ASHI636: Advanced Topics in Neuroscience

Dr. Olav E. Krigolson

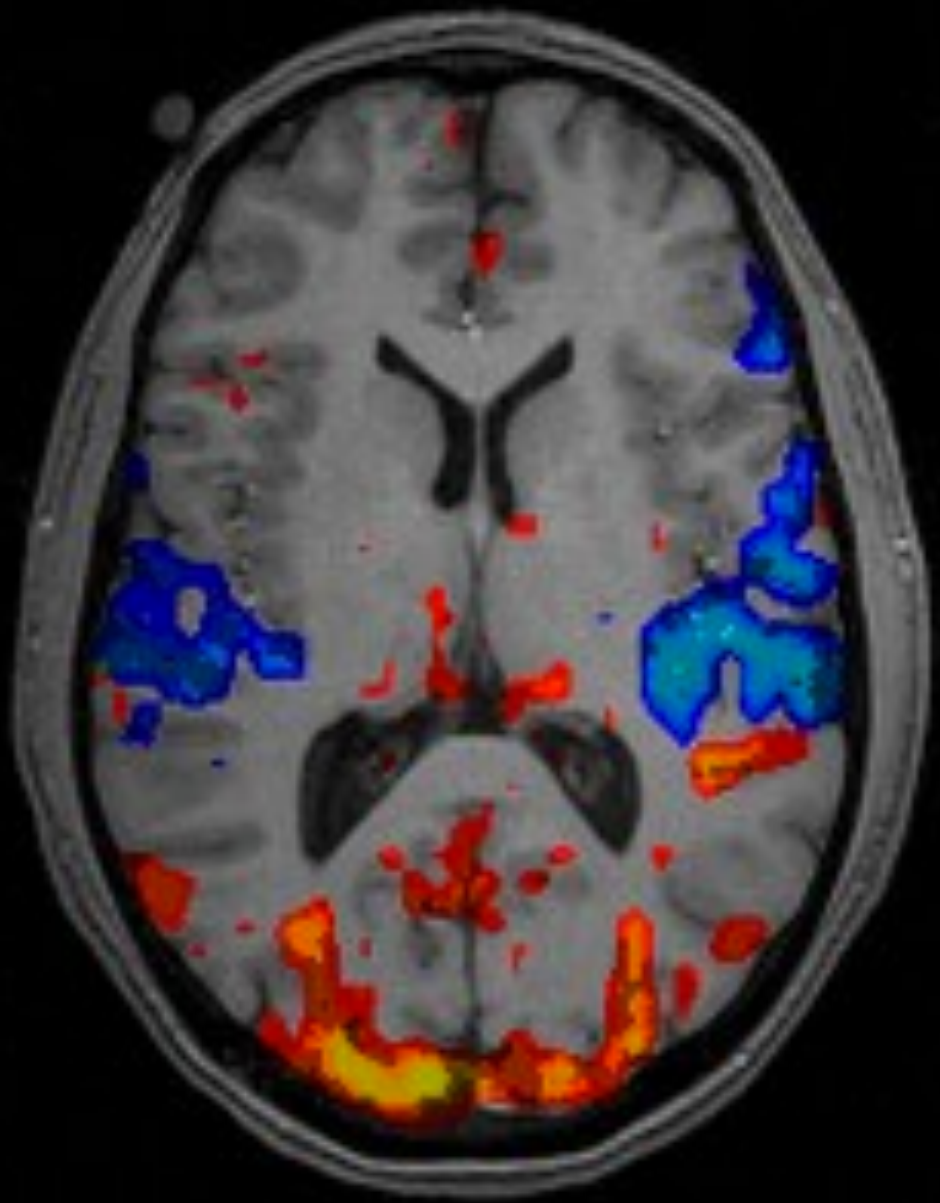

Lecture 1
Learning

Our Course

Jan 28th	Learning
Feb 4th	Decision-Making
Feb 11th	Executive Control
Feb 18th	NO CLASS (READING BREAK)
Feb 25th	Emotion, Good, and Evil
Mar 3rd	Consciousness
Mar 10th	Computational Neuroscience OR Neuroscience Laboratory Experience

Disclaimers

1. Scope of course
   - Eric Kandel
   - Principles of Neuroscience
   - 67 Chapters, 7 Appendices, 1720 pages

2. Level of detail
   - Background
   - Time

3. Medicine

4. What having a PhD really means

Quick Review

What is Neuroscience?

- Cellular Level
- Neuron Level
- Systems Level

What did we cover in Introductory Neuroscience?

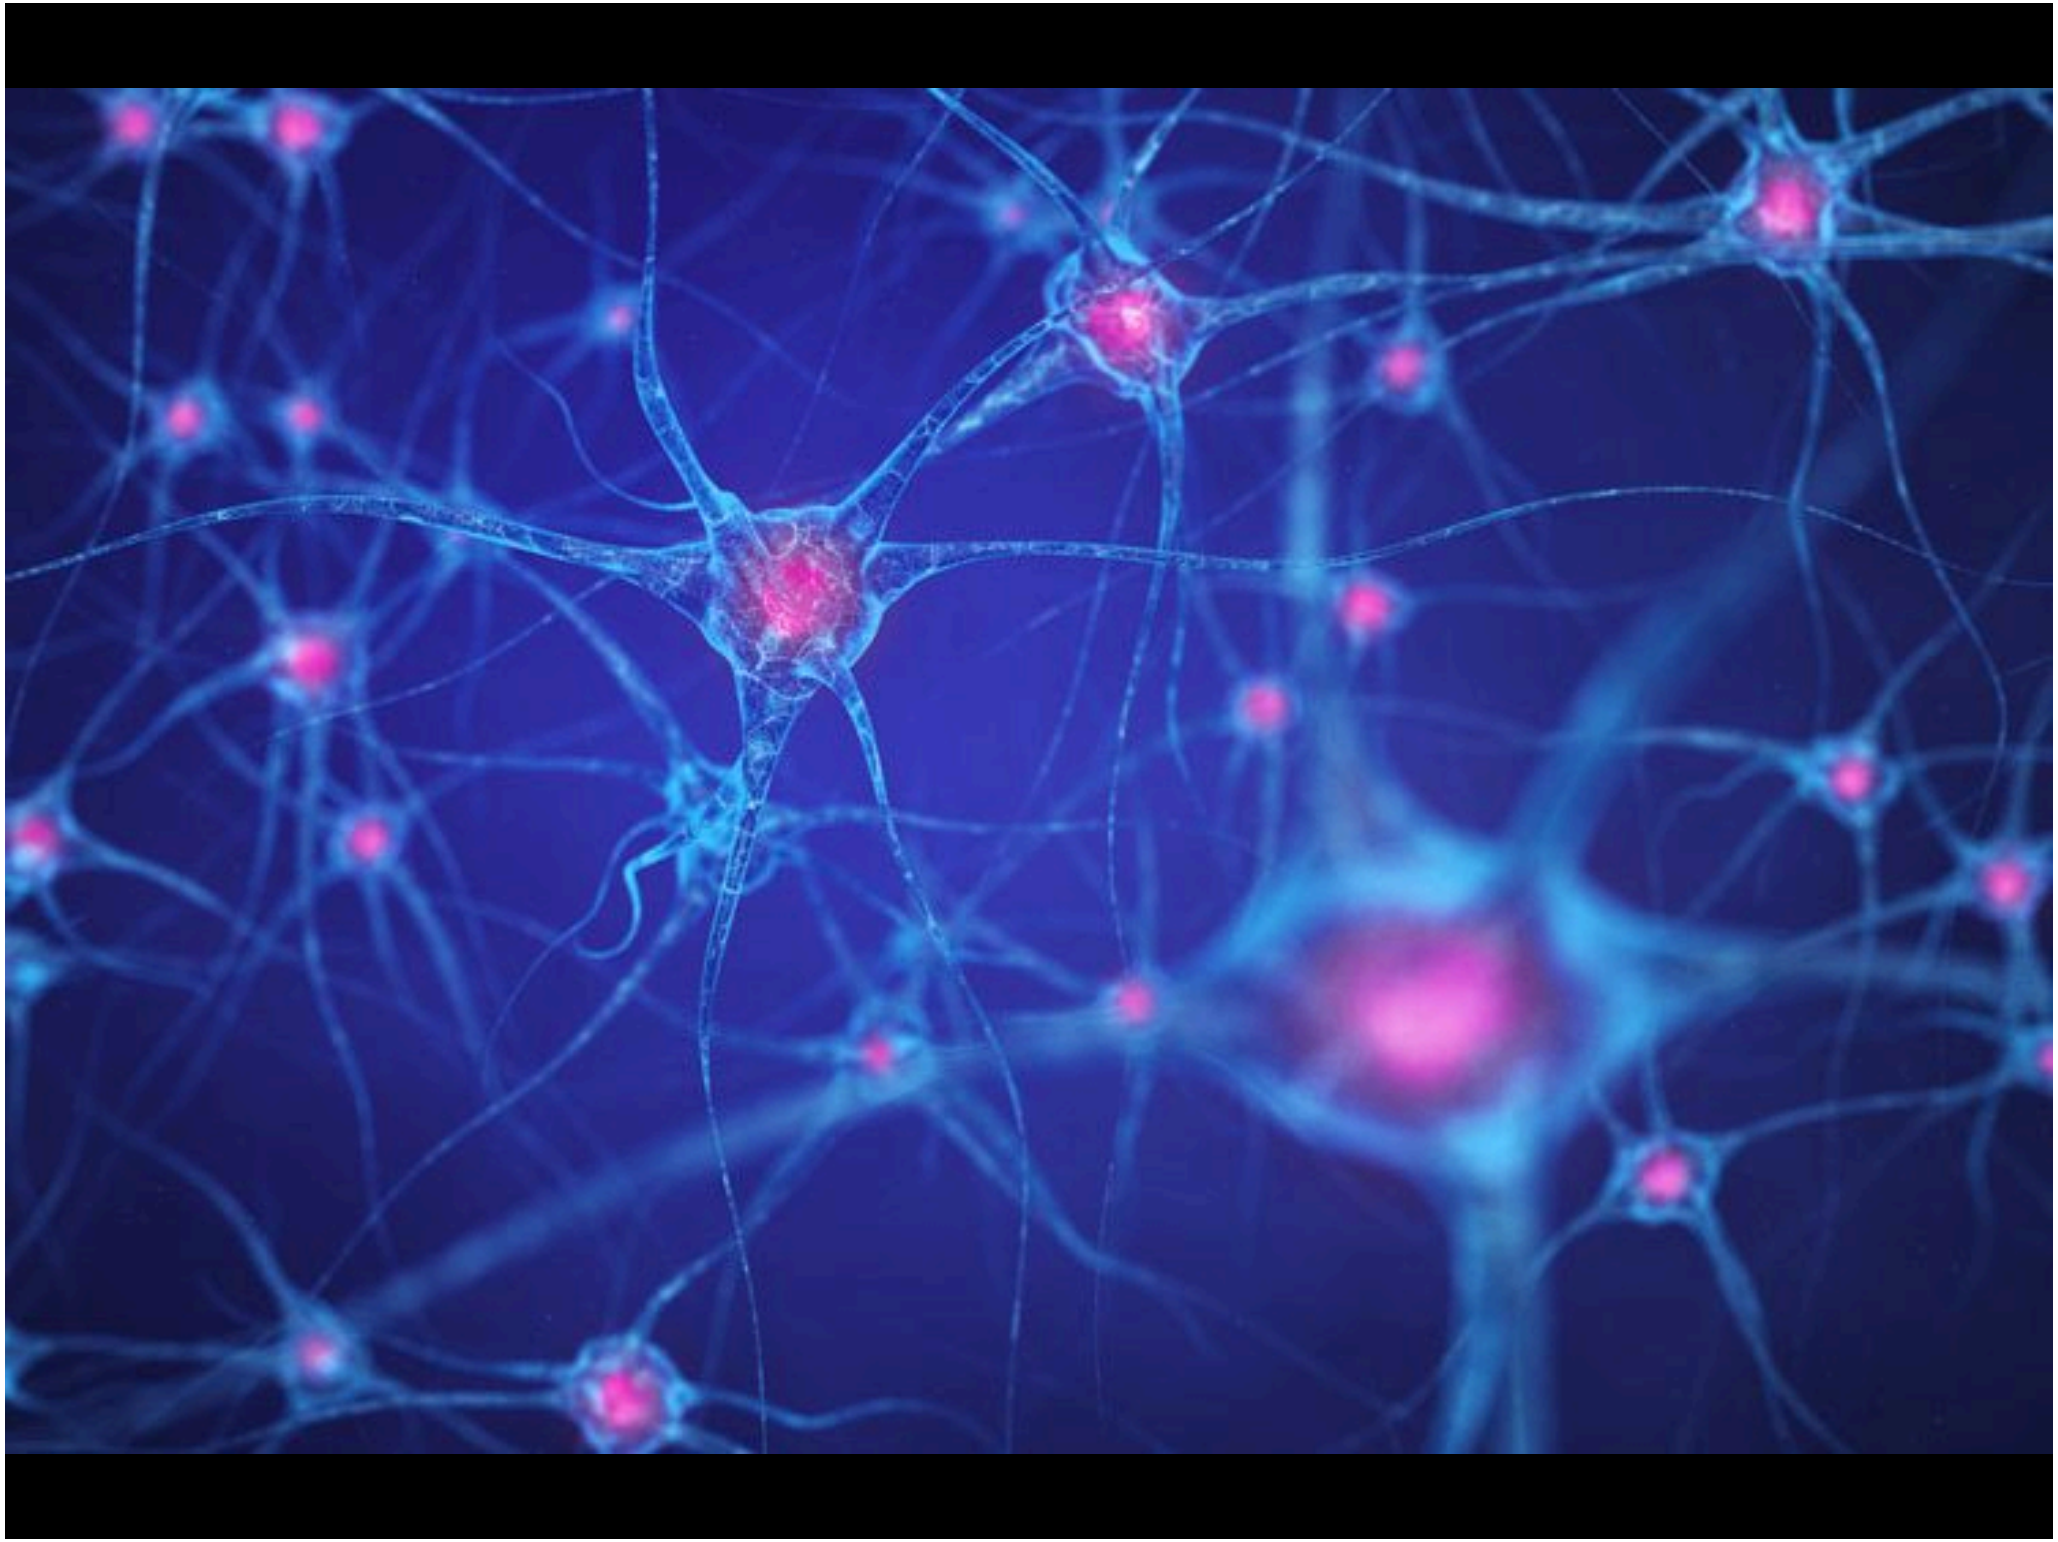

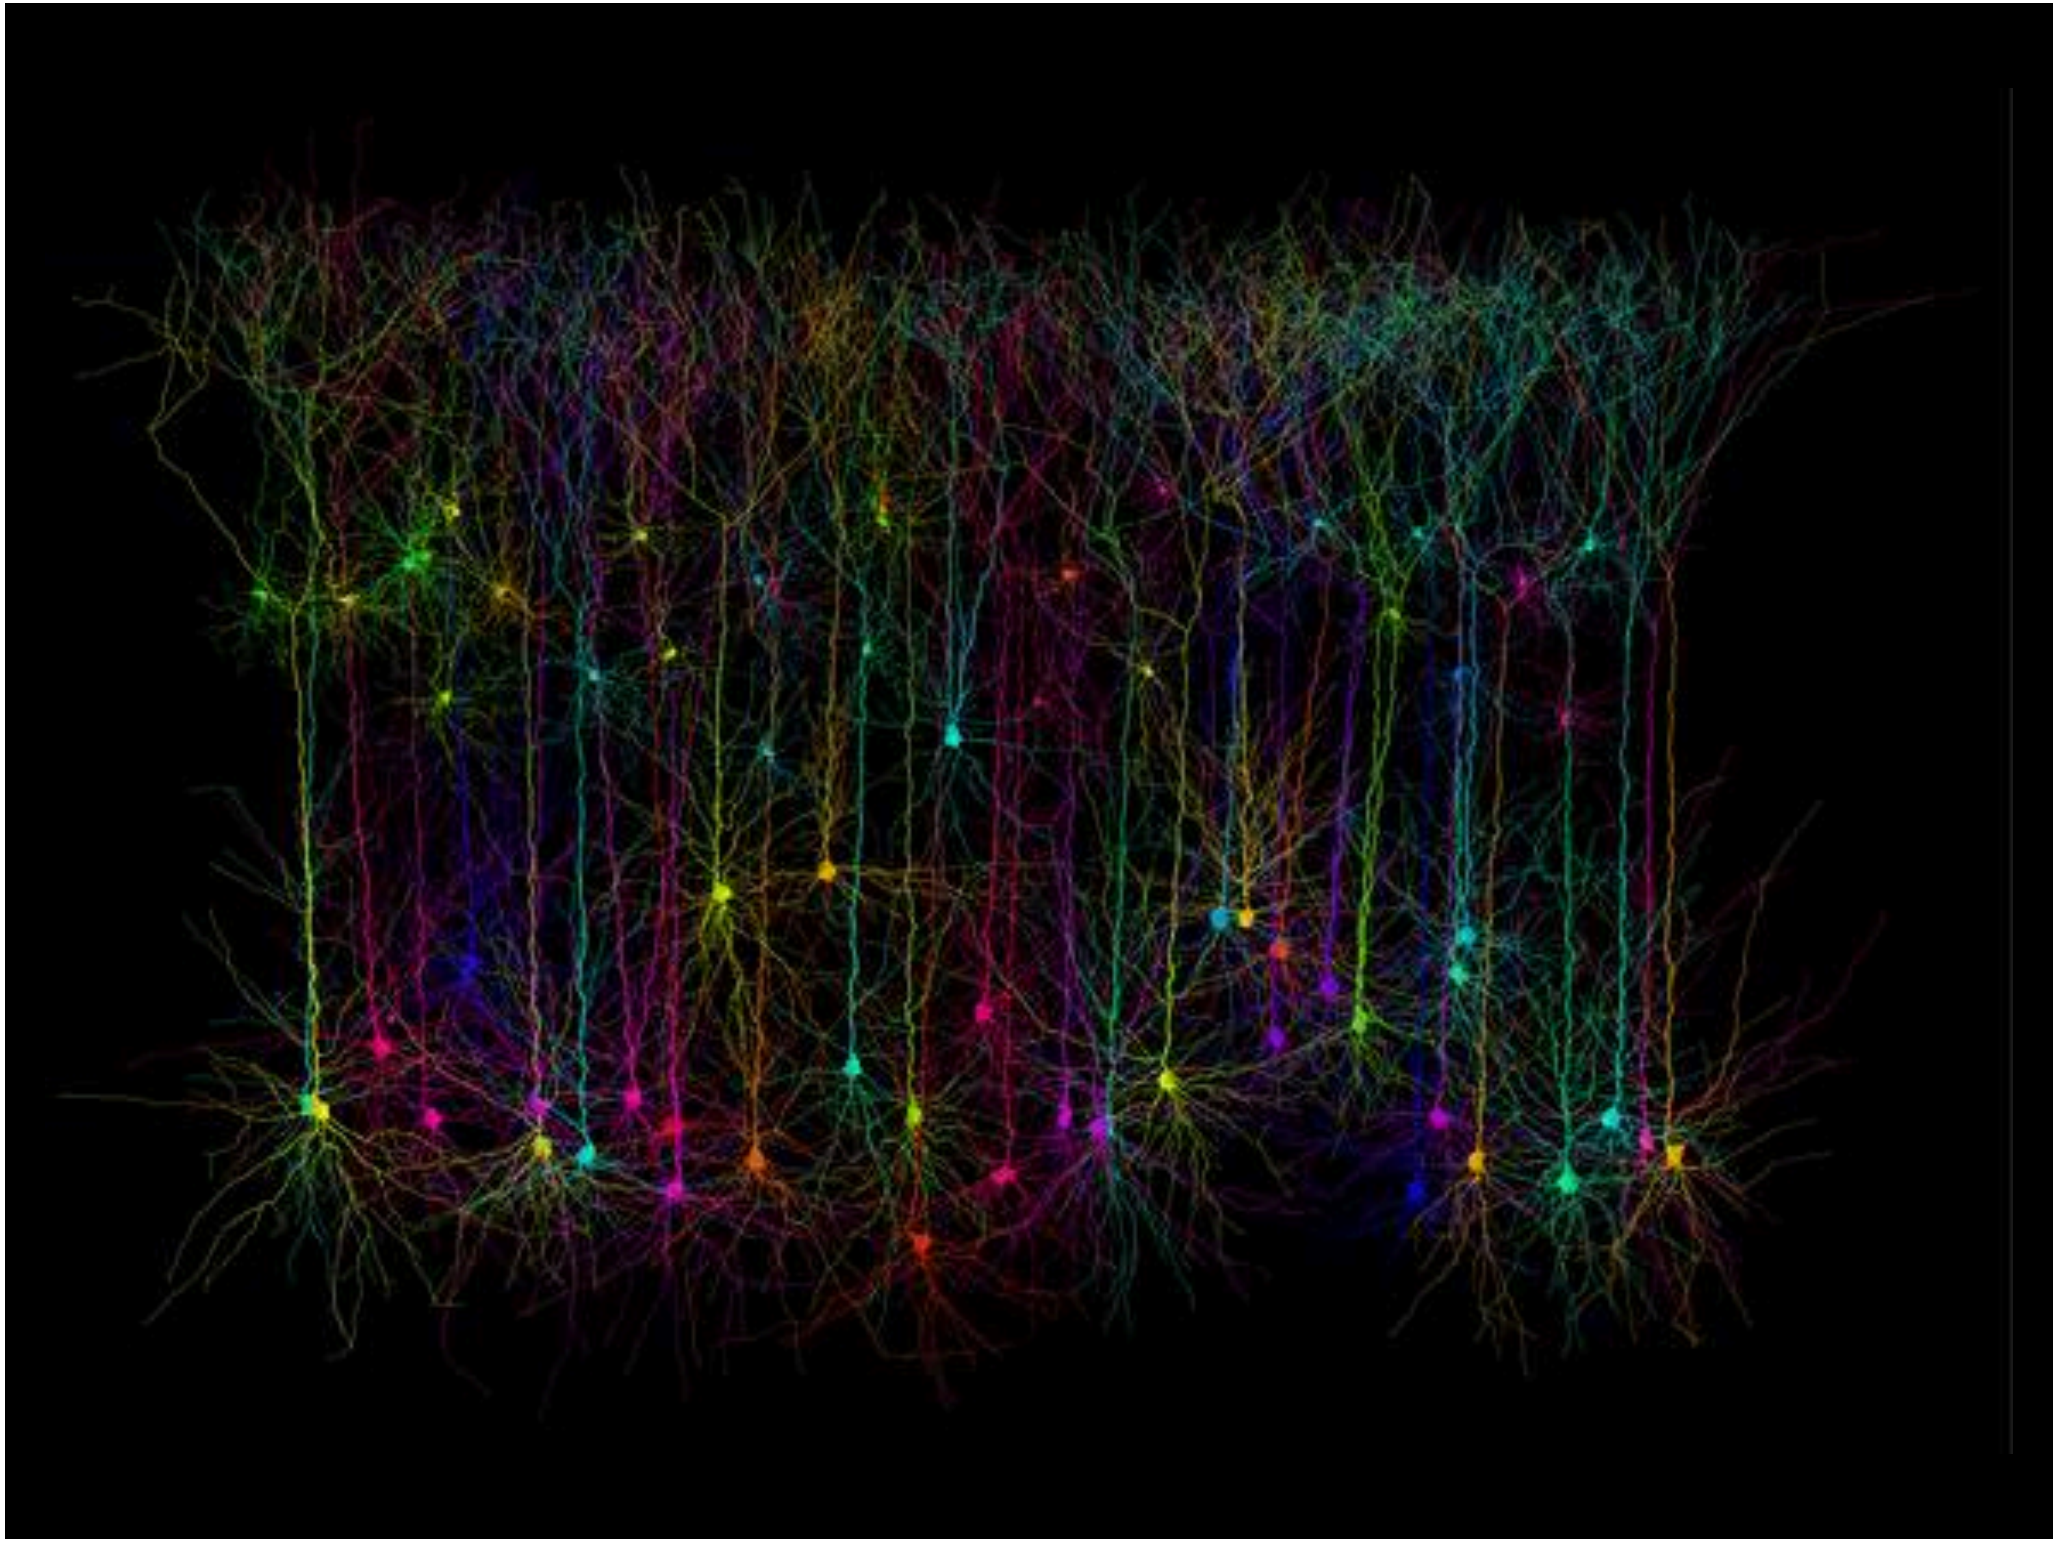

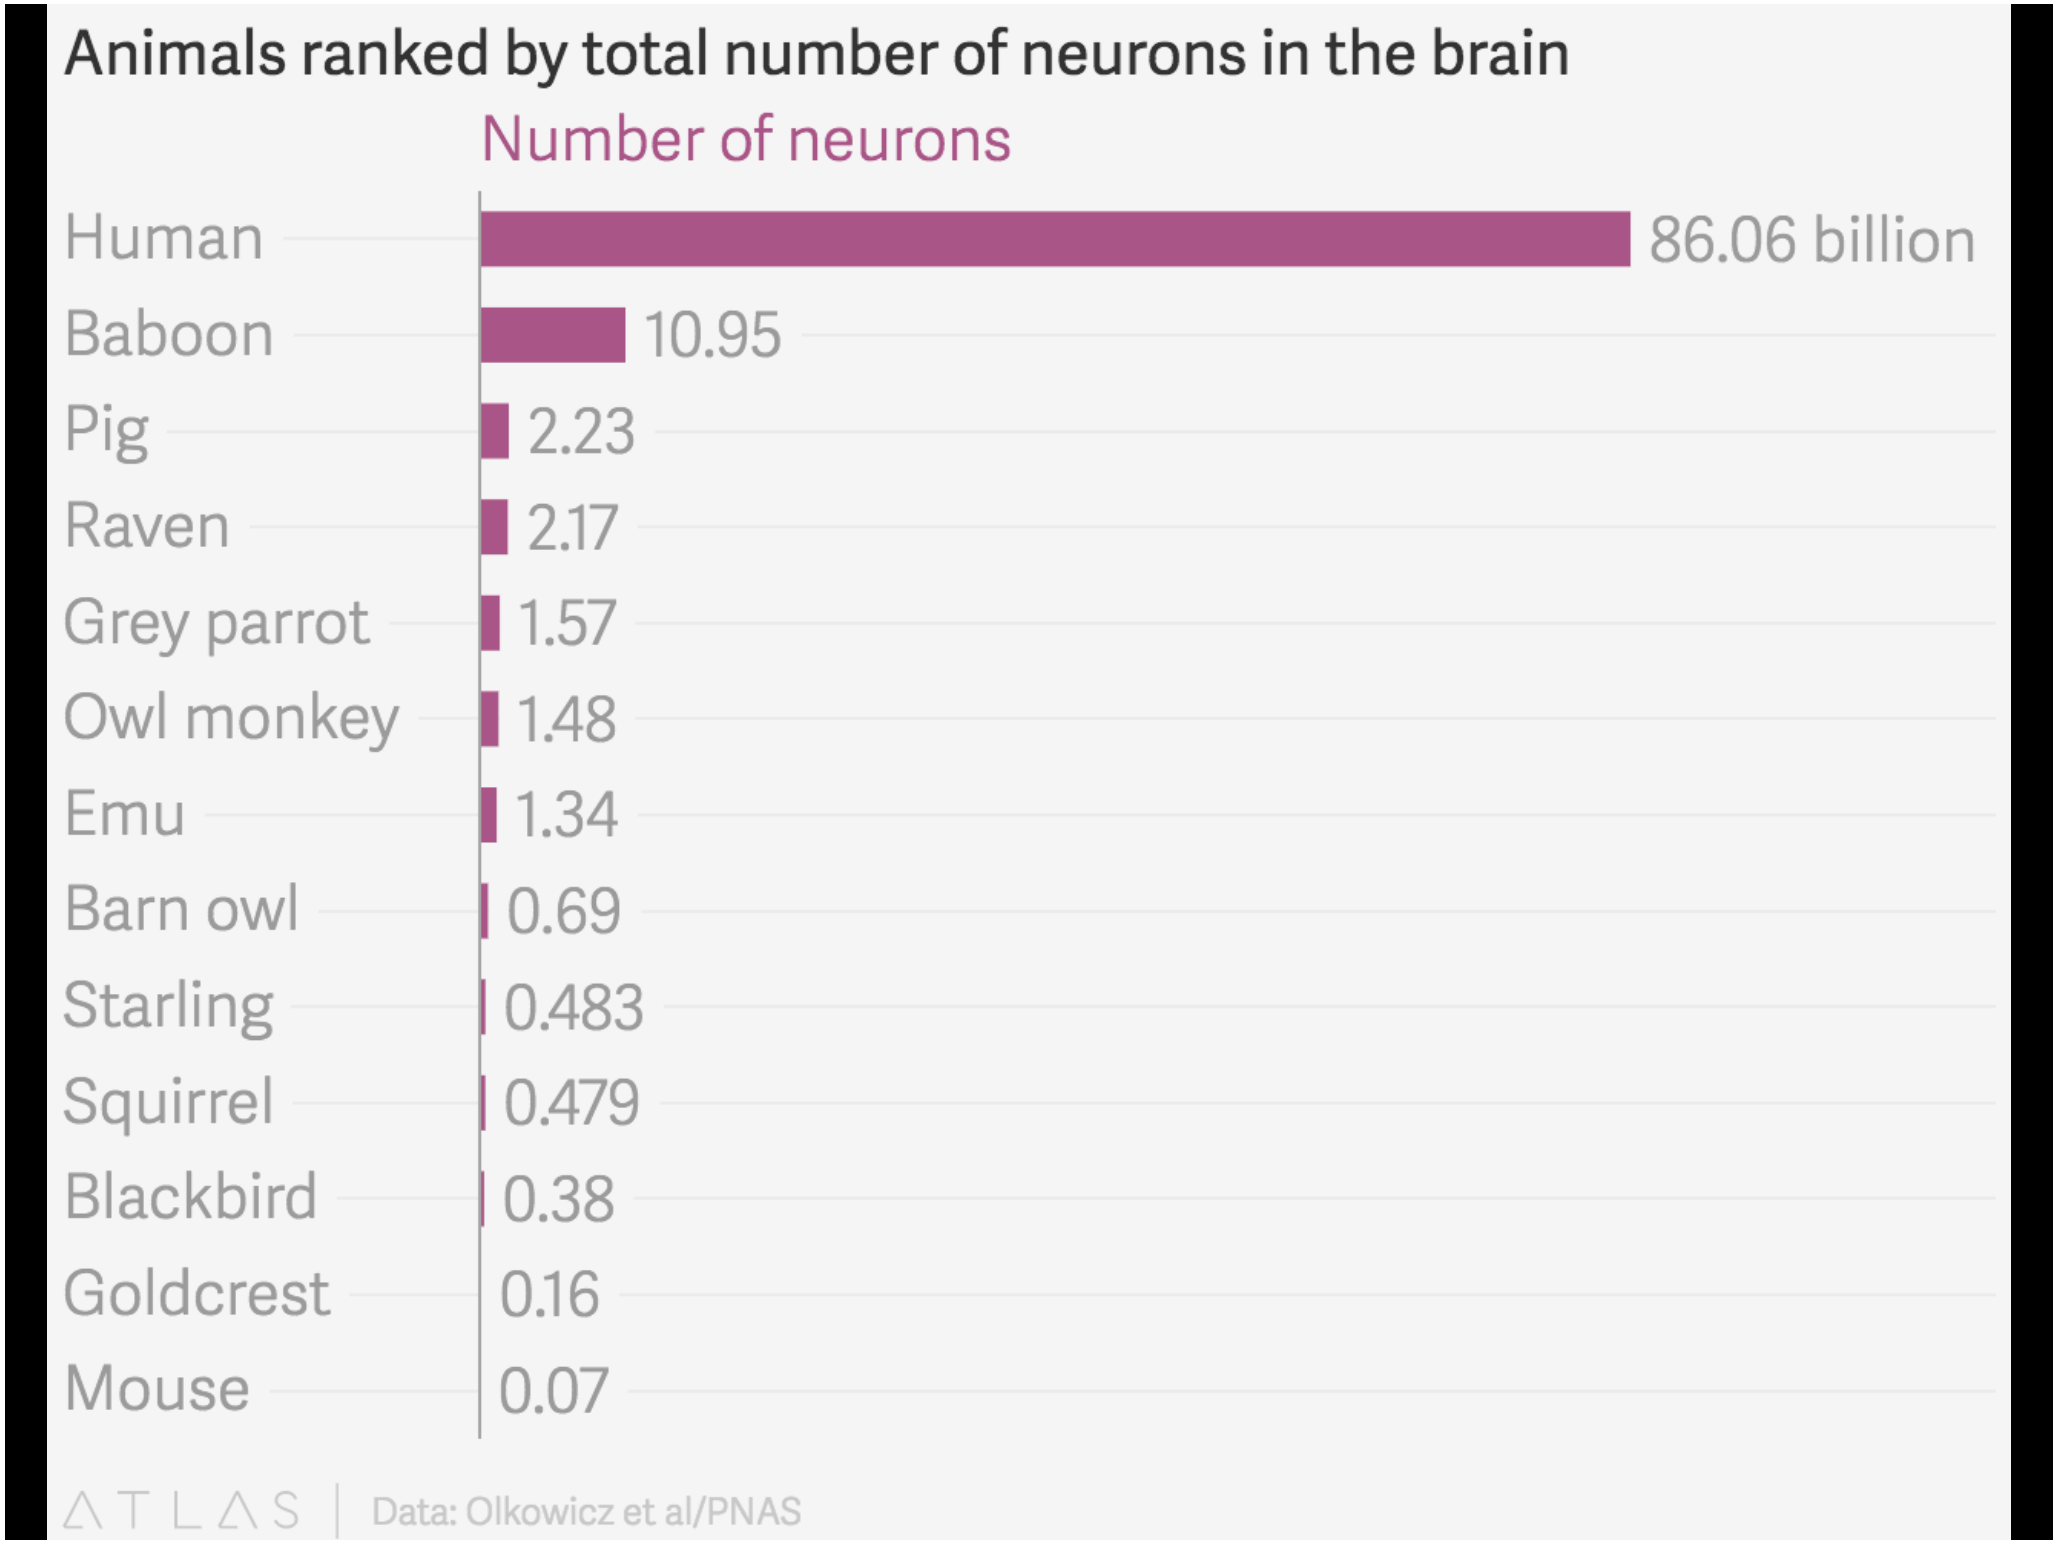
Animals ranked by total number of neurons in the brain
Number of neurons
Human
86.06 billion
Baboon
10.95
Pig
2.23
Raven
2.17
Grey parrot
1.57
Owl monkey
1.48
Emu
1.34
Barn owl
0.69
Starling
0.483
Squirrel
0.479
Blackbird
0.38
Goldcrest
0.16
Mouse
0.07
ATLAS
Data: Olkowicz et al/PNAS

So what happens when we learn?

Part I

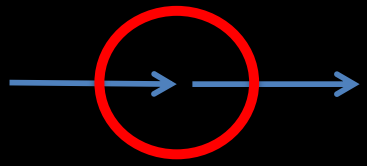

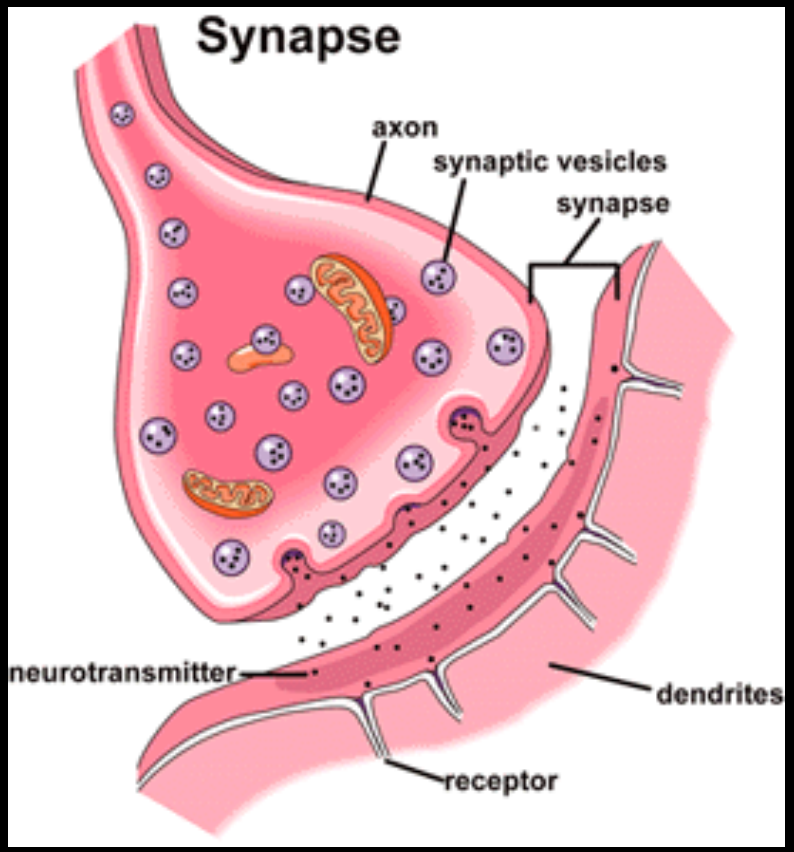

Increased neurotransmitter release
Increase receptors
Structural changes

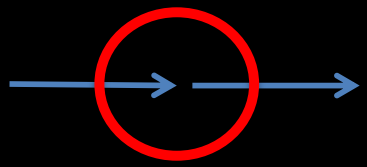

In practice, we talk about the strength of a connection in terms of a “weight” or a “value”.

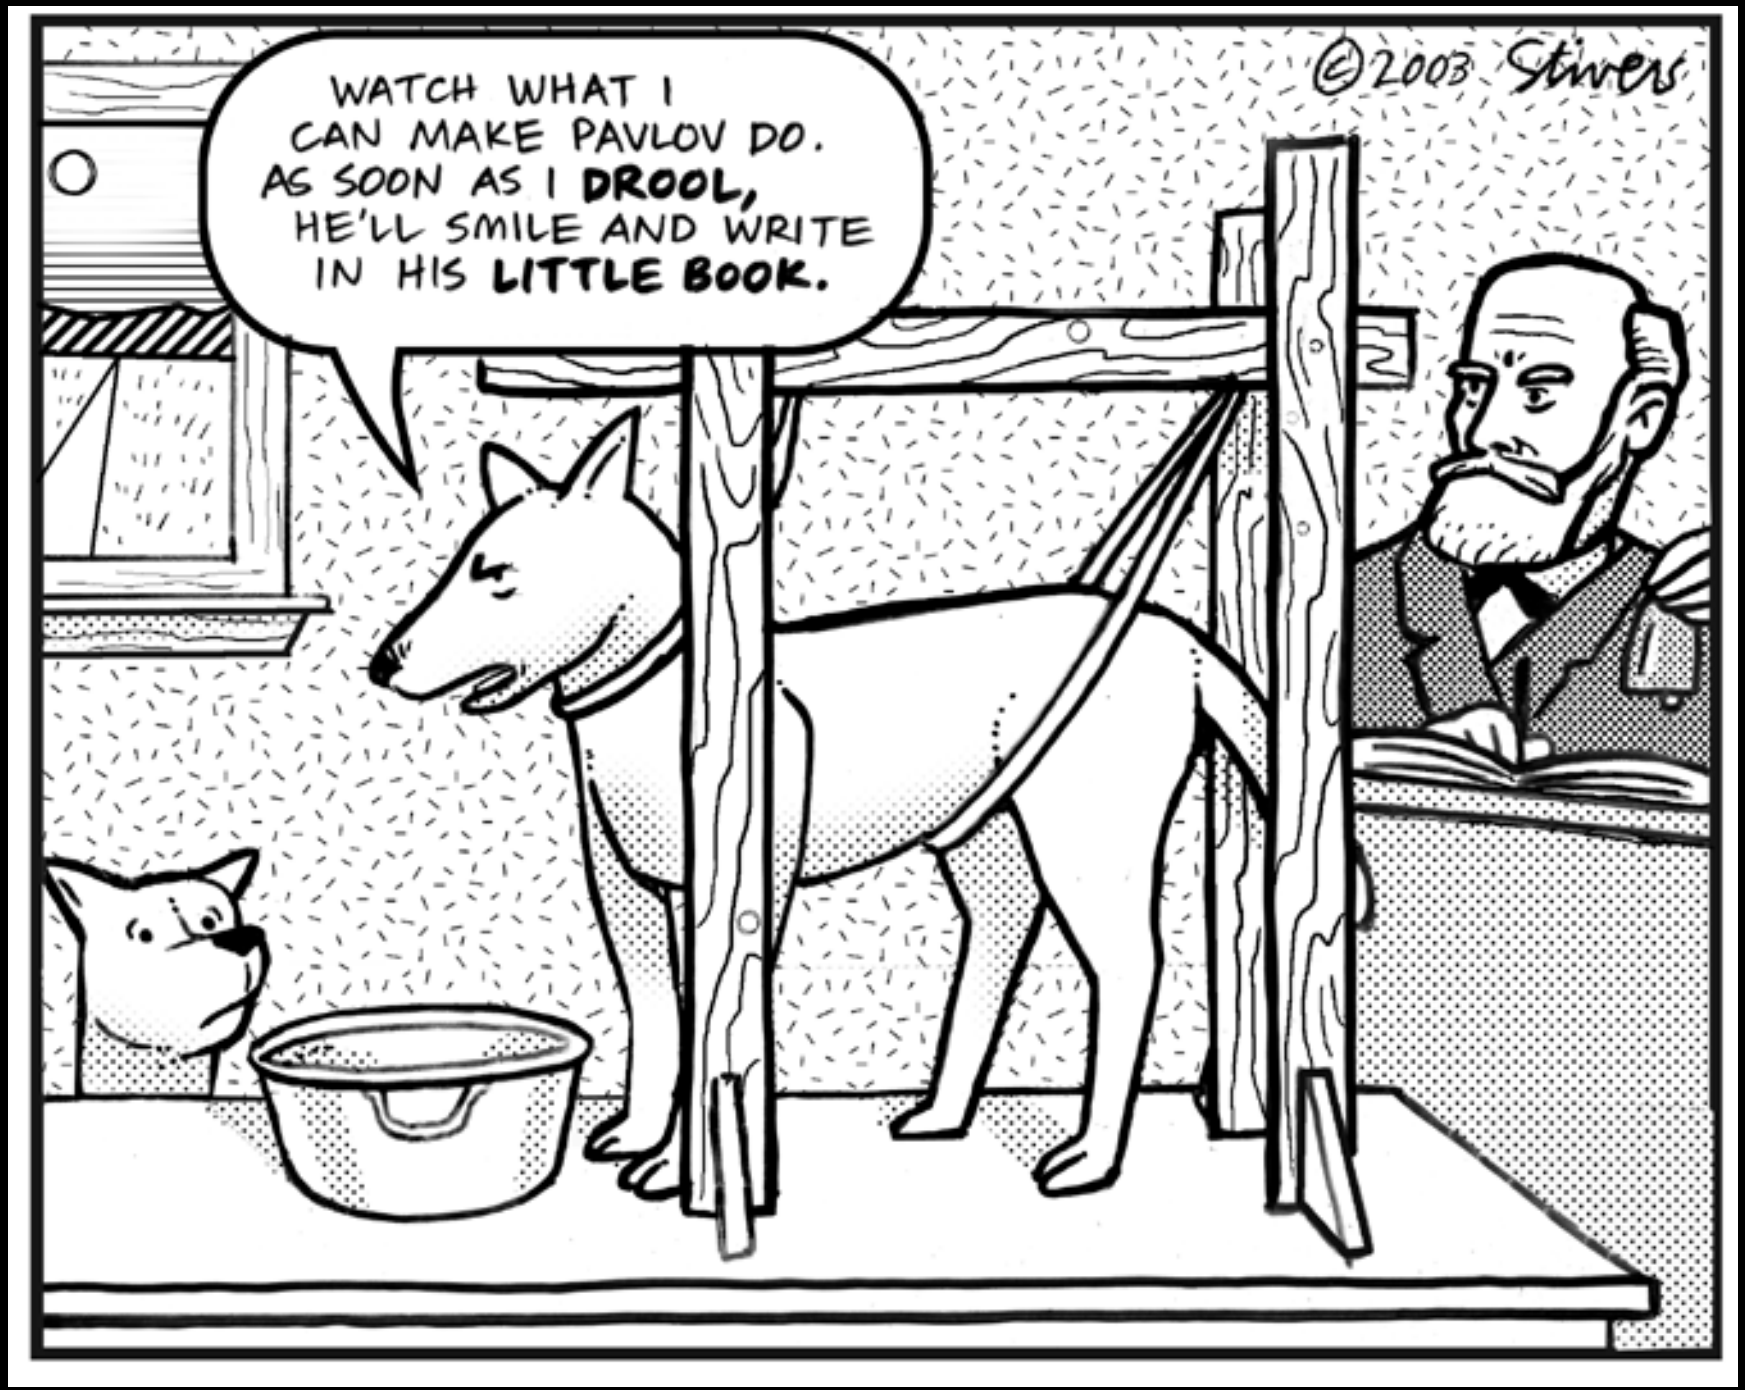
WATCH WHAT I CAN MAKE PAVLOV DO. AS SOON AS I DROOL, HE'LL SMILE AND WRITE IN HIS LITTLE BOOK.
©2003 Stivers

1. Before conditioning

response

Food

Salivation

Unconditioned stimulus

Unconditioned response

1. Before conditioning

Food

Unconditioned stimulus

response

Salivation

Unconditioned response

2. Before conditioning

Tuning fork

Neutral stimulus

response

No salivation

No conditioned response

1. Before conditioning

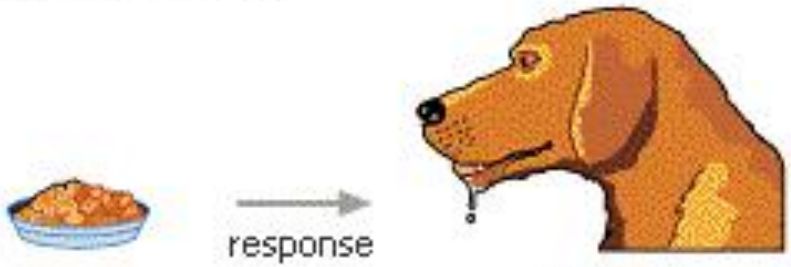

Food

response

Salivation

Unconditioned stimulus

Unconditioned response

2. Before conditioning

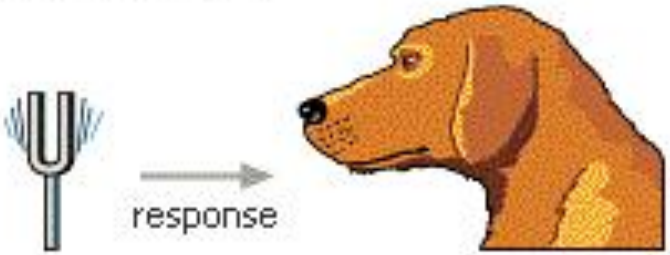

Tuning fork

response

No salivation

Neutral stimulus

No conditioned response

3. During conditioning

Tuning fork

Food

response

Salivation

Unconditioned response

1. Before conditioning

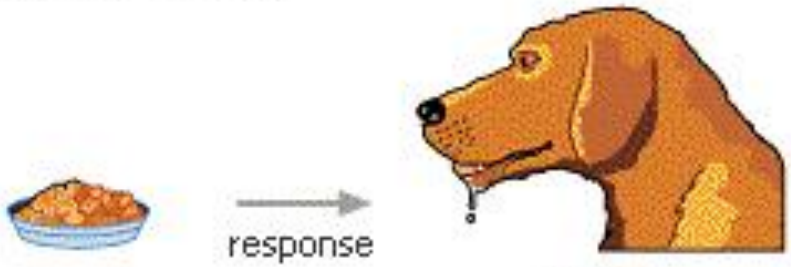

Food

response

Salivation

Unconditioned stimulus

Unconditioned response

2. Before conditioning

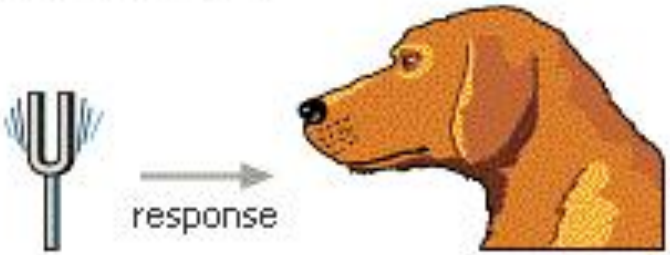

Tuning fork

response

No salivation

Neutral stimulus

No conditioned response

3. During conditioning

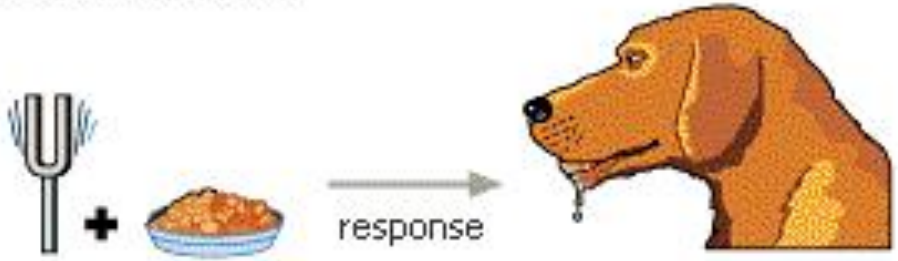

Tuning fork

Food

response

Salivation

Unconditioned response

4. After conditioning

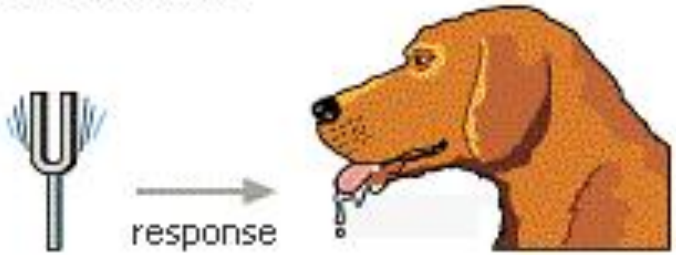

Tuning fork

response

Salivation

Conditioned stimulus

Conditioned response

Rescorla - Wagner

Expectancy

Expectancy

Expectancy

Expectancy

Expectancy

Prediction Error

When the actual outcome is different than the expected outcome

Value

V_{reward}

$V_{punishment}$

V_{tone}

States

Current State

Previous State

Current State

Prediction Error

The difference in VALUE between the current state and preceding state

$$PE = (V_{reward} - V_{cue})$$

$$PE = (V_{current\ state} - V_{preceding\ state})$$

Learning IS ALWAYS a two step process.

