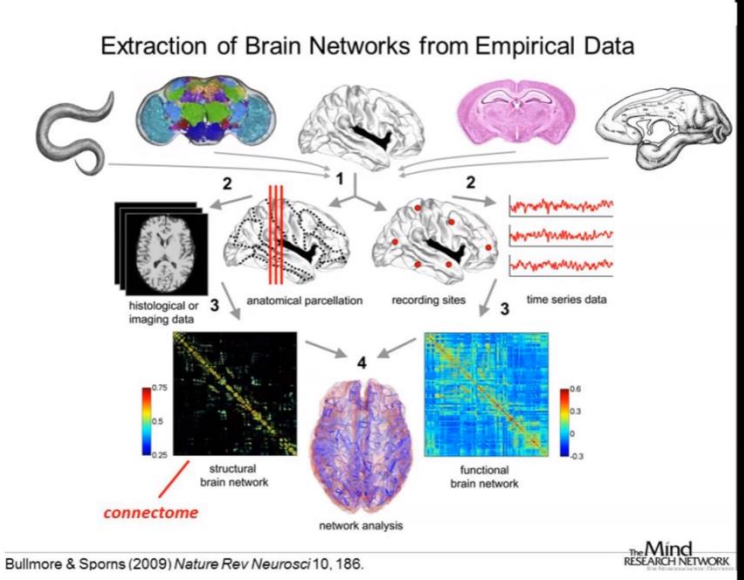

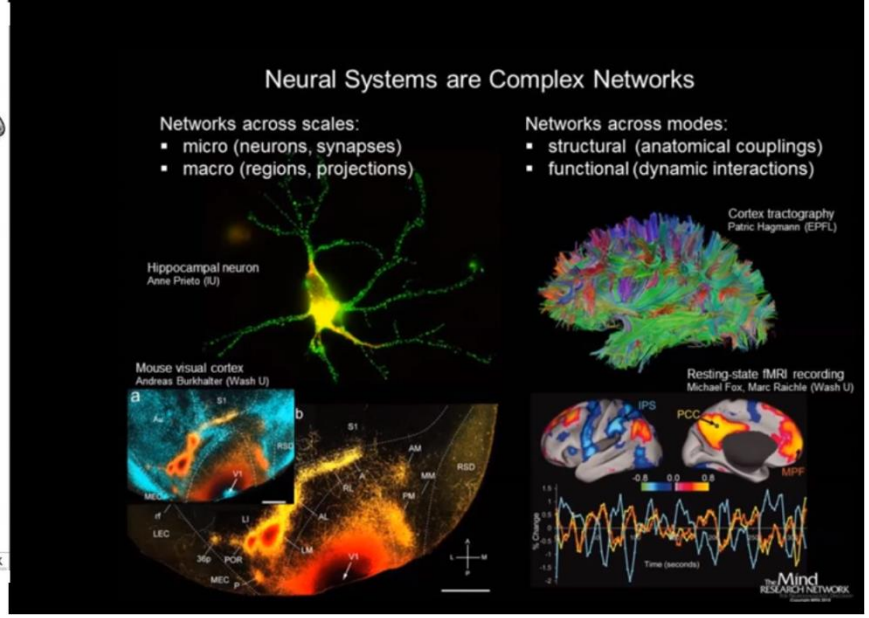

Introductory Lecture on Neuronal Models

Prof. Maria Papadopouli

CS – 588 Analysis and Modeling of Brain Networks

Department of Computer Science
University of Crete

Greek Diaspora Fellowship Program

ΙΣΝ SNF
ΙΔΡΥΜΑ ΣΤΑΥΡΟΣ ΝΙΑΡΧΟΣ
STAVROS NIARCHOS FOUNDATION

Resources

- Biological Neuron Models Tutorial 1 Videos 1 & 2 by Neuromatch Academy Richard Naud University of Ottawa

  https://www.youtube.com/watch?v=rSExvwCVRYg
- MIT 9.49/9.490/6.S076, Instructor: Professor Ila Fiete

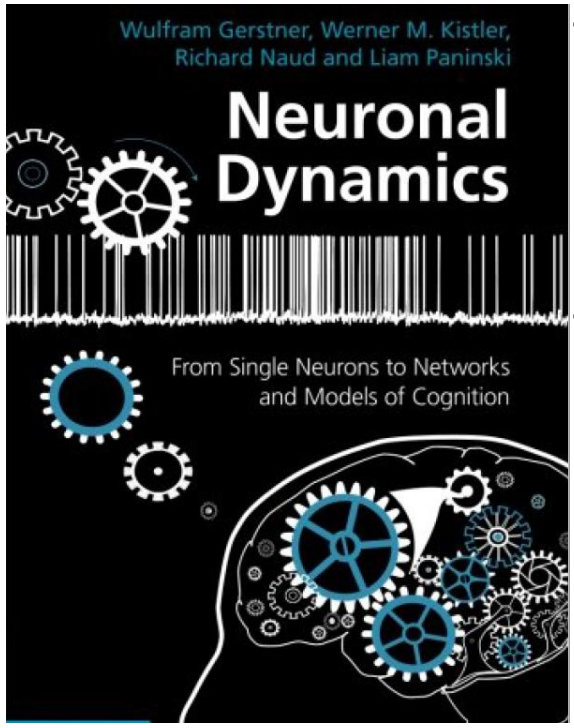

- The Hodgkin-Huxley theory of the action potential, Michael Häusser https://www.nature.com/articles/nn1100_1165.pdf
- Neuronal Dynamics: From Single Neurons to Networks and Models of Cognition By Wulfram Gerstner, Werner M. Kistler, Richard Naud, Liam Paninski (available as PDF online)

Roadmap

- Leaky Integrator and Fire Model
- Spiking vs Firing-rate Models

Modeling the Interaction at Different Scales

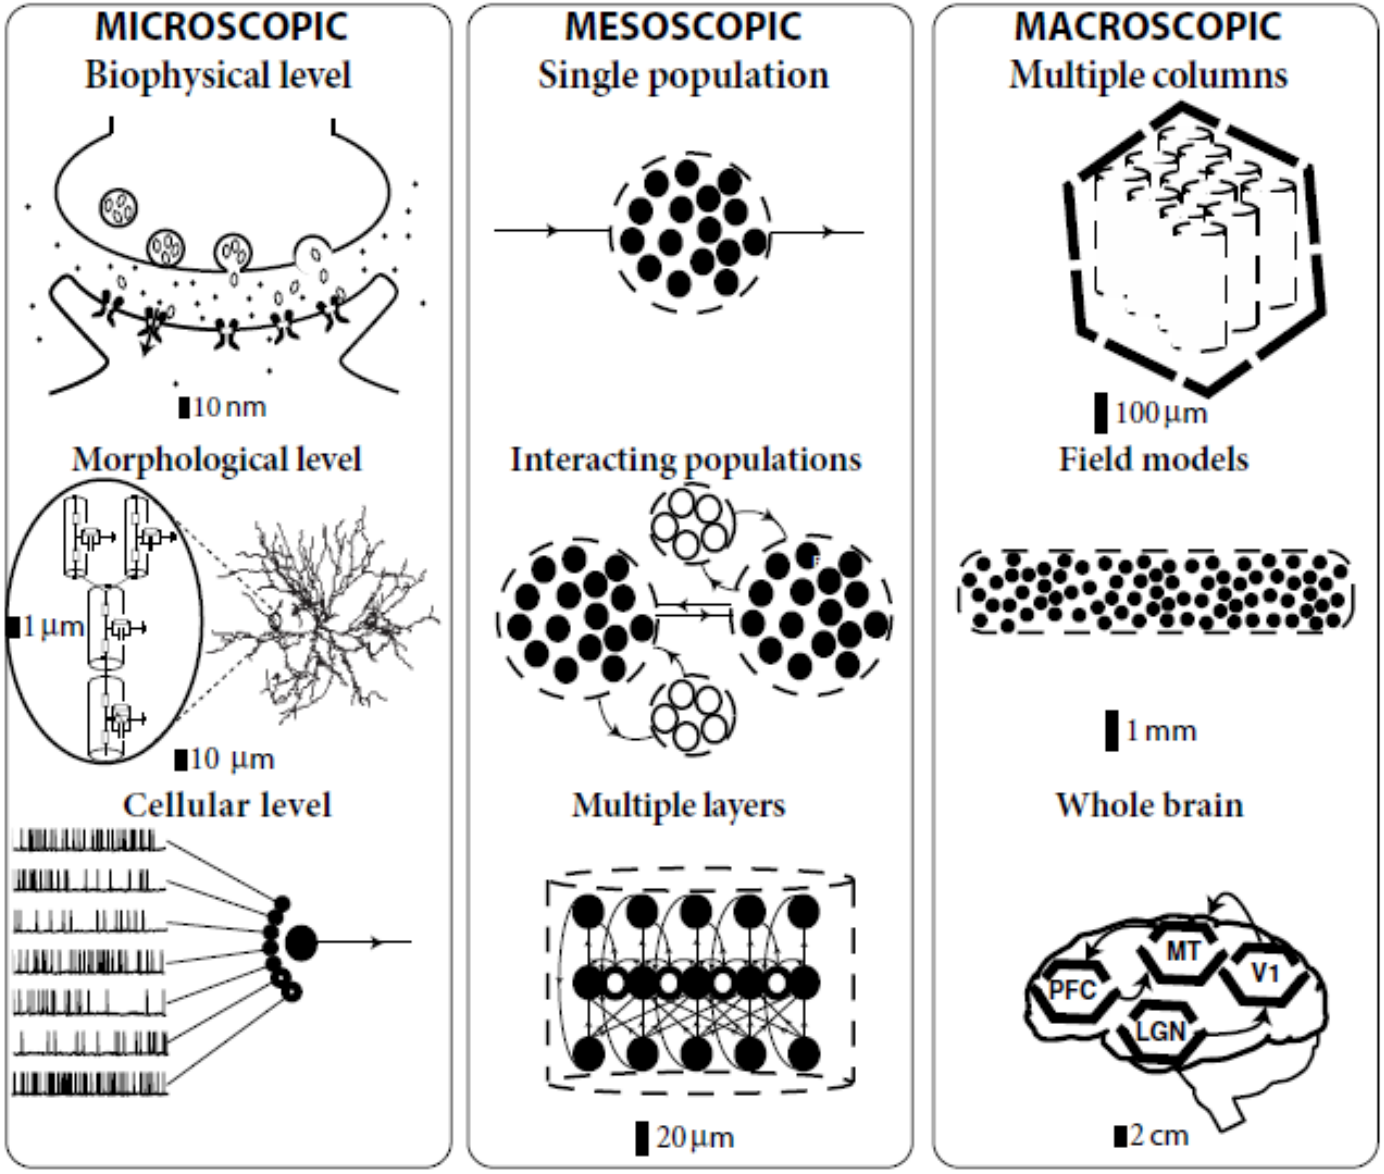

If the experimenter moves the electrode vertically down from the cortical surface to deeper layers, the **location of the receptive field and its preferred orientation** does ***not*** change substantially.

Arrows on the left indicate inter-layer connection probabilities between excitatory neurons

Arrows on the right show intra-layer connection probabilities between excitatory neurons

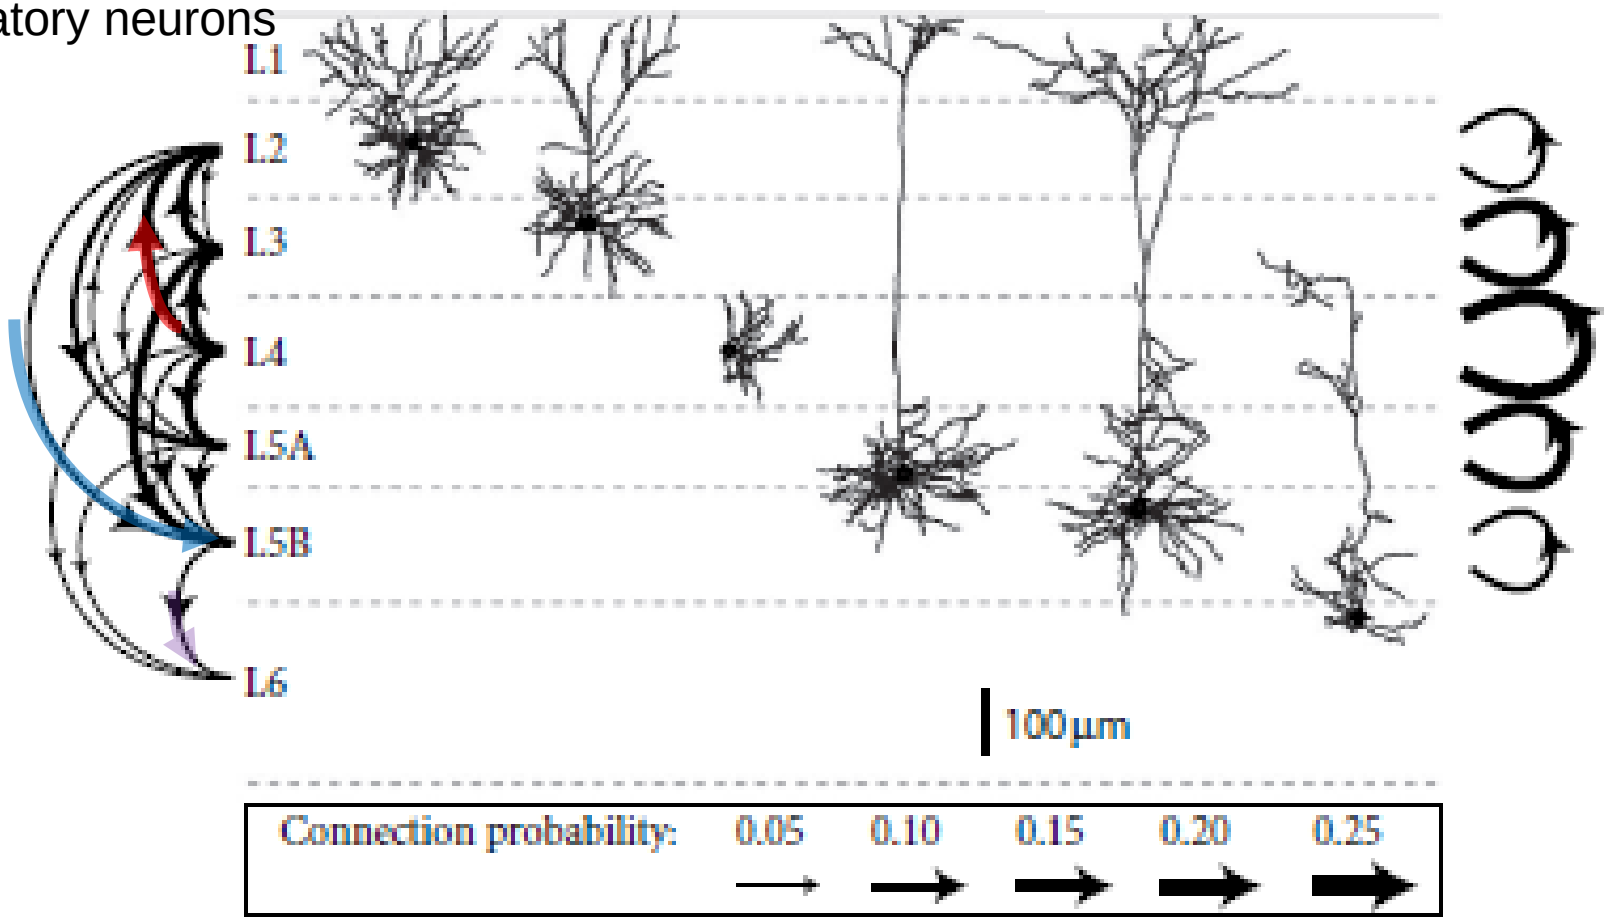

Connectivity patterns inside one column. Examples of shapes of excitatory neurons in different layers.
Data from a barrel column in the somatosensory cortex of the mouse. After Lefort *et al*. (2009)

- The cortex is a rather thin sheet of cells.
- Cortical columns extend vertically across the sheet.
- The connection probability within a column depends on the layer where pre- and postsynaptic neurons are located.
- In addition to this vertical connectivity, neurons make many horizontal connections to neurons in other, cortical columns in the same, but also in other, areas of the brain.

- Within the same brain area, the probability of making a connection is often modeled as distance dependent.

  Note that distance dependence is a rather coarse feature, because the actual connectivity depends also on the function of the pre- and postsynaptic cell.

  In the primary visual area, pyramidal neurons with a preferred orientation for horizontal bars are more likely to make connections to other columns with a similar preferred orientation (Angelucci and Bressloff, 2006).

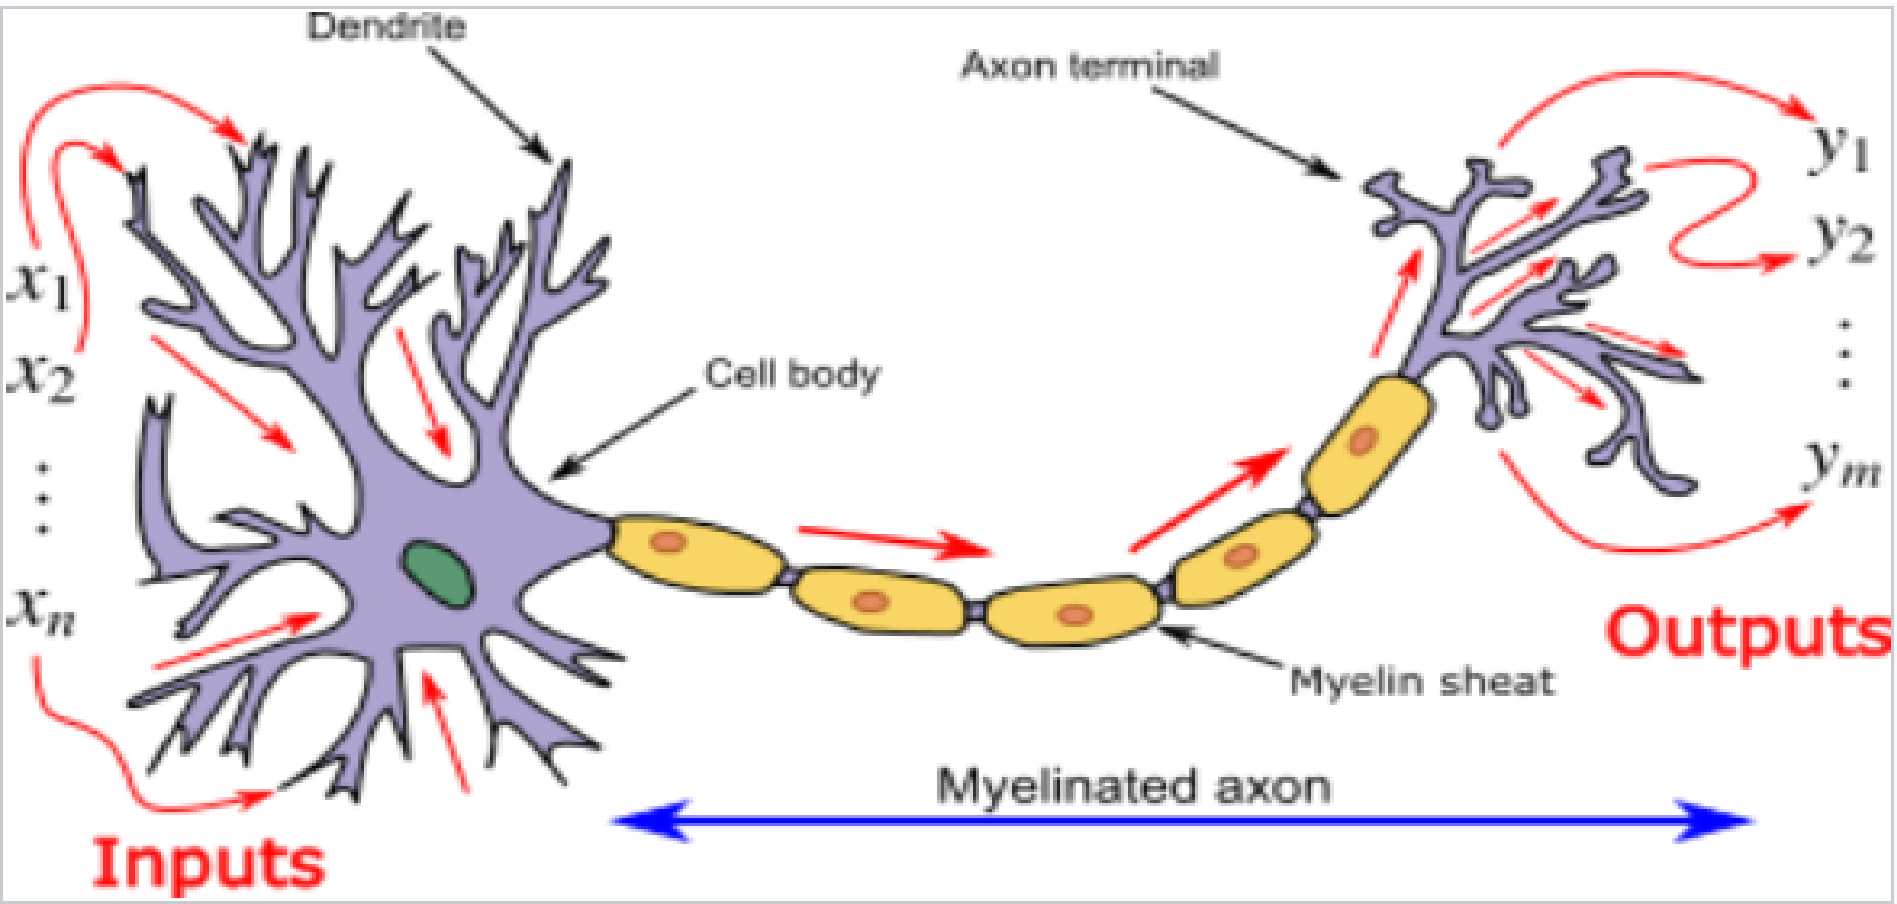

Fig. 1. Neuron and myelinated axon, with signal flow from inputs at dendrites to outputs at axon terminals. The signal is a short electrical pulse called action potential or 'spike'.

A typical neuron in the mammalian **neocortex receives *thousands* of synaptic inputs**

Neuronal Excitability

Steps of information processing:

- **Synapses**: connection between neurons
- **Dendrites**: receive inputs
- **Cell body**: sums currents from dendrites
- **Axon**: sends to action potentials

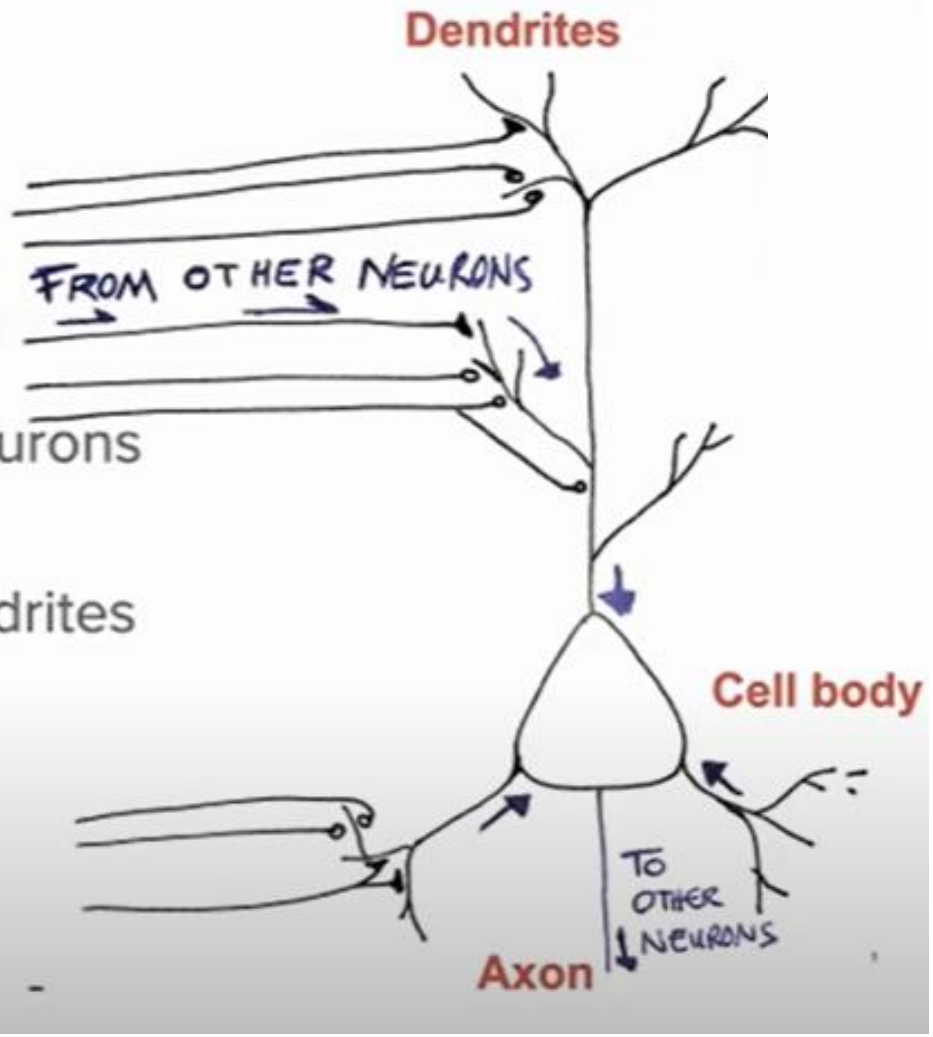

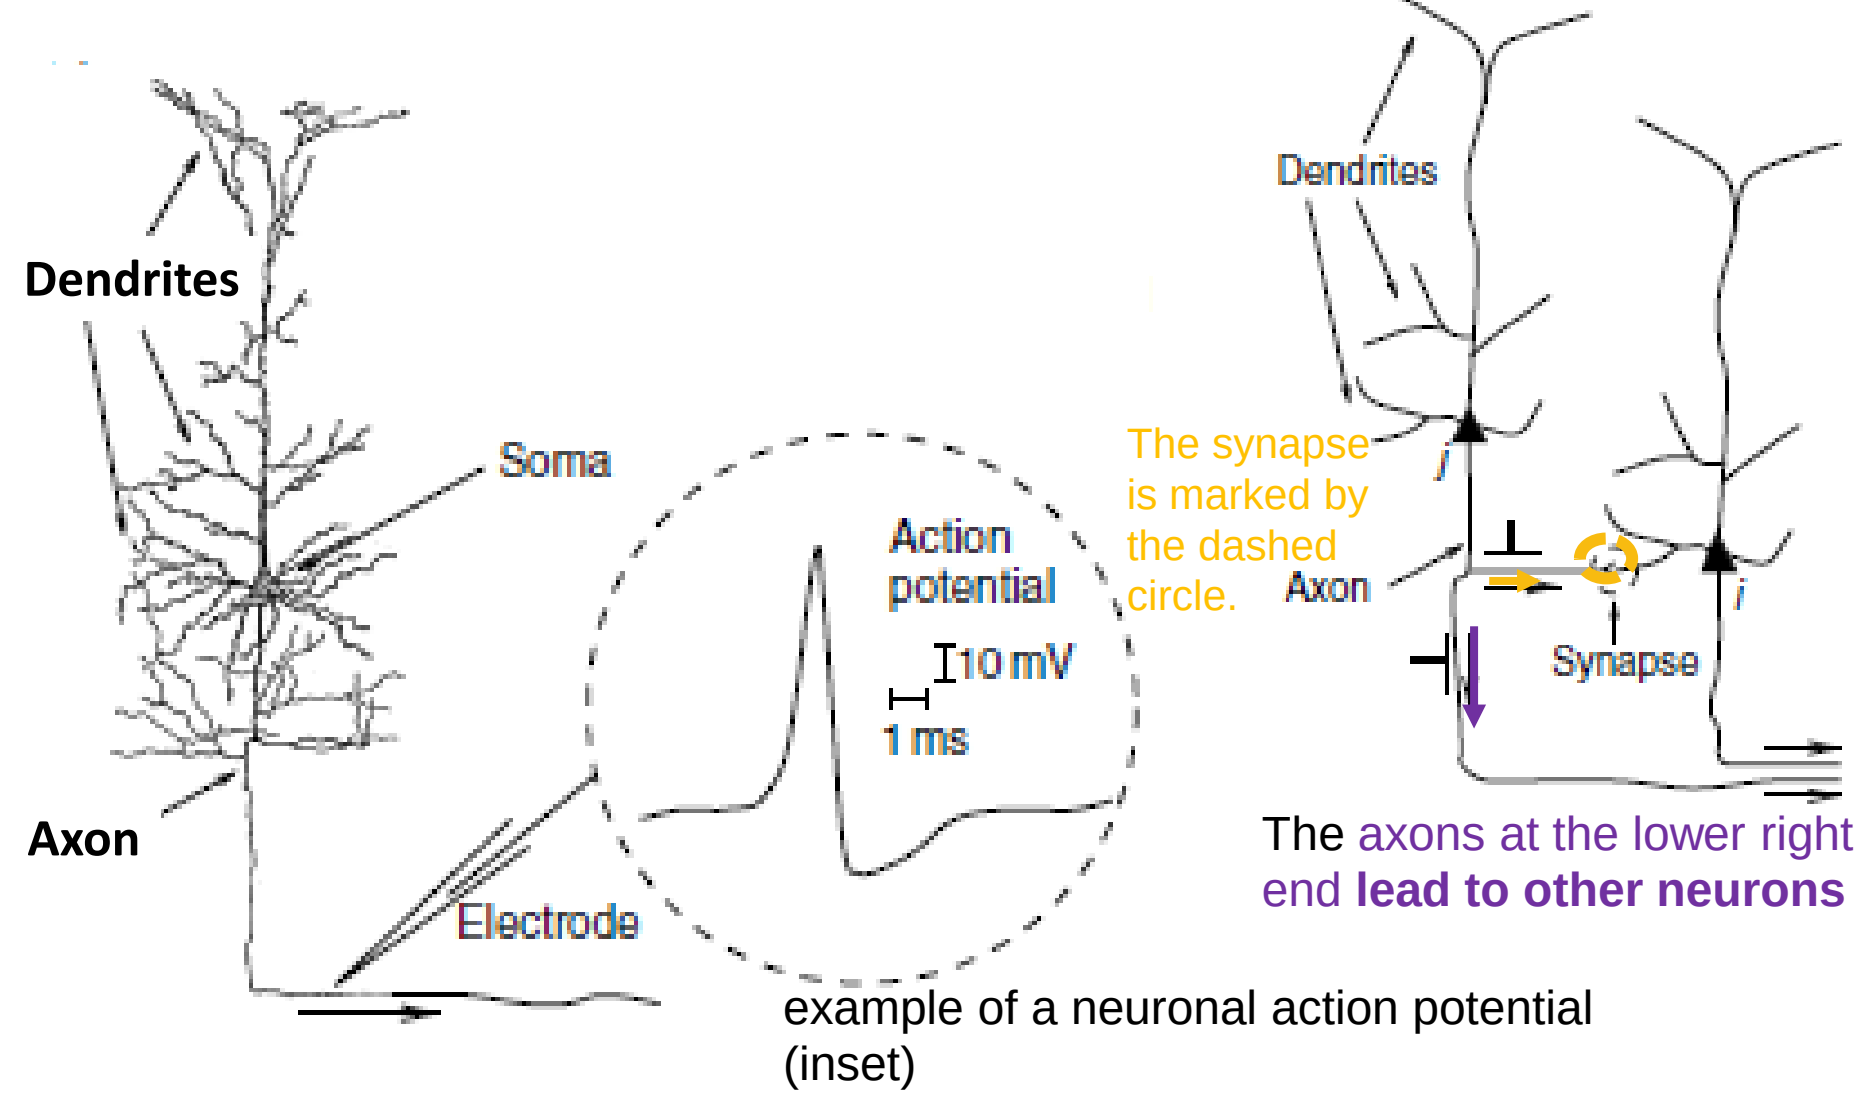

example of a neuronal action potential (inset)

- Single neuron in a drawing by Ram´on y Cajal. Dendrites, soma, and axon can be clearly distinguished.
- The action potential is a ***short voltage pulse of 1–2 ms*** **duration** and **an amplitude of about 100mV.**

Mathematical Model for Predicting the Spikes based on the Input Current

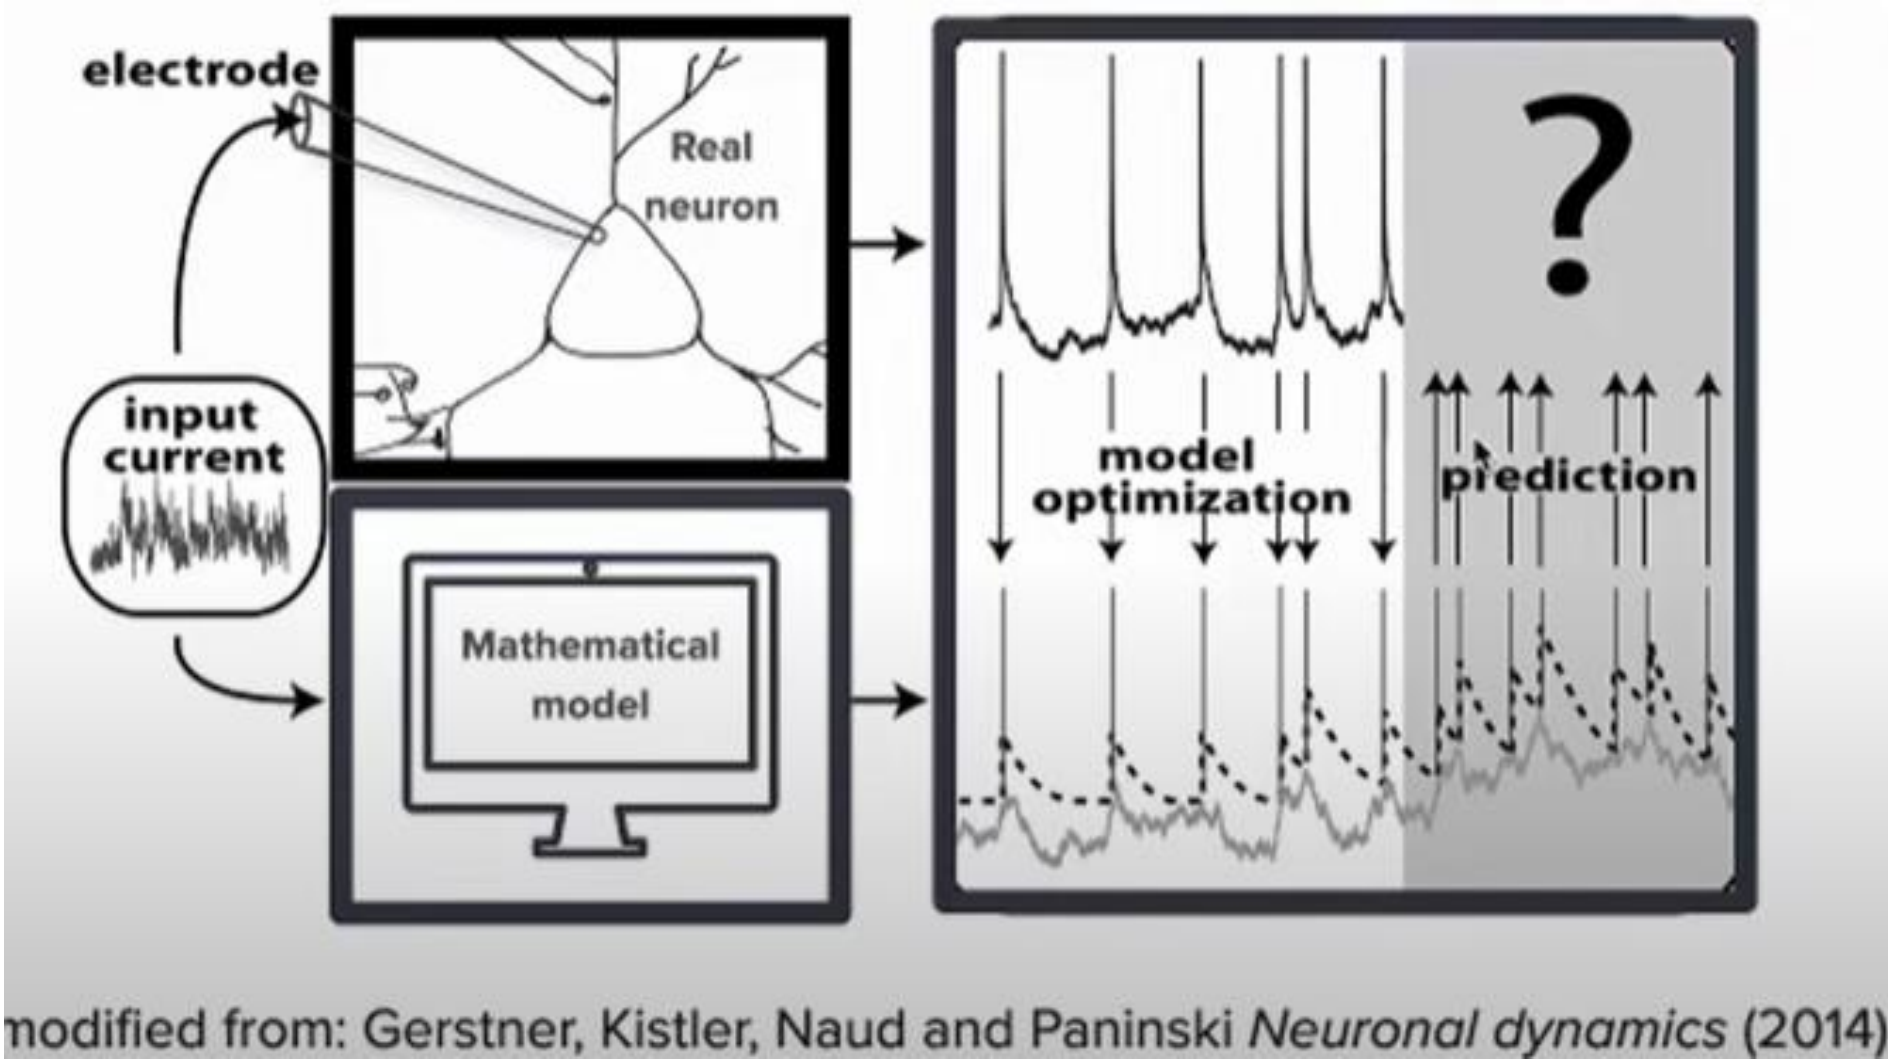

If a neuron is a core **computational unit**, similar to the ***transistor,*** what is its **input and output function?**

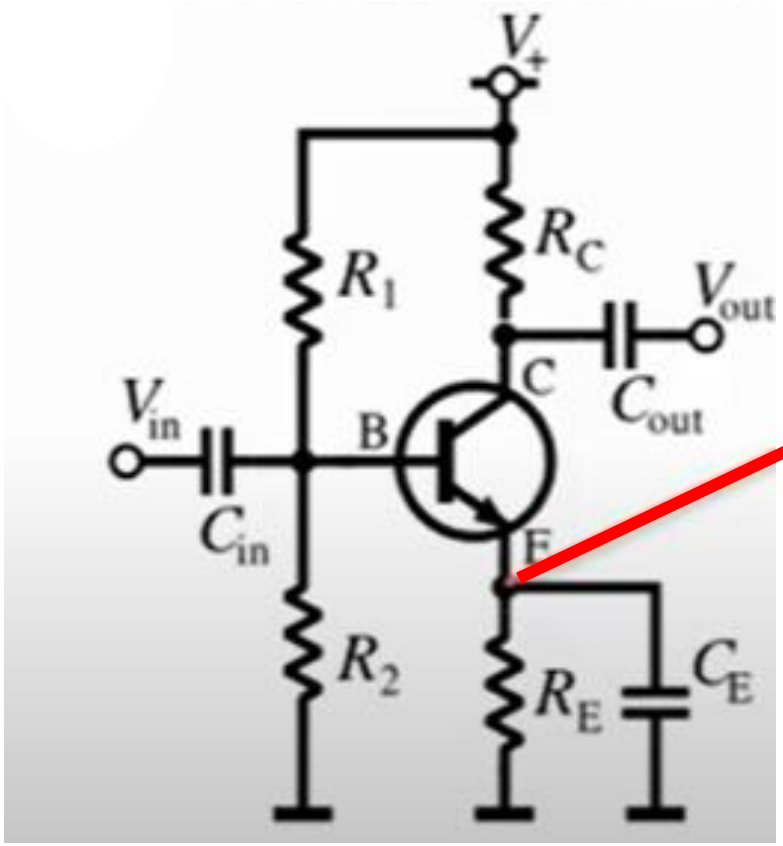

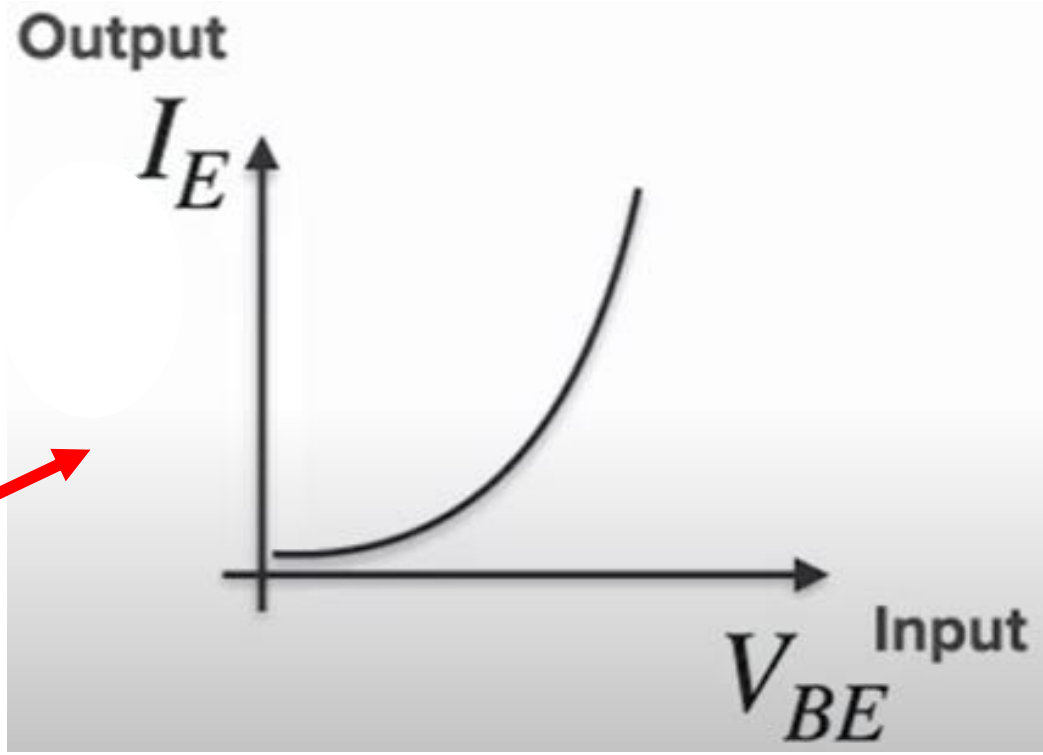

Βασικά Κυκλωμάτων

- **Πυκνωτής** (capacitor): Αποθηκεύει ηλεκτρικό φορτίο και επομένως ηλεκτρική ενέργεια

  Q_c: το **φορτίο (current load)** του θετικά φορτισμένου οπλισμού
- **Τάση πυκνωτή V (voltage)**: η **διαφορά δυναμικού** μεταξύ των οπλισμών ενός φορτισμένου πυκνωτή (Vc)
- **Χωρητικότητα** (capacity) του πυκνωτή **C** = Q / **V**
- **Αντίσταση R (resistance)**: δυσκολία διέλευσης του ρεύματος

  R = **V** / **I**, τάση του υλικού (V), προς το **ρεύμα I**

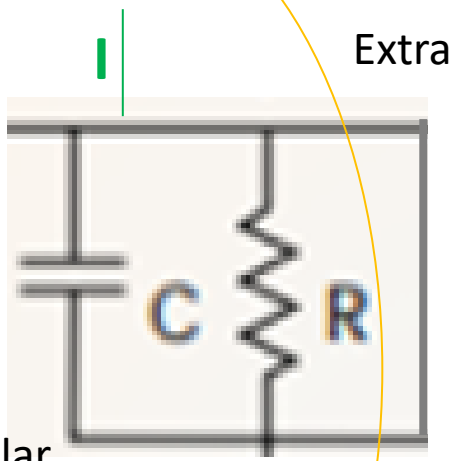

Ohm's law states that the current through a conductor between two points is directly proportional to the voltage across the two points.

