VS 265 - Neural Computation

Bruno A. Olshausen, Instructor
baolshausen@berkeley.edu

Sophia Sanborn, GSI
sophia.sanborn@gmail.com

Class meets TTH 3:30-5 online

Challenge problem assignments (60% of grade)

Final Project (40% of grade)

Readings:

- Handouts
- Hertz, Krogh & Palmer, *Introduction to the Theory of Neural Computation*
- Dayan & Abbott, *Theoretical Neuroscience*
- MacKay, *Information Theory, Inference and Learning Algorithms*
- Sterling & Laughlin*, Principles of Neural Design.*

All reading materials and assignments on website at

http://redwood.berkeley.edu/courses/vs265

Piazza discussion forum

Readings for this week
(available on the website)

Today:

- Sterling & Laughlin, Chapter 1
- Bell, A.J. *Levels and loops: the future of artificial intelligence and neuroscience*. Phil Trans: Bio Sci. **354**:2013--2020 (1999)
- Dreyfus, H.L. and Dreyfus, S.E. *Making a Mind vs. Modeling the Brain: Artificial Intelligence Back at a Branchpoint*. Daedalus, Winter 1988.

Next week:

- Sterling & Laughlin, Chapters 2-4
- Solari & Stoner (2011) *Cognitive Consilience.*

- Mead*, Analog VLSI and Neural Systems, Chapter 1: Introduction and Chapter 4: Neurons*
- Carandini M, Heeger D (1994) *Summation and division by neurons in primate visual cortex.*

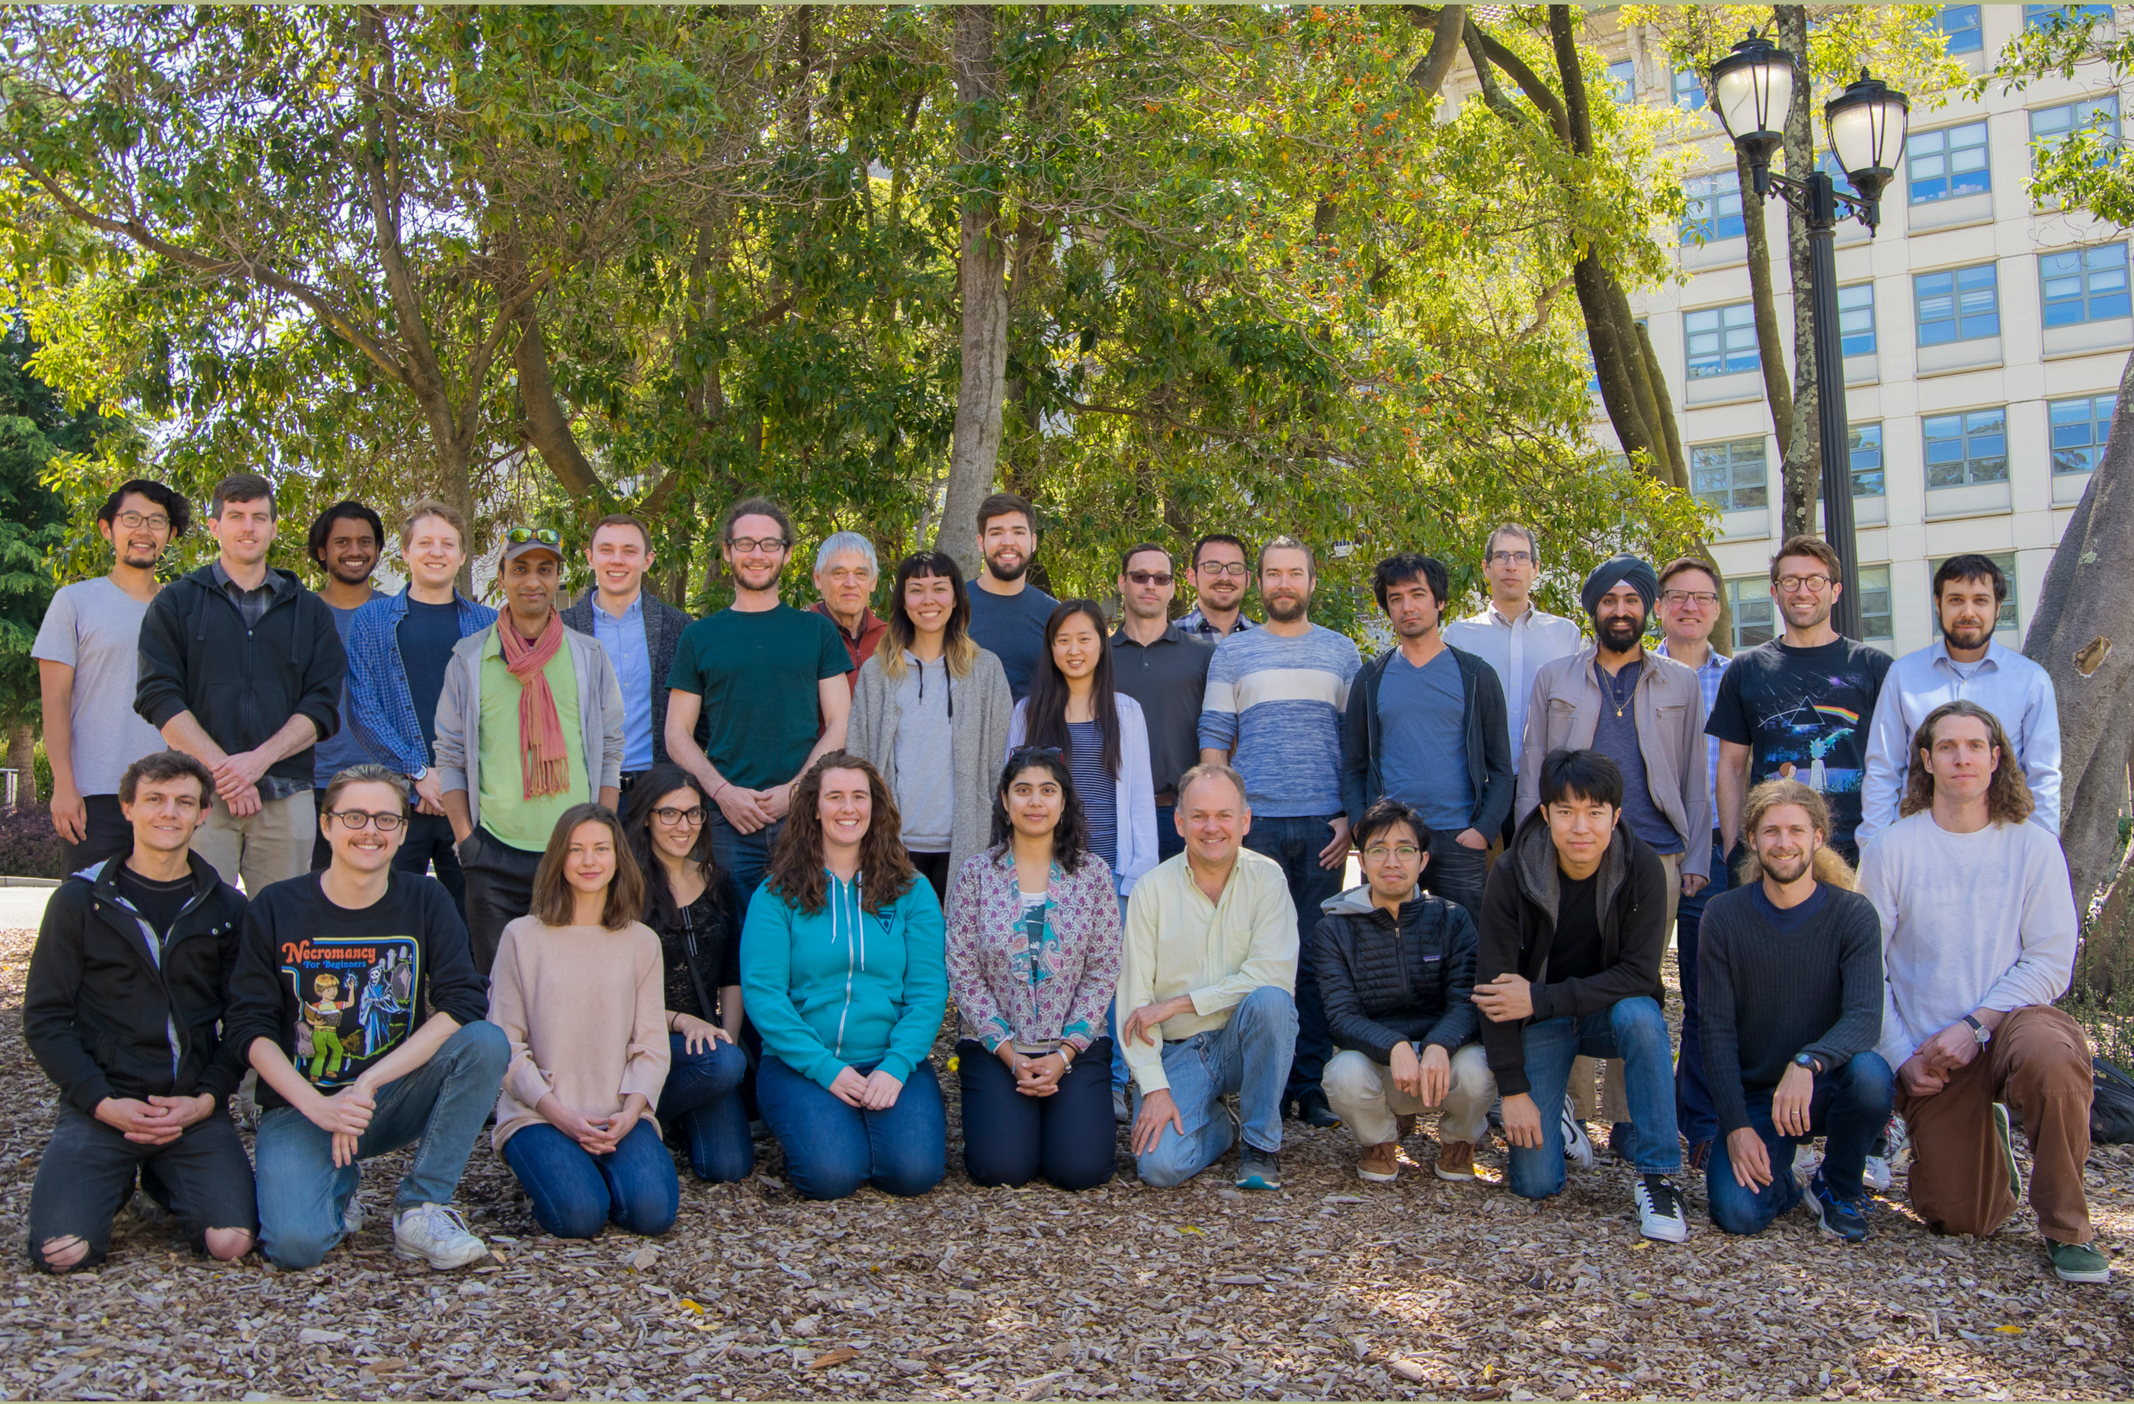

Redwood Center for Theoretical Neuroscience - April 2018

What have brain scans and single-unit recording taught us about the computations underlying perception and cognition?