At each point in time we:

1) Calculate a prediction error
2) Update the previous value

$$V_{\text{cue new}} = V_{\text{cue old}} + PE$$

$V_{tone} = 0$

$V_{reward} = 100$

$PE = (100 - 0)$

$PE = 100$

$$V_{\text{cue new}} = V_{\text{cue old}} + PE$$

$$V_{\text{cue new}} = 0 + 100$$

$V_{tone} = 100$

$V_{reward} = 100$

$PE = (100 - 100)$

$PE = 0$

$$V_{cue\ new} = V_{cue\ old} + PE$$

$$V_{cue\ new} = 100 + 0$$

Learning Rates

$$V_{cue\ new} = V_{cue\ old} + PE * \mathbf{LR}$$

$$V_{cue\ new} = 0 + 100 * \mathbf{0.2}$$

$$V_{cue\ new} = 0 + 20$$

LR = 0.2

Trial	V_{cue}	V_{reward}	PE	PE x LR
1	0	100	100	20
2	20	100	80	16
3	36	100	64	12.8
4	48.8	100	51.2	10.24
5	59.04	100	40.96	8.192

V_{reward}

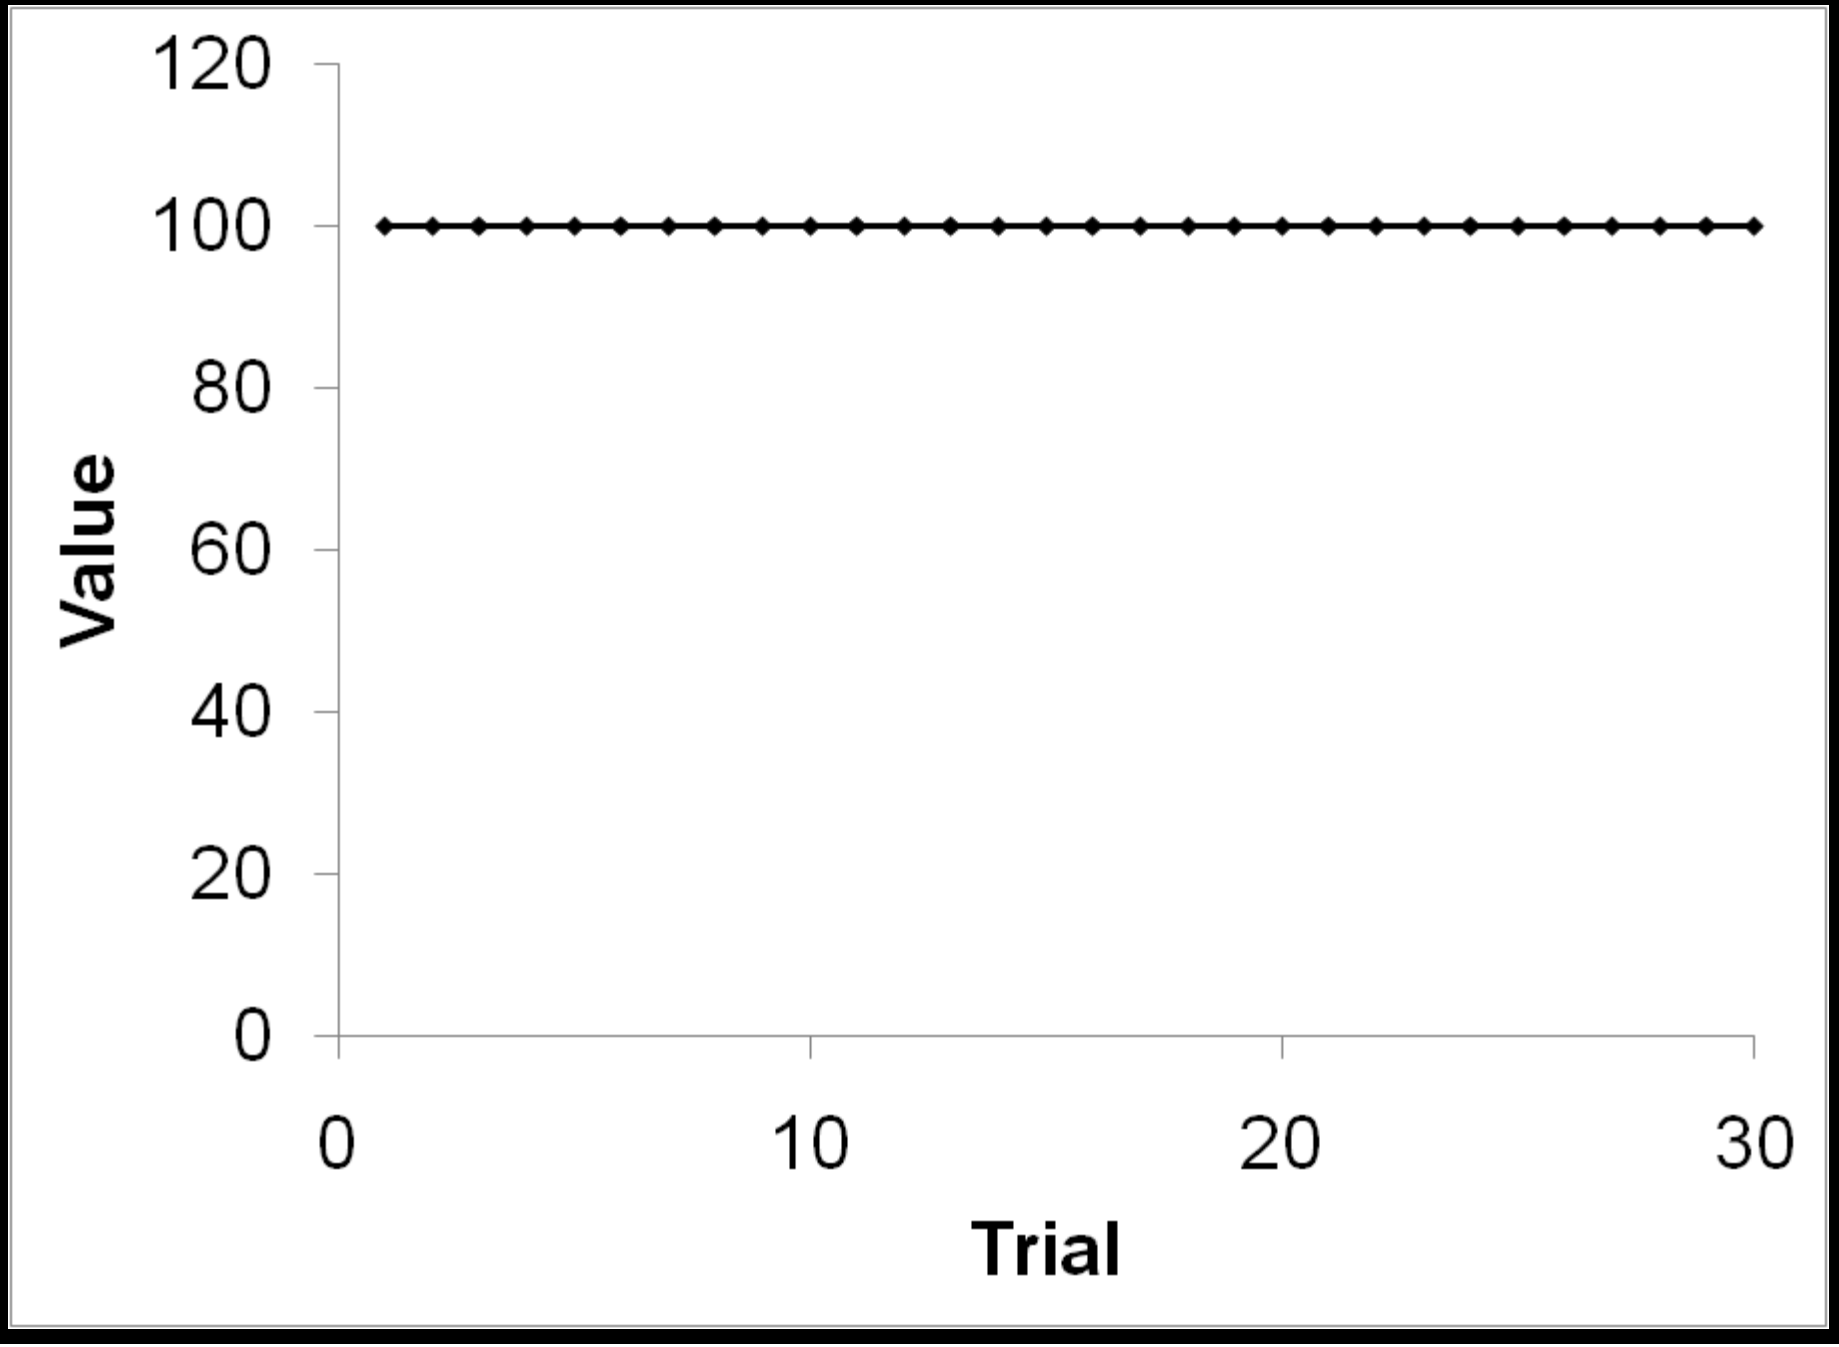

V_{cue}

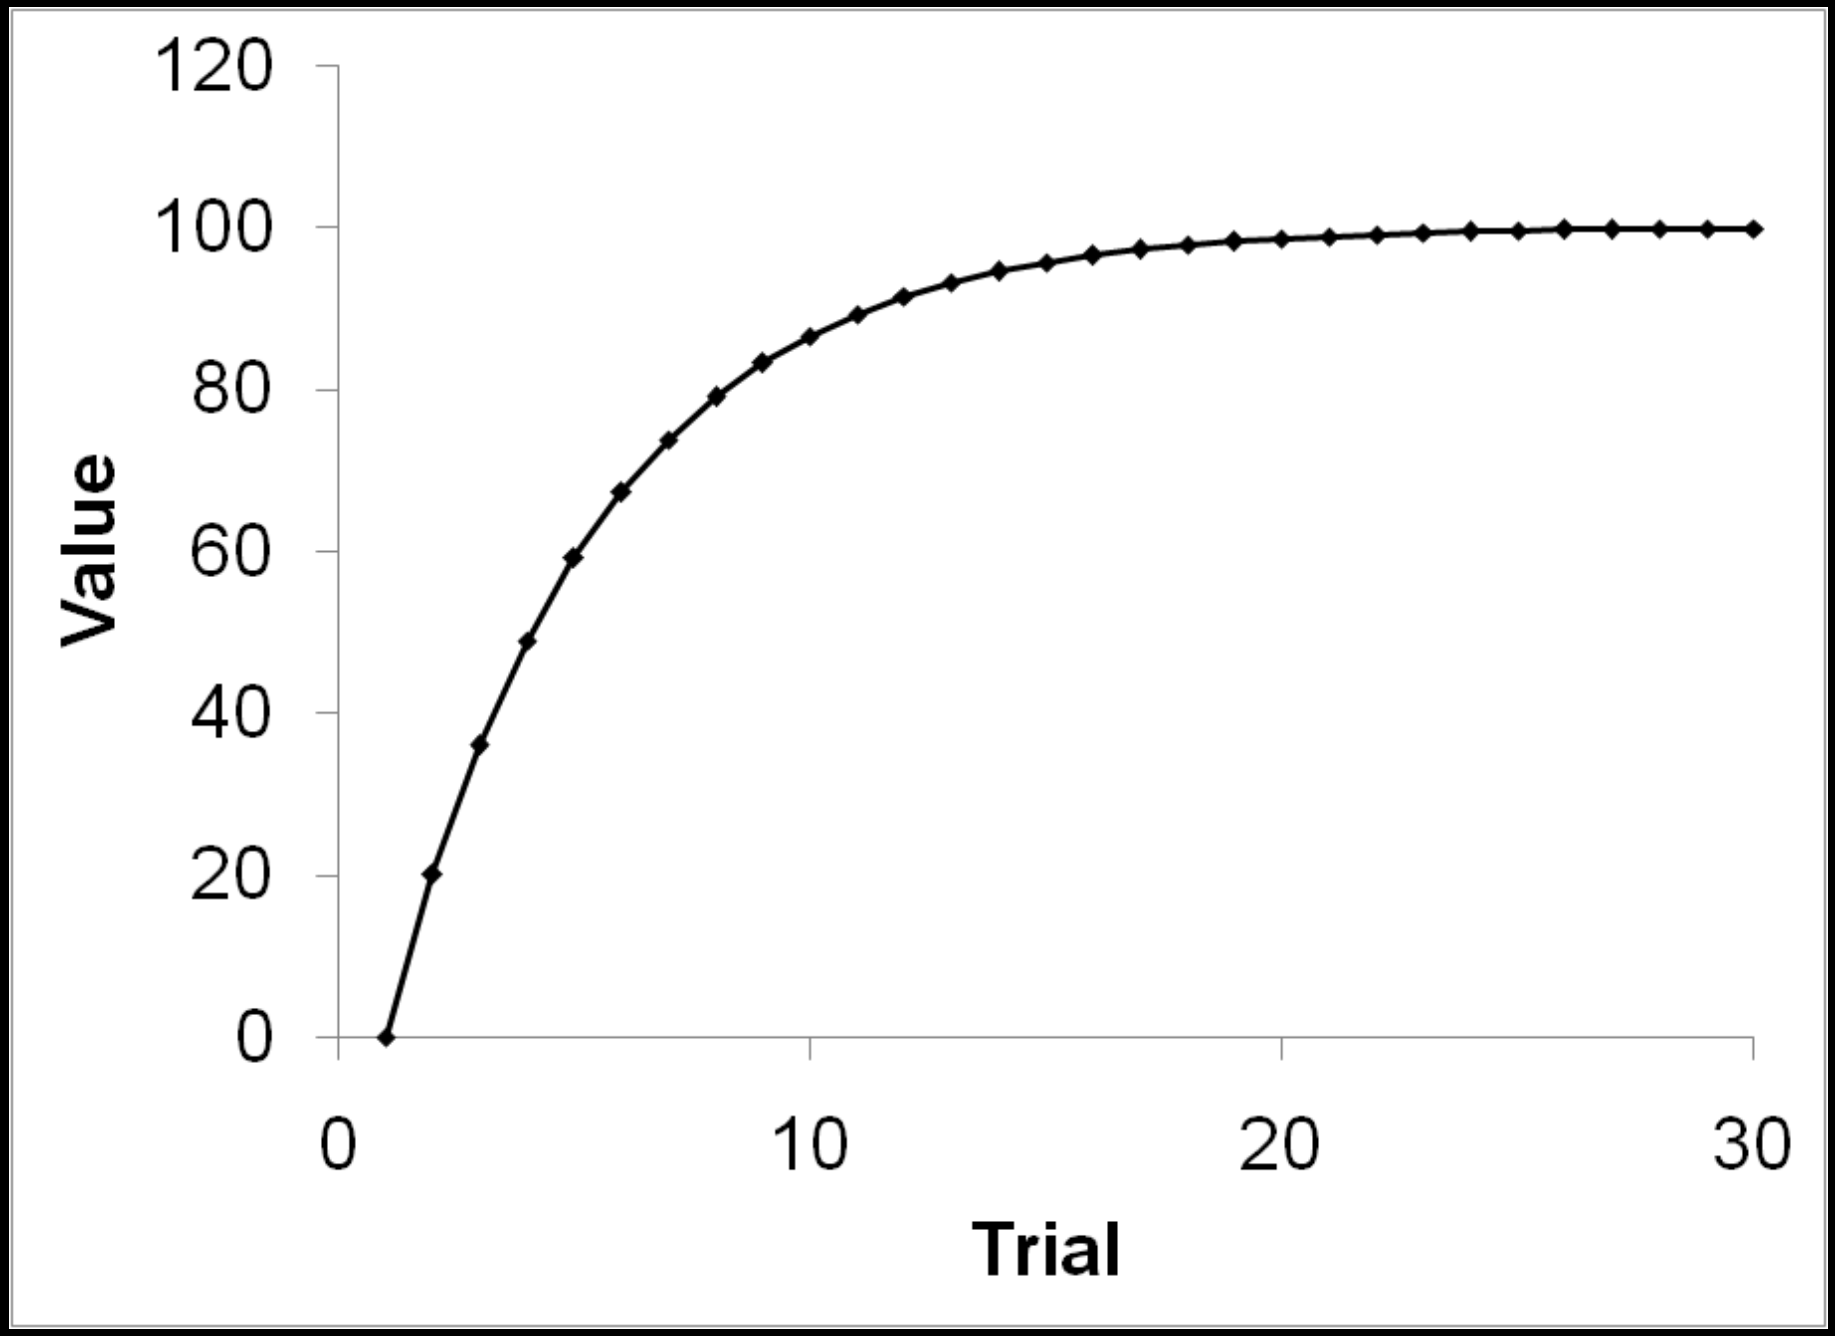

Prediction Error x Learning Rate

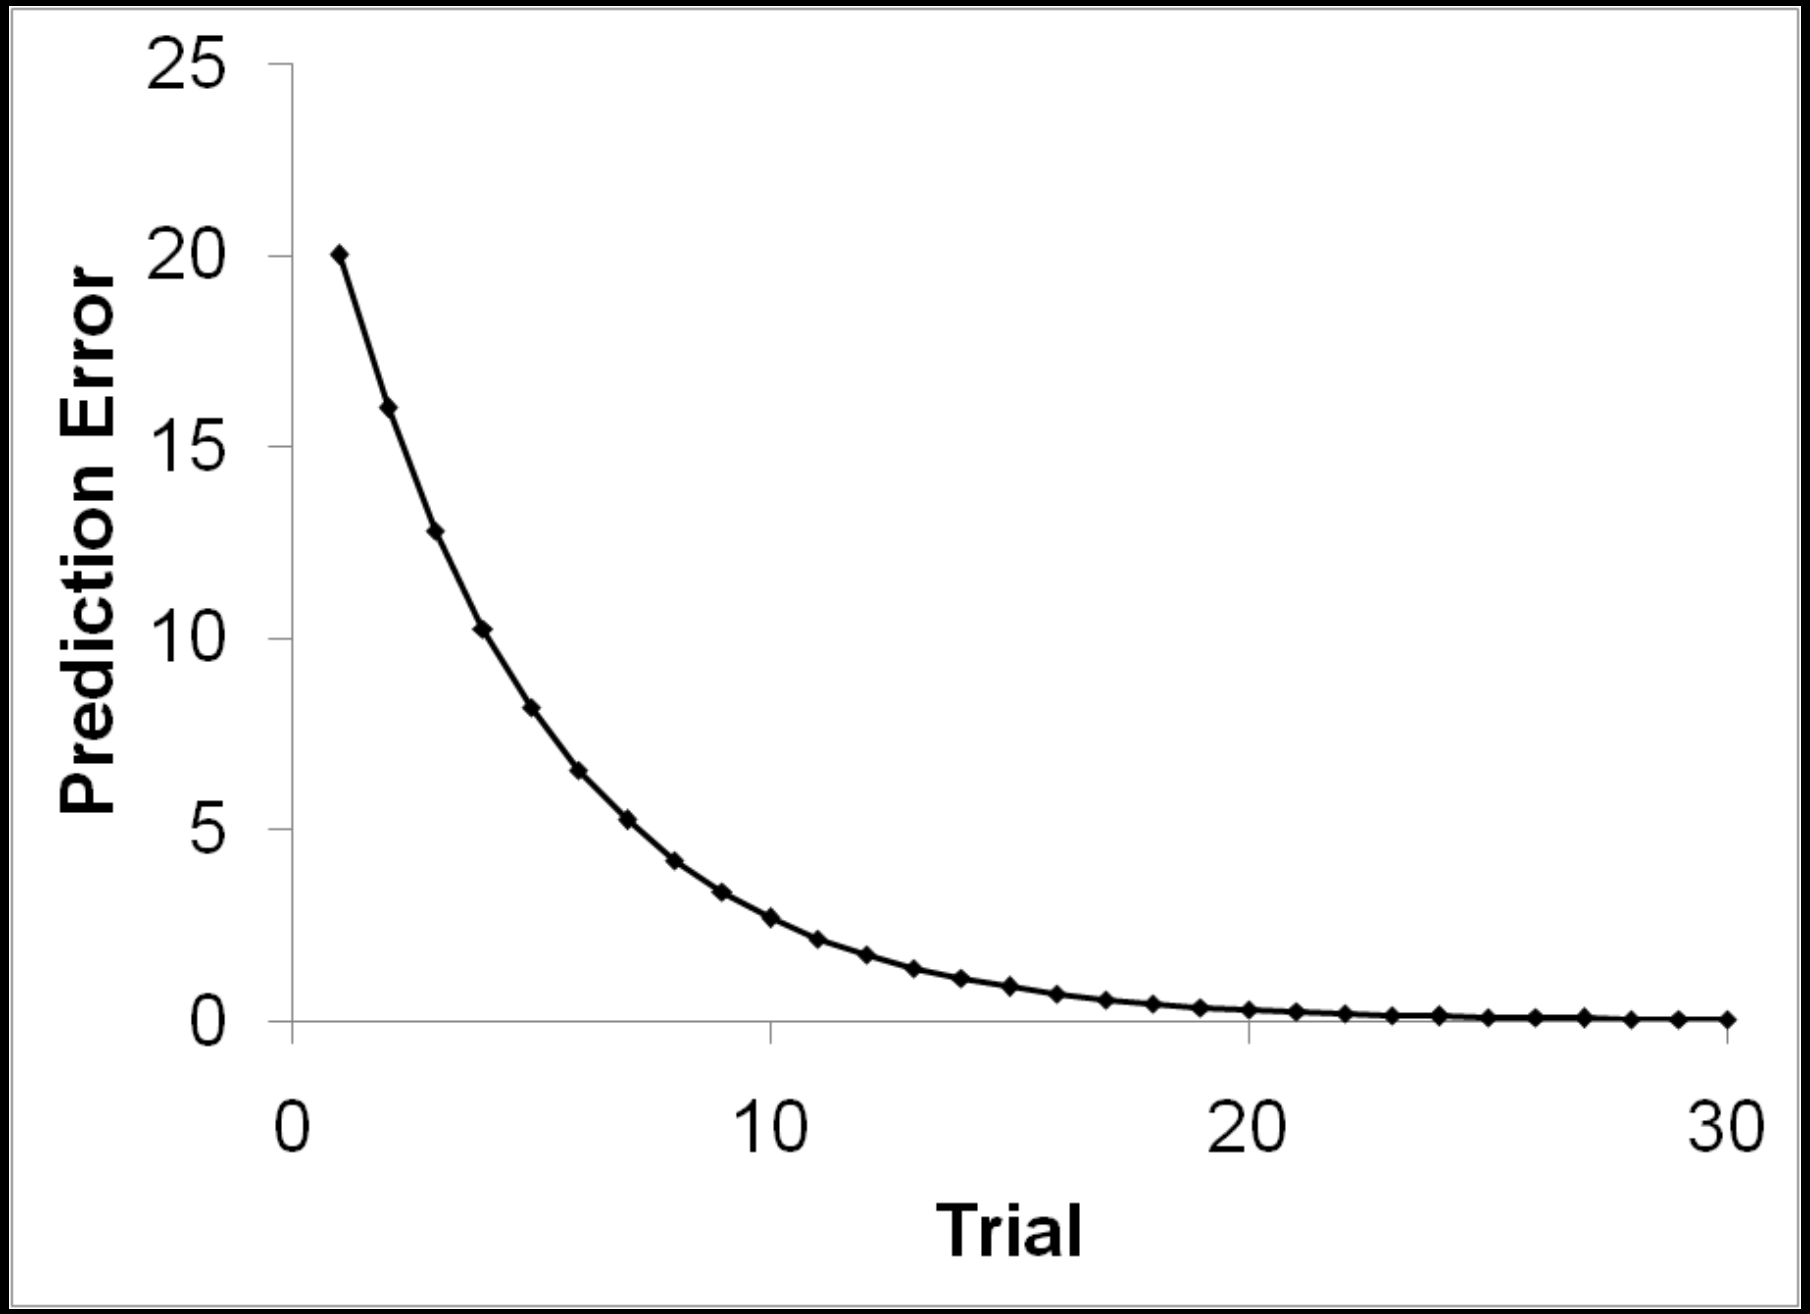

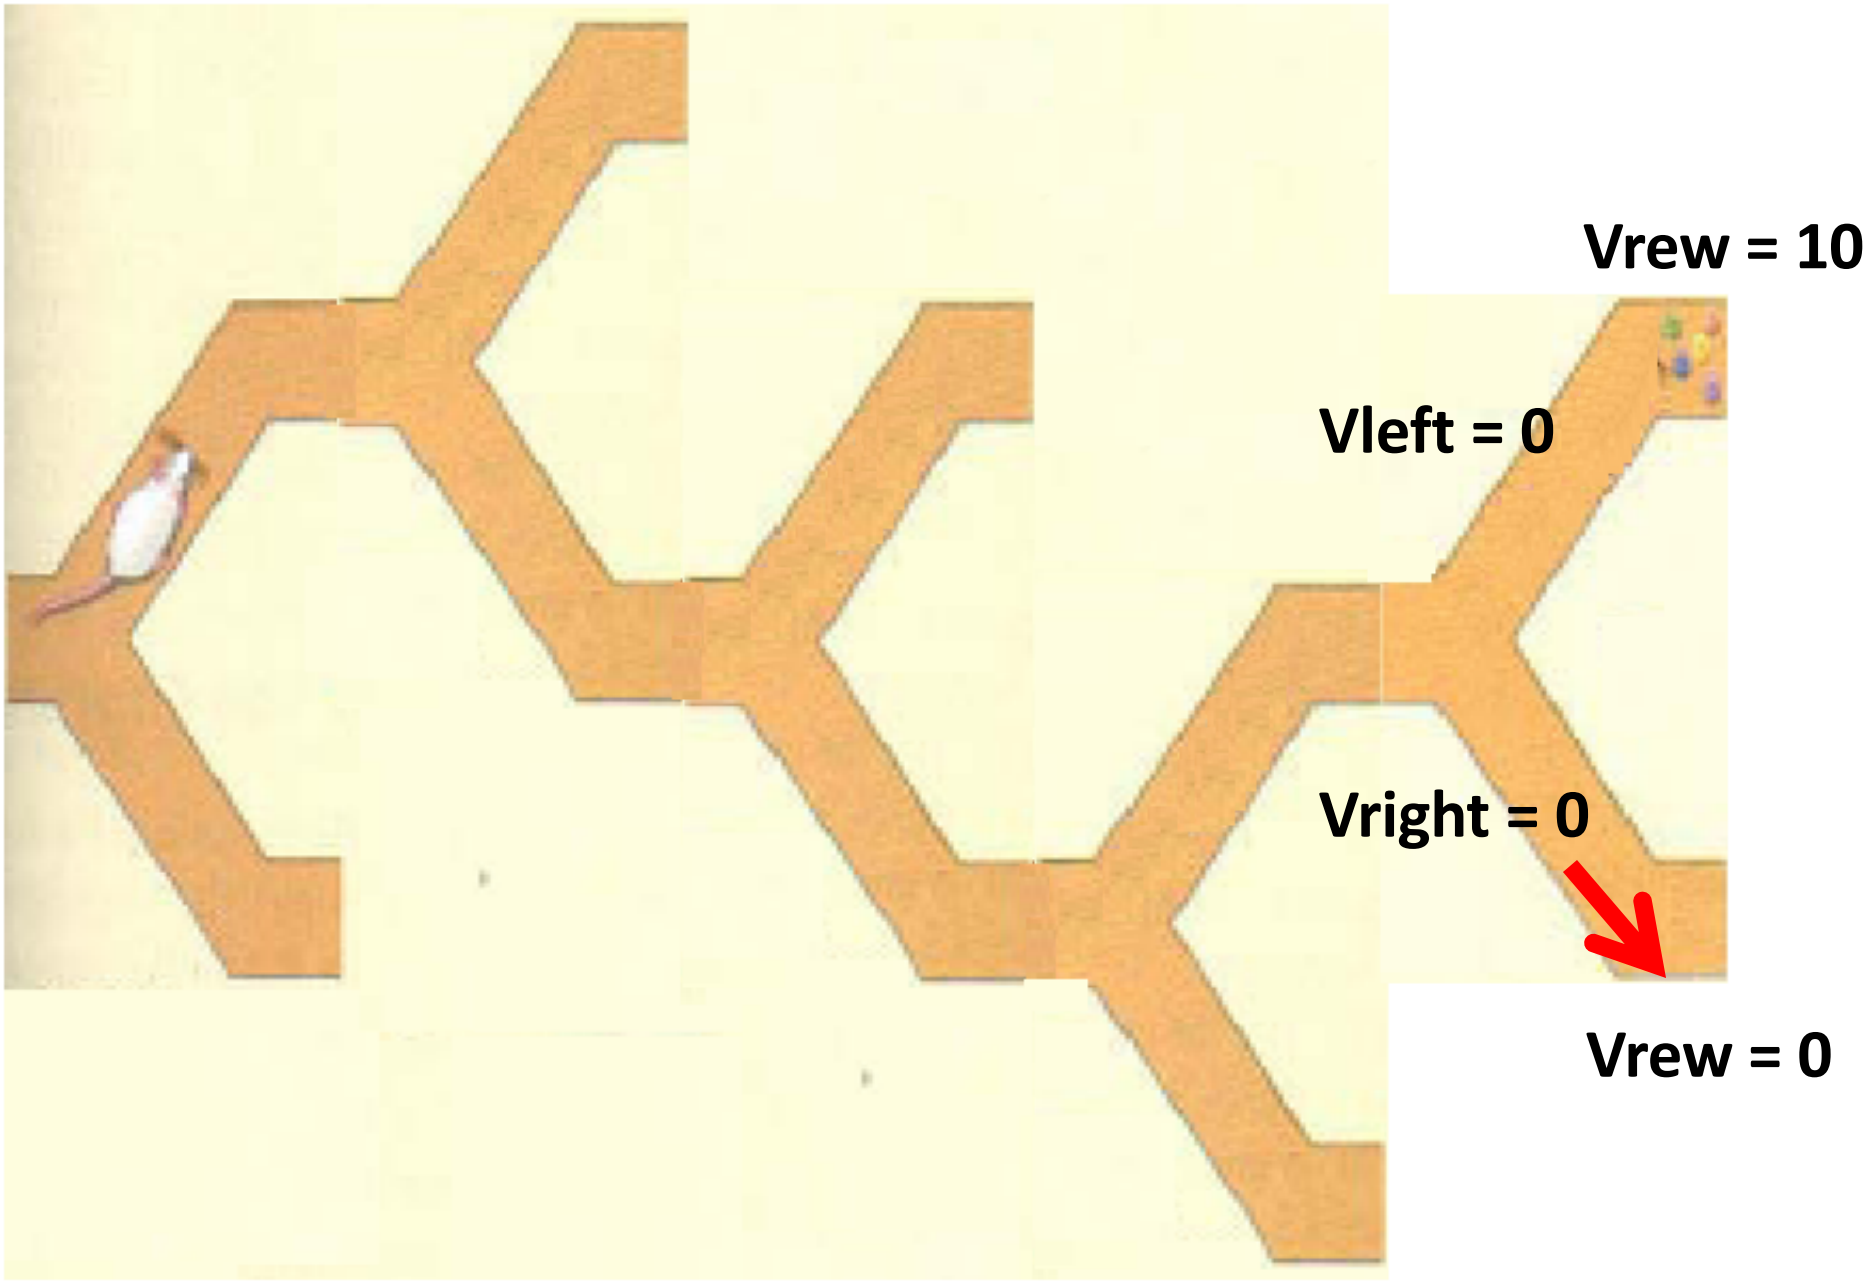
Vrew = 10
Vleft = 0
Vright = 0
Vrew = 0

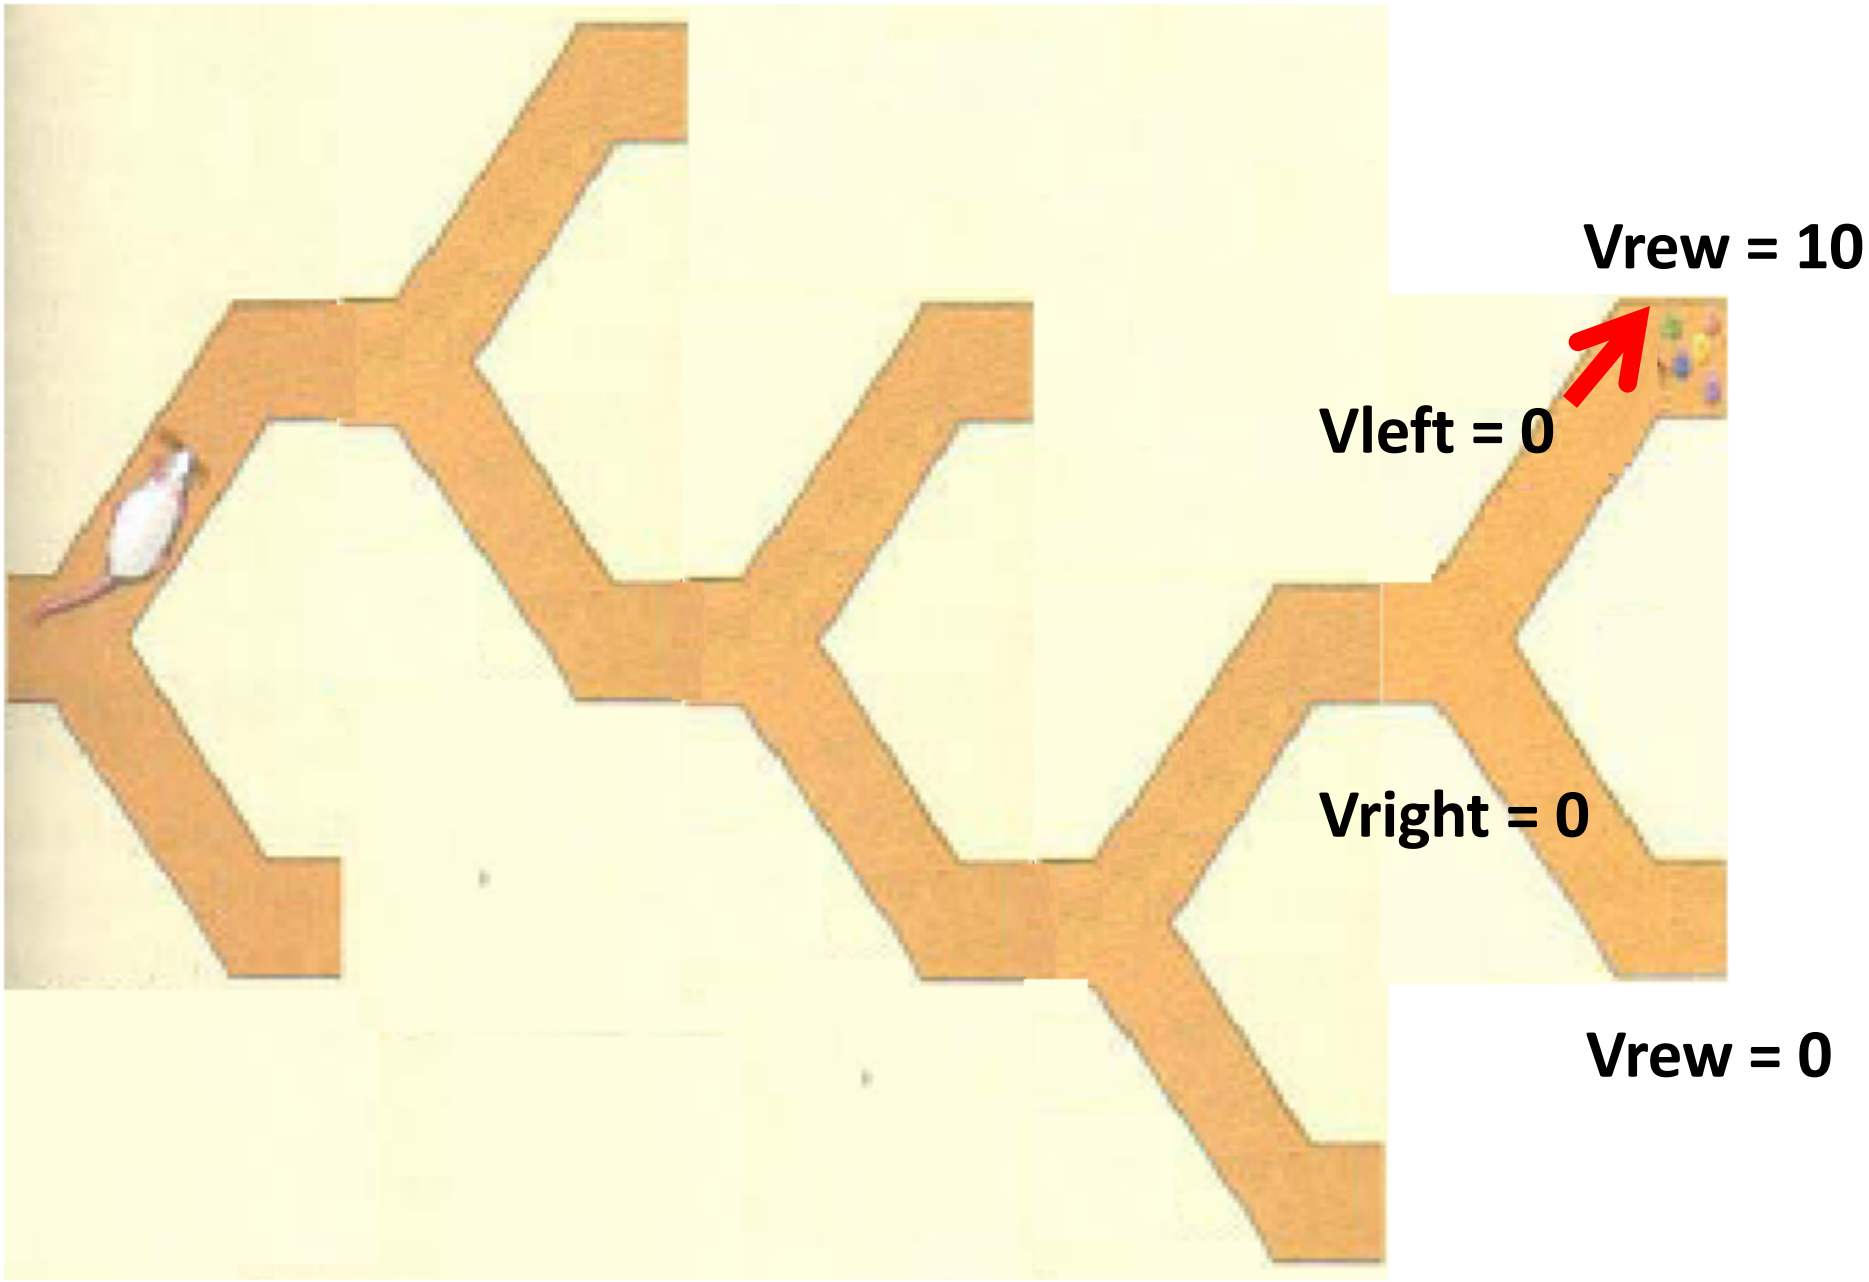
Vrew = 10
Vleft = 0
Vright = 0
Vrew = 0

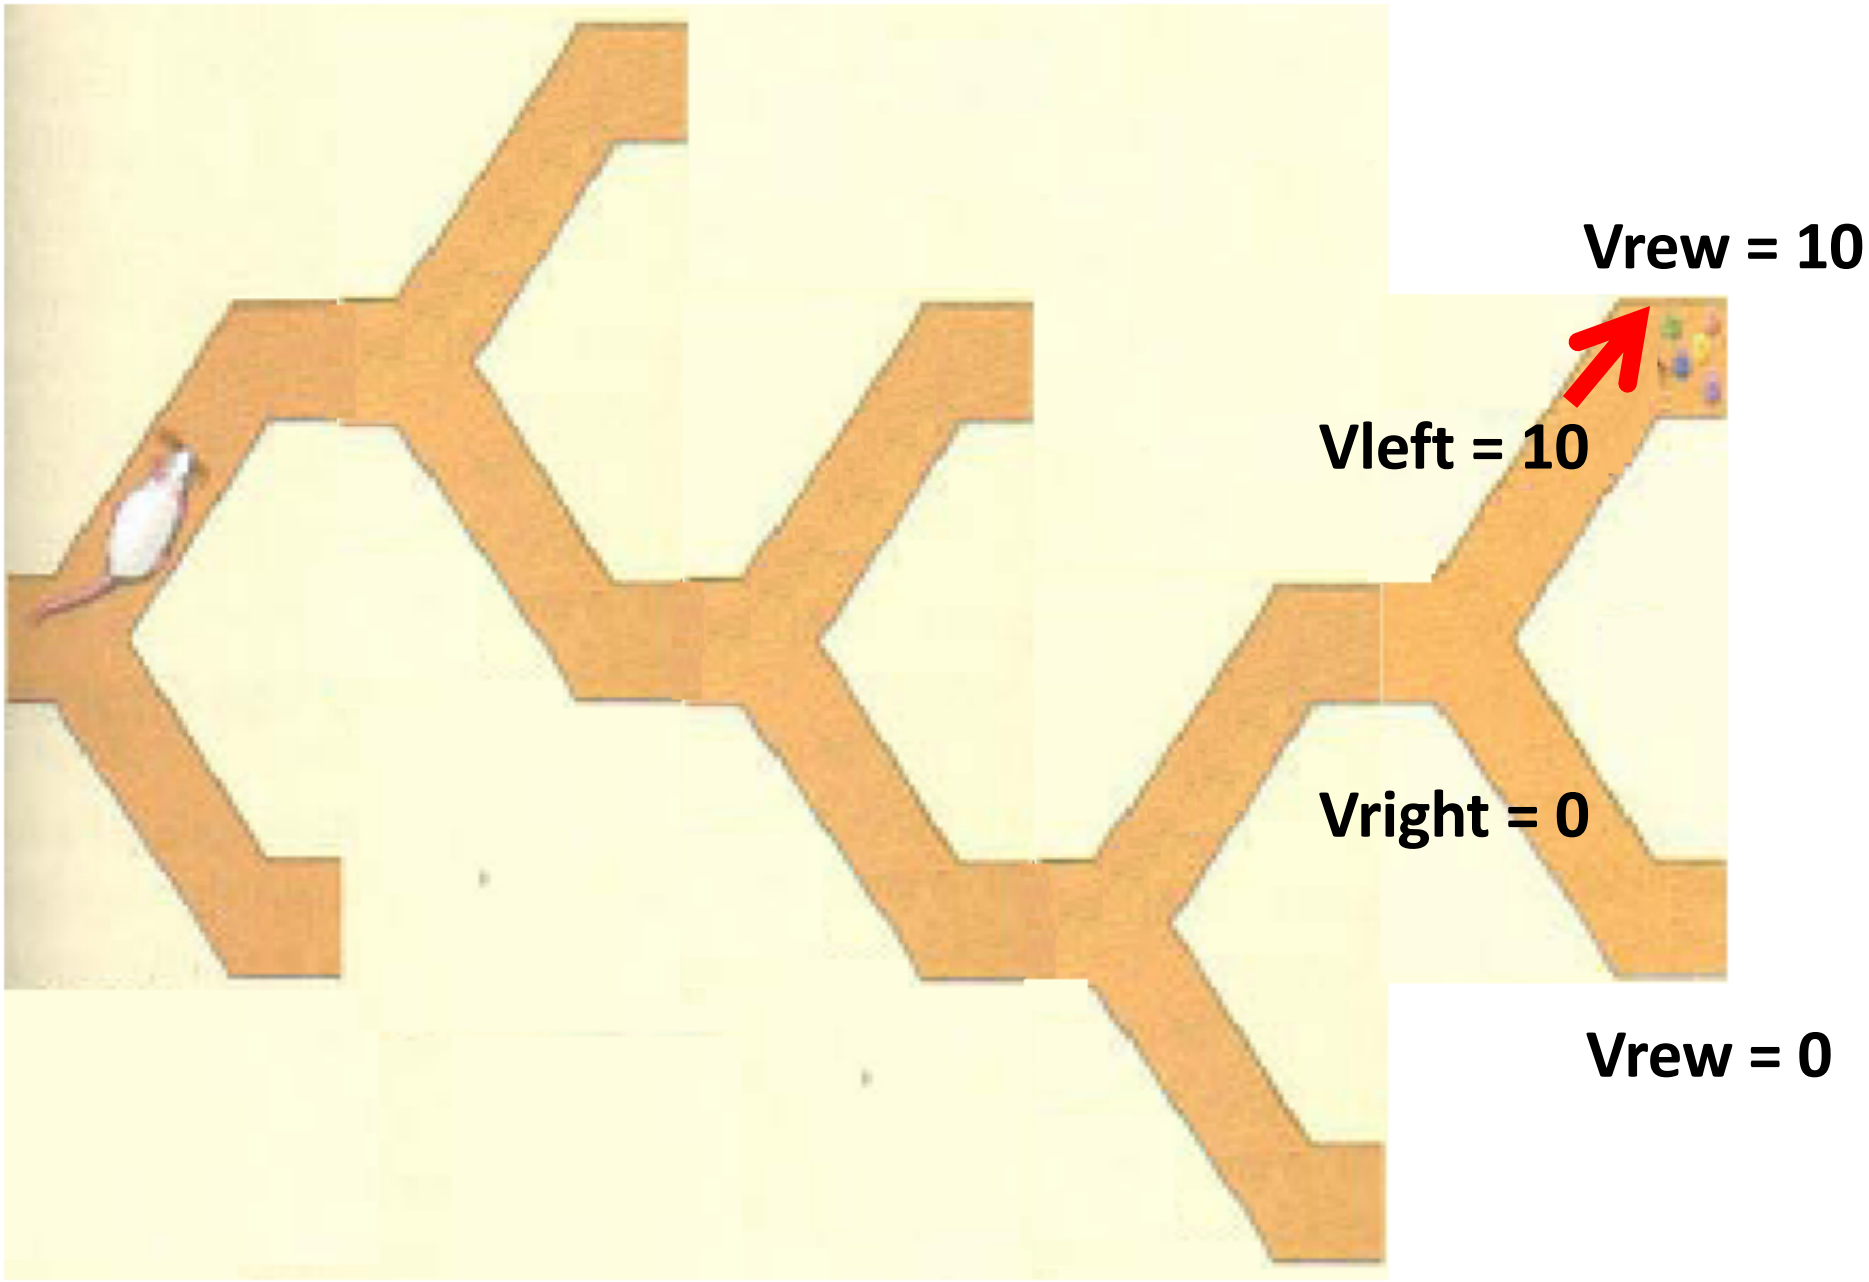
Vrew = 10
Vleft = 10
Vright = 0
Vrew = 0

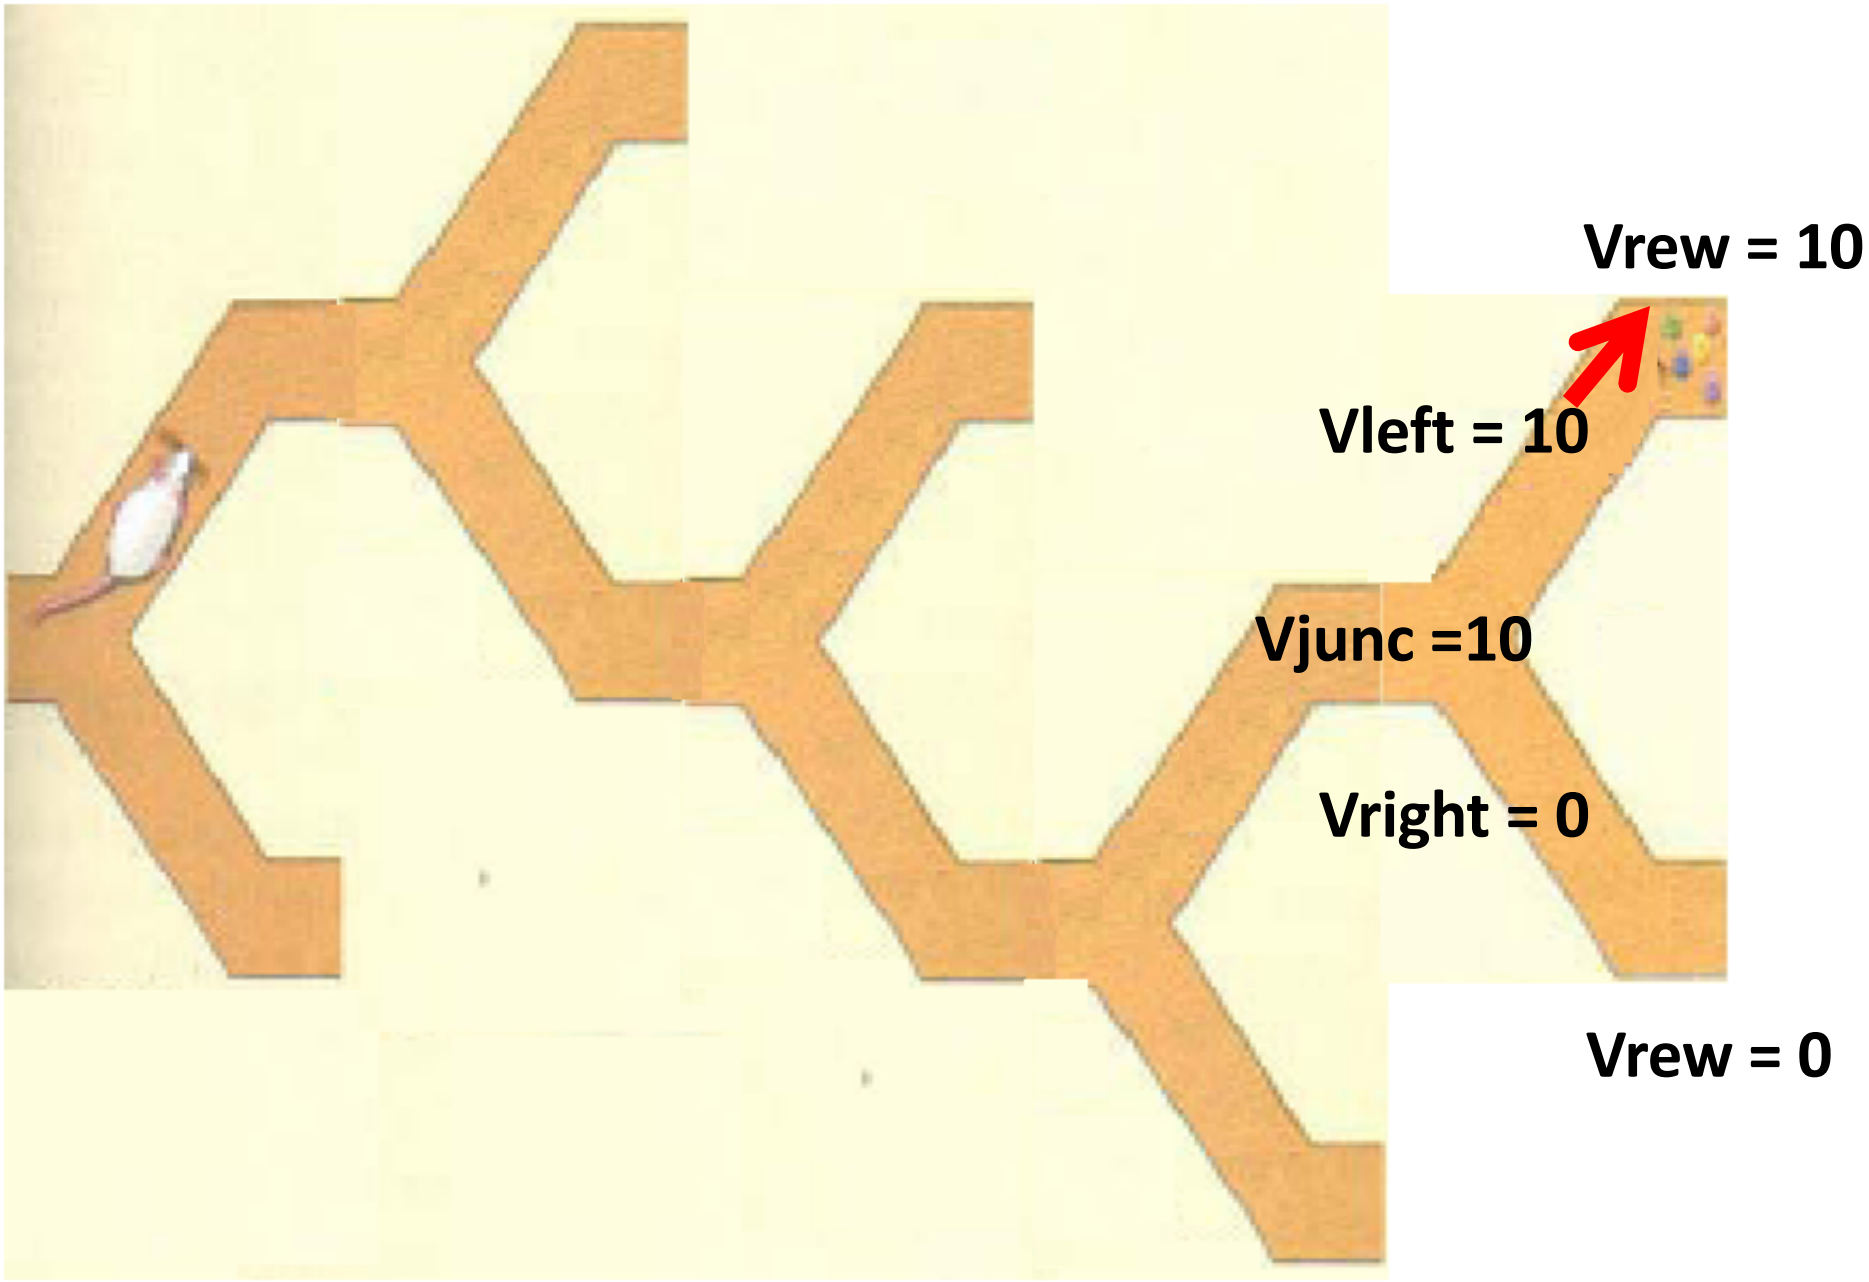
Vrew = 10
Vleft = 10
Vjunc =10
Vright = 0
Vrew = 0

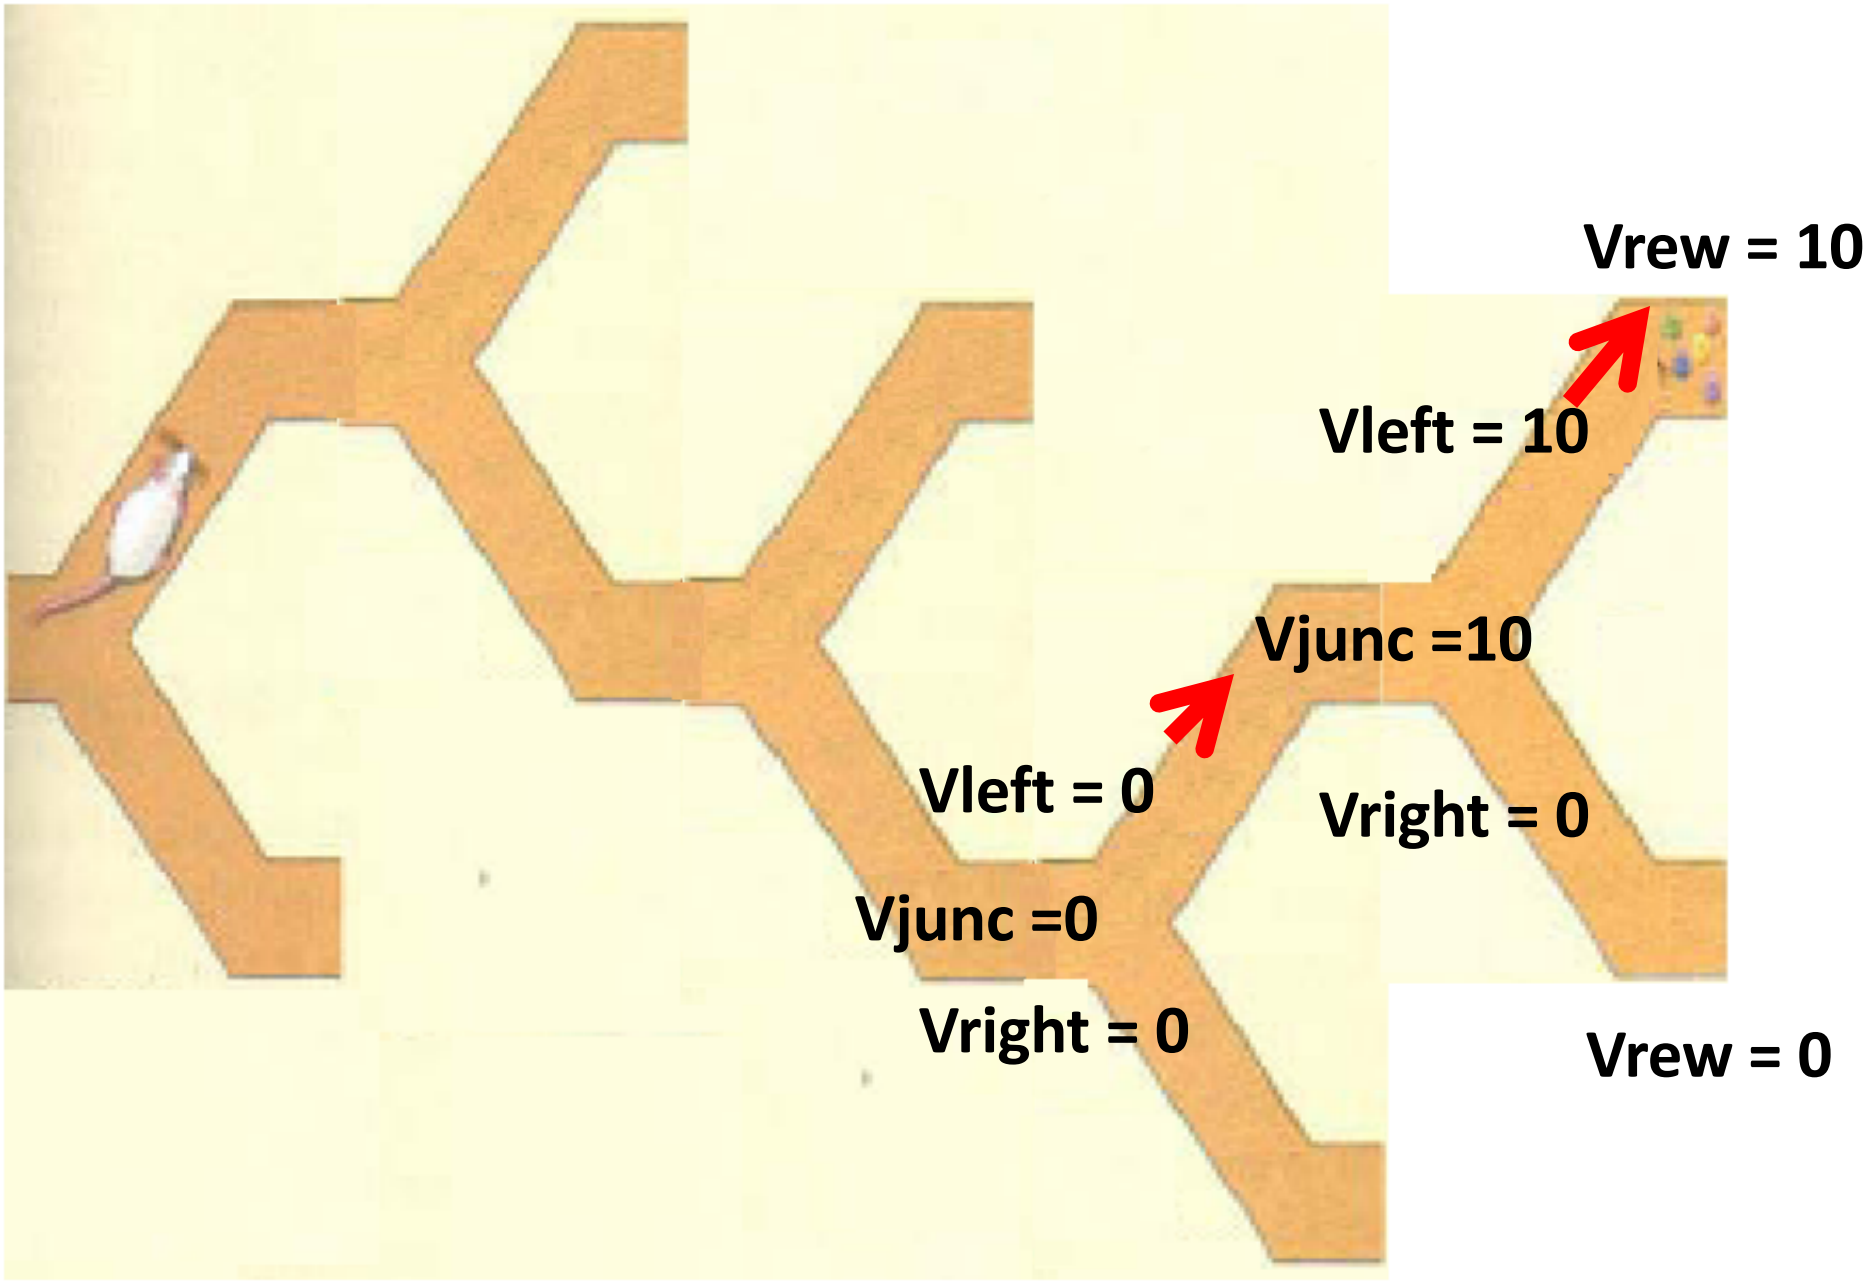
Vrew = 10
Vleft = 10
Vjunc =10
Vleft = 0
Vright = 0
Vjunc =0
Vright = 0
Vrew = 0

