Cerebellar Timing and Classical Conditioning

Computational Models of Neural Systems
Lecture 2.4

David S. Touretzky
September, 2023

Feedback vs. Feedforward Control

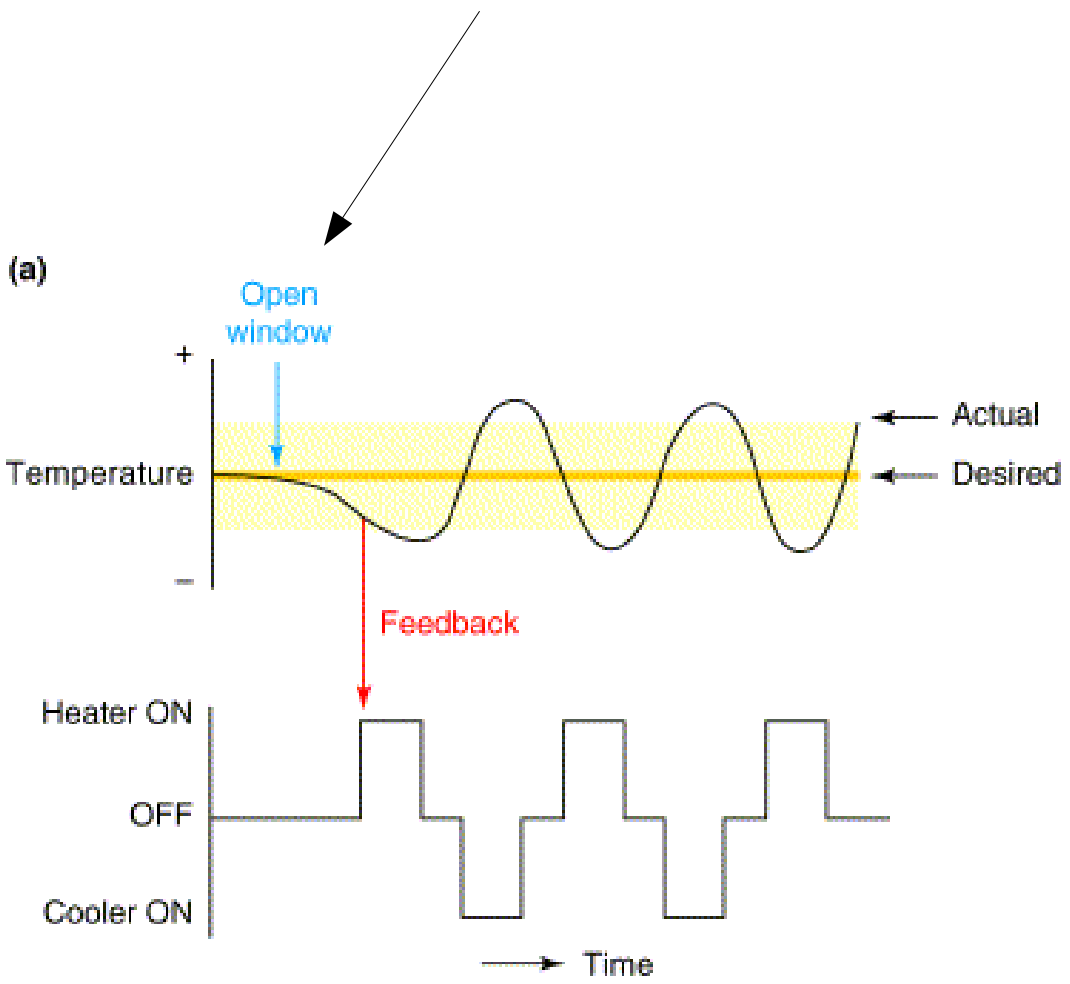

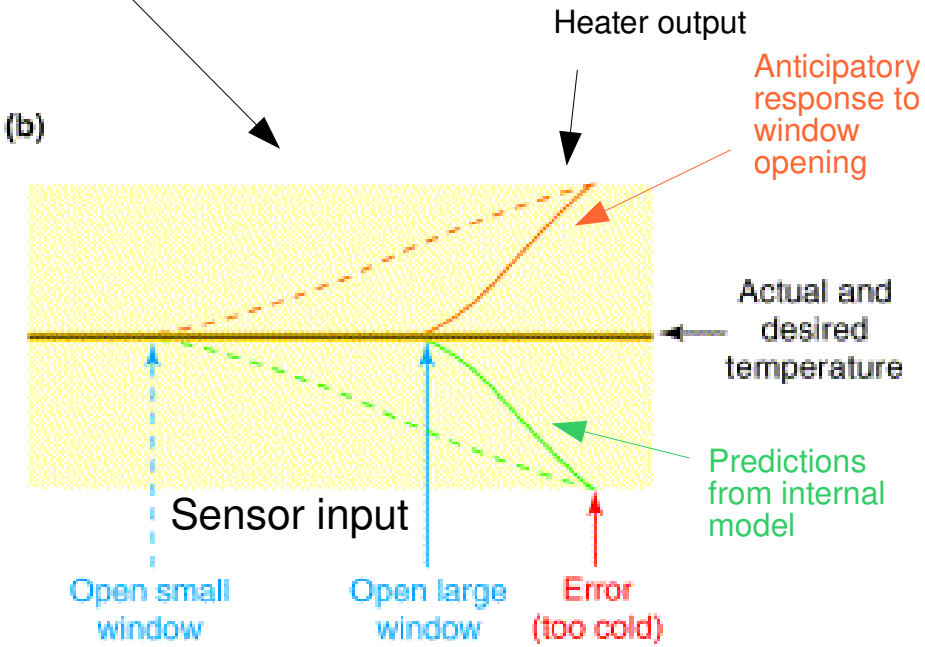

TRENDS in Neurosciences

High latency
Residual errors
Subject to oscillations if gain too high

Anticipatory response gives low latency
Better accuracy (lower error)
Sensors tell us the system state
Control requires an internal model that includes timing information.

Pavlovian Eyeblink Conditioning

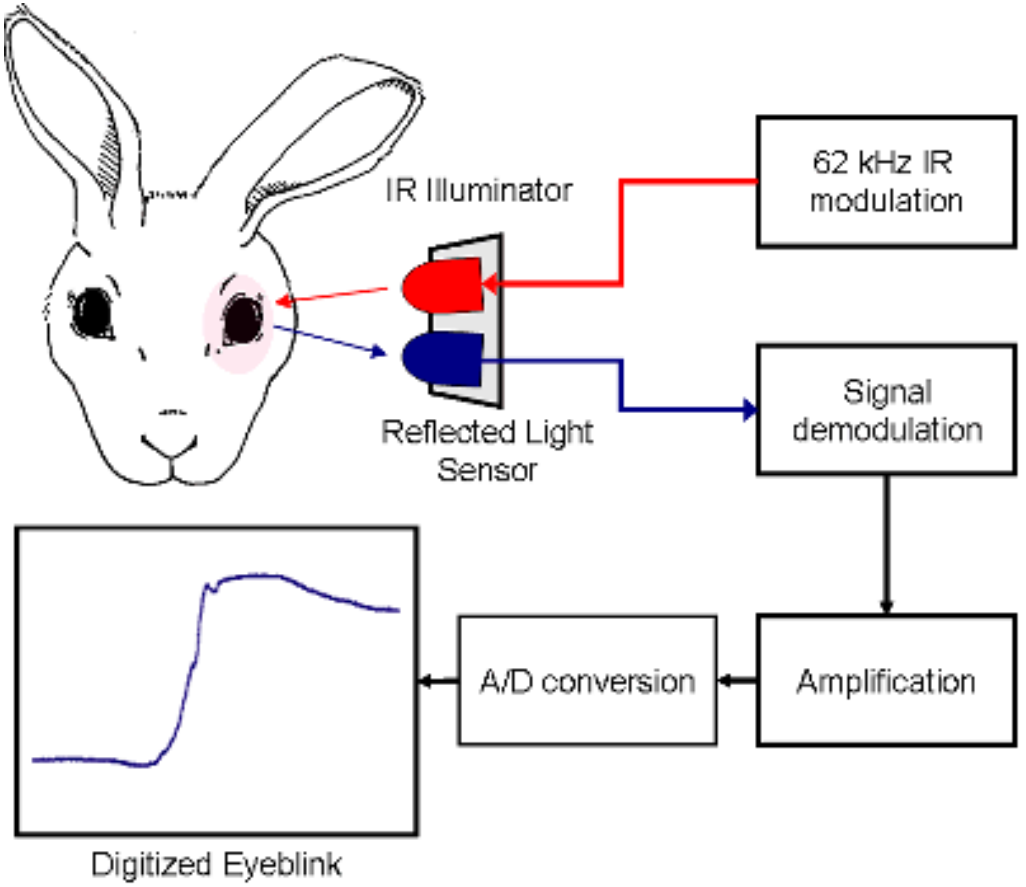

Eyeblink Conditioning in Humans

Eyeblink Conditioning System
from San Diego Instruments

- Measure cognitive development
- Impaired by mental disorders:
  - Schizophrenia
  - OCD
- Fetal alchohol syndrome

Delay vs. Trace Conditioning

- **Delay conditioning**: CS stays on until US arrives (up to 4 secs)

- **Trace conditioning**: CS comes on and then goes off again. US must be associated with the *memory trace* of the CS. Trace can be up to 2 secs in duration.

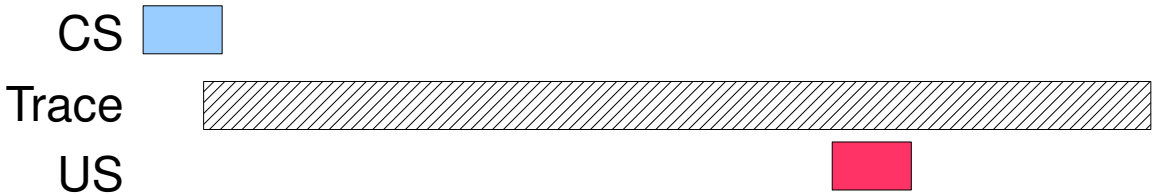

- Trace conditioning takes about 5x as many trials to learn.
- Trace conditioning (but not delay conditioning) is disrupted by lesions of hippocampus or medial prefrontal cortex.

Effect of Inter-Stimulus Interval (ISI)

- ISI must be 100-3000 msec (ideal is 200-500 msec)
- The learned CR (blink) is timed to just precede the US (air puff).
- Several hundred trials required for long ISIs
- Long ISIs also generate a broader response
- ITI (Inter-Trial Interval) is the time between successive trials. Should be long and somewhat random.