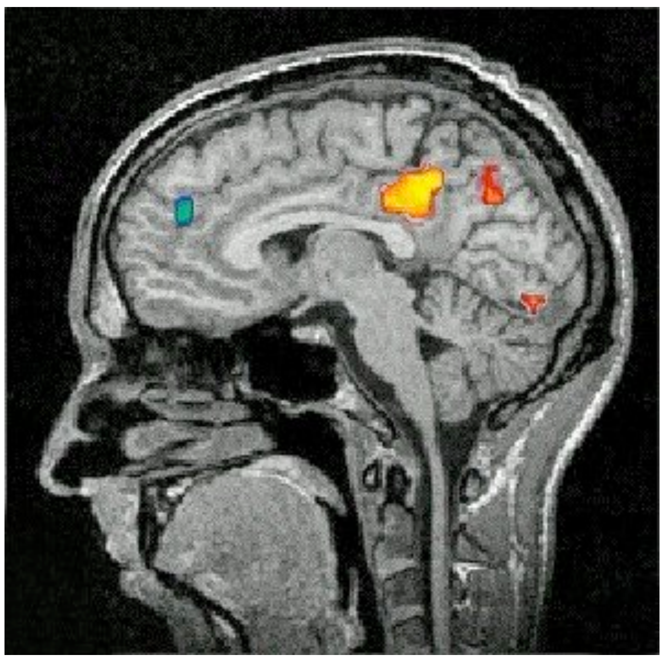

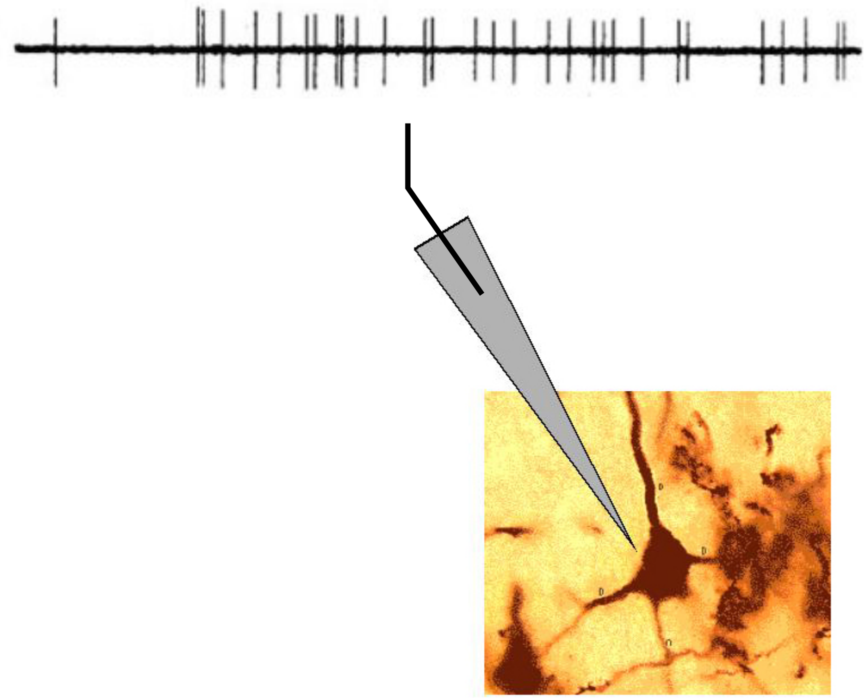

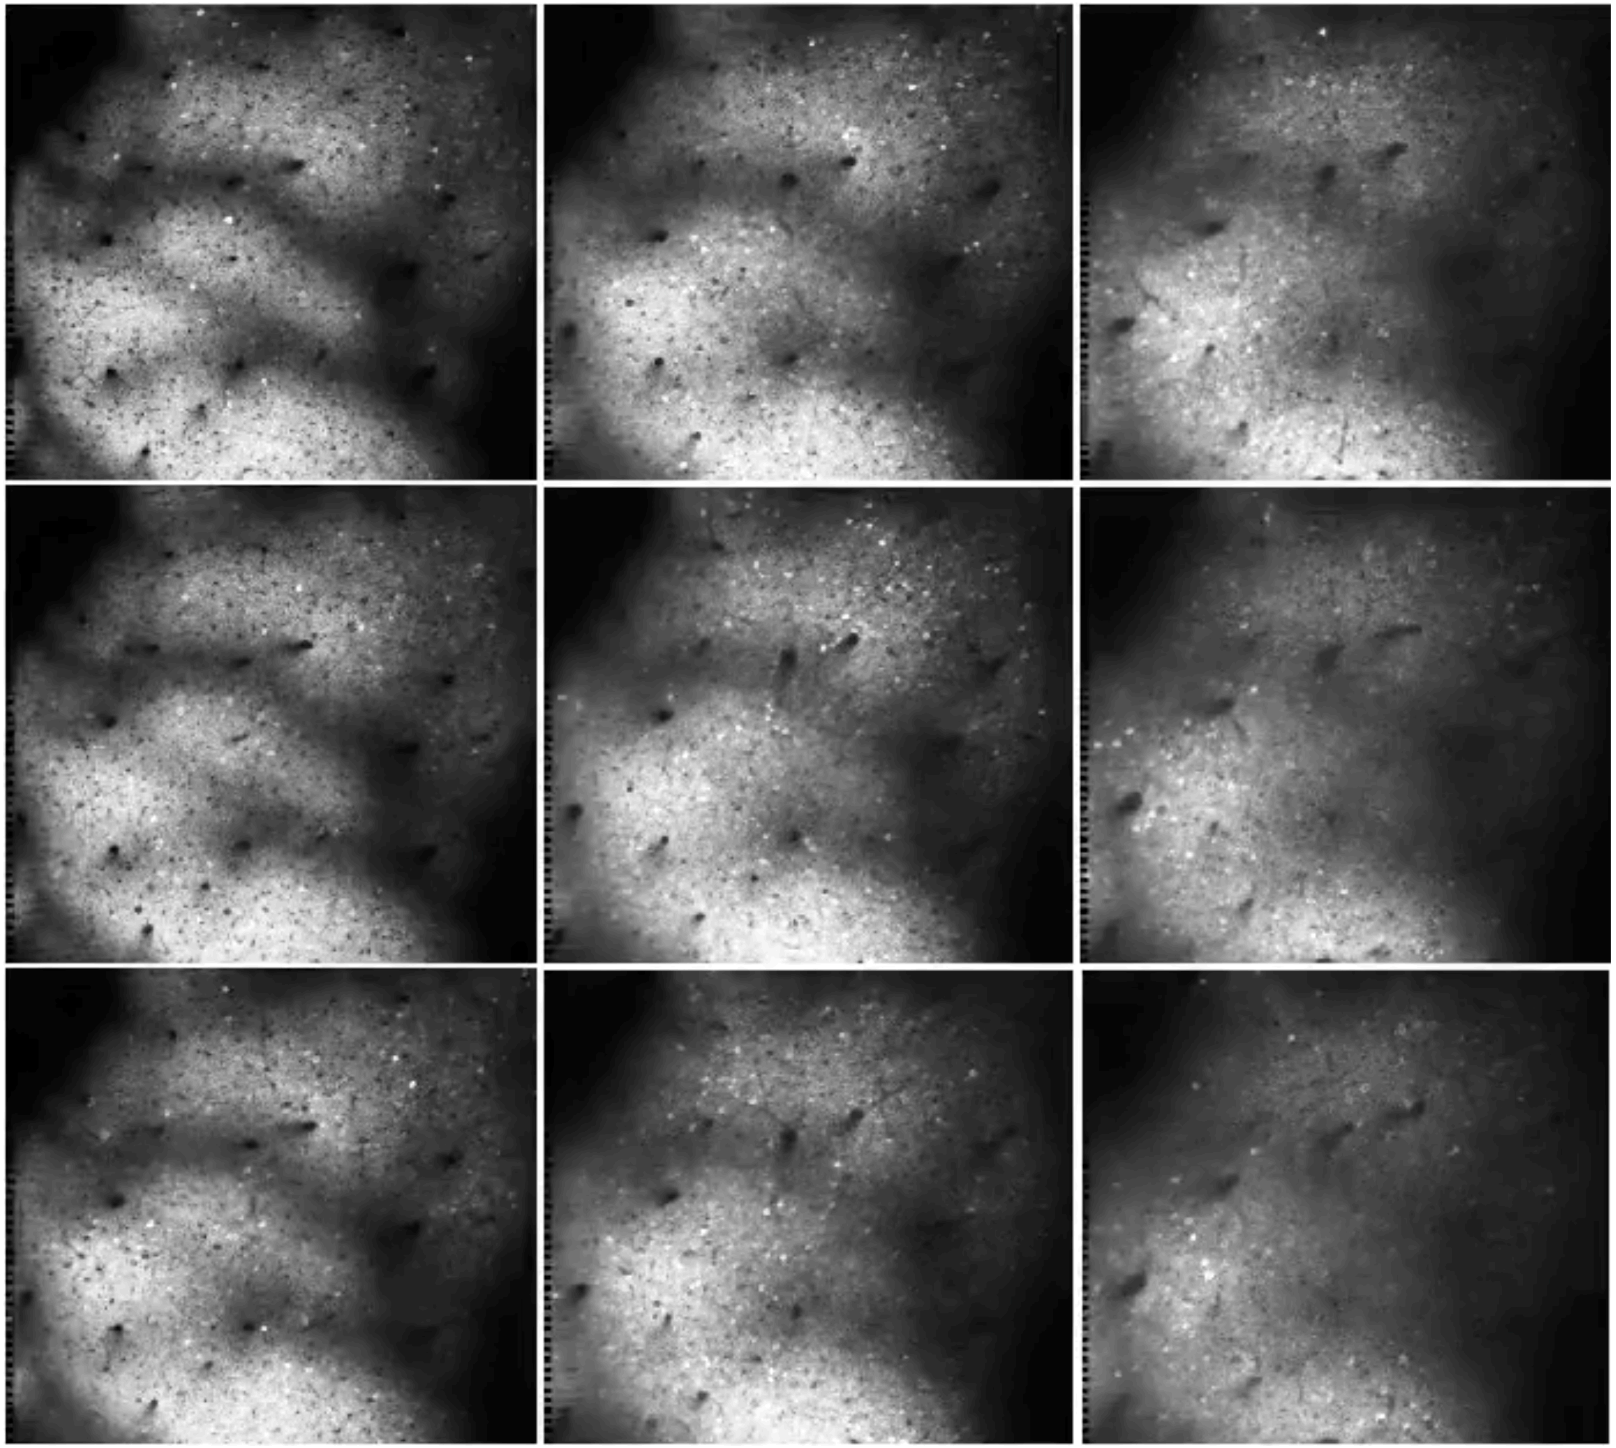

(from Stringer et al., 2019)

Why hasn't machine intelligence scaled with Moore's law?

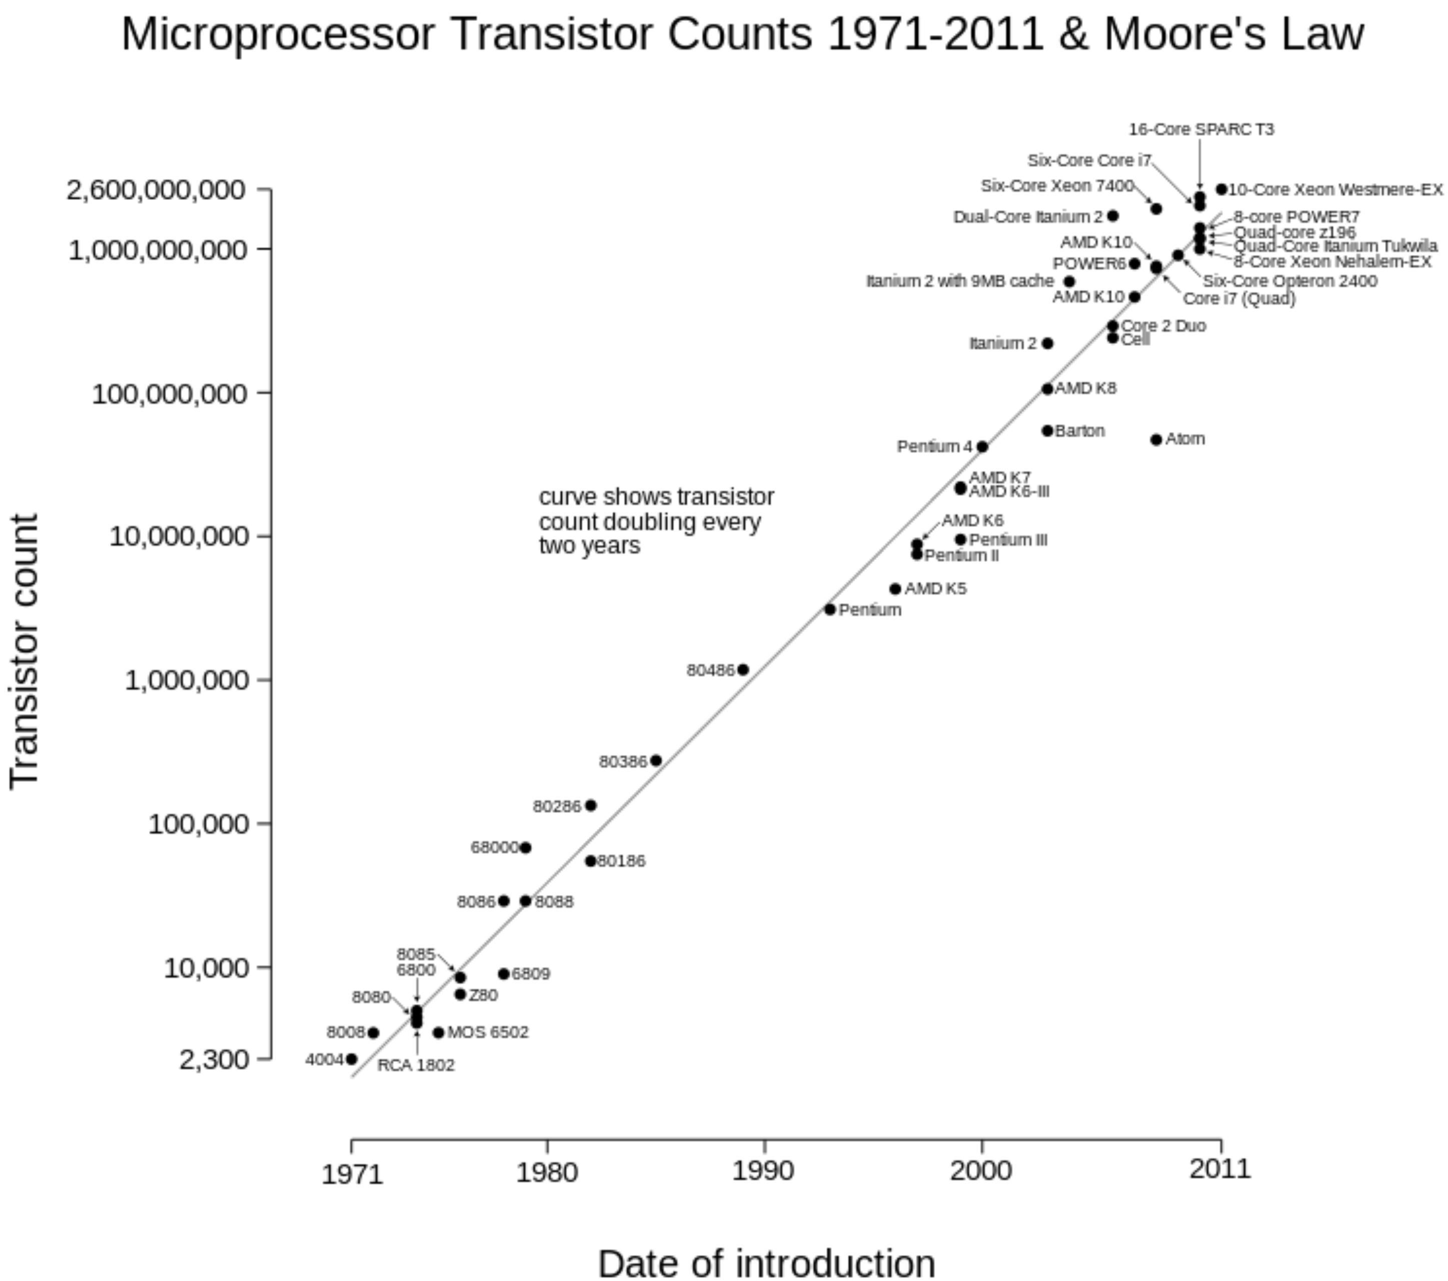

After **50 years** of concerted research efforts...

- there is little understanding of how neurons interact to process sensory information or to control actions.

- machines are still incapable of solving simple perceptual or motor control tasks.

We are missing something fundamental on both fronts: *we are ignorant of the underlying principles governing perception and action.*

How did we get here?

Artificial Intelligence

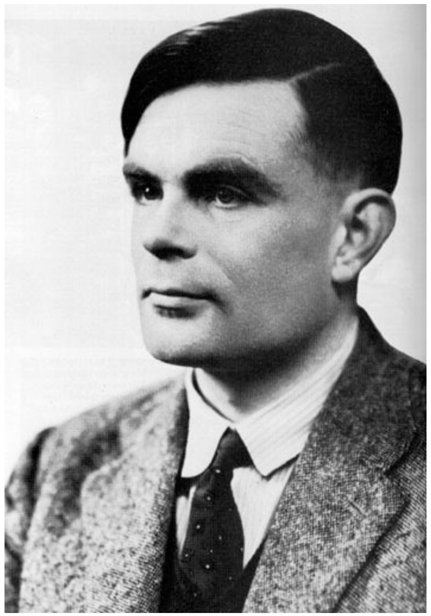

Alan Turing

John von Neumann

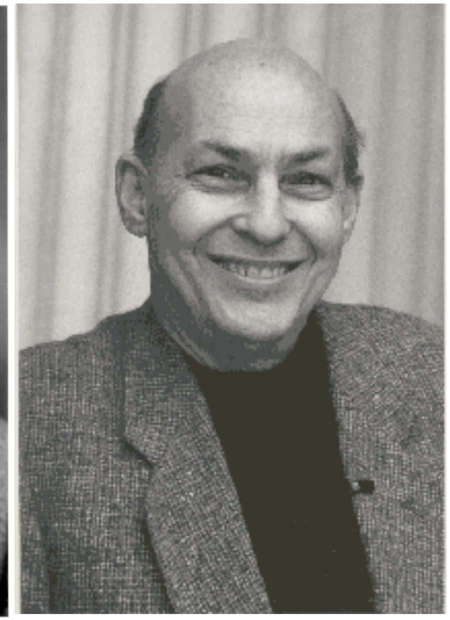

Marvin Minsky

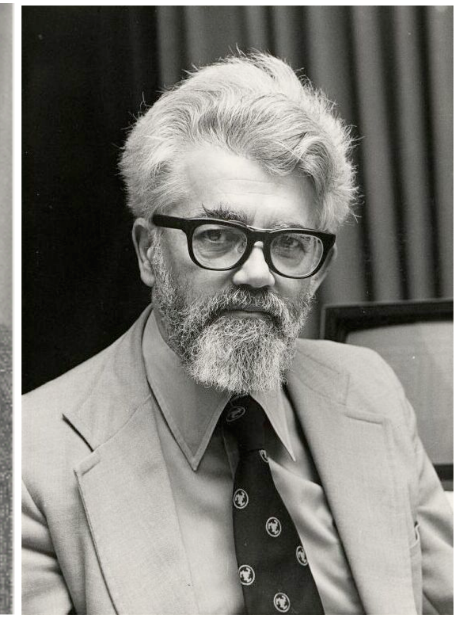

John McCarthy

Among the most challenging scientific questions of our time are the corresponding analytic and synthetic problems: How does the brain function? Can we design a machine which will simulate a brain?
-- *Automata Studies*, 1956

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

PROJECT MAC

Artificial Intelligence Group July 7, 1966
Vision Memo. No. 100.

THE SUMMER VISION PROJECT

Seymour Papert

The summer vision project is an attempt to use our summer workers effectively in the construction of a significant part of a visual system. The particular task was chosen partly because it can be segmented into sub-problems which will allow individuals to work independently and yet participate in the construction of a system complex enough to be a real landmark in the development of "pattern recognition".

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

PROJECT MAC

<u>Subgoal for July</u>

Analysis of scenes consisting of non-overlapping objects from the following set:

balls

bricks with faces of the same or different colors or textures

cylinders.

Each face will be of uniform and distinct color and/or texture. Background will be homogeneous.

<u>Extensions for August</u>

The first priority will be to handle objects of the same sort but with complex surfaces and backgrounds, e.g. cigarette pack with writing and bands of different color, or a cylindrical battery.

Then extend class of objects to objects like tools, cups, etc.

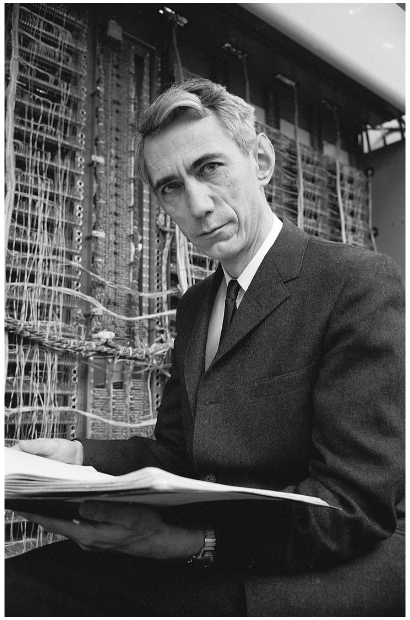

I confidently expect that within a matter of 10 or 15 years, something will emerge from the laboratory which is not too far from the robot of science fiction fame.
— Claude Shannon, 1961

Machines will be capable, within twenty years, of doing any work that a man can do.
— Herbert Simon, 1965

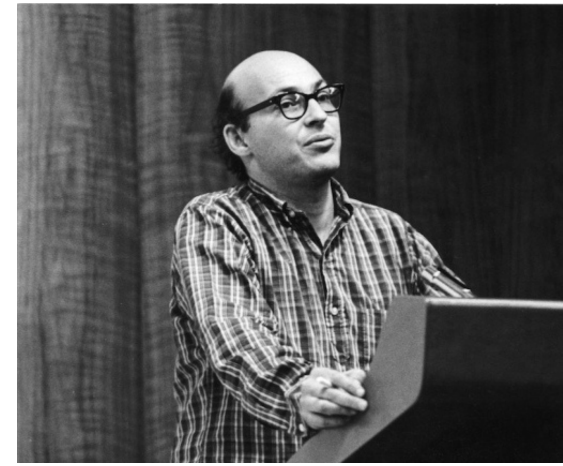

Within a generation...the problem of creating 'artificial intelligence' will be substantially solved.
— Marvin Minsky, 1967

The Lighthill debate (1973)

http://www.aiai.ed.ac.uk/events/lighthill1973/

Sir James Lighthill

vs.

