EPFL

NX-414: Brain-like computation and intelligence

Review session

Alexander Mathis
alexander.Mathis@epfl.ch

Lecture 14, May 28

I think;
therefore
I am.

Descartes

Understanding the brain, a challenge for the 21st Century!

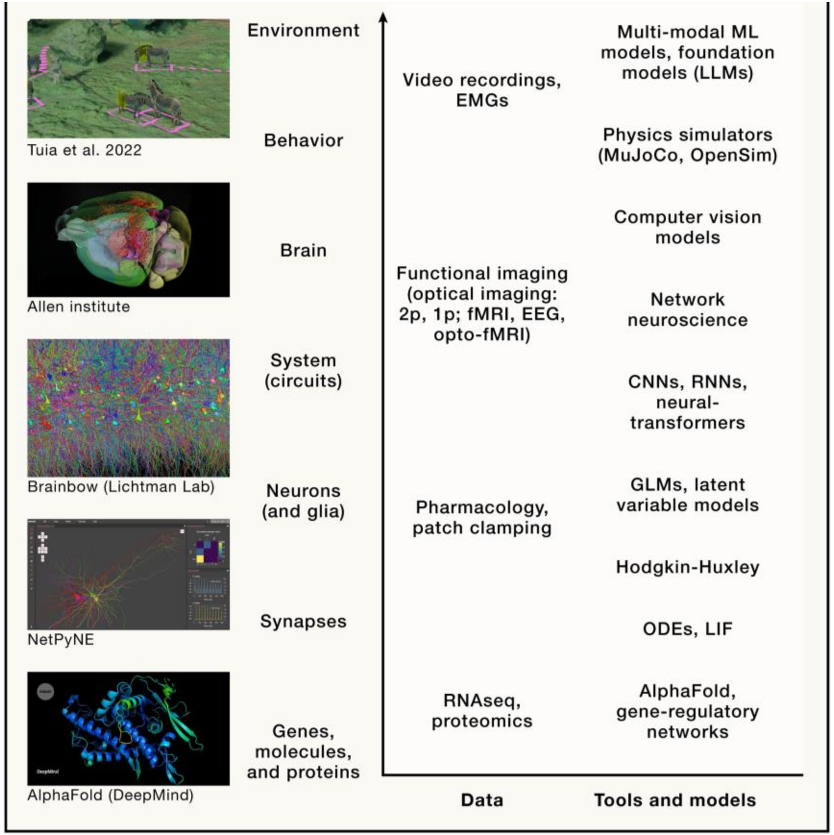

M. Mathis, Perez Rotondo, Cheng, Tolias, Mathis Cell 2024

EPFL

https://en.wikipedia.org/wiki/Human_brain

How does the brain compute?

vs. Turing Machines, von Neumann machines, …

EPFL

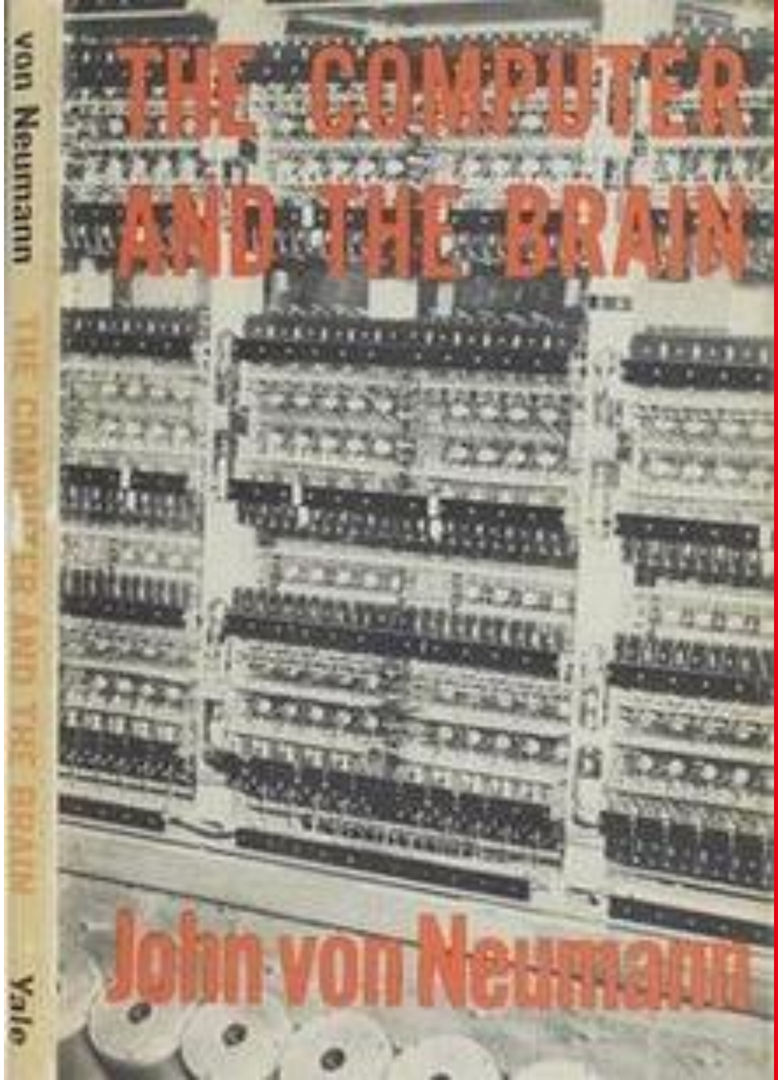

How does the brain compute?

vs. Turing Machines, von Neumann machines, …

https://en.wikipedia.org/wiki/The_Computer_and_the_Brain

How does the brain compute?

- Parallel processing
- Distributed representation and computation
- (Mostly) using action potentials/spikes (energy efficient)
- Broadly projecting neuromodulators
- Highly **plastic**
  - Synaptic plasticity (Hebbian learning, STDP…)
  - Dynamic gating (e.g. selective attention)
  - Critical period/Development
  - …

Modeling approaches:

- **Classical modeling**
  Sparse coding, plasticity theory, attractor models, low-dimensional visualization, …

- **Top-down modeling**
  Task-driven models optimized for an ecological behavioral objective. E.g., object recognition, next-word prediction, optimal feedback control, …

- **Bottom-up modeling**
  Start from biophysical, anatomical, behavior observations and check how they impact the model
  E.g., SDS, …

How does the brain compute?

EPFL NX-414 overview

Modeling the brain

- **Classic modeling**
  - Lecture 1 neural code
  - Lecture 2 normative models
  - Lecture 3 Bayes and attractor models
- **Top-down modeling** (task-driven)
  - Lecture 4 path integration and vision I
  - Lecture 5 vision II
  - Lecture 6 vision III and audition
  - Lecture 7 proprioception
  - Lecture 8 language I
  - Lecture 9 language II
  - Lecture 10 motor control and OFC
  - Lecture 11 language III and cognition
  - Lecture 12 learning to control
- **Bottom-up modeling**
  - Lecture 13 brain-inspired skill learning

