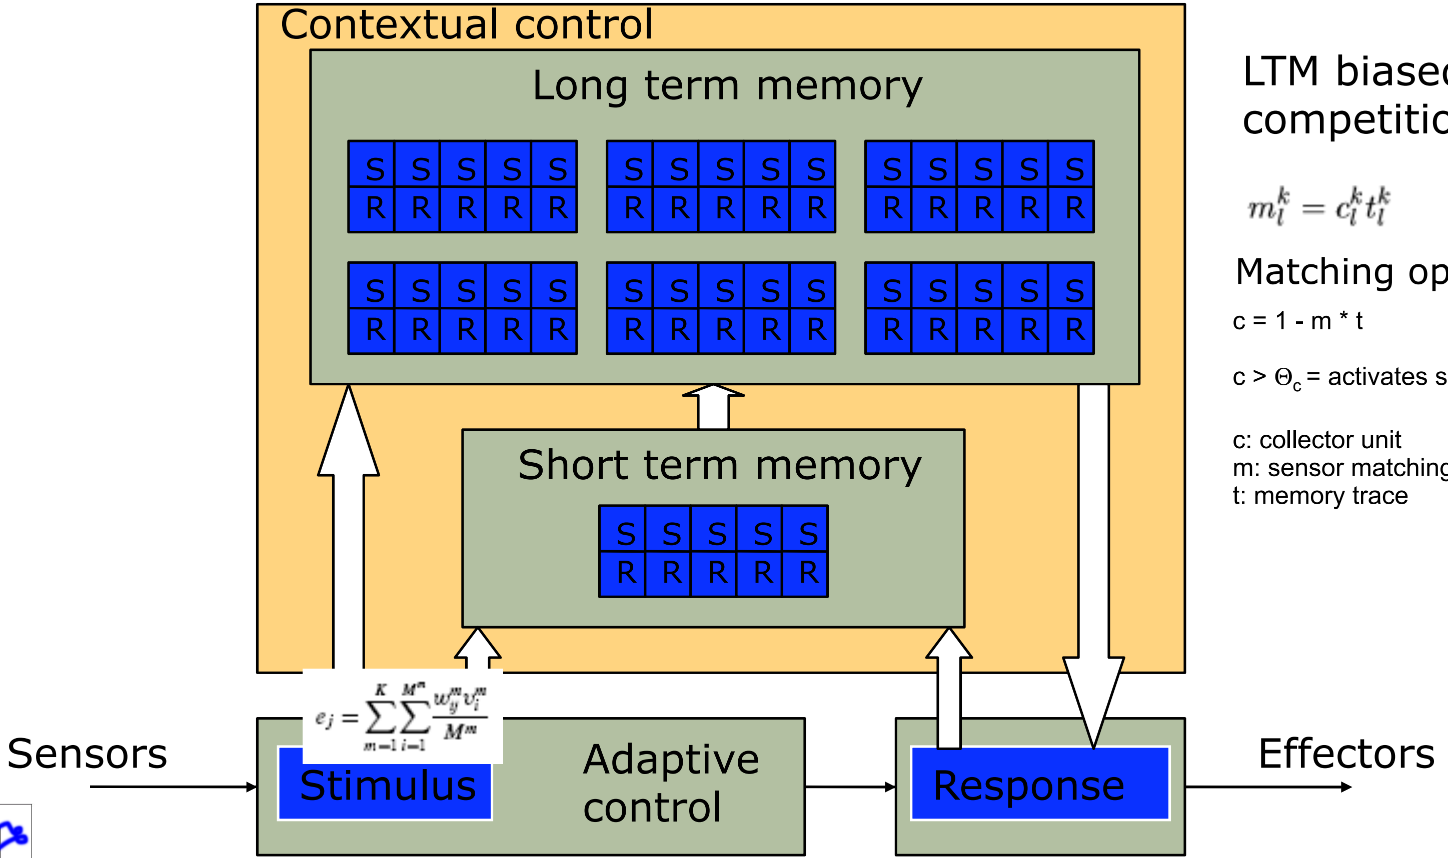

LTM biased competition:

$$m_l^k = c_l^k t_l^k$$

Matching optimizes:

c = 1 - m * t

c > Θ_c = activates segment

c: collector unit
m: sensor matching
t: memory trace

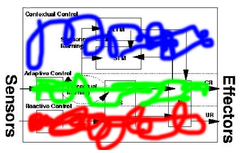

Verschure & Voegtlin (1998) Neur.Netw. Verschure & Althaus (2003) Cog.Sci., 27: 561-590

Optimal Bayesian decision making requires a simpler model (DAC5)

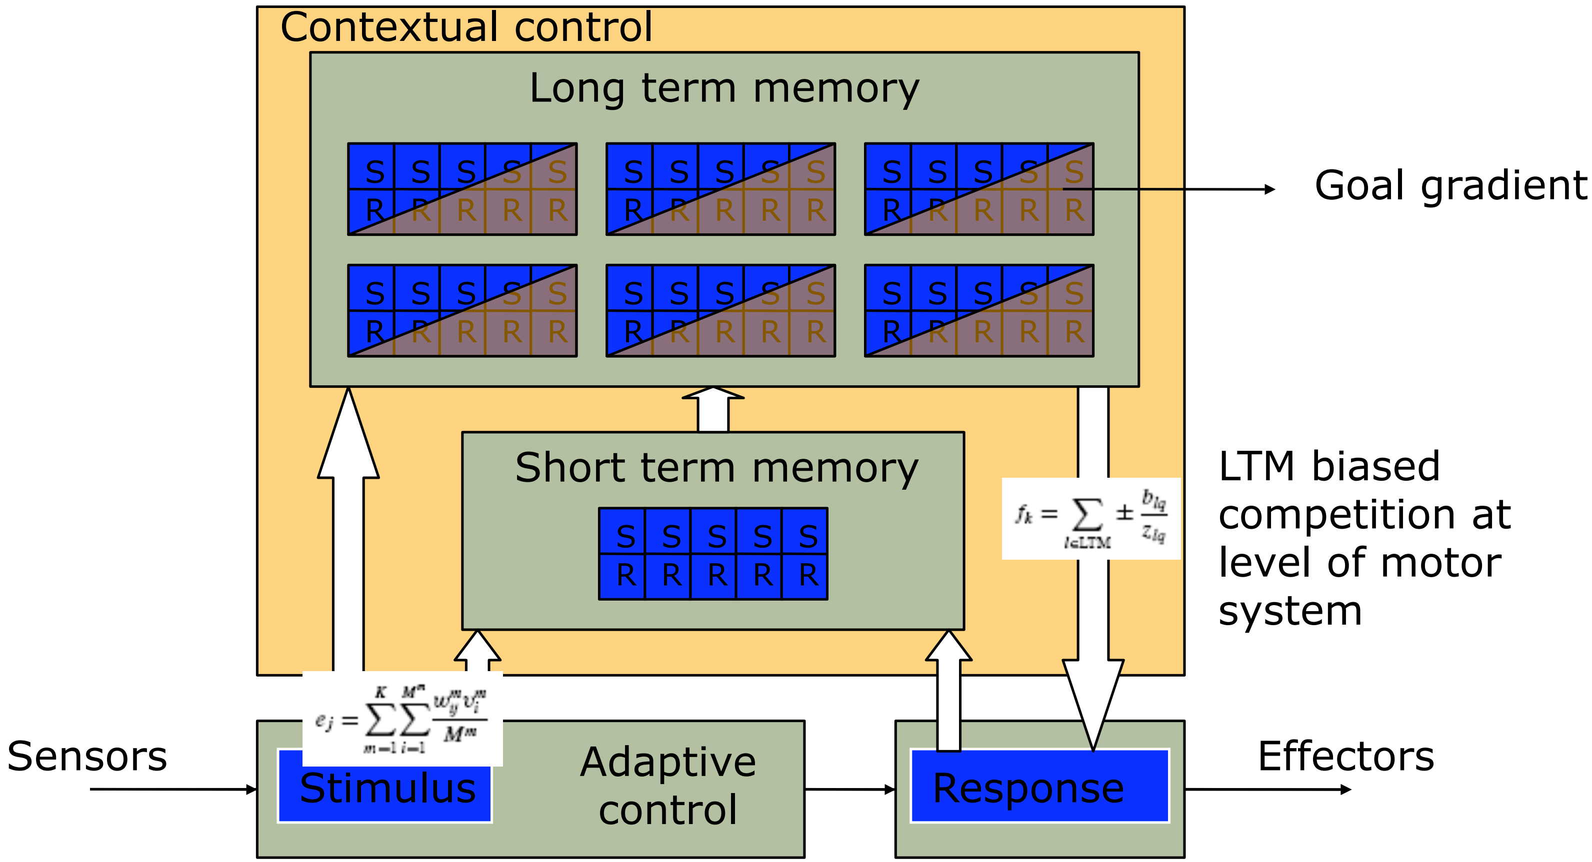

Verschure & Althaus (2003) Cog.Sci., 27: 561-590

Contextual control performance

LTM use

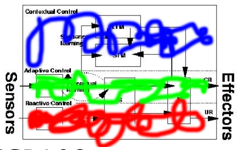

Adaptive (Disabled) VS Contextual (Enabled)

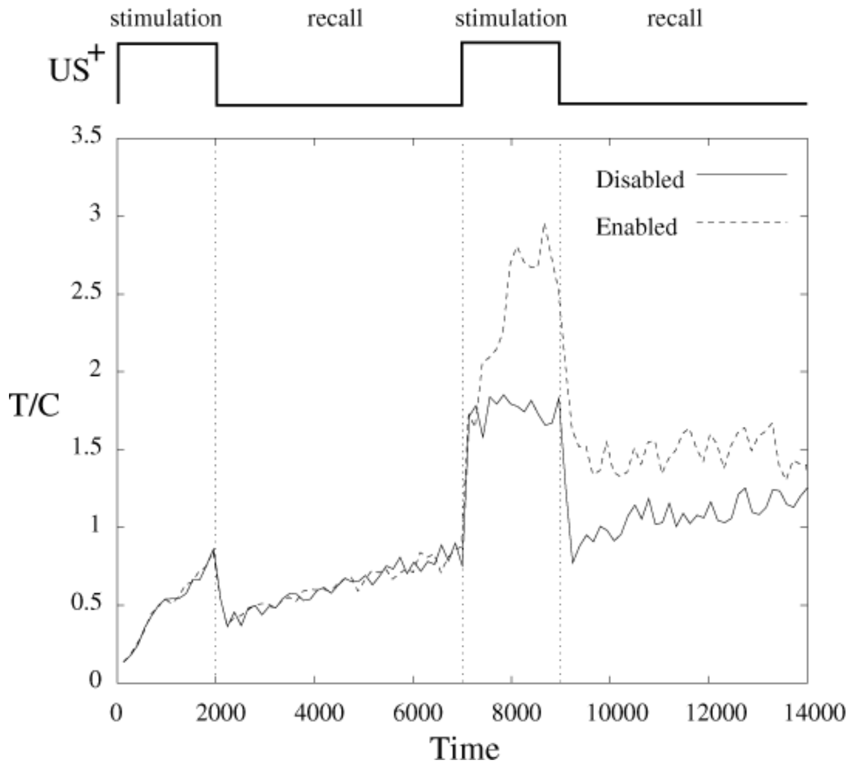

BugWorld: 1000 robots per condition

Verschure et al (2003) Nature

DAC: Adaptive VS Contextual II

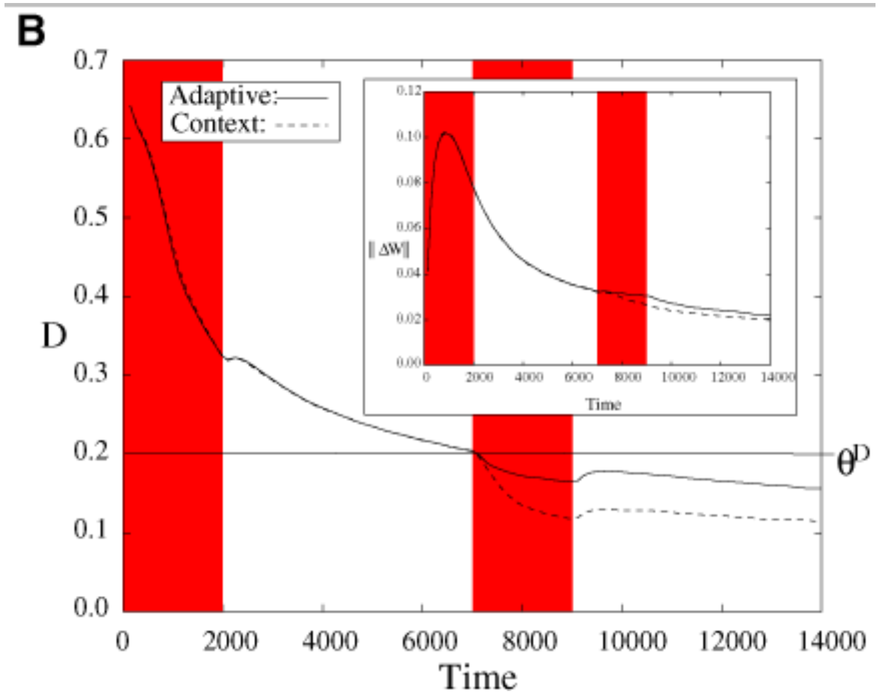

BugWorld: 1000 robots per condition

$$d(u, e) = \frac{1}{M^{CS}} \sum_{j=1}^{M^{CS}} \left| \frac{u_j}{\max_{1 \le j \le N} u_j} - \frac{e_j}{\max_{1 \le j \le N} e_j} \right|$$

Verschure et al (2003) Nature

$$D(t+1) = \alpha^D D(t) + (1 - \alpha^D) d(u, e)$$

BUT No internal feedback from CC to AC

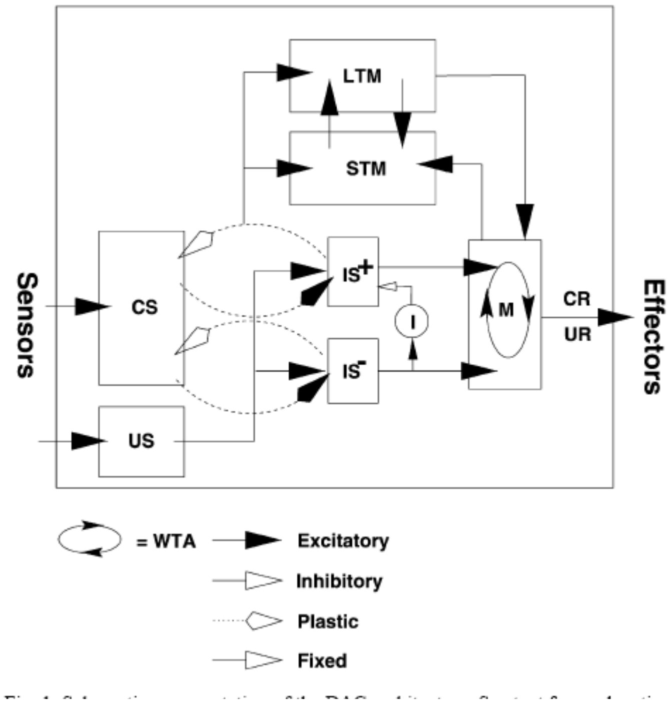

Reorganization of the adaptive layer must be the result of changes in overt behavior

