Modelling the Dynamics of Sensory Neural Responses

For low-level modelling of the brain

Ulysse Rançon
ulysse.rancon@cnrs.fr

22nd Oct. 2024

My background

now !

2016 — Engineering school — 5 Years — Lyon

2021 — PhD #1 (quit) — 1 Year — Bordeaux

2022 — PhD #2 (now) — 3 Years — Toulouse

2025

INSA INSTITUT NATIONAL DES SCIENCES APPLIQUÉES LYON

université de BORDEAUX

cnrs

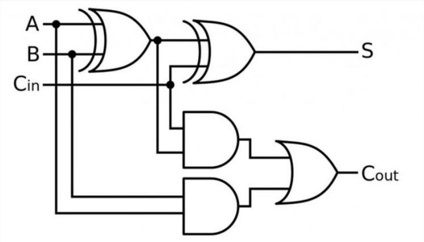

Outline

I. What is a model ?

II. Models of neurons / Spiking Neural Networks (SNNs)

III. Neuromorphic computing and sensing

IV. Conclusion

Little to no maths

Large overview

Main keywords

Starting point for self-study !

What is modelling (“modélisation”) ?

Taking a step back

A very difficult question : *What is a model ?*

A very difficult question : *What is a model ?*

Essentially a **description of reality**
A “model” can be… just ***anything*** !

A very difficult question : *What is a model ?*

Essentially a **description of reality**
A “model” can be… just ***anything*** !

“The more I drink coffee, the less I sleep”

even just ***an idea****…!*

A very difficult question : *What is a model ?*

Essentially a **description of reality**
A “model” can be… just ***anything*** !

$$\Delta U = Q - W$$

*a **physics formula**…*

*a **live animal**…*

“The more I drink coffee, the less I sleep”

*even just **an idea**…!*

***a computer simulation** of many small agents thinking they have a life on their own…*

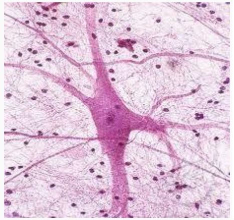

***in vitro neurons** in a petri dish, extracted from animals (therefore killing them)…*

A very difficult question : *What is a model ?*

Essentially a **description of reality**
A “model” can be… just ***anything*** !

“The more I drink coffee, the less I sleep”

even just ***an idea****…!*

a computer simulation *of many small agents thinking they have a life on their own…*

A very difficult question : *What is a model ?*

Essentially a **description of reality**
A "model" can be… just ***anything*** !

$$\Delta U = Q - W$$

*a **physics formula**…*

> ***"The more I drink coffee, the less I sleep"***

*even just **an idea**…!*

__a computer simulation__ of many small agents thinking they have a life on their own…

A very difficult question : *What is a model ?*

Essentially a **description of reality**
A “model” can be… just ***anything*** !

“The more I drink coffee, the less I sleep”

*even just **an idea**…!*

$$\Delta U = Q - W$$

*a **physics formula**…*

*a **live animal**…*

***a computer simulation** of many small agents thinking they have a life on their own…*

A very difficult question : *What is a model ?*

Essentially a **description of reality**

A “model” can be… just ***anything*** !

“*The more I drink coffee, the less I sleep*”

*even just **an idea**…!*

$$\Delta U = Q - W$$

*a **physics formula**…*

*a **live animal**…*

*a **computer simulation** of many small agents thinking they have a life on their own…*

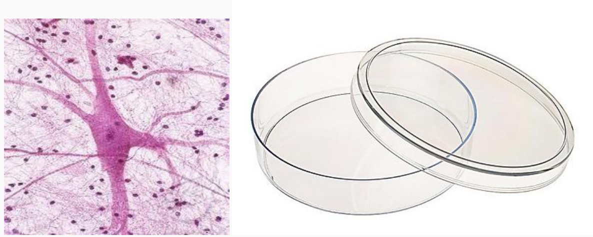

***in vitro neurons** in a petri dish, extracted from animals (therefore killing them)…*

How to choose my model ?

Depends on what you are working on…

"The more I drink coffee, the less I sleep"

*even just **an idea**…!*

NO !

$$\Delta U = Q - W$$

*a **physics formula**…*

How to choose my model ?

Depends on what you are working on…

"The more I drink coffee, the less I sleep"

even just ***an idea****…!*

NO !

$$\Delta U = Q - W$$

a ***physics formula****…*

A GOOD model:

- *provides a* ***quantitative description*** *of a feature, a phenomenon*
- *a* ***qualitative*** *description*
- ***inside and outside the scope*** *of what served to conceive it*

It must be also:

- *as* ***simple*** *as possible (Occam's razor)*
- *easily* ***expandable***

Abstraction levels when modelling

…And how much detail you want to capture !

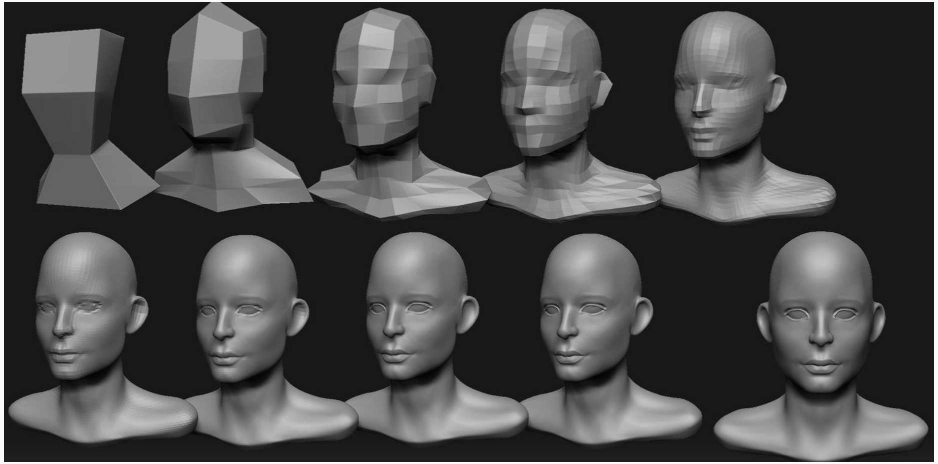

"The more I drink coffee, the less I sleep"

↓

Because of caffeine

↓

Theory of how caffeine is digested →

Depends on many factors:

- ***amount in the coffee cup***
- ***person constitution***
- ***time of the day/night***
- ***…***

…

…

The essence of modelling: Qiestion

"All models are wrong, but some are __(?)__"

George Box (1919-2013)

"What I cannot __(?)__, I cannot understand"

Richard Feynman (1918-1988)

The essence of modelling: Answers

*“All models are wrong, but some are **useful**”*

George Box (1919-2013)

*“What I cannot **create**, I cannot understand”*

Richard Feynman (1918-1988)

Models of Biological Neurons /

Spiking Neural Networks (SNNs)

According to you, what does this mean ?

A short history of computational neuroscience


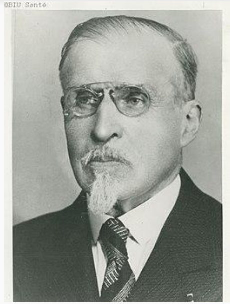
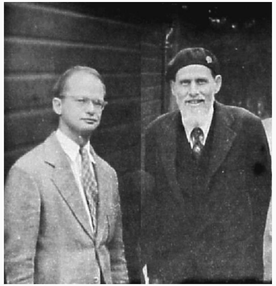
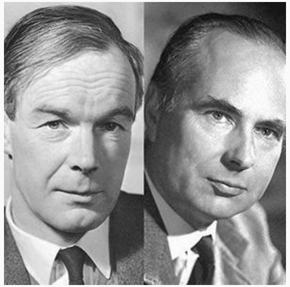

