Layered Architectures for Distributed Control

Terrence Sejnowski

Salk Institute
University of California, San Diego

Levels of Investigation

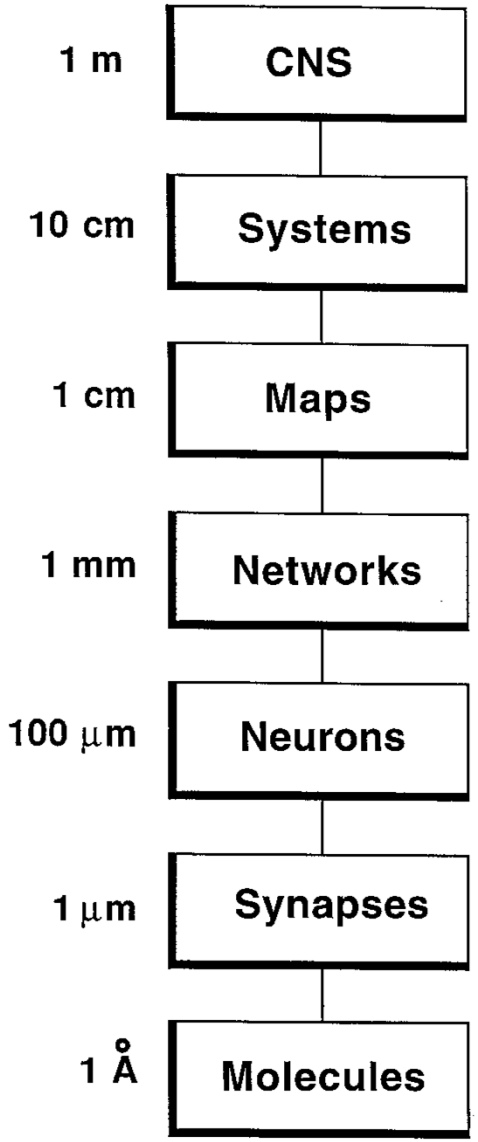

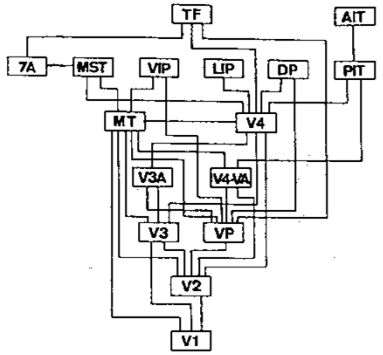

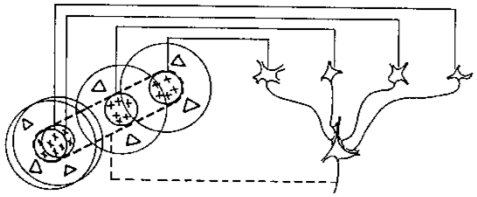

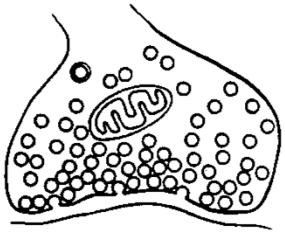

Marr's Levels of Analysis

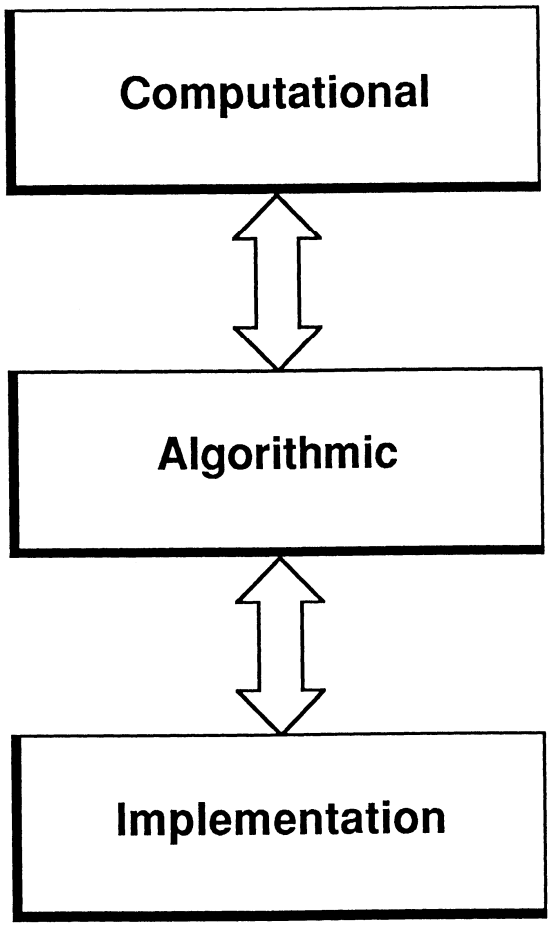

Churchland, Ramachandran and Sejnowski
A Critique of Pure Vision, 1994

Marr's Levels of Analysis

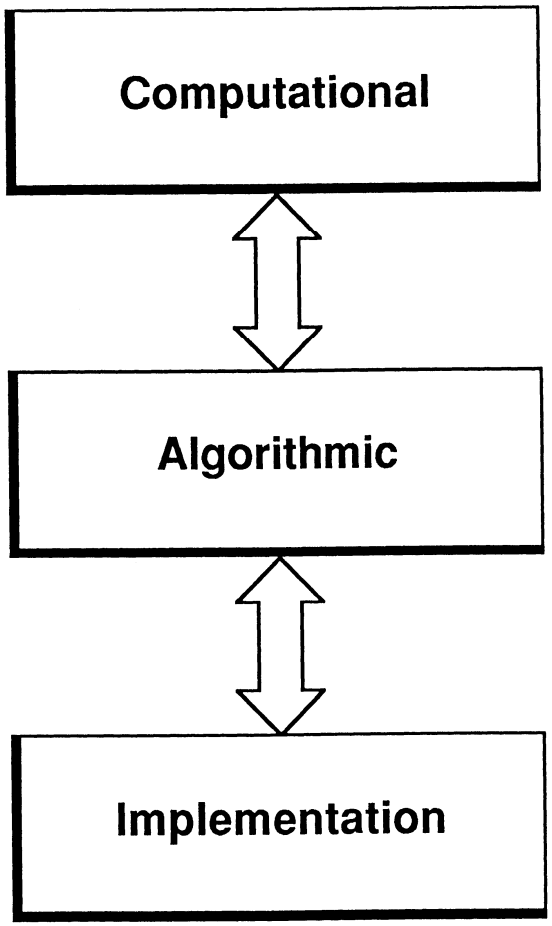


Time is Honey

