1-HOUR GUIDED TUTORIAL

From ion channels to memories

How the right representation changes when the question changes scale.

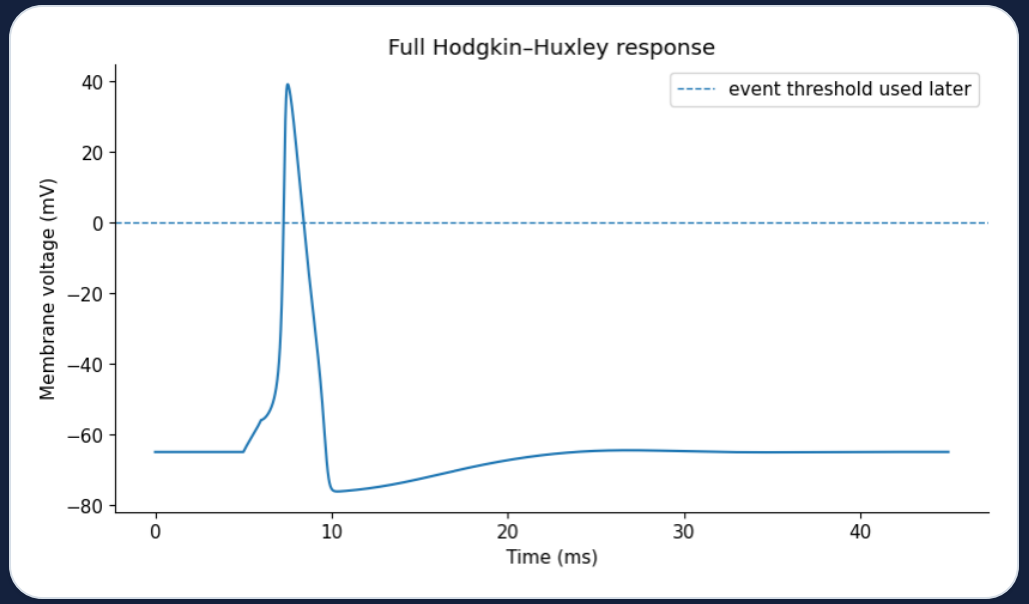

One neuron. Four states.
Then: one event.

A neuron carries voltage — then passes the signal on

Within a neuron, the signal travels along the axon. Between neurons, it crosses a synapse.

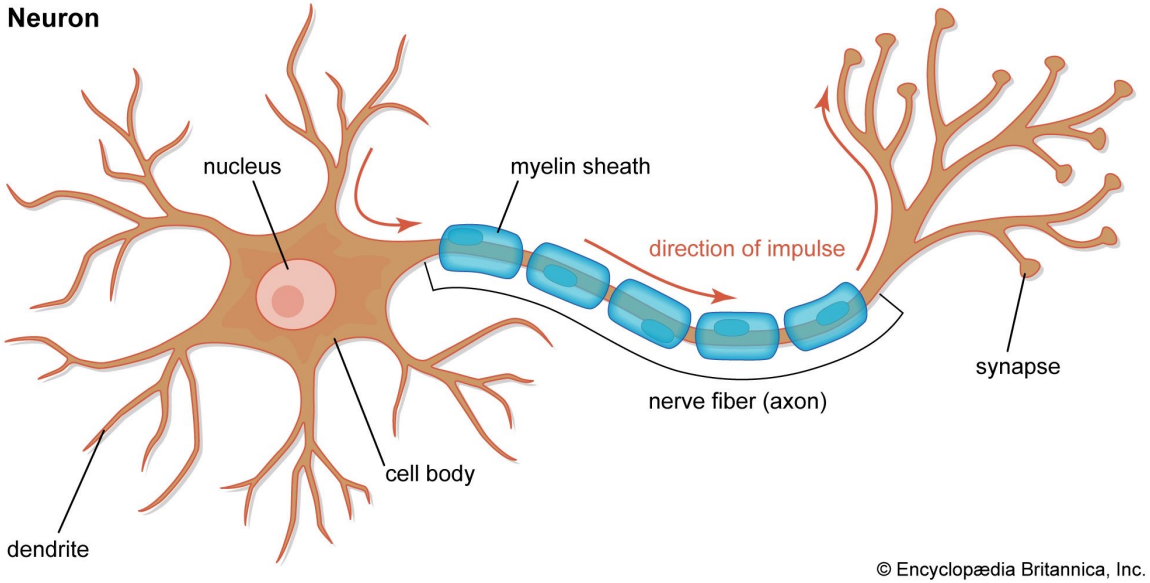

along one neuron

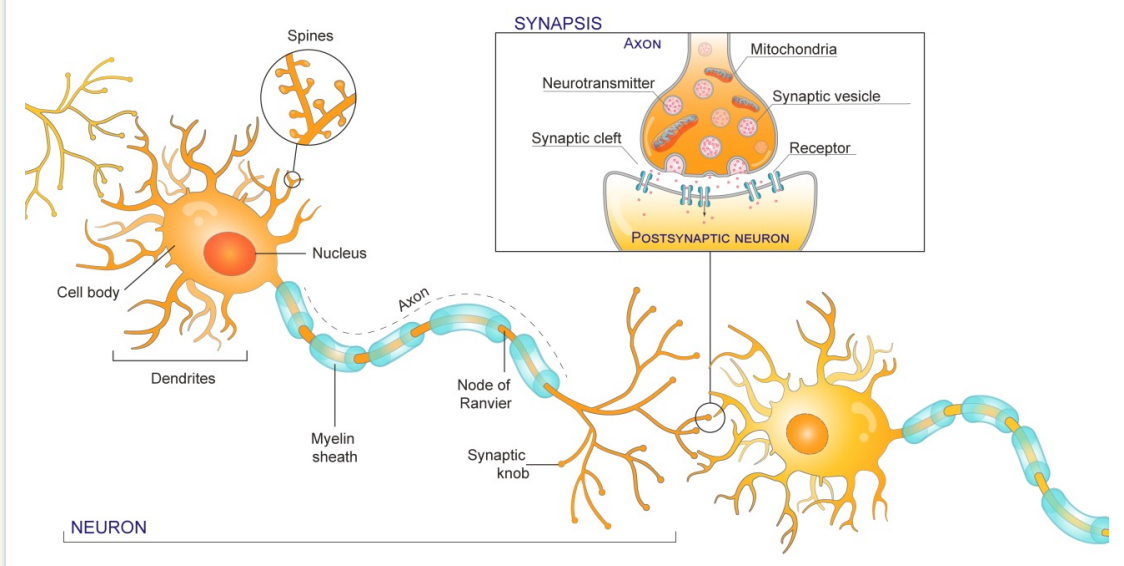

from one neuron to the next

Action potential = a quick change in voltage across the membrane

Here is where we are going

First simplify one neuron; then let its events interact and change a network.

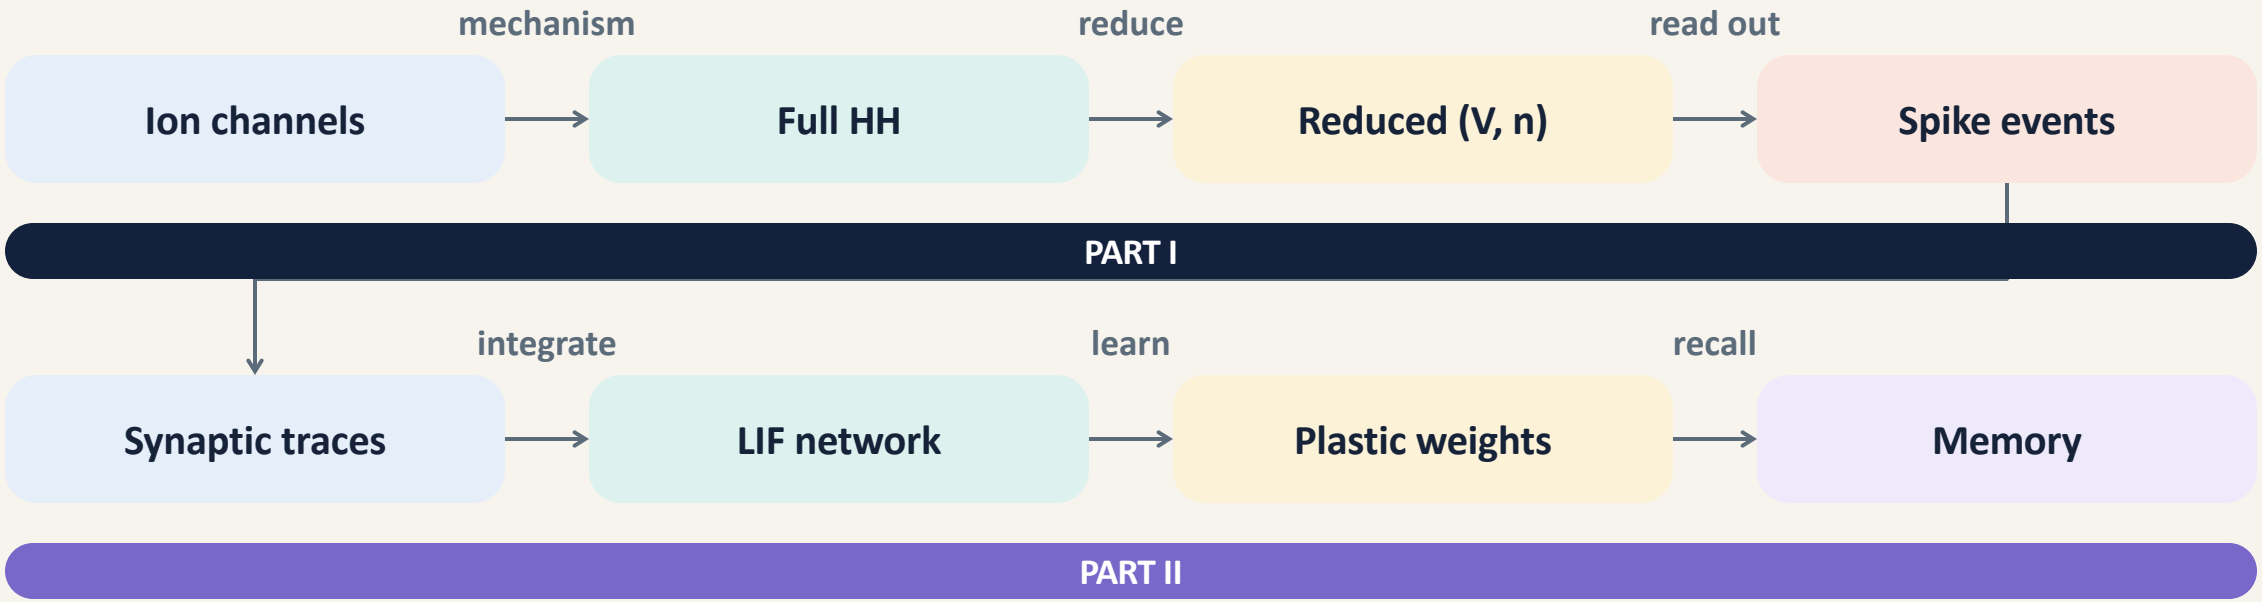

Each arrow is not just simplification. It is a change in what information is needed.

What can we safely leave out?

At each step, keep what the next question needs.

KEEP

What must still work?

DROP

What no longer matters?

CHECK

When does the shortcut break?

1. Capacitance turns net current into voltage change

The membrane stores separated charge, so voltage cannot jump instantaneously.

$$Q = C_m V$$

stored charge = capacitance × voltage

because current is the rate of change of charge: I = dQ/dt

$$C_m \frac{dV}{dt} = I_{\text{ext}} - I_{\text{Na}} - I_{\text{K}} - I_{\text{L}}$$

the unbalanced current charges or discharges the membrane

Causality: net current changes membrane charge; membrane charge sets voltage.