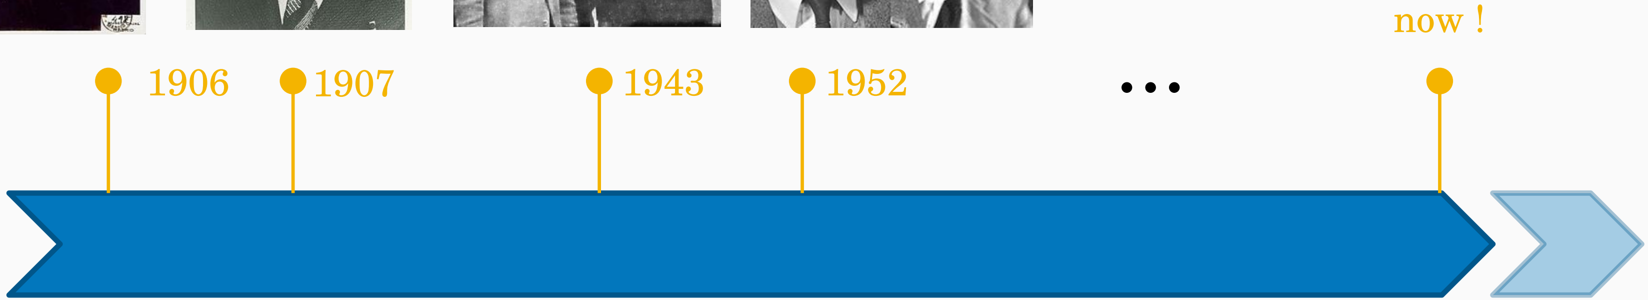

Ramon y Cajal
discovery of neurons

Louis Lapicque
Integrate-and-Fire
(IF) model

McCulloch & Pitts
the Perceptron

Hodgkin & Huxley
Eponimous model

The McCulloch & Pitts model

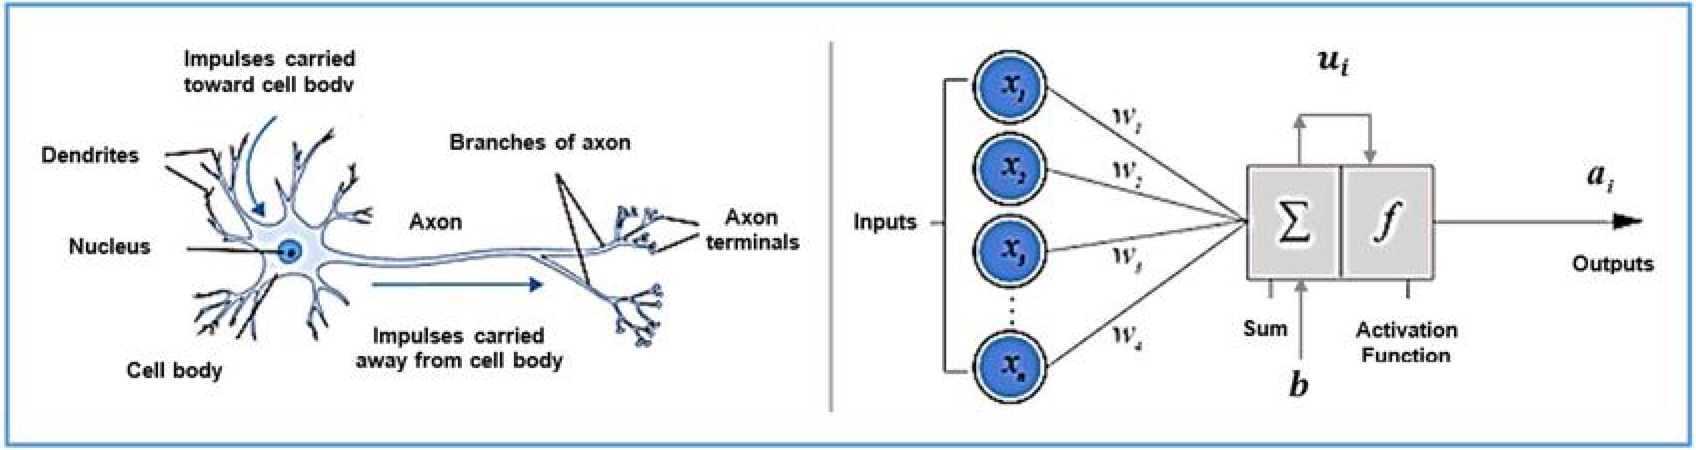

Supposed to mimick how actual neurons work:

- receive **weighted contributions** of several presynaptic inputs (electrical currents)
- **integrates** (i.e. sum) them
- outputs a (nonlinear) activation

The McCulloch & Pitts model

Q: Who can write the equation for this model ?

The McCulloch & Pitts model

Q: Who can write the equation for this model ?

$$\boxed{a = f(\sum_i w_i x_i + b)}$$

The McCulloch & Pitts model

Q: Who can write the equation for this model ?

$$a = f\left(\sum_i w_i x_i + b\right)$$

in previous classes / in deep learning, f was / is a **sigmoid** or **ReLU**

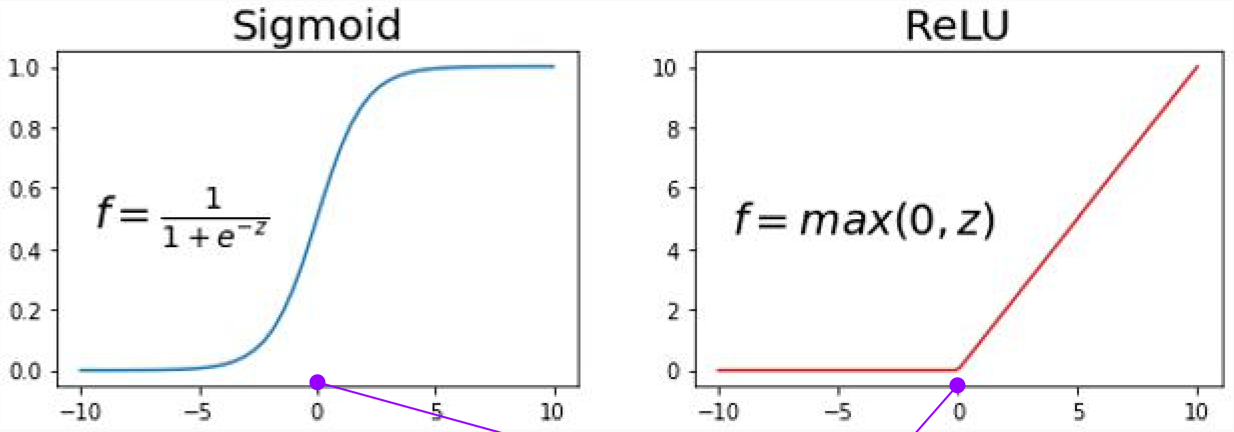

change of behavior

The McCulloch & Pitts model

Q: Who can write the equation for this model ?

$$a = f(\sum_i w_i x_i + b)$$

in the MCP model, it is a **Heaviside** function !

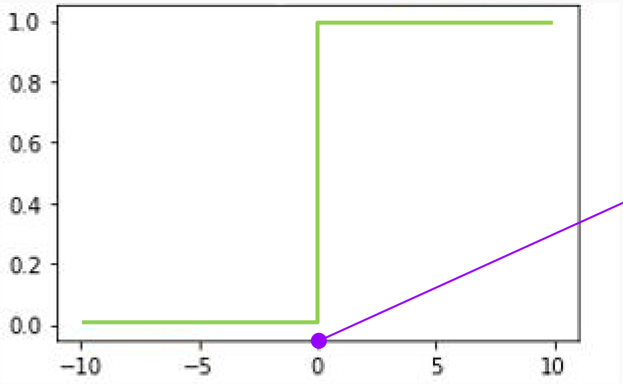

change of behavior (again) at a certain **threshold** θ (here: 0)

The output is **All or nothing** / 1 or 0 !

Spikes / Action potentials

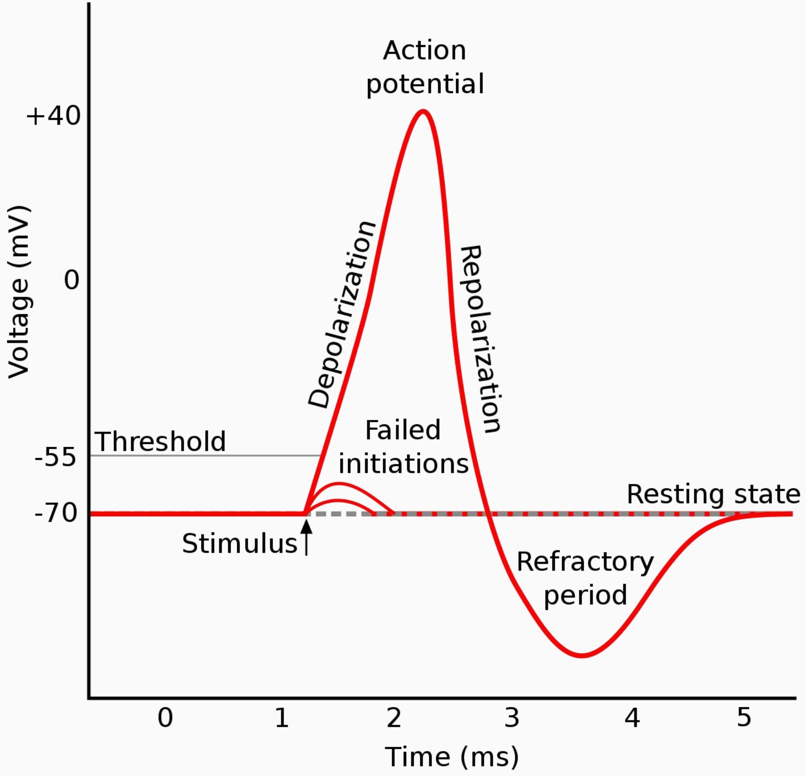

Real biological neurons emit **spikes**: stereotyped electrical impulses ("all-or-none") emitted when a neuron is <u>sufficiently</u> stimulated

Spikes propagate along the **axons** without attenuation ("active propagation").