Parallel, modular processing

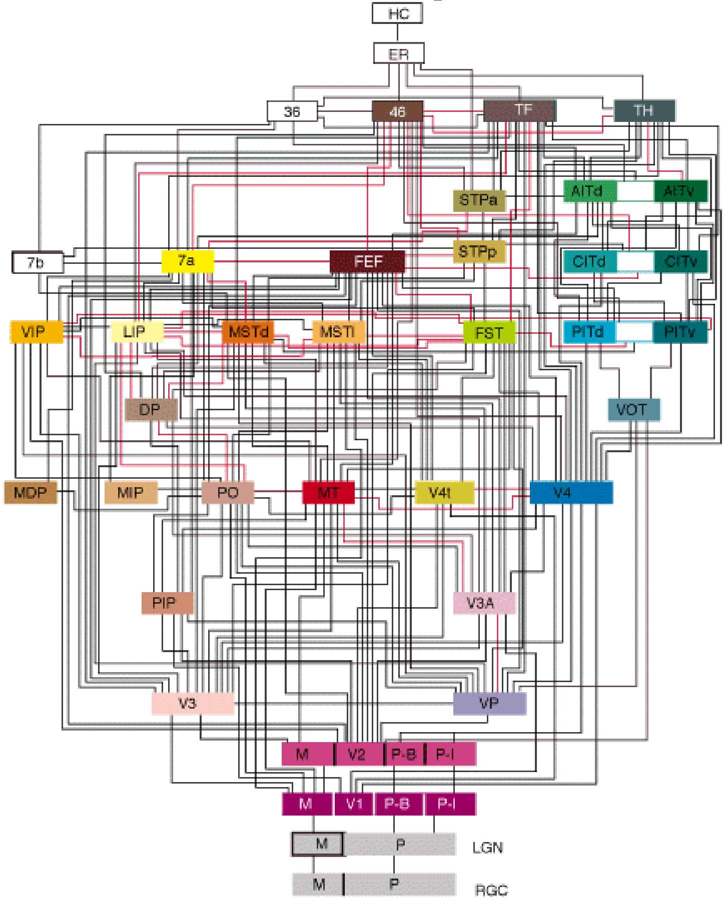

Felleman and Van Essen *Cerebral Cortex* 1991

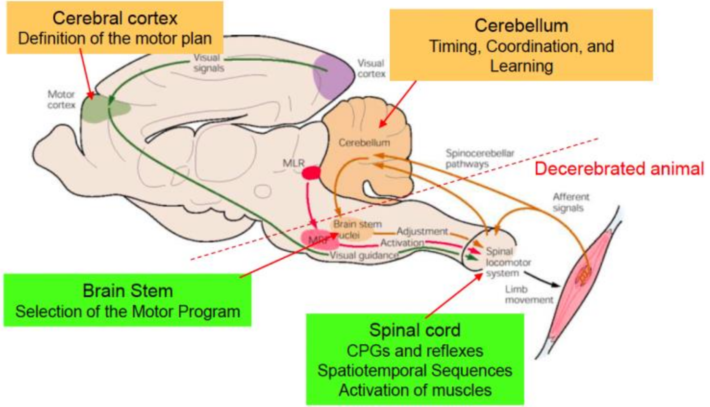

. Kandel et al. 4th edition Principles of Neural Science

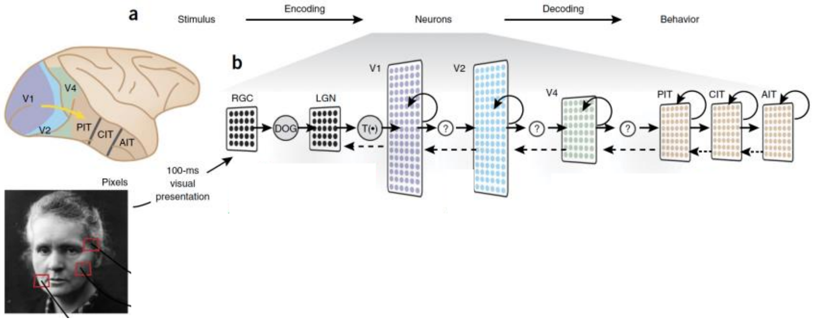

Yamins and Di Carlo, Nat Neuro 2016

Distributed, energy-efficient representations

Results of training sparse coding model on 16 x 16 patches

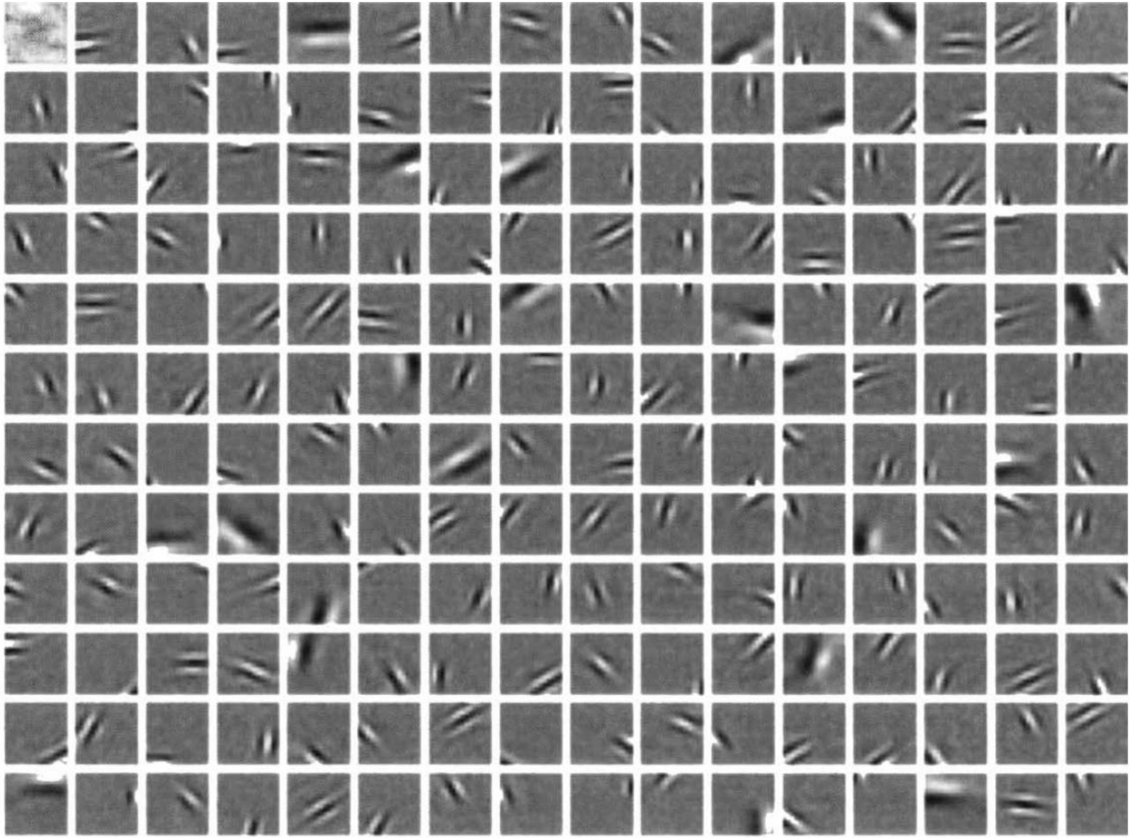

Olshausen & Field, 1996 Nature

Distributed representations

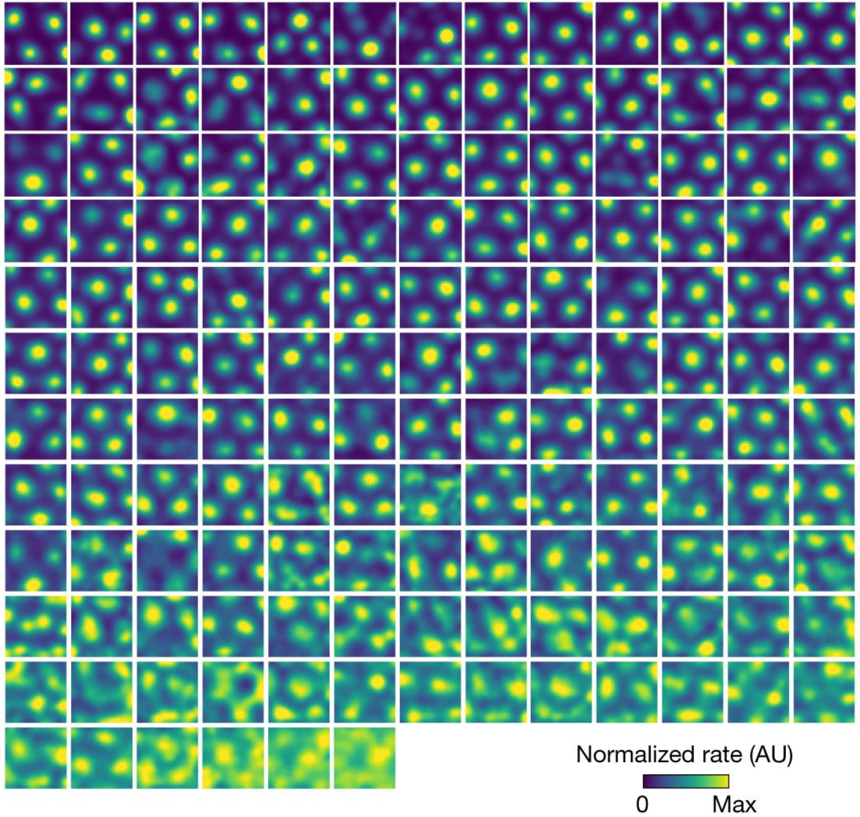

Gardner et al. 2022 Nature

Elegant, distributed representations

$$\left(\frac{f(\varphi)}{f_{max}}\right)_i = [cos(\varphi - \varphi_i)]_+$$

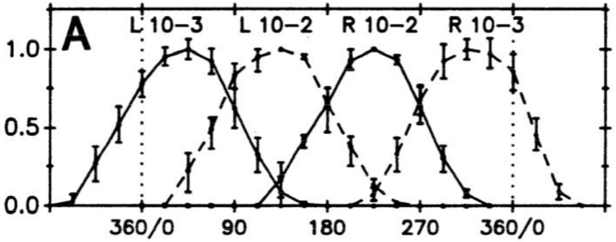

Miller, Jacobs, Theunissen, J Neurophysiology 1991

$$v_{est} = \sum_i \left(\frac{k}{f_{max}}\right)_i \cdot \mathrm{v}_i$$

Neural Plasticity

Donald Hebb: "What fires together, wires together."

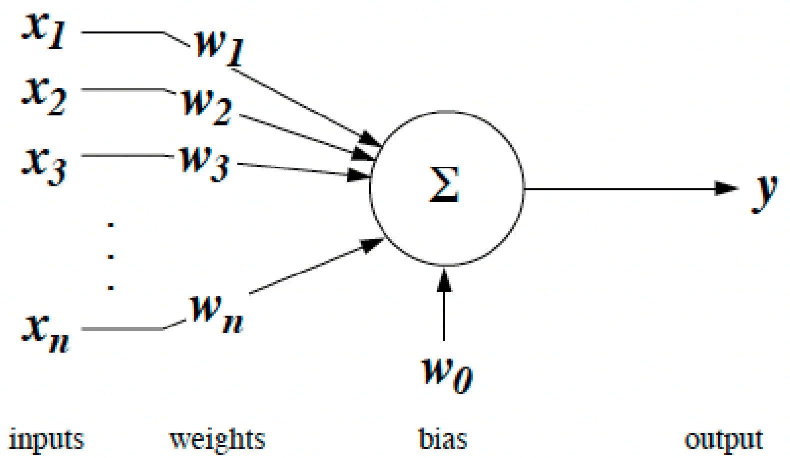

Donald Hebb (1904 –1985)
Wikipedia

Plasticity in the brain

Computational level:

- Unsupervised learning
- Supervised learning
- Reinforcement learning
- Transfer learning
- Curriculum learning
- Lifelong learning

Mechanistic level:

- Synaptic plasticity
  - Long-term potentiation
  - Short-term depression
- Intrinsic plasticity
- Homeostatic plasticity
- Metaplasticity
- Neurogenesis
- …

What is intelligence?

Definitions

“Viewed narrowly, there seem to be almost as many definitions of intelligence as there were experts asked to define it.” — R. J. Sternberg

“We shall use the term ‘intelligence’ to mean the ability of an organism to solve new problems . . . ” W. V. Bingham

Intelligence measures an agent’s ability to achieve goals in a wide range of environments. NOTE: it thus needs to be adaptive!

Universal Intelligence:
A Definition of Machine Intelligence

Shane Legg

IDSIA, Galleria 2, Manno-Lugano CH-6928, Switzerland
shane@vetta.org www.vetta.org/shane

Marcus Hutter

RSISE @ ANU and SML @ NICTA, Canberra, ACT, 0200, Australia
marcus@hutter1.net www.hutter1.net

What forms of intelligence did we cover?

- Visual intelligence (object recognition)
- Bodily intelligence (proprioception)
- Language
- Adaptive motor control
- Skill learning
- Path integration
- …

Some brain-like computations

Attractor models I

Path integration in calculus:
(on a v. Neuman Machine)

$$\int_{start}^{now} d\,\gamma$$

Path integration in the brain:
(collective computation)

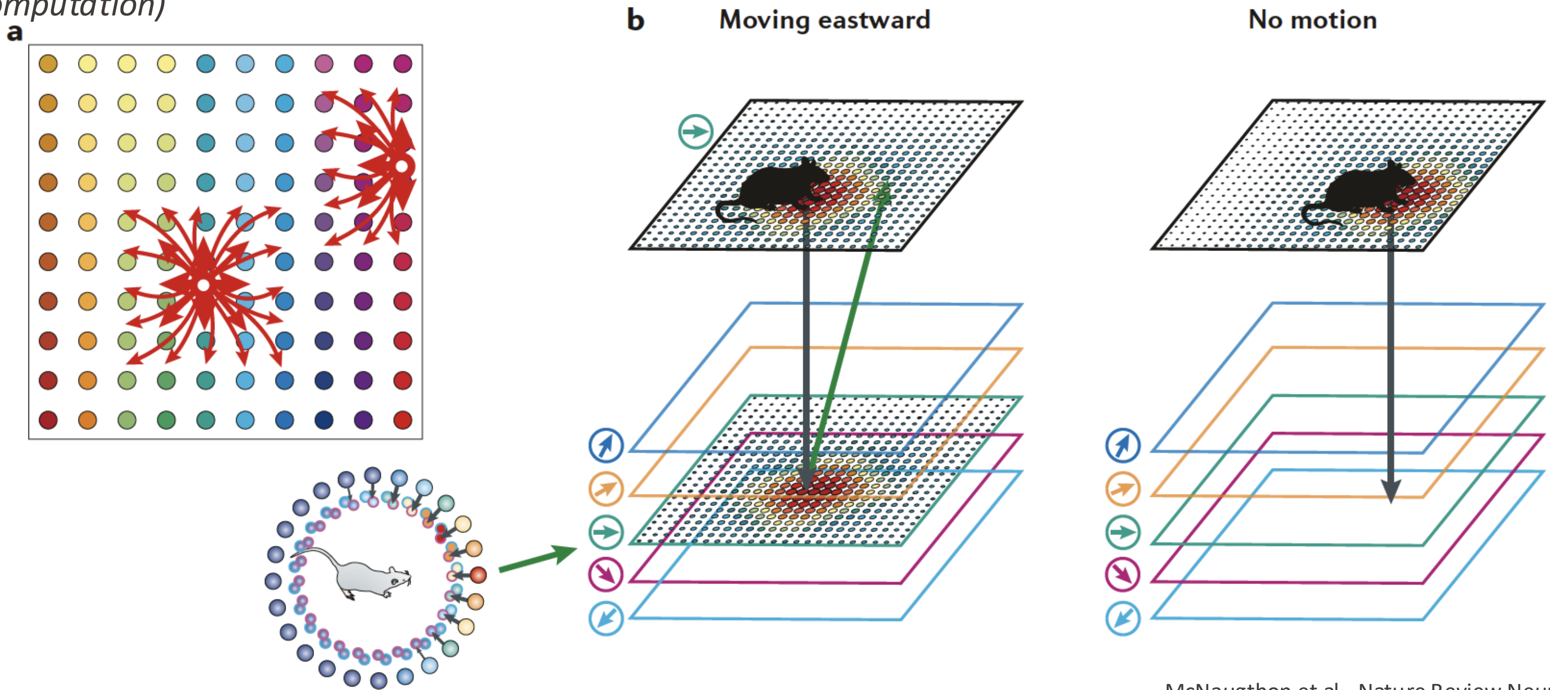

McNaugthon et al., Nature Review Neuroscience 2006

Attractor models II

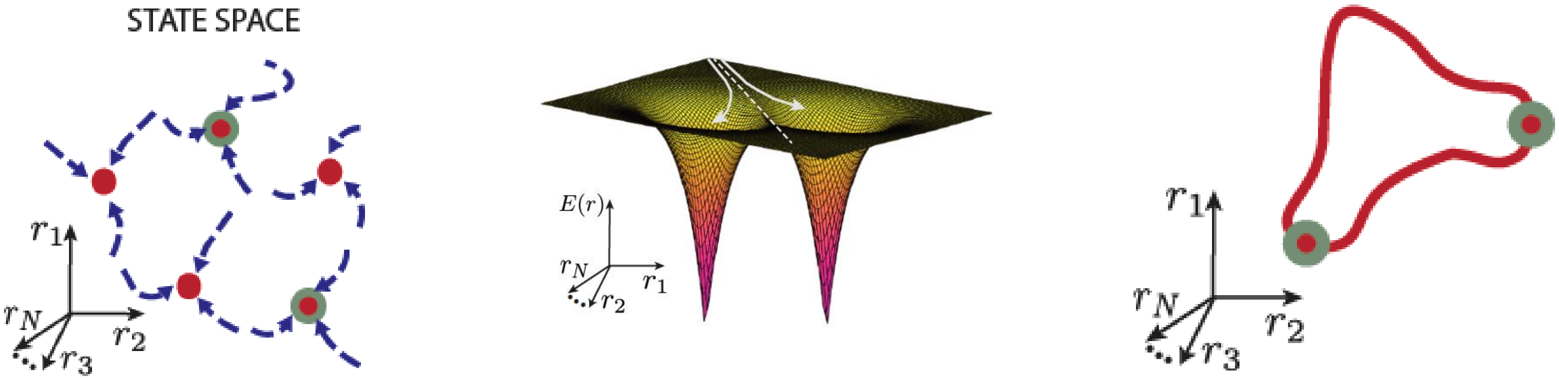

Functions:

