Can a single cell learn?

Ben Gaskin

University of Sydney

Introduction

How did we get here?

Methodology

- We will take a **broad functional definition of learning** as our point of departure, to be used in a **scoping review of relevant evidence**
  - Per Gershman et al. **"any persistent and adaptive modification of an organism's behavior as a function of its experience"**
- We will review the **evidence for reports of such processes in single cells**
  - We will further consider this by types: **habituation, associative conditioning**
  - As well as experimental evidence, we will look at **proposed mechanisms**
- Taken together, we will review the common structure of these cases and **propose a distinction** between this phenomenon and **learning proper**
  - Namely, **to distinguish between learning and adaptive behavioural plasticity**
  - We suggest a **substrate-neutral definition** in terms of **spatiotemporal patterns**

Historical controversy

- **Habituation** to mechanical stimulation in **Stentor roeselli** (Jennings, 1906)
  - This was interpreted as habituation in a simple organism
  - Some elements have been controversial, especially reports of ‘decision-making’
    - Later work failed to replicate, but used a different species (Reynierse & Walsh, 1967)
    - More recent work replicated successfully with same species (Dexter et al., 2019)
- **Associations** between wire and food in **paramecia** (Gelber, 1952)
  - Wire was coated in bacteria (attractive) and inserted into the medium
    - Paramecia congregate around the wire, drawn by presence of bacteria
    - When wire is later presented without bacteria, paramecia congregate
      - This does not occur when the process is conducted in the dark (Gelber, 1952)
  - This was interpreted as associative learning in a simple organism
    - Controversy ensued, claims of contaminated medium (Jensen, 1957)
    - Later replication found stirring eliminated effect (Katz & Deterline, 1958)

“Reconsidering the evidence” (Gershman et al.)

- This paper seems to have three aims—
  1. To **review the evidence for unicellular learning**
  2. To **assess the possibility of conserved mechanisms**
  3. To emphasise and rehabilitate Gelber’s contributions
- We will **focus on the first** of these, but keep in mind the second
  - Single cell learning as hinting the possibility of **“an evolutionary solution to the information processing challenges common to all living systems”**
  - The presence of learning in single cells is obviously a prerequisite here

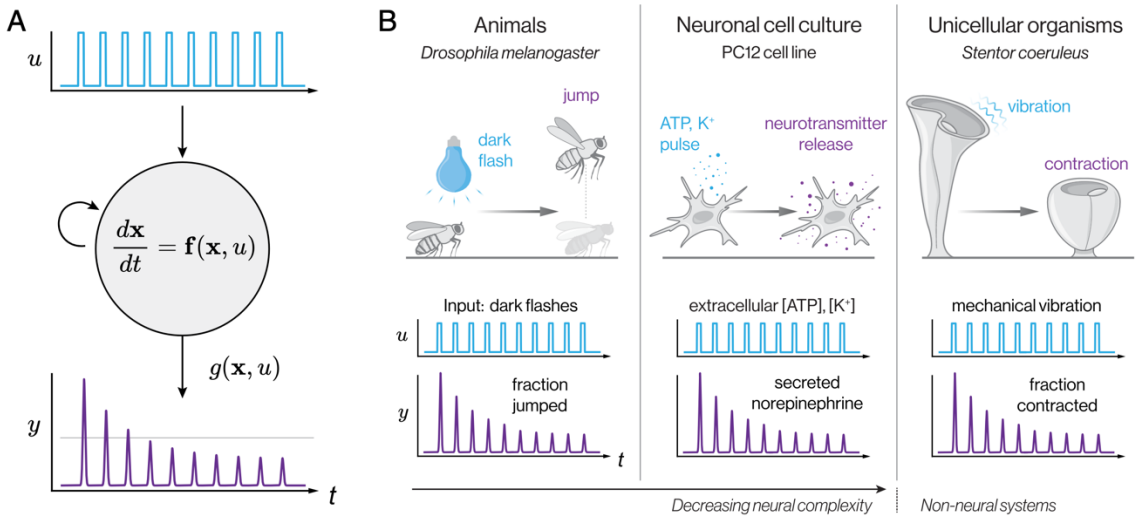

Habituation

Where response decreases with repeated application of a given stimulus.

Chemotaxis in E. coli

- They use “run and tumble” taxis
  - Stimulus decreases tumble-rate
  - Adaptation to levels, biased walk
- This adaptation is habituation
  - Phosphorylation cascade (CheA → CheY) promotes tumbling
  - Attractant binding inhibits cascade, reducing tumbling
  - This disrupts CheR–CheB balance, leads to increased methylation
  - Receptor sensitivity decreases with increase in methylation
- This lets them integrate state over time through negative feedback

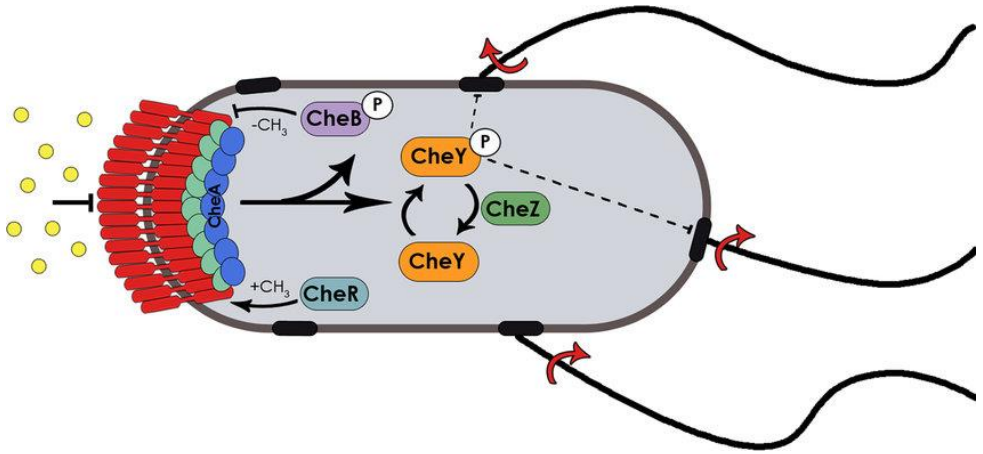

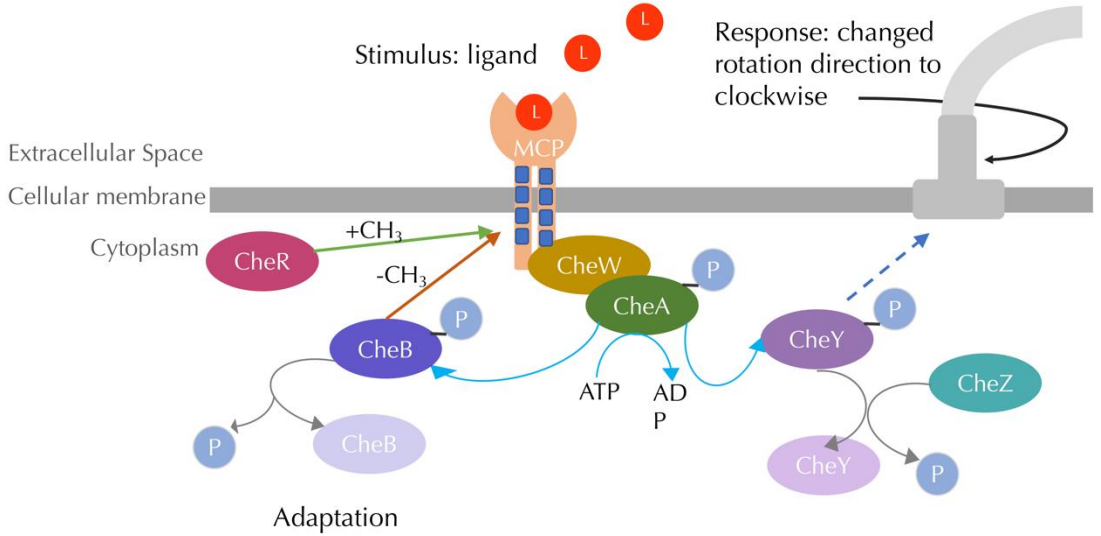

Habituation in Stentor roeselli

- Complex hierarchical response to stimulation (Dexter et al., 2019)
  - Bending
  - Ciliary alterations
  - Contraction
  - Detachment
- Response shows habituation (Wood, 1988)
  - Escape behaviours impede feeding
  - Learning allows discriminate response