2. Equilibrium potential defines zero ionic current

The Nernst equation balances the concentration gradient against the electric field.

$$E_{\text{ion}} = \frac{RT}{zF} \ln\left(\frac{[\text{ion}]_{\text{out}}}{[\text{ion}]_{\text{in}}}\right)$$

voltage at which this ion has no net electrochemical drive

distance from equilibrium

$$\text{driving force} = V - E_{\text{ion}}$$

concentration gradient chemical tendency to diffuse

membrane voltage electrical force on charge z

at V = Eion the two tendencies balance

V = Eion → no net ionic current
|V – Eion| → strength of the drive
sign of V – Eion → current direction

$$I_{\text{ion}} = \bar{g}_{\text{ion}}\, p_{\text{open}}\, (V - E_{\text{ion}})$$

Conductance sets capacity; V – Eion sets current direction and strength.

3. A gate switches between two coarse-grained states

For activation gates, C is closed/non-permissive and O is open/permissive.

$$C \underset{\beta_x(V)}{\overset{\alpha_x(V)}{\rightleftarrows}} O$$

C

closed / non-permissive
population fraction 1 − x

$\alpha_x(V)$

forward transition rate

$\beta_x(V)$

reverse transition rate

O

open / permissive
population fraction x

$$\text{flow into O} = \alpha_x(V)(1-x), \qquad \text{flow out of O} = \beta_x(V)x$$

4. Two-state kinetics becomes a gating ODE

For a gating variable x ∈ [0, 1], these two forms are equivalent.

$$\frac{dx}{dt} = \alpha_x(V)(1-x) - \beta_x(V)x$$

opening minus closing

$$\frac{dx}{dt} = \frac{x_\infty(V) - x}{\tau_x(V)}$$

target divided by timescale

$$x_\infty(V) = \frac{\alpha_x(V)}{\alpha_x(V) + \beta_x(V)}$$

target gate value at this voltage

$$\tau_x(V) = \frac{1}{\alpha_x(V) + \beta_x(V)}$$

how quickly the gate approaches that target

This same relaxation equation will justify our first reduction.

What do the gating variables represent physically?

Microscopic picture

Membrane voltage V

↓

membrane electric field

↓

charged voltage-sensor motion

↓

channel conformational change

Charged sensor domains feel the membrane electric field.
Changing V changes conformational energies and can favour opening, closing, activation or inactivation.

HH model

$$C \underset{\beta_x(V)}{\overset{\alpha_x(V)}{\rightleftarrows}} O$$

x is not the sensor's physical position.

x is a probability: the fraction of gates in the permissive state.

$$\frac{dx}{dt} = \alpha_x(V)(1-x) - \beta_x(V)x$$

$\alpha_x(V)(1-x)$: gates entering the permissive state
$\beta_x(V)x$: gates leaving that state

charged-domain motion → $\alpha_x(V)$, $\beta_x(V)$ → gating probability x → ionic conductance

KEY DISTINCTION

Charged-domain motion is the microscopic mechanism. HH does not model that motion explicitly; it compresses it into voltage-dependent rates and a gating probability.

5. Put the currents together

The spike appears on its own once fast excitation and slow recovery interact.

$$
\begin{aligned}
C_m \frac{dV}{dt} &= I_{\text{ext}}(t) - \bar{g}_{\text{Na}} m^3 h (V - E_{\text{Na}}) \\
&\quad - \bar{g}_{\text{K}} n^4 (V - E_{\text{K}}) - g_{\text{L}} (V - E_{\text{L}}), \\
\frac{dm}{dt} &= \alpha_m(V)(1 - m) - \beta_m(V) m, \\
\frac{dh}{dt} &= \alpha_h(V)(1 - h) - \beta_h(V) h, \\
\frac{dn}{dt} &= \alpha_n(V)(1 - n) - \beta_n(V) n.
\end{aligned}
$$

FAST + **V ↑ → m ↑ → inward Na^+ current ↑ → V ↑**

Regenerative excitation

SLOW – **V ↑ → n ↑ and h ↓**

K^+ exit strengthens while Na^+ entry weakens

STATE: (V, m, h, n)

INTERACTIVE

6. Let’s see whether the model spikes

Choose the pulse, make a prediction, then run the notebook cell.

1 Amplitude

How hard do we push?

2 Start time

When does the pulse arrive?

3 Duration

How long does it last?

Prediction first: below threshold, one spike, or repeated spikes?

INTERACTIVE OUTPUT

A brief pulse produces a full action potential

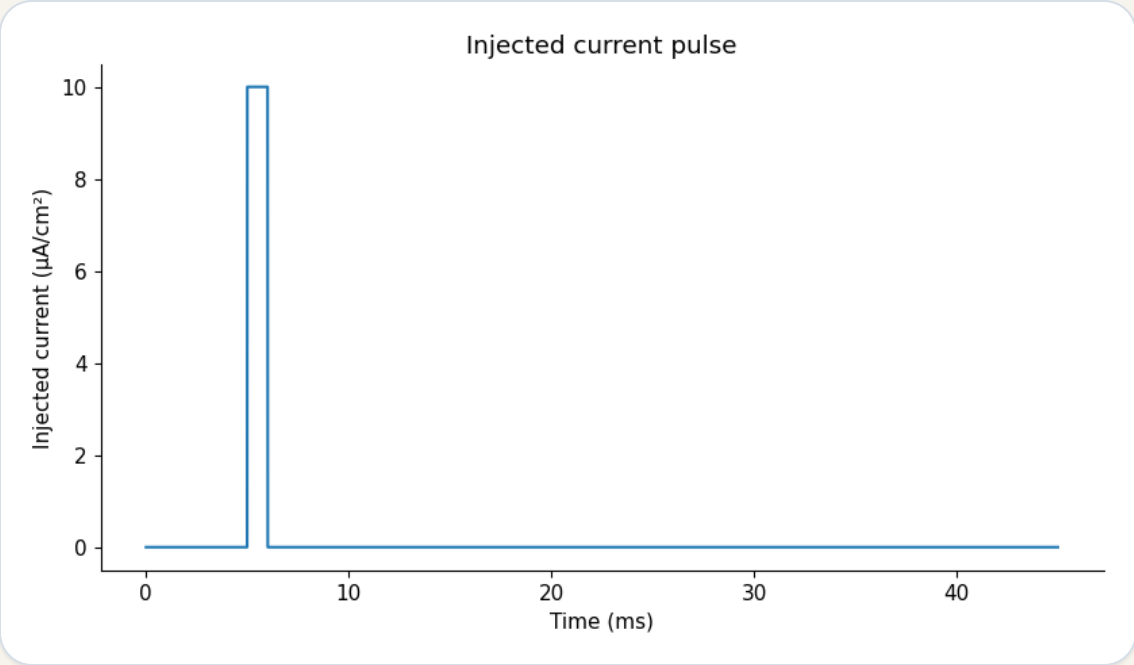

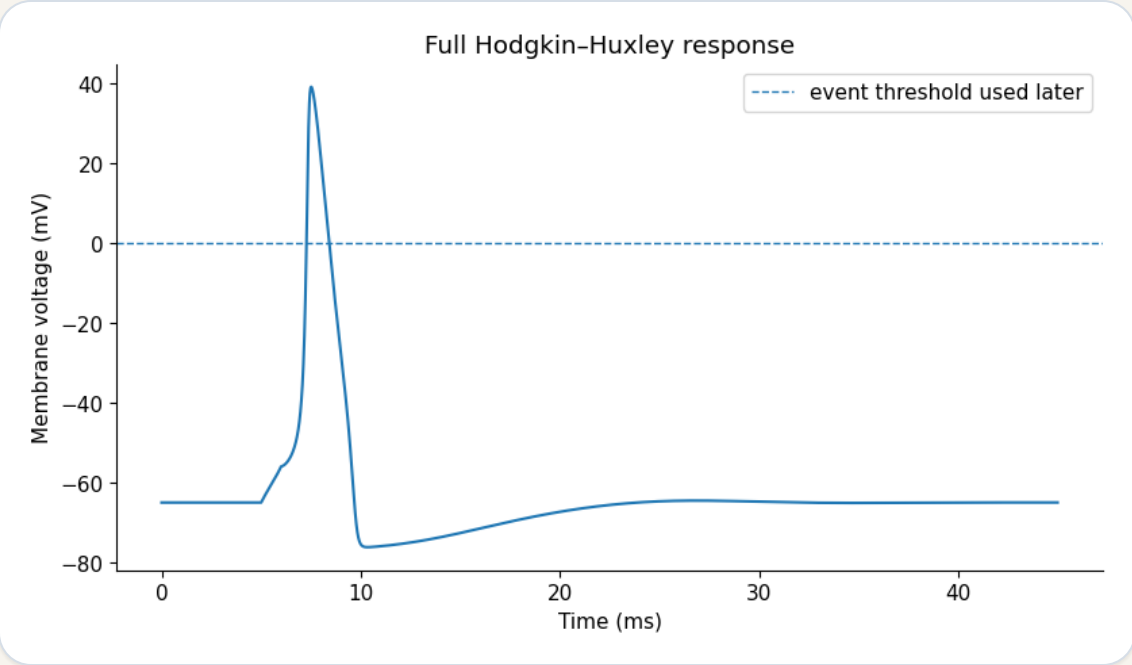

INPUT IS IMPOSED

VOLTAGE IS PREDICTED

The waveform is an output of the mechanism — not a scripted template.

Now look at what caused it

The currents and gates are mechanism variables hidden behind the voltage output.