The Lighthill debate (1973)

http://www.aiai.ed.ac.uk/events/lighthill1973/

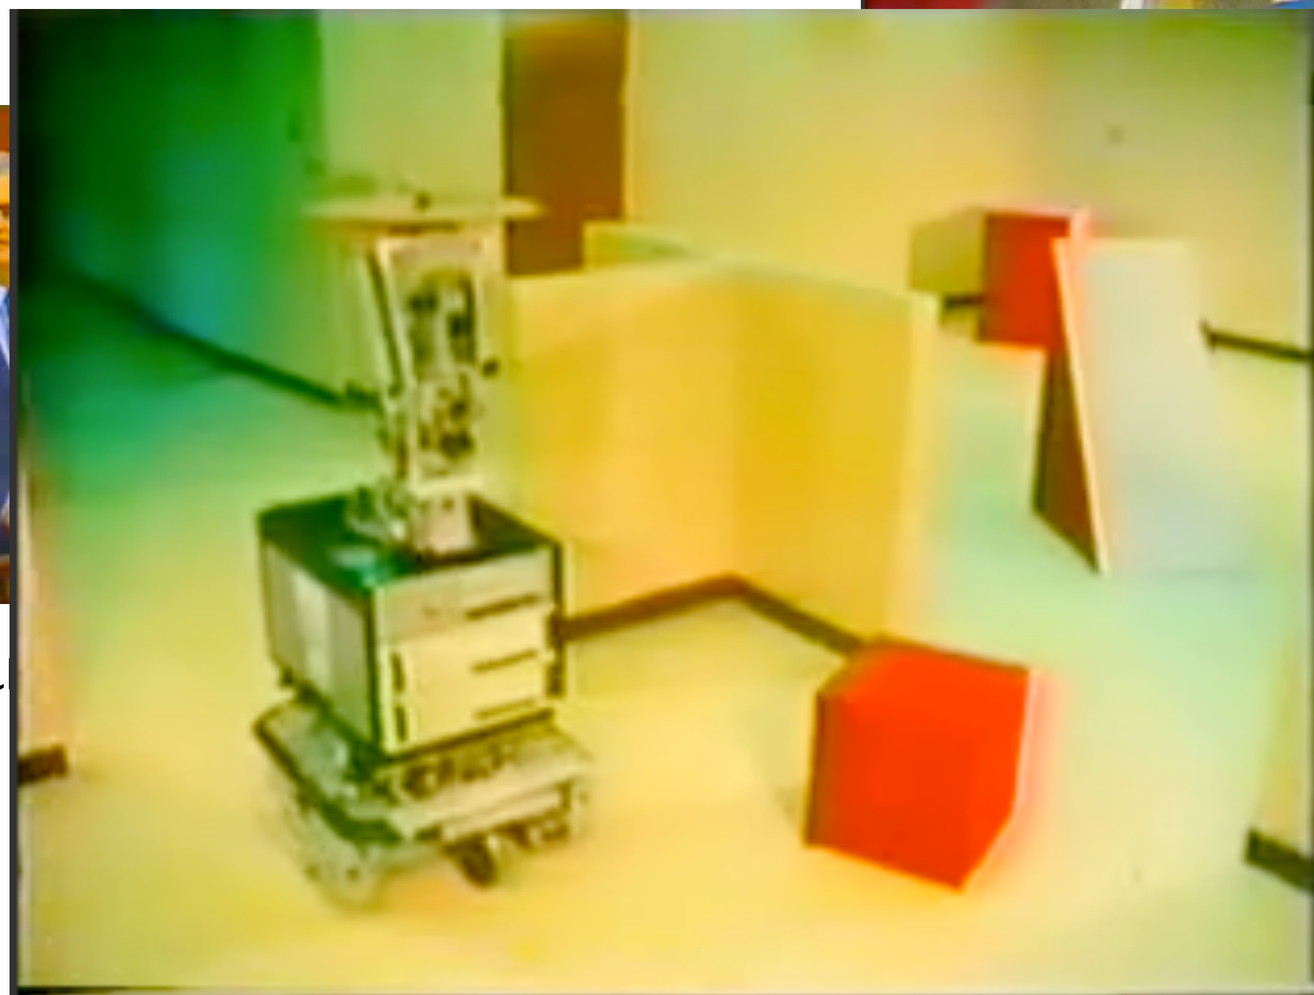

Sir Ja

Our first foray into Artificial Intelligence was a program that did a credible job of solving problems in college calculus. Armed with that success, we tackled high school algebra; we found, to our surprise, that it was much harder. Attempts at grade school arithmetic, involving the concept of numbers, etc., provide problems of current research interest. An exploration of the child's world of blocks proved insurmountable, except under the most rigidly constrained circumstances. It finally dawned on us that the overwhelming majority of what we call intelligence is developed by the end of the first year of life.

--Minksy, 1977

Even 'simple' nervous systems can exhibit profound visual intelligence

jumping spider

sand wasp

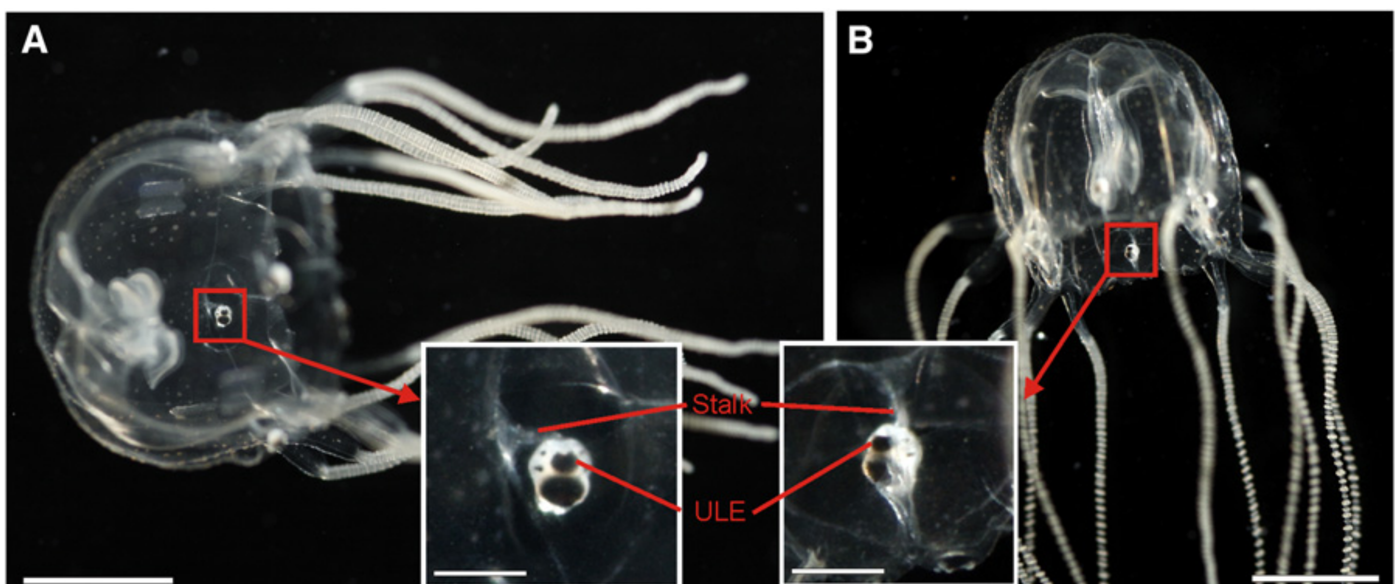

box jellyfish

Cybernetics/neural networks

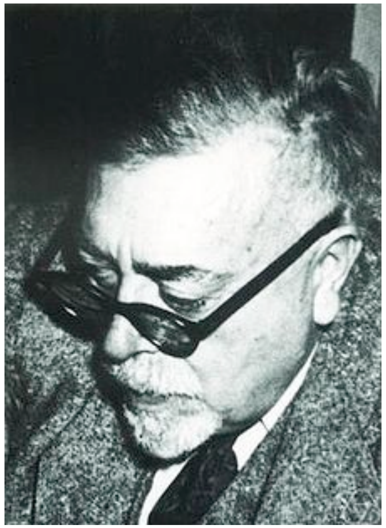

Norbert Wiener

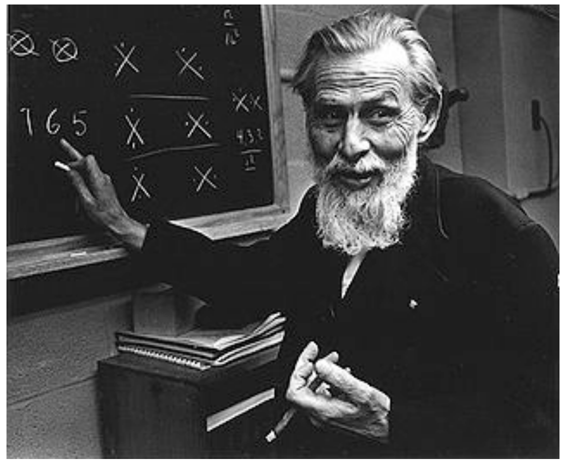

Warren McCulloch & Walter Pitts

Frank Rosenblatt

"The theory reported here clearly demonstrates the feasibility and fruitfulness of a quantitative statistical approach to the organization of cognitive systems. By the study of systems such as the perceptron, it is hoped that those fundamental laws of organization which are common to all information handling systems, machines and men included, may eventually be understood." -- Frank Rosenblatt

The Perceptron: A Probabilistic Model for Information Storage and Organization in the Brain. In, *Psychological Review*, Vol. 65, No. 6, pp. 386-408, November, 1958.

Perceptron model
(Rosenblatt, ca. 1960)

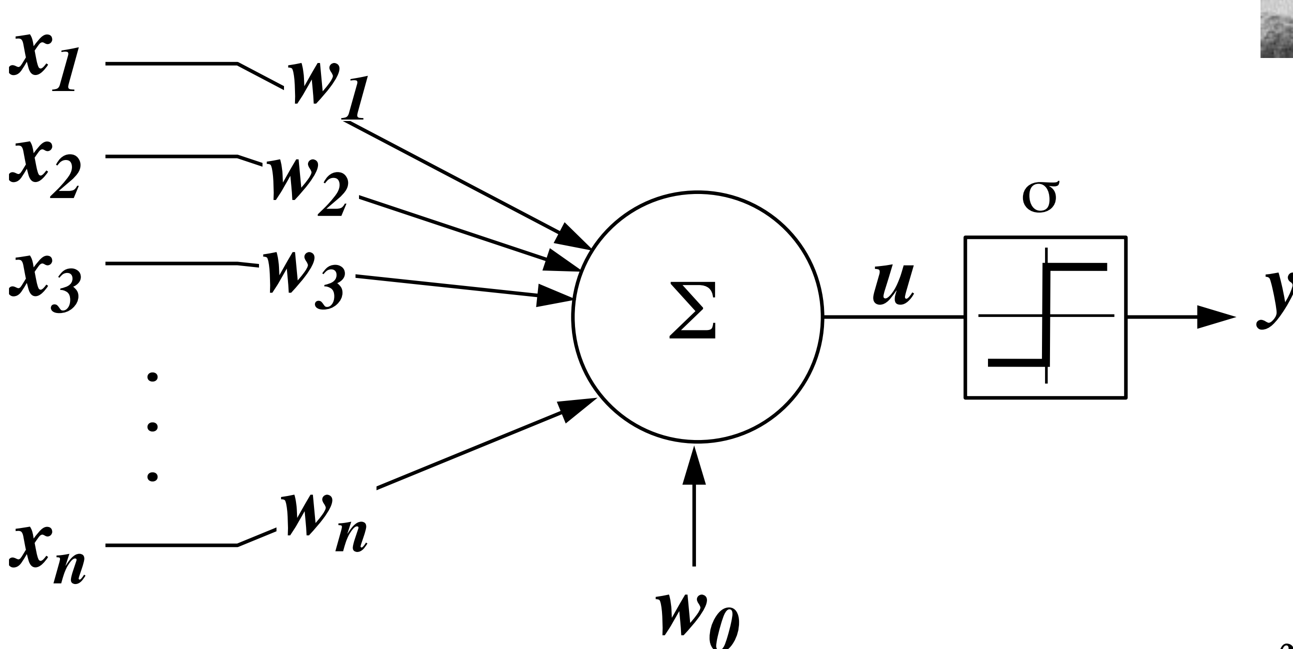

$$u = w_0 + \sum_{i=1}^{n} w_i \, x_i$$

$$y = \sigma(u)$$

Hubel & Wiesel (1962, 1965)

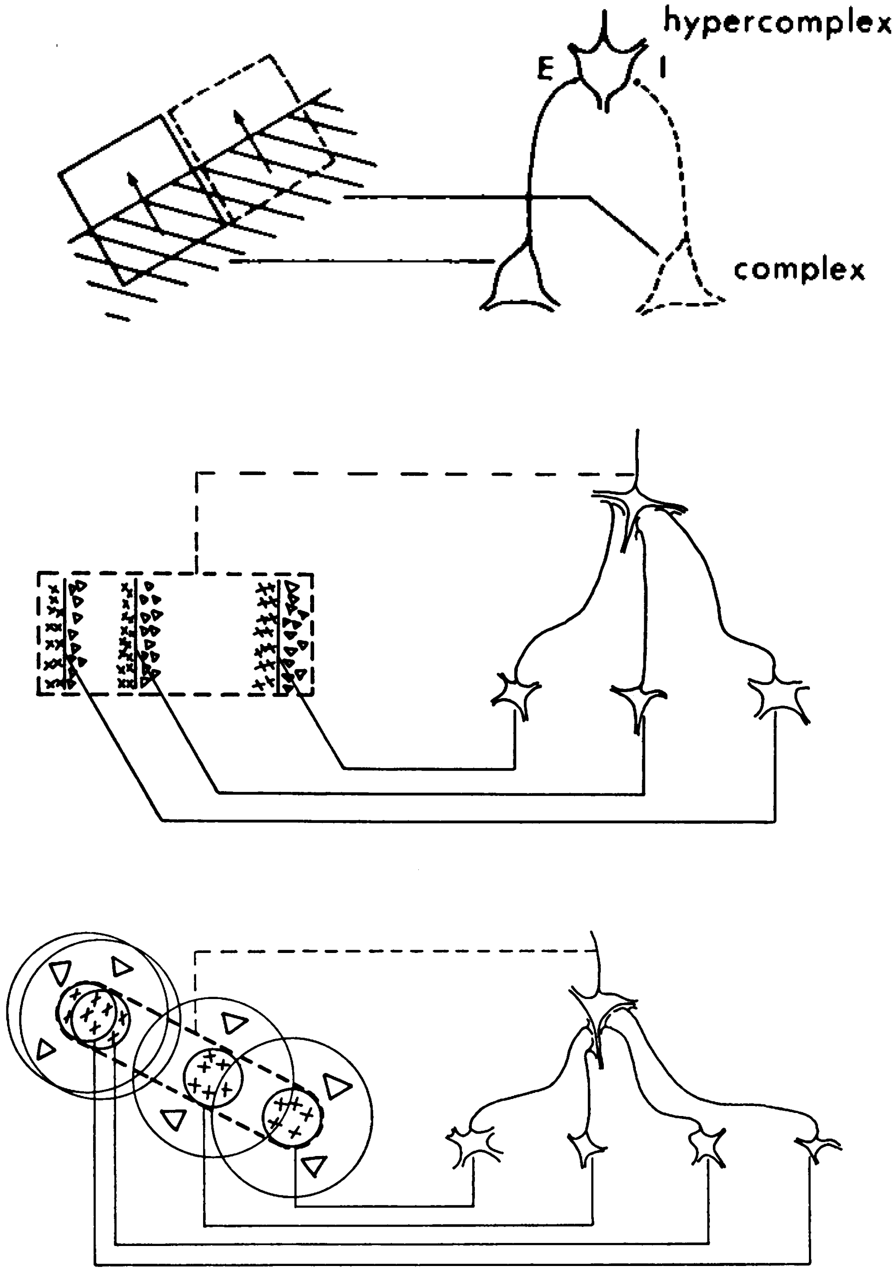

Hypercomplex

Complex

Simple

Neocognitron
(Fukushima 1980)