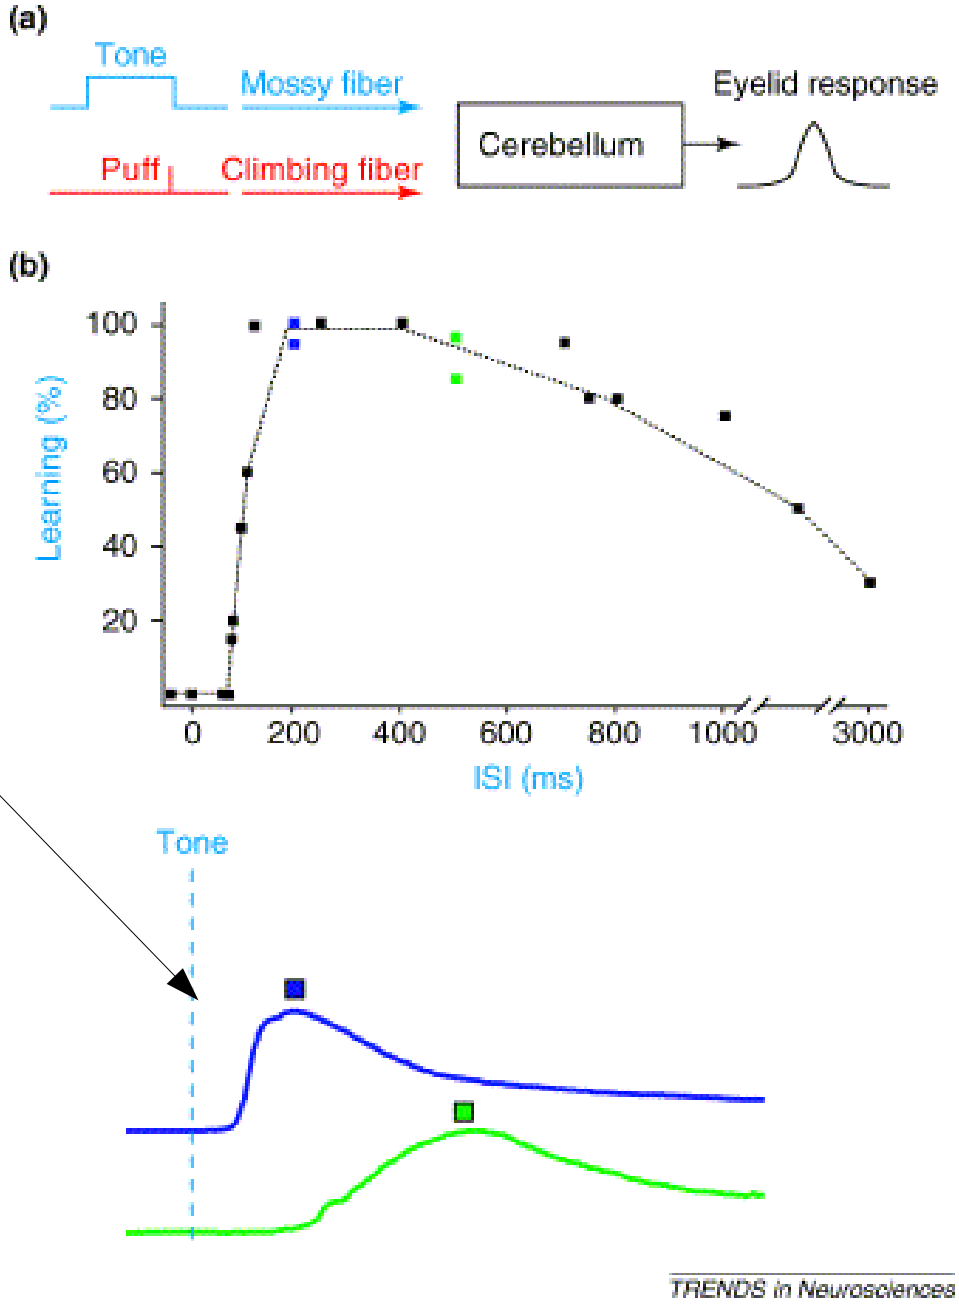

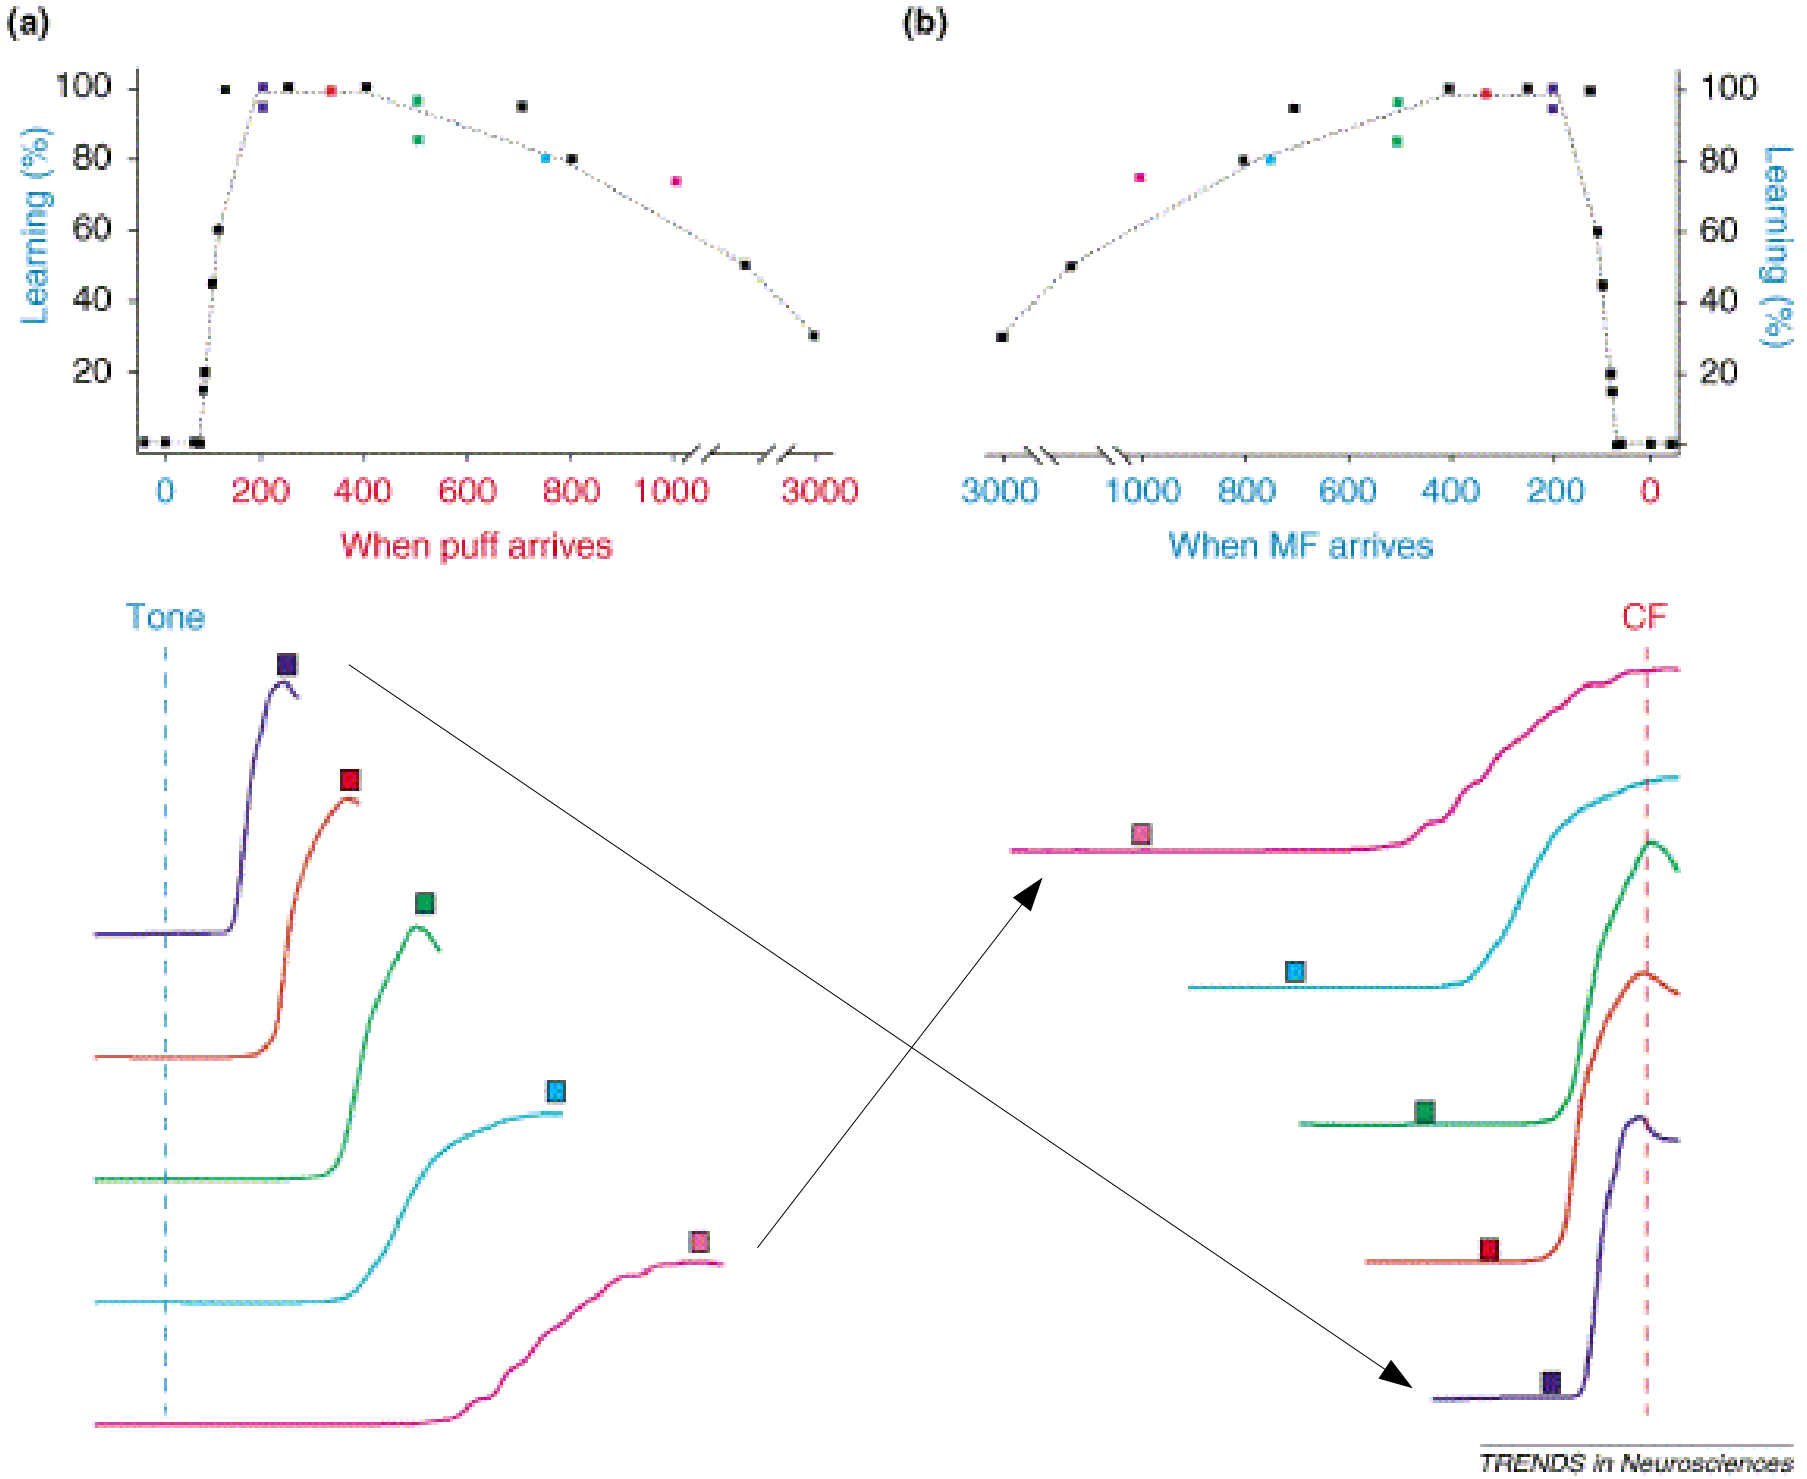
(a)
Learning (%)
100
80
60
40
20
0
200
400
600
800
1000
3000
When puff arrives
(b)
Learning (%)
When MF arrives
Tone
CF
TRENDS in Neurosciences

Mixing 200 ms and 700 ms ISI Trials

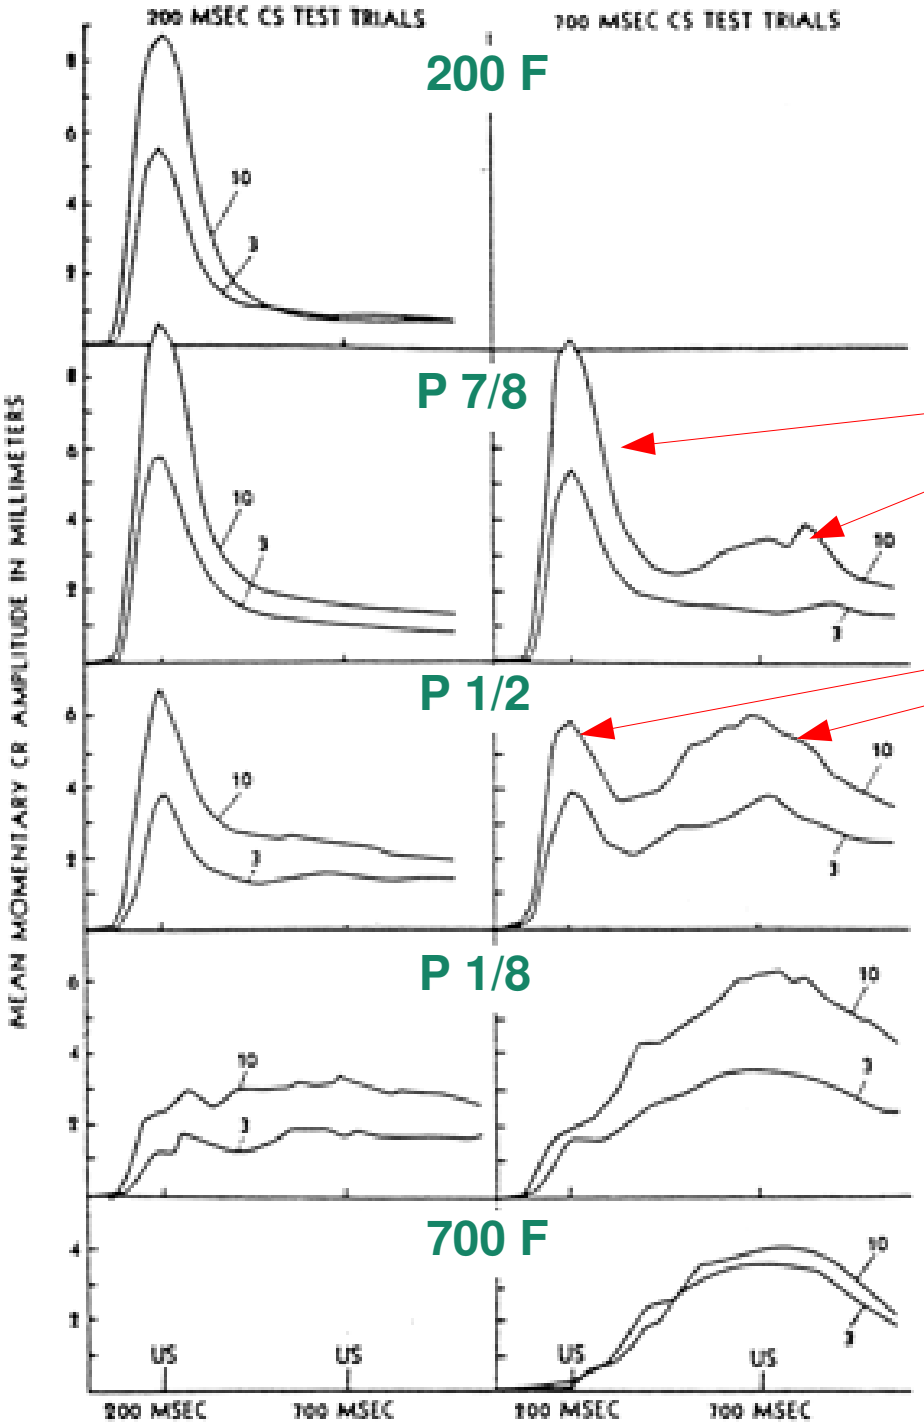

Two responses

Two responses

Eyelid Conditioning Circuitry

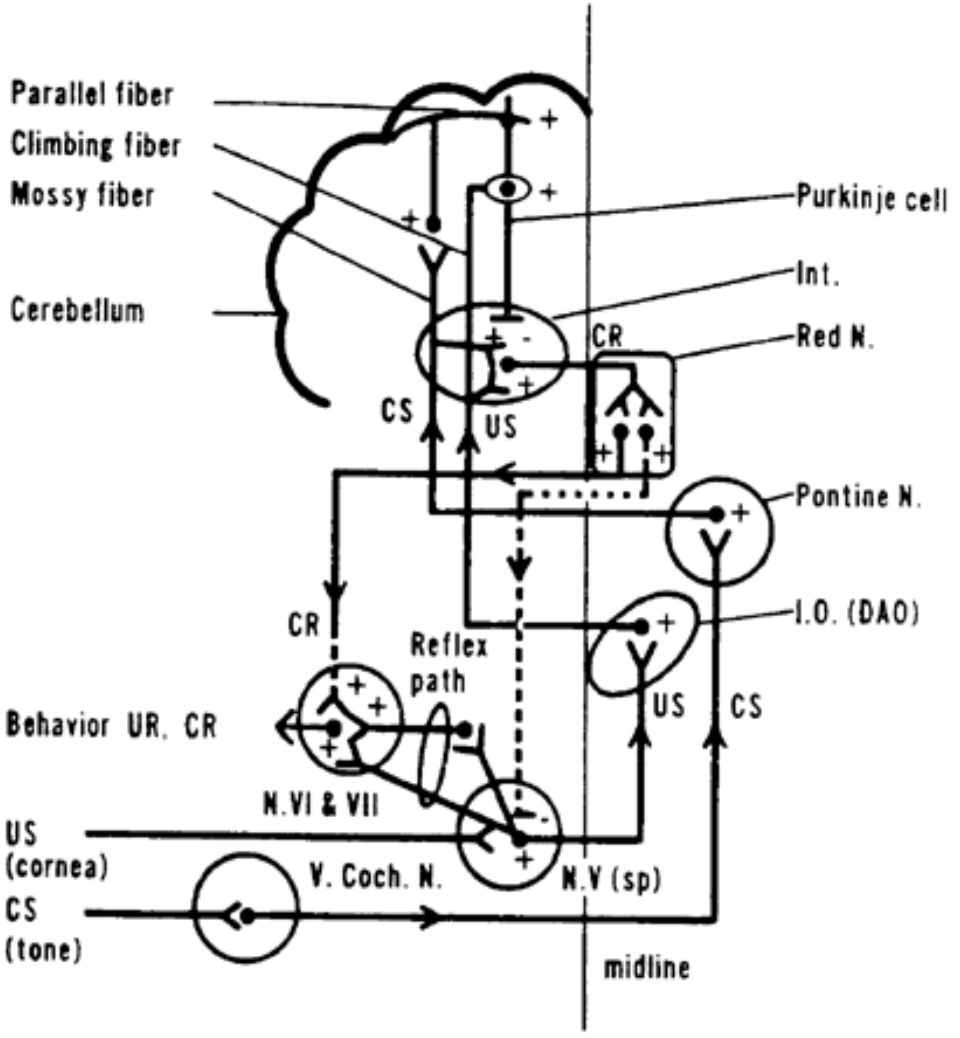

Effects of Lesions

- Lesioning the cerebellar cortex disrupts response timing but does not abolish the response entirely.
- Associative learning can still occur, but responses have very short latency (timing is off).