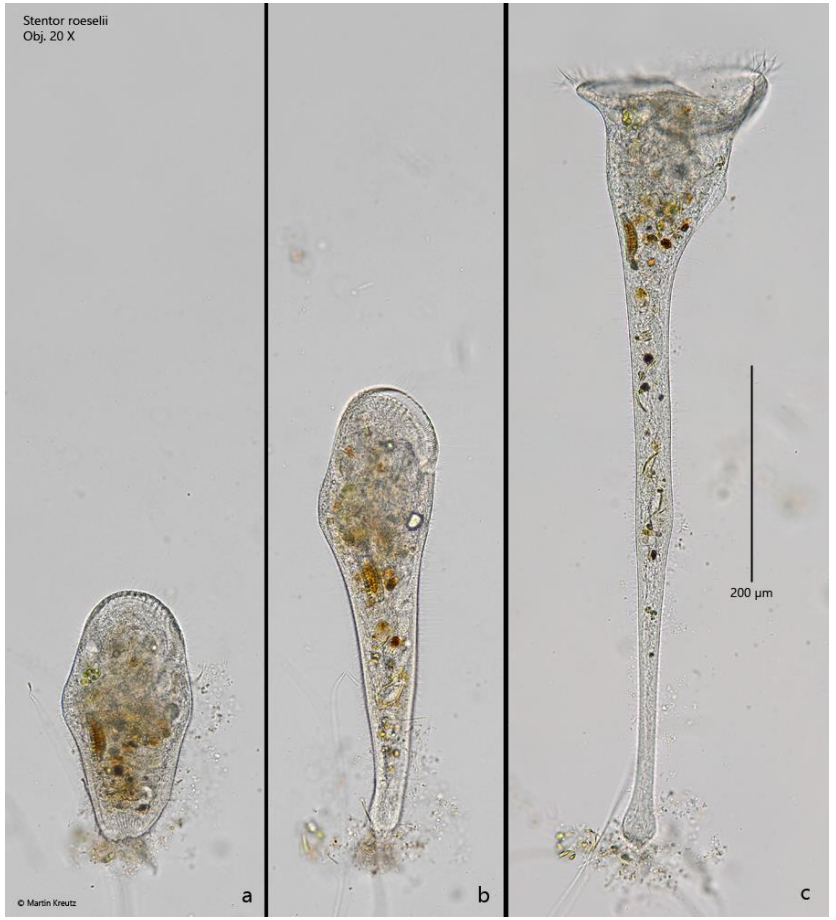

Molecular mechanisms for habituation

- Two-state model
  - States: responsive, non-response
  - Transition probabilities, forward increased when stimulus applied
- Receptor inactivation model
  - Stimulus degrades the receptor
  - Continuous receptors produce step-like population behaviour
- Concatenated IFF model
  - Nested motifs, incoherent feed-forward; stimulus produces response + memory
  - Memory inhibits response, decays slower

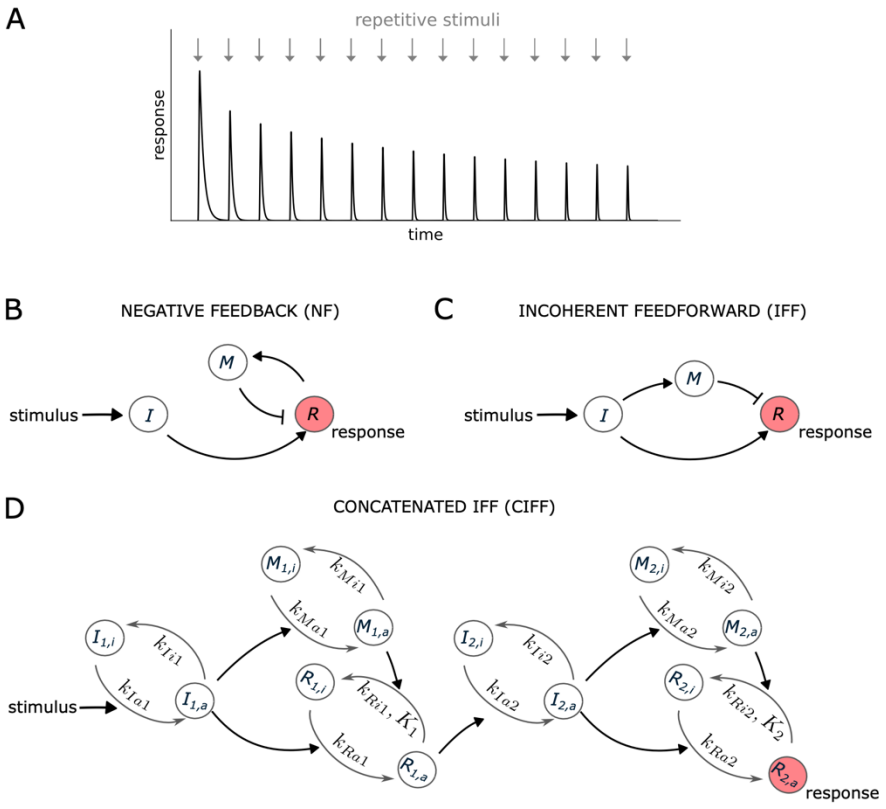

The hallmarks of habituation

Hallmark	Stentor roeselli	Two-State Model	Receptor Inact. Model	CIFF Model
1. Basic habituation	✅	✅	✅	✅
2. Recovery	✅	❓	✅	✅
3. Potentiation	✅	❓	✅	✅
4. Frequency sensitivity	✅	❓	✅	✅
5. Intensity sensitivity	✅	❓	✅	✅
6. Subliminal accumulation	✅	❓	✅	✅
7. Stimulus specificity	❓	❓	✅	❓
8. Dishabituation	❓	❓	❓	❓
9. Habituation of dishabituation	❓	❓	❓	❓
10. Long-term habituation	❓	❓	❓	❓

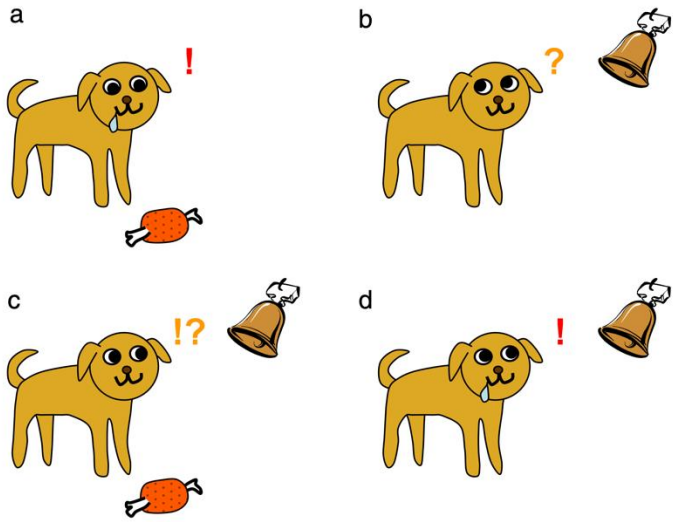

Associative conditioning

An initially neutral stimulus comes to elicit a conditioned response when it reliably predicts the occurrence of an attractive or aversive stimulus.

Paramecia: wire + bacteria (Gelber)

- Platinum wire inserted, either clean or coated with bacteria
  - Paramecia **congregated at wire when bacteria-coated**
  - **Post-conditioning, congregate even at clean wire**
    - Conditioned subjects **did not congregate in the dark** (1956)
    - Effect **endured up to 3 hours** after spaced training (1958)