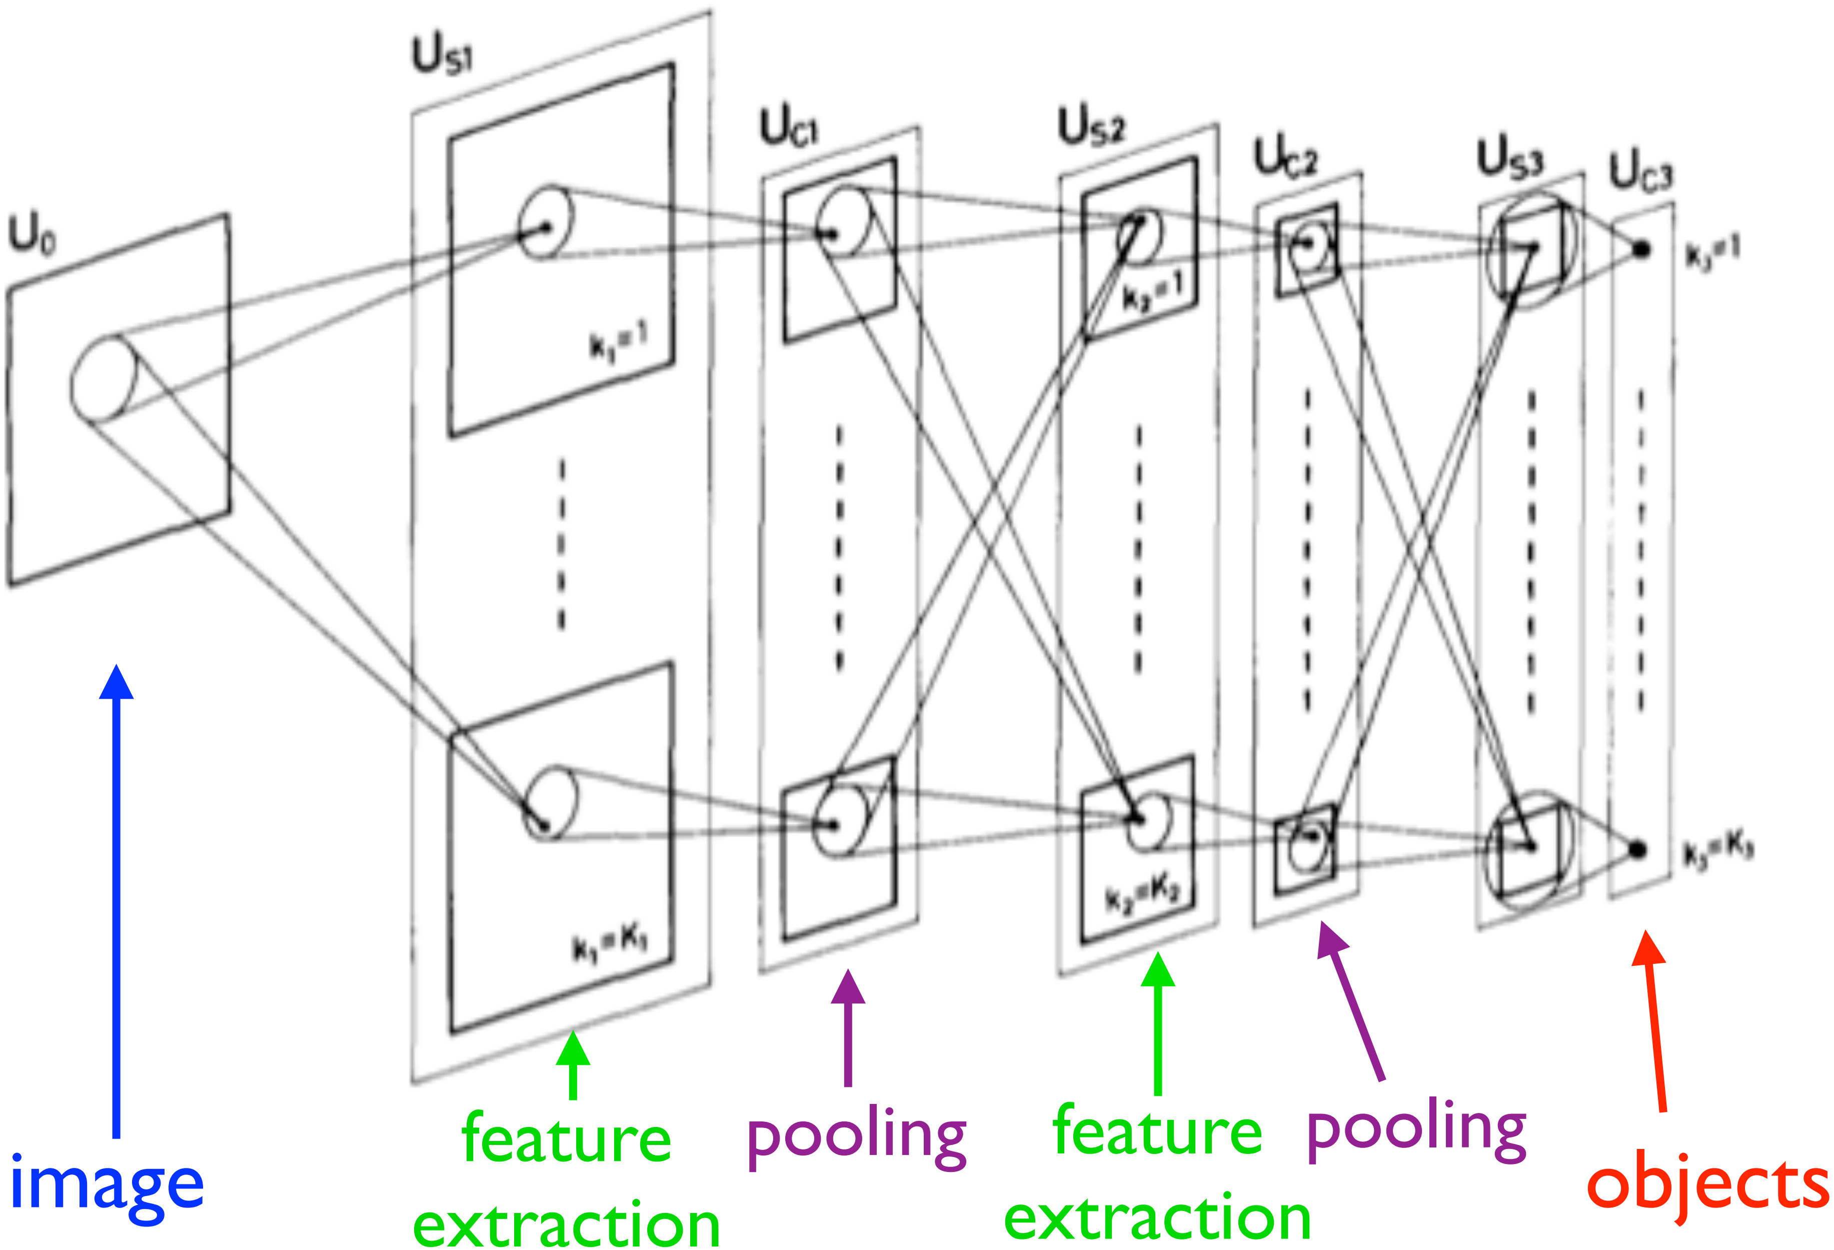

THE WALL STREET JOURNAL.

An AI Breaks the Writing Barrier

A new system called GPT-3 is shocking experts with its ability to use and understand language as well as human beings do

ILLUSTRATION: WREN MCDONALD

By David A. Price
Aug. 22, 2020 12:01 am ET

Word has been making its way out from the technology community: The world changed this summer with the rollout of an artificial intelligence system known as GPT-3. Its ability to interact in English and generate coherent writing have been startling hardened experts, who speak of "GPT-3 shock."

I copied and pasted the first paragraph of George Washington's 1796 Farewell Address: "The period for a new election of a citizen to administer the executive government of the United States being not far distant, and the time actually arrived when your thoughts must be employed in designating the person who is to be clothed with that important trust, it appears to me proper, especially as it may conduce to a more distinct expression of the public voice, that I should now apprise you of the resolution I have formed, to decline being considered among the number of those out of whom a choice is to be made."

GPT-3 gave me its translation: "I am not going to run for president." Take a bow, HAL 9000.

The approach of David Marr (ca. 1980)

Computational theory	Representation and algorithm	Hardware implementation
What is the goal of the computation, why is it appropriate, and what is the logic of the strategy by which it can be carried out?	How can this computational theory be implemented? In particular, what is the representation for the input and output, and what is the algorithm for the transformation?	How can the representation and algorithm be realized physically?

Figure 1–4. The three levels at which any machine carrying out an information-processing task must be understood.