- Associative memory
- Decision making
- Path integration
- …

Khona & Fiete, Nat Review Neuro 2021

Hierarchical, convolutional neural networks

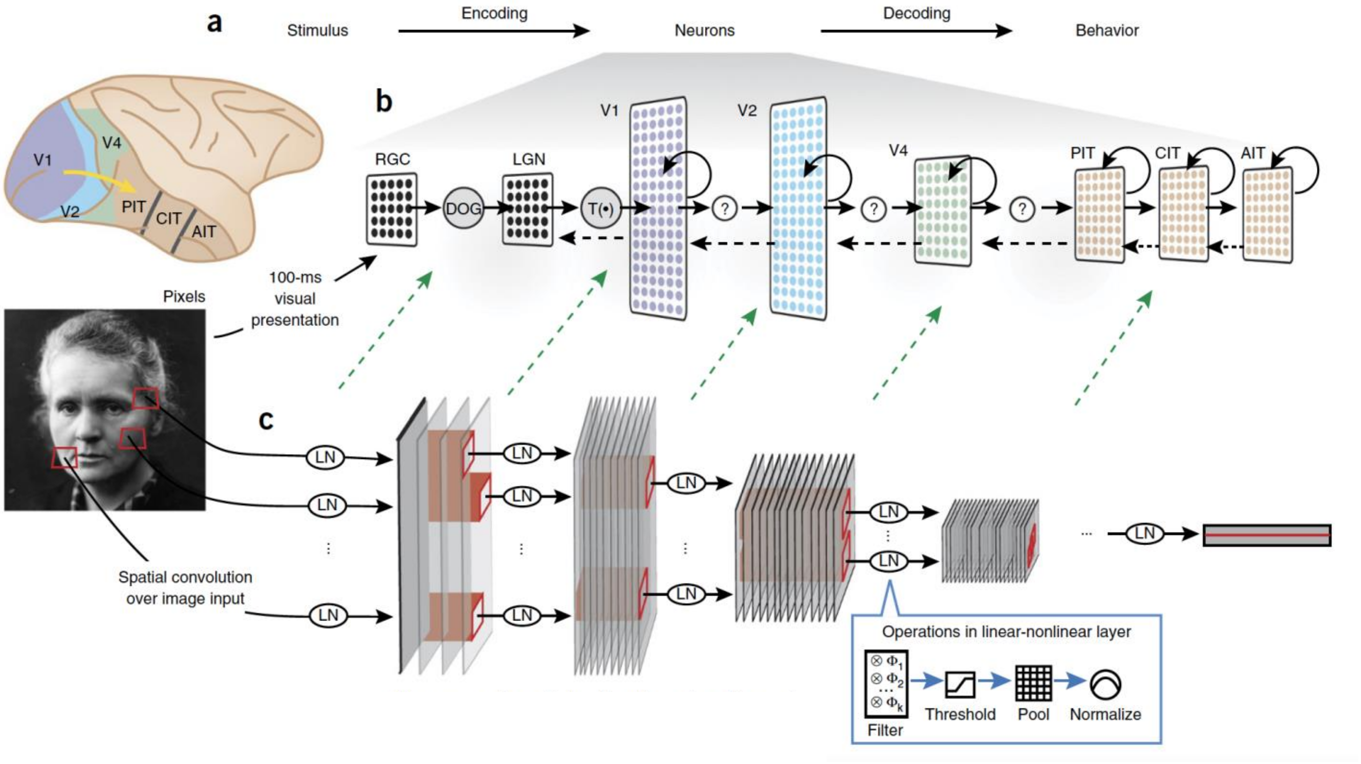

Yamins and Di Carlo, Nat Neuro 2016

DMAP's brain inspired architecture

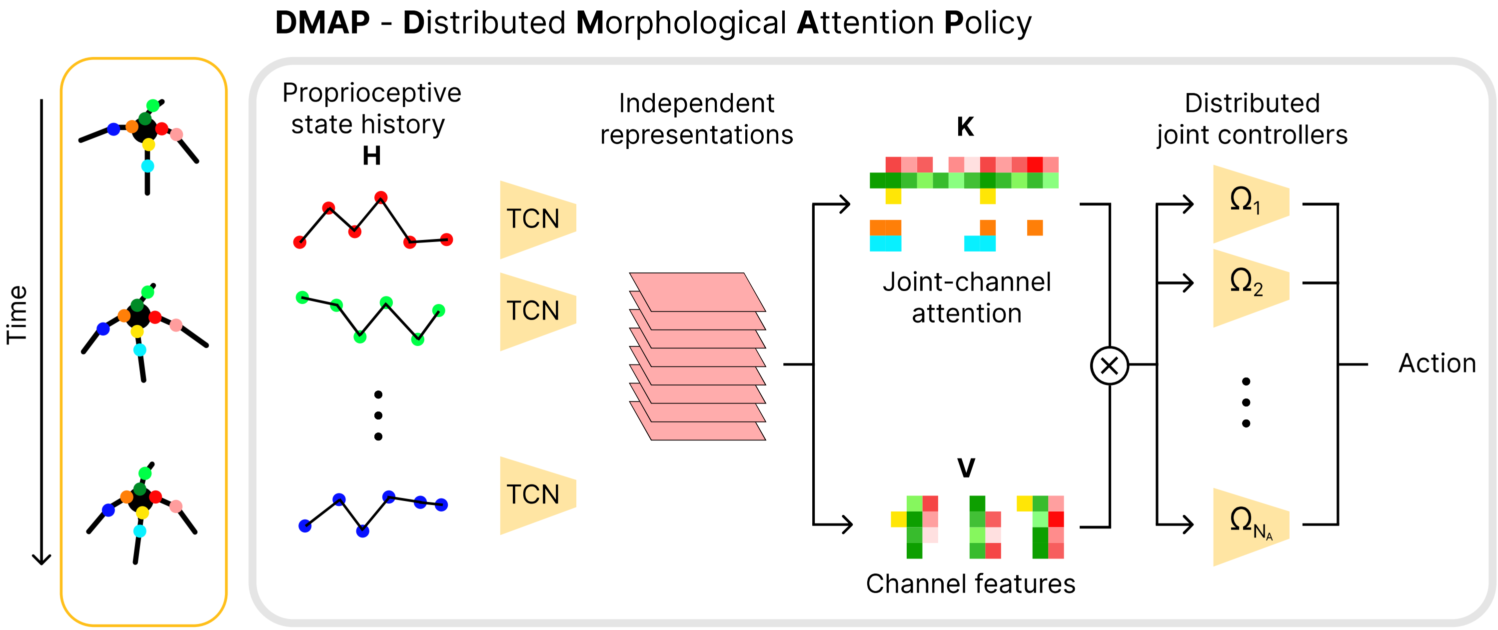

Chiappa, Marin Vargas, Mathis, Neurips 2022 /arxiv

Cortical column

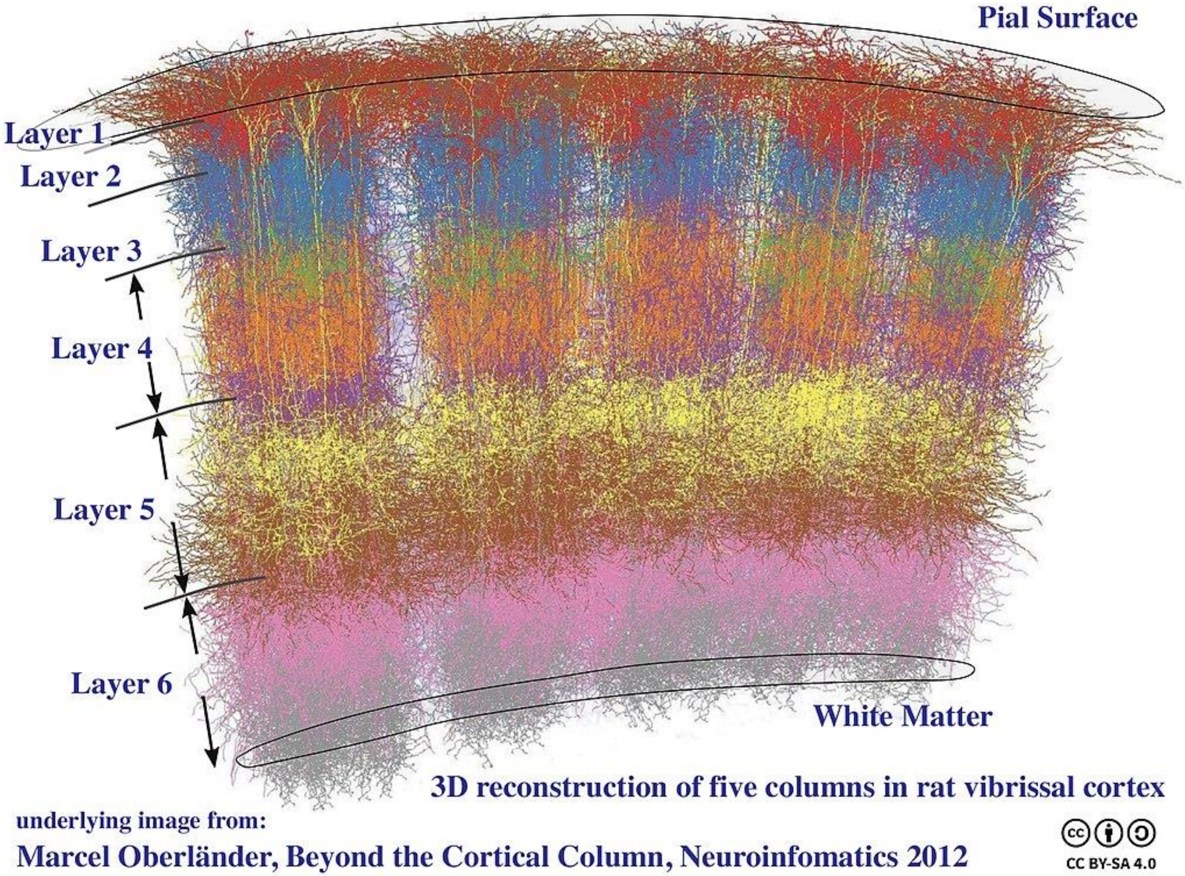

EPFL

A petavoxel fragment of human cerebral cortex reconstructed at nanoscale resolution

“We found a previously unrecognized class of directionally oriented neurons in deep layers (see figure, panel J) and very powerful and rare multisynaptic connections between neurons throughout the sample (see figure, panel K).”

Browse it online!