Characterization of the input-output function of a neuron

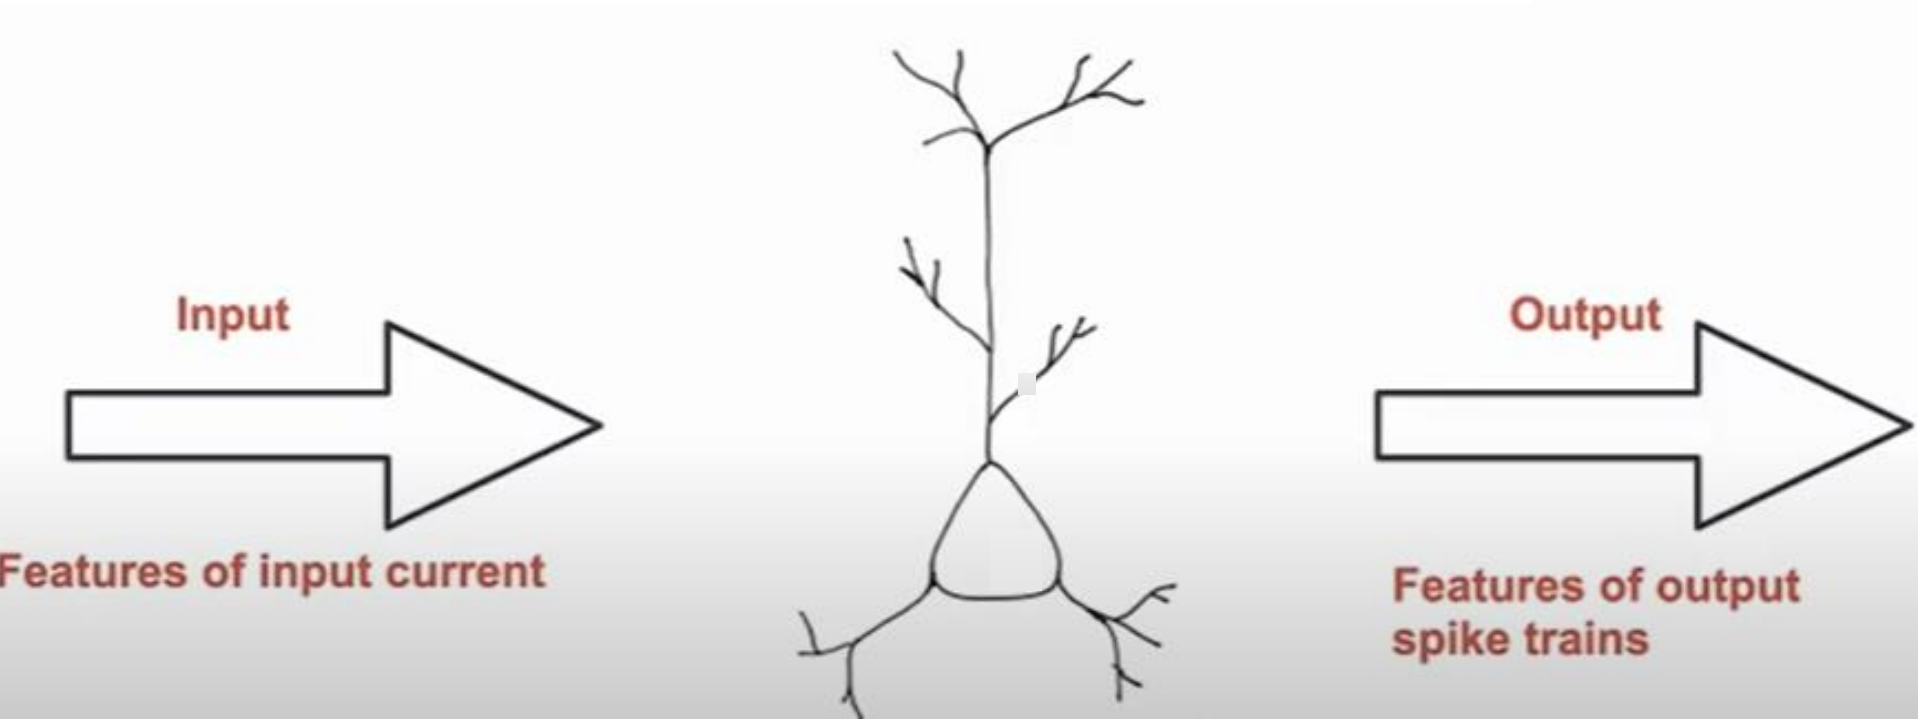

Electrical properties of neurons: the passive membrane.

A neuron receives a (positive) **input current** ***I*(*t*)** which increases the electrical charge inside the cell.

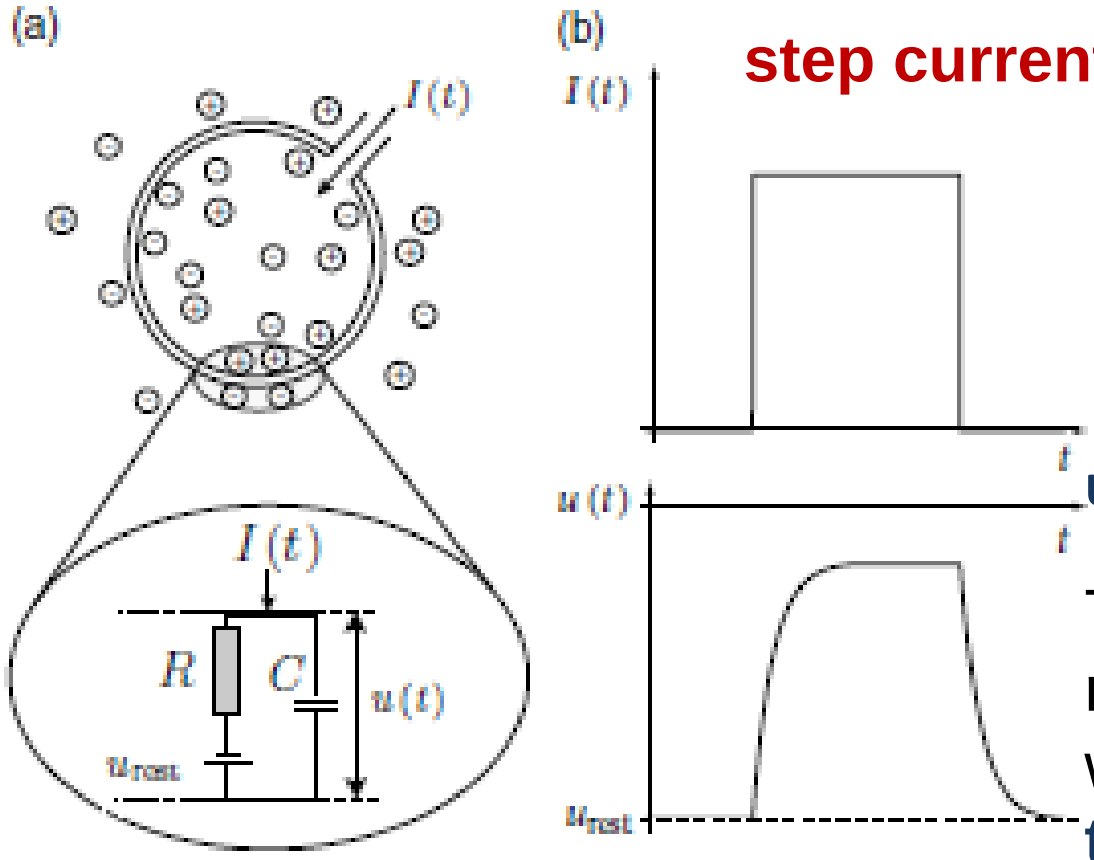

The **cell membrane** reacts to a step current with a smooth **voltage trace**

The **cell membrane acts like a capacitor** in parallel with a **resistor** which is in line with a **battery of potential u_{rest}**

Use the law of current conservation and split the driving current into two components:

$I(t) = I_R + I_C$

From Ohm's law: $\boldsymbol{I_R = u_R / R}$, where $u_R = u - u_{rest}$

C = q / u

$I_C = dq / dt = C\, du/dt$

Capacity C
Charge q
Voltage u
Voltage through resistor u_R
Capacitive current I_C
Resistive current I_R

$$\tau_m \frac{du}{dt} = -[\, u(t) - u_{rest}\,] + R\, I(t)$$

Leaky Integrator term $\tau_m = R\, C$

Decay of Membrane Potential

- In the absence of input, the **membrane potential decays exponentially** to its resting value
- Characteristic time of decay: membrane time constant $\tau_m = RC$
- For a typical neuron, it is **~10 ms**, long compared to the **duration of a spike (~1ms)**

Leaky Integrator Equation for Output

Basic model of a neuron

Input from **multiple synapses**

v: **membrane potential**

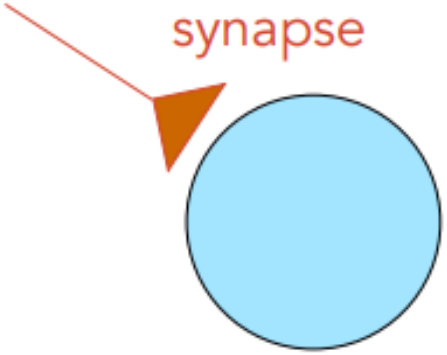

$$\tau_n \frac{dv}{dt} = -v + Input$$

change in firing rate over time

intrinsic decay

synaptic input

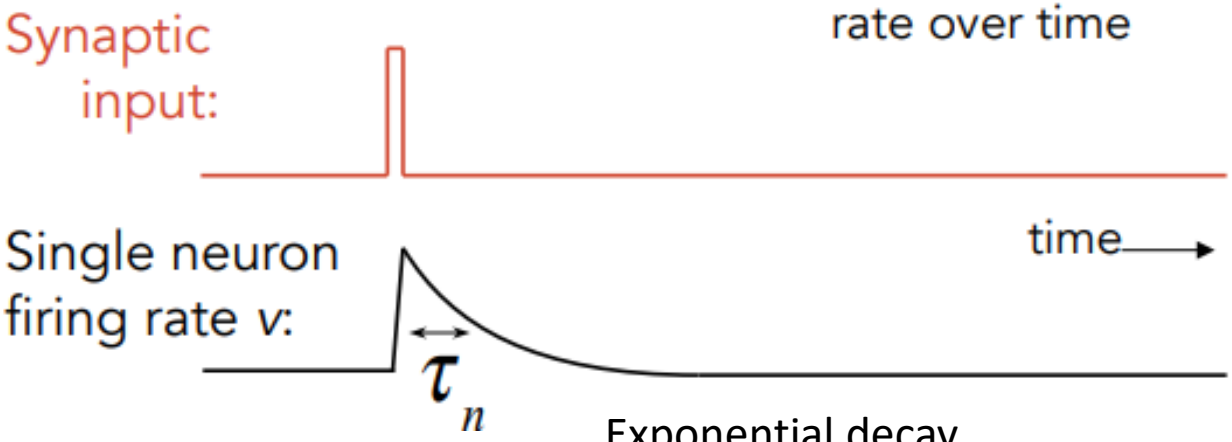

τ_n: depends on **neuron n**
rate that the neuron **leaks**

Tractable model

Electrical Input–Output Membrane Voltage Models

- Produce a **prediction for membrane output voltage** as a function of electrical stimulation given as current or voltage input
- Different **functional relationships** between the input current & output voltage and in the **level of details**

Examples of models:

- Predict the moment of occurrence of output spike (also known as "**action potential**")
- Account for **sub-cellular processes** and can be either deterministic or probabilistic

Natural Stimulus or Pharmacological Input Neuron Models

- Connect the ***input* stimulus** (e.g., pharmacological, natural) to the **probability of a spike event**
- Input stage is not electrical but rather has either pharmacological (chemical) concentration units or physical units that characterize an external stimulus, e.g., light, sound, physical pressure
- Output stage represents the probability of a spike event and not an electrical voltage

Biophysical description

I(t) Current impinging on excitable membrane patch

V(t) Membrane potential (διαφορά δυναμικού)

C Capacitance of the membrane (χωρητικότητα)

g_L Conductance of the membrane (αγωγιμότητα)

E_L Equilibrium potential of "leak"

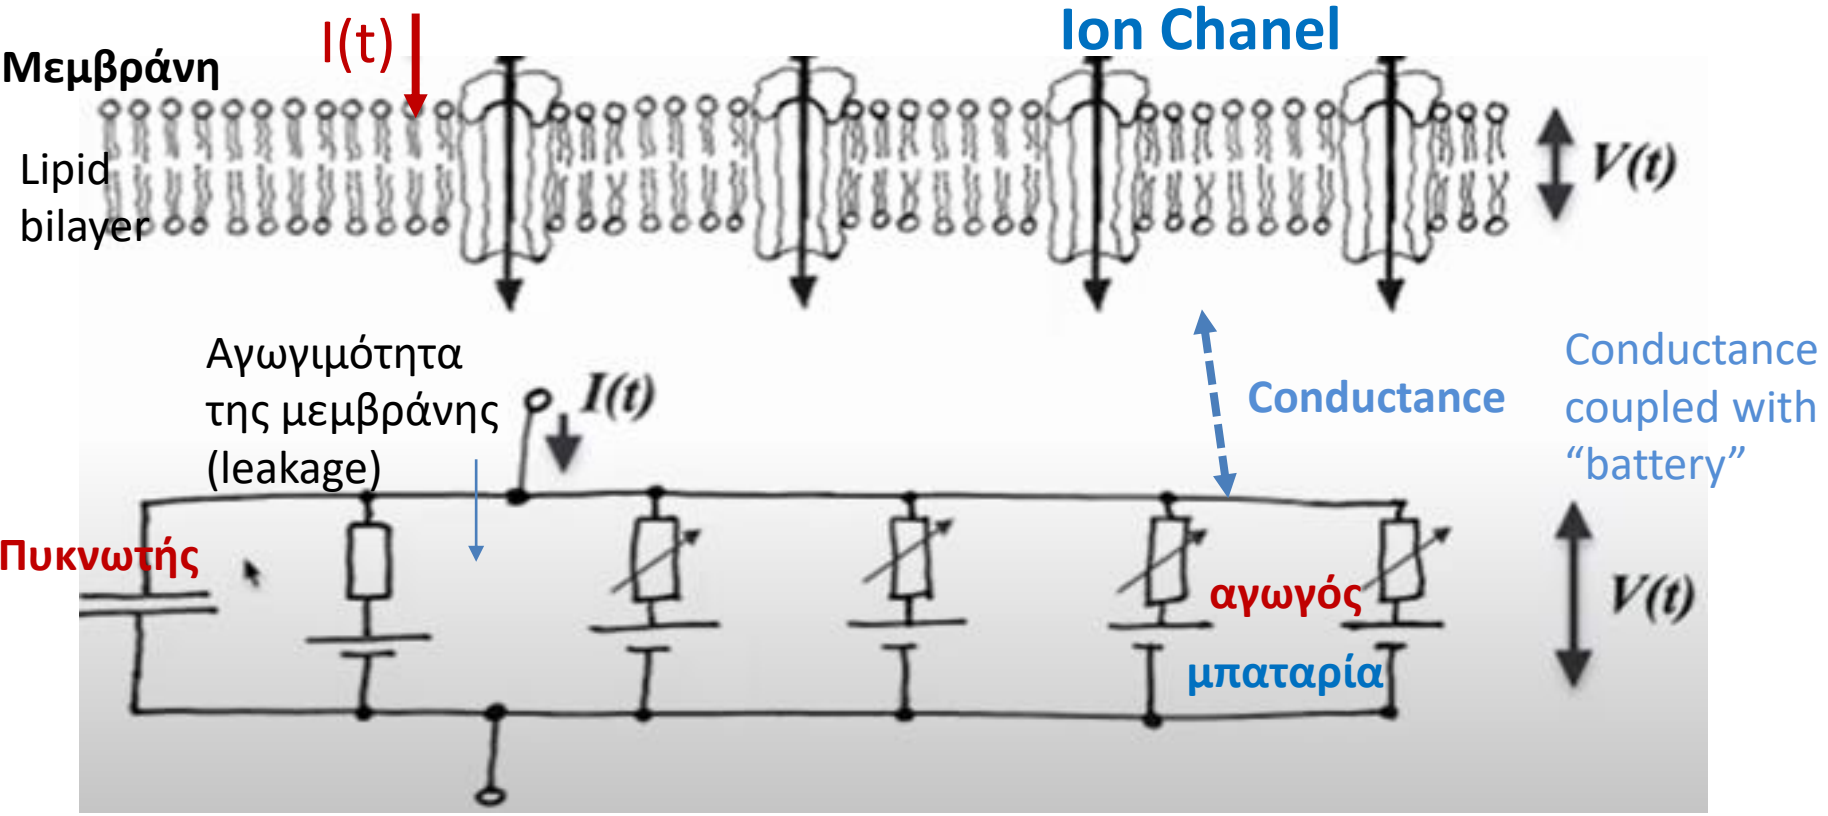

Leaky Integrate and Fire Model

Machinery of generation of action potential

2) Replace by a threshold for ***spike emission*** Followed by a ***reset*** to a fixed potential

Action potential of ions channels *other than* the *stereotypical generation of spikes*

1) We will ignore them…

The impinging **current** will either **charge the capacitor** or **leak through the membrane or flow through the different ion channels**

$$I = C\frac{dV}{dt} + g_L(V - E_L) + I_{Na} + I_K + I_H + I_{AHP}$$

I(t) Current impinging on excitable membrane patch

V(t) Membrane potential

C Capacitance of the membrane

g_L Conductance of the membrane

E_L Equilibrium potential of "leak"

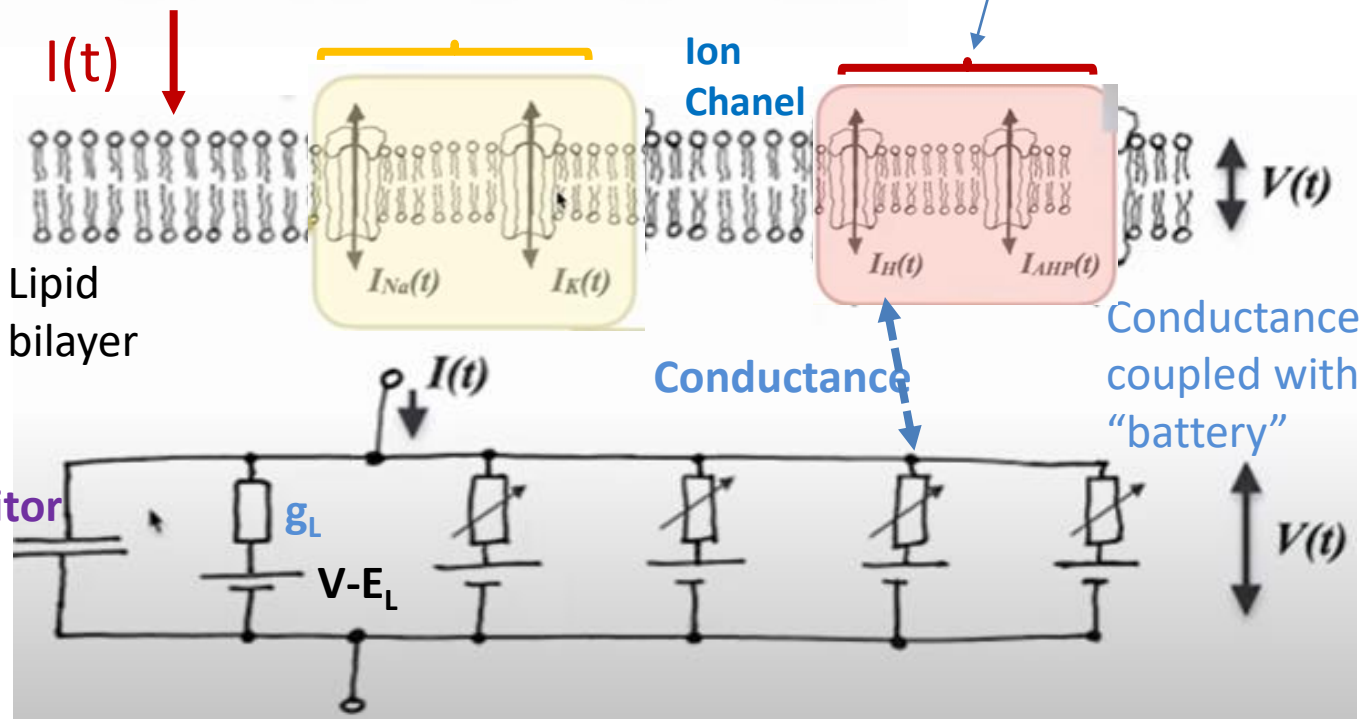

Leaky Integrate-and-Fire (LIF)

$$C_m \frac{dV}{dt} = -g_L(V - E_L) + I$$

If $V(t) = V_{th}$ then $V(t + \Delta) = E_L$

Δ: the time it takes for the action potential to be generated, **1-2ms** (***refractory period)***

When the action potential reaches a threshold, a **spike is generated (*fire*)** and ***stops*** **the dynamics for time Δ**

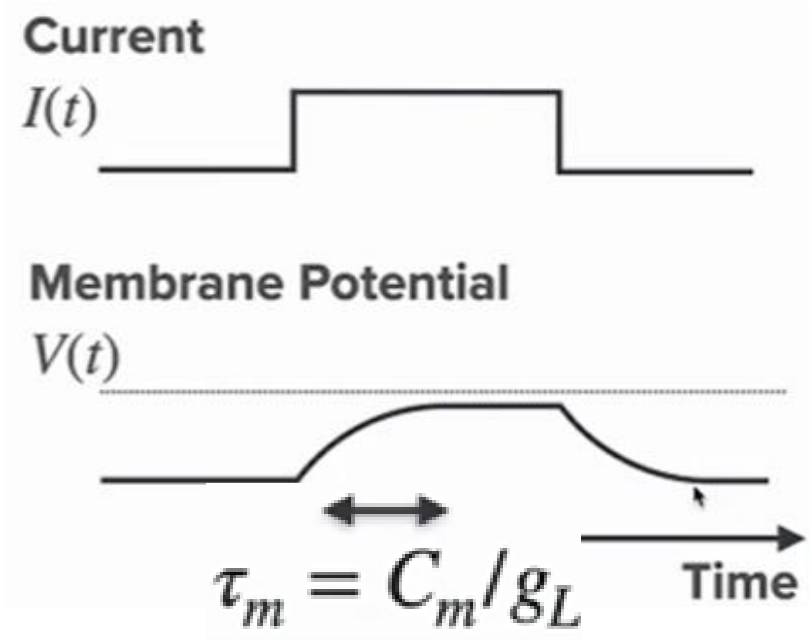

Subthreshold current step: Exponential relaxation to a steady-state.

$$V(t) = (\frac{I}{g_L} + E_L)[1 - e^{-t/\tau_m}]$$

$$C_m \frac{dV}{dt} = -g_L(V - E_L) + I$$

If $V = V_{th}$ then $V(t + \Delta) = E_L$

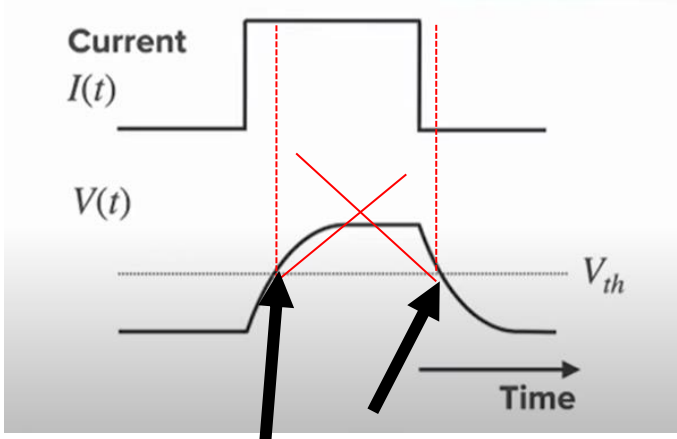

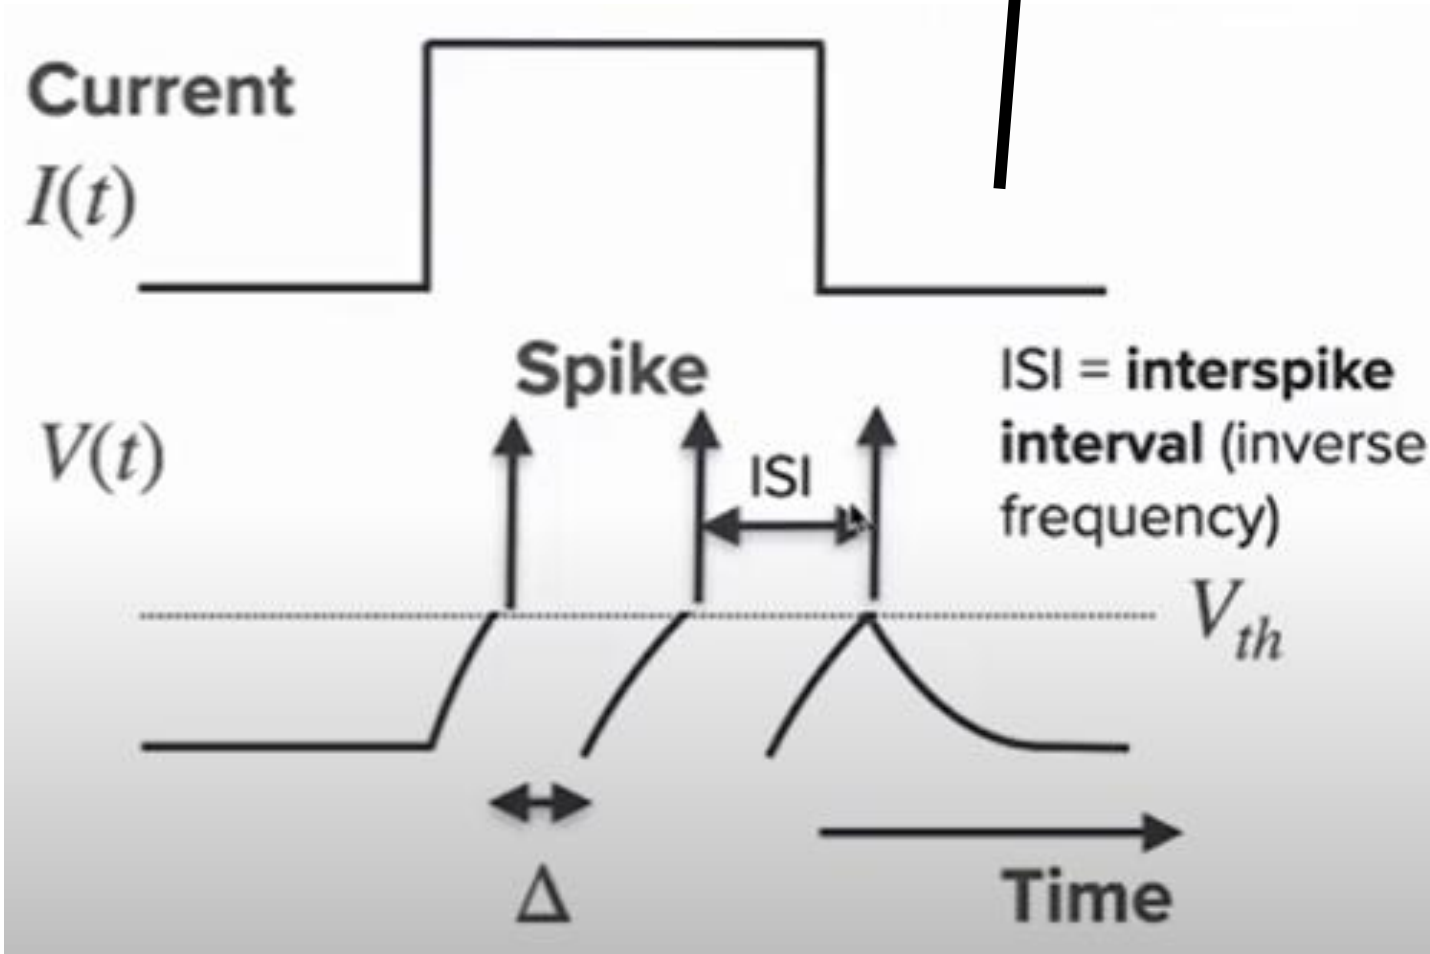

Suprathreshold current step

Regular firing

Leaky Integrate and Fire Model

The firing frequency depends on the magnitude of the current being injected

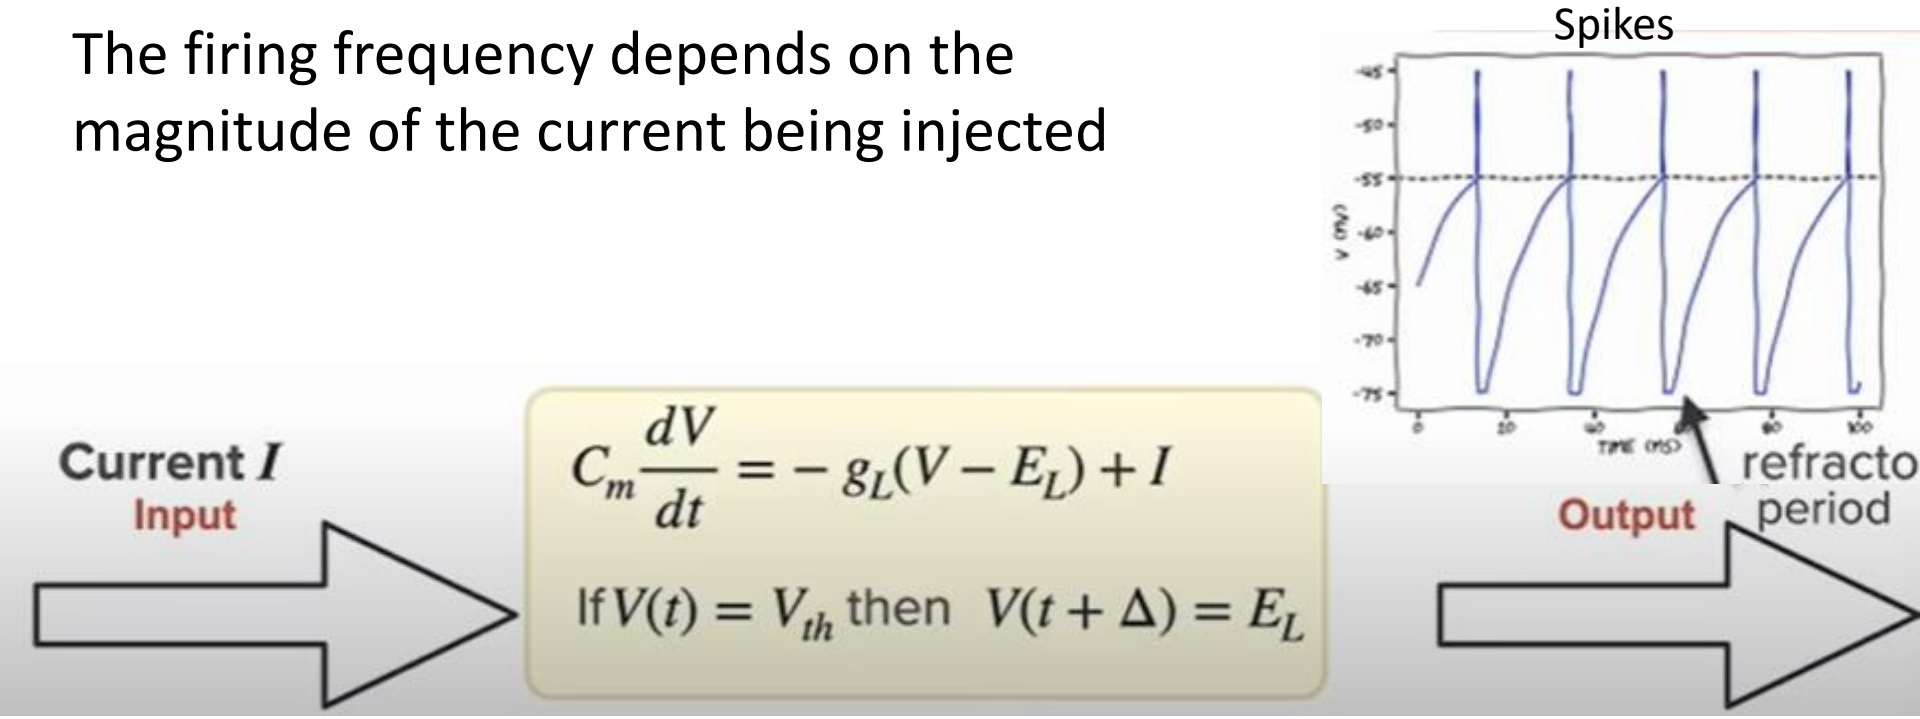

The **refractory period** of a neuron is the time in which a nerve cell is unable to fire an action potential

Do spikes always have the same shape?