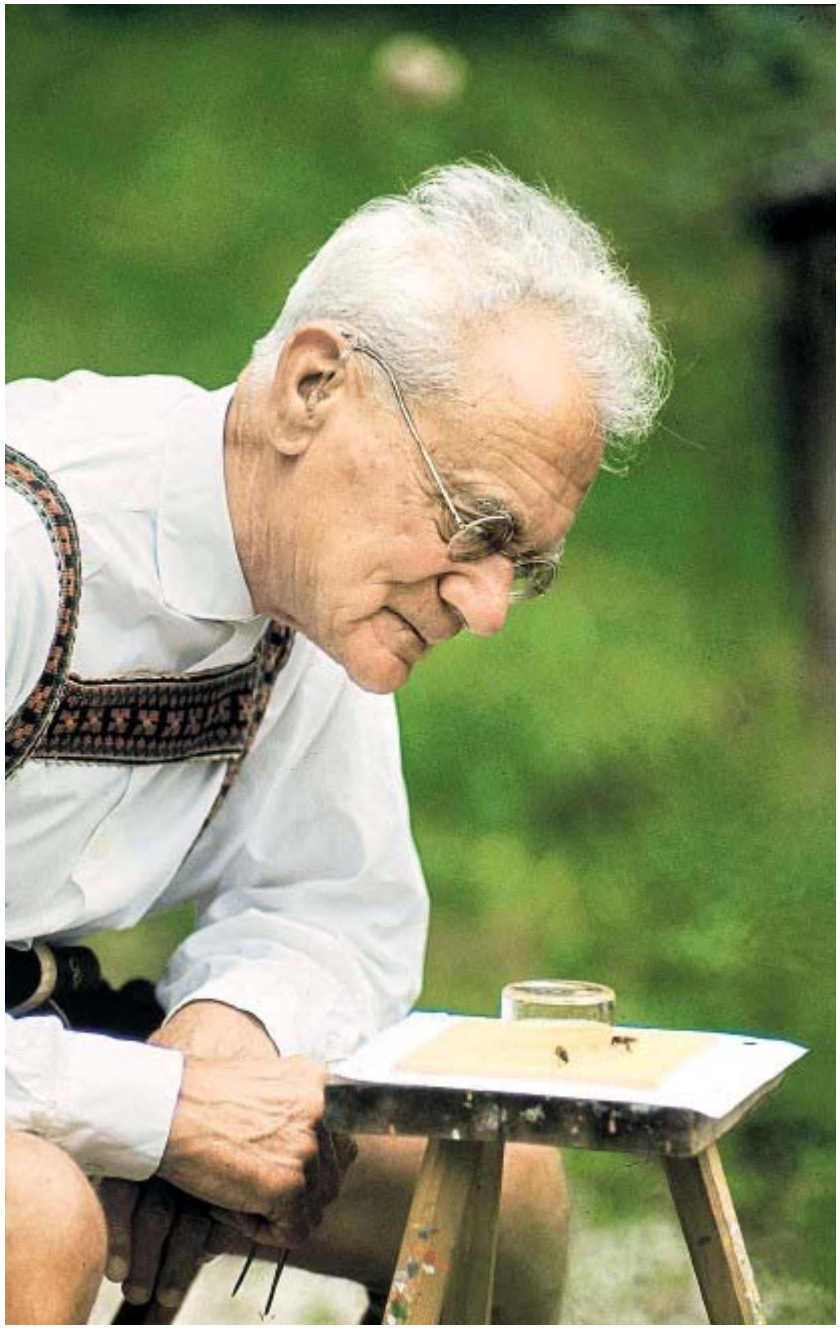

Karl von Frisch

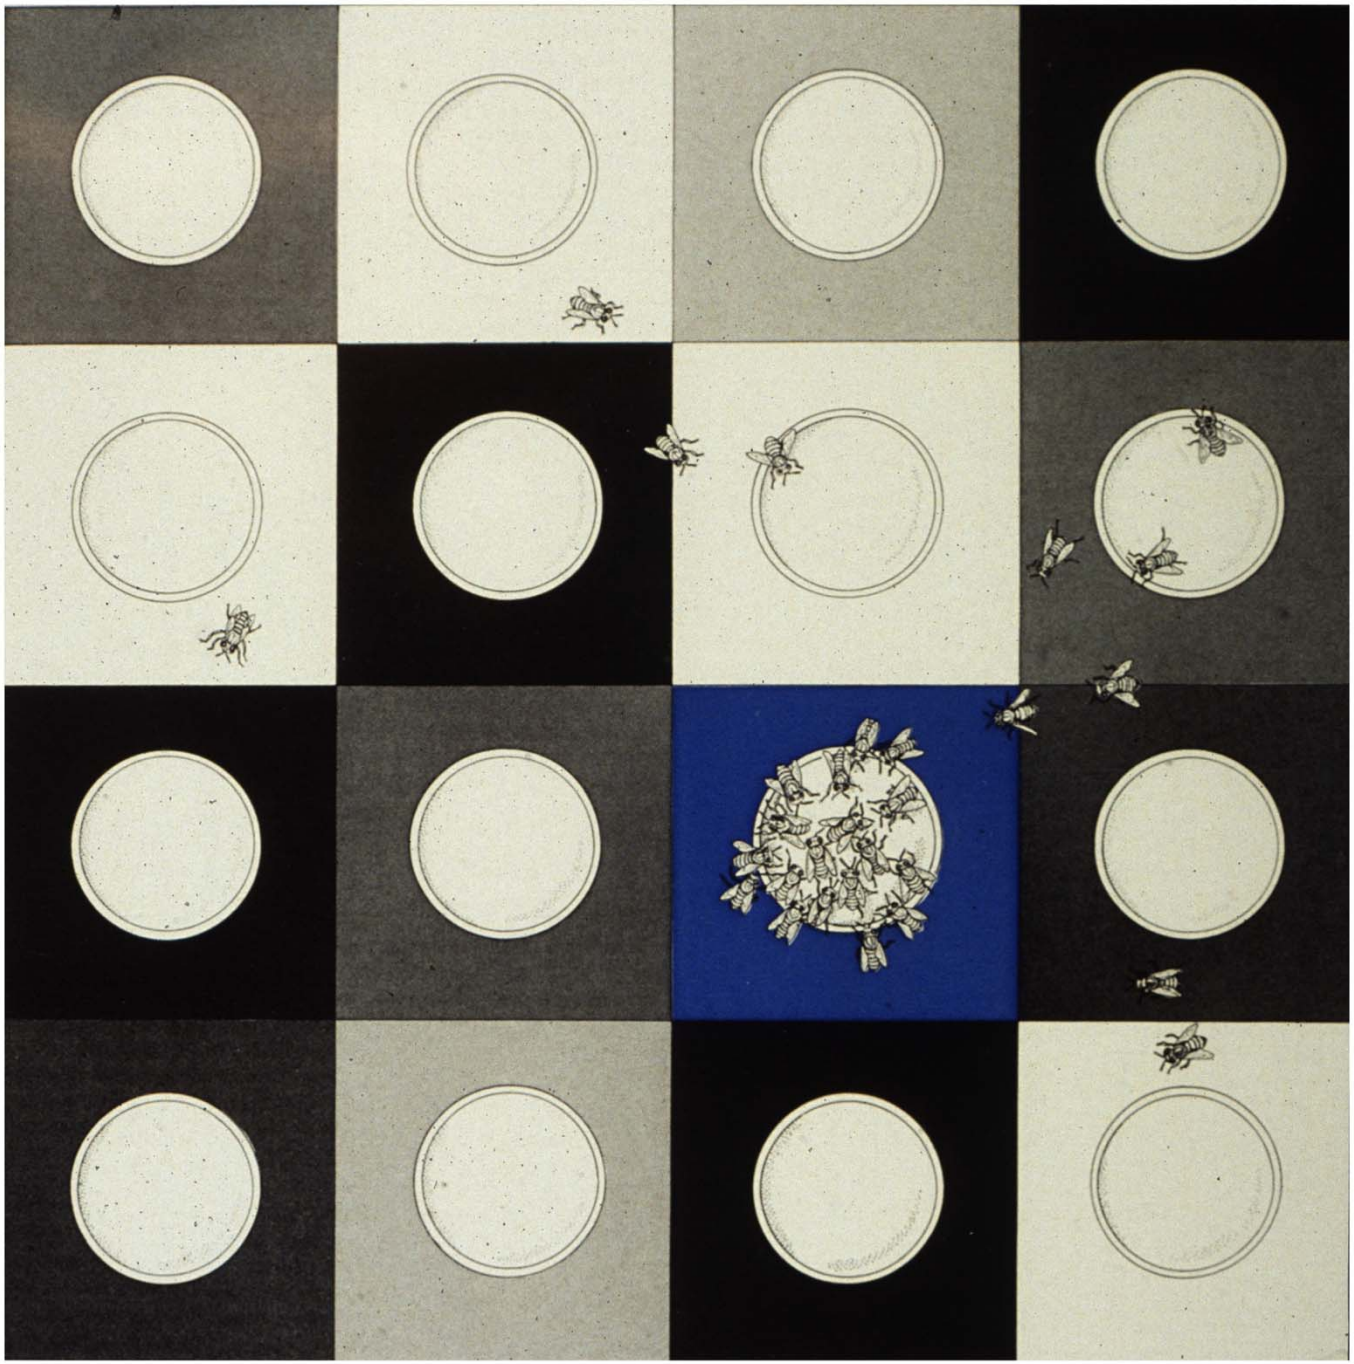

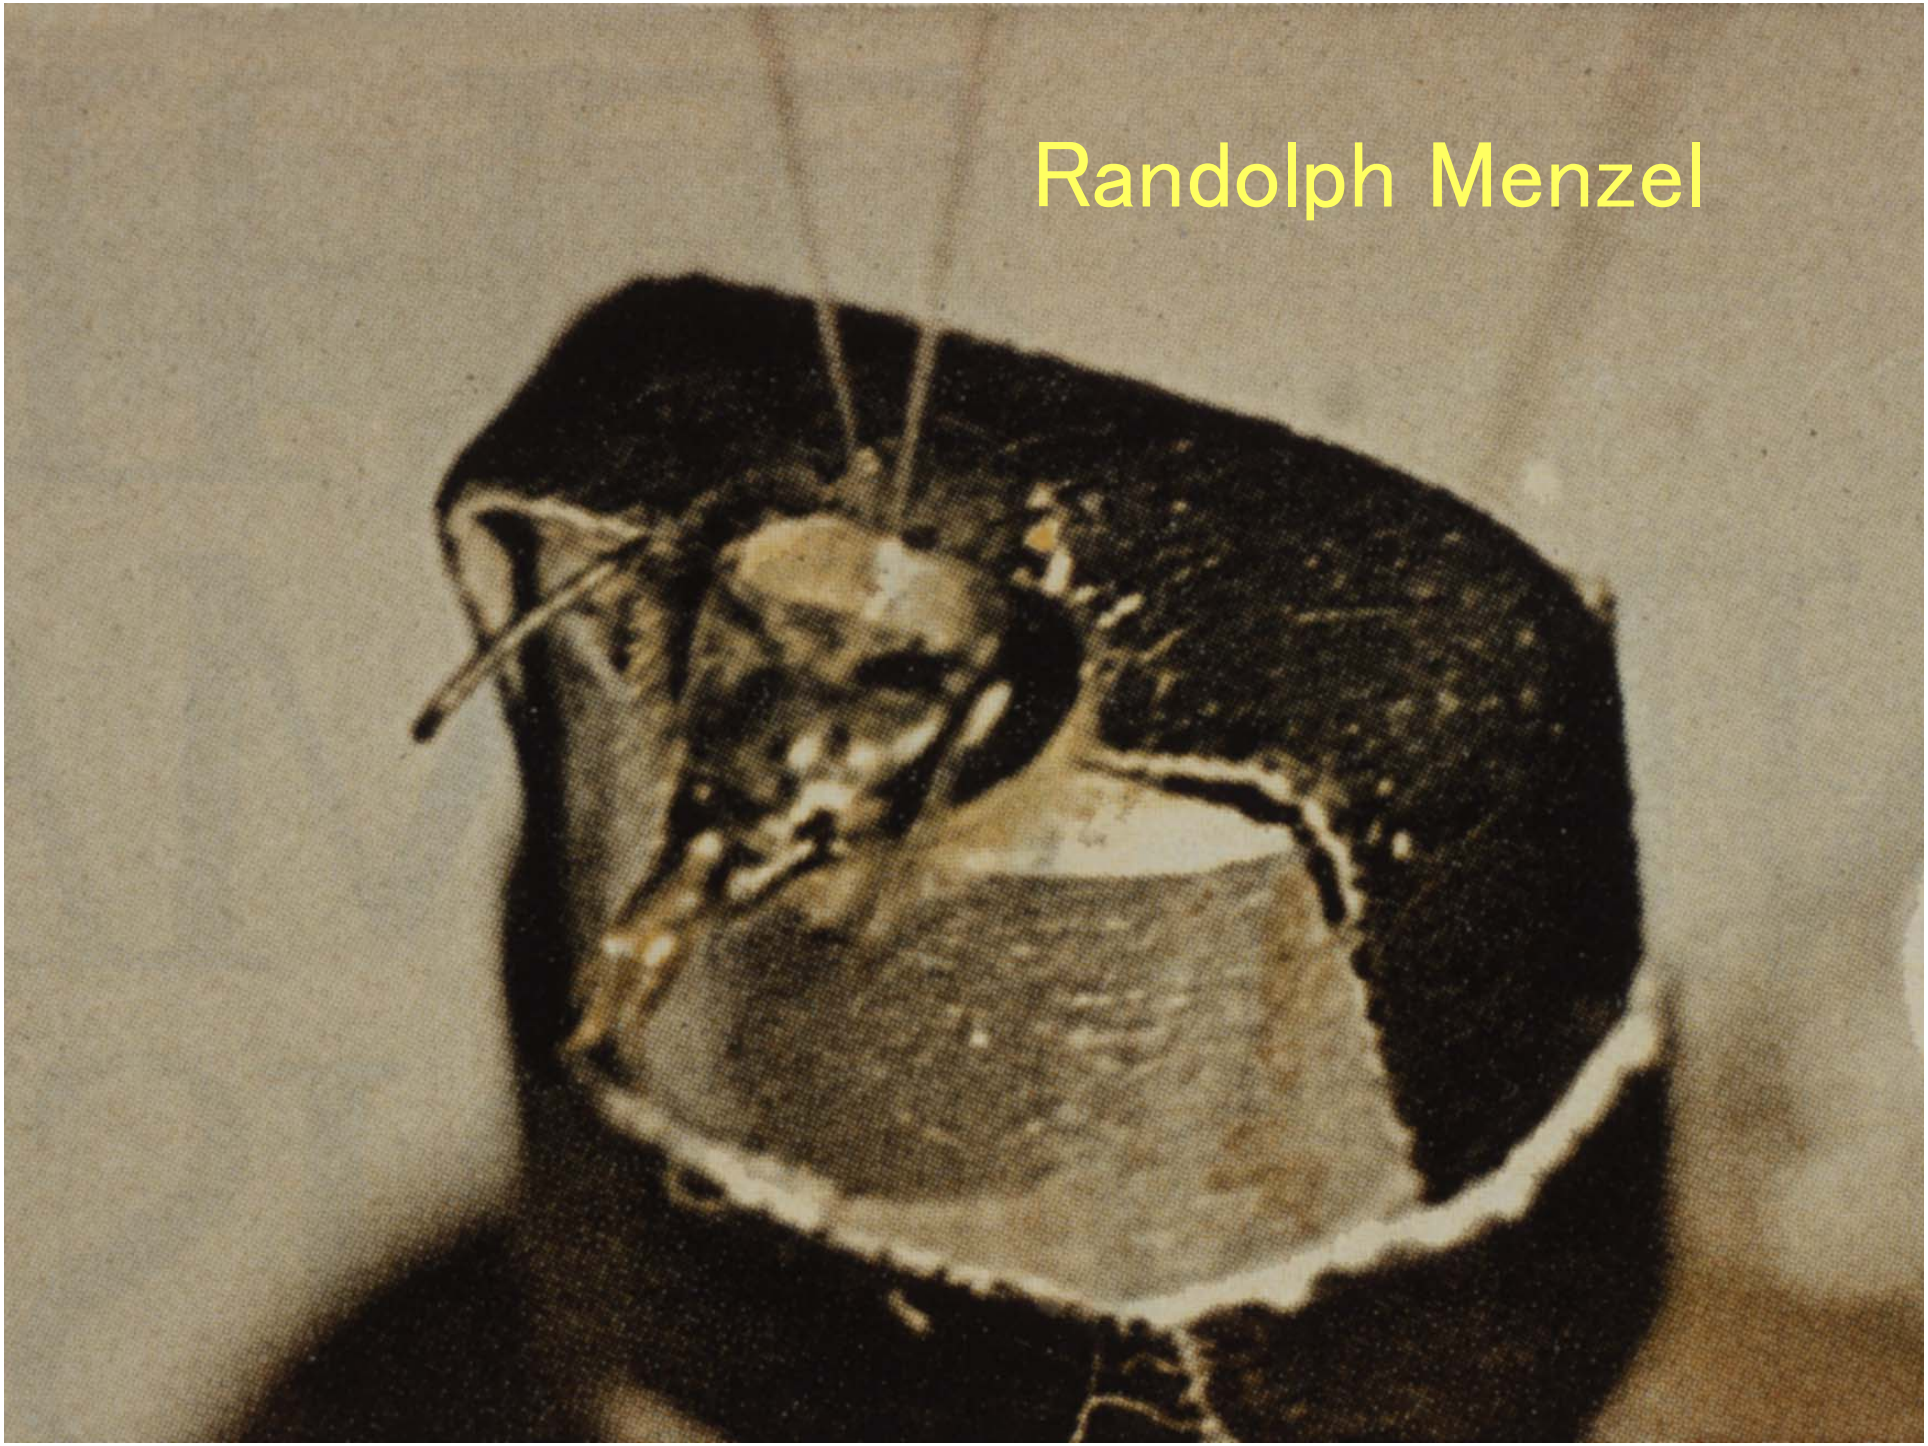
Randolph Menzel

VUMmx1 - Octopamine

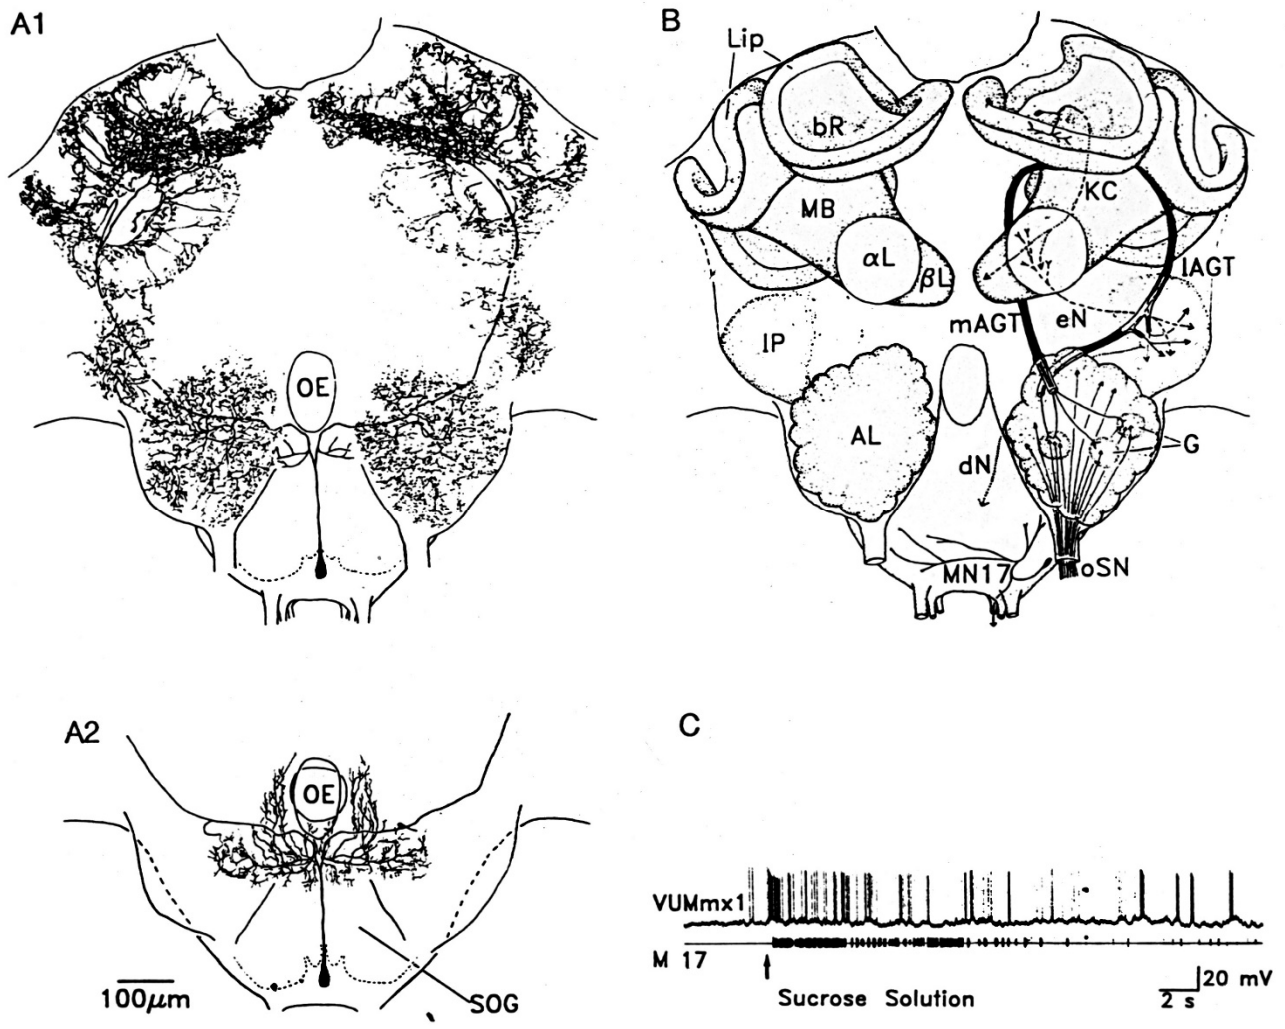

Hammer and Menzel, 1997

Temporal Difference Learning

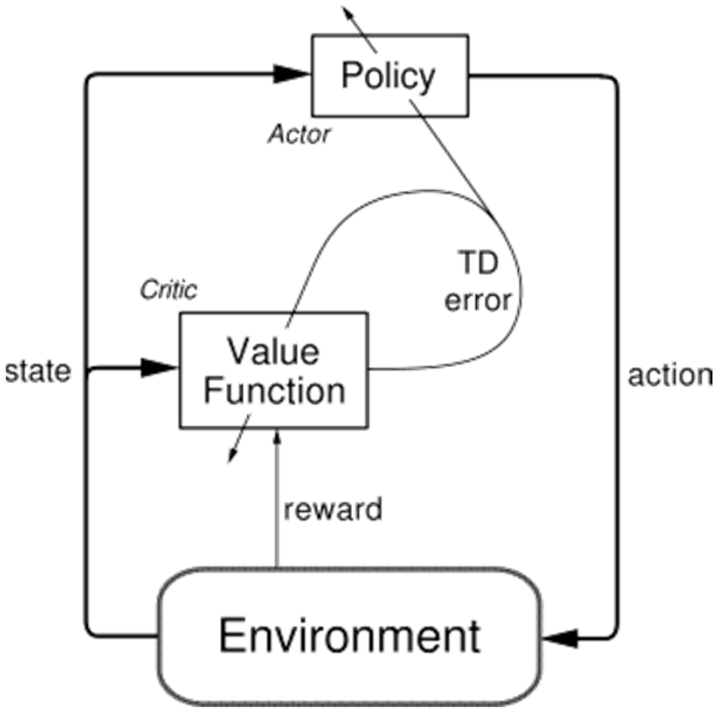

Sutton and Barto, 1988

TD - error :

$$\delta_t = r_{t+1} + \gamma V(s_{t+1}) - V(s_t)$$

Actions are determined by preferences :

$$\pi_t(s,a) = \Pr\{a_t = a \,|\, s_t = s\} = \frac{e^{p(s,a)}}{\sum_b e^{p(s,b)}},$$