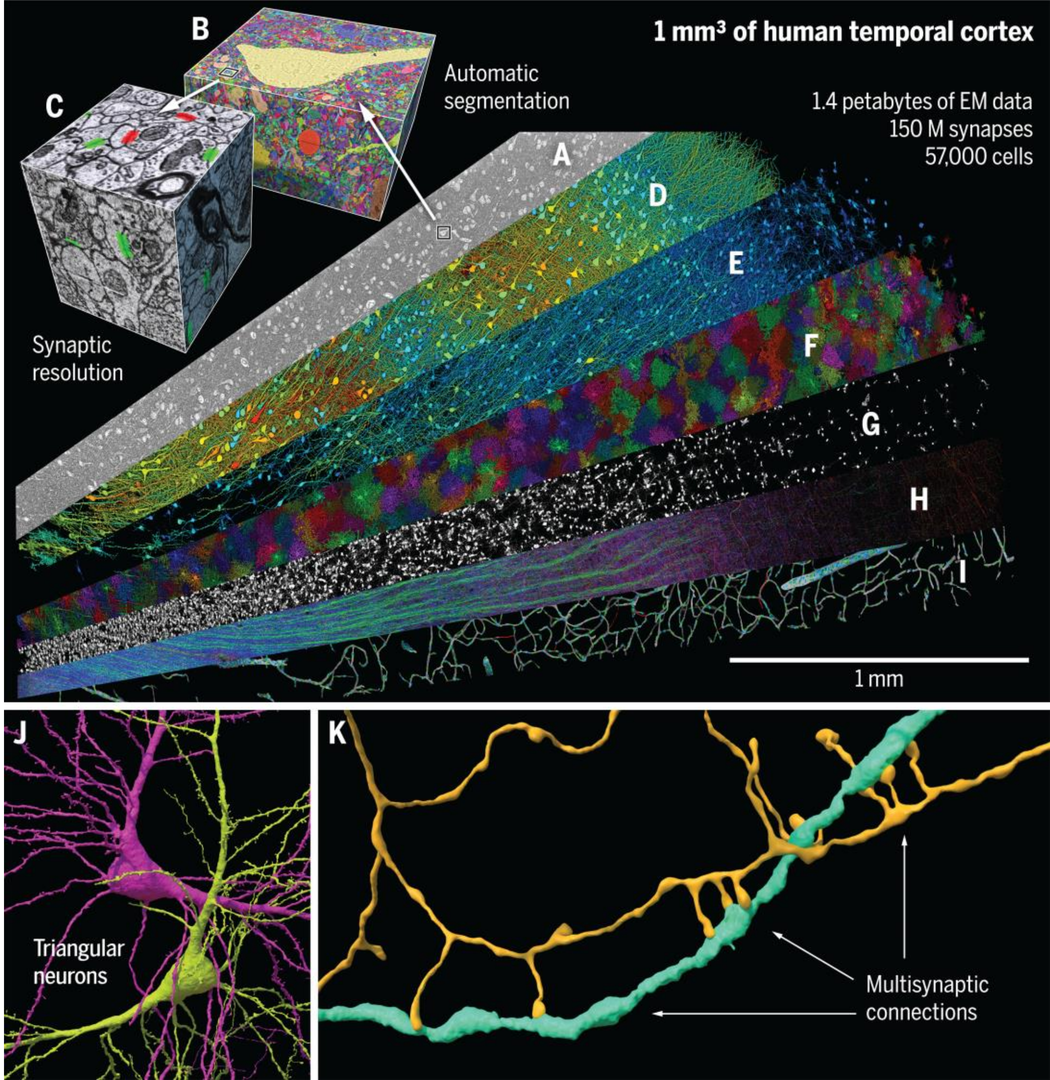

Shapson-Coe, … Lichtman, Science 2024

From: **Large Language Models and the Reverse Turing Test**

Neural Comput. 2023;35(3):309-342. doi:10.1162/neco_a_01563

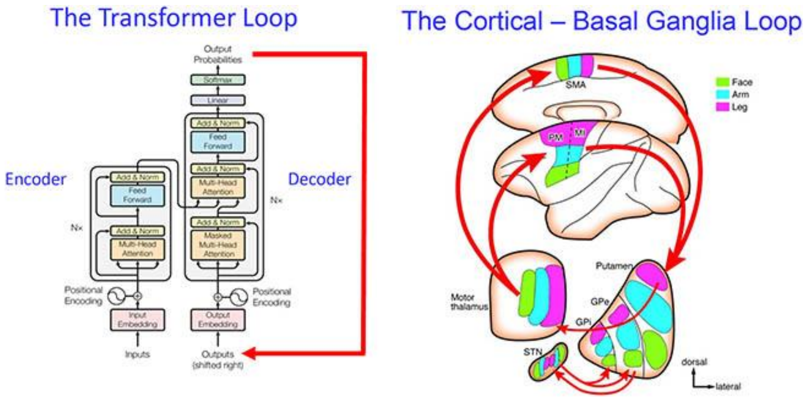

Dopaminergic circuits and reward prediction errors

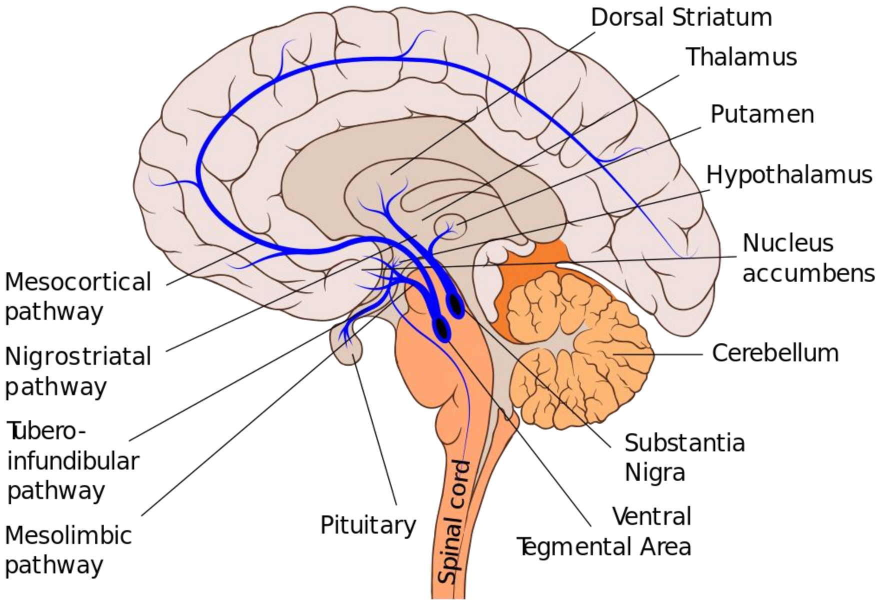

Wikipedia

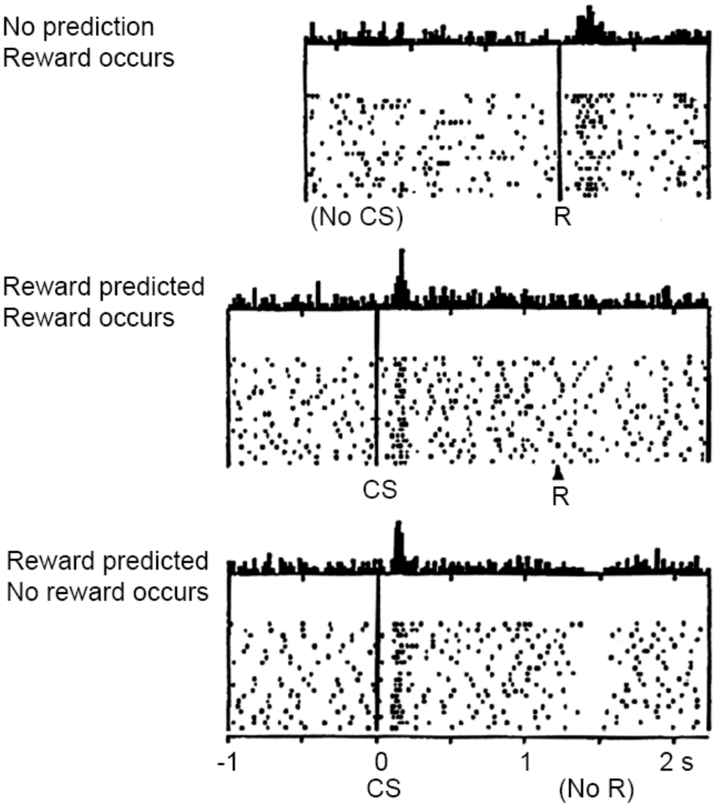

Schultz, Dayan, Montague, 1997

A Neural Substrate of Prediction and Reward

Wolfram Schultz, Peter Dayan, P. Read Montague*

The capacity to predict future events permits a creature to detect, model, and manipulate the causal structure of its interactions with its environment. Behavioral experiments suggest that learning is driven by changes in the expectations about future salient events such as rewards and punishments. Physiological work has recently complemented these studies by identifying dopaminergic neurons in the primate whose fluctuating output apparently signals changes or errors in the predictions of future salient and rewarding events. Taken together, these findings can be understood through quantitative theories of adaptive optimizing control.

$$V(s_t) \leftarrow V(s_t) + \alpha_t \underbrace{(R_t + \gamma V(s_{t+1}) - V(s_t))}_{\text{Dopamine?}}$$

Schultz, Dayan, Montague, Science 1997

Reinforcement learning scales really well!

- The computational and memory requirements (even) for games is enormous!

Chess: 10^{120}

Go: 3^{361}

Parameters θ_i

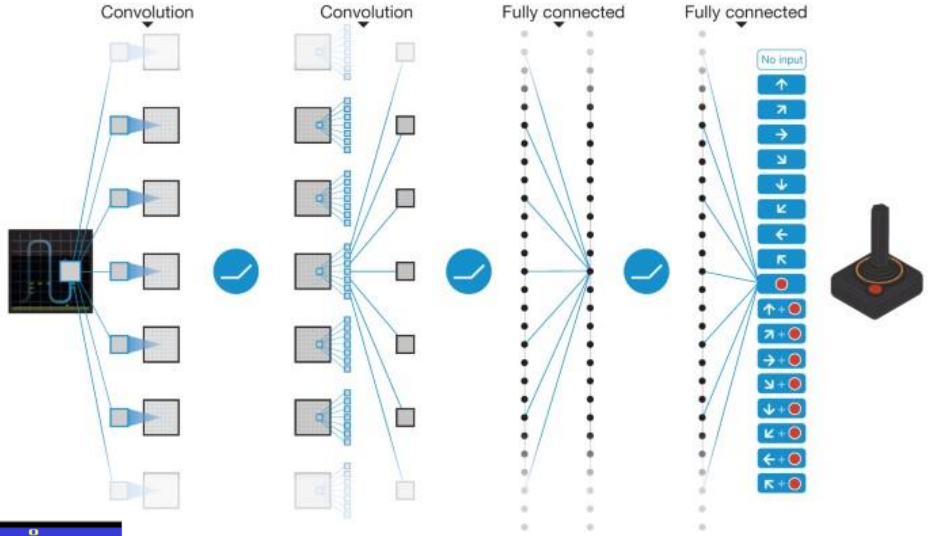

$$L_i\left(\theta_i\right) = \mathbb{E}_{s,a\sim\rho(\cdot)}\left[\left(y_i - Q\left(s,a;\theta_i\right)\right)^2\right]$$

$$\nabla_{\theta_i} L_i\left(\theta_i\right) = \mathbb{E}_{s,a\sim\rho(\cdot);s'\sim\mathcal{E}}\left[\left(r + \gamma \max_{a'} Q(s',a';\theta_{i-1}) - Q(s,a;\theta_i)\right)\nabla_{\theta_i} Q(s,a;\theta_i)\right]$$

-

Mnih et al., Nature 2013 (Deep Mind)

Building brain-like models (top-down)

Biological Intelligence ⟷ Artificial Intelligence

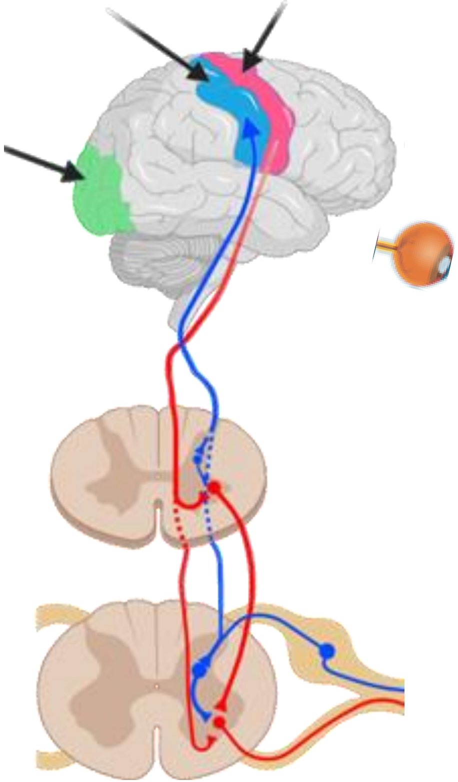

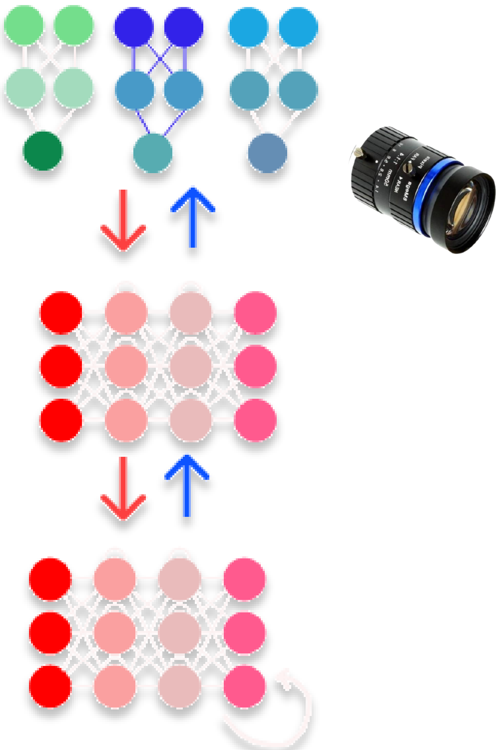

Hausmann*, Marin-Vargas*, Mathis**, Mathis**, Curr op. in Neuro., 2021

Normative frameworks

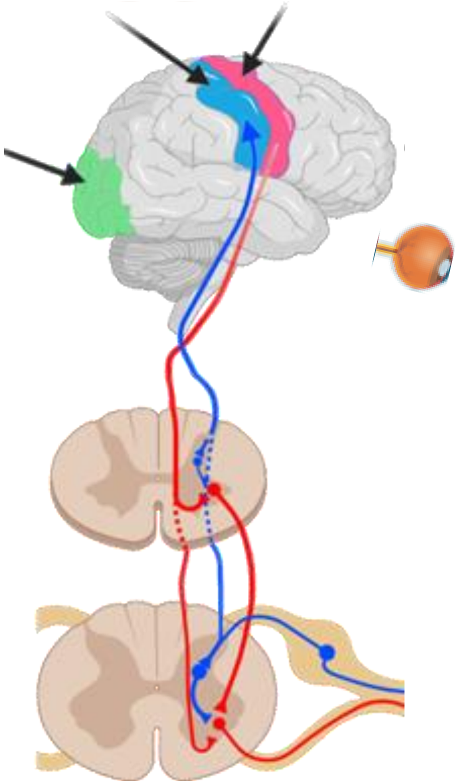

Information theoretic

e.g. sparse coding,
redundancy reduction,
mutual information …

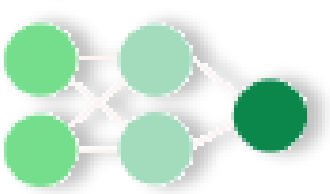

Utilitarian

e.g. recognize objects,
chase prey, navigate, path
integration, localize body
parts, control limbs …

Example: proprioception

Ingredient 3:
Putative goals (of proprioception)

Ingredient 2:
ANNs

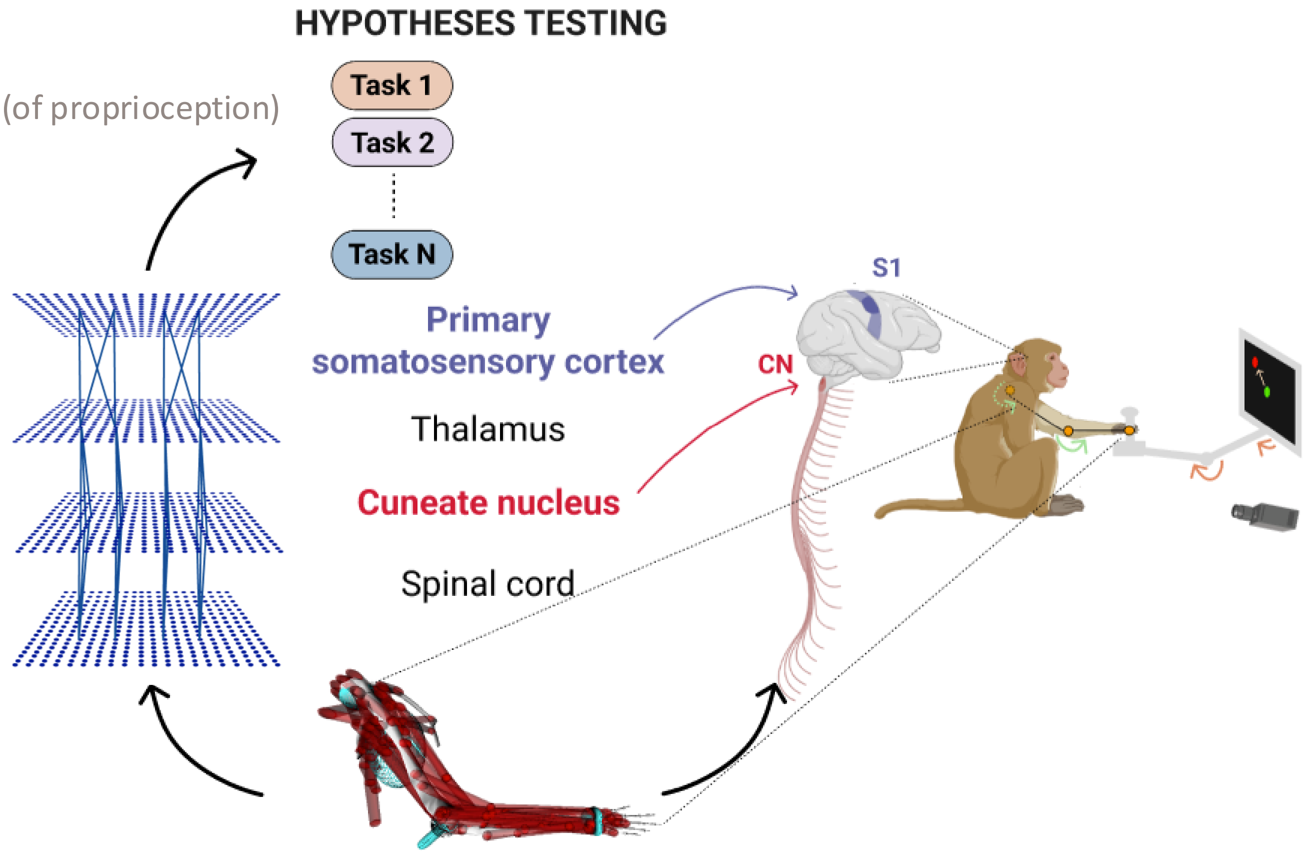

Ingredient 1: simulating spindle dynamics at scale

Marin Vargas A., Bisi* A., Chiappa A., Versteeg C., Miller L., Mathis A. Cell 2024*

Task-performance and neural predictability are correlated

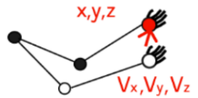

Hand position and velocity task (HP & HV)