Another Example…

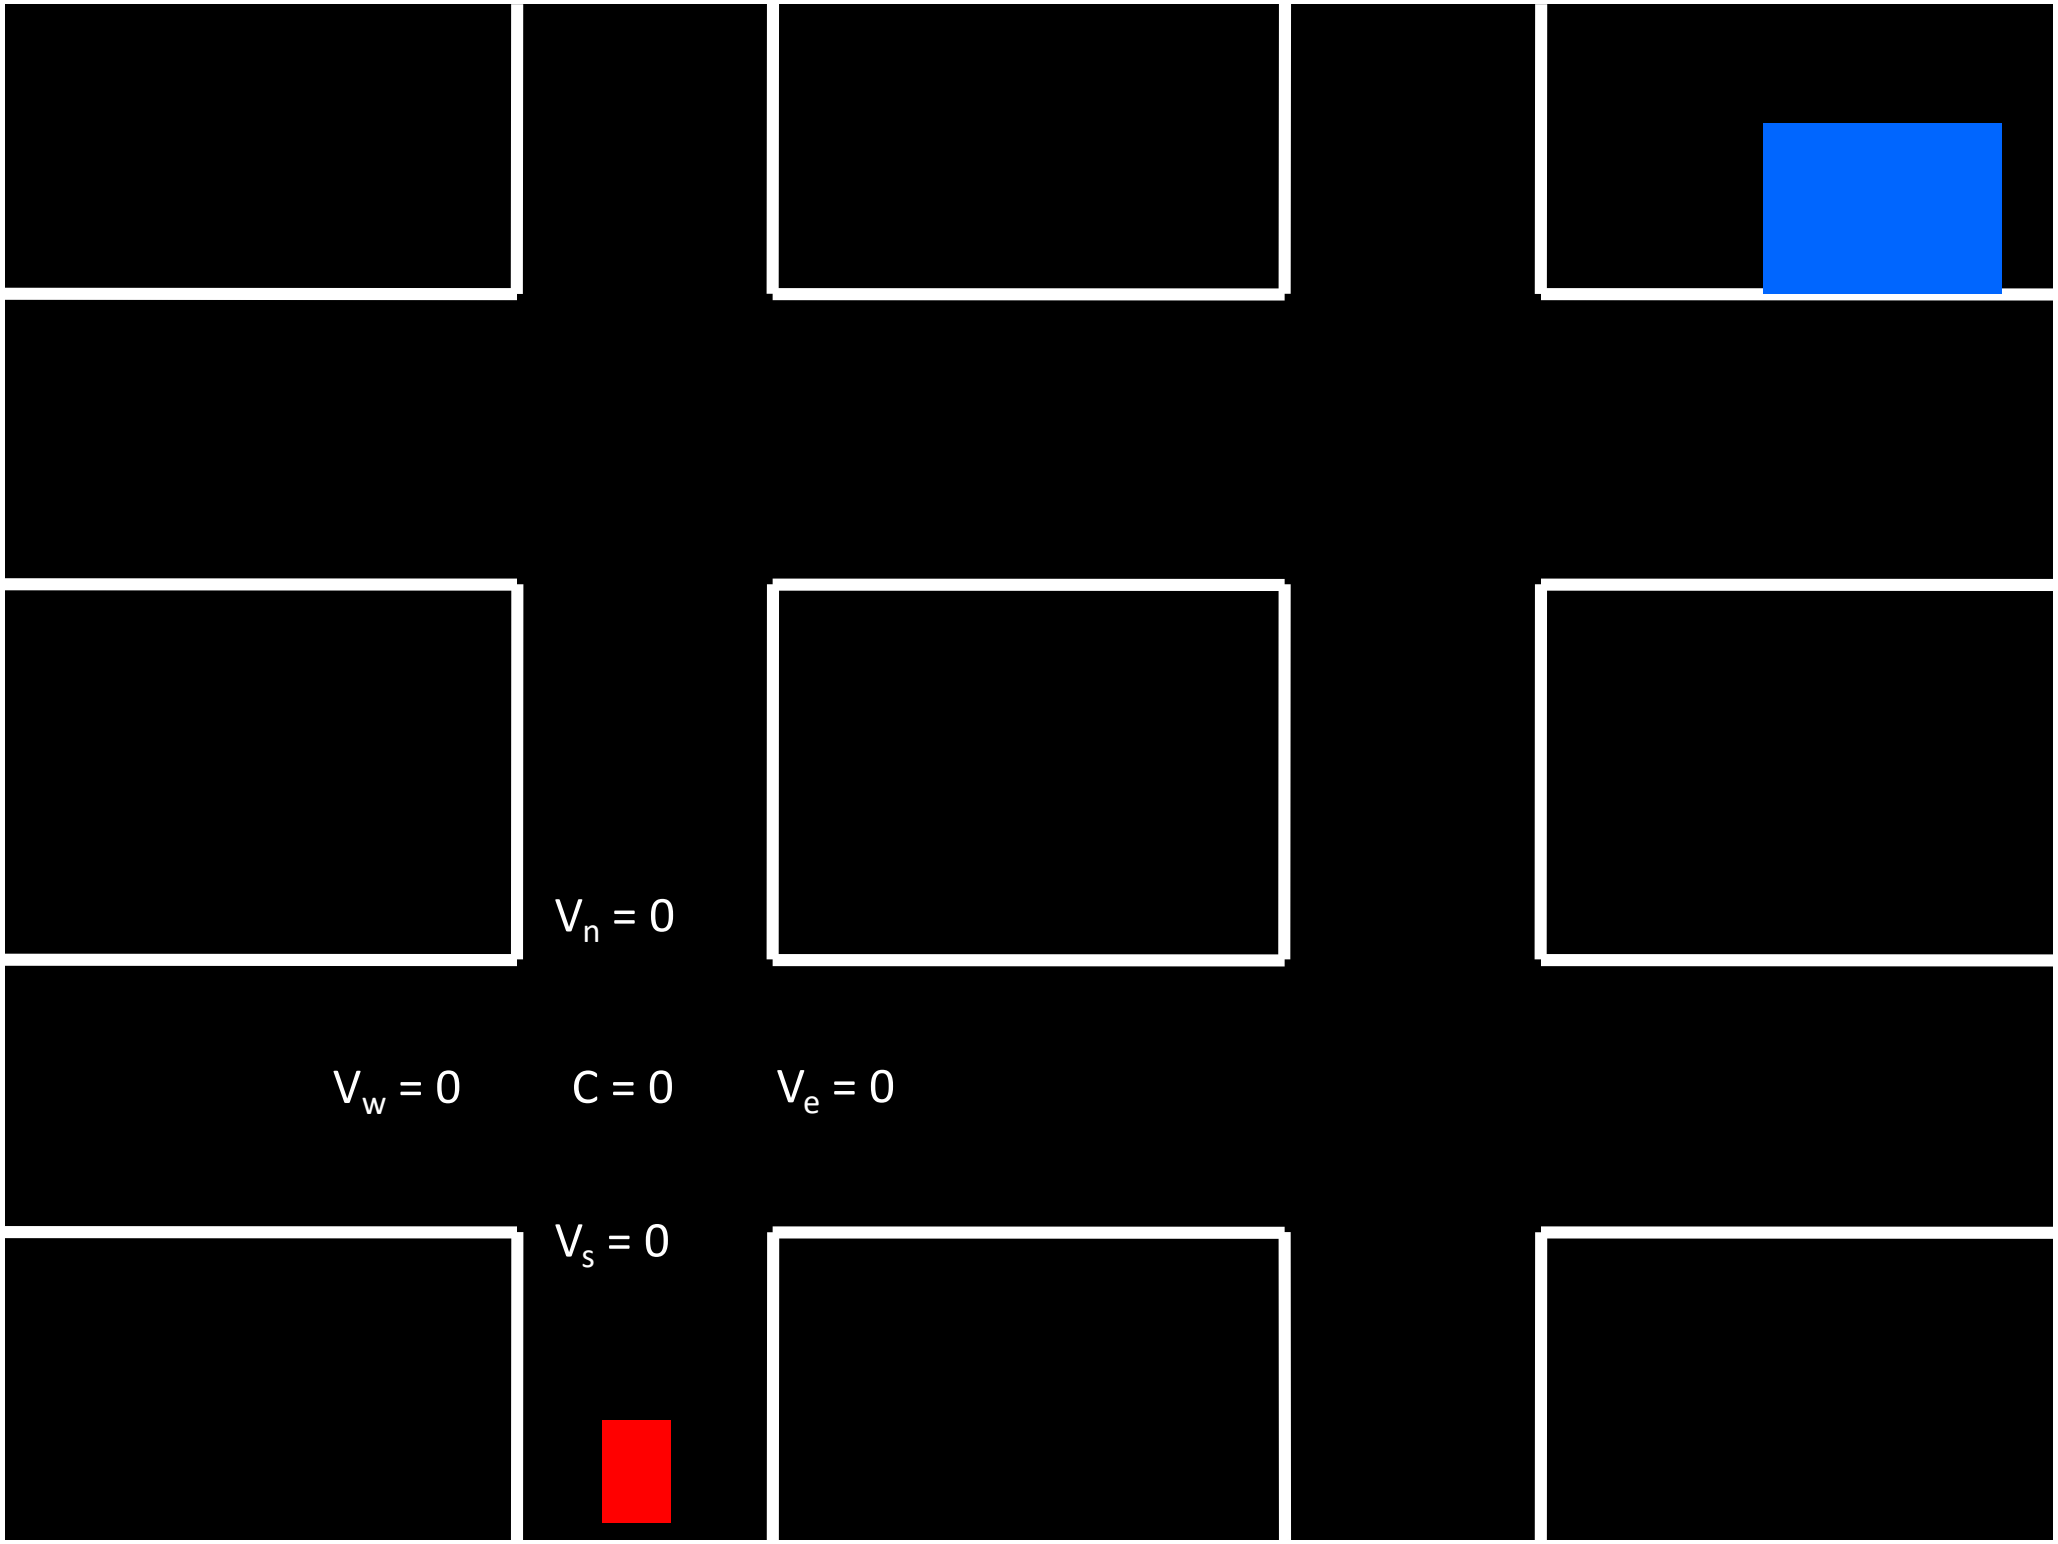
$V_n = 0$
$V_w = 0$
$C = 0$
$V_e = 0$
$V_s = 0$

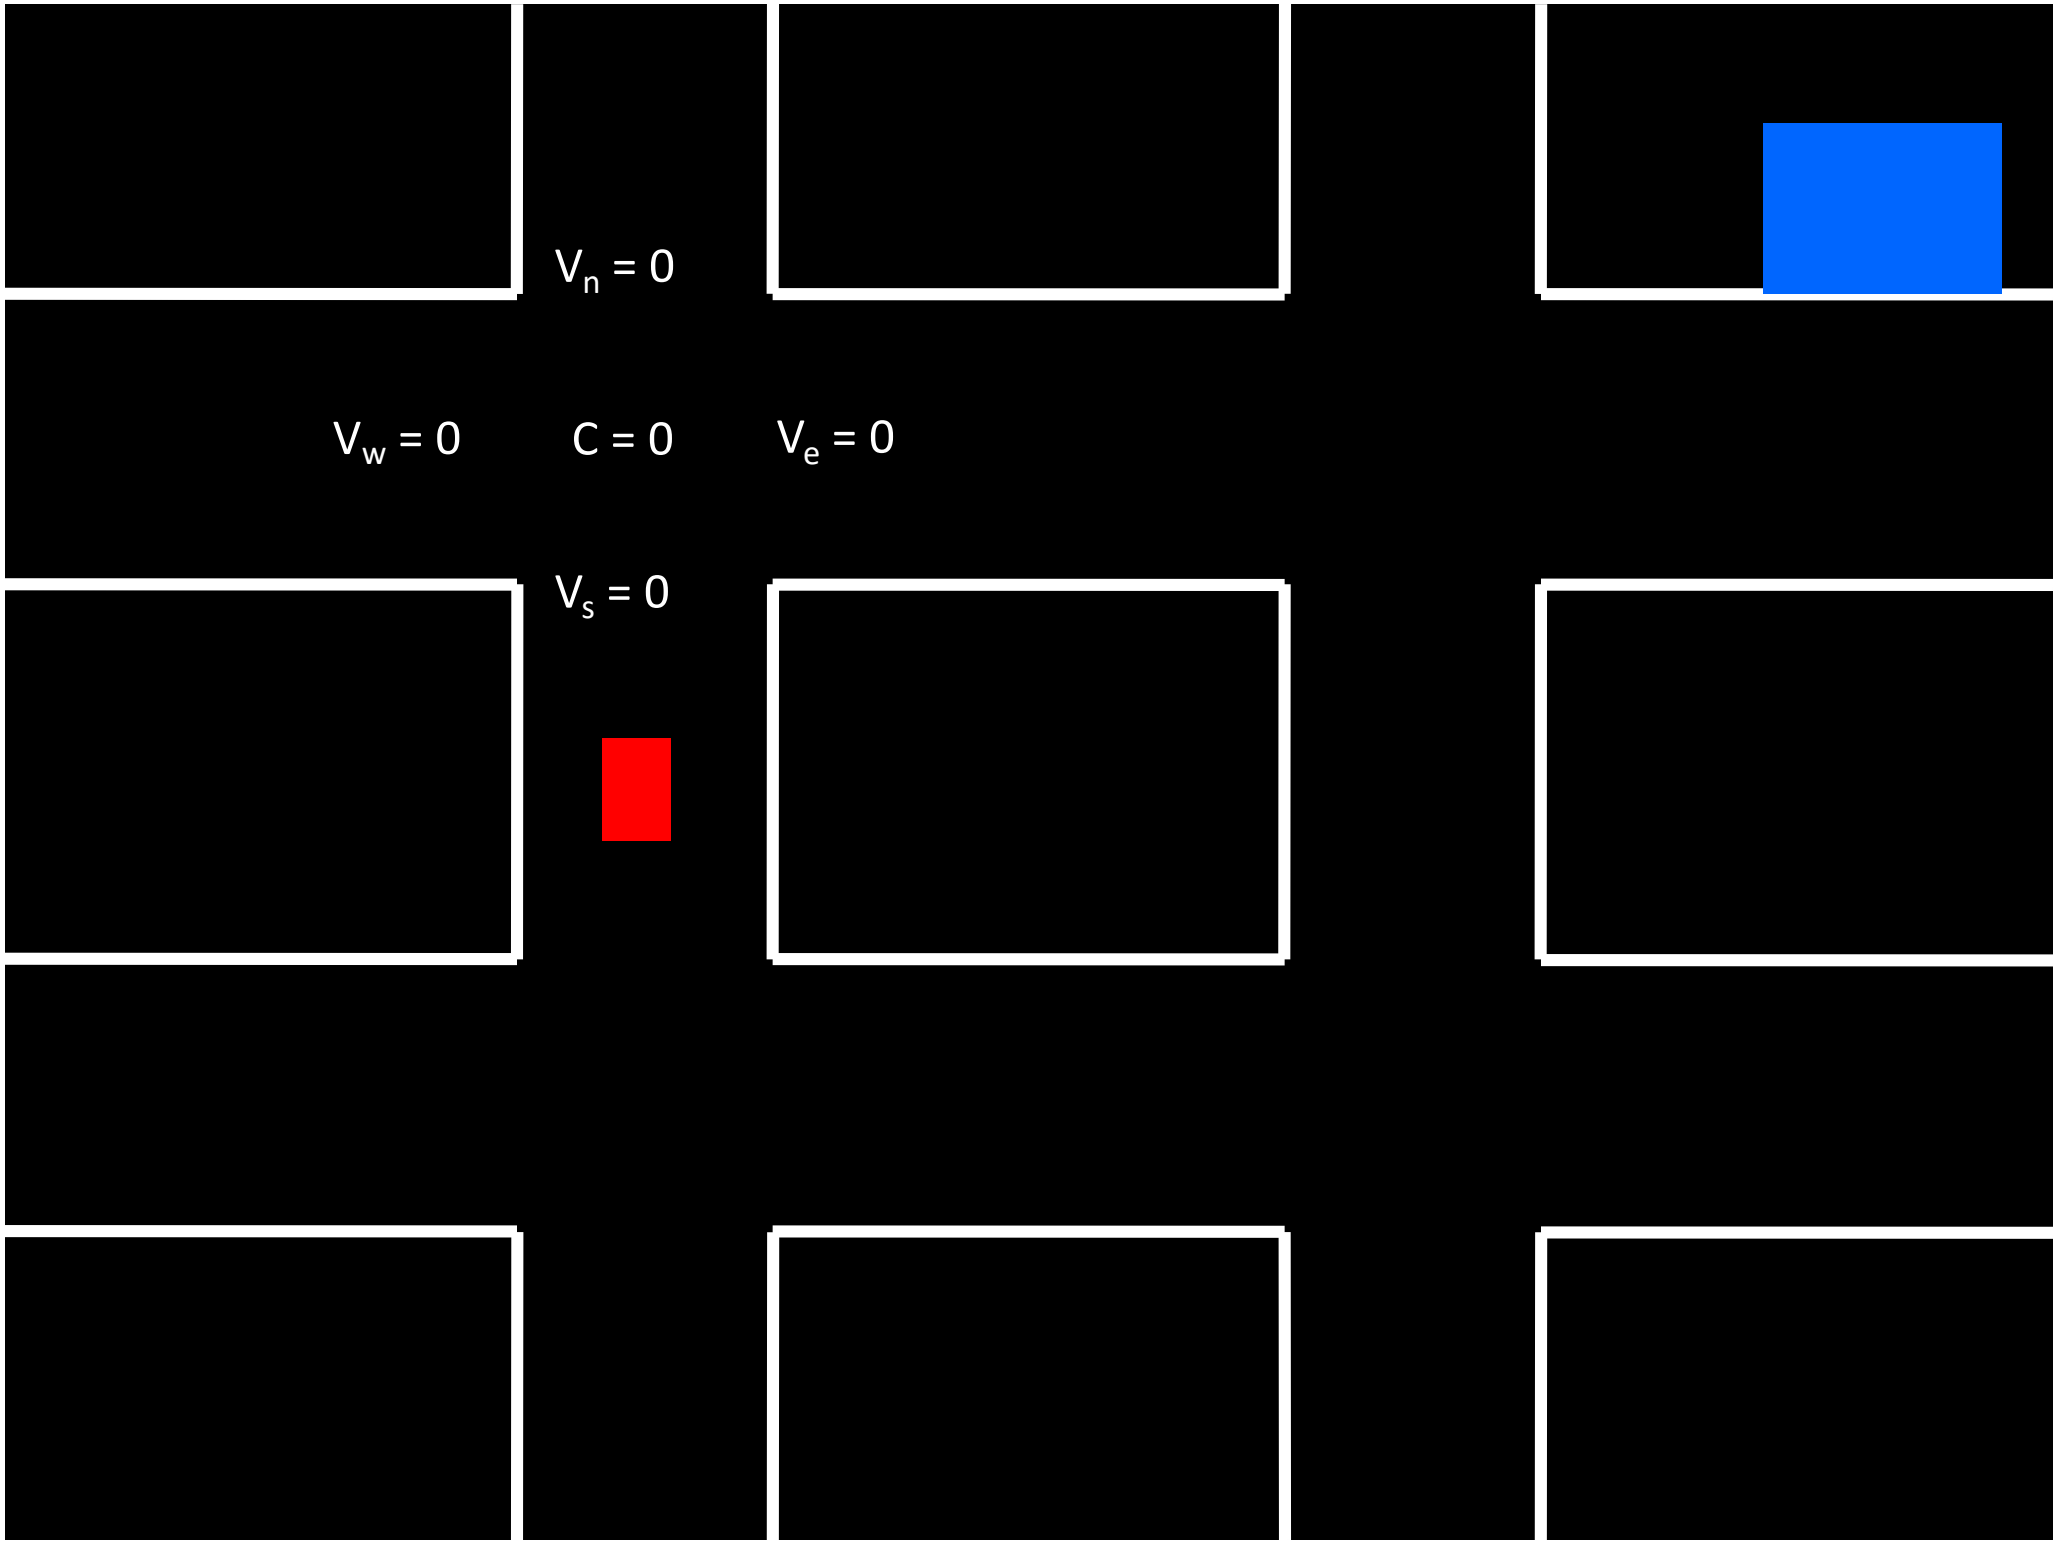
$V_n = 0$
$V_w = 0$
$C = 0$
$V_e = 0$
$V_s = 0$

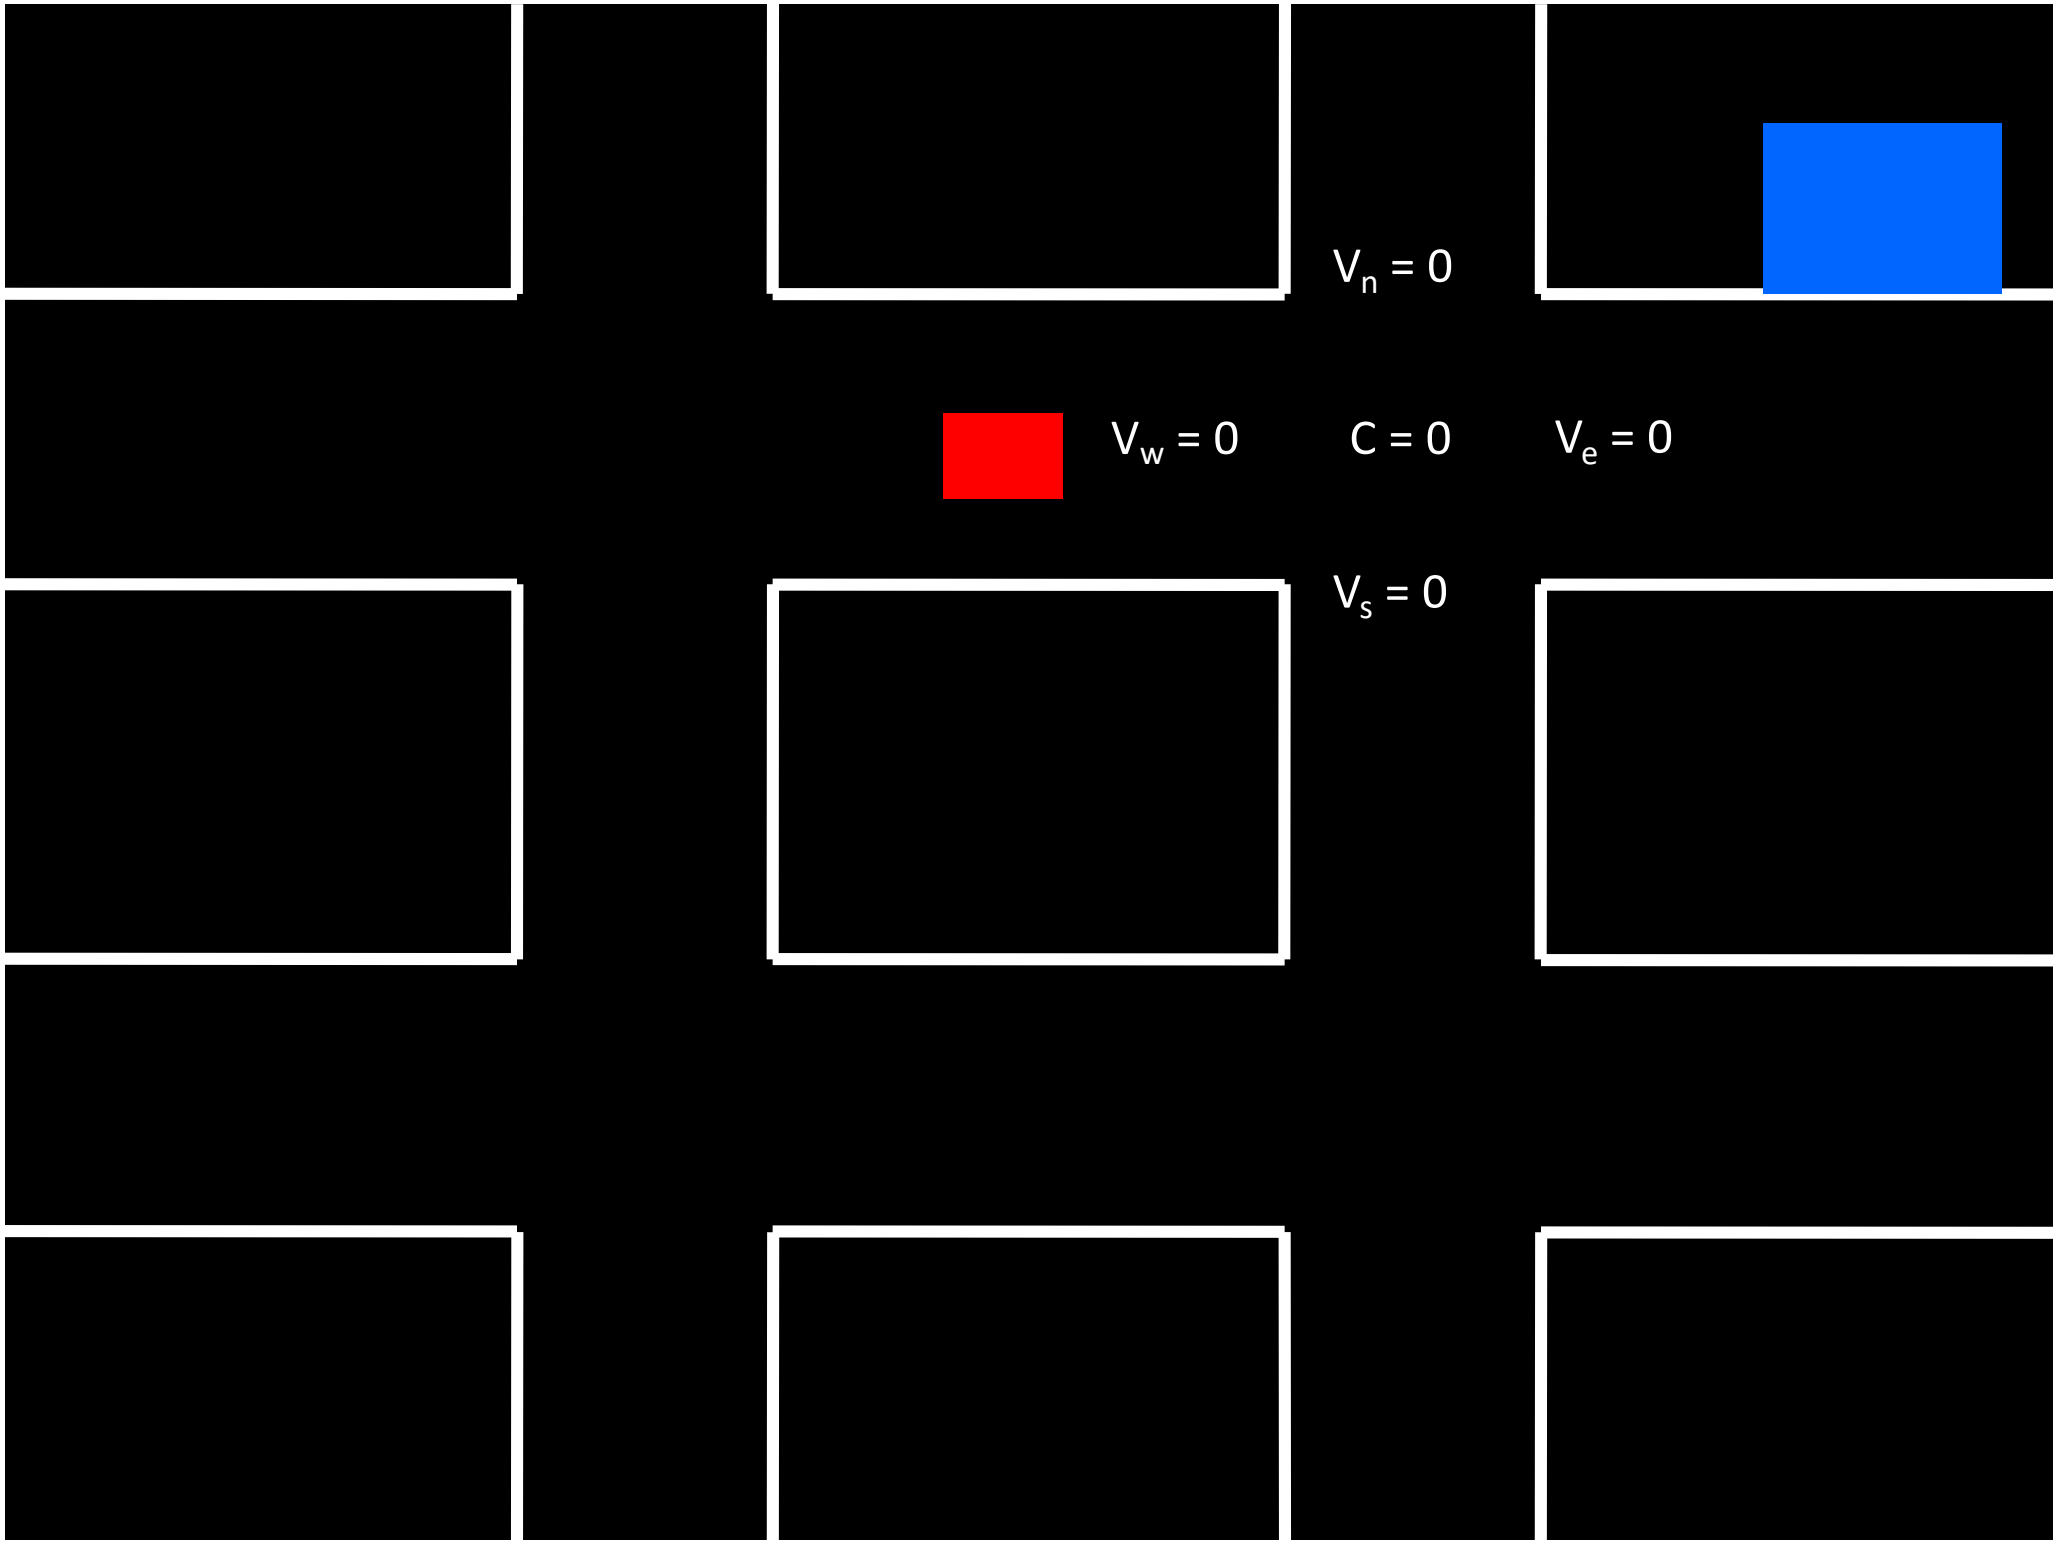
$V_n = 0$
$V_w = 0$
$C = 0$
$V_e = 0$
$V_s = 0$

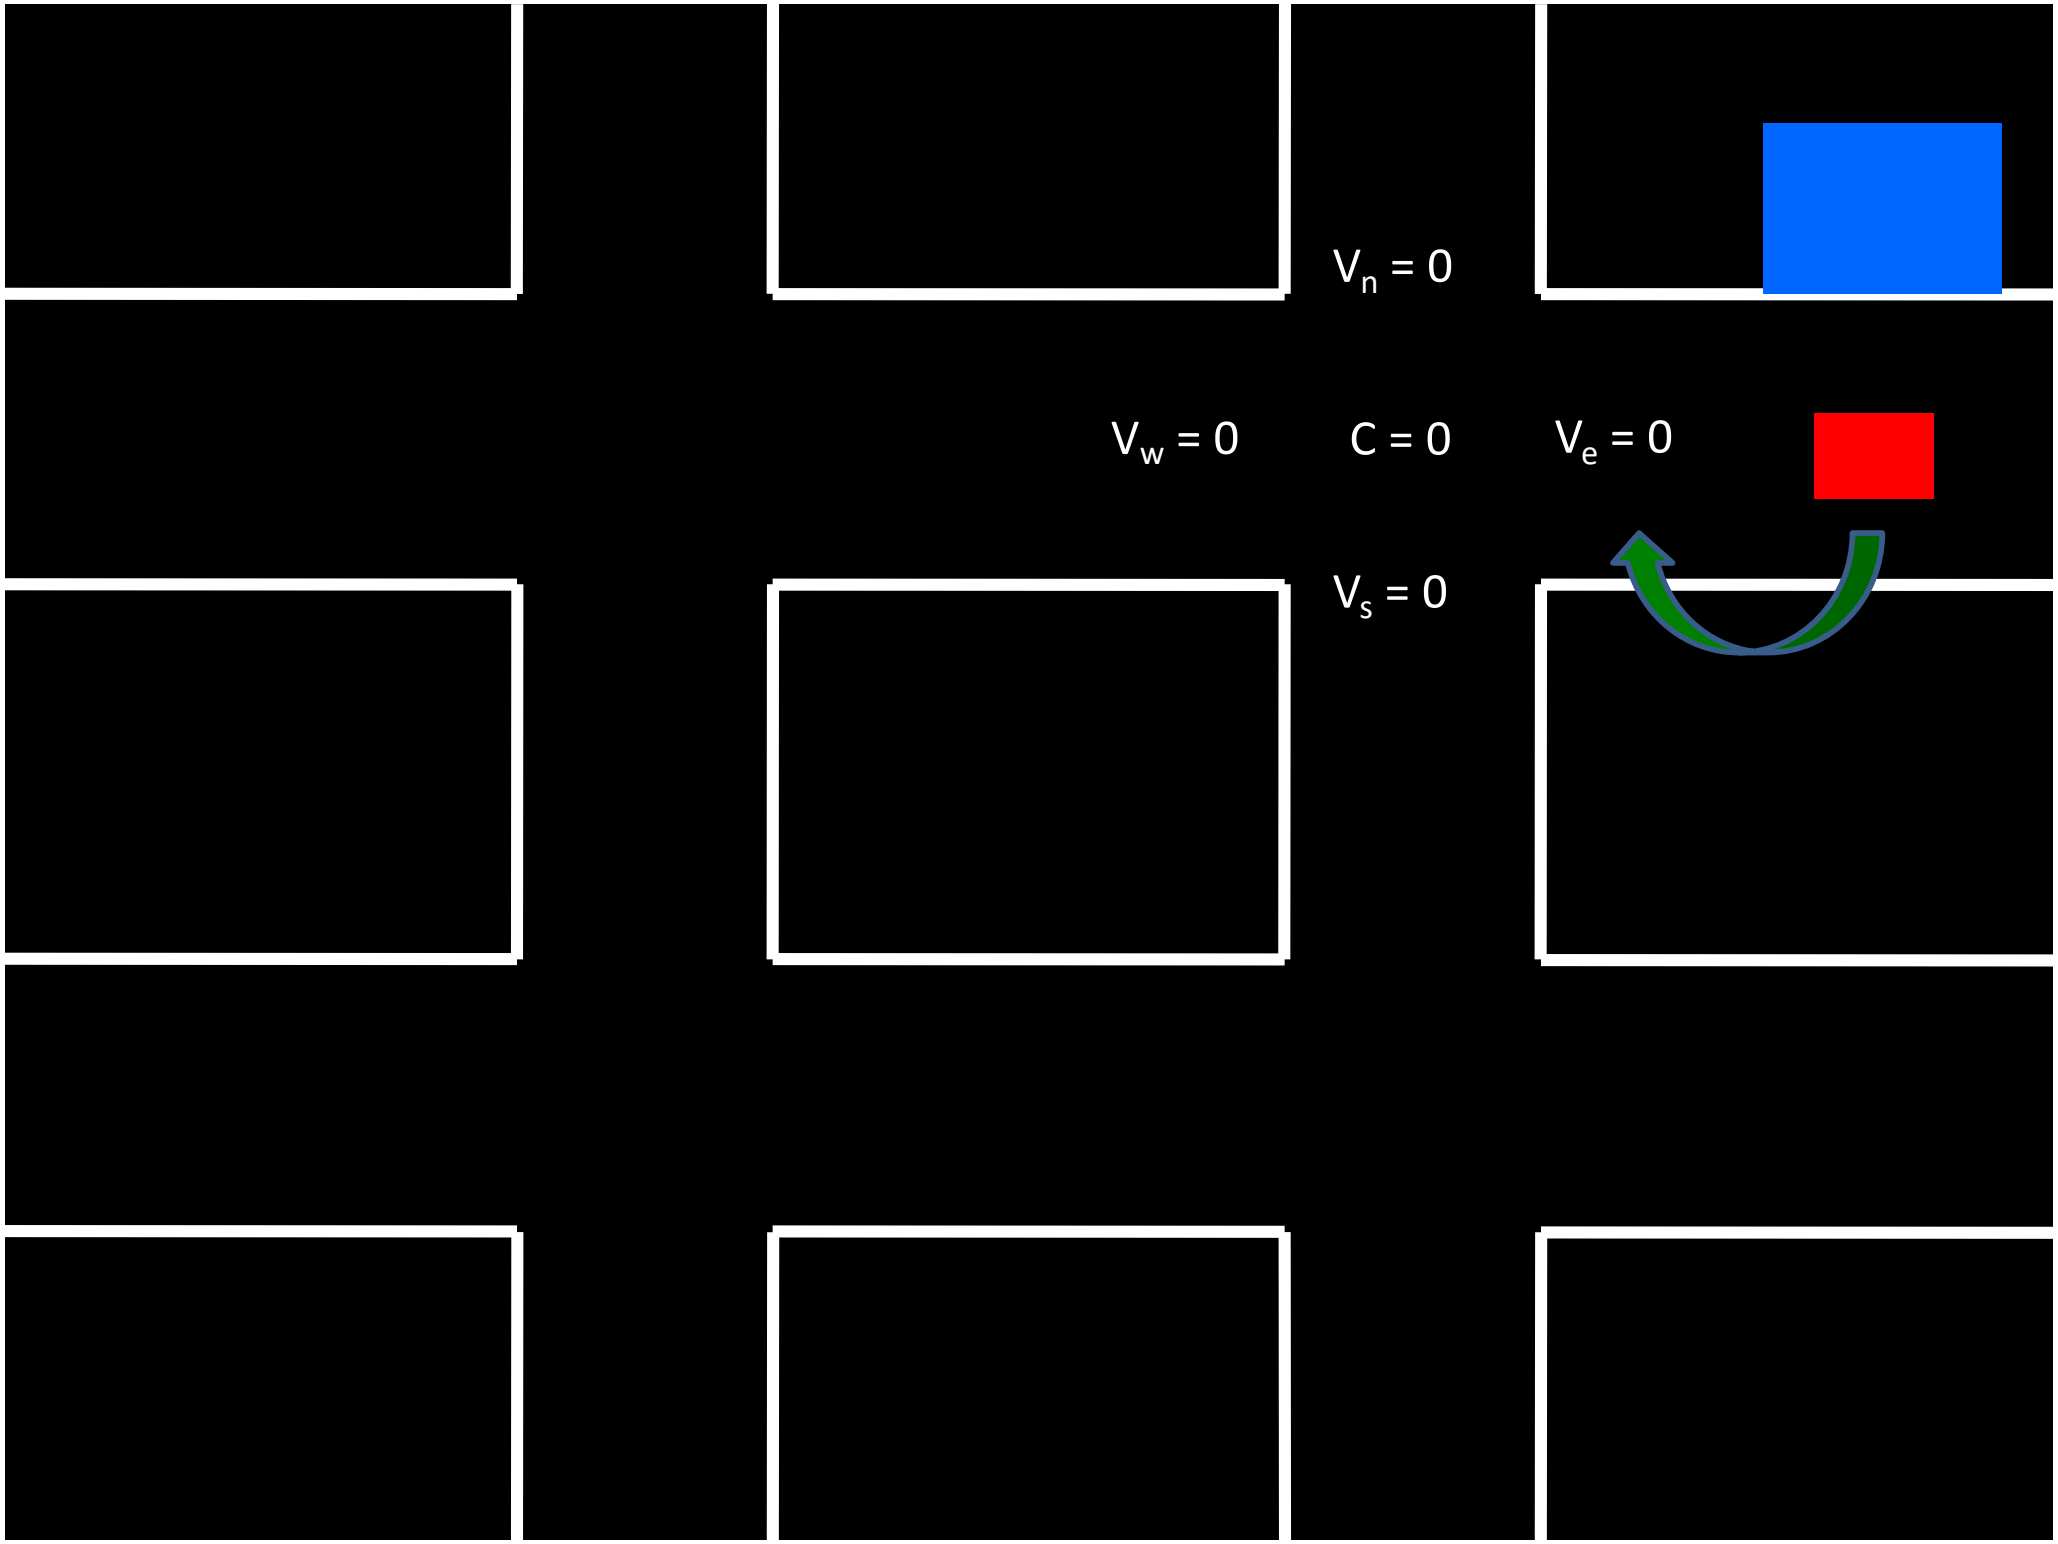
$V_n = 0$
$V_w = 0$
$C = 0$
$V_e = 0$
$V_s = 0$

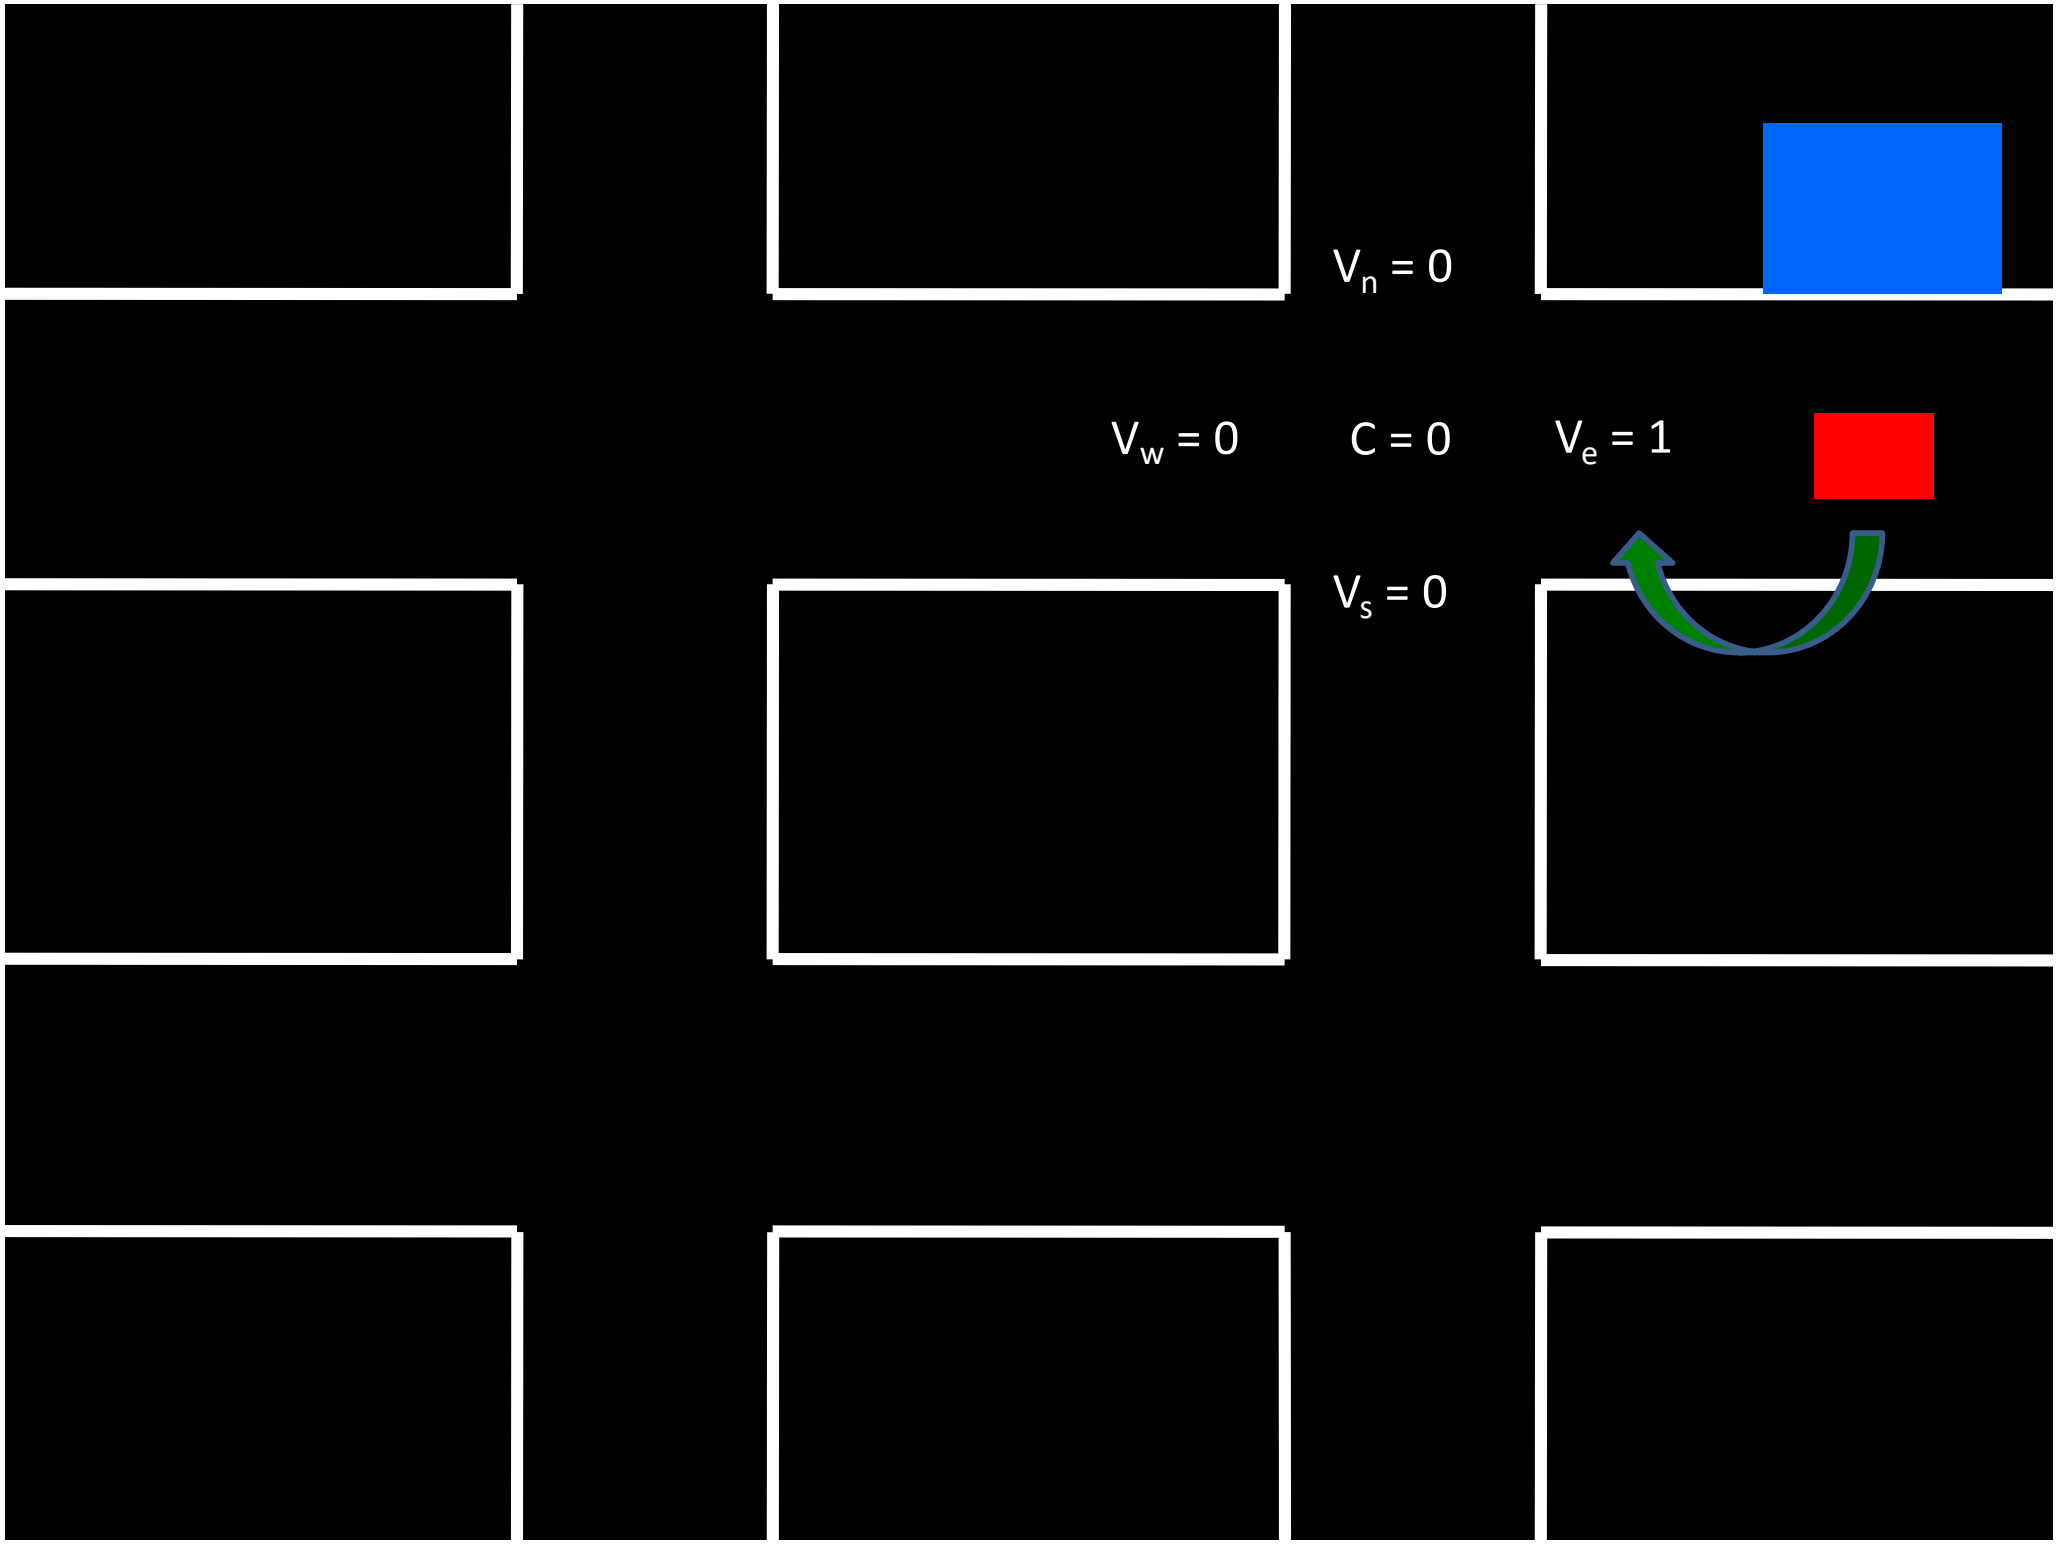
$V_n = 0$
$V_w = 0$
$C = 0$
$V_e = 1$
$V_s = 0$