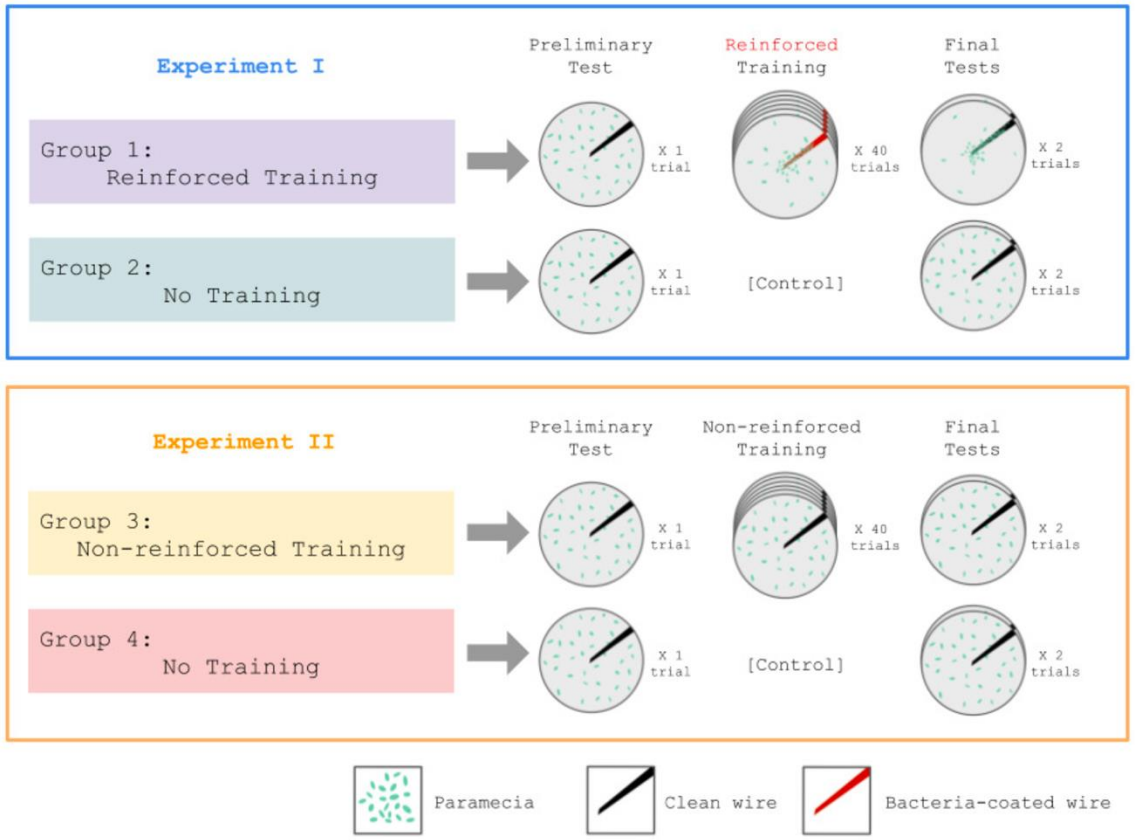

Paramecia: shock + vibration (Hennessey et al.)

- Conditioning by a four second **vibration**, ending with a two second **shock**
  - **Initially vibration alone, no reversal**
  - After training**, reliably induced reversal**
- Control procedures—
  - **Unpaired stimulus** (explicitly unpaired)
  - **Stimulus specificity** (different frequencies)
  - **Extinction** (to prove reversibility)
  - **Truly random** (contingency degradation)

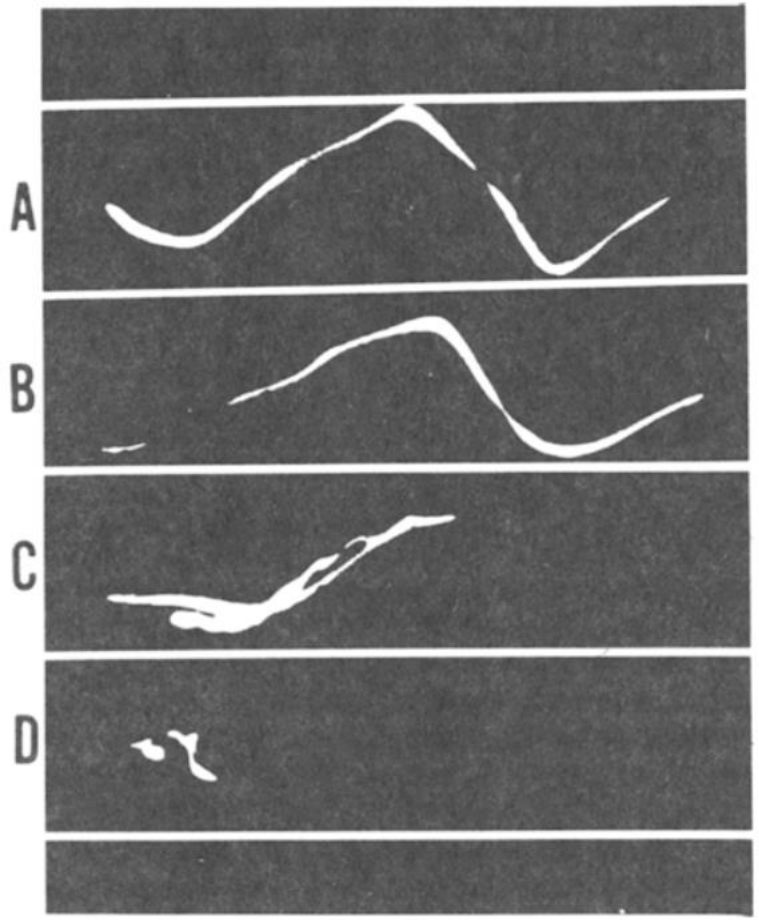

Amoeba: peptide + anode (Carrasco-Pujante et al.)

- Three species of amoeba
  - **Negative galvanotaxis**
  - **Positive chemotaxis**
- Conditioning presents **galvanotaxis and chemotaxis together**
  - Those that move to the anode are taken to test
  - **Robust reversal of prior**