Update the preferences :

$$p(s_t, a_t) \leftarrow p(s_t, a_t) + \beta\delta_t$$

The value function update :

$$V(s_t) \leftarrow V(s_t) + \alpha\delta_t$$

Temporal Difference Learning

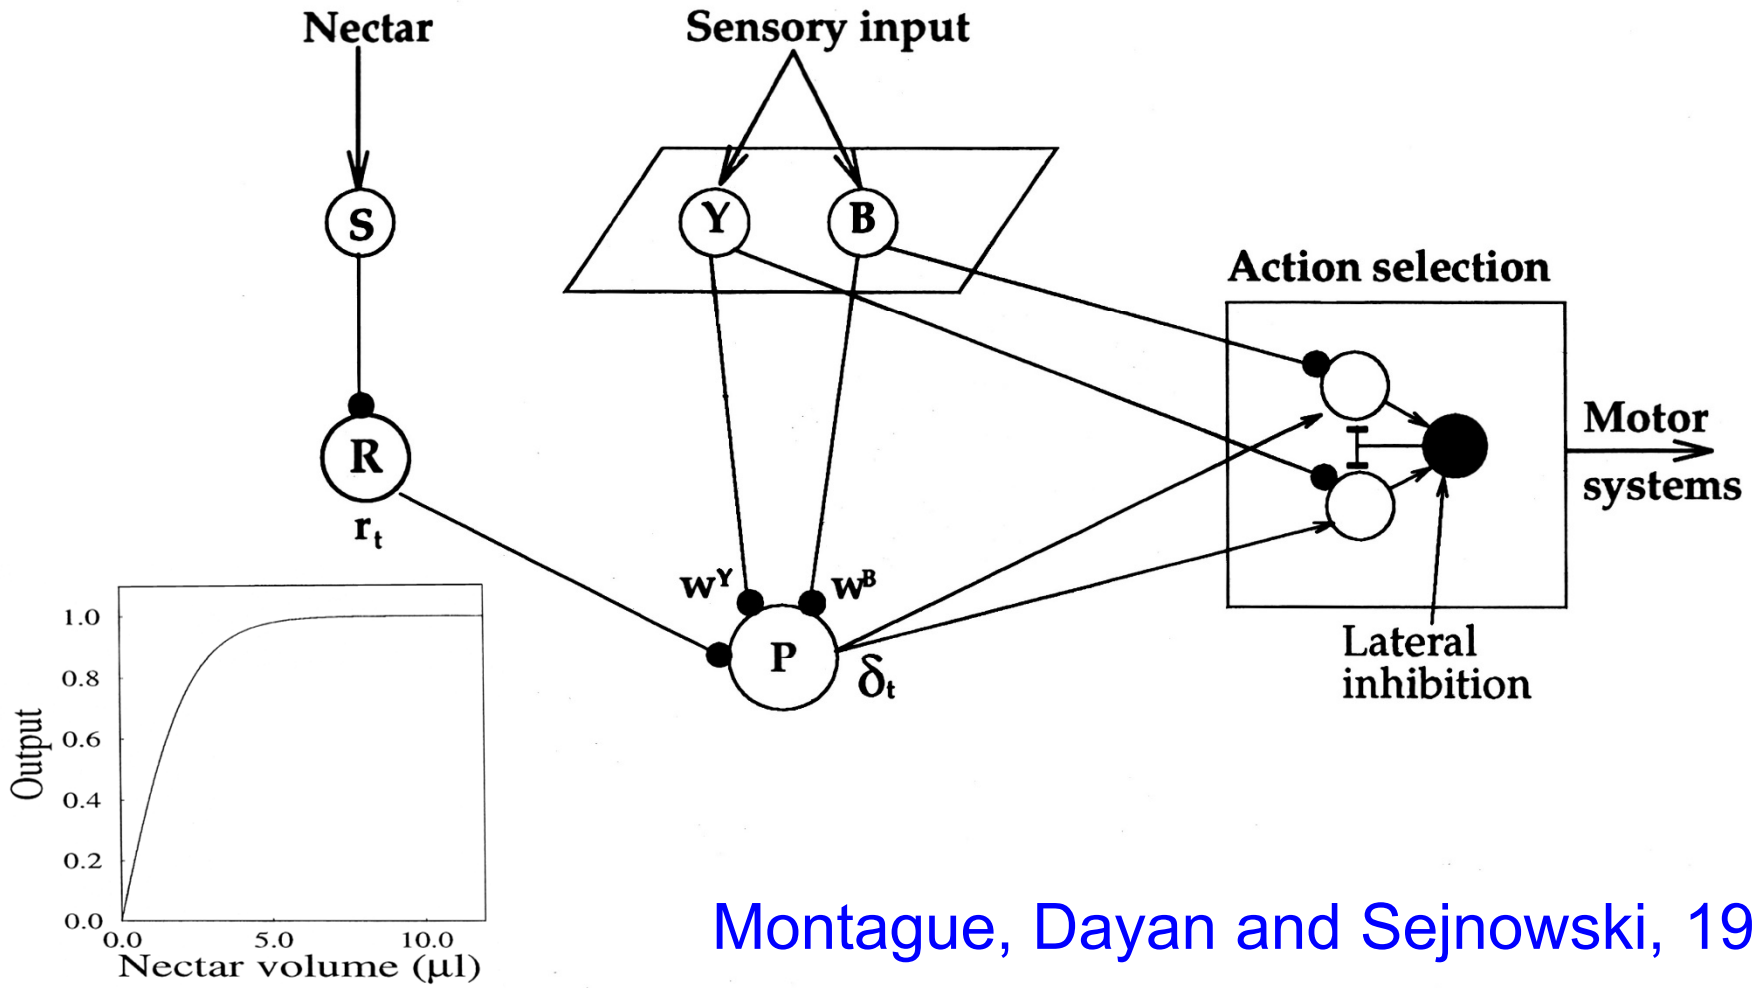

Montague, Dayan and Sejnowski, 1994

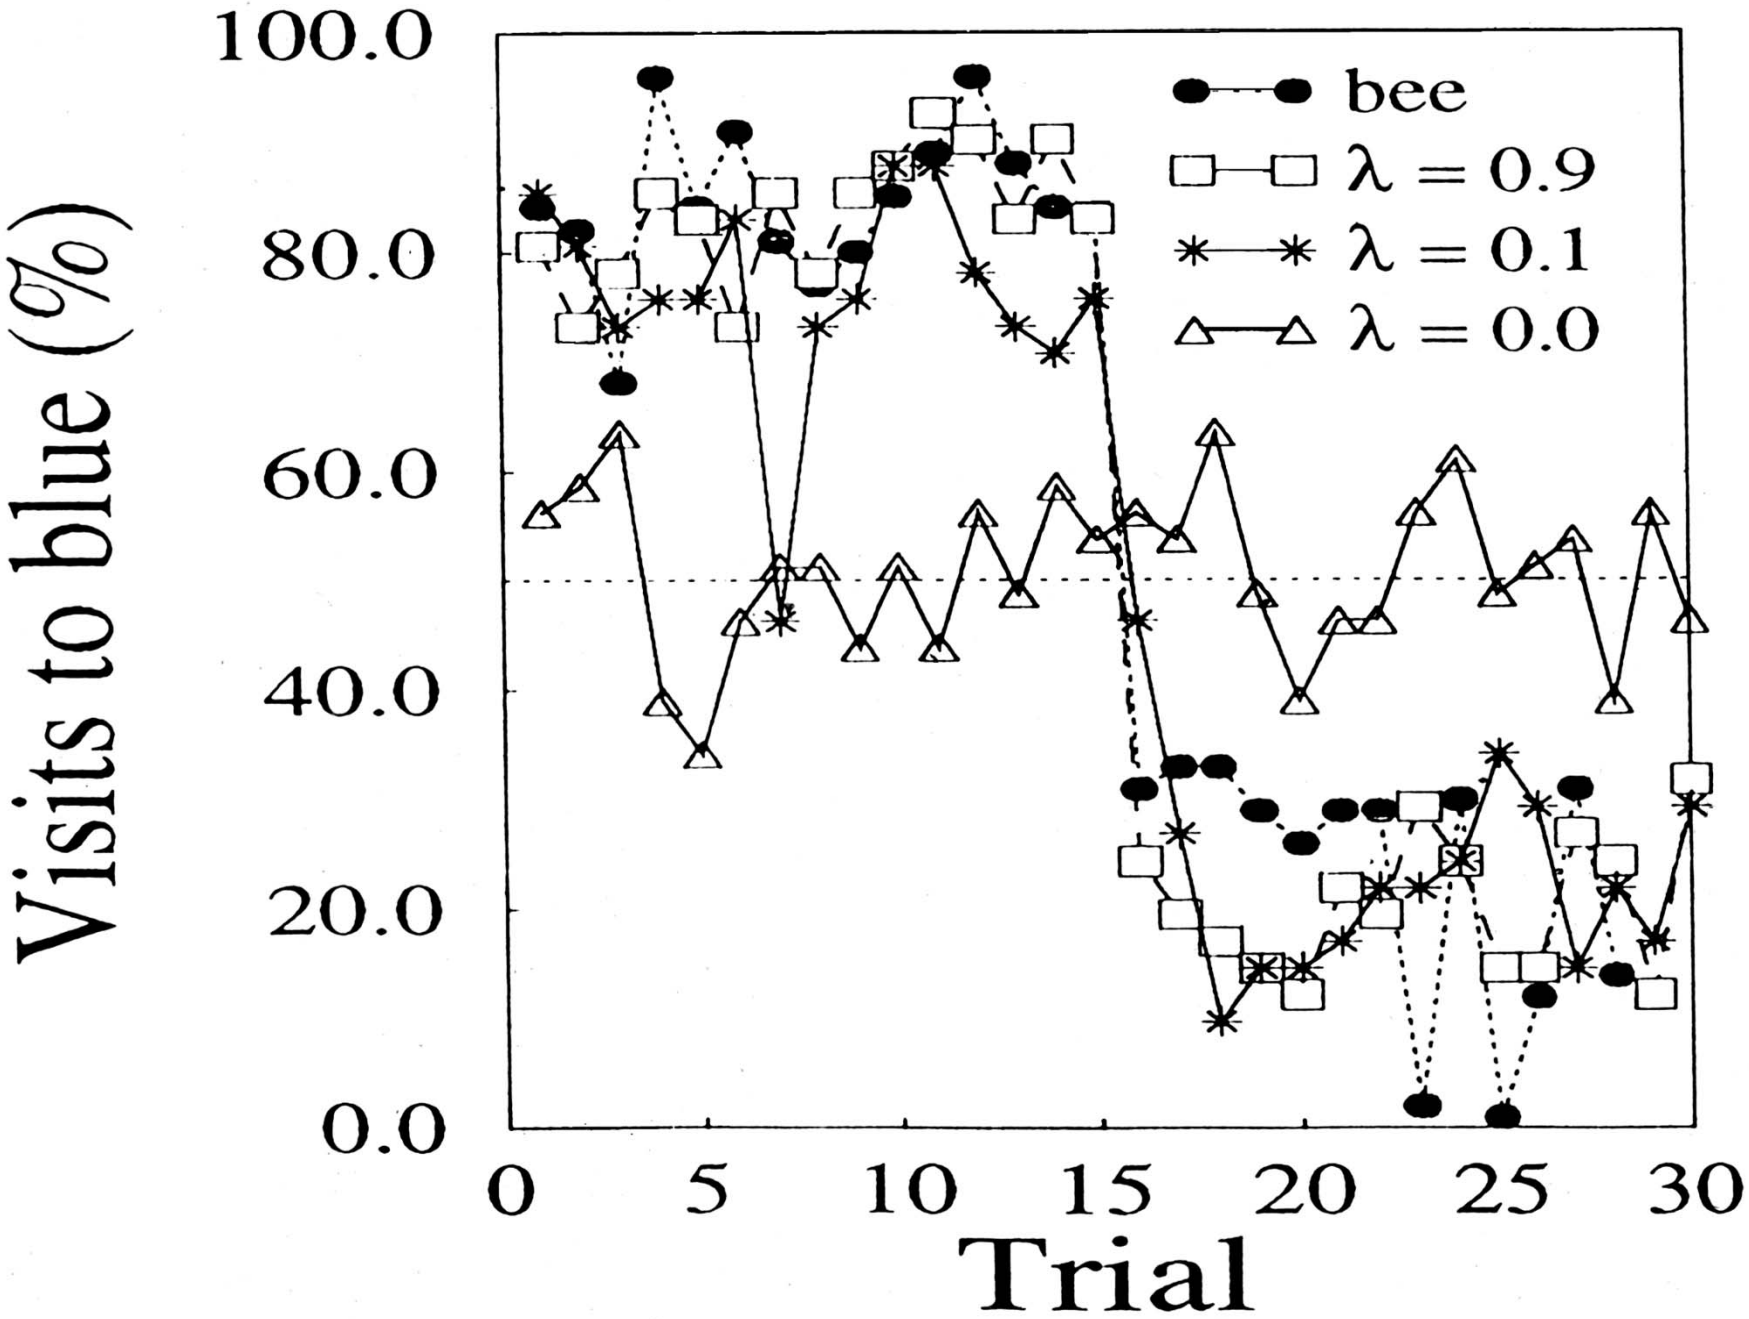
100.0
80.0
60.0
40.0
20.0
0.0
Visits to blue (%)
0
5
10
15
20
25
30
Trial
bee
λ = 0.9
λ = 0.1
λ = 0.0

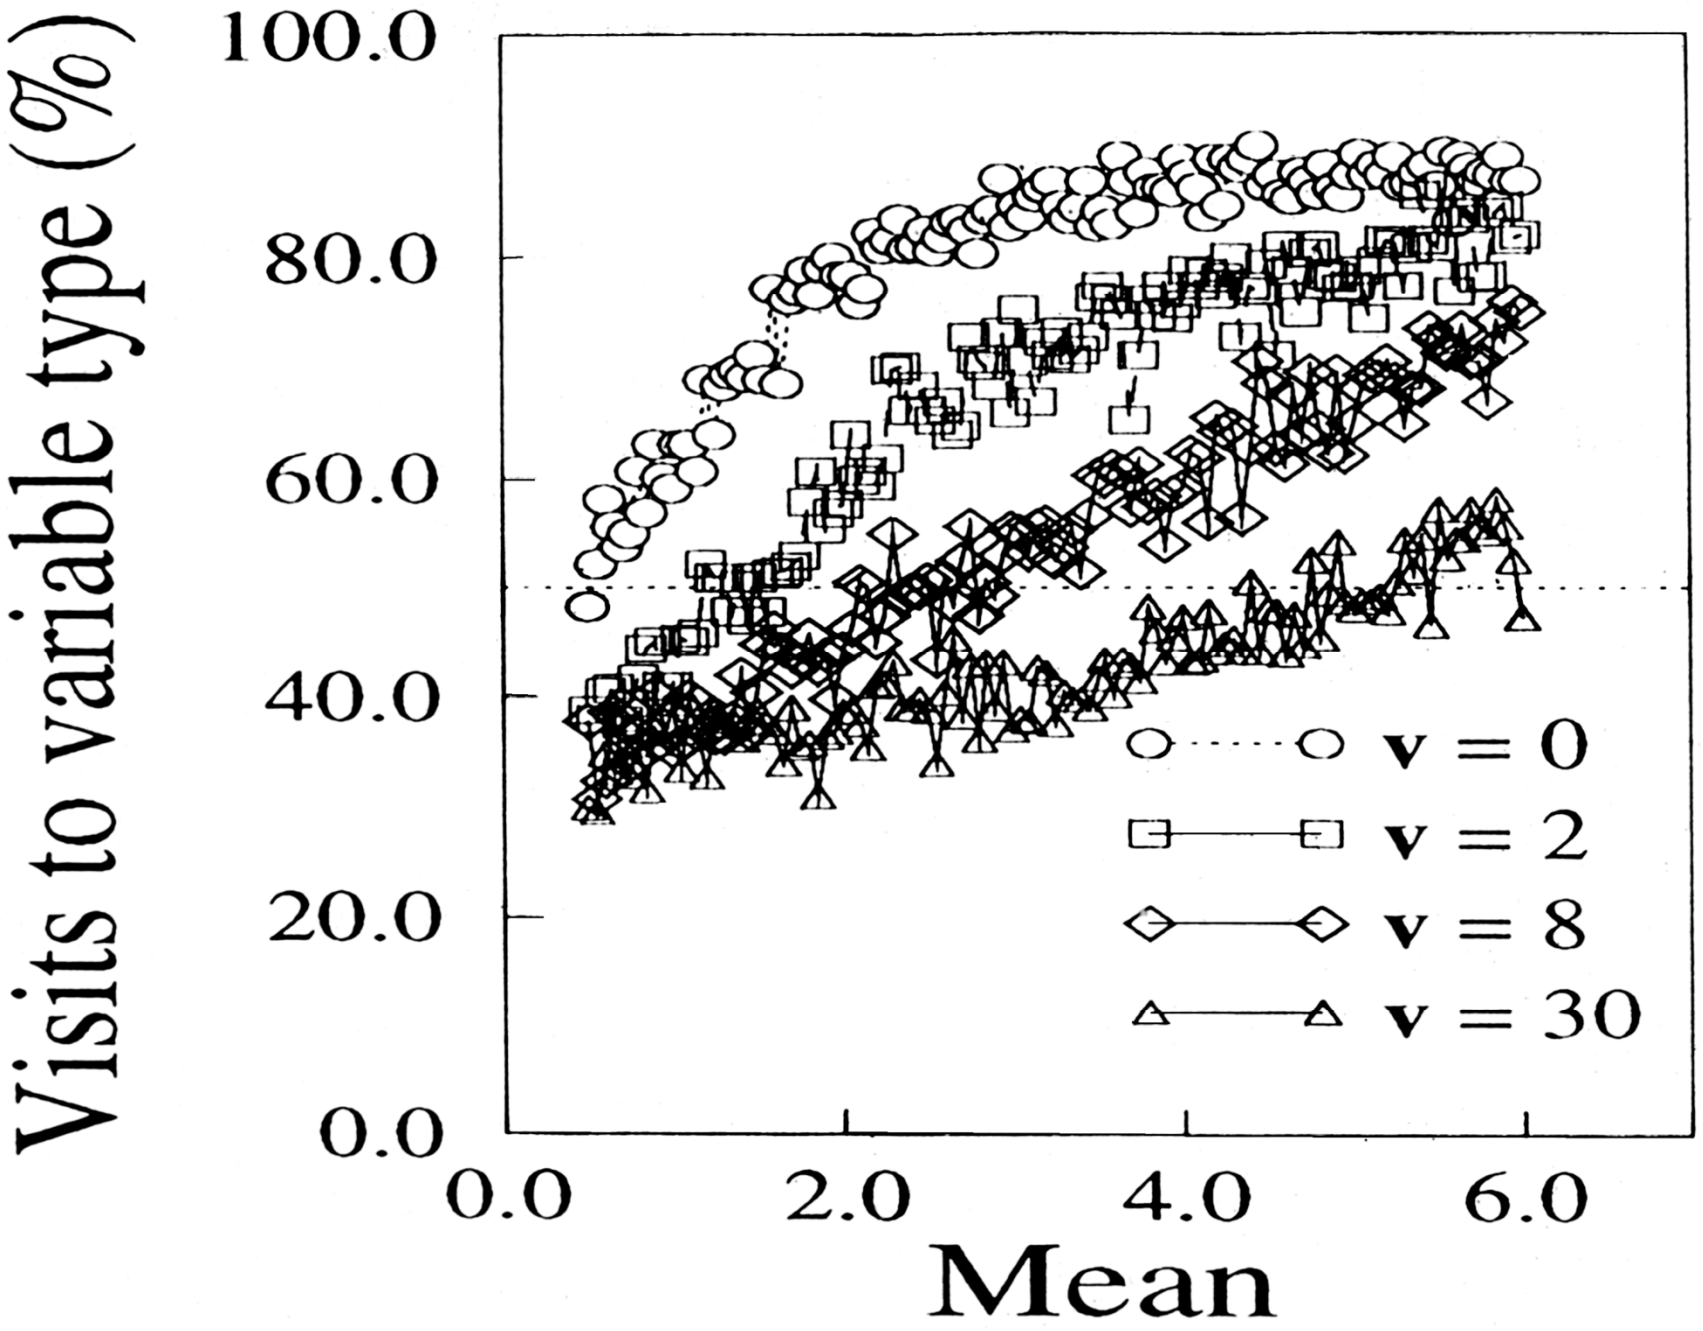
Visits to variable type (%)
100.0
80.0
60.0
40.0
20.0
0.0
0.0
2.0
4.0
6.0
Mean
v = 0
v = 2
v = 8
v = 30