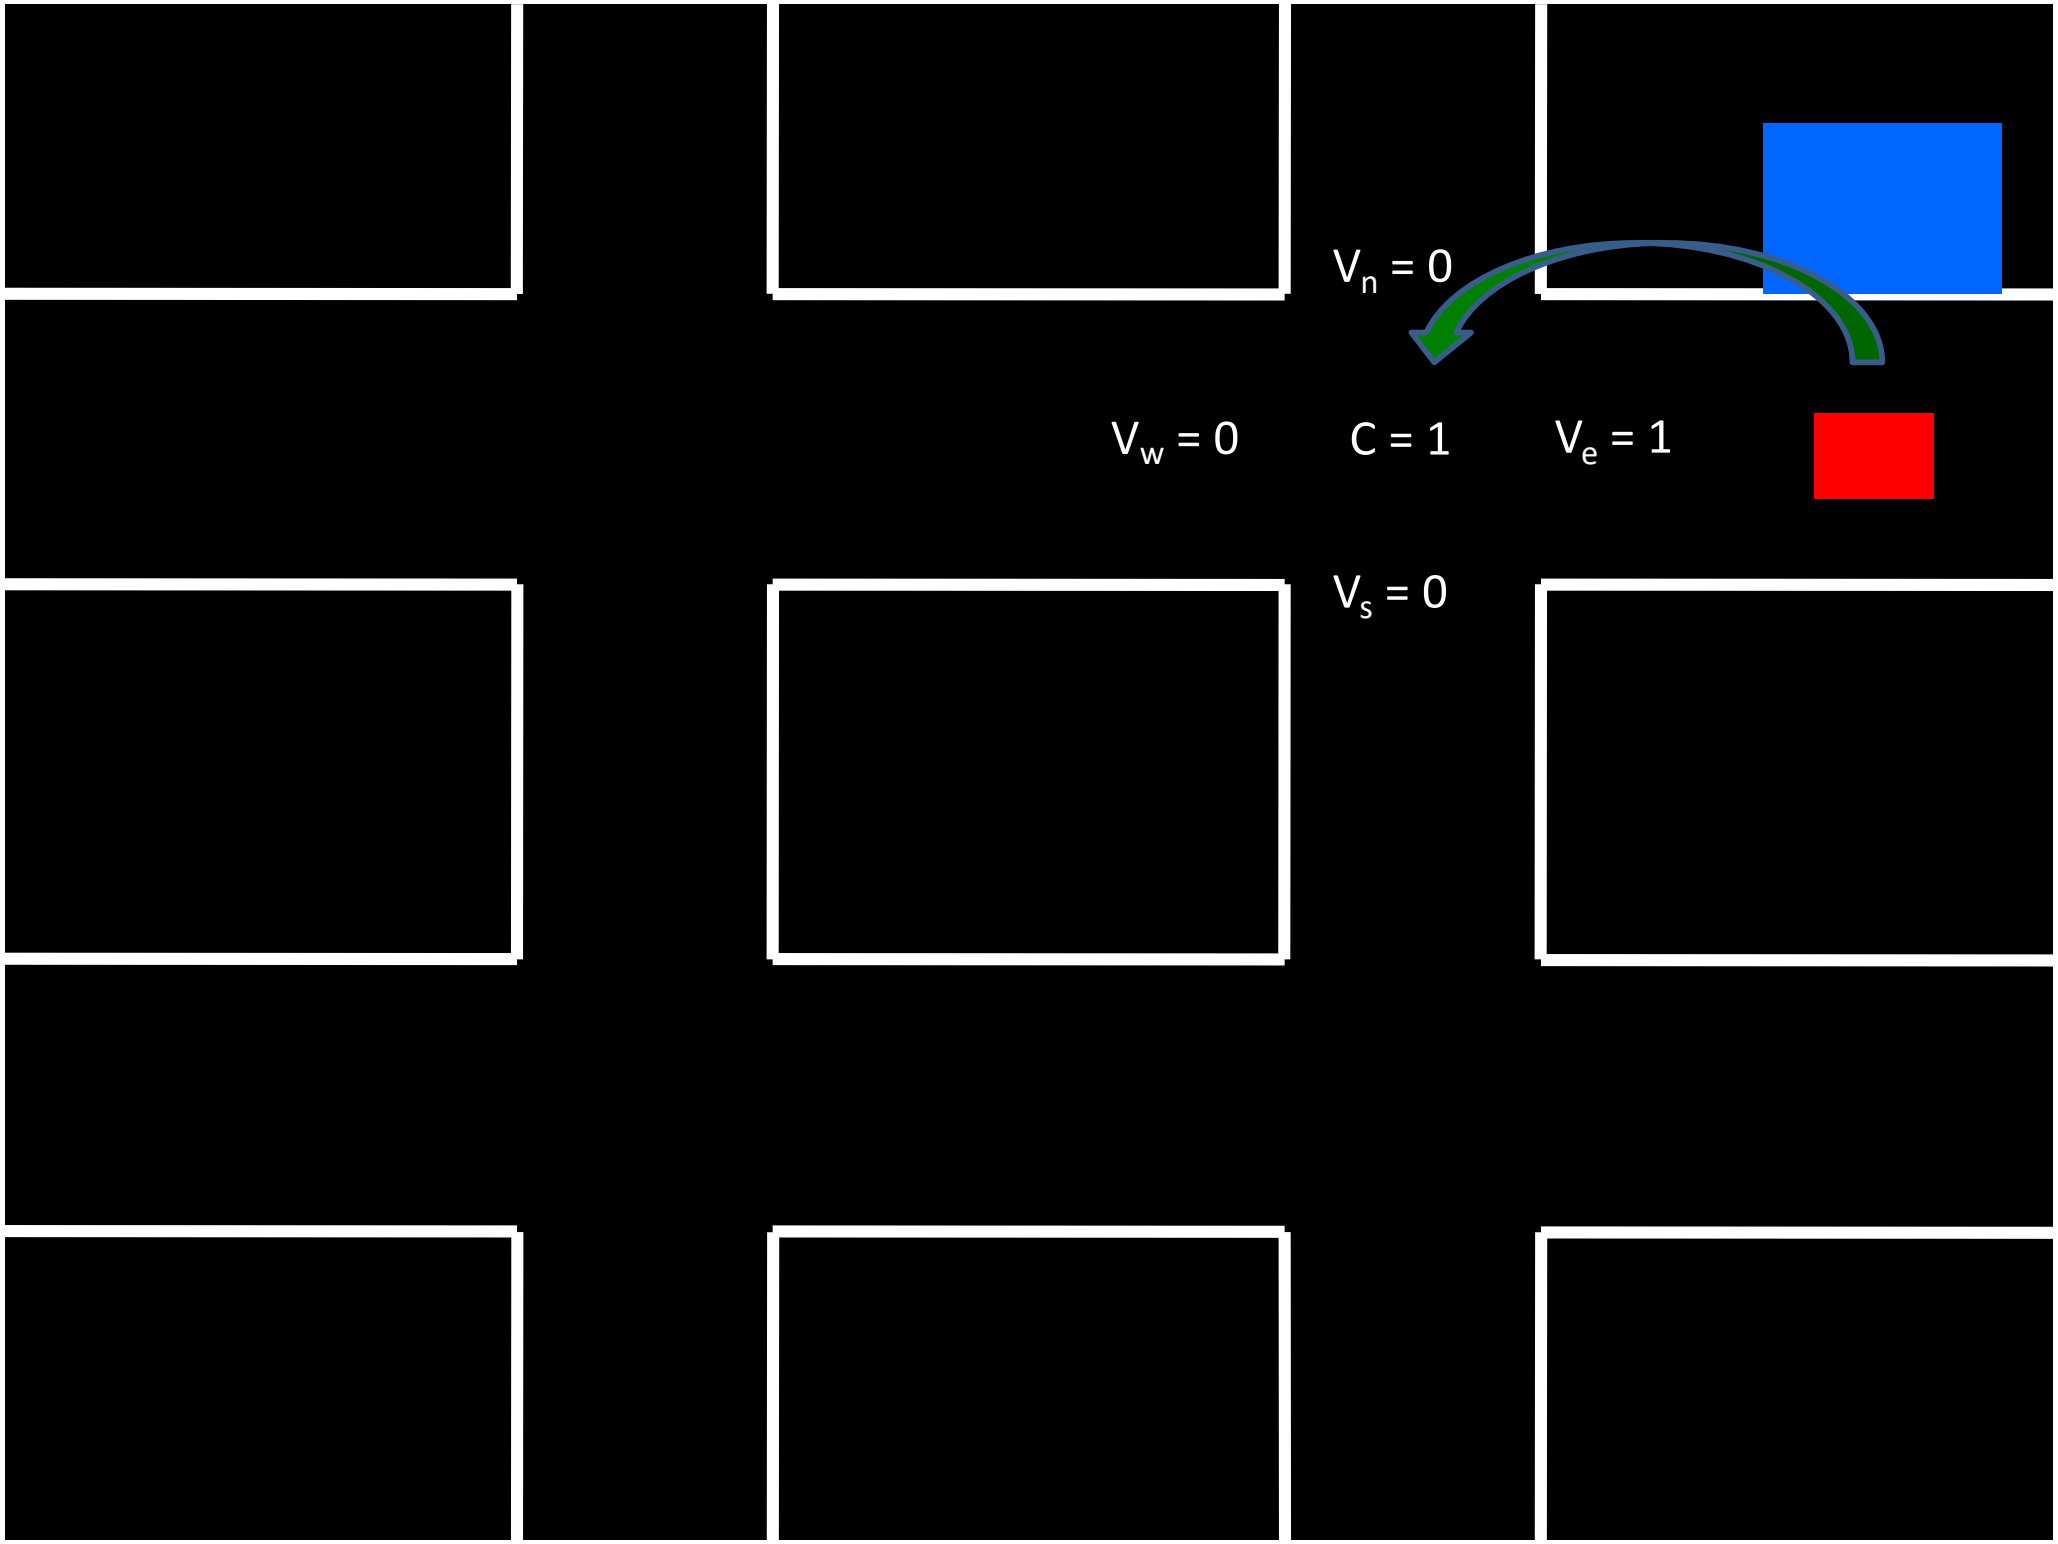
$V_n = 0$
$V_w = 0$
$C = 1$
$V_e = 1$
$V_s = 0$

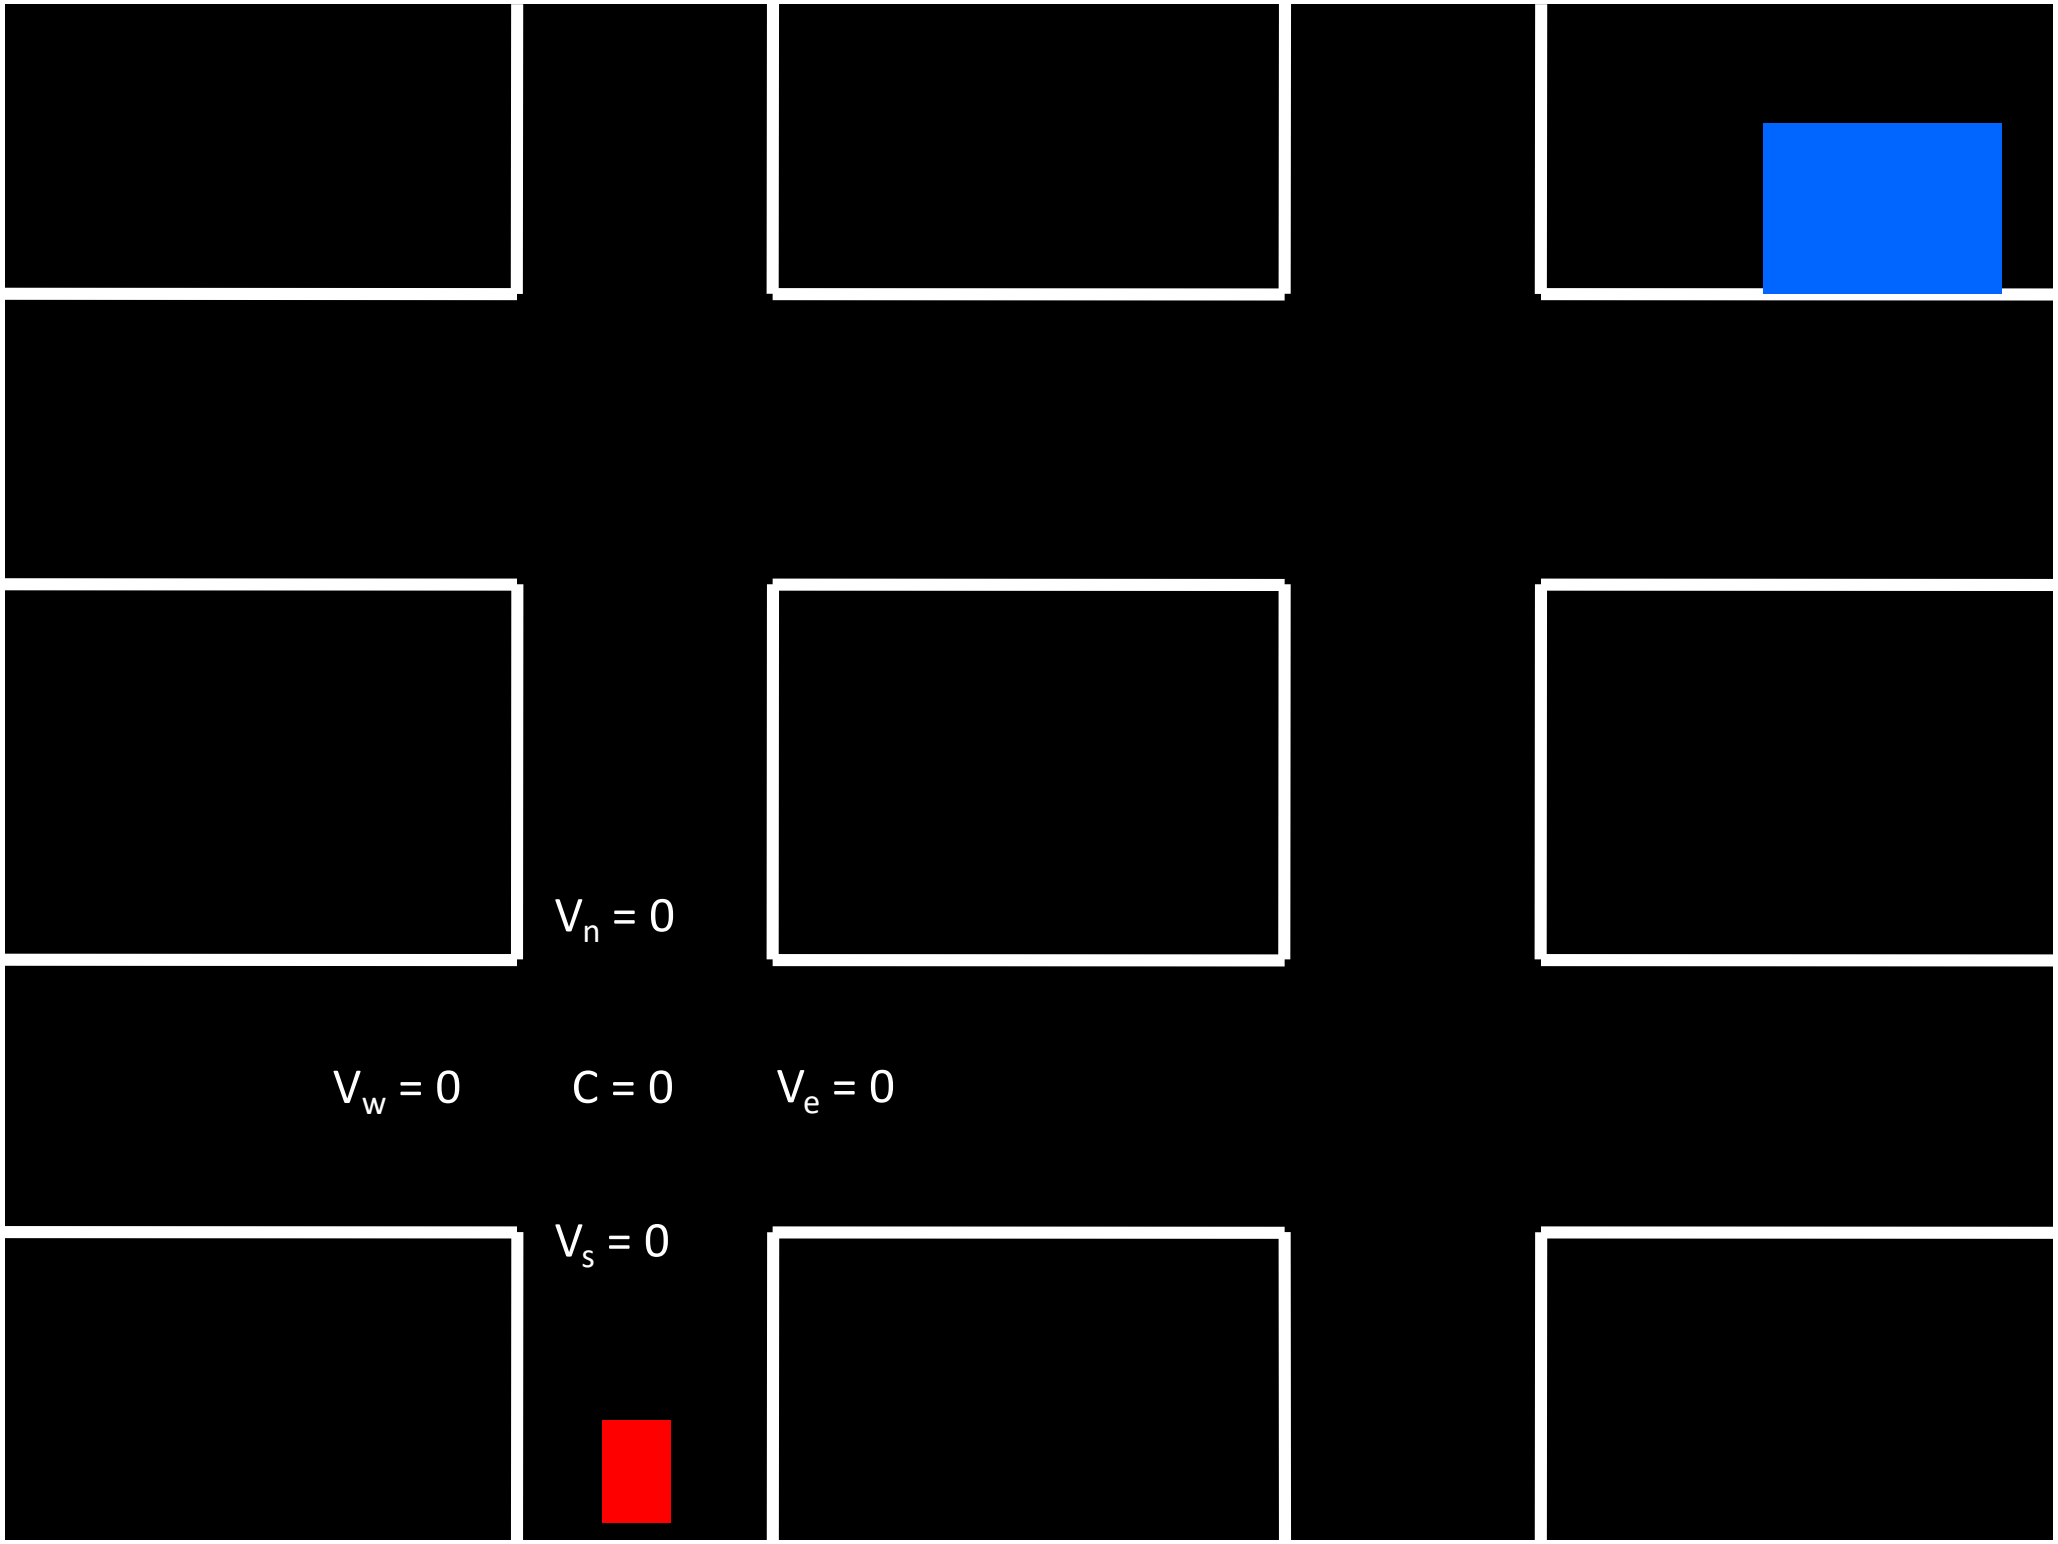
$V_n = 0$
$V_w = 0$
$C = 0$
$V_e = 0$
$V_s = 0$

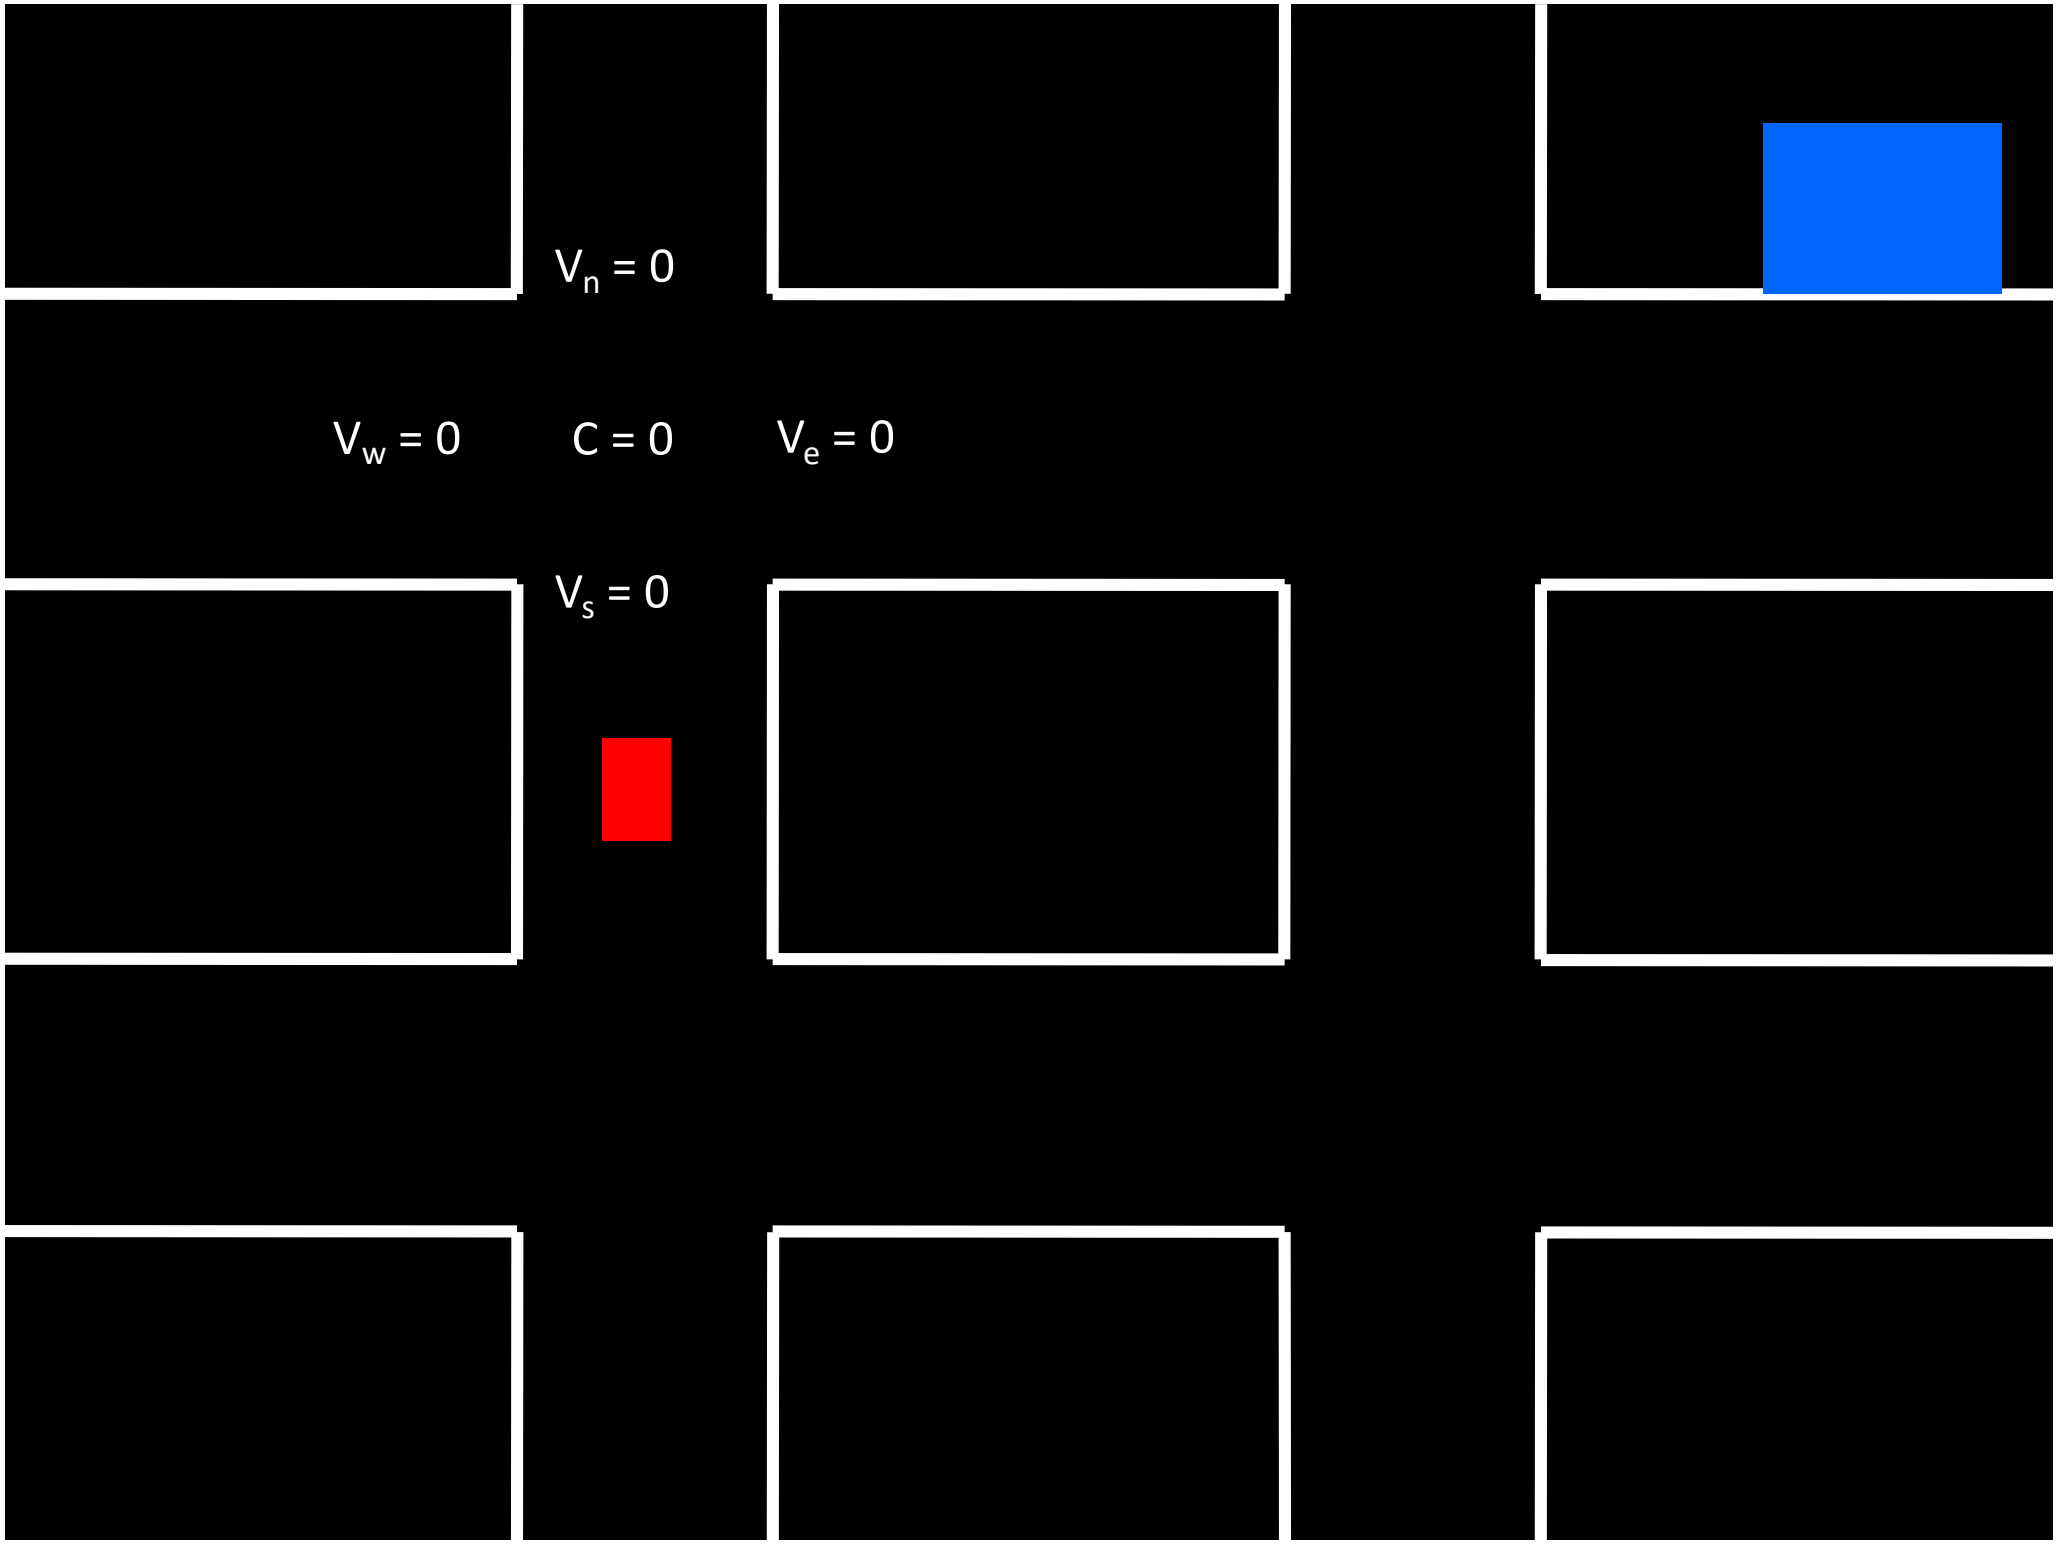
$V_n = 0$
$V_w = 0$
$C = 0$
$V_e = 0$
$V_s = 0$

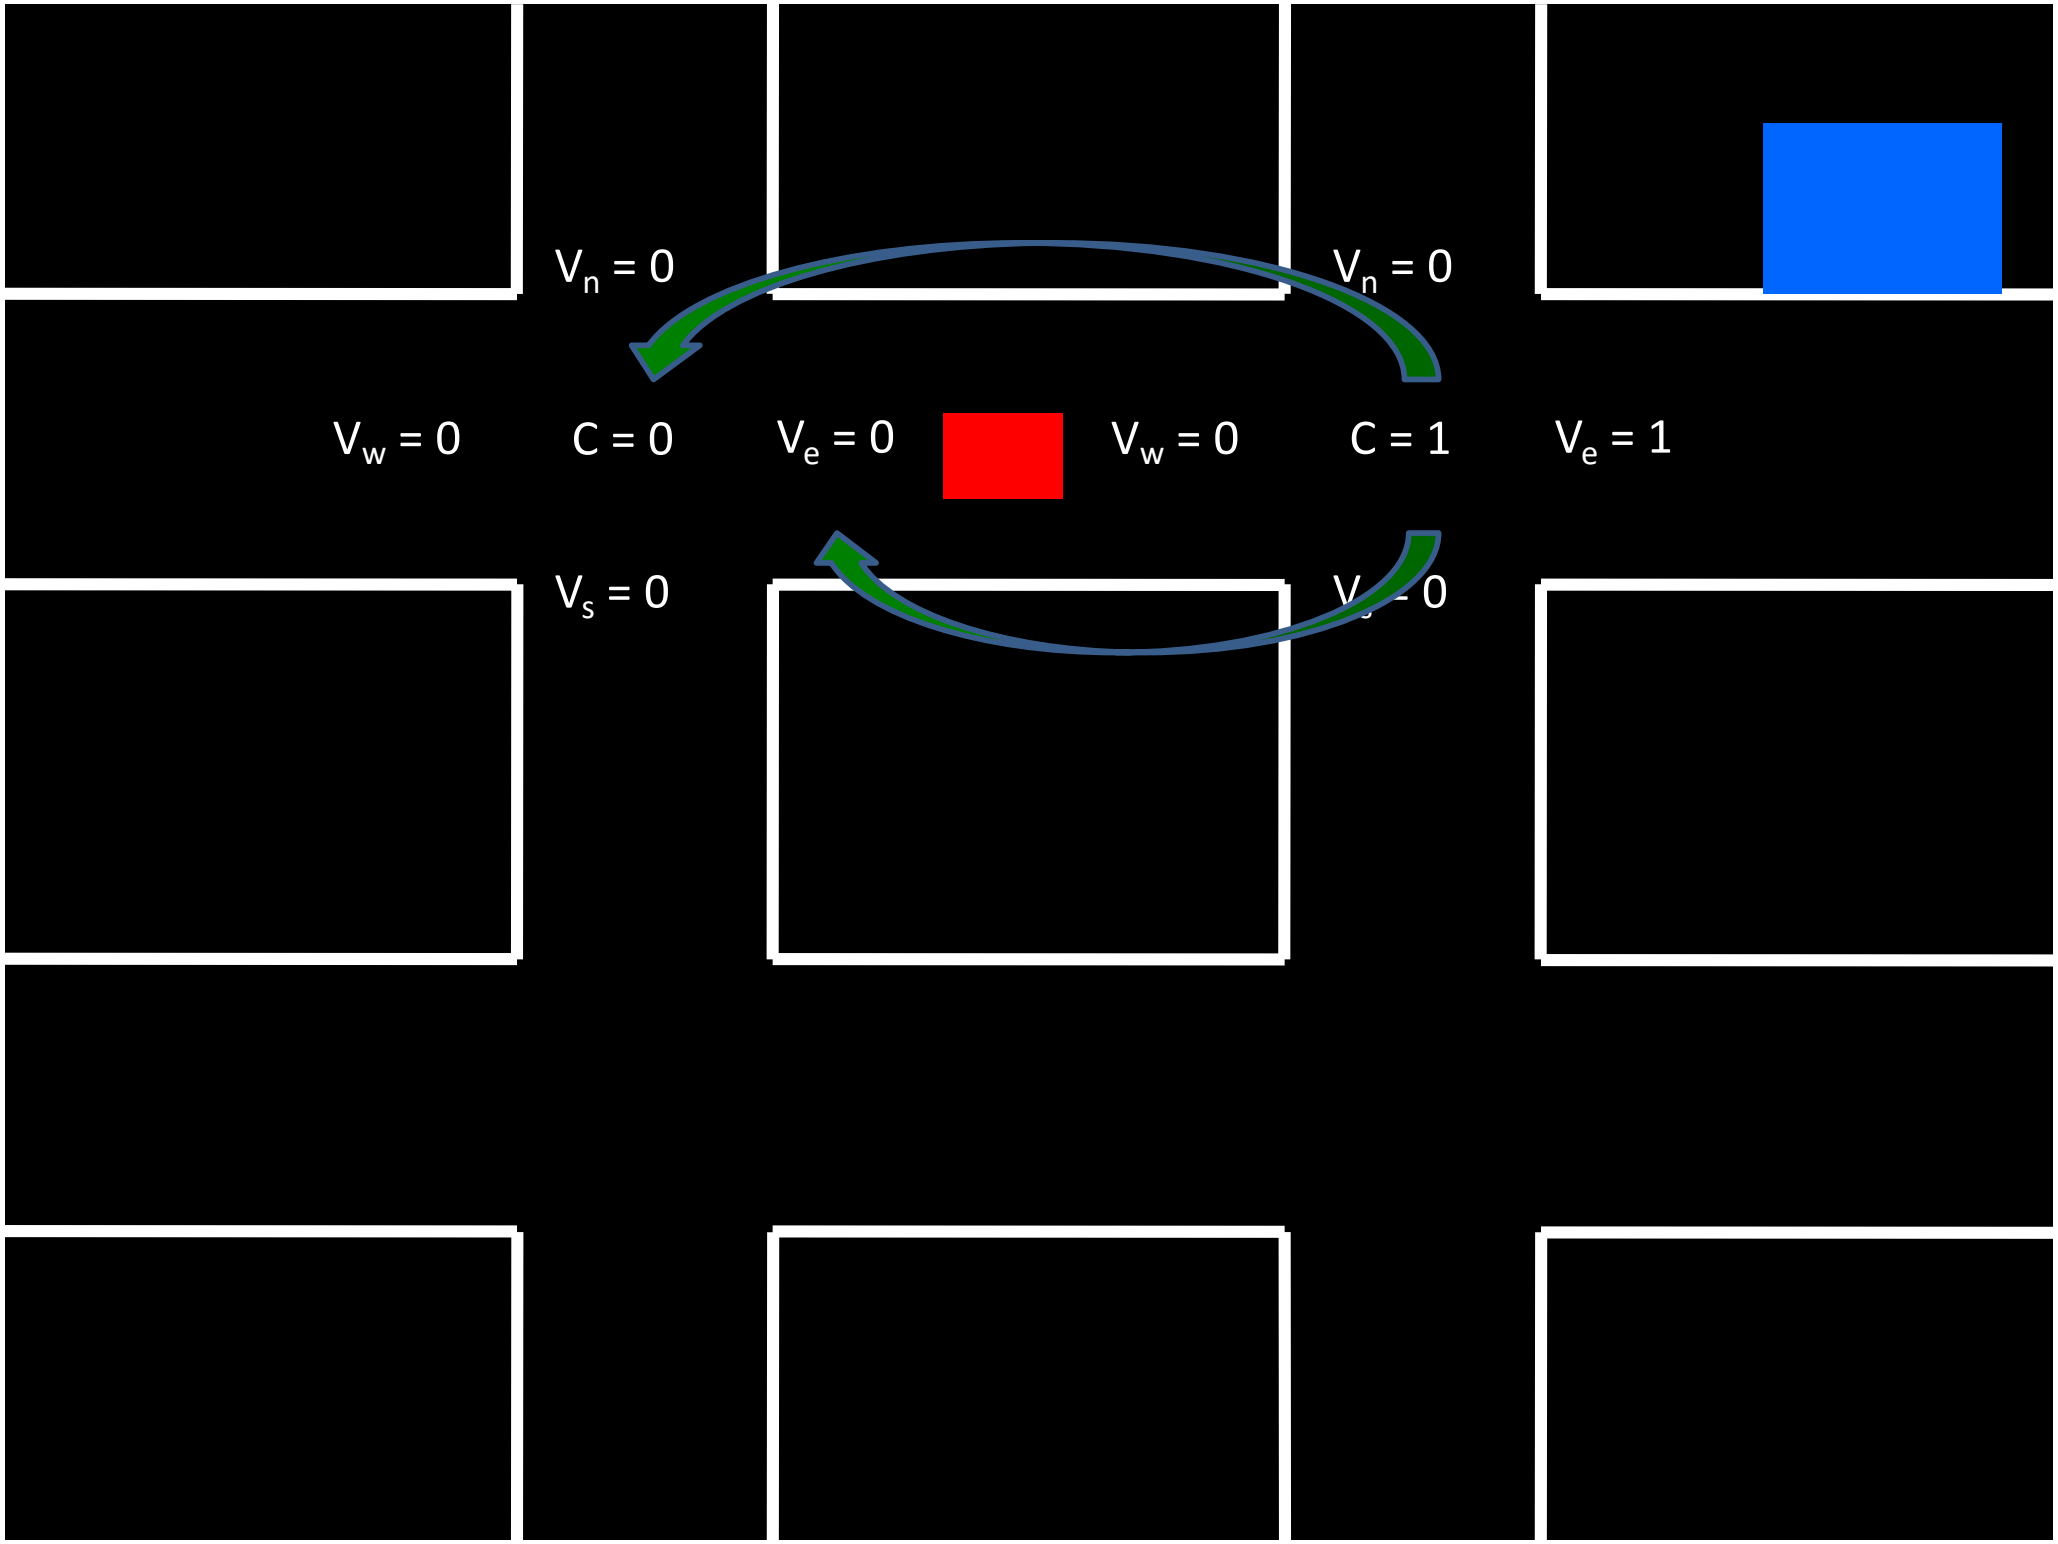
$V_n = 0$
$V_n = 0$
$V_w = 0$
$C = 0$
$V_e = 0$
$V_w = 0$
$C = 1$
$V_e = 1$
$V_s = 0$

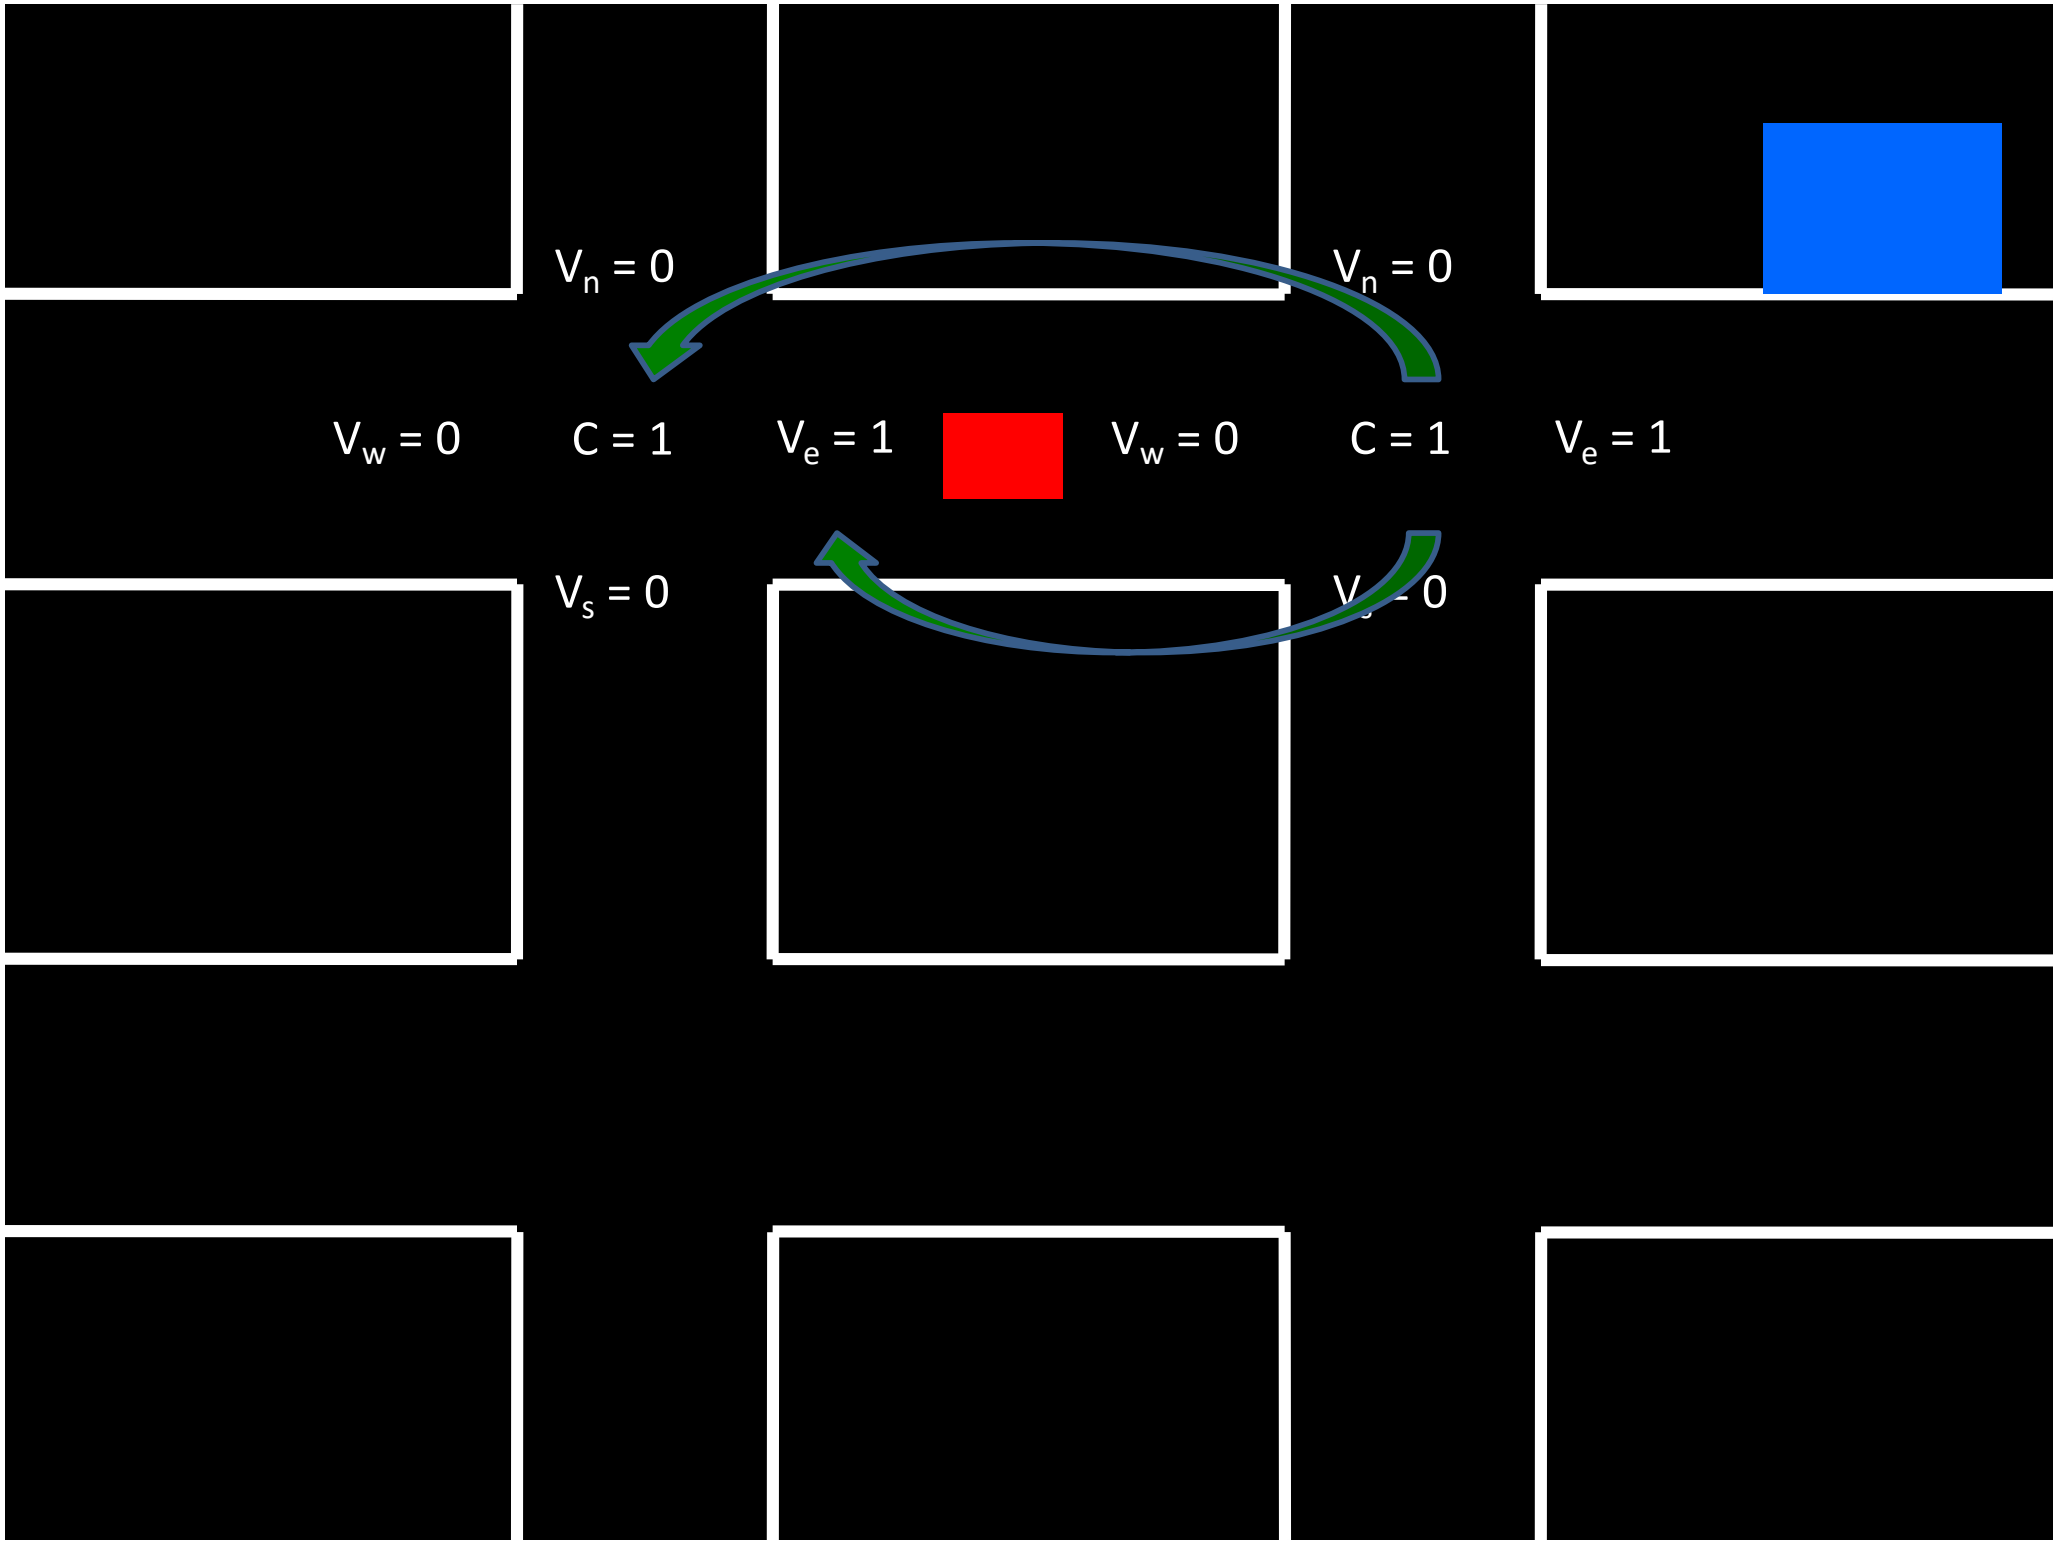
$V_n = 0$
$V_n = 0$
$V_w = 0$
$C = 1$
$V_e = 1$
$V_w = 0$
$C = 1$
$V_e = 1$
$V_s = 0$

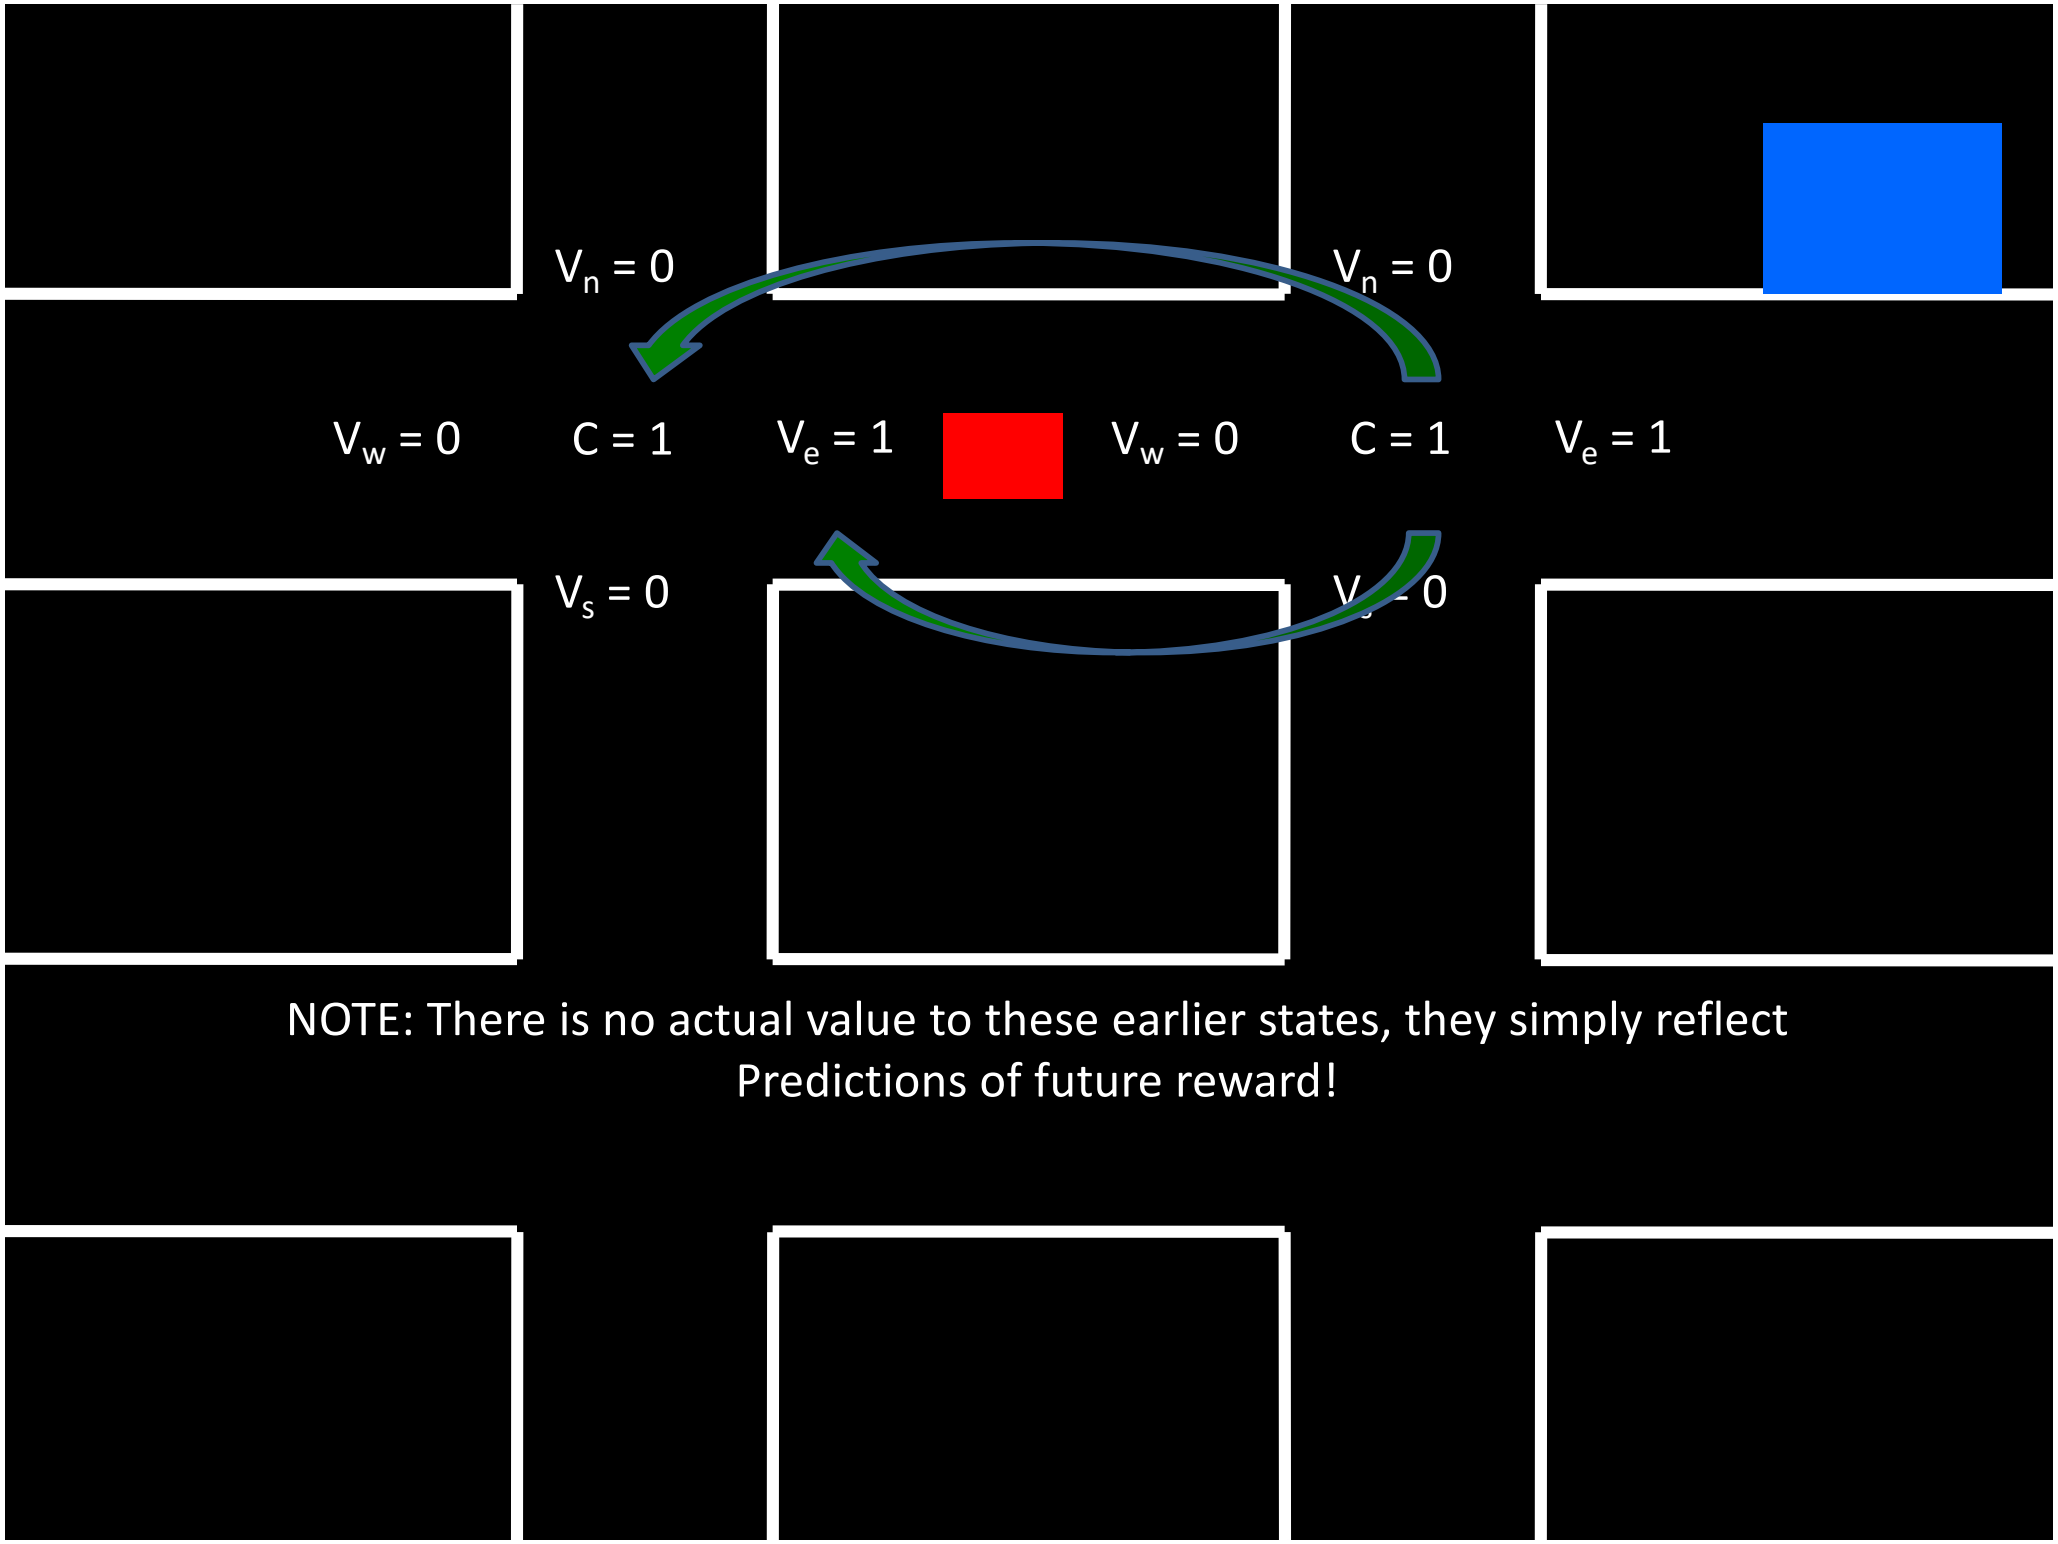

NOTE: There is no actual value to these earlier states, they simply reflect Predictions of future reward!