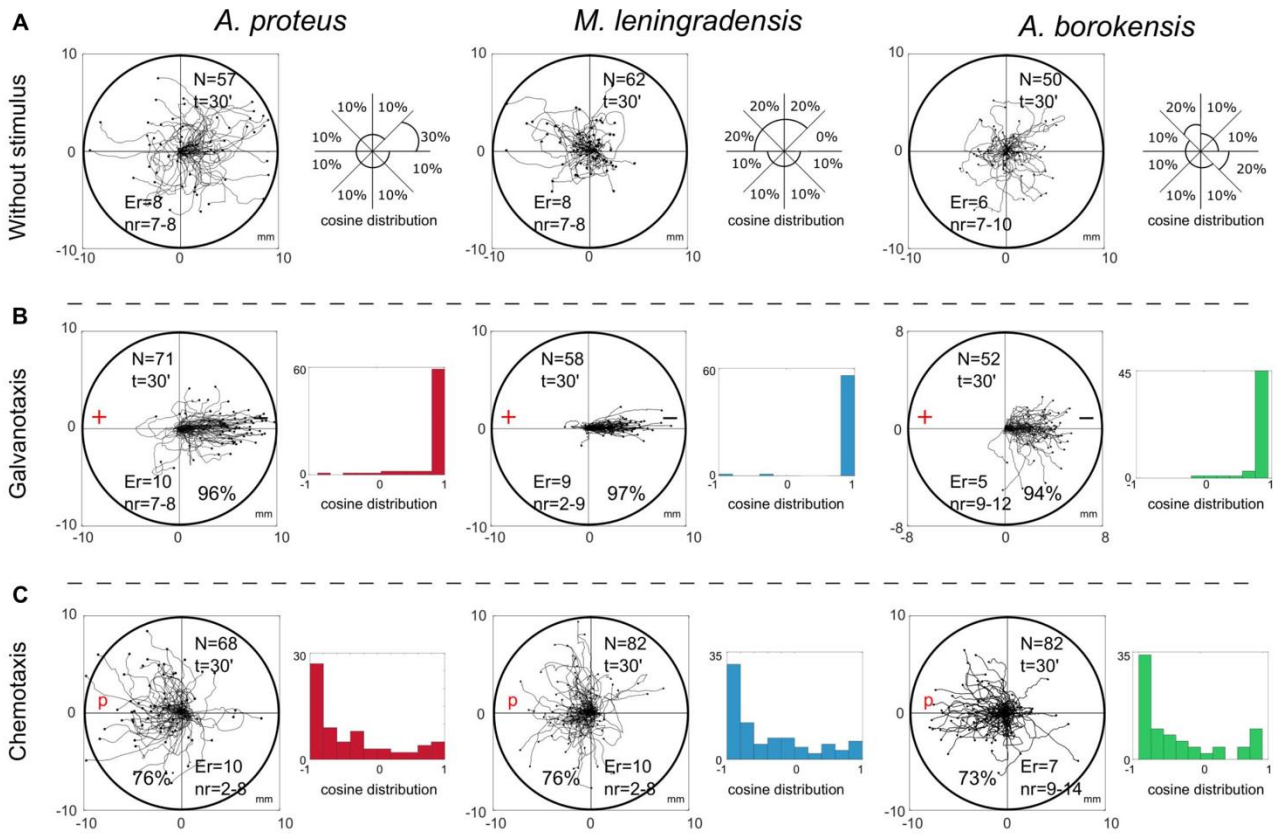

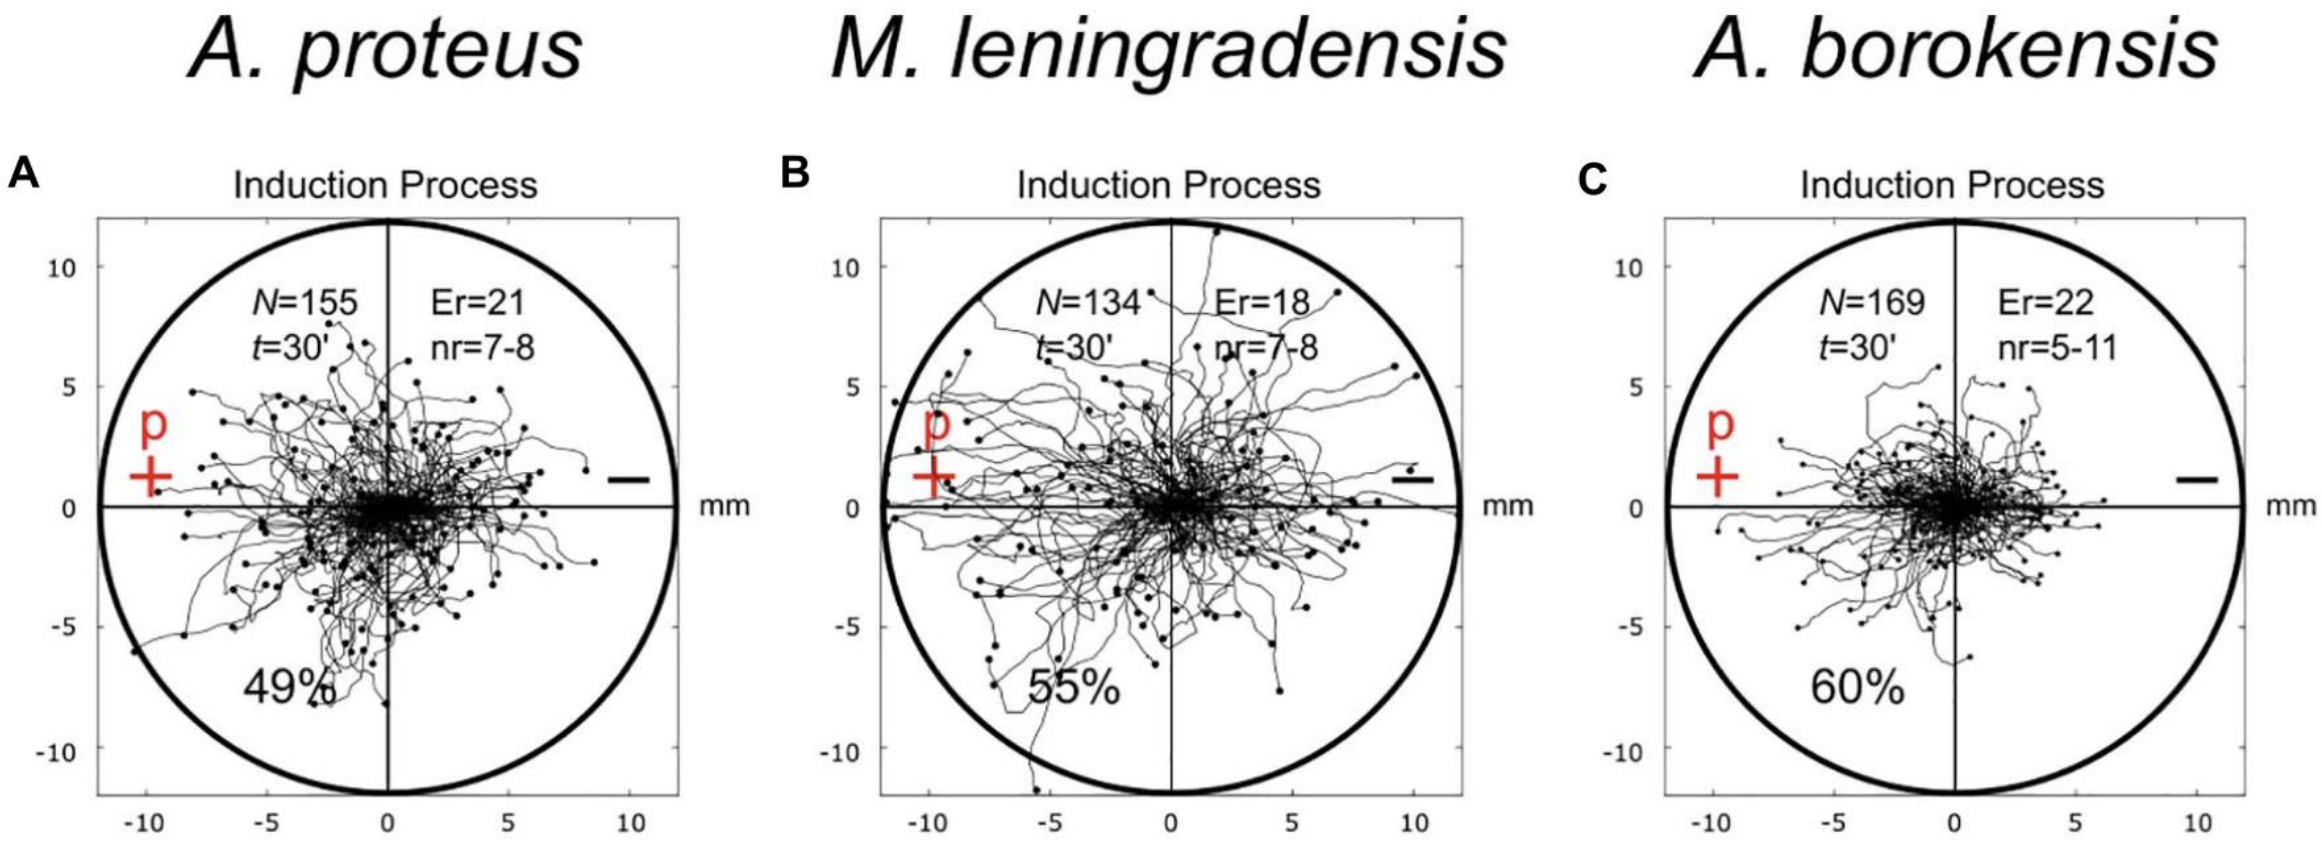
A. proteus
M. leningradensis
A. borokensis
A
B
C
Induction Process
Induction Process
Induction Process
N=155
t=30'
Er=21
nr=7-8
N=134
t=30'
Er=18
nr=7-8
N=169
t=30'
Er=22
nr=5-11
p
+
−
49%
55%
60%
mm
10
5
0
-5
-10