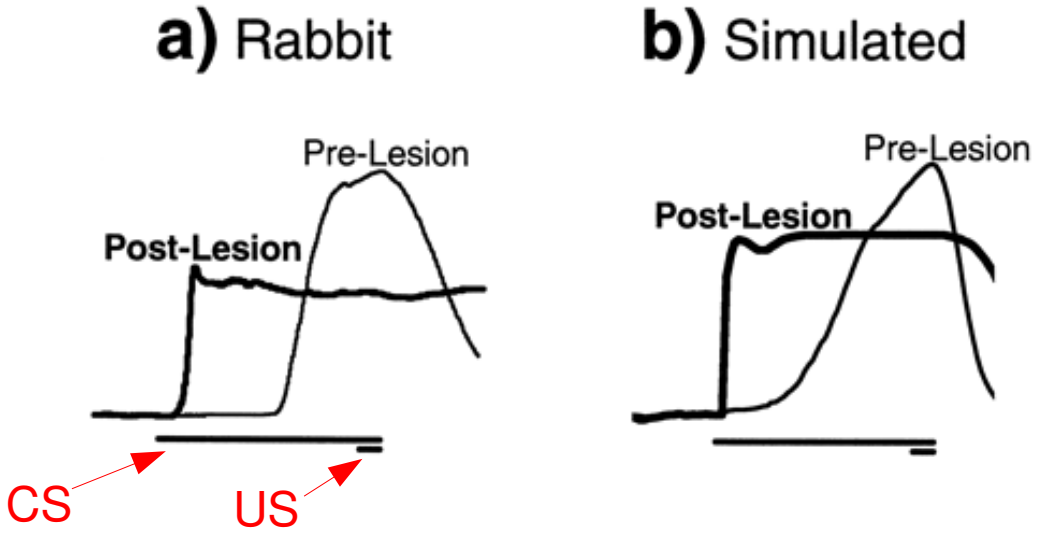

- Conclusion: two sites of Pavlovian learning in the cerebellum:
  - Interpositus nucleus learns to respond to the CS (mf → nuc)
  - Cerebellar cortex fine tunes the temporal response (pf → Pk)

Theories of Cerebellar Response Timing

a) Tapped delay lines

b) Spectral timing models
   i) PCs with fixed timing
   ii) PCs w/adjustable timing

c) Conjunctions of oscillators

d) State machines:
   i) Mauk & colleagues
   ii) liquid state machines

e) Selectable “timing units”

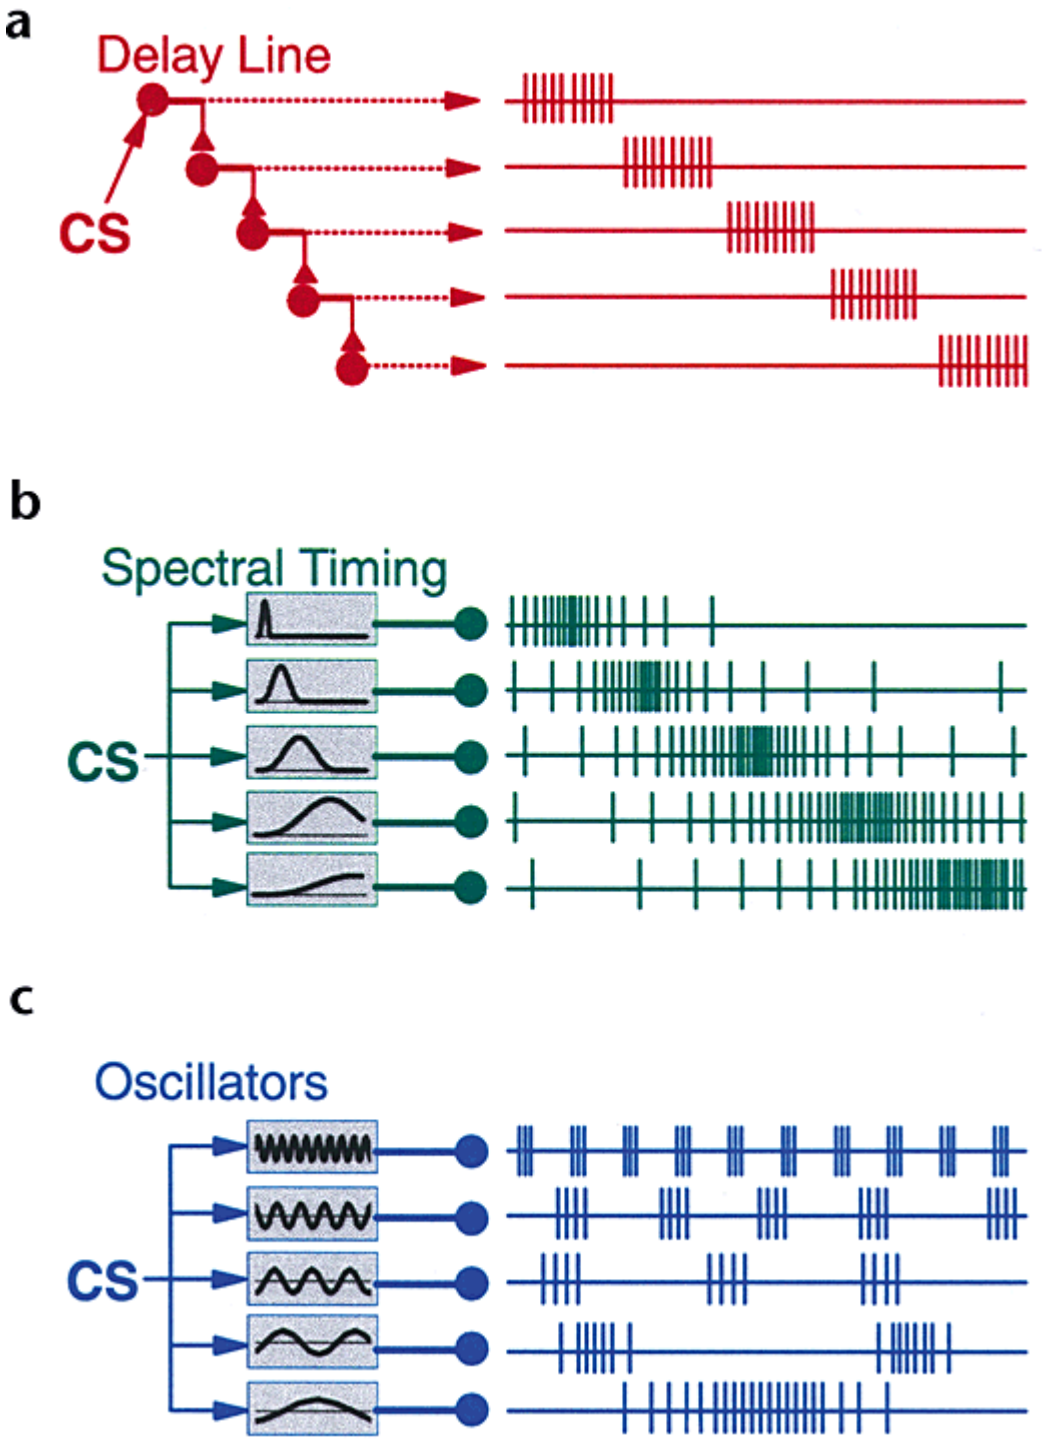

Medina & Mauk (2000) Simulation

600 mossy fibers
10,000 granule cells
900 Golgi cells
60 basket/stellate cells
20 Purkinje cells
6 nucleus cells
> 300,000 synapses

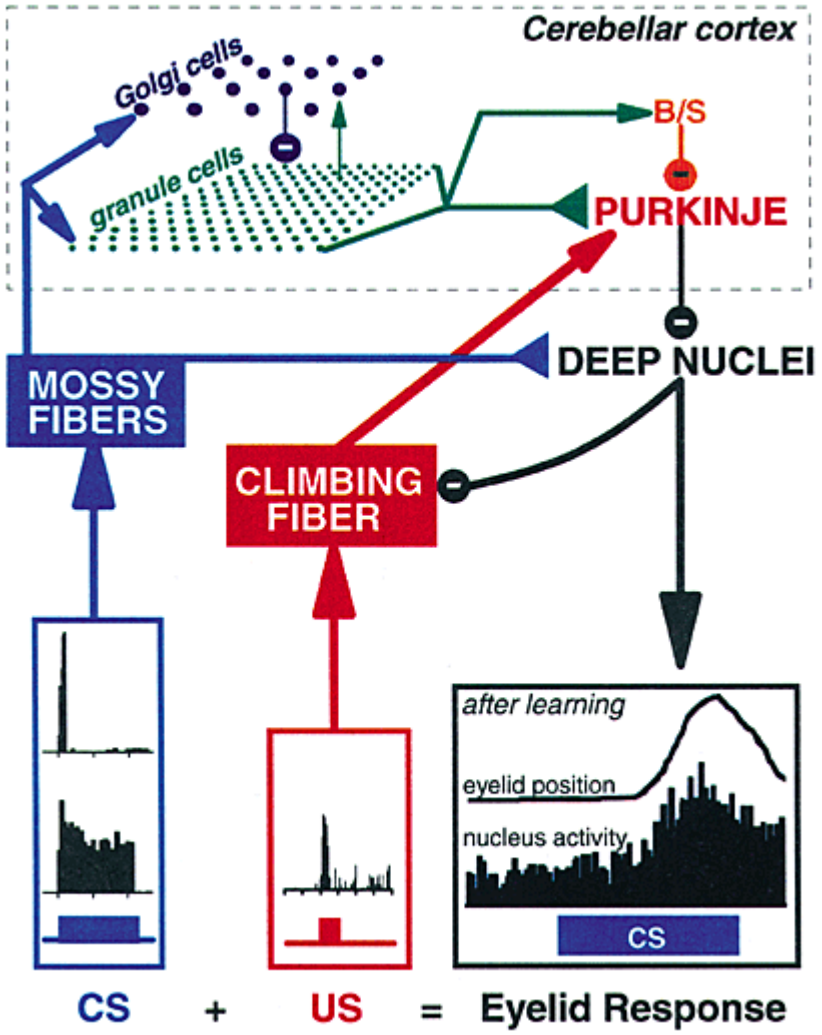

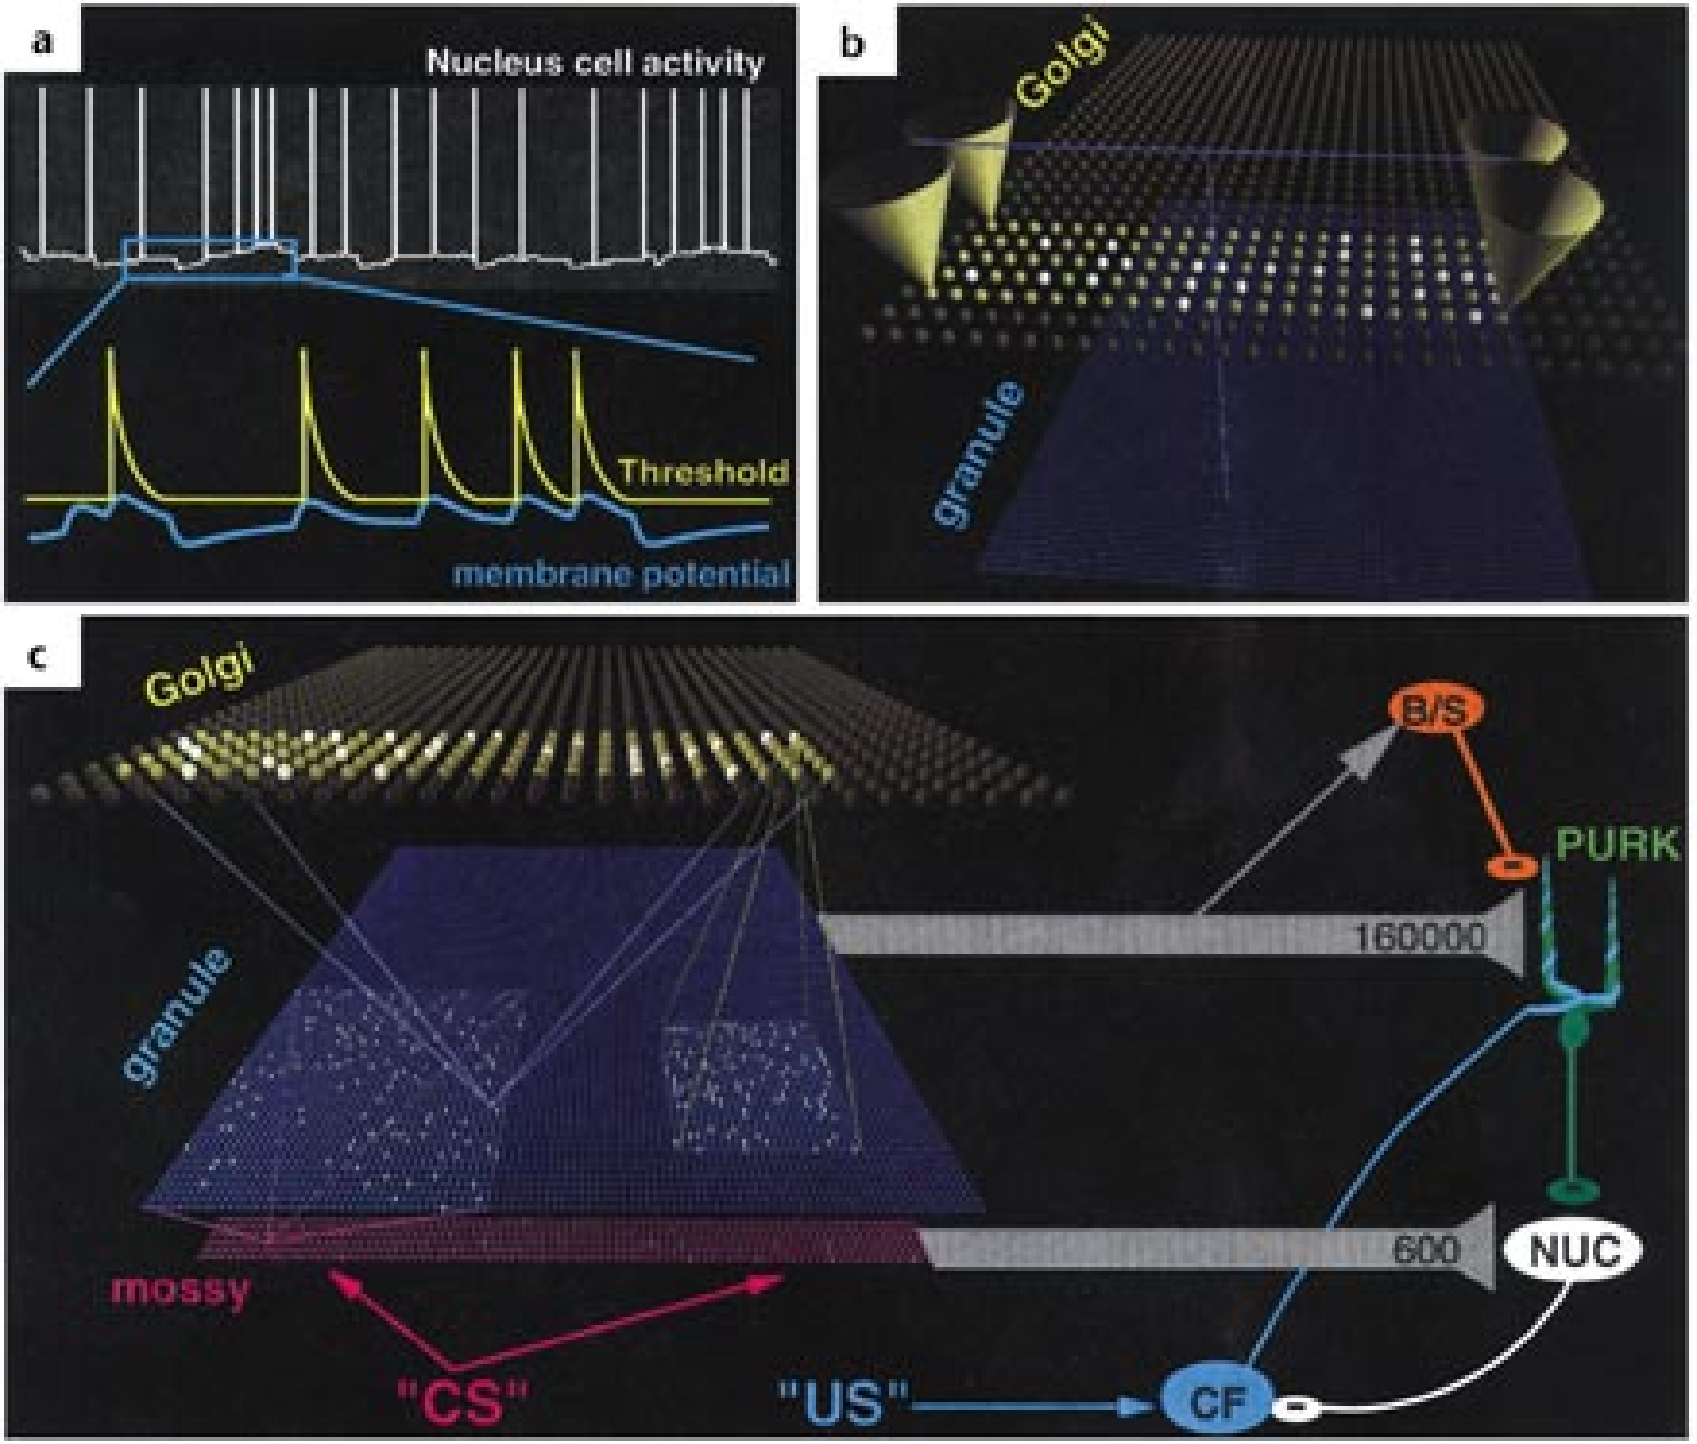
a
Nucleus cell activity
Threshold
membrane potential
b
Golgi
granule
c
Golgi
granule
B/S
PURK
160000
600
NUC
mossy
"CS"
"US"
CF