Quantifying the effect of behavior on perception

Adaptive

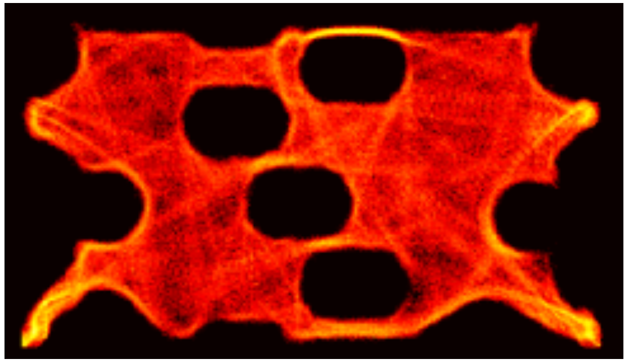

Contextual

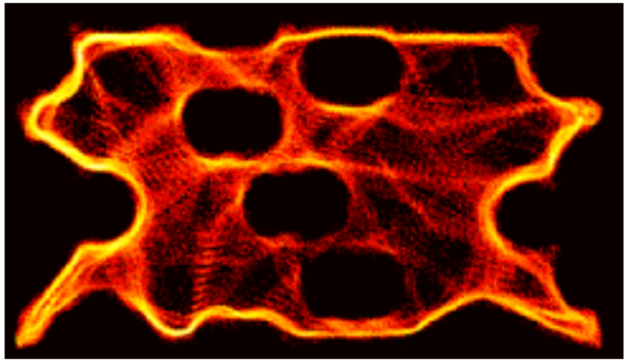

BugWorld: 1000000 timesteps per condition

Verschure et al (2003) Nature

Behavioral and stimulus sampling entropy

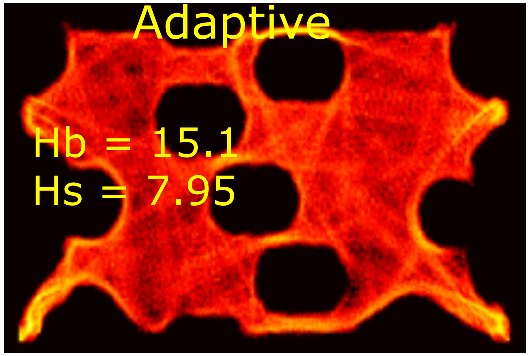

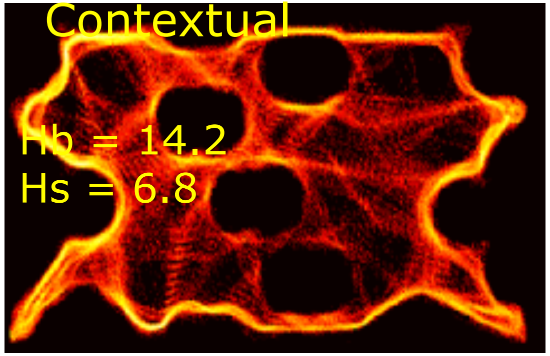

Reactive:
Hb = 15.4
Minimal:
Hb = 11.2

$$H = -\sum_{a \in E} p(a)\ \log_2\ p(a)\ \text{with}\ \sum_{a \in E} p(a) = 1$$

BugWorld

Verschure et al (2003) Nature

Generalization to the real world: Khepera recall tests

Adaptive

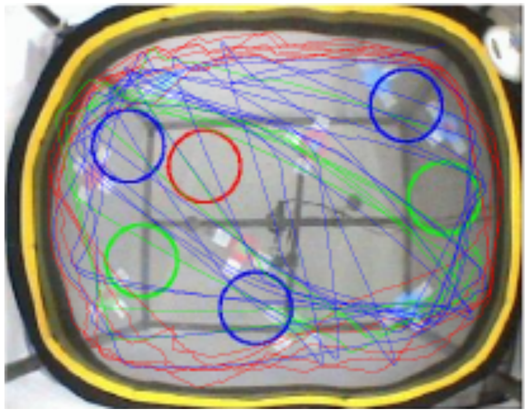

Contextual

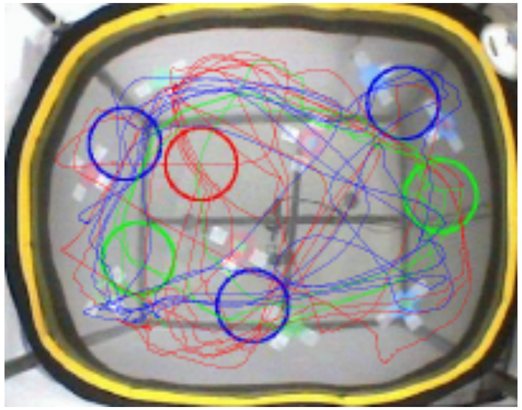

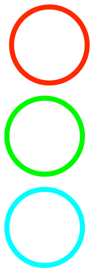

Khepera-IQR421

Verschure et al (2003) Nature

Quantification of behavioral structuring in the real world: Markov model

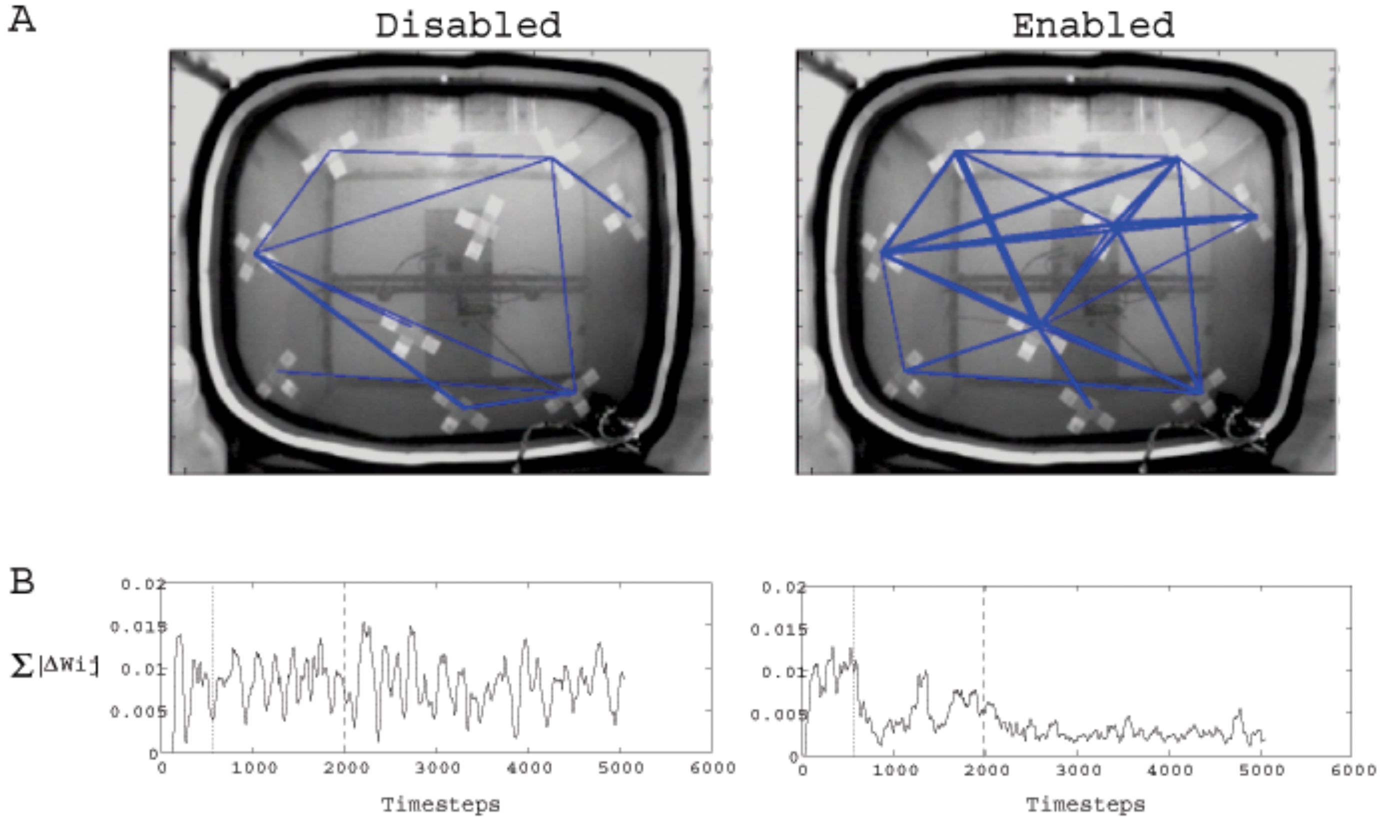

Verschure et al (2003) Nature

Behavioral feedback

- Behavioral feedback affects neuronal organization:
  - Perceptual learning (AC) -> Contextual control
  - Behavioral learning (CC) -> Structures behavior
  - Change input sampling -> Reduced input space
  - Adaptation perceptual structures to behavior

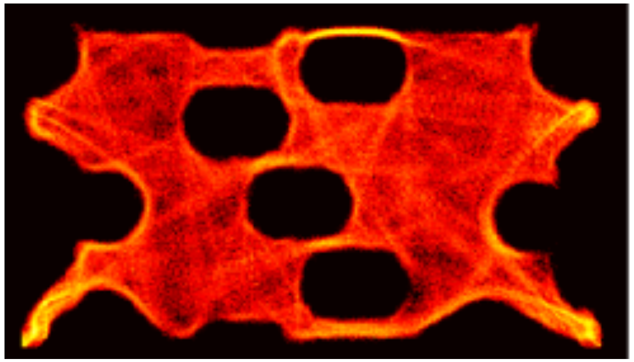

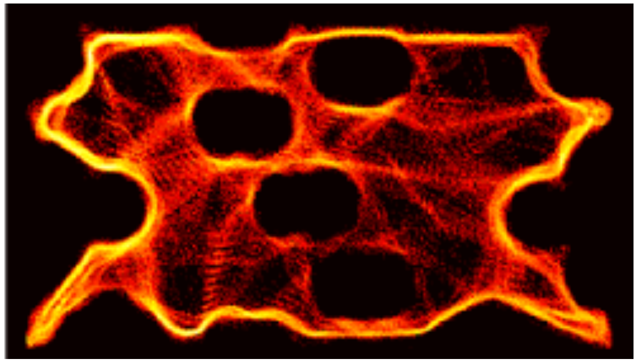

Macroscopic behavioral properties of a real-world system can cause changes of its microscopic neuronal organization

Behavioral feedback

Bottom line (s)

- DAC proposes a framework how multiple systems in the brain work together to generate perception, cognition and behavior
- DAC has identified a novel feedback loop in the structuring of the neuronal substrate: behavioral feedback

Adaptive layer

- a model of classical conditioning

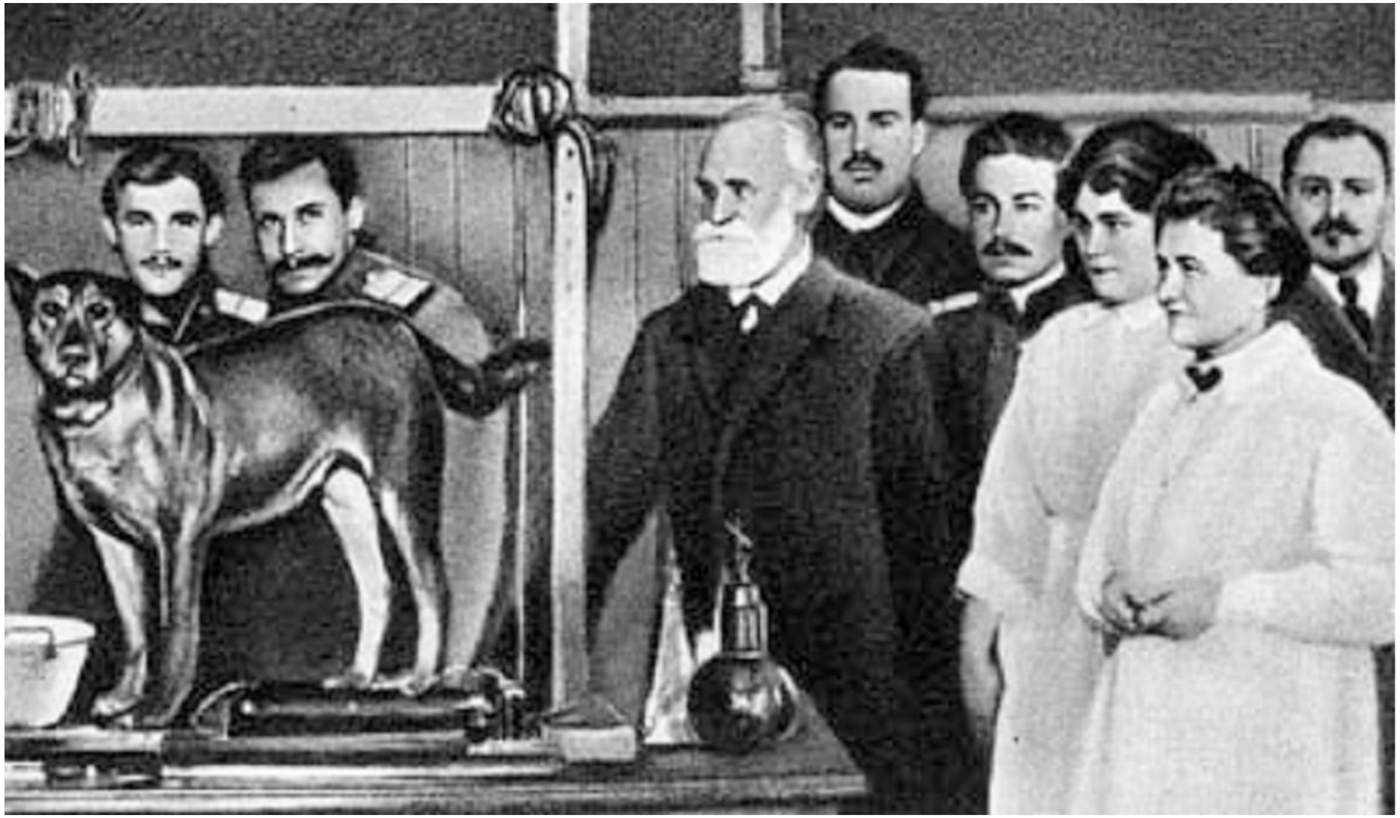

Ivan Pavlov (1849-1936)

From DAC to the brain

Adaptive layer

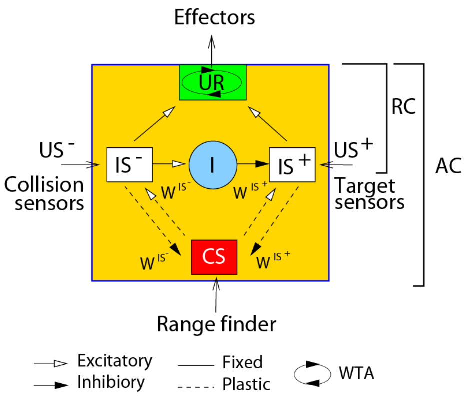

Classical conditioning

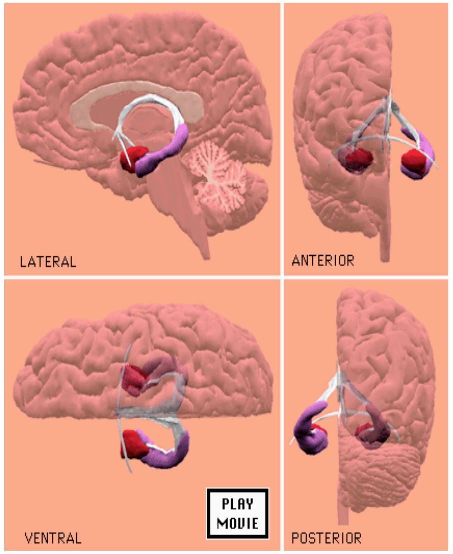

Laws of conditioning

- Rescorla & Wagner, 72

$$V_{ab} = V_a + V_b$$

$$\Delta V_i = \alpha_{cs}\, \gamma_{us}\, (\lambda - \Sigma_j V_j)$$

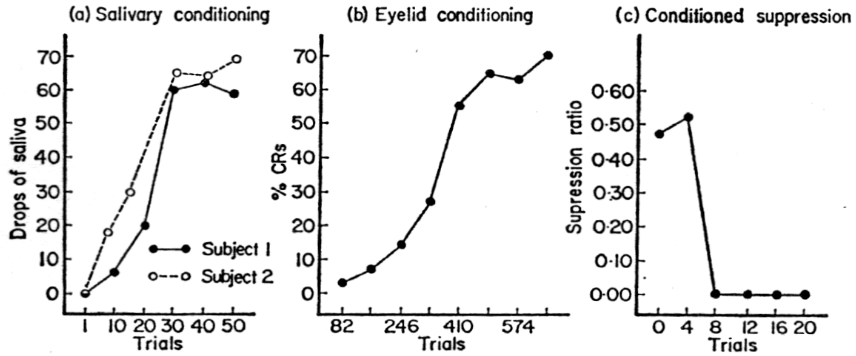

non-specific specific

Konorski's 2-phase theory of conditioning distinguishes a fast non-specific learning system from a slow specific one

Animals only learn when events violate their expectations
BUT
How do neurons develop and express expectations?

The non-specific learning system

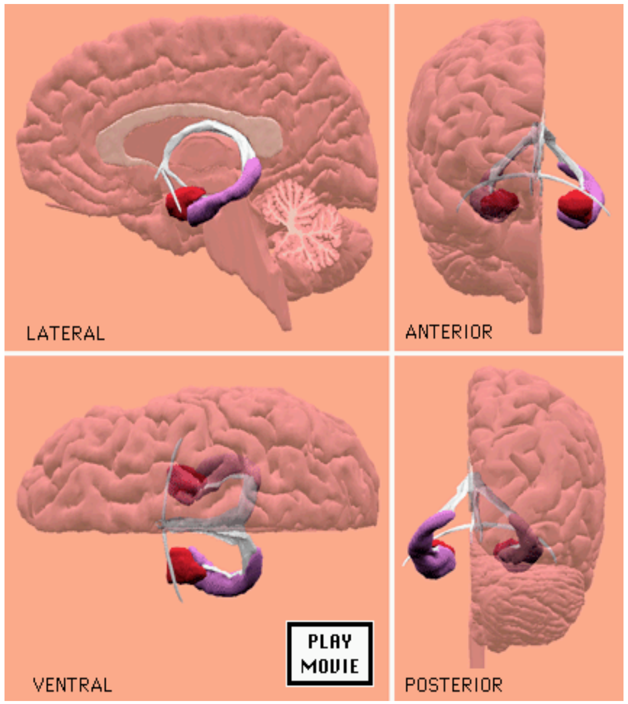

Amygdala circuit

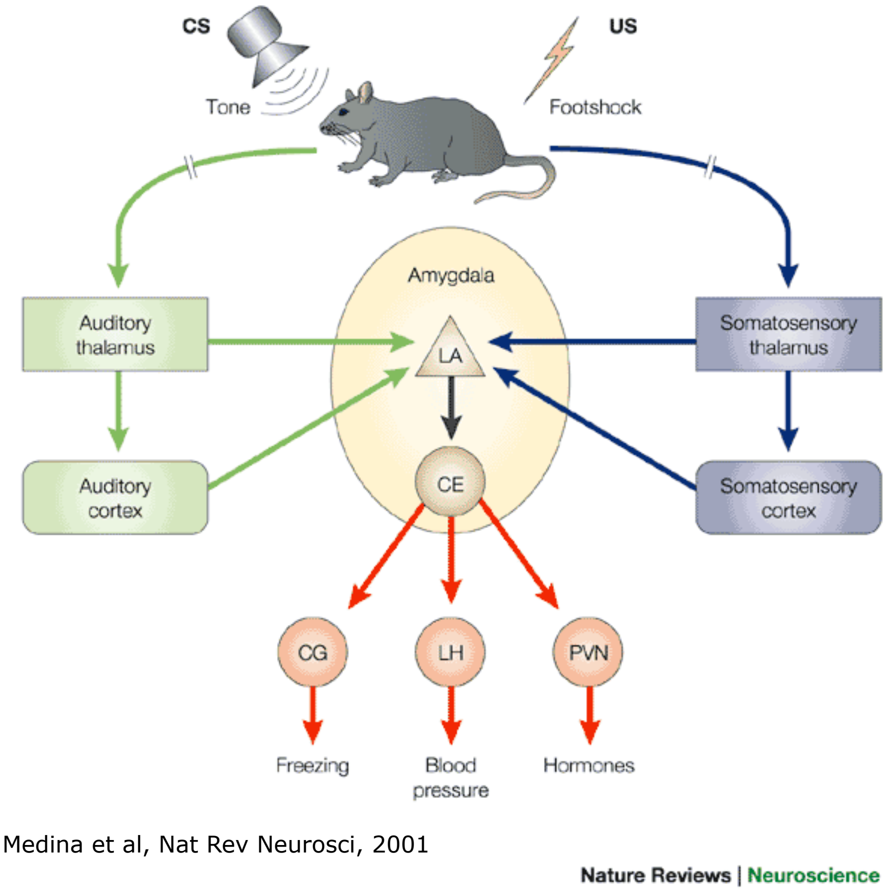

Medina et al, Nat Rev Neurosci, 2001

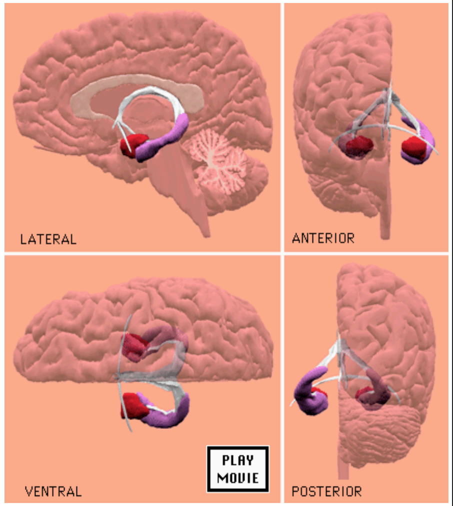

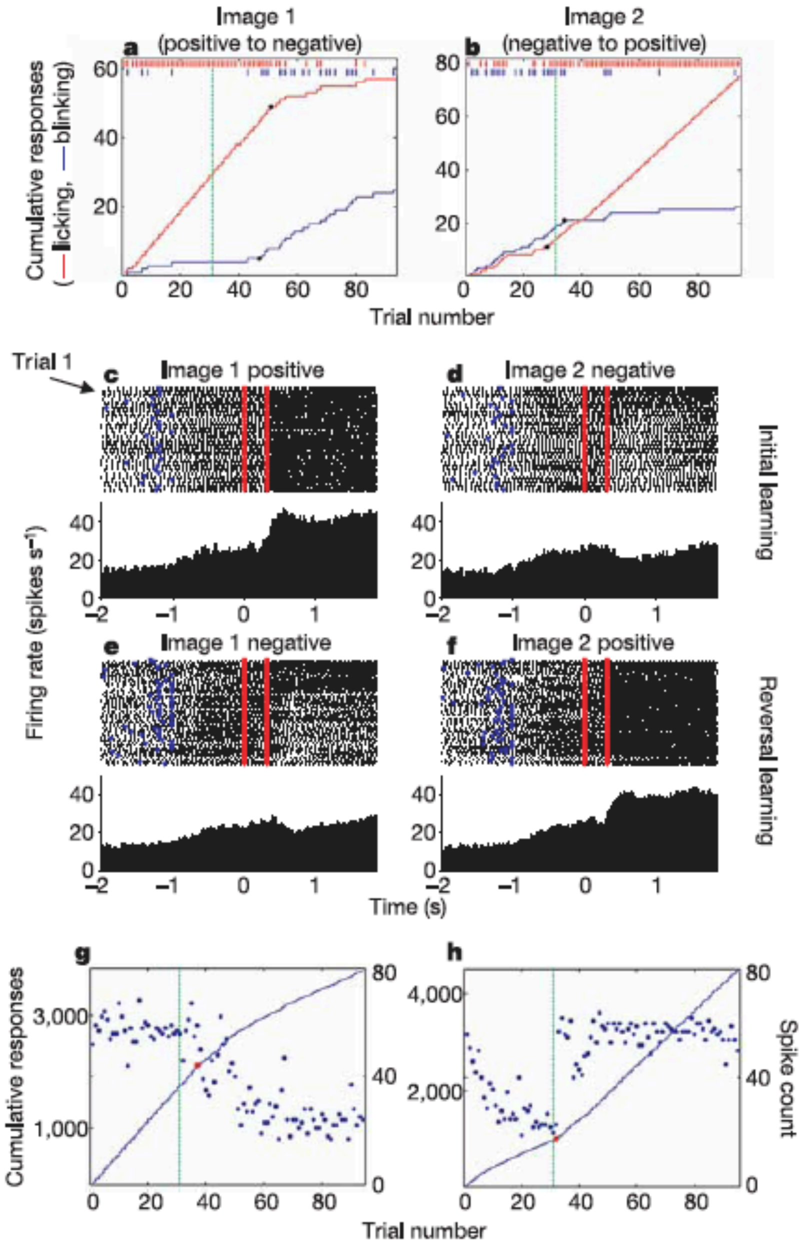

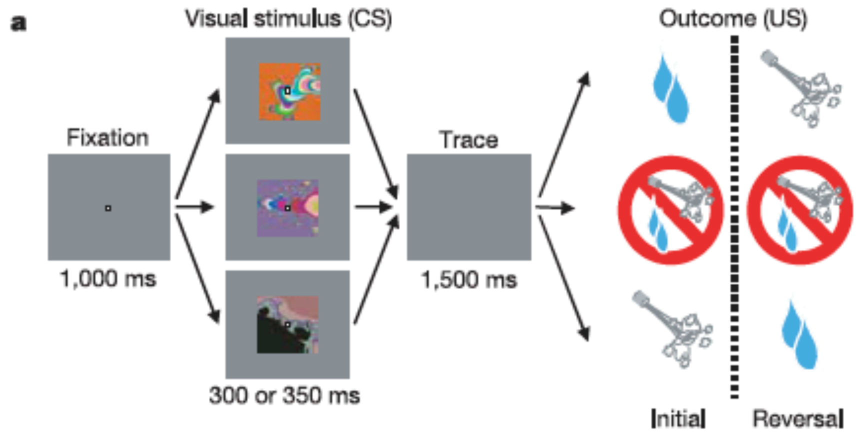

Vol 439|16 February 2006|**doi:10.1038/nature04490**

natur

LETTERS

The primate amygdala represents the positive and negative value of visual stimuli during learning