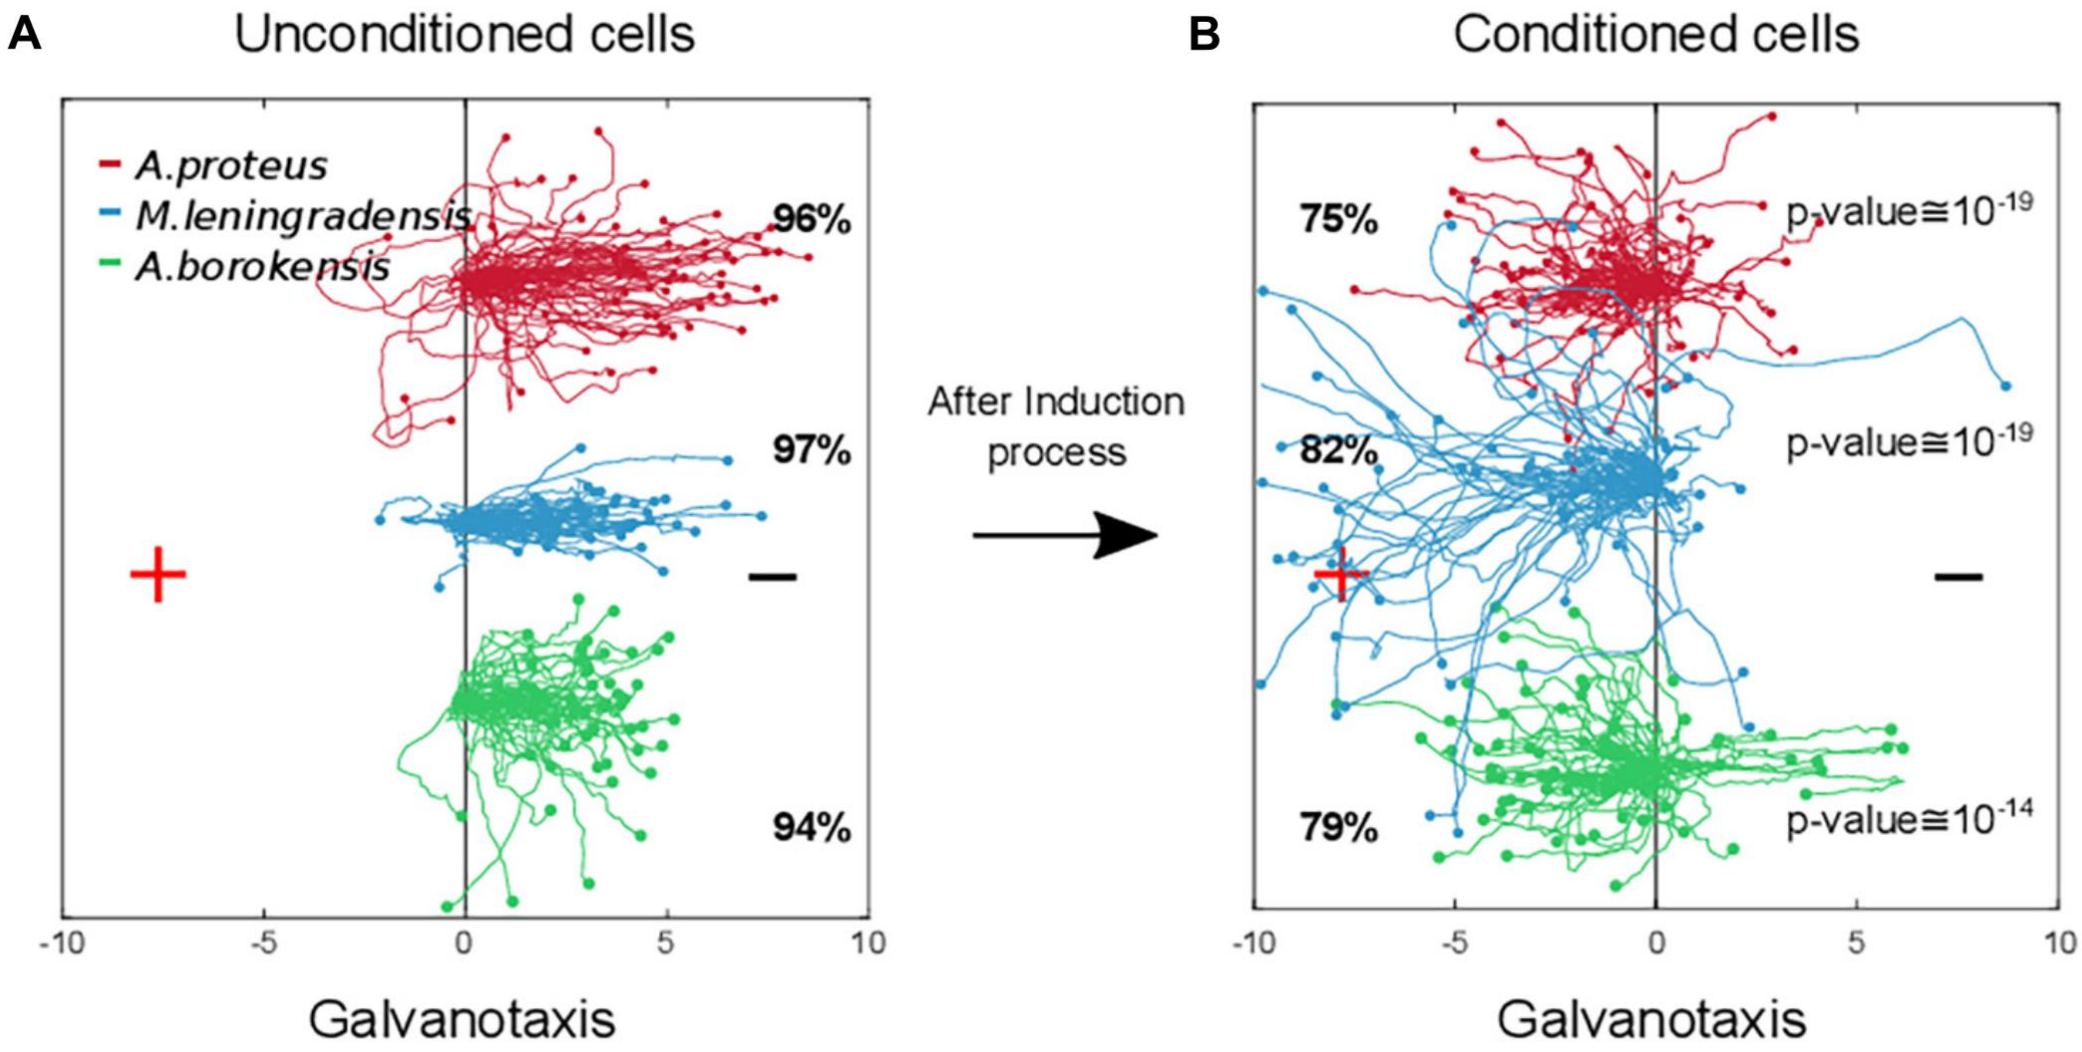
A
Unconditioned cells
A.proteus
M.leningradensis
A.borokensis
96%
97%
94%
-10
-5
0
5
10
Galvanotaxis
After Induction process
B
Conditioned cells
75%
82%
79%
p-value≅10-19
p-value≅10-19
p-value≅10-14
-10
-5
0
5
10
Galvanotaxis

Molecular mechanisms for associative learning

- Fernando et al. (2009), provide a model for **how single cells might in principle accomplish associative learning** via **GRNs** or **phosphorylation cycles**
  - The u_1 and u_2 operate by removing repressors r_1 and r_2, which allows w_1 and w_2 to express p
  - Here w_1 has a moderate baseline, while w_2 has a nil baseline
  - Both w_1 and w_2 are generated by p, with this being gated by the repressors r_1 and r_2
- This is limited to u_1 and u_2, which are **linear intensities of specific molecules**
  - These may be downstream of allosteric gates, but are still expressed as scalar concentrations

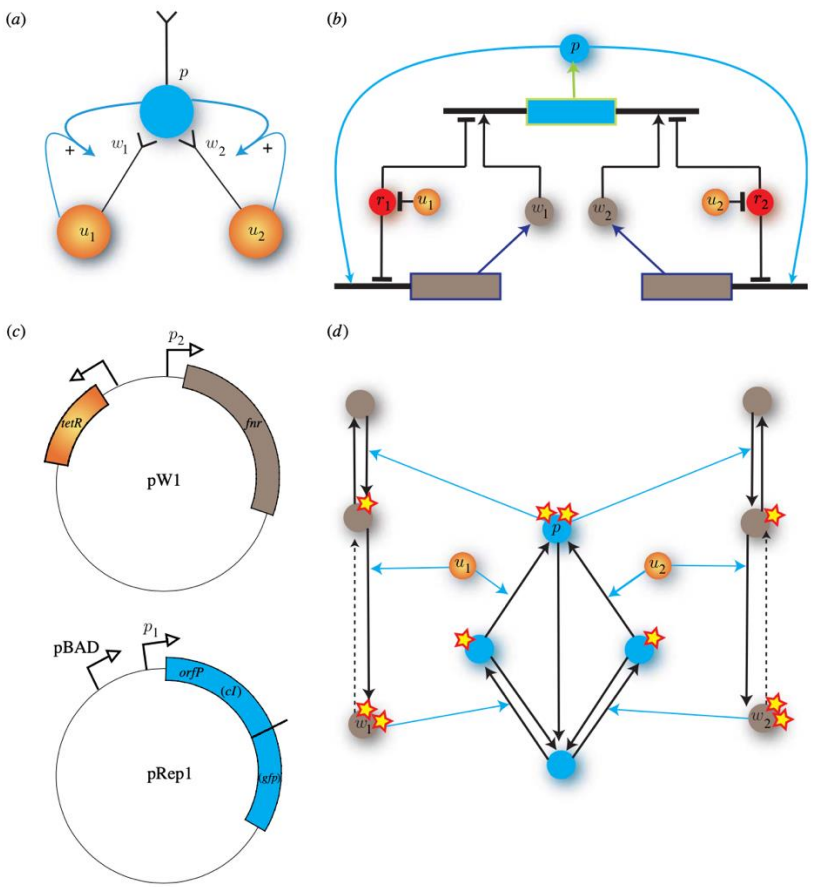

Rescorla's truly random control

- Rescorla **criticises classic control procedures** (explicitly unpaired, etc.)
  - Confound non-associative with removal of contingency
  - Create negative contingency rather than eliminating contingency
- He proposes the **"truly random control"** (1967)
  - CS and US occur **independently with no contingency between them**
  - Distinguishes two views: pairing and contingency
- These are later tested in experiments with rats (1968)
  - Equal shock probability in / out of CS (random), versus only during CS (gated)
    - **Gated showed conditioning, random did not**
  - Manipulated shock probabilities in / out of CS across groups
    - **Contingency predicted conditioning better than pairing quantity**