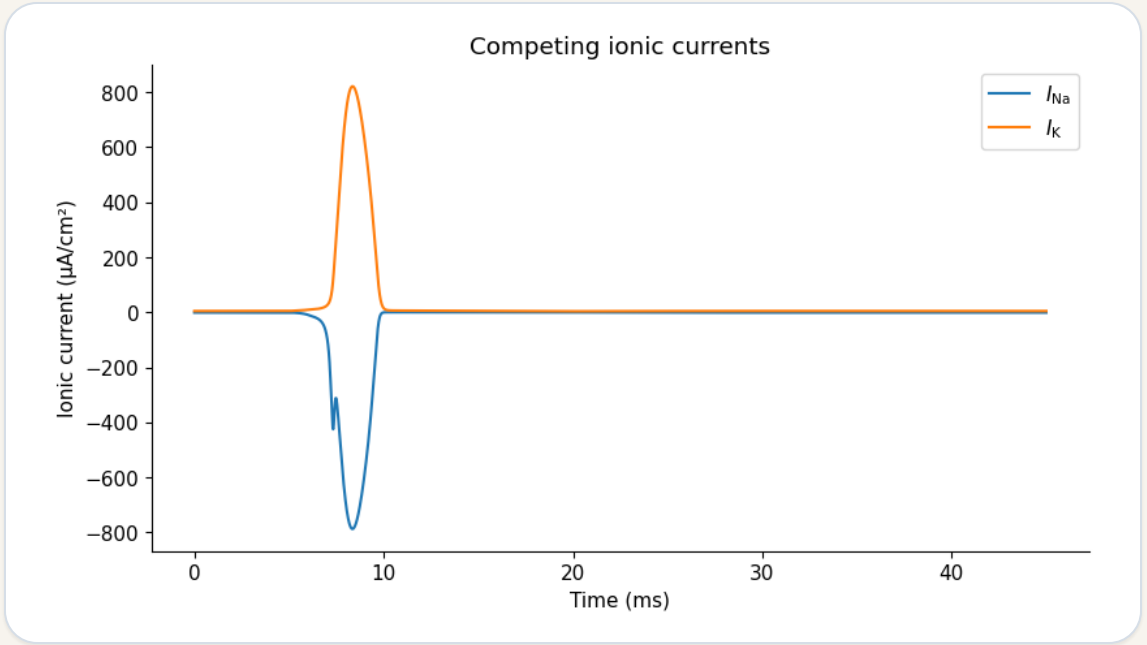

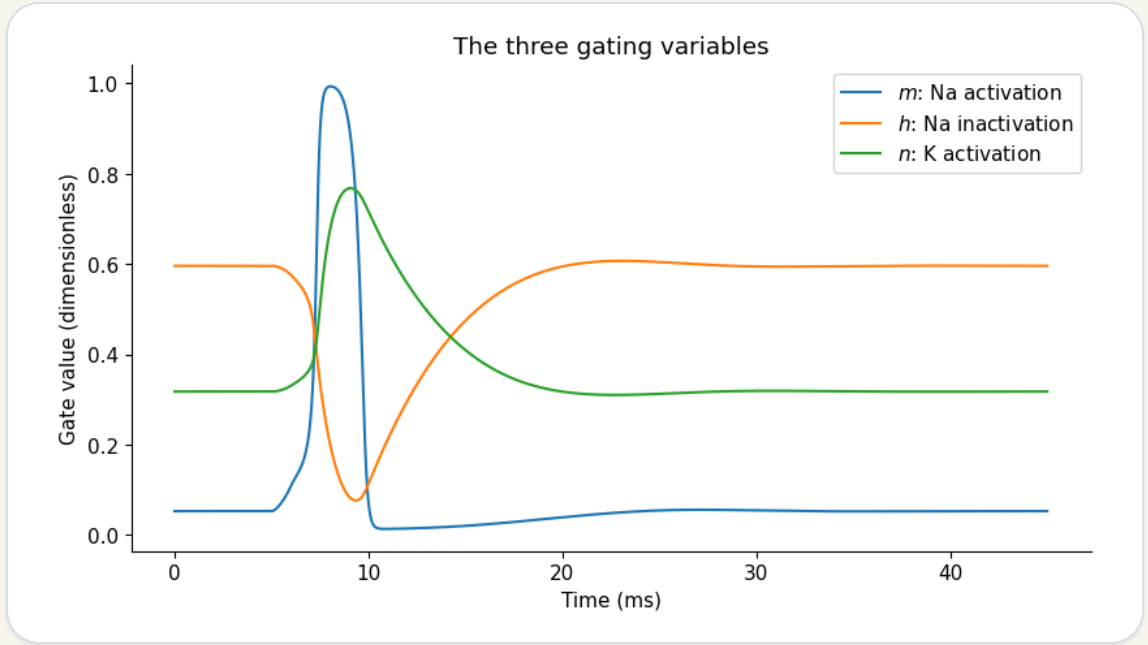

Negative I_{na} = inward current under this sign convention.

I_k is delayed and outward.

m, h and n are hidden state variables — not the voltage output.

CHECKPOINT

Can you explain the spike in words?

Discuss with someone nearby before we simplify the model.

1 Positive feedback

Why does sodium current amplify depolarization?

2 Timescale

Why must recovery be slower than sodium activation?

3 Mechanism vs output

Which variables describe hidden mechanism, and which one is visible?

We will check that every reduction keeps this basic story.

Now simplify — but give a reason each time

STEP	REASON	KEEP	DROP	CHECK
1	**Timescale**	fast Na^+ activation	finite m lag	Does it still spike?
2	**Correlation**	one recovery coordinate	independent h dynamics	Does excitability survive?
3	**Readout**	whether and when	waveform + channels	Do event times suffice?

A reduction is an argument plus a validation test.

m is fast enough to follow voltage almost instantly

$$\frac{dm}{dt} = \frac{m_\infty(V) - m}{\tau_m(V)}$$

If m catches its moving target much faster than h and n…

We replace one ODE with an instantaneous function of voltage.

$$m(t) \approx m_\infty(V(t))$$

RETAIN

Fast sodium positive feedback

DISCARD

Finite response time of m

At rest, m is much faster than h or n

Representative standard HH time constants at V = −65 mV.

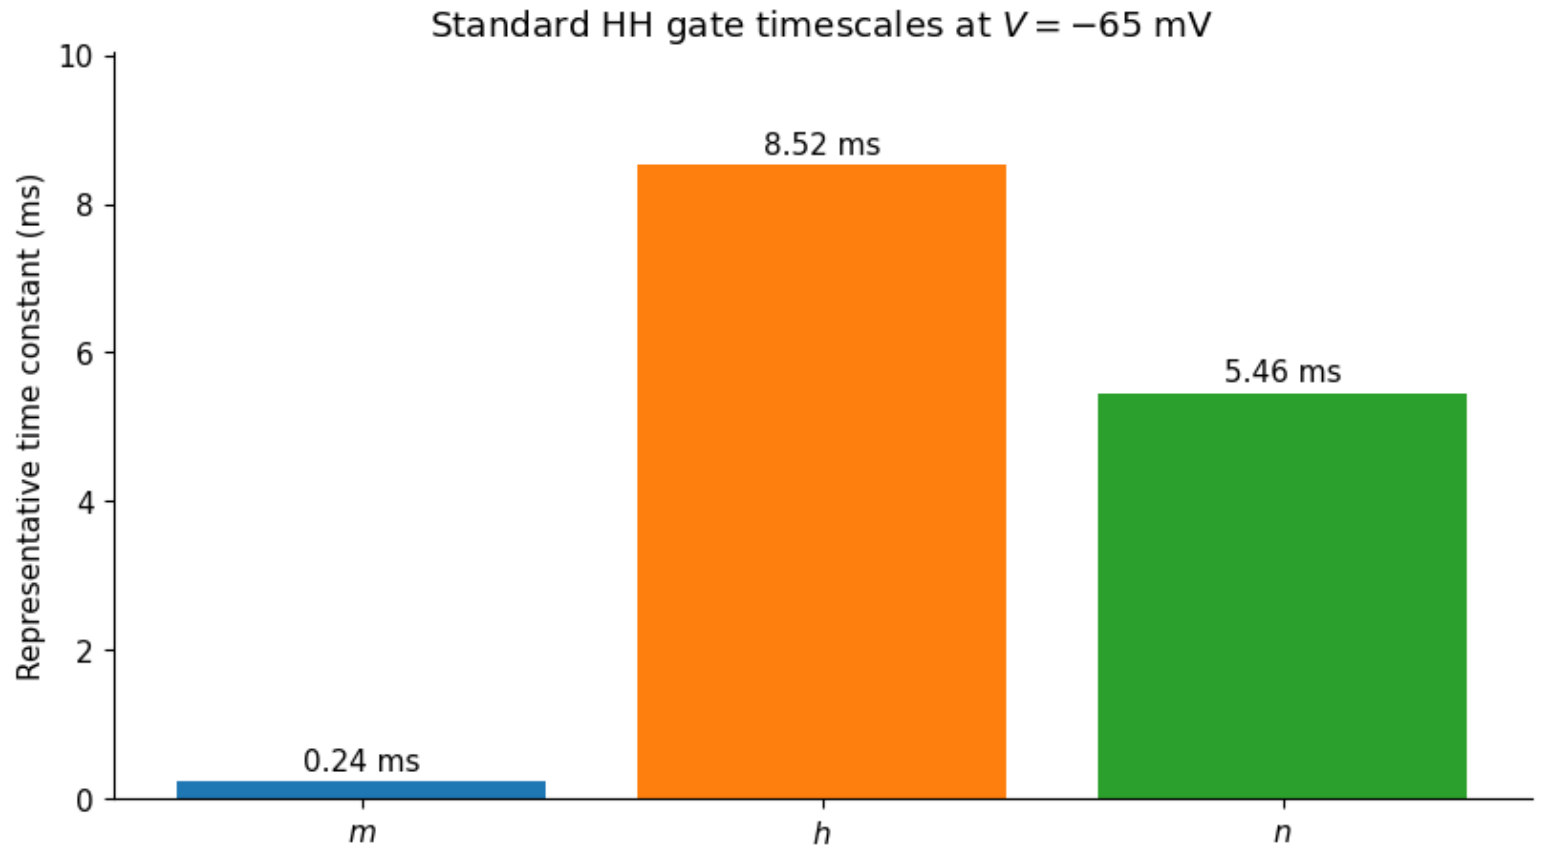

Evidence for replacing the m ODE by m∞(V)

One equation, three very different speeds

For a fixed target, the gate has an exponential relaxation solution.

$$x(t) = x_\infty + (x(0) - x_\infty)e^{-t/\tau_x}$$

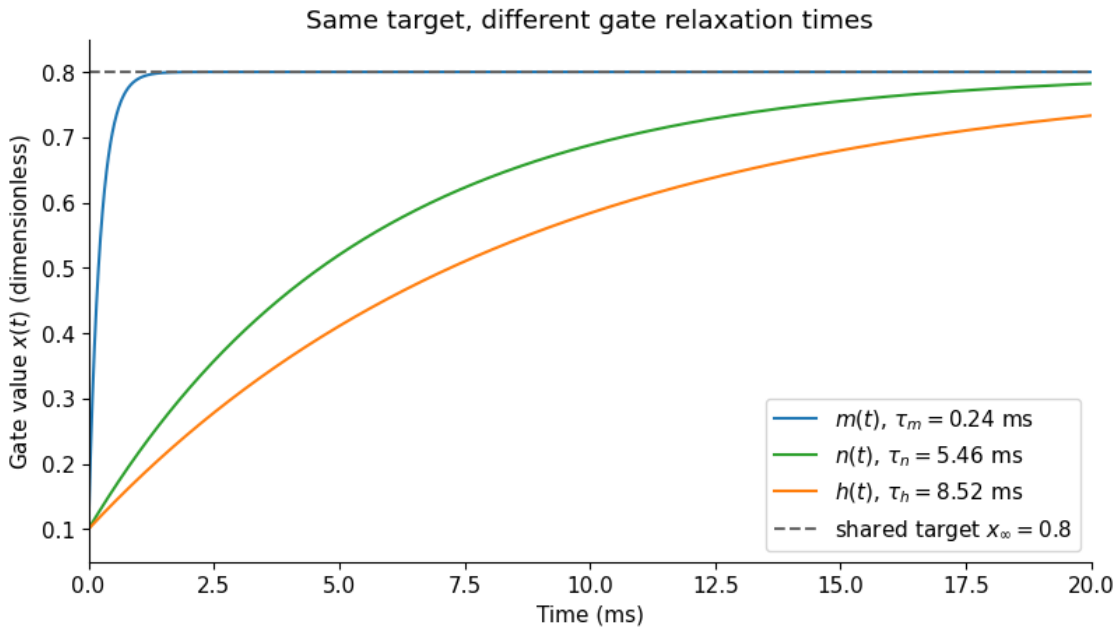

m reaches the target fast enough to use m(t) ≈ m∞(V(t))

The shortcut still gives us a spike

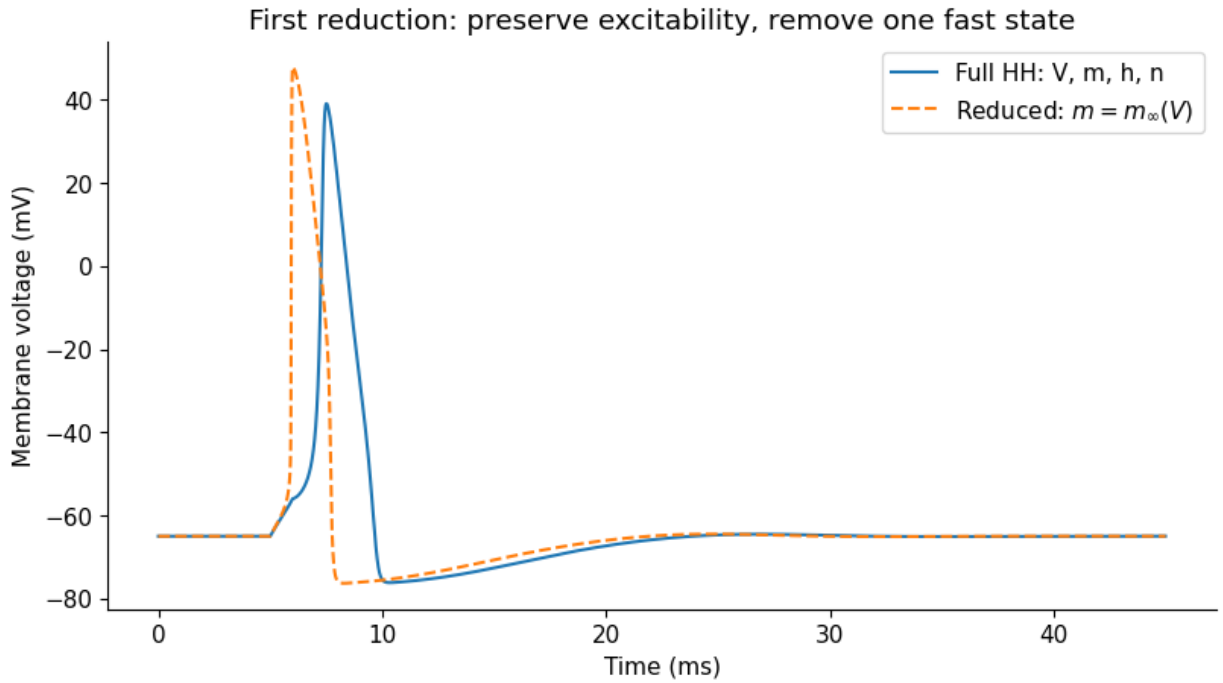

Validation: replacing the m ODE by m∞(V) keeps the qualitative spike.