Nervous systems are difficult to observe and manipulate

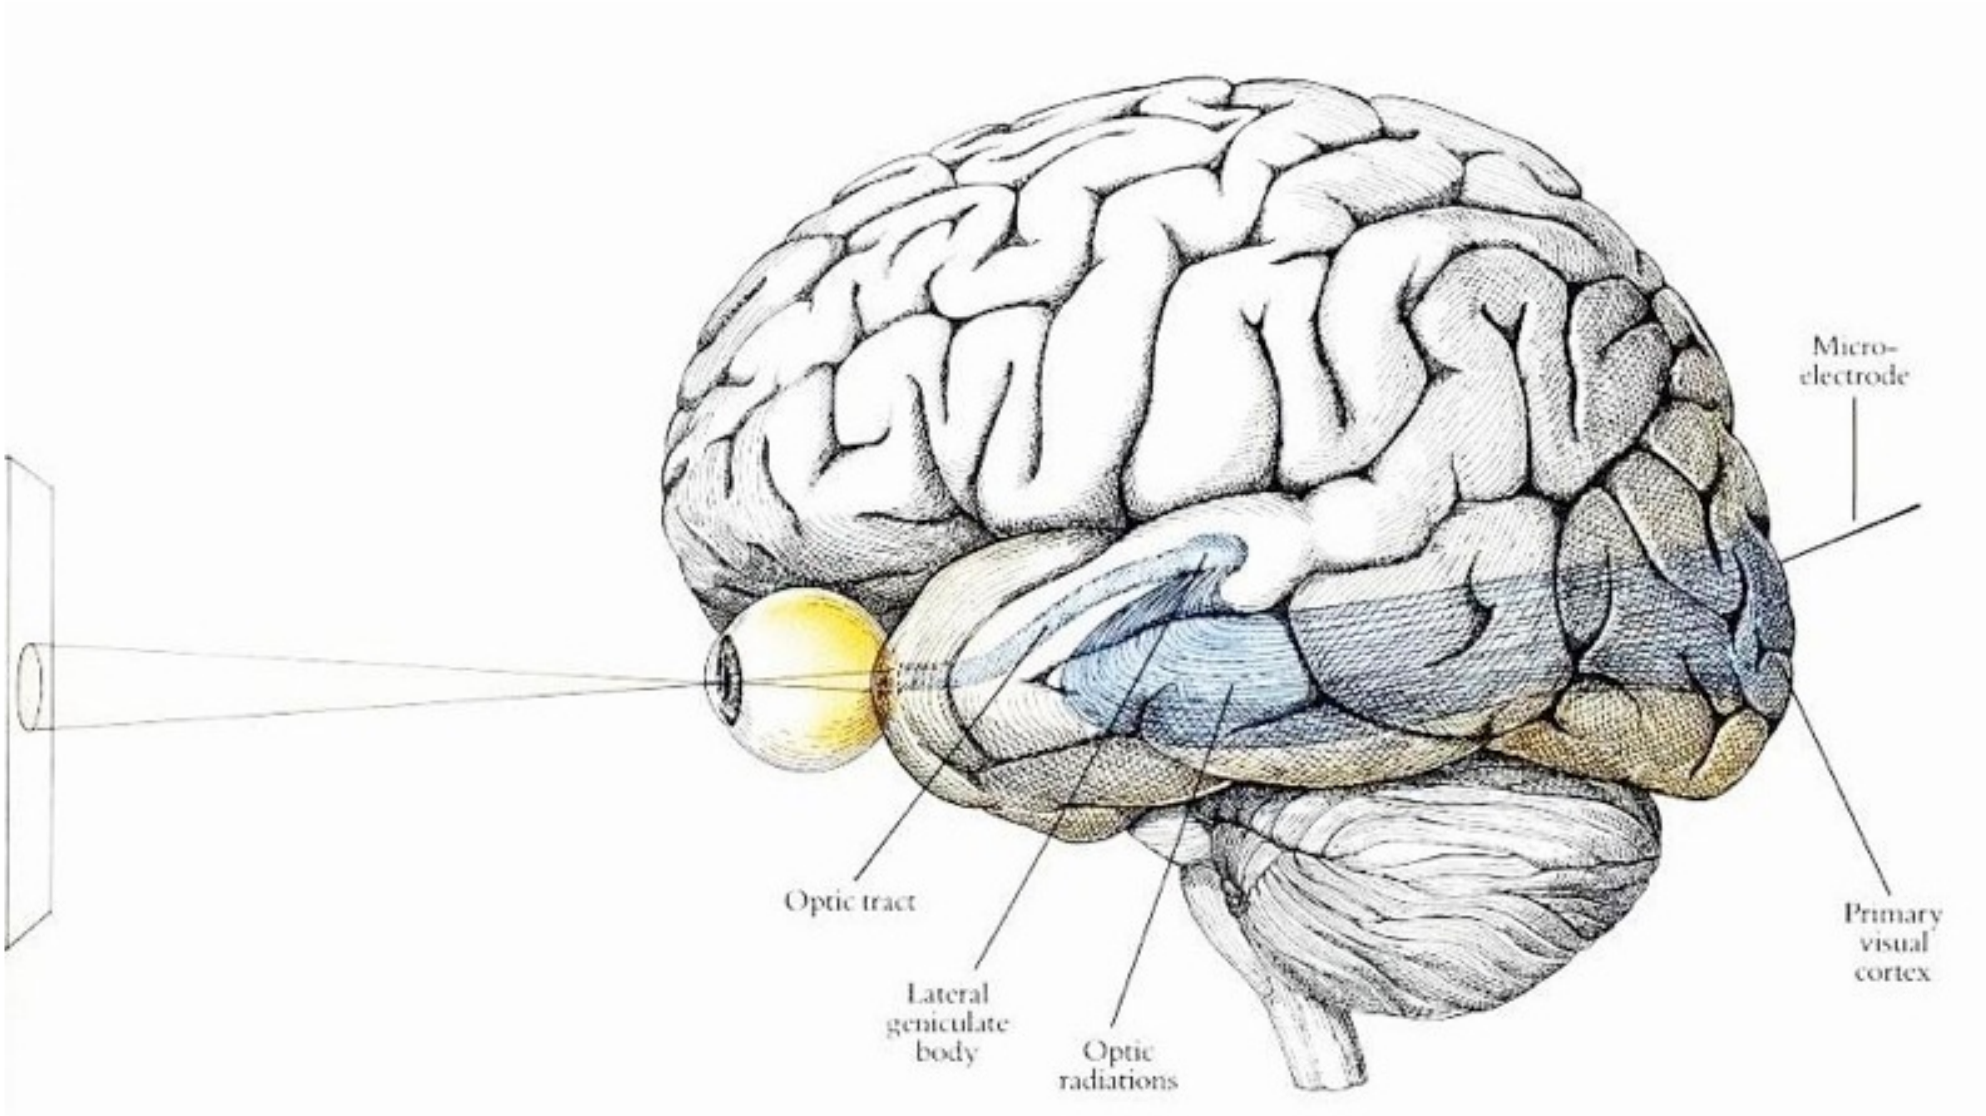

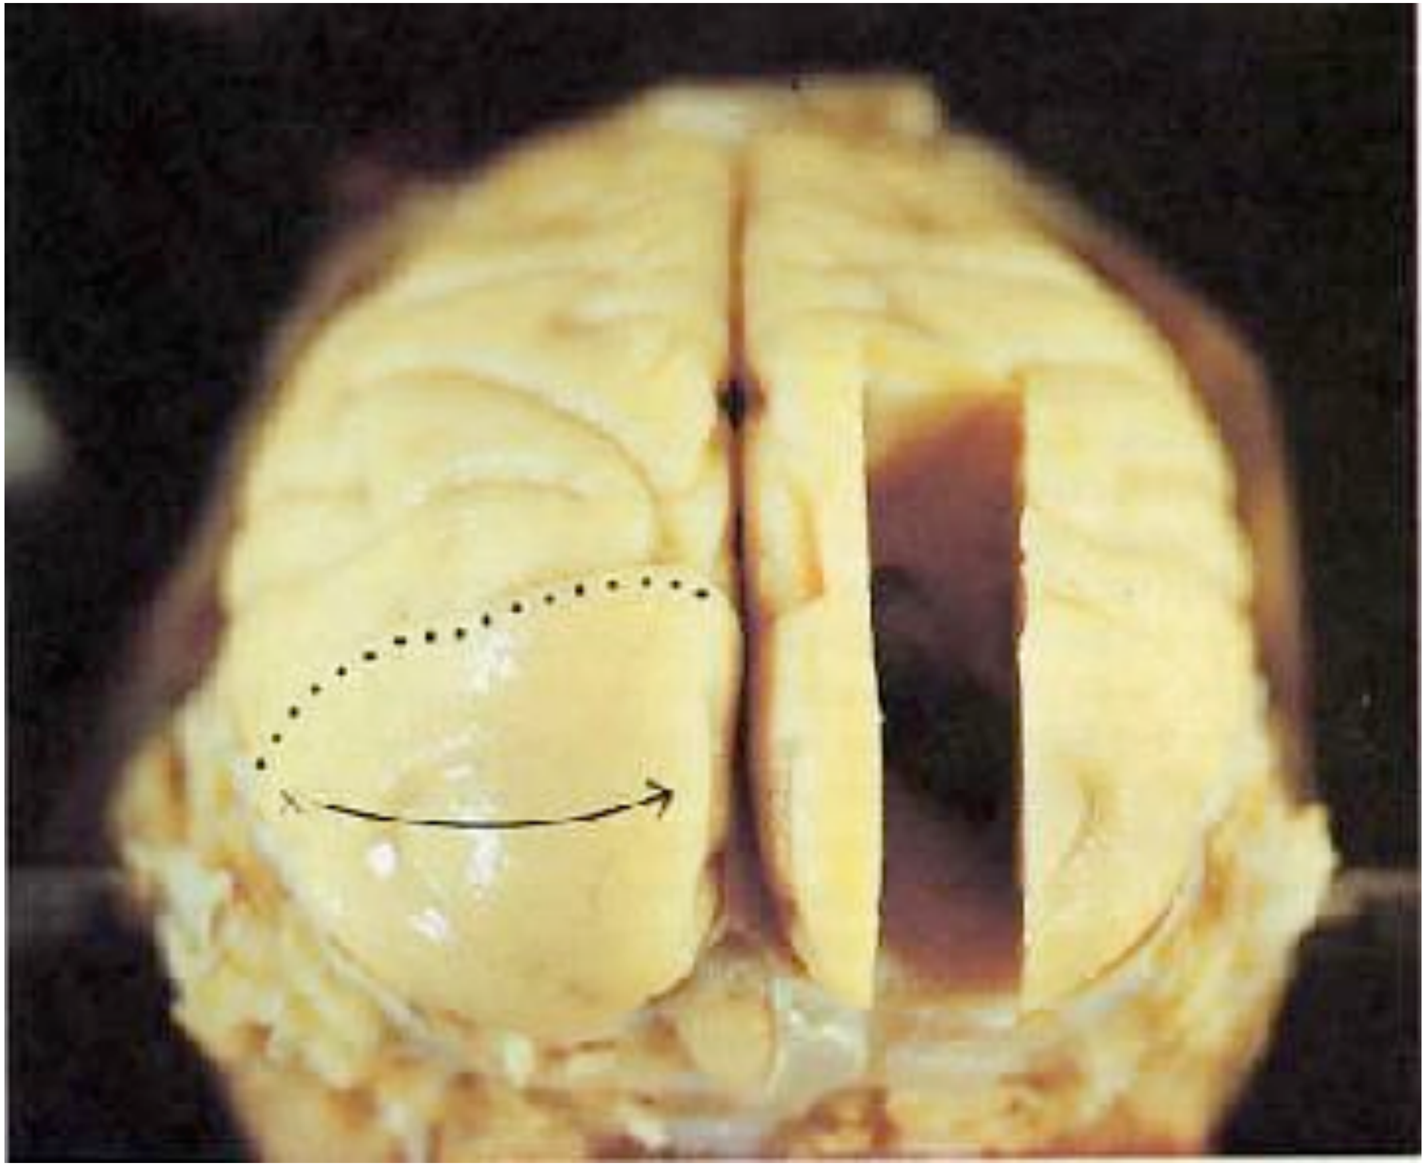

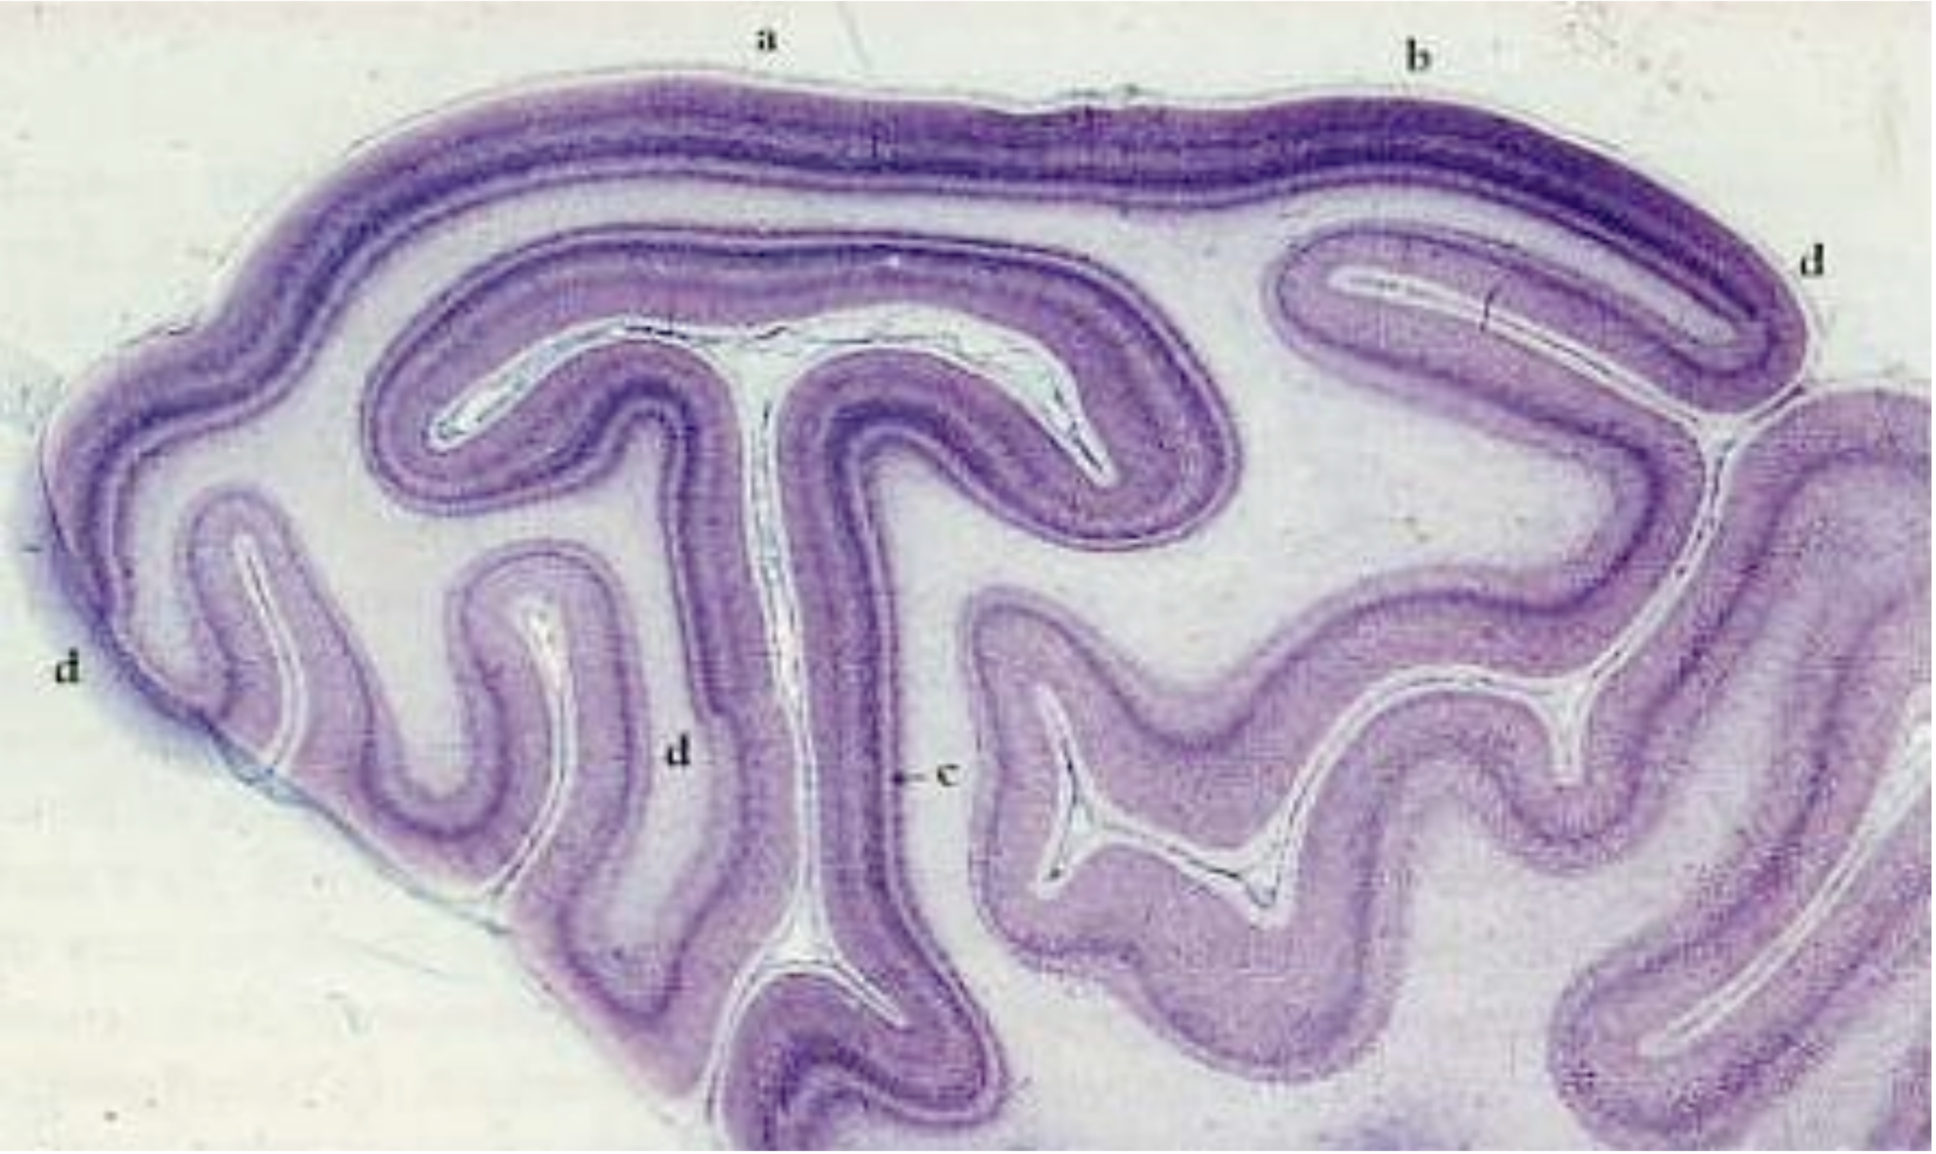
a
b
d
d
d
c

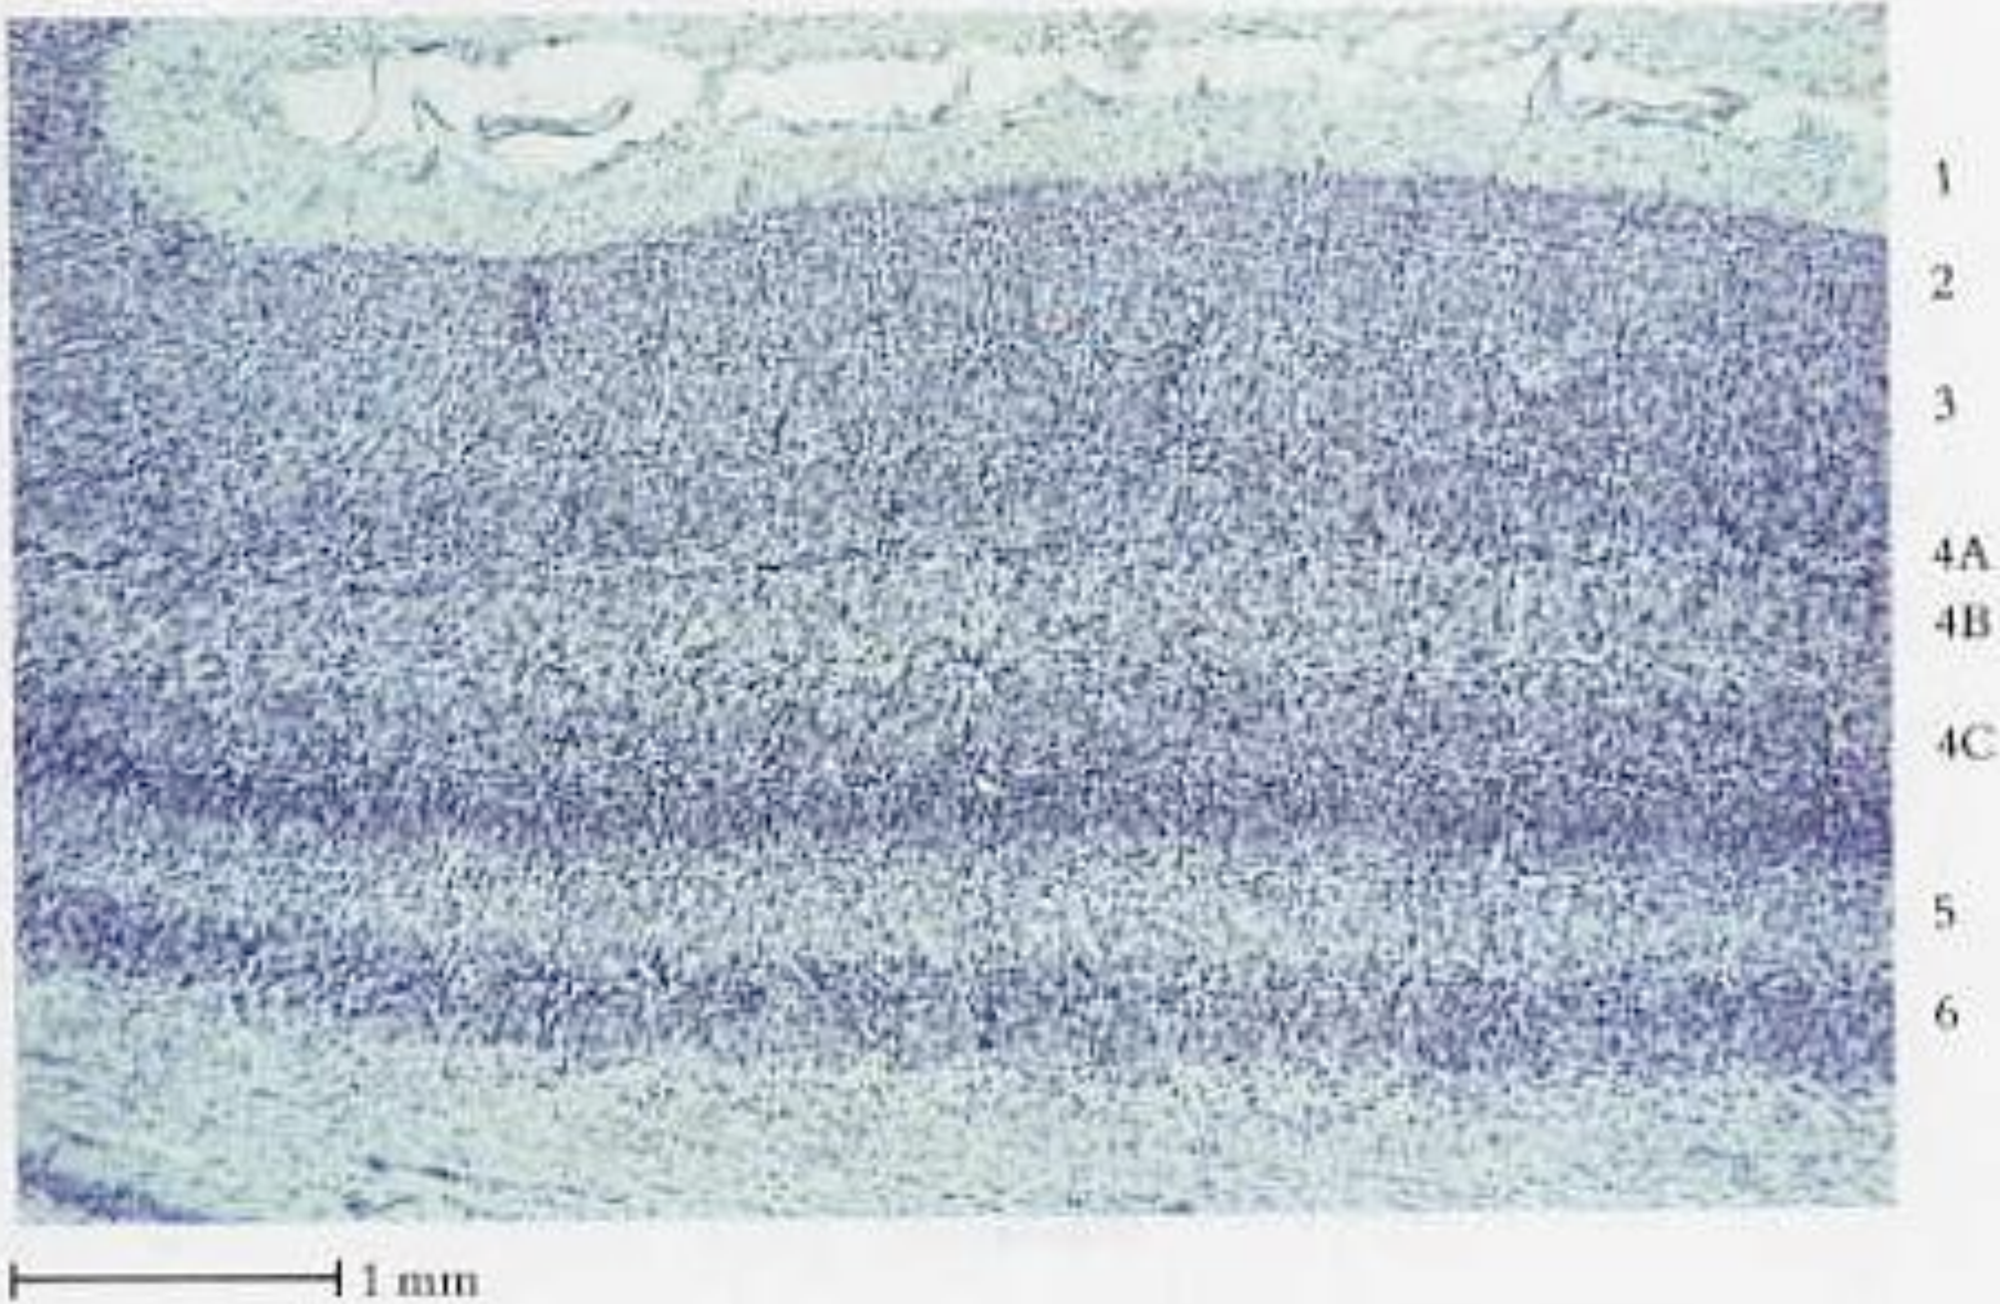
1
2
3
4A
4B
4C
5
6
1 mm

1 mm^2 of cortex contains 100,000 neurons

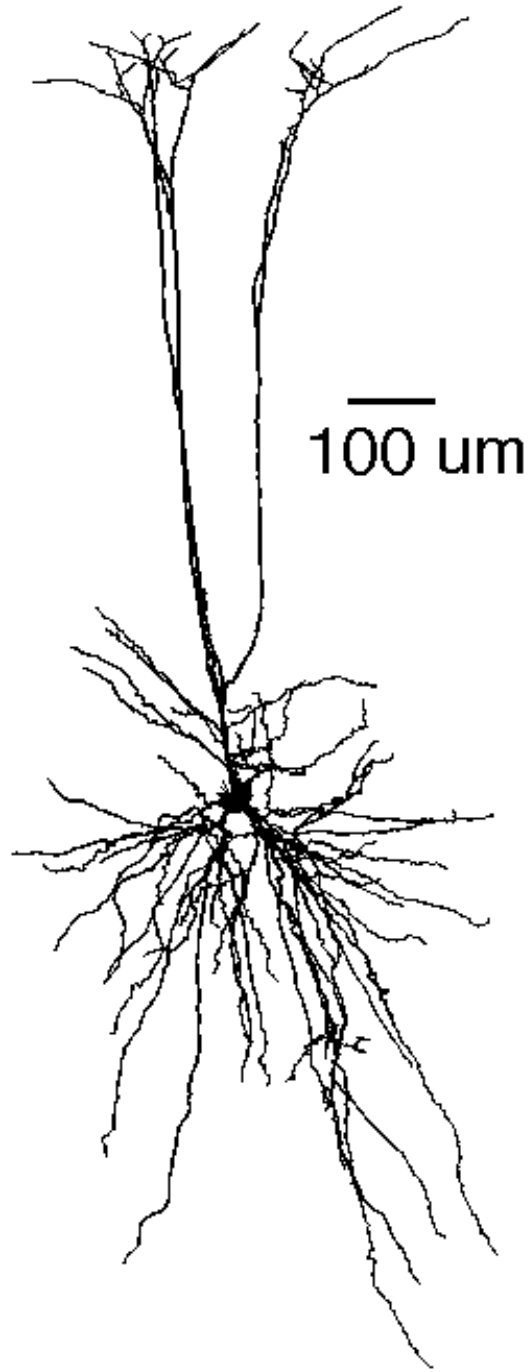
100 um

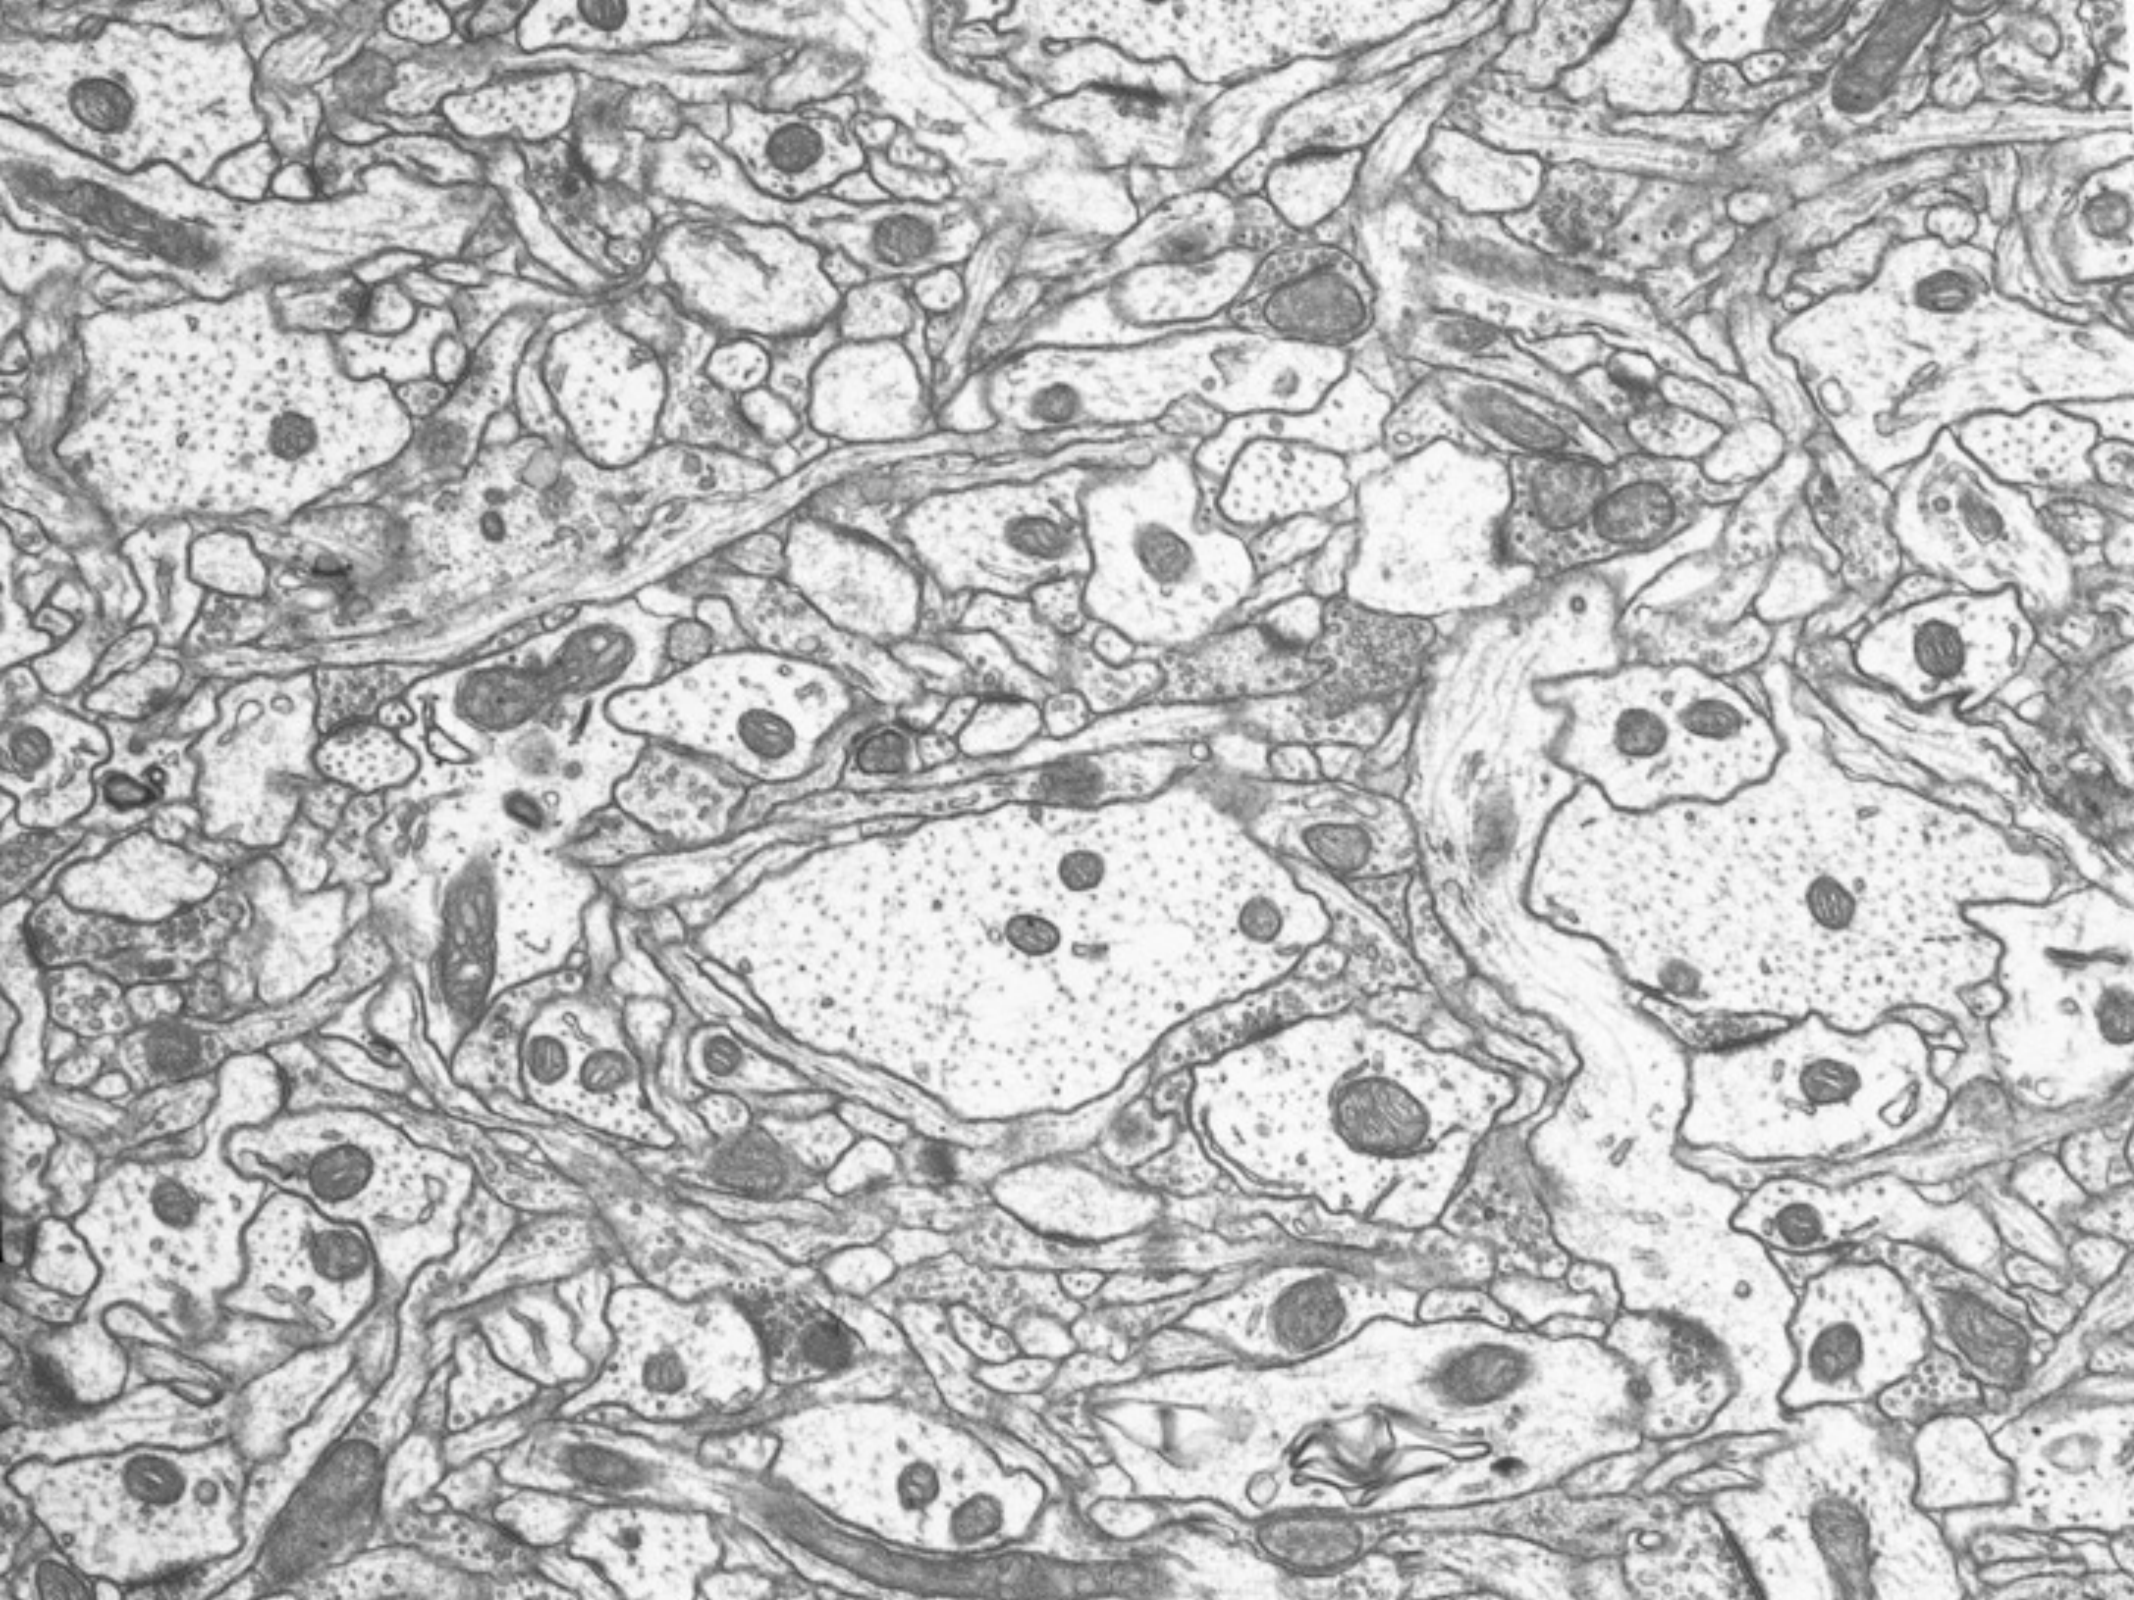

Anatomy of a synapse

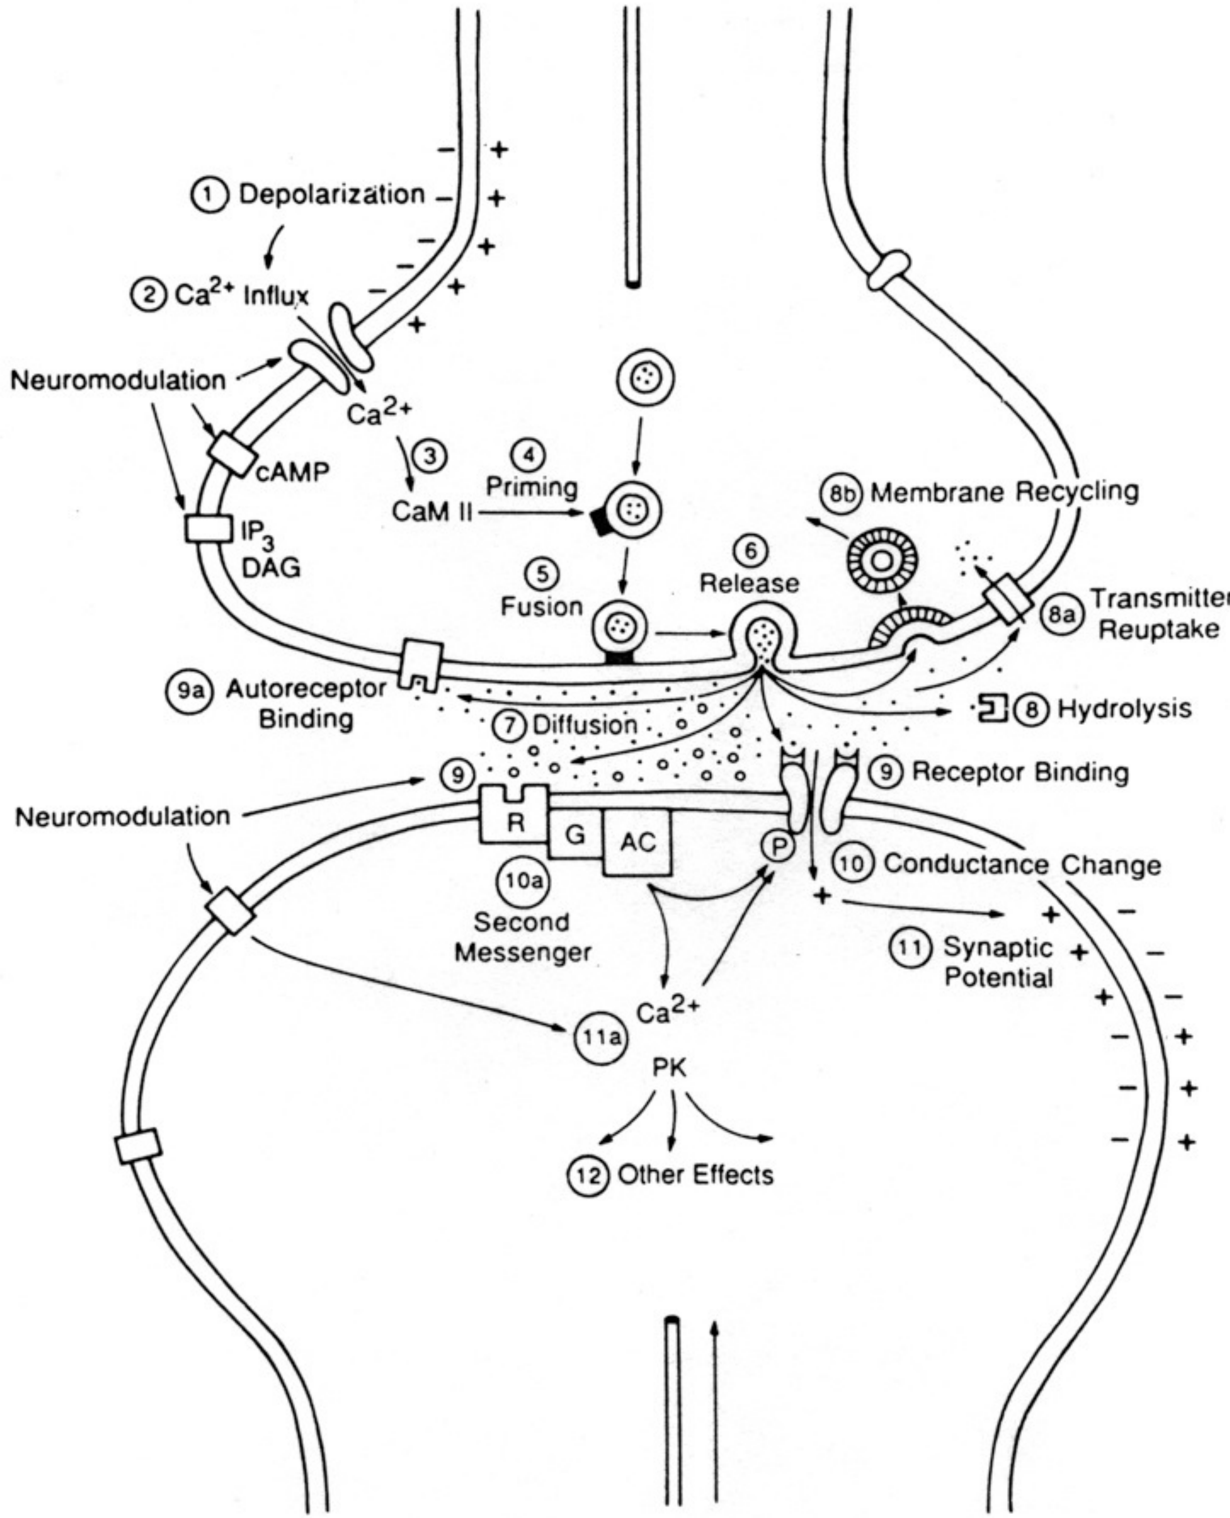

Are there principles?

"God is a hacker"
– Francis Crick

"Individual nerve cells were formerly thought to be unreliable… This was quite wrong, and we now realise their apparently erratic behavior was caused by our ignorance, not the neuron's incompetence."