Joseph J. Paton[1]*, Marina A. Belova[1]*, Sara E. Morrison[1] & C. Daniel Salzman[1,2,3]

Conditioning in primary Auditory cortex (A1)

Naive Trained

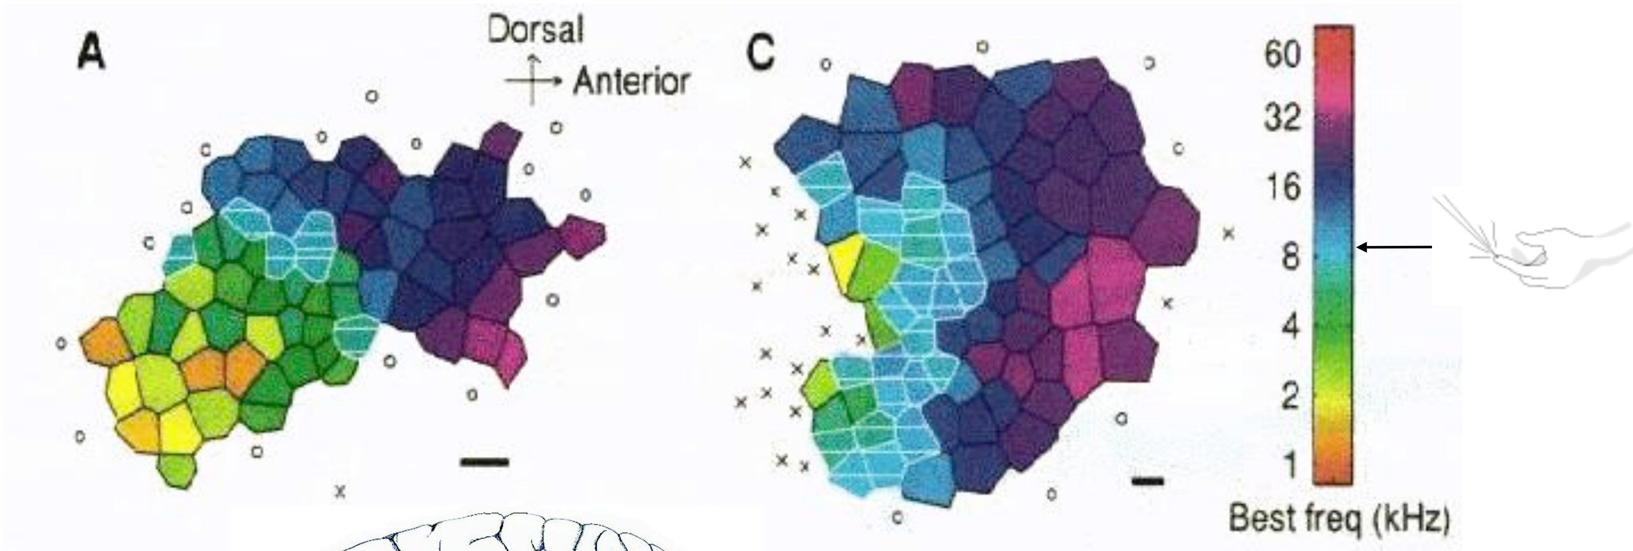

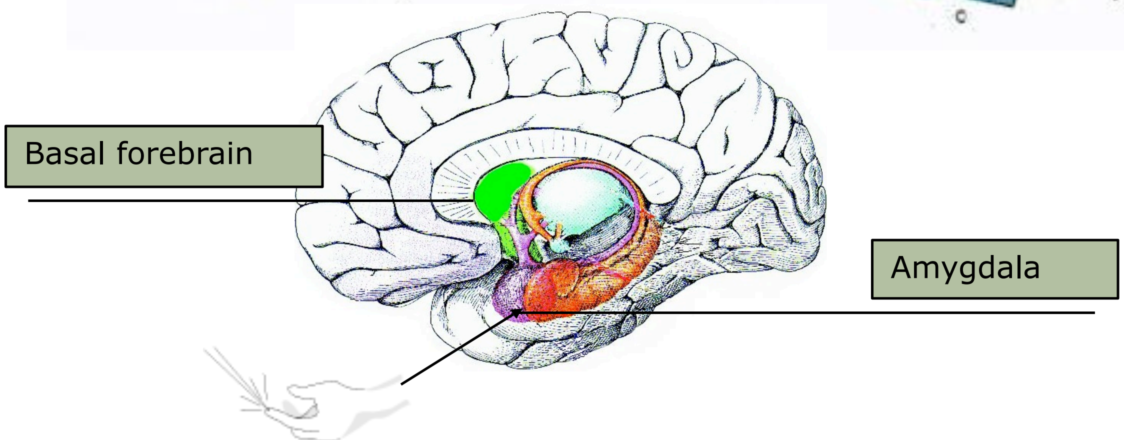

(Kilgard & Merzenich, 1998)

A1: Neuron model

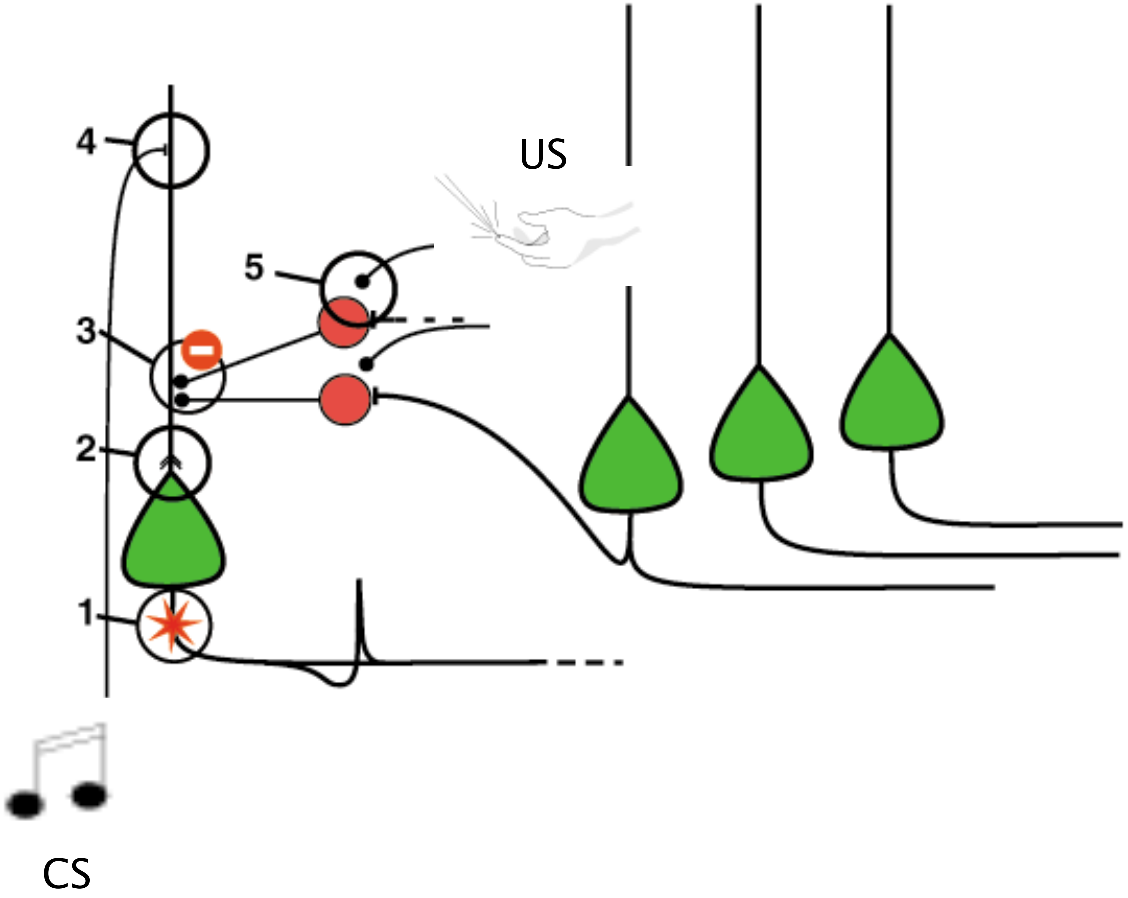

1: action potential (AP)
2: back-propagating AP
3: Shunting inhibition can prevent BAP to reach into the dendrite
4:synaptic plasticity requires coincidence pre-synaptic AP and post-synaptic BAP

Sanchez-Montanes et al (2000/2002) Neural Computation/IEEE N.N.

A1 mode real-world evaluation: Hardware setup

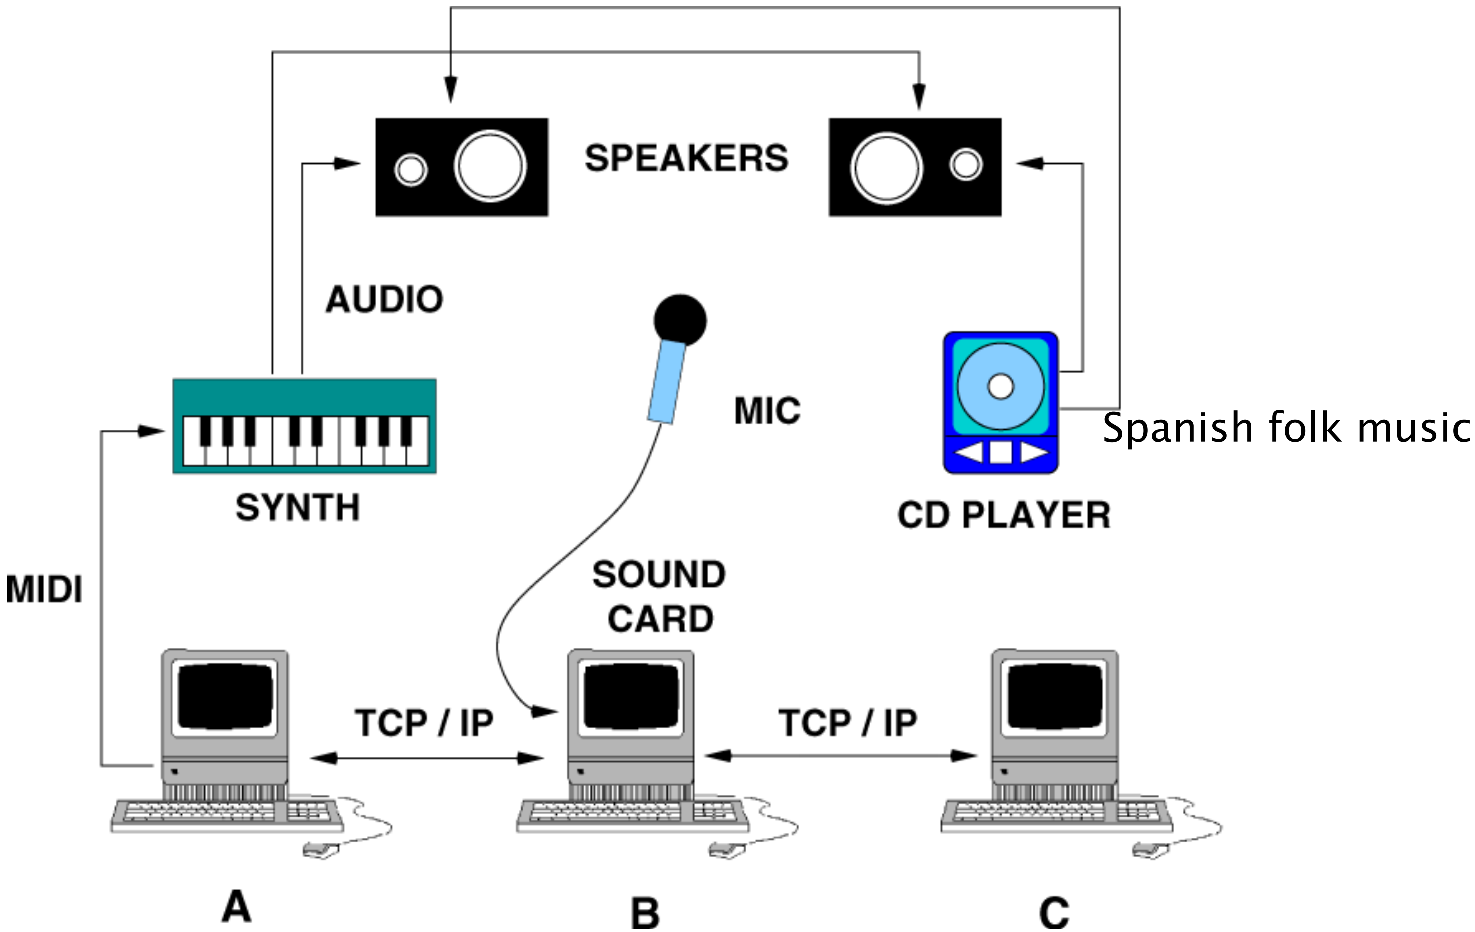

Sanchez-Montanes et al (2000/2002) Neural Computation/IEEE N.N.

A1: Map reorganization

0.74 kHz

2.96 kHz

Naive

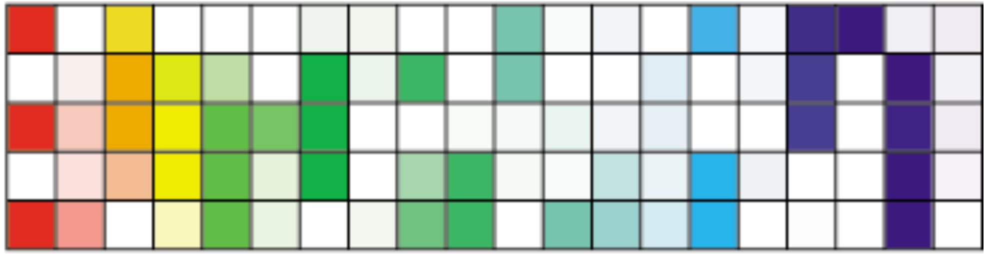

Trained

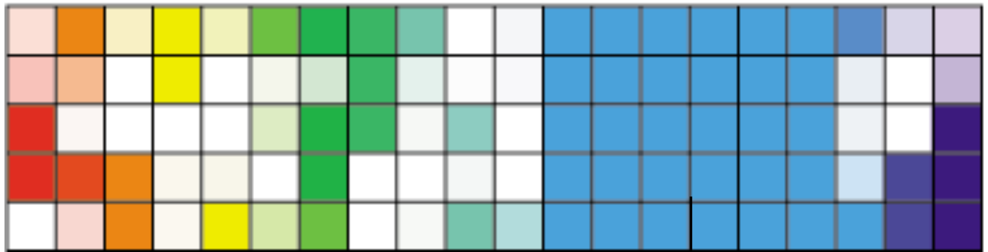

22 CS-US trials

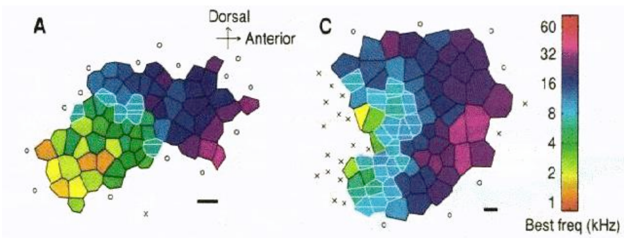

Sanchez-Montanes et al (2000/2002) Neural Computation/IEEE N.N.

