VS 265 - Neural Computation

Bruno A. Olshausen, Instructor
570 Evans
baolshausen@berkeley.edu

Shariq Mobin, GSI
577 Evans
shariq.mobin@gmail.com

Class meets TTH 3:30-5 in 560 Evans

Weekly lab 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*

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:

- 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.

For Tuesday:

- Mead, C. Chapter 1: Introduction and Chapter 4: Neurons from *Analog VLSI and Neural Systems*, Addison-Wesley, 1989.
- Jordan, M.I. An Introduction to Linear Algebra in Parallel Distributed Processing in McClelland and Rumelhart, *Parallel Distributed Processing*, MIT Press, 1985.
- Linear neuron models (handout)
- Linear time-invariant systems and convolution (handout)
- Simulating differential equations (handout)
- Carandini M, Heeger D (1994) Summation and division by neurons in primate visual cortex. *Science, 264*: 1333-1336.

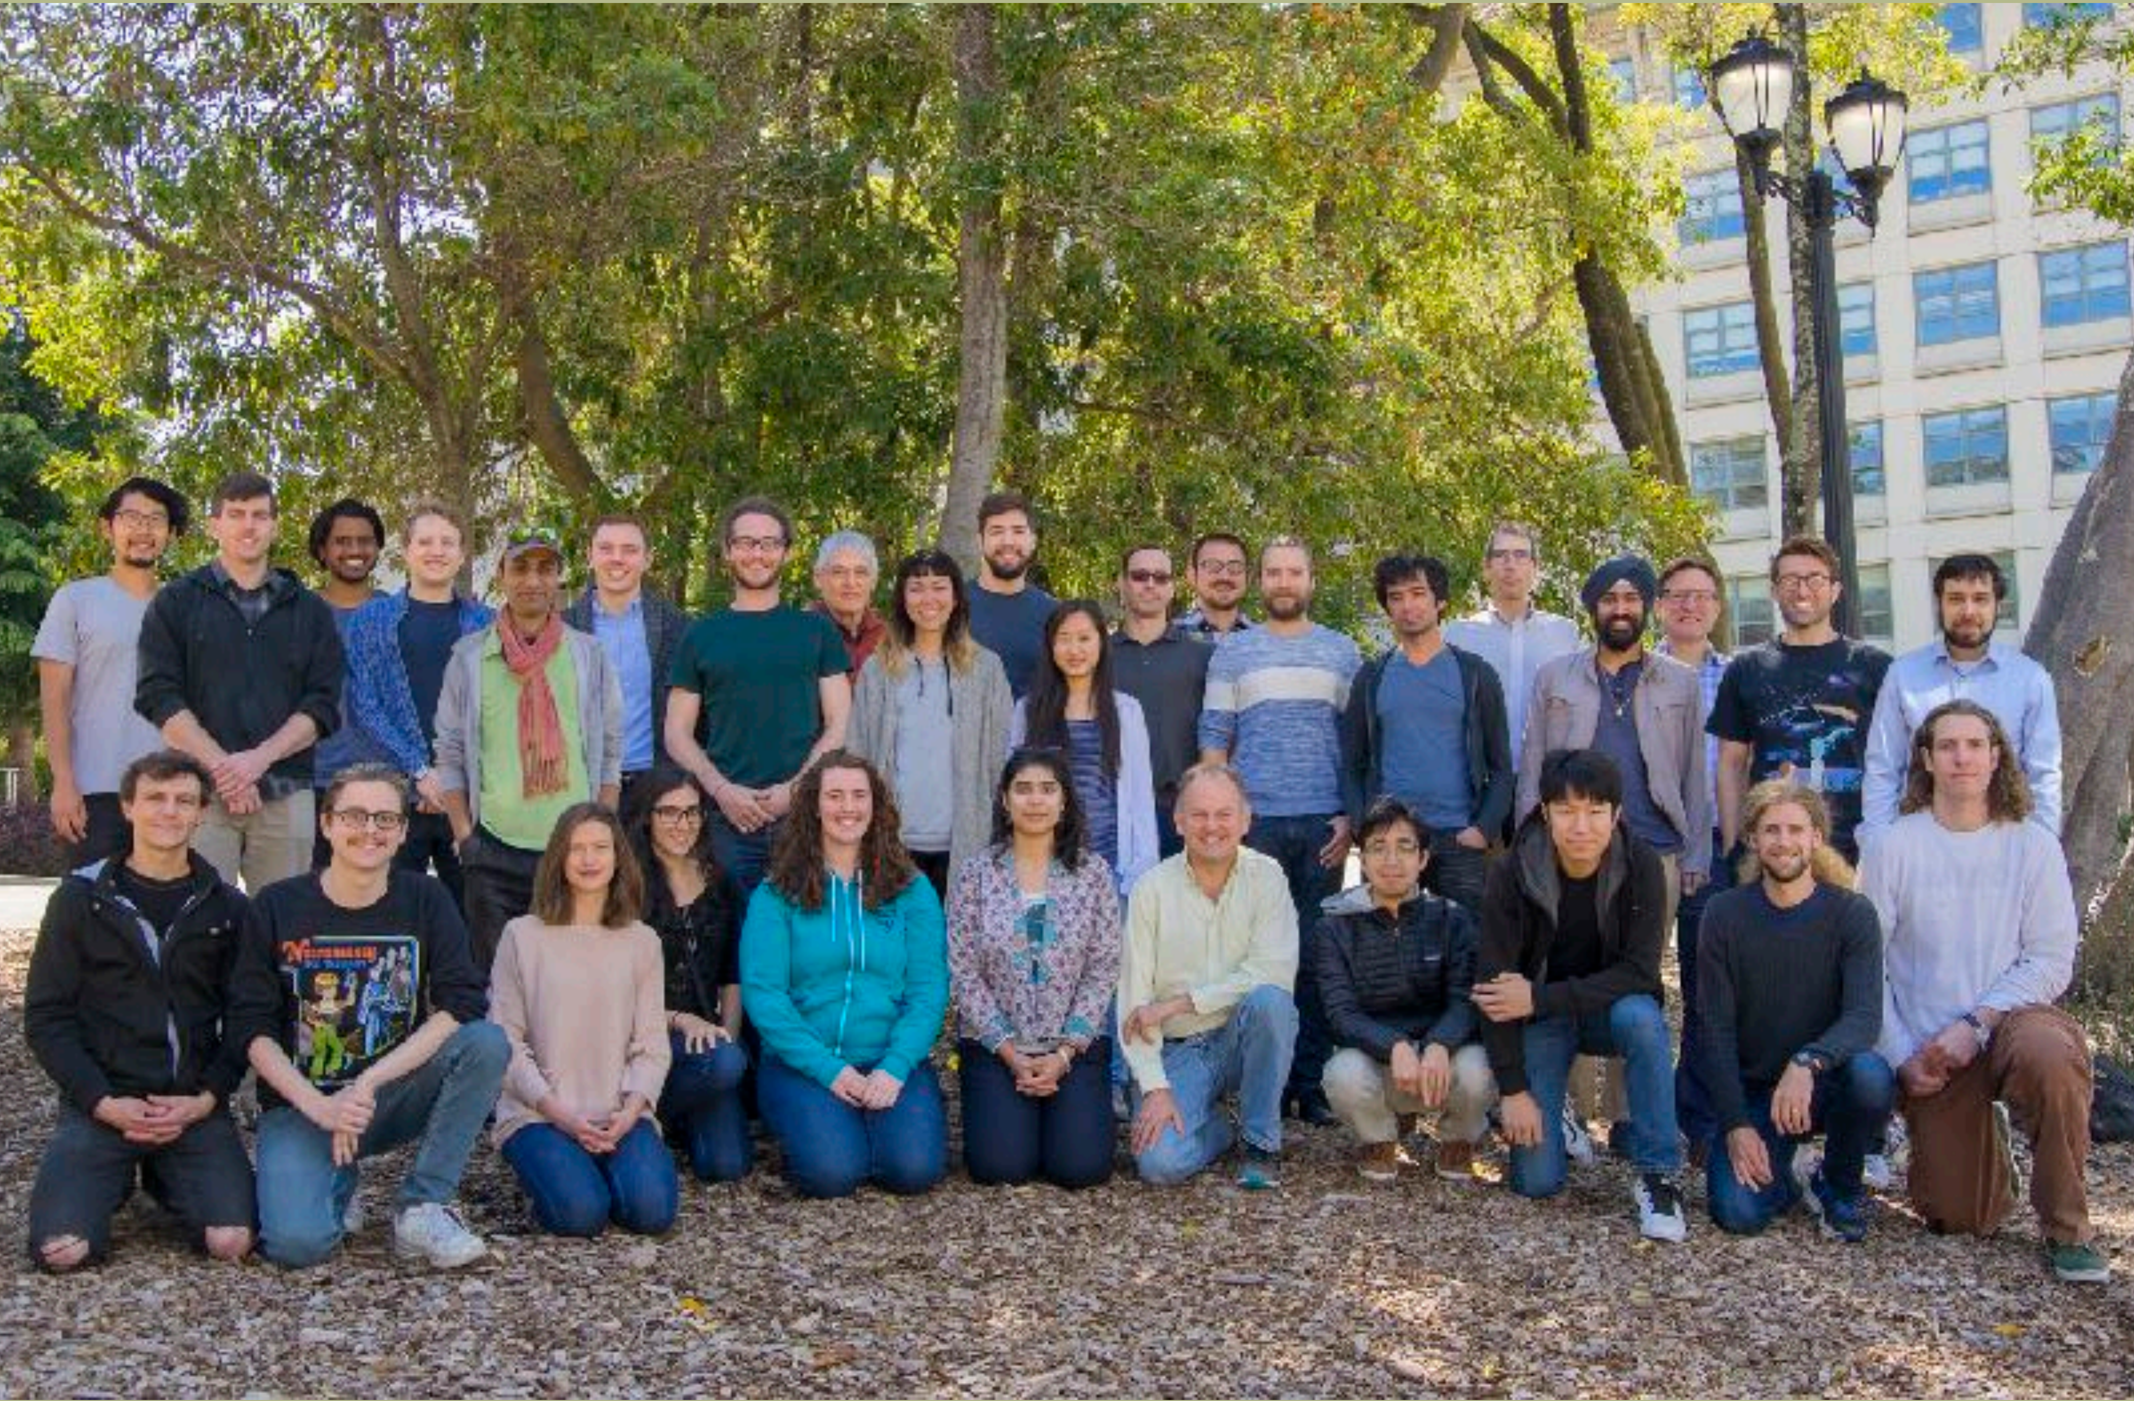

Redwood Center for Theoretical Neuroscience - April 2018

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

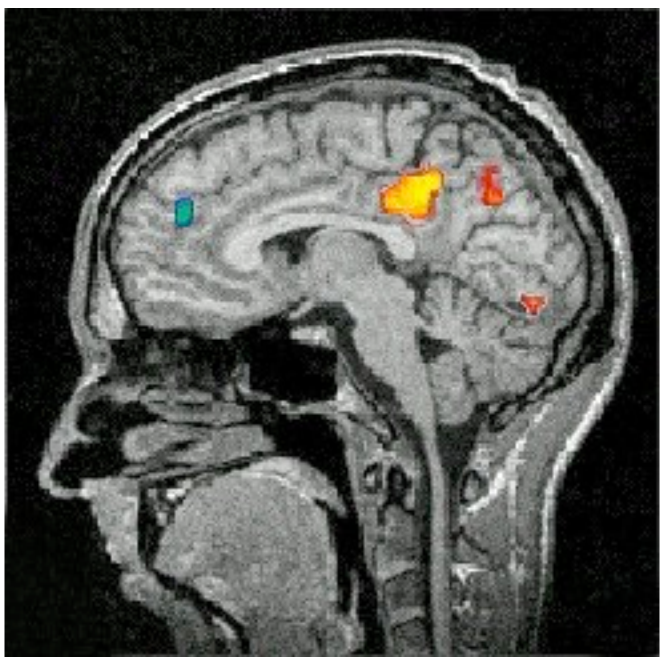

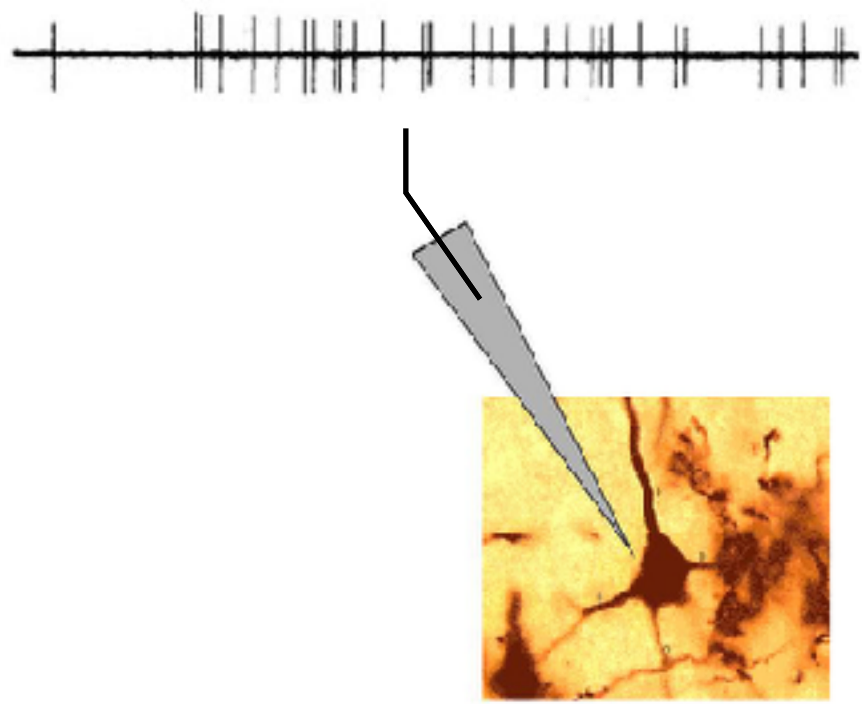

Why hasn’t machine intelligence scaled with Moore’s law?