– H.B. Barlow (1972)

Otto Lilienthal experiments with flight (1890's)

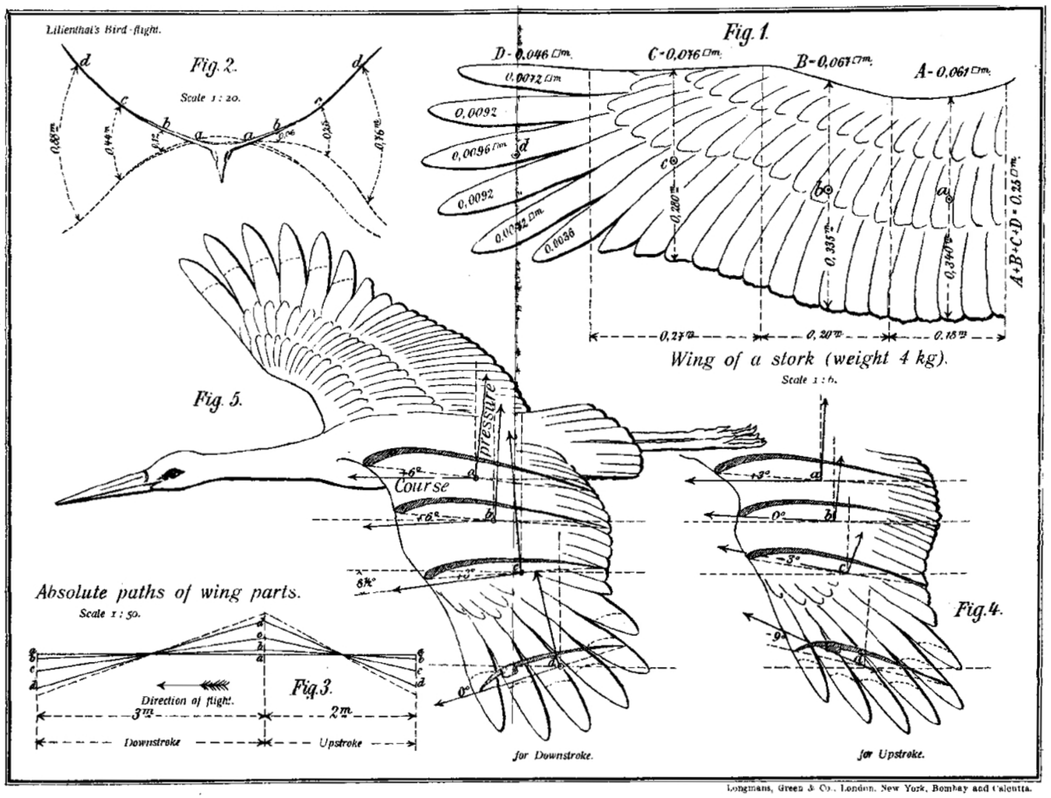

Der Vogelflug als Grundlage der Fliegekunst (1889)

Wright Flyer (1903)

ORVILLE WRIGHT
DAYTON, OHIO

December 27, 1941.

Mr. Horace Lytle, President,
The J. Horace Lytle Company,
Dayton, Ohio.

Dear Mr. Lytle:-

Your letter of November 26th was duly received, but having become buried among other papers, it has just come to my attention again.

I can not think of any part bird flight had in the development of human flight excepting as an inspiration. Although we intently watched birds fly in a hope of learning something from them I can not think