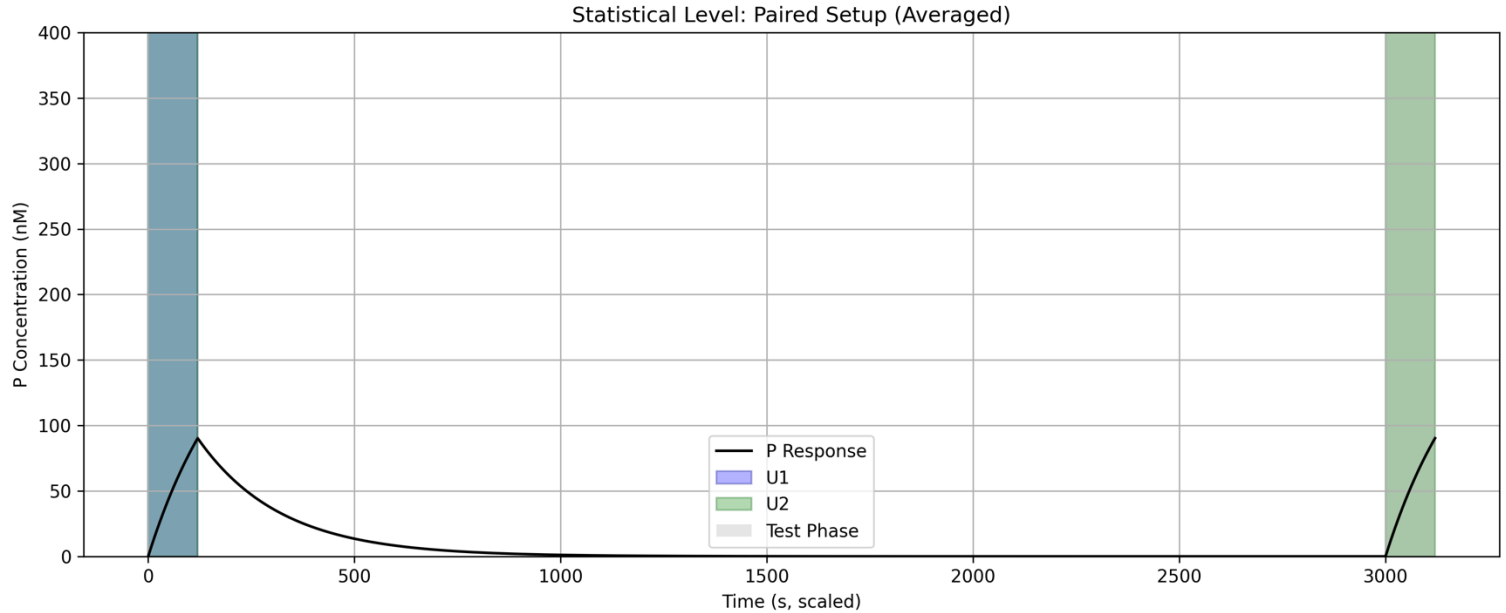
Statistical Level: Paired Setup (Averaged)
P Concentration (nM)
Time (s, scaled)
P Response
U1
U2
Test Phase
0
50
100
150
200
250
300
350
400
0
500
1000
1500
2000
2500
3000

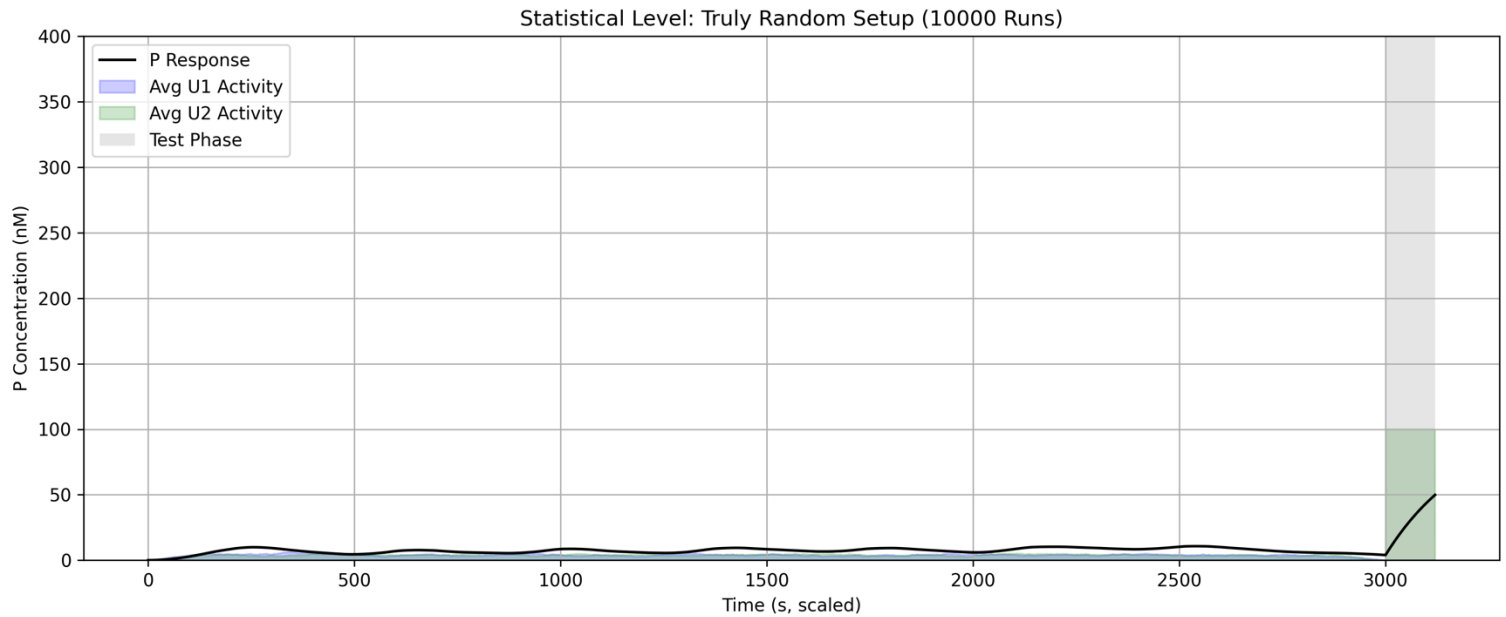
Statistical Level: Truly Random Setup (10000 Runs)
P Concentration (nM)
Time (s, scaled)
P Response
Avg U1 Activity
Avg U2 Activity
Test Phase
0
50
100
150
200
250
300
350
400
0
500
1000
1500
2000
2500
3000

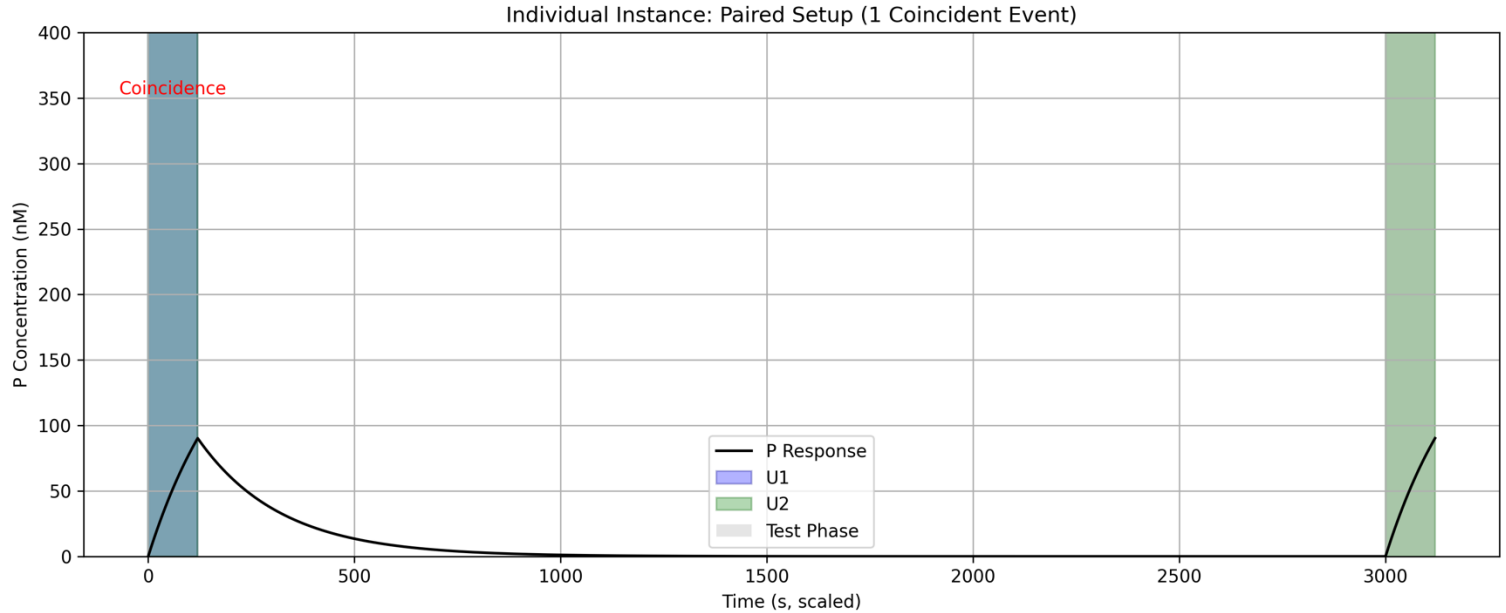
Individual Instance: Paired Setup (1 Coincident Event)
Coincidence
P Concentration (nM)
Time (s, scaled)
P Response
U1
U2
Test Phase
0
50
100
150
200
250
300
350
400
0
500
1000
1500
2000
2500
3000