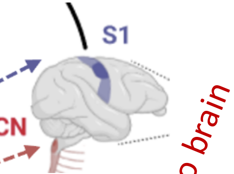

Generalization to brain

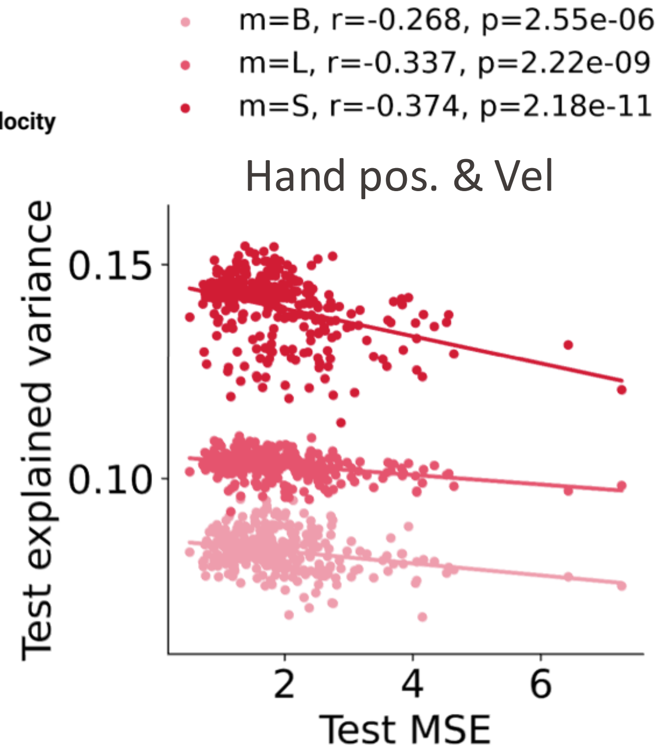

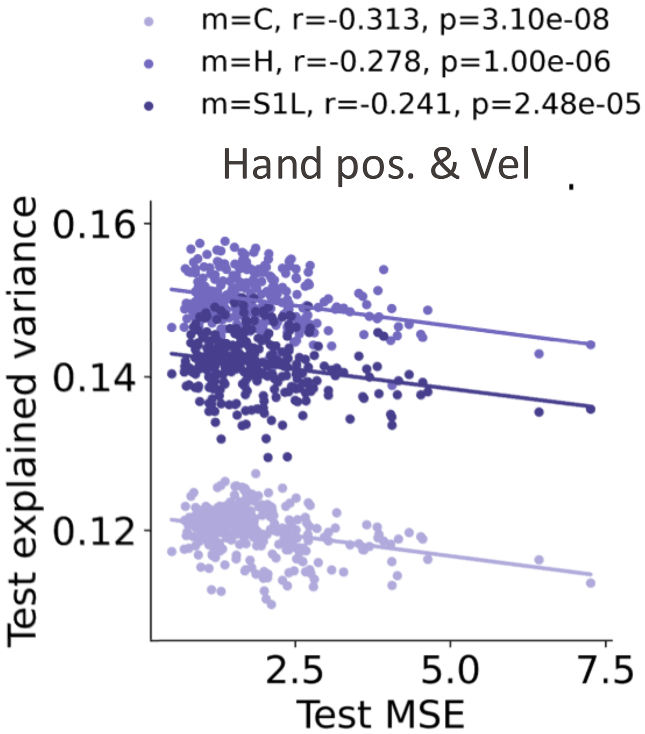

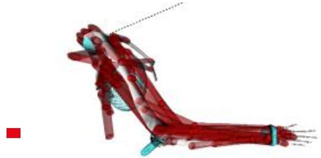

Optimized on biomechanics (body)

Different model architectures

Example: language

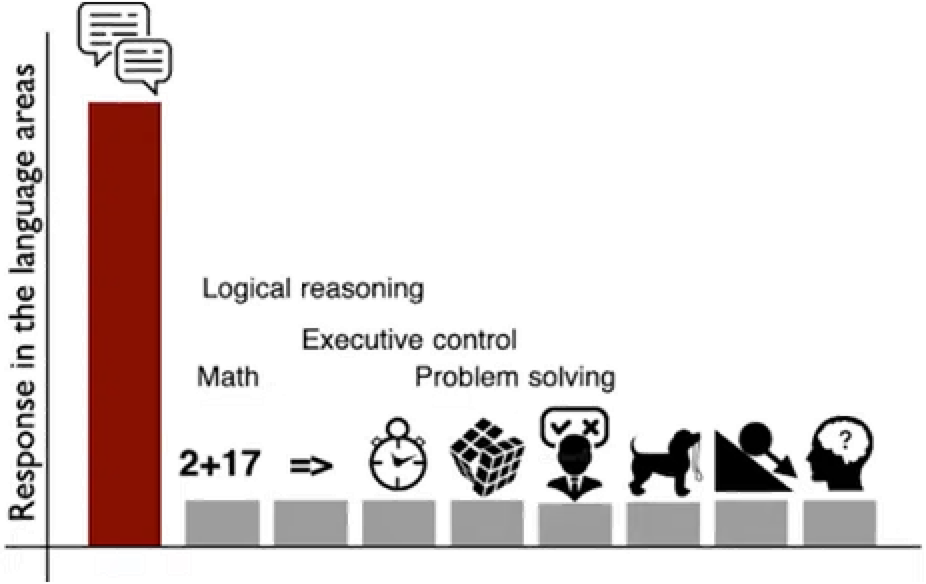

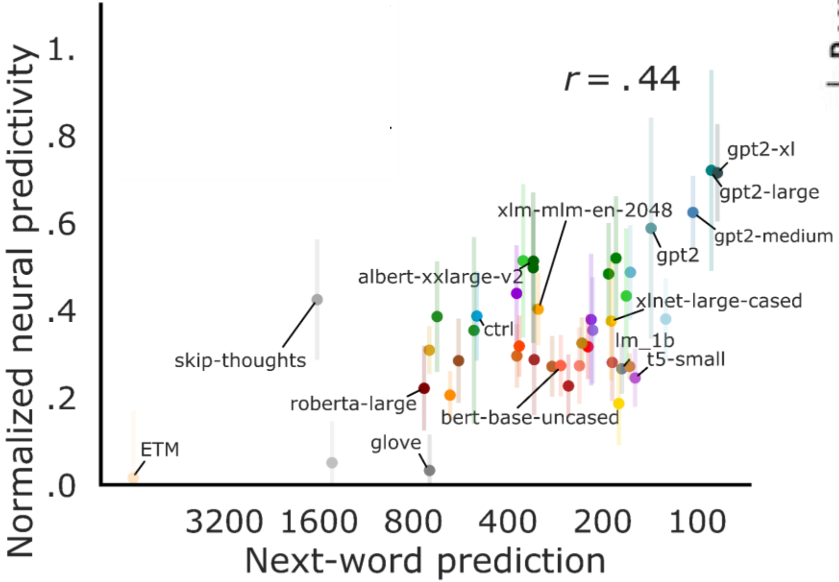

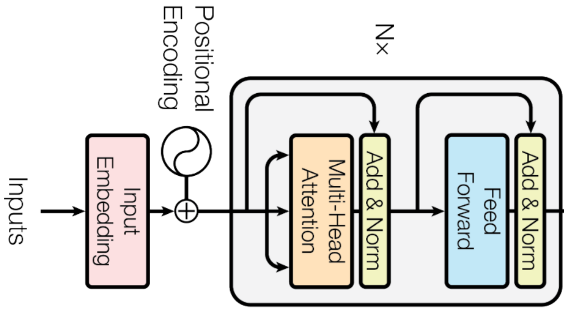

Schrimpf et al. PNAS 2021

Using deep neural networks as goal-driven models of a system



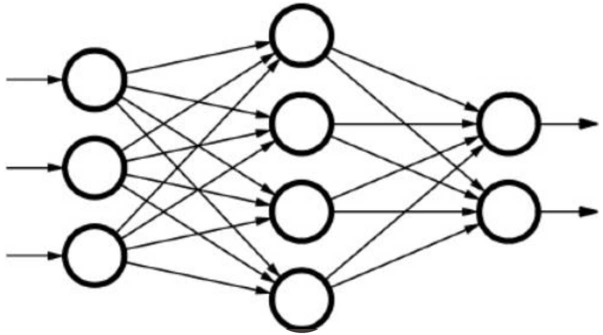
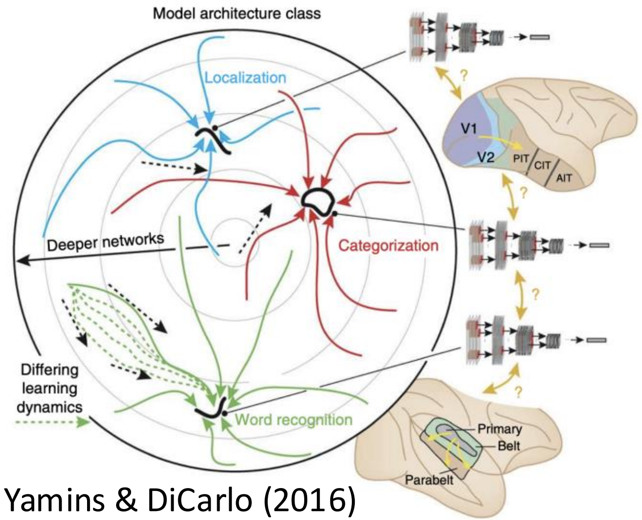

Yamins & DiCarlo (2016)


Vision: object recognition. Yamins & Hong et al. (2014), Schrimpf & Kubilius et al. (2018)


Audition: speech recognition, speaker & sound identification. Kell et al. (2018), …


Somatosentation: shape recognition. Zhuang et al. (2017)


Language: next-word prediction. Schrimpf et al. (2021), …


Decision making: context-dependent choice. Mante & Sussilo et al. (2013),


Proprioception: action recognition. Sandbrink et al. (2023)

Statistical models

- At this point, we can build powerful models (task-driven, data-driven)
- We can compare hypotheses and scaling of models (*why* questions)
- **Enables causal experiments**

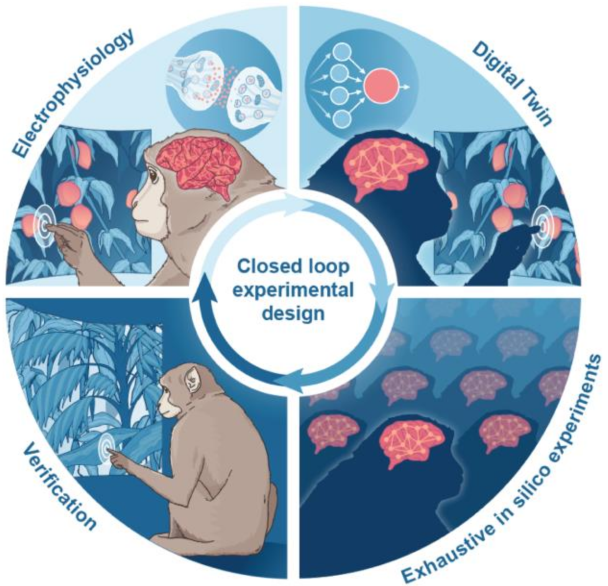

M. Mathis, Perez Rotondo, Cheng, Tolias, Mathis Cell 2024

Emergence of brain-like tuning in path integration and motor control

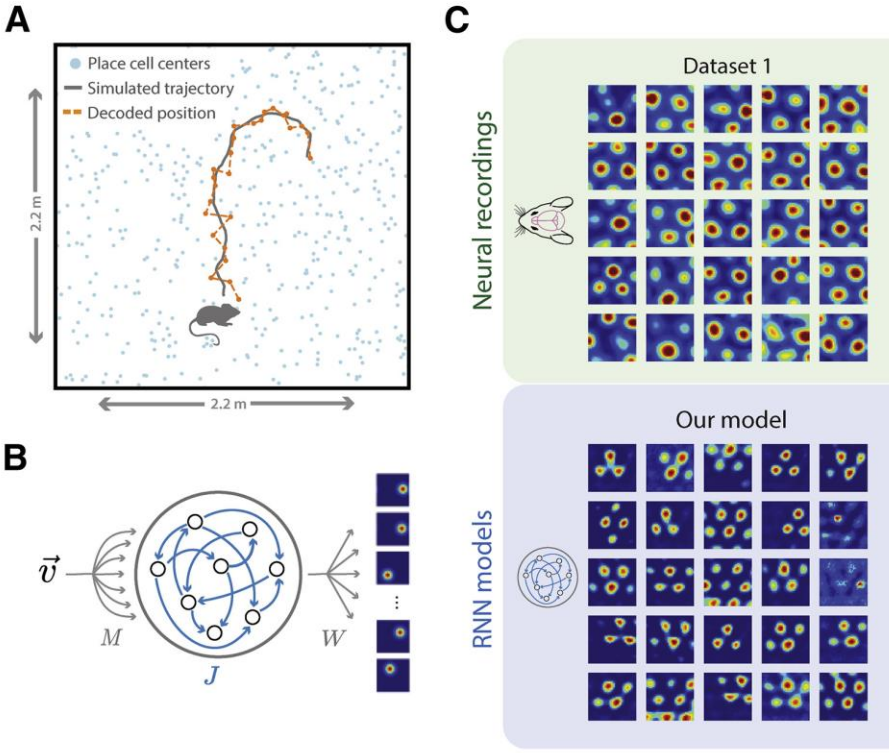

Sorscher et al. Neuron 2022

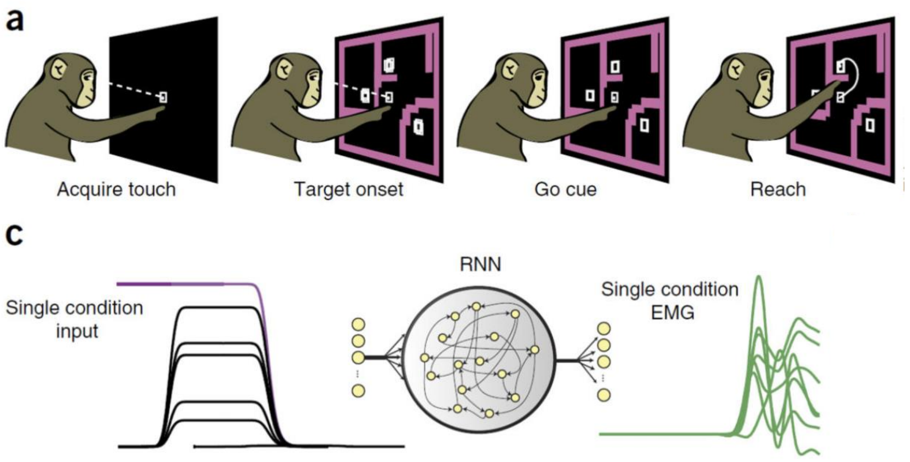

Susillo et al. Nature Neuro 2015

Inverse models of perception and motor control

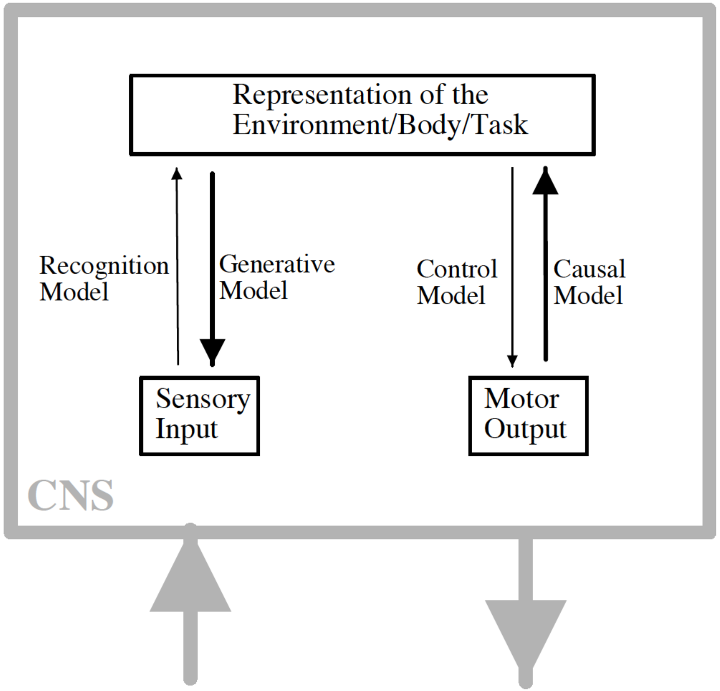

$P(G|M)$ causal model

$P(M)$ Movement prior

$$P(M|G)$$

Objective: find motor commands with high posterior probability!