RETAIN

Fast–slow spiking mechanism, qualitatively

DISCARD

Exact spike timing and waveform agreement

h and n carry similar recovery information

During a spike

n rises → K^+ current grows
h falls → Na^+ current weakens
both support recovery

$$h \approx a - bn$$

This is a trajectory-specific fit from full HH data — not a universal biochemical law.

RETAIN Fast excitation + one slow recovery coordinate

DISCARD Independent h dynamics and exact waveform

Fit h from n using the full HH trajectory

Trajectory-specific approximation: useful here, not a universal biochemical identity.

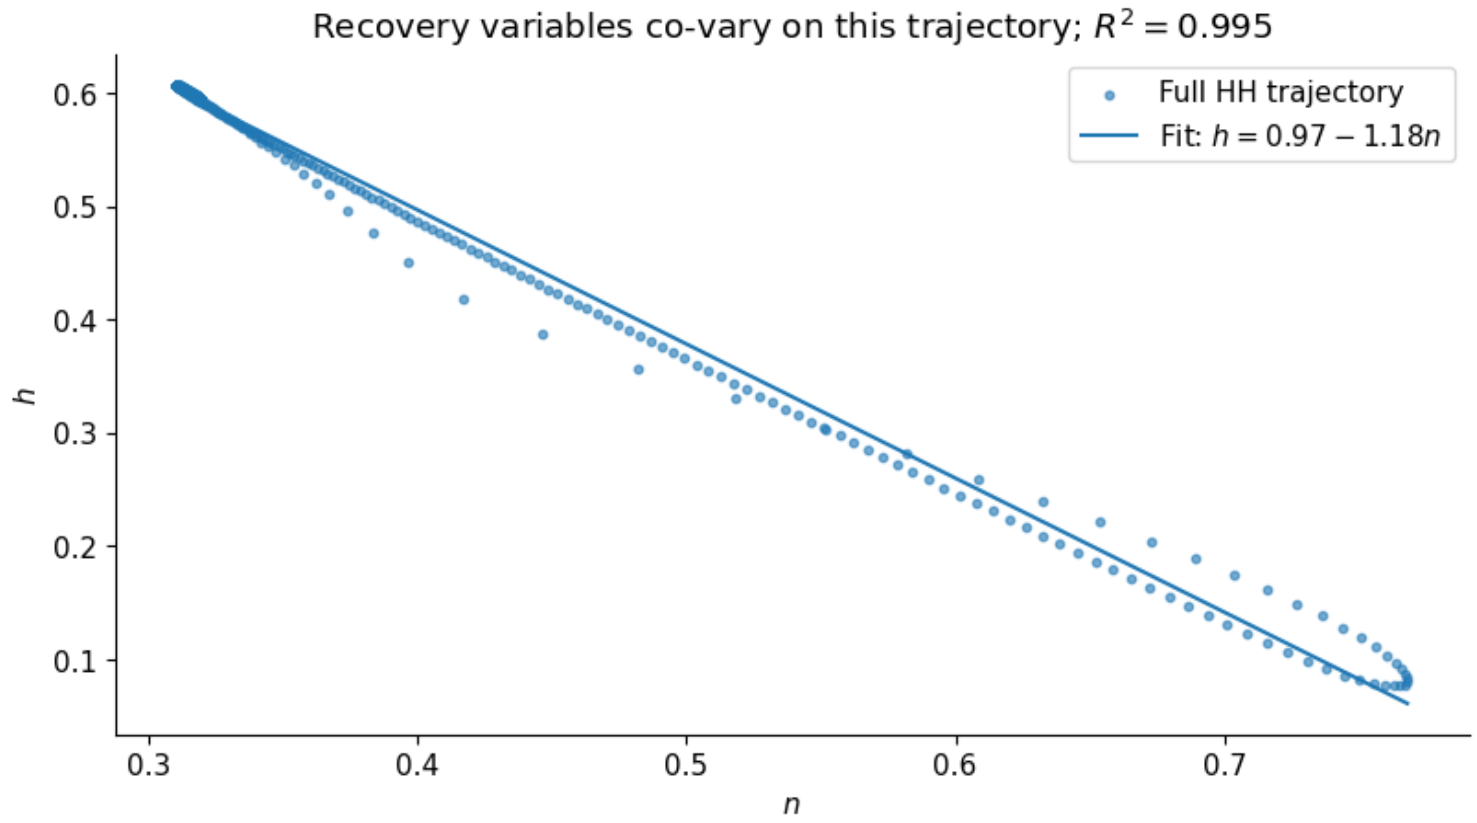

Two state variables are enough to preserve excitability

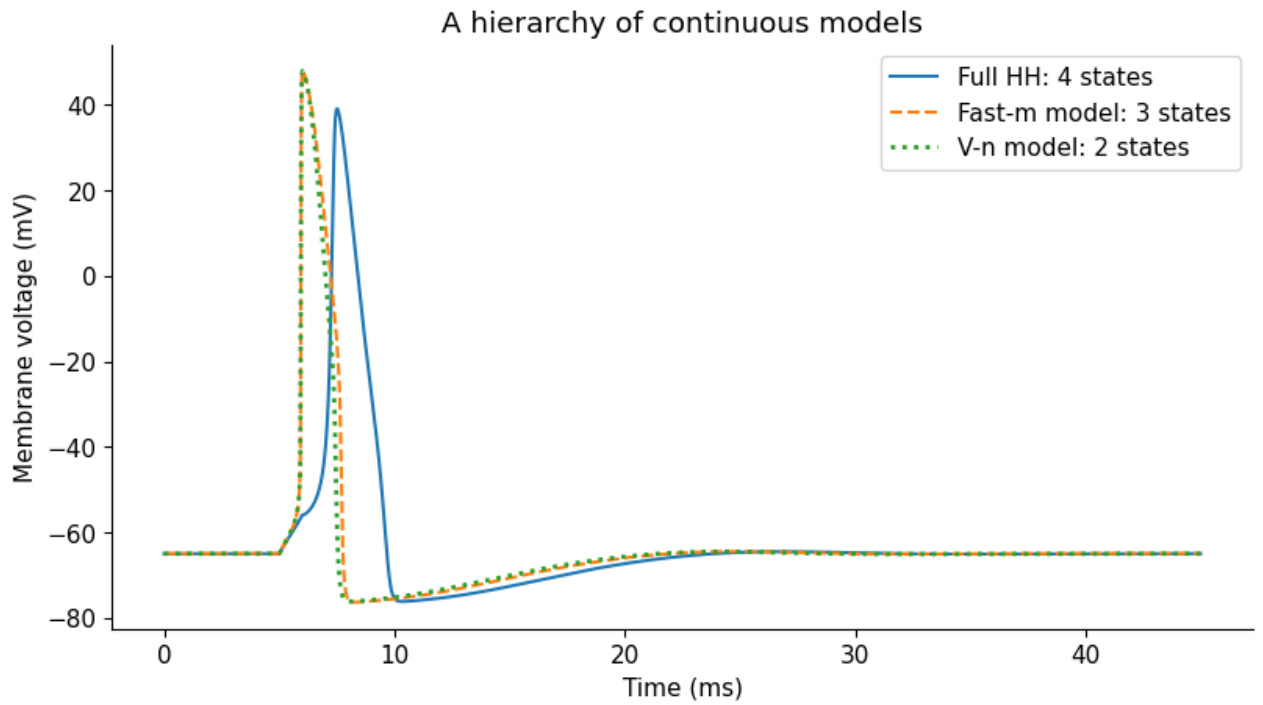

RETAIN

Threshold-like response and continuous voltage

DISCARD

Separate h dynamics; exact shape and timing

For a network, the spike time may be enough

Downstream synapses often only need to know whether and when a spike occurred.

$$V(t^-) < V_{\text{event}}, \qquad V(t^+) \geq V_{\text{event}}$$

an upward threshold crossing

$$V(t) \longrightarrow \{t_1, t_2, \ldots\}$$

Here: threshold reads out an HH event. Later: the LIF threshold generates a new event.

The waveform has become a timestamp

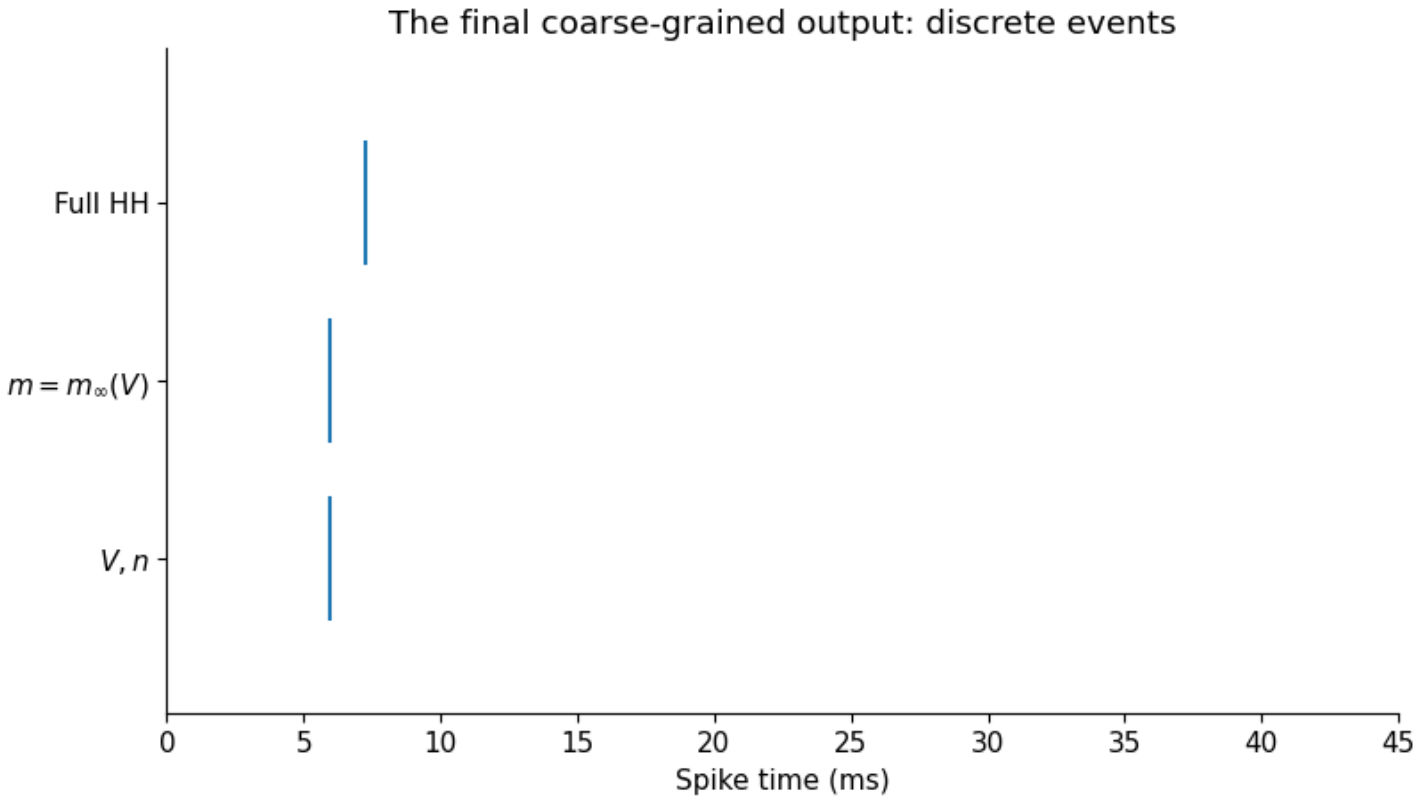

Useful for network timing

Not enough for channel questions

Use the simplest model that answers your question

FULL HH	V, m, h, n	ionic mechanism + waveform
FAST-M HH	V, h, n	excitation + two recovery processes
TWO-STATE HH	V, n	fast excitation + slow recovery
EVENTS	t_1, t_2, …	whether and when spikes occur

Model choice depends on the question — reduced means matched to a different question.