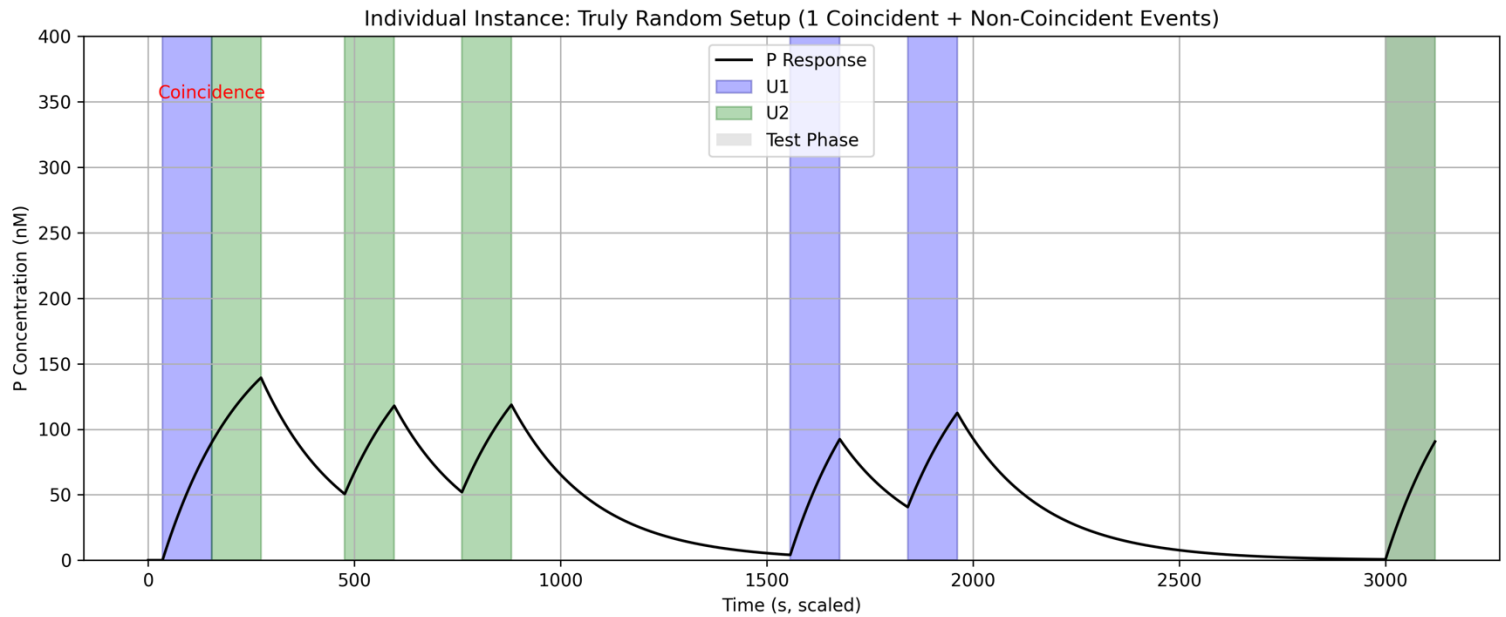
Individual Instance: Truly Random Setup (1 Coincident + Non-Coincident Events)
Coincidence
P Concentration (nM)
Time (s, scaled)
P Response
U1
U2
Test Phase
0
50
100
150
200
250
300
350
400
0
500
1000
1500
2000
2500
3000

Unicellular learning in multicellular organisms?

- Habituation—
  - PC12 from rats (McFadden & Koshland, 1990)
  - Human kidney cells (Bonzanni et al., 2020)
- Associative conditioning?
  - Negative results in human macrophages (Nilsonne et al., 2011)

Refractory periods as habituation

- **When a neuron has recently fired, it is less likely to fire again**
  - **Absolute** refractory period (voltage-gated sodium channels are briefly inactivated)
  - **Relative** refractory period (hyperpolarised membrane requires stronger stimulus)
- In either case, these represent a form of **habituation**: when the occurrence of an event makes this less likely to occur after
  - This varies across neuron types (bursting, chittering, etc.) but is ubiquitous throughout

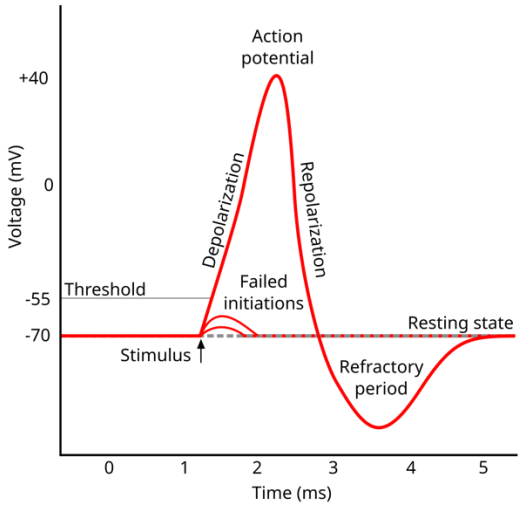

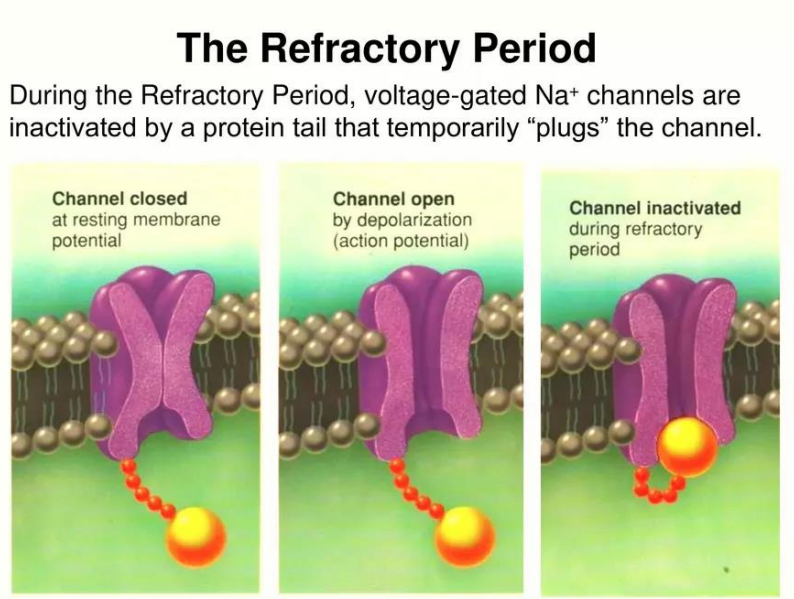

Spike-timing-dependent plasticity

- Temporally-precise version of Hebb's law:
  - "When an axon of cell A is near enough to excite a cell B and **repeatedly or persistently takes part in firing it**, some growth process or metabolic change takes place in one or both cells such that **A's efficiency, as one of the cells firing B, is increased**" (Hebb, 1949).
    - From "**fire together, wire together**" to "**out of sync, lose your link**" (Shatz, 1992)
- Bidirectional plasticity—
  - Long term potentiation (LTP)
  - Long term depression (LTD)