This also applies to motor skills!

Recall that a skill is a collection of neurons being activated.

PE's can be used in principle to strengthen the connections between these neurons to "learn" the correct movement pattern.

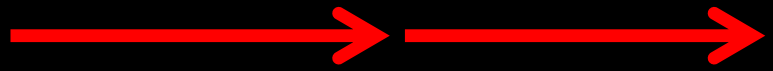

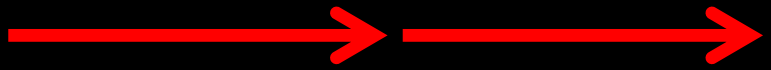

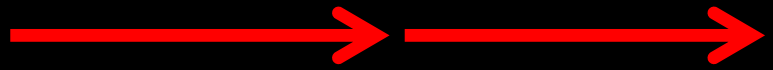

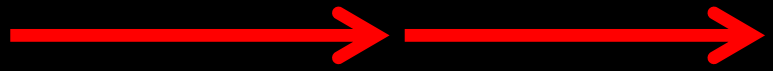

$PE = V_{outcome} - V_{action}$

$PE = +$

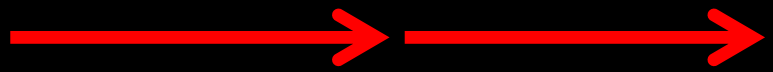

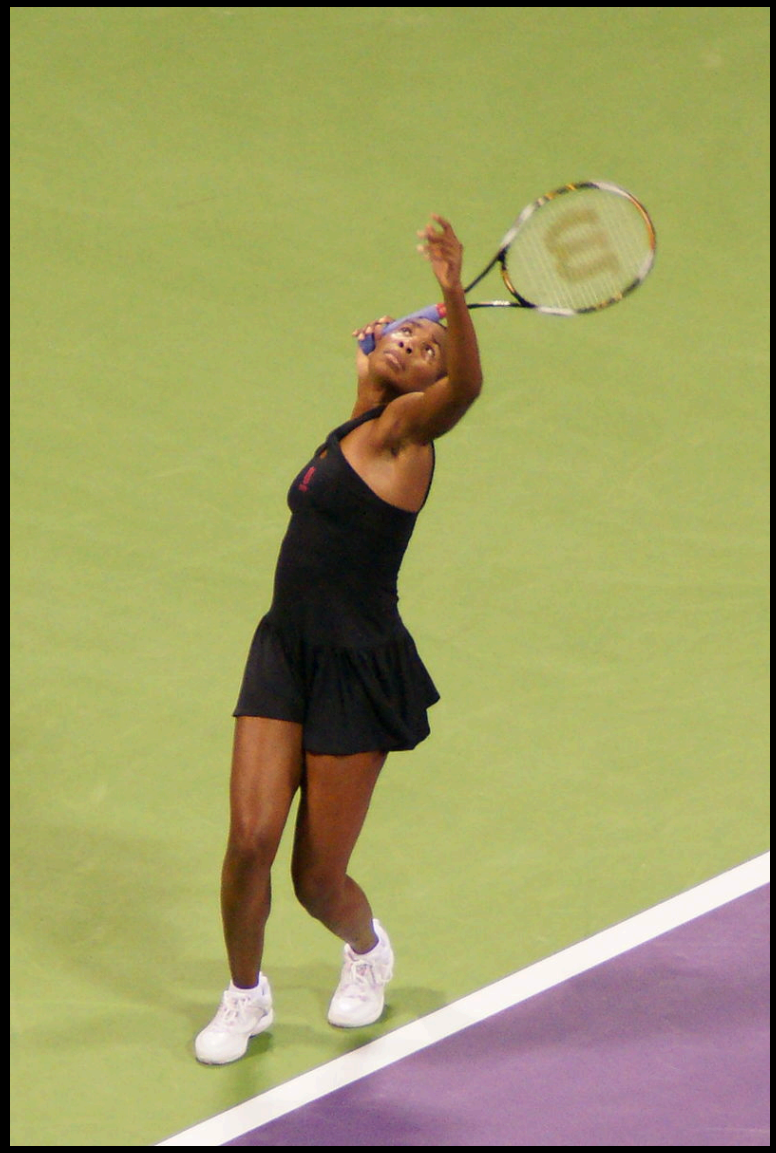

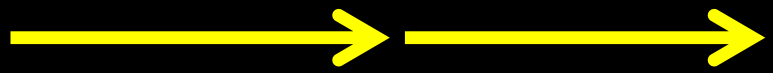

$V_{action} = V_{action} + PE * LR$

$V_{action} \uparrow$

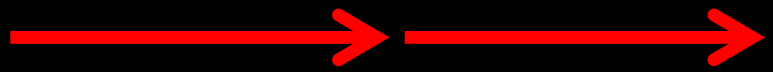

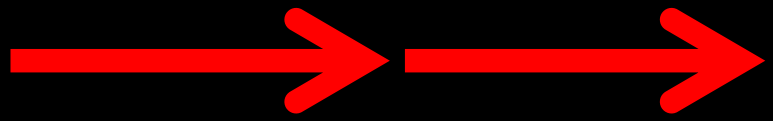

Thus, we choose the top set of neurons because this choice has a higher VALUE

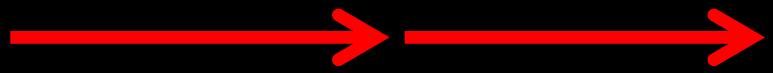

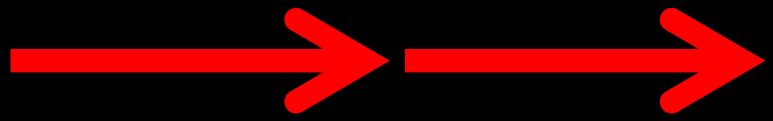

This is a basic principle of decision making – always choose the highest value option

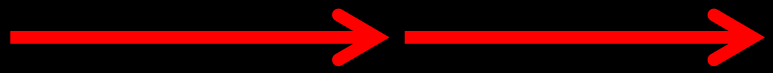

Part II

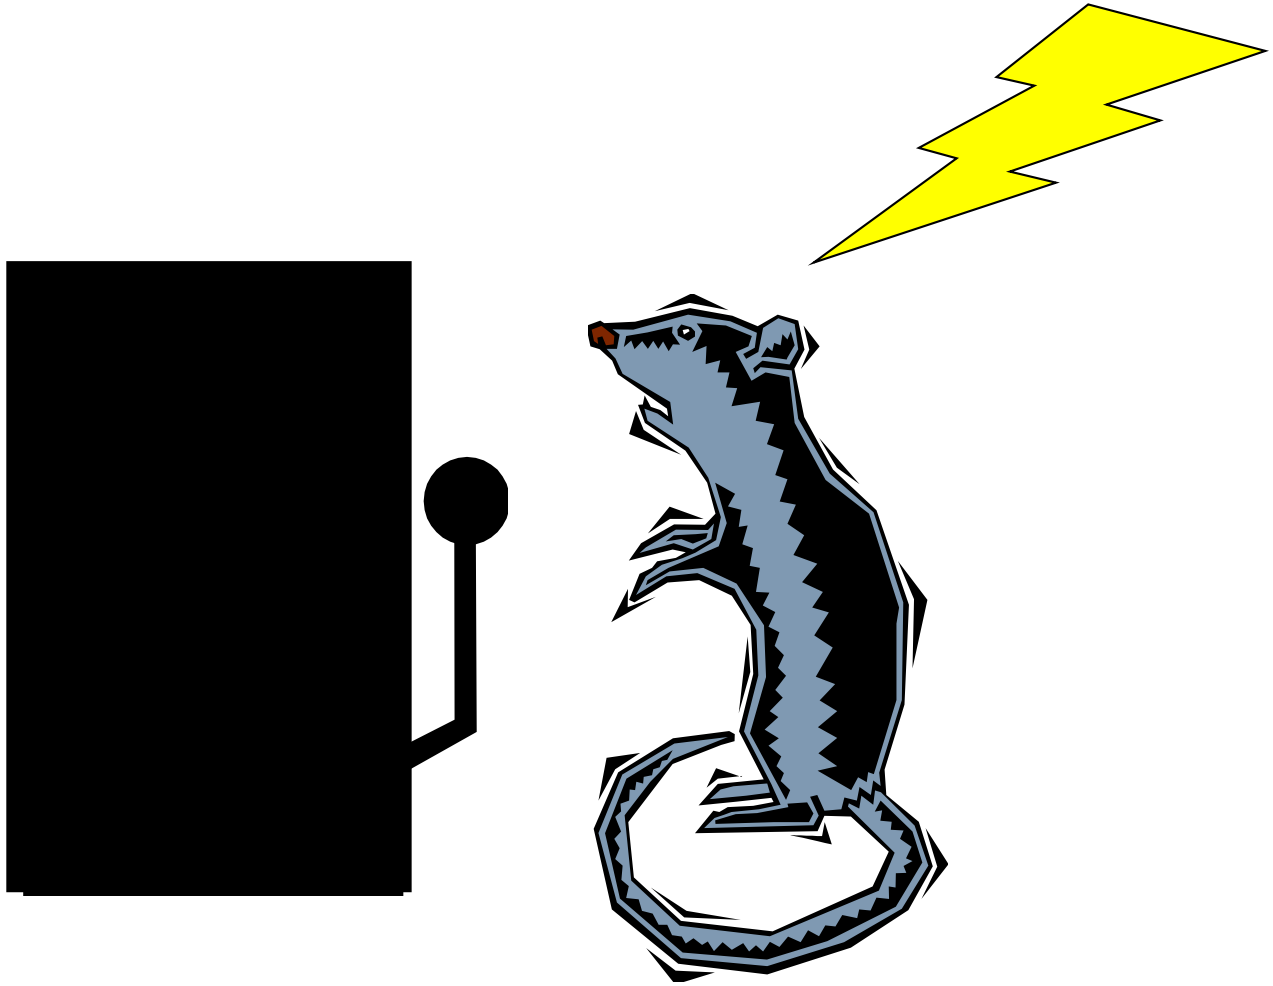

Olds & Milner, 1954
Routtenberg & Malsbury, 1969
Crow, 1972

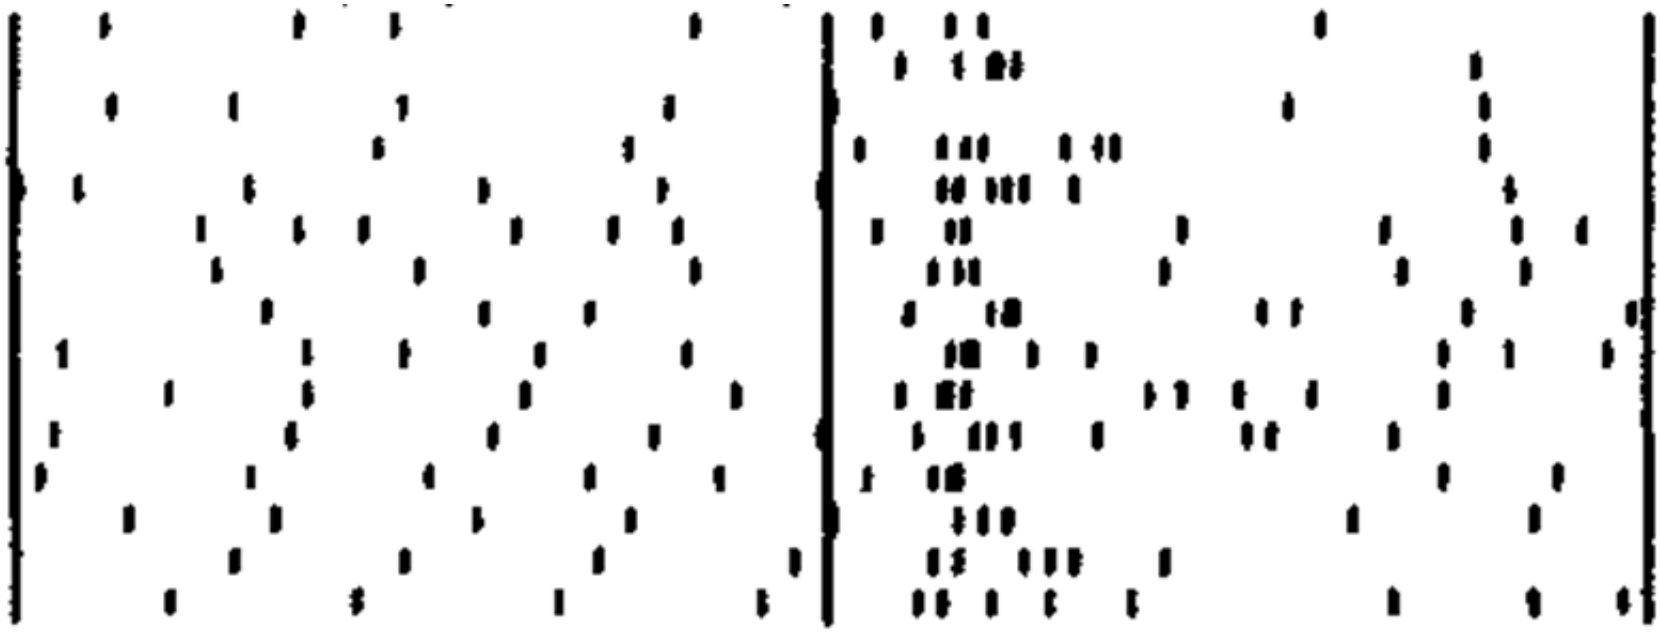


Juice

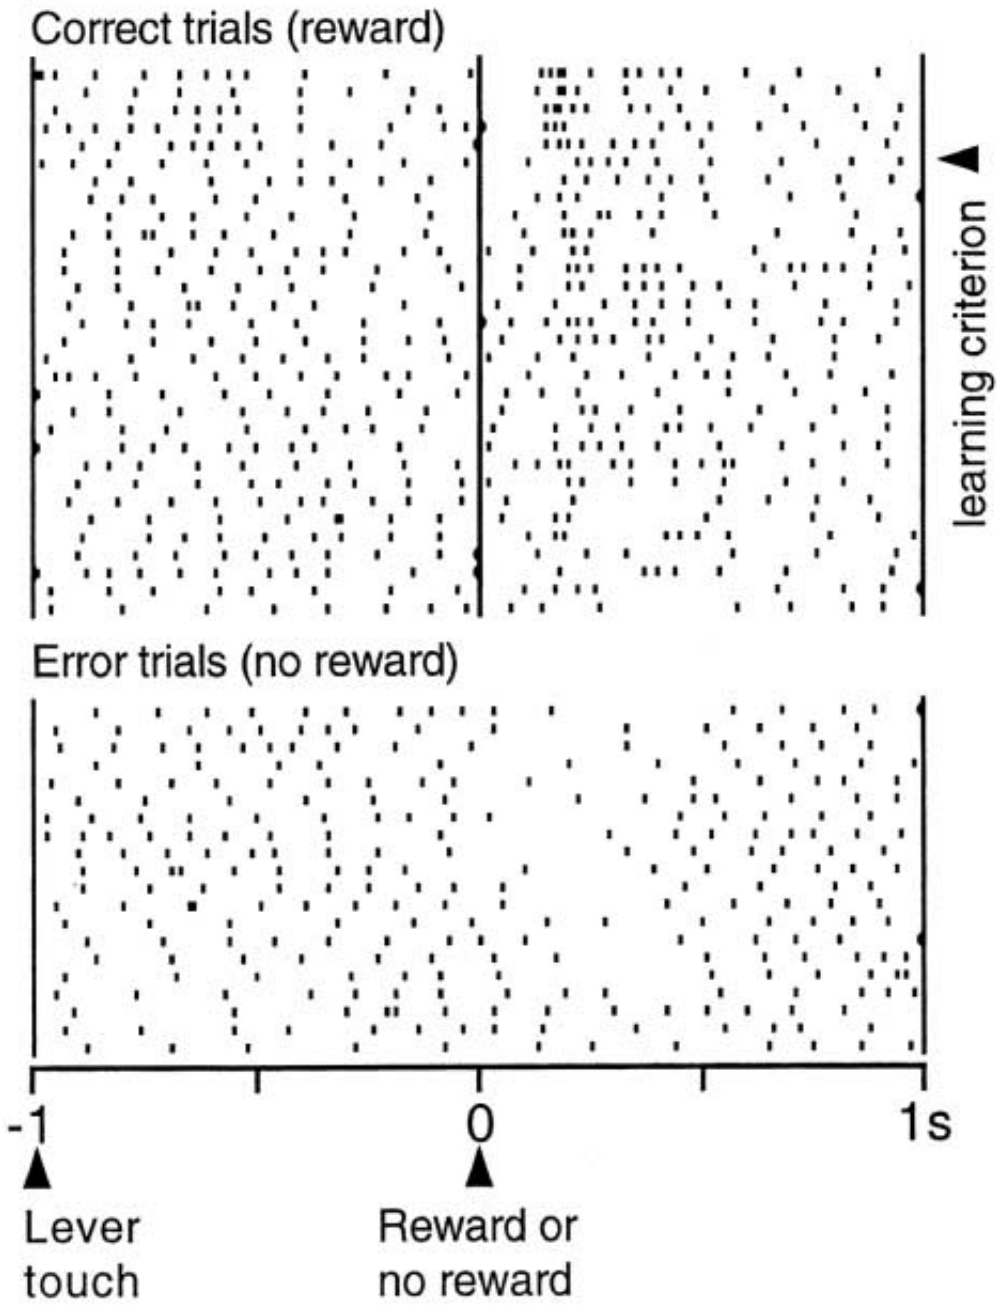
Correct trials (reward)
learning criterion
Error trials (no reward)
-1
0
1s
Lever
touch
Reward or
no reward

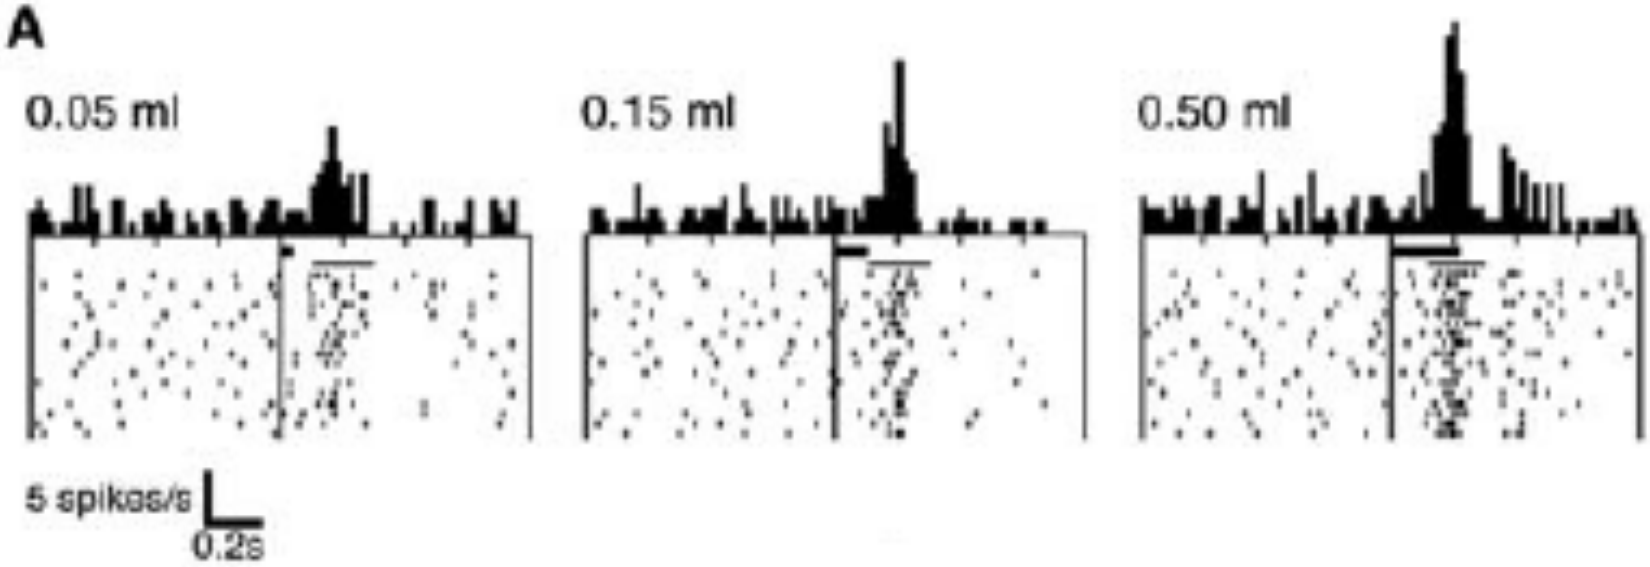
A
0.05 ml
0.15 ml
0.50 ml
5 spikes/s
0.2s

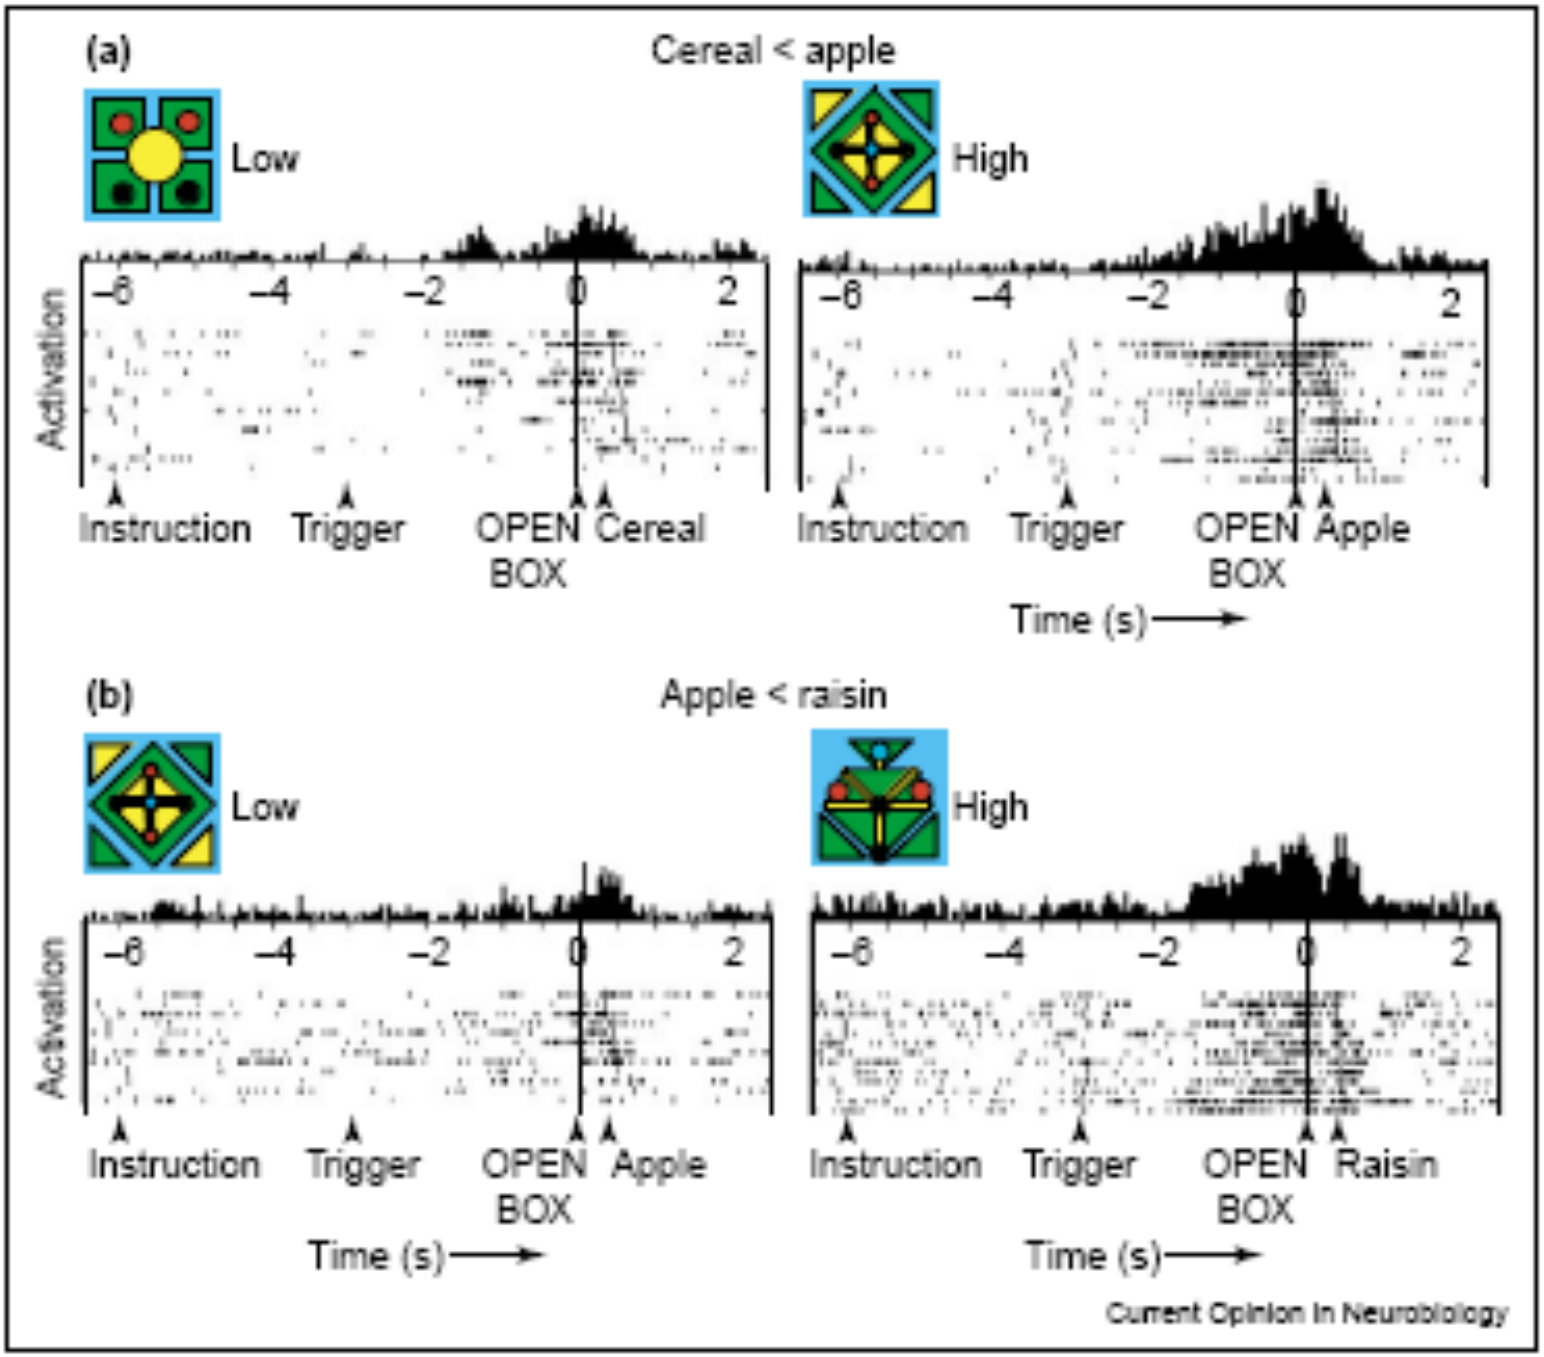
(a)
Cereal < apple
Low
High
Activation
–6
–4
–2
0
2
Instruction
Trigger
OPEN BOX
Cereal
Apple
Time (s)
(b)
Apple < raisin
Low
High
Activation
–6
–4
–2
0
2
Instruction
Trigger
OPEN BOX
Apple
Raisin
Time (s)
Current Opinion in Neurobiology

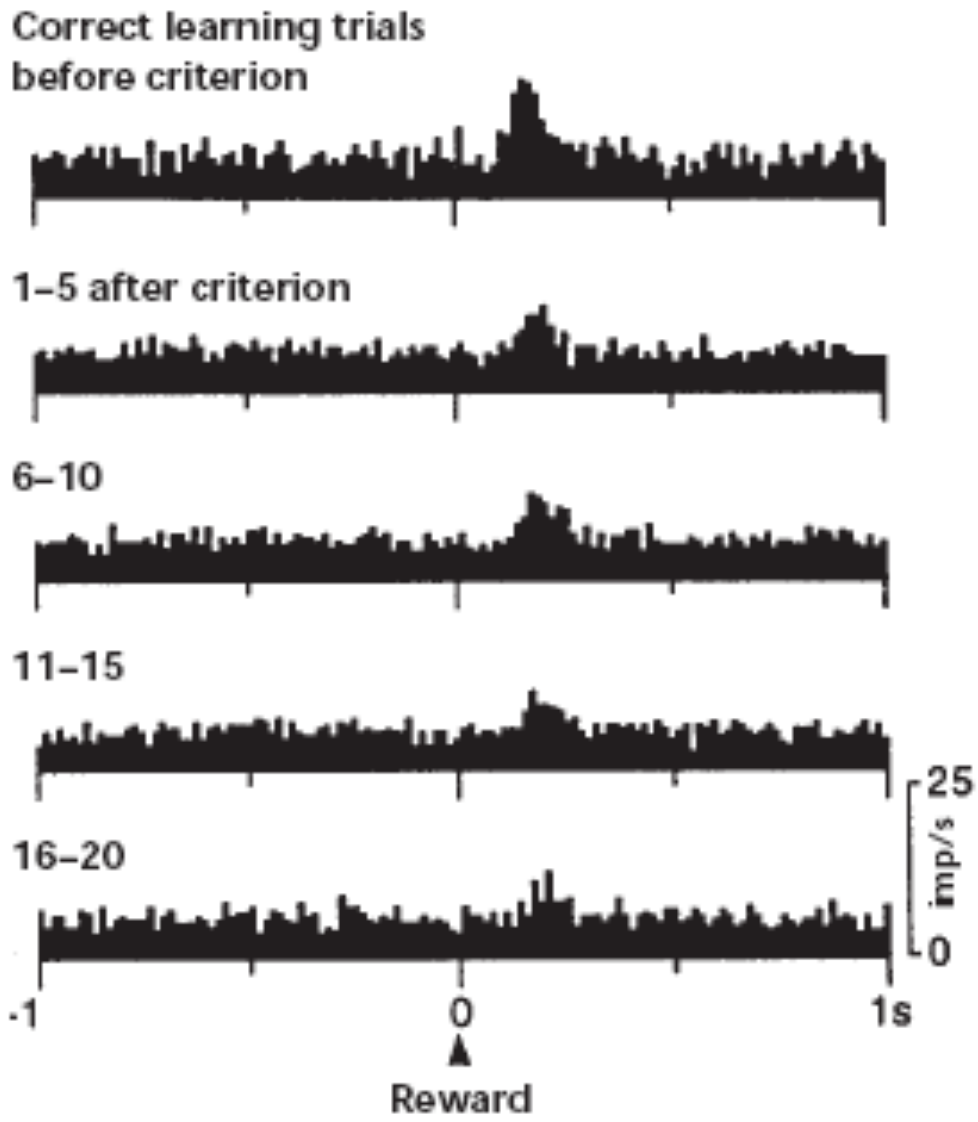
Correct learning trials
before criterion
1–5 after criterion
6–10
11–15
16–20
25
imp/s
0
-1
0
1s
Reward

Classical Conditioning

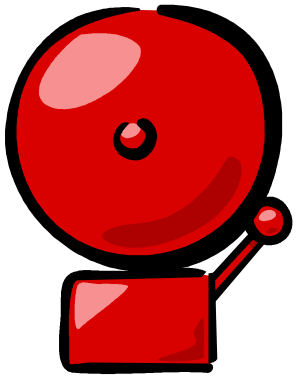

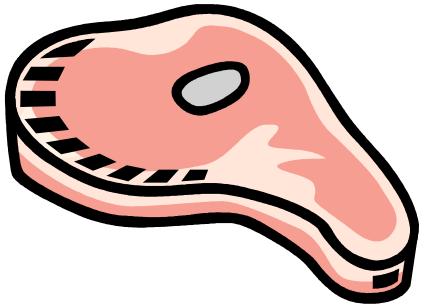

Classical Conditioning

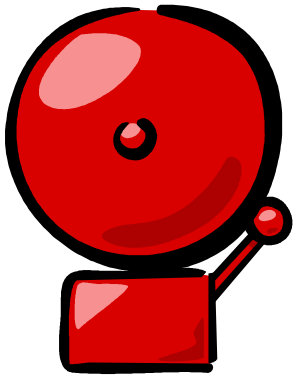

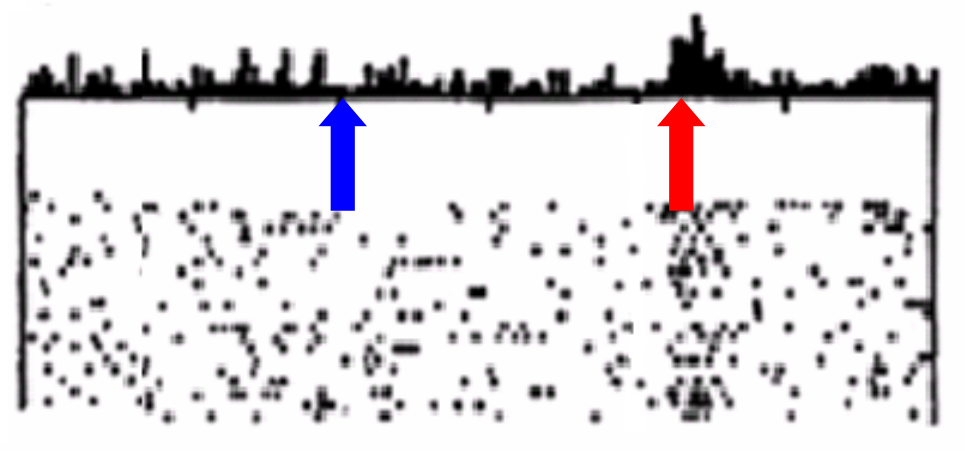

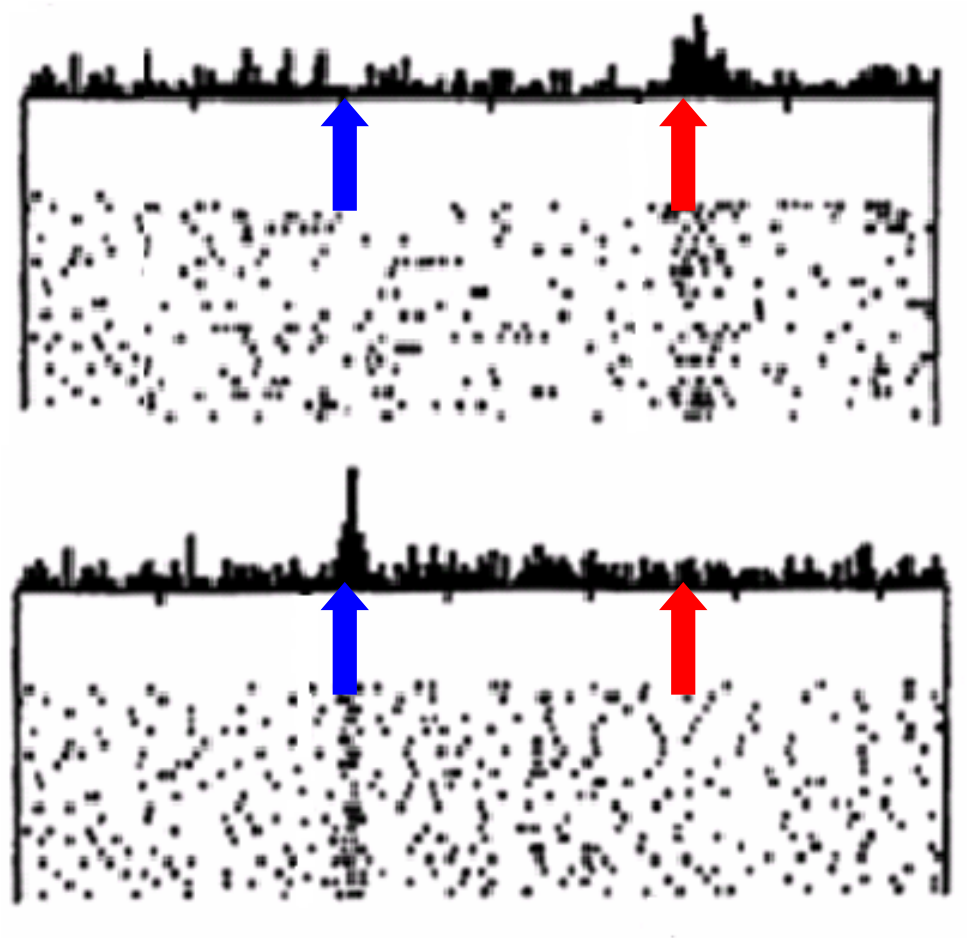

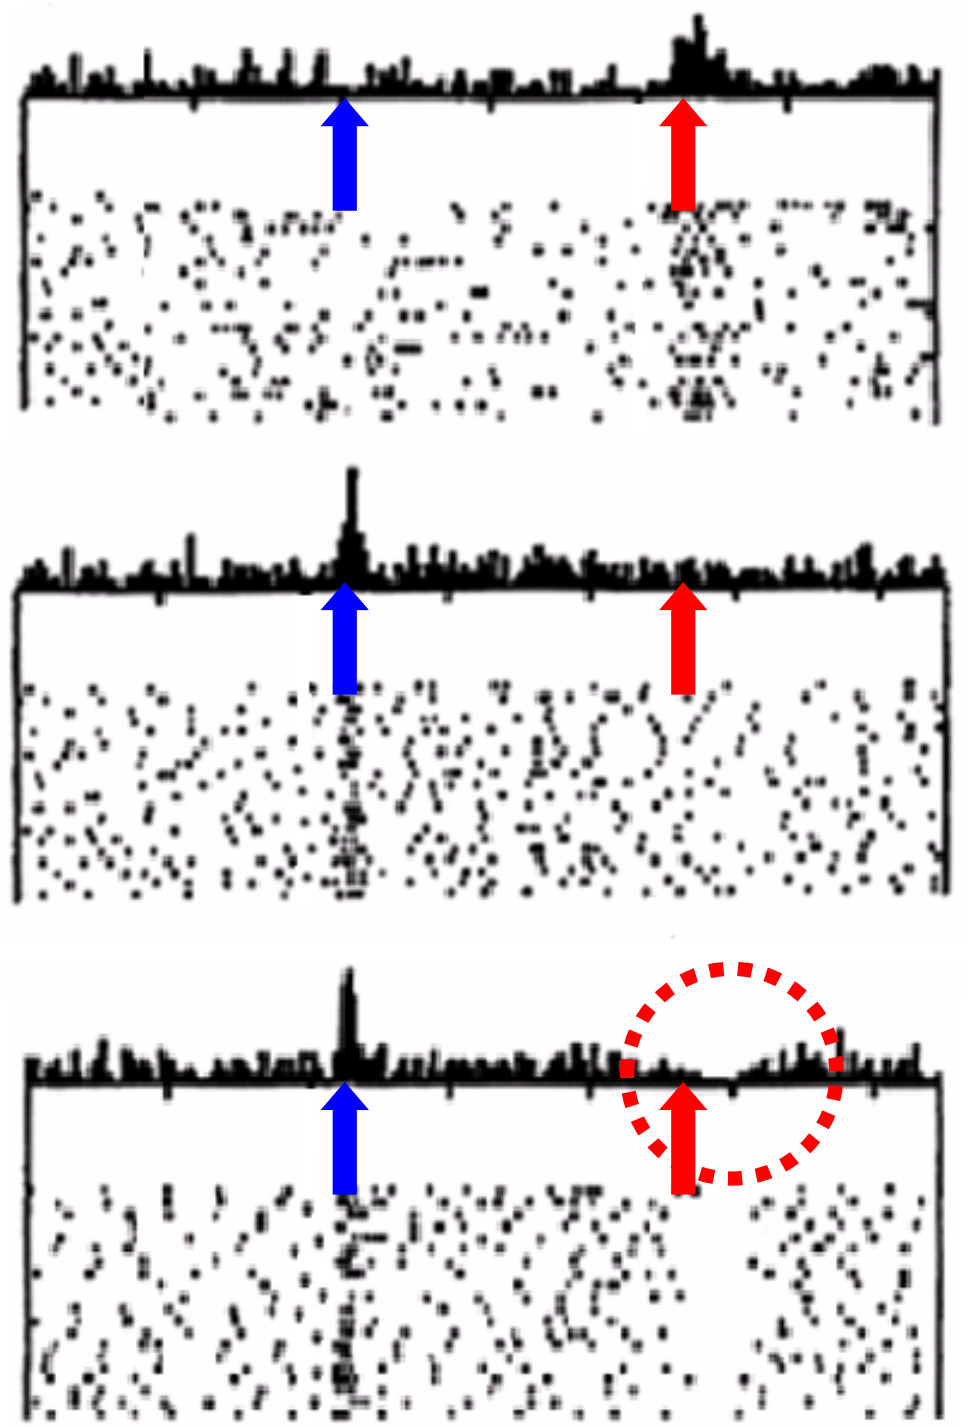

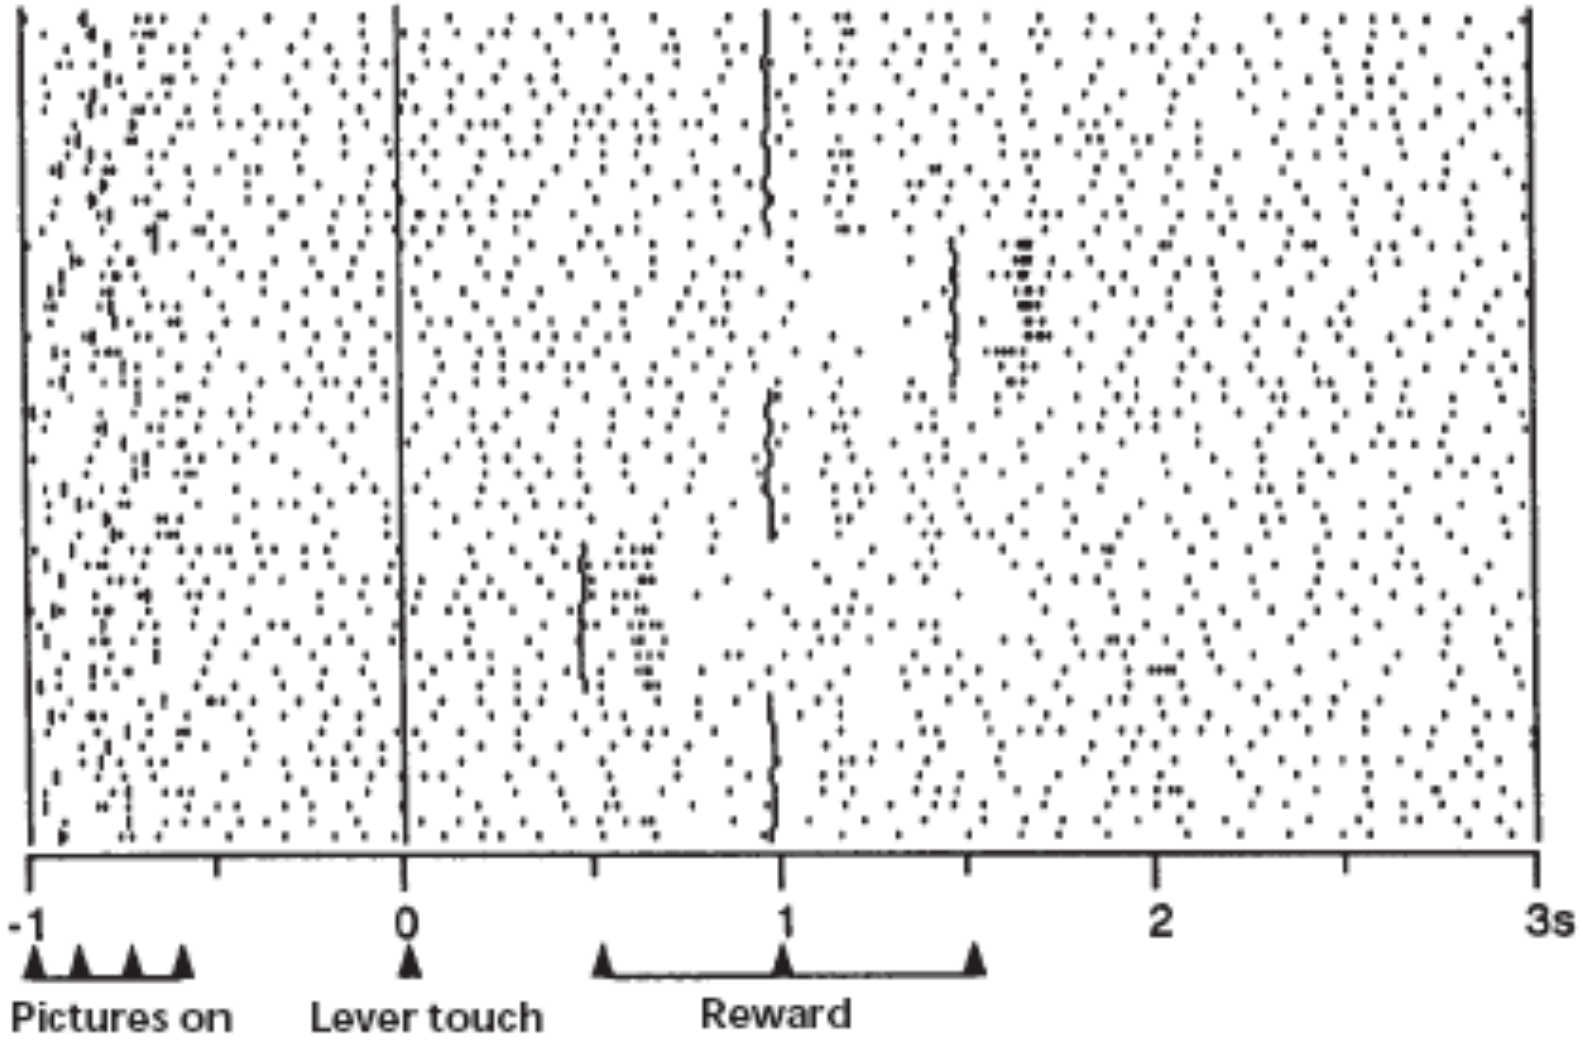
-1
0
1
2
3s
Pictures on
Lever touch
Reward

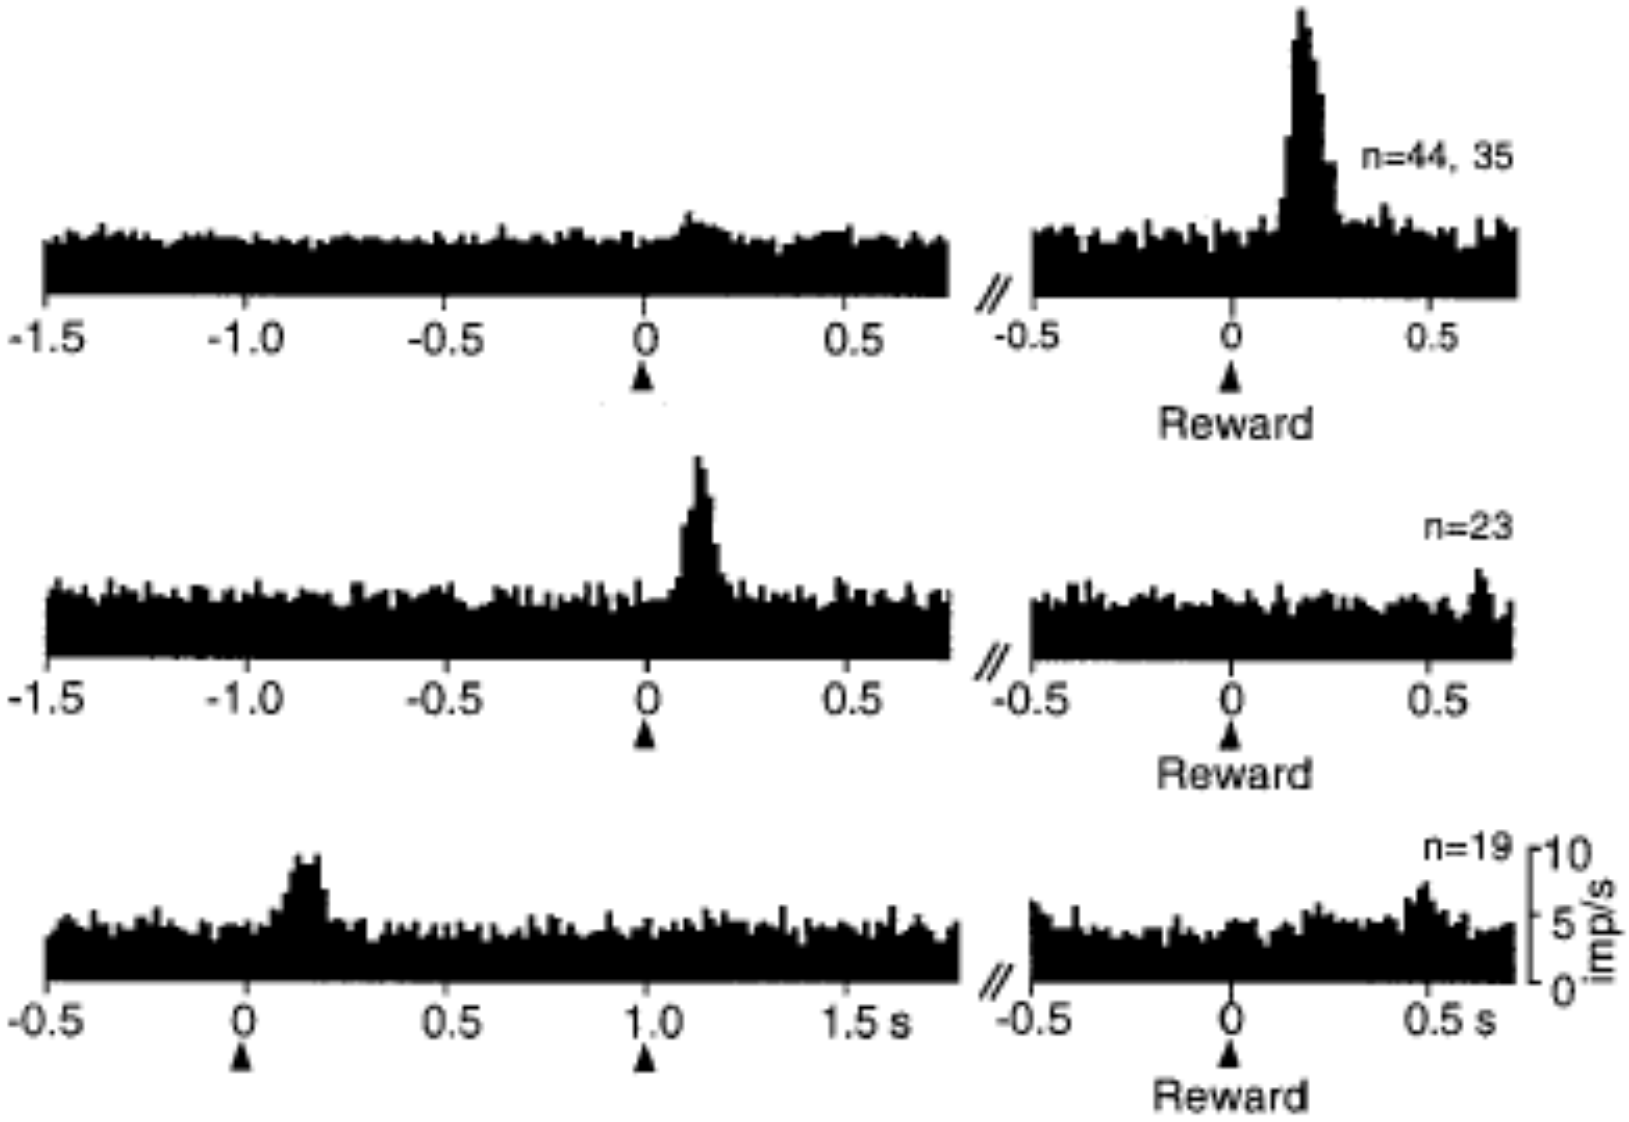
n=44, 35
-1.5
-1.0
-0.5
0
0.5
-0.5
0
0.5
Reward
n=23
-1.5
-1.0
-0.5
0
0.5
-0.5
0
0.5
Reward
n=19
10
5
0
imp/s
-0.5
0
0.5
1.0
1.5 s
-0.5
0
0.5 s
Reward