Now push the shortcut until it breaks

Same input; different models; compare spike count and timing.

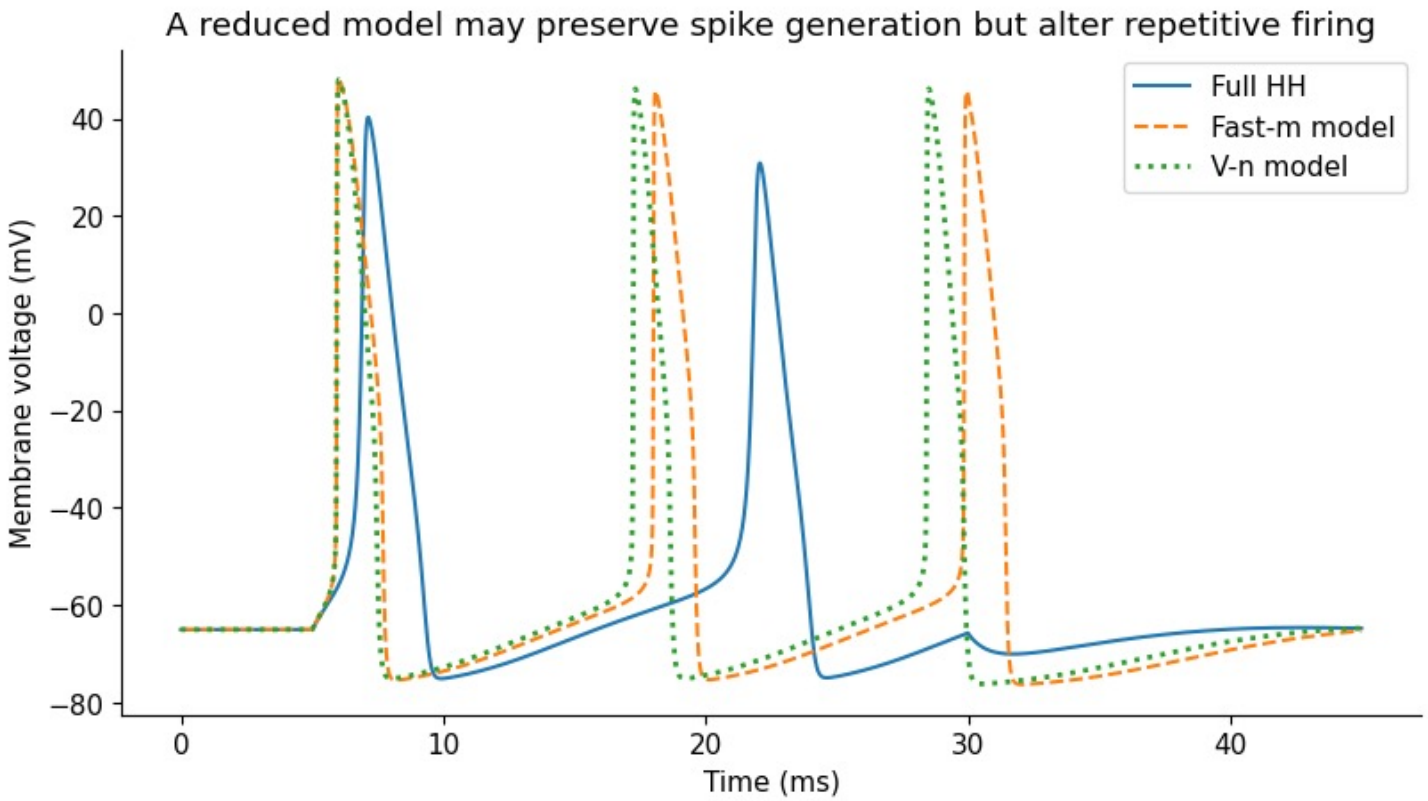

Same qualitative mechanism — different spike timing or spike count.

A good reduction

keeps the answer you care about

and forgets the rest — as long as you know when the shortcut fails.

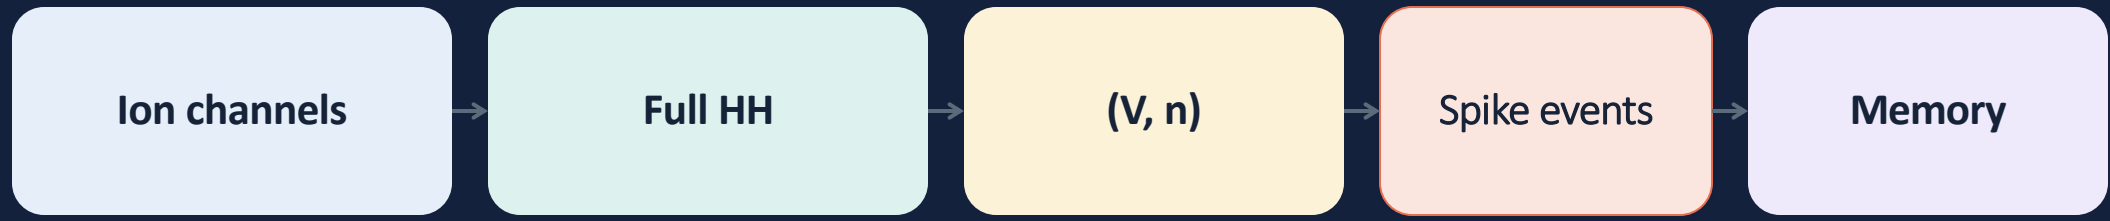

NEXT: events become synaptic input

PART II — Events become network dynamics

Part I explained one spike. Now events must interact, propagate and change the network.

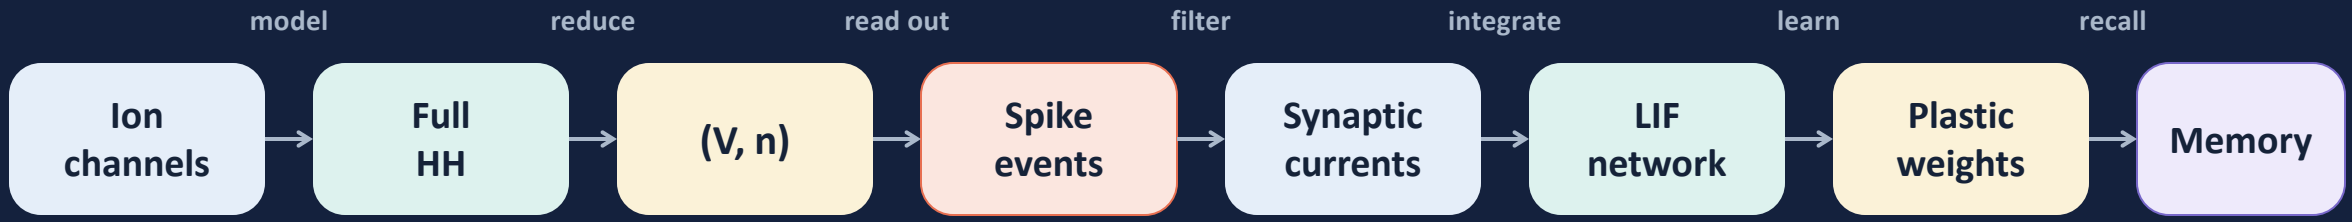

A new scale needs new state variables.

We compressed an action potential into an event time. Next we reconstruct only the information a network needs.

Spike times alone are not enough

A downstream neuron needs a brief input signal, not just a timestamp.

$$t_j^{(k)} \longrightarrow r_j(t) \longrightarrow I_i^{\mathrm{syn}}(t) \longrightarrow V_i(t) \longrightarrow t_i^{(k)}$$

$t_j^{(k)}$	$r_j(t)$	$w_{ij}\, r_j(t)$	$V_i(t)$	$t_i^{(k)}$
presynaptic event time	synaptic trace left by that event	weighted postsynaptic input	effective integration variable	new output event time

The missing interface is a continuous synaptic trace that lasts for a few milliseconds.

WHY
timestamps cannot be integrated directly

KEEP
who fired and when

DROP
the presynaptic waveform

NEXT
build a short-lived trace

A discrete event creates a continuous synaptic effect

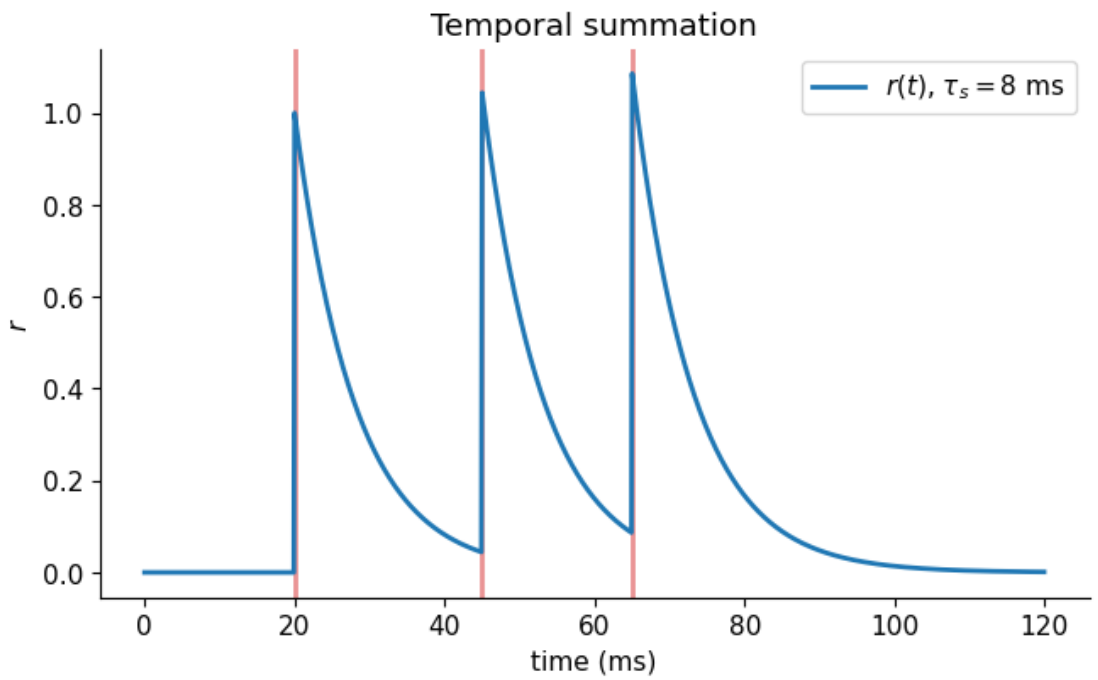

$$\tau_s \frac{dr_j}{dt} = -r_j$$

at a spike: $r_j \leftarrow r_j + 1$

WHAT TO NOTICE

red lines are event times
blue $r_j(t)$ persists and decays
τ_s sets how long the effect lasts