Yes! Spikes follow stereotypical time course within 1-2 ms of onset
Notable **exception**: spikes late in a high frequency burst

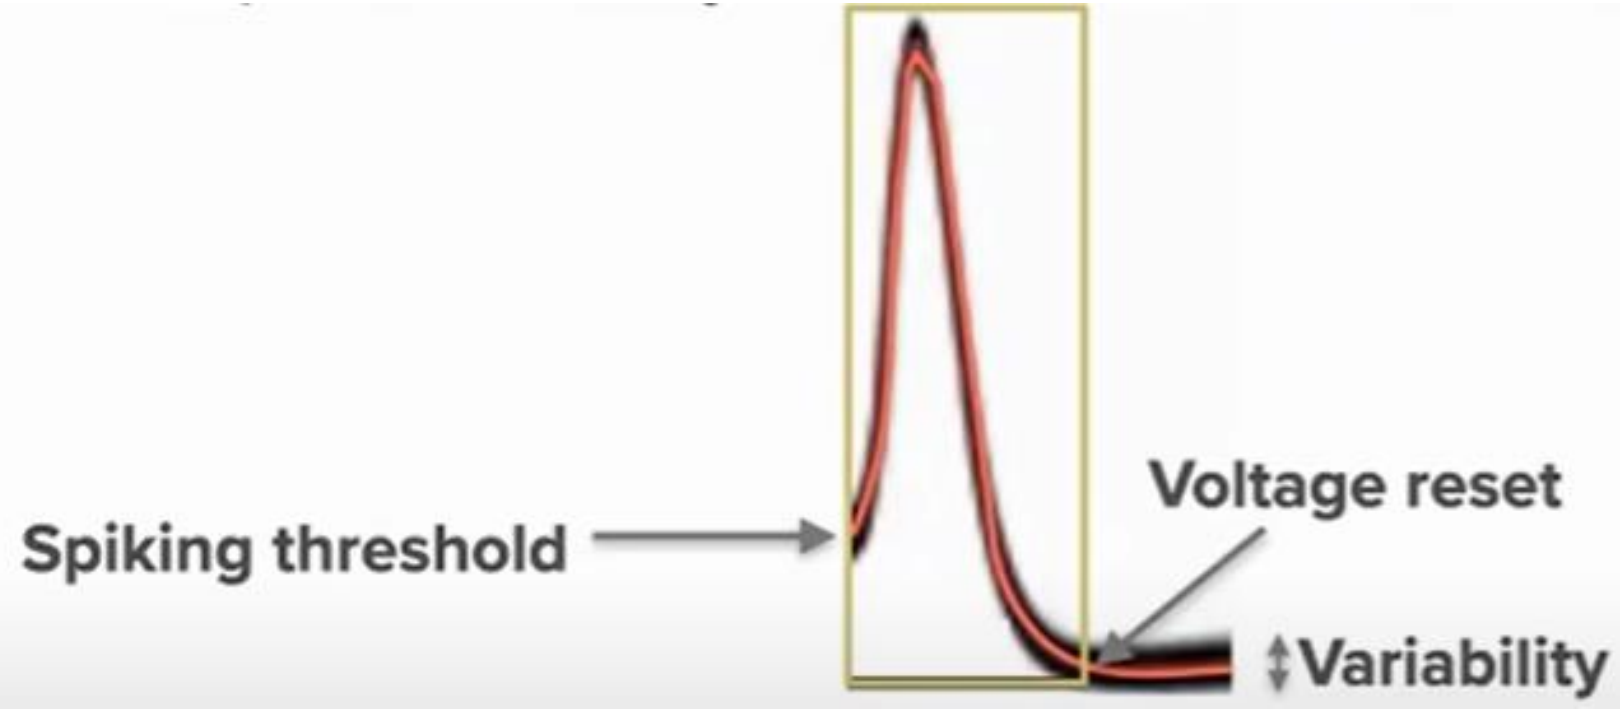

Across > 100 spikes

If the shape of an action potential is always the same, the ***shape cannot be used to transmit information:***
Rather information is “carried” with the presence or absence of a spike
Therefore action potentials are reduced to “events” that happen at a precise moment in time

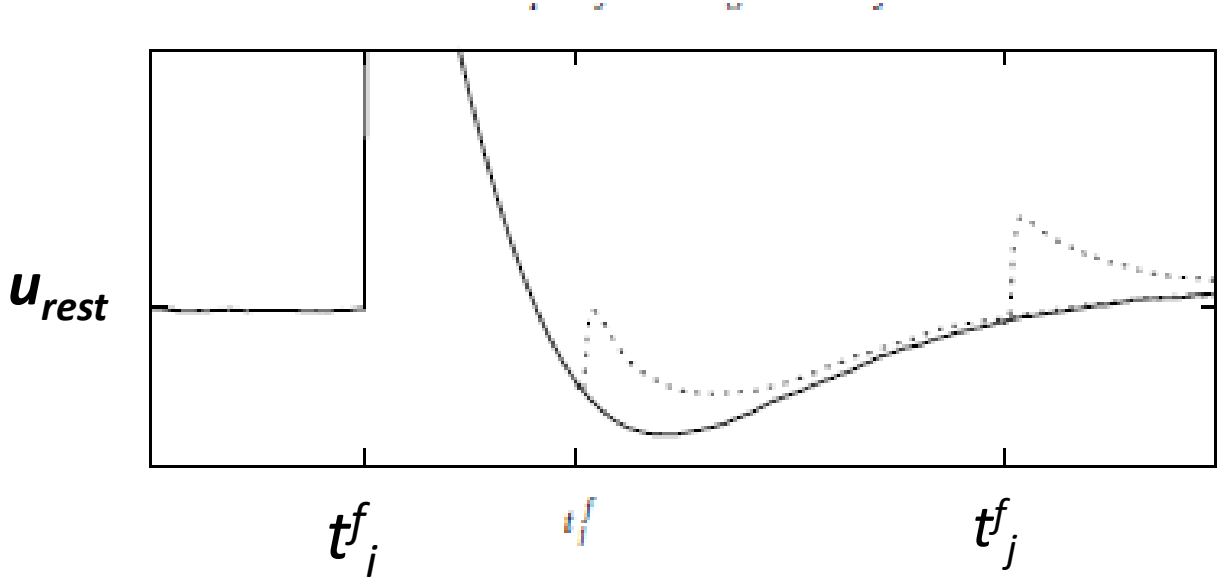

The shape of postsynaptic potentials (dashed lines) depends on the time $t - t^f_i$ that has passed since the last output spike of **neuron *i***.
The postsynaptic spike has been triggered at time t^f_i
A presynaptic spike that arrives at time t^f_j shortly after the spike of the postsynaptic neuron has a
smaller effect than a spike that arrives much later.
(Data is courtesy of Thomas Berger. Berger *et al*., 2009).

- If the shape of an action potential is always the same, the shape cannot be used to transmit information:
- Rather information is “carried” with the presence or absence of a spike
- Therefore action potentials are reduced to “events” that happen at a precise moment in time

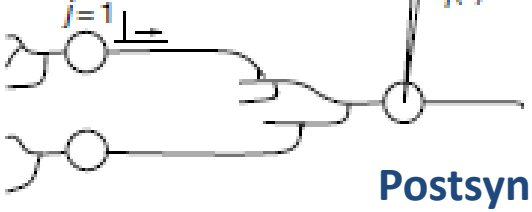

Postsynaptic neuron i

Presynaptic neurons j=1,2

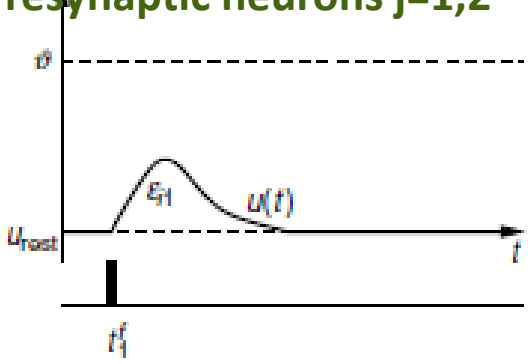

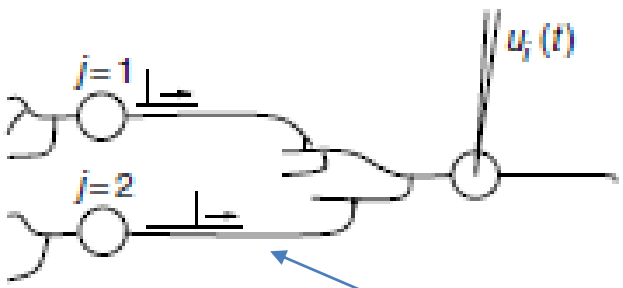

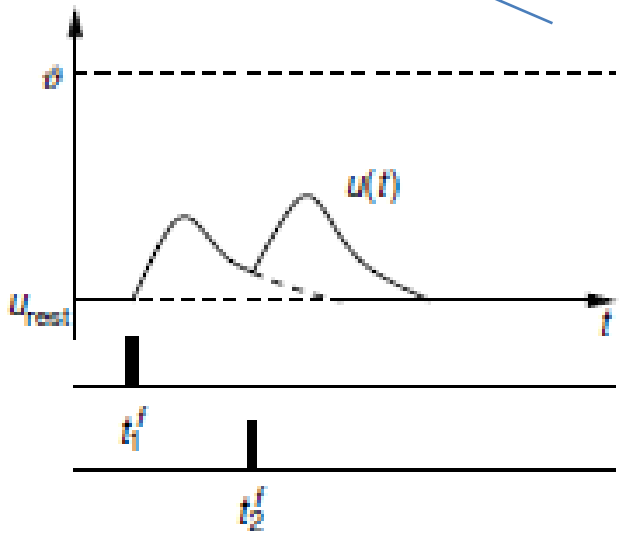

An **input spike** from presynaptic neuron $j = 2$ that arrives shortly after the spike from neuron $j = 1$ causes a second postsynaptic potential that adds to the first one.

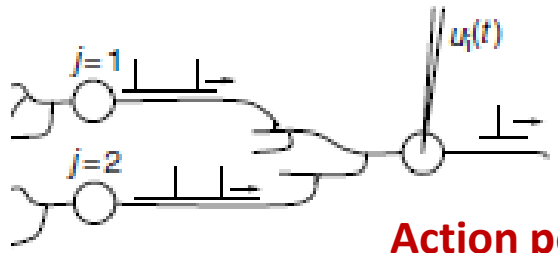

Action potential is triggered

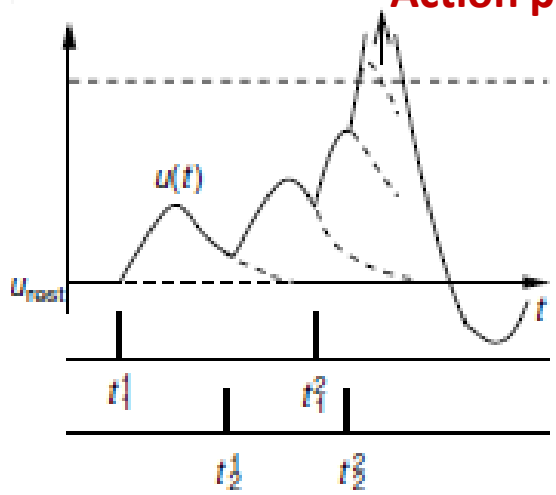

Single EPSPs have amplitudes ~ 1mV.
The **critical value for spike initiation** is ~ 20 to 30mV above the resting potential.
In most neurons, four spikes are thus ***not sufficient*** to trigger an action potential. Instead, about **20–50 presynaptic spikes** have **to arrive *within a short time window*** to trigger a postsynaptic action potential.

Simple RC model for subthreshold voltage

Well below “AP threshold”, cell membrane dynamics well-modeled by a simple RC circuit.

Membrane capacitance

Membrane conductance

Trans-membrane voltage drop

Resting voltage

Other currents (inputs, input-triggered/ voltage-dependent…)

$$C_m \frac{dV}{dt} = -g_m(V(t) - V_m) + I$$

Equivalent RC circuit:

$$V(t) = V_{in}(t) - V_{out}$$

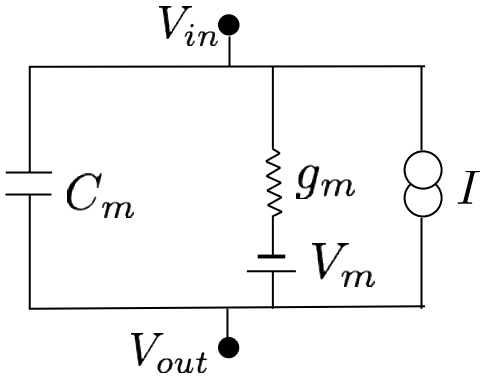

With appropriate choice of I, this includes HH and other models.

Single voltage variable V(t) in model: ignoring spatial dynamics

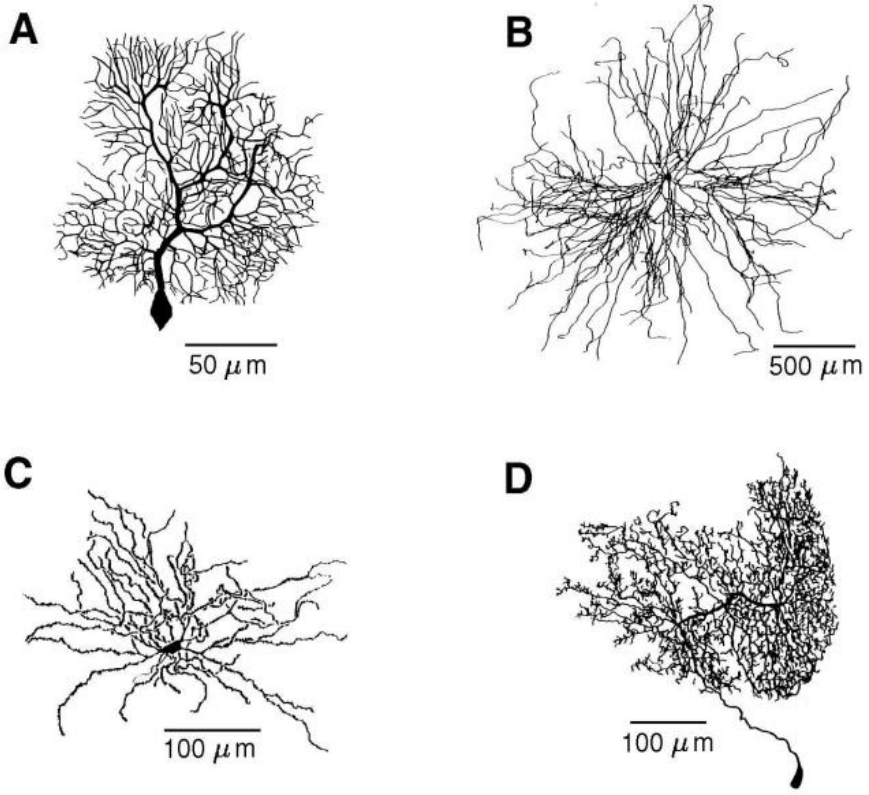

Simplest spatial models: multiple discrete equi-voltage compartments, resistively coupled.

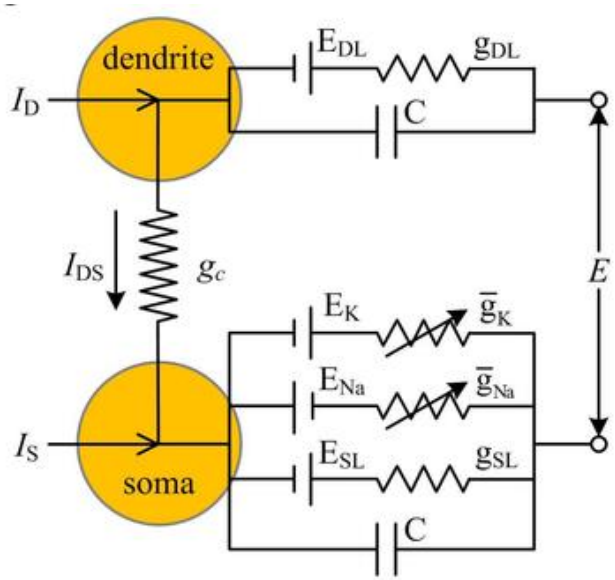

Image: from Genesis project

Modeling software for biophysically detailed and spatially extended neurons: NEURON.

Numerical integration of subthreshold voltage

Membrane capacitance

Membrane conductance

Trans-membrane voltage drop

Resting voltage

Input currents

$$C_m \frac{dV}{dt} = -g_m(V(t) - V_m) + I_{app}$$

$$\tau_m = C_m / g_m \sim 10\ ms$$

$$\bar{V} = V_m + \frac{I_{app}}{g_m}$$

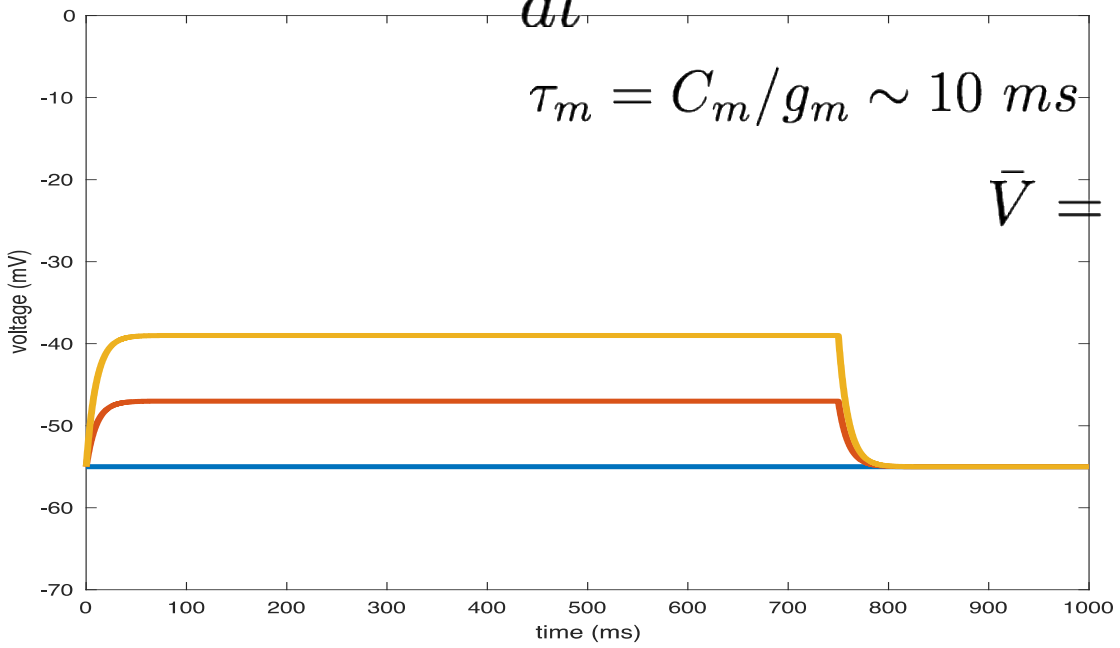

Leaky integrate-and-fire (LIF) model

Replace complex, detailed AP currents with a simple reset condition

$$C_m \frac{dV}{dt} = -g_m(V(t) - V_m) + I_{app}$$

+ spike-and-reset condition

When $V \nearrow V_{th}$
then reset $V \rightarrow V_{reset}$
and consider that the cell has spiked

Single voltage variable V(t) in model: ignoring spatial dynamics

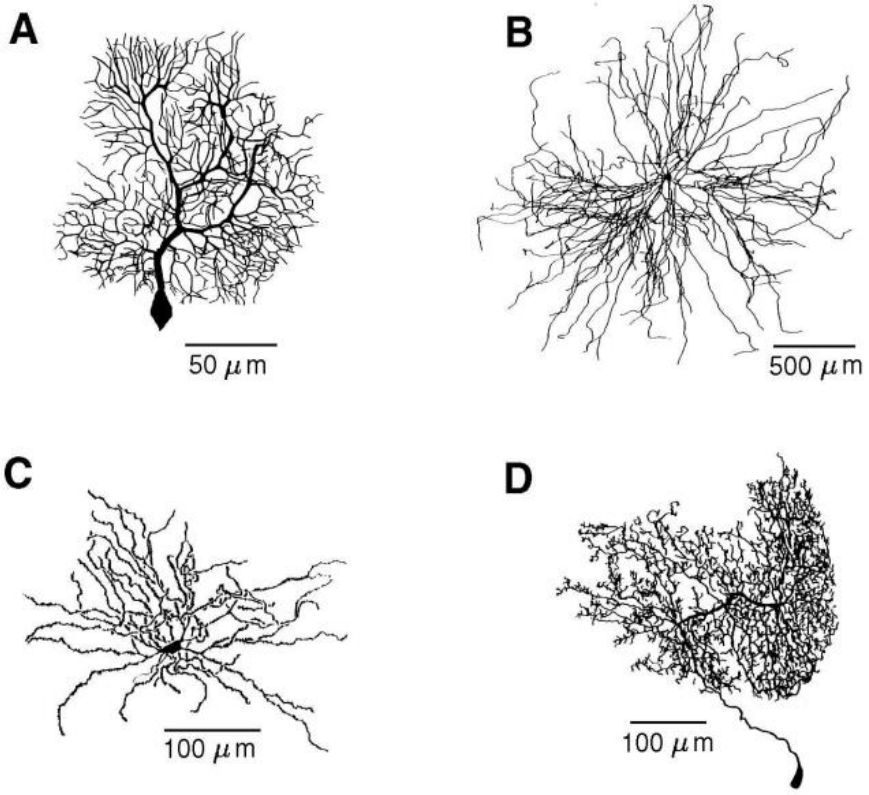

Simplest spatial models: multiple discrete equi-voltage compartments, resistively coupled.

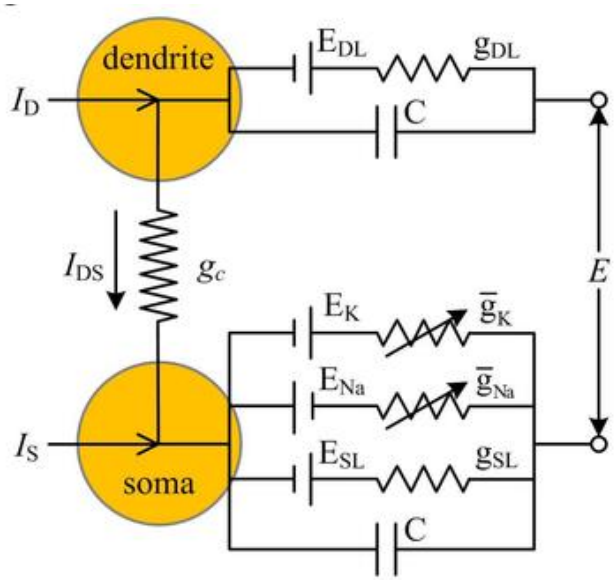

Image: from Genesis project

Modeling software for biophysically detailed and spatially extended neurons: NEURON.

Single voltage variable V(t) in model: ignoring spatial dynamics

Simplest spatial models: multiple discrete equi-voltage compartments, resistively coupled.

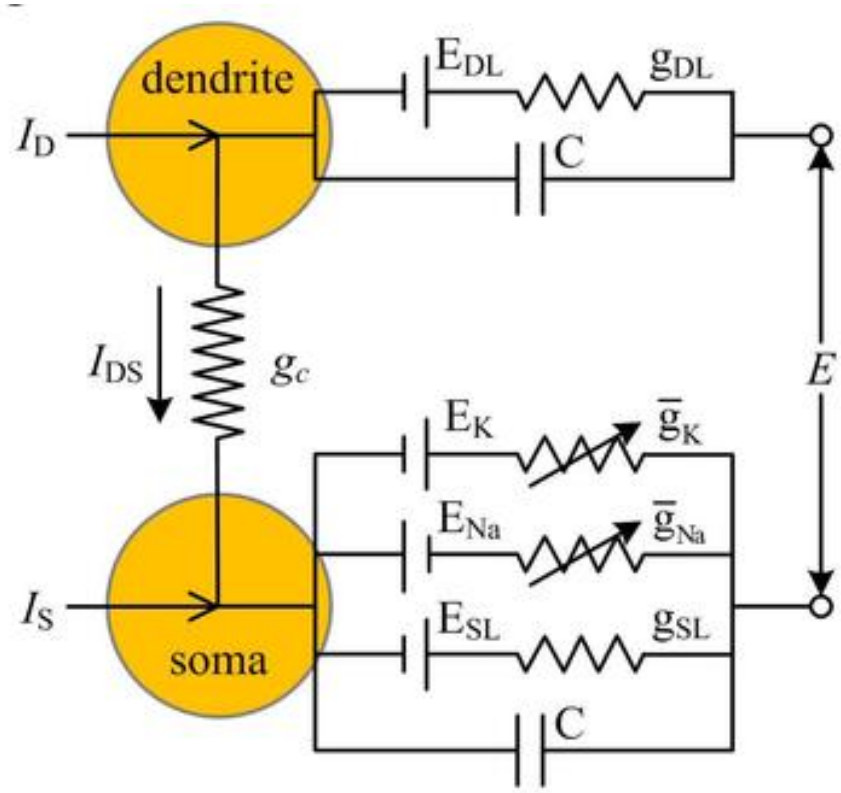

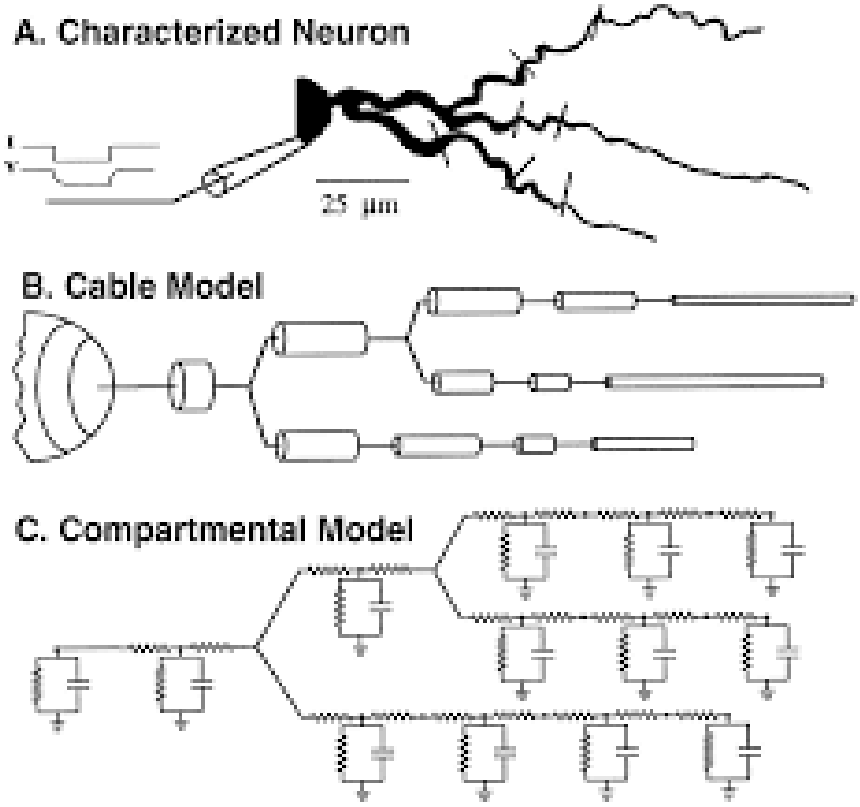

Image: f

Modeling software for biophysically detailed and spatially extended neurons: NEURON.

Simple RC model for subthreshold voltage

Membrane capacitance

Membrane conductance

Trans-membrane voltage drop

Resting voltage

other currents

$$C_m \frac{dV}{dt} = -g_m(V(t) - V_m) + I$$

Equivalent RC circuit:

$V(t) = V_{in}(t) - V_{out}$

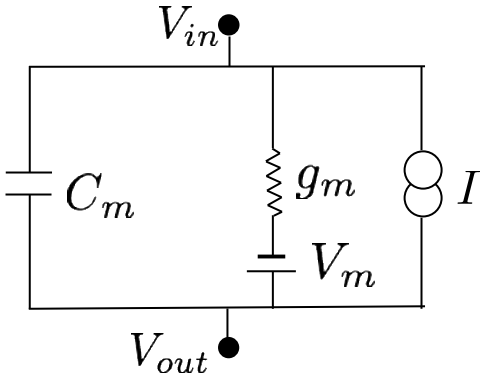

$$C_m = 1\ \mu F/cm^2$$
$$1/g_m \sim 10000\ \Omega\ cm^2$$
$$\Rightarrow \tau_m = C_m/g_m \sim 10\ ms$$

Take note of the short single-neuron time-constant (memory of single cells).

Add spike mechanism

Hodgkin-Huxley model for both subthreshold voltage and AP generation

$$C_m \frac{dV}{dt} = -g_m(V(t) - V_m) + I_{app} + I_{spk}(V(t), t)$$

Complex, nonlinear voltage-dependent currents for AP generation (see Hodgkin-Huxley model for details)

Equivalent RC circuit:

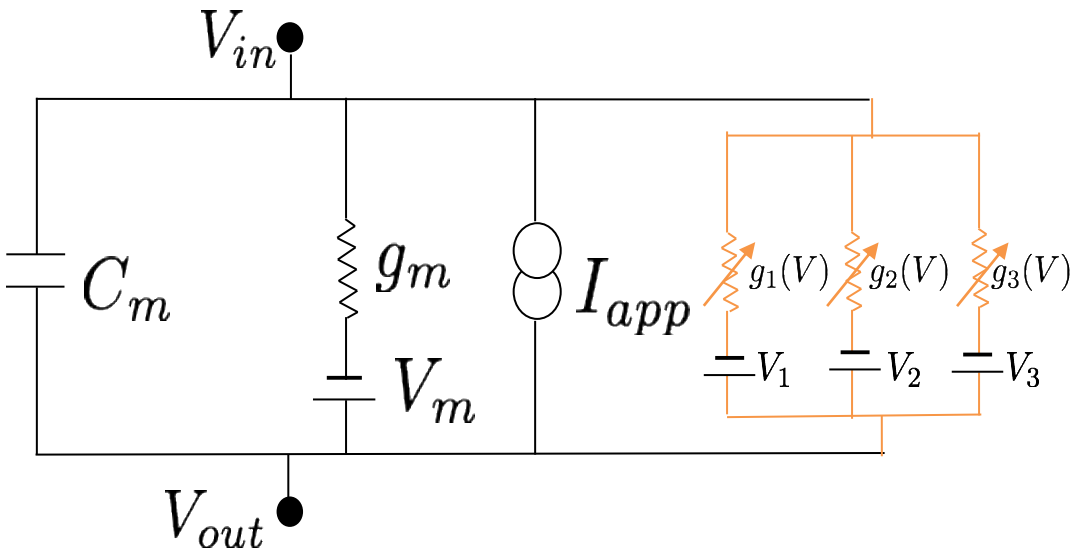

Leaky integrate-and-fire (LIF) model

Replace complex, detailed AP currents with a simple reset condition

$$C_m \frac{dV}{dt} = -g_m(V(t) - V_m) + I_{app}$$

+ spike-and-reset condition

When $V \nearrow V_{th}$
then reset $V \rightarrow V_{reset}$
and consider that the cell has spiked

Leaky integrate-and-fire (LIF) model

Replace complex, detailed AP currents with a simple reset condition

$$C_m \frac{dV}{dt} = -g_m(V(t) - V_m) + I_{app}$$

When $V \nearrow V_{th}$
then reset $V \rightarrow V_{reset}$
and consider that the cell has spiked

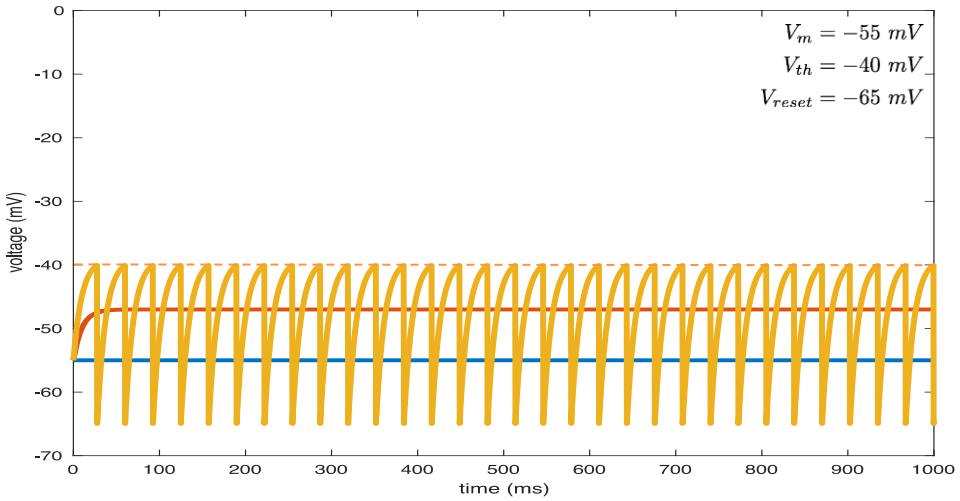

As I_{app} increases, firing rate will increase

Synaptic activation model

Each **synapse is a linear, low-pass filter of the presynaptic neuron's spikes**; activation is a dimensionless variable than can be thought of as "fractional activity"

$$\frac{ds}{dt} = -\frac{s}{\tau_{syn}} + \beta \sum_{\alpha} \delta(t - t_{spk,\alpha})$$

Simple exponential decay

Upward increment whenever there is a spike (when $t = t_{spk,\ \alpha}$)

Aside: Dirac delta function

$$\int_{-\epsilon}^{\epsilon} \delta(x)dx = 1$$

→ has units of the inverse of its argument

$$\int^{I} f(x)\delta(x-a)dx = f(a) \quad \text{if } a \in I$$

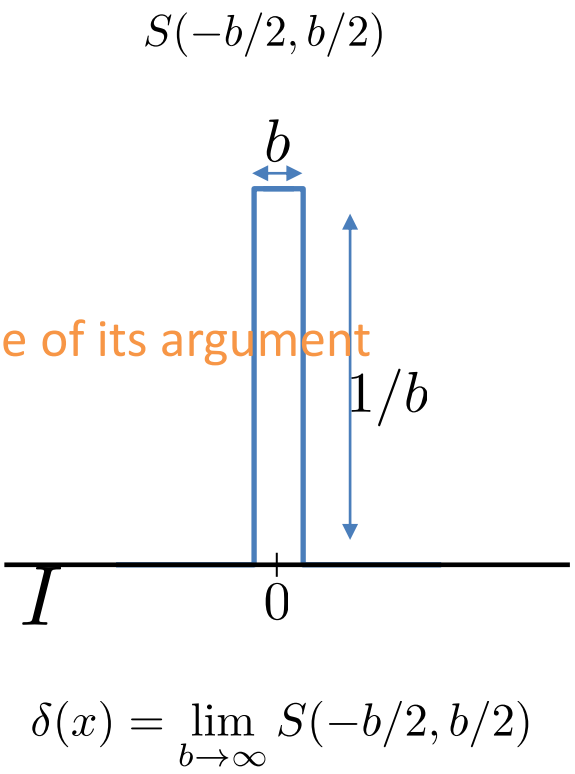

$$\int^{I} f(x)\delta(x-a)dx = 0 \quad \text{if } a \notin I$$

Aside: Kronecker delta

$$\delta_{ij} = \begin{cases} 1 & \text{if } i = j \\ 0 & \text{if } i \neq j \end{cases}$$

Synaptic Activation Model Numerical Integration

$$\frac{ds}{dt} = -\frac{s}{\tau_{syn}} + \beta \sum_{\alpha} \delta(t - t_{spk,\alpha})$$

Discretize equation in time

$$\frac{s(t+\Delta t) - s(t)}{\Delta t} = -\frac{s(t)}{\tau_{syn}} + \frac{\beta}{\Delta t} \int_{t}^{t+\Delta t} dt' \sum_{\alpha} \delta(t' - t_{spk,\alpha})$$

$$= -\frac{s(t)}{\tau_{syn}} + \frac{\beta}{\Delta t} \sum_{\alpha} \delta_{t,t^{b}_{spk,\alpha}}$$

where the b superscript in $t^{b}_{spk,\alpha}$ indicates the spike time bin in place of the precise spike time

$$s(t+\Delta t) = (1 - \frac{\Delta t}{\tau_{syn}}) s(t) + \beta \sum_{\alpha} \delta_{t,t^{b}_{spk,\alpha}}$$

Synaptic Activation Model

Each synapse is a linear, low-pass filter of the presynaptic neuron's spikes; activation is a dimensionless variable than can be thought of as "fractional activity"

$$\frac{ds}{dt} = -\frac{s}{\tau_{syn}} + \beta \sum_{\alpha} \delta(t - t_{spk,\alpha})$$

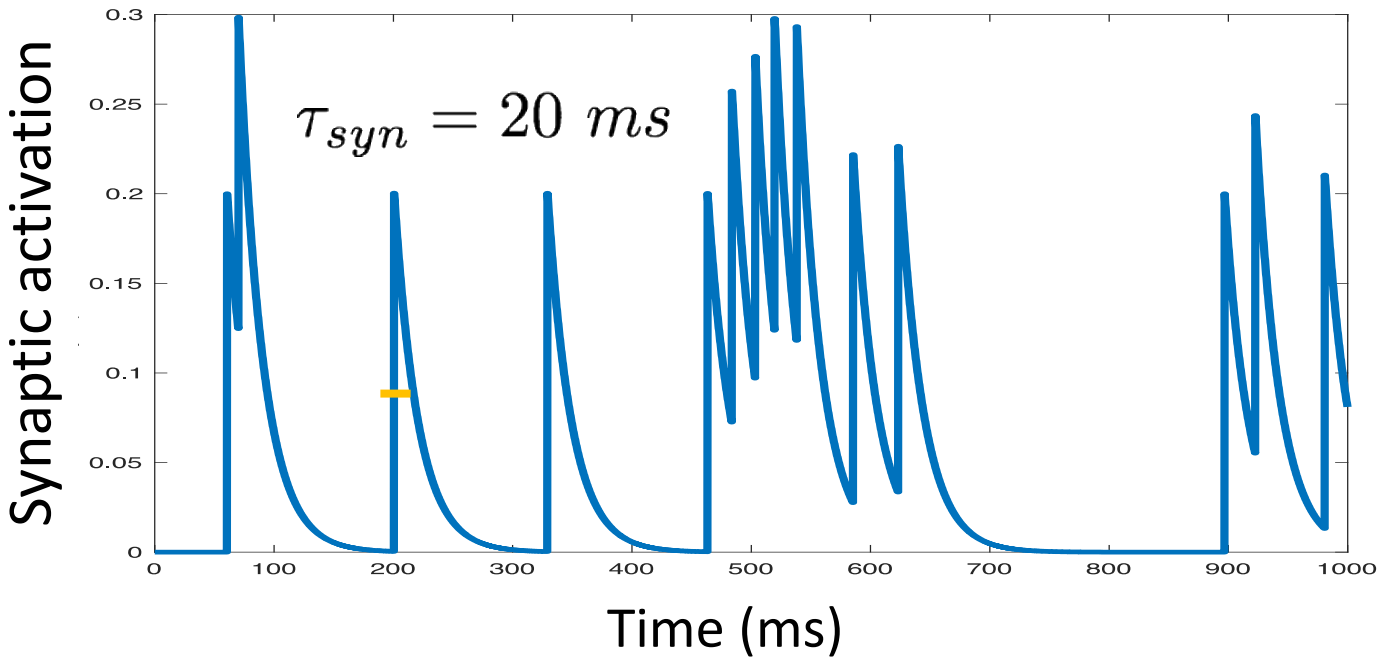

Conductance-based Model

The efficacy of synaptic input depends on postsynaptic neuron voltage

$$C_m \frac{dV}{dt} = -g_m(V(t) - V_m) + I_{app} + I_{spk}(V(t), t)$$

Input to neuron i

Output of neuron j

$$I_{i,app} = \sum_j W_{ij} s_j(t)$$

Current-based model

OR:

$$I_{i,app} = \sum_{j \in E} W_{ij} s_j(t)(V_i(t) - V_E) + \sum_{j \in I} W_{ij} s_j(t)(V_i(t) - V_I)$$

The Hodgkin-Huxley (H-H) theory of the action potential

- The **Hodgkin-Huxley (H-H) theory** of the action potential, formulated 50 years ago, remains one of the great success stories in biology, and ranks among the most significant conceptual breakthroughs in neuroscience.
- Together with the **artificial neural networks of McCulloch and Pitts**, the **quantal theory of Katz**, and the **cable theory of Rall**, all developed at around the same time, the H-H theory provided the foundation for modern computational neuroscience

The History of the H-H Theory

1. Cole and Curtis demonstrated that the action potential is associated with a large increase in membrane conductance
2. Hodgkin and Huxley made the **first intracellular recording of an action potential**
3. Hodgkin and Katz explained the **overshooting action potential** by showing that it results from an increase in sodium permeability (validating the neglected work of Overton)
4. Hodgkin, Huxley and Katz (following Cole and Marmont) developed a **voltage-clamp circuit to enable quantitative measurement of ionic currents from squid axon**
5. Hodgkin and Huxley then showed that step depolarizations of the squid axon trigger an inward current followed by an outward current.
6. Using ionic substitution, they demonstrated that this net current could be separated into **two distinct components**, a ***rapid inward current*** **carried by Na+ ions**, and **a *more slowly activating outward* current carried by K+ ions.**

From experiments using ingenious voltage-clamp protocols, they concluded that these ***two currents result from independent permeability mechanisms*** for Na+ and K+ with conductances changing as a function of time and membrane potential.

H-H model

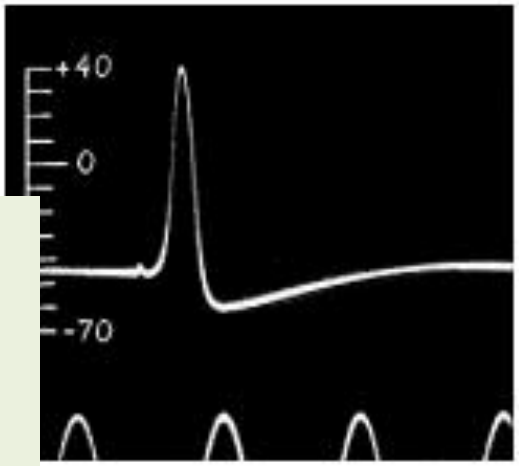

The first intracellular recording of an action potential, from squid axon. Time calibration, 2 ms.

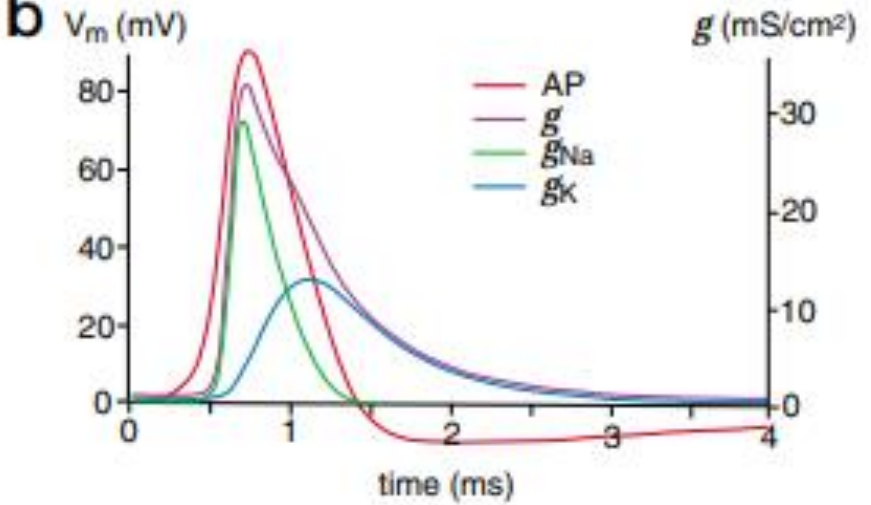

Separation of **ionic conductances** underlying the action potential (AP) in the **H-H model**.