The thin arrows correspond to the the directions that are desirably but *harder to implement*!

The thick arrows correspond to well-defined (relatively simpler transformations);
e.g., - generative model of vision: given the state of the world, predicting the retinal image (Optics,…)

- - causal model: given a motor command we can predict how it will change the world (Newtonian physics, ..)

Todorov, 1998

Optimal feedback control (OFC) theory

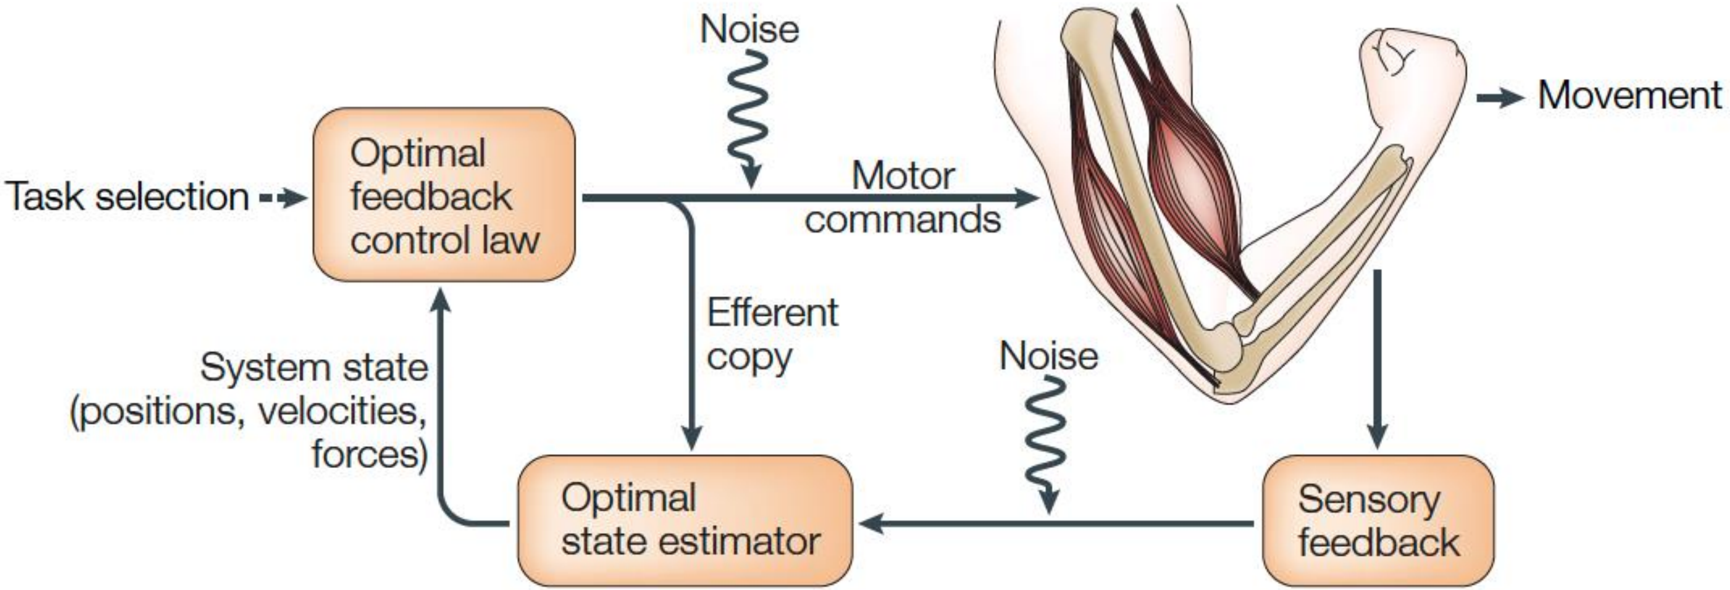

Scott, Nature Reviews Neuroscience 2004

OFC predicts many behavioral features of motor control

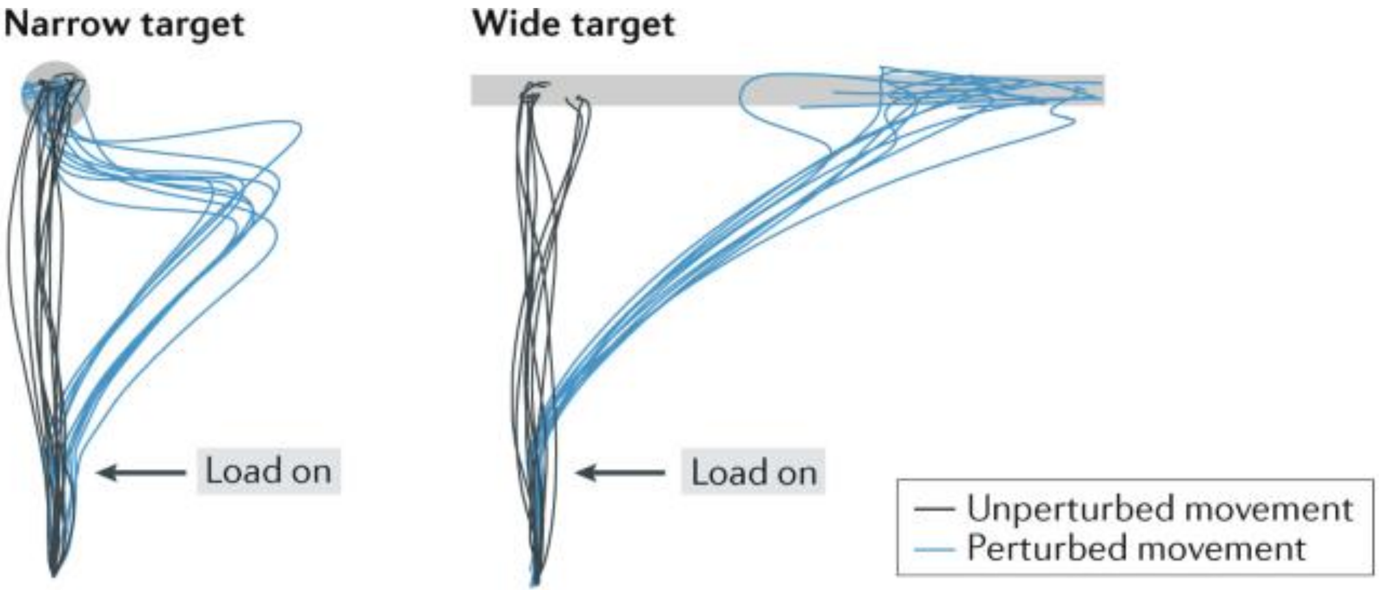

Nashed, Crevecoeur & Scott J. Neuro 2012

Building brain-like models (bottom-up)

DMAP's brain inspired architecture

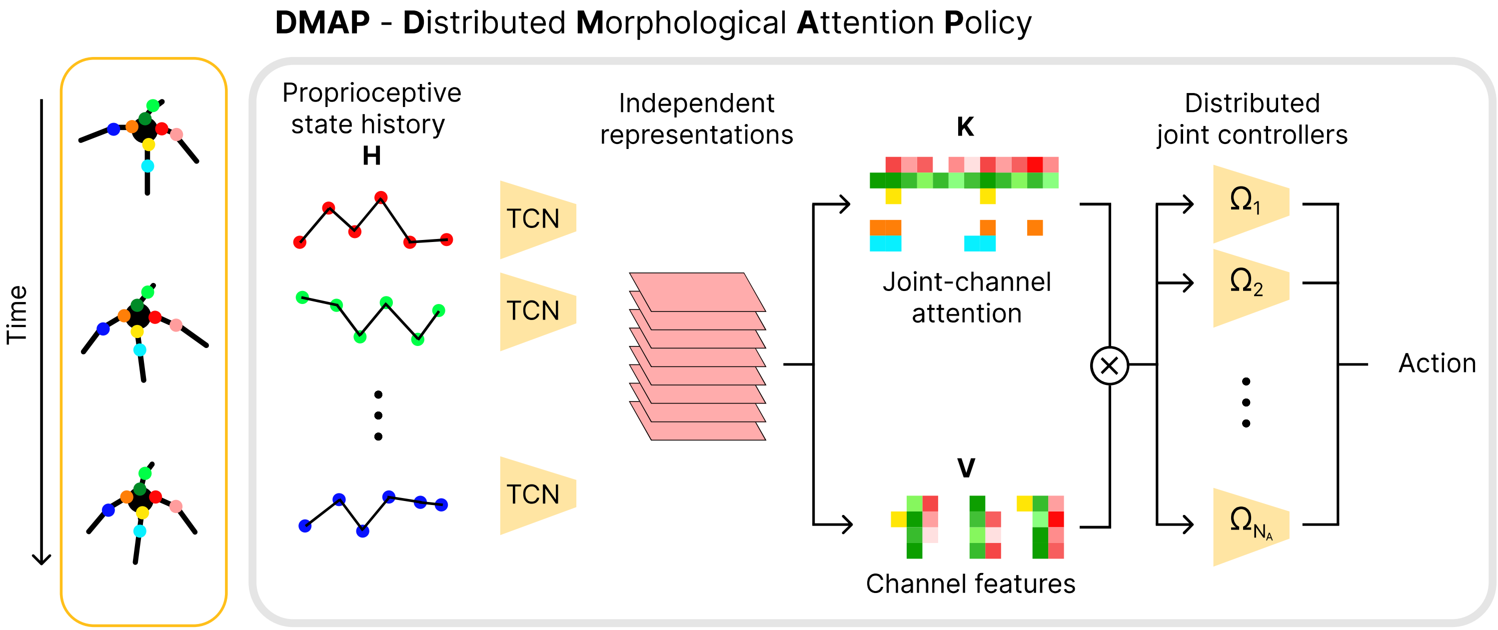

Chiappa, Marin Vargas, Mathis, Neurips 2022 /arxiv

Natural language instructions induce compositional generalization in networks of neurons

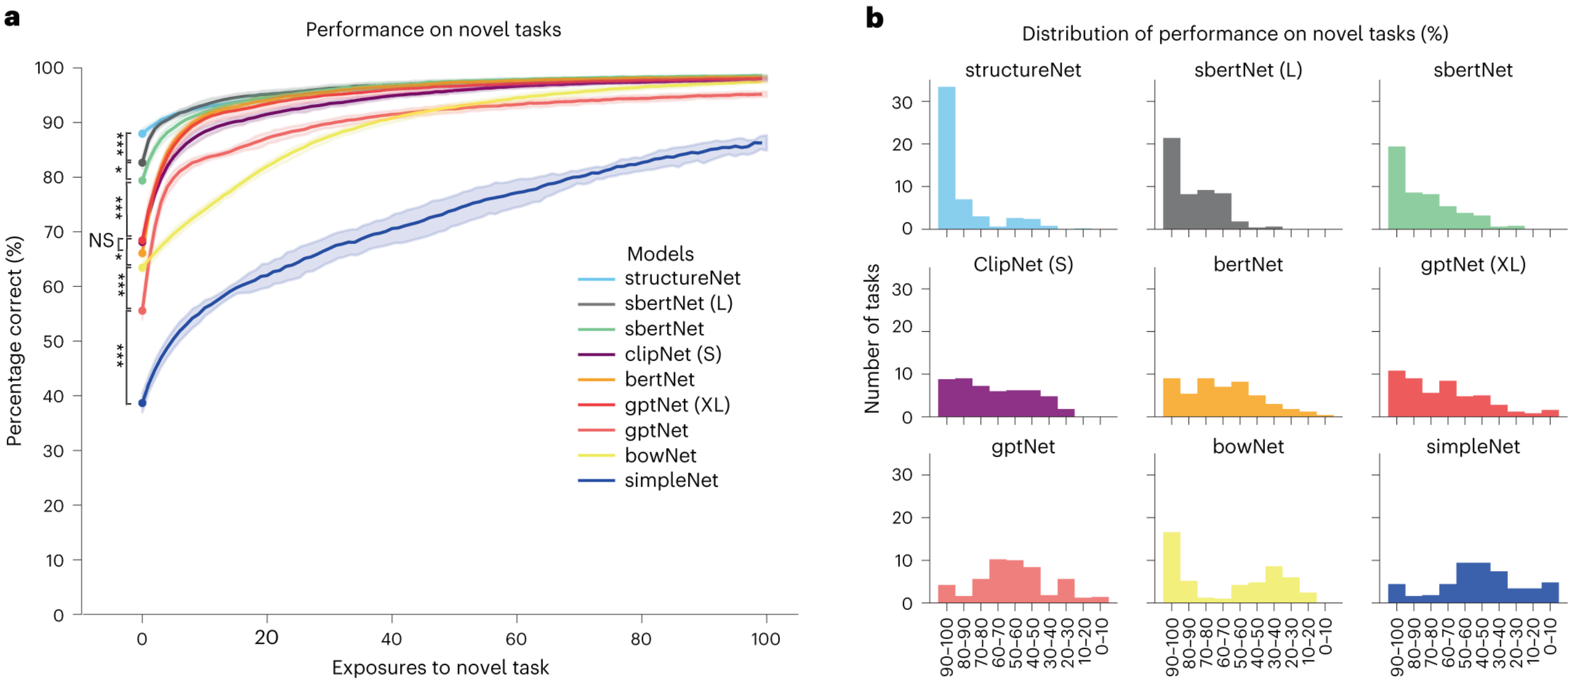

Riveland & Pouget, Nature Neuro 2024

Curriculum learning

- Static to Dynamic Stability (SDS)
  - SDS creates stability at desired states *before* learning a policy that reaches them
  - A curriculum gradually transforms static stability into dynamic movement motifs