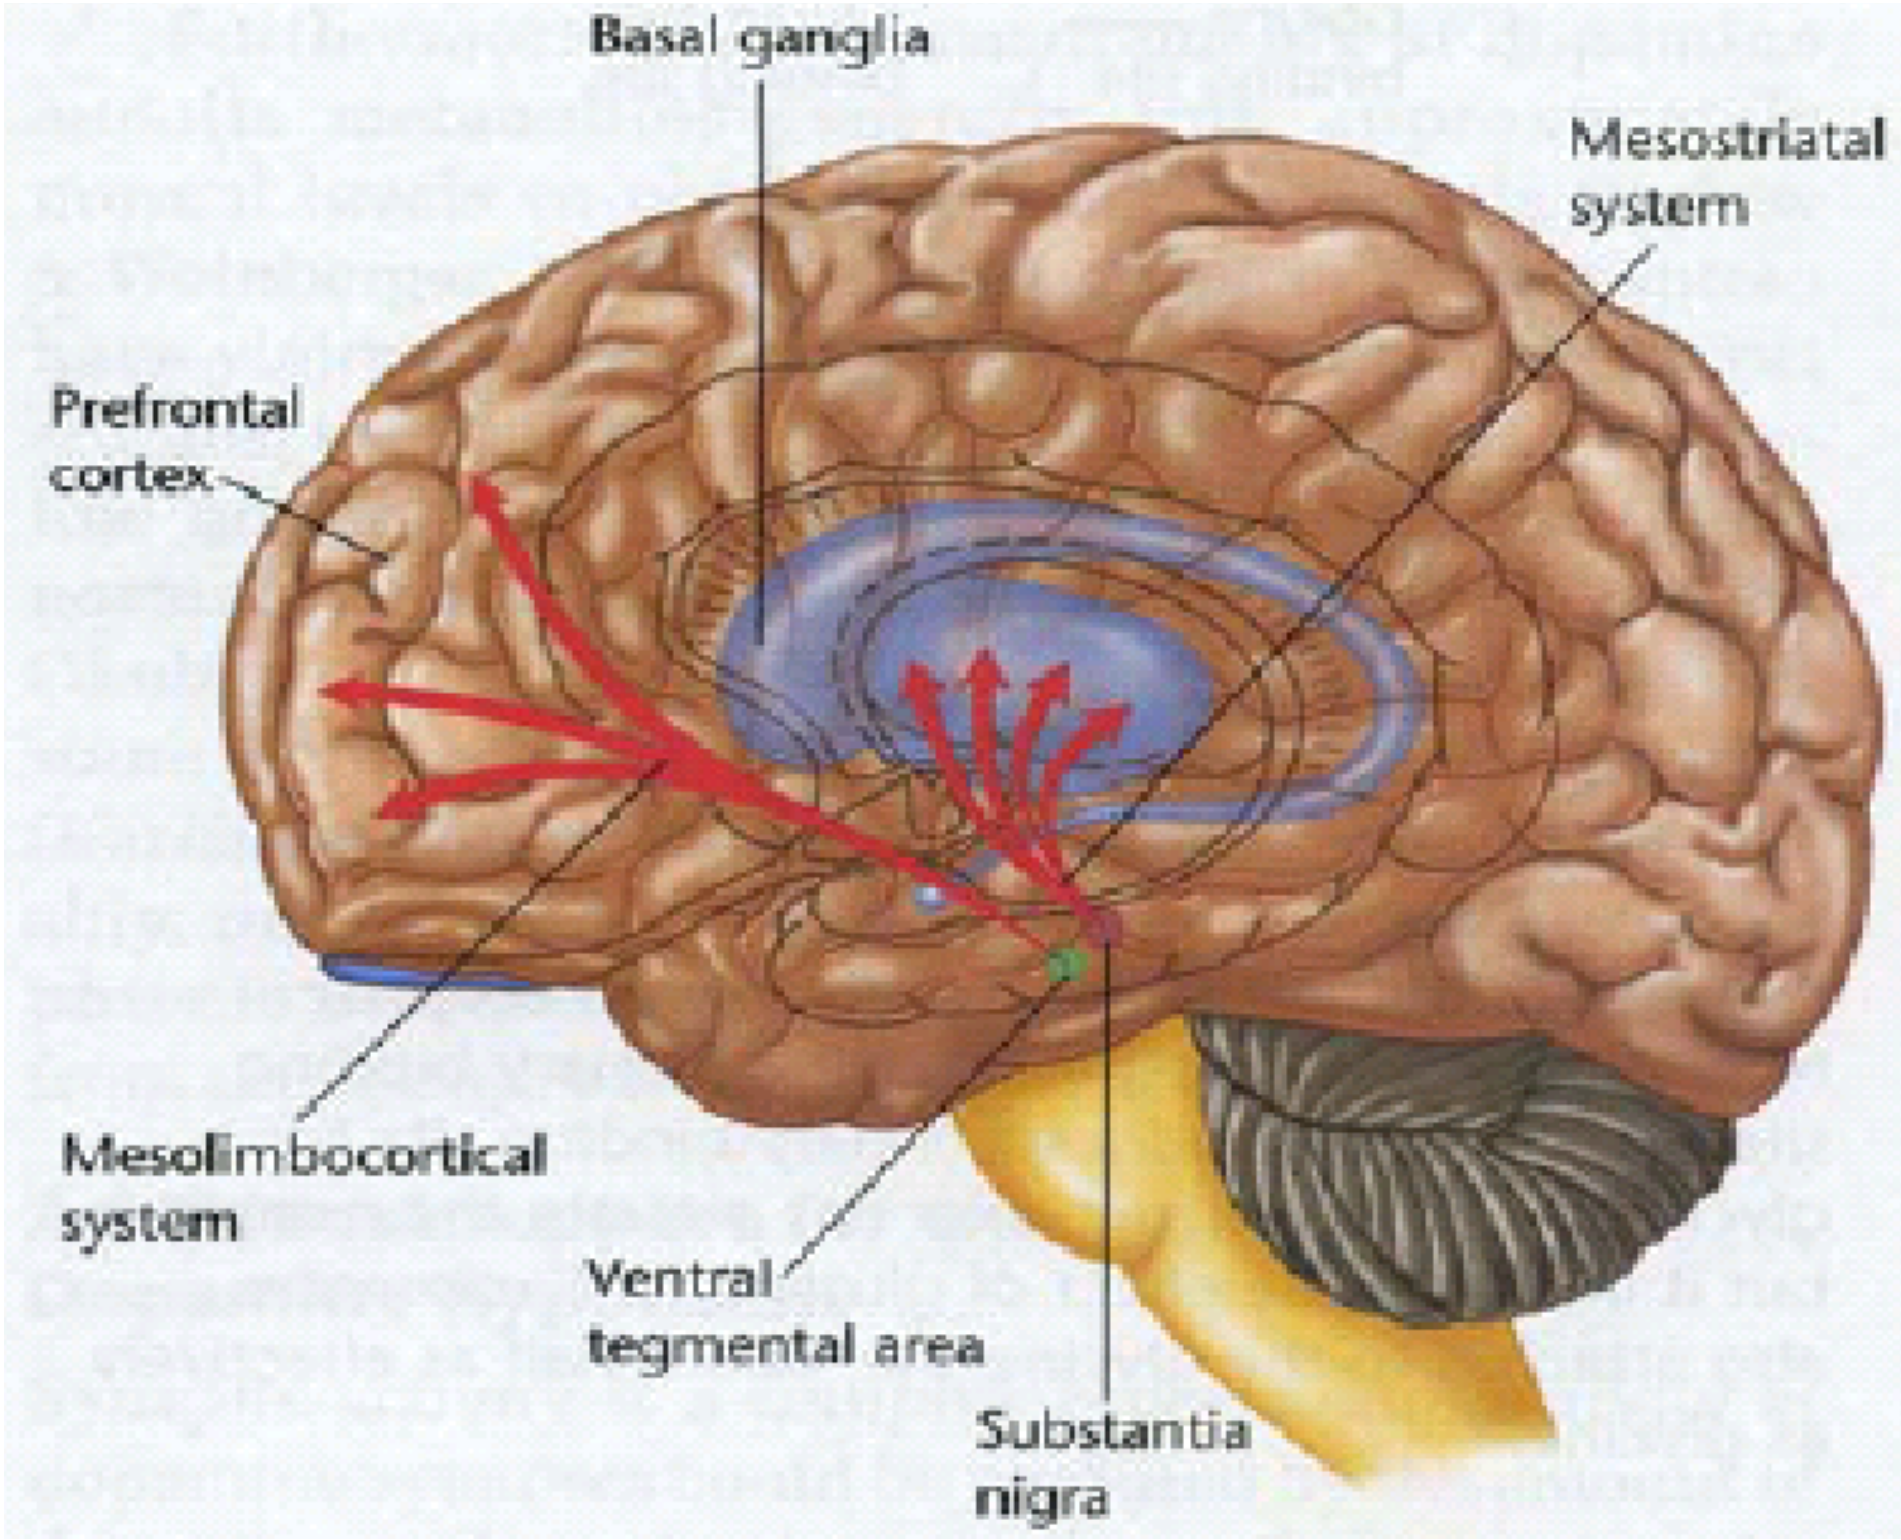
Basal ganglia
Mesostriatal system
Prefrontal cortex
Mesolimbocortical system
Ventral tegmental area
Substantia nigra

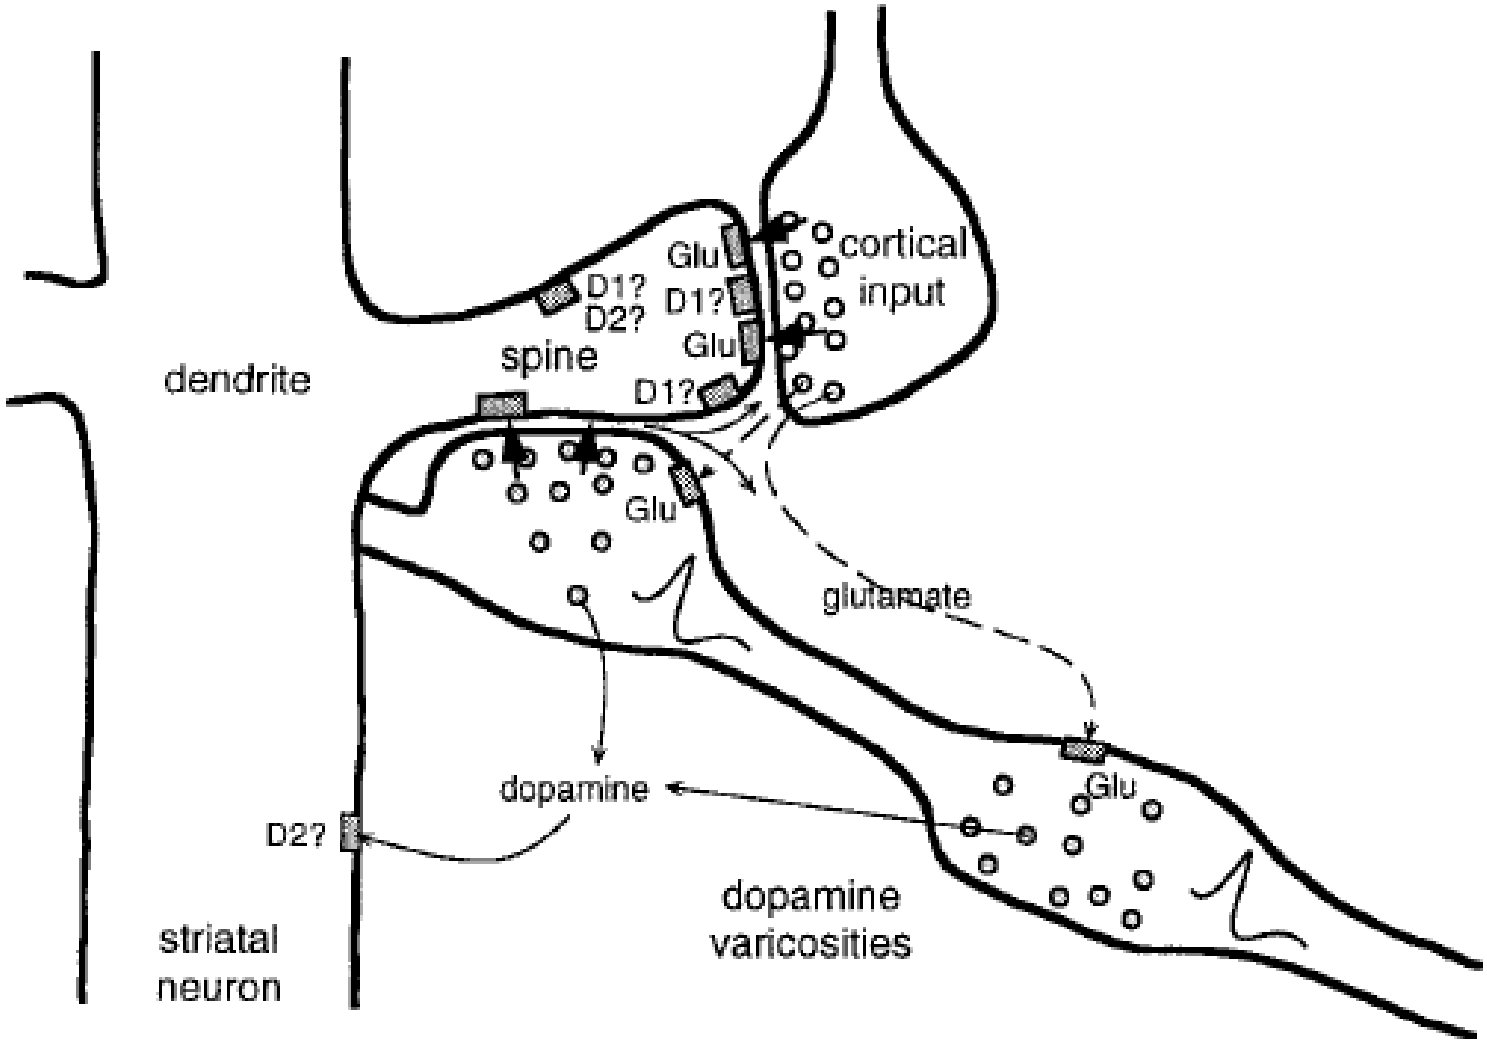
Glu
cortical
input
D1?
D2?
D1?
Glu
spine
dendrite
D1?
Glu
glutamate
dopamine
Glu
D2?
dopamine
varicosities
striatal
neuron

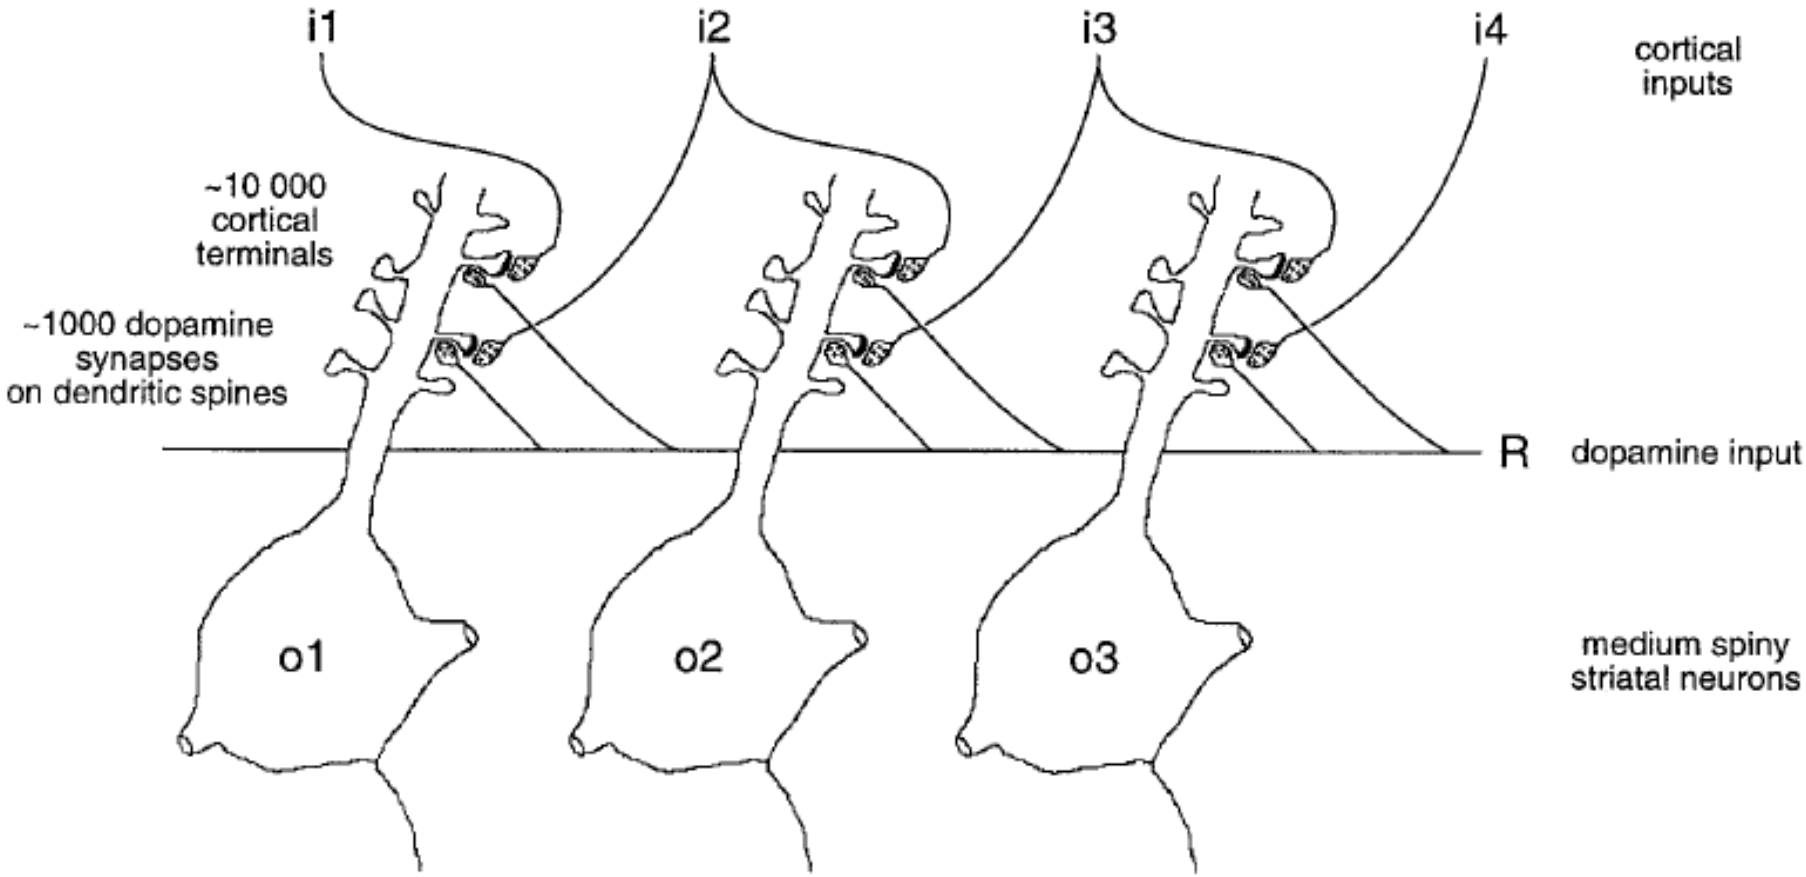
i1
i2
i3
i4
cortical inputs
~10 000 cortical terminals
~1000 dopamine synapses on dendritic spines
R
dopamine input
o1
o2
o3
medium spiny striatal neurons

Thorndike's Law of Effect

Actions that are followed by feelings of satisfaction have a greater likelihood of being generated again in the future, whereas actions that are followed by feelings of dissatisfaction have a lesser likelihood of being generated again in the future.

S **R**

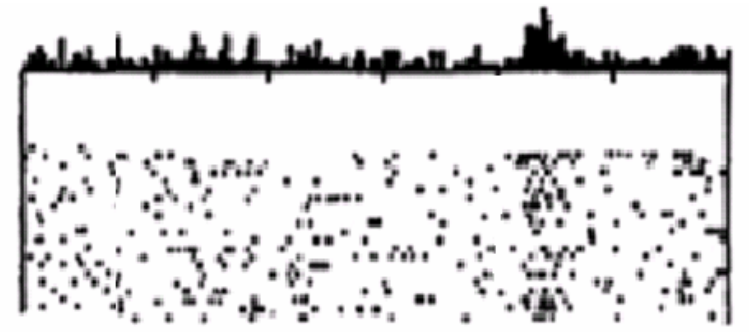

Early in learning, $V_S = 0$

Thus, when juice is given:

$$\delta(t) = r(t) + \gamma V(t) - V(t-1)$$

$$\delta(t) = 1 + 0 - 0$$

S **R**

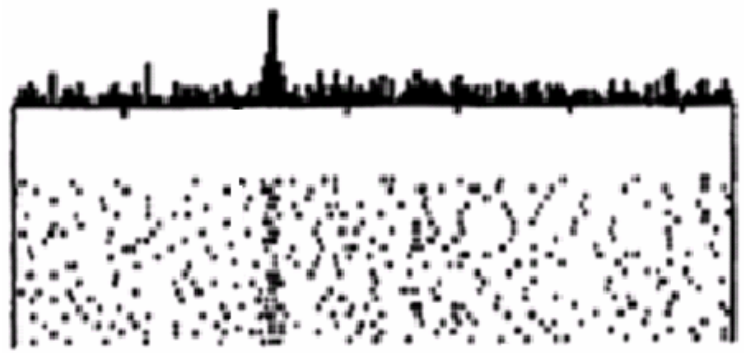

After learning, $V_S > 0$

Thus, when juice is given:

$$\delta(t) = r(t) + \gamma V(t) - V(t-1)$$

$$\delta(t) = 1 + 0 - 1$$

$$\delta(t) = 0.0$$

S R

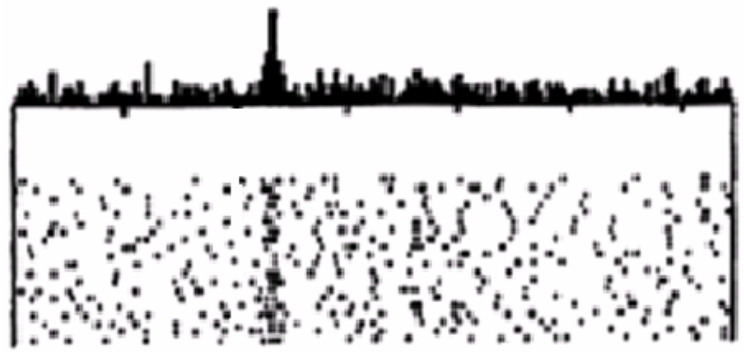

However, when the stimulus is encountered

$$V_{S-1} = 0$$

$$\delta(t) = r(t) + \gamma\, V(t) - V(t-1)$$

$$\delta(t) = 0 + 1 - 0$$

$$\delta(t) = 1.0$$

And when a predicted reward does not occur:

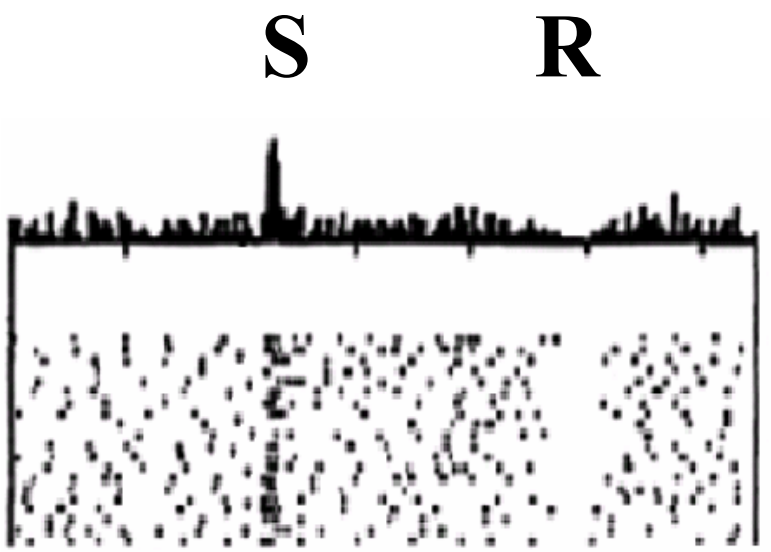

$$\delta(t) = r(t) + \gamma V(t) - V(t-1)$$

$$\delta(t) = 0 + 0 - 1.0$$

$$\delta(t) = -1.0$$

Part III

H H H H H

H H S H H

The **E**rror-**R**elated **N**egativity

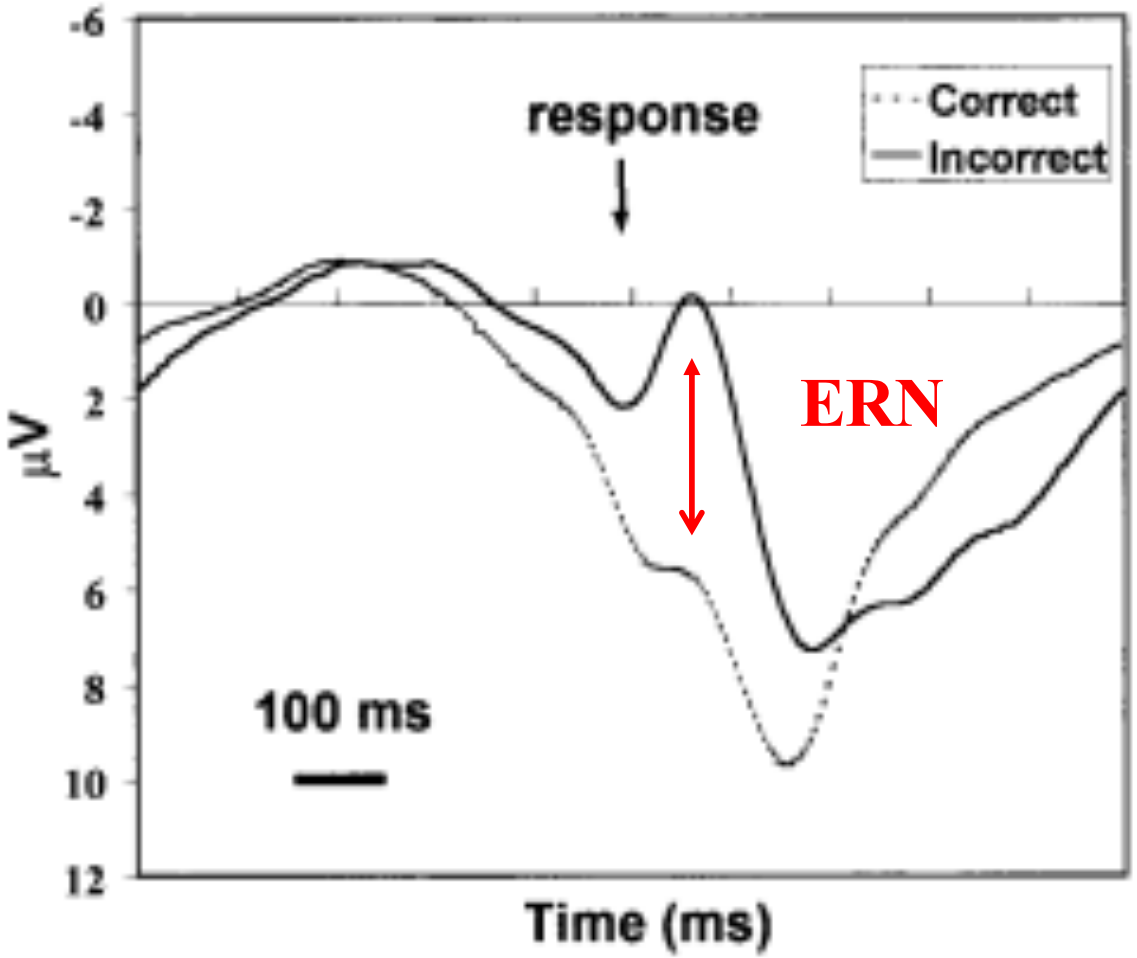

Beep!

Correct!

The **Feedback** Related Negativity

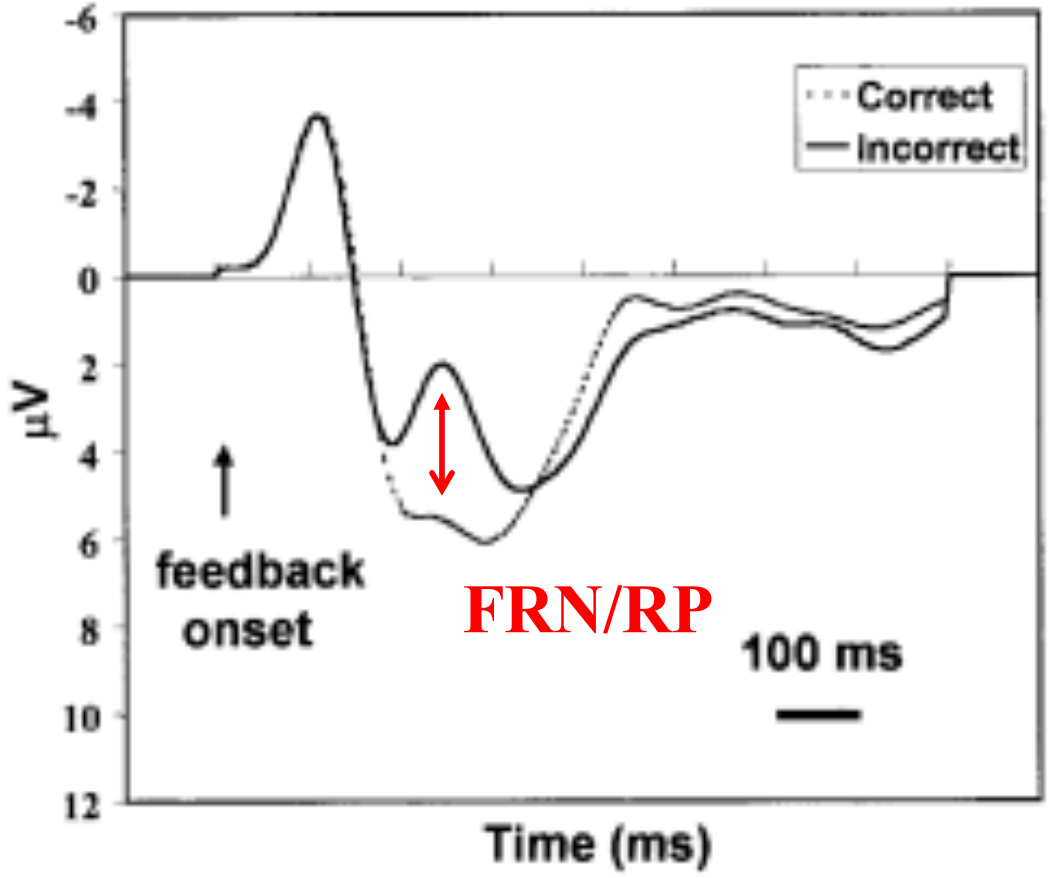

Miltner et al., 1997; Holroyd et al., 1998

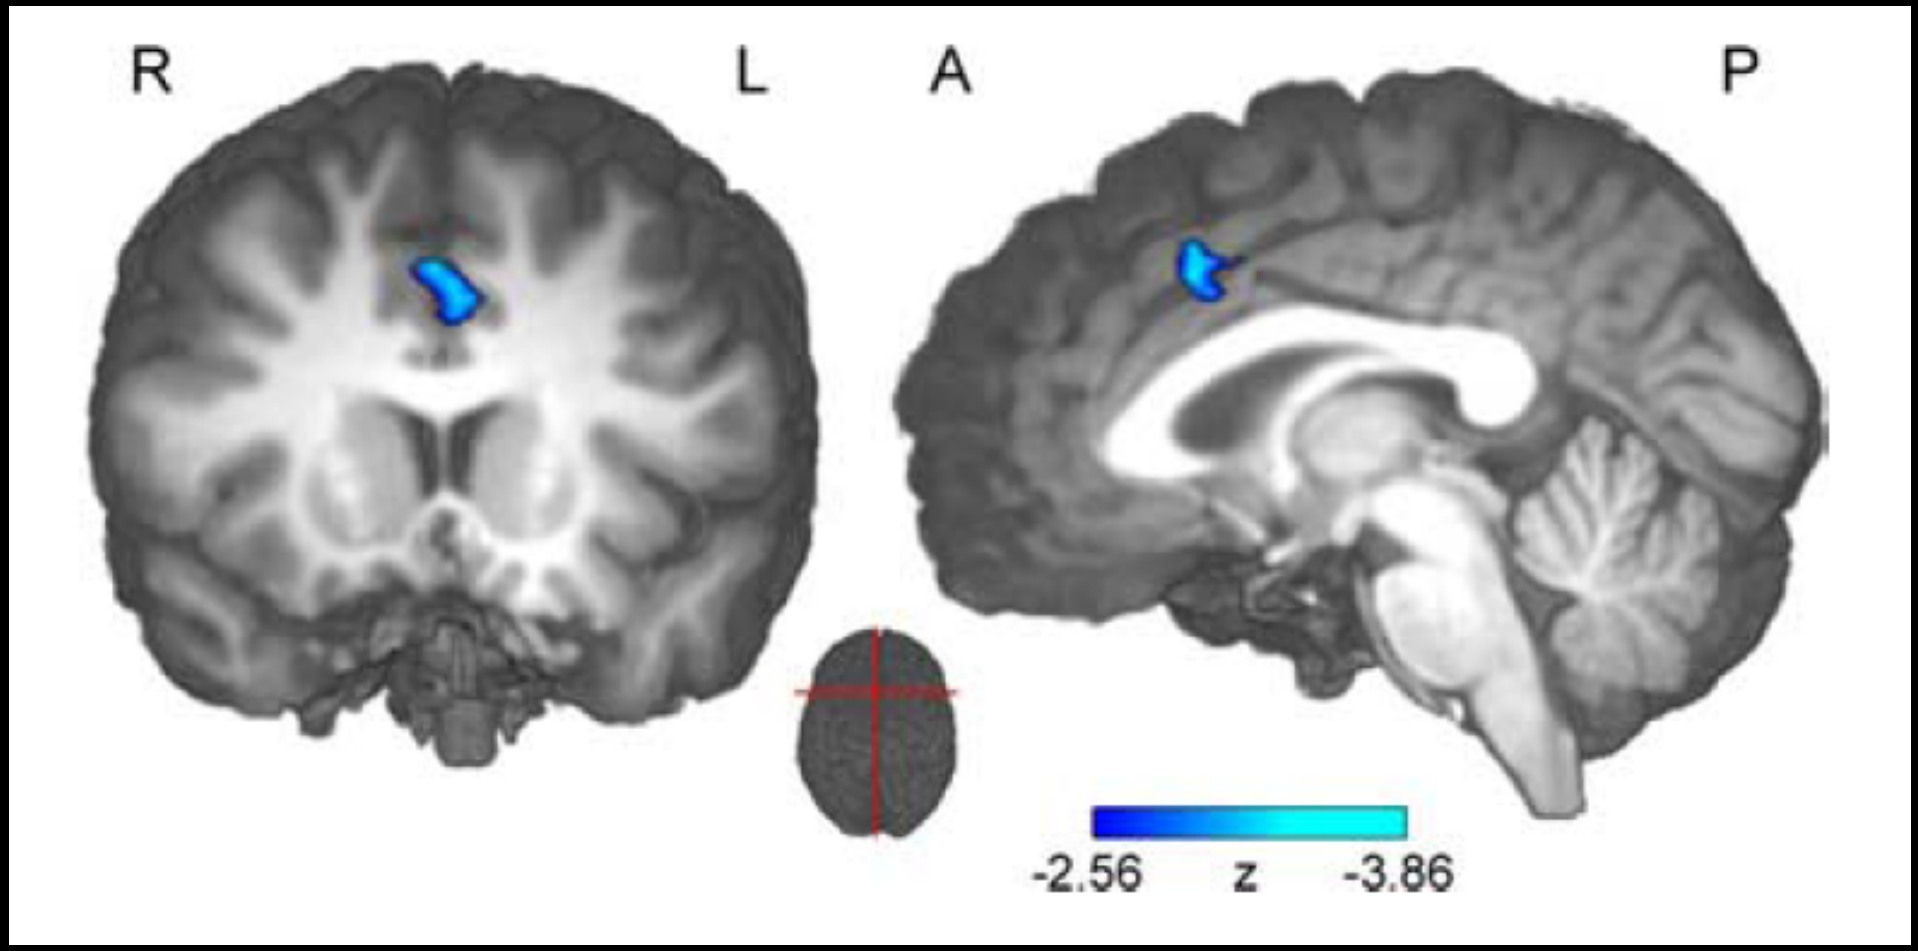

Anterior Cingulate Cortex is the probable source of the ERN

EEG: Debener et al., 2005; Holroyd et al., 2004 Monkey: Ito et al., 2003

(Holroyd & Coles, 2002)

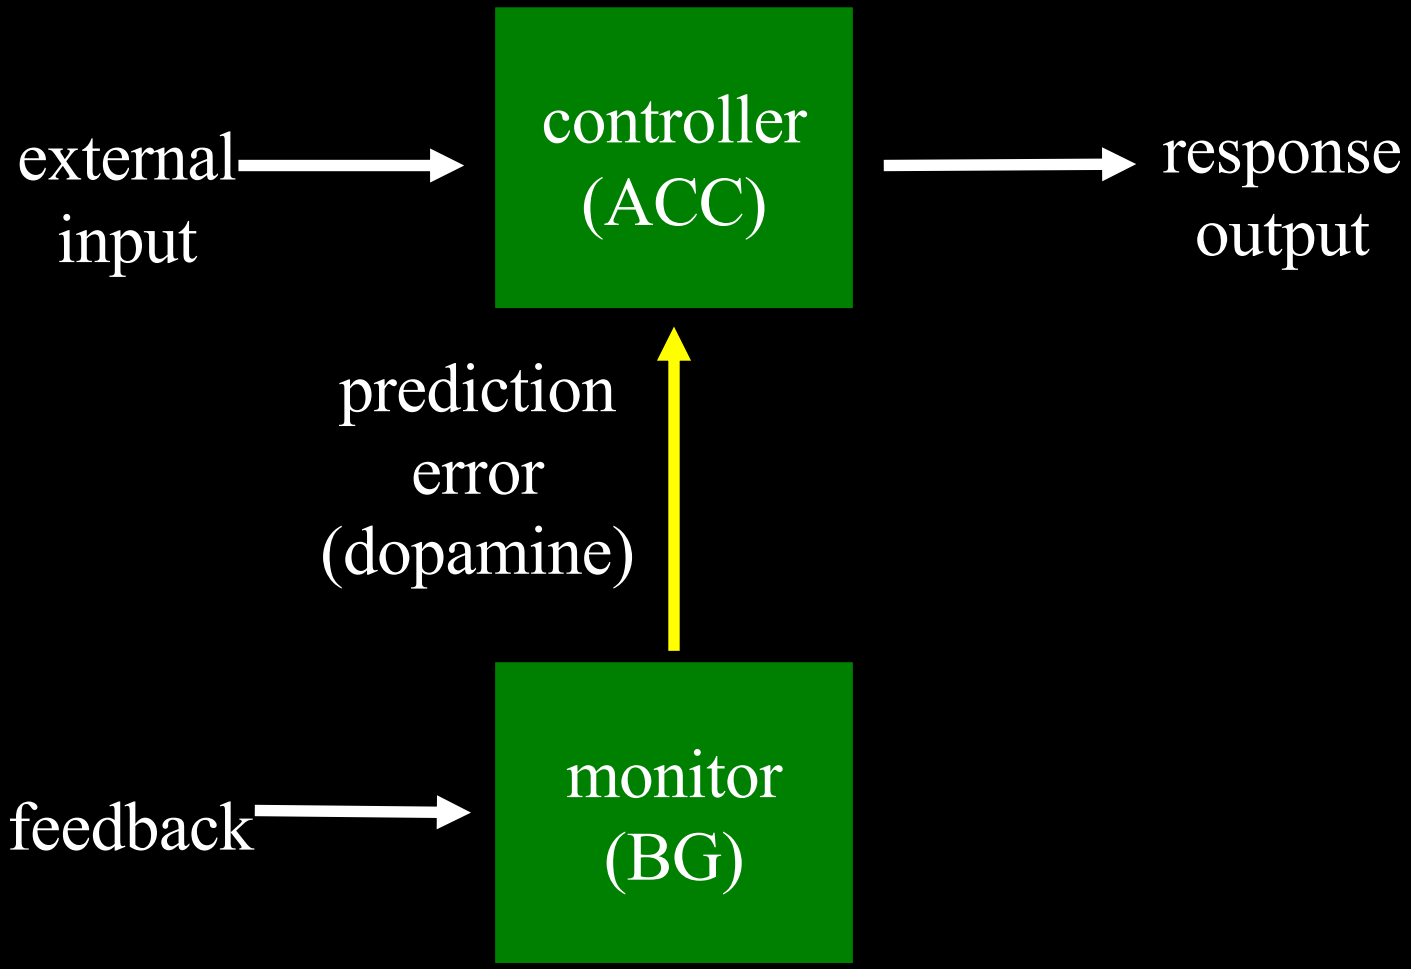

(Holroyd & Coles, 2002)

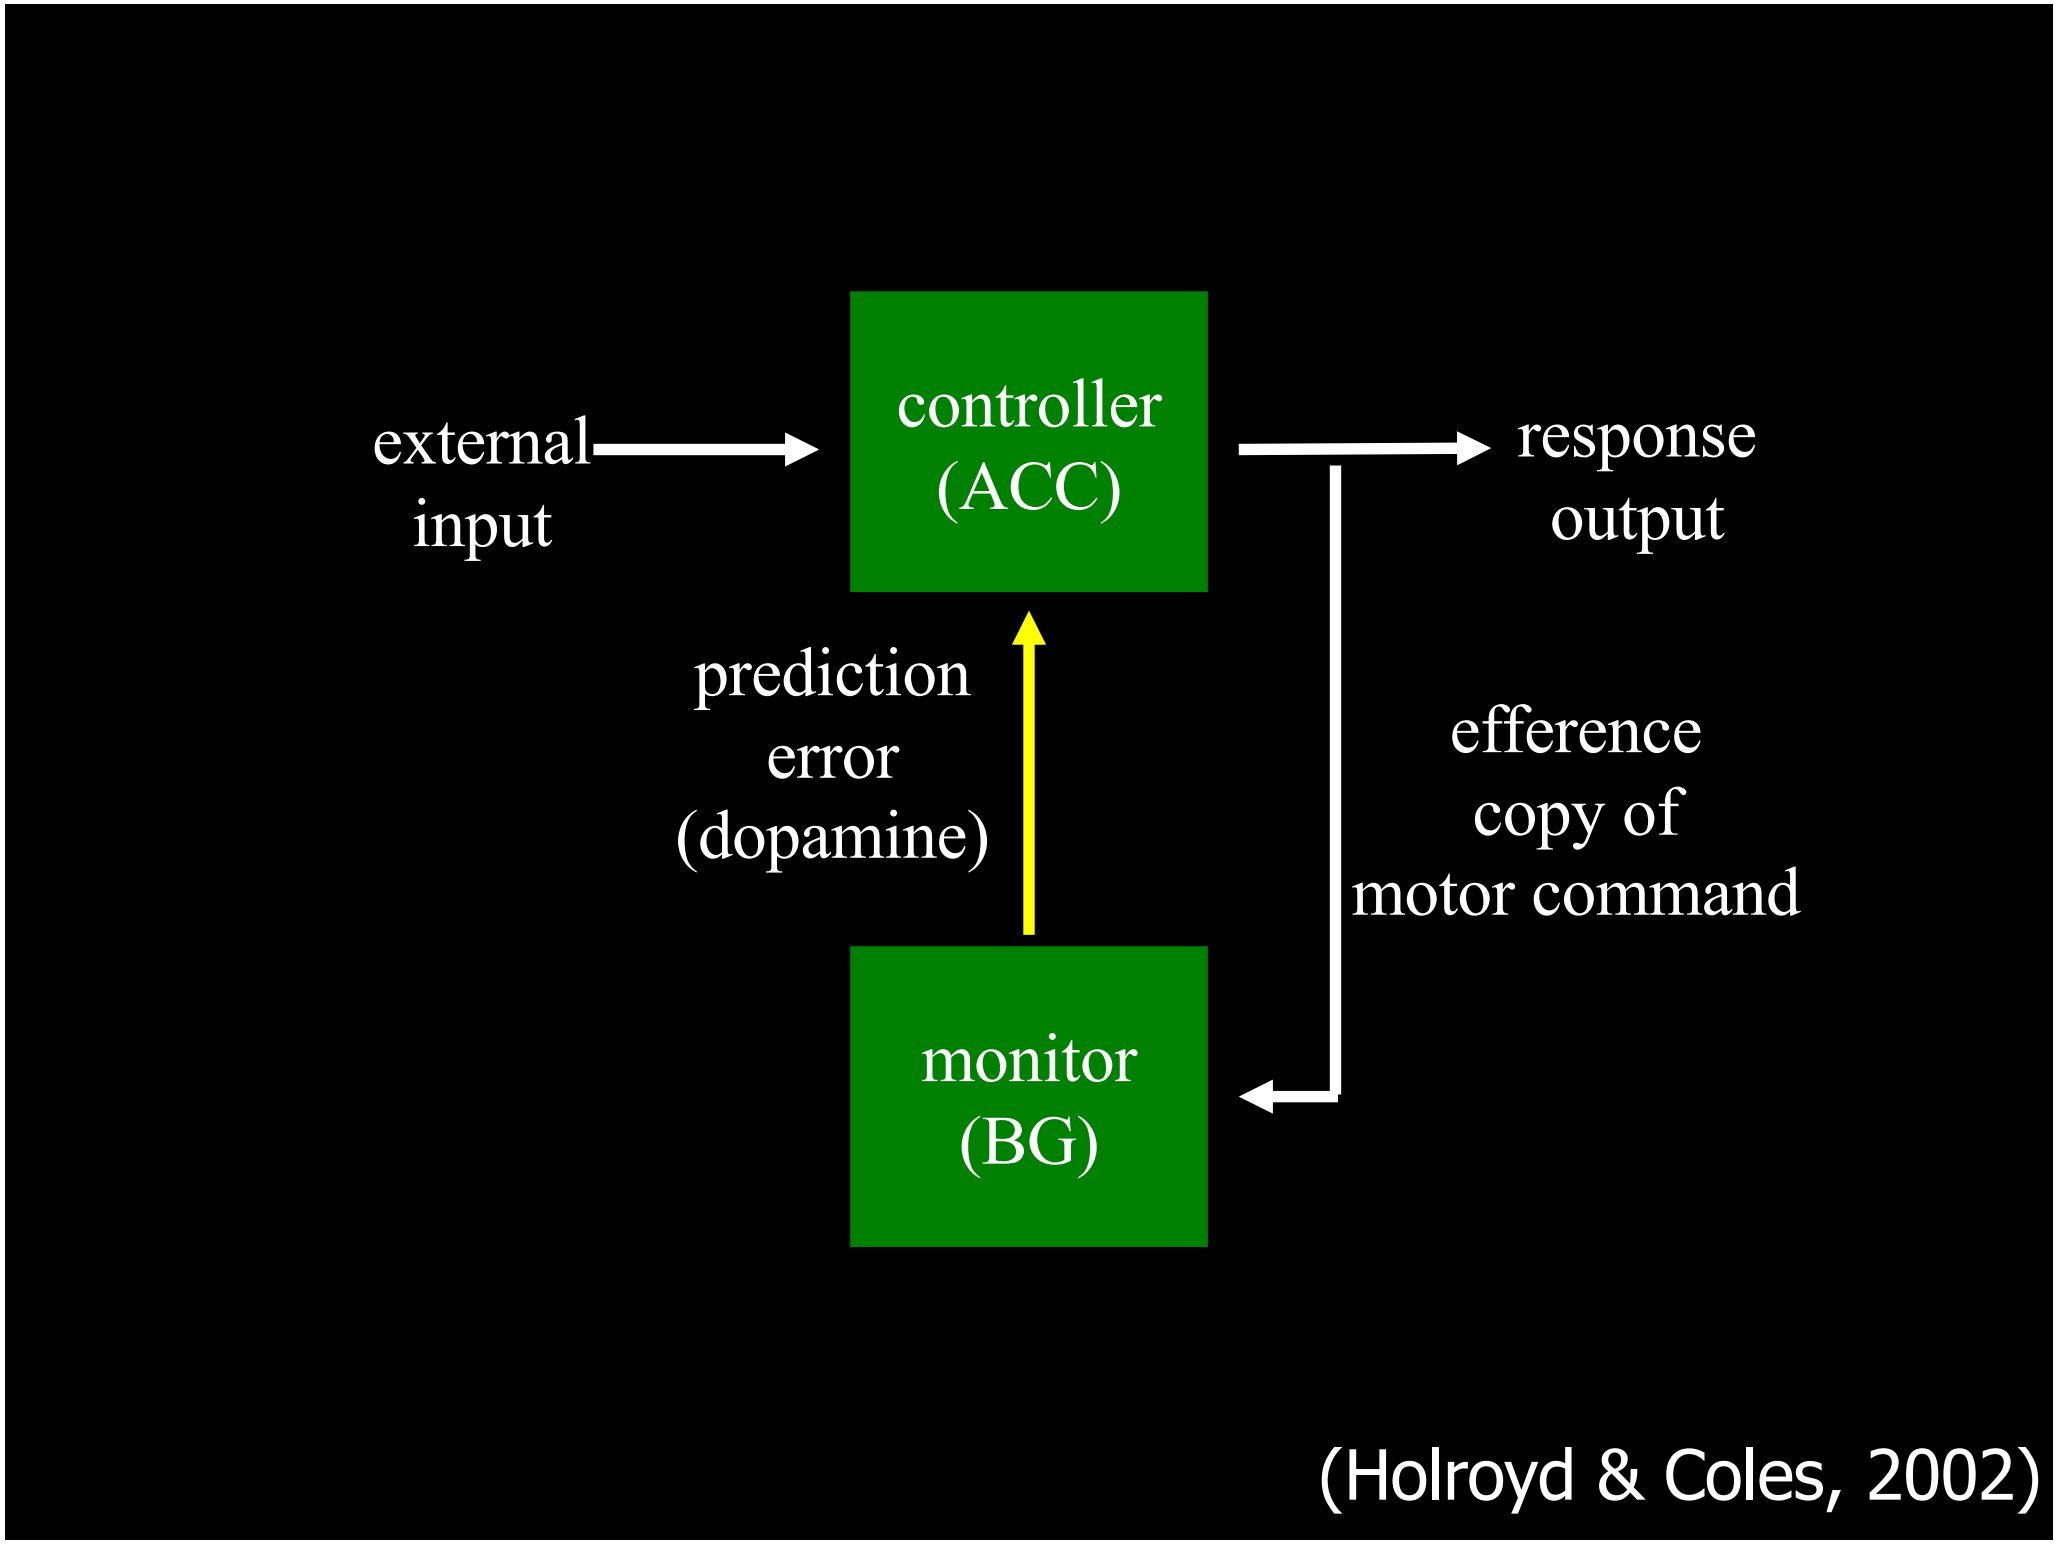
external
input
controller
(ACC)
response
output
prediction
error
(dopamine)
efference
copy of
motor command
monitor
(BG)
(Holroyd & Coles, 2002)