Neurons can only exchange information via <u>binary spikes</u>

Spikes / Action potentials

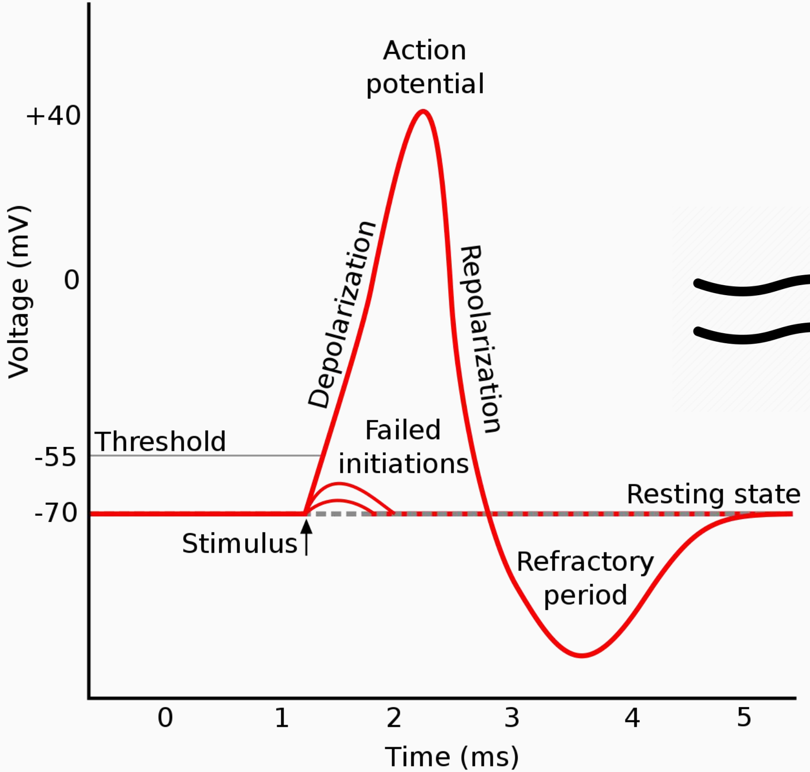

≈

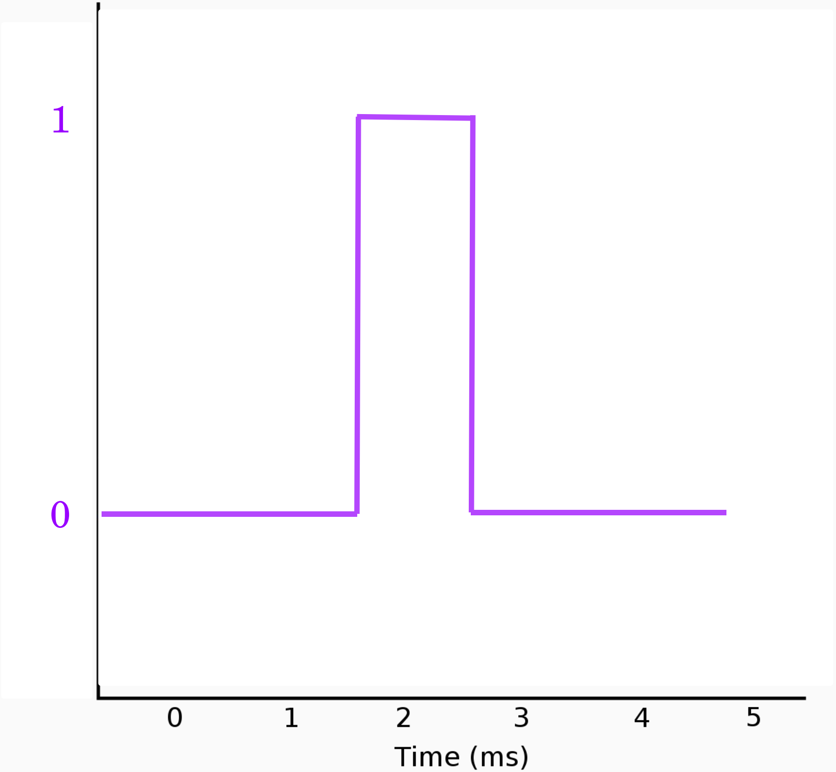

Spikes / Action potentials

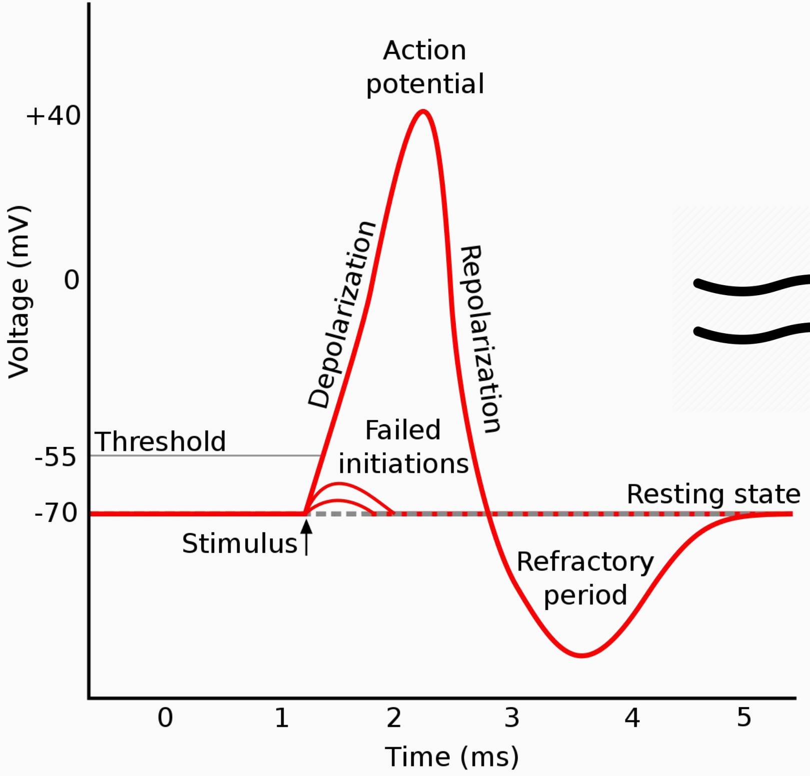

≈

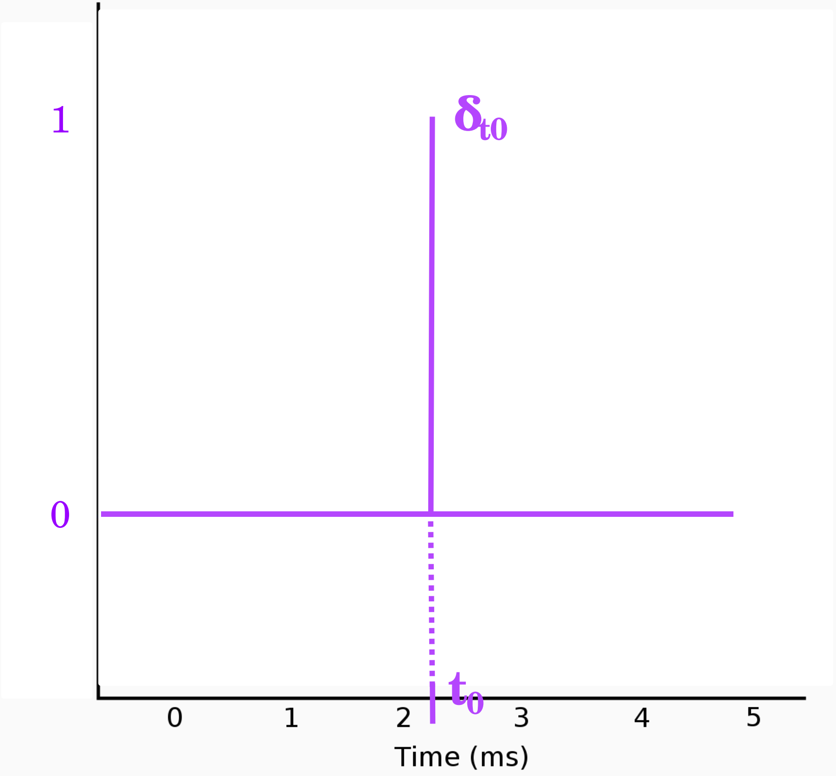

The Integrate and Fire (IF) neuron model

What was lacking in the MCP model, is a **memory trace**

The state of a neuron at a given instant directly depends on its previous state (makes sense, no ?) , i.e. they are **recurrent**

In reality, biological neurons are characterized by the **membrane potential** $\boldsymbol{V}$ which has a **resting value** $\boldsymbol{V_{rest}}$

Let's complexify the MCP model a bit:

(Potential update) $$V_t = V_{t-1} - \sum_i w_i x_i$$

(Spike at time t) $$S_t = H(V_t - \theta)$$

(Reset condition) $$V_t > \theta \longrightarrow V_t = V_{rest}$$

Integrate-and-Fire (IF) spiking model (discrete time equation)