More Simulation Details in the J.Neurosci. Paper

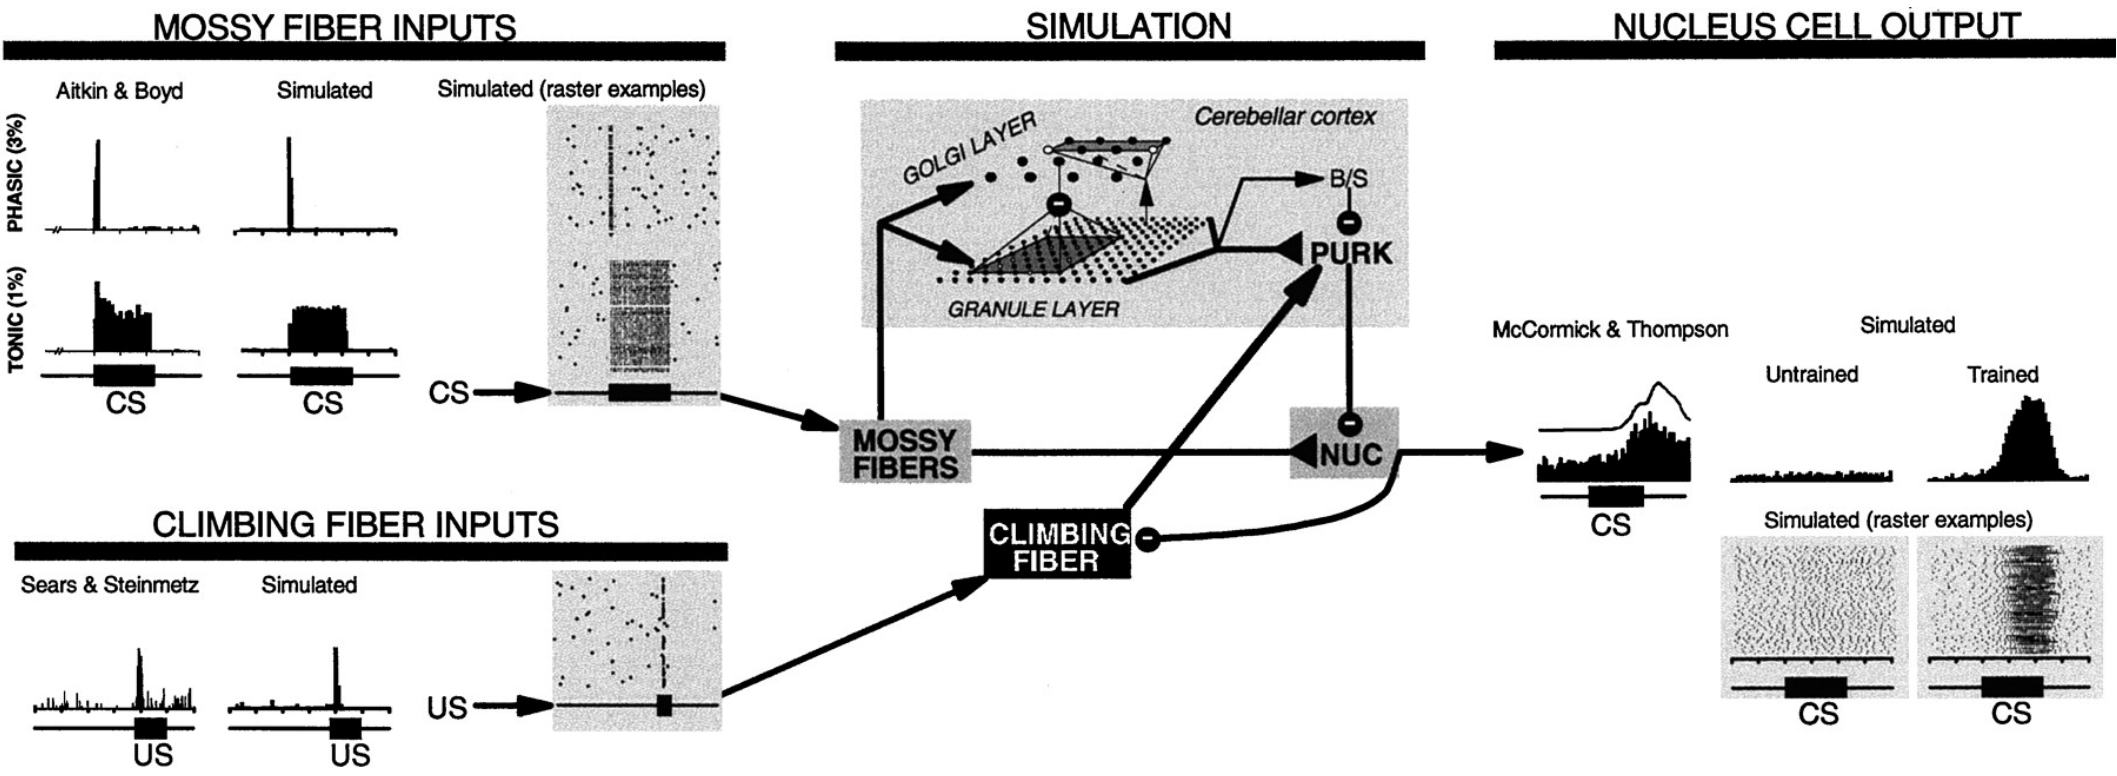

Realistic mossy and climbing fiber inputs based on experimental data.

Response Timing in the Model

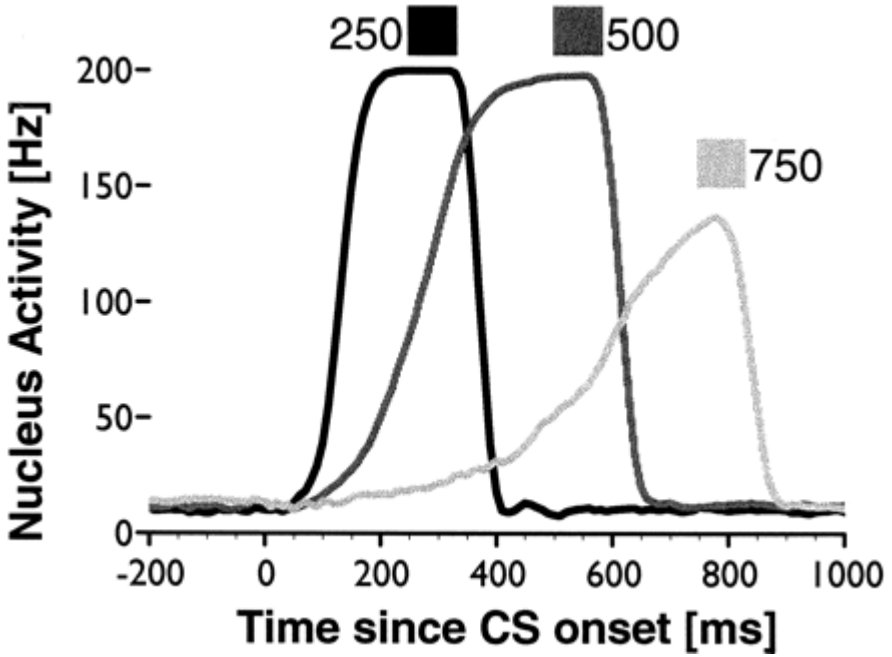

ISI values of 250, 500, and 750 msec.`

LTP + LTD

- Granule cells exhibit a variety of broad temporal responses
- LTD alone produces an overly broad CR (right).
- But LTP + LTD together produces a precisely timed response by combining inputs from multiple Purkinje cells to keep the DCN inhibited until just before the US is expected to arrive.