Total ionic current: sum of separate **Na+, K+ & leak currents**

$$I = \bar{g}_{Na} m^3 h (V - E_{Na}) + \bar{g}_K n^4 (V - E_K) + \bar{g}_{leak} (V - E_{leak})$$

where separate equations for the gating variables ***m*** and ***h*** (for activation and inactivation of $\mathbf{g_{Na}}$) or ***n*** (for activation of $\mathbf{g_K}$) describe all the smoothly varying voltage and time dependence of the kinetics. Thus, the **H-H model links the *microscopic* level of ion channels to the *macroscopic* level of currents and action potentials.**

Integrate and fire models are useful approximations of the real neuronal dynamics

Hodgkin-Huxley (H-H)

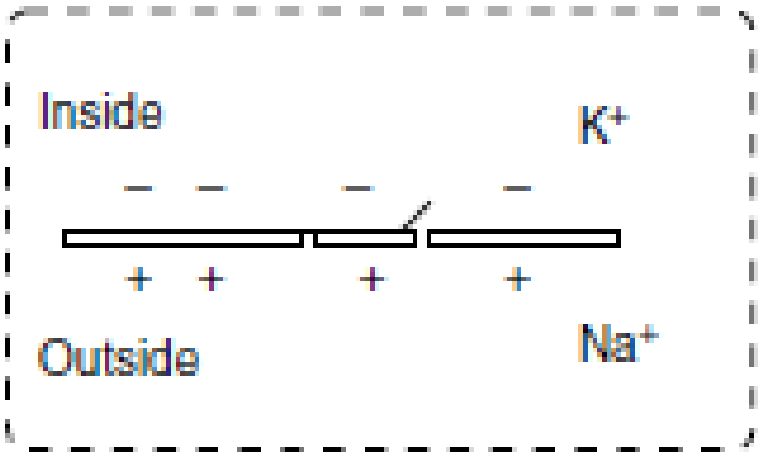

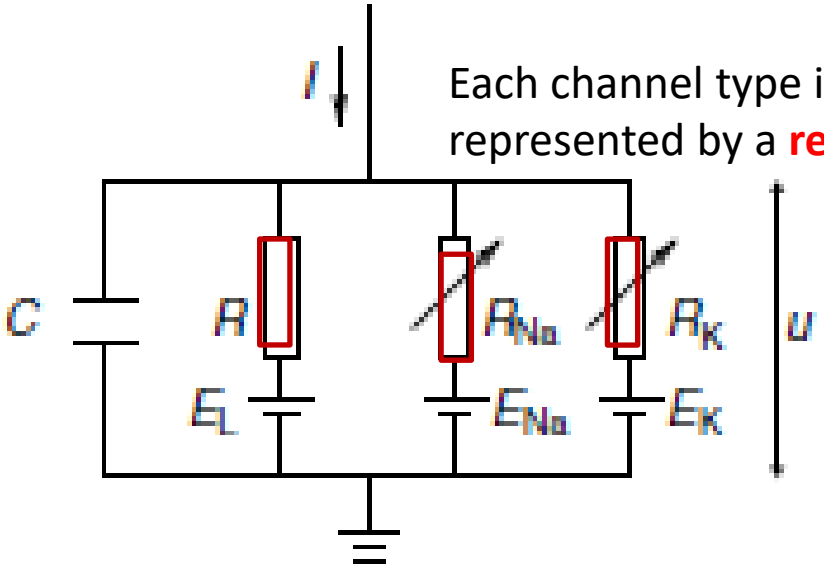

Each channel type is represented by a **resistor**

The value of the resistance is not fixed but changes depending on whether the ion channel is open or closed. Because of **active ion transport through the cell membrane**, the ion concentration inside the cell is different from that in the extracellular liquid.

Nernst potential generated by the **difference in ion concentration**

- The unspecific channel has a ***leak resistance R***, the **sodium channel a resistance R_{Na}** and the **potassium channel a resistance R_K**.
- **Separate batteries** for sodium, potassium &the unspecific third channel, with battery voltages E_{Na}, E_K and E_L, respectively (since the Nerst potential differs depending the ion type)

Leakage channel: described by a **voltage-independent conductance**

$g_L = 1/R$

u: total voltage across the cell membrane

E_L **: voltage of the battery**

Leak current: $I_L = g_L(u - E_L)$

If all channels are open, they transmit currents with a maximum conductance g_{Na} or g_K, respectively. Normally, however, some of the channels are blocked.

The breakthrough of Hodgkin and Huxley was that they succeeded in measuring how the **effective *resistance of a channel* changes as a function of *time and voltage*.**

Nernst potential

From thermodynamics: the **probability of a molecule taking a *state of energy E*** is proportional to the Boltzmann factor **$P(E) \sim \exp(-E / kT)$,** where ***k:* Boltzmann constant**, ***T* temperature**.

Consider **positive ions with charge *q*** in a static electrical field.

Their energy at **location *x:*** $E(x) = q\,u(x)$, where ***u*(*x*) the potential at *x*.**

Probability of finding an ion around *x* is proportional to $\exp[-q\,u(x) / k\,T]$

Since the number of ions is huge, we may interpret the probability as an **ion density**.
For ions with positive charge $q>0$, the ion density is higher in regions with low potential u.

***n*(*x*):** ion density at point *x*

A **difference in electrical potential** $\Delta u = u(x1) - u(x2)$ generates a ***difference in ion density***;
Since this is a statement about an equilibrium state, the reverse must also be true. **A difference in ion density generates a difference Δ*u* in the electrical potential**

Consider two regions of ions with concentration ***n*1 and *n*2**, respe

At equilibrium, the **concentration difference generates a voltage**

$$\Delta u = \frac{kT}{q} \ln \frac{n_2}{n_1}$$

Ion Concentration of Potassium

- The ion concentration of potassium is ***higher inside the cell*** (≈ 140 mM) **than in the extracellular liquid** (≈ 5mM)
- **Potassium ions have a single *positive charge q*** = 1.6×10−19 C
- Application of the Nernst formula yields $\boldsymbol{E_K \approx -83mV}$ at room temperature (with the Boltzmann constant k = 1.4×10−23 J/K
- The reversal potential for **K+ ions** is therefore negative

Sodium ions (Na+) & Reversal Potential

At equilibrium:

the difference in concentration causes a **Nernst potential** E_{Na} **~ +67mV.**

the interior of the cell has a positive potential w.r.t the surround.

The interior of the cell and the surrounding liquid are in contact through ion channels where Na+ ions can pass from one side of the membrane to the other.

If $\Delta u < E_{Na}$, **more Na+ ions *flow into* the cell** so as to decrease the concentration difference.

If $\Delta u > E_{Na}$, ions would ***flow out*** the cell.

Thus, the direction of the current is reversed when the voltage Δu passes E_{Na}.

For this reason, E_{Na} is called the **reversal potential**.

Ionic Currents

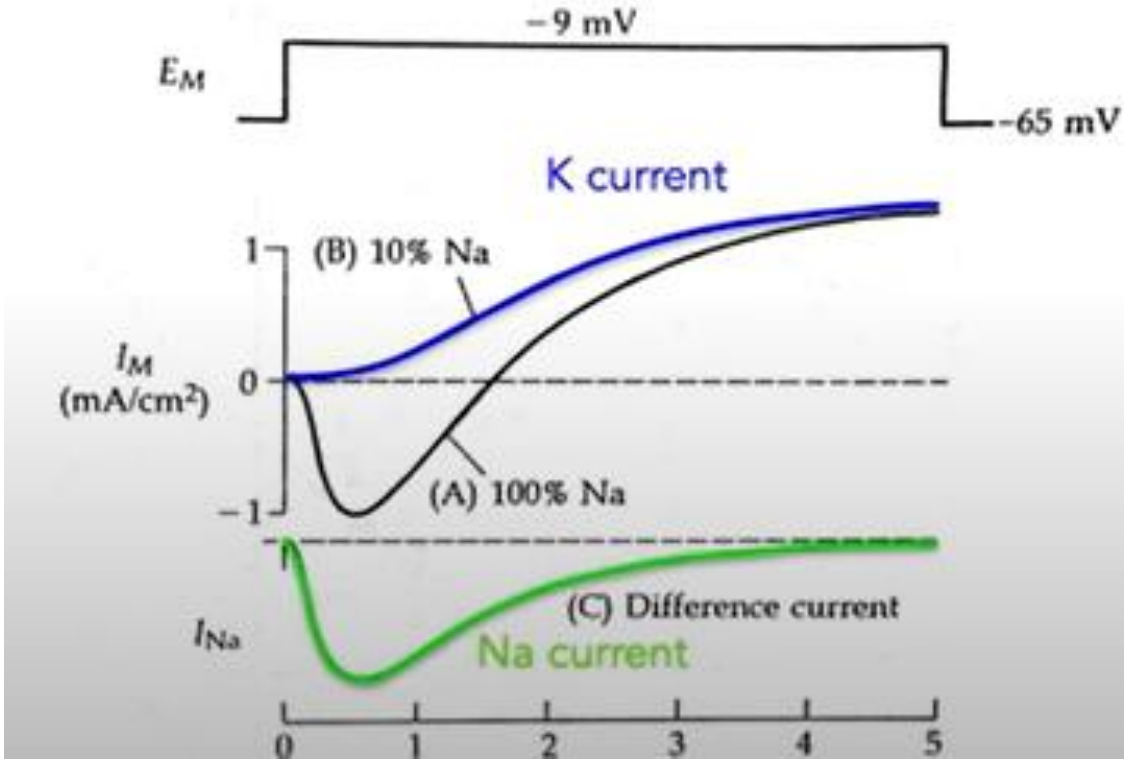

Membrane ionic current

$$I_L = \frac{V}{R_L} = G_L V$$

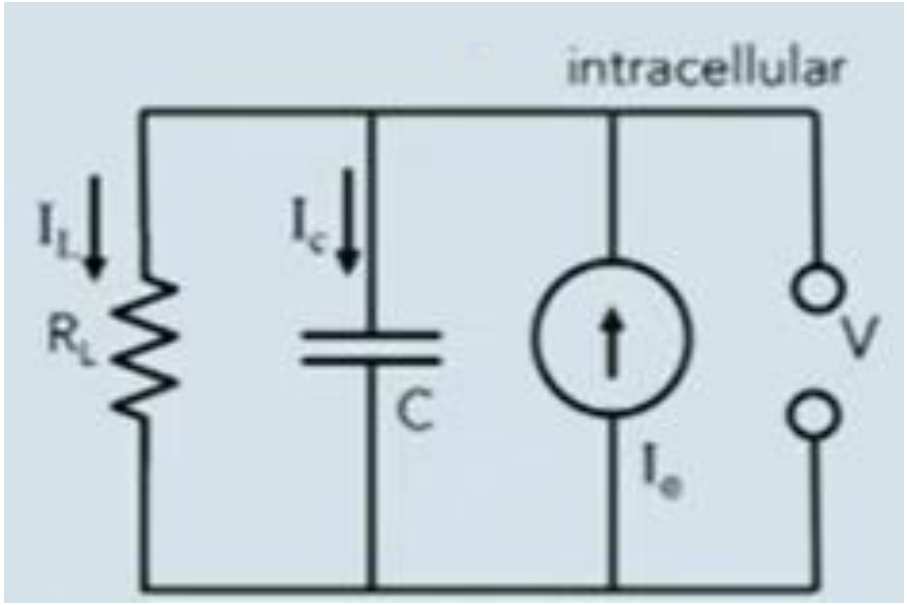

Our equation for our model becomes:

$$I_L + I_c - I_e = 0$$

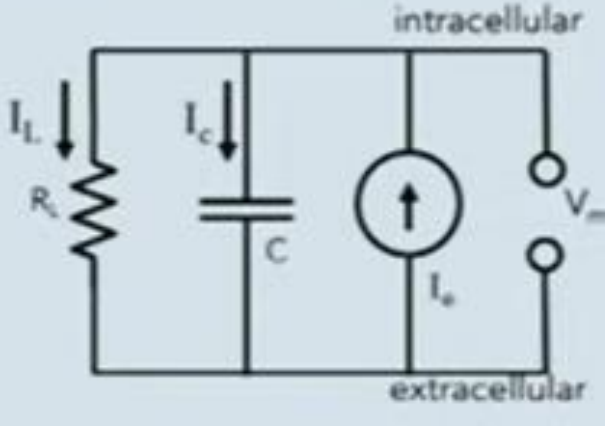

Our equation for our model becomes:

$$I_L + I_c - I_e = 0$$

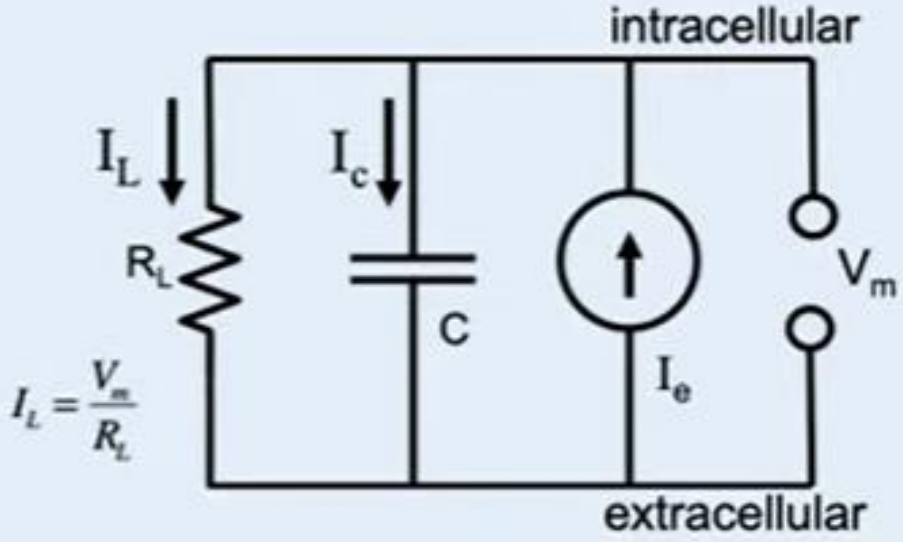

$$I_L + C\frac{dV}{dt} = I_e$$

$$\frac{V}{R_L} + C\frac{dV}{dt} = I_e$$

$$V + R_L C\frac{dV}{dt} = R_L I_e$$

Rewrite as:

$$V + \tau\frac{dV}{dt} = V_\infty \qquad \text{where} \quad \tau = R_L C \qquad V_\infty = R_L I_e$$

Simplified Hodgkin–Huxley Model

I_i: current from the **i-th pre-synaptic neuron**
V: membrane **voltage**

- Relationship between the ***flow* of ionic currents** across the neuronal cell membrane and the membrane voltage of the cell
- Set of **nonlinear differential equations** describing the behaviour of ion channels that permeate the cell membrane

Cell membrane of *capacity* C_m

$$C_{\mathrm{m}} \frac{dV(t)}{dt} = -\sum_i I_i(t, V).$$

Time derivative of the law of capacitance, Q = CV where the **change of the total charge** must be explained as the **sum over the currents**.

Hodgkin-Huxley model of action potential generation

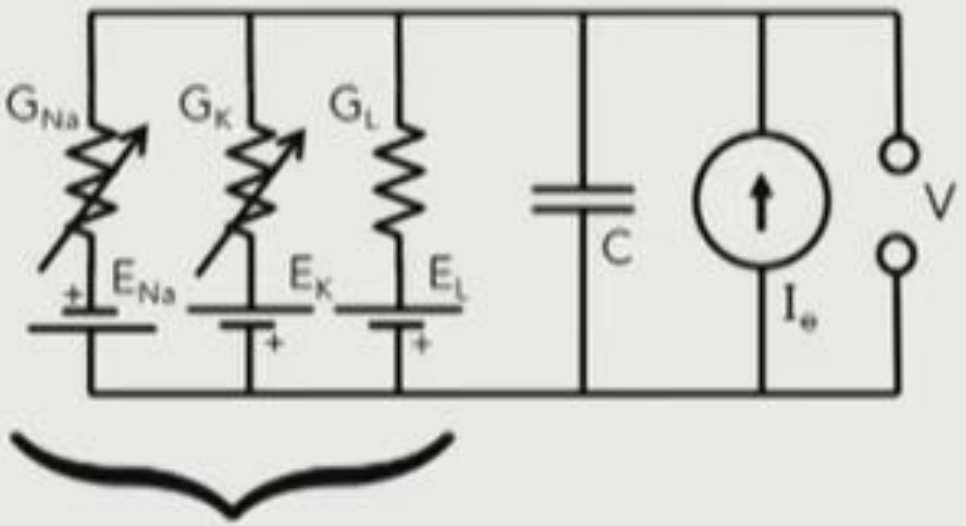

$$I_m = I_{Na} + I_K + I_L \qquad I_m(t) + C\frac{dV(t)}{dt} = I_e(t)$$

$$I_{Na} = G_{Na}(V,t)(V - E_{Na}) \qquad I_K = G_K(V,t)(V - E_K) \qquad I_L = G_L(V - E_L)$$

The algorithm that the neuron uses to generate a spike

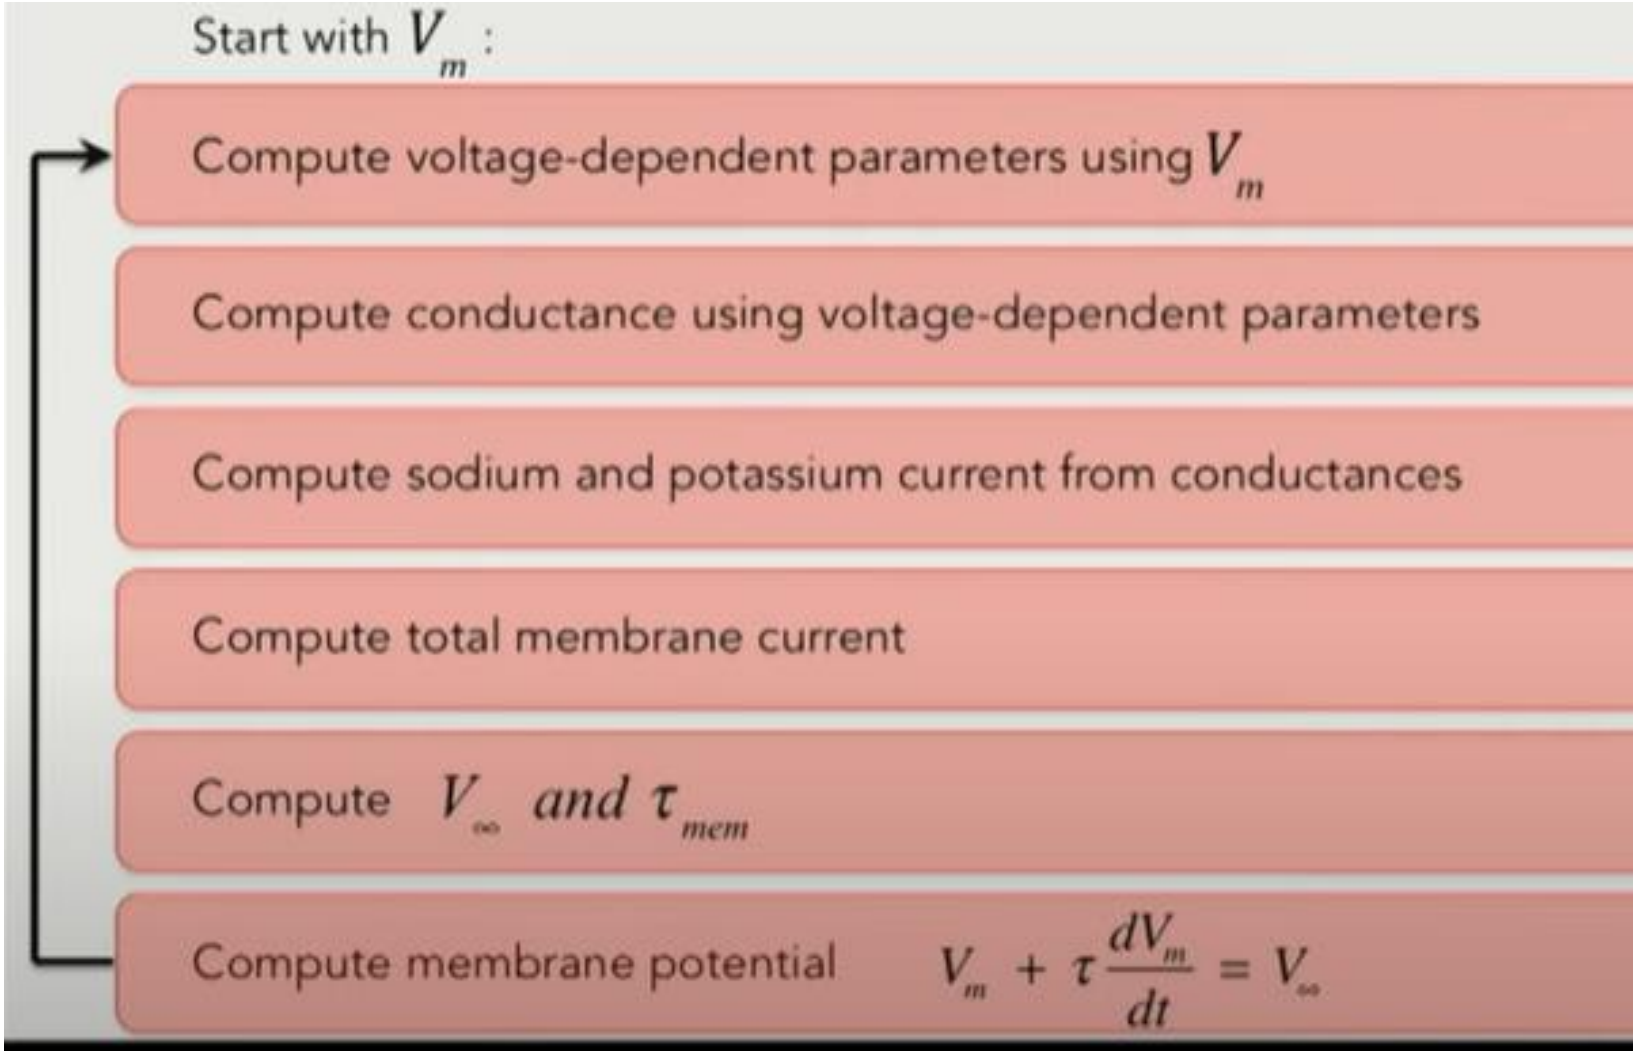

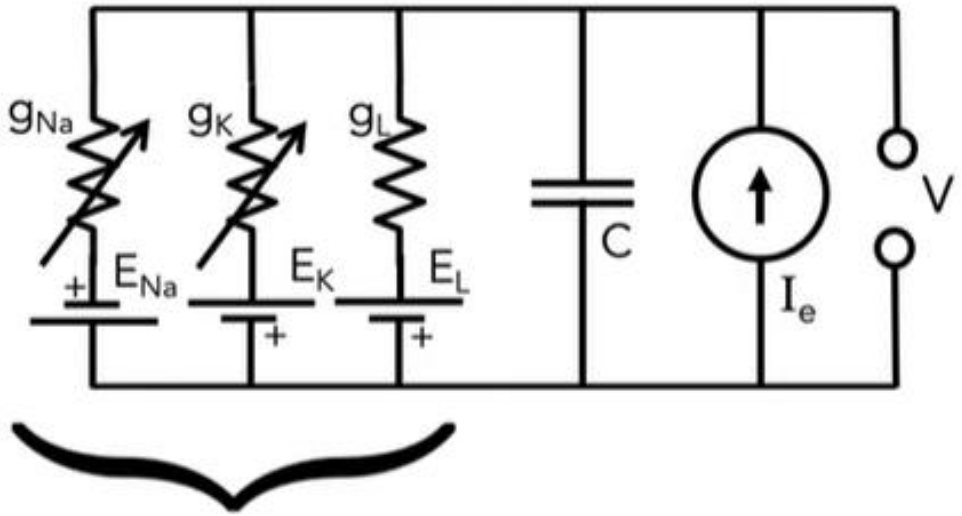

$$I_m = I_{Na} + I_K + I_L$$

$$I_X = G_X \cdot (V - E_X)$$

$$I_{Na} = G_{Na}(V,t)(V - E_{Na})$$

$$I_K = G_K(V,t)(V - E_K)$$

$$I_L = G_L(V - E_L)$$

The sodium conductance is time-dependent and voltage-dependent

The potassium conductance is time-dependent and voltage-dependent

The leak conductance is neither time-dependent nor voltage-dependent

Electrical Circuit Idioms for Modeling Neurons

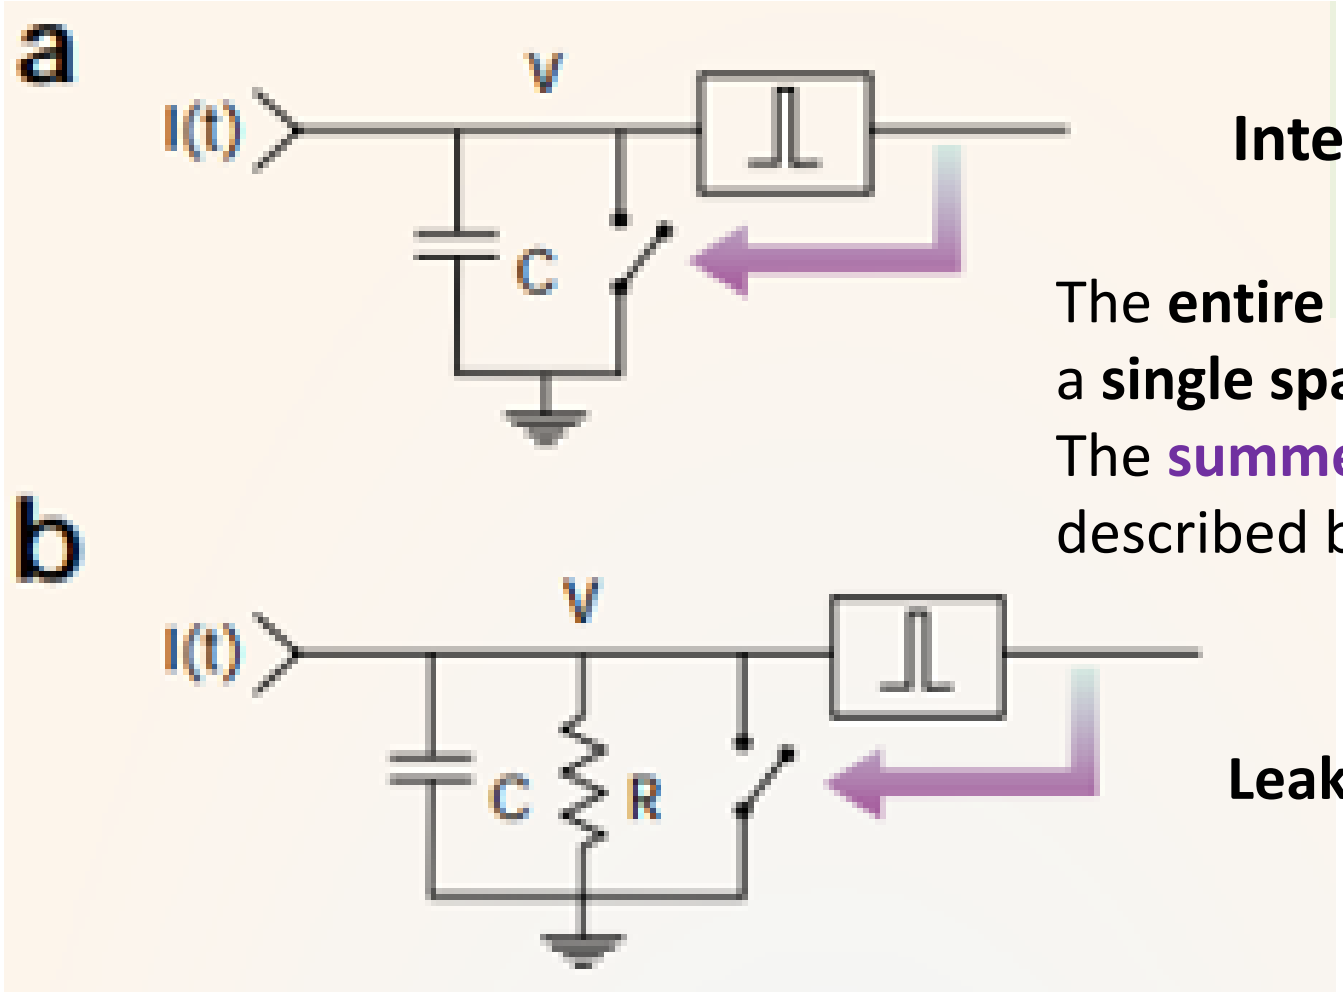

Integrate-and-fire

The **entire neuron** is reduced to a **single spatial compartment**. The **summed synaptic** input is described by a ***net current I(t).***

Leaky integrate-and-fire

The **state of the neuron** is given by the voltage across a capacitance, with each **synaptic input** adding to or subtracting from the charge accumulating across the membrane

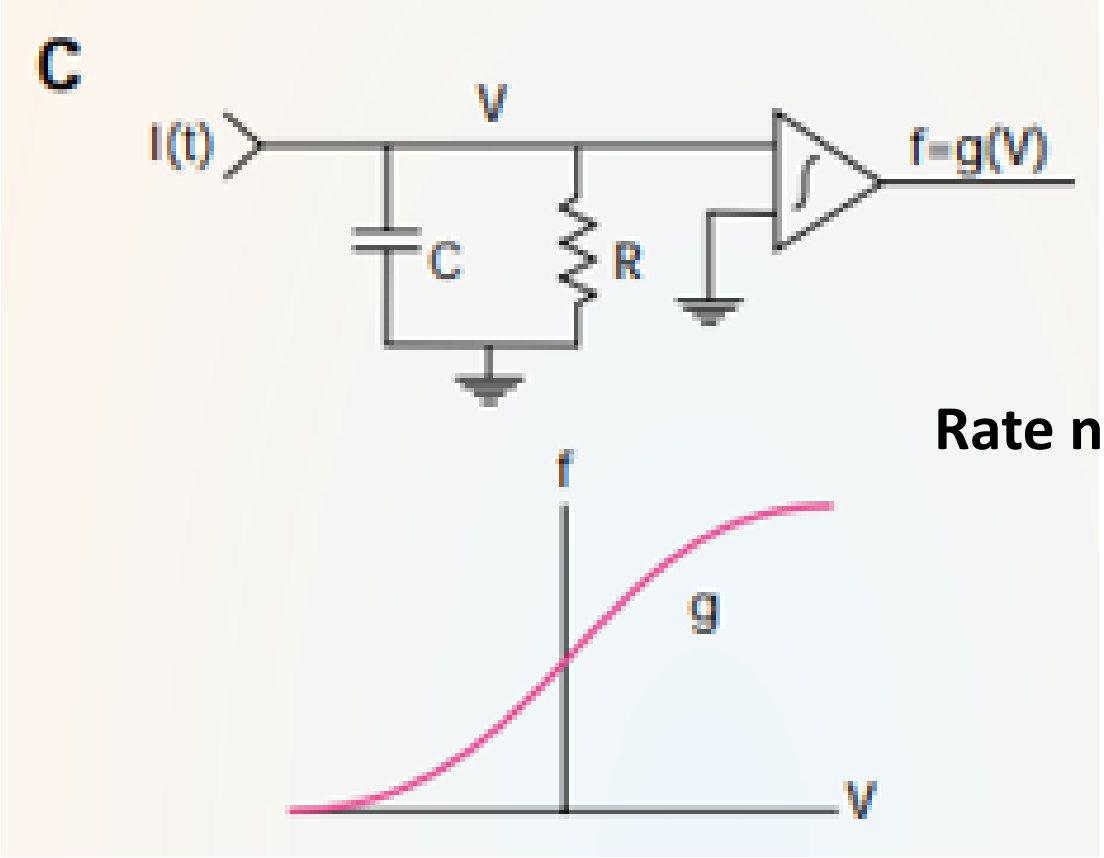

Rate neuron

In rate neuron: these discrete pulses are replaced by a continuous output rate.

The **monotonically increasing relationship** between **V & output rate**

f = g(V) can be thought of as the discharge function of a population of spiking cells.

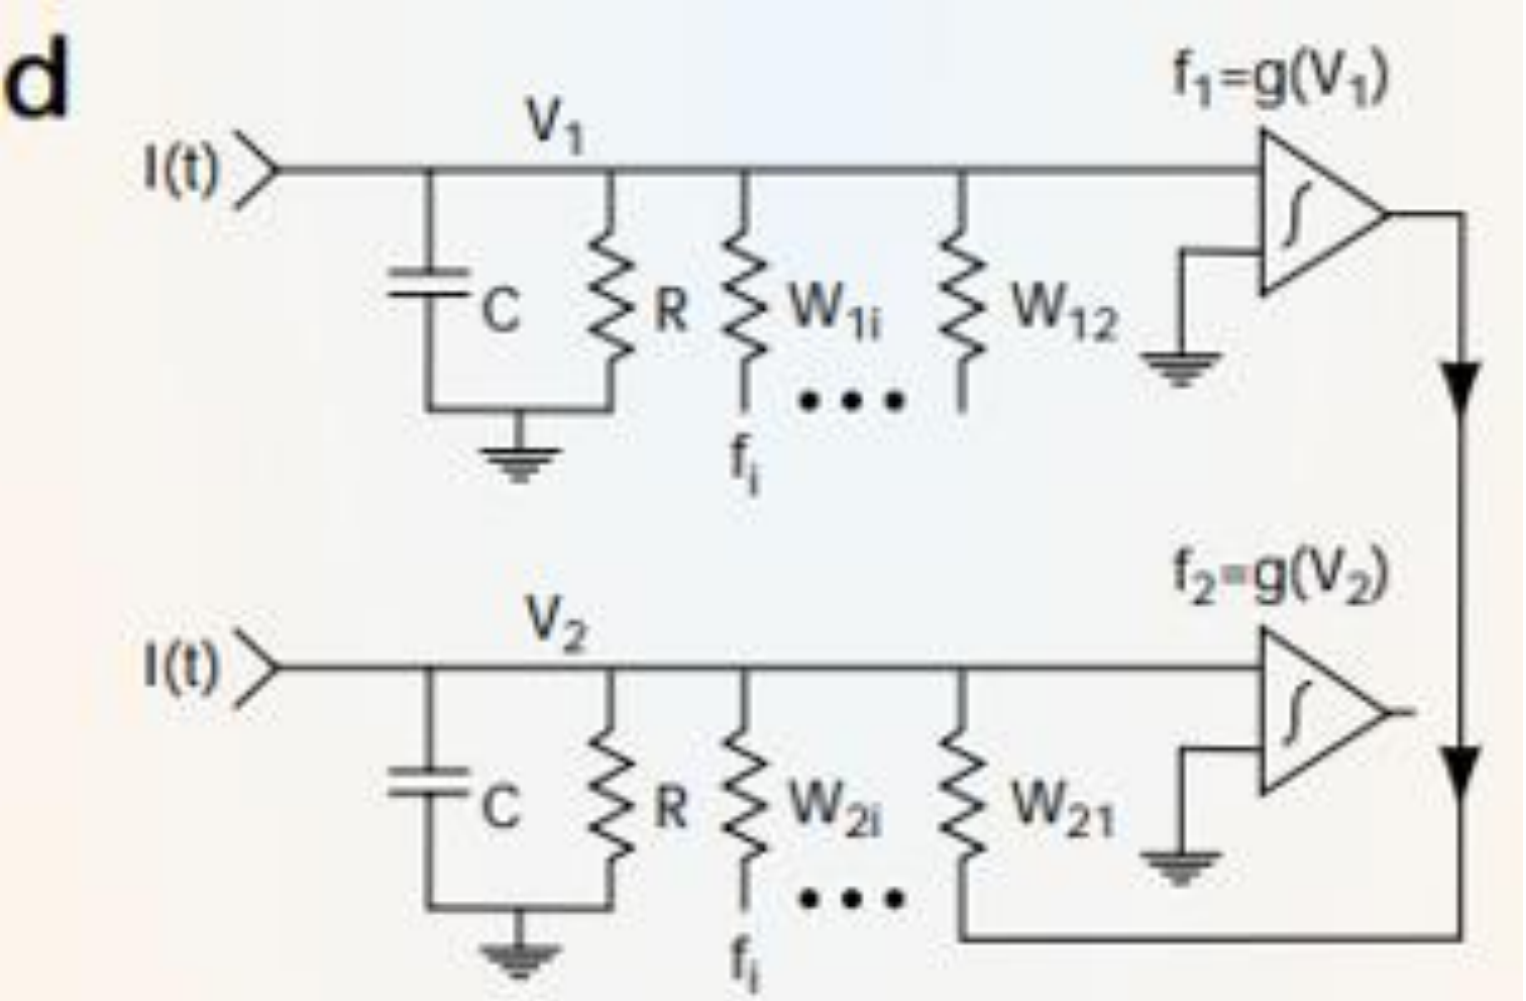

The summed synaptic input is described by a net current **I(t)**.

d) In most neural networks, interactions within neurons are linear. The necessary nonlinearity is provided by the ***sigmoidal*** g(V) function. Here, the output of **neuron 1** is unidirectionally connected to **neuron 2** with **synaptic weight w_{21}**.

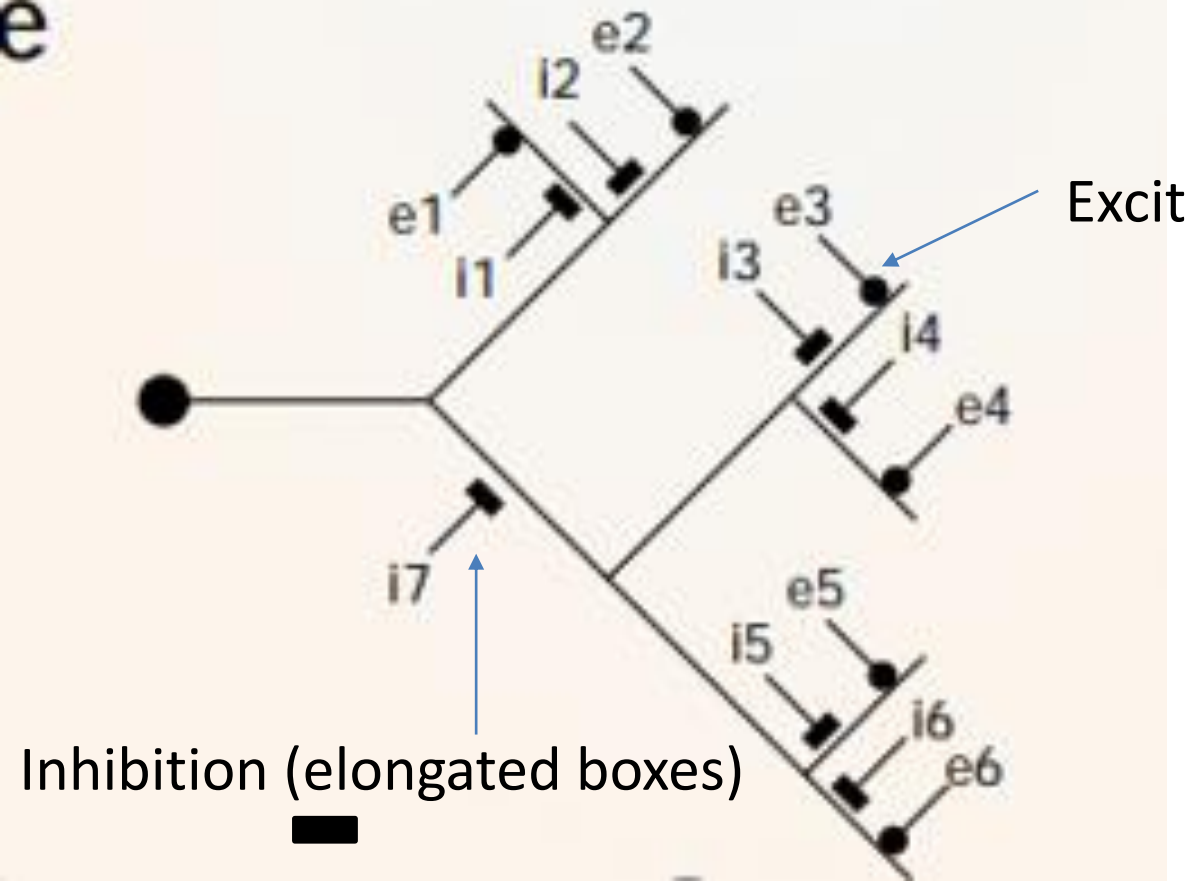

Non-linear saturating interactions can be mediated in a **passive dendritic tree by synapses** that increase the postsynaptic conductance. The interaction between excitation and inhibition of the shunting type is of the **AND-NOT** type and is specific in space and in time, e.g., the **inhibitory synapse i_7 vetos excitation e_3 or e_6** but has **only a negligible effect on e_1**.