of anything that was first learned in that way. After we had thought out certain principles, we then watched the bird to see whether it used the same principles. In a few cases we did detect the same thing in the bird's flight.

Learning the secret of flight from a bird was a good deal like learning the secret of magic from a magician. After you once know the trick and know what to look for you see things that you did not notice when you did not know exactly what to look for.

Sincerely yours,

Orville Wright

Principles of optics govern the design of eyes

THE EVOLUTION OF EYES

Michael F. Land *Russell D. Fernald*

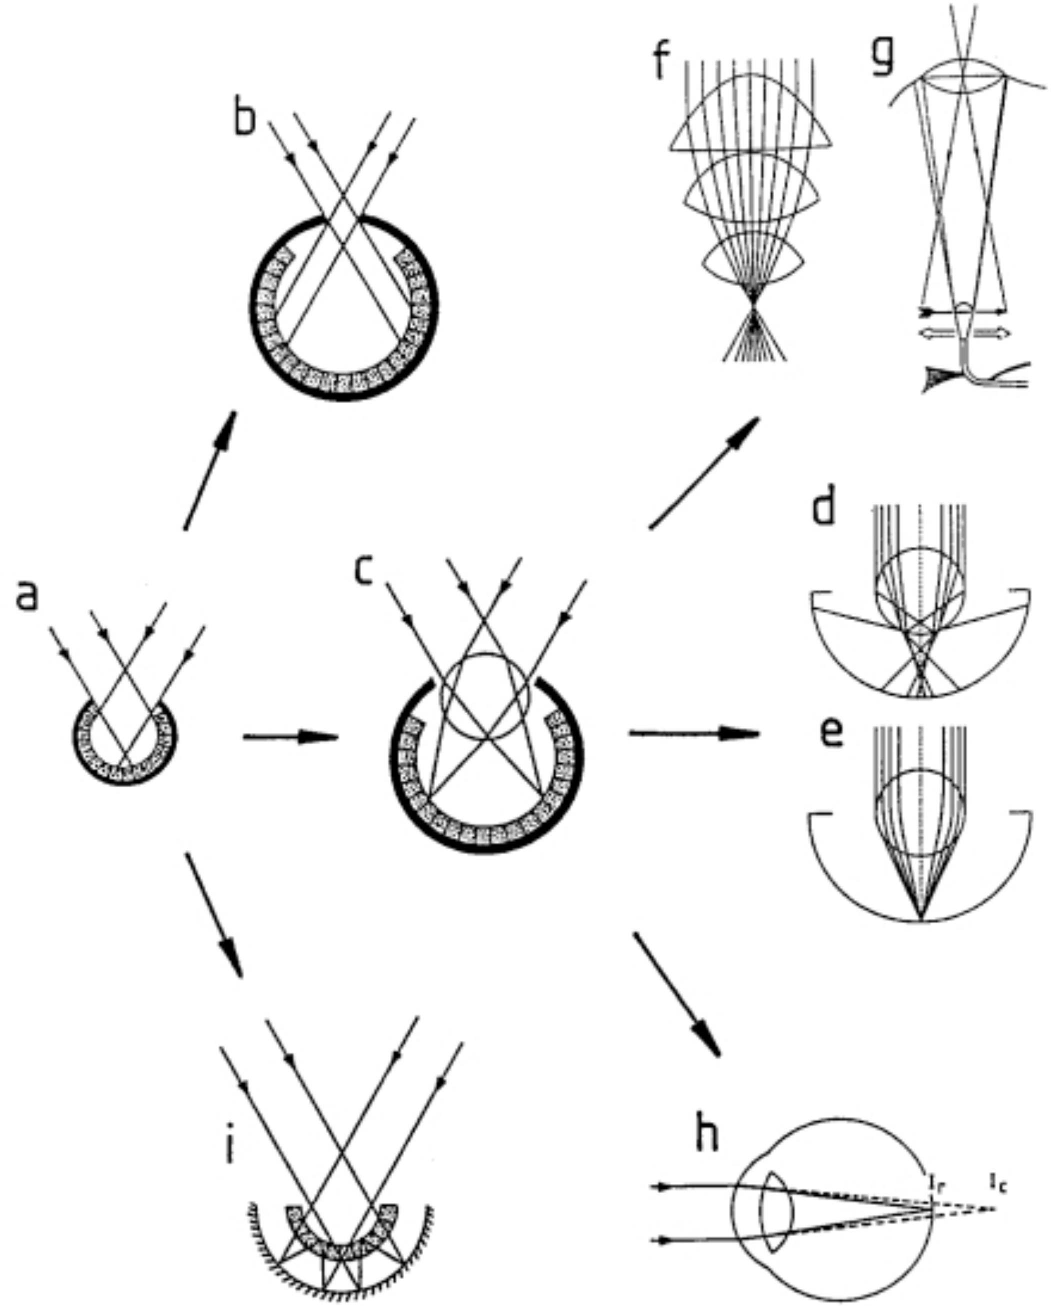

Zhang & Sejnowski (2000)