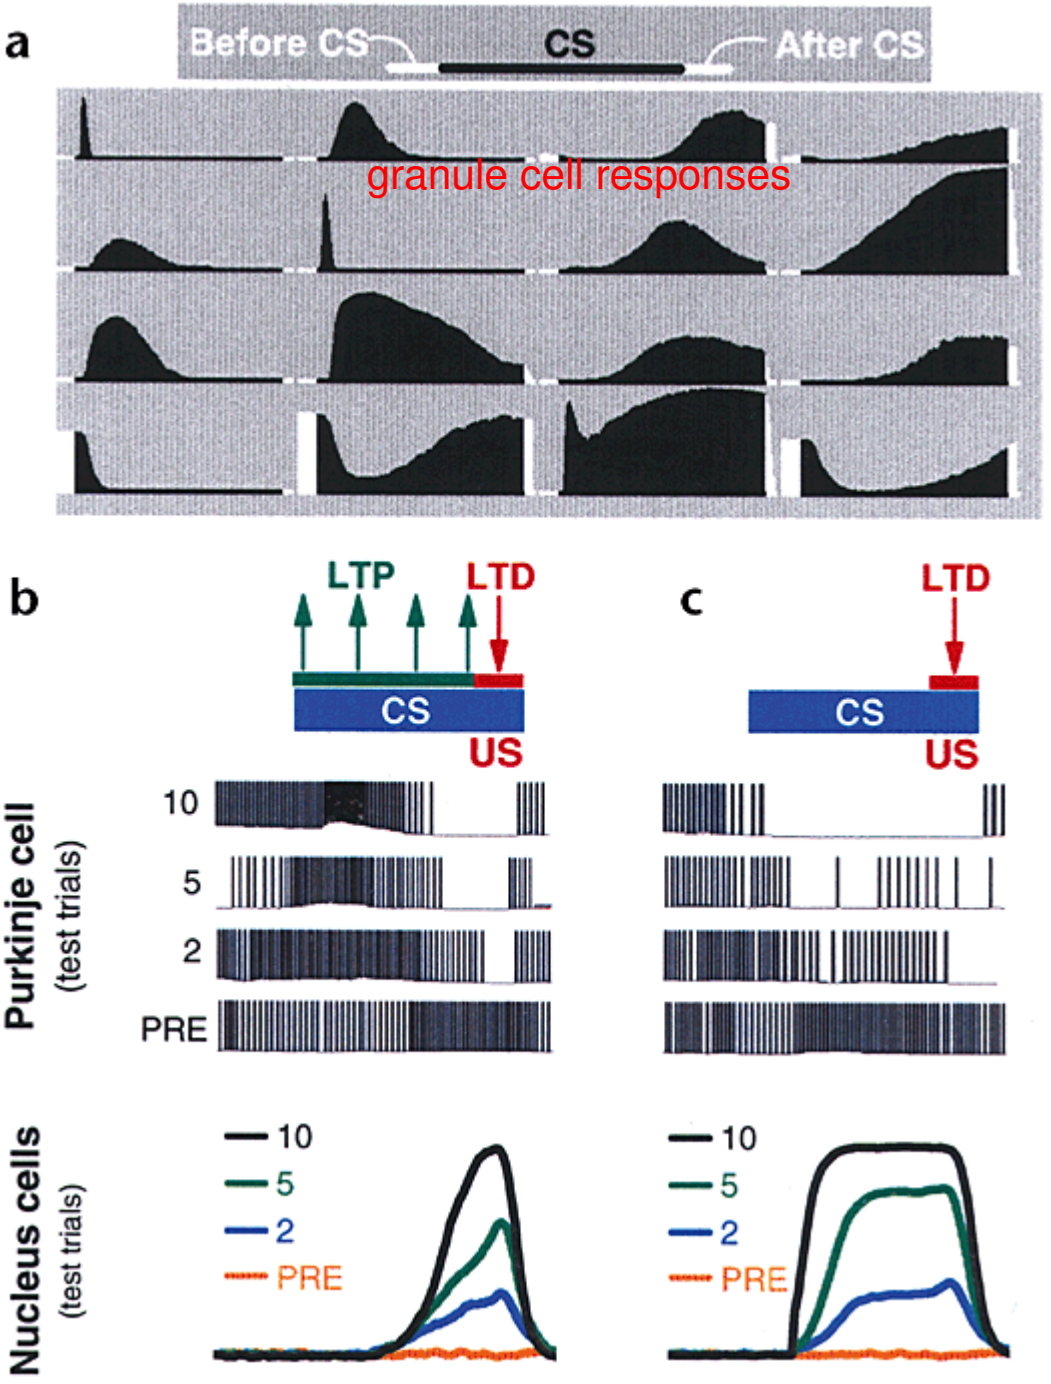

Time Course of Learning and Response Shaping

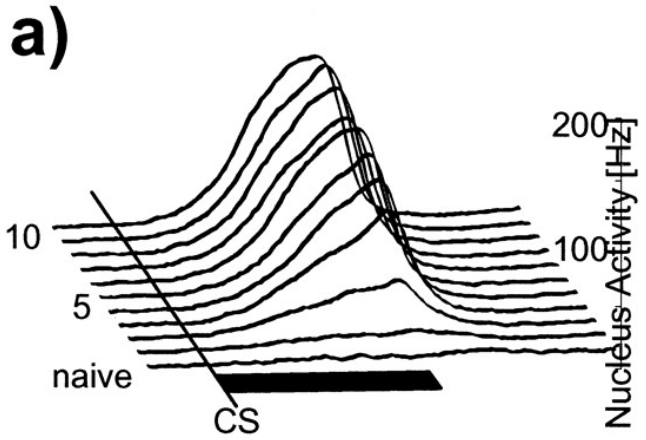

Nuclear cell

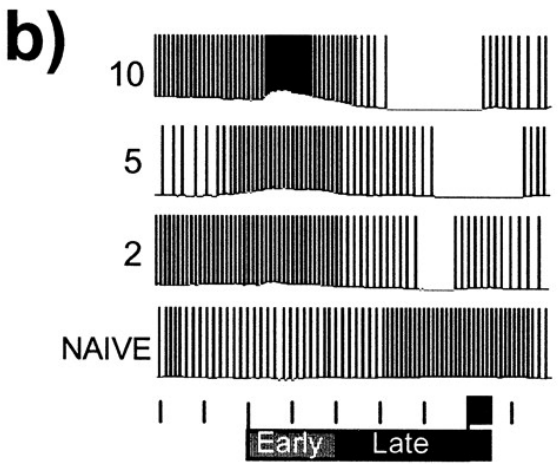

Simulated Purkinje cell

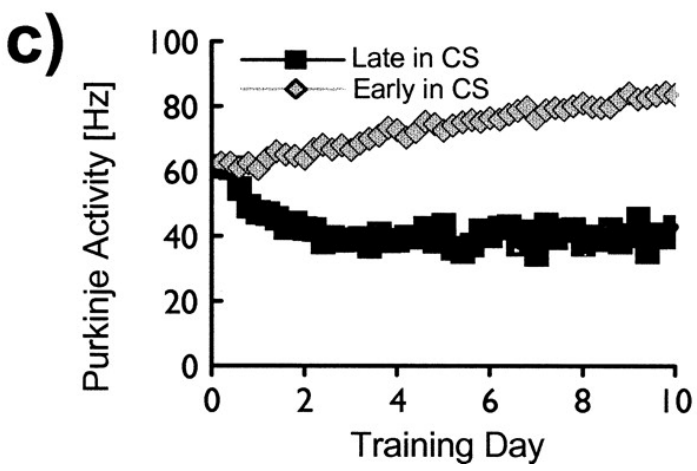

Early LTP + Late LTD

Learning With LTP Disengaged: Response Timing is Poor

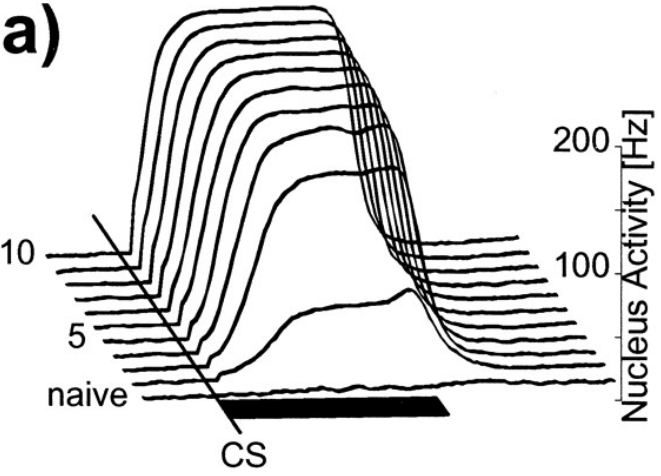

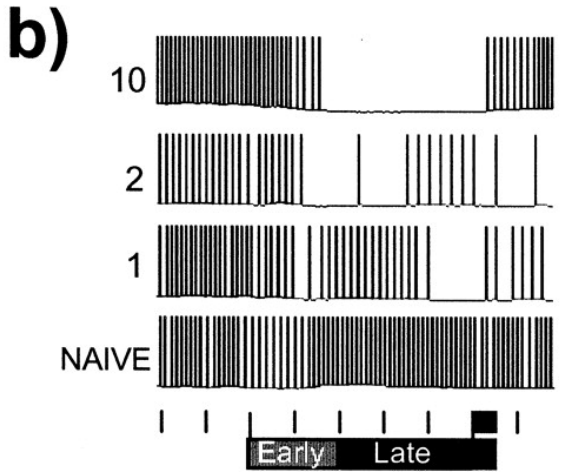

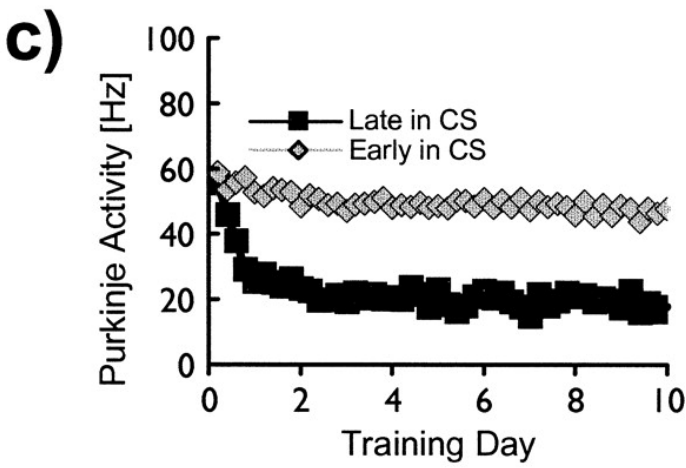

Recovery After Partial Lesion to Cerebellar Cortex

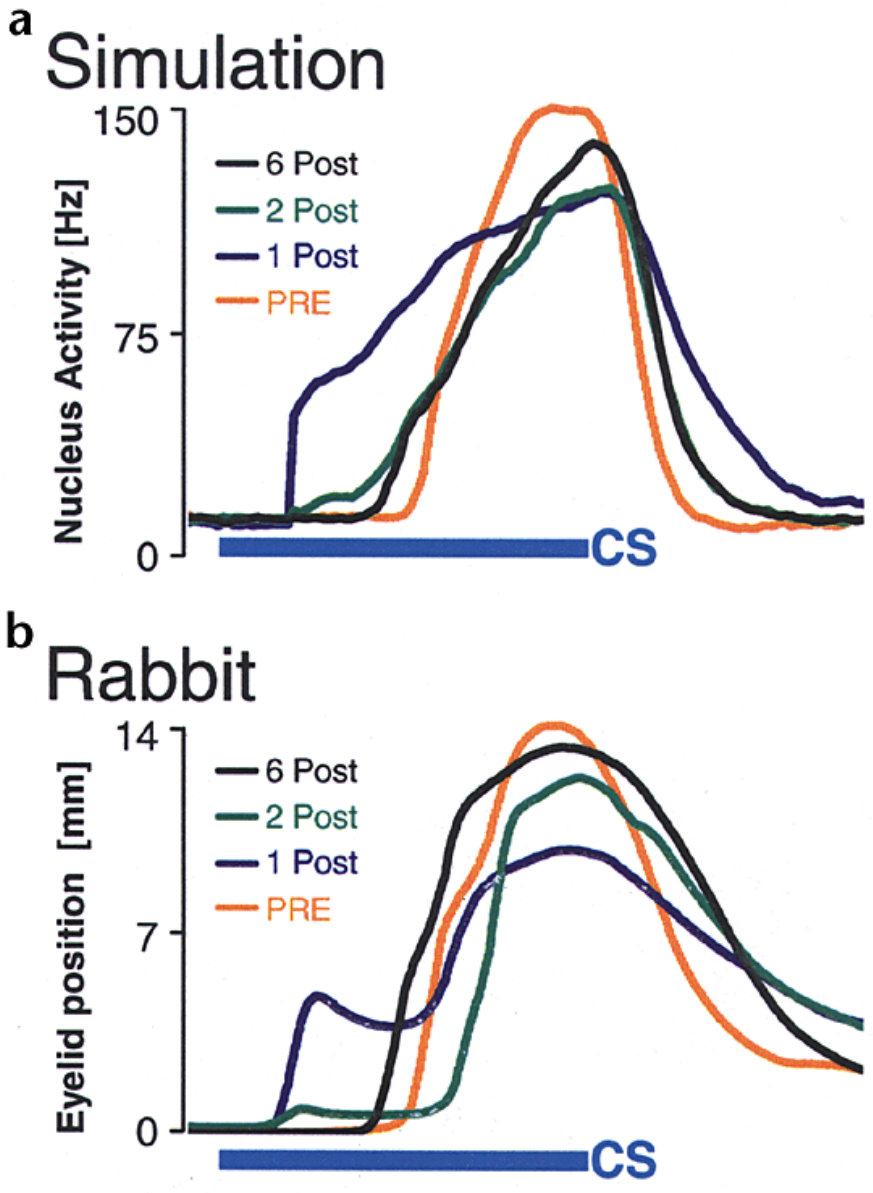

Recovery After Lesioning Cerebellar Cortex

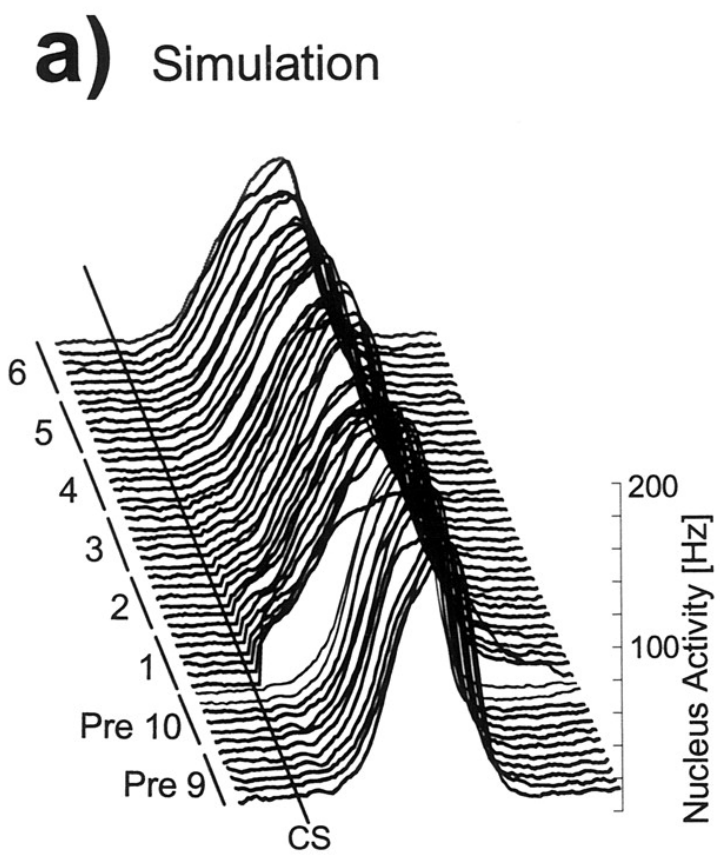

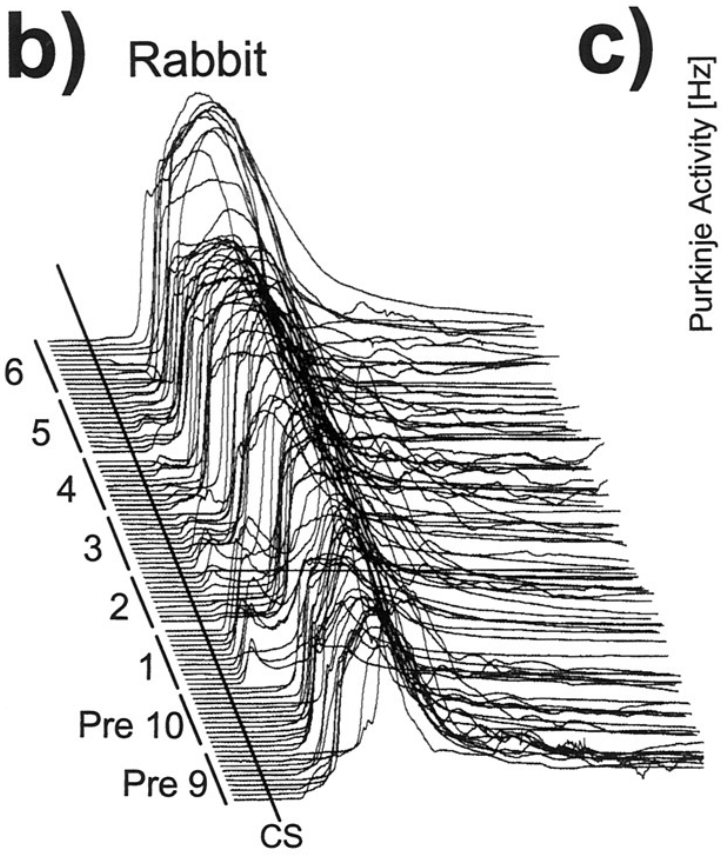

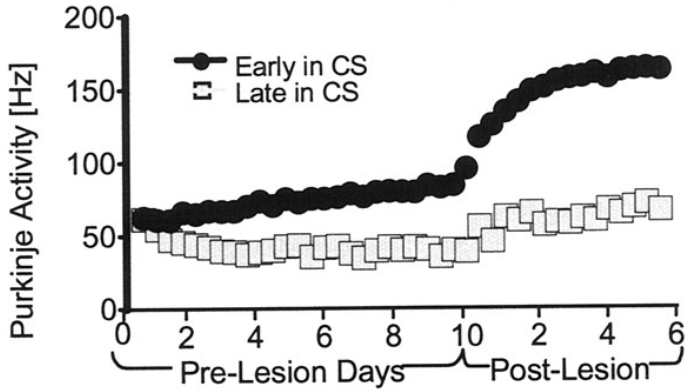

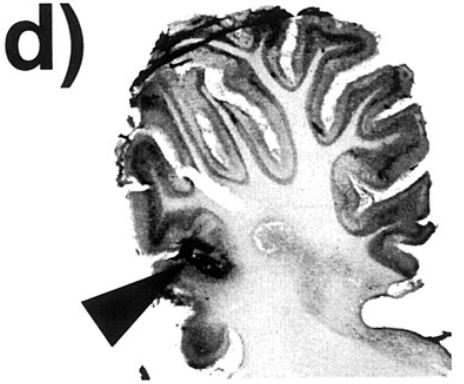

Why Do Long ISIs Prevent Learning?
Hypothesis: Too Much LTP Overwhelms LTD

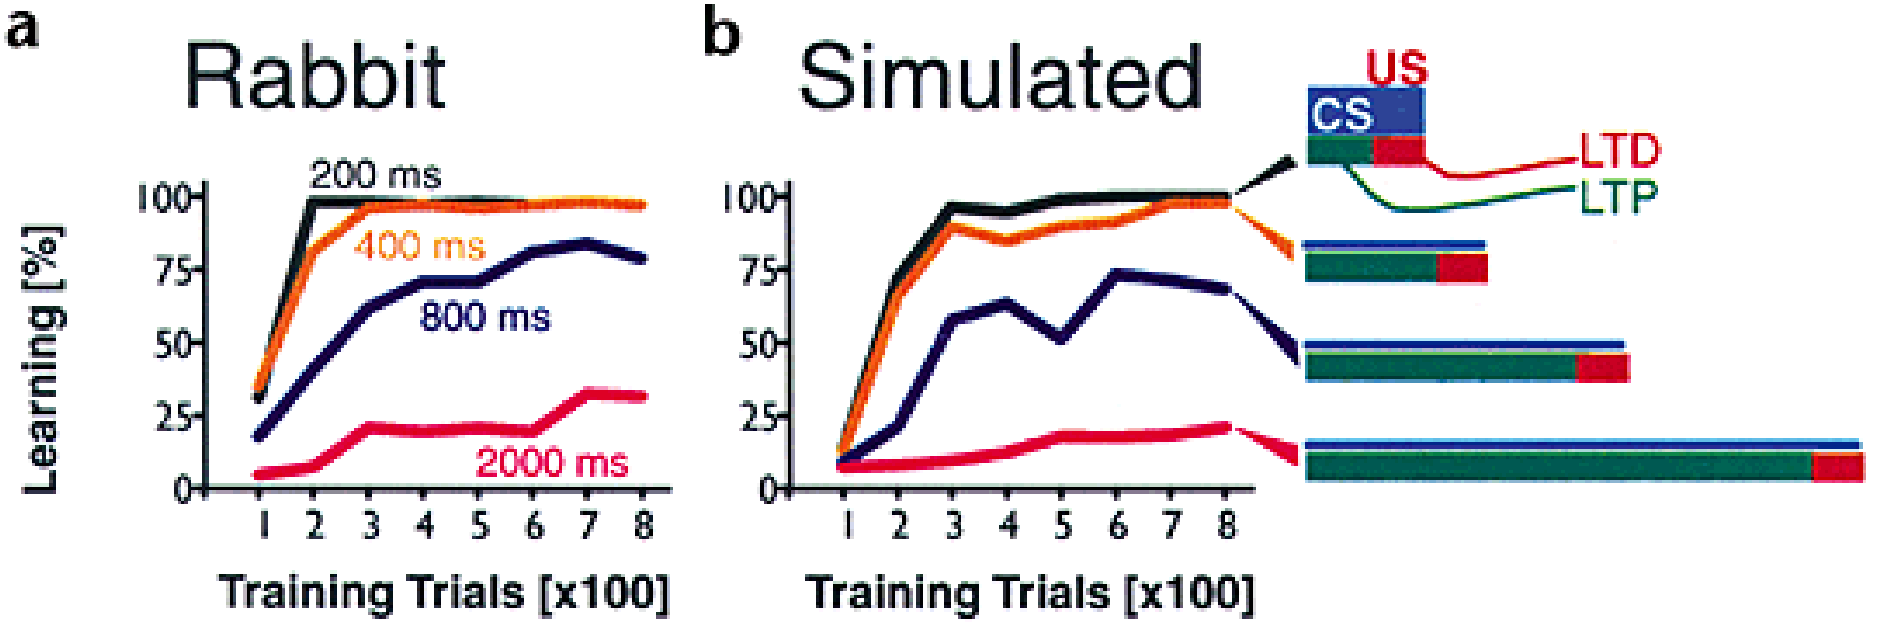

Scaling Up to 1 Million Granule Cells

- Li et al. (2013) scale up model using a GPU (NVIDIA GTX 580).
  - 1024 mossy fibers
  - 2^{20} = 1,048,567 granule cells (vs. 50 billion in humans)
  - 32 Purkinje cells (each with 32,768 granule cell synapses)
  - 128 basket cells, 512 stellate cells; 1024 Golgi cells
- Results for eyeblink:
  - Original model couldn't handle 1000 msec ISI
  - New model can (sort of) handle 1000 msec ISI
  - New model still can't handle 1150 msec ISI
- Results for cart-pole balancing task:
  - Mossy fibers encode pole angle, angular velocity, and acceleration
  - Two groups of opposed output cells, for left and right motion
  - Sort of works, with no special tuning