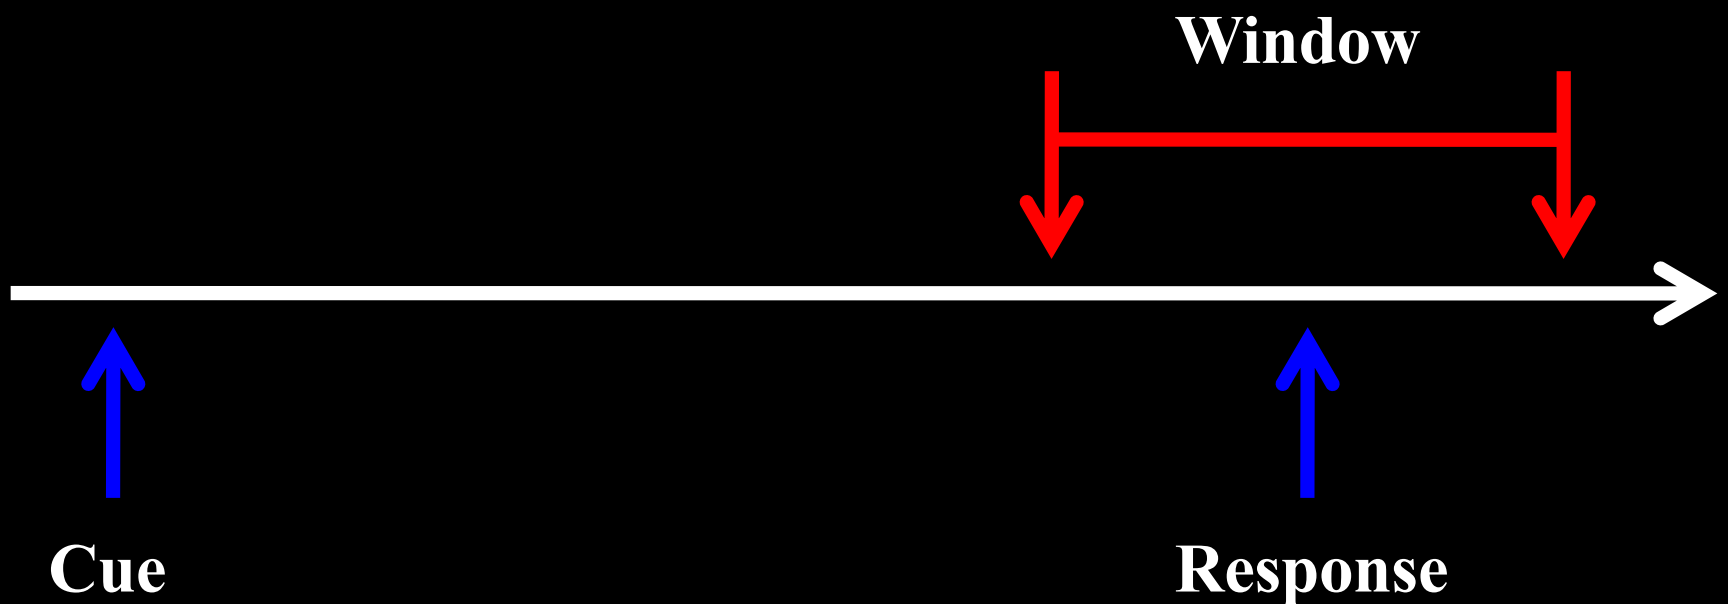
Window
Cue
Response

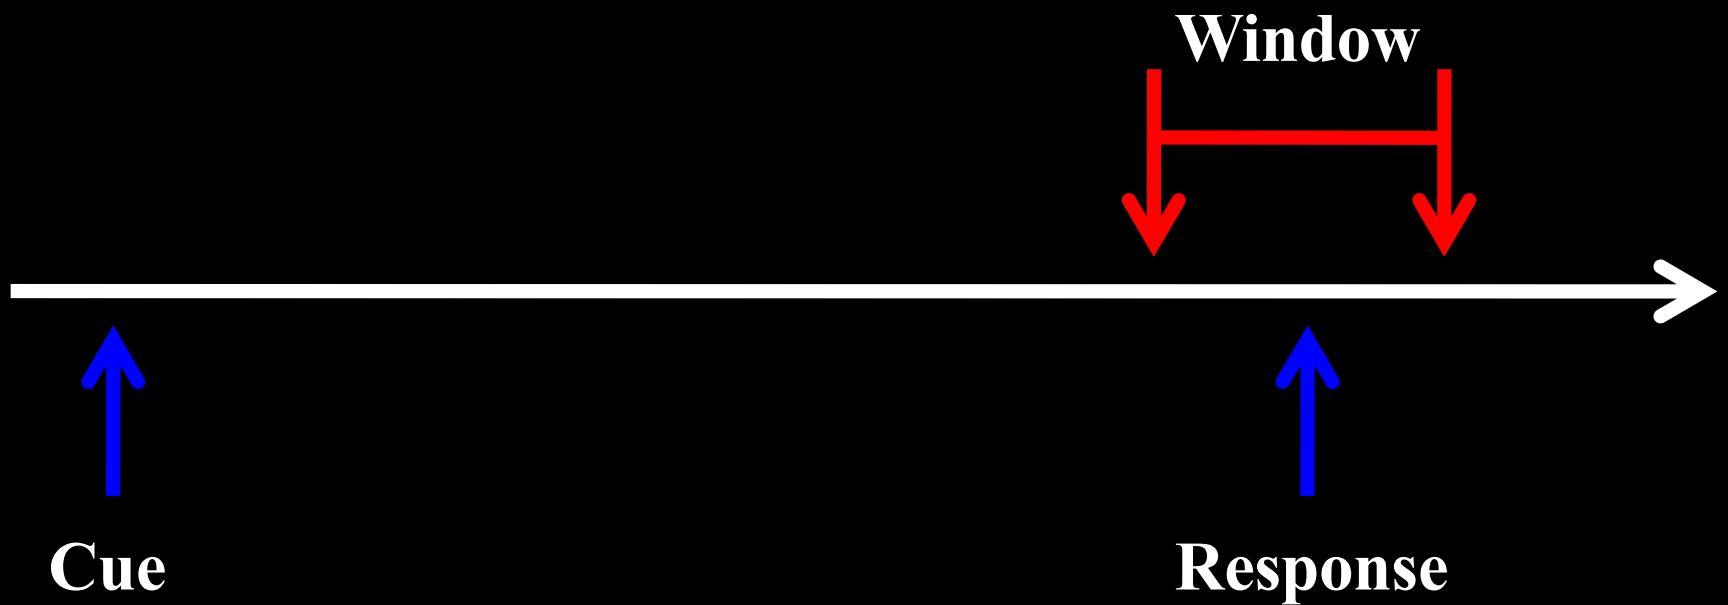
Window
Cue
Response

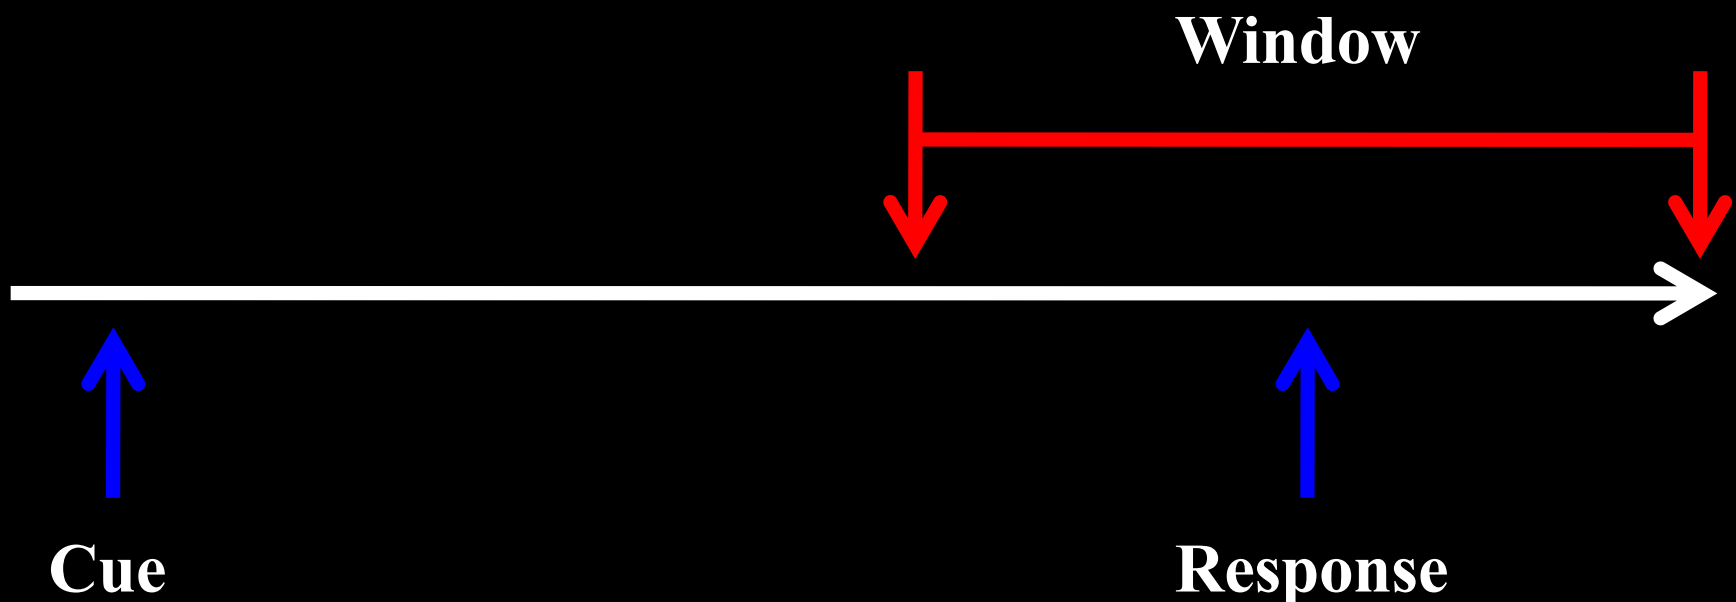
Window
Cue
Response

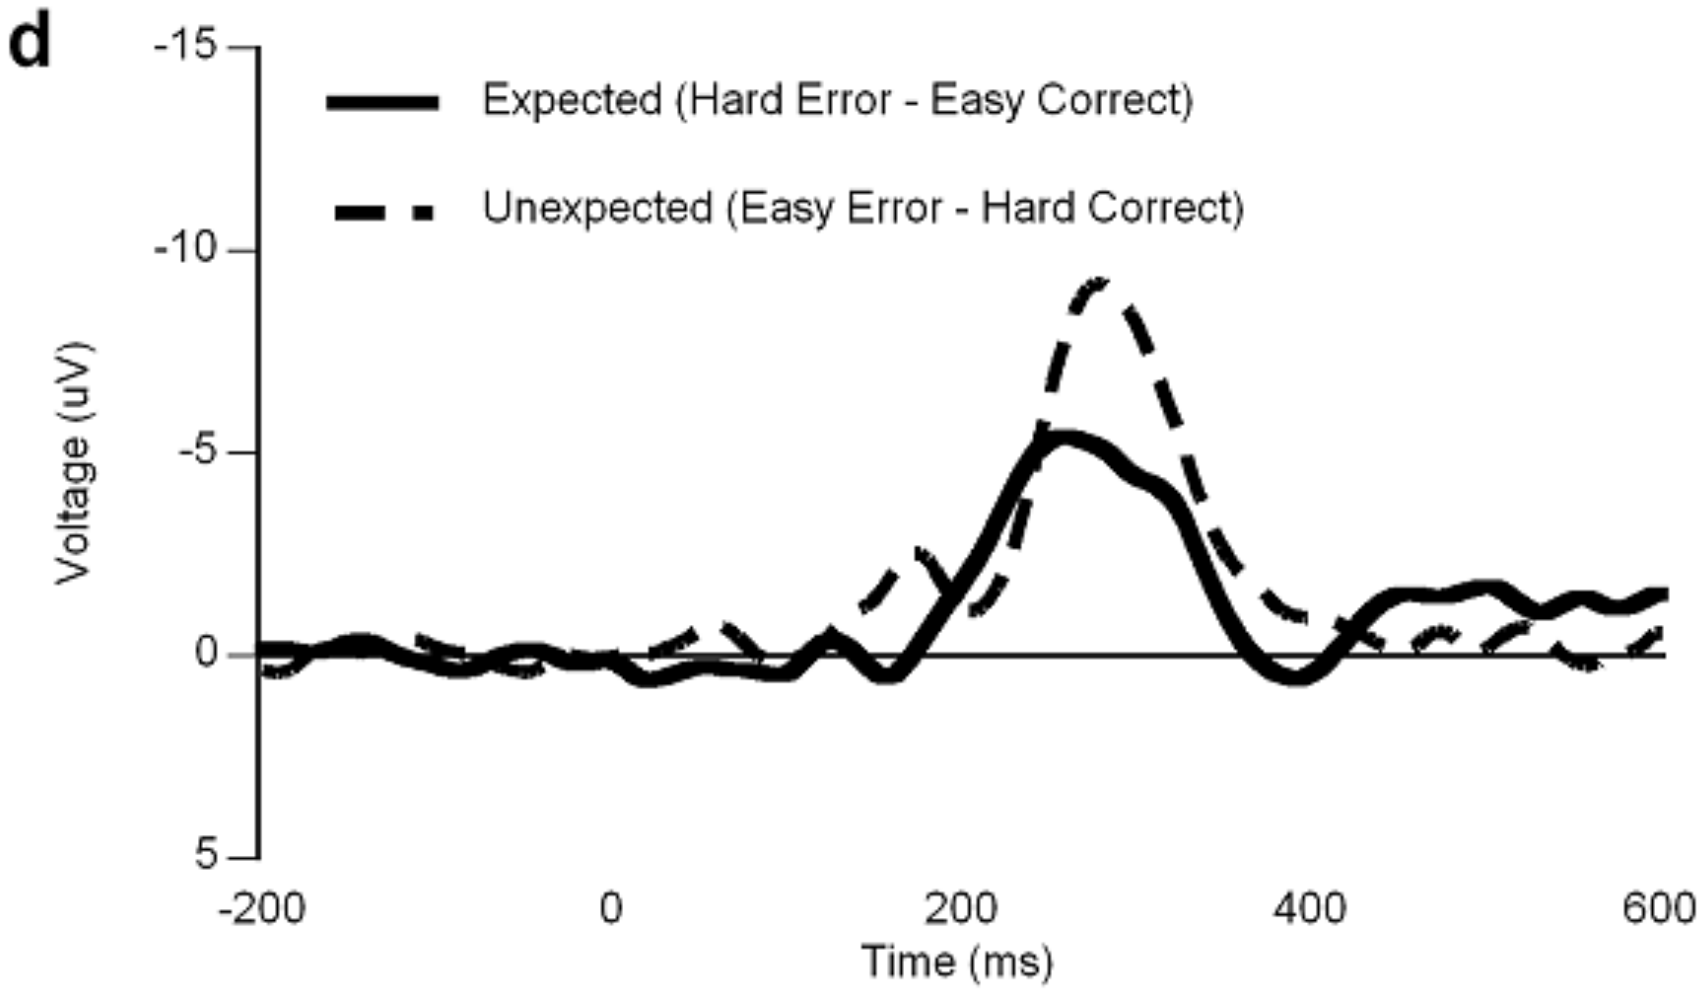

Holroyd & Krigolson, 2007

The Role of Medial-Frontal System in the Acquisition of Perceptual Expertise

Krigolson, Pierce, Tanaka, & Holroyd, 2009

A

Correct!

A Family Blobs

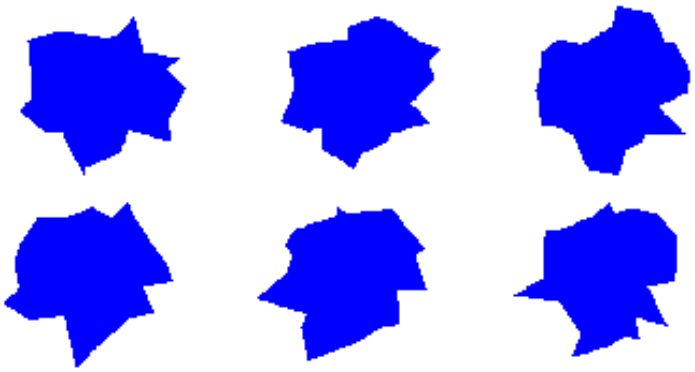

B Family Blobs

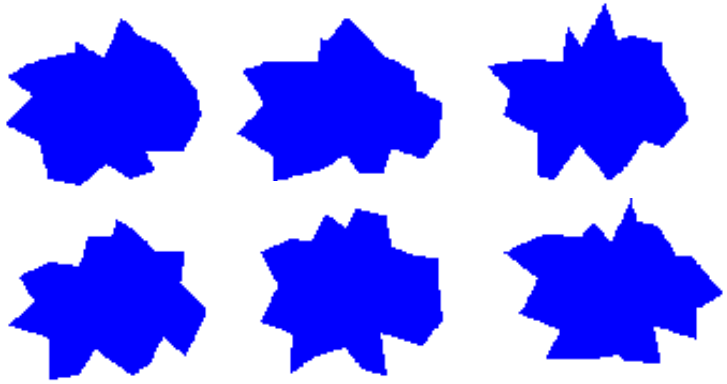

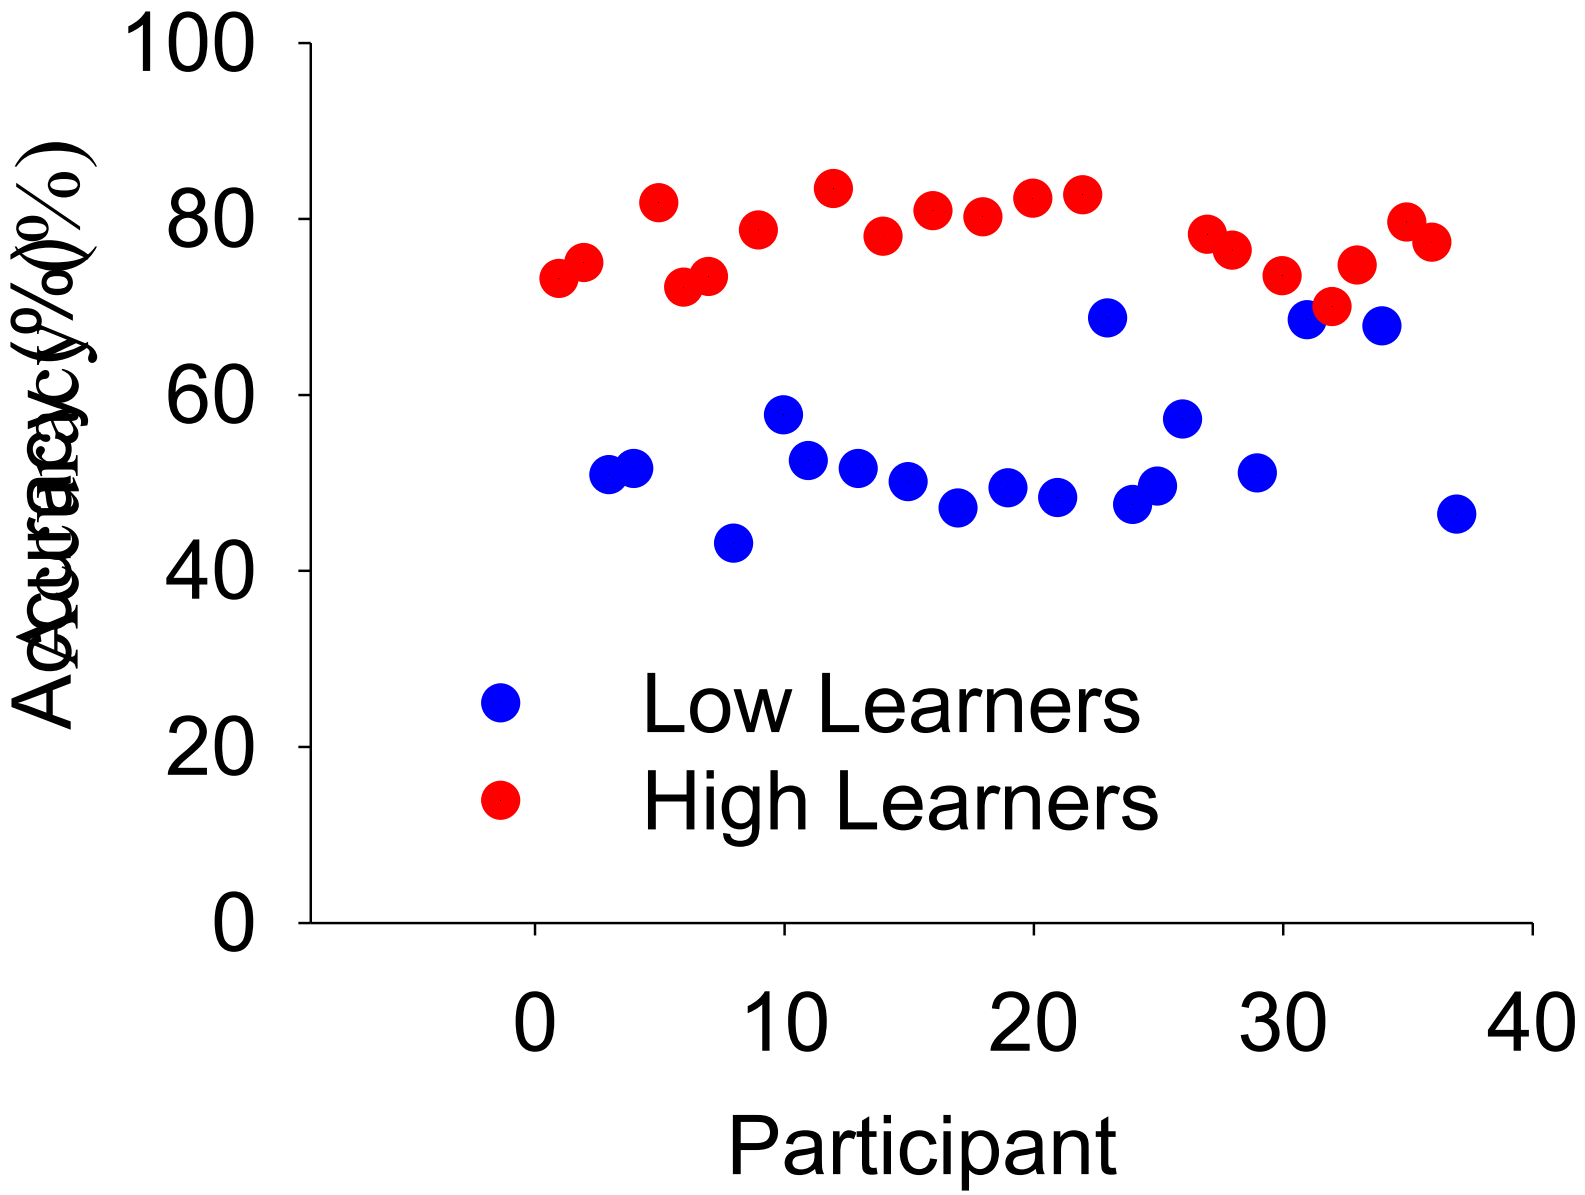
100
80
60
40
20
0
Accuracy (%)
0
10
20
30
40
Participant
Low Learners
High Learners

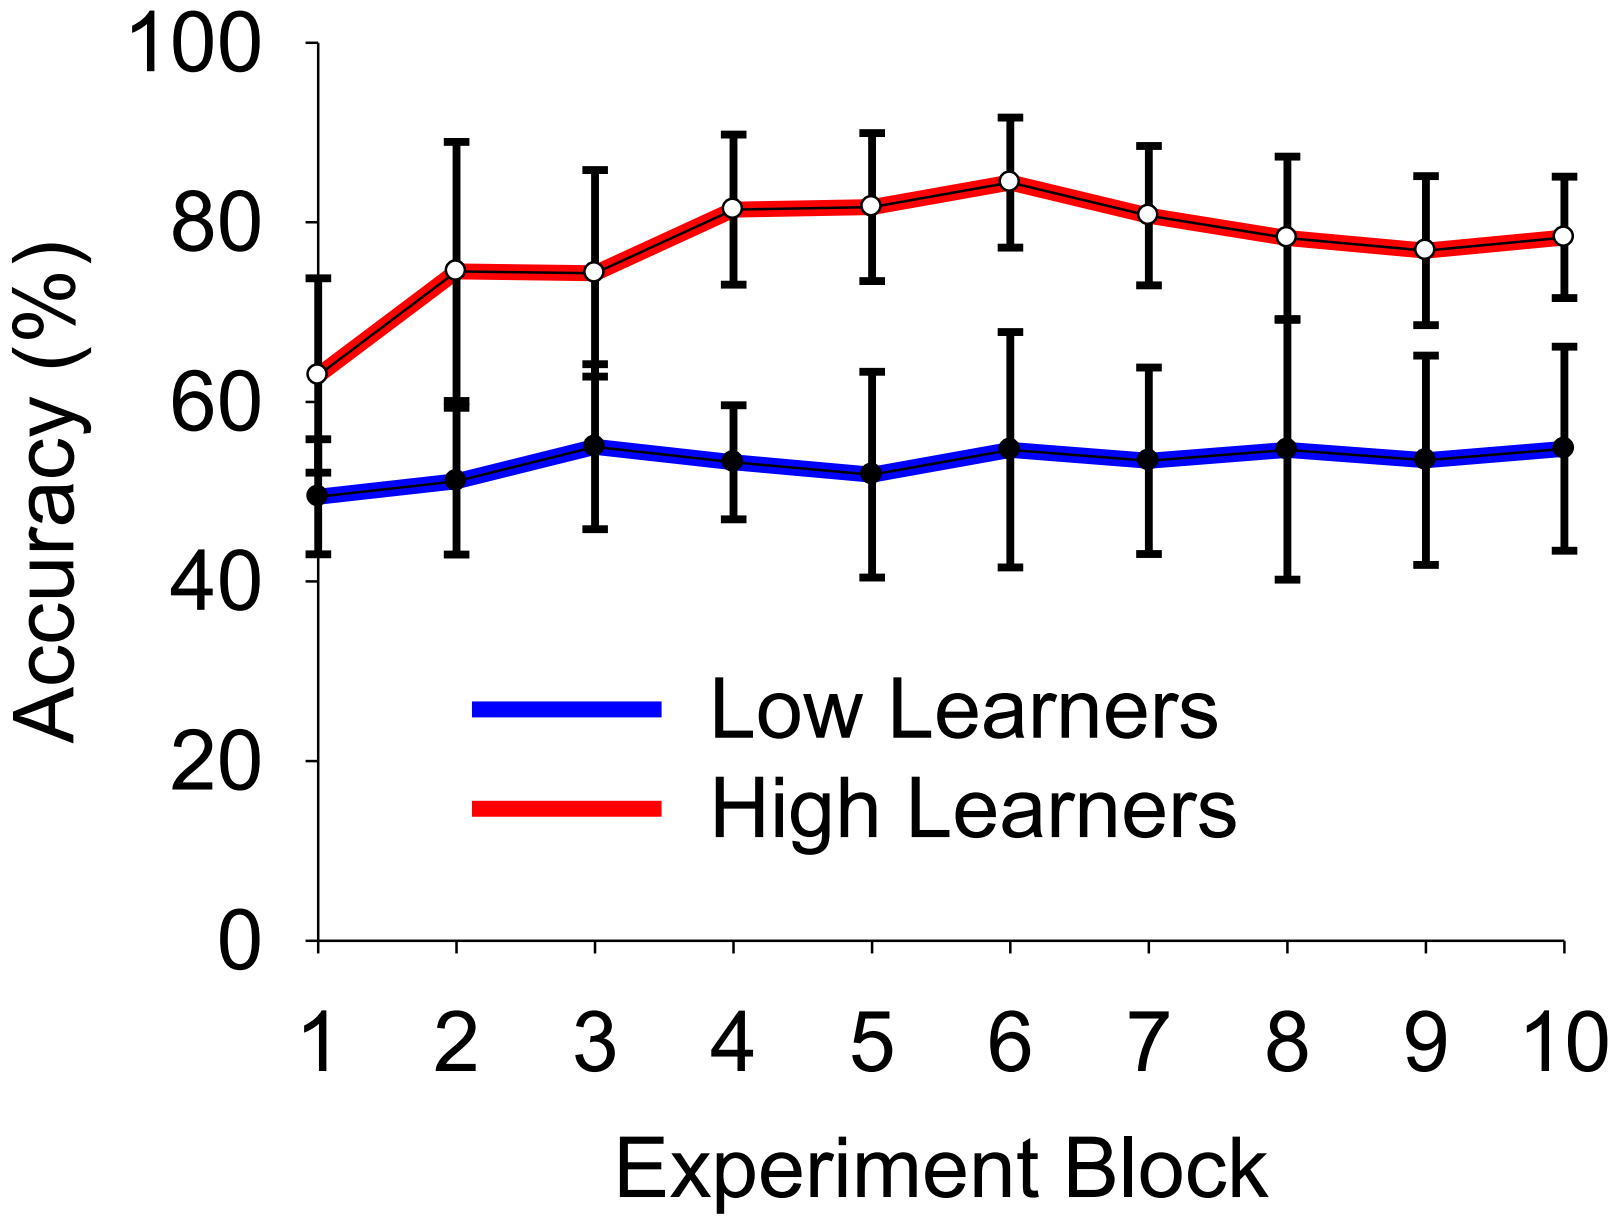
100
80
60
40
20
0
Accuracy (%)
Low Learners
High Learners
1
2
3
4
5
6
7
8
9
10
Experiment Block

Correct!