Complexifying the IF with a refractory period

Real neurons have a **refractory period**:
their membrane potential is maintained to V_{rest} *during T time-steps after a spike*

$$V_t = V_{t-1} - \sum_i w_i x_i$$

$$S_t = H(V_t - \theta)$$

$$V_t > \theta \longrightarrow V_t = V_{rest}$$

(Refractoriness) $\boxed{t - t_{\text{last spike}} < T \longrightarrow V_t = V_{rest}}$

refractory IF

Towards the Leaky-Integrate and Fire (LIF) neuron model

In real neurons, the membrane potential **leaks** back to its rest value :
their membrane potential is maintained to V_{rest} during T time-steps after a spike

refractory LIF

$$V_t = V_{t-1} + \sum_i w_i x_i \boxed{- \frac{1}{\tau} V_{t-1}}$$

$$S_t = H(V_t - \theta)$$

$$V_t > \theta \longrightarrow V_t = V_{rest}$$

$$t - t_{\text{last spike}} < T \longrightarrow V_t = V_{rest}$$

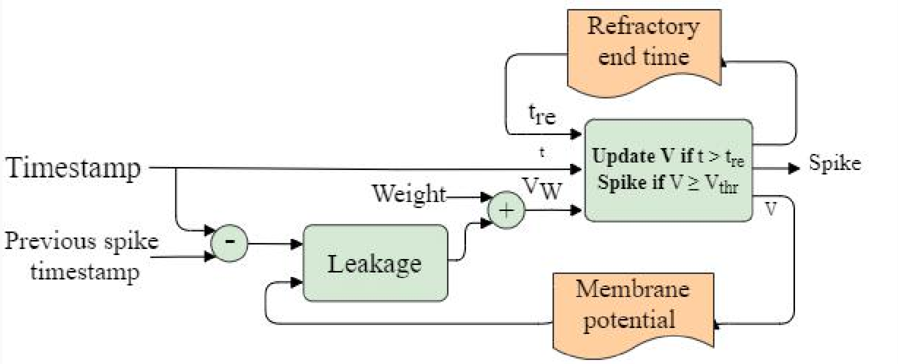

Example LIF behavior

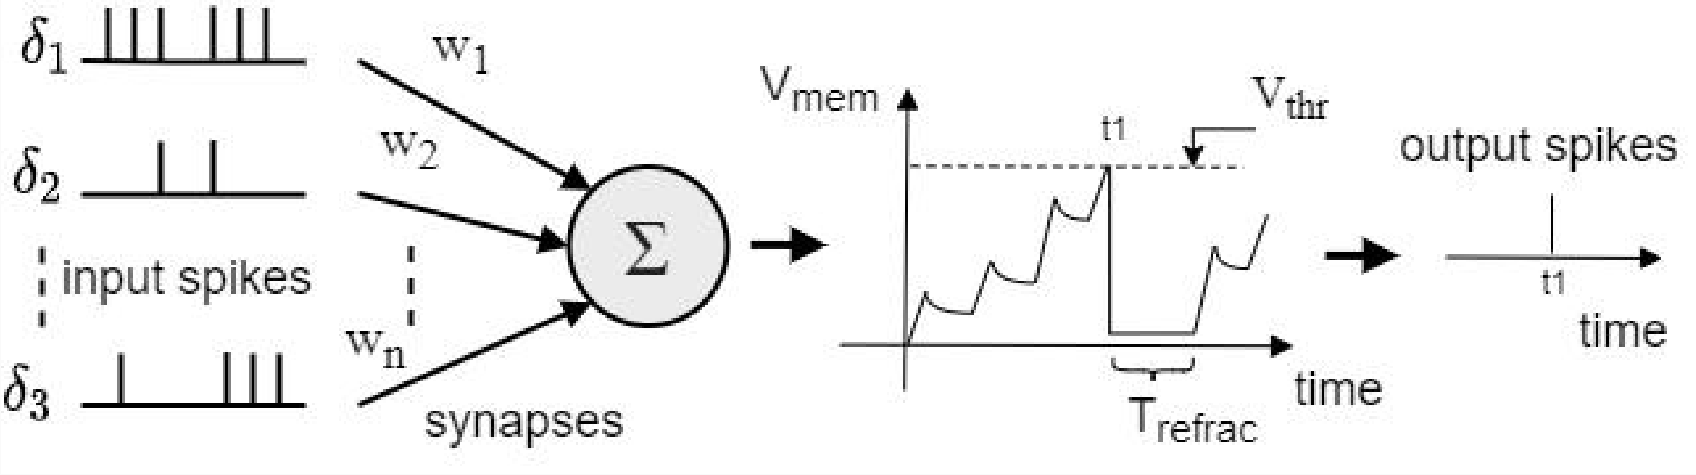

Electrical analogy

The LIF model comes from **biophysics** !

- Biological neurons are **electrically charged**: difference in potential between the inside and the outside of the cell = membrane potential
- **Ion channels** act as a Resistance
- Act as a **RC circuit** (Resistive & Capacitive)

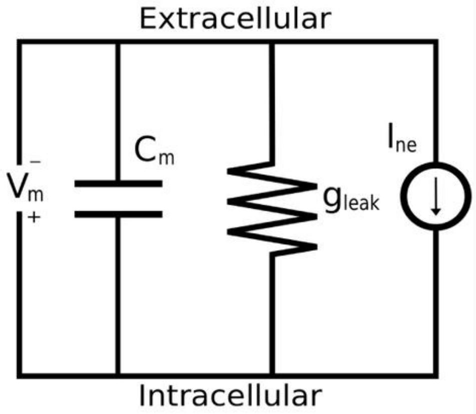

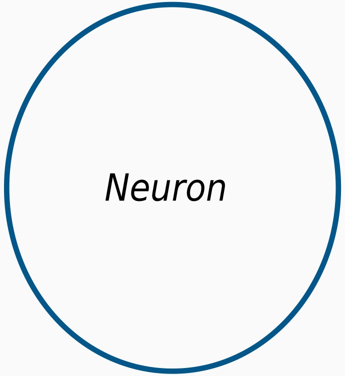

The Hodgkin-Huxley model

1952:
- H&H experiment on the giant squid axon
- Elaboration of a model with 3 ionic channels

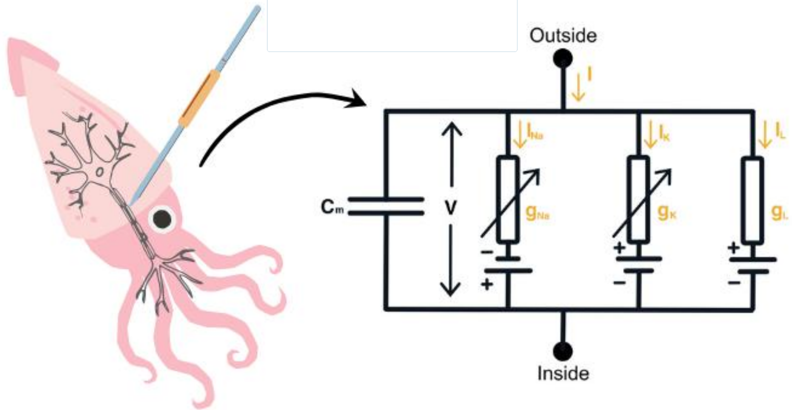

https://perceptron.blog/hodgkin-huxley/

1963:
Nobel prize in Physiology or Medicine

The Hodgkin-Huxley model

$$C\frac{dV}{dt} = -G_L(V - V_L) - G_{Na}m^3h(V - E_{Na}) - G_K n^4(V - E_K) + I_e, \tag{13}$$

where V is the membrane potential, I_e is an external current and the gating variables m, n and h obey the first-order ODEs

$$\frac{dm}{dt} = \alpha_m(V)(1-m) - \beta_m(V)m, \tag{14}$$

$$\frac{dh}{dt} = \alpha_h(V)(1-h) - \beta_h(V)h, \tag{15}$$

$$\frac{dn}{dt} = \alpha_n(V)(1-n) - \beta_n(V)n, \tag{16}$$

with transition rates $\alpha(V)$ and $\beta(V)$ given by

$$\alpha_m(V) = \frac{0.1(V+40)}{1-e^{-0.1(V+40)}}\ ; \quad \beta_m(V) = 4e^{-0.0556(V+65)}, \tag{17}$$

$$\alpha_h(V) = 0.07e^{-0.05(V+65)}\ ; \quad \beta_h(V) = \frac{1}{1+e^{-0.1(V+35)}}, \tag{18}$$

$$\alpha_n(V) = \frac{0.01(V+55)}{1-e^{-0.1(V+55)}}\ ; \quad \beta_n(V) = 0.125e^{-0.0125(V+65)}. \tag{19}$$

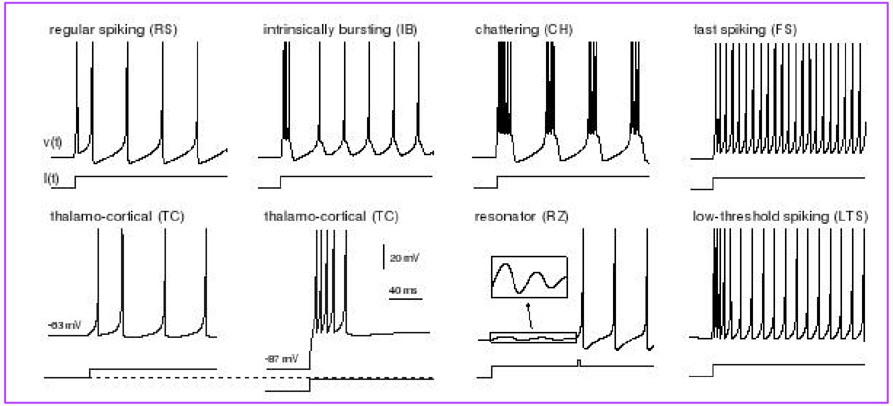

A very complex and precise model…

…capable of much more than the LIF !