Bottom lines

- A bio-physically constrained model of the cerebral cortex can replicate the receptive field changes observed in the primary auditory cortex as a result of classical conditioning
- The model develops both a tonotopic map and represents the CS by dynamically recruiting more neurons to represent the reinforced frequency

The specific learning system of the cerebellum

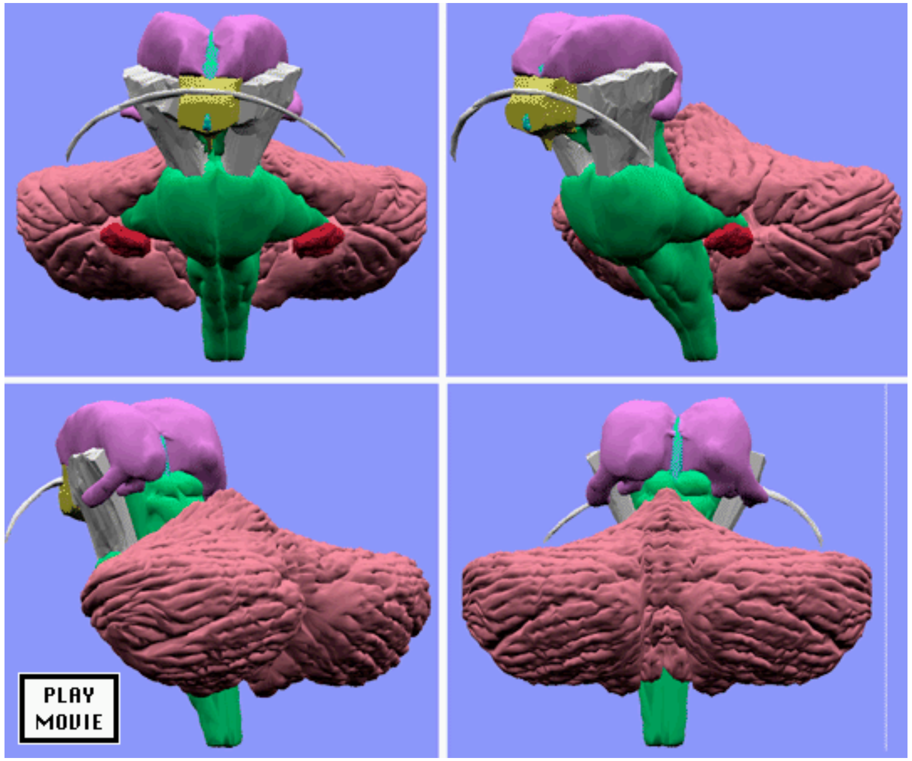

Cerebellar cortex

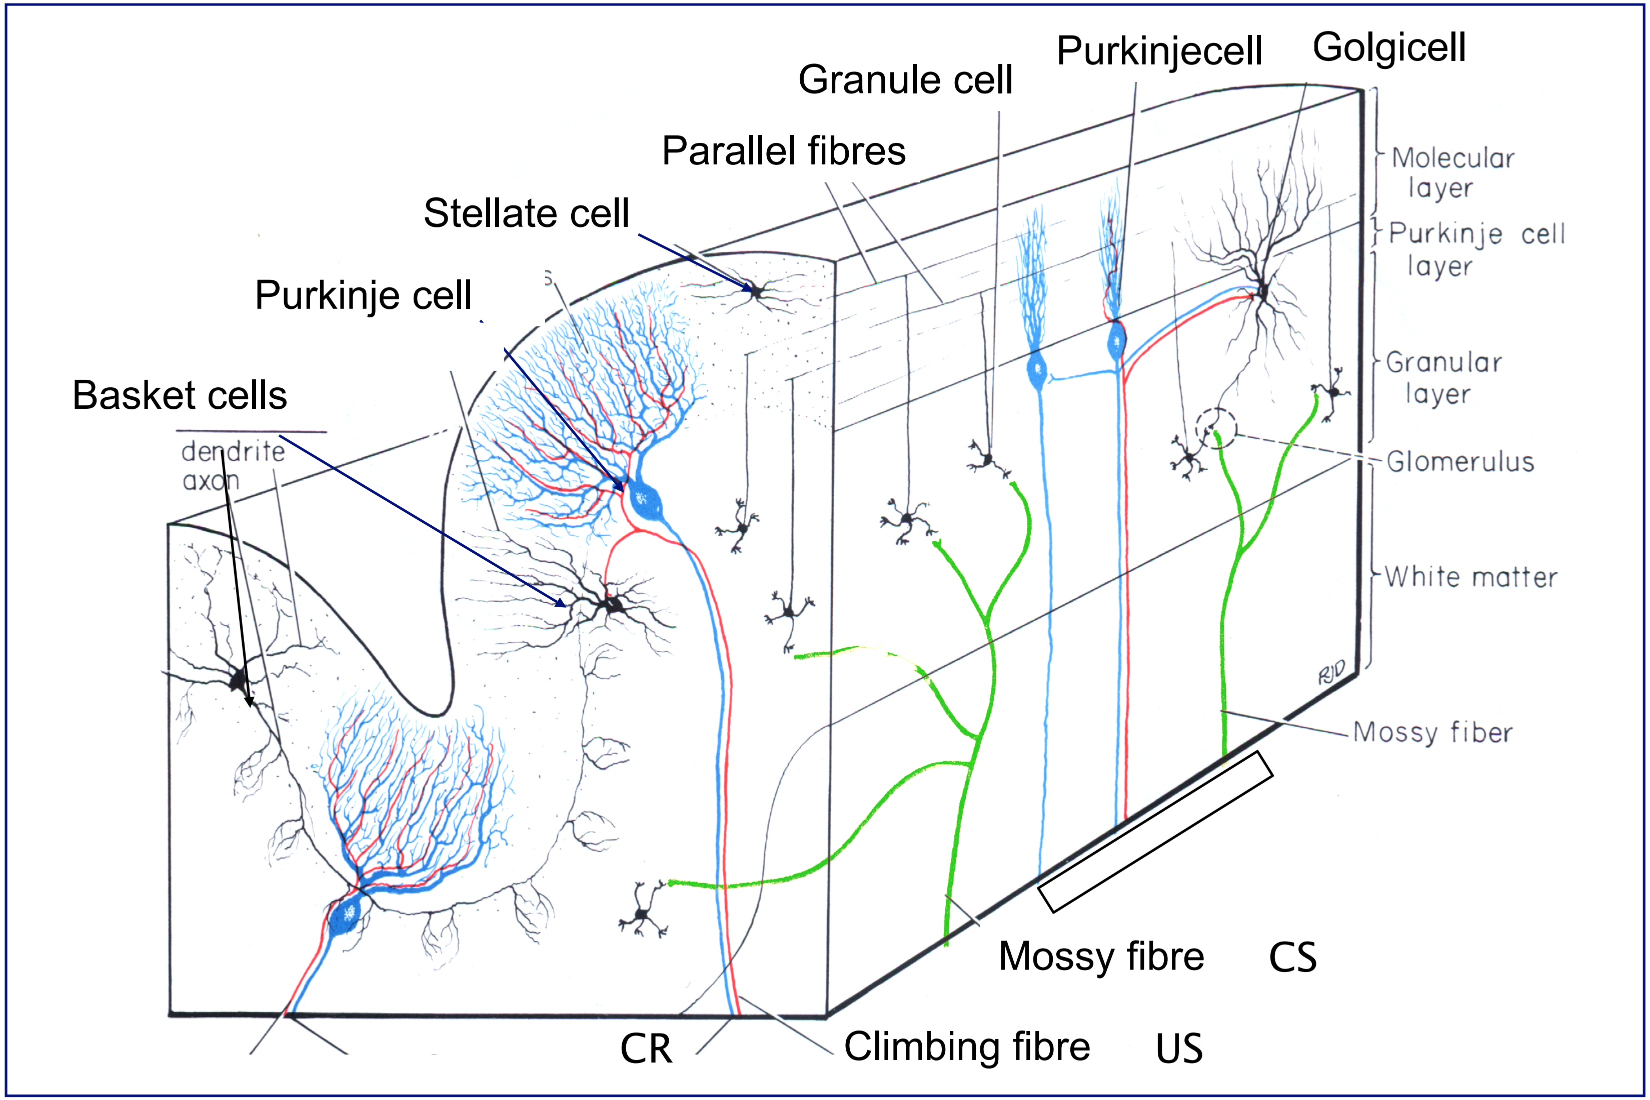

Cerebellar hypothesis

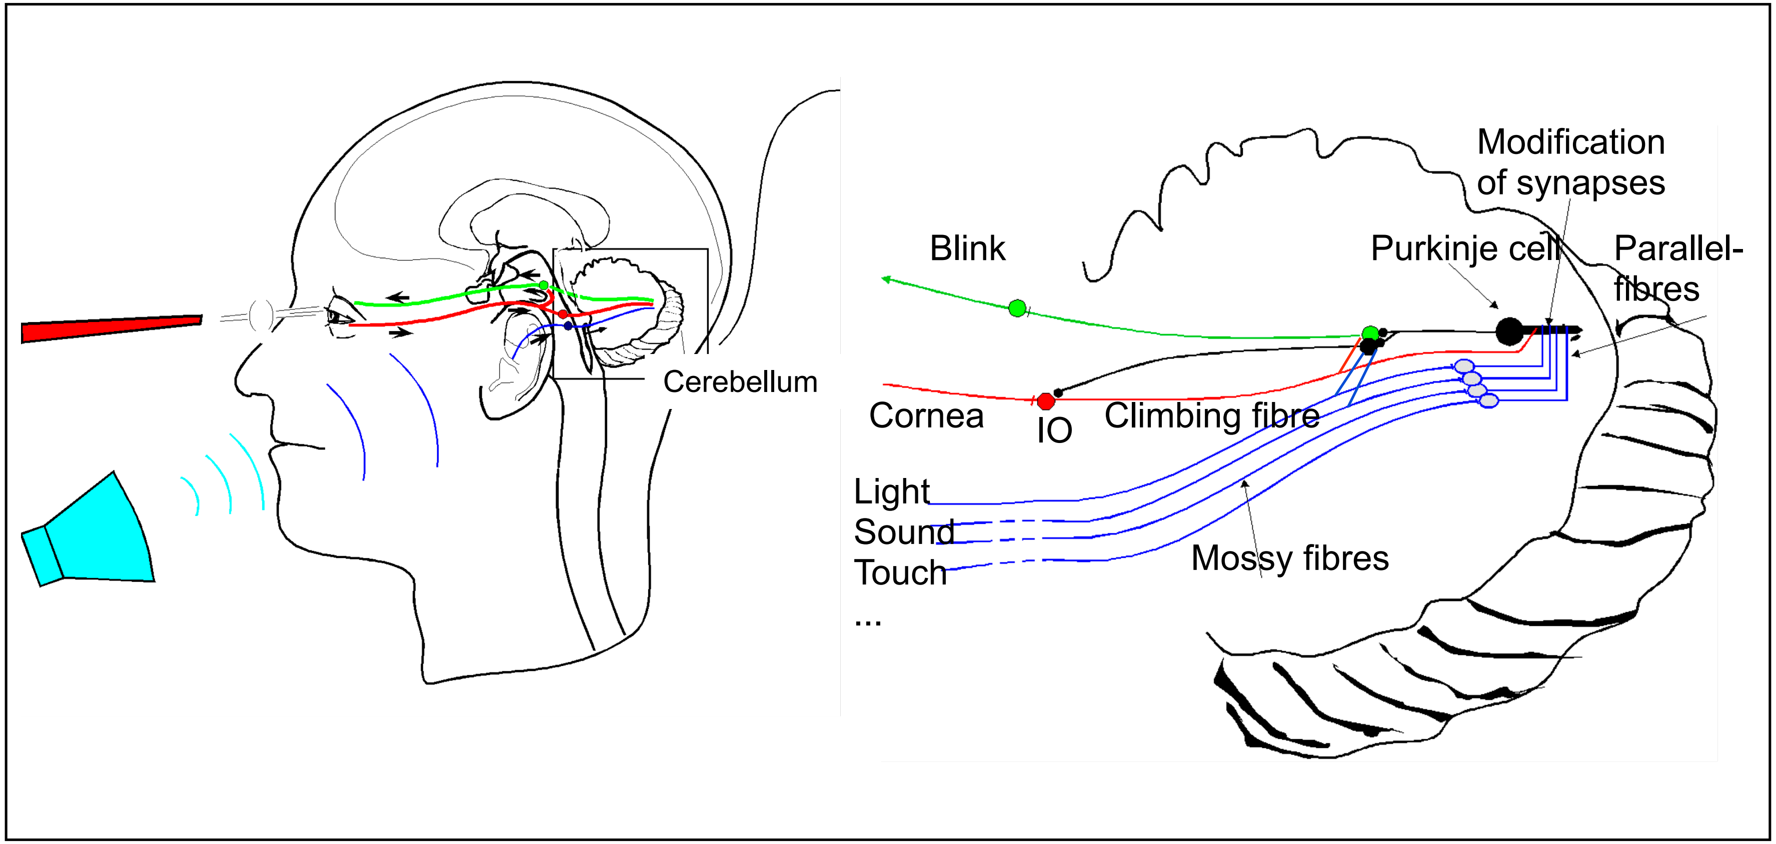

Adapted from Hesslow

Cerebellar module – microcircuit

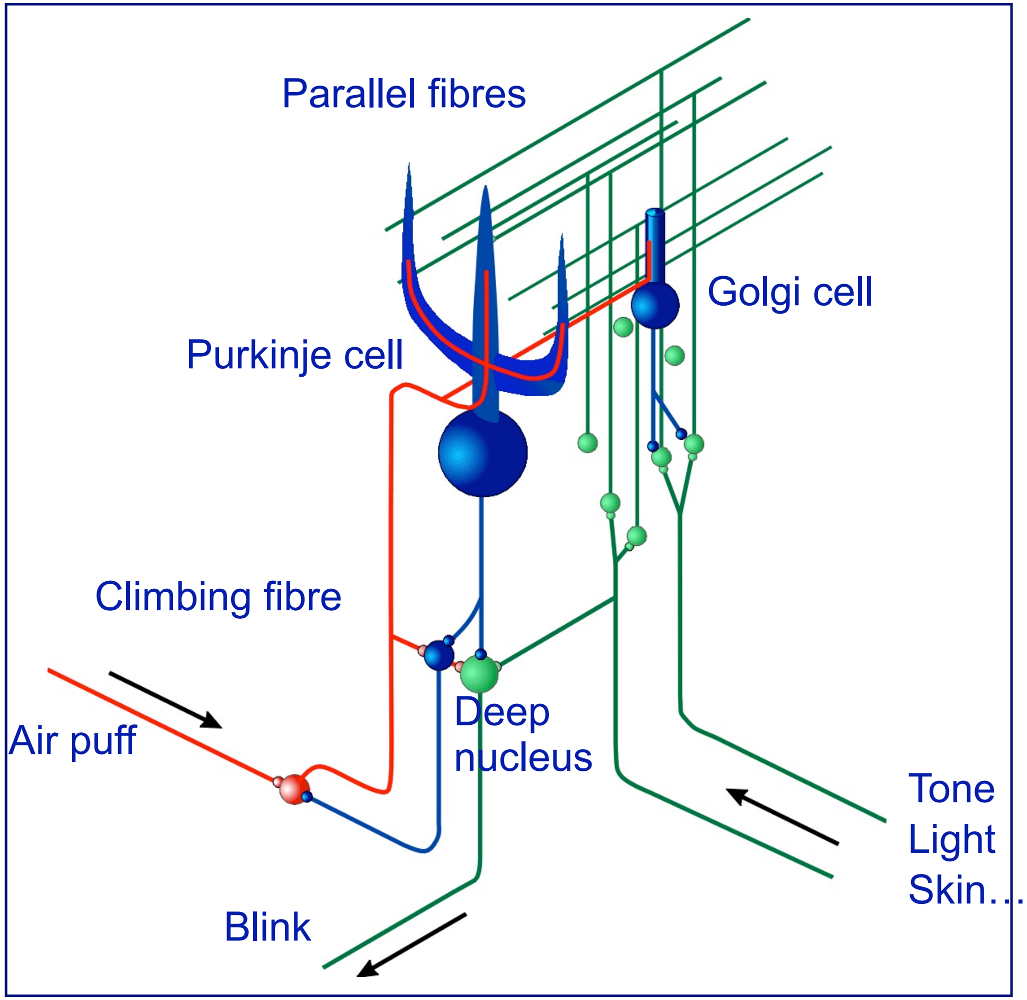

US is conveyed through climbing fibre

CS is conveyed through parallel fibre

CR is triggered through the deep nucleus

Learning depends on long-term depression of the parallel fibre-Purkinje cell synapse induced by coincident CF & PF activity

Adapted from Hesslow

Cerebellum: Physiology I

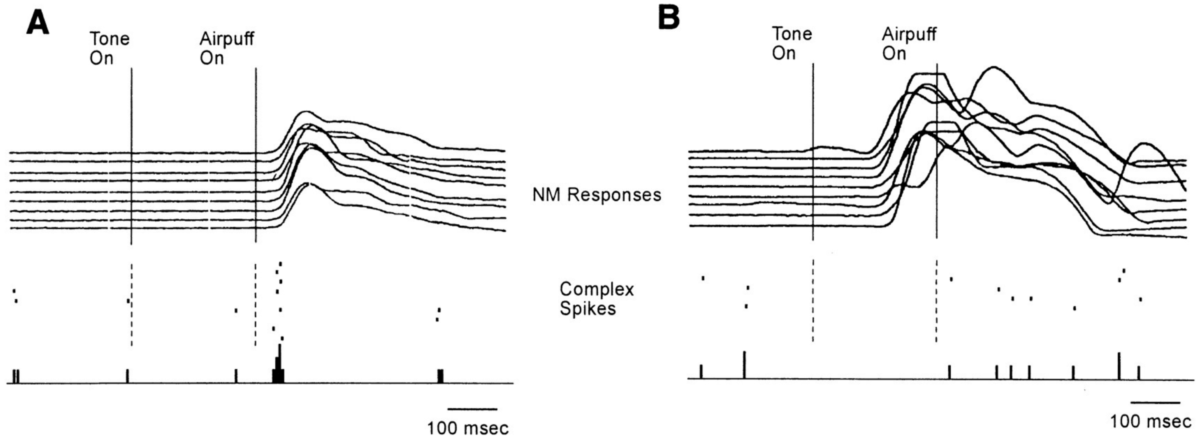

Thompson et al, 1998

Learning implies the absence of complex spikes in the Purkinje cell

Simple and complex spike activity in a trained Purkinje

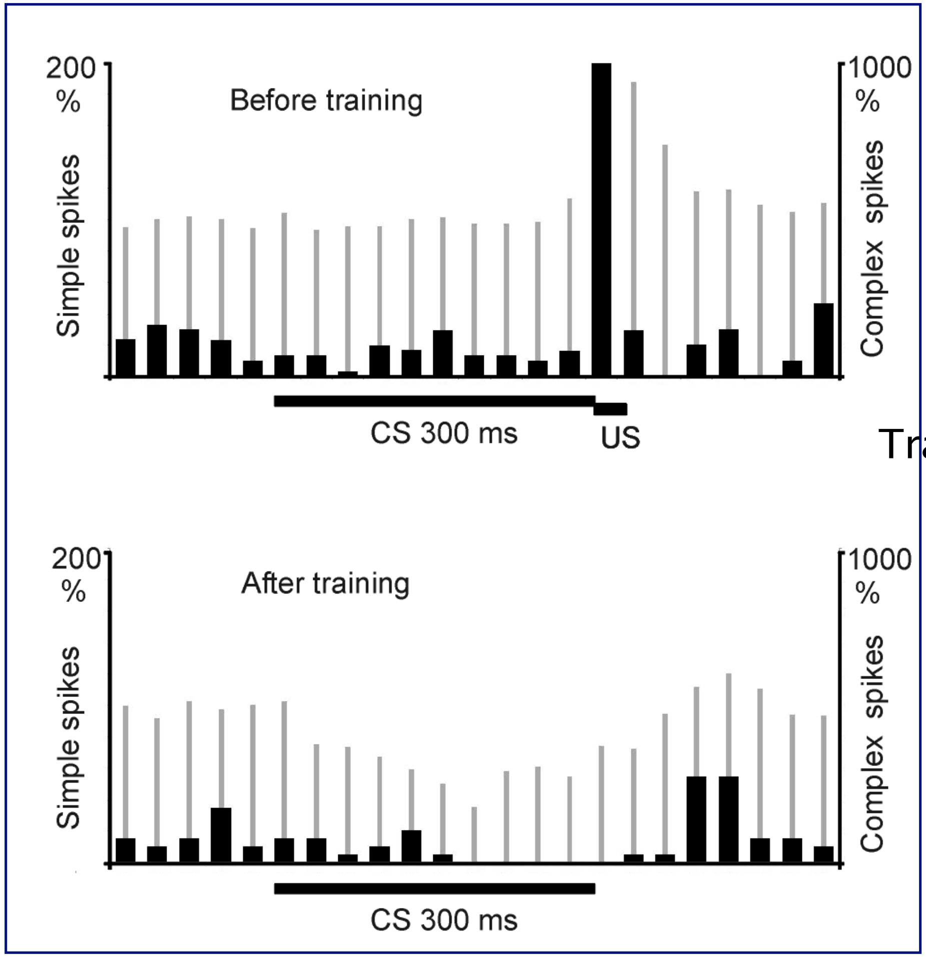

Training means the cessation of spiking

Hesslow, Jirenhed, Rasmussen et al., *Cerebellum*. 2008

Deep nucleus I

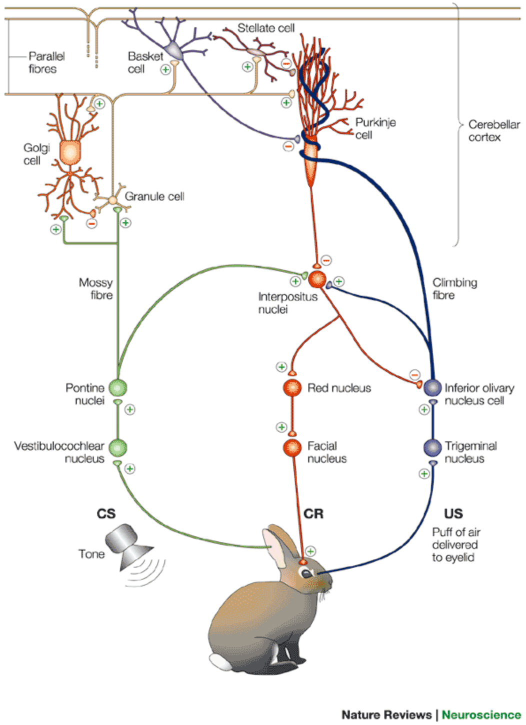

Medina et al, Nat Rev Neurosci, 2001

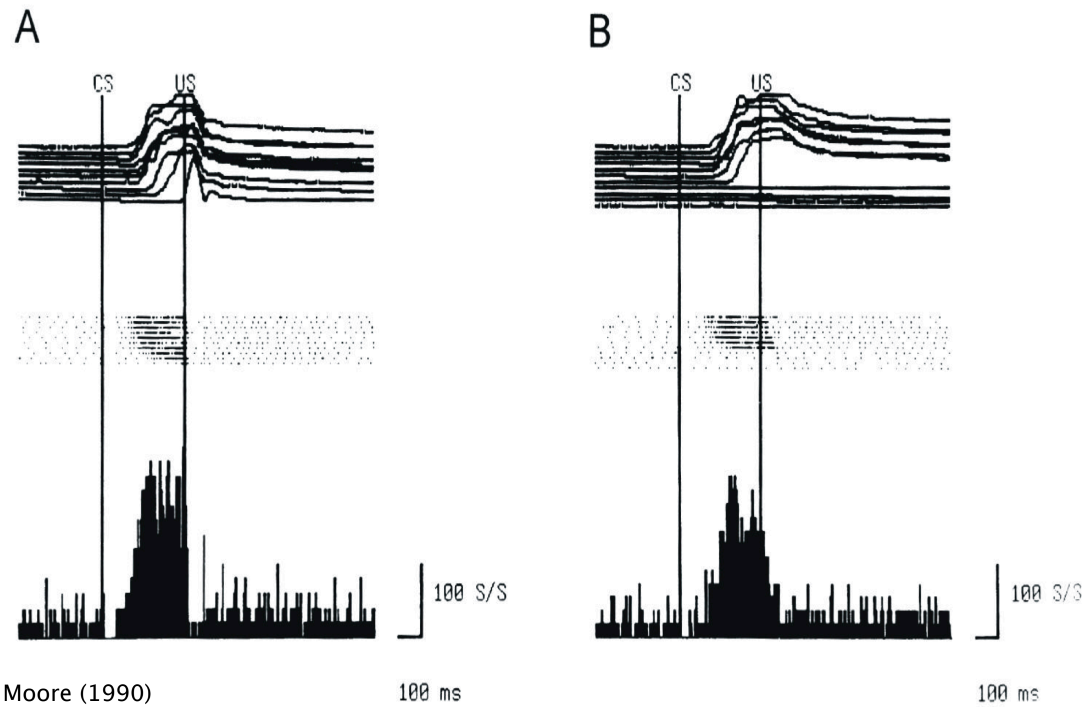

Moore (1990)