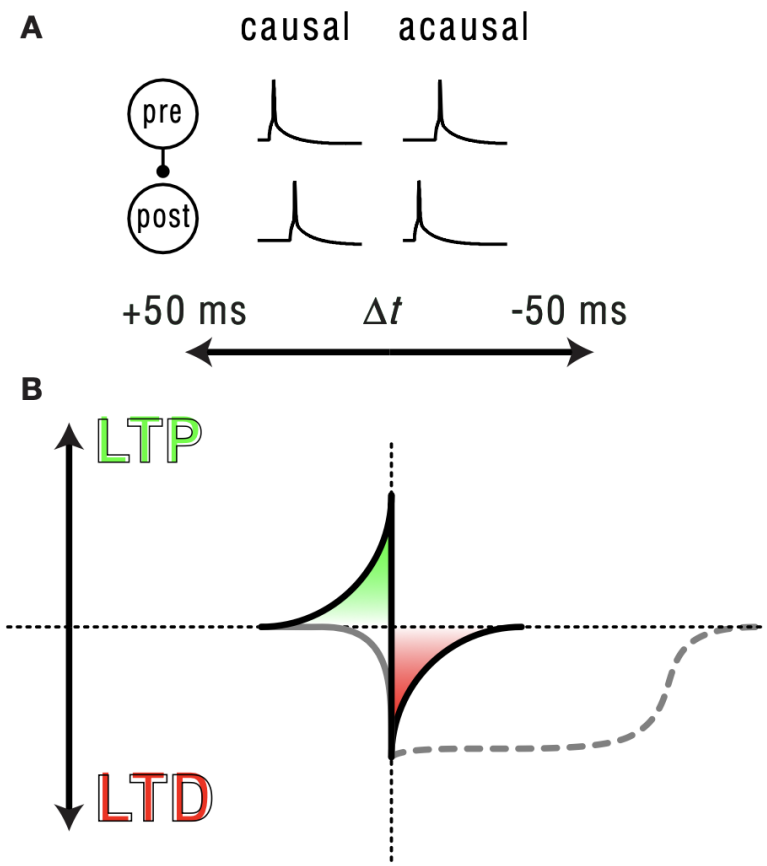

STDP as associative learning

- If we take a neuron with **two synapses, hence two inputs**—
  - Suppose that **one** of them is weighted such that **any presynaptic spike will cause a firing** in the postsynaptic neuron
  - The **other** synapse starts out such that **a presynaptic spike will not trigger the firing** of the postsynaptic neuron
- This setup will further **satisfy Rescorla's truly random control**:
  - Where they fire in sync, LTP
  - Where they fire out of sync, LTD
- Rescorla's truly random control requires bidirectional plasticity
  - This gives the statistical quality which he sees as its essential character

Paramecium, the swimming neuron

- Hennessey et al. (1979) **satisfy Rescorla's truly random control**
  - This, however, has **not since been replicated**
  - Nevertheless, **suggests bidirectional plasticity**
- Paramecium as the "swimming neuron" (Brette, 2021)
  - Avoidance reaction (as in Hennessey et al.) **triggered by action potential**
    - Similar effect triggers contraction in Stentor roeselli
  - Limited evidence to suggest bidirectional plasticity, **otherwise only STDP**
- This may suggest **some phylogenetically conserved mechanism**
  - Perhaps of the sort that would satisfy Gershman et al.
  - It may be as much the role of action potentials as associative mechanisms

Phylogenetically conserved mechanisms

- Recall, this was the secondary interest in Gershman et al.
  - Word on the street, Gershman doesn't believe in the paper anymore
  - Subsequent paper (2023) reconciles connectionist and computational
- STDP can provide roughly what was sought here—
  - **Conserved mechanism underlying multicellular memory and learning**
  - Exceeds most associative learning demonstrated in unicellular organisms
  - This obviously depends on a molecular instantiation of weighting
    - The line between synaptic and molecular mechanisms of learning is unclear
- The trouble with these cases is the **identity of signal and medium**
  - Most of these mechanisms have **little possibility of flexibility or generality**

Conclusion

Where are we going?

What is learning?

- The stimuli in these cases (even neurons) are **simple linear intensities**
- They involve the **constrained modulation of strictly specified pathways**
  - Fernando et al., "capable of learning only associations between N pre-defined dimensions of conditional stimuli, u_n and an unconditioned stimulus."
- Fernando et al. note this specific distinction, give three responses—
  - First, that u_n can be downstream of a range of signal transduction cascades, hence we might consider this as a "perceptual class" amenable to learning
    - What would it take to generate a novel perceptual class within a lifetime?
  - Second, that there is no absolute generality of association in nervous systems
  - Lastly, that **"nervous systems allow greater capacity for constructing novel pathways between arbitrary stimuli, because, instead of the network being defined by unchanging nucleotide and amino acid sequences, the spatial location of a synapse defines the pathway and 'meaning' of a signal."**

What can networks do that cells alone cannot?

- Birch, Ginsburg, Jablonka (2010) address precisely this, albeit in a prelude to the Unlimited Associative Learning framework: **“what is the added cognitive value of the neural type of learning?”**
- They provide four elements, of which only one concerns a truly fundamental addition (rather than an enhancement) of learning:
  - **“Crucially, the action potential is the same whether carried by sensory or motor nerves, and all modes of sensory stimuli—photons, chemicals, heat, sound waves and other mechanical types of energy—are translated into this common communication currency, the electrical impulse. This allows information from various sources to be integrated and is the basis of mapping within the nervous system of patterns of stimuli emanating from the world and the body.”**

What exactly happens in networks?

- **Distributed dynamical computation** (Gong & Leeuwen, 2009)
  - **Spatiotemporal patterns propagate through the network**
  - These are **modified throughout by the synaptic weightings**
  - The collision of these waves lead **to interference patterns**
  - **Interference with thresholding gives rise to logical operations**
- Through STDP, networks **detect hidden causes** (Nessler et al., 2009)
  - Lateral inhibition, here implemented through winner-takes-all mechanism
  - **Higher-dimensional inputs are reduced to essential statistical features**
- This allows features of associative learning (Jablonka & Ginsburg)—
  - Novel stimuli, due to flexibility of **receptors arrayed as sensory surfaces**
  - Compound stimuli, due to **compression of patterns across deep networks**

What would this mean for unicellular learning?

- **Behavioural plasticity ✓**
  - Habituation ✓
    - Basic habituation ✓
    - Several hallmarks ✓
  - Associative conditioning
    - Coincident ✓
    - Contingent ?
- **Learning proper ✗**
  - Depends upon molecular mechanisms of behavioural plasticity ✓
  - Further leverages networked structures to provide for flexibility ✗
    - Requires specific receptor and effector structures, sensorimotor surfaces ✗

What do we mean by plasticity and learning?

- **Plasticity—**
  - The behaviour of a system is modified by **scalar molecular concentrations**
- **Learning—**
  - The behaviour of a system is modified by **complex spatiotemporal patterns**
- This is a **substrate-neutral distinction**, theoretically learning of this sort could also be implemented in single cells (and may well be)
  - Physarum polycephalum, for instance, may use **oscillations** combined with its **multinucleate structure** to learn temporal patterns (Saigusa et al., 2008)

What is the value of this distinction?

- **Stimulus generalisation curves** are not easy to explain in a molecular framework
  - Hennessey et al. specifically found **paramecia discriminated between tone frequencies**
    - This **may be confounded by contingency learning**
    - We would **encourage replications to focus here**
  - This is especially prominent for **natural kinds**
    - Visual concepts, as **Herrnstein's pigeon studies**
    - Phobias, as in the apocryphal case of **Little Albert**
  - These can be explained, however, via the **self-assembly of hierarchical network features**
- **Temporal patterns** are also a difficulty here
  - This may require more (recurrence, oscillations, etc.)

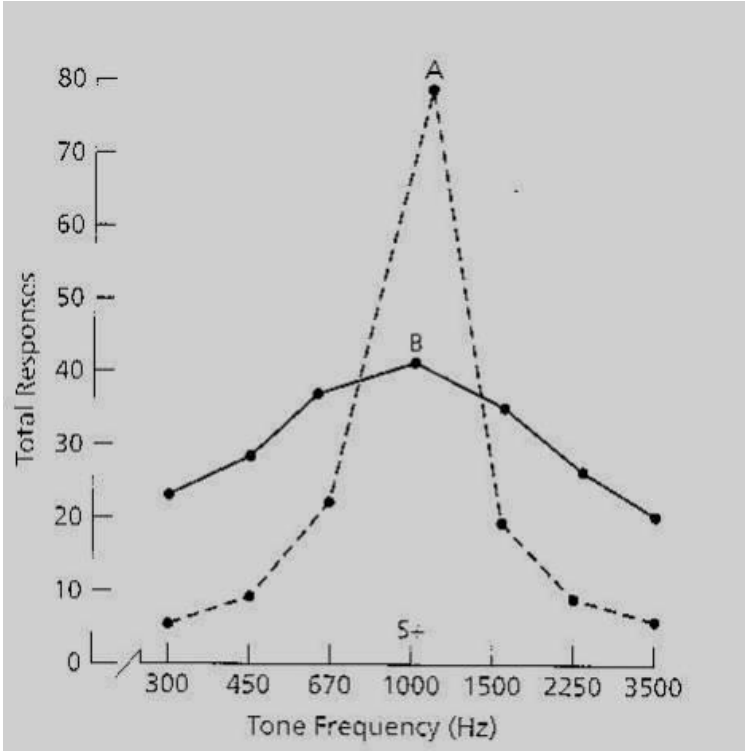

FIGURE 7.4 Generalization Gradients. Examples of steep (A) and shallow (B) generalization gradients after training on a 1000-Hz tone.