Point vs Compartmental models

Compartmental models take into account the neuron's ***2d or 3d morphology****.*

- **discretize** the axon and dendrites into N small compartments
- presynaptic neurons can be anywhere along these axes

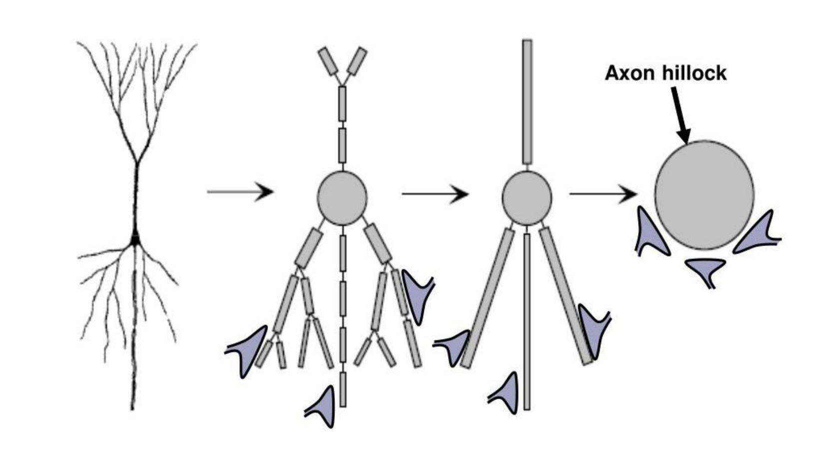

Point vs Compartmental models

*Compartmental models take into account the neuron's **2d or 3d morphology**.*

- **discretize** the axon and dendrites into N small compartments
- presynaptic neurons can be anywhere along these axes
- apply the **laws of physics** (wave equation, telegrapher's equations)

--> propagation of the AP !!!

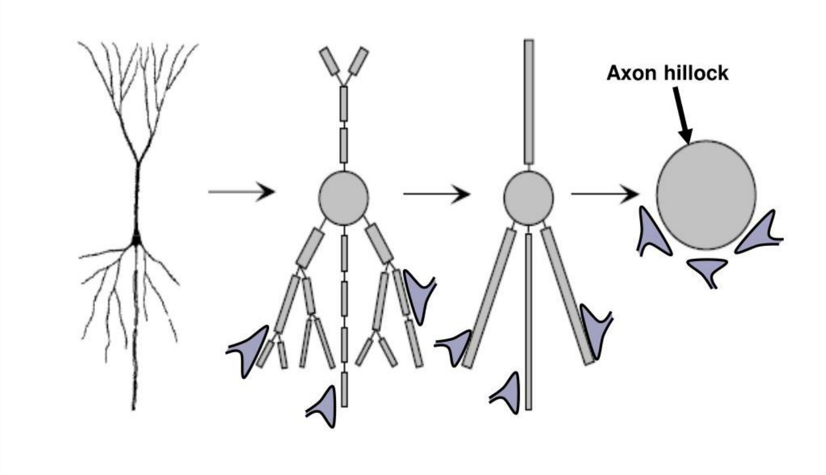

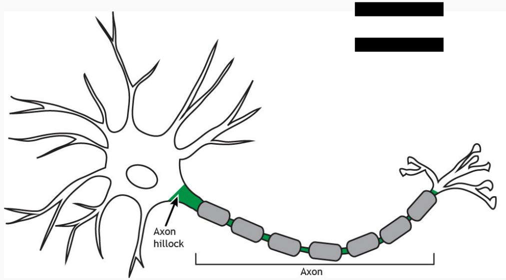

Raster plots and derivatives

<u>Raster plot</u>
To represent the simultaneous spiking activity (spike trains) of:
- several neurons in a population
- a single neurons over repeated trials of the same condition

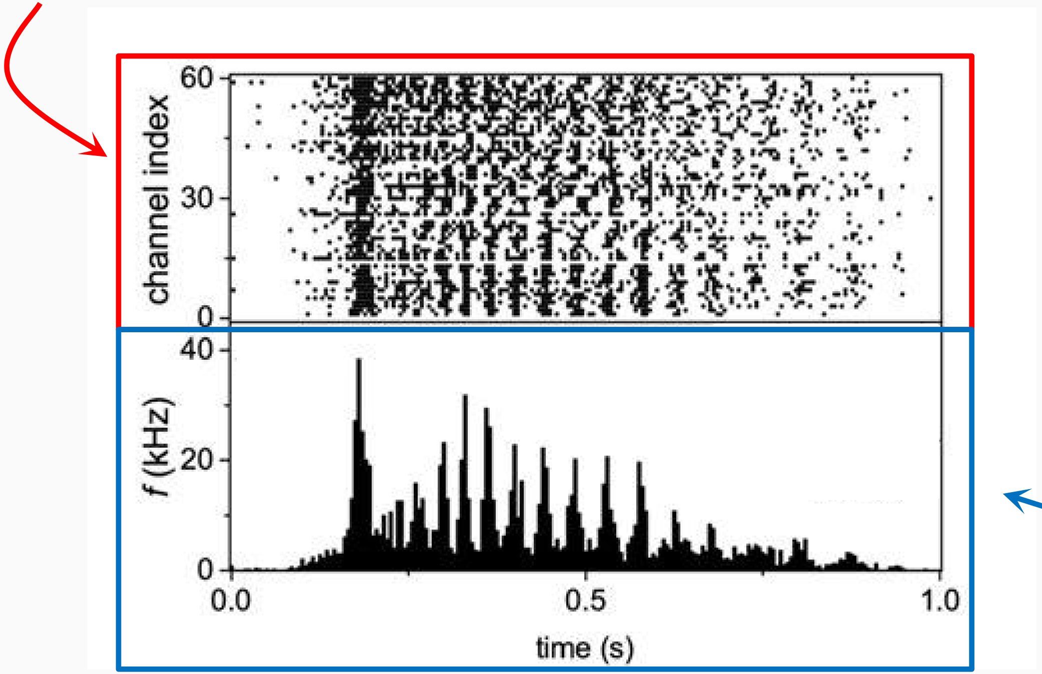

<u>Spike histogram</u>
Sum or average spikes over channels to give:
- population activity
- the target neuron's firing probability or firing rate

Coincidence detectors / Spike patterns

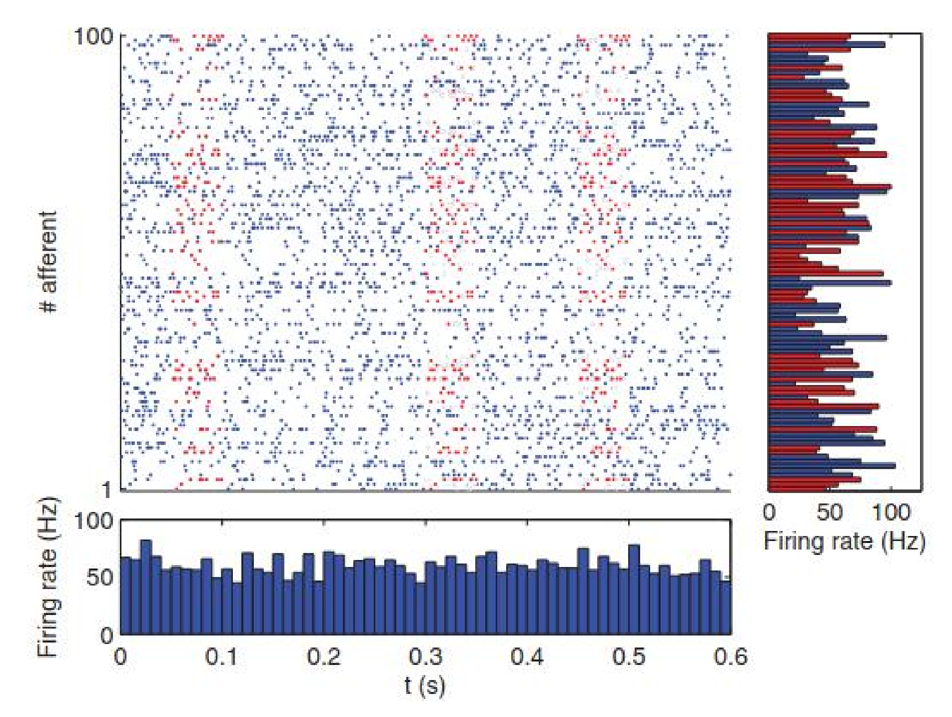

For a given action / thought / stimulus, a **pattern of spikes** are repeated, with more or less **variability**.

Downstream (post-synaptic) neurons learn to **detect** (spike in response to) these patterns.