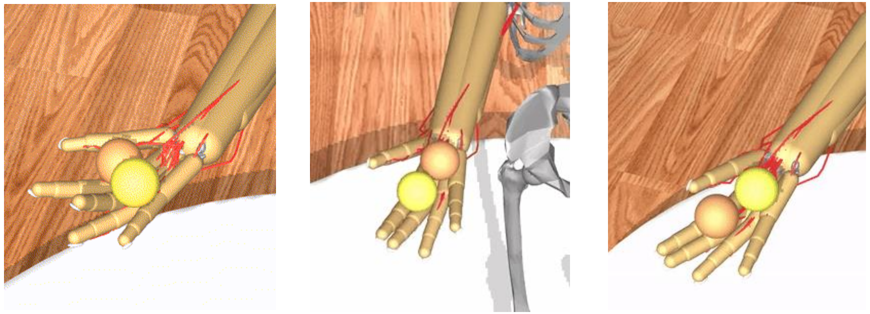

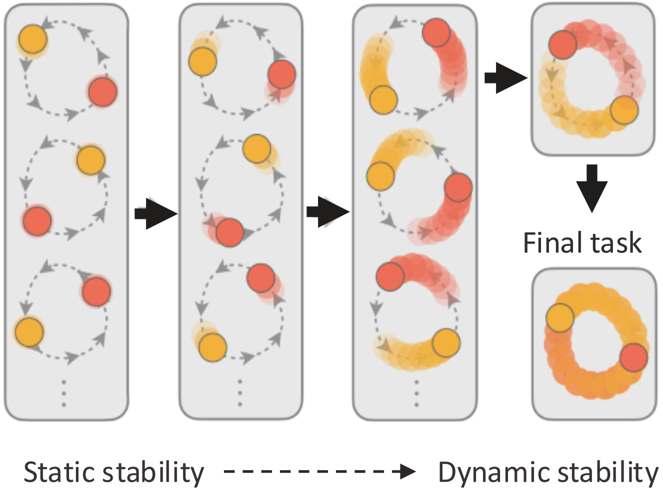

Chiappa, A. S., Tano*, P., Patel*, N., Ingster, A., Pouget, A., & Mathis, A. Neuron 2024*
Caggiano et al. Proceedings of the NeurIPS 2022 Competitions Track, PMLR 220:233-250

Artificial and biological intelligence

The rise of artificial intelligence…

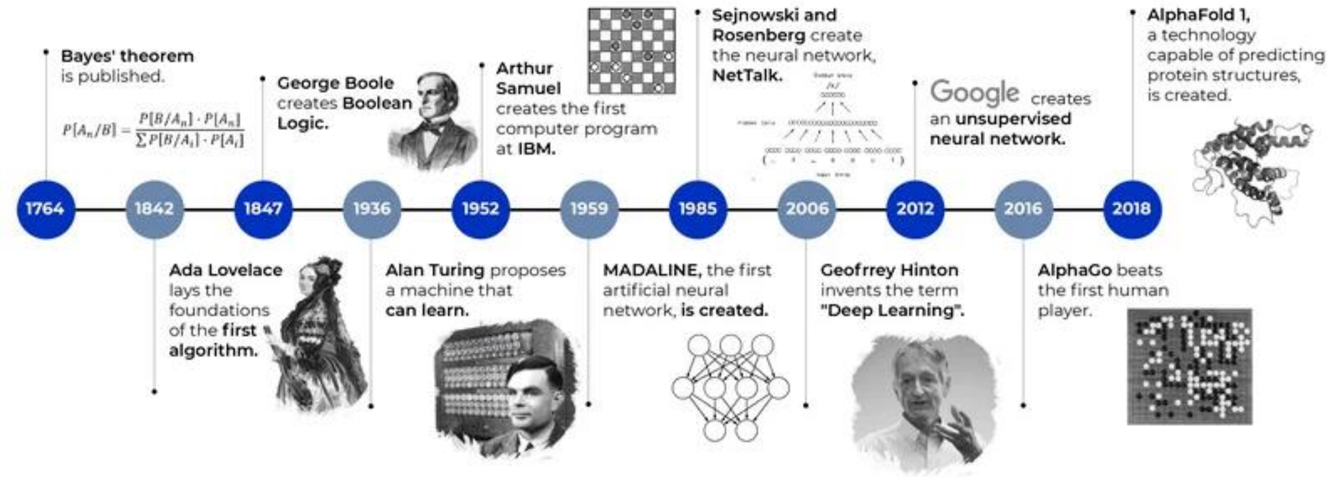

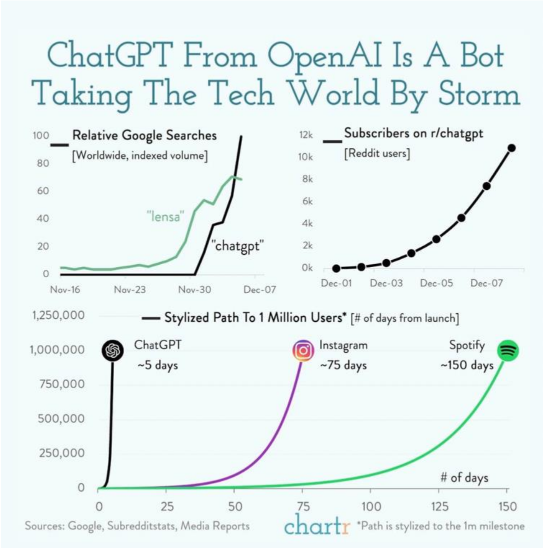

Image source: https://www.algotive.ai/blog/machine-learning-what-is-ml-and-how-does-it-work

Remember, the bitter lesson…

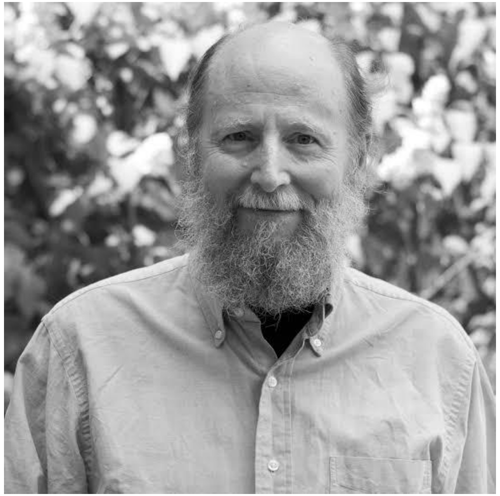

Richard Sutton

"The bitter lesson is based on the historical observations that 1) AI researchers have often tried to build knowledge into their agents, 2) this always helps in the short term, and is personally satisfying to the researcher, but 3) in the long run it plateaus and even inhibits further progress, and 4) breakthrough progress eventually arrives by an opposing approach based on scaling computation by search and learning. *The eventual success is tinged with bitterness, and often incompletely digested, because it is success over a favored, human-centric approach."*

http://www.incompleteideas.net/IncIdeas/BitterLesson.html

The human brain has a lot of parameters….

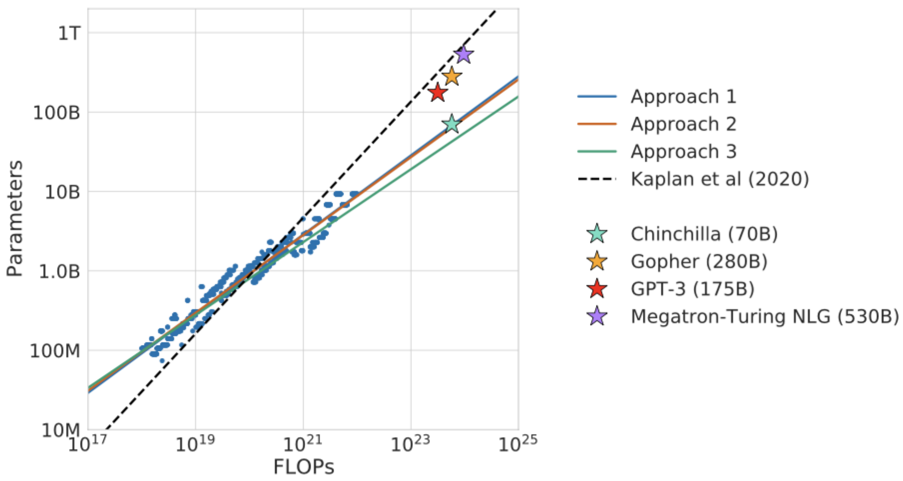

https://sebastianraschka.com/blog/2023/llm-reading-list.html

GPT-3 has 175-billion parameter

Human cortex has

~250 million synapses/mm3

Sherwood et al. Cerebral Cortex 2020

Is scale is all you need?

$$\text{PF-day} = 10^{15} \times 24 \times 3600 = 8.64 \times 10^{19} \text{ floating point operations.}$$

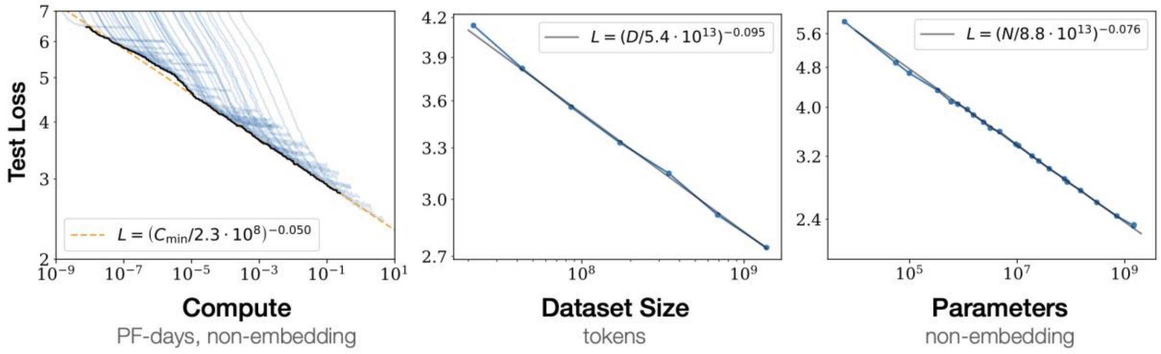

Figure 1 Language modeling performance improves smoothly as we increase the model size, datasetset size, and amount of compute[2] used for training. For optimal performance all three factors must be scaled up in tandem. Empirical performance has a power-law relationship with each individual factor when not bottlenecked by the other two.

Kaplan et al. 2020

Schaeffer, Miranda & Koyejo, NeurIPS 2023

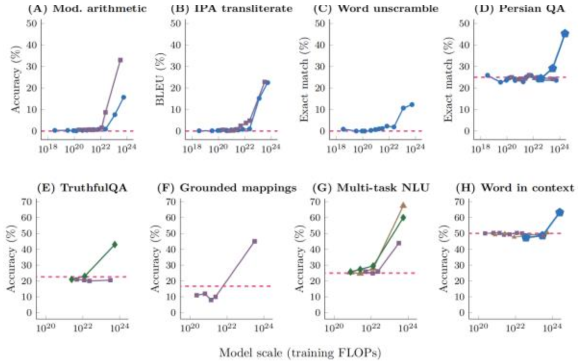

Figure 1: **Emergent abilities of large language models**. Model families display *sharp* and *unpredictable* increases in performance at specific tasks as scale increases. Source: Fig. 2 from [33].

Read this paper, they have a great critique of what the emergent abilities mean!

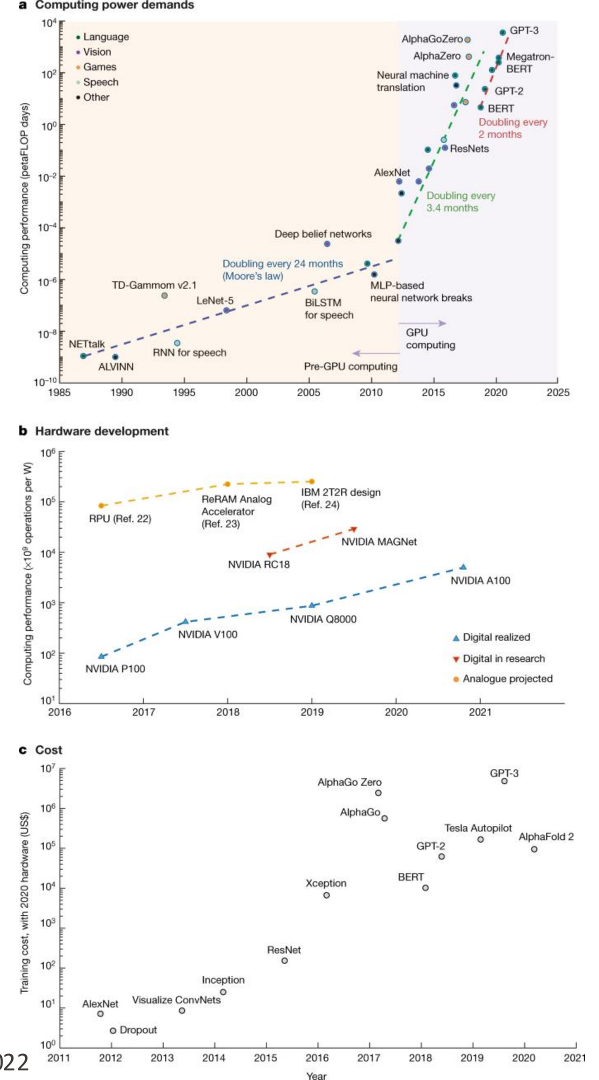

Mehonic & Kenyon Nature 2022

Scale seems to be all you need...