Electrical Circuit Idioms for Modeling Neurons

The **entire neuron** is reduced to a **single spatial compartment**.
The **summed synaptic** input is described by a **net current I(t).**

Integrate-and-fire

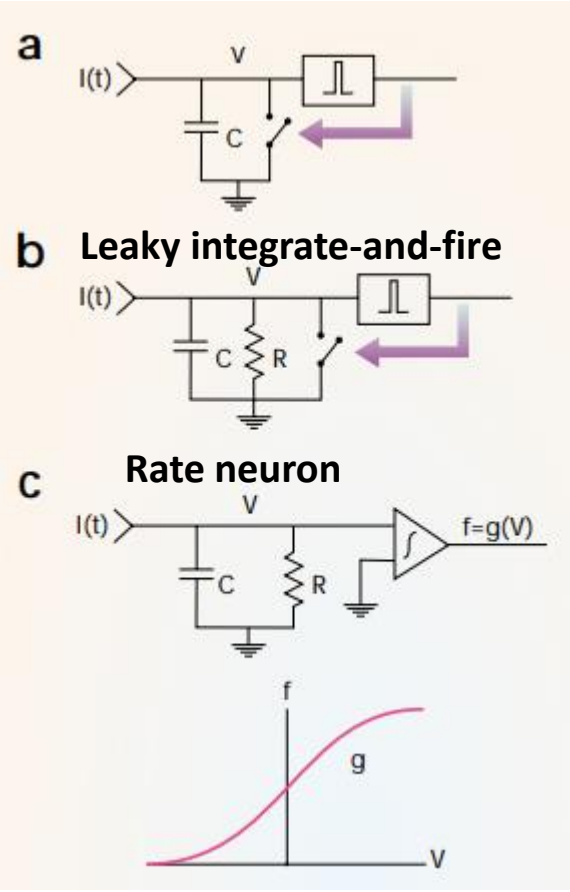

(a) If **the voltage V > a fixed threshold**, a unit pulse is generated, and all charge on the capacitance is removed by resetting V to zero (solid arrow).

(b) **Leaky integrate-and-fire model**: **charge leaks away** with a time constant (given by the product of the capacitance C and the resistance R) is a series of asynchronous spikes.

(c**) In rate neuron**: these discrete pulses are replaced by a continuous output rate.
The monotonically increasing relationship between V & output rate f = g(V) can be thought of as the discharge function of a population of spiking cells.

The **state of the neuron** is given by the voltage across a capacitance, with each **synaptic input** adding to or subtracting from the charge accumulating across the membrane

Neural Models: Spiking vs Firing-rate Models

- **Spiking models** involve **dynamics over time scales** ranging from *channel openings* (**less than 1 ms**) to ***collective network* processes** that may be **several orders of magnitude slower.**
- **Firing-rate models** avoid the **short-time scale dynamics** required to simulate action potentials and thus are much easier to simulate on computers.
- Firing-rate models also allow us to present analytic calculations of some aspects of network dynamics that could not be treated in the case of spiking neurons.
- Spiking models tend to have more free parameters than firing-rate models, and setting these appropriately can be difficult.

Modeling the Interaction at Different Scales

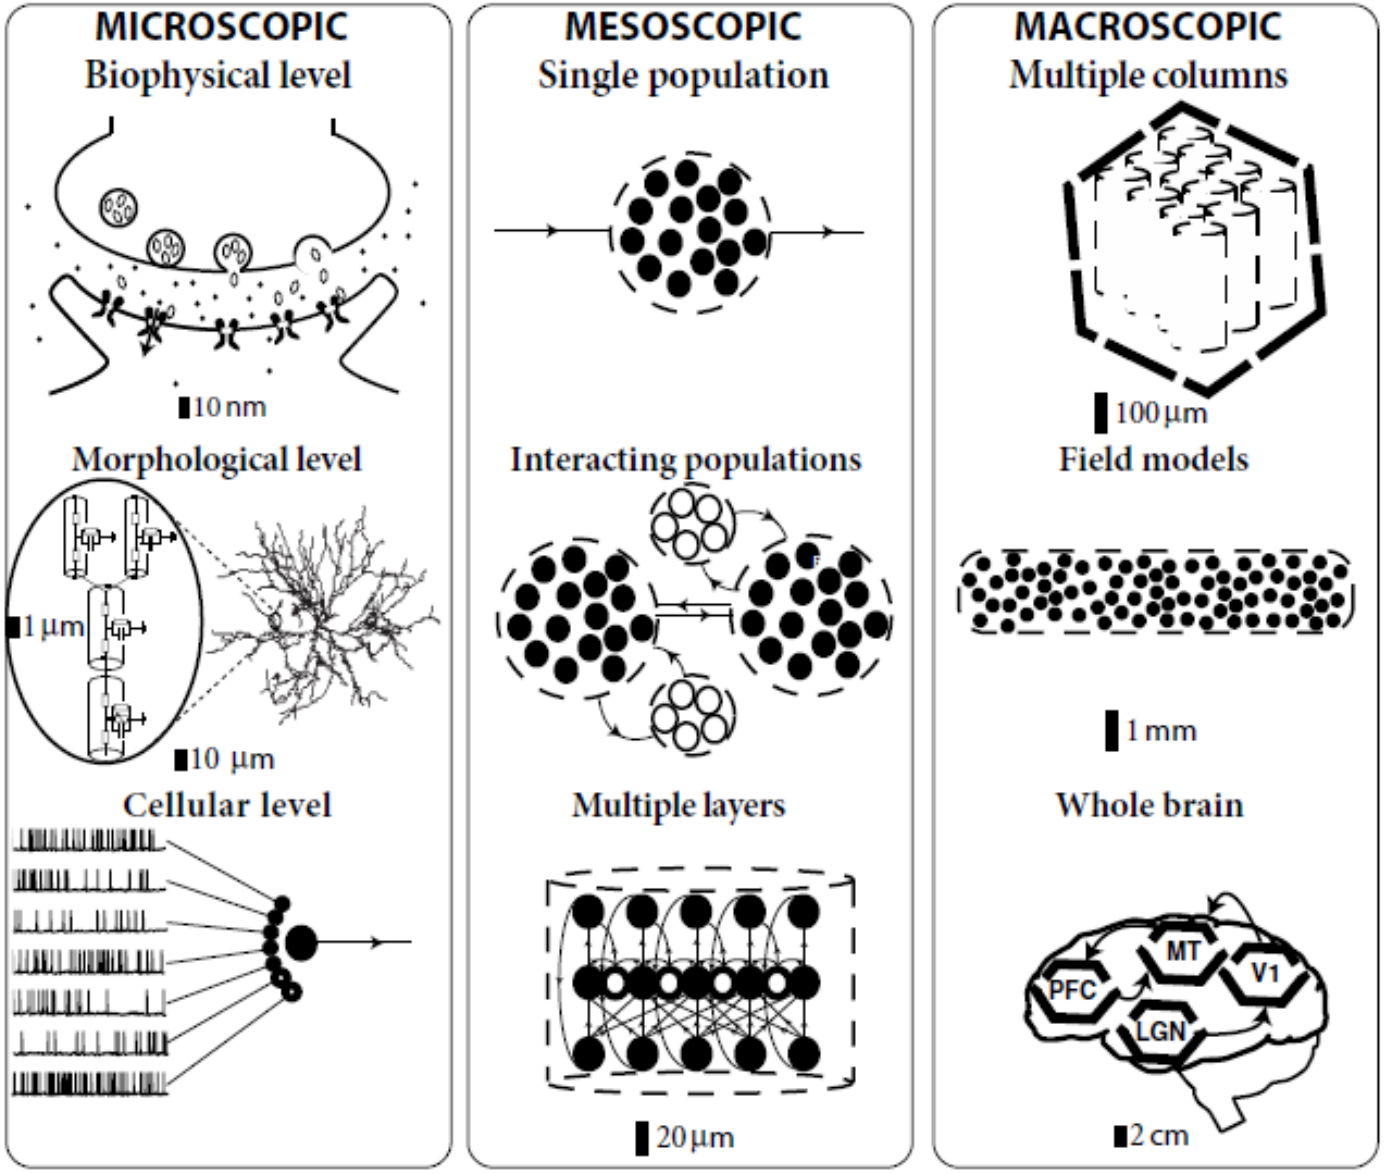

If the experimenter moves the electrode vertically down from the cortical surface to deeper layers, the **location of the receptive field and its preferred orientation** does ***not*** change substantially.

Arrows on the left indicate inter-layer connection probabilities between excitatory neurons

Arrows on the right show intra-layer connection probabilities between excitatory neurons

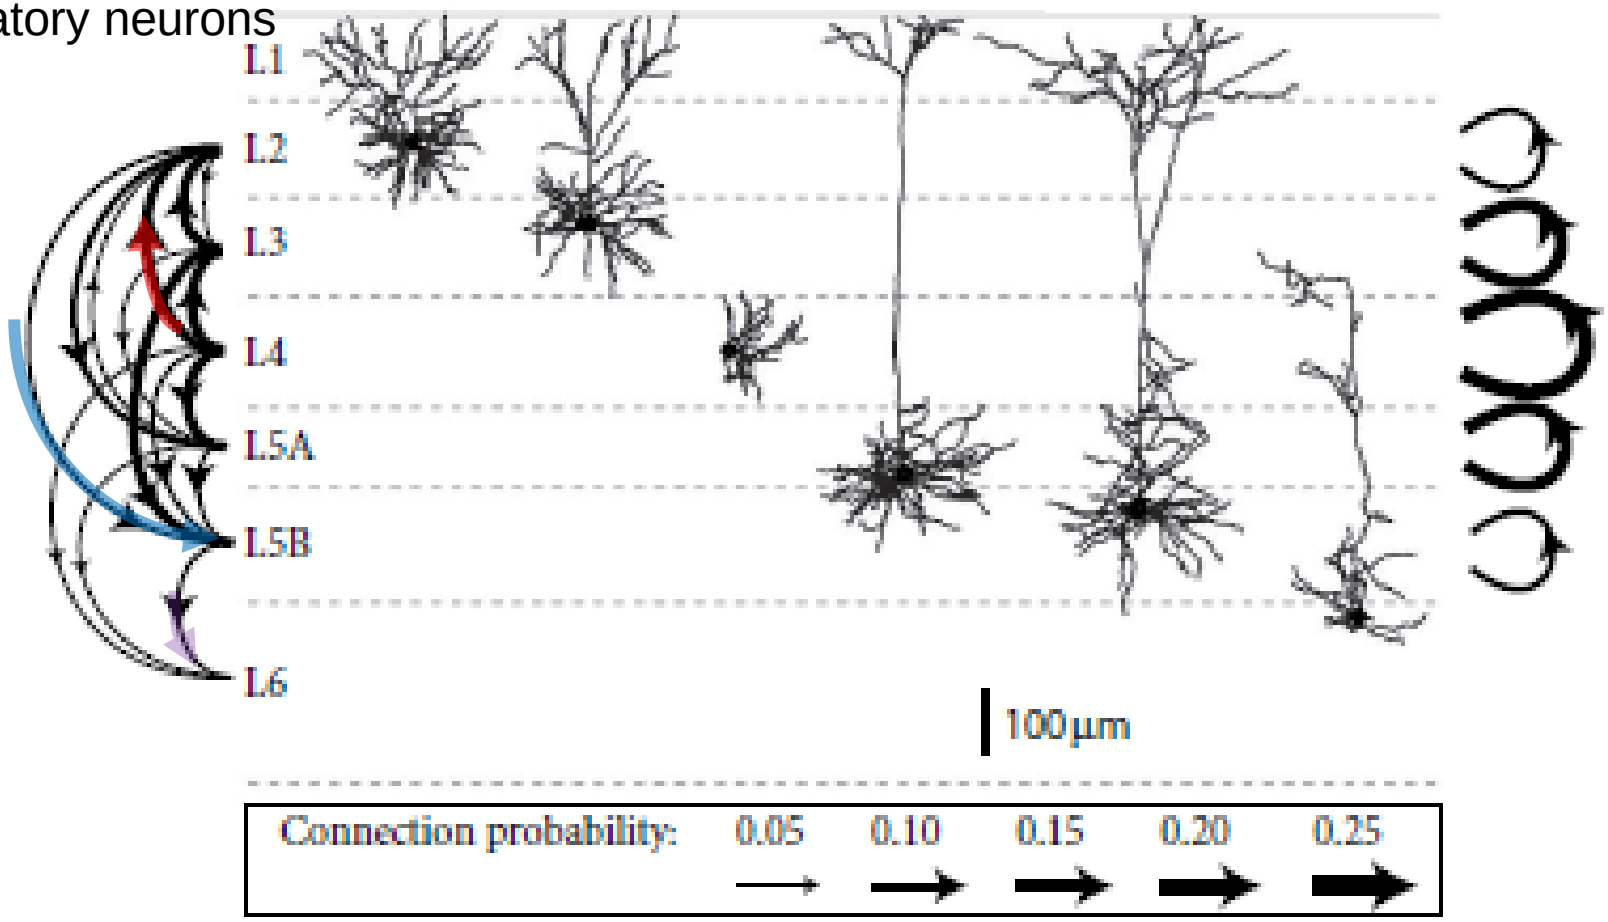

Connectivity patterns inside one column. Examples of shapes of excitatory neurons in different layers.
Data from a barrel column in the somatosensory cortex of the mouse. After Lefort *et al*. (2009)

- The cortex is a rather thin sheet of cells.
- Cortical columns extend vertically across the sheet.
- The connection probability within a column depends on the layer where pre- and postsynaptic neurons are located.
- In addition to this vertical connectivity, neurons make many horizontal connections to neurons in other, cortical columns in the same, but also in other, areas of the brain.

- Within the **same brain area, the probability of making a connection** is often modeled as **distance dependent**.

  Note that distance dependence is a rather coarse feature, because the actual connectivity depends also on the function of the pre- and postsynaptic cell.

  In the primary visual area, pyramidal neurons with a preferred orientation for horizontal bars are more likely to make connections to other columns with a similar preferred orientation (Angelucci and Bressloff, 2006).

Interconnections in Cortex - Columns

- In the neocortex, which forms the convoluted outer surface of the human brain, neurons lie in **six vertical layers** highly coupled within cylindrical columns.
- Such ***columns*** "act" as basic functional units, cortical columns and stereotypical patterns of connections both **within a column** and **between columns** are repeated across cortex.

Top view onto the surface of the visual cortex.

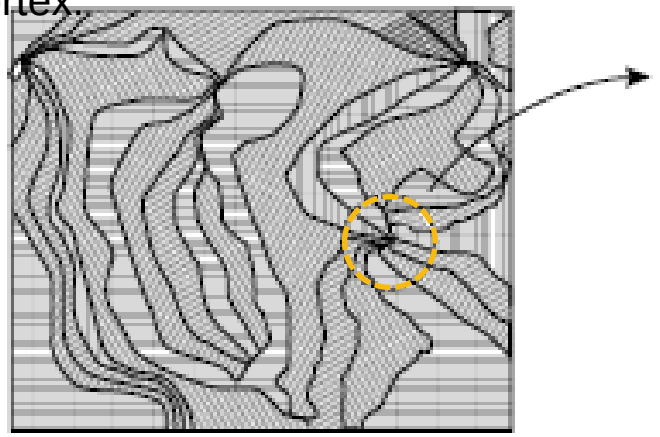

Neurons that are optimally activated by a moving grating with an orientation of, say, 60◦, form bands. The direction of the hash-line texture indicates the preferred orientation. Iso-orientation contour lines converge to form pinwheels. One of the pinwheels is

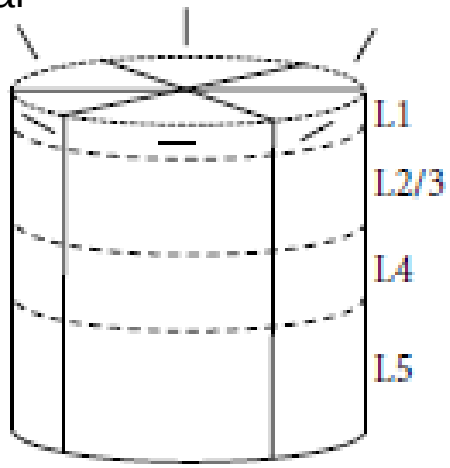

Side view of a pinwheel (dashed circle)
Orientation selectivity is indicated by thick bars. Neurons with the same orientation form vertical columns.
Schematic representation following experimental data shown in Bressloff and Cowan (2002).

Interconnections in Cortex (cont'd)

- **Feed-forward connections** bring input to a **given region from another region located at an *earlier stage*** along a particular processing pathway
- **Recurrent connections** interconnect neurons ***within a particular region*****,** considered to be at the same stage along the processing pathway, e.g., connections within a cortical column as well as connections between both nearby and distant cortical columns **within a region**
- **Top-down connections** carry signals back from areas located at ***later stages***

Neurons within a given region send top-down projections back to the areas from which they receive feed-forward input, and receive top-down input from the areas to which they project feedforward output

Pathways & Assemblies

- The sensory pathways create **abstract representations**
- Top-down attention can modify these representations
- Higher areas selectively represent task-relevant information
- Information often is coded ***sparsely*** and ***dynamically***

Donald Hebb (1949) introduced the notion of ***neuronal assemblies***, i.e., groups of cells which get activated together so as to represent a mental concept such as the preparation of a movement of the right arm toward the left.
Assembly: group of neurons, **distributed across one or several brain areas** (thus ***not*** necessarily a local group); homogeneous population that is activated **whenever the corresponding mental concept is evoked**.
Important: the assignment of a neuron to a population is not fixed but can depend on the stimulus.

Prominent Examples of Columns

- In the somatosensory and visual cortex (Mountcastle, 1957; Hubel and Wiesel, 1962)
- Pools of motor neurons (Kandel *et al*., 2000).
- Given the large number of neurons within such a column or pool, it is sensible to describe the mean activity of the neuronal population rather than the spiking of individual neurons

Mean-Field Concept

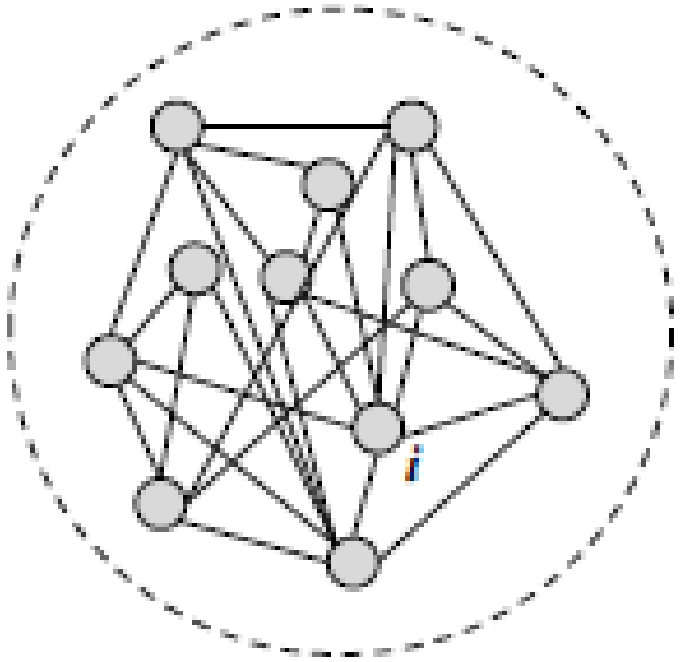

A fully connected population of neurons
(not all connections are shown).

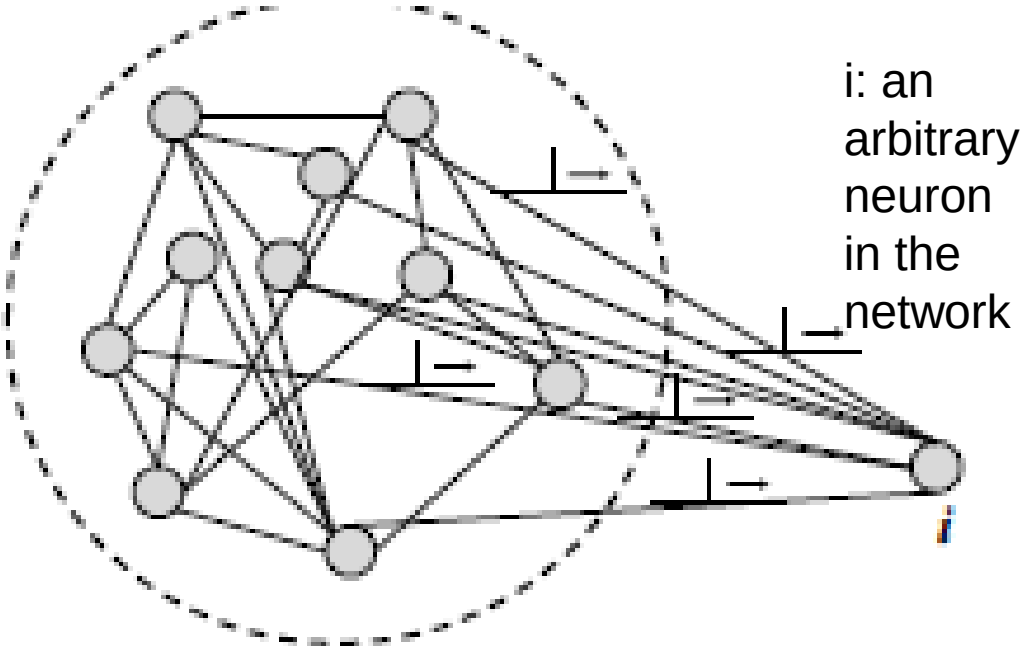

i receives input spikes from the whole population. Hence it is driven by the population activity $A(t)$. The same is true for all other neurons.

Homogeneous Networks

- **all neurons** $1 \le i \le N$ **are identical**
- all neurons receive the **same external input** I_{ext} $I(t)=I_{ext}(t)$
- the **interaction strength** $w_{i,j}$ for the connection between (j, i)

of pre- & post-synaptic neurons is "**statistically uniform**" connections **inside the population** as being either absent or "roughly the same," $w_{i,j} \approx w_0$ (w_0 is a parameter)

$w_0 = 0$: ***independent*** **neurons**

$w_0 > 0$ ($w_0 < 0$): **excitatory (inhibitory) coupling**

each **input spike** generates a postsynaptic current with some course $\alpha(t - t_j^f)$

input current I_i *includes both synaptic coupling and external drive*

$$I_i = \sum_{j=1}^{N} \sum_f w_{ij} \alpha(t - t_j^f) + I^{ext}(t)$$

Heterogeneous Populations

- Strongly heterogeneous population should be split until (nearly) homogeneous groups remain.

e.g., split the population into two populations, one with all neurons with parameters θ_1 & the other with all neurons with parameters θ_2

Stationary firing rate $\nu_i = g_{\theta_i}(I)$

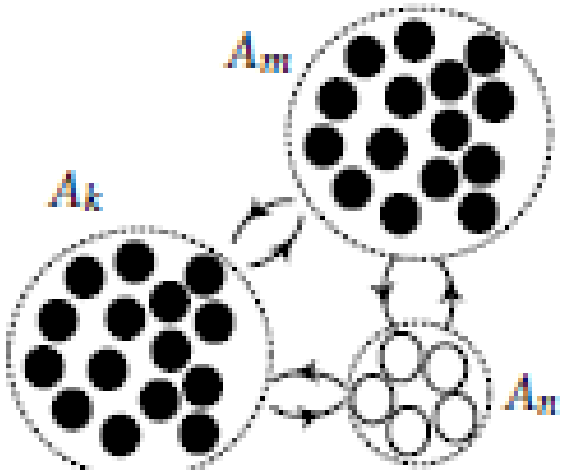

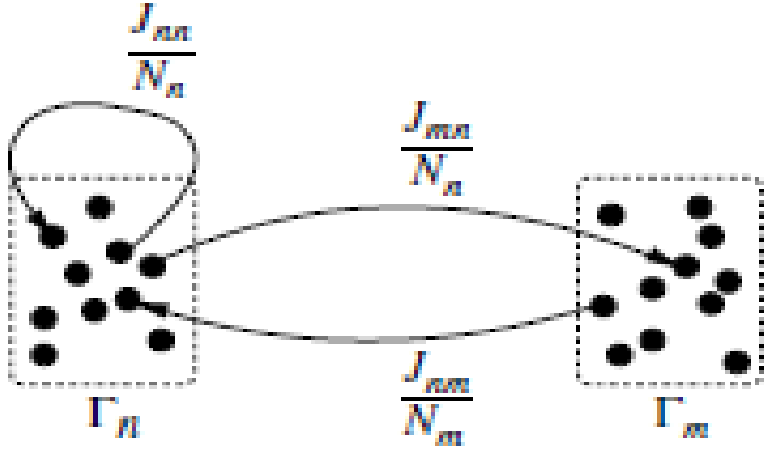

All neurons in group Γ_n are coupled with synaptic efficacy $w_{ij} = J_{nn}/N_n$.

Plasticity

- Synaptic plasticity is the ability of synapses to strengthen or weaken over time in response to increases or decreases in their activity
- Plastic change often results from the alteration of the number of neurotransmitter receptors located on a synapse

as well as changes in how effectively cells respond to those neurotransmitters

Long-term plasticity: from minutes to hours

Short-term plasticity: tens of milliseconds to a few minutes

$[Ca^{2+}]_i$ Concentration of calcium

Synaptic weight of the i-th input axon Wi

$$\frac{dW_i(t)}{dt} = 1/(\tau\,[Ca^{2+}]_i)\;(\mathbf{\Omega}\,([Ca^{2+}]_i) - W_i)$$

Ω: function of the concentration of calcium that depends linearly on the number of receptors on the membrane of the neuron

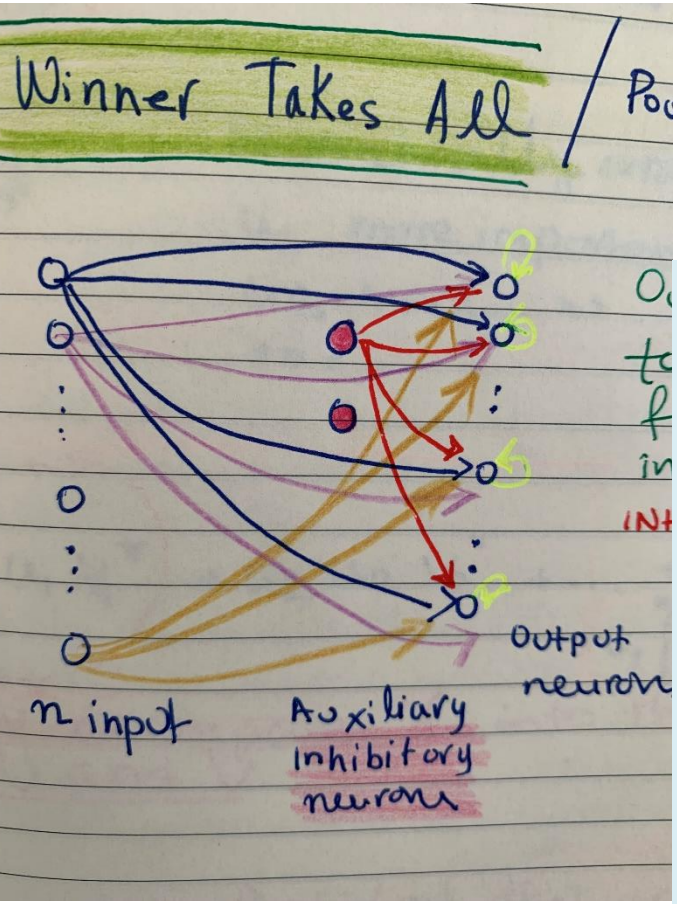

Winner-Takes-All

a powerful computational paradigm

Outputs that fire in response to stimulation from their firing inputs **excite two inhibitors which in turn inhibit *all the outputs***

When more than one outputs fires, both inhibitors get excited

This leads to a high level of inhibition, casing firing outputs to stop firing and drop out of the WTA competition

When exactly one output fires, just one of the inhibitors (known as ***stability inhibitor***) is excited

This inhibitor is responsible for **maintaining a WTA steady state**: once a single output fires at a time state, it becomes the winner of the network

It has a positive feedback self-loop that allows it to keep firing at subsequent times, while **all other outputs do *not* fire** due to inhibition from the stability inhibitor

Winner-Takes-All

A powerful computational paradigm

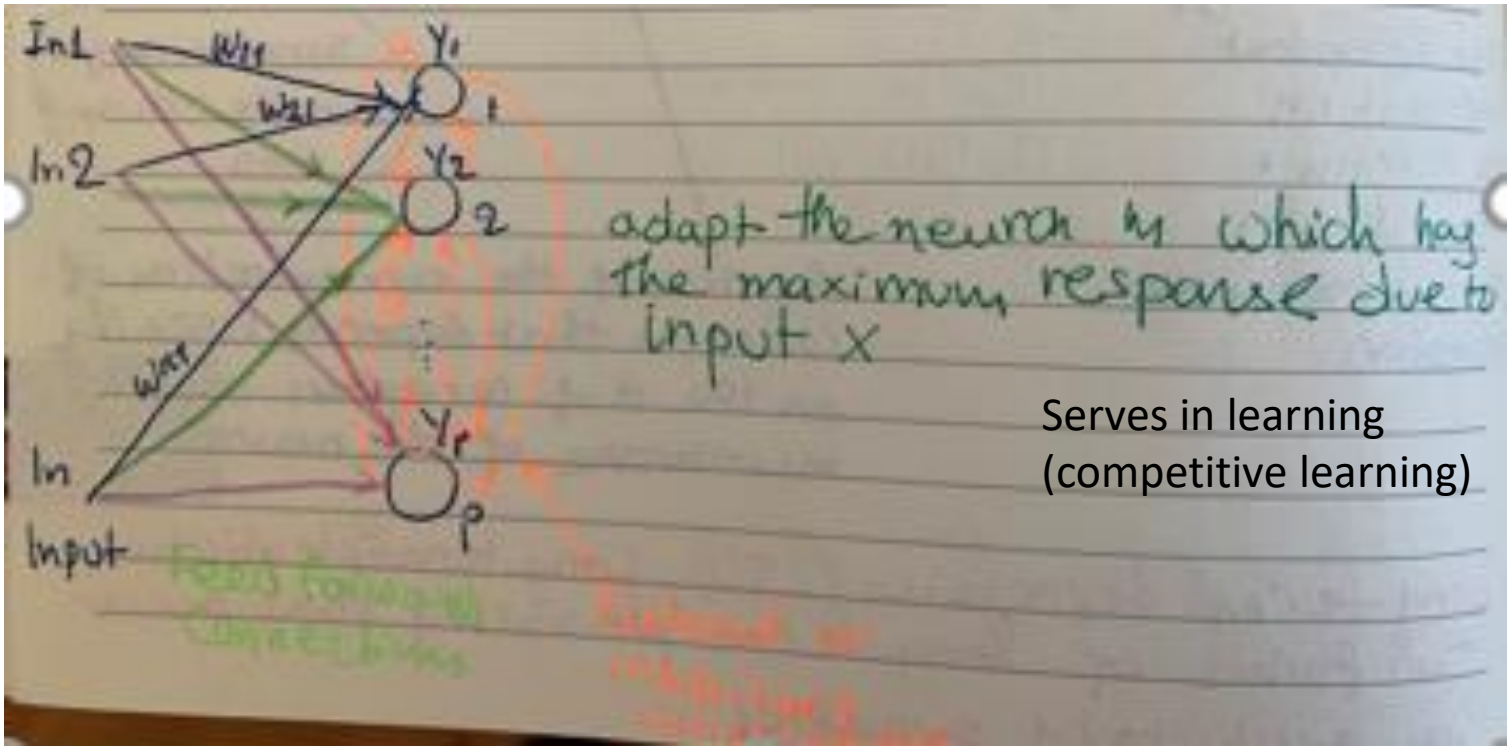

- The output nodes in the network ***mutually inhibit each other***, while simultaneously activating themselves through **reflexive connections**
- After some time, only one node in the output layer will be active, namely the one corresponding to the strongest input
- The **winne**r is the node with the ***largest response***

Fluorescence Imaging

Fluorescence illumination of a single point

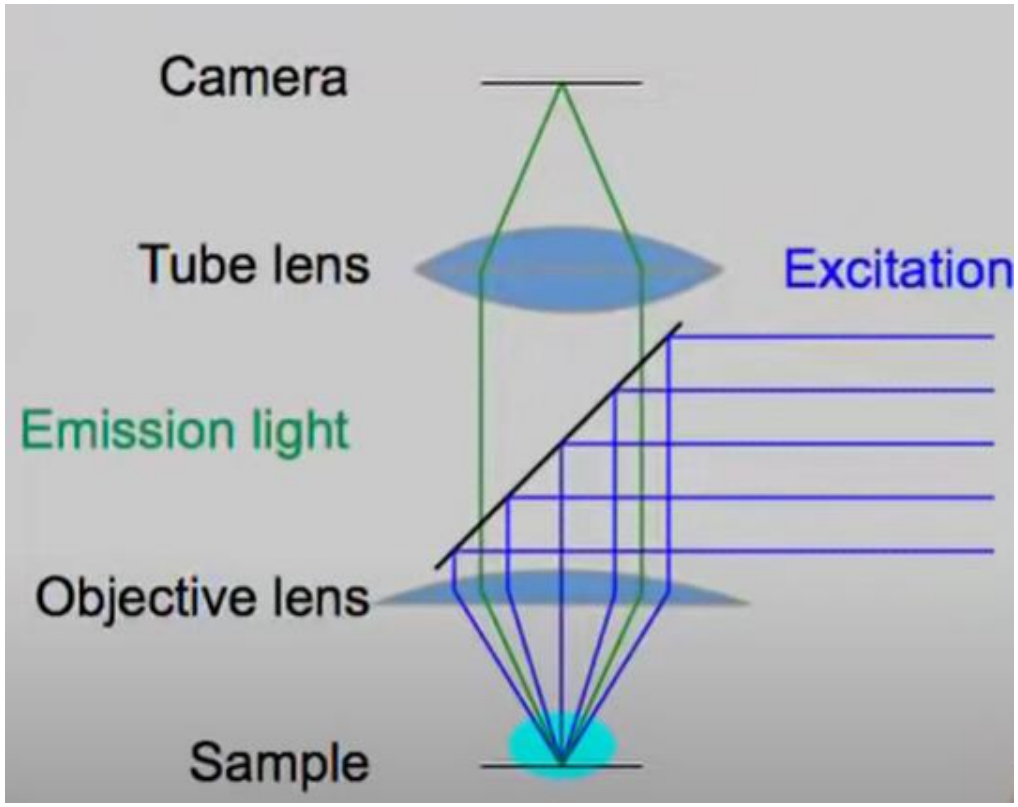

Problem: fluorescence is emitted along entire illuminated cone, not just at the focus

The confocal microscope uses a pinhole to **block out-of-focus light**

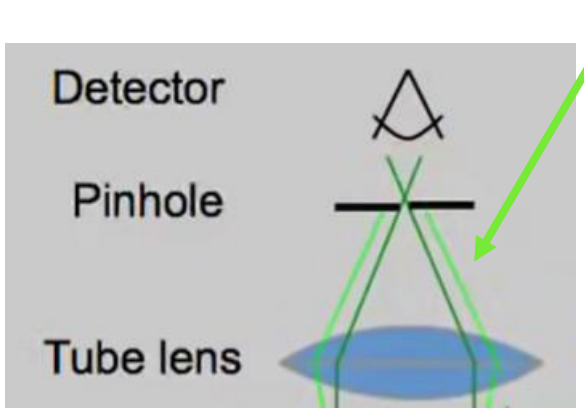

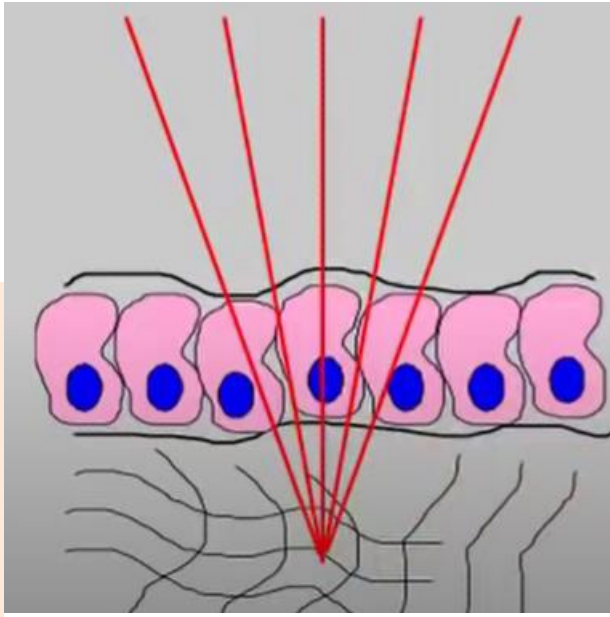

Limitations of tissue penetration depth:

Absorption

Scattering

- Imaging in **near-infrared** results in lower scattering and minimizes absorption