WHY

an event needs duration

KEEP

event time and recent history

DROP

channel-scale waveform

NEXT

sum traces into current

Synapses turn events into summed input

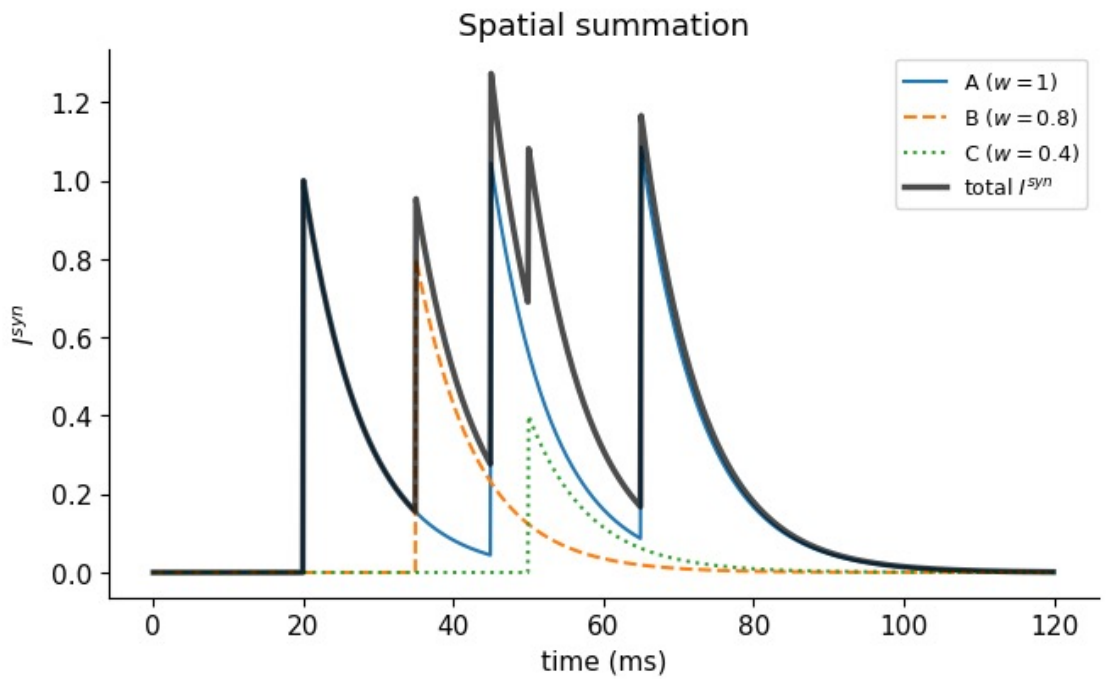

$$I_i^{\mathrm{syn}}(t) = \sum_j w_{ij} r_j(t)$$

READ THE GRAPH

x-axis: time; y-axis: synaptic input
coloured curves: weighted sources A, B and C
black curve: total input to neuron i
close spikes add in time; sources add in space

WHY

many events converge on one cell

KEEP

timing × effective weight

DROP

individual channel currents

NEXT

integrate total input

LIF is the simplest event processor

We have changed scale: this is an effective model for many neurons, not a strict derivation from HH.

$$\tau_m \frac{dV_i}{dt} = -(V_i - E_L) + R_m I_i^{\mathrm{syn}}(t)$$

$$V_i \geq V_{\mathrm{th}} \Rightarrow \text{emit spike and reset}$$

In Part I, voltage explained the spike waveform.
Here, V_i only tracks how close the unit is to firing again.

V_i	effective distance to threshold
τ_m	integration window
E^L	leak / resting level
V_{th}	generative event threshold
reset	coarse post-spike recovery

WHY
simulate many event-processing units

KEEP
integration, threshold and reset

DROP
HH waveform and channel mechanism

NEXT
close the recurrent loop

Events in, events out: the recurrent loop

Every output event can become another neuron's input event.

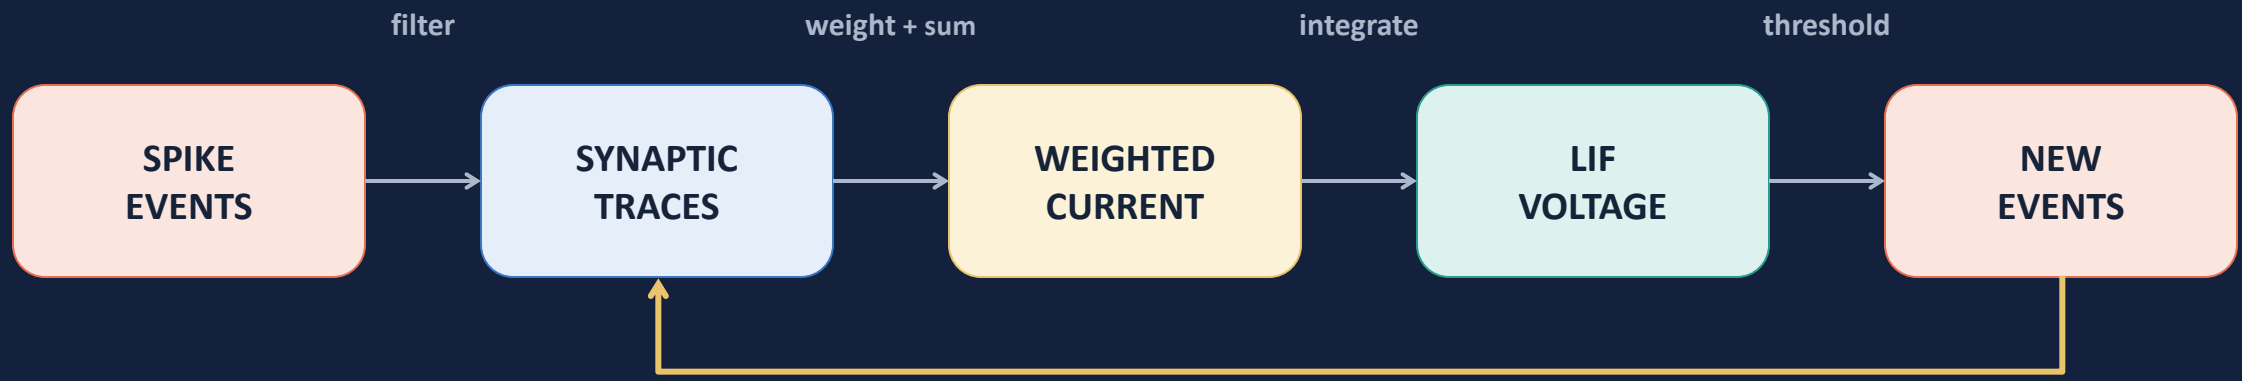

recurrent connections feed events back into the network

events ↔ continuous dynamics

What’s in your mind?

The molecular mechanism of memory and learning — plasticity

Within a neuron, the signal travels along the axon. Between neurons, it crosses a synapse.

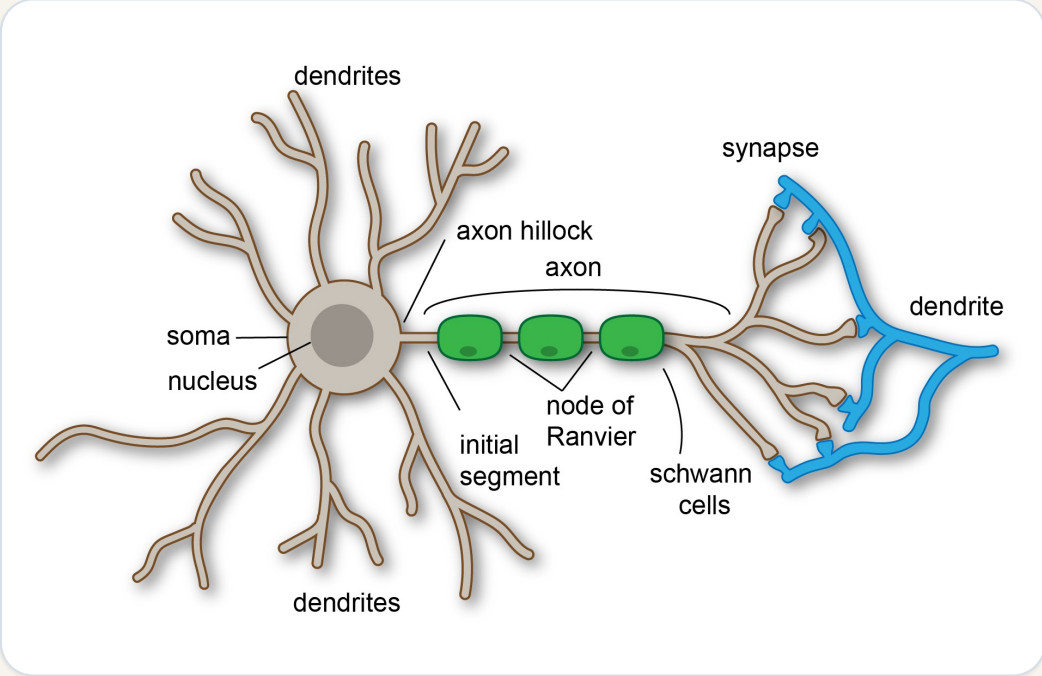

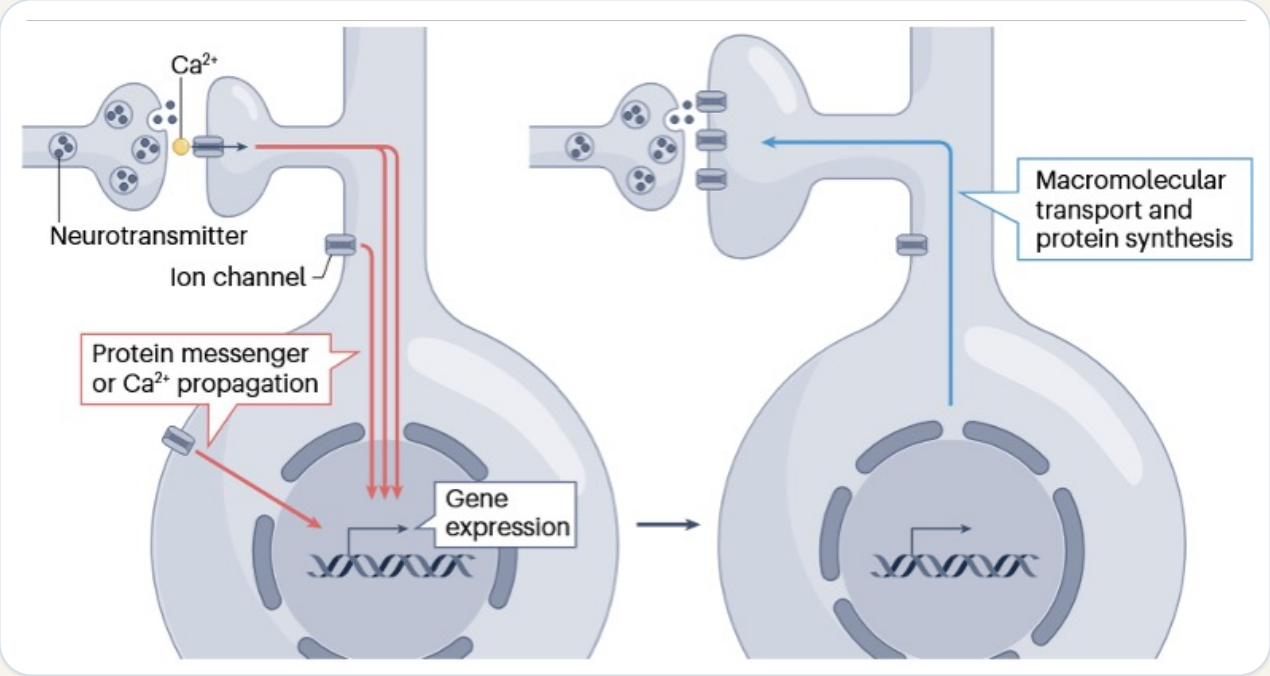

A memory is not a picture stored in a cell. It is a change in how cells talk to each other.