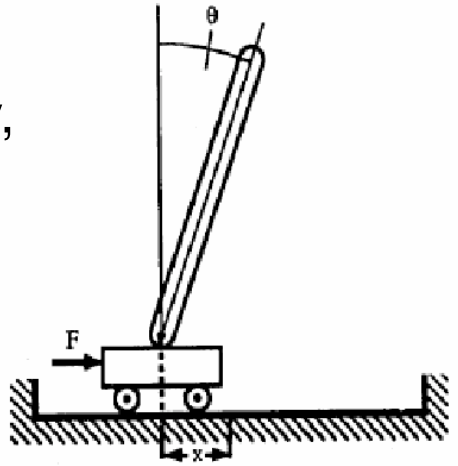

Cerebellar Cortex As a Liquid State Machine

Yamazaki and Tanaka, *Neural Networks* ,20(3):290-297, April 2007

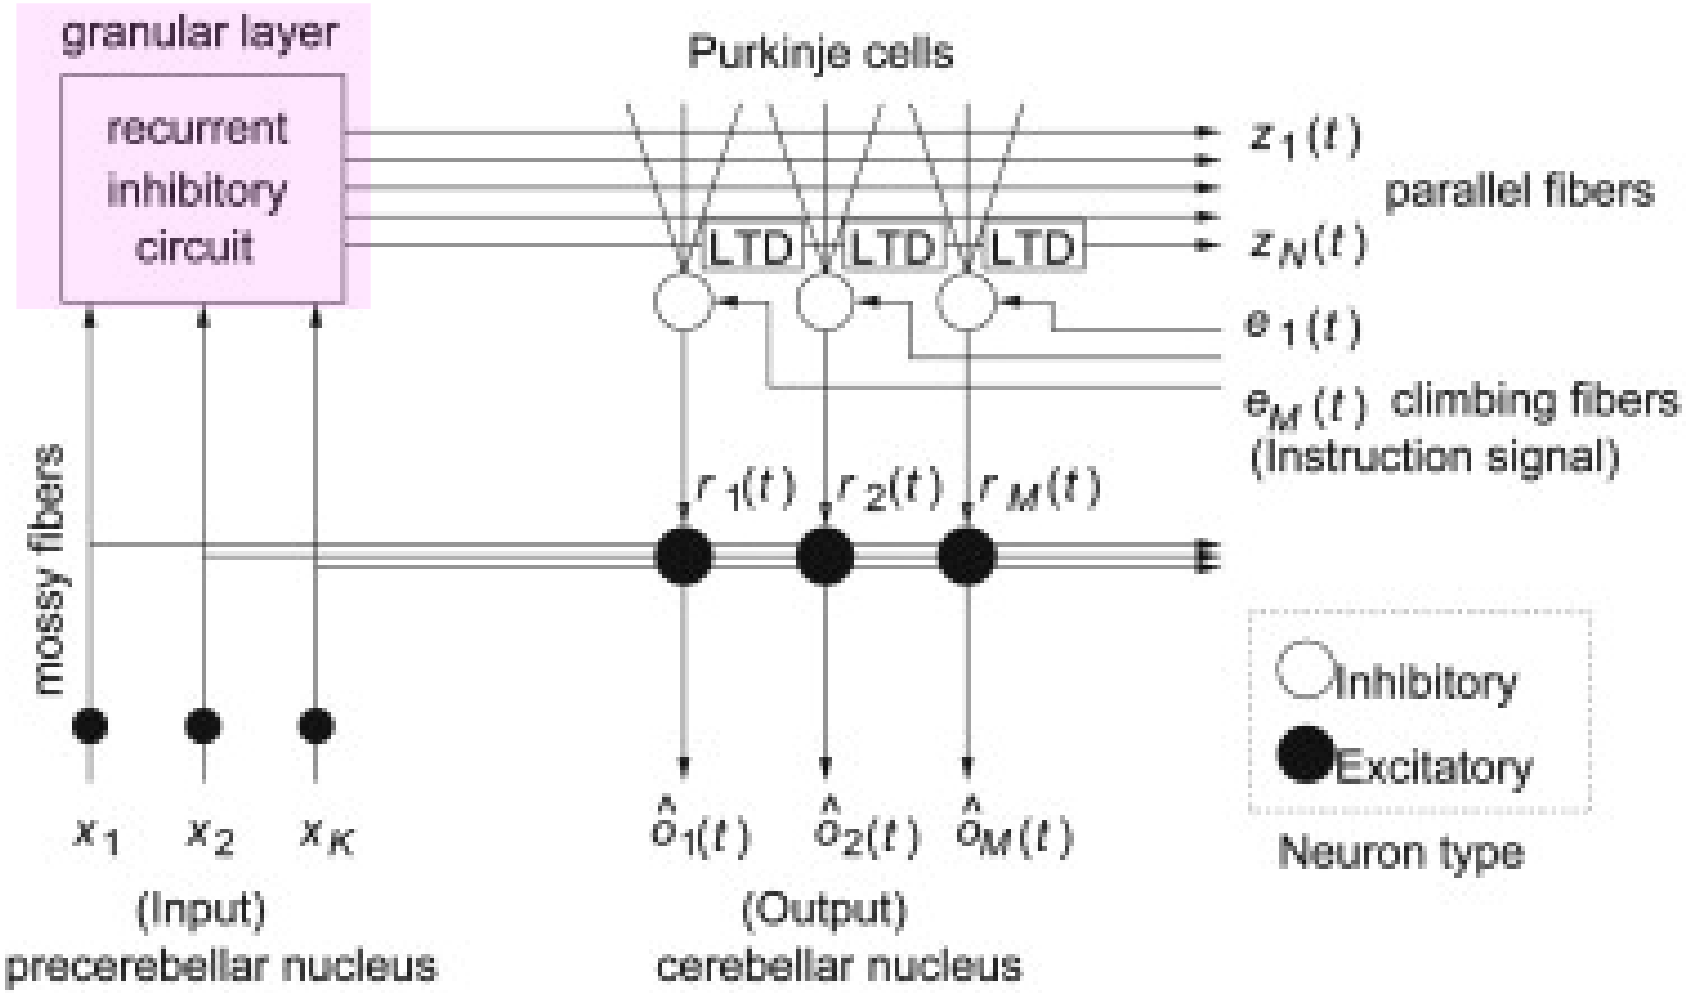

Rich Variety of Granule Cell Activity Patterns (Medina & Mauk Noted This Too)

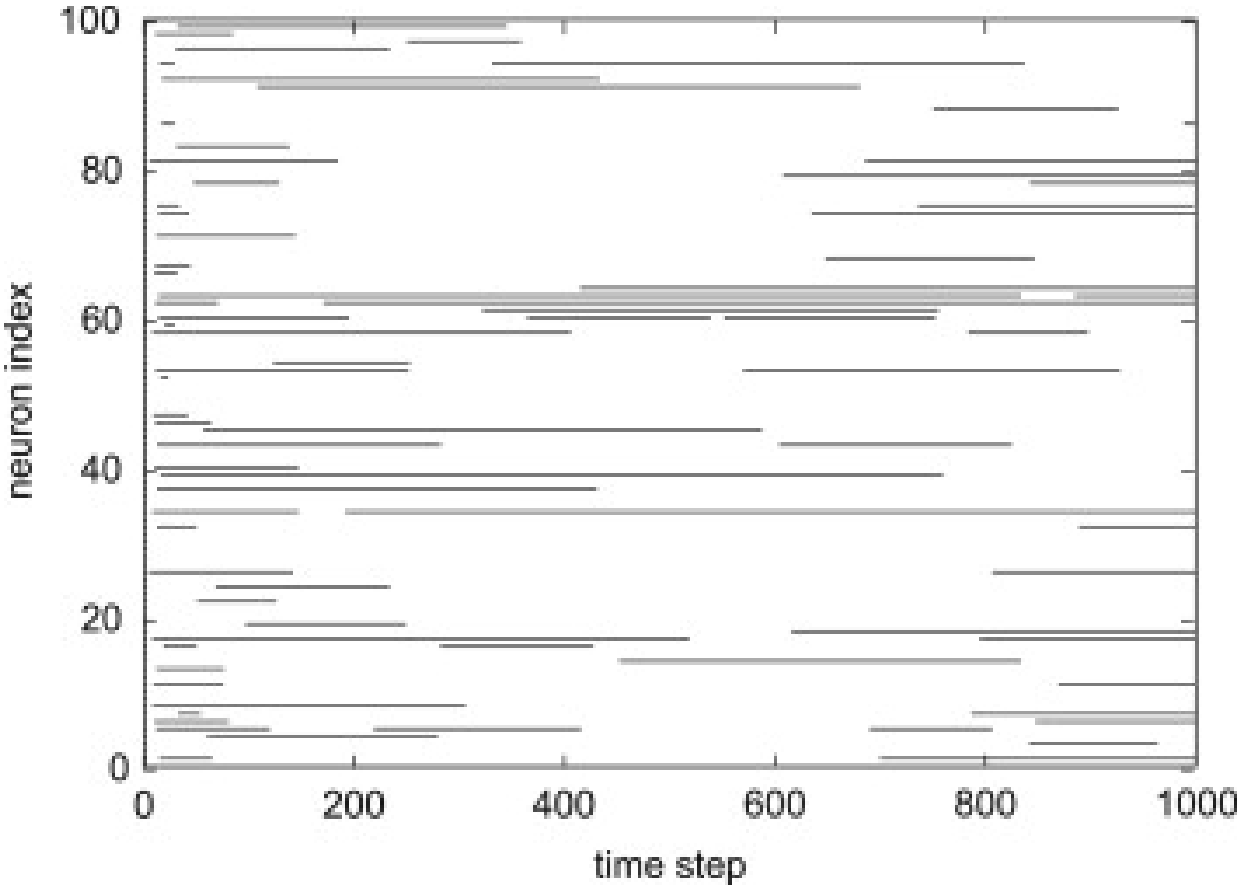

Similarity Index: Granule Cell Activity Patterns Evolve Over Time

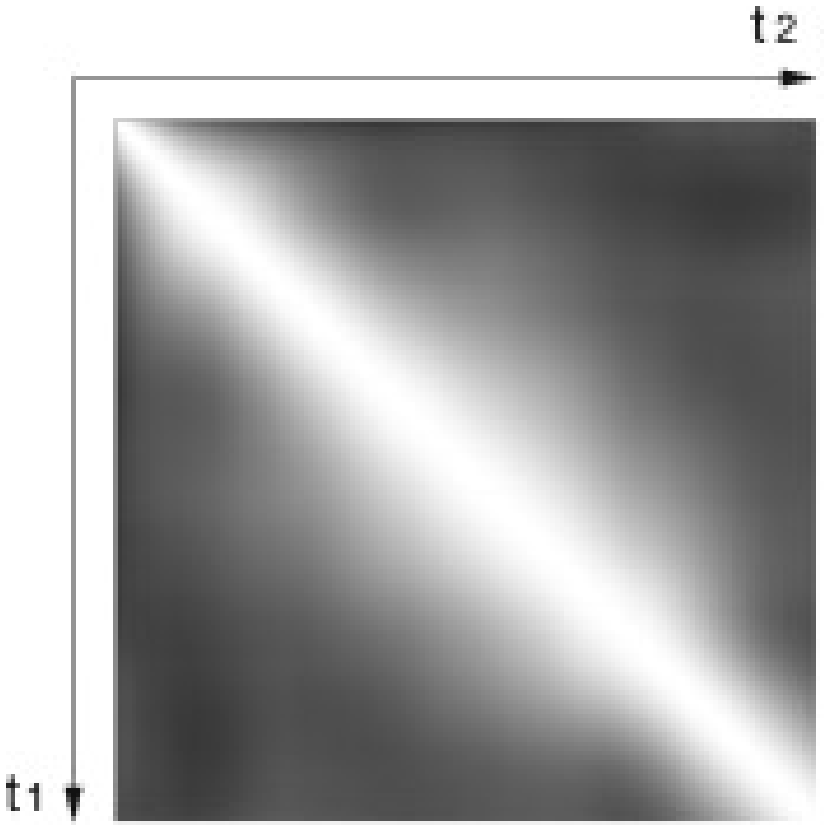

Correlation of LSM activity patterns at times t_1 and t_2.

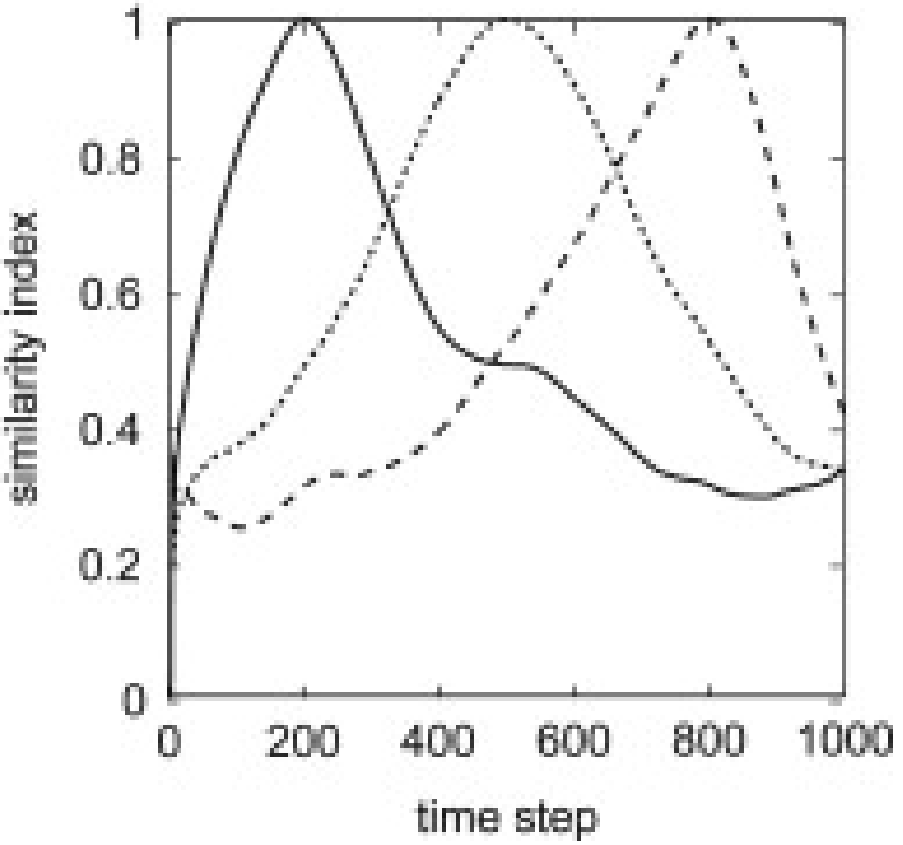

Slices through the graph at left at t=200, t=500, and t=800 show that similarity changes smoothly.