Masquelier et al. (2008) *Spike Timing Dependent Plasticity Finds the Start of Repeating Patterns in Continuous Spike Trains, PLoS One*

Coincidence detectors / Spike patterns

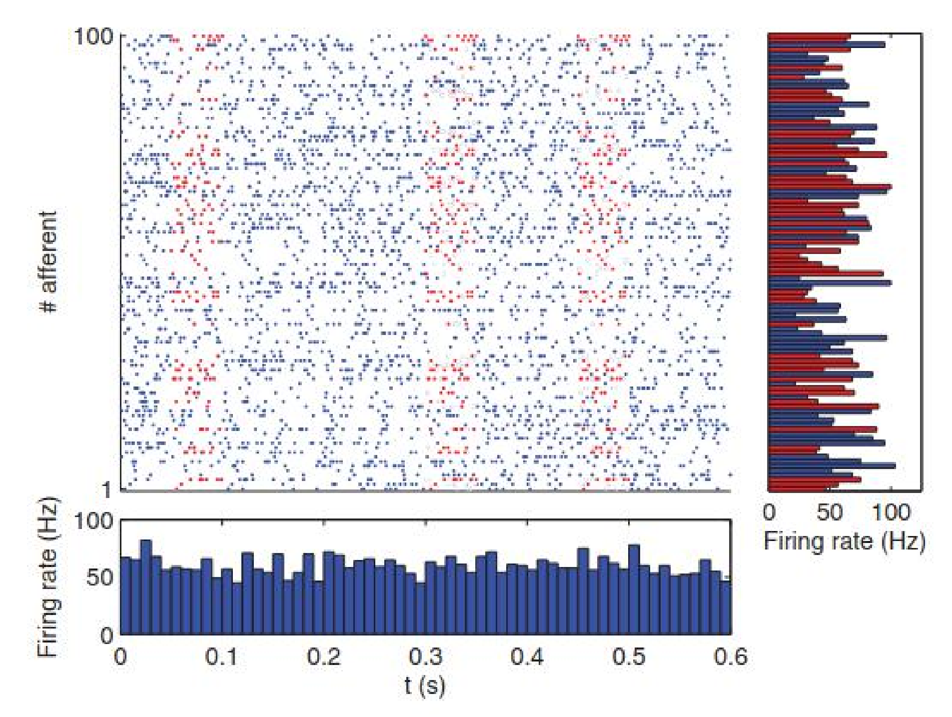

For a given action / thought / stimulus, a **pattern of spikes** are repeated, with more or less **variability**.

Downstream (post-synaptic) neurons learn to **detect** (spike in response to) these patterns.

But how ?

Masquelier et al. (2008) *Spike Timing Dependent Plasticity Finds the Start of Repeating Patterns in Continuous Spike Trains, PLoS One*

Spike-Timing Dependent Plasticity (STDP)

A local, **Hebbian** learning rule:

"Neurons that fire together, wire together"

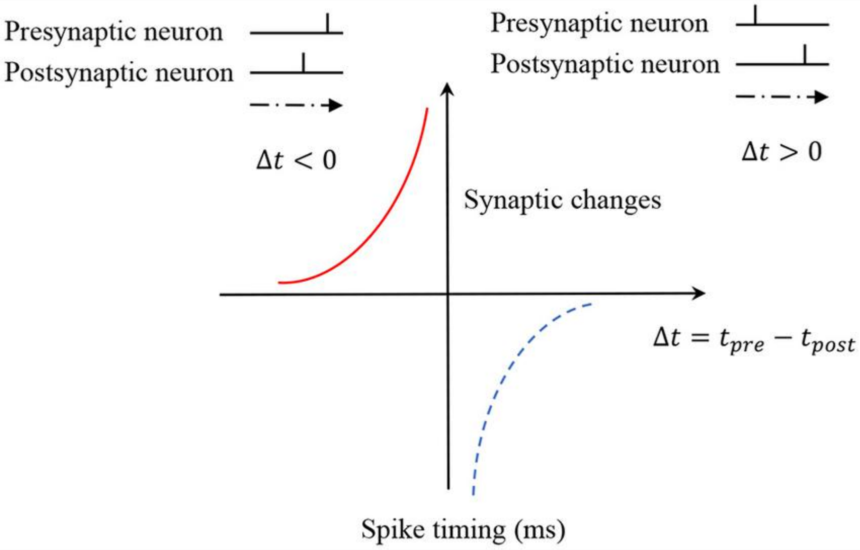

Features:

- local
- unsupervised
- biologically-plausible

Spike-Timing Dependent Plasticity (STDP)

A local, **Hebbian** learning rule:

"Neurons that fire together, wire together"

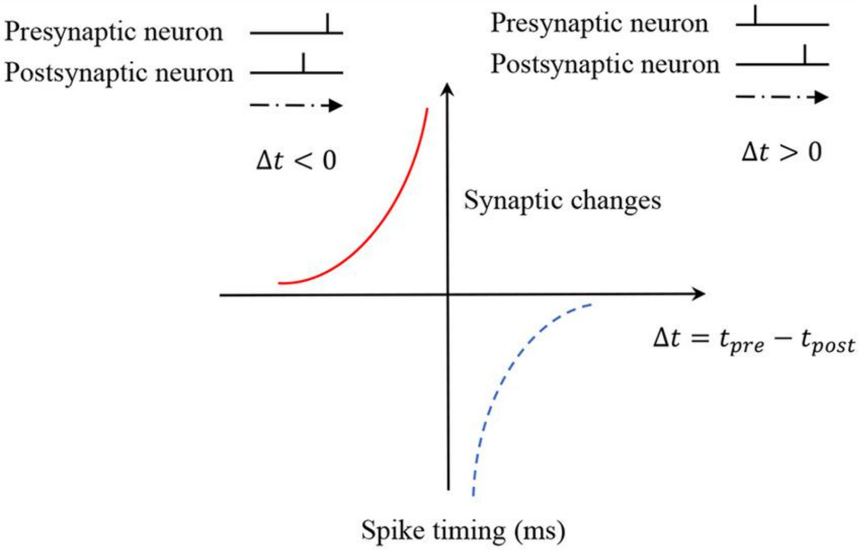

between directly connected neurons only, macroscopic behaviors emerge from microscopic changes

no label / groundtruth

Features:

- local
- unsupervised
- biologically-plausible

observed in the brain, experimental evidence

Sparse coding

Generating action potentials consumes substantial energy.

To limit overconsumption, neurons in the brain use patterns with very few spikes: **sparse coding**

This is a form of **homeostasis**, under **evolutionary pressure**

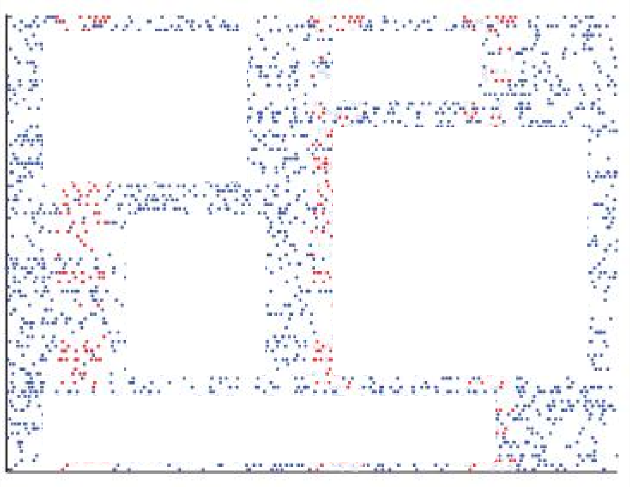

Temporal vs Rate coding

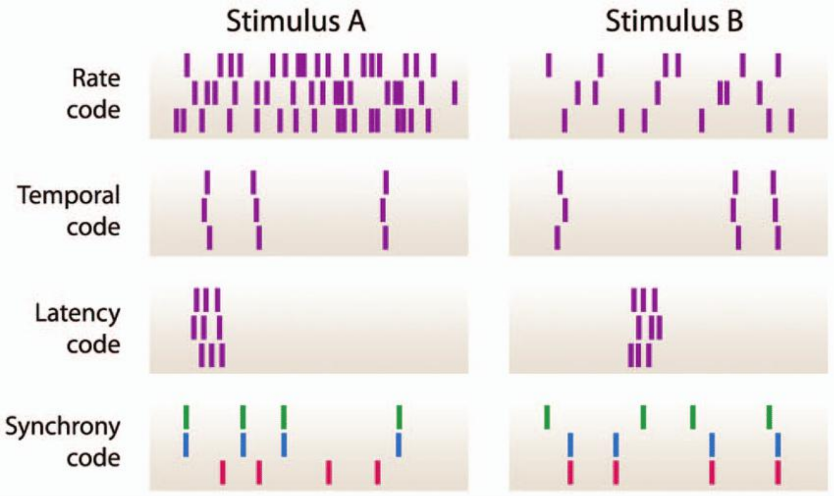

There many ways to encode (the same) information.

Biological neurons use 2 main types:

Temporal coding:
Information is encoded in the relative timing between presynaptic spikes

Rate coding:
Information is encoded in the firing rate (i.e. the number of incoming spikes)

SNNs vs ANNs

SNNs originated as low-level models of the brain.

Contrarily to ANNs, SNNs only output binary spikes, not real values.