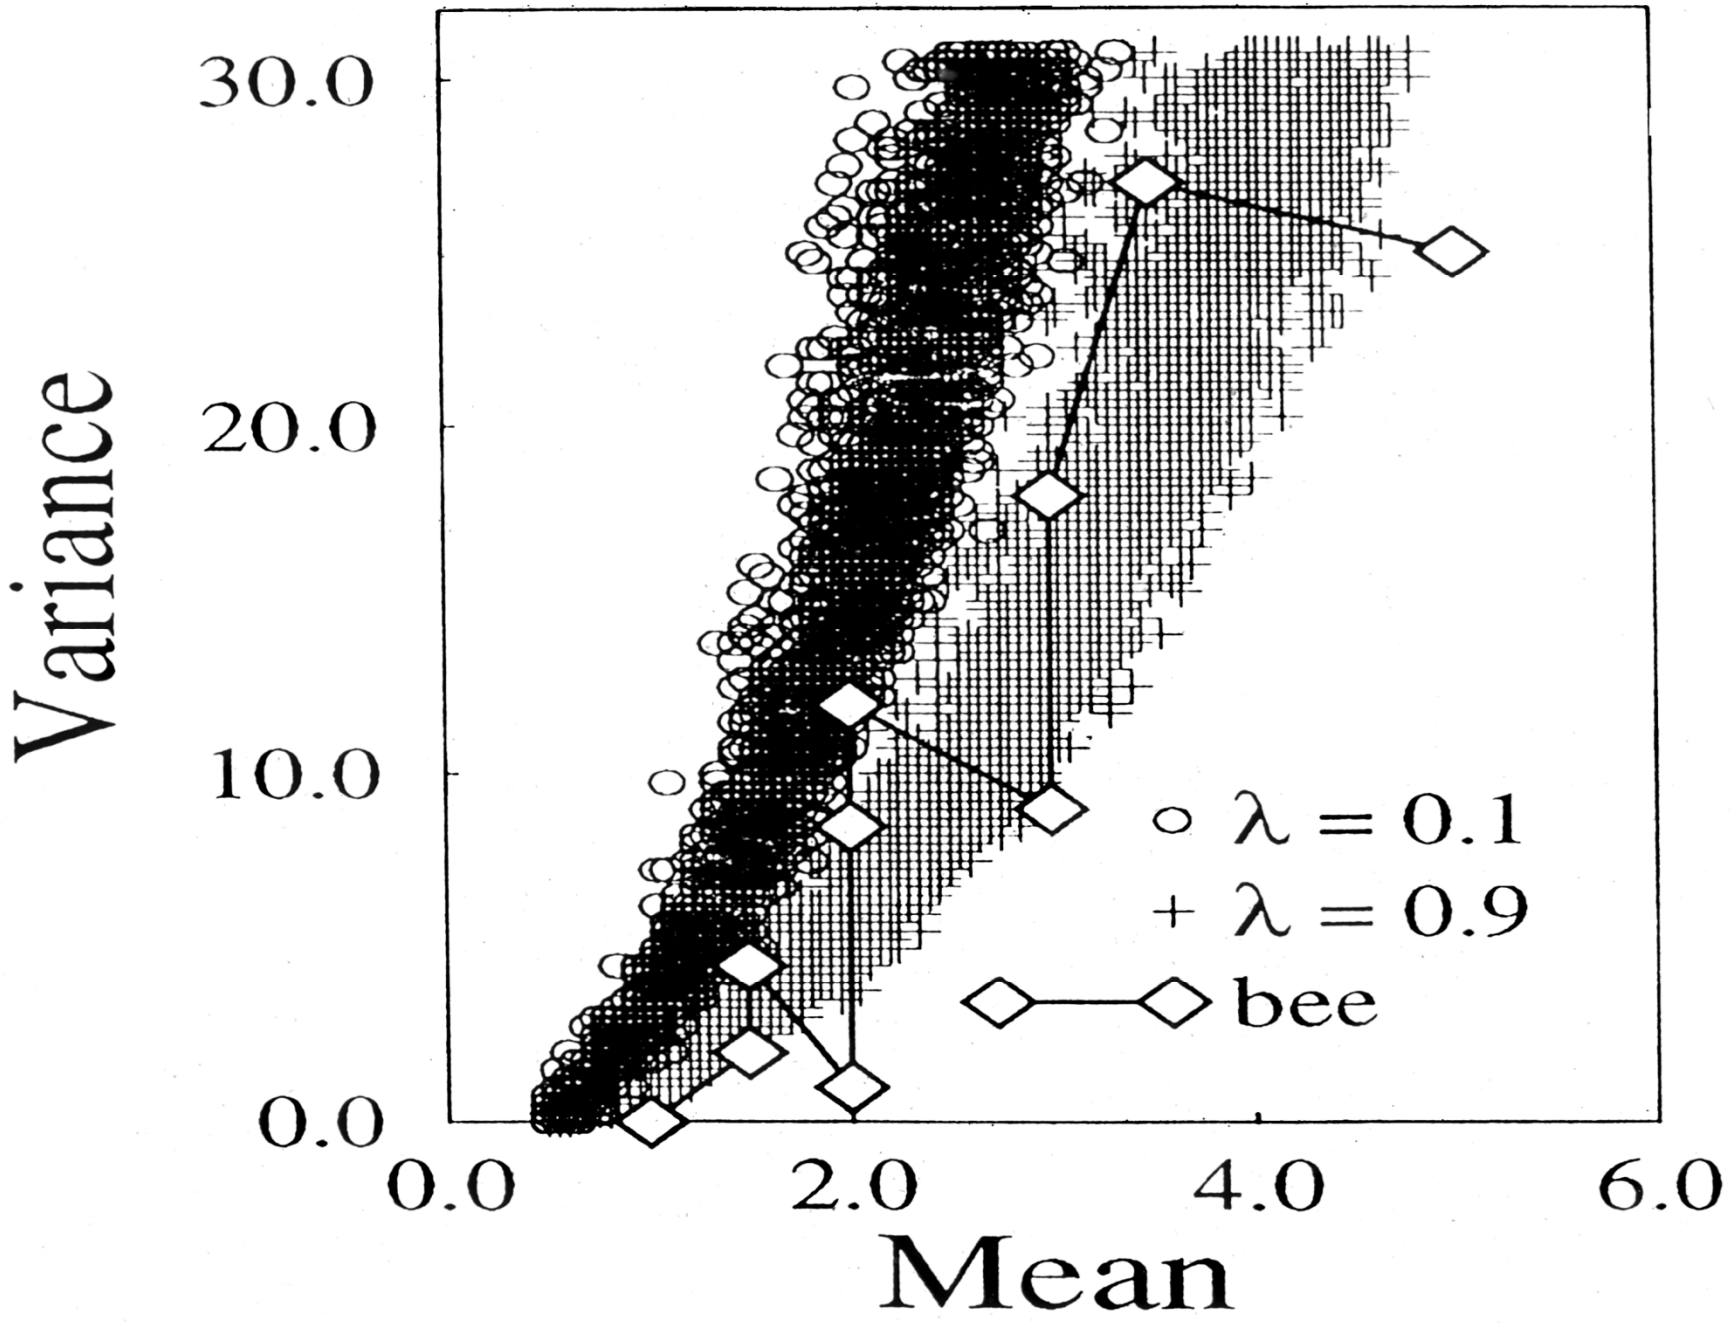
Variance
30.0
20.0
10.0
0.0
0.0
2.0
4.0
6.0
Mean
λ = 0.1
λ = 0.9
bee