MA H. Nature Reviews Neuroscience, 2023.

Long term potentiation — first discovered in 1966

Induced by trains of stimulation on the **presynaptic neuron**: 10 s, 15 Hz.

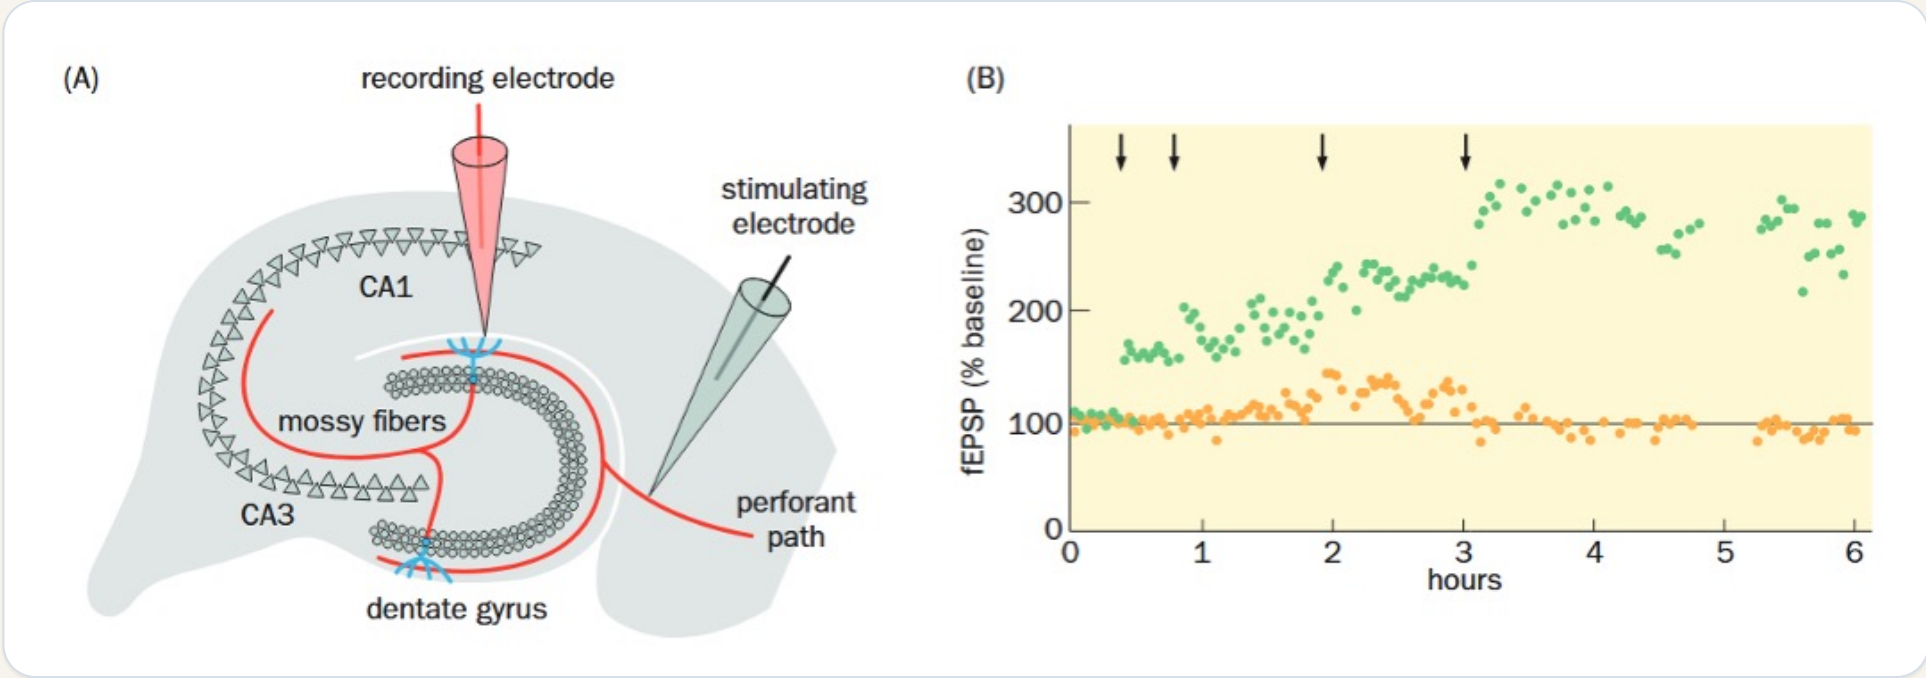

Enhanced synaptic strength and the enhancement can last for hours to days

LUO L. Principles of neurobiology[M]. 2020.

Learning begins when spike timing changes weights

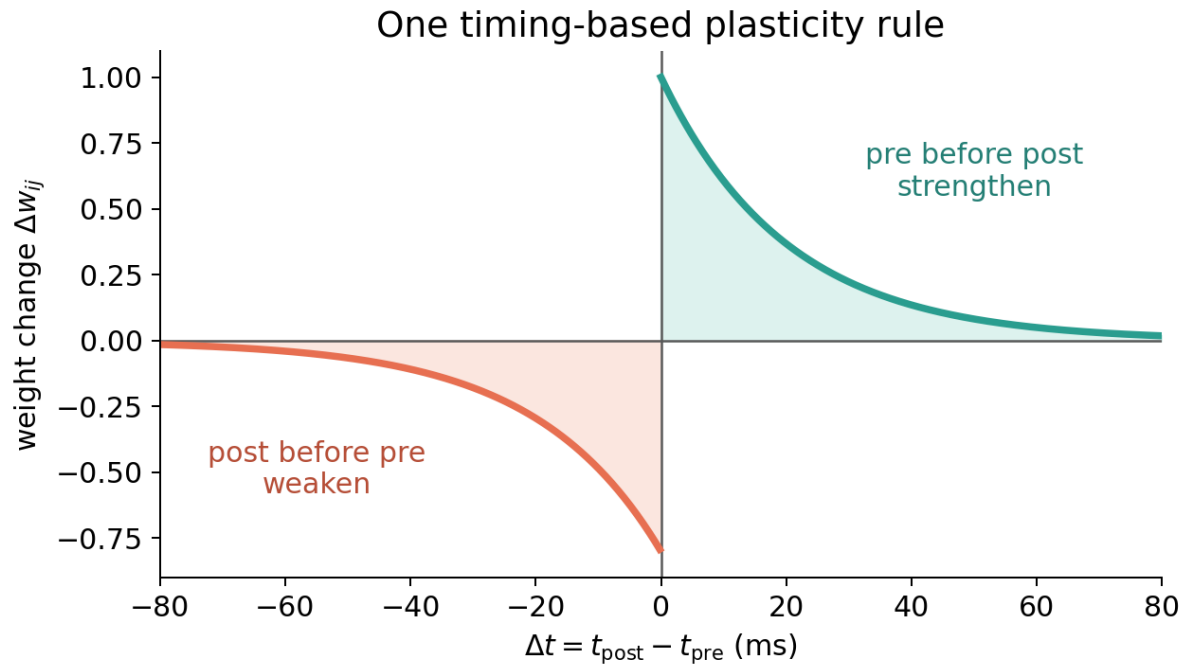

$$\Delta t = t_{\text{post}} - t_{\text{pre}}$$

$$\Delta w_{ij} = \begin{cases} +A_+ e^{-\Delta t/\tau_+}, & \text{pre before post} \\ -A_- e^{\Delta t/\tau_-}, & \text{post before pre} \end{cases}$$

STDP is one timing-based rule that can produce potentiation or depression.

WHY

past events should change future propagation

KEEP

relative pre/post timing

DROP

molecular LTP mechanism

NEXT

store association in W

Repeated co-activity builds a cell assembly

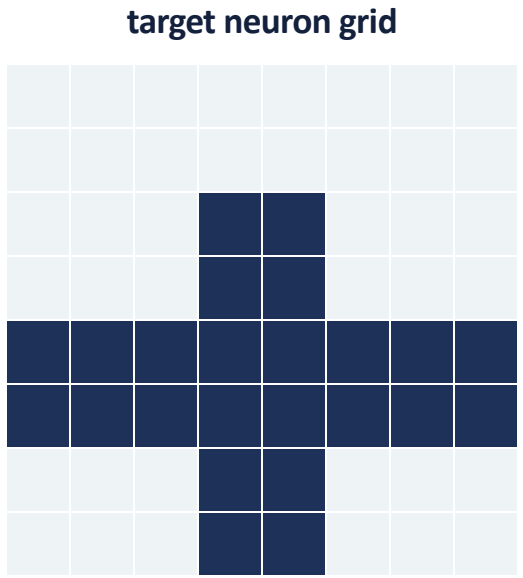

1 external input activates the target group

↓

2 target neurons fire together

↓

3 STDP strengthens recurrent connections

↓

4 the strongly connected group becomes an assembly

$$w_{ij} \uparrow, \qquad i, j \in \text{same pattern}$$

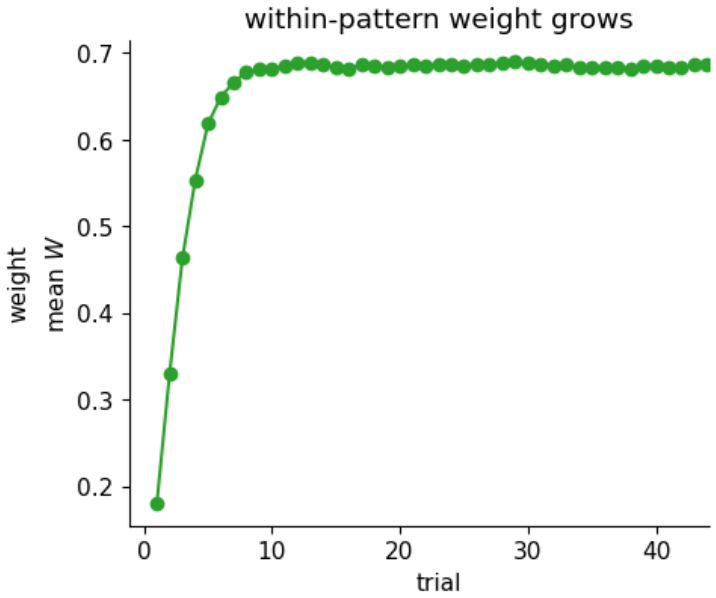

What to notice: the within-pattern weight rises quickly, then saturates.