Purkinje cell pause leads to rebound in DN that correlates with CR

Stimulation of the deep nucleus – olivary pathway can reduce the climbing fibre field potentials

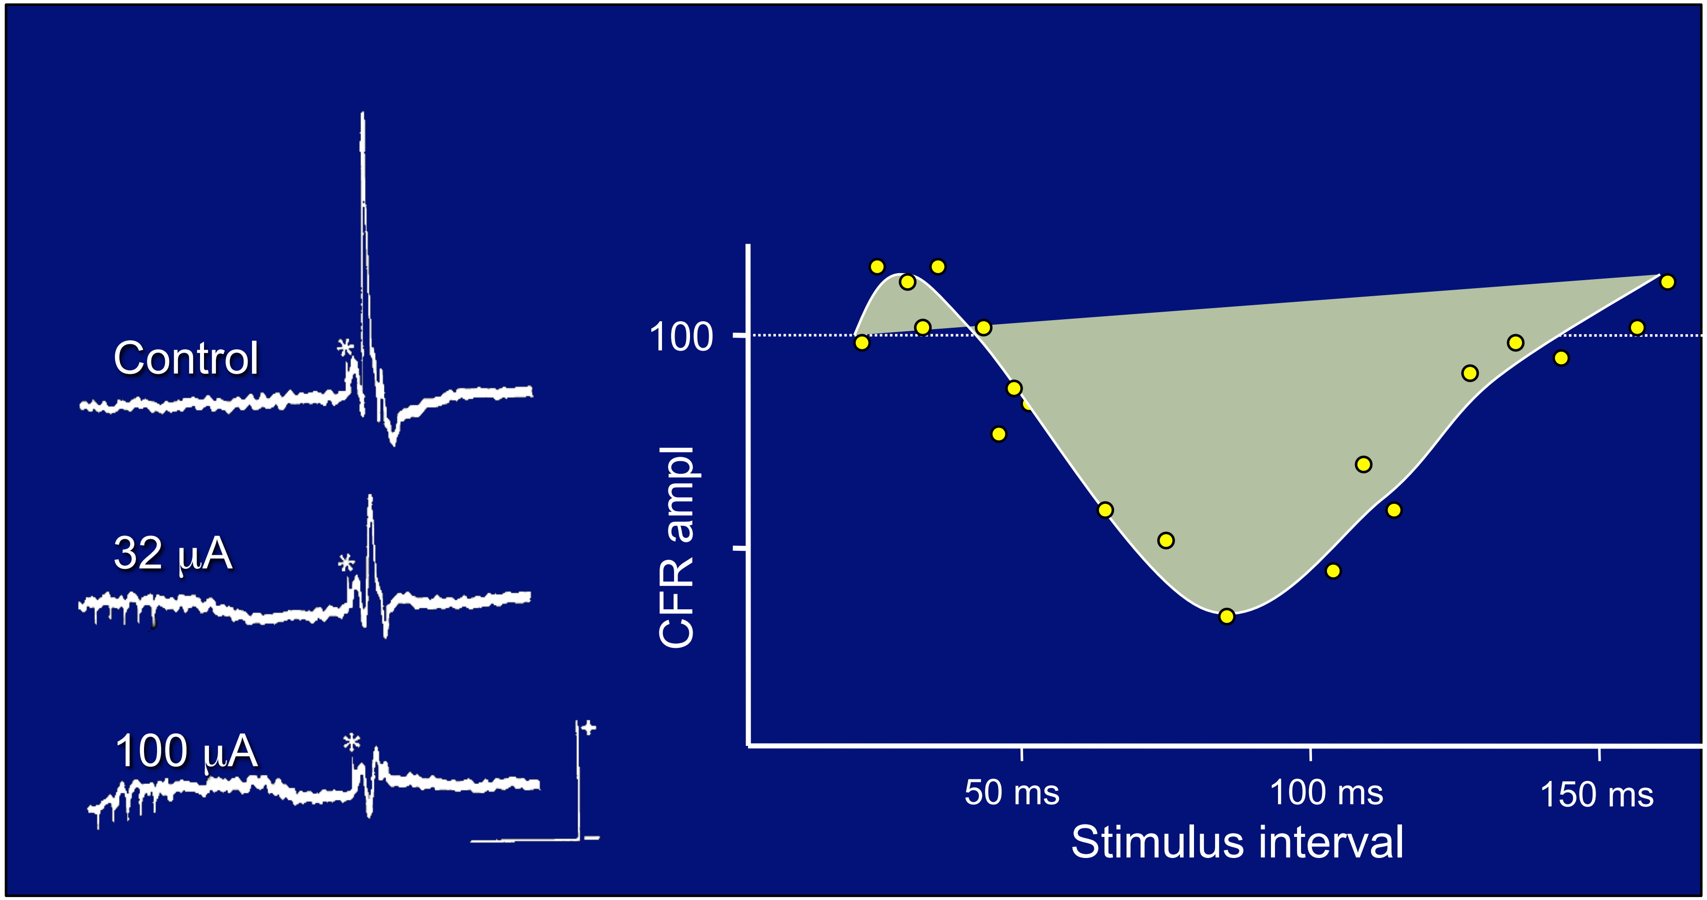

Hesslow, *Neurosci Lett*, 1986

Blocking of the deep nucleus – inferior olive fibres leads to the return of complex spikes in Purkinje cells