Cerebellum = Liquid State Machine + Perceptron

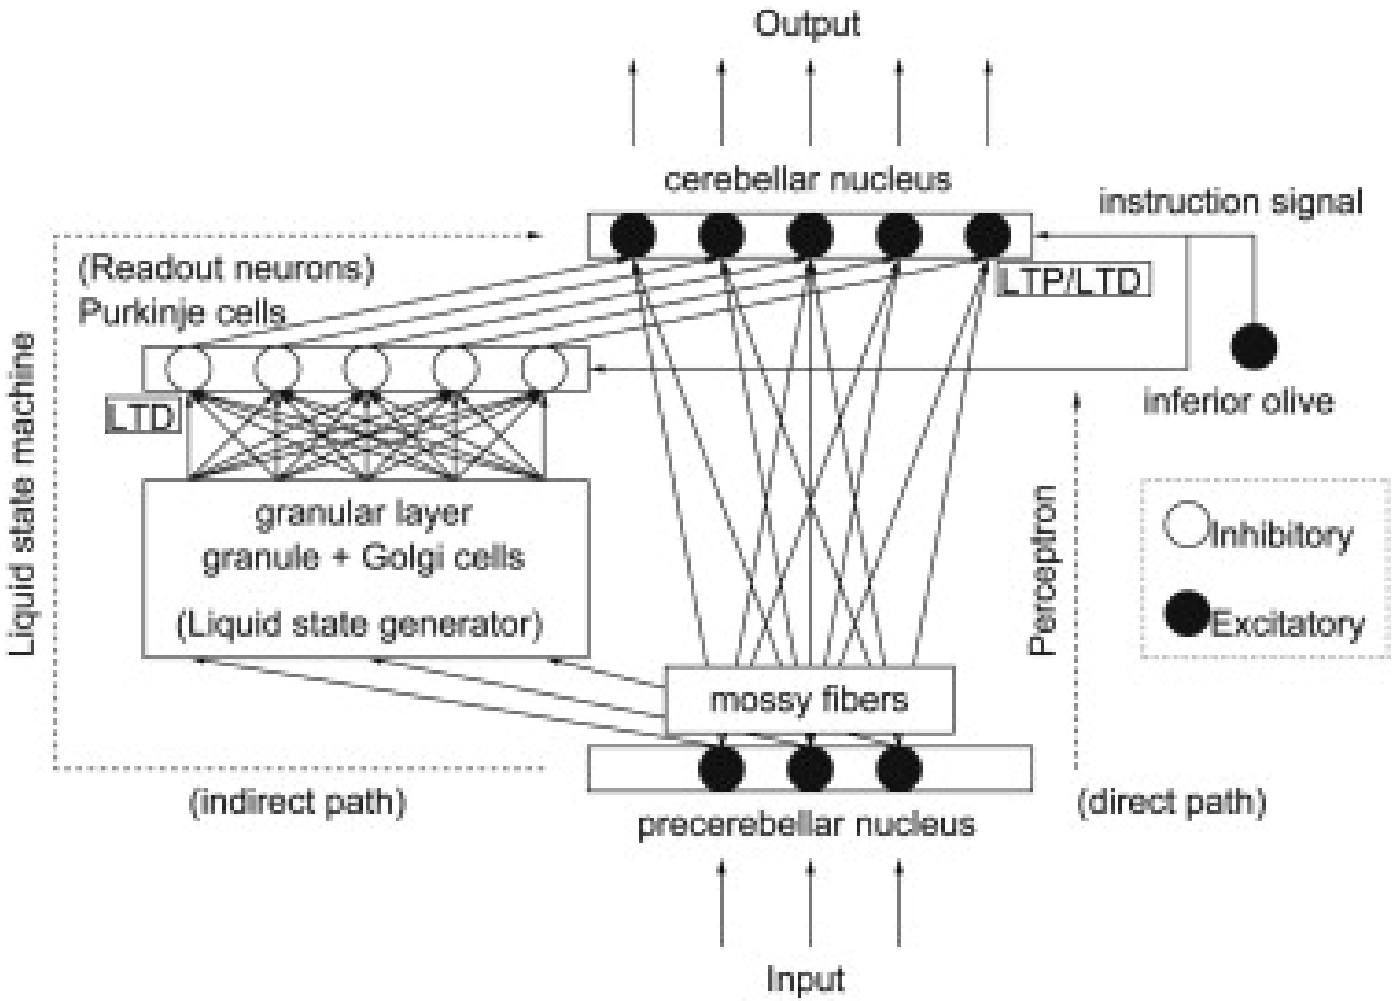

Fiala et al. Spectral Timing Model

Fiala, Grossberg, and Bullock, *J. Neurosci.* 16(11):3760-3774, 1996

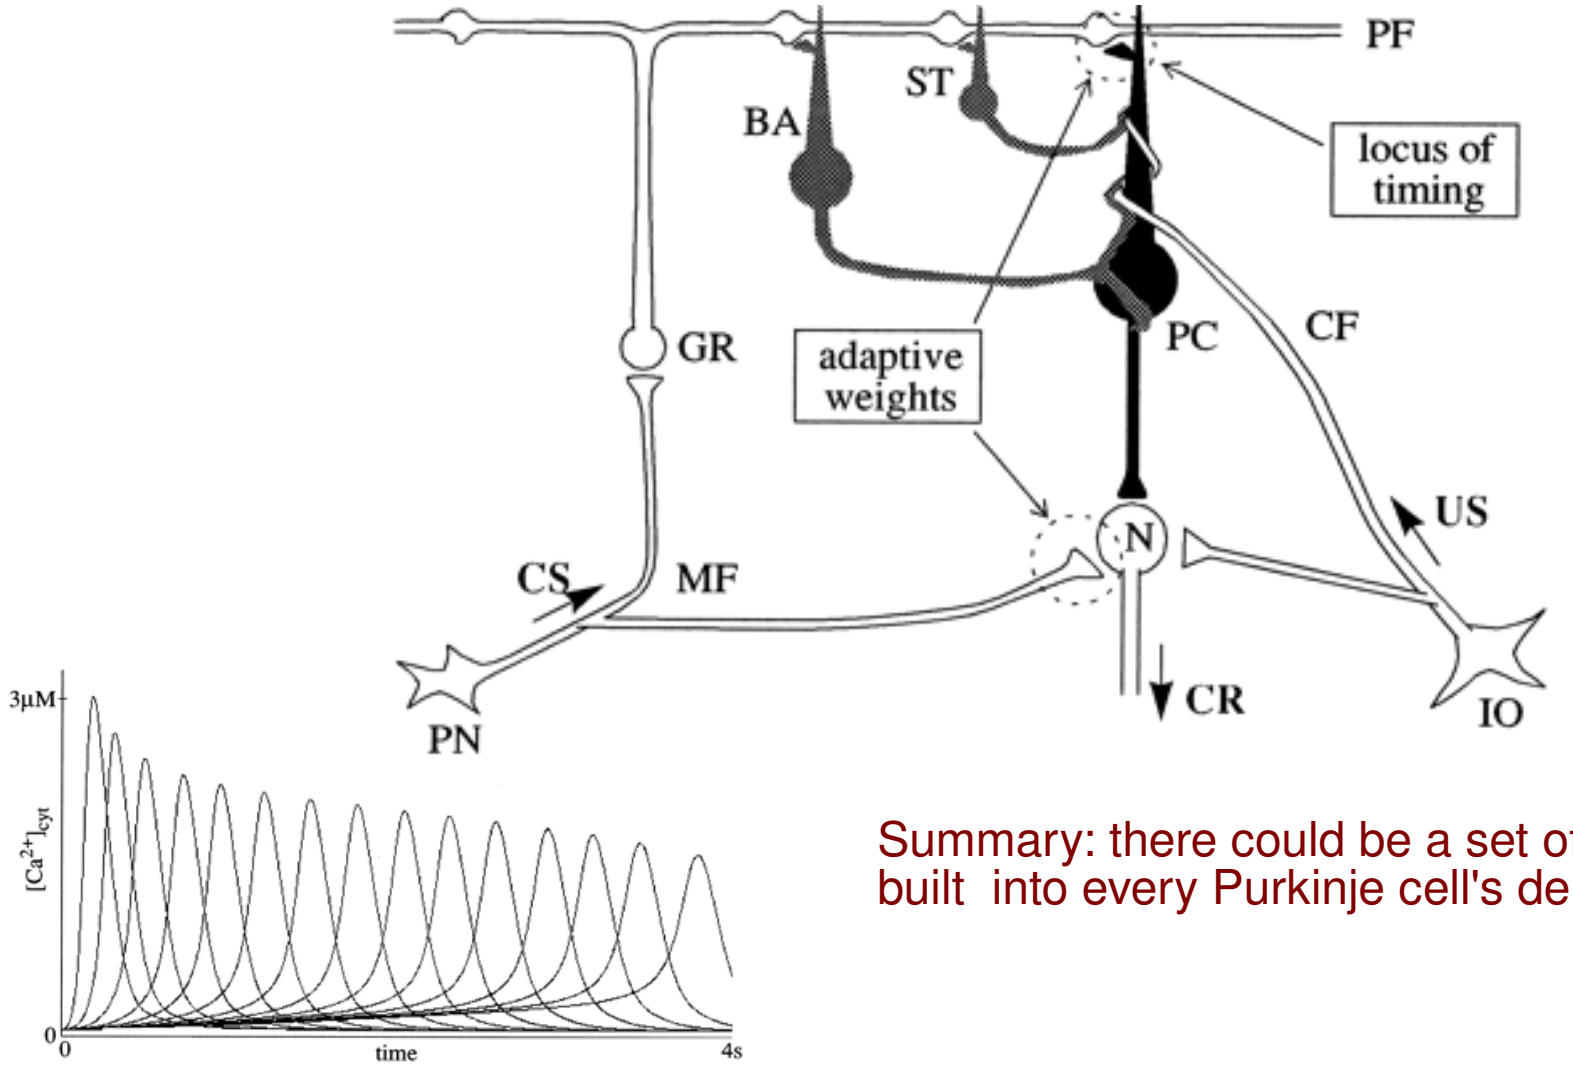

Summary: there could be a set of delay lines built into every Purkinje cell's dendritic tree.

Metabolic Transmission Pathway in Purkinje Cell Dendrites

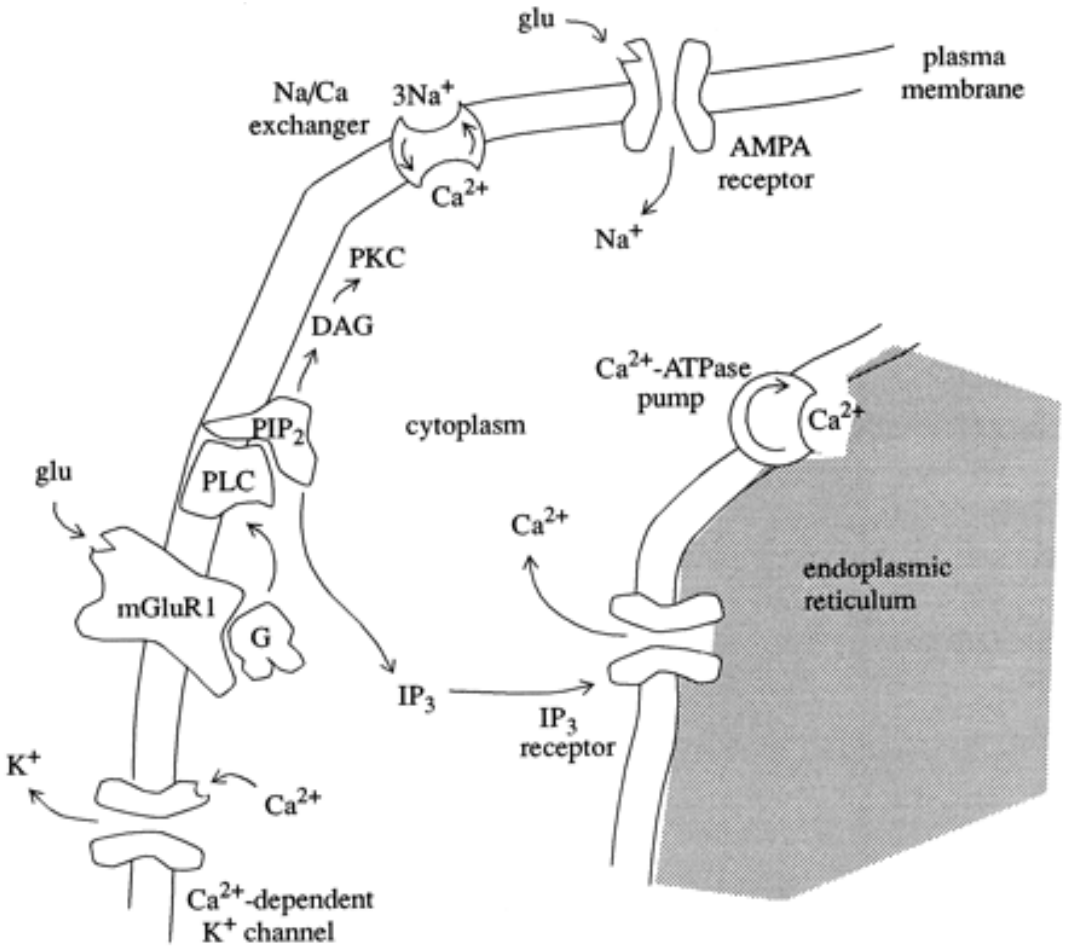

DAG = diacylgycerol
G = guanine nucleotide-binding protein
mGluR1 = metab. glutamate receptor
PKC = phospholipase C
PIP_2 = phosphatidylinositol 4,5-biphosphate
IP_3 = inositol 1,4,5-triphosphate, a second messenger
IP_3R = IP_3 receptor

Basic Story

- Glutamate binds to *mGluR1* receptors, causing second messenger IP_3 to bind to IP_3R receptor.
- IP_3R receptor causes release of calcium from storage in the endoplasmic reticulum (ER).
- Ca^{2+} activates calcium-dependent potassium channels, hyperpolarizing the dendrite and pausing the cell.
- When Ca^{2+} concentration gets too high, the IP_3R receptor closes again.

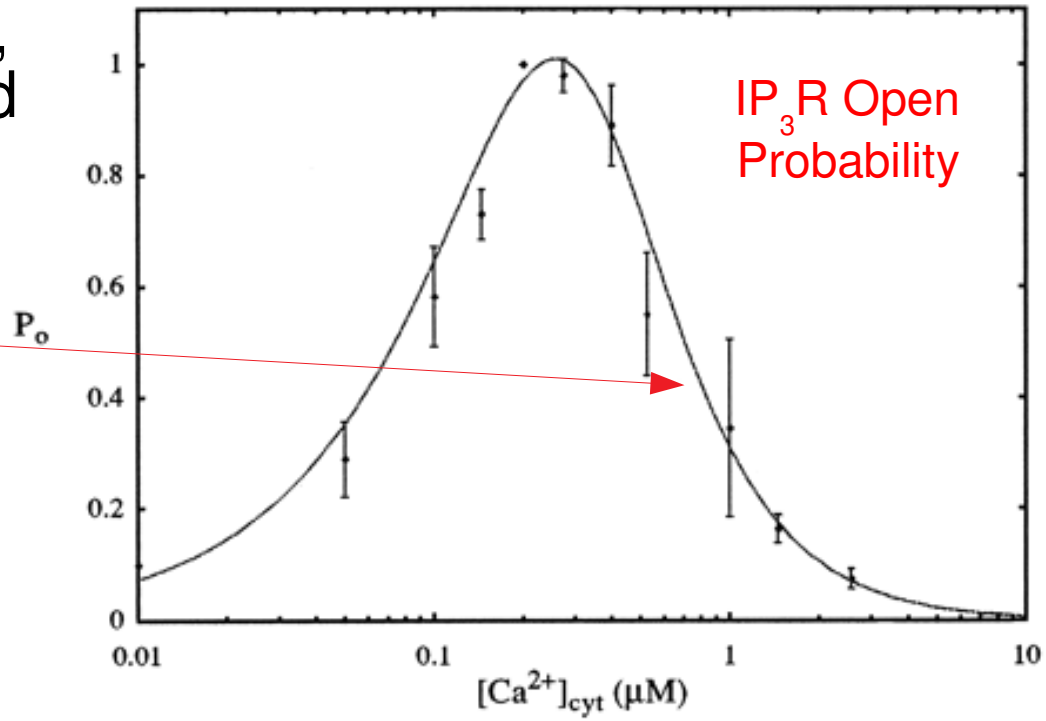

Spectral Timing

- Calcium level in the dendrite builds slowly as IP_3 accumulates.
- Positive feedback on IP_3 production and IP_3R channel opening results in a rapid rise in calcium level.
- But when Ca^{2+} level high enough, IP_3R channels close again.
- The speed at which this happens depends on the number of *mGluR1* receptors in the synapse.
- Different concentrations of *mGluR1* receptors produce different timing characteristics.
- High calcium level hyperpolarizes the dendrite through calcium-dependent potassium channels and inhibits firing.