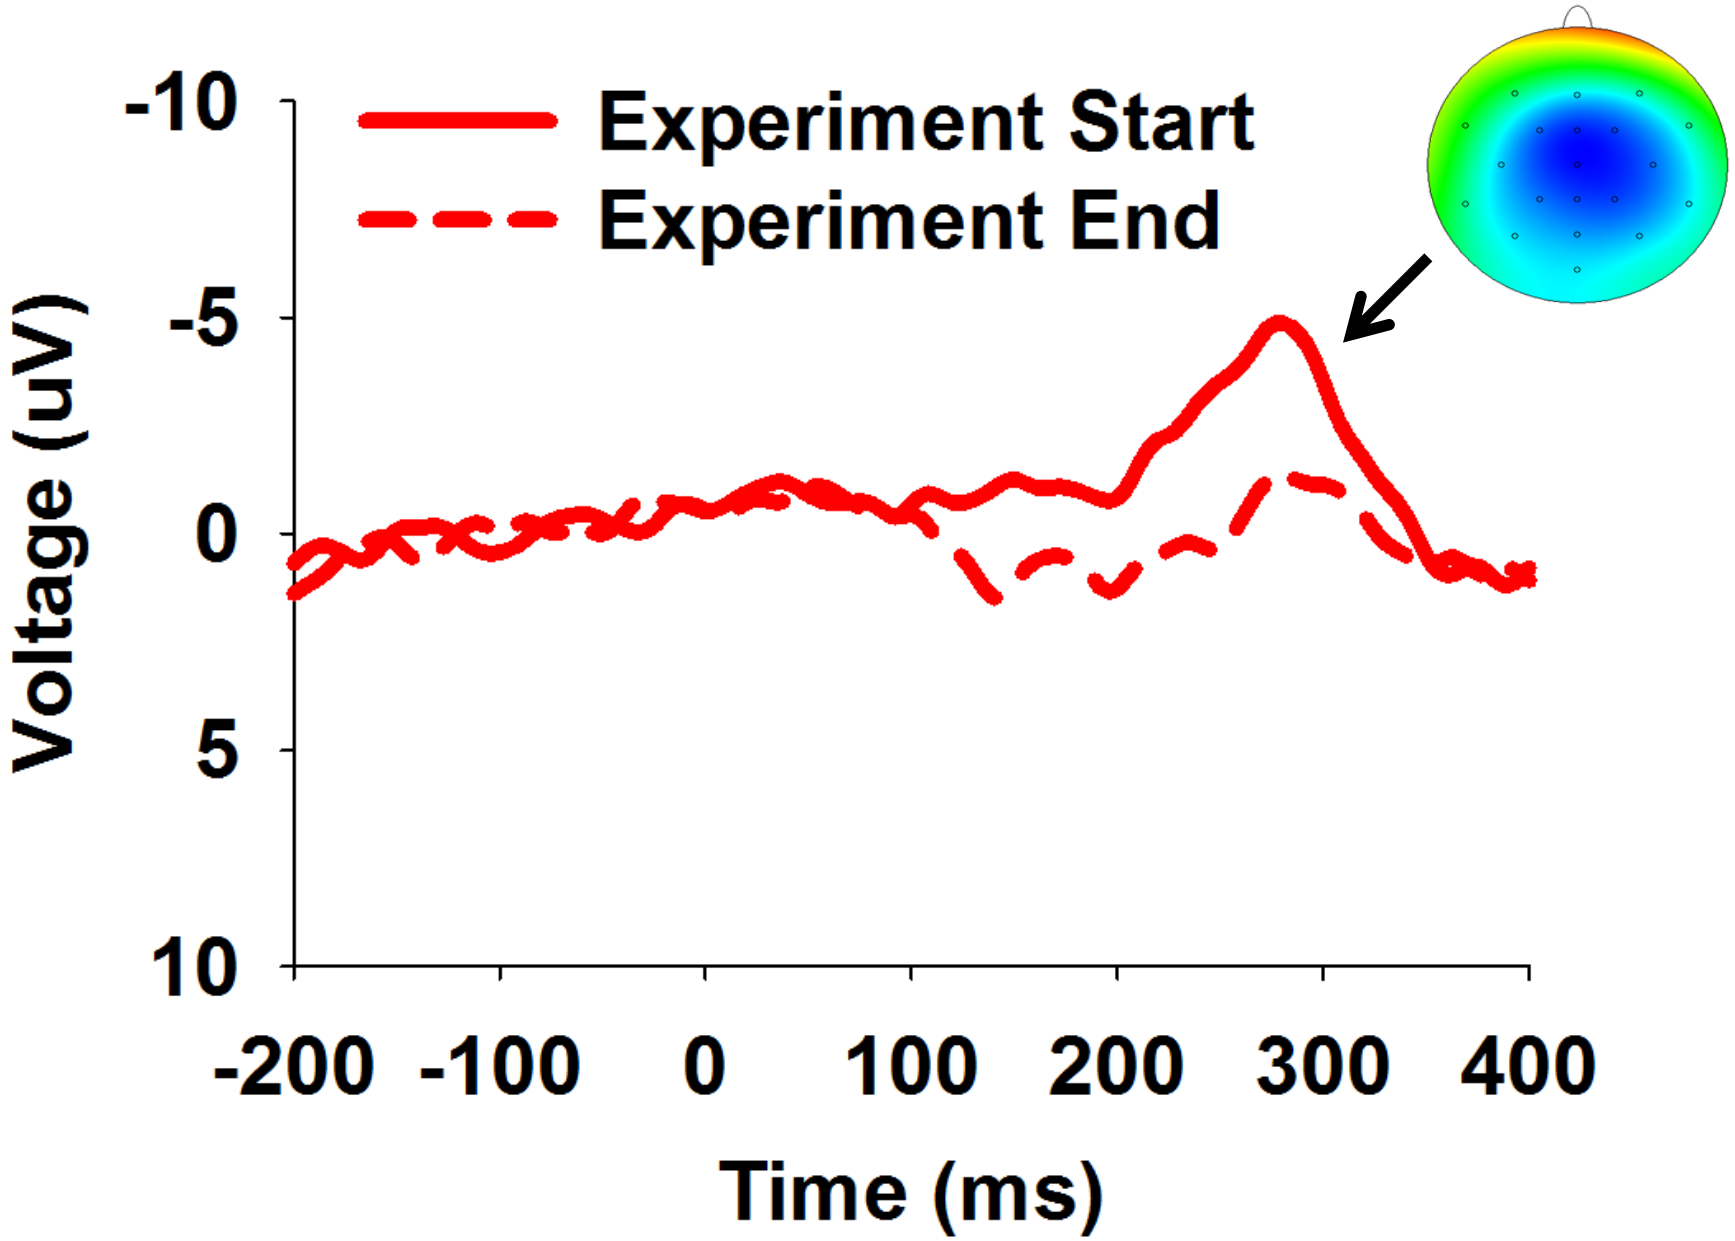

High Learners

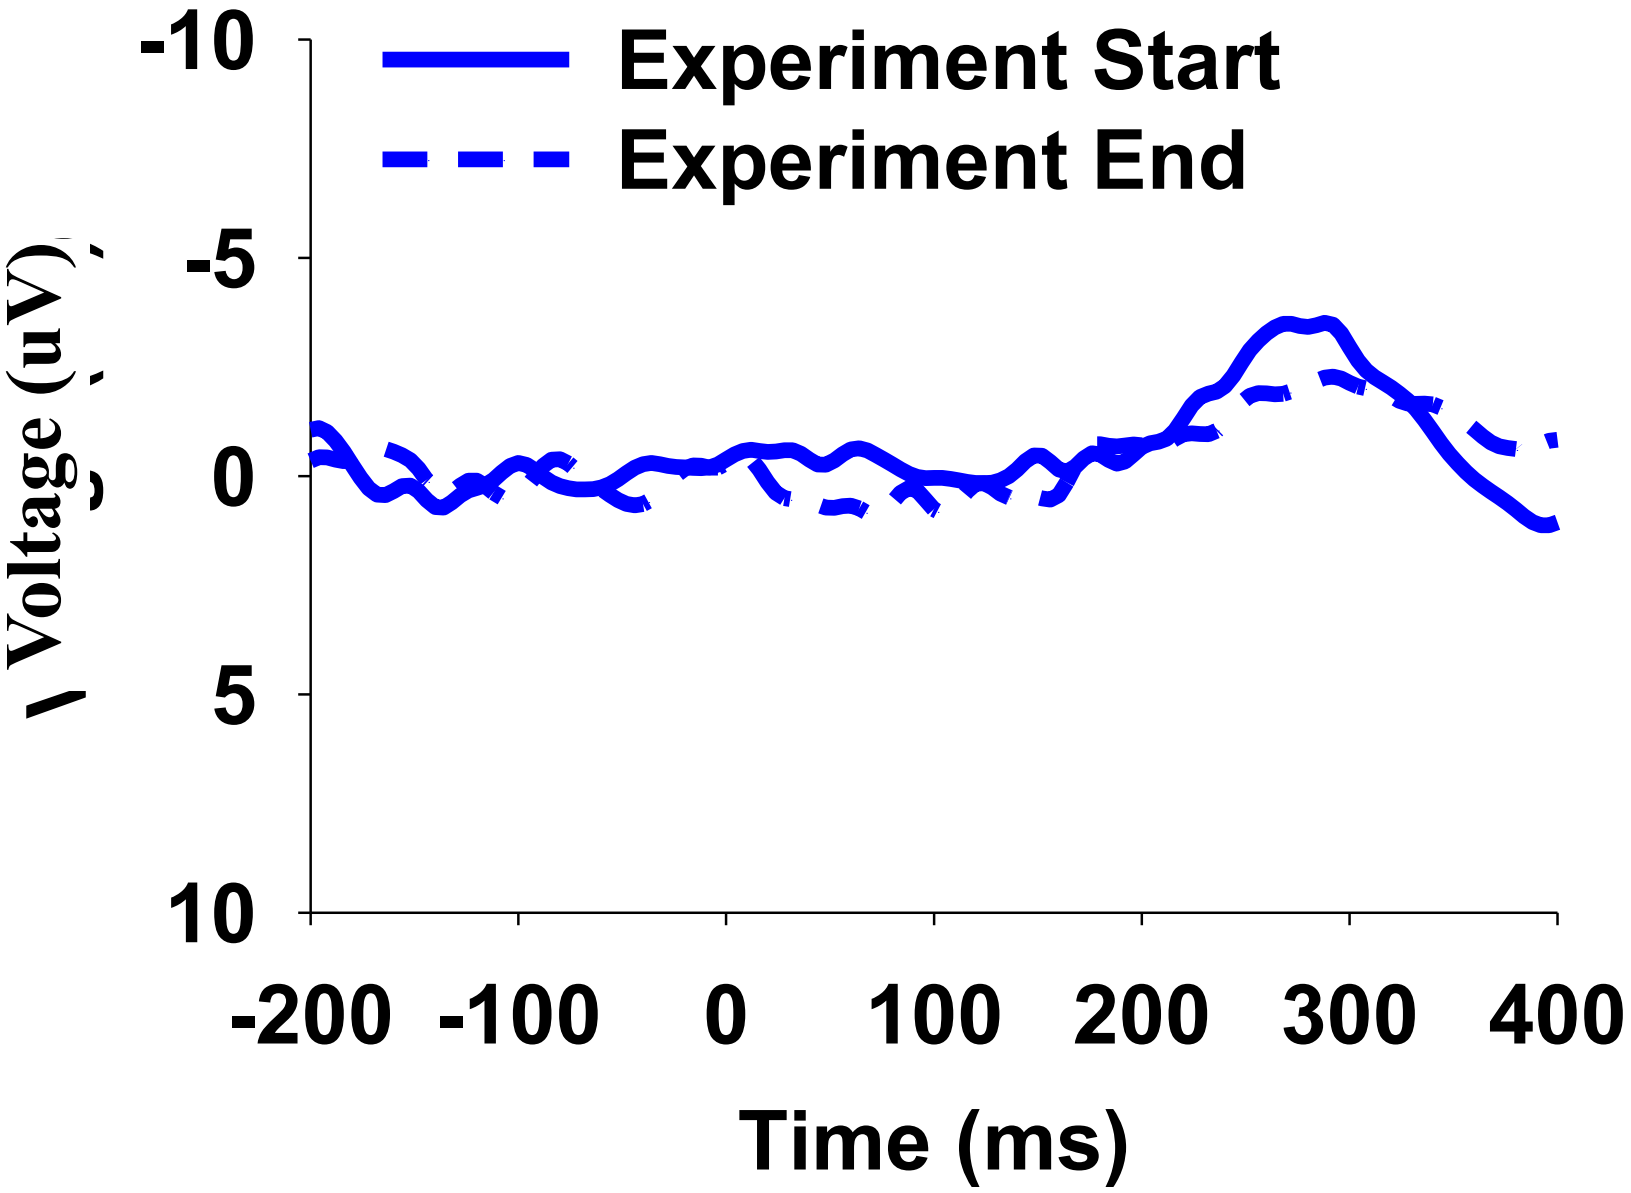

Low Learners

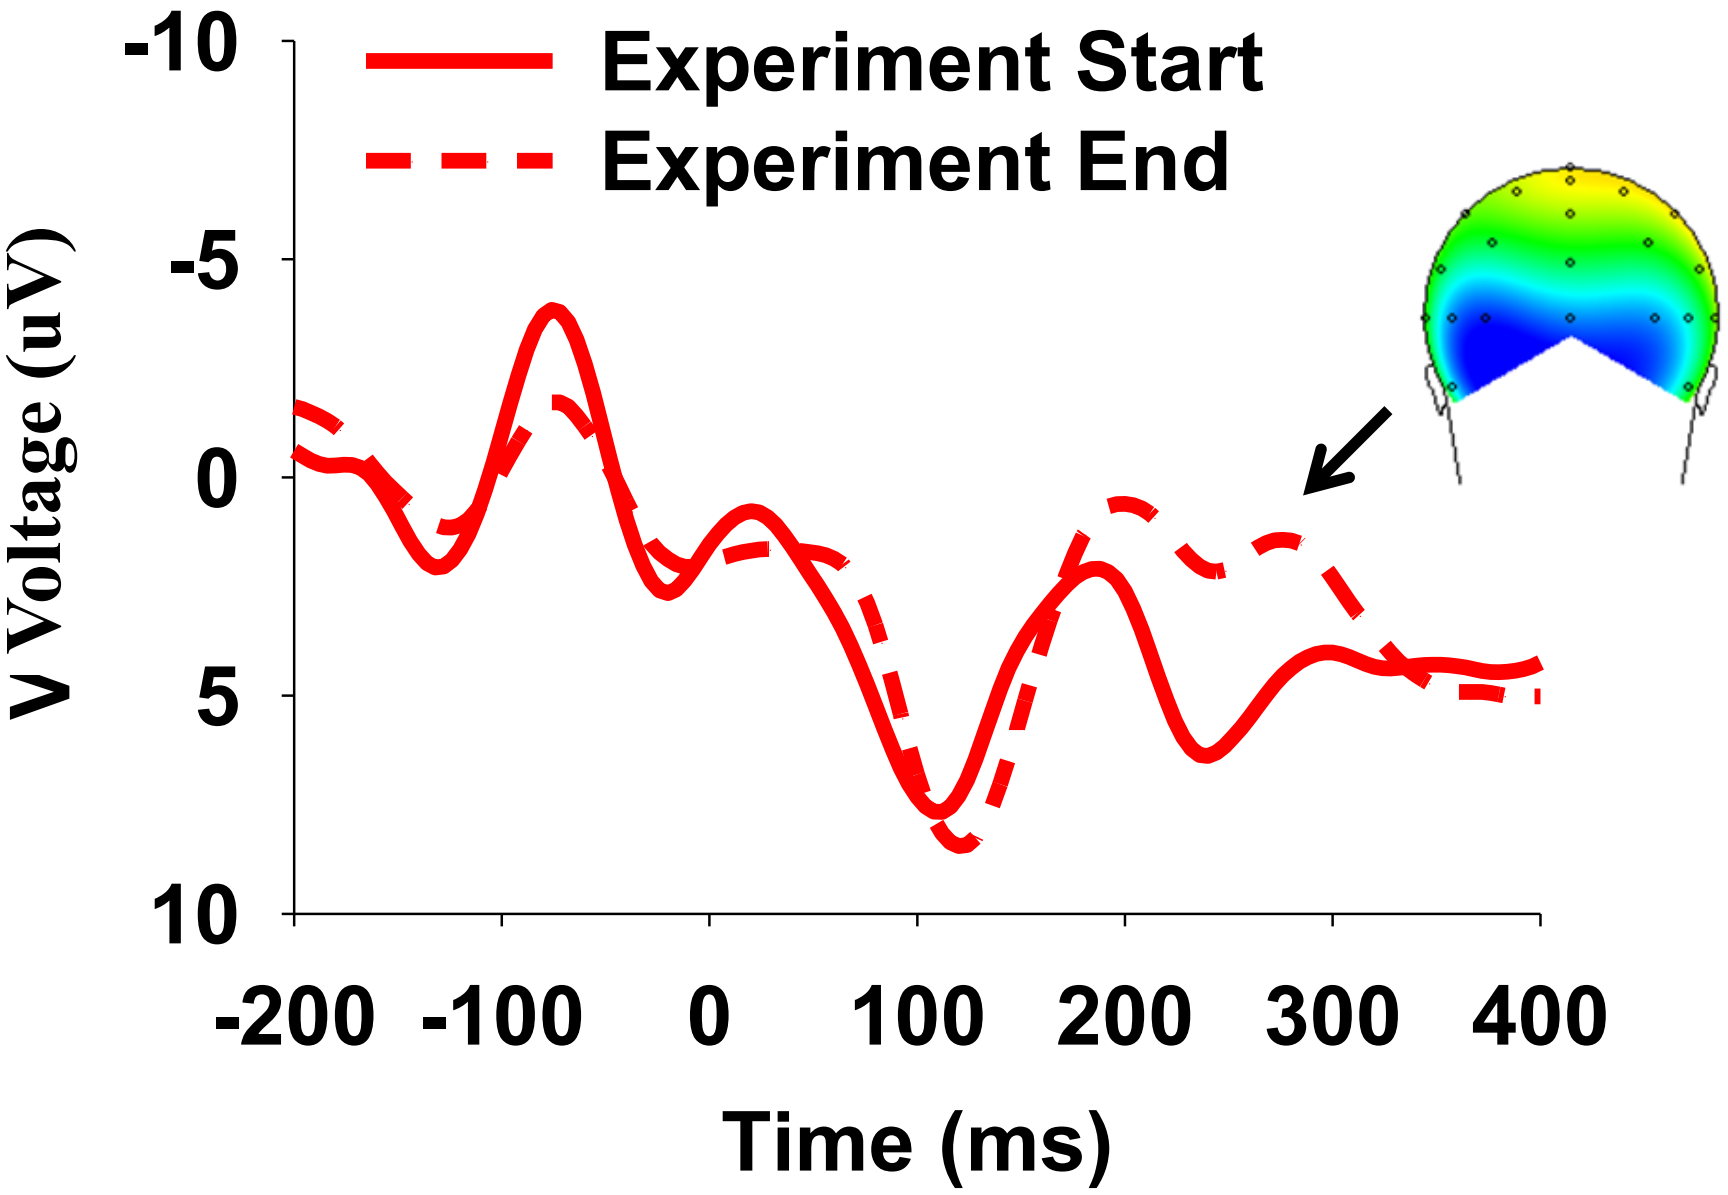

High Learners

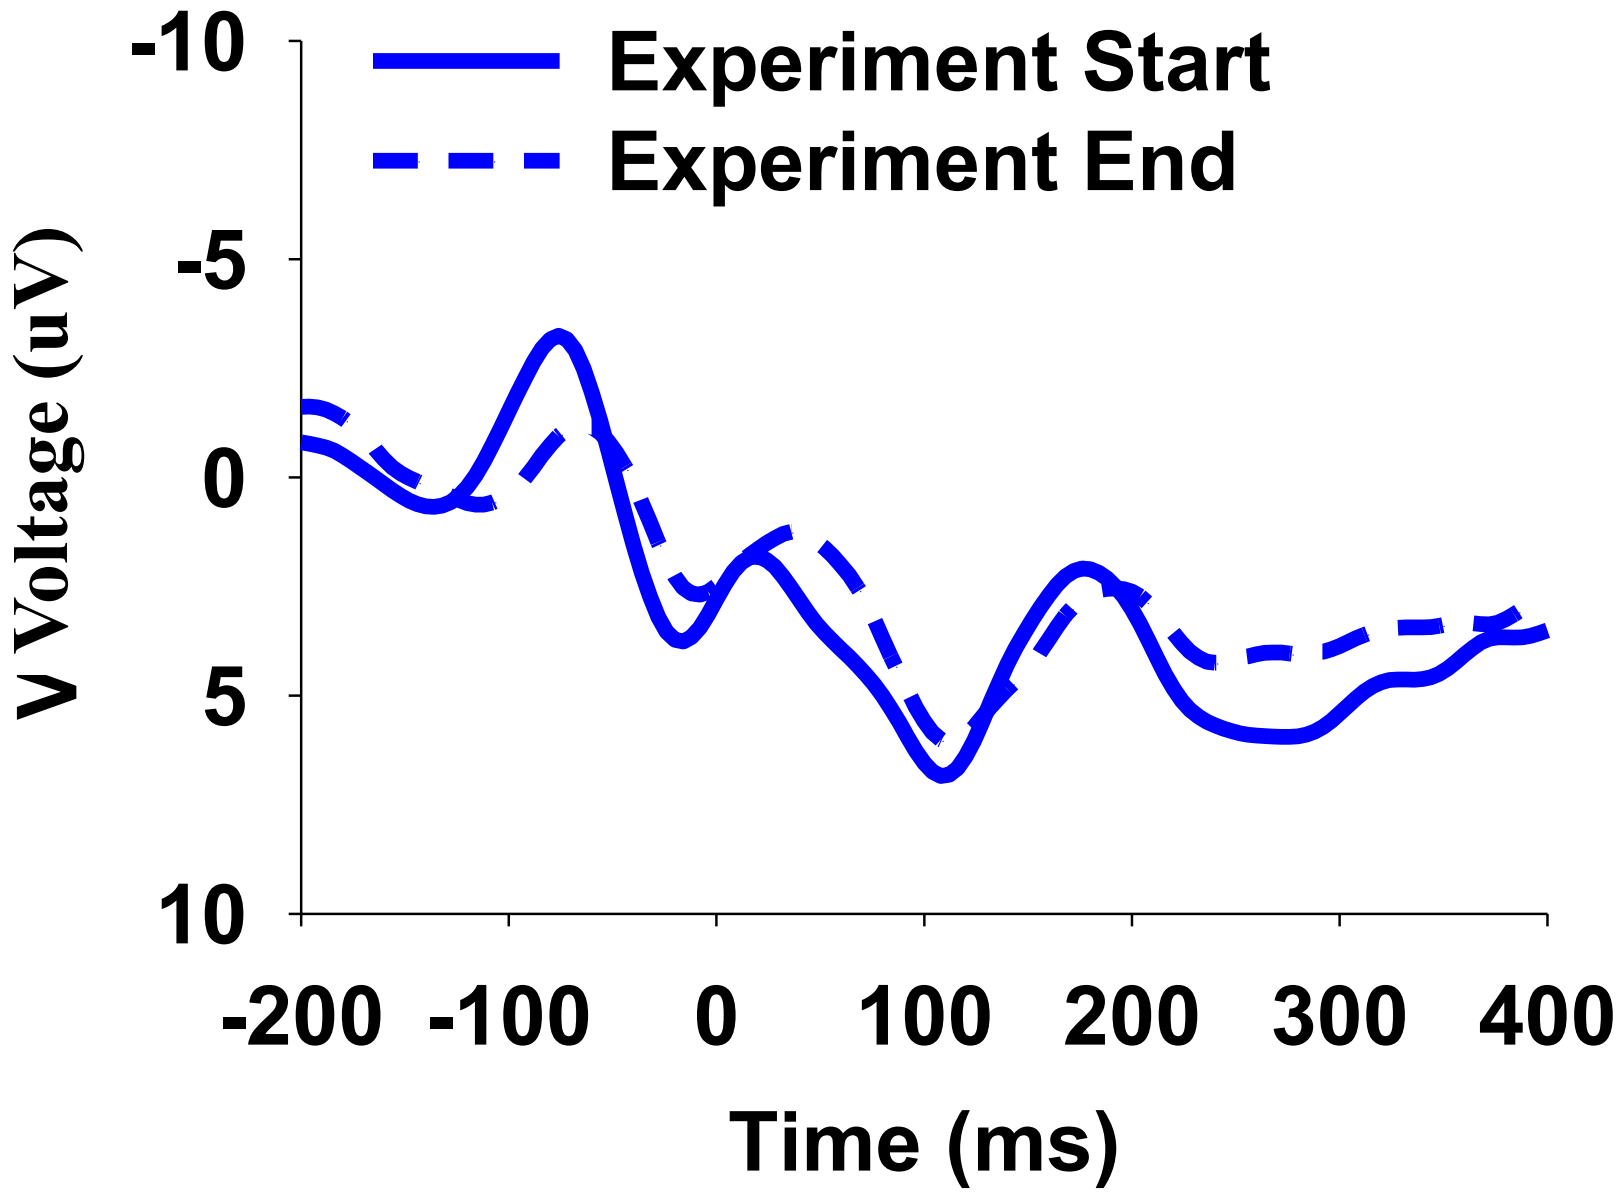

Low Learners

A

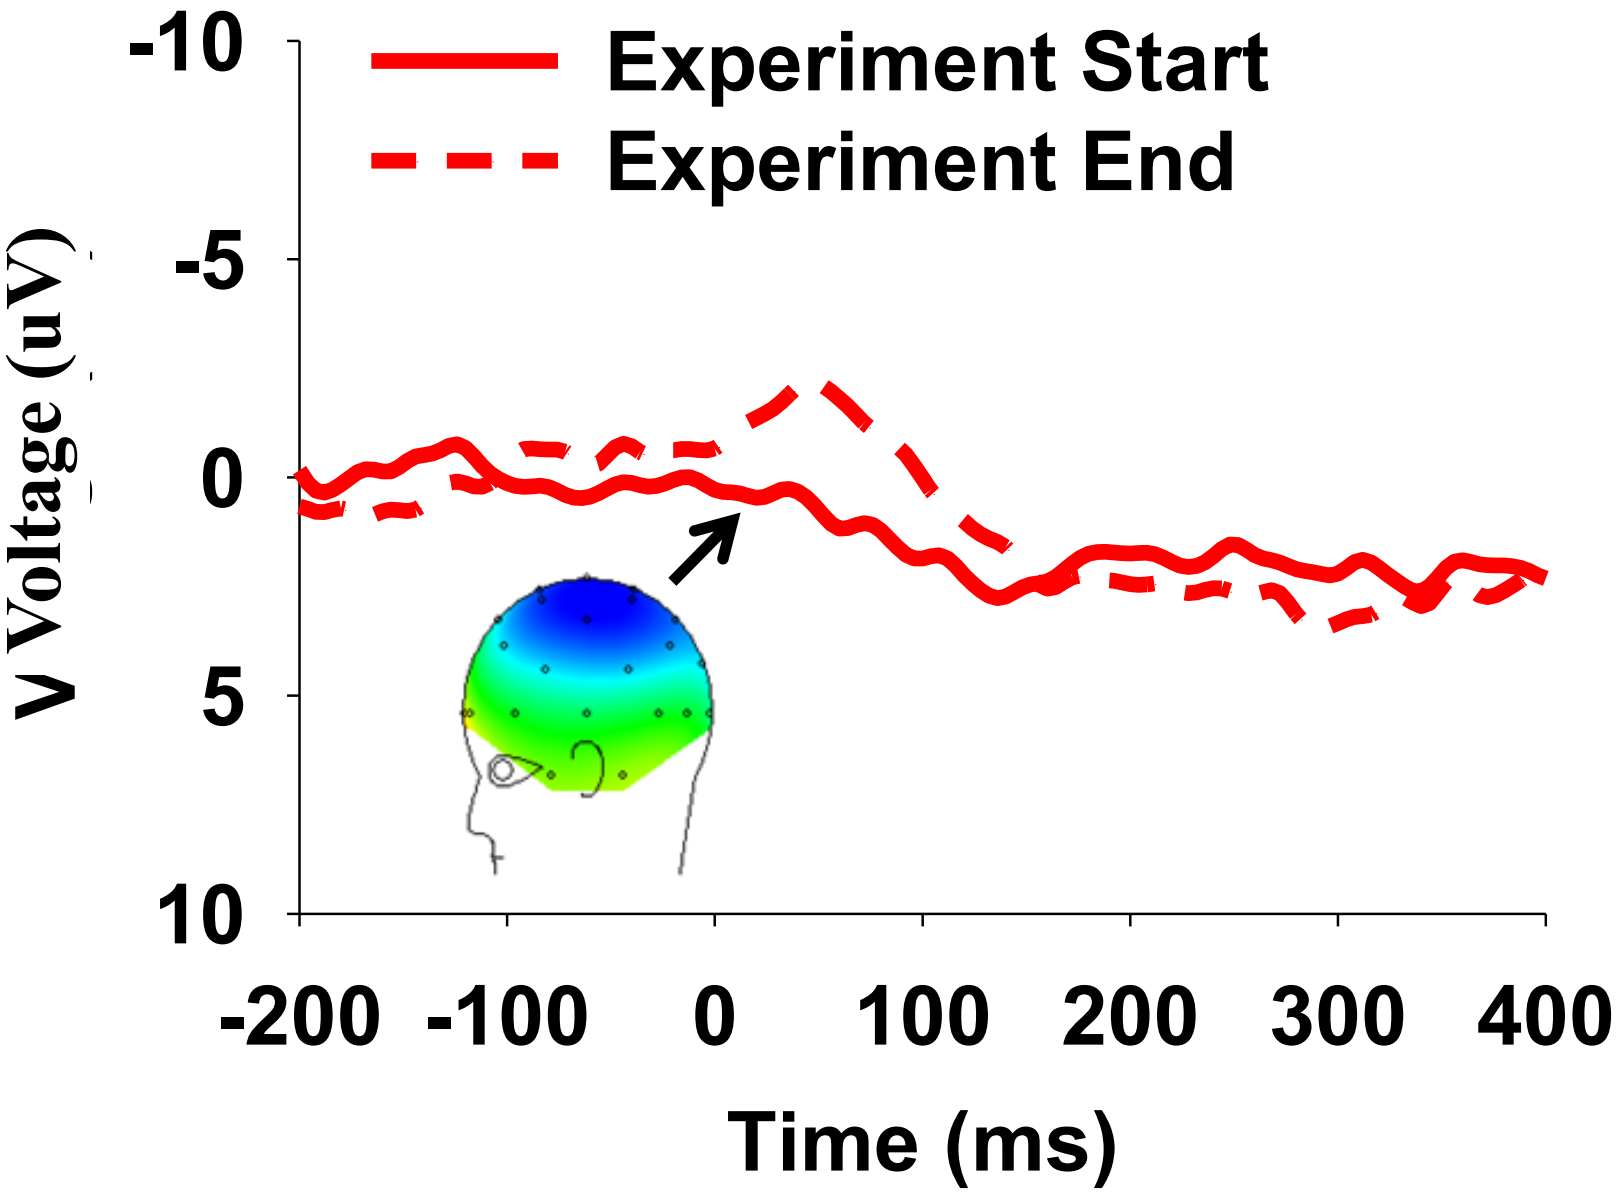

High Learners

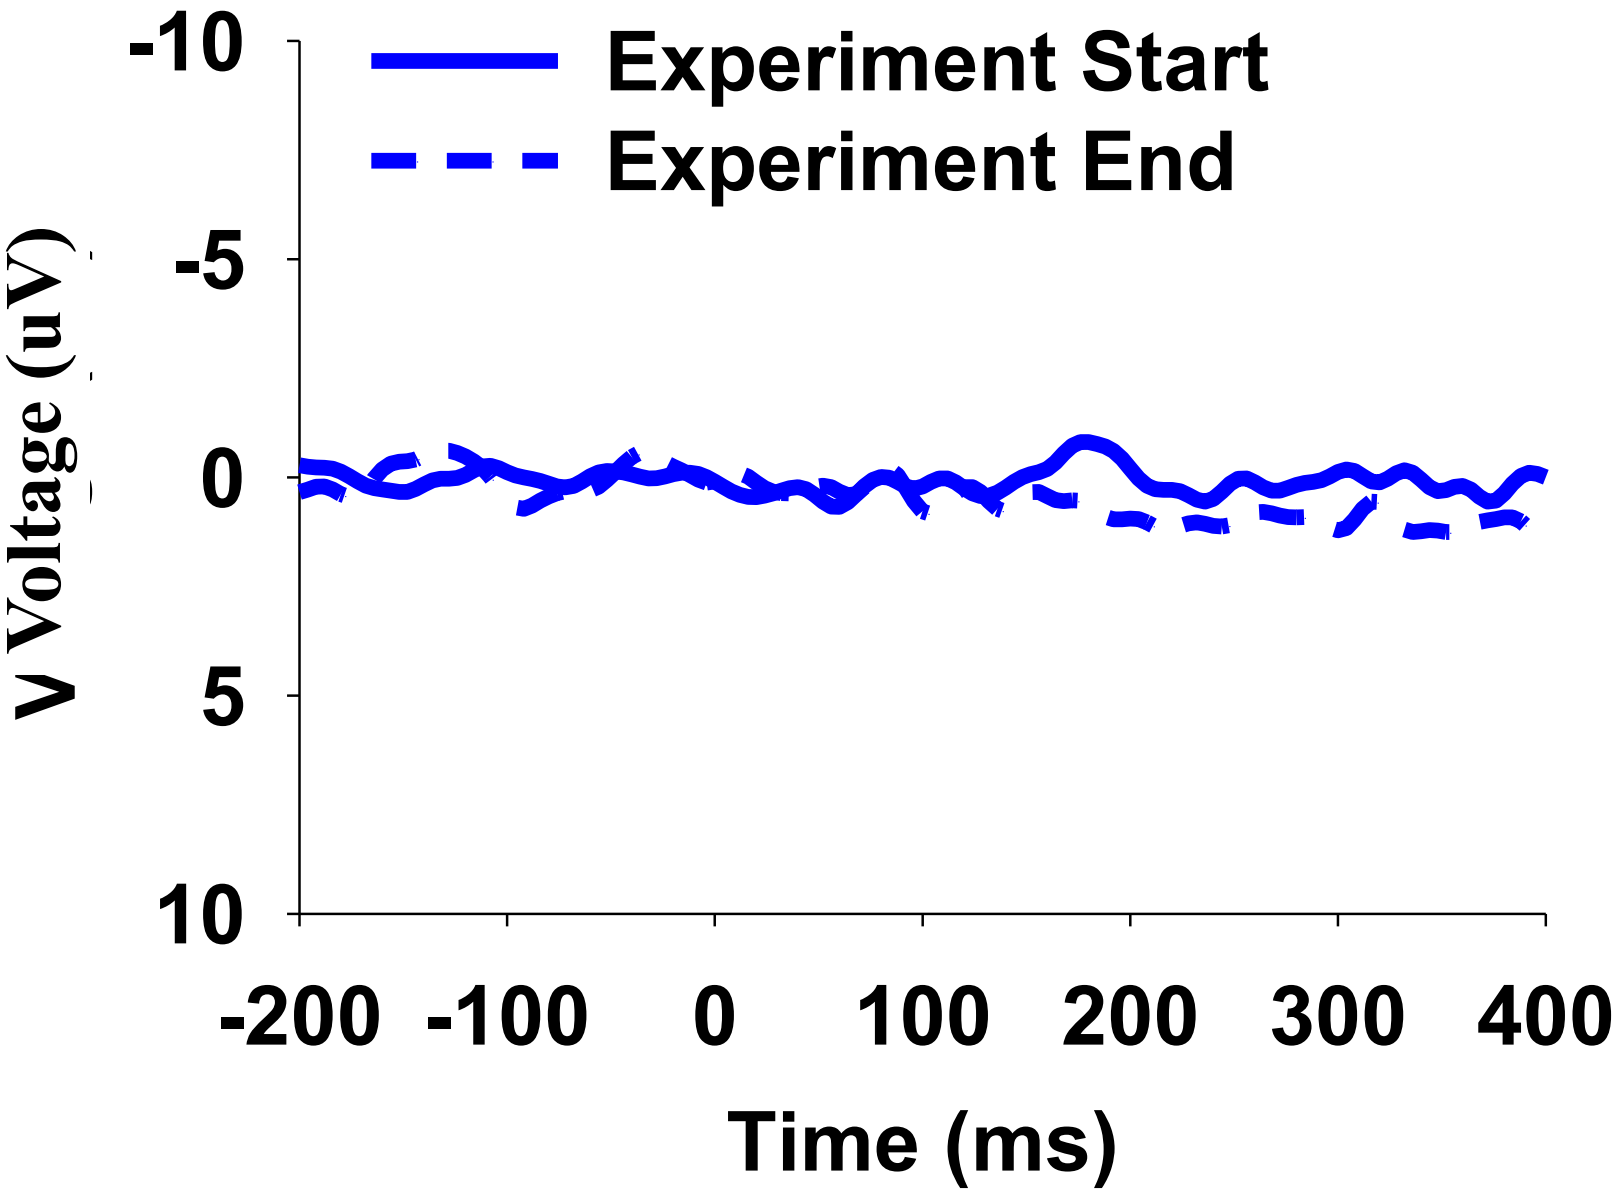

Low Learners

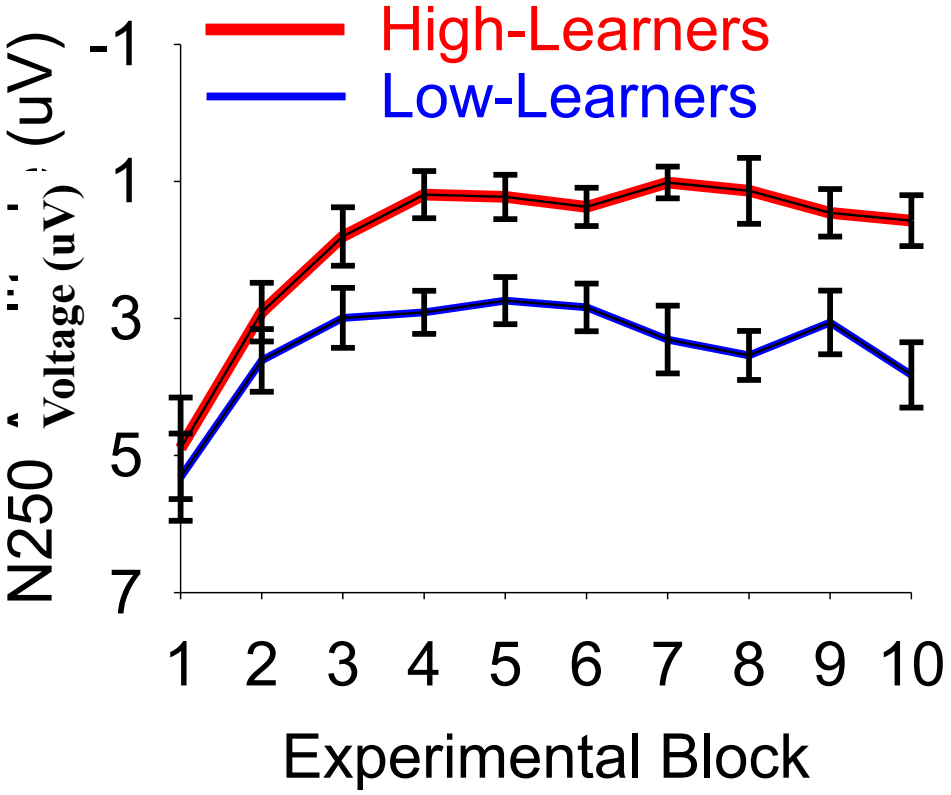
High-Learners
Low-Learners
-1
1
3
5
7
Voltage (uV)
1 2 3 4 5 6 7 8 9 10
Experimental Block

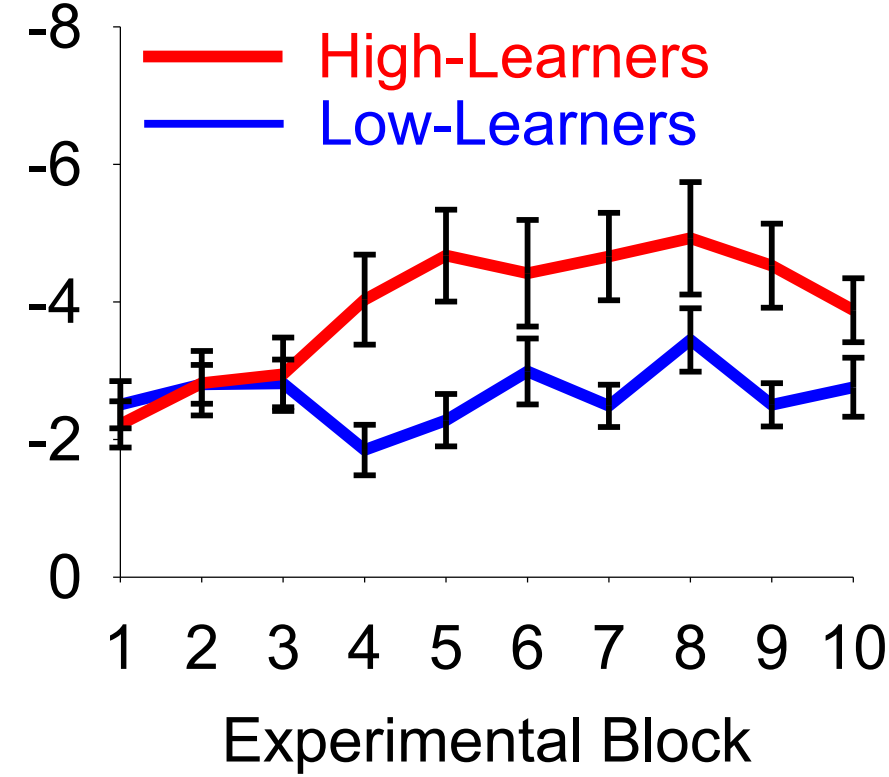
High-Learners
Low-Learners
-8
-6
-4
-2
0
Voltage (uV)
1 2 3 4 5 6 7 8 9 10
Experimental Block

Propagation of the Prediction Error

S **R**

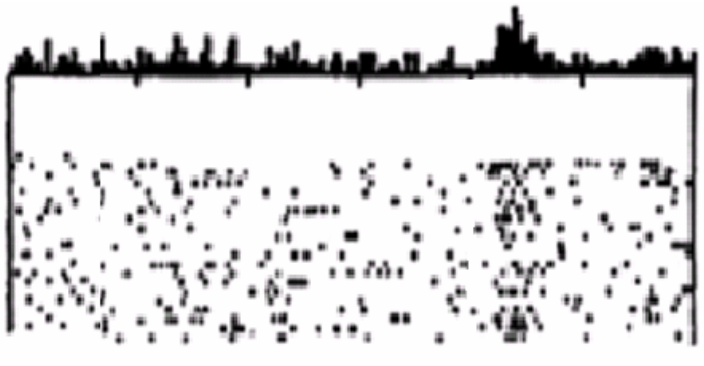

Before Learning

S **R**

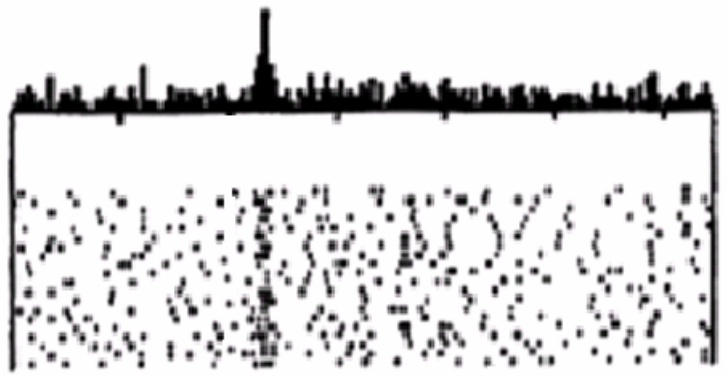

Learned Stimulus/
Reward Mapping

+

Krigolson et al., In press

Krigolson et al., 2014

$

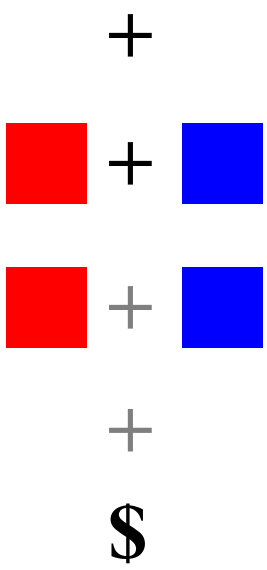

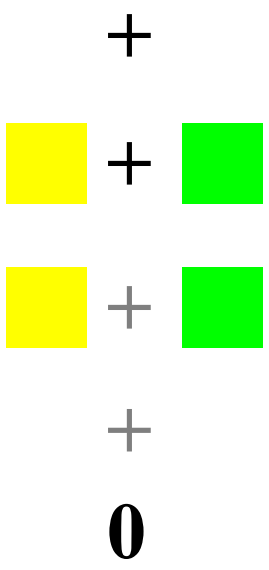

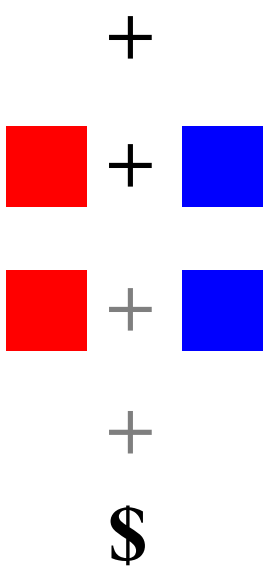

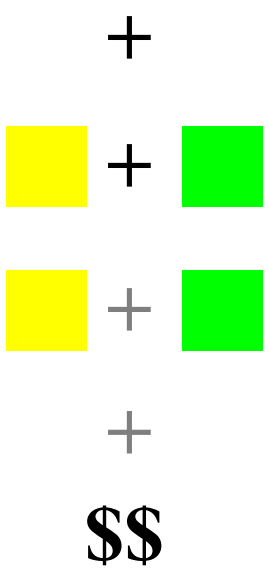

$$P_k = \frac{e^{w_k / t}}{\sum_{k=1}^{2} e^{w_k / t}}$$

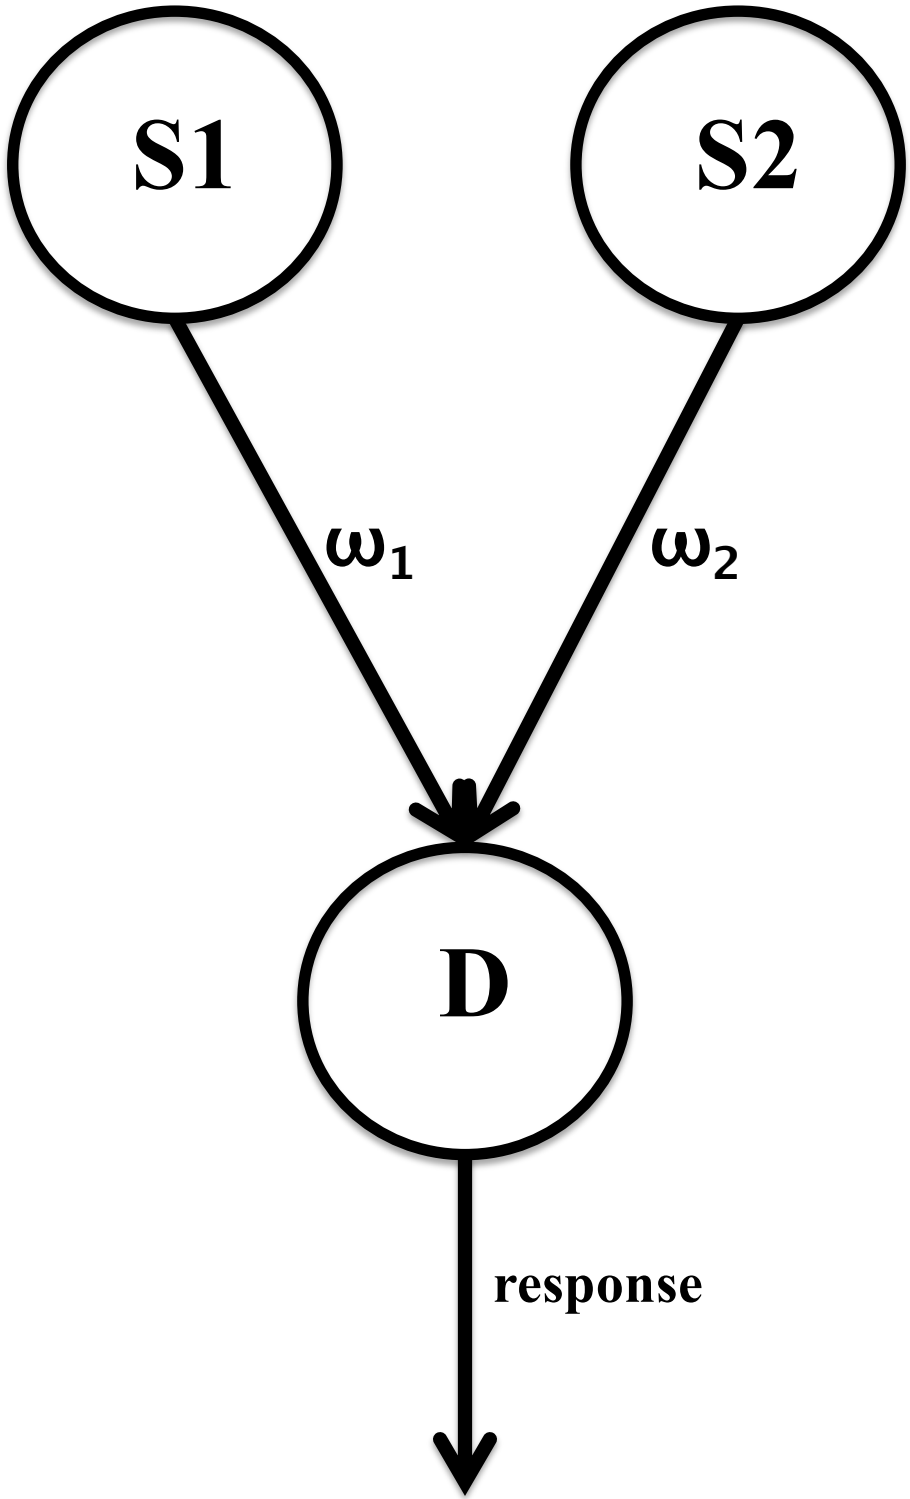

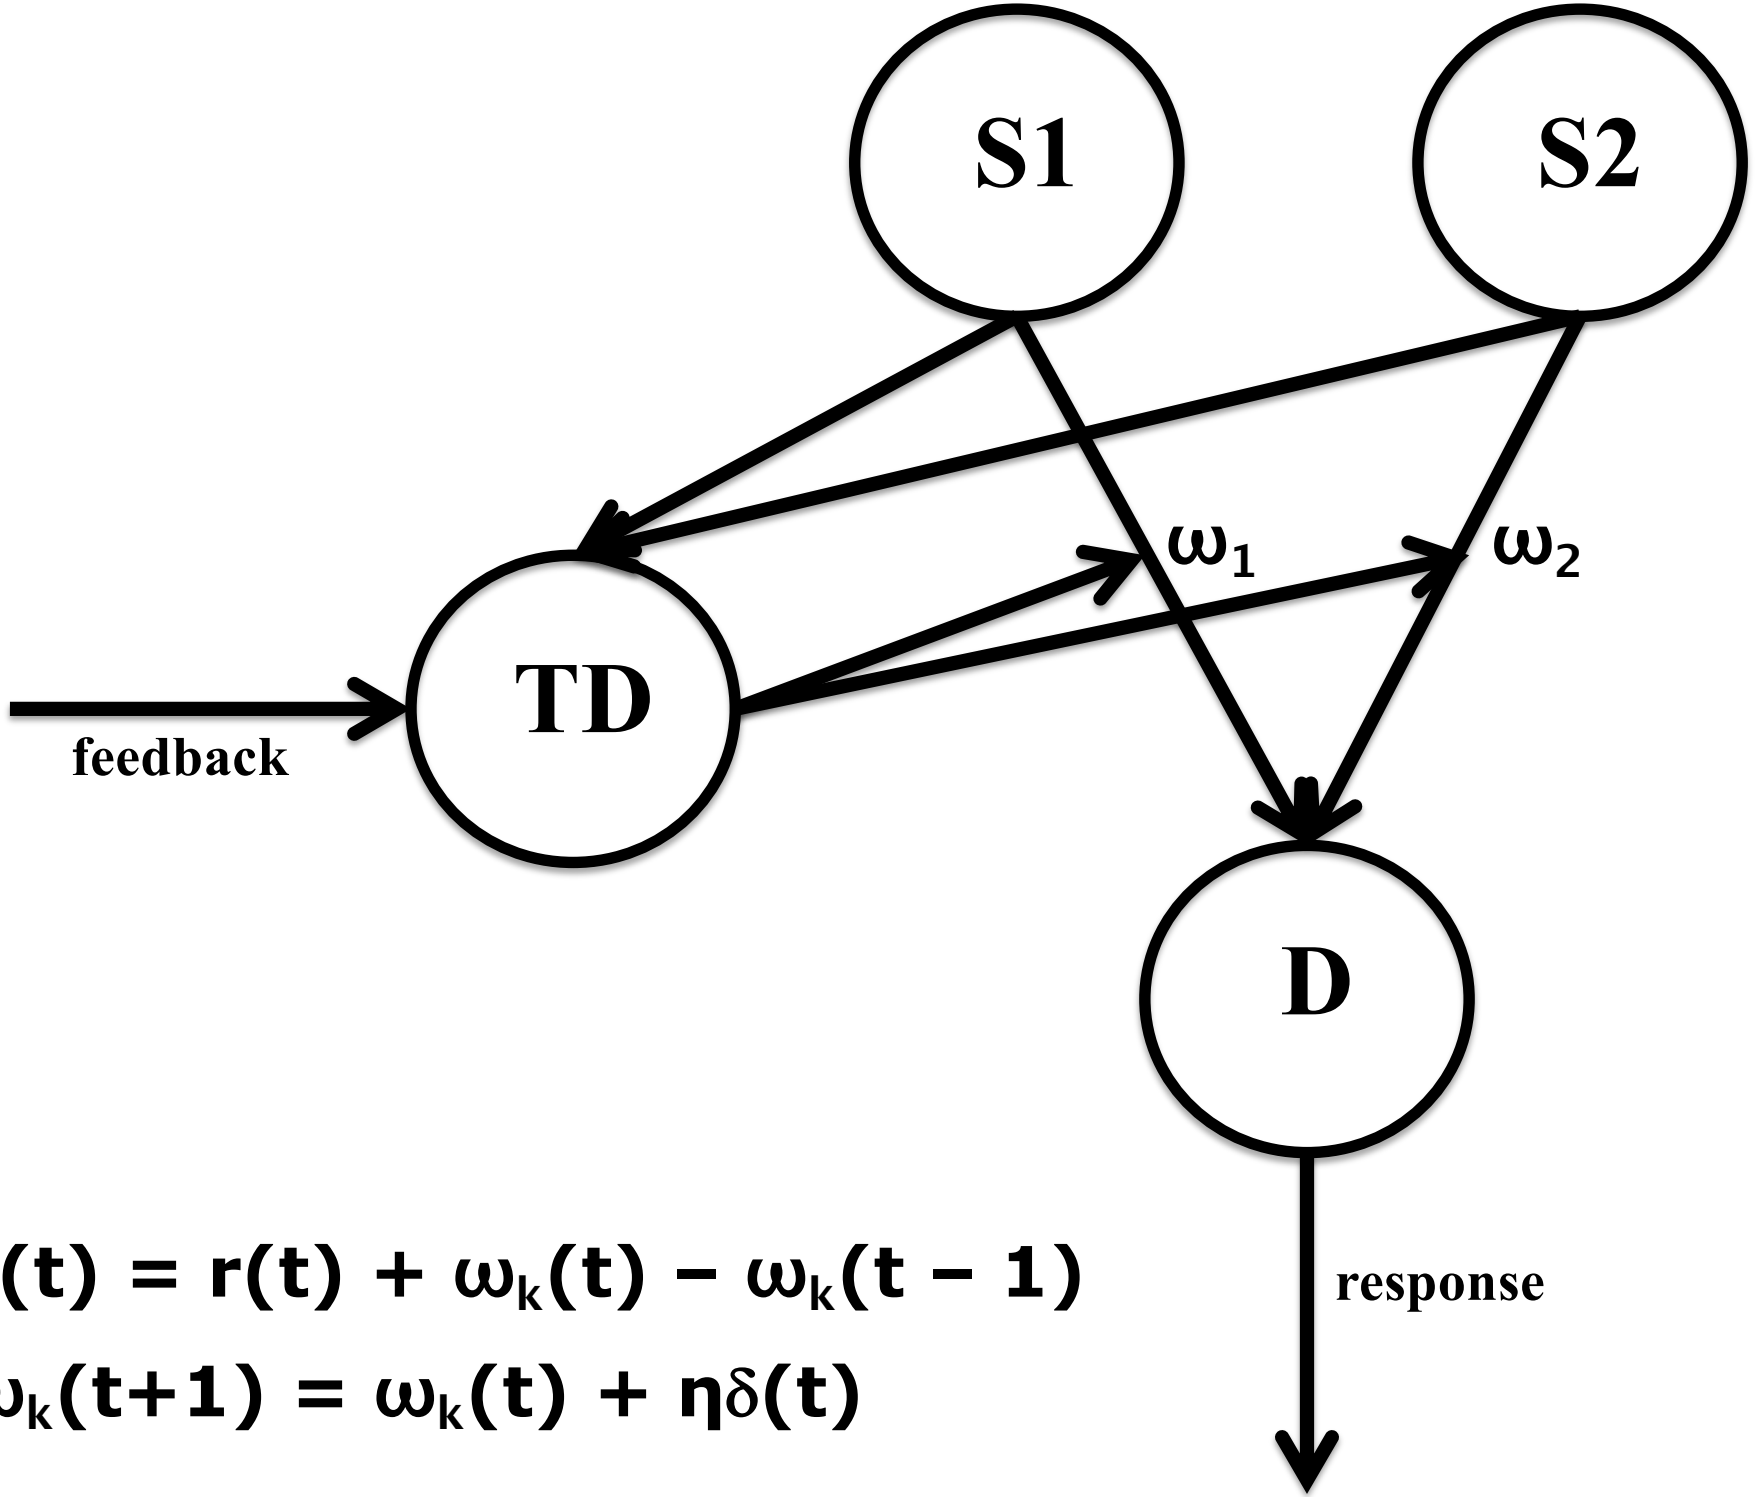

$$\delta(t) = r(t) + \omega_k(t) - \omega_k(t-1)$$

$$\omega_k(t+1) = \omega_k(t) + \eta\delta(t)$$

Behavioral Accuracy

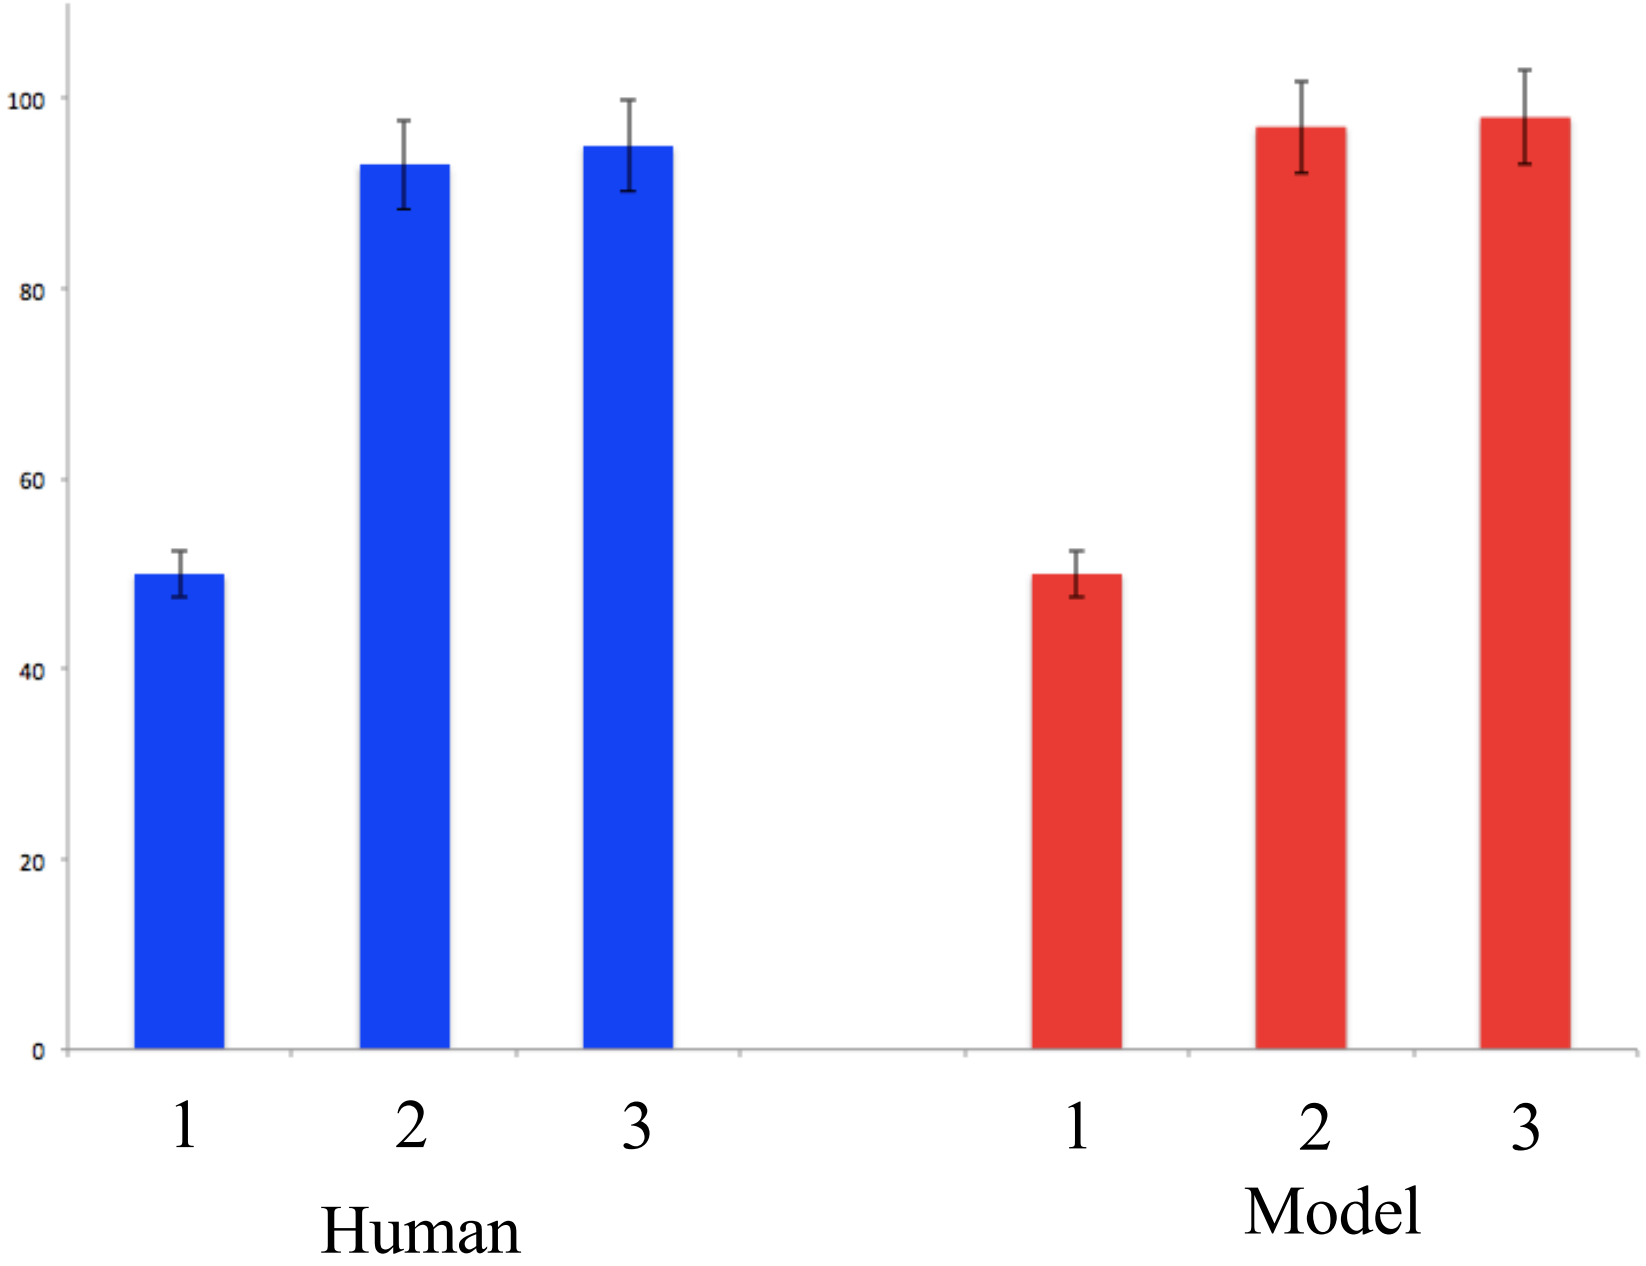

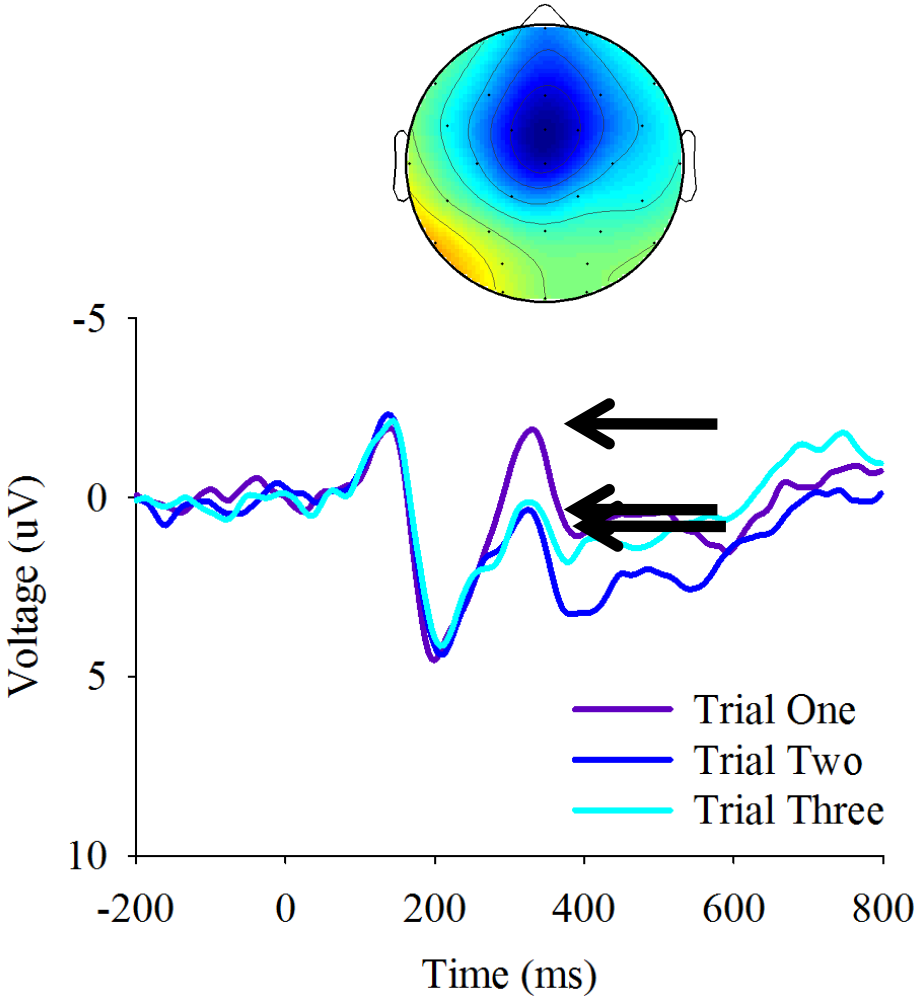

Stimulus Averaged ERP Waveforms

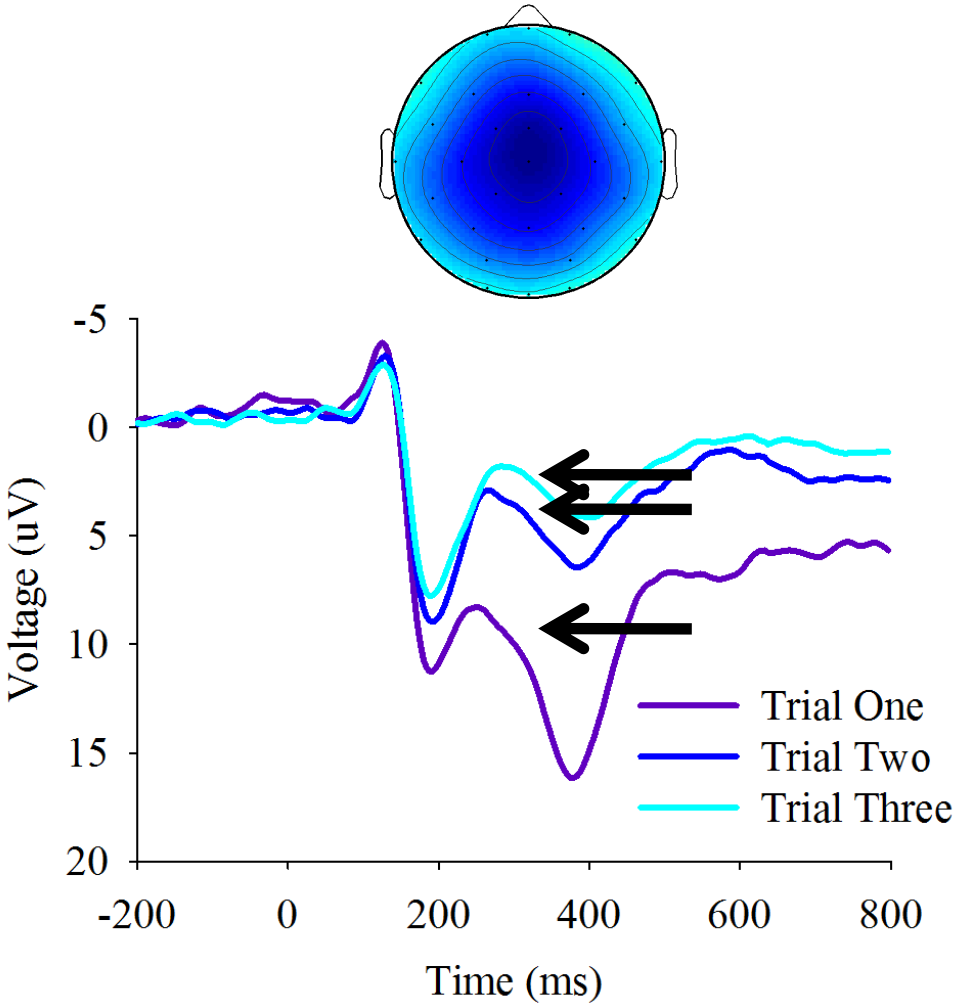

Feedback Averaged ERP Waveforms

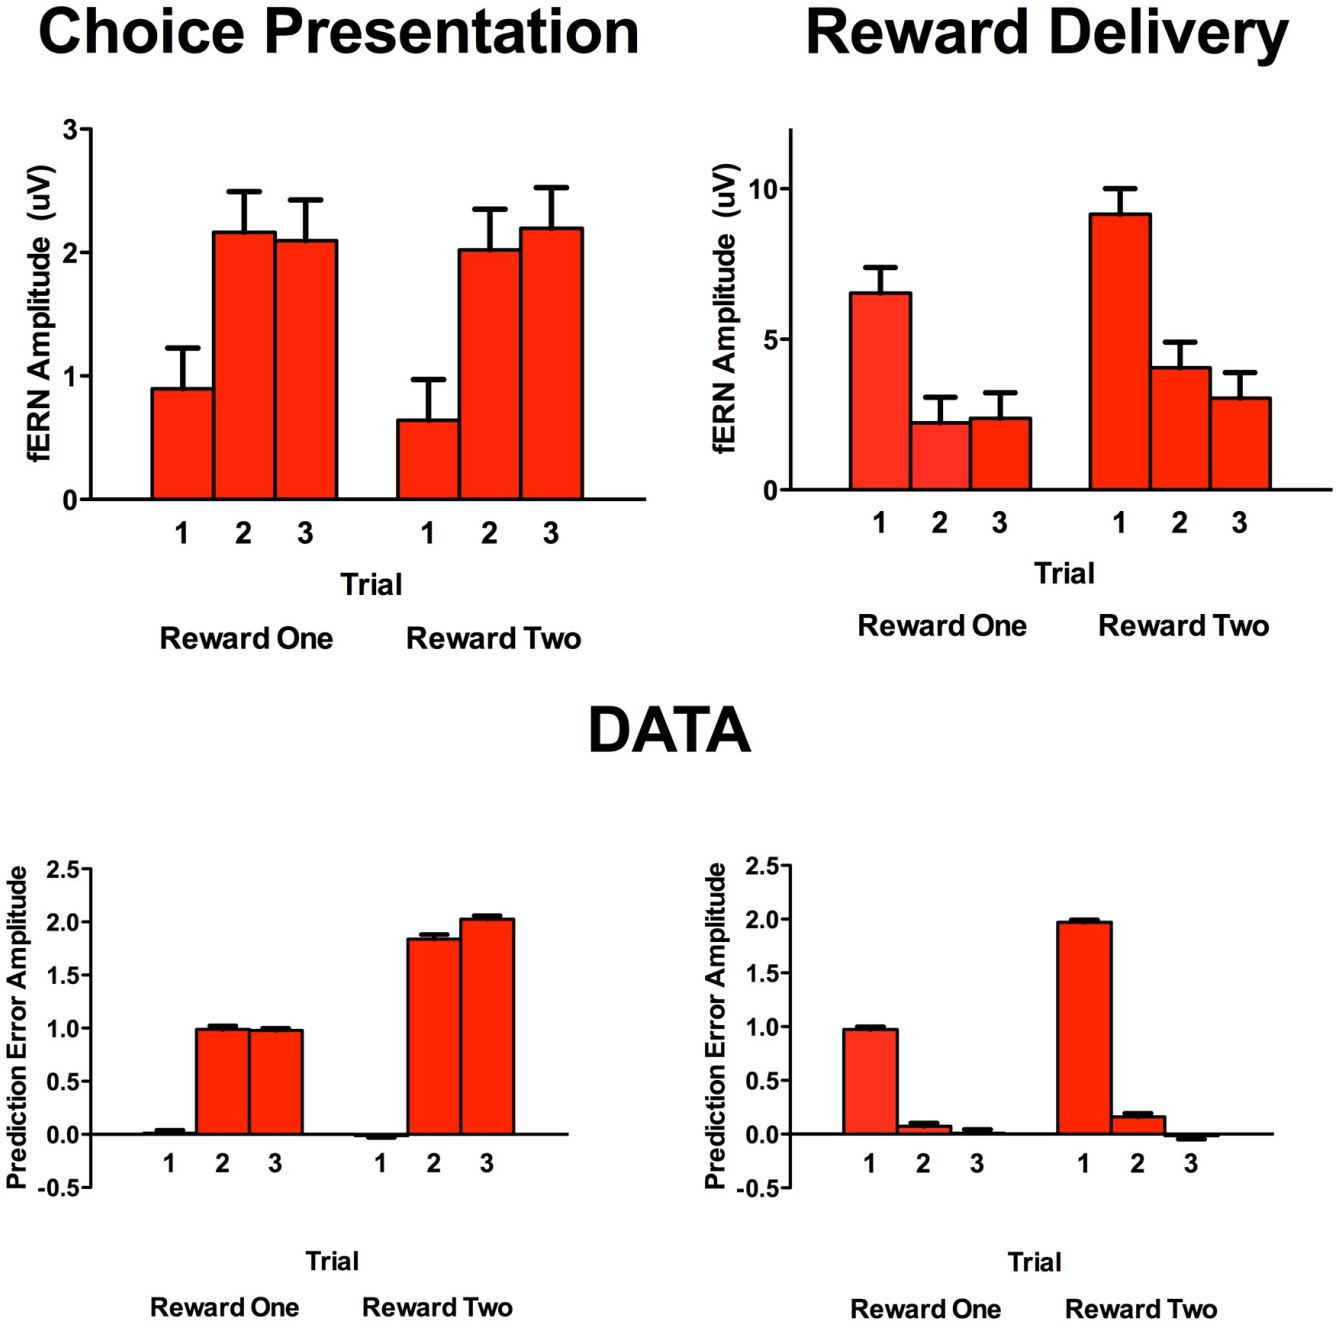
Choice Presentation
Reward Delivery
fERN Amplitude (uV)
0
1
2
3
1 2 3 1 2 3
Trial
Reward One
Reward Two
fERN Amplitude (uV)
0
5
10
1 2 3 1 2 3
Trial
Reward One
Reward Two
DATA
Prediction Error Amplitude
-0.5
0.0
0.5
1.0
1.5
2.0
2.5
1 2 3 1 2 3
Trial
Reward One
Reward Two
Prediction Error Amplitude
-0.5
0.0
0.5
1.0
1.5
2.0
2.5
1 2 3 1 2 3
Trial
Reward One
Reward Two
MODEL

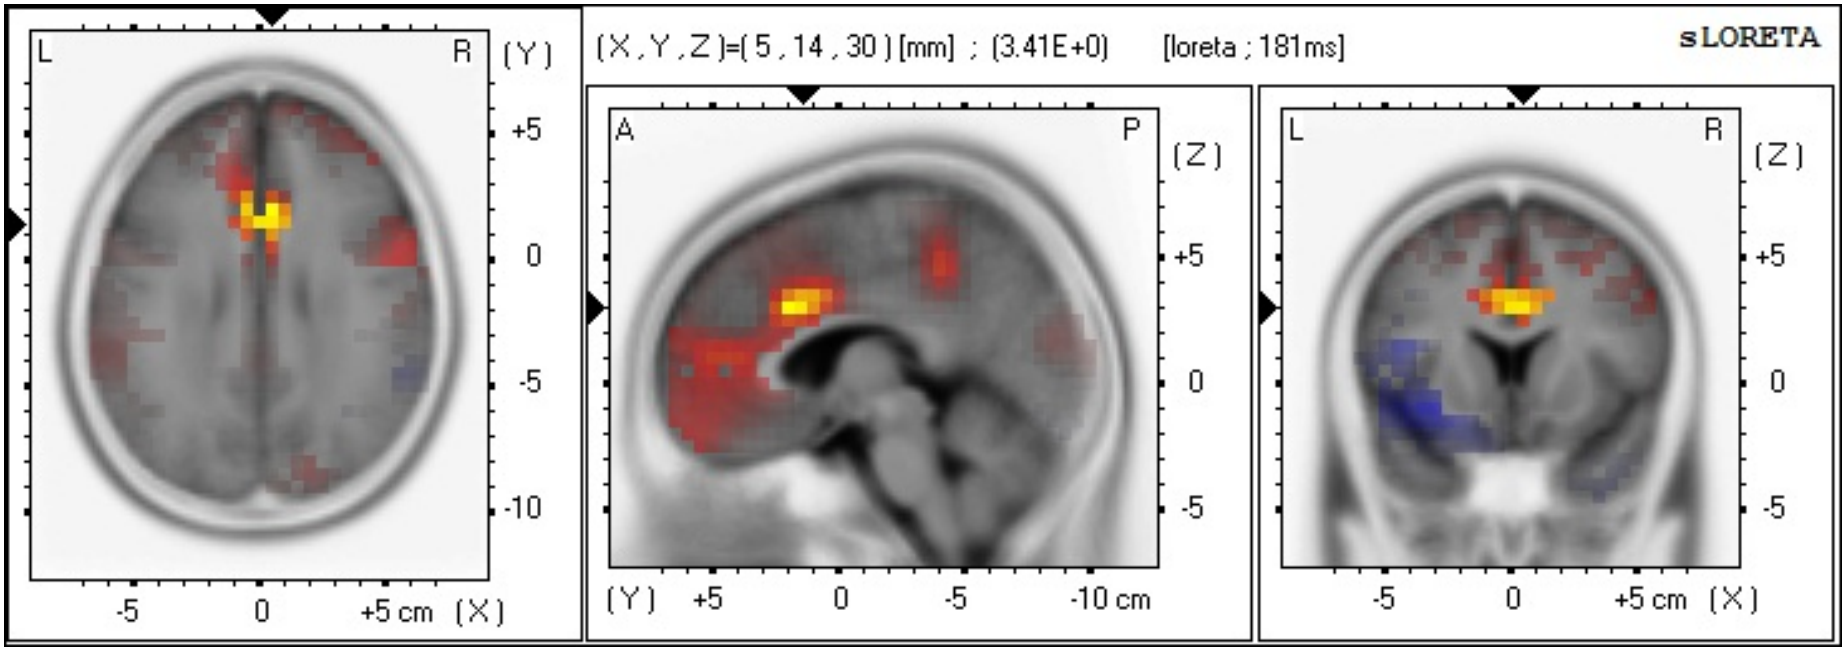
(X,Y,Z)=(5 , 14 , 30) [mm] ; (3.41E+0)
[loreta ; 181ms]
sLORETA
L
R
(Y)
+5
0
-5
-10
-5
0
+5 cm
(X)
A
P
(Z)
+5
0
-5
(Y)
+5
0
-5
-10 cm
L
R
(Z)
+5
0
-5
-5
0
+5 cm
(X)

Krigolson et al., 2013

Krigolson et al., 2013

Krigolson et al., 2013

Krigolson et al., 2013

Win!

Krigolson et al., 2013

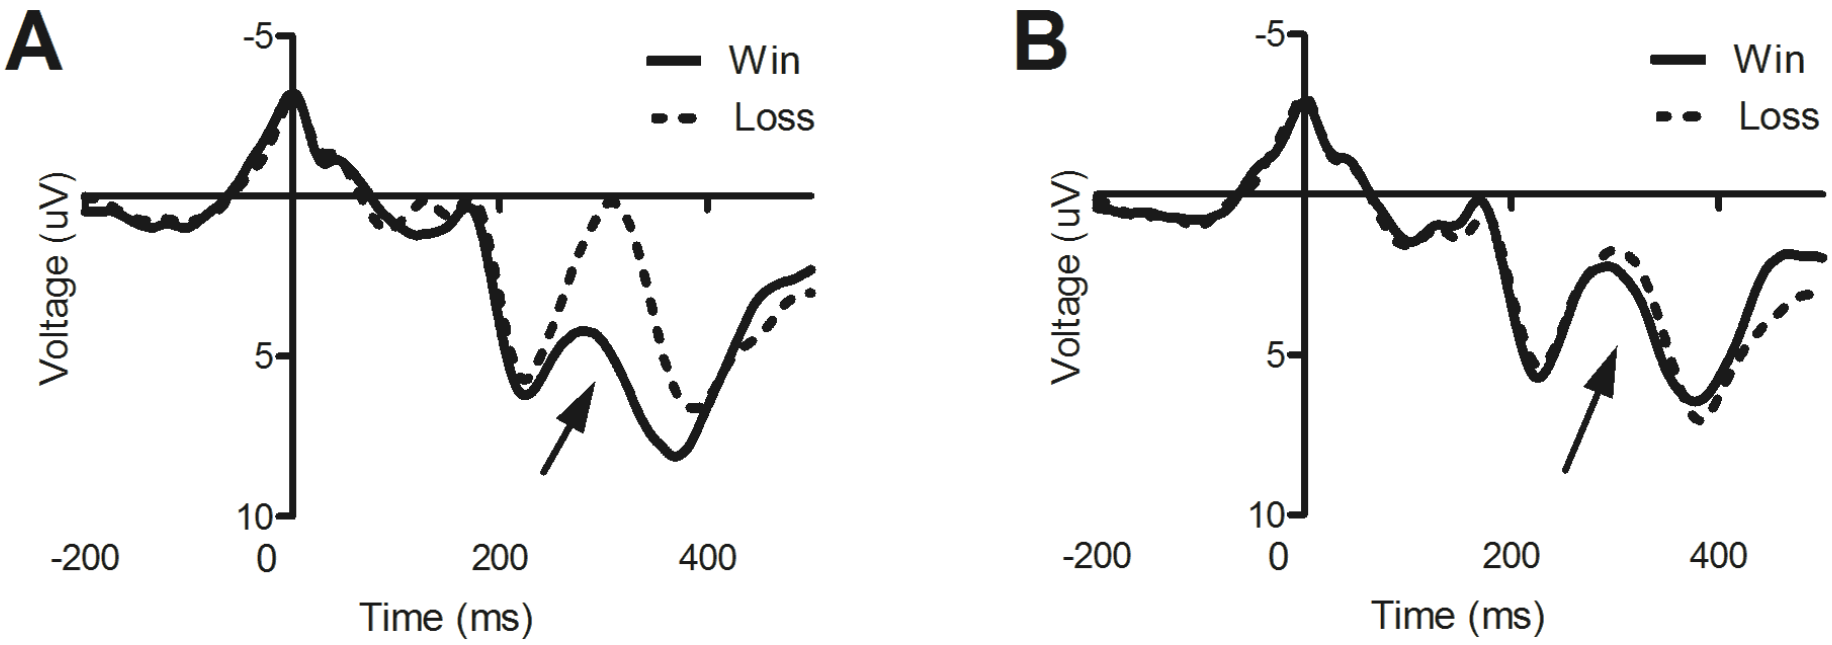

Krigolson et al., 2013