WHY

bind co-active neurons together

KEEP

recurrent association

DROP

trial-by-trial spike details

NEXT

read memory from W

The memory trace is visible in W

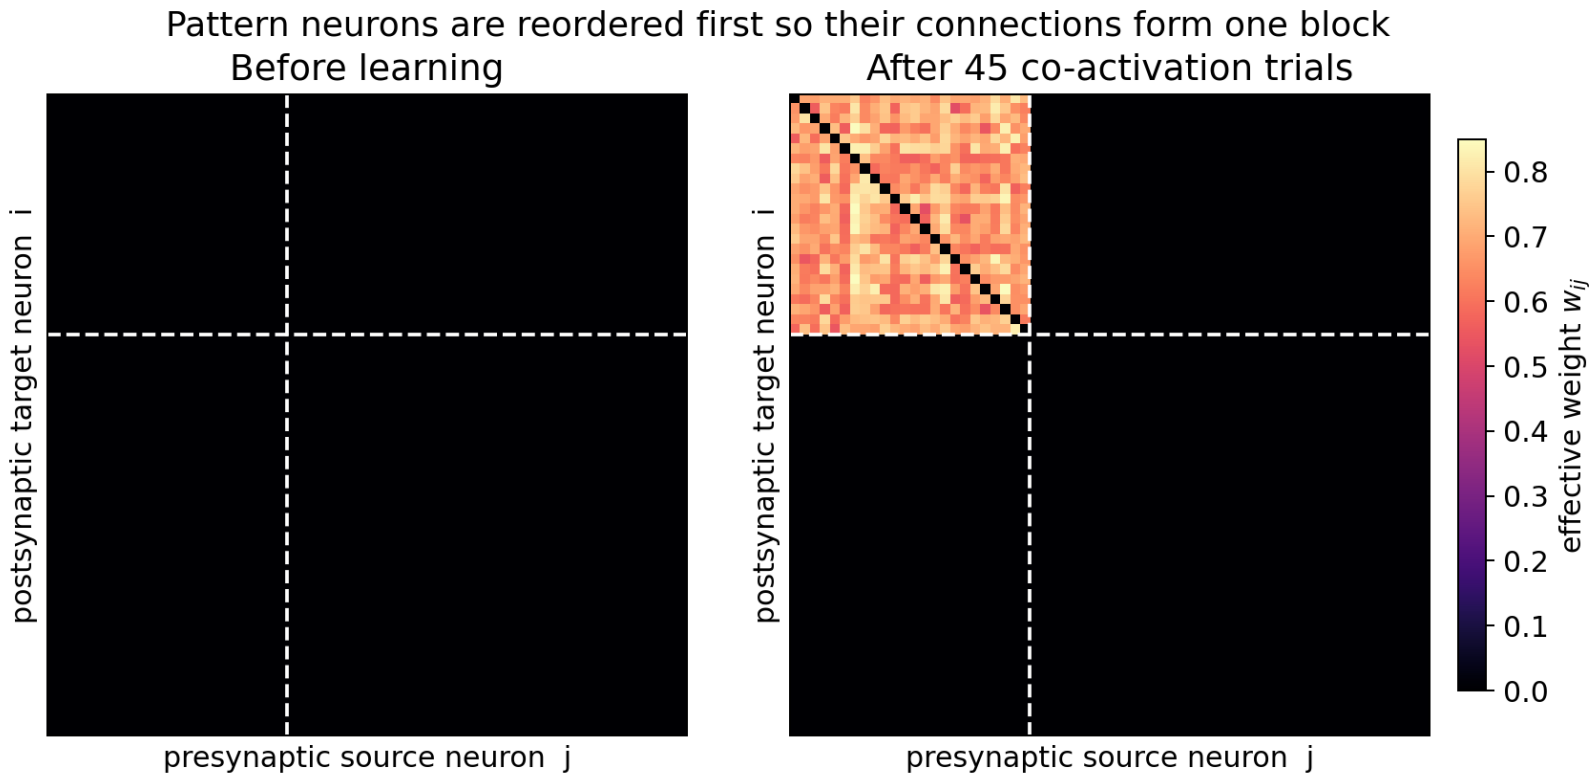

HOW TO READ W

rows: postsynaptic target i
columns: presynaptic source j
brighter means stronger w_{ij}
the bright block links assembly neurons
the neurons did not change; connectivity did

The matrix is reordered only for display: pattern neurons are grouped first.

WHY

make the learned association explicit

KEEP

who strongly excites whom

DROP

the training trajectory

NEXT

test recall from a cue

A partial cue recalls the full pattern

Strengthened recurrent weights recruit the missing members of the assembly.

FULL LEARNED PATTERN

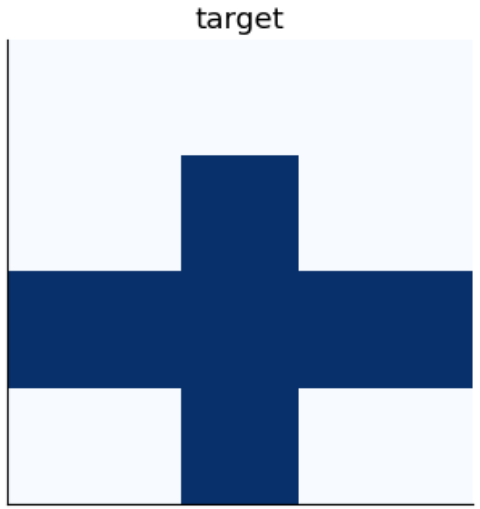

show half

PARTIAL CUE

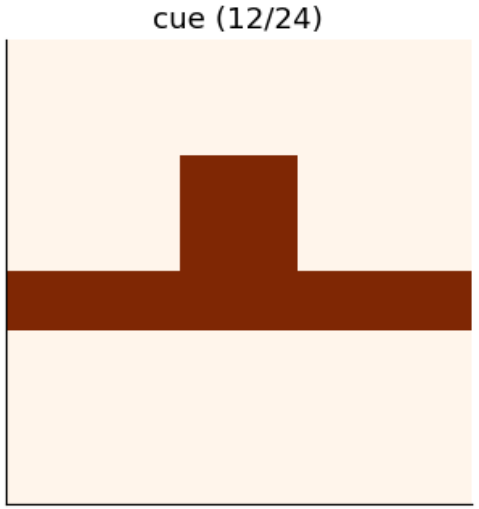

recurrent dynamics

RECALLED ACTIVITY

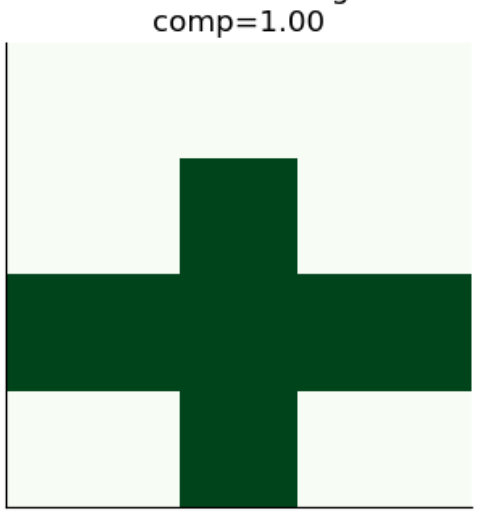

Associative / content-addressable memory: the cue selects the stored assembly.

WHY

complete a partial pattern

KEEP

assembly activity + W

DROP

training order

NEXT

stress-test with noise

Noise tests whether the memory is selective

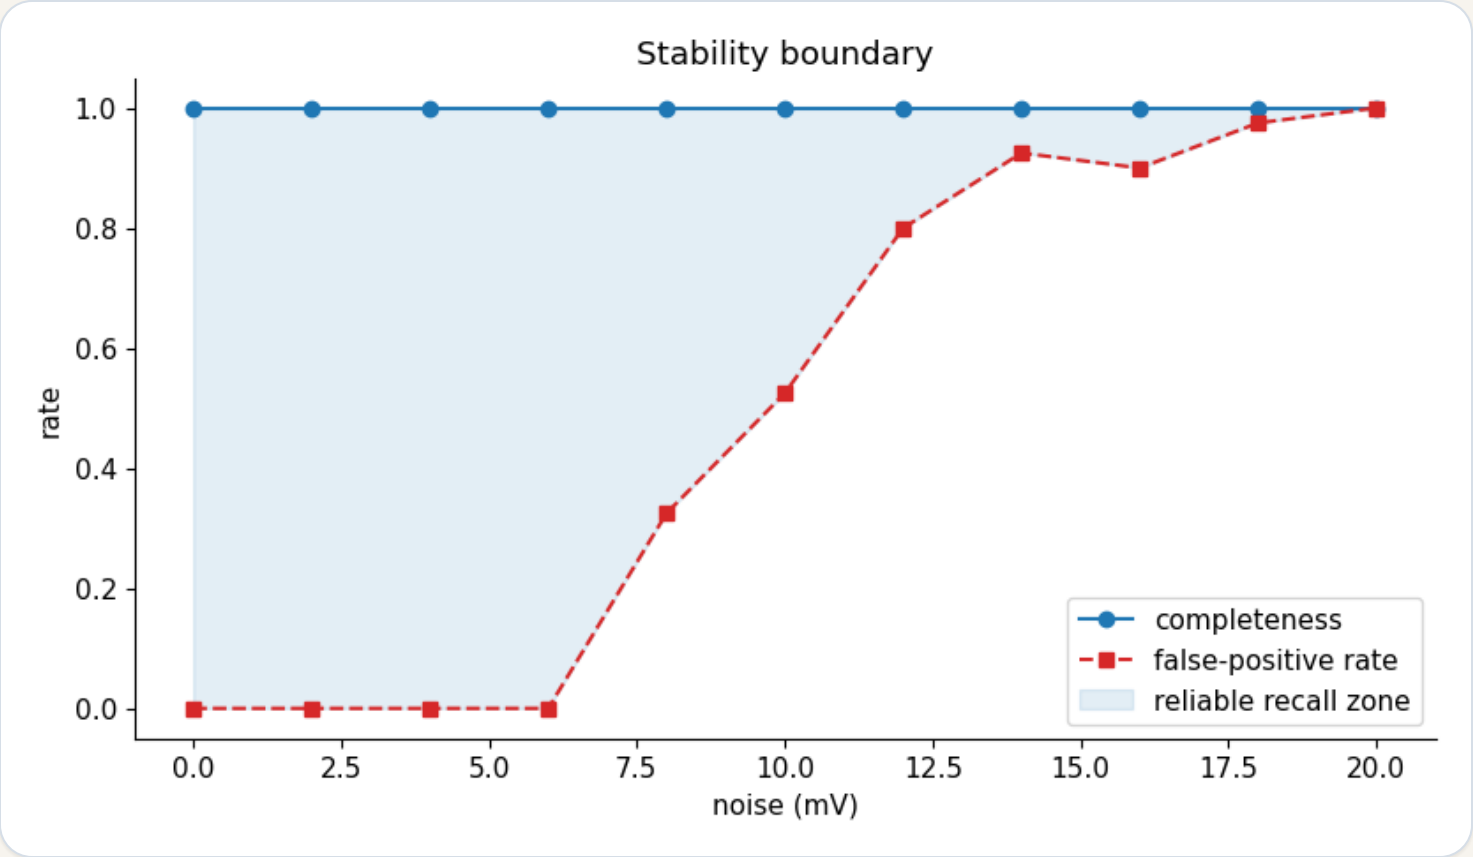

Noise perturbs fast neuronal dynamics during recall.
It is not added to the image or to learned weights.

WHAT TO NOTICE

blue: target neurons recovered
red: non-target neurons firing
low noise: recall is selective
high noise: the target remains active, but irrelevant neurons join

This sweep probes one slice of the attractor basin — not the full basin.

At high noise the failure is selectivity collapse, not necessarily loss of the target memory.

From ion channels to memory

One chain, several representations, one question passed to the next scale.

$$(V, m, h, n) \to (V, n) \to t_i^{(k)} \to r_i(t) \to I_j^{\mathrm{syn}}(t)$$
$$\to V_j(t) \to t_j^{(k)} \to \Delta w_{ij} \to W$$

HH explains where spikes come from

Reduction keeps event times for the next scale

Synapses turn events into continuous input

LIF turns input into new events

STDP turns relative timing into Δw_{ij}

Recurrent W stores association and supports recall

What should a model keep?

ION-CHANNEL SCALE

explain spike mechanism

keep V, m, h, n

NETWORK SCALE

explain communication + learning

keep spike times and synaptic weights

MEMORY SCALE

explain recall at population level

keep population activity and recurrent connectivity

Multiscale modelling is not just making models smaller.
It is choosing the right representation for the question at the next scale.