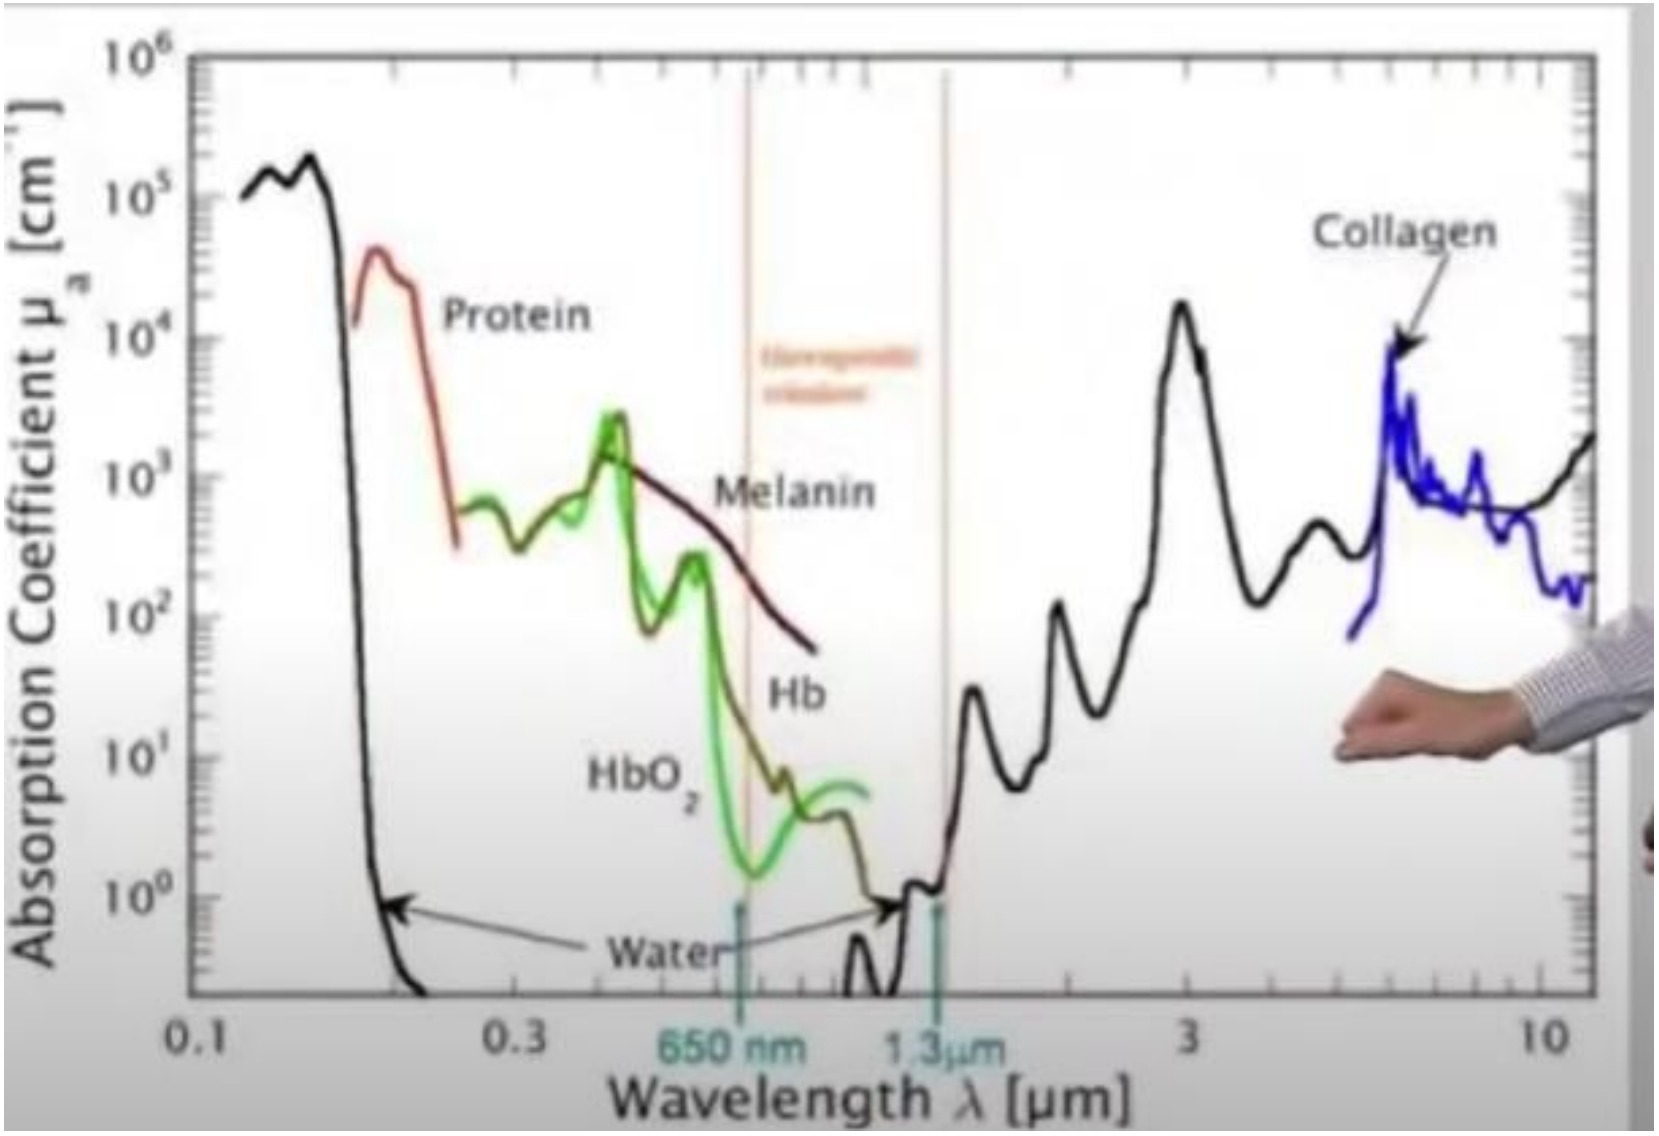
Absorption Coefficient μ_a [cm^{-1}]
10^6
10^5
10^4
10^3
10^2
10^1
10^0
Protein
Melanin
Hb
HbO_2
Water
Collagen
0.1
0.3
650 nm
1.3μm
3
10
Wavelength λ [μm]

Fluorescence Imaging

- **Form of luminescence** that results from matter emitting light of a certain wavelength after absorbing electromagnetic radiation
- **Absorb light** in a certain color and emit light in a different color
- ***Fluorophores***: Molecules that re-emit light upon absorption of light

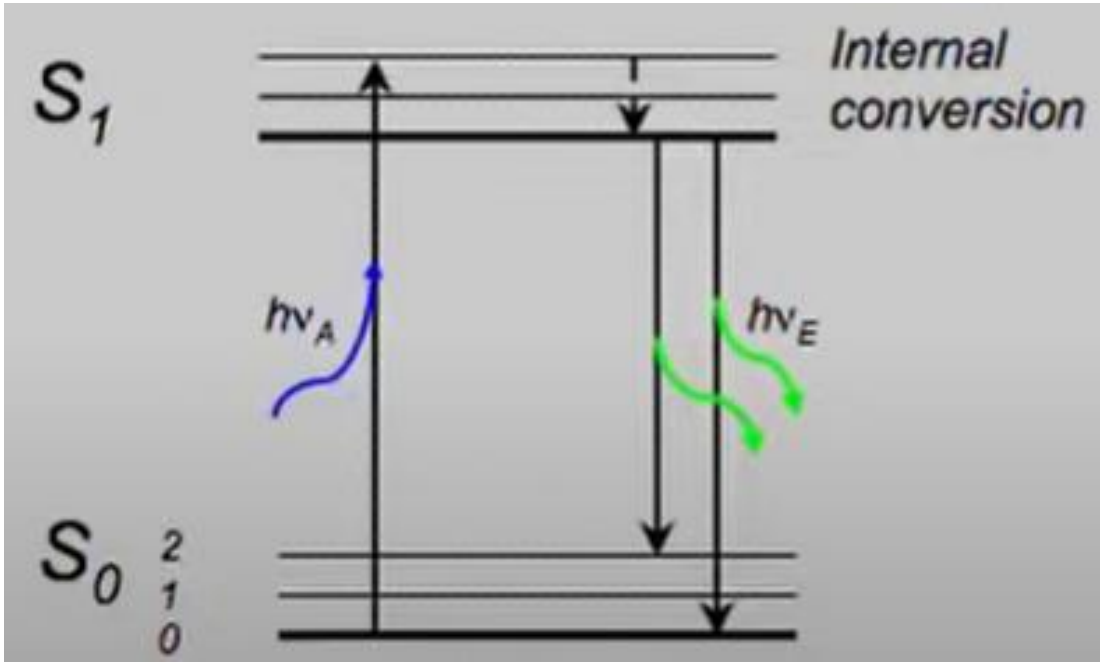

Basic Concepts

- When a certain molecule absorbs light, the **energy of the molecule** is **briefly raised to a higher excited state** (μεταβάσεις ηλεκτρονίων)
- The subsequent return to ground state results in ***emission of fluorescent light*** that can be **detected and measured**
- The emitted light, resulting from the absorbed photon of **energy h_v, has a specific wavelength**
- The measuring device needs to know this wavelength to detect light production
- Fluorescent dyes: when the bind to proteins they become more easily detectable
- Other molecules may aborbe light, however the light they emit is in different frequency and thus are not get detected

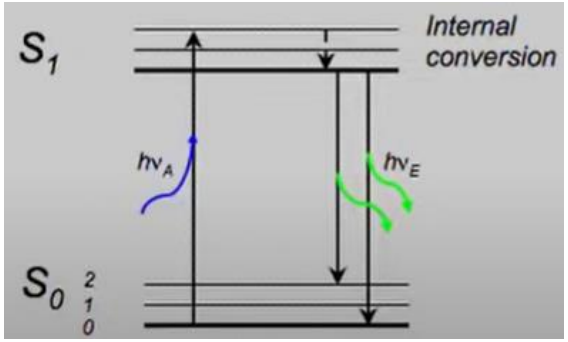

Photon Florescence Microscopy

One-photon excitation

Linear process, i.e., if you double the laser intensity, you will double the fluorescence intensity

1P excitation

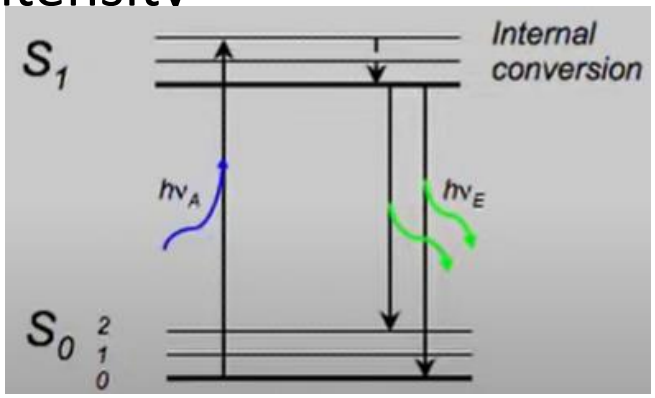

2P excitation

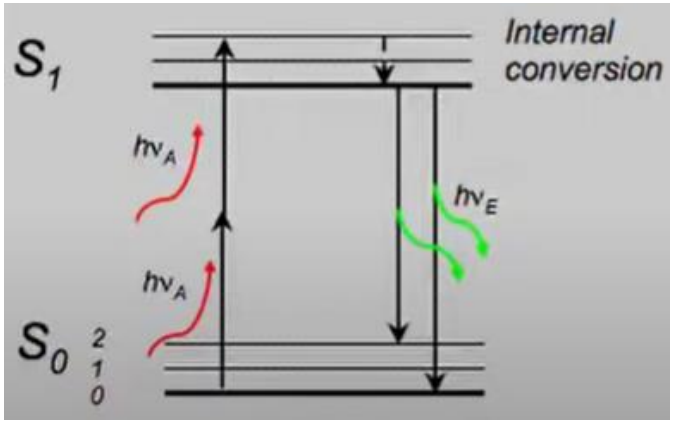

Two-photon excitation

- Emission of two laser photons
- **Non-linear process**

The absorption rate depends on the second power of the light intensity

In a focus laser, the intensity is highest in the vicinity of the focus and ***drops off quadratically*** with ***distance above and below***

As a result, fluorophores are **excited almost exclusively in a tiny diffraction-limited focal volume**

GCaMP a genetically encoded calcium indicator

When bound to Ca^{2+}, GCaMP fluoresces **green** with a peak excitation wavelength of 480 nm and a peak emission wavelength of 510 nm

In order to identify the cells that fire, cells that are genetically engineered with GCaMP dye, when the density of calcium ions increases, the light from the 2 photons from the laser gets absorbed and results in the emission of a light in frequency that can be detected and estimated from the microscope from the cells that have fired

Για την επιτυχή έκλυση νευροδιαβιβαστή

Απαιτείται η έκλυση δυναμικού ενέργειας από το προσυναπτικό κύτταρο

- Το δυναμικό ενέργειας **προκαλεί την *ενεργοποίηση διαύλων ασβεστίου***
- Αύξηση ***ενδοκυττάριας συγκέντρωσης ασβεστίου***
- Το **μέγεθος του εισερχόμενου ρεύματος ασβεστίου** επηρεάζει **την ποσότητα νευροδιαβιβαστή που θα εκλυθεί**, και κατά επέκταση το μέγεθος του μετασυναπτικού δυναμικού

- Ο ***<u>νευροδιαβιβαστής</u>*** είναι **αποθηκευμένος σε κυστίδια** είτε κοντά στην μεμβράνη του τερματικού του άξονα είτε προσδεμένα πάνω στο κυτταρο-σκελετό
- Συγκεκριμένες πρωτεΐνες σχηματίζουν πόρους συγχώνευσης των δύο μεμβρανών δηλ. του κυστιδίου και του κυττάρου, ώστε να γίνει η ένωση τους & η **απελευθέρωση του νευροδιαβιβαστή στο εξωκυττάριο** περιβάλλον

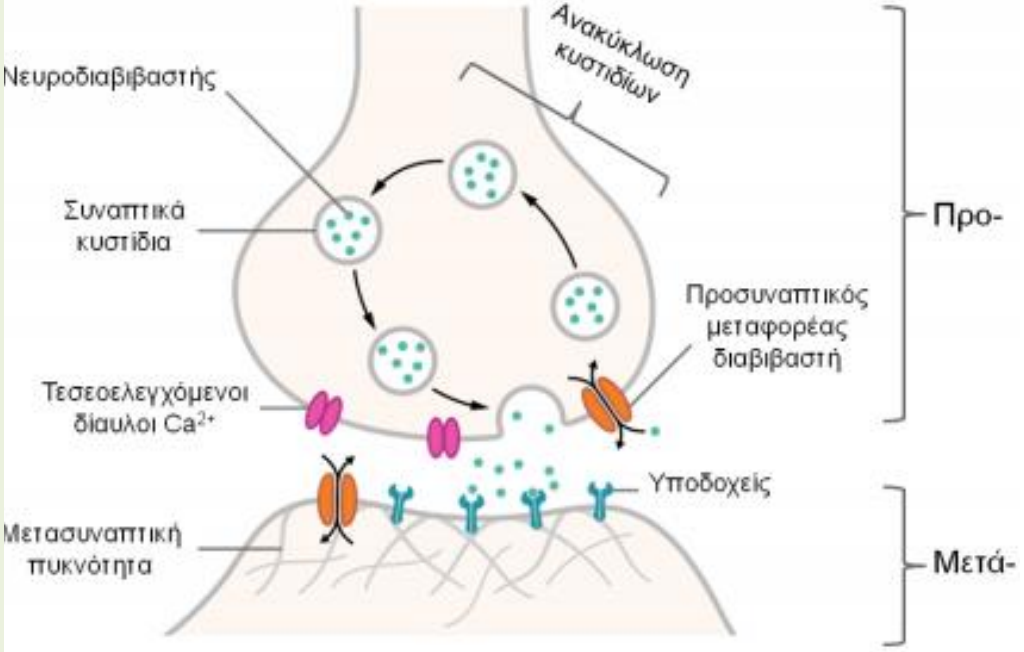

Neural Population Decoding

- Neural decoding predict stimuli/behavior

f(neural activity) → stimulus

Decoding has been used for 30 years

Georgopoulos *et al* 1986

Neural Population in Primate Motor Cortex

- Although individual neurons in the arm area of the primate motor cortex are only ***broadly tuned*** to a **particular direction in 3D-space**, the animal can very precisely control the movement of its arm.
- The ***direction of movement*** was found to be **uniquely predicted by the action of *a population of motor cortical neurons***.
- When individual cells were represented as vectors that make weighted contributions along the axis of their preferred direction (according to changes in their activity during the movement under consideration) the resulting vector sum of all cell vectors (population vector) was in a direction congruent with the direction of movement.
- This population vector can be monitored during various tasks, and similar measures in other neuronal populations could be of heuristic value where there is a neural representation of variables with vectorial attributes.

Developing the classifier

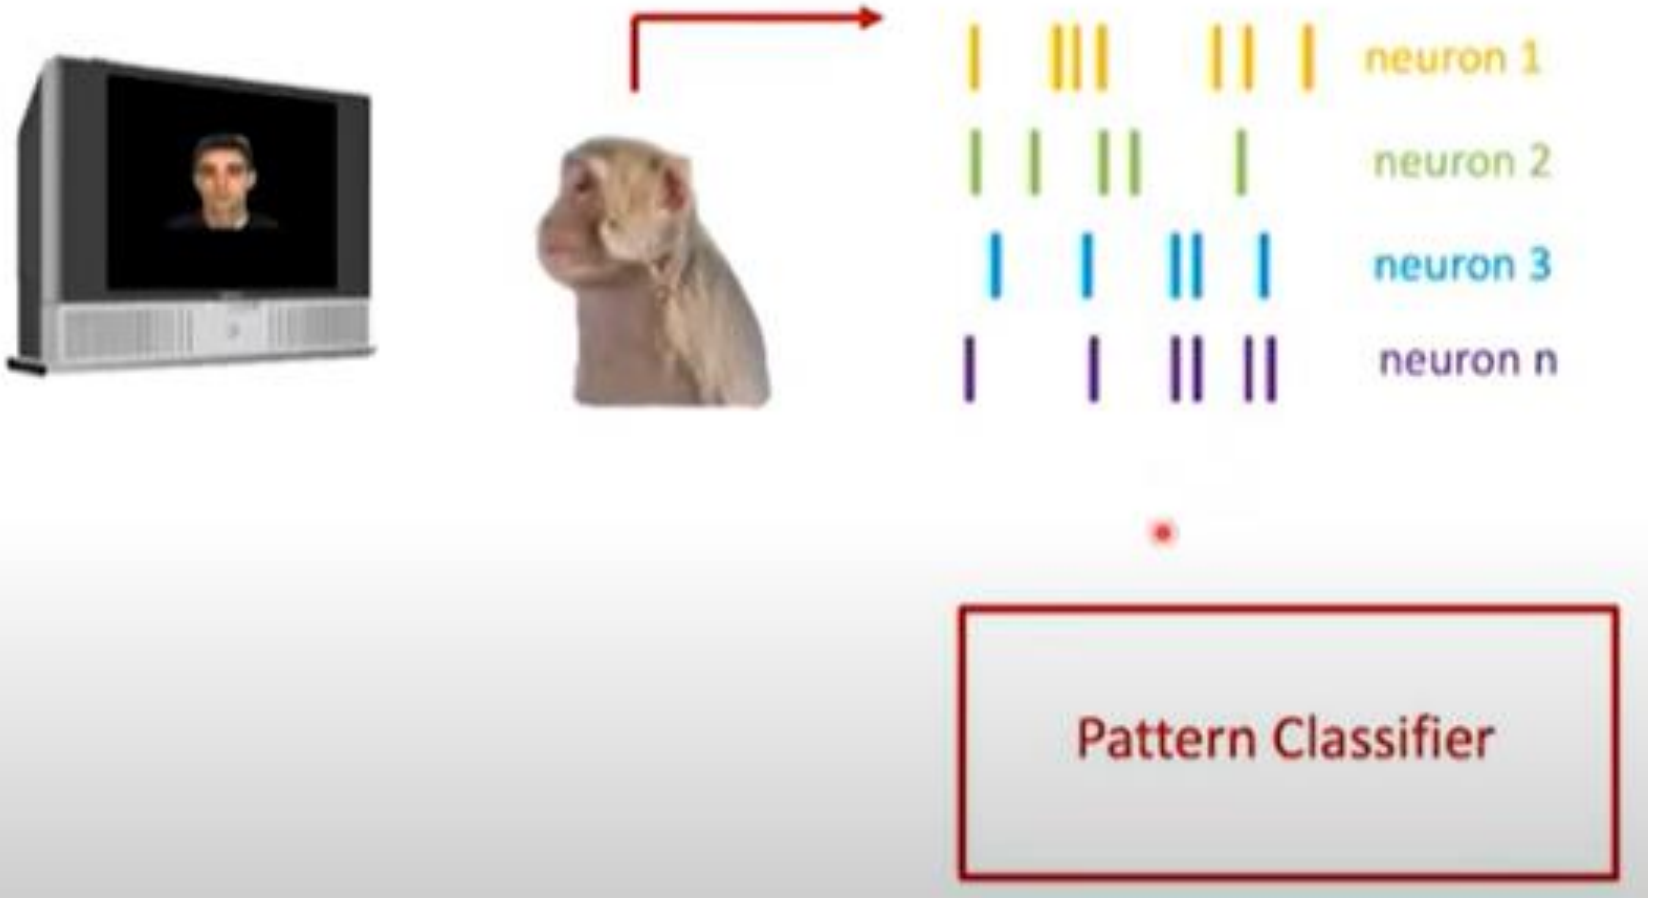

Decoding basics: a simple example

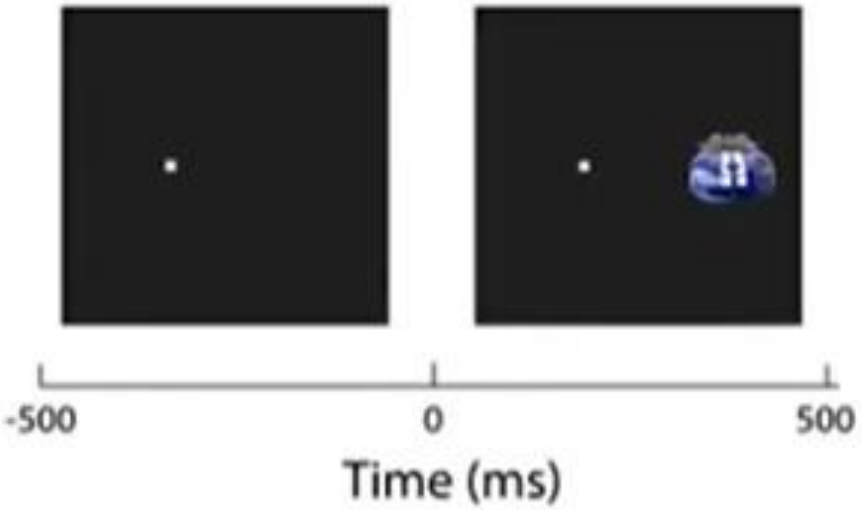

132 neurons recorded from IT

Seven objects:

Zhang, Meyers, Bichot, Serre, Poggio, and Desimone, PNAS, 2011

Face identification invariant to head pose

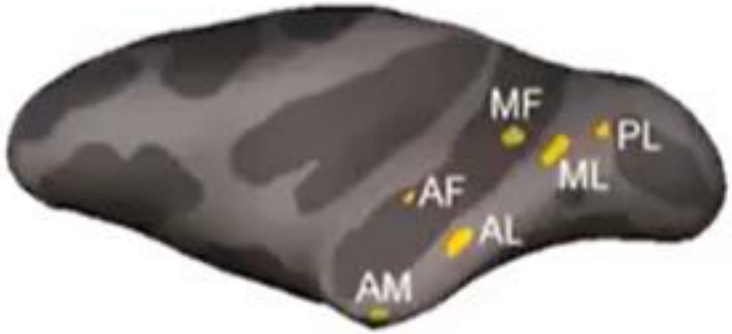

Meyers, Borzello, Freiwald, Tsao, J Neurosci, 2015

Face identification invariant to head pose

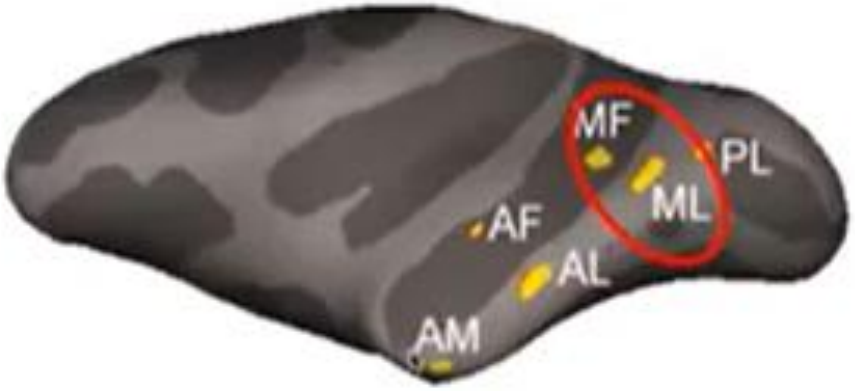

Stimulus set: 25 individuals, 8 head poses per individual

Face identification invariant to head pose

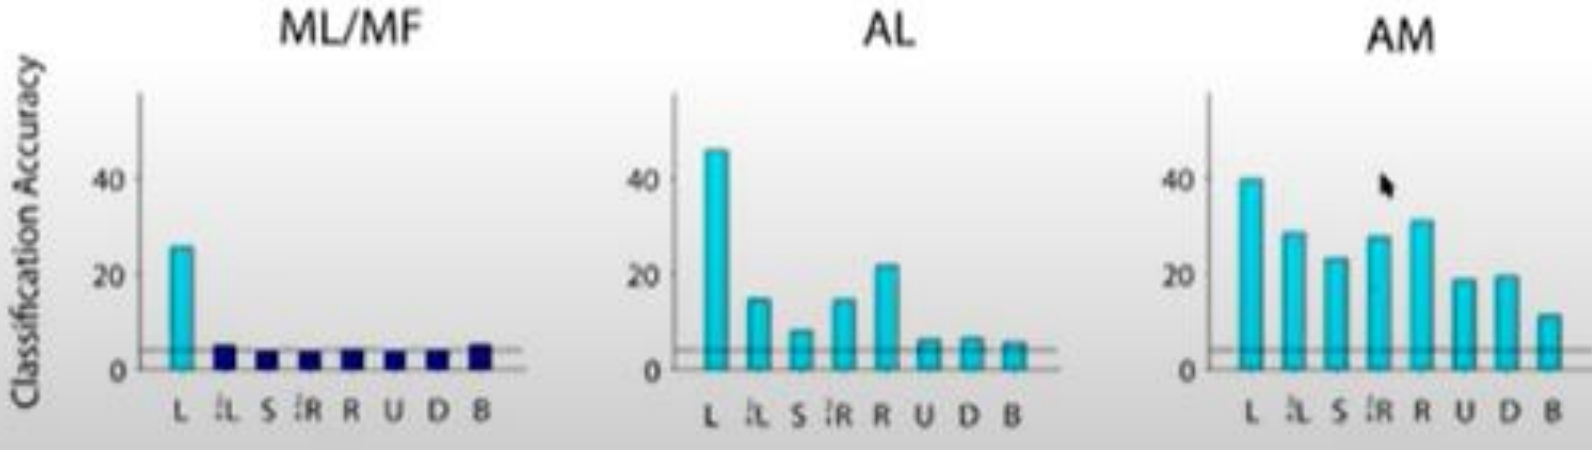

Is information contained in a dynamic population code?

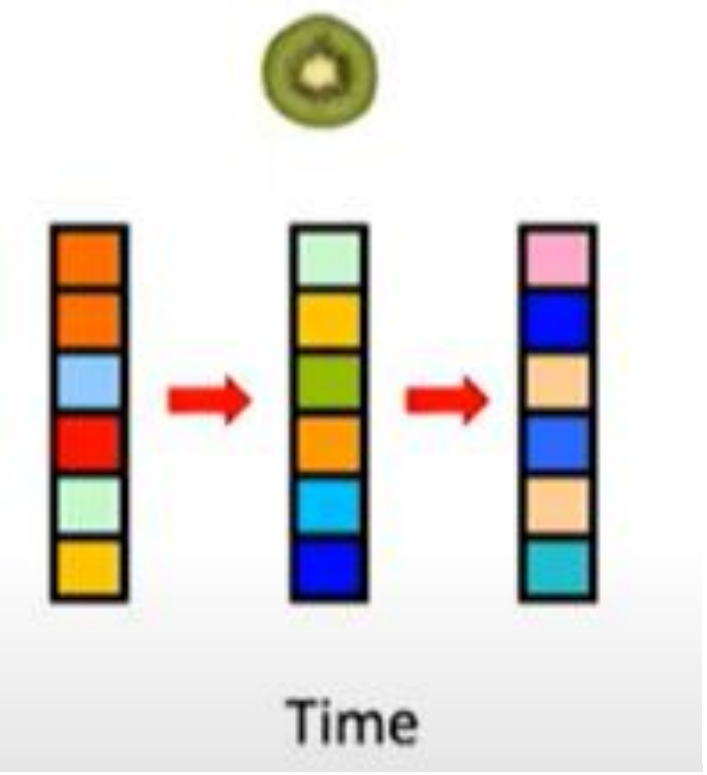

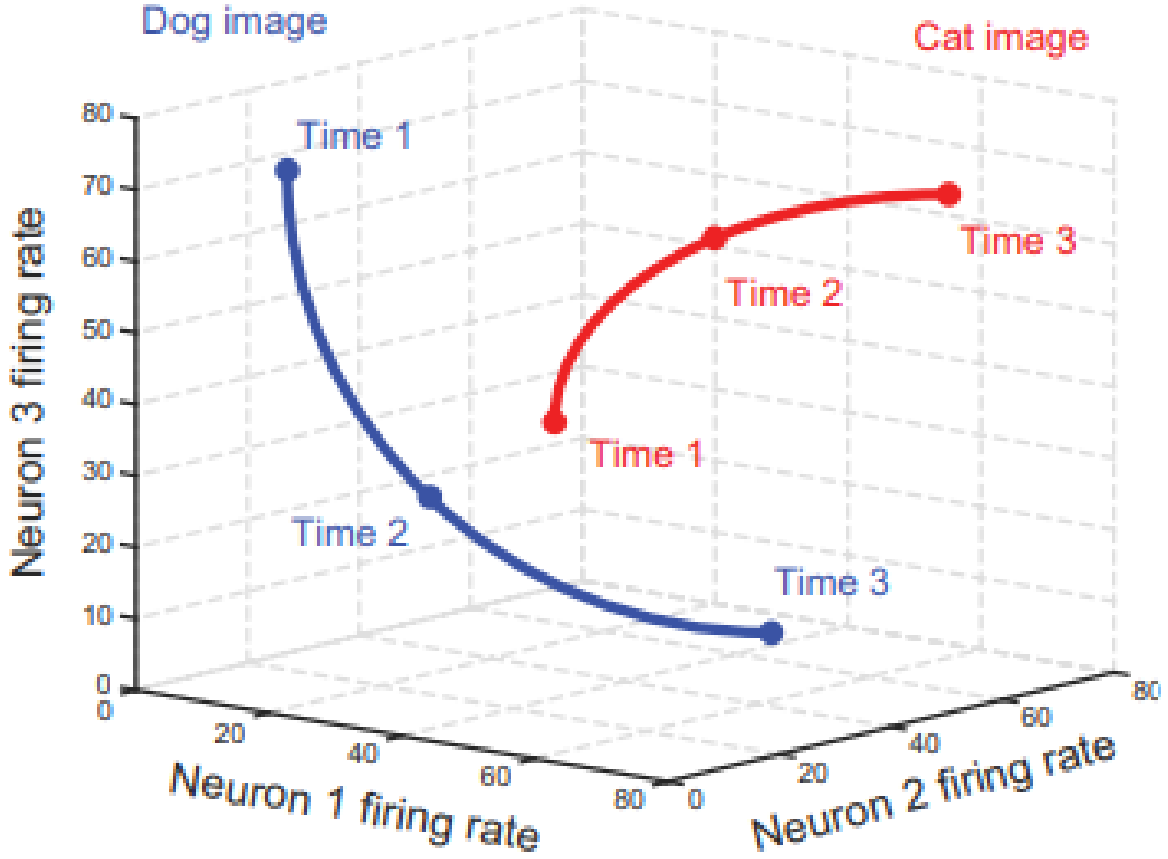

Meyers et al 2008, King and Dehaene 2014, Meyers 2018

Introduction of Neural Codes

- How do neurons process information received from a **stimulus**
- Measures electrical activity between neurons and how it carries information
- How is information encoded in a series of action potentials?

Neural Coding: Which **features of neural activity** carry this information

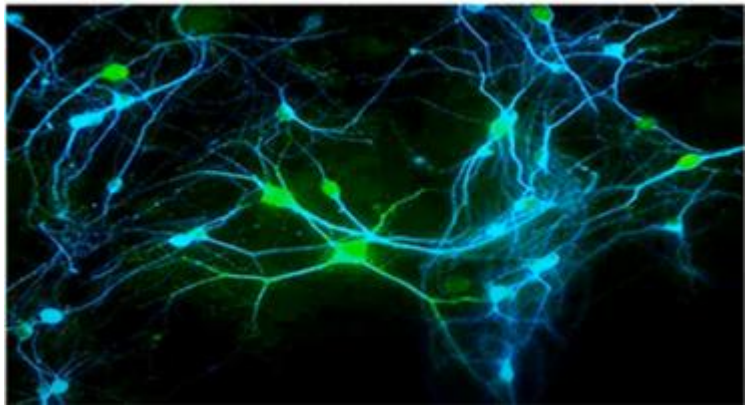

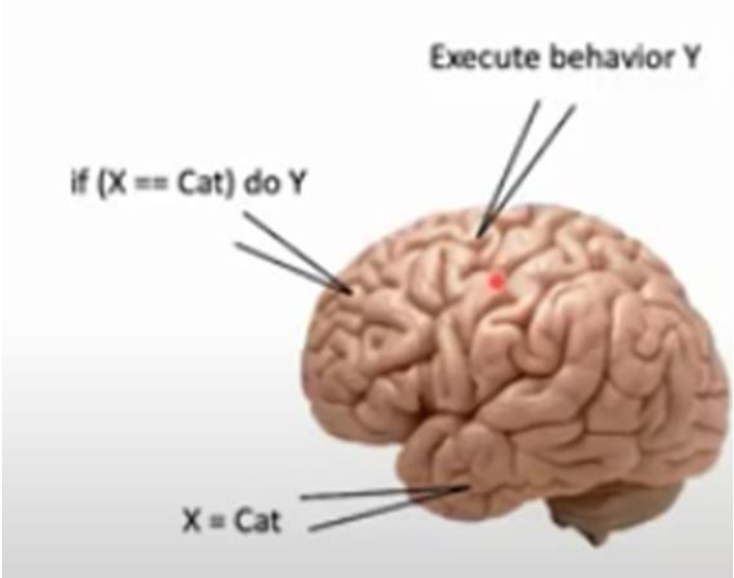

Predict the stimuli/behavior from the neuronal activity

f(neuronal activity) → ***stimuli***

Population Coding: An Overview

"The quantitative study of which algorithms or representations are used by the brain to combine together and evaluate the messages carried by different neurons" [2]

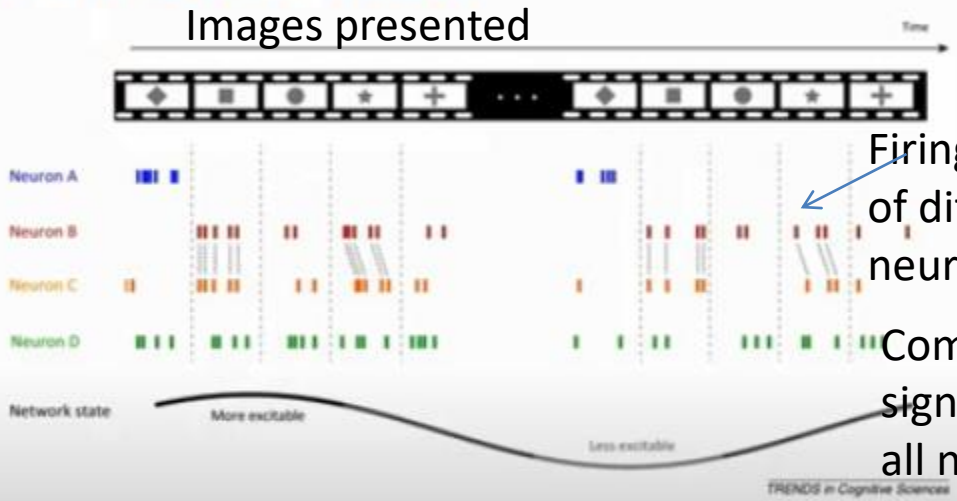

Figure 2: A population of four neurons are exposed to different stimuli and the resulting action potentials are used to determine fluxuations in the state of the network [3]

Understanding neural algorithms

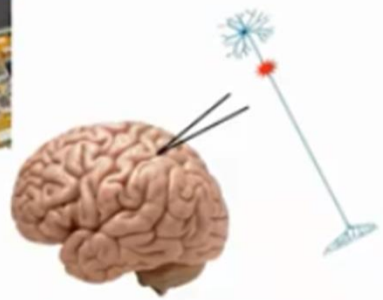

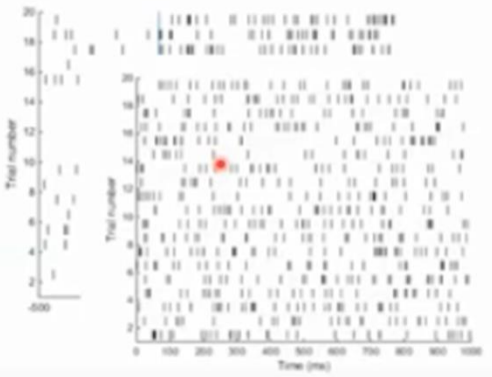

How can we convert noisy data into useful information?

From the talk: https://www.youtube.com/watch?v=m3OQwz9PhcA

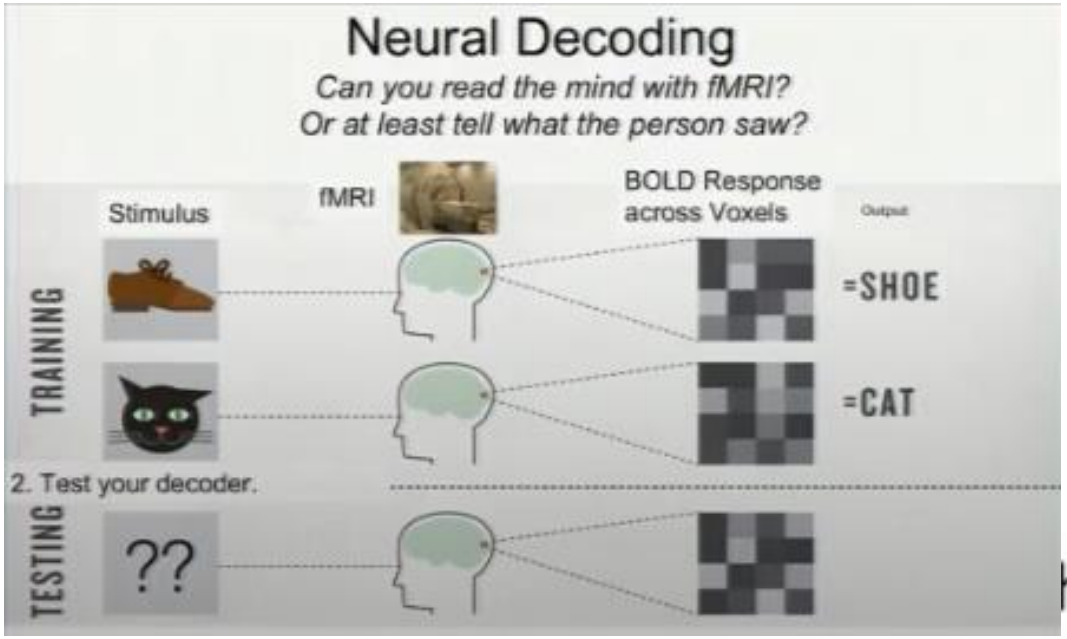

he classifier

Pattern Classifier

Population Coding: Why Use It?