Spectral Timing: Calcium Concentration Profiles

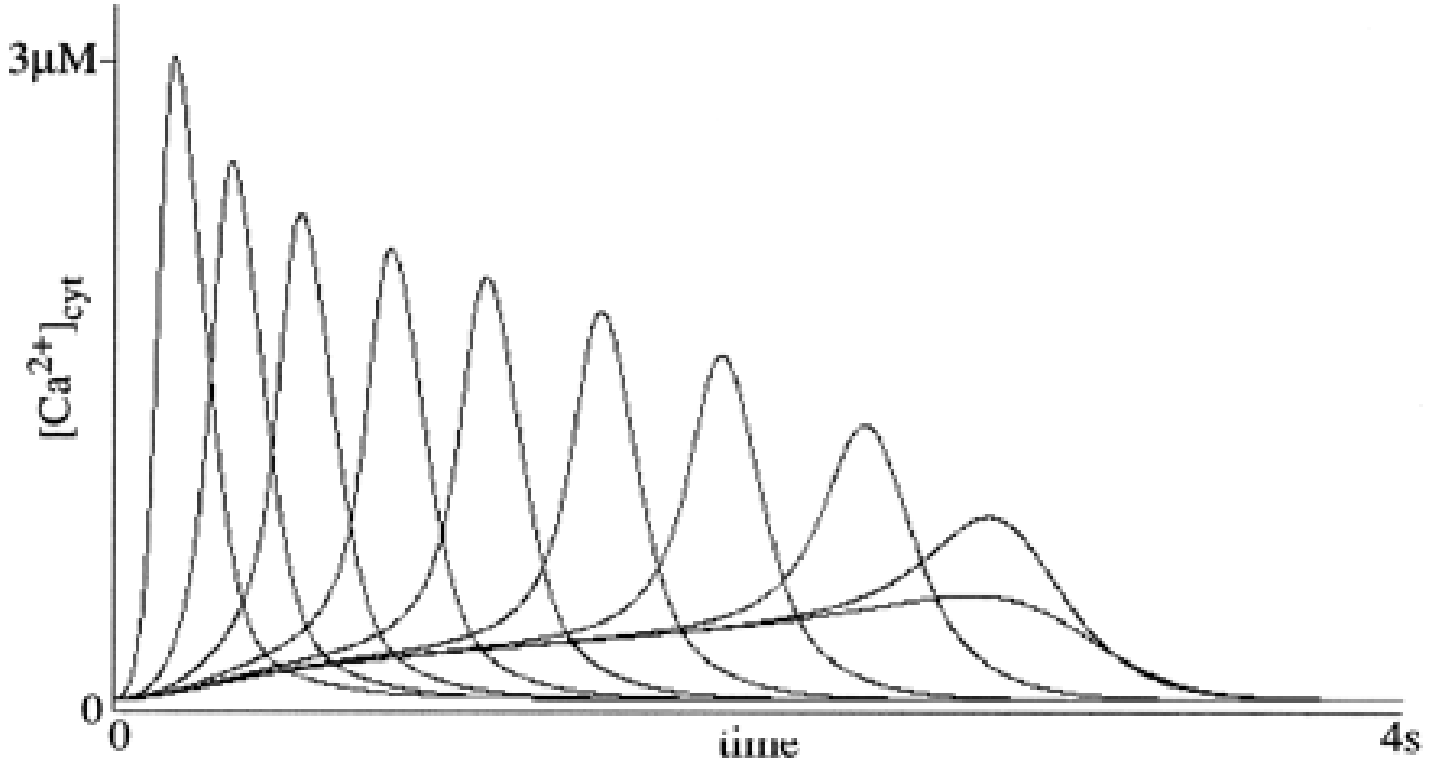

Fiala et al. simulation: responses to 50 msec glutamate application produced by varing B_{max} parameter.

Learning Performance of the Model Using a Population of Purkinje Cells

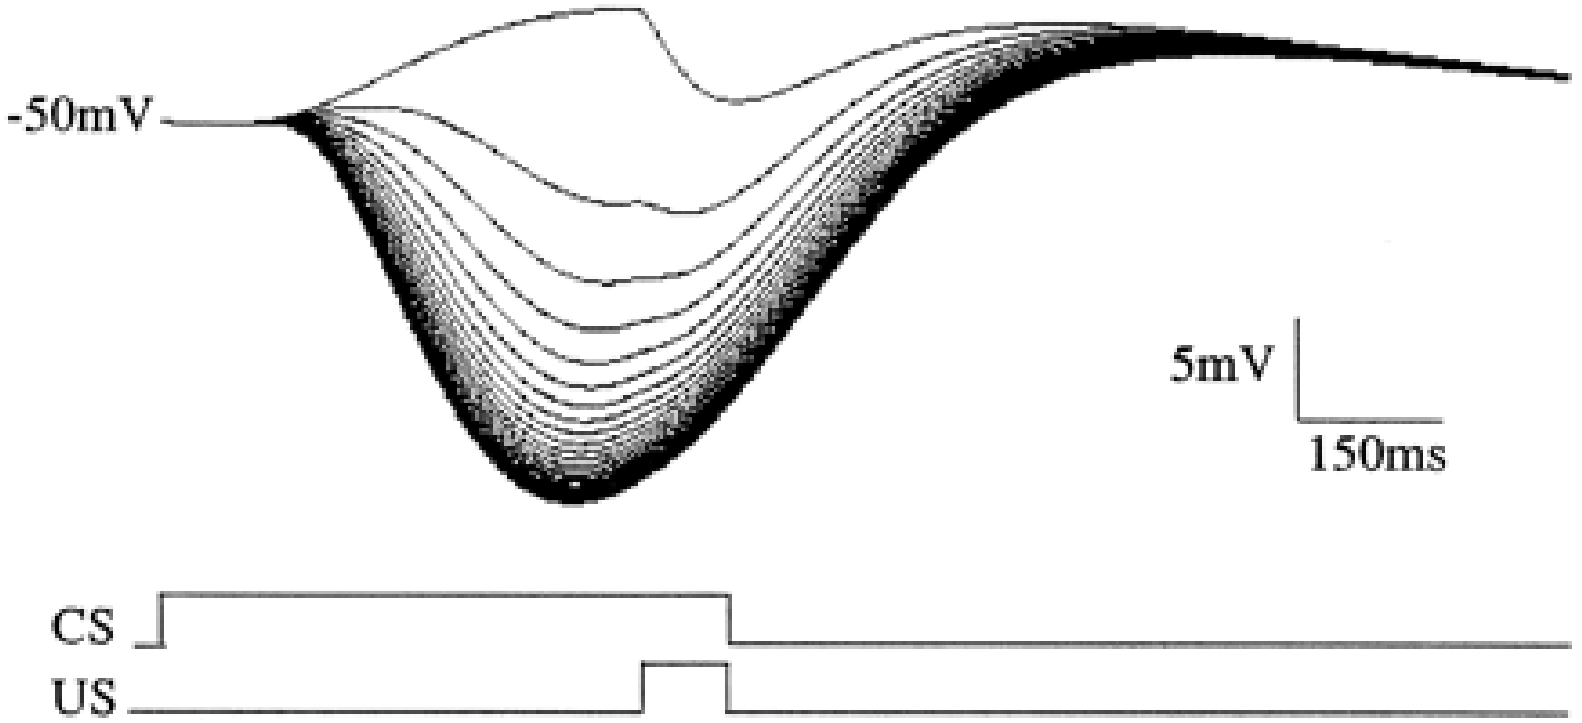

30 trials; ISI = 500 msec

Learning in Purkinje Cell Dendrites

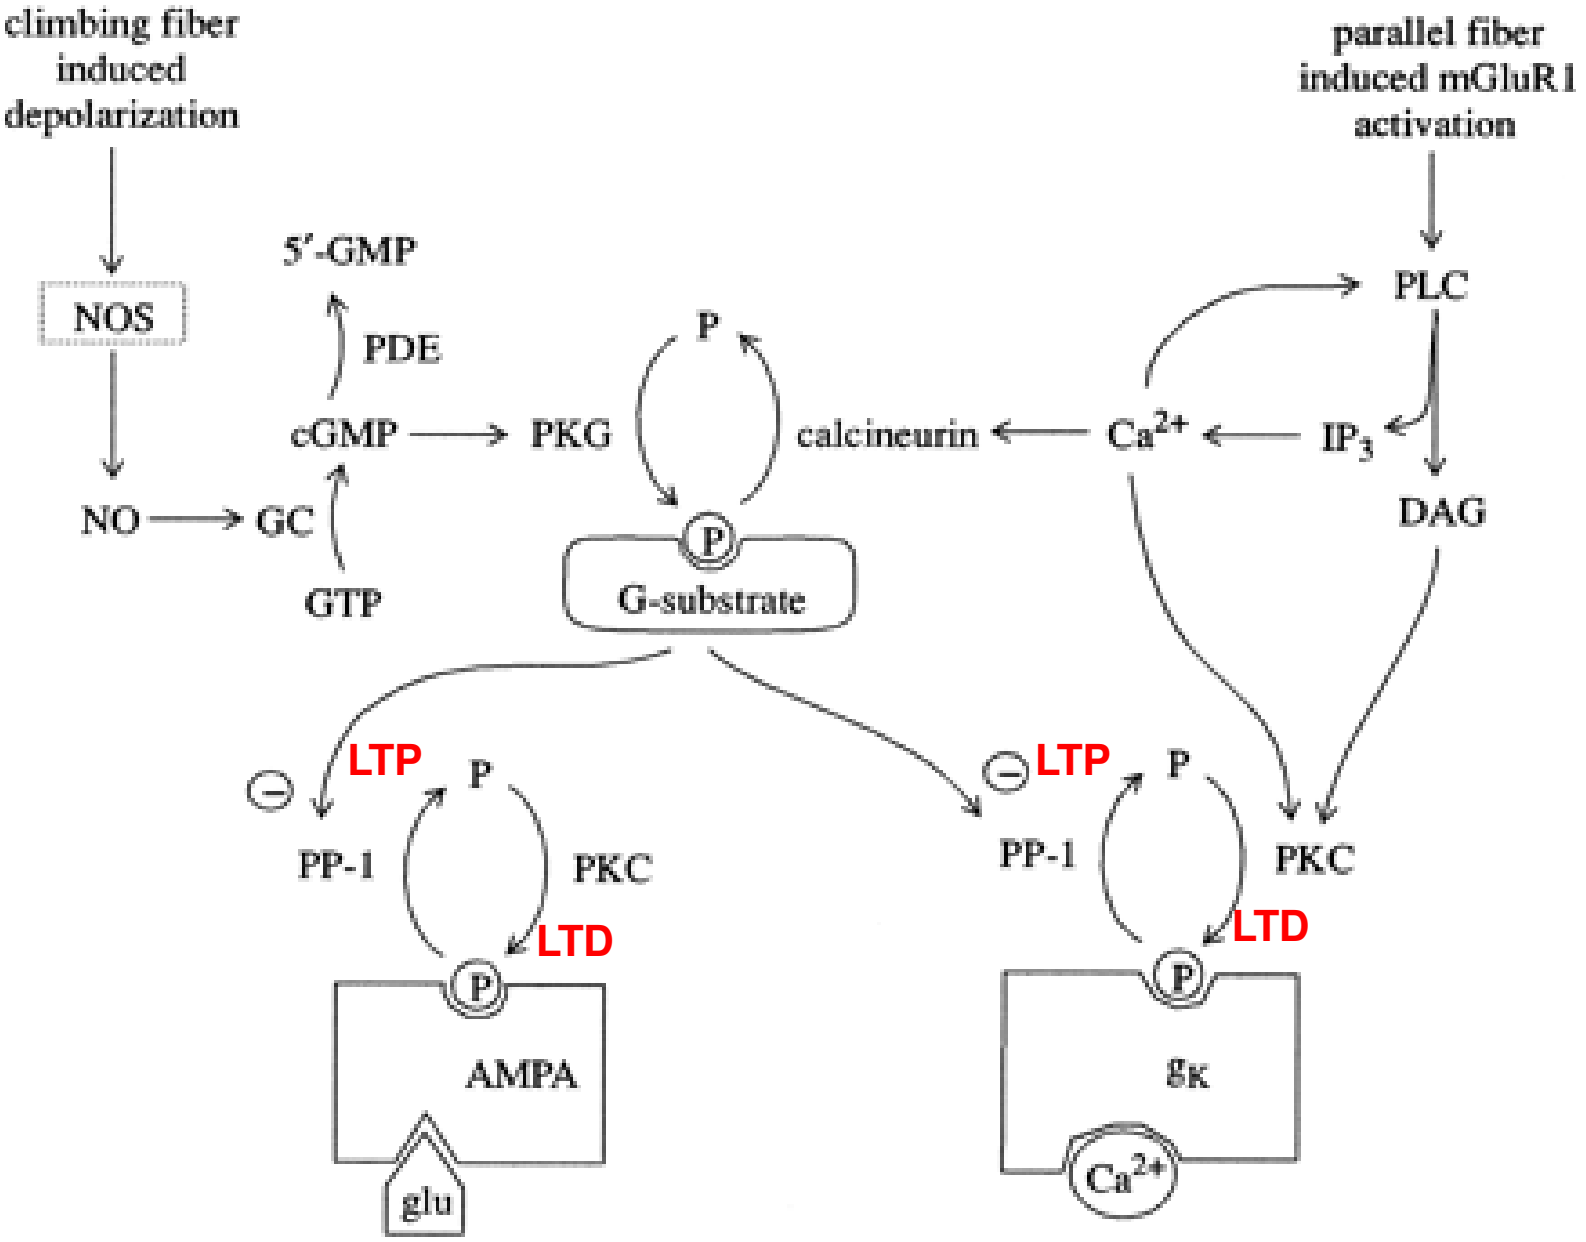

Problems with Spectral Timing Models

- Fiala et al. assume that each Purkinje cell (or each dendrite) has a fixed number of mGluRs, giving a fixed latency value.
  - But Jirenhed & Hesslow (2011) show that any Purkinje cell can learn any CS-US interval.
- Alternative model by Steuber and Willshaw (2004) assumes that learning modulates the number of mGluRs. This predicts that CR latency should decrease as learning proceeds.
  - But Jirenhed et al. (2007) found that while CR magnitude increases with learning, CR latency remained constant.
  - Changing the CS-US interval should cause a gradual shift in latency, but experiments show simultaneous extinction and acquisition.
  - Model can't account for double peak CRs seen in animals.

Summary

- Two sites of cerebellar learning for eyeblink conditioning:
  - Cells in interpositus nucleus learn to respond to tone CS
  - Purkinje cells in cerebellar cortex learn timing of the response
- Purkinje cells require both LTP and LTD to produce temporally accurate responses.
- Granule cells have diverse response profiles
- Multiple hypotheses about how the cerebellum keeps time: delay lines, spectral timing, oscillators, liquid state machines
- Two hypotheses for why learning fails at long ISIs:
  - Medina et al: long period of LTP overwhelms LTD
  - Medina & Mauk recurrent network (= LSM) model: granule cell activity sequence gradually diverges

Are All These Models Wrong?

- Hesslow et al. (2013) find problems with all existing models:
  - Purkinje cells have an intrinsic spiking mechanism that does not depend on parallel fiber input, so LTD of the pf→Pk synapse should not be sufficient to silence the cell.
  - The time course of LTD does not agree with that of eyeblink conditioning. (But in vitro slice experiments aren't a direct match for behavioral experiments.)
  - Granule cells may not have the rich variety of temporal responses these models assume.
  - A single Purkinje cell can learn a range of CS-US timings, so spectral timing models that assign a specific delay to each Purkinje cell cannot be correct.
  - Models that learn by adapting a cell's delay value cannot account for dual-peak responses, or for the fact that changing the ISI after training simultaneously extinguishes the old CR latency and potentiates a new one; it does not gradually shift the latency.

Hesslow et al.'s Proposal (2013)

- Each Purkinje cell has a family of “timer units” with different latencies.
- Learning CR timing is done by selecting the units with the correct latency value.
- Once a timer is activated (by parallel fiber input), it runs autonomously and triggers hyperpolarization with its characteristic latency.
- Double-peak responses are explained by having more than one set of timer units selected.

Lots of open questions:
- *What is the neurophysiological basis of timer units?*
- *How do timer units become selected?*
- *How do timers become activated?*

A New Proposal

- Hesslow et al. (2013) theorize about a new mechanism:
  - Doesn't depend on parallel fiber input timing.
  - Mechanism is intrinsic to the Purkinje cell or interneurons.
  - After training, pf input activates a molecular mechanism with a particular constant time delay that turns on a hyperpolarizing response for a specific duration.
  - The delay is fixed, not adjustable.
  - There is a family of these “timer units”, and the learning process *selects* the appropriate timer to use.
  - Once a timer has been activated, it runs its course independent of further inputs, so extending the duration of the CS will not affect the CR.

Recent Advances

- Suvrathan et al. (2016) showed that Purkinje cells in the flocculus, involved in learning the VOR, have a preferred PF-CF interval of about 120 msec for LTD. But Purkinje cells in the vermis, which implement a variety of different behaviors, have preferred interval varying from 0 to 150 msec.
  - Conclusion: different regions of cerebellum have different timing characteristics based on the behavior being controlled.

- Boele et al. (2018) show that two mechanisms contribute to learned Purkinje cell responses: (1) LTD of PF-to-PC synapses, and (2) inhibition from molecular layer interneurons (stellate and basket cells).
  - Both mechanisms must be knocked out by genetic manipulation in order to severely impair eyeblink conditioning.