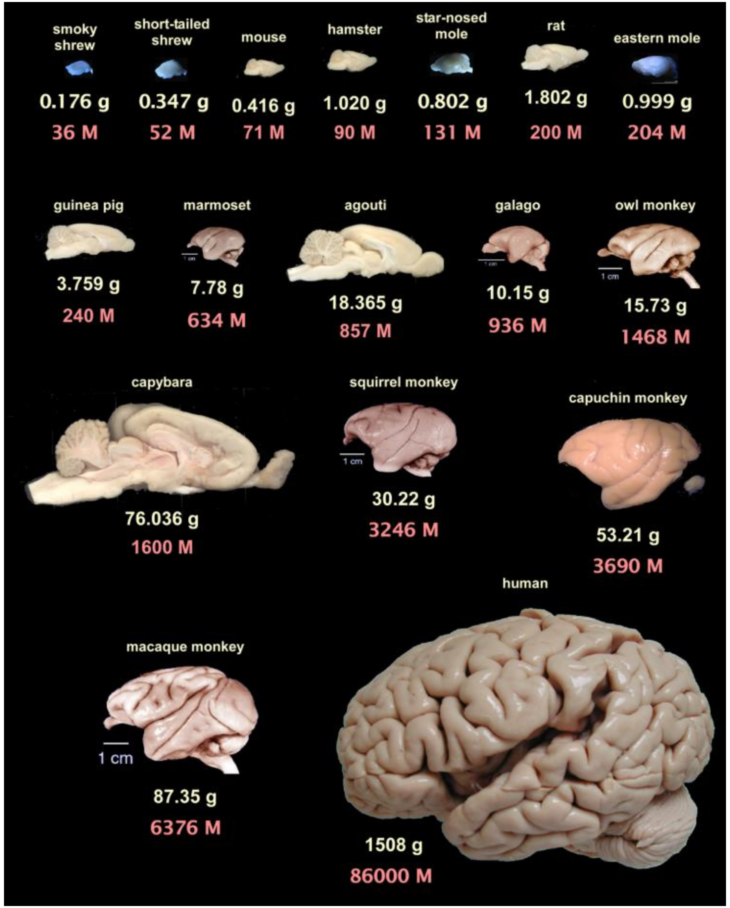

Note: numbers of neurons increase faster in the cerebral cortex and cerebellum than in the remaining brain areas...

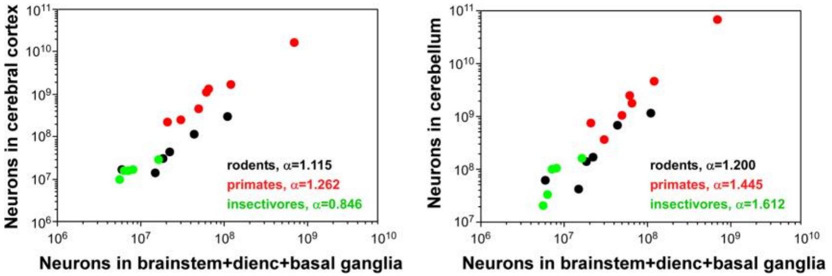

Herculano-Houzel, Frontiers Human Neurosci, 2009

Is it just a question of scaling AI up?

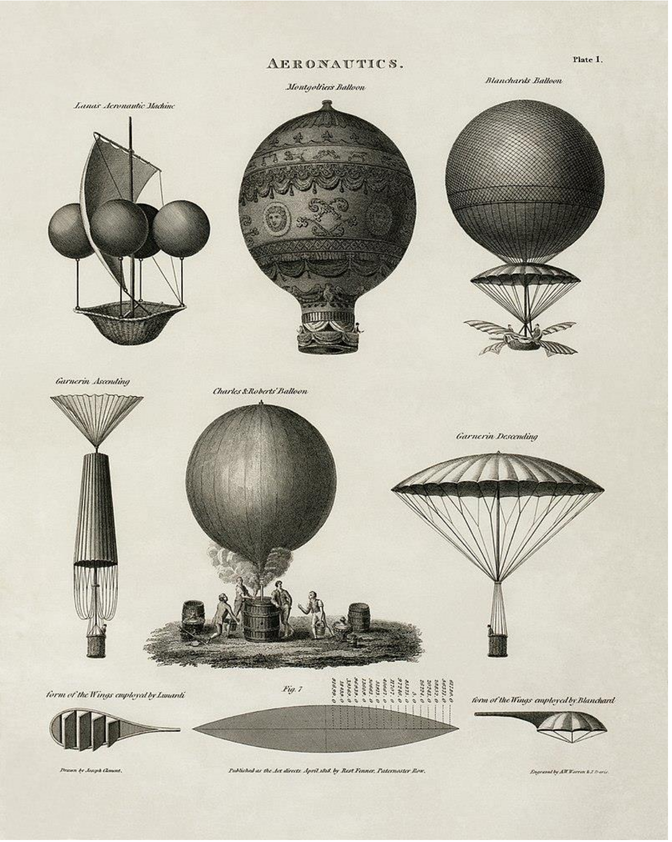

This 1818 technical illustration shows early balloon designs – Wikipedia.

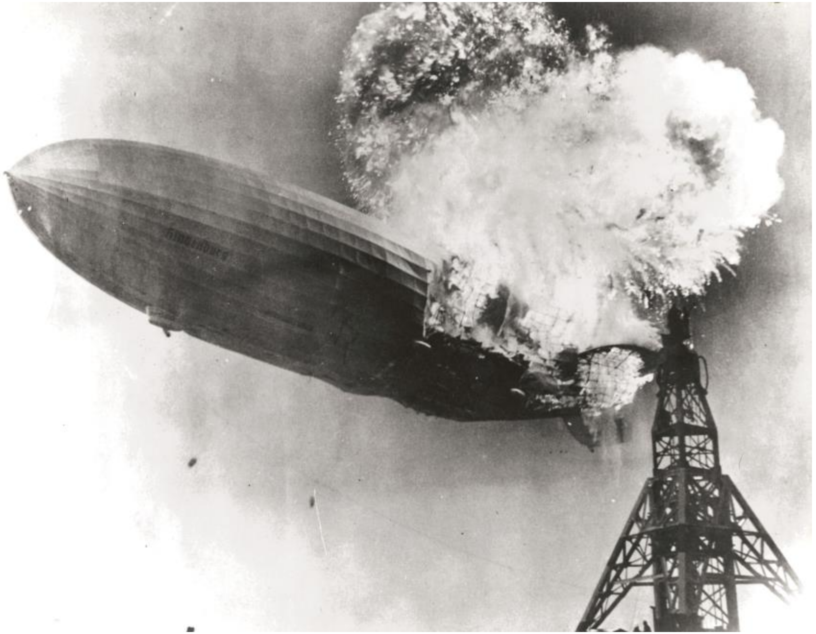

Hindenburg disaster, 1937

[Analogy from great talk by McCleland: “Can we capture intelligence in a neural network? Professor Jay McClelland “]()

What is missing to understand biological intelligence?

Ingredients for closing the gap that we discussed

- Internal models
- Inductive biases (innate architecture)
- Better exploration
- Baked in reward functions (which we don’t know…)
- Using language
- Curriculum learning
- Deliberate practice
- ….

EPFL

What was this course about?

Biological Intelligence ⟷ Artificial Intelligence

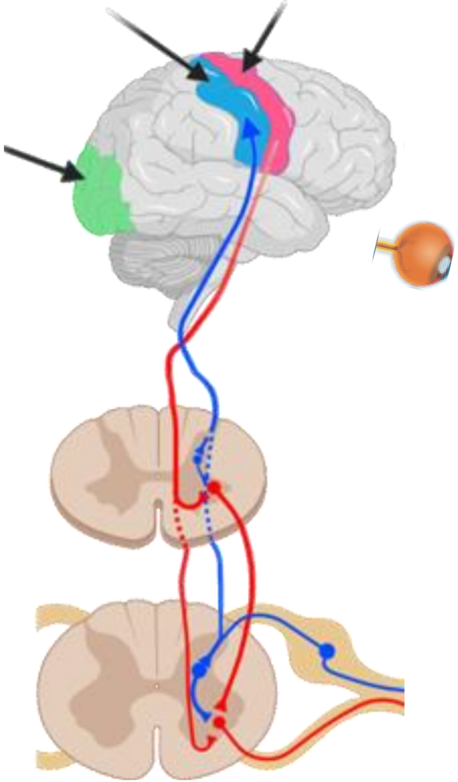

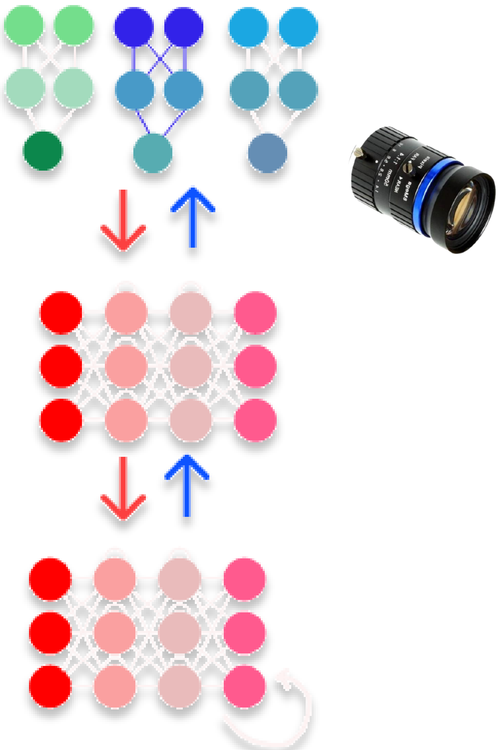

Hausmann*, Marin-Vargas*, Mathis**, Mathis**, Curr op. in Neuro., 2021

Artificial and biological intelligence

Neuroscience and machine learning have a “long, intertwined history”[1]….

…. and exciting future … join the journey!

E.g., check out this review that used this phrase: Neuroscience-Inspired Artificial Intelligence by Hassabis et al., Neuron 2017

EPFL

Biological intelligence and natural intelligence are *in* vogue …

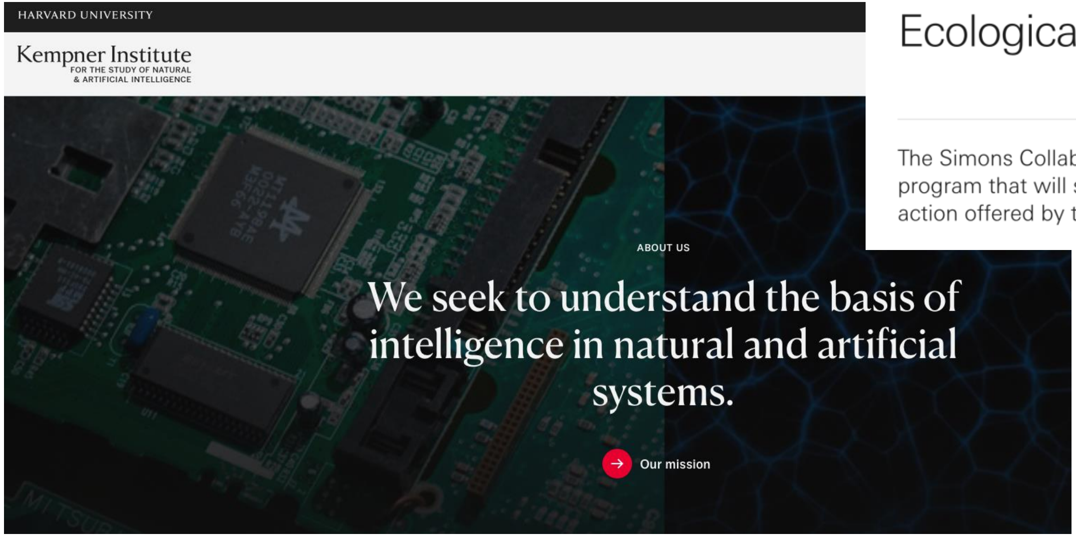

Simons Foundation Launches Collaboration on Ecological Neuroscience

The Simons Collaboration on Ecological Neuroscience (SCENE) is a 10-year program that will support projects aimed at uncovering how opportunities for action offered by the world shape representations in the mind and the brain.

MIT Quest for Intelligence

MAX PLANCK INSTITUTE
FOR BIOLOGICAL INTELLIGENCE

Logistics

Grading

Assessment methods

- The final mark is a combination of three evaluations:
  - *problem sets (25%)* → *scores will be on Moodle by next week!*
  - *quizzes (25%)* → *scores are on Moodle!*
  - final exam (50%)
    - You can bring 1 A4 page cheat sheet (both sides can be used)

What did we cover?

Class	Date	Topic
1	19/02/2025	Introduction & neural code
2	26/02/2025	Normative models
3	05/03/2025	Bayes and Brain-like circuits
4	12/03/2025	Task-driven models (Path integration)
5	19/03/2025	Task-driven models (Vison)
6	26/03/2025	Task-driven (Unsupervised, Audition, metamers, optimal stimuli)
7	02/04/2025	Task-driven Somatosensation
8	09/04/2025	Language modeling in the brain I
9	16/04/2025	Language modeling in the brain II
10	23/04/2025	EPFL Easter break 🥚🌸
11	30/04/2025	Motor control
12	07/05/2025	Language modeling in the brain III (language in the service of cognition)
13	14/05/2025	Reinforcement learning
14	21/05/2025	Skill learning
15	28/05/2025	Review

Thanks to the team!

Martin Schrimpf

Thanks to our TAs:

Abdulkadir Gokce (PhD student)

Merkourios Simos (PhD Student)

Hossein Mirzaei (PhD Student)

Michael Hauri (NX student)

Class feedback, please fill it out on moodle!

This was the third time, we would greatly appreciate comments!

https://moodle.epfl.ch/my/courses.php

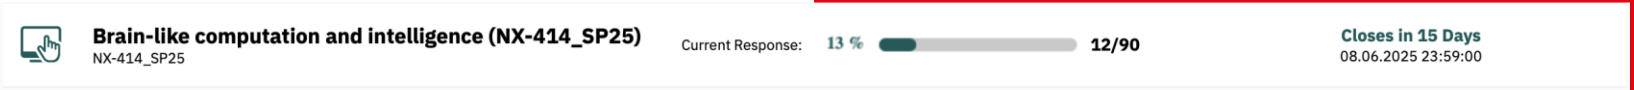

Thank you for attending!

Thanks to our TAs:

Abdulkadir Gokce (PhD student)

Merkourios Simos (PhD Student)

Hossein Mirzaei (PhD Student)

Michael Hauri (NX student)