To obtain real values and to convey finer information, need:

- time as an additional dimension,
- to average across other neurons in the population (space)

BUT they can implement sparse coding for higher energy efficiency

→ *take inspiration for engineering ?*

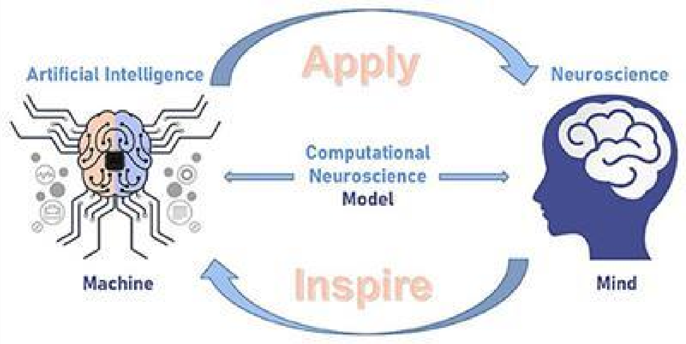

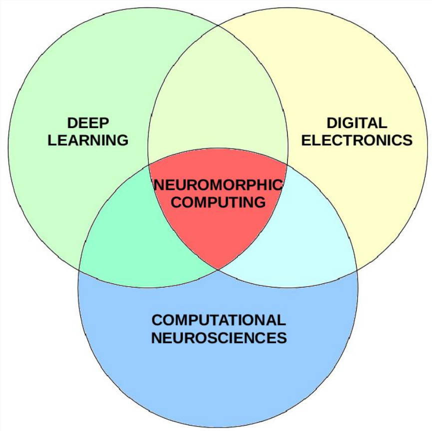

Neuromorphic computing and sensing

Brain-inspired AI

The energy efficiency of the brain

? W (Watt)

70 W

40 W

Q: How much is the brain's estimated power consumption ?

The **incredible** energy efficiency of the brain

20 W (Watt)

70 W

40 W

Neuromorphic chips

 +  = 

Neuromorphic chips

+

=

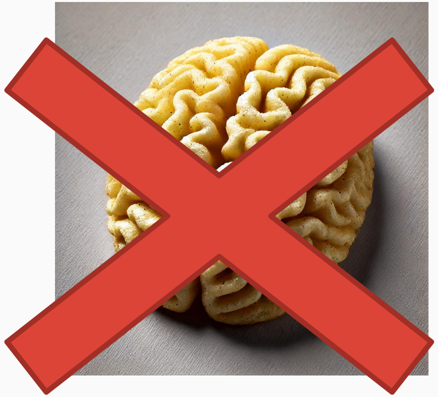

Neuromorphic chips

+

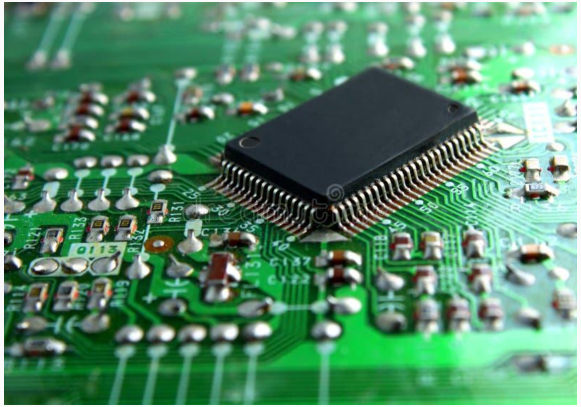

=

Electronic chips ("puces") specially designed to emulate spiking neural networks (SNNs)

GPU / Graphics card --> gaming
Neuromorphic chips --> SNNs

Neuromorphic chips: main principle

Electronic chips ("puces") specially designed to emulate spiking neural networks (SNNs)

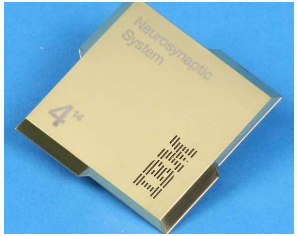

IBM TrueNorth

Only consume power when generating a spike (like the brain)
Few spikes (sparse coding) ==> little energy consumption

Advantages (over regular computers/chips):

- ultra **low power** consumption
- very **fast**
- **asynchronous**

Intel Loihi 2

Industrial applications:

- Embedded systems
- IoT
- BCIs
- …

HBP SpiNNaker

Brainchip Akida

Neuromorphic Retina: pros and cons

Also known as the **Dynamic Vision Sensor (DVS)** or **Event Camera**

Mimics how actual retinas work:
emit a spike only on big changes of luminance at a given pixel (location)

https://youtu.be/MjX3z-6n3iA?si=IHiwIEPFFPPMySS6

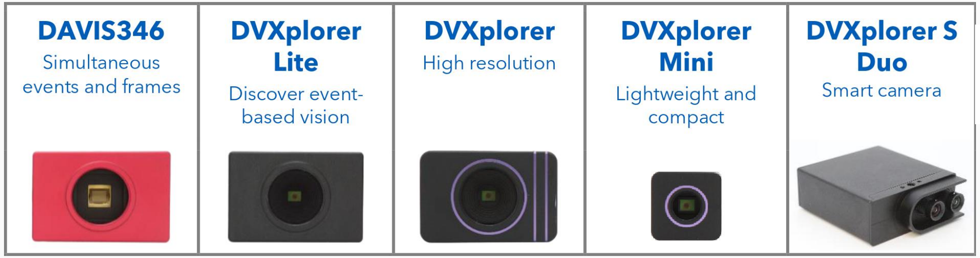

The DV suite from Inivation,
https://inivation.com/

DVS event streams contain a lot of information !

Proof: *quality grayscale images can be reconstructed from events thanks to deep neural networks !*

These **reconstructed videos inherit the benefits of event data**:

- events are not subject to motion blur → reconstructions cleaner than frame based cameras
- events have a higher dynamic range → same
- events have a high temporal resolution → reconstruction / information retrieval possible at high FPS rates

Many types of information can be retrieved:

- **intensity** (this work)
- **optical flow** (Javier's)
- **depth** (StereoSpike)

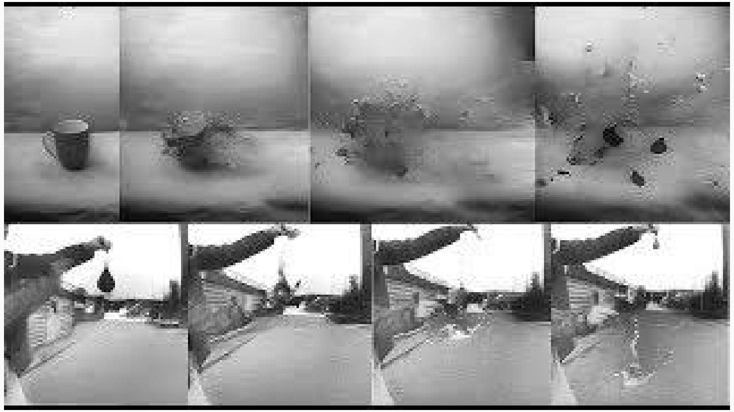

H. Rebecq, R. Ranftl, V. Koltun and D. Scaramuzza, "**High Speed and High Dynamic Range Video with an Event Camera**," in *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 43, no. 6, pp. 1964-1980, 1 June 2021, doi: 10.1109/TPAMI.2019.2963386

Neuromorphic Retina: pros and cons

Feature \ Camera type	Event-based	Frame-based
Power consumption	Very low (~ mW)	~W
Temporal resolution	Very high (~ µs)	60 FPS >~ 10 ms
Spatial Resolution	Low (<1 MPx), best 1280x720	Very High (>10 MPx), e.g. Huawei P20
Dynamic Range	Very high (120 dB)	~50 dB
Data efficiency	Sparse events, compact representation	Redundant information
Mode	Asynchronous	Synchronous
Price	High (several k €)	Low

Energy-efficient

No motion blur !

Dim / Bright light ok !

Fast

Neuromorphic Cochleas

Equivalent of the DVS but for audio: **Dynamic Audio Sensor (DAS)**

show raster plots

Conclusion

—

Outline

I. What is a model ?

II. Models of neurons / Spiking Neural Networks (SNNs)

III. Neuromorphic computing and sensing

IV. Conclusion

Little to no maths

Large overview

Main keywords

Starting point for self-study !

Going further… (Some exercises)

Exercise #1:

Code by yourself, in Python, a LIF neuron in discrete time: with thresh=1, V_{rest}=0, and tau=10
It should receive a constant input current I=0.3 between time-steps 10 and 30.
Simulate during 50 time-steps and plot the membrane potential and spikes.
You will iterate through time-steps with a for loop.

Exercise #2:

Add a refractory period of T=5dt after each spike, during which potential is maintained to V_{rest} and launch a new simulation. Do you notic any differences ?

Exercise #3:

Play with the input current ! Change it to a ramp, and/or to a spike train, and see the difference in behavior on your neuron.

If you don't know where to start, you can use ChatGPT… BUT it's better if you do it yourself, and make sure you understand its answer !