Decisions

Decisions

Reward Modulates Visual Responses

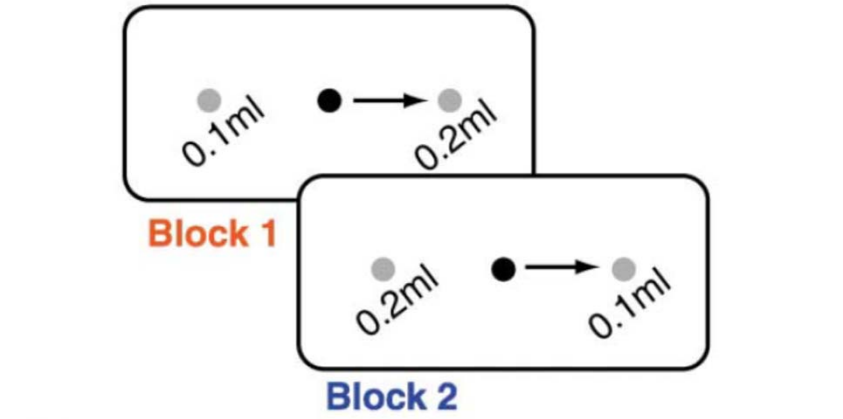

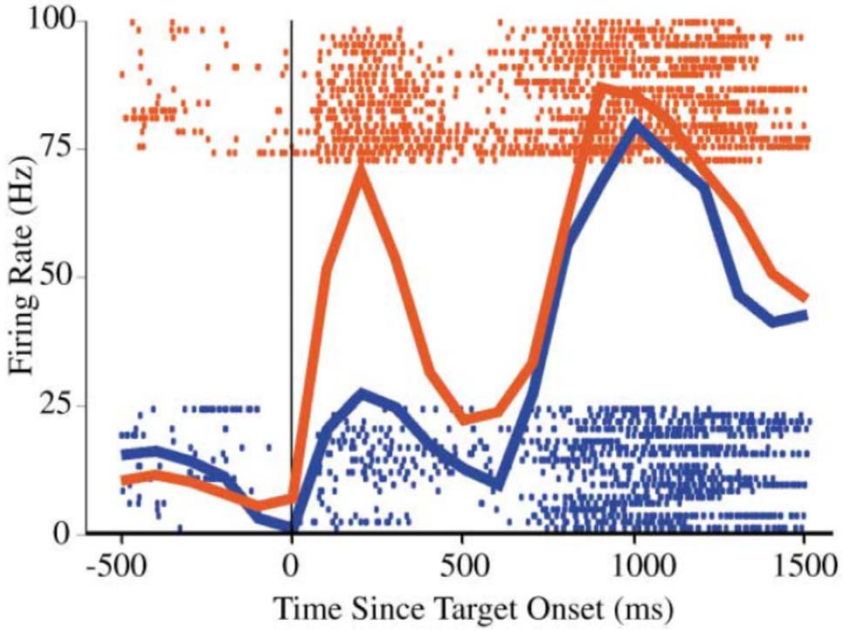

Glimcher, 2002

Actor Critic Model

Dopamine Neurons

Cerebral Cortex

Basal Ganglia

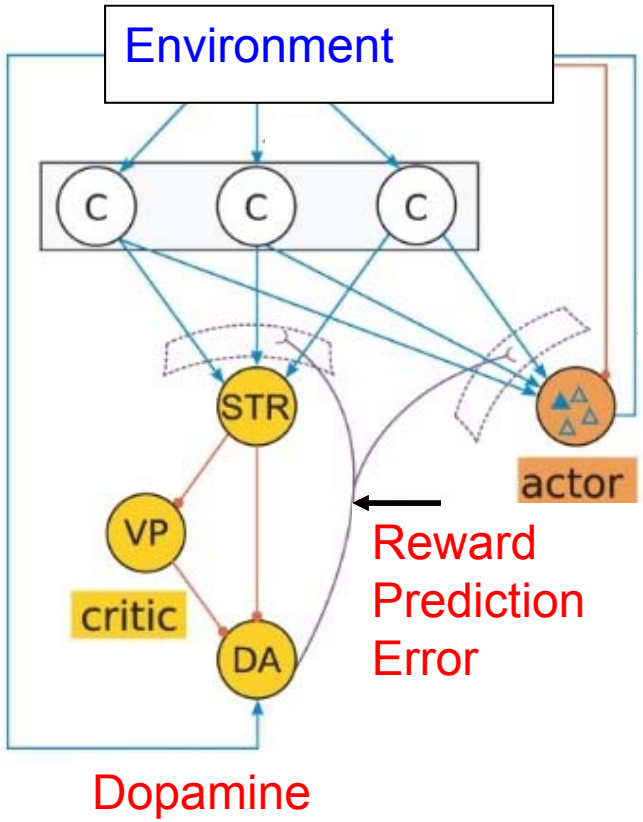

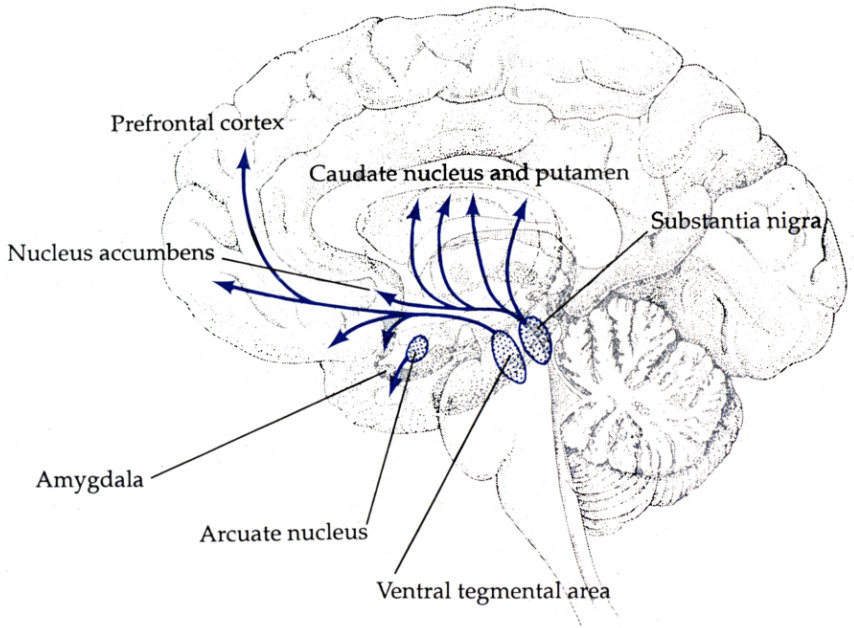

Montague, Dayan and Sejnowski, 1996

Temporal Differences Revealed

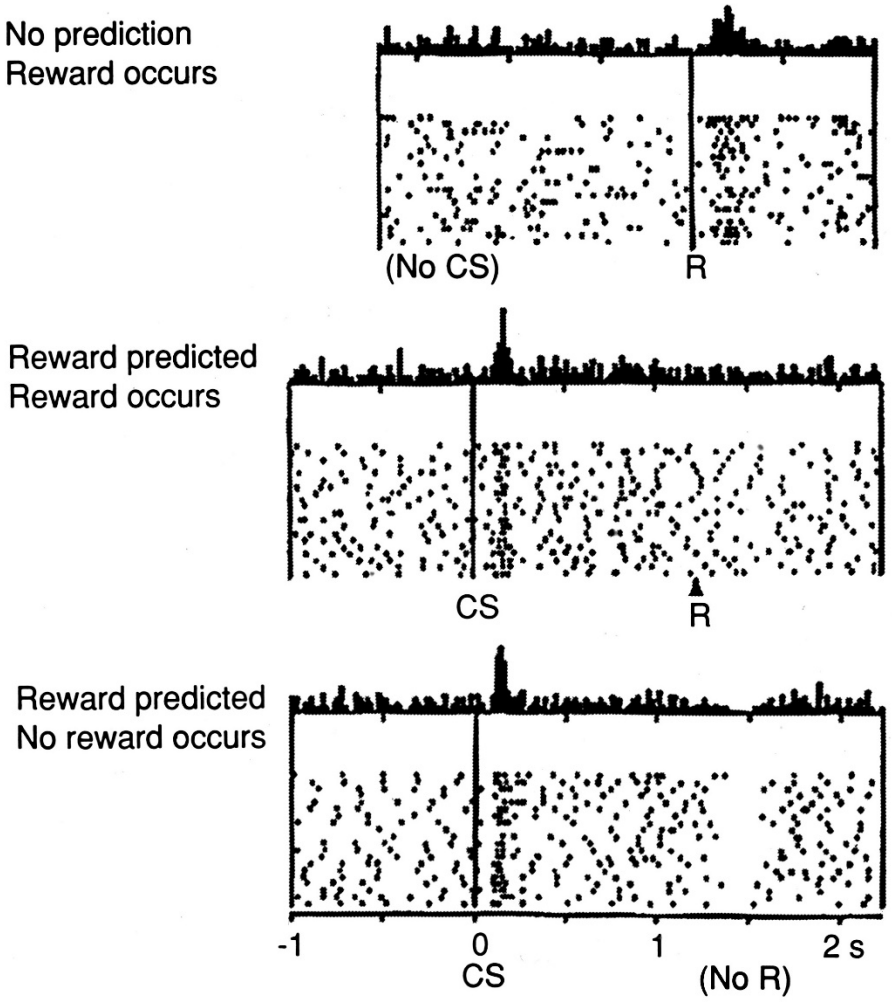

Schultz, Dayan, Montague, 1997

Temporal Difference Learning

Cerebral Cortex

Basal Ganglia

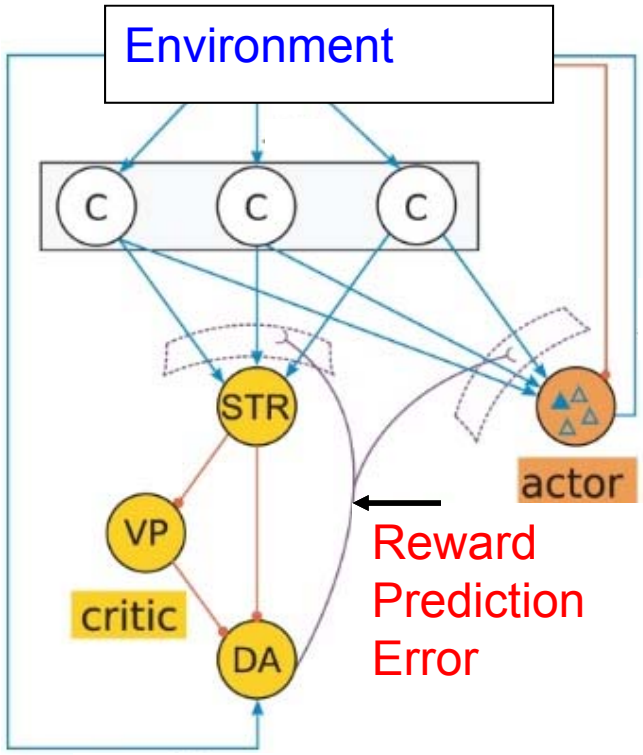

Go Defeat, 2017

DeepMind

Speed-Accuracy Tradeoff

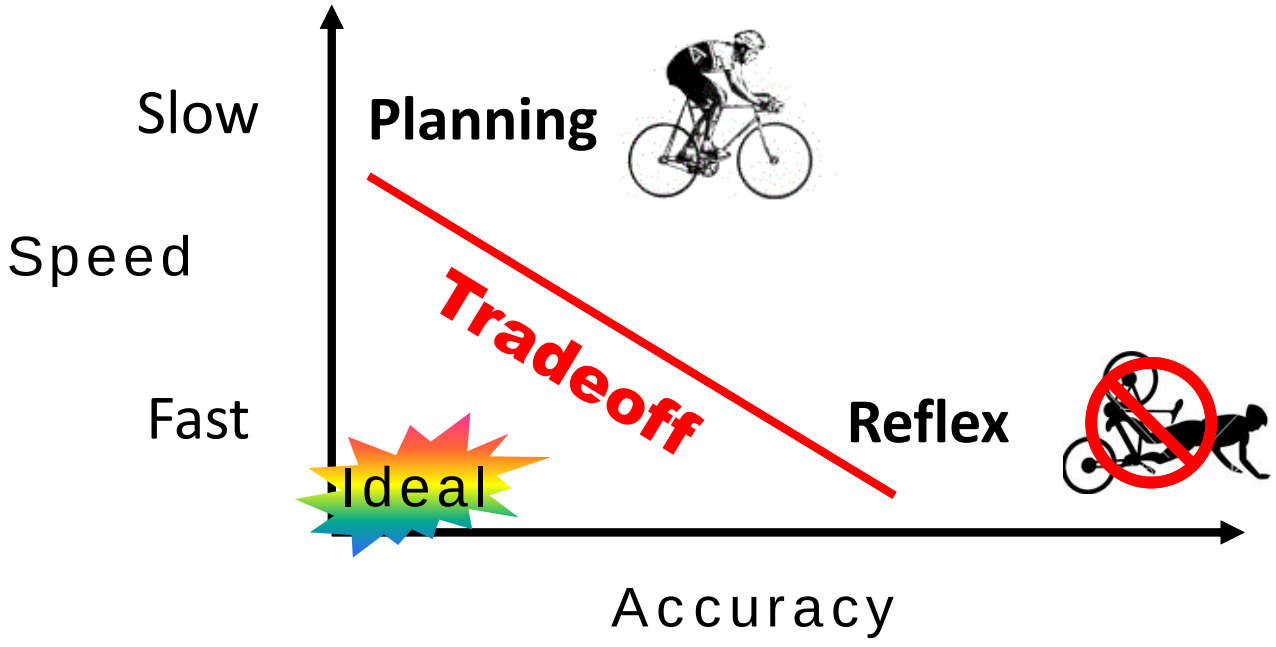

Accurate	Inaccurate
Flexible	Rigid
Centralized	Localized, distributed
Conscious	Unconscious
Deliberate	Automatic
Stable virtual	Unstable real dynamics

John Doyle

Controlling Smooth Eye Movements

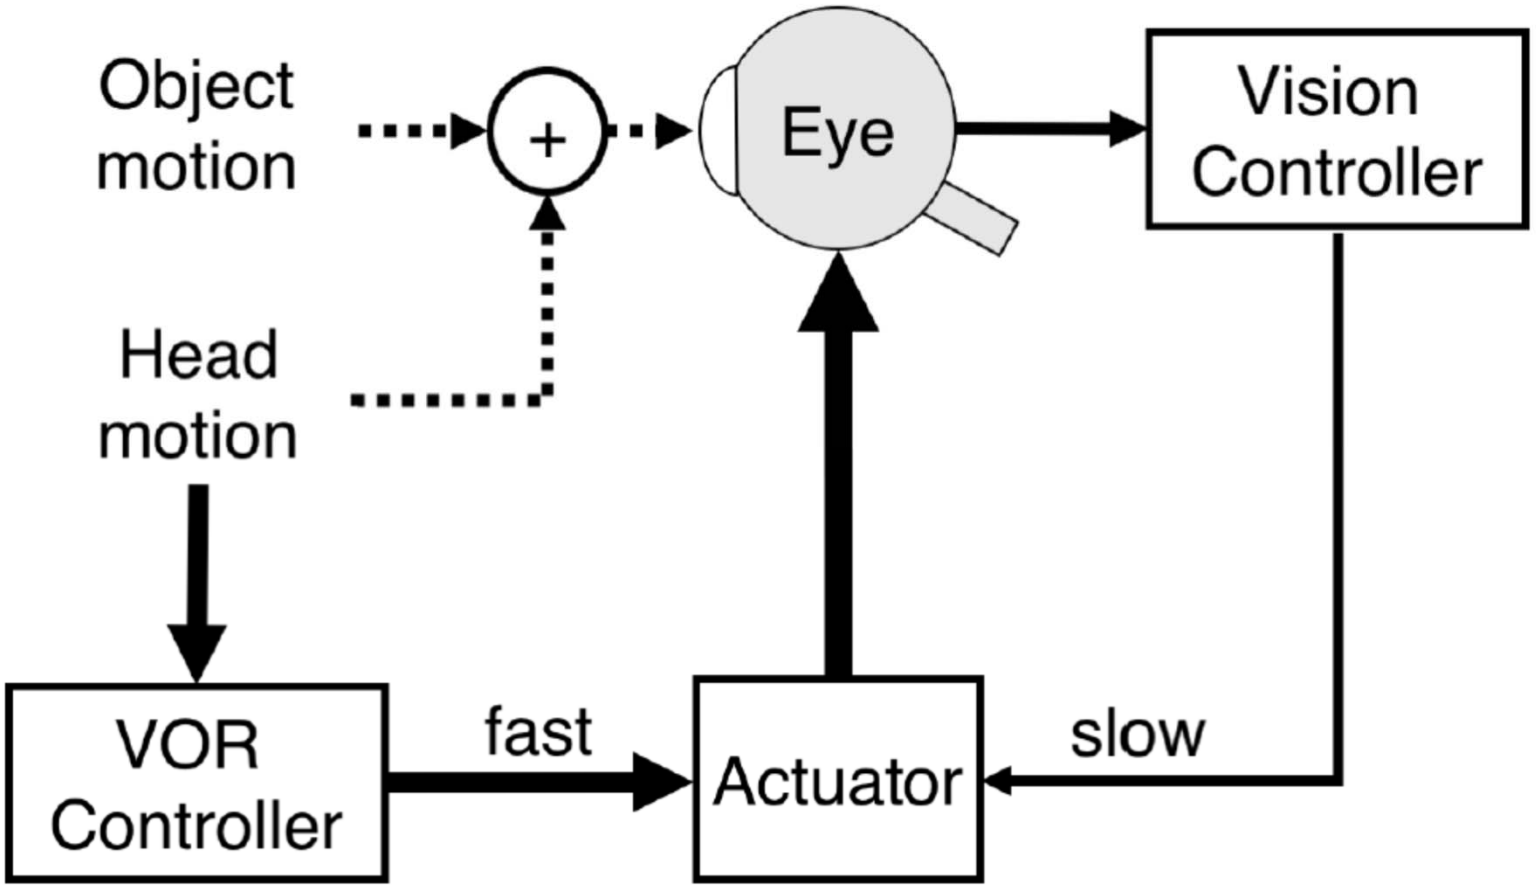

Sensorimotor Control System

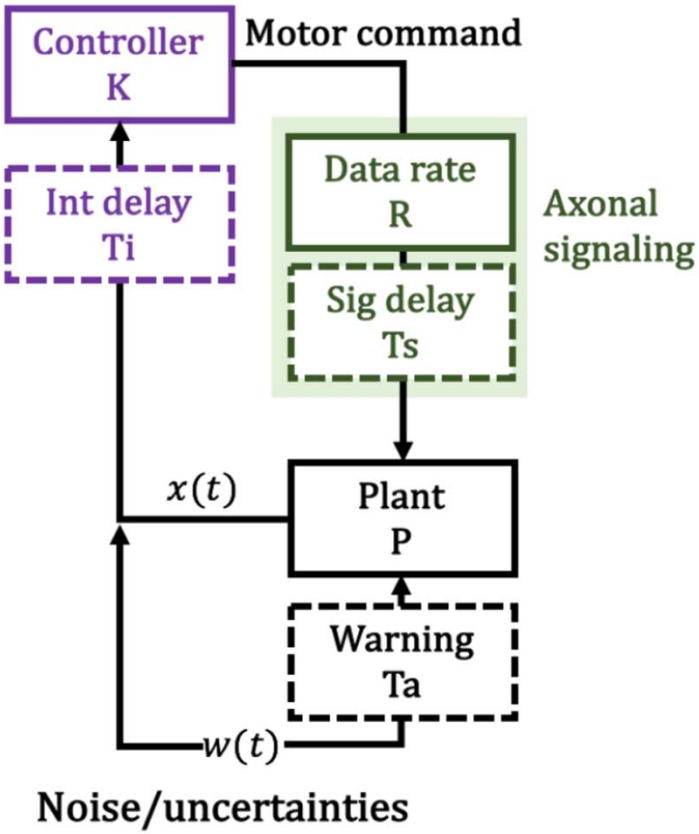

$x(t)$	error at time t
$\mathcal{K}_R$	A controller that uses the causal information of the state, disturbance and control input to compute a suitable control action
R	The maximum information (bits) that can be transmitted per unit time (referred to as the signaling rate)
$T_s \geq 0$	The signaling delay
$T_a \geq 0$	The advanced warning
$T_i \geq 0$	The internal delay
$T = T_s + T_i - T_a$	The net delay (when $T \geq 0$) or warning (when $T < 0$), representing the total time elapsed from the moment $w(t)$ is revealed to the controller until the moment a control action can act against it)