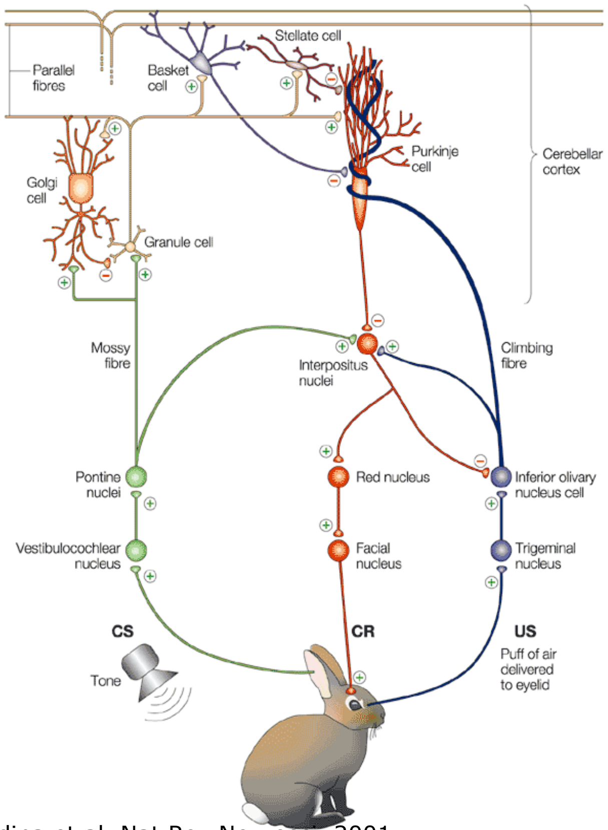

Medina et al, Nat Rev Neurosci, 2001

Nature Reviews | Neuroscience

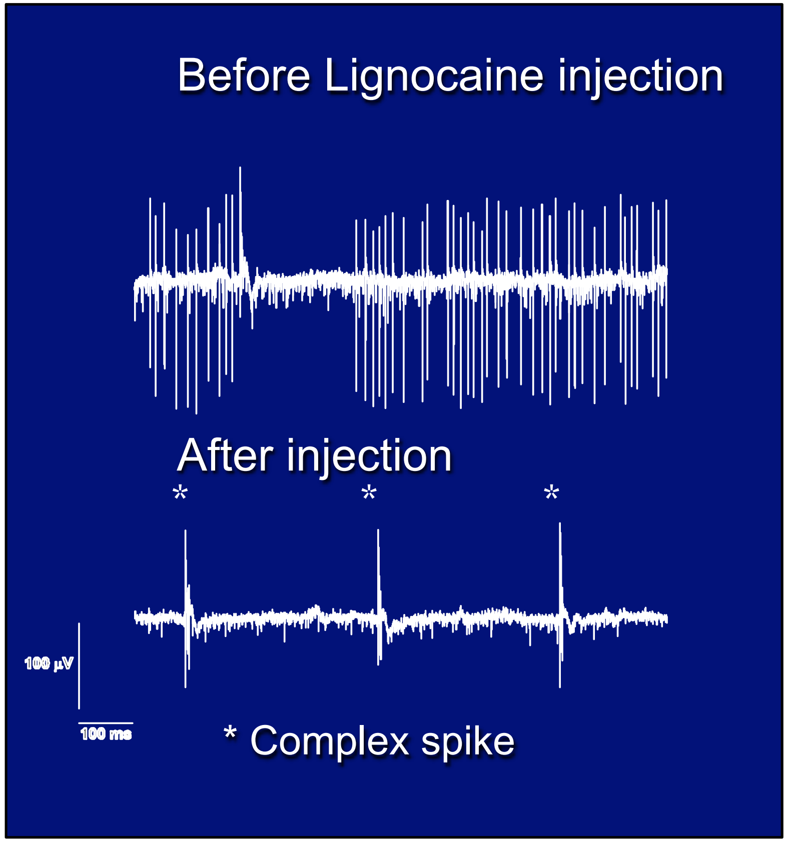

Bengtsson et al. EJN 2004

Extinction is induced by stimulation of the deep nucleus – inferior olive

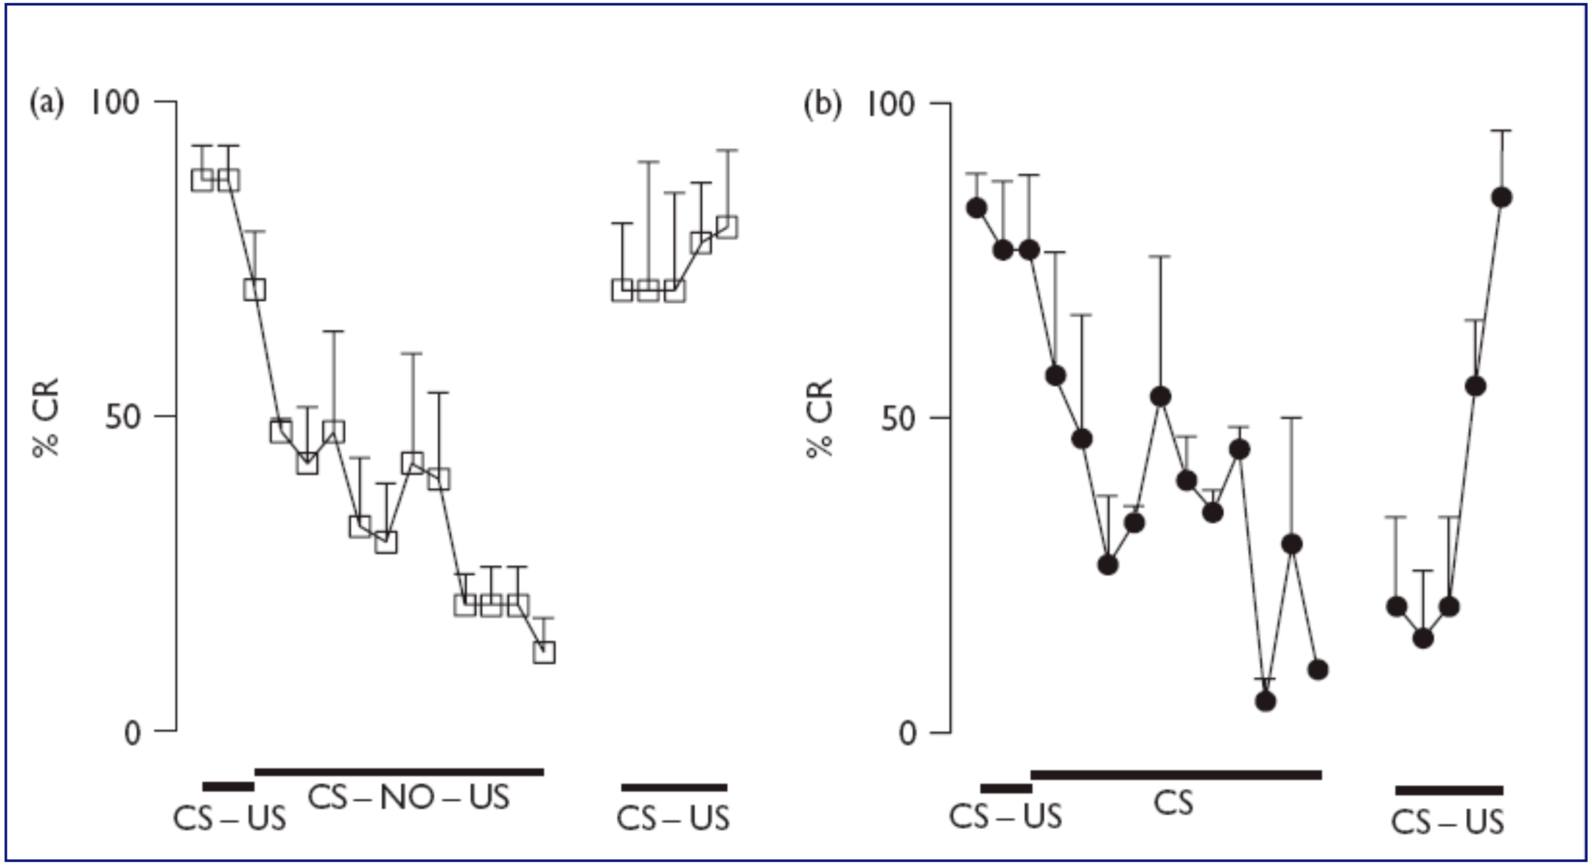

Bengtsson et al., *Neuroreport* 2007

Specific learning and the cerebellum

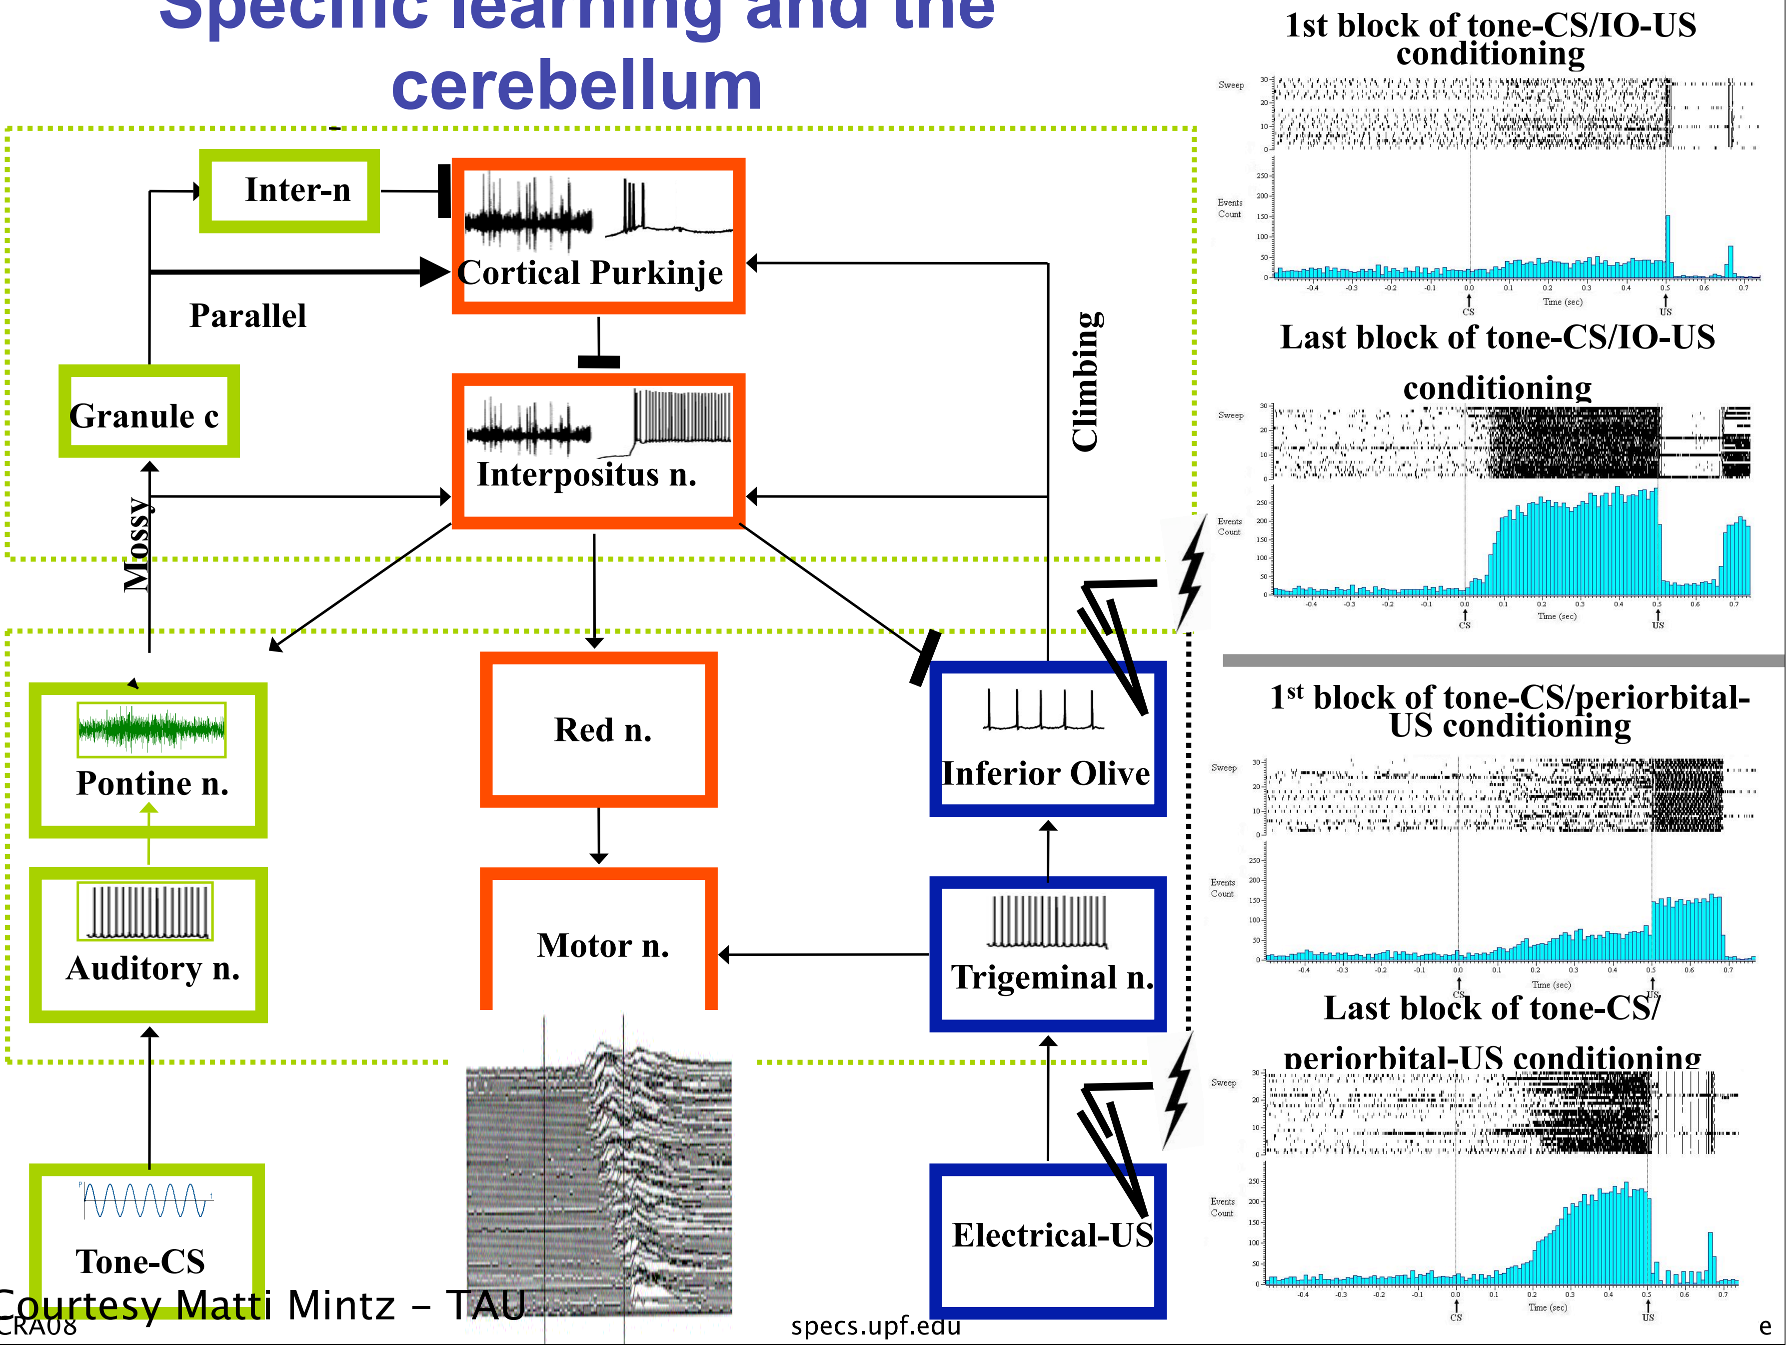

Courtesy Matti Mintz – TAU

The model circuit

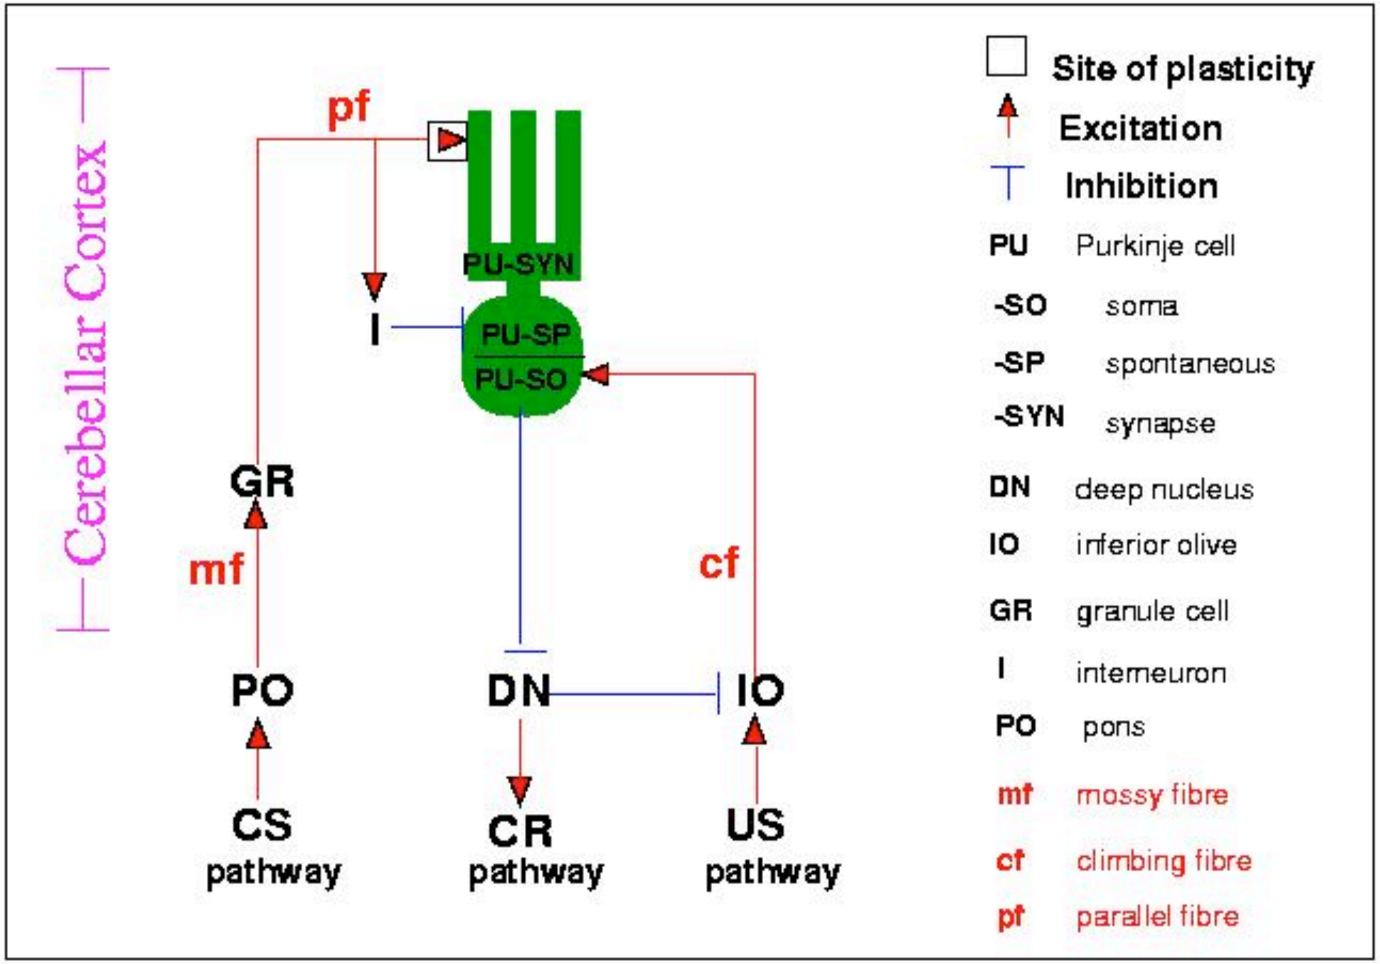

Verschure & Mintz (2001) Comp. Neur., Hofstotter et al (2003) Eur.J.neurosci

Cerebellum: Functional interpretation II

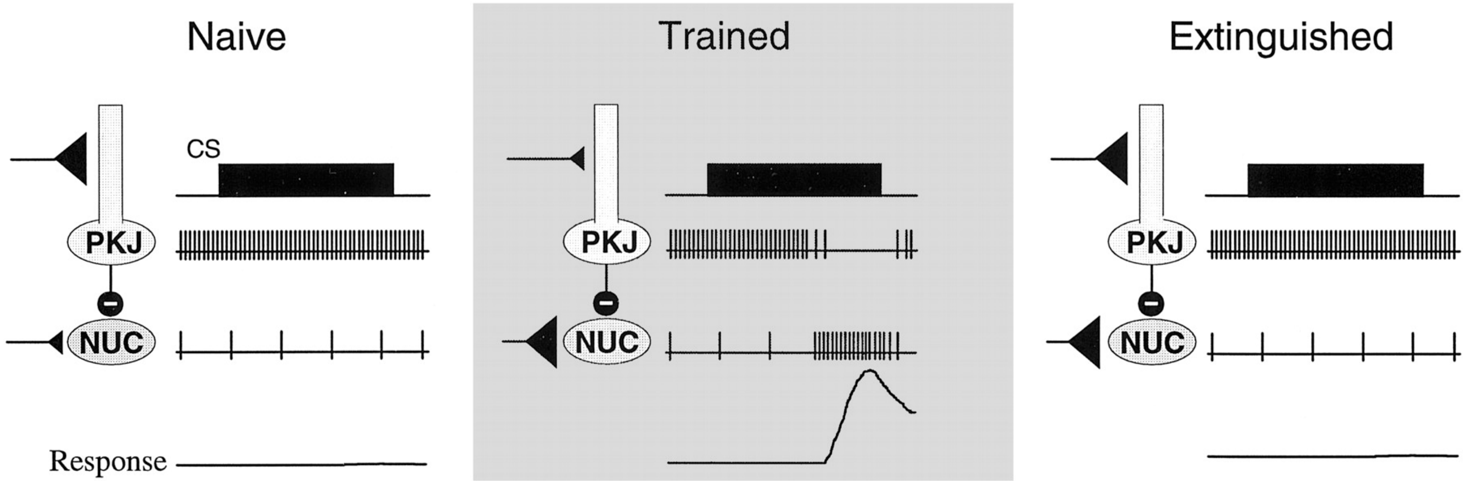

Learning leads to a cessation of Purkinje cell activity which releases the deep nucleus from inhibition leading to a CR through rebound polarization

Garcia et al, 1999
Verschure & Mintz, 2001

Generic integrate and fire neuron model

$$E_i(t) = \gamma^E \sum_{j=0}^{N} A_j(t) w_{ij}(t)$$

$$I_i(t) = -\gamma^I \sum_{j=0}^{N} A_j(t) w_{ij}(t)$$

$$V_i(t+1) = \beta V_i(t) + E_i(t) + I_i(t)$$

$$A_i(t) = H(V_i(t) - \theta^A)$$

Verschure & Mintz (2001) Comp. Neur., Hofstotter et al (2003) Eur.J.neurosci

Exceptions to the I&F model

Pu-syn is a linear threshold unit:

$$A_i(t) = H(V_i(t) - \theta^A) \cdot V_i(t)$$

DN neurons show rebound polarization:

$$V_i(t+1) = \beta V_i(t) + [H(V_i(t) - \theta^R)H(\theta^R - V_i(t-1)) \cdot \mu] + I_i(t)$$

Verschure & Mintz (2001) Comp. Neur., Hofstotter et al (2003) Eur.J.neurosci

Equations for LTD/LTP

$$w_{ij}(t+1) = \begin{cases} \varepsilon w_{ij}(t) & \text{if } E_i \in [E_{min}^{LTD}, E_{max}^{LTD}] \\ w_{ij}(t) & \text{otherwise} \end{cases}$$

$$w_{ij}(t+1) = \begin{cases} w_{ij}(t) + \eta(w_{ij}^{max} - w_{ij}(t)) & \text{if } E_i \in [E_{min}^{LTP}, E_{max}^{LTP}] \\ w_{ij}(t) & \text{otherwise} \end{cases}$$

Verschure & Mintz (2001) Comp. Neur., Hofstotter et al (2003) Eur.J.neurosci

negative feedback & LTD

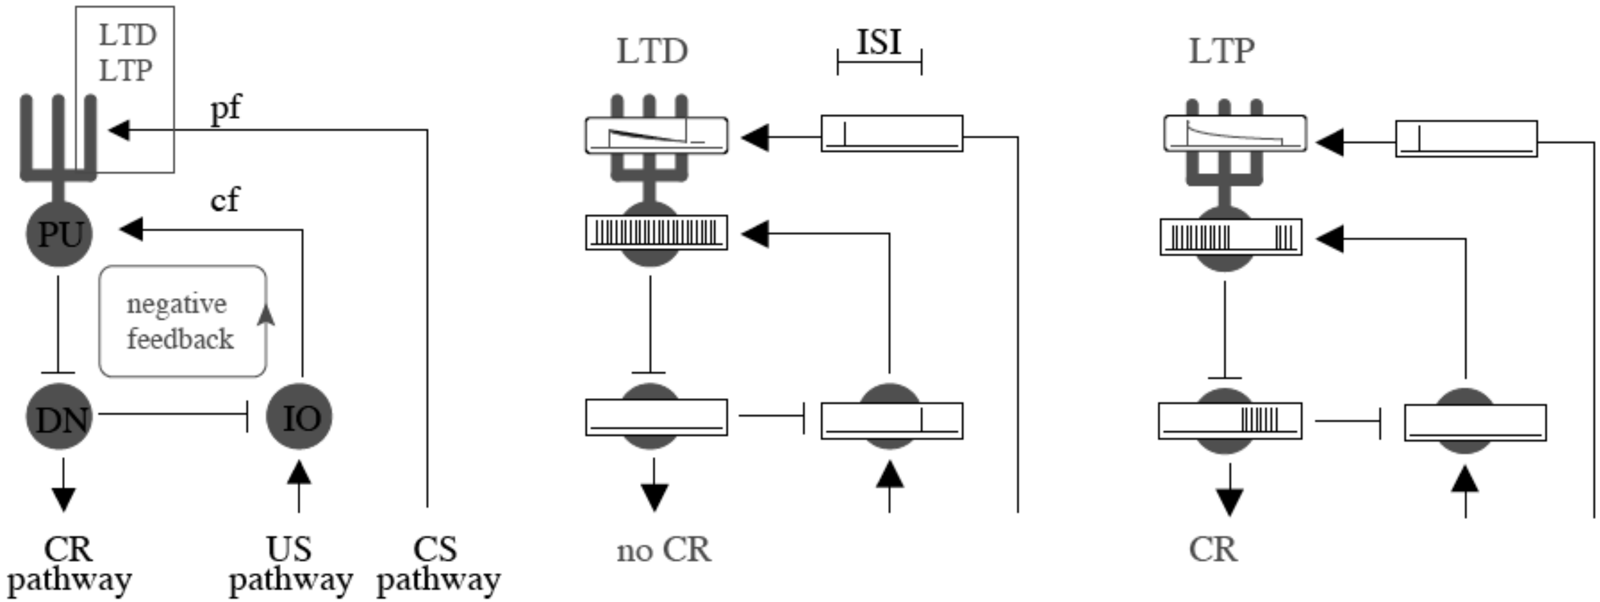

Negative feedback through the DN->IO inhibitory projection prevents CF activity for USs that are well predicted by the circuit