Going further... (Programming / coding / modelling)

The Brian2 spiking neuron simulator
(Python, with tutorials)

https://brian2.readthedocs.io/en/stable/

The NEST simulator
(Python, tutorials)

https://nest-simulator.readthedocs.io/en/stable/

Going further… (Selected materials 1/2)

Free book and online course by Wulfram Gerster (EPFL)

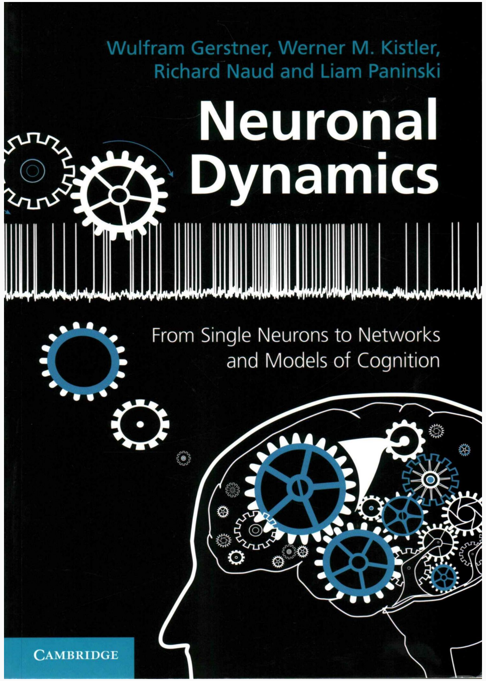

https://neuronaldynamics.epfl.ch/

Artem Kirsanov's YT channel

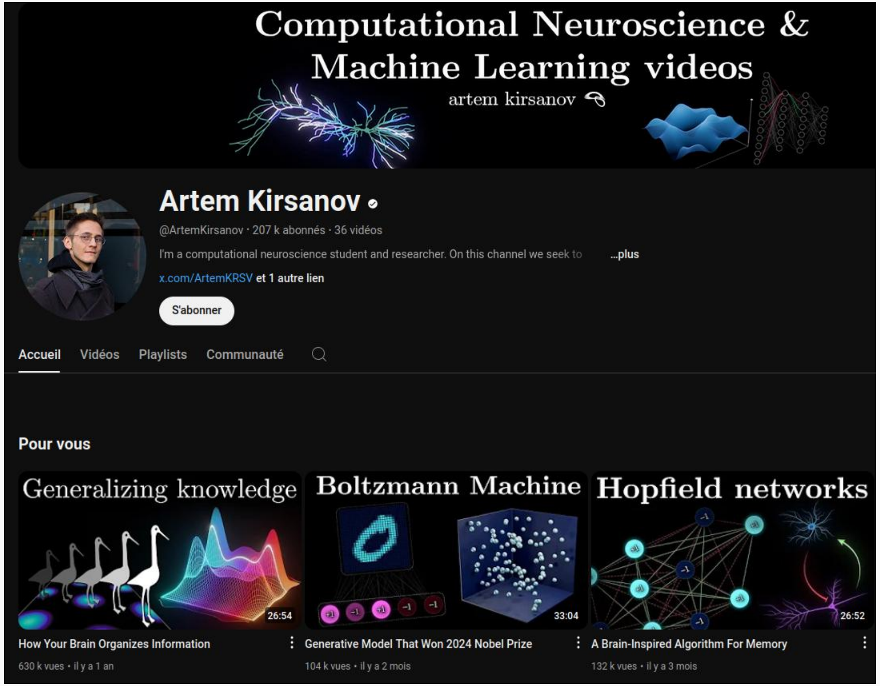

https://www.youtube.com/@ArtemKirsanov

Going further… (Selected materials 2/2)

Eugene Izhikevich's textbook

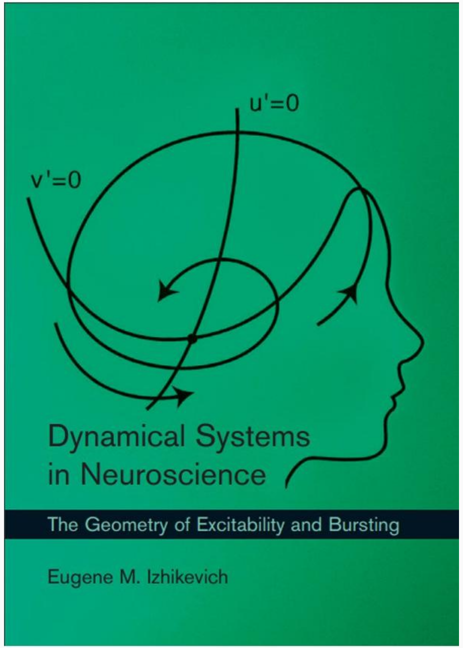

Cosyne 2022 Tutorial on Spiking Neural Networks
(Dan Goodman, YouTube)

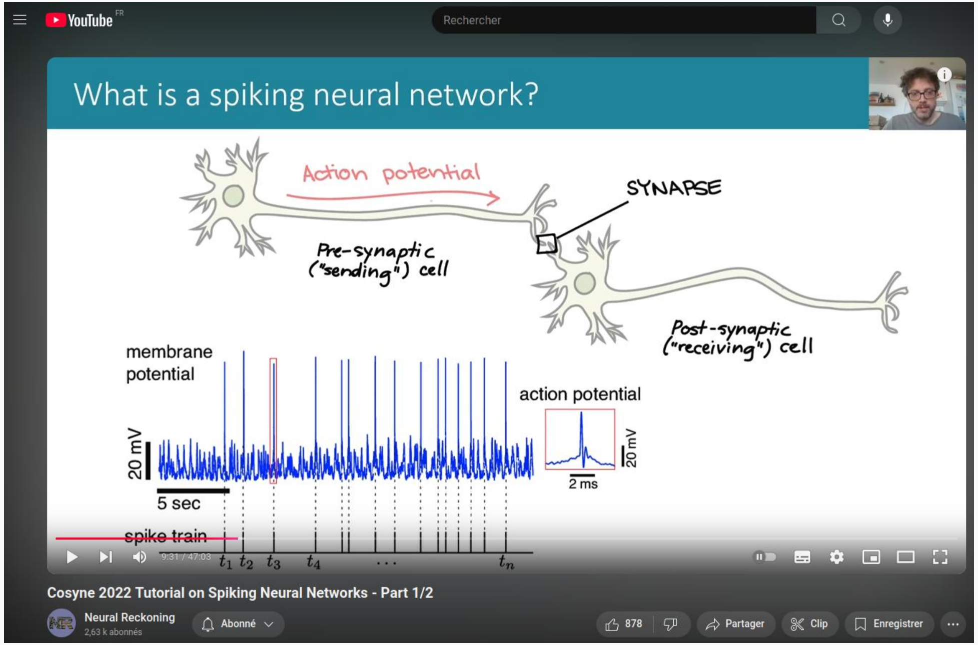

https://youtu.be/GTXTQ_sOxak?si=5BDwA02QHEjatV7t

Thank you for your attention !

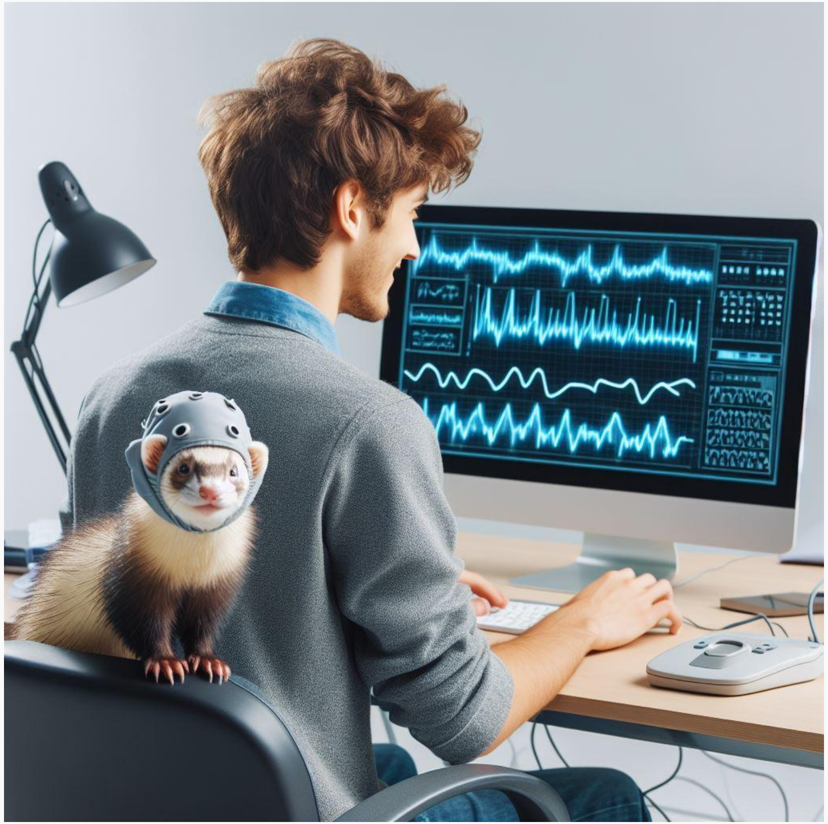

Ask your questions now, or send me an email (**ulysse.rancon@cnrs.fr**)