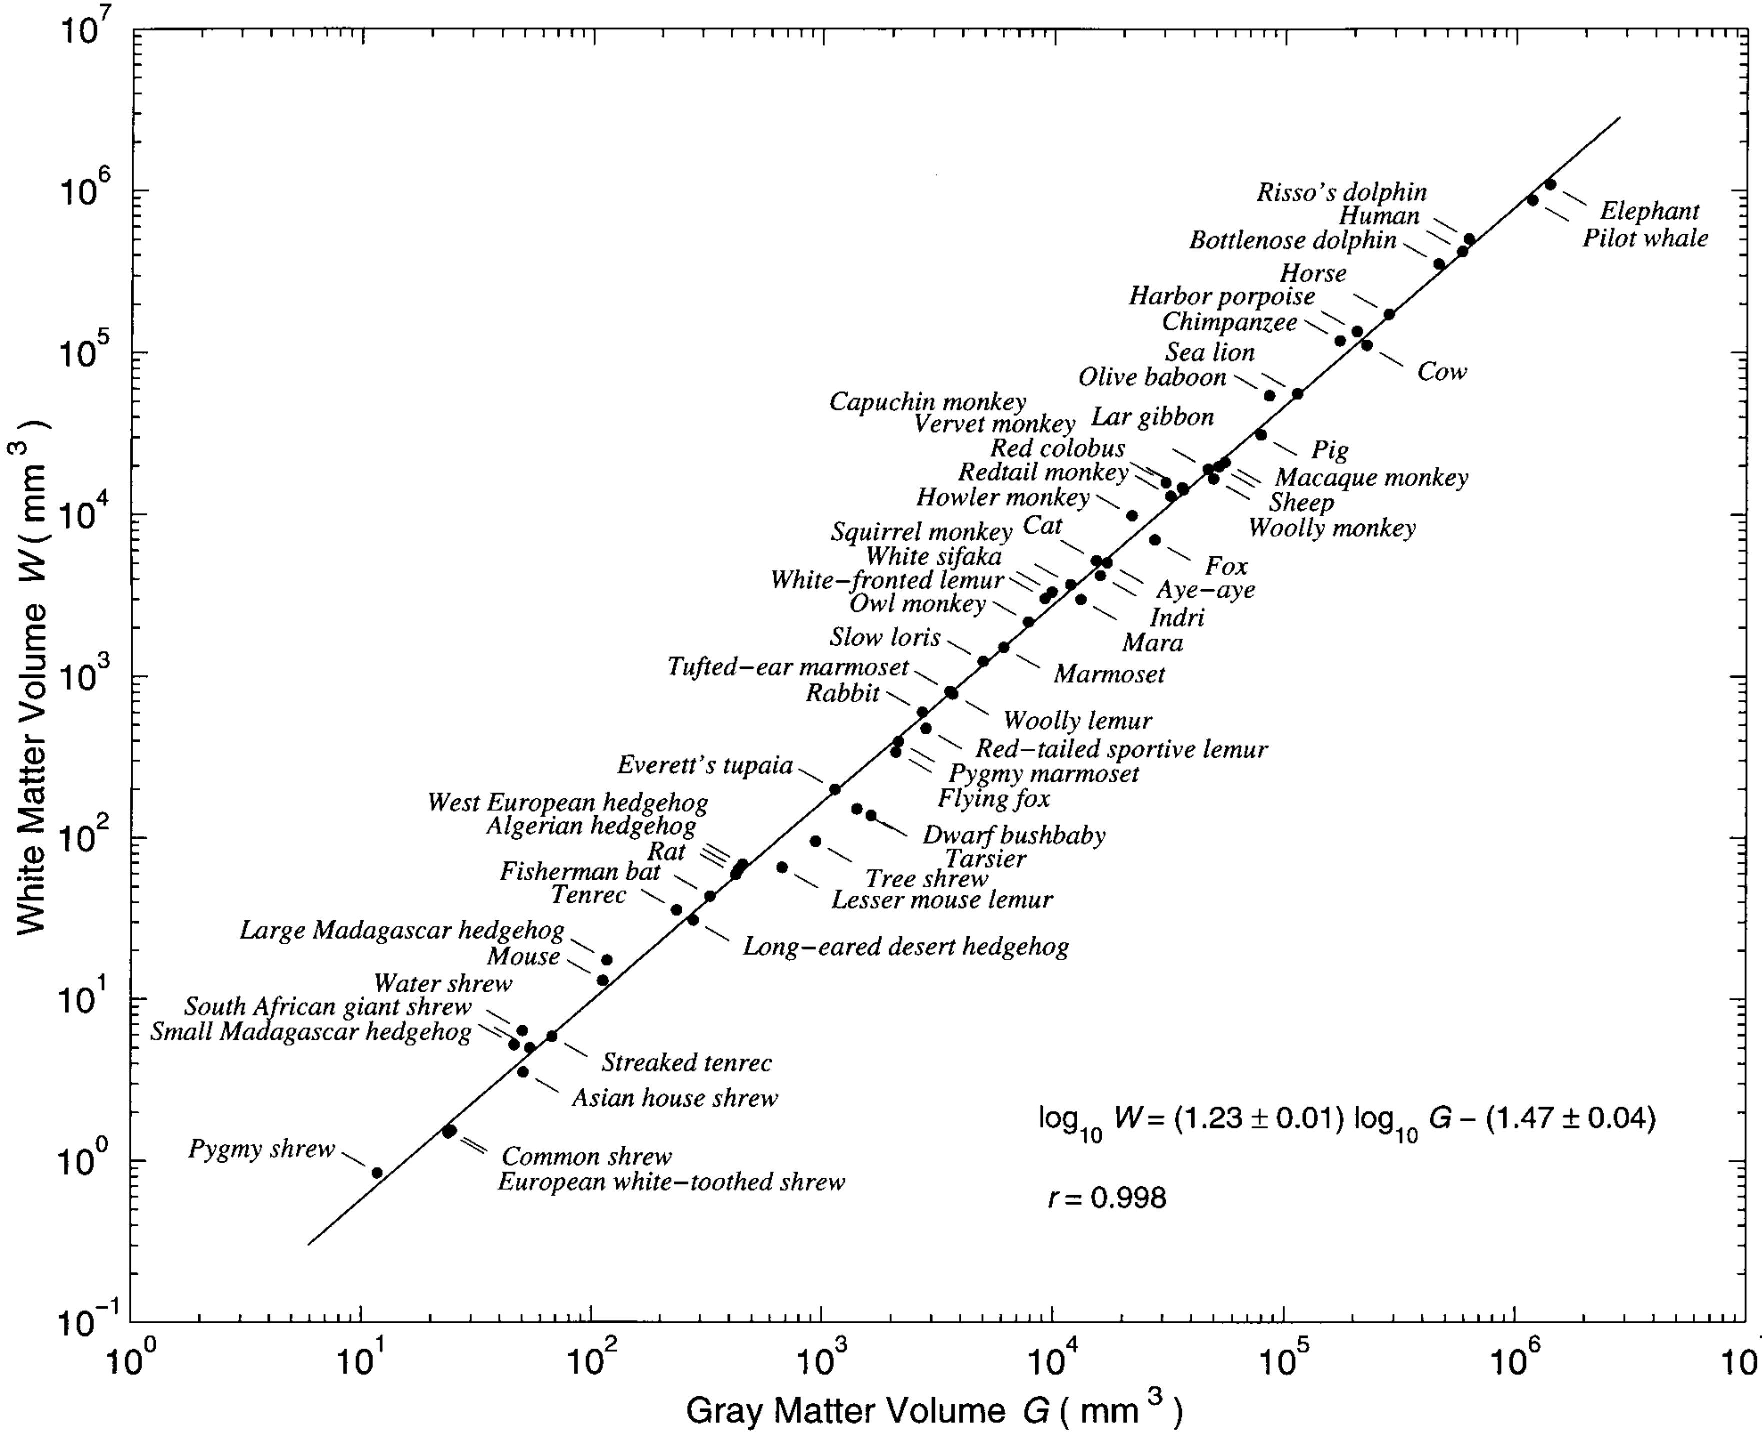

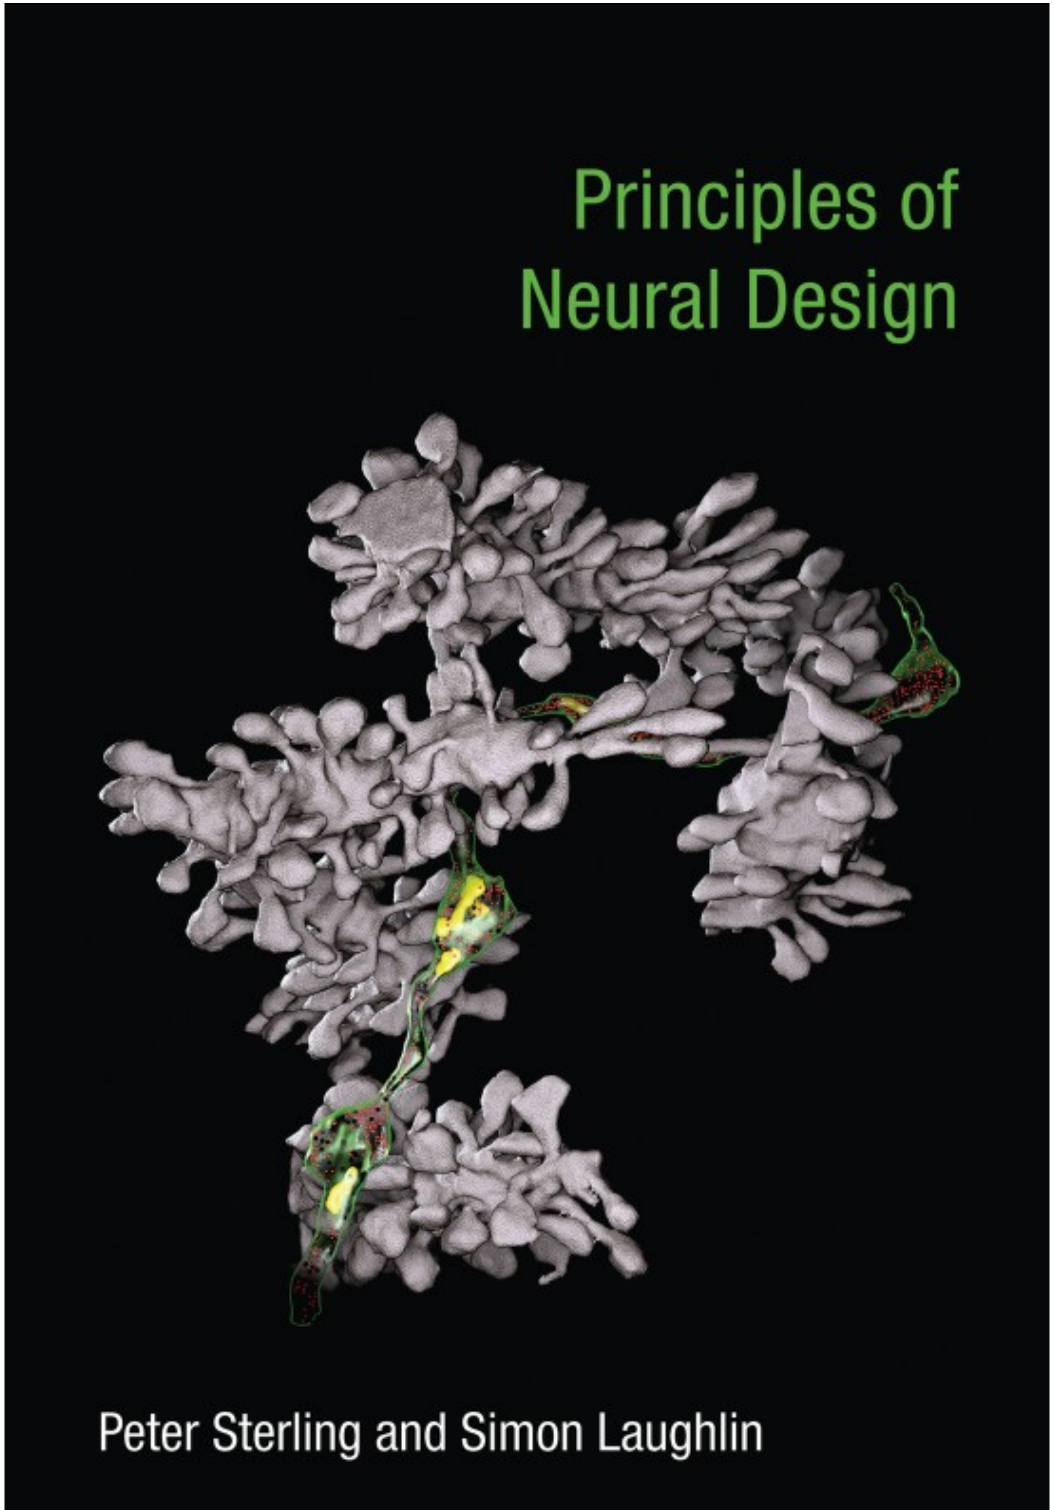

Principles

Compute with chemistry

Compute directly with analog primitives

Combine analog and pulsatile processing

Sparsify

Send only what is needed

Send at the lowest acceptable rate

Minimize wire

Make neural components irreducibly small

Complicate

Adapt, match, learn, and forget

Computational principles

- Efficient coding
- Unsupervised learning
- Bayesian inference
- Dynamical systems
- Prediction
- High-dimensional vector arithmetic
- Computing with waves

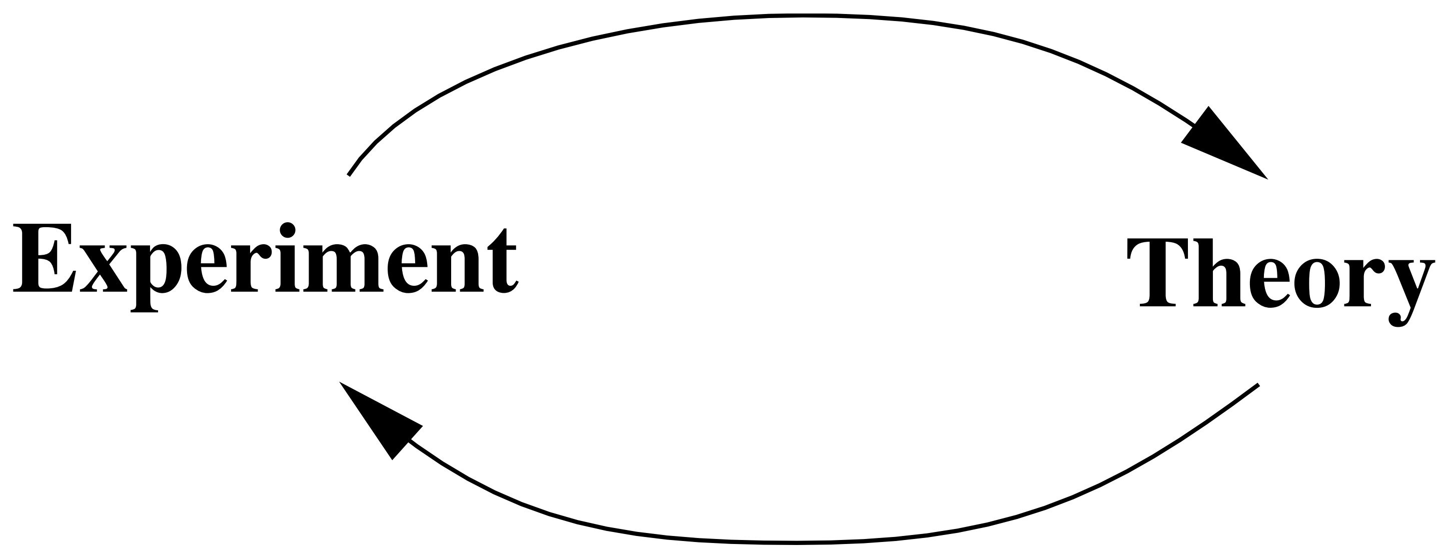
Experiment
Theory

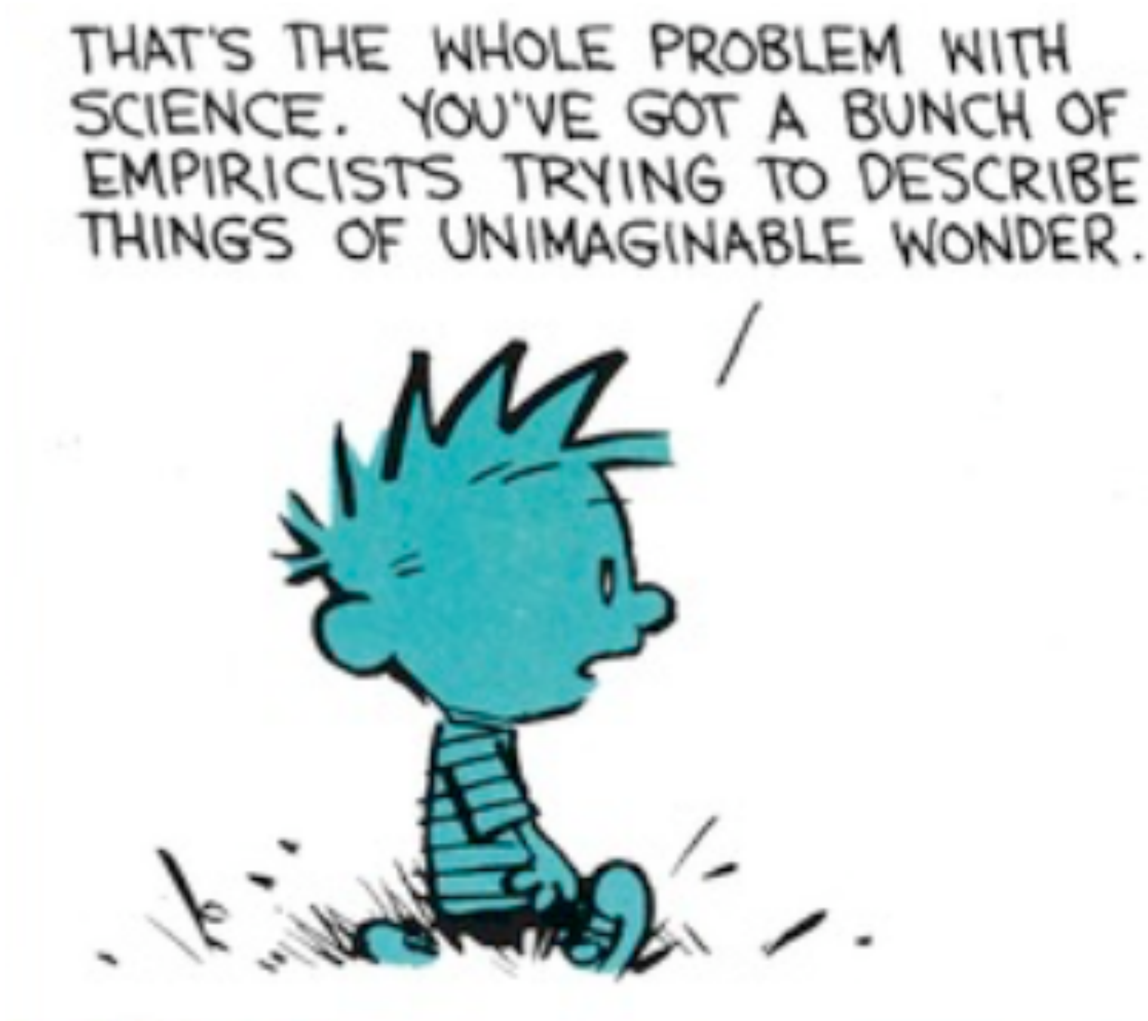
THAT'S THE WHOLE PROBLEM WITH SCIENCE. YOU'VE GOT A BUNCH OF EMPIRICISTS TRYING TO DESCRIBE THINGS OF UNIMAGINABLE WONDER.