- Difficult to differentiate information from a single neuron
- Better explains behavioral performance
- Responses of many neurons may be combined to determine some value about the inputs

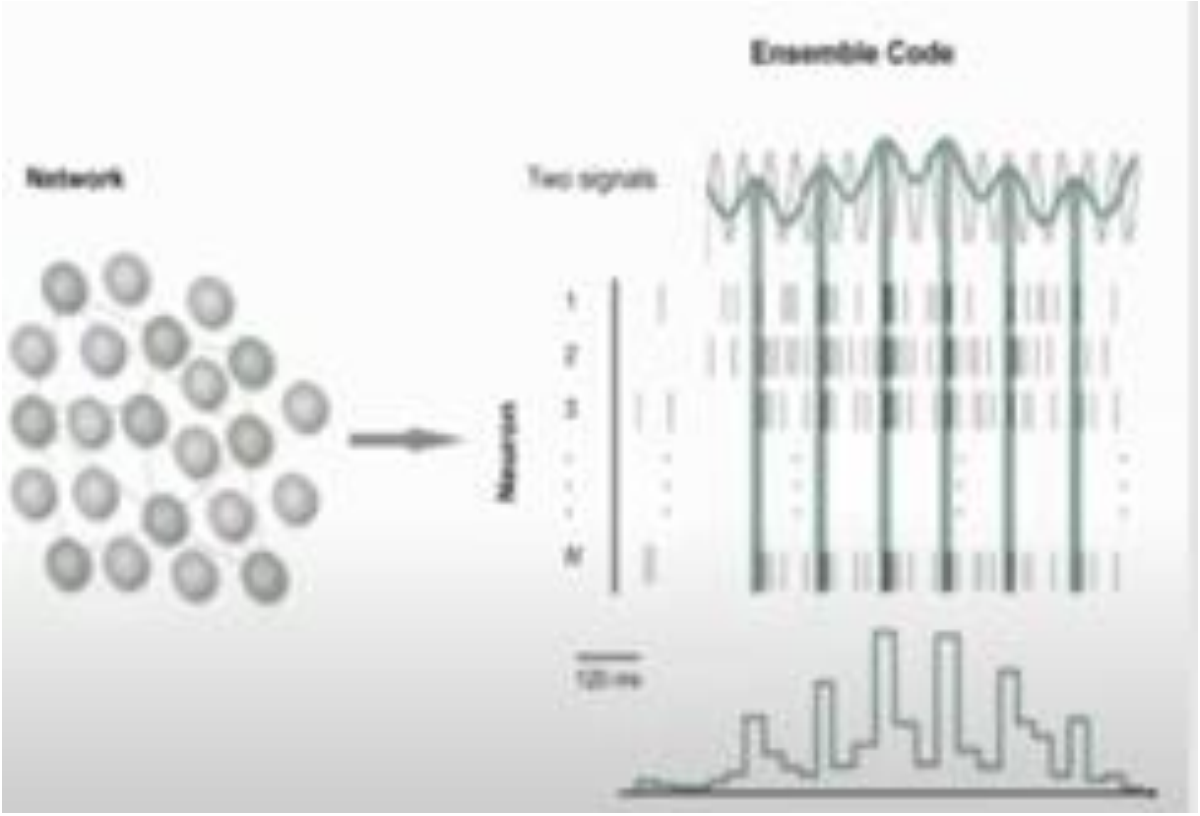

Population Coding: Primary Visual Cortex (V1)

- One of the best area to examine the population coding, since the **relationships between the characteristics of stimuli and activation of neurons is *well documented***
- ***Similar structure*** **across mammalian species**
- Pyramidal neurons: tuned to retinal location, orientation, contrast, speed, spatial and temporal frequency

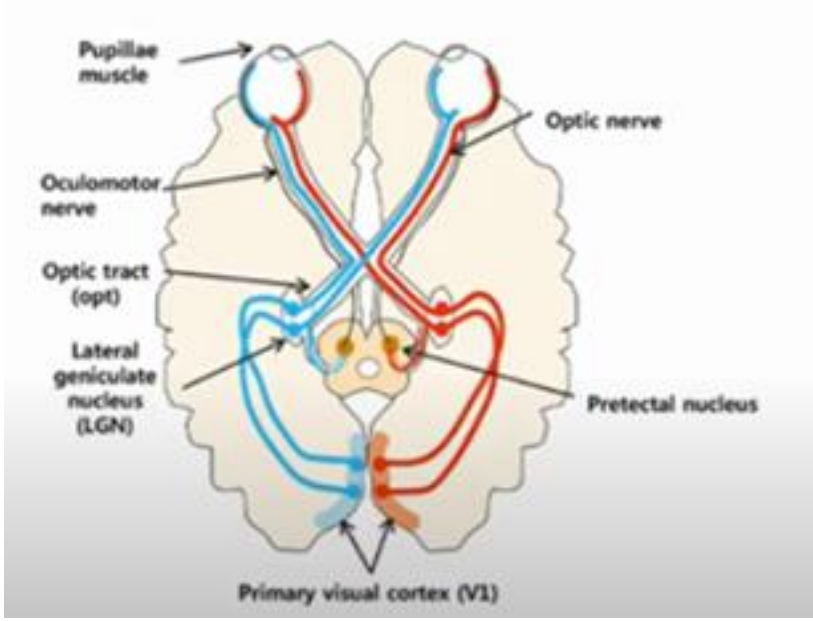

Motivation for Population Coding

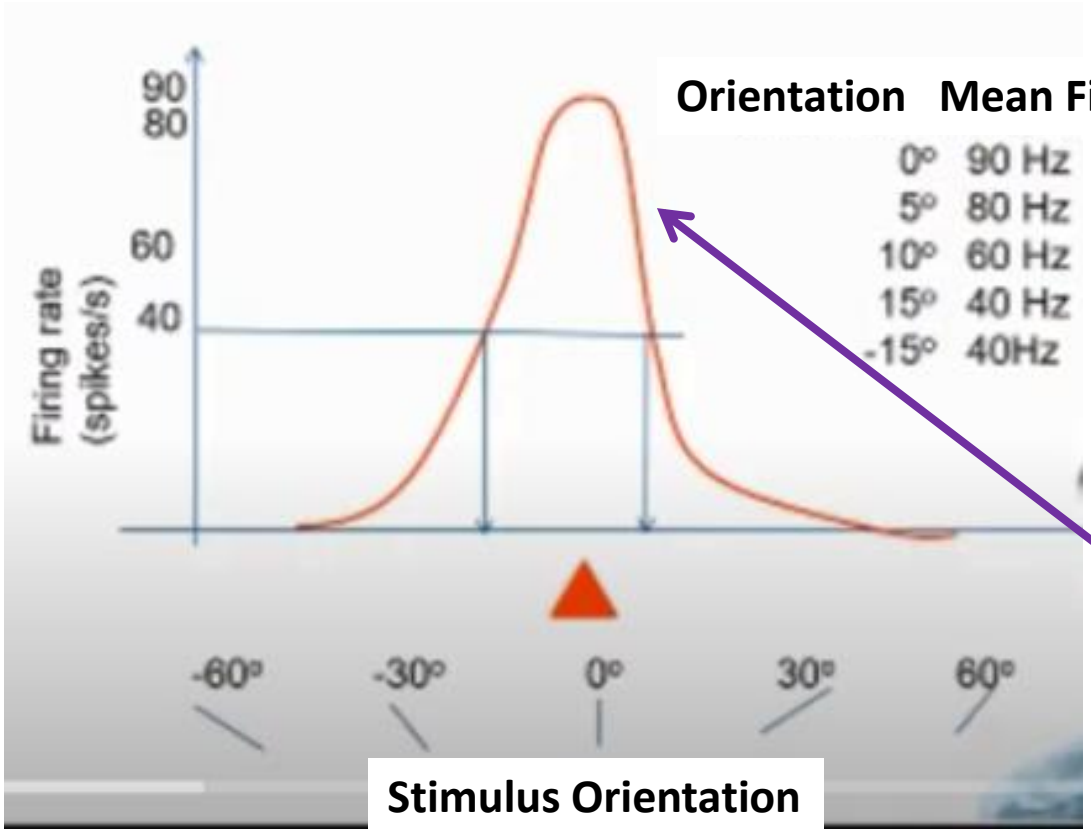

When we have the tuning curve of a ***single*** neuron, and record its firing rate, we ***cannot*** estimate precisely the stimulus orientation (e.g., the edge present in the visual field of the neuron)

Examples in our datasets

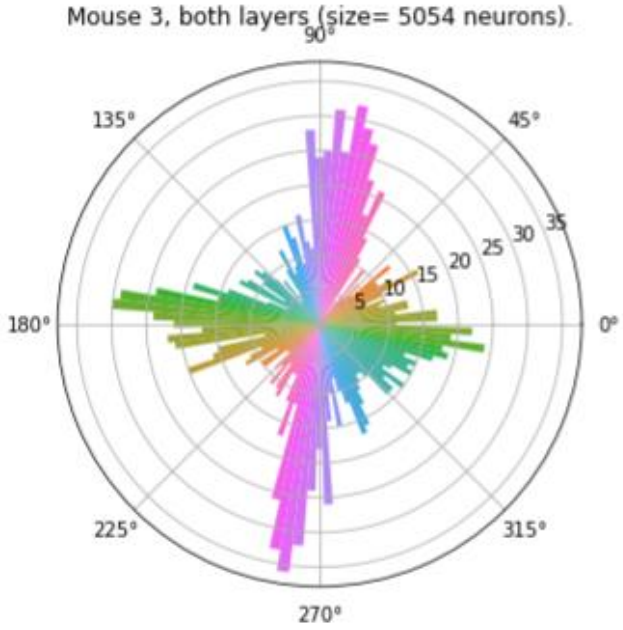

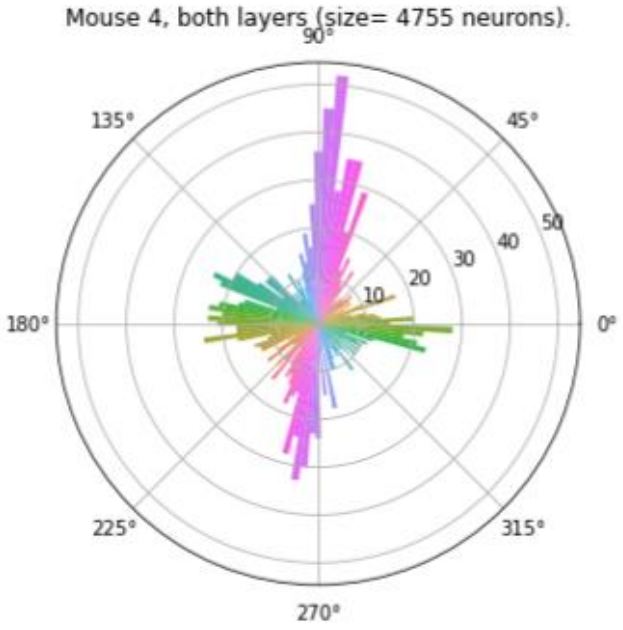

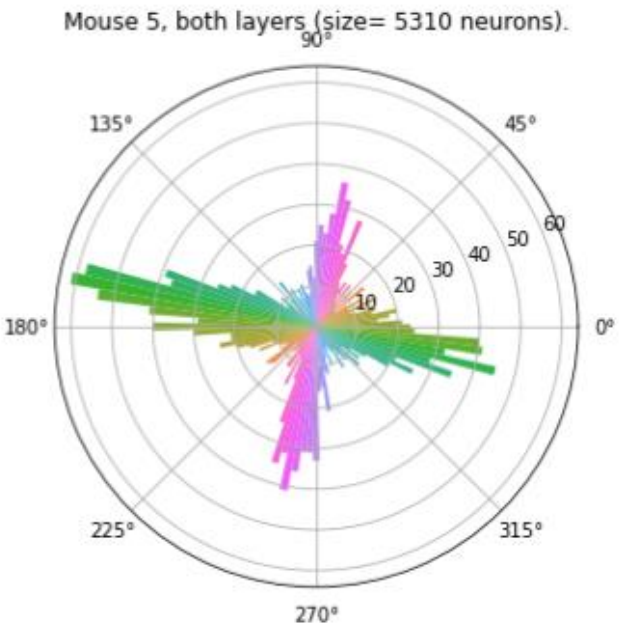

Motivation for Population Coding

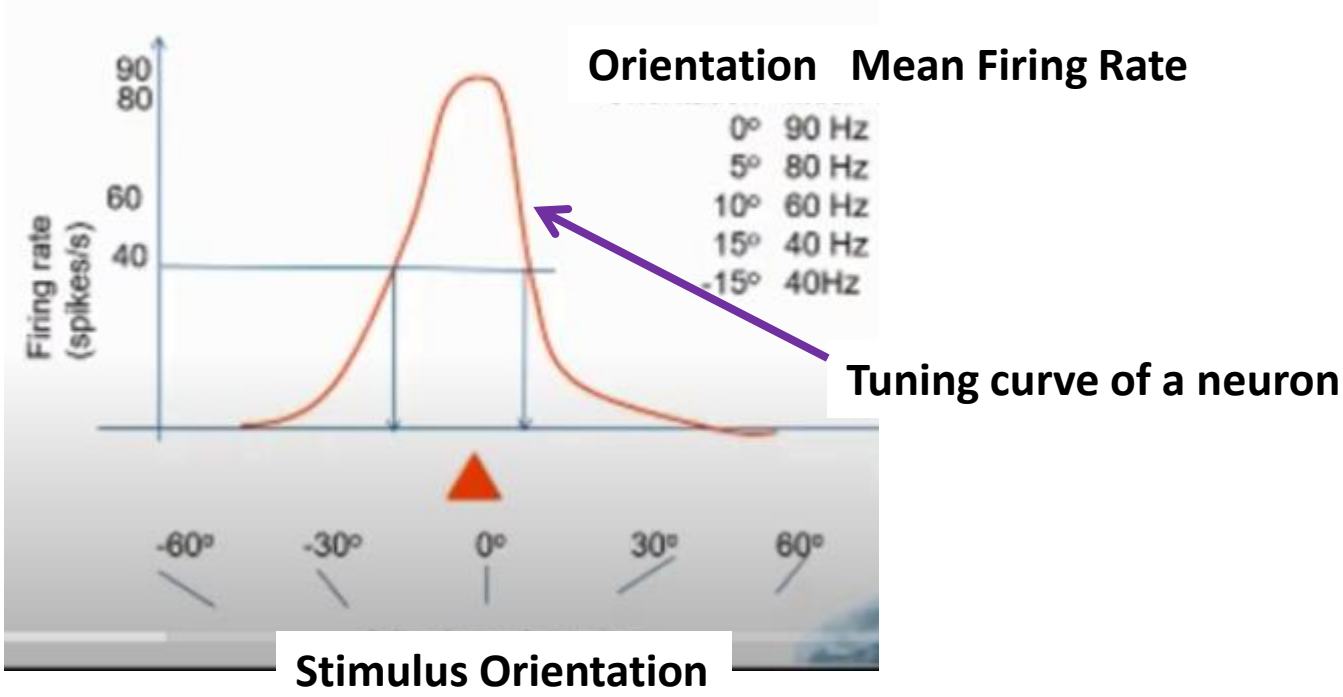

When we have the tuning curve of a ***single*** neuron, and record its firing rate, we ***cannot*** estimate precisely the stimulus orientation (e.g., the edge present in the visual field of the neuron)

Motivation for Population Coding

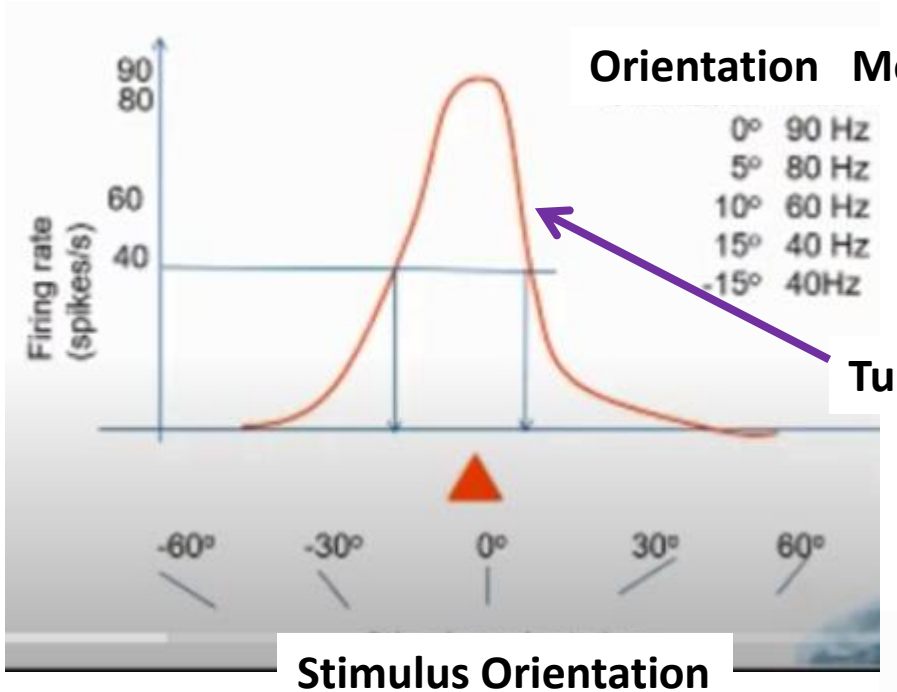

Using a number of neurons to resolve the ambiguity

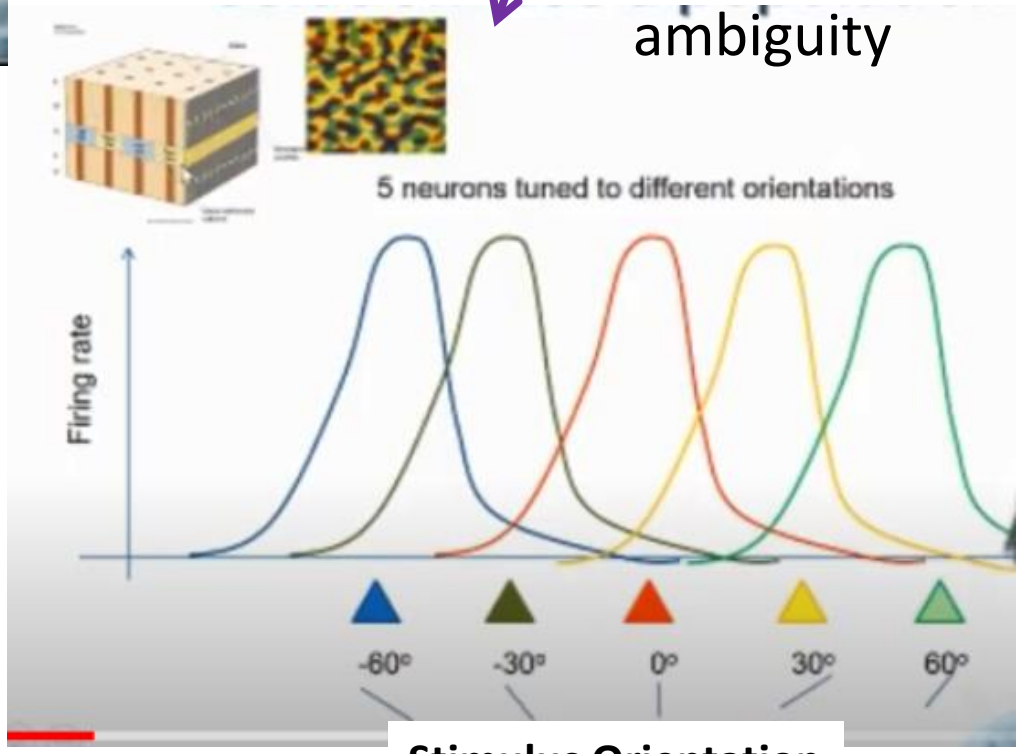

When we have the tuning curve of a ***single*** neuron, and record its firing rate, we ***cannot*** estimate precisely the stimulus orientation (e.g., the edge present in the visual field of the neuron)

Example of Population Coding using Five Neurons

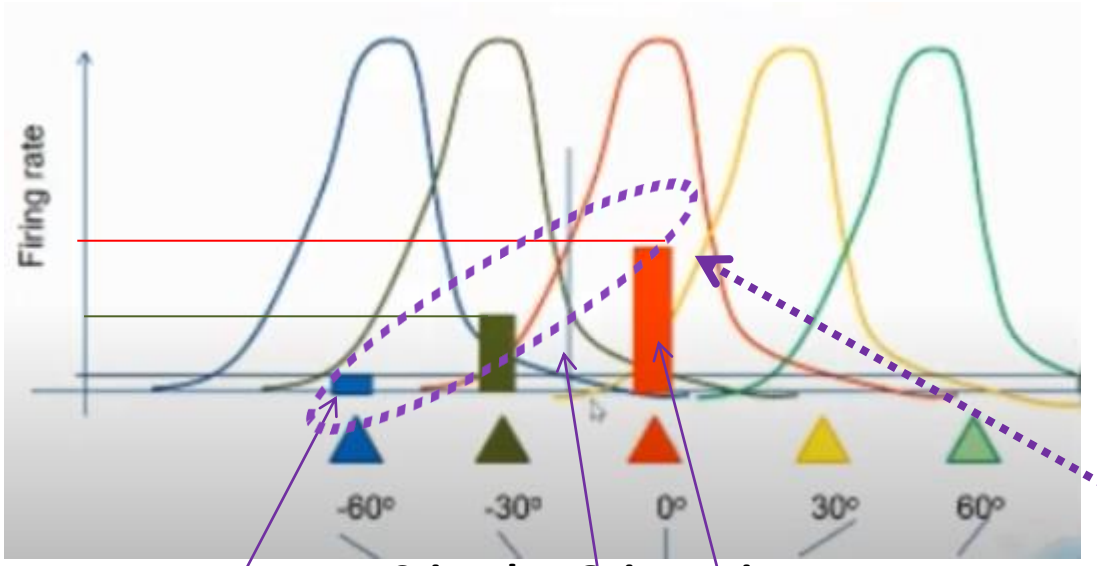

Retinal Stimulus

Firing rate of the "blue" neuron

Stimulus

Firing rate of the "red" neuron

If the brain gets these firing rates from each neuron, it can then ***deduce*** the orientation of the stimulus without ambiguity

Each tuning curve corresponds to a different neuron (indicated with blue, red, dark green, yellow, and light green color)

The wide bars indicate how much the corresponding neurons respond to the presentation of the stimulus.

Example of Population Coding using Five Neurons (con't)

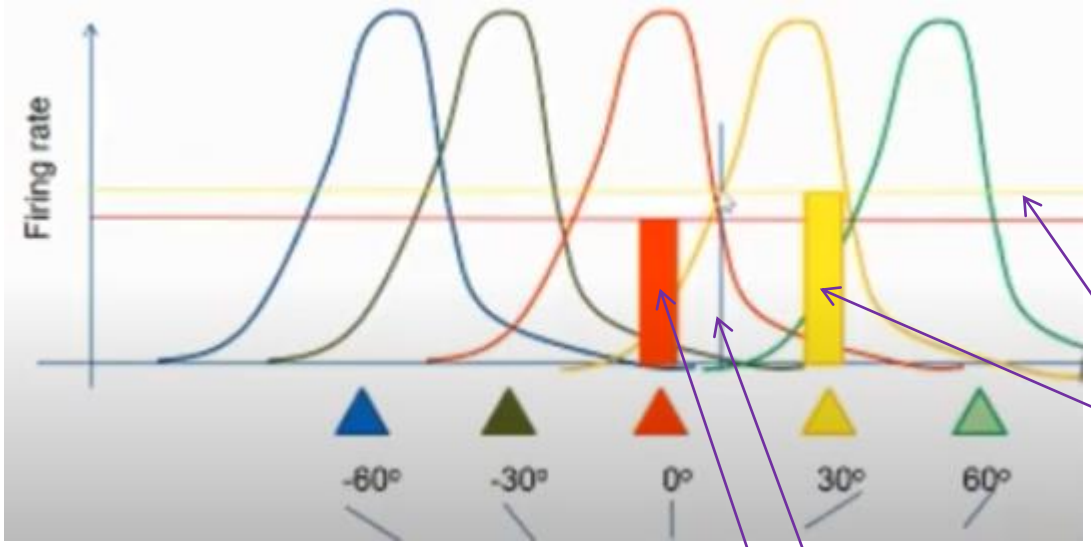

Stimulus Orientation

Retinal Stimulus

Firing rate of the "yellow" neuron, **which now fires more**

Stimulus

Firing rate of the "red" neuron

If the brain gets these firing rates from each neuron, it can then deduce the orientation of the stimulus without ambiguity

Each tuning curve corresponds to a different neuron (indicated with blue, red, dark green, yellow, and light green color)

The wide bars indicate the firing rate of the corresponding neuron at the presentation of the stimulus.

Example of Population Coding using Five Neurons (con't)

Double peaks

Now there is a stimulus with two orientation (+60 and -30)

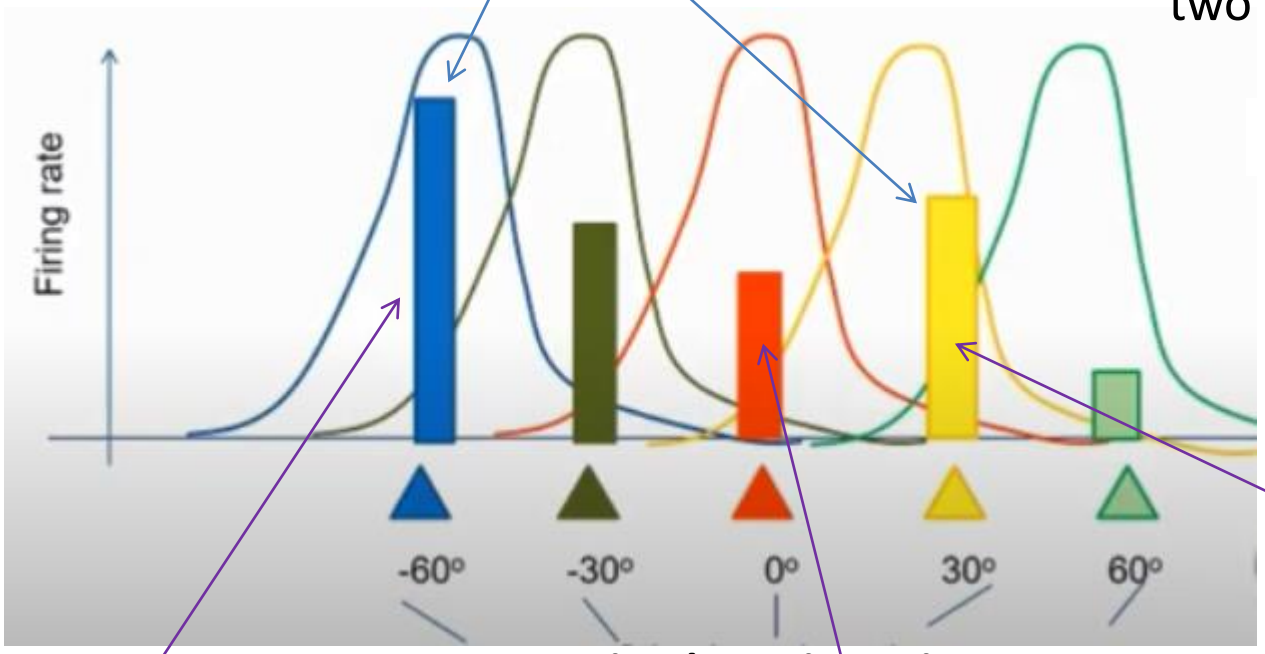

Stimulus Orientation

Retinal Stimulus

The "yellow" neuron still fires

Now neuron "blue" fires more

Firing rate of the "red" neuron

If the brain gets these firing rates from each neuron, it can then deduce the orientation of the stimulus without ambiguity

Spontaneous Events Outline the Realm of Possible Sensory Responses in Neocortical Populations

Artur Luczak,[1,2] Peter Barthó,[1] and Kenneth D. Harris[1,*]

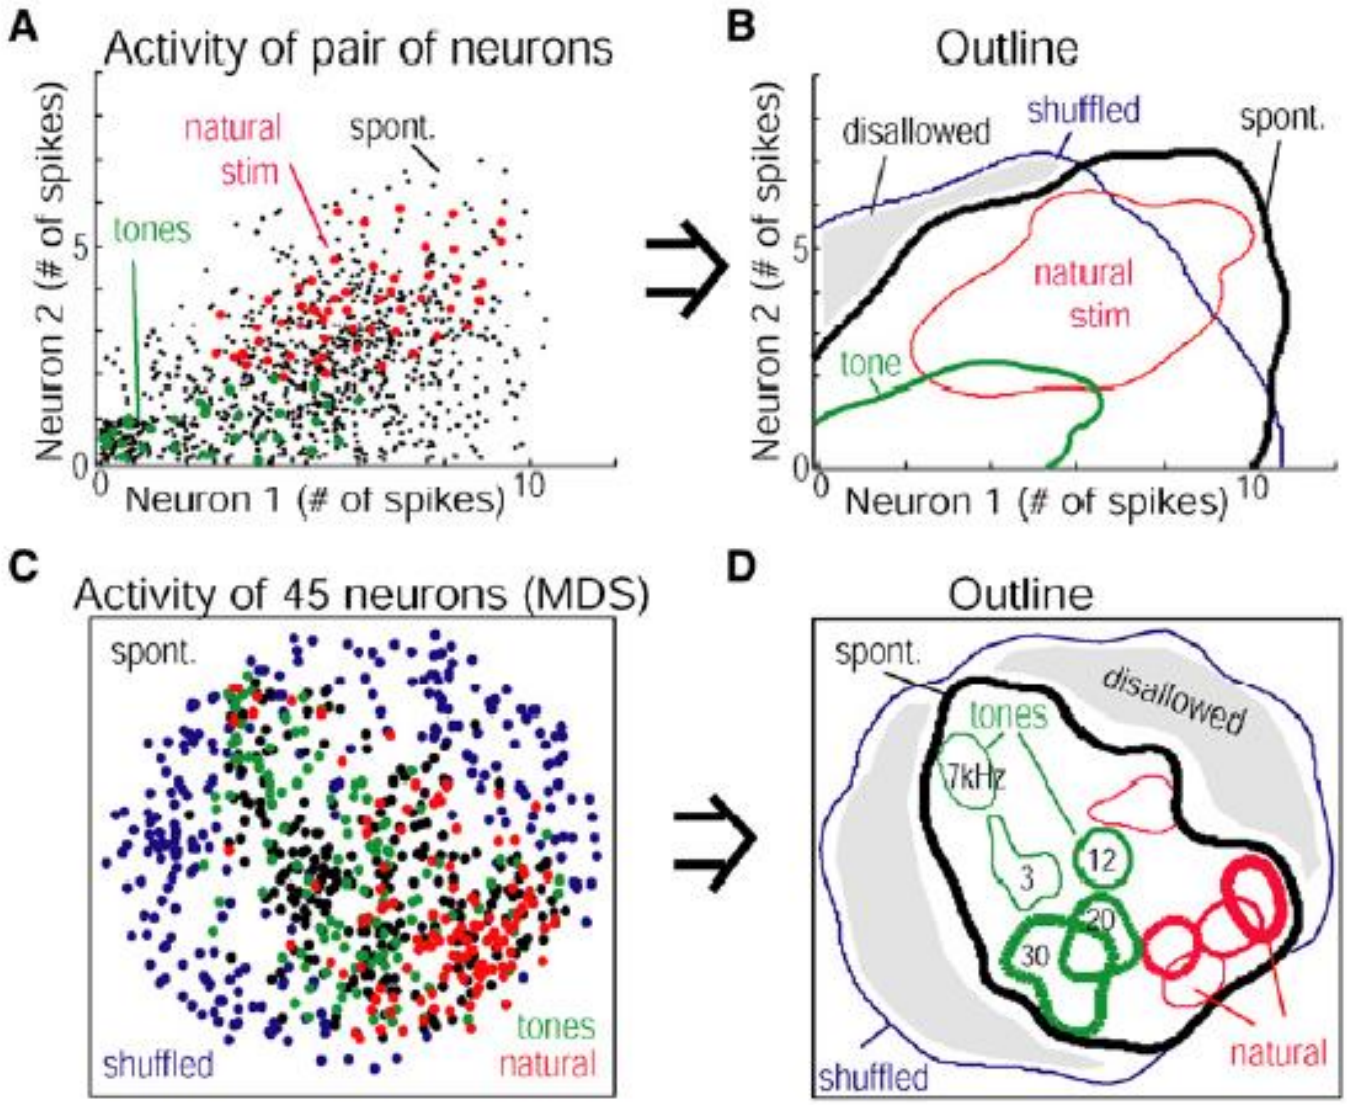

D. Contour plot derived with responses to individual stimuli marked separately. **Sensory-evoked responses again lie within the realm outlined by spontaneous events.**

A. Spike counts of two neurons (recorded from separate tetrodes) during the first 100 ms of **spontaneous upstates** (black), **responses to a tone** (green), and **responses to a natural sound** (magenta). Data were jittered to show overlapping points.

Regions occupied by responses to the sensory stimuli differ but are both contained in the **realm outlined by spontaneous patterns.**

How are assemblies created?

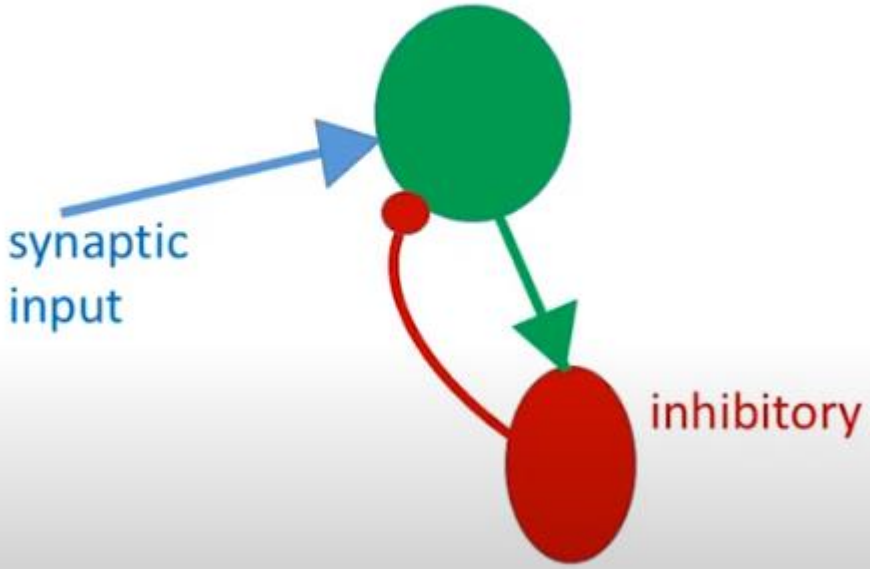

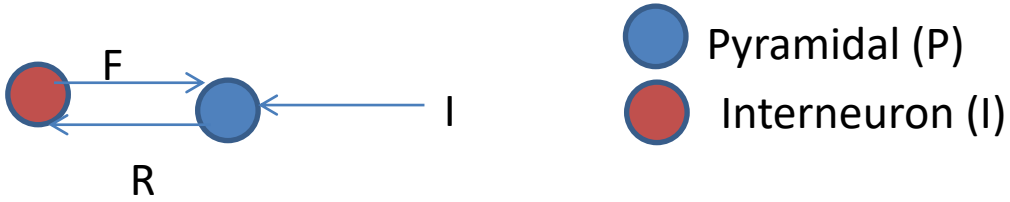

Input I: arrival rate of spikes
R: input firing rate of the pyramidal to the Interneuron
F: input firing rate of the interneuron to the pyramidal
At each neuron, we have an **Integrate-and-Fire model**
The firing of the **interneuron** causes **suppression of the Pyramidal neuron**

State of a neuron: its **membrane potential**

Assume: $N_i(t)$: counting process $\{N_i(t), t>=0\}$ number of spikes that have arrived at **neuron i** by time t. Poisson Process of **rate λ_i**

If the Interneuron fires, it will instantaneously silence the Pyramidal
Model the system as discreet Markov-Chain
What is the **equilibrium states?** What are its stable solution?
Compute its limiting probabilities

State of the system: (State of P, State of I), where the state of the Pyramidal neuron (P) is the number of spikes that arrive within Dt, and the state of the Interneuron (I) is 1 (fire) or 0 (not firing)

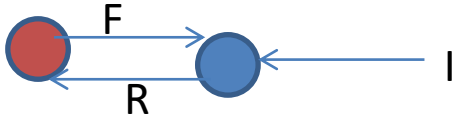

Pyramidal (P) Interneuron (I)

Input **I**: **arrival rate of spikes**
R: input firing rate of the **Pyramidal to the Interneuron**
F: input firing rate of the **Interneuron to the Pyramidal**
At each neuron, we have an **Integrate-and-Fire model**
The firing of the Interneuron causes suppression of the Pyramidal neuron

Assume: $\mathbf{N_i(t)}$: **counting** process $\{N_i(t), t>=0\}$ number of spikes that have arrived at neuron i by tir
t. Poisson Process of rate λ_i

If the Interneuron fires, it will instantaneously silence the Pyramidal
Model the system as **Discreet Markov-Chain**
What is the **equilibrium states?** Which are its stable solutions? Compute its limiting probabilities

State of the system (**State of P**, **State of I**): state of the Pyramidal neuron (P) is # spikes that arrive
within Dt & the state of the Interneuron (I) is 1 (fire) or 0 (not firing);

The **state of the I** is the number of spikes that arrive within Dt (from the Pyramidal) and that **I fire**
with **prob.** $\mathbf{f_I}$

$\text{Prob}[(i, 0) \rightarrow (0, 1)] = f_I$ for i> 1, $\quad\quad$ $\text{Prob}[(i, 0) \rightarrow (0, 0)] = 1- f_I$ for i>1,

$\text{Prob}[(0, 0) \rightarrow (1, 0)] = \lambda_m \Delta t\, e^{-\lambda_m Dt}$, $\quad\quad$ $\text{Prob}[(0, 0) \rightarrow (\mathbf{n}, 0)] = ((\lambda_m \Delta t)^n / n!)\, e^{-\lambda_m Dt}$

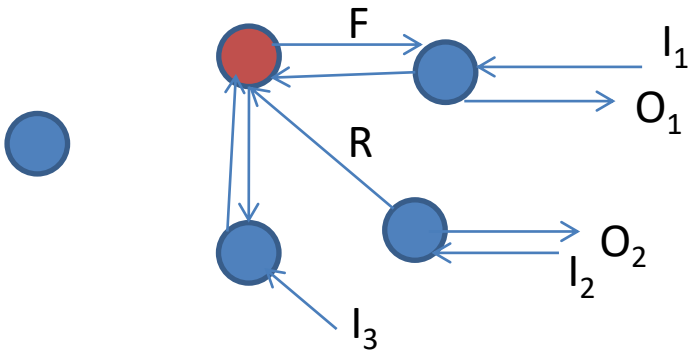

Input I: firing rate Output O: firing rate
R: input firing rate of the pyramidal to the Interneuron
F: input firing rate of the interneuron to the pyramidal
At each neuron, we have an Integrate-and-Fire model
The firing of the interneuron causes a suppression

State of a neuron is the membrane potential

What is the equilibrium state?

A neuron may give output to another layer… Not all neurons give output to another layer

Networks with Attractor States to Model **Associative Memory**

- Synaptic connectivity in a recurrent neural network (RNN) is set up in such a way that the network dynamics have multiple **attractor states**

Each attractor state:

- represents a **particular item**, stored in memory.
- is a specific pattern of activity of the network that is correlated with the state of the network when the particular item is presented through external inputs.

The **attractor property** means that the **network converges** to the stored pattern even when the **external inputs are correlated but not identical to the pattern** (necessary requirement for an associative memory model)

This is a *potential* way that the brain works to "resolve" the noise.

Learning & Retrieval in RNNs with Unsupervised Hebbian Learning Rules

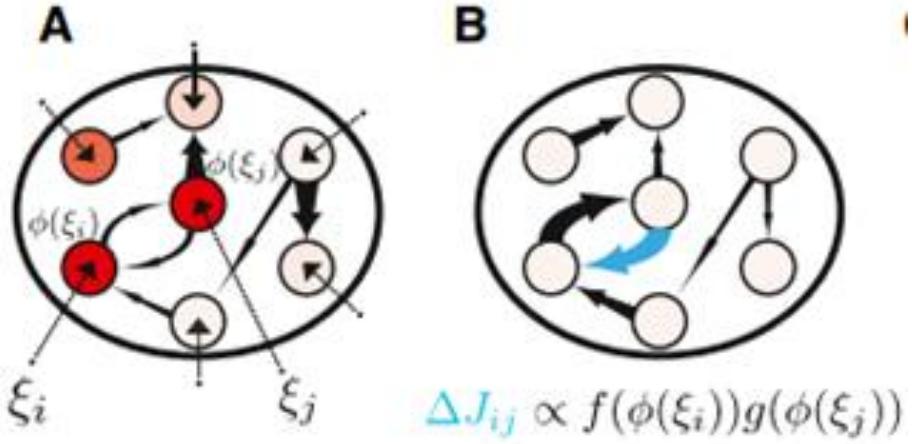

ξ: synaptic inputs to each neuron in the network

(B) The firing rate pattern produced by the synaptic input currents modifies the network connectivity according to **an unsupervised Hebbian Learning rule.**

The connection strength is represented by the thickness of the corresponding arrow (the thicker the arrow, the stronger the connection).

Synaptic inputs elicit firing rates through the static transfer function **ϕ(ξ)**

Some neurons **respond strongly** (red circles), **others weakly** (white circles)

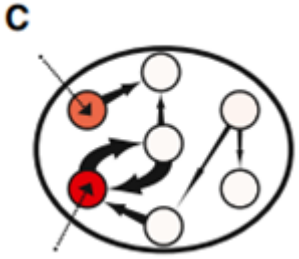

After learning, a pattern of synaptic inputs that is correlated but not identical to the stored pattern is presented to the network.