Speed-Accuracy Tradeoff

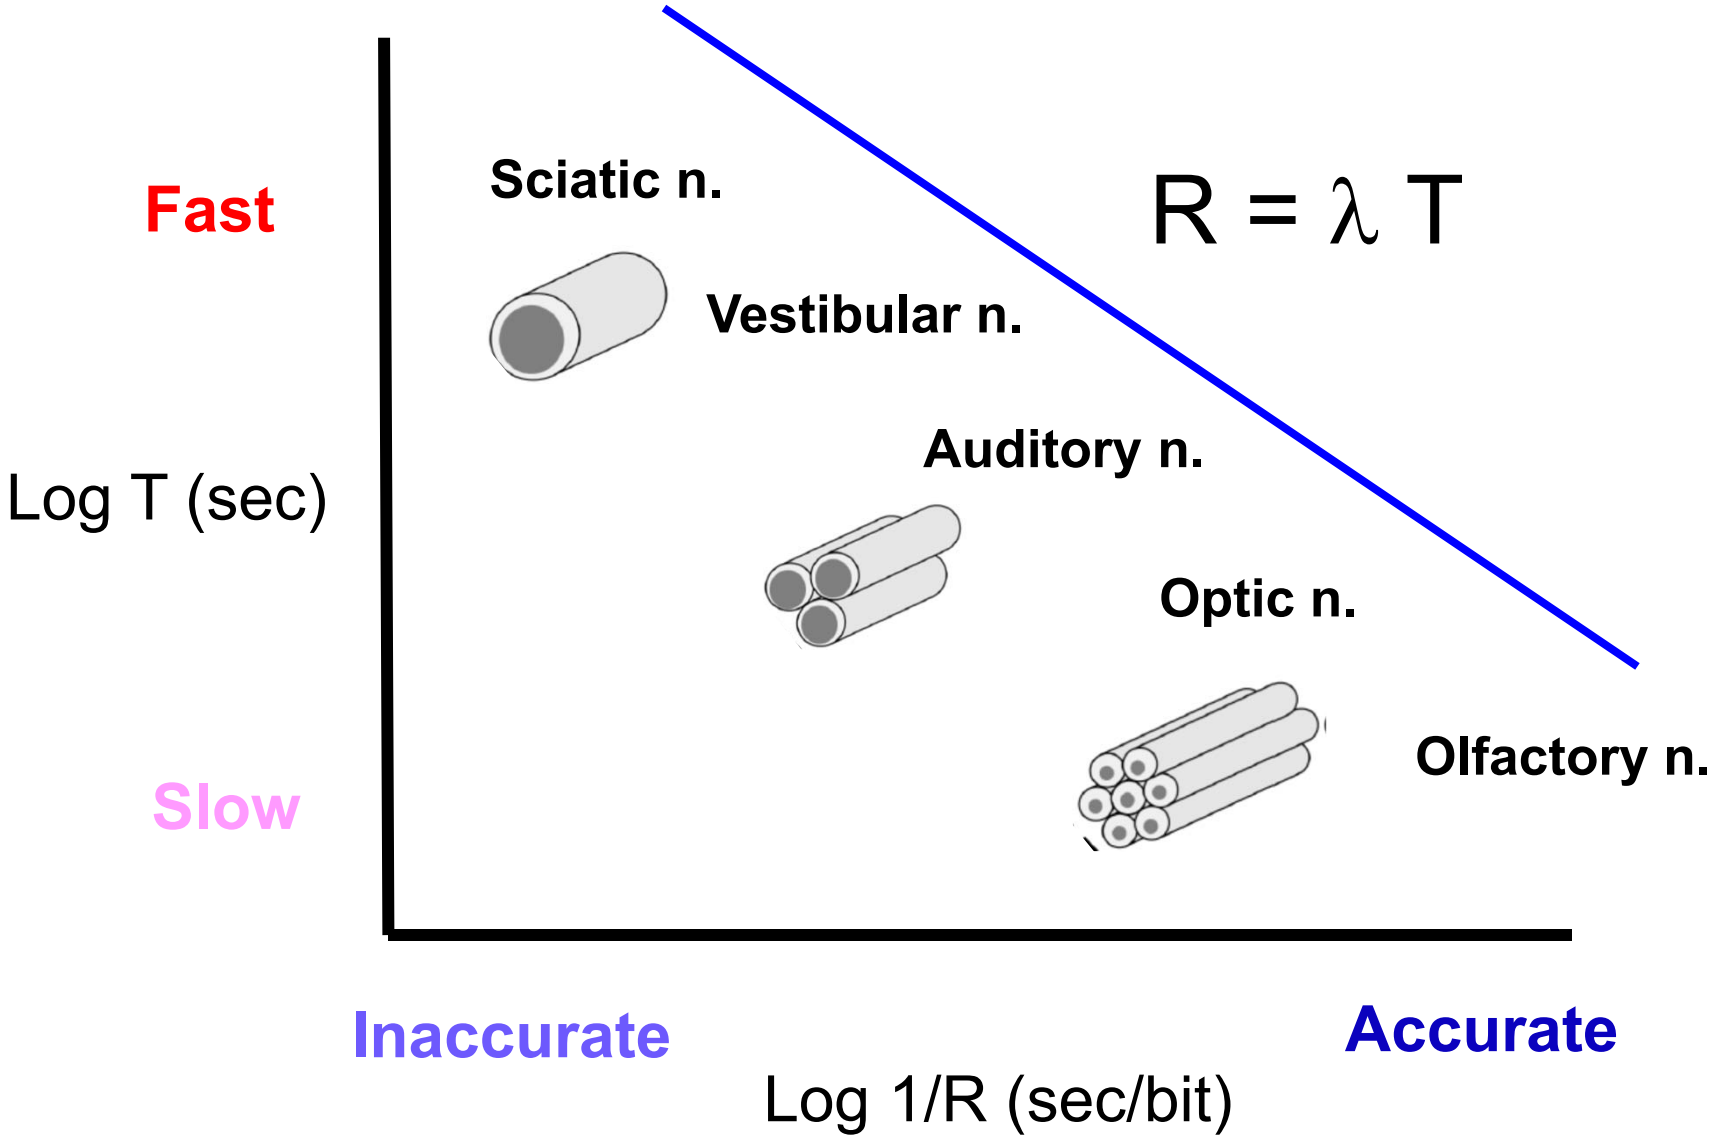

Controlling Smooth Eye Movements

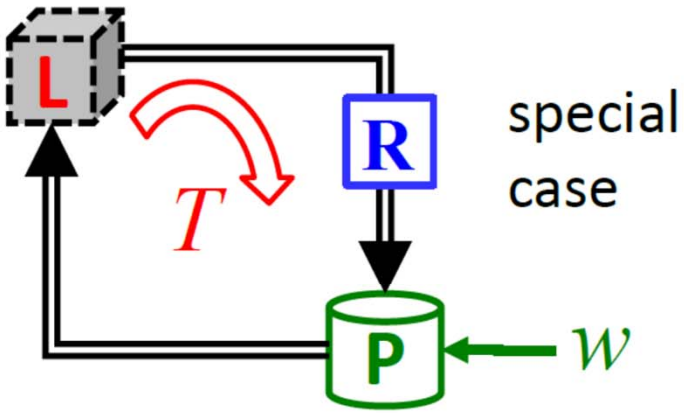

$$R = \lambda_\alpha T$$

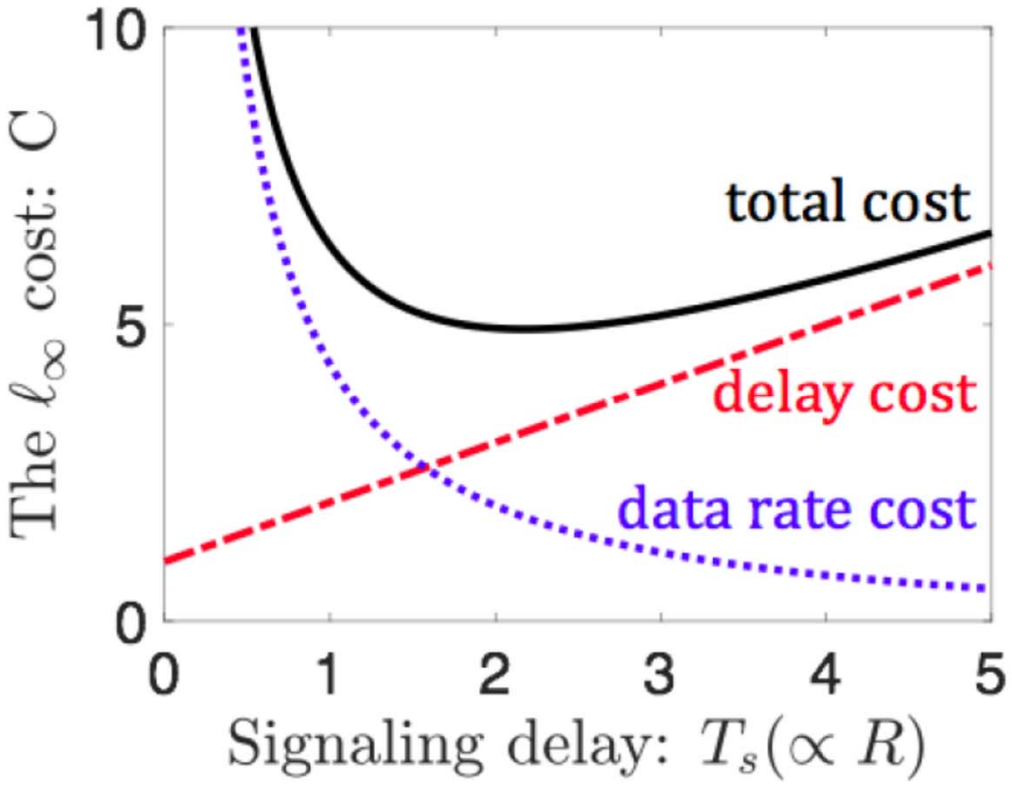

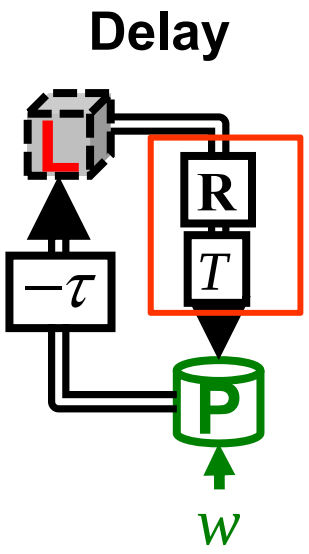
Delay
L
R
T
$-\tau$
P
w

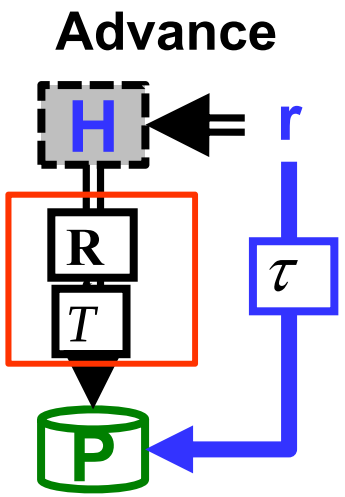
Advance
H
r
R
T
τ
P

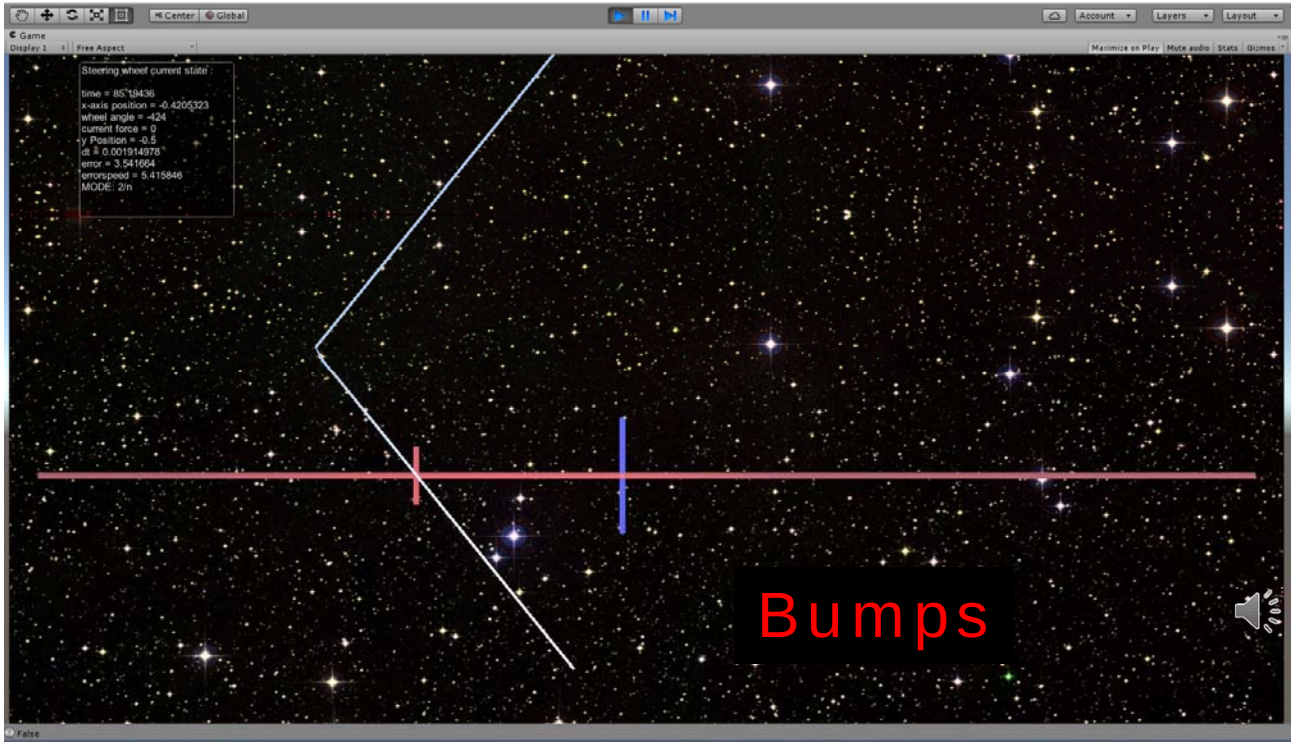
Steering wheel current state :
time = 85.18436
x-axis position = -0.4205323
wheel angle = -424
current force = 0
y Position = -0.5
dt = 0.001914978
error = 3.541664
errorspeed = 5.415846
MODE: 2in
Bumps