Verschure & Mintz (2001) Comp. Neur., Hofstoder et al (2003) Eur.J.neurosci

Model: Learning dynamics

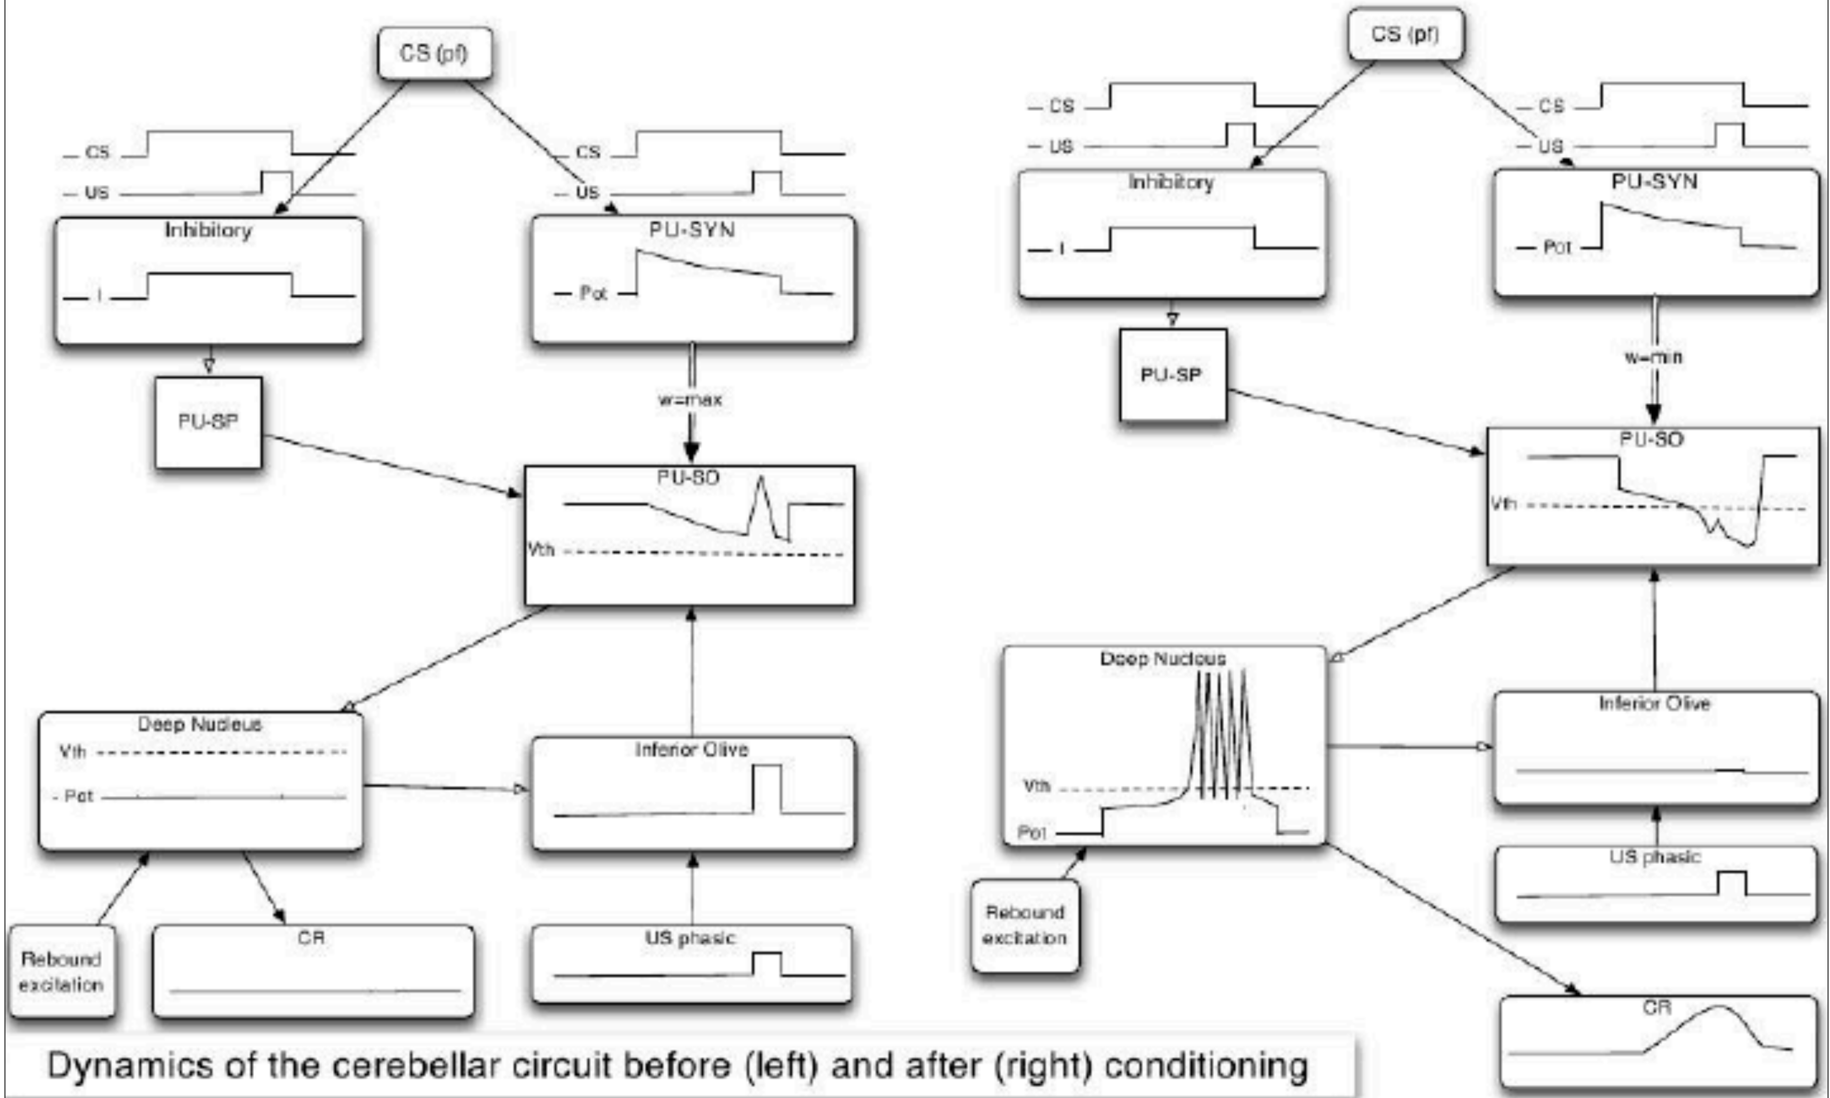

Dynamics of the cerebellar circuit before (left) and after (right) conditioning

Verschure & Mintz (2001) Comp. Neur., Hofstotter et al (2003) Eur.J.neurosci

Simulated Conditioning Experiments

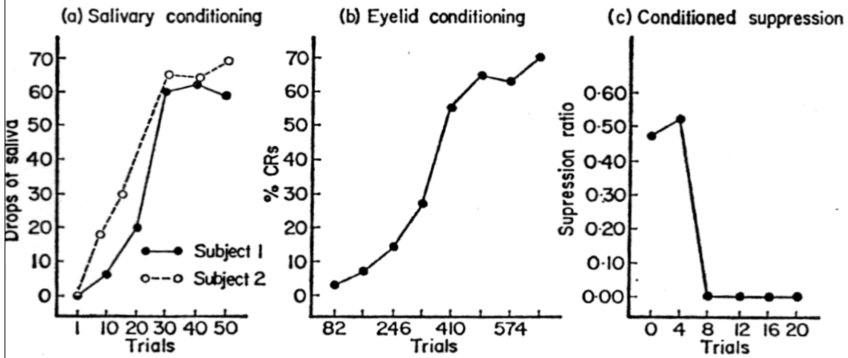

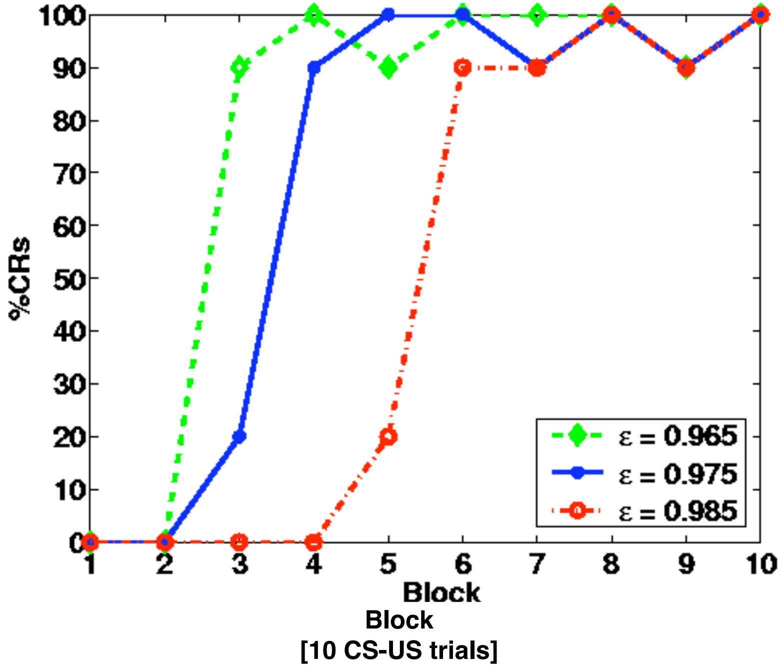

Acquisition training:
LTD factor, ε, determines
speed of acquisition.

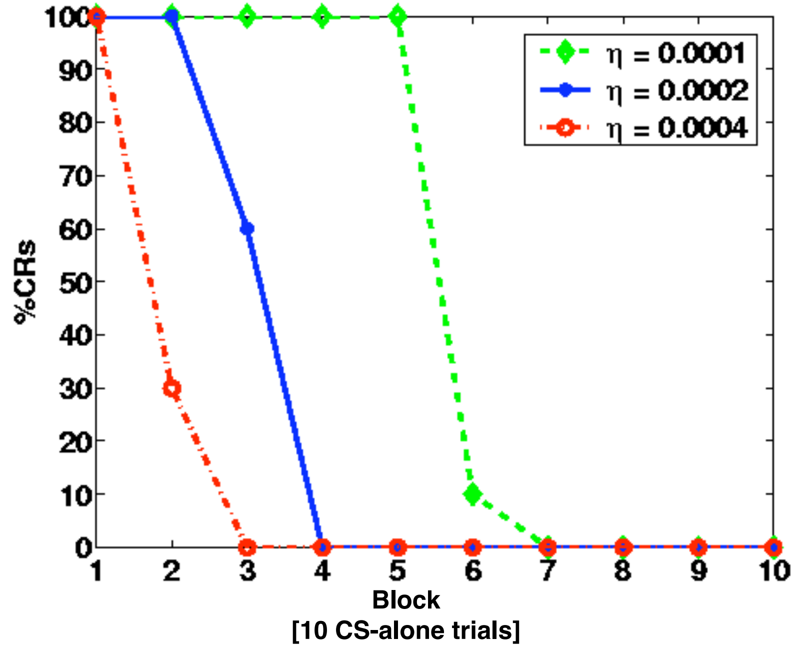

Extinction training:
LTP factor, η, determines
speed of acquisition.