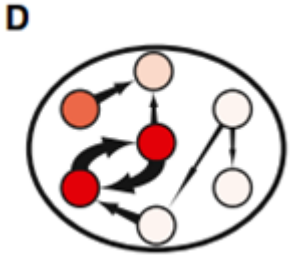

(D) Following the presentation, the network goes to an attractor state that strongly overlaps the stored pattern (compare with A), which indicates retrieval of the corresponding memory

a learning rule that changes the synaptic connectivity matrix by a factor $\Delta J_{ij} \propto f\left[\phi\left(\xi_i^\mu\right)\right] g\left[\phi\left(\xi_j^\mu\right)\right]$ when a pattern μ is presented to the network, starting from an initial *tabula rasa* $J_{ij} = 0$, and neglecting the contributions of recurrent connections during learning. This rule is a generalization of Hebbian rules used in classic models, such as the Hopfield model (Hopfield, 1982) or the Tsodyks-Feigel'man model (Tsodyks and Feigel'Man, 1988), with two important differences: patterns have a Gaussian distribution instead of binary, and the dependence of the rule on firing rates is non-linear instead of linear. In the following, the patterns that have shaped the connectivity matrix will be termed "familiar," whereas all other random patterns presented to the network will be termed "novel."

Non-linear functions, f & g, that characterize the dependence of the learning rule on the post-synaptic rate (f) and pre-synaptic rate (g), respectively.

Final connectivity after learning

$$J_{ij} = \frac{Ac_{ij}}{cN} \sum_{k=1}^{p} f\left[\phi\left(\xi_i^k\right)\right] g\left[\phi\left(\xi_j^k\right)\right]$$

where c_{ij} is a sparse random (Erdos-Renyi) structural connectivity matrix ($c_{ij} = 1$ with probability c, $c_{ij} = 0$ with probability $1 - c$, $c \ll 1$

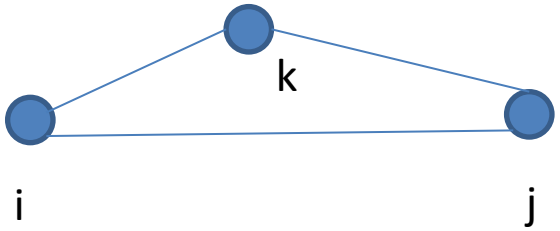

Non-linear functions, f & g, that characterize the dependence of the learning rule **on the post-synaptic rate (f)** and **pre-synaptic rate (g), respectively.**

Probabilistic Brains: Known & Unknowns

- An efficient, and under some circumstances optimal, way to perform tasks involving uncertainty is to **represent knowledge with probability distributions** and to **acquire new knowledge by following the rules of probabilistic inference**.

- Cox's theorem tells us that probability theory provides the only sensible and coherent way to reason under uncertainty.

- Experiments have shown that human behavior is **highly consistent with probabilistic reasoning** not only in the sensory domain, but also in the motor and cognitive domains.

- Although it is well-established that humans and monkeys (and other animals) perform probabilistic inference, it is less clear **how inference is implemented at the level of neural circuits.**

Pouget, A., Beck, J., Ma, W. et al. Probabilistic brains: knowns and unknowns. **Nat Neurosci** 16, 1170–1178 (**2013**).

One may even say, strictly speaking, that almost all our knowledge is only probable; and in the small number of things that we are able to know with certainty, the principle means of arriving at the truth -induction and analogy- are based on probabilities.

Pierre Simon Laplace. Theorie analytique des probabilites. 1825.

Cue Integration. Independent visual and haptic measurements (left) support to different degrees the three possible interpretations of object identity (middle). Integrating these sources of information according to their respective uncertainties provides an optimal probabilistic estimate of the correct object (right).

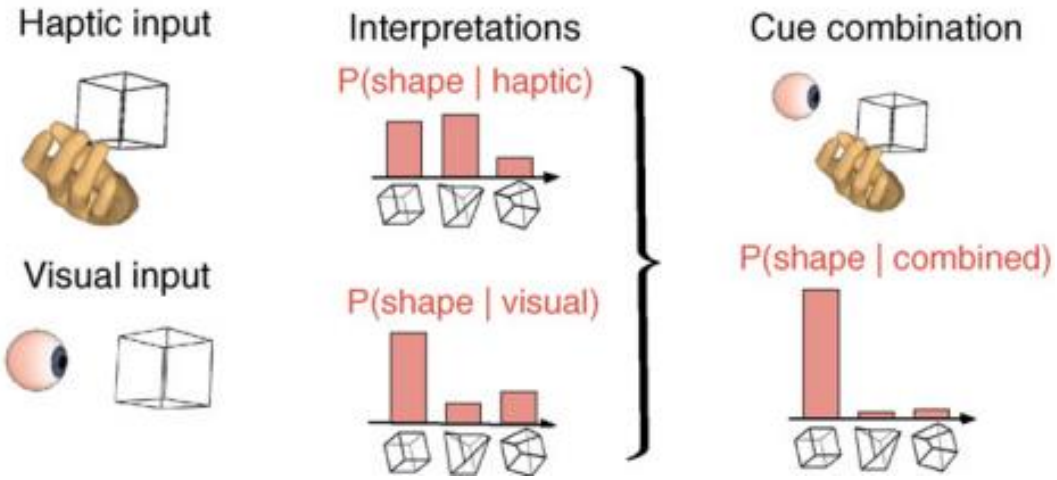

Statistically optimal perception and learning: from behavior to neural representations

József Fiser [1, 2] ✉, Pietro Berkes [1], Gergő Orbán [1, 3], Máté Lengyel [4]

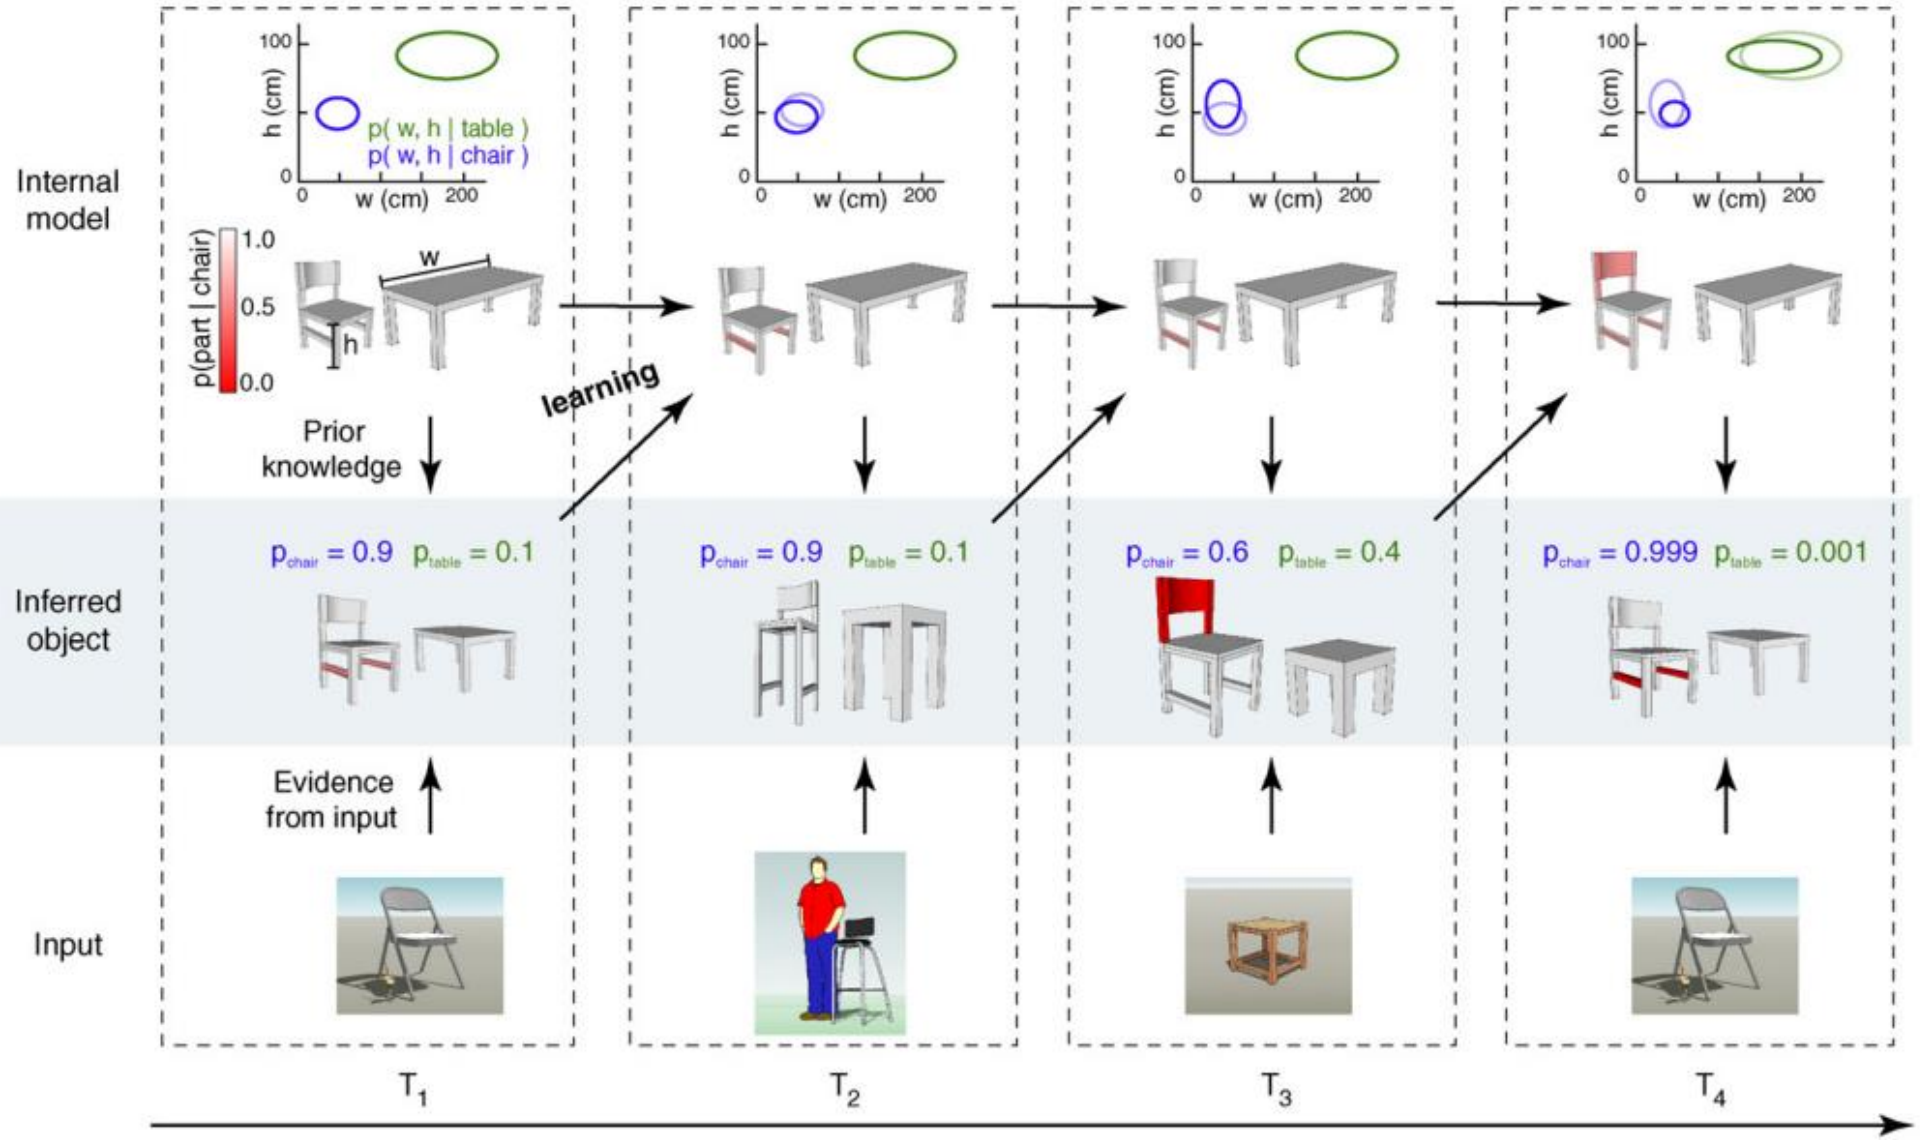

Statistically optimal perception and learning: from behavior to neural representations

Likelihood: the function specifying the probability $p(x|y,M)$ of observing a particular stimulus x for each possible state of the environment, y, under a probabilistic model of the environment, M.

Marginalization: the process by which the distribution of a subset of variables, y_1, is computed from the joint distribution of a larger set of variables, $\{y_1, y_2\}$: $p(y_1) = \int p(y_1, y_2)\, dy_2$. (This could be important if, for example, different decisions rely on different subsets of the same set of variables.) Importantly, in a sampling-based representation, in which different neurons represent these different subsets of variables, simply "reading" (e.g. by a downstream brain area) the activities of only those neurons that represent y_1 automatically implements such a marginalization operation.

Maximum *a posteriori* (or MAP) estimate: in the context of *probabilistic inference*, it is an approximation by which instead of representing the full *posterior* distribution, only a single value of y is considered that has the highest probability under the *posterior*. (Formally, the full *posterior* is approximated by a Dirac-delta distribution, an infinitely narrow Gaussian, located at its maximum.) As a consequence, *uncertainty* about y is no longer represented.

Posterior: the probability distribution $p(y|x,M)$ produced by *probabilistic inference* according to a particular probabilistic model of the environment, M, giving the probability that the environment is in any of its possible states, y, when stimulus x is being observed.

Prior: the probability distribution $p(y|M)$ defining the expectation about the environment being in any of its possible states, y, before any observation is available according to a probabilistic model of the environment, M.

Probabilistic inference: the process by which the posterior is computed. It requires a probabilistic model, M, of stimuli x and states of the environment y, containing a *prior* and a *likelihood*. It is necessary when environmental states are not directly available to the observer: they can only be inferred from stimuli through inverting the relationship between y and x through Bayes' rule: $p(y|x,M) = p(x|y,M)\ p(y|M)/Z$, where Z is a factor independent of y, ensuring that the *posterior* is a well-defined probability distribution. Note, that the *posterior* is a full probability distribution, rather than a single estimate over environmental states, y. In contrast with approximate inference methods, such as *maximum likelihood* or *maximum a posteriori* that compute single best estimates of y, the posterior fully represents the uncertainty about the inferred variables.

Probabilistic Inference for Multisensory Integration

Problem: Estimate the ***width*** of the object by touching it and via visual inspection

Assuming that the **noise** corrupting the visual & tactile measurements is **independent**

$$p(w|w_v, w_t) = \frac{p(w_v, w_t|w)p(w)}{p(w_v, w_t)} = \frac{p(w_v|w)p(w_t|w)p(w)}{p(w_v, w_t)}$$

Bayesian rule

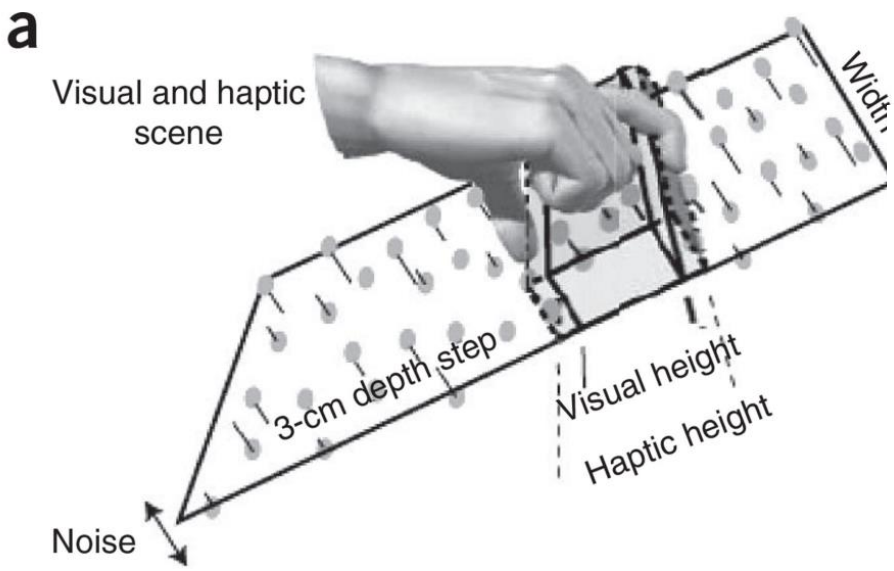

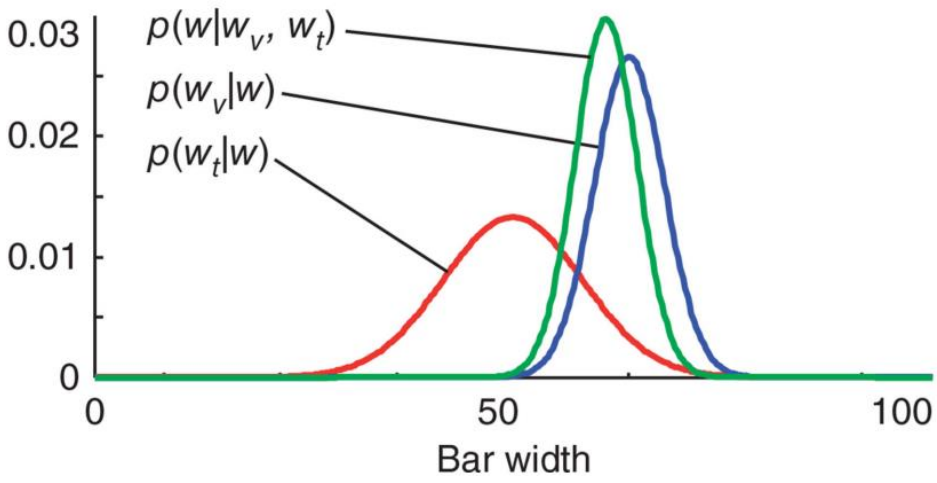

b) The posterior distribution over the width ($p(w|w_v, w_t)$, green curve) is proportional to the product of the visual ($p(w_v|w)$, blue curve) and haptic ($p(w_t|w)$, red curve) likelihood functions. Note that the posterior distribution is shifted toward the more reliable cue (the one with the smaller variance; in this case, vision).

Our intuition: the **mean of the posterior distribution** is a compromise between the mean obtained from **vision** & the mean obtained from **touch**, weighted by the inverse of the variance (that is, the precision) of each cue

$$\mu_{vt} = \frac{1/\sigma_v^2}{1/\sigma_v^2 + 1/\sigma_t^2} w_v + \frac{1/\sigma_t^2}{1/\sigma_v^2 + 1/\sigma_t^2} w_t$$

Combined variance is smaller than both the visual and the tactile variance—as it should, given that combining cues increases the information

$$\sigma_{vt}^2 = \frac{\sigma_v^2 \sigma_t^2}{\sigma_v^2 + \sigma_t^2}$$

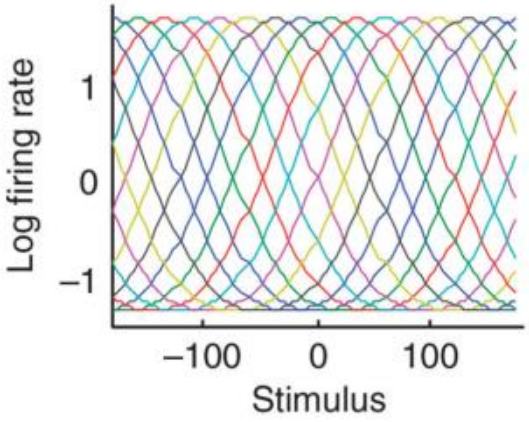

Basis functions: the log of the tuning curve, of **15 neurons,** to a periodic stimulus with value from −180 to 180

Pattern of spike counts: calculated over a 200-ms interval, across the same neuronal population in **response to a stimulus** with value 0

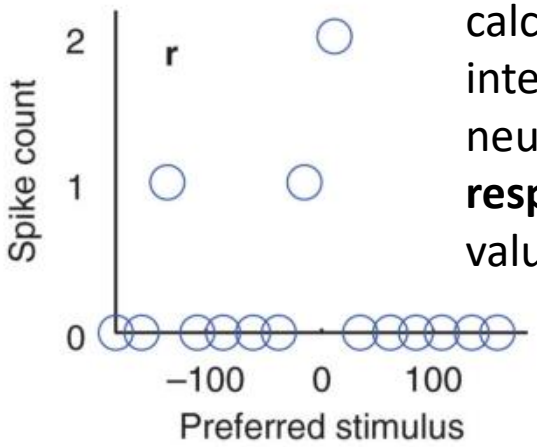

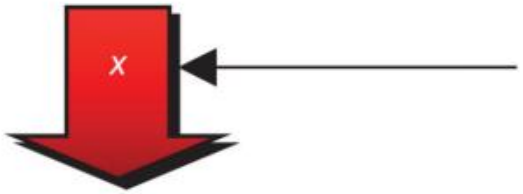

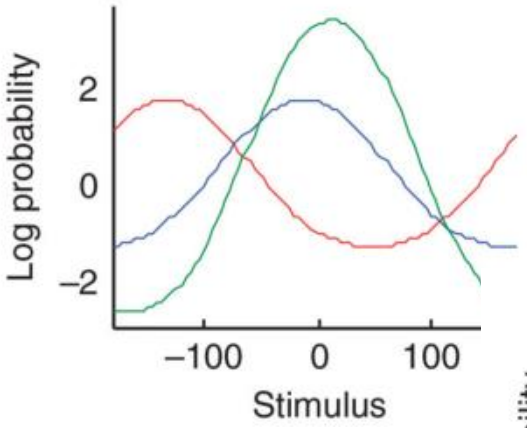

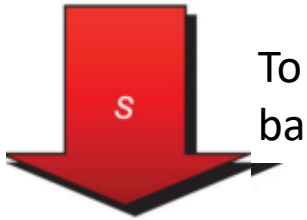

To turn **spike counts into log prob**, multiply each basis function by its corresponding spike count.

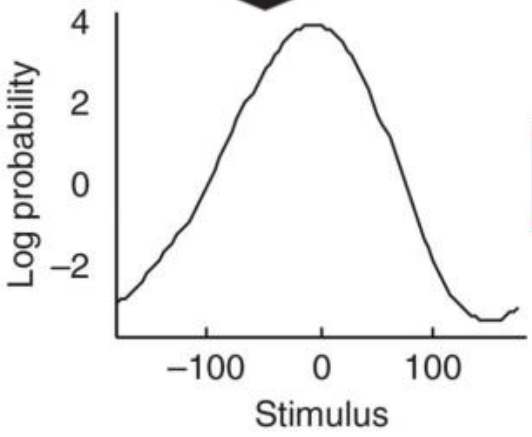

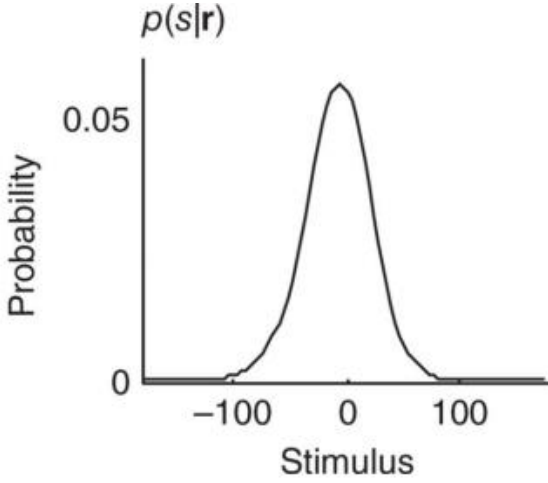

$$p(w|w_v, w_t) = \frac{p(w_v, w_t|w)p(w)}{p(w_v, w_t)} = \frac{p(w_v|w)p(w_t|w)p(w)}{p(w_v, w_t)}$$

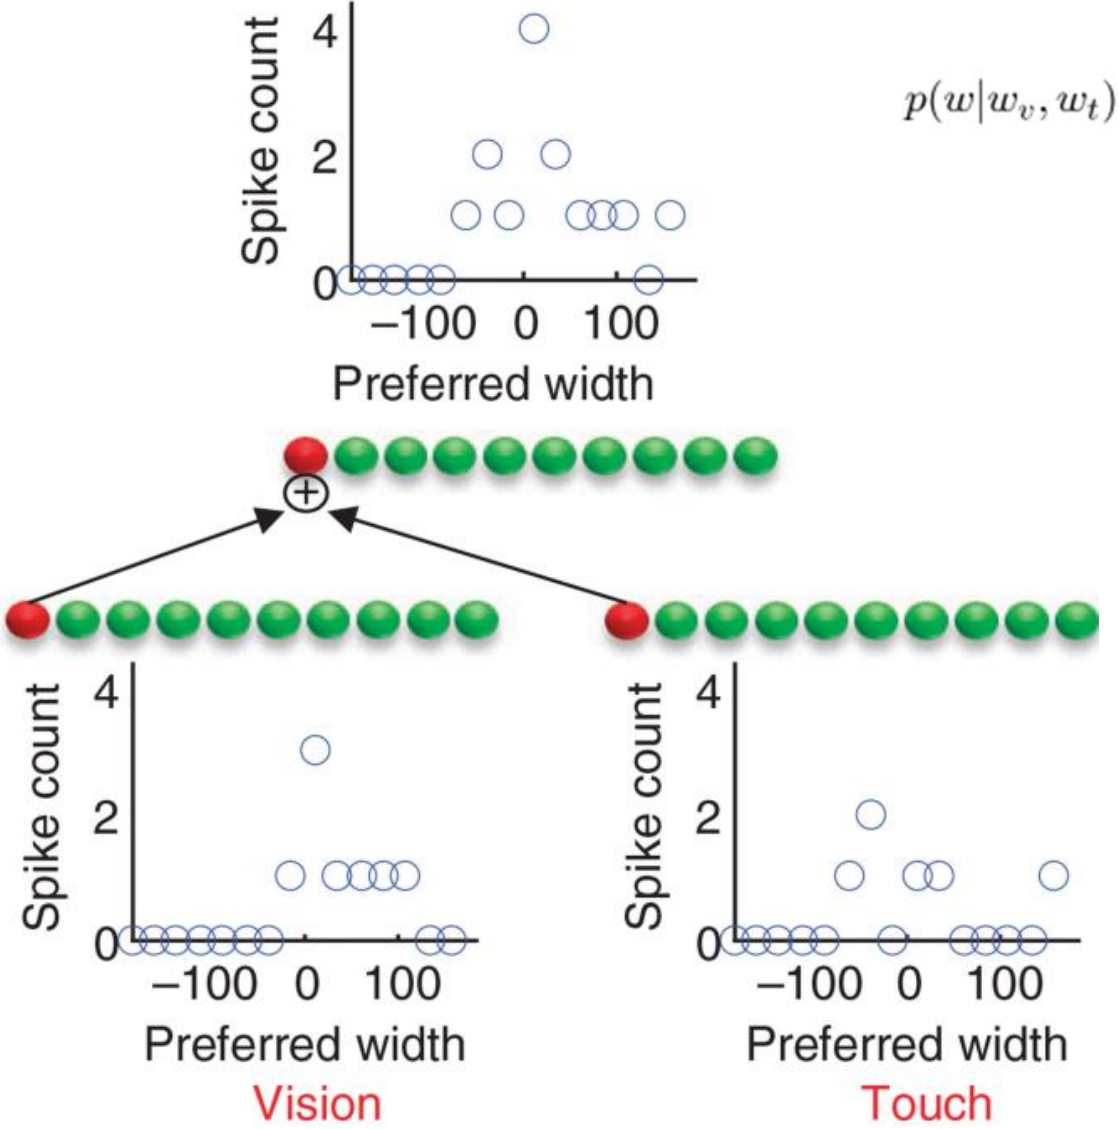

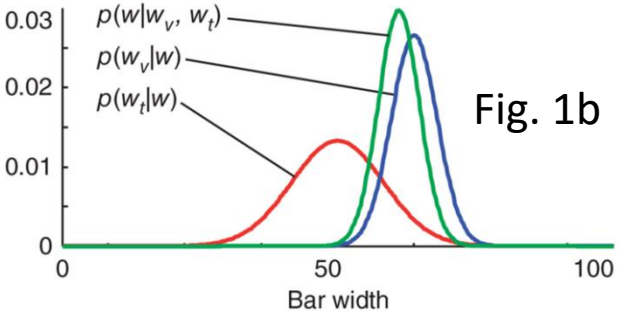

Fig. 1b

Taking a product of likelihood functions with probabilistic population codes. Bottom panels, probabilistic population codes for the two likelihoods shown in Figure 1b (the blue and red curves). Summing the two population codes (neuron by neuron) yields a population code (top) for the product of the two likelihoods (the green curve in Fig. 1b), as required for optimal multisensory integration (equation (1)).

Multisensory Integration – Probabilistic Inference

- Compute probability distributions over variables of interest **s** given **sensory measurements I** and **prior knowledge p(s).**
- In probabilistic models, the **variable s** is referred to as a **latent** variable (the width of the bar in the example)

Probabilistic inference starts with the generative model, a statistical model of how the measurements, I, are generated (which has to be learned by the animal). The generative model consists of a prior distribution $p(s)$ and a distribution $p(I|s)$ (known as the likelihood function when viewed as a function of s; Box 1). In the previous example, the prior, $p(w)$, was assumed to be flat, and the likelihood functions corresponded to the functions $p(w_v|w)$ and $p(w_t|w)$. Bayes' rule then provides a recipe for formulating beliefs about s, in the form of the posterior distribution

$$p(s|I)=\frac{p(I|s)p(s)}{p(I)} \quad (4)$$

Bayesian logic

The denominator, $p(I)$, ensures that the posterior distribution integrates to 1.

Encoding Probabilities with Neurons

- Several groups have proposed that neural activity encodes functions of latent variables, as opposed to single values.
- In the probabilistic framework, these functions are either probability distributions or likelihood functions. In that case, neural computations must manipulate whole functions, and must do so according to the rules of probabilistic inference.

Proposed Models of Neuronal Codes (1/2)

- The **response of a neuron** tuned to a particular image feature (e.g., the orientation of a contour) is proportional to the ***log* of the probability that the feature is present in the neuron's receptive field** (Barlow)
- Neuronal responses are proportional to the **probability rather than to its log** (Anastasio)
- Neuron codes the **log probability that a feature takes on a *particular value***
- Probability distributions are ***functions***, and, as such, can be encoded using a variety of techniques, e.g., as a sum of other functions, where the **coefficients would be encoded by neural activity**

Note: For a code that uses probability, **adding probabilities** is easy, whereas, for one that uses **log probabilities, multiplying them** is easy.

As both addition & multiplication are key steps in probabilistic inference, neither code has an obvious advantage over the other

Basis Functions

- A common one is to express functions as the sum of other functions (called ***basis functions*** in this context)

  e.g., radial basis functions

  in Fourier analysis, a function is expressed as a linear combination of sines & cosines.

- With the basis function approach, probability distributions would be represented as a **set of coefficients** and the coefficients would be encoded by neural activity

$h_i(s)$ are the basis functions and the constant is needed to ensure proper normalization

Linear Probabilistic Population Codes

Experimental data[51,52] suggest that $p(\mathbf{r}|s)$ belongs to a family of distributions known as the exponential family with linear sufficient statistics[47], leading to the code shown in equation (5) if the prior is flat. Thus, linear probabilistic population codes have the advantage that they are consistent with the statistics of neural responses. Moreover, as the $h_i(s)$ can be any set of functions of s, equation (5) can represent virtually any posterior distribution, $p(s|\mathbf{r})$.

$$\log p(s|\mathbf{r}) = \sum_i r_i h_i(s) + \text{constant} \qquad (5)$$

p(r|s) is the distribution of **neural variability**: the variability in spike counts in response to repeated presentations of the same stimulus **(s)**

Assuming a flat prior, Bayes' rule tells us that $p(s|\mathbf{r}) \propto p(\mathbf{r}|s)$

Proposed Models of Neuronal Codes (2/2)

- Brain may represent probability distributions by the values of a set of samples drawn from the encoded distribution

  e.g., spikes represent samples from a distribution over **binary** random variables (r.v.), whereas the membrane potential values represent samples from a probability distribution over **real-valued** r.v.

- Whether this type of code is mutually exclusive or complementary to other codes is still being debated

Neural Implementation of Probabilistic Inference

1. **Combining multiple sources of information**

e.g., in the multisensory experiment, the posterior distribution over the width of the bar is the product of the visual & haptic likelihood functions

- Can be generalized to the problem of accumulating evidence ***over time*** (in decision-making), instead of ***across modalities***
- Consistent with responses of neurons in areas, e.g., lateral intraparietal cortex, when they are accumulating information about direction of motion

2. **Marginalization** Recovering the **distribution over a variable x, p(x),** from **a *joint* distribution over x and other variables**, e.g., p(x, y, z)

- Involves sums of probabilities and is implemented by adding neural activities

Marginals: e.g., p(x | y, z)

Note: It is easier to compute p(x |y, z) from samples than p(x, y, z) (dimensionality problem)

3. **Estimation of the maximum a posteriori estimate**. Given a posterior distribution p(s|r), estimate the **value of s** corresponding to the peak of this distr. (i.e., the most probable value of s given the neural activity)

Implemented using an attractor network

Probabilistic computation	Linear Probabilistic Population codes	Codes proportional to probabilities	Sampling-based codes
Evidence integration: Cue combination, temporal accumulation of evidence for decision making	Linear: sums across populations or over time	Nonlinear: products	Nonlinear: products of histograms of samples
Estimation: Maximum likelihood	Nonlinear: attractor dynamics	Nonlinear: Winner Take All	Nonlinear: avg of samples
Kalman filtering Motor control, visual object tracking	Non-linear: quadratic nonlinearity with divisive normalization	Nonlinear	Nonlinear: particle filters

	Neural implementation		
robabilistic computation	**Linear probabilistic population codes**	**Codes proportional to probabilities**	**Sampling-based codes**
vidence integration (for example, cue mbination, temporal accumulation of idence for decision-making)	Linear: sums across populations[47] or over time[56]	Nonlinear: products	Nonlinear: products of histograms of samples[5]
timation (for example, maximum kelihood)	Nonlinear: attractor dynamics[61,62]	Nonlinear: winner take all	Nonlinear: average of samples[34,53]
alman filtering (for example, for otor control, visual object tracking)	Nonlinear: quadratic nonlinearity with divisive normalization[58]	Nonlinear[63,64]	Nonlinear: particle filte
mple marginalization (for example, ear coordinate transforms)	Nonlinear: quadratic nonlinearity with divisive normalization[58]	Linear[63,64]	Linear: sums over histo
corporating prior knowledge	Nonlinear: bias current[47]	Nonlinear: products	Nonlinear: products of histograms of samples[5]
pproximate high dimensional ference (for example, olfactory ocessing)	Nonlinear: for example, divisive normalization[60]	Nonlinear: products and sums[32]	Nonlinear: Monte Carlo sampling[53]

Related Lectures

https://www.youtube.com/watch?v=KqqHJrs74_c

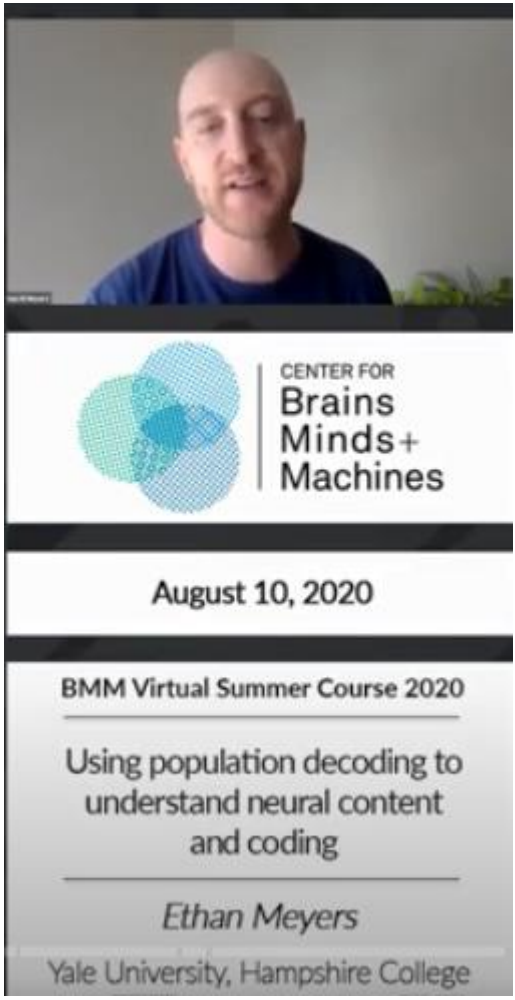

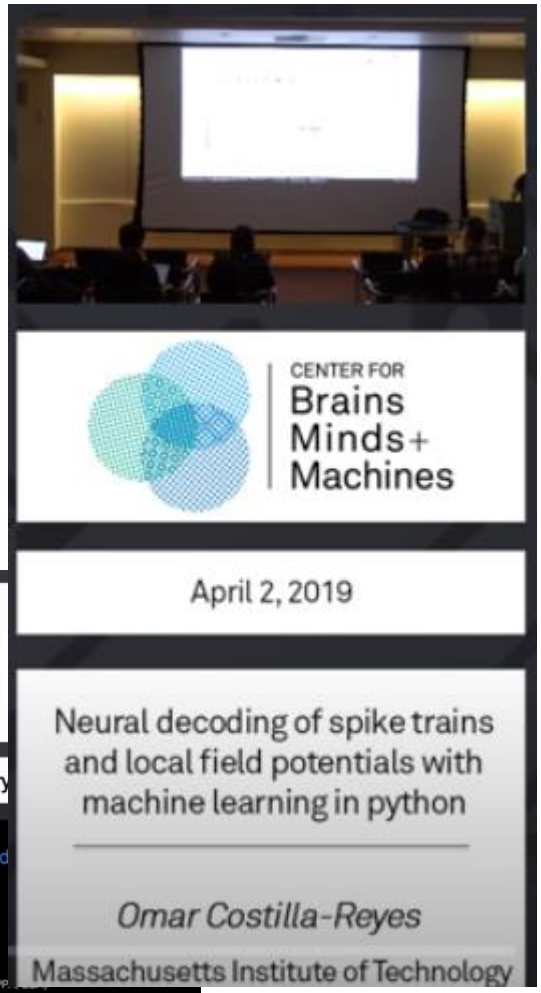

https://www.youtube.com/watch?v=OYDIy5KgNKo

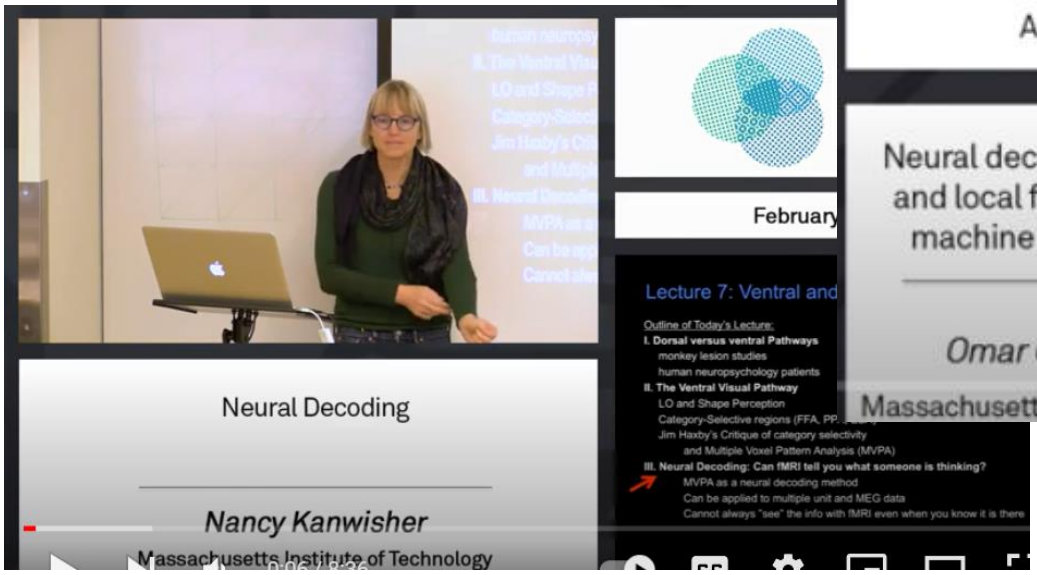

https://www.youtube.com/watch?v=HDk1hczPky4