Bumps and Trails

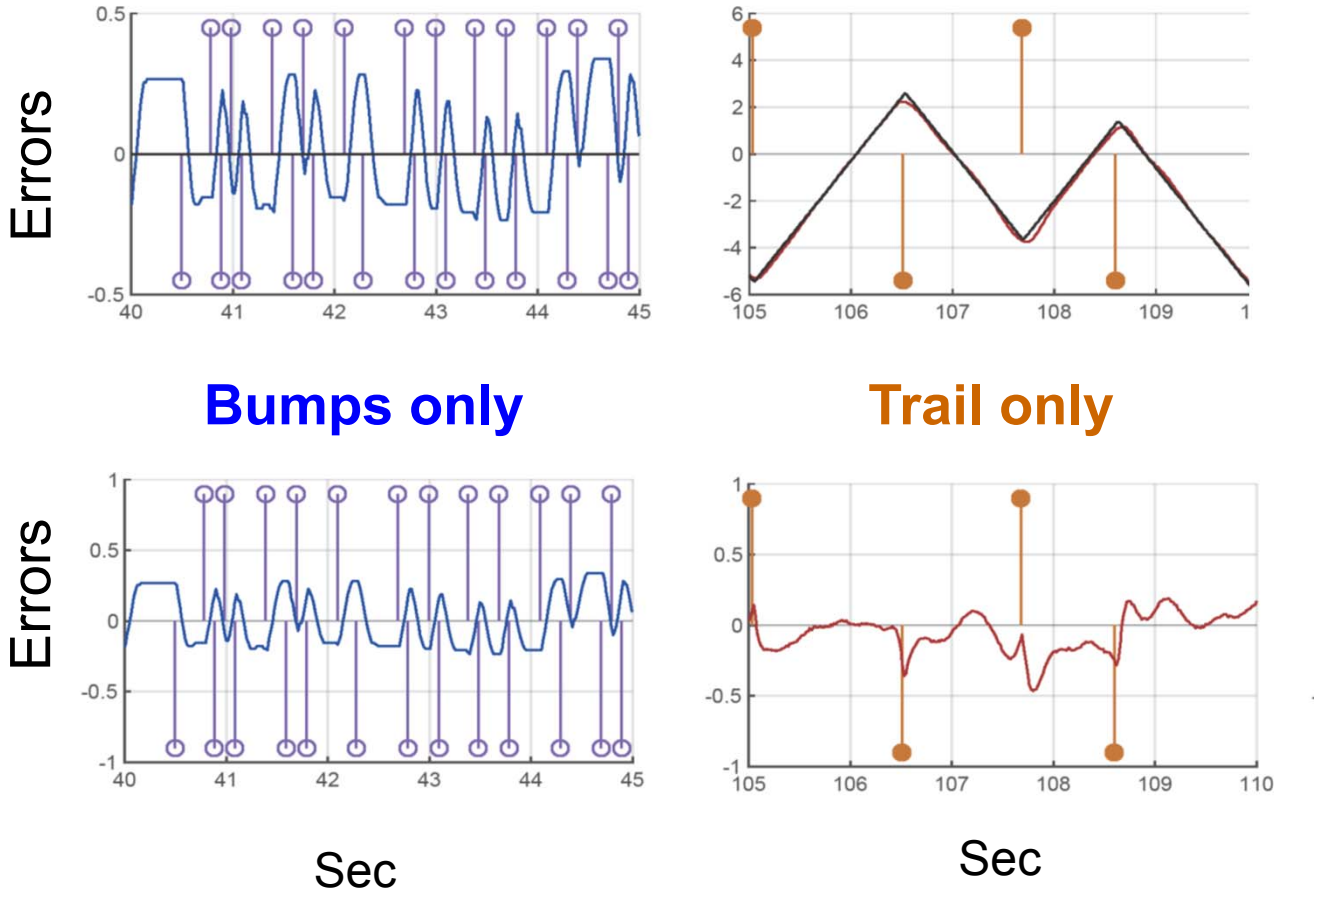

Delays in Vision and in Action

- Vision delay
- Action delay

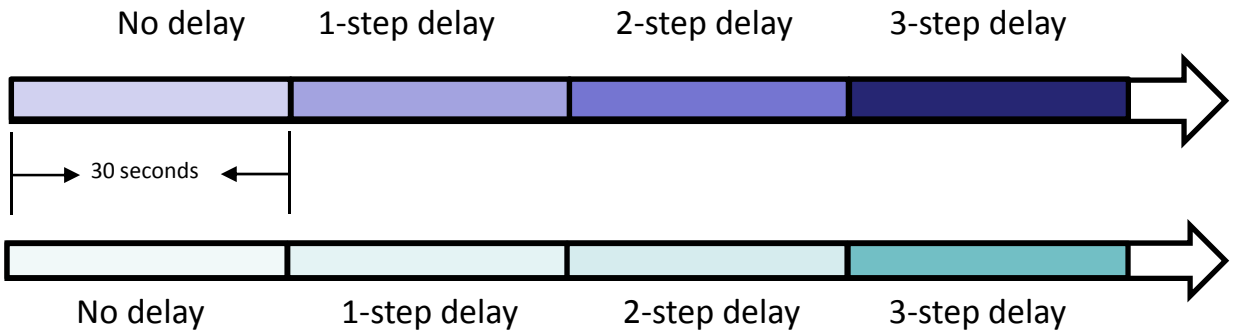

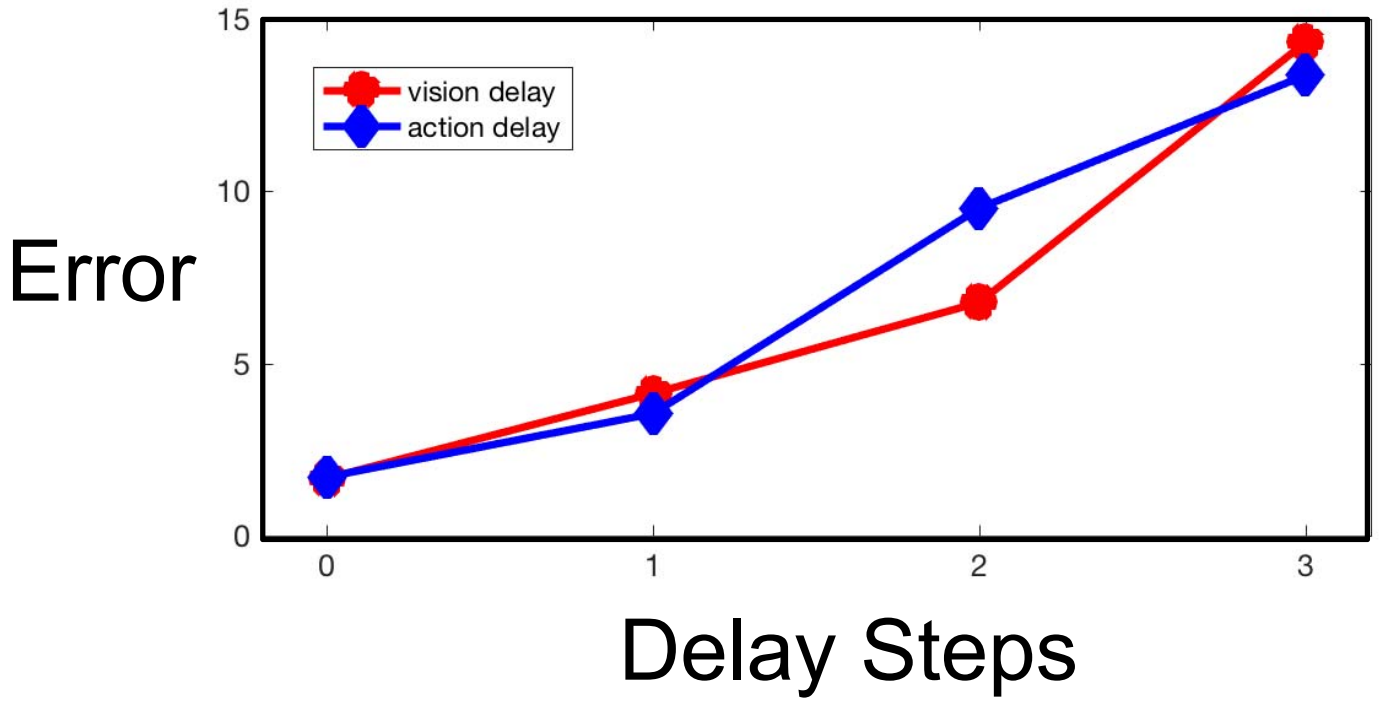

Combining Bumps and Trails

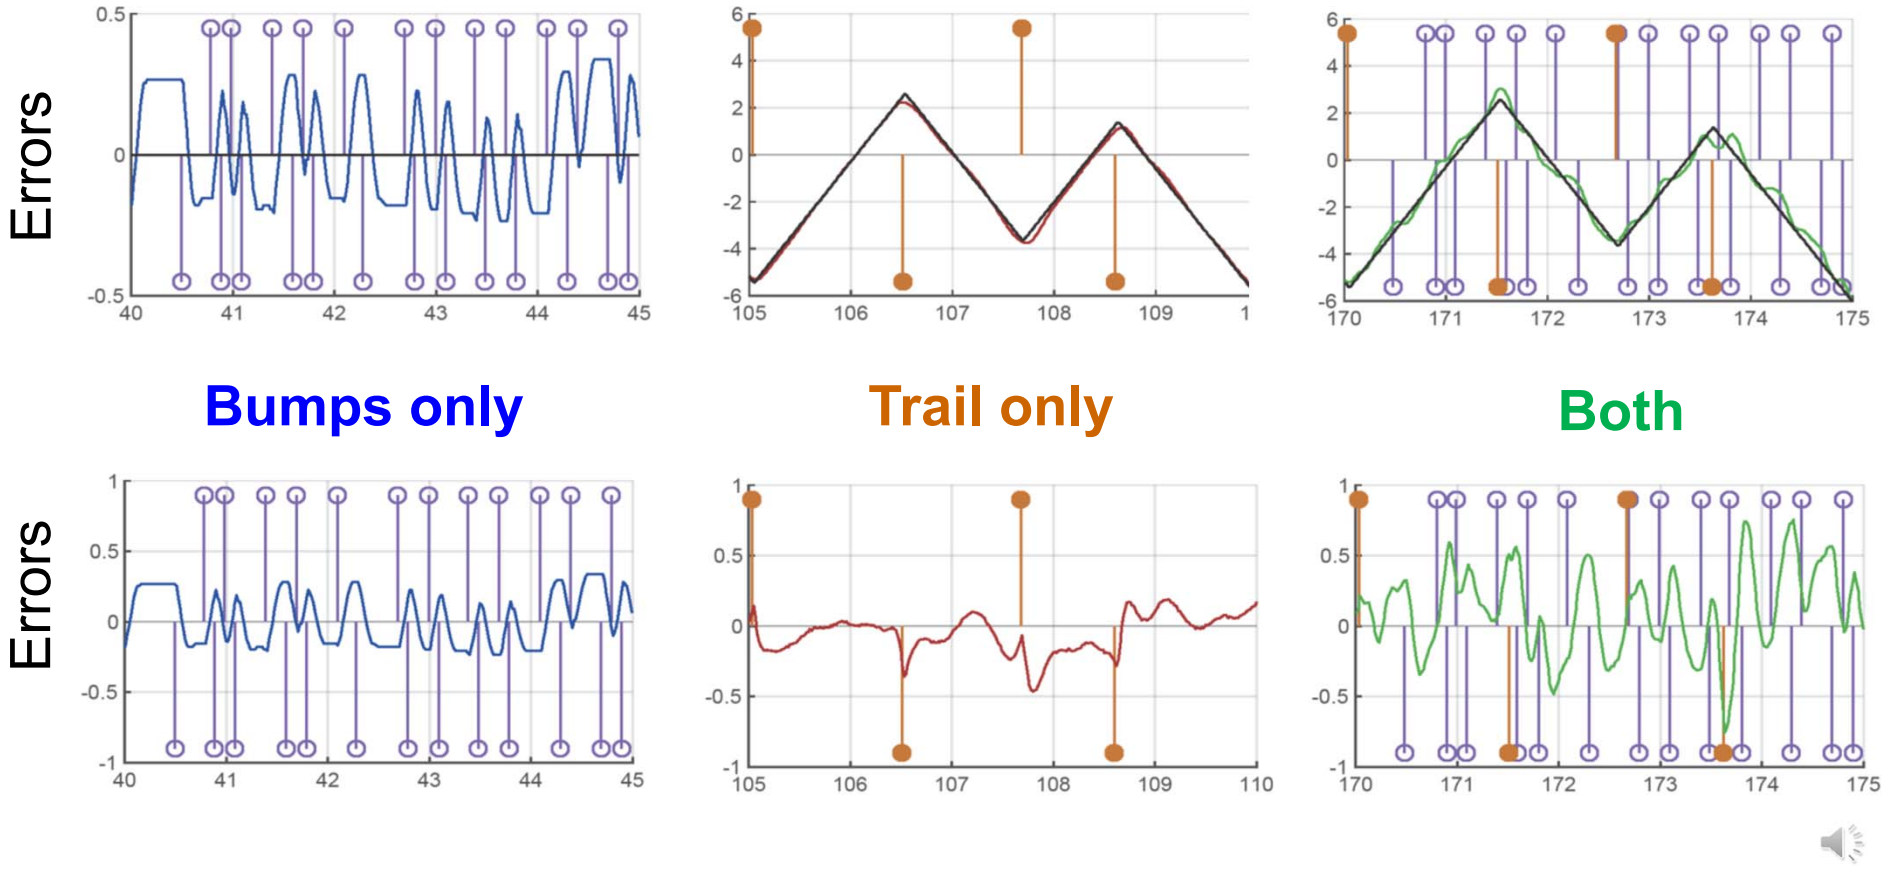

Speed-Accuracy Tradeoff

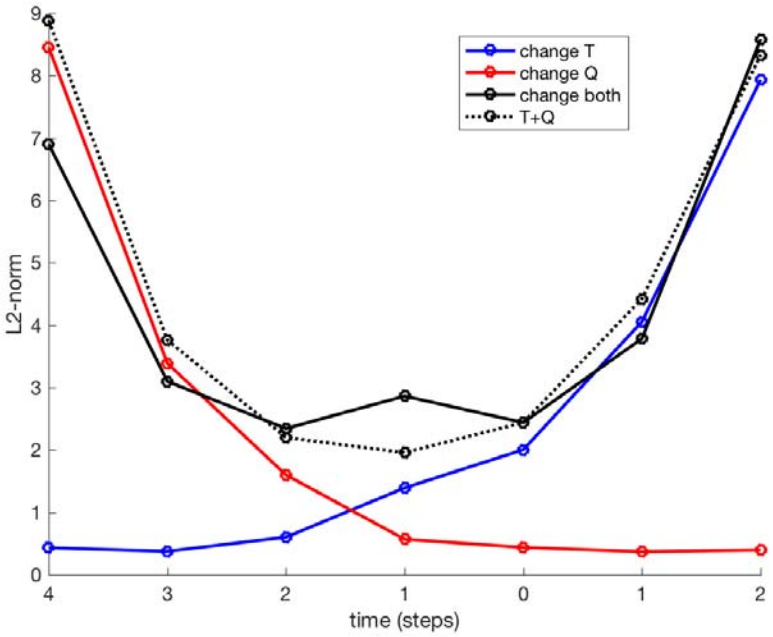

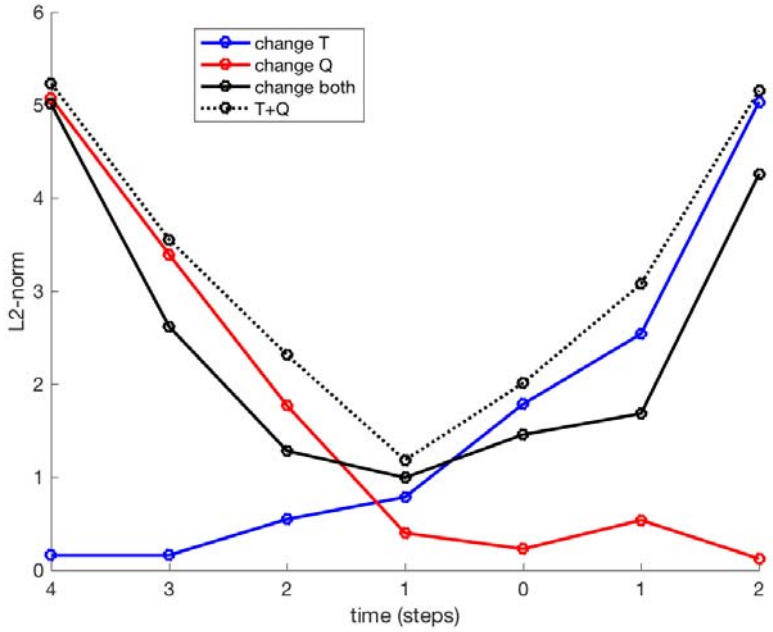

- Delay session: 30 second for each setting (T = -4, -3, -2, -1, 0, 1, 2)
- Quantizer session: 30 second (L = 1, 2, …, 7)
- Delay and Quantizer session: 30 second for each pair of settings