LTD/LTP and Pu response duration

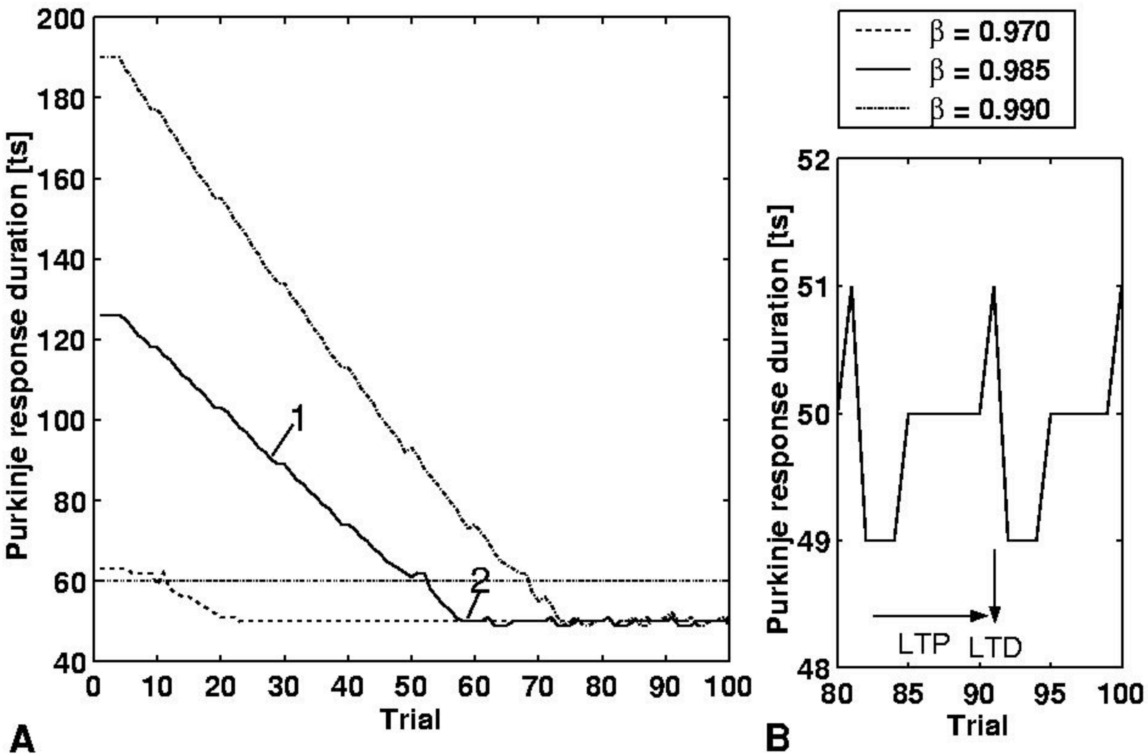

Verschure & Mintz (2001) Comp. Neur., Hofstoder et al (2003) Eur.J.neurosci

Observed Purkinje cell responses confirm the model's learning hypothesis

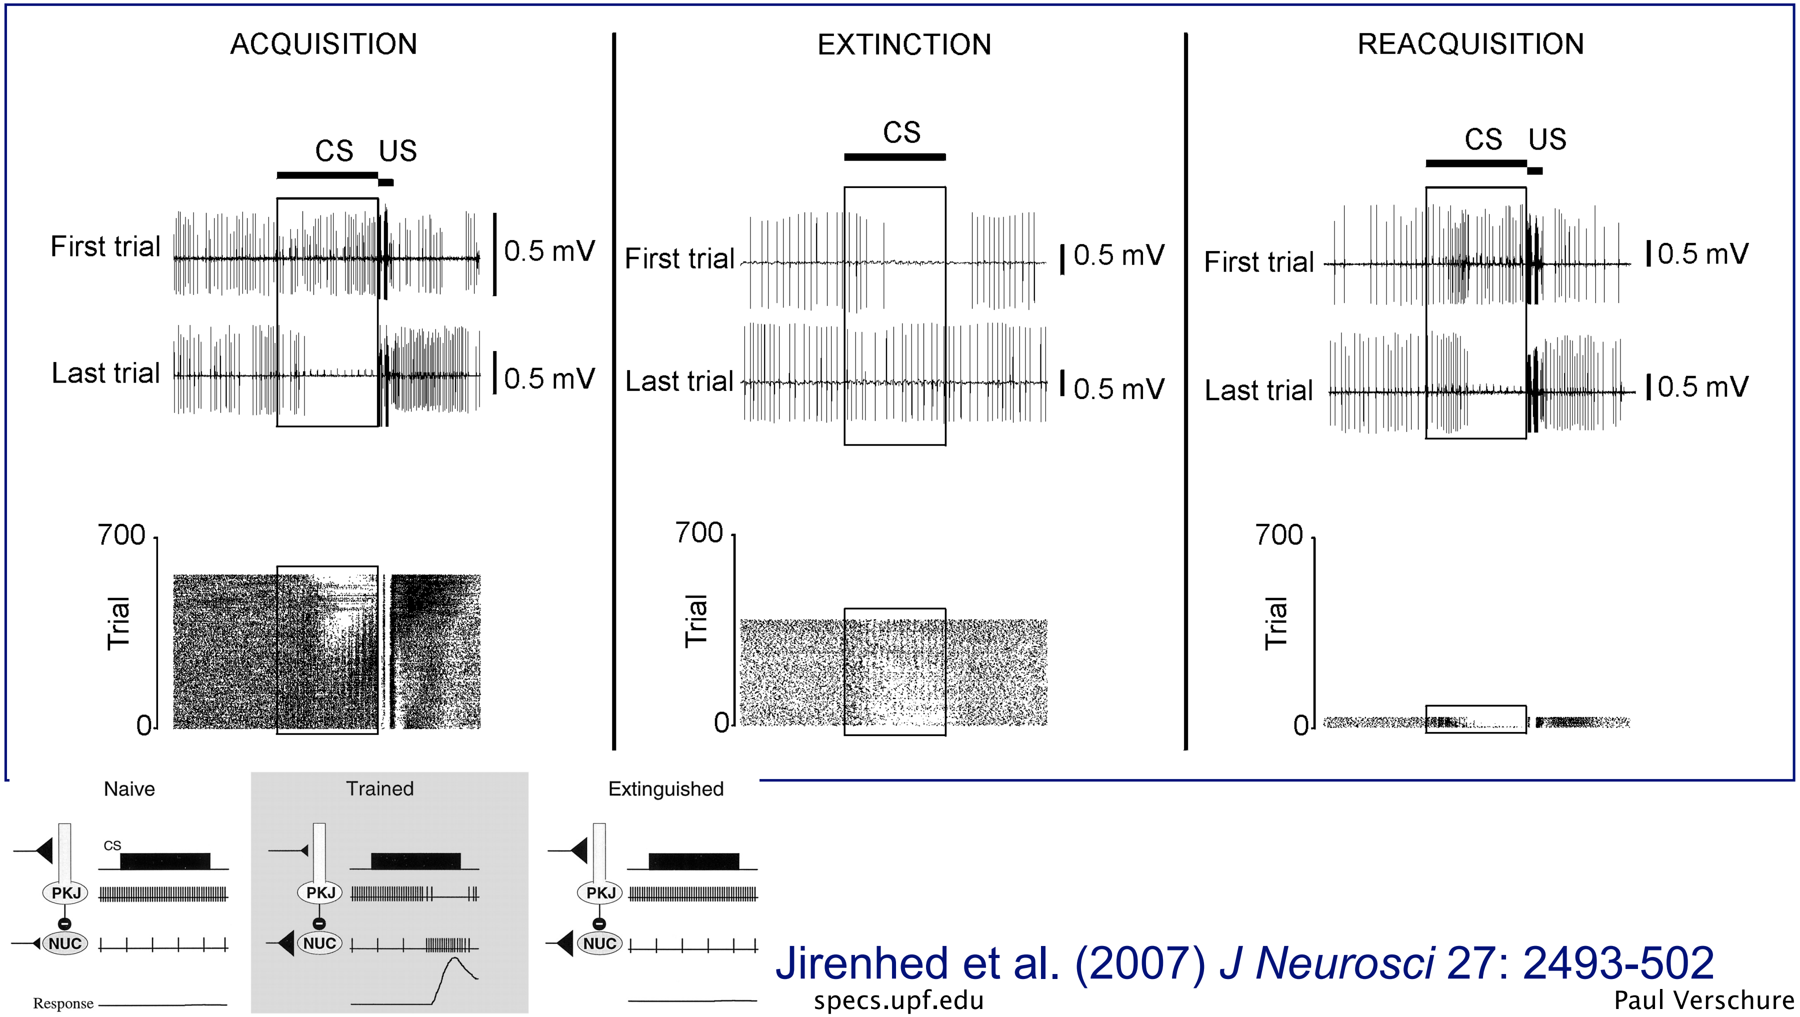

Jirenhed et al. (2007) *J Neurosci* 27: 2493-502

The silicon cerebellum

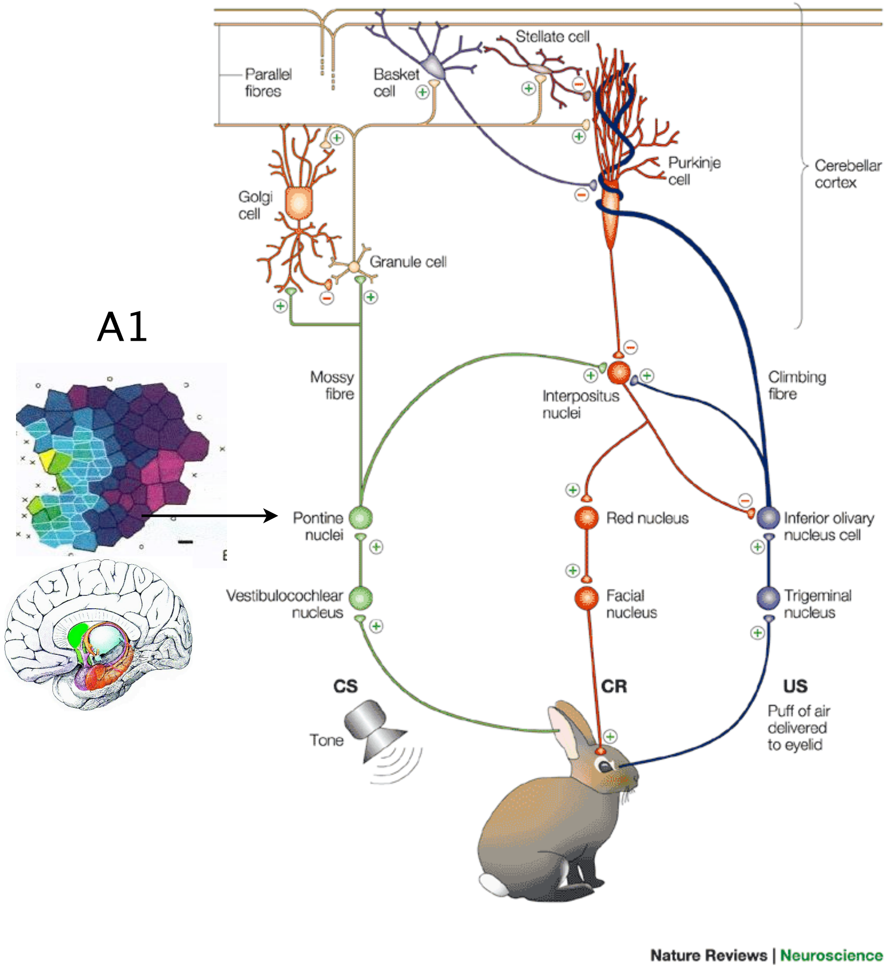

Medina et al, Nat Rev Neurosci, 2001

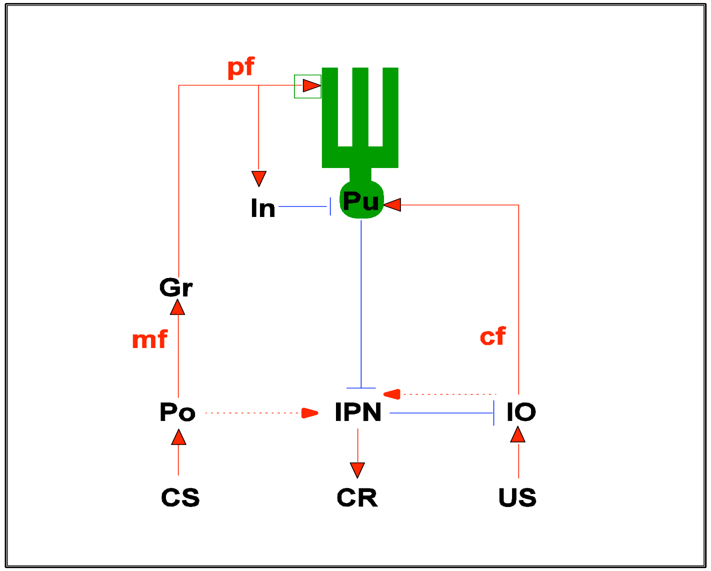

Hofstotter et al (2003) Eur. J. Neurosci

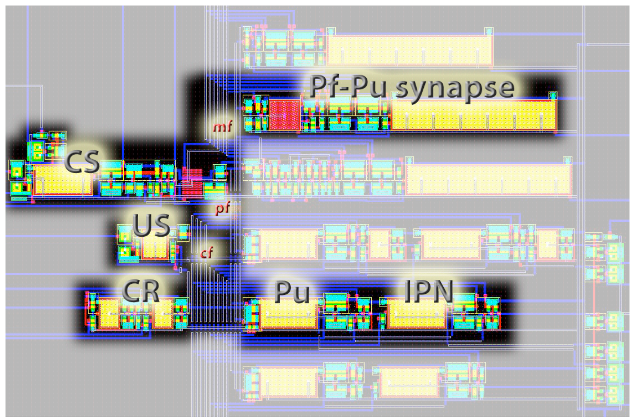

Hofstotter et al (2005) NIPS

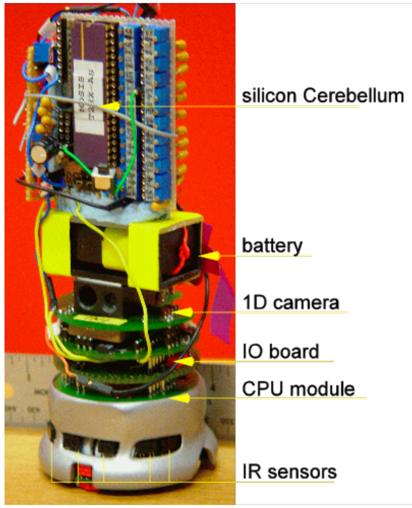

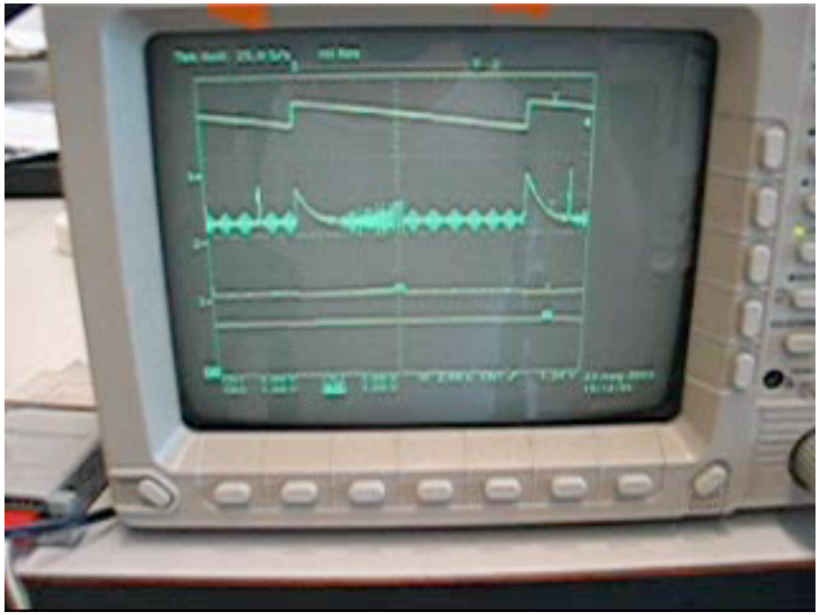

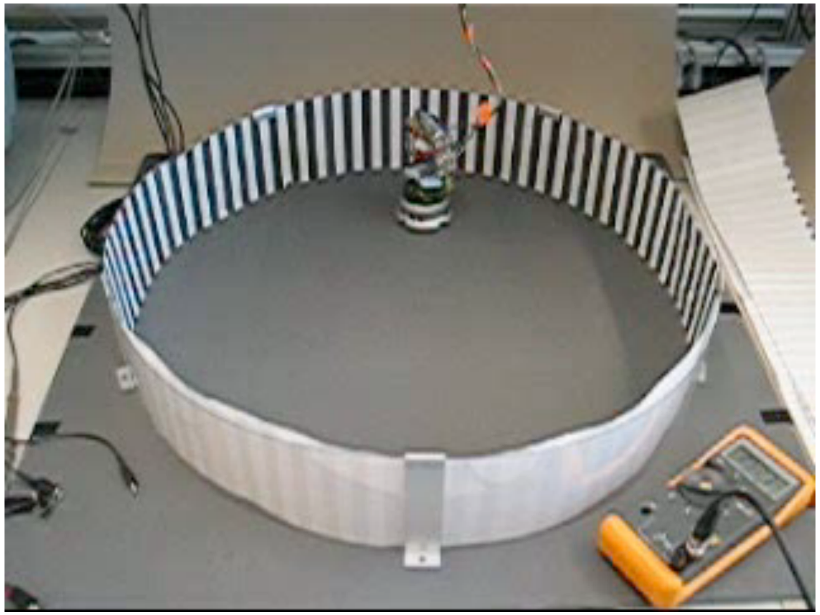

Hofstotter et al (2005) NIPS

Silicon Cerebellum dynamics

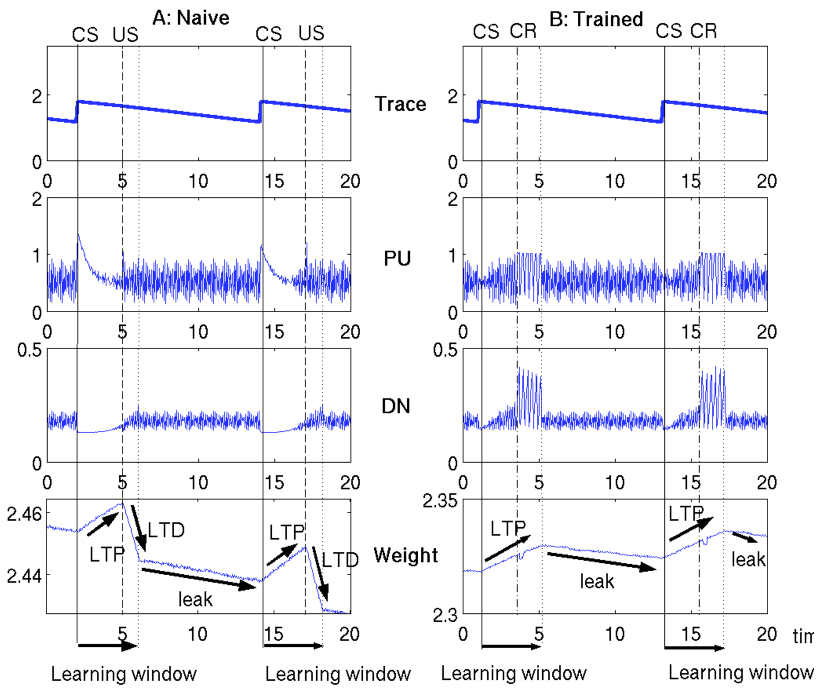

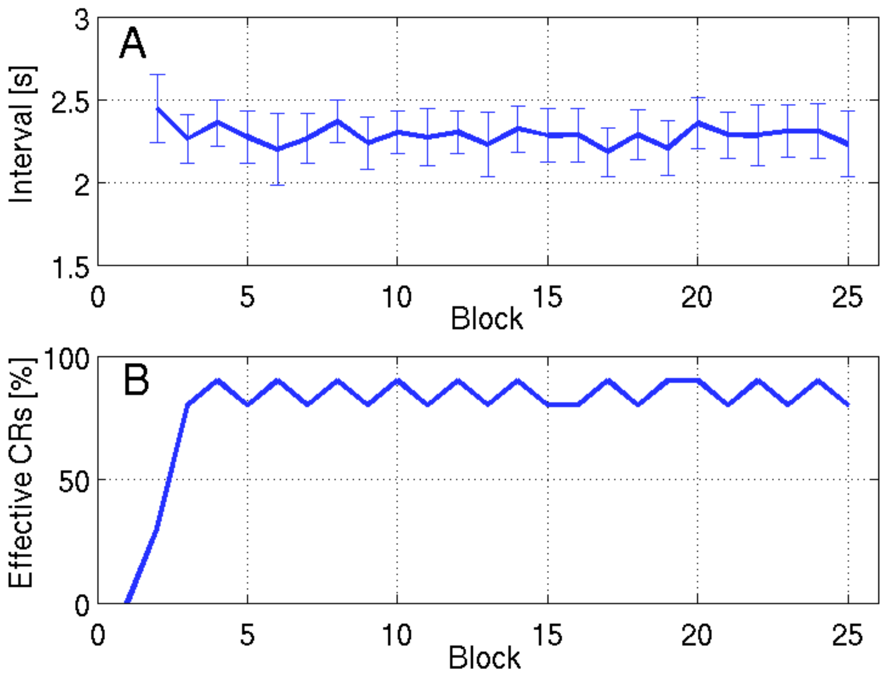

Hofstotter et al (2005) NIPS

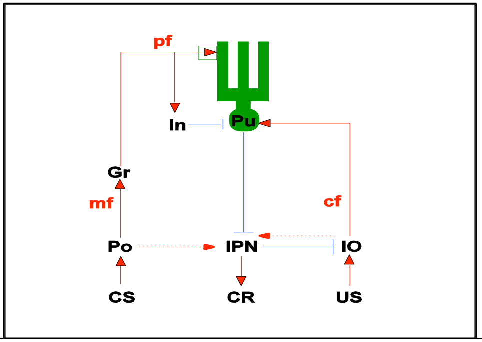

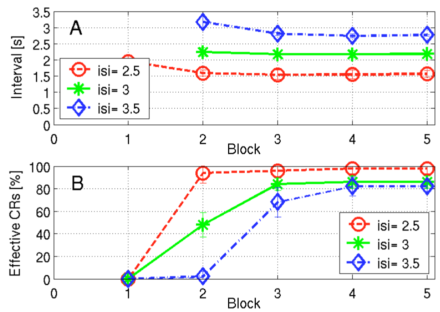

ReNaChip: replacing a discrete microcircuit of the cerebellum by an artificial system

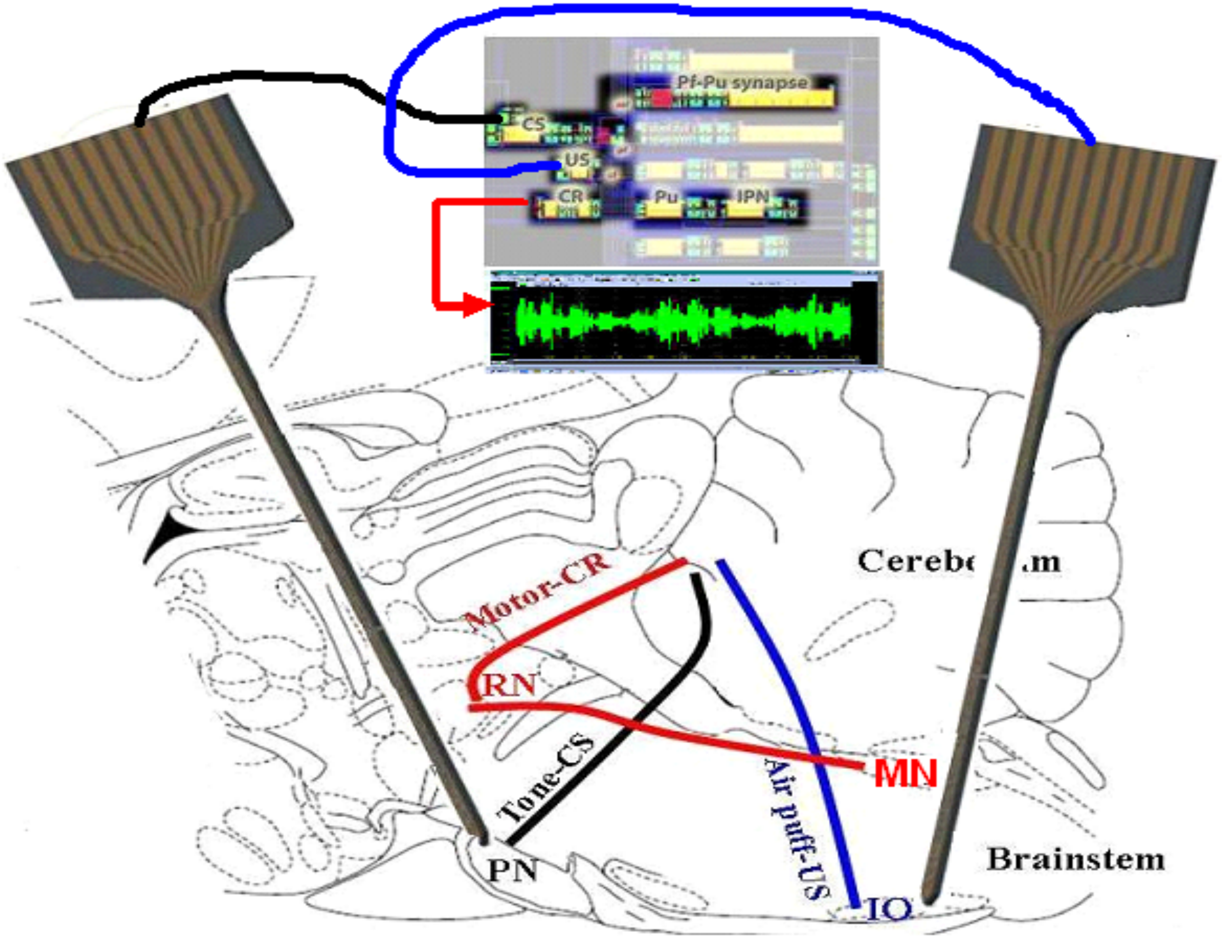

The cerebellum and prediction in motor control: Do we need forward models?

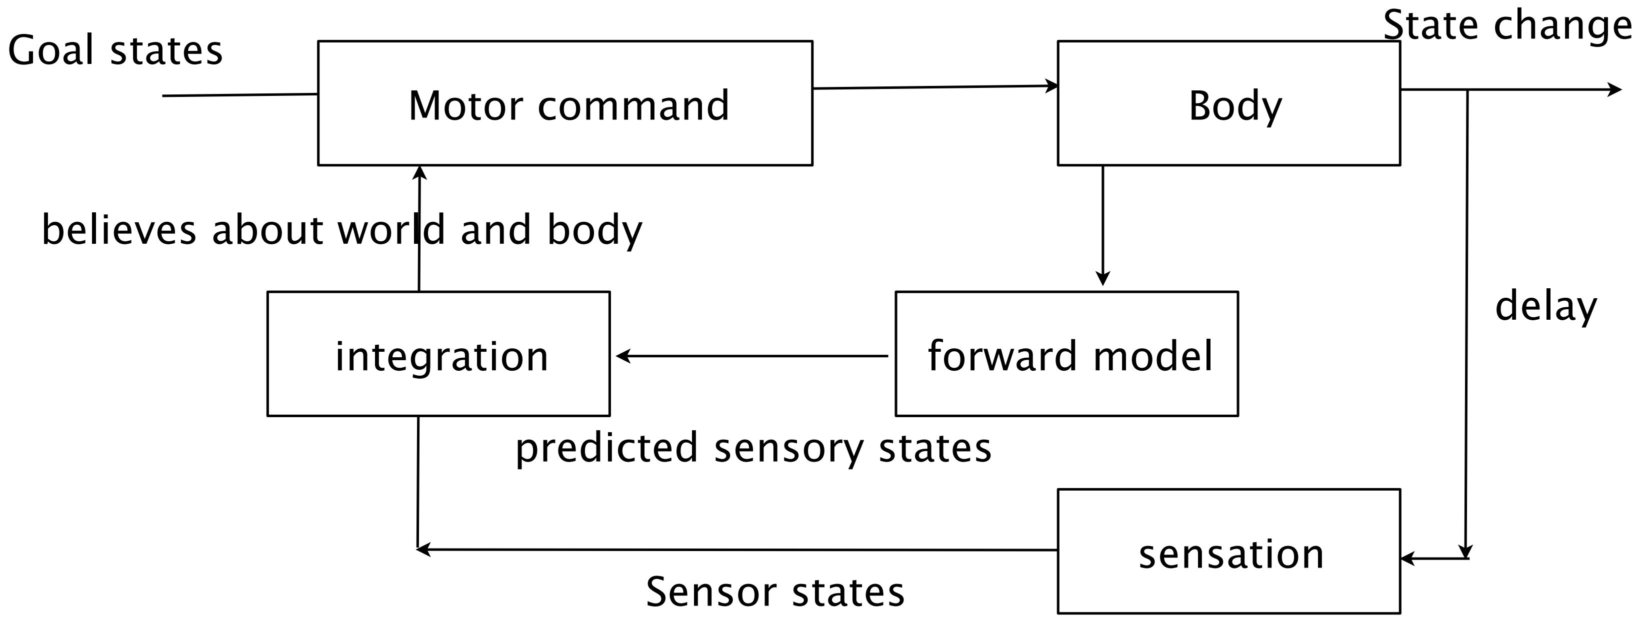

The cerebellum predicts the occurrence of events to trigger well timed outputs
BUT
It does not predict the state of the world

Bottom line (s)

- We have generalized the adaptive layer of DAC to a detailed model of the 2-phase theory of classical conditioning
- The cerebellar model of the non-specific learning system provides a substrate for response timing, cognition & prediction
- The model has been transformed into a silicon implementation: The silicon cerebellum
- The cerebellum questions the validity of localized self-contained forward models

Can we build a cyborg?

Hardware replacement

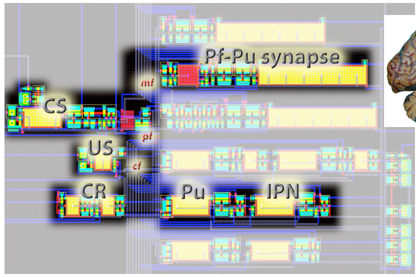

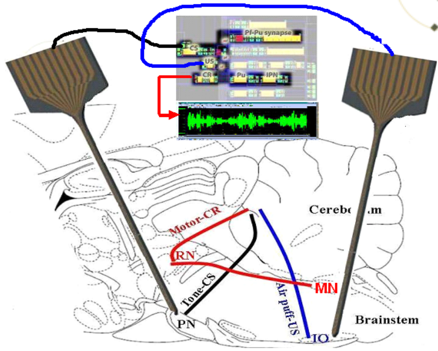

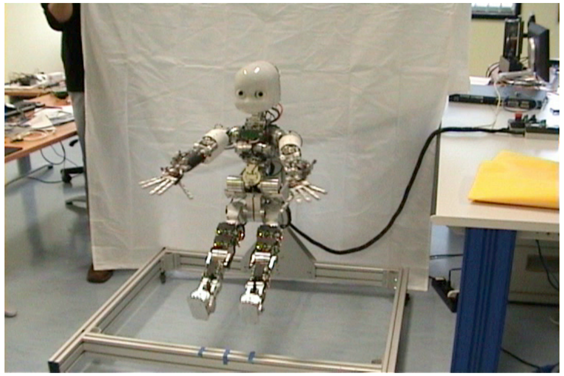

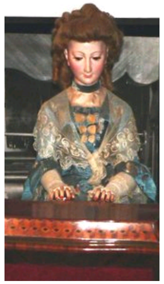

Sensory stimulation

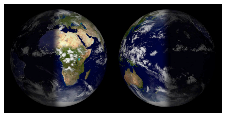

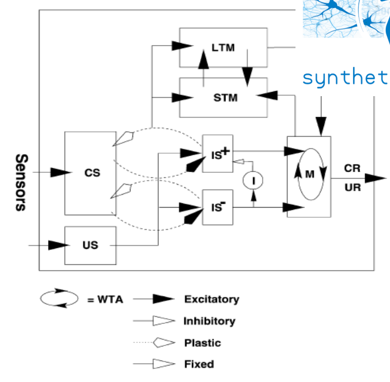

specs.upf.edu

YES!