Parallel Control Systems

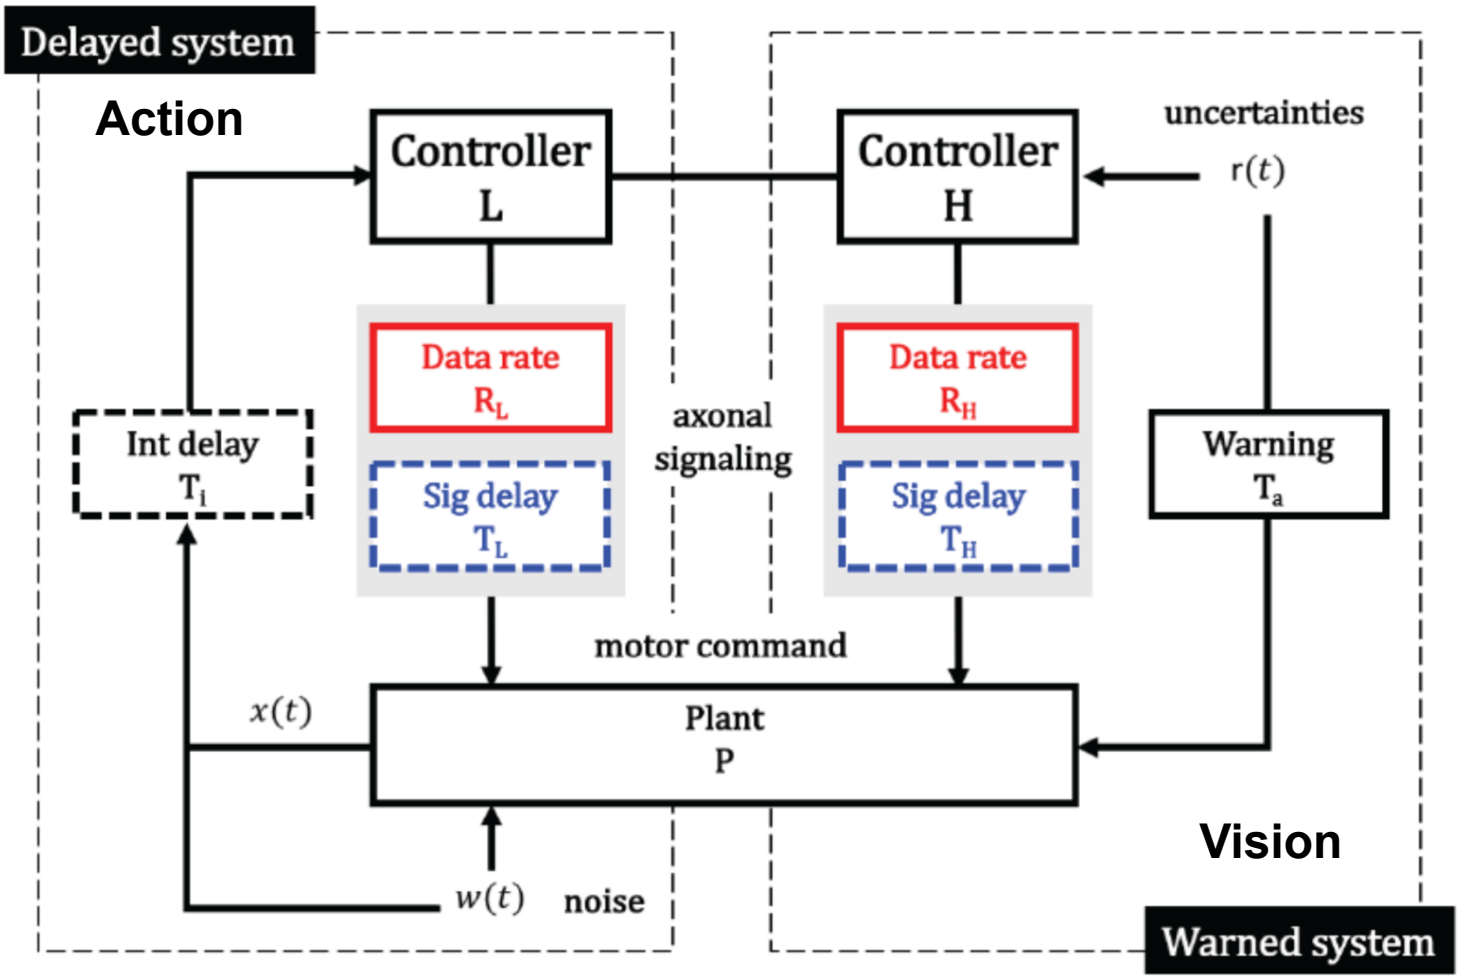

Scalability of Control Systems

Sparse,
distributed &
localized control
w/ delayed
communications

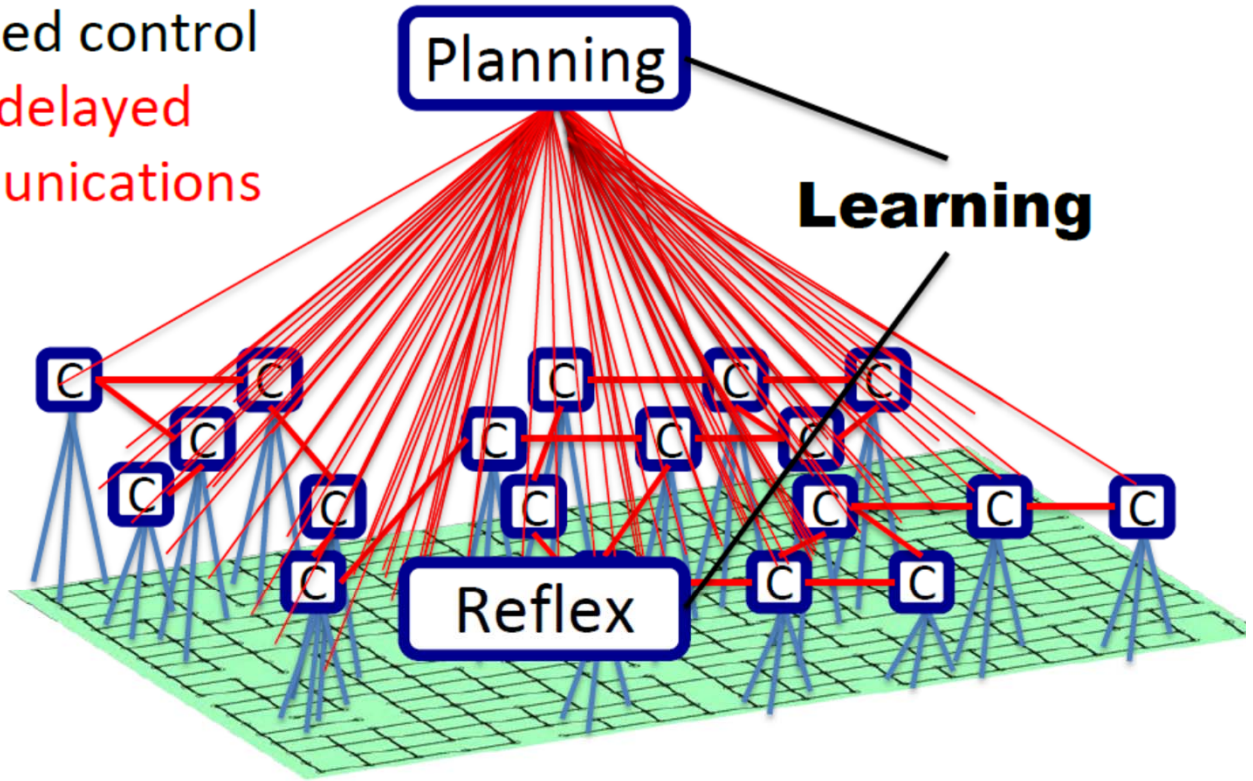

Physical plant

Sparsely connected

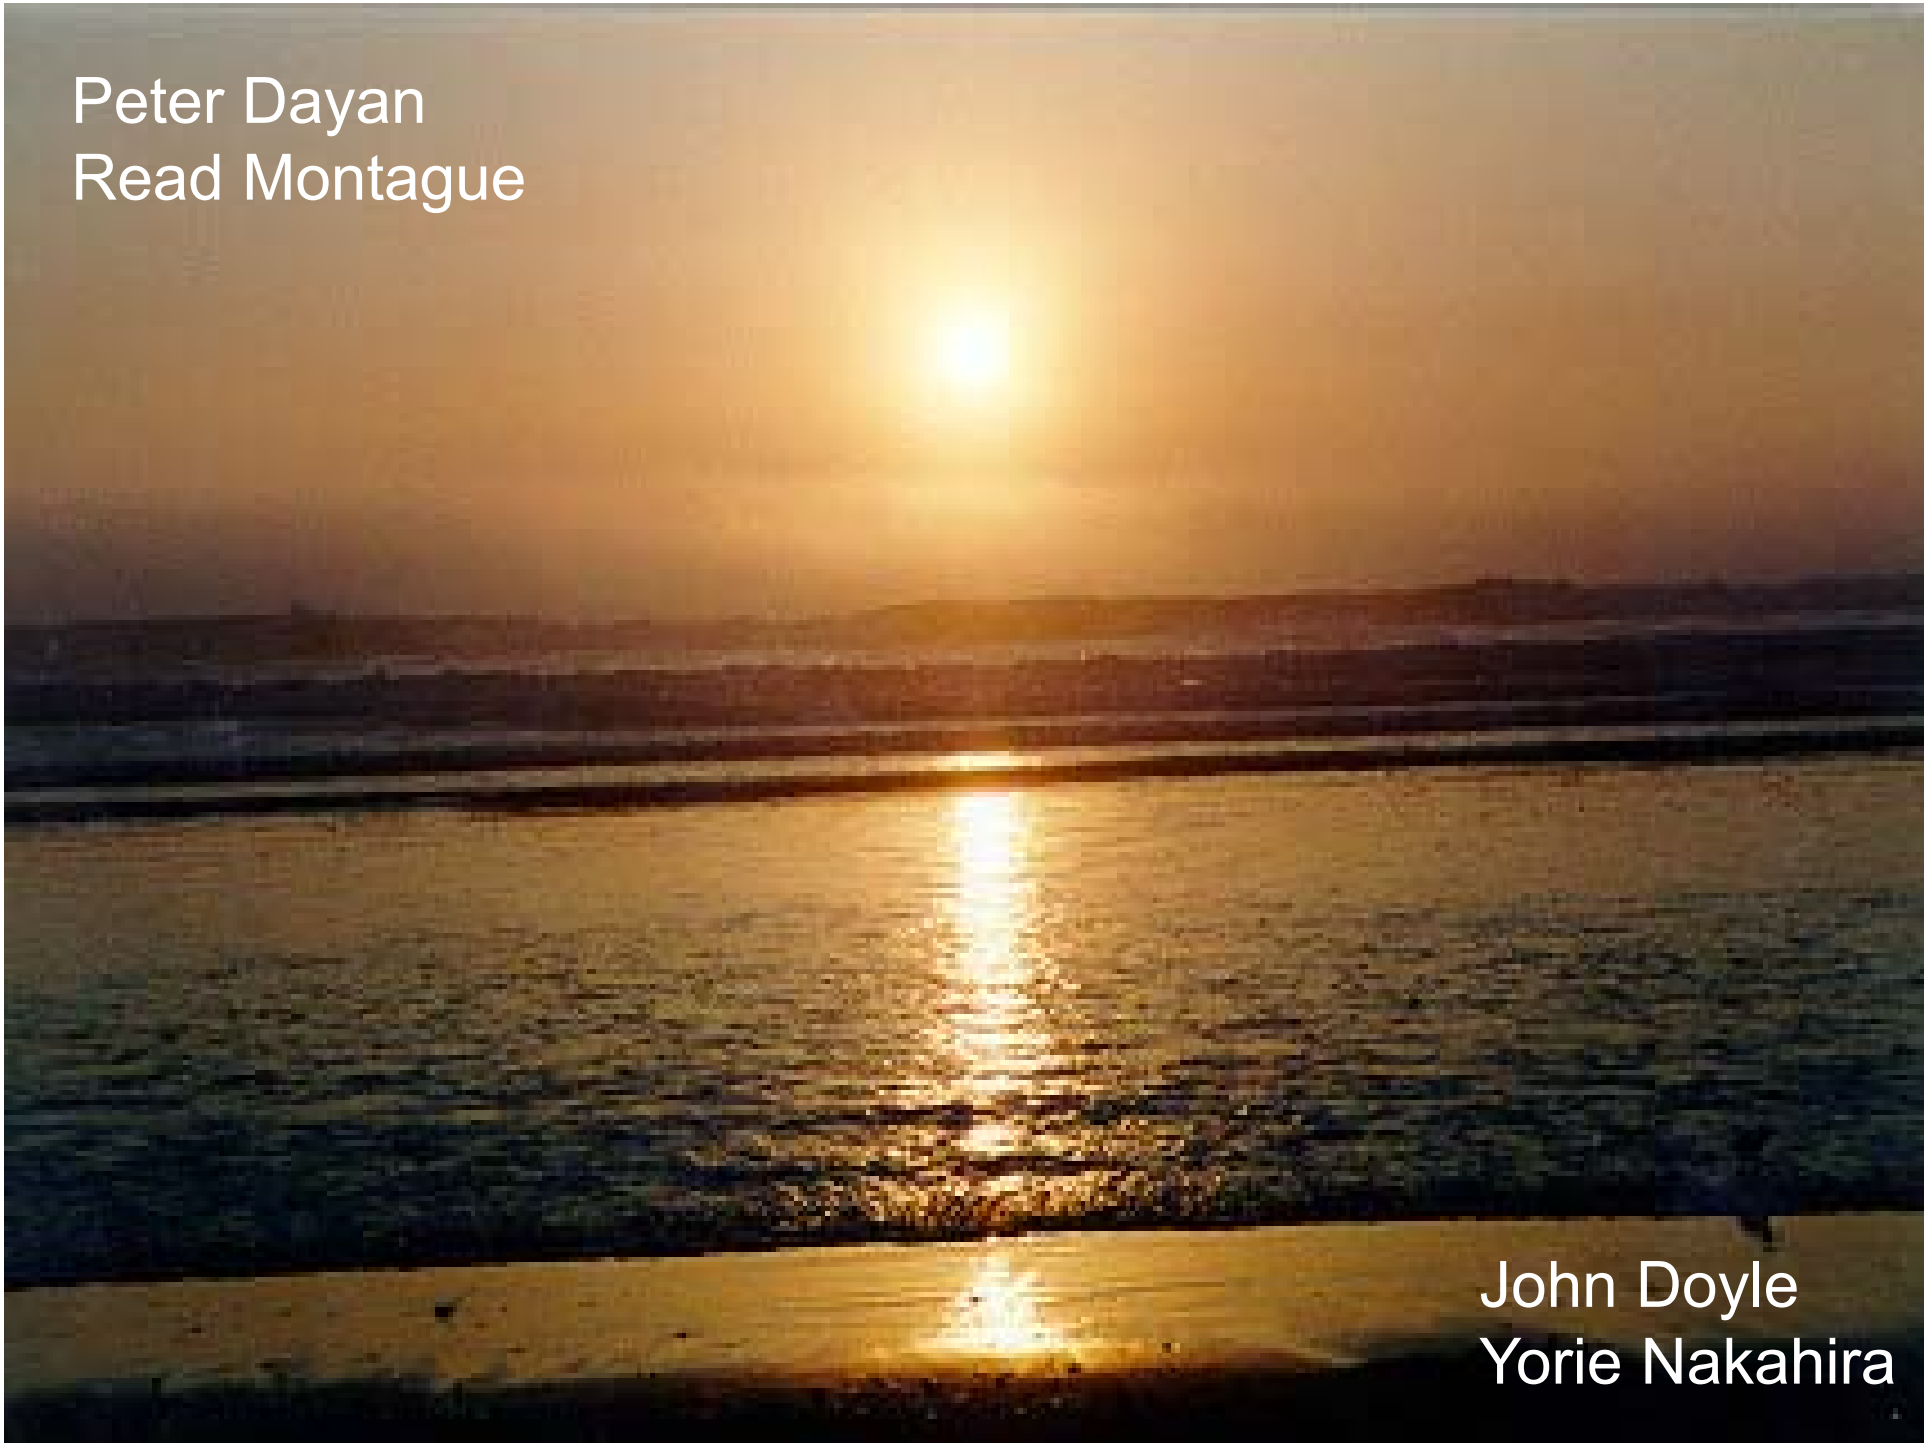
Peter Dayan
Read Montague
John Doyle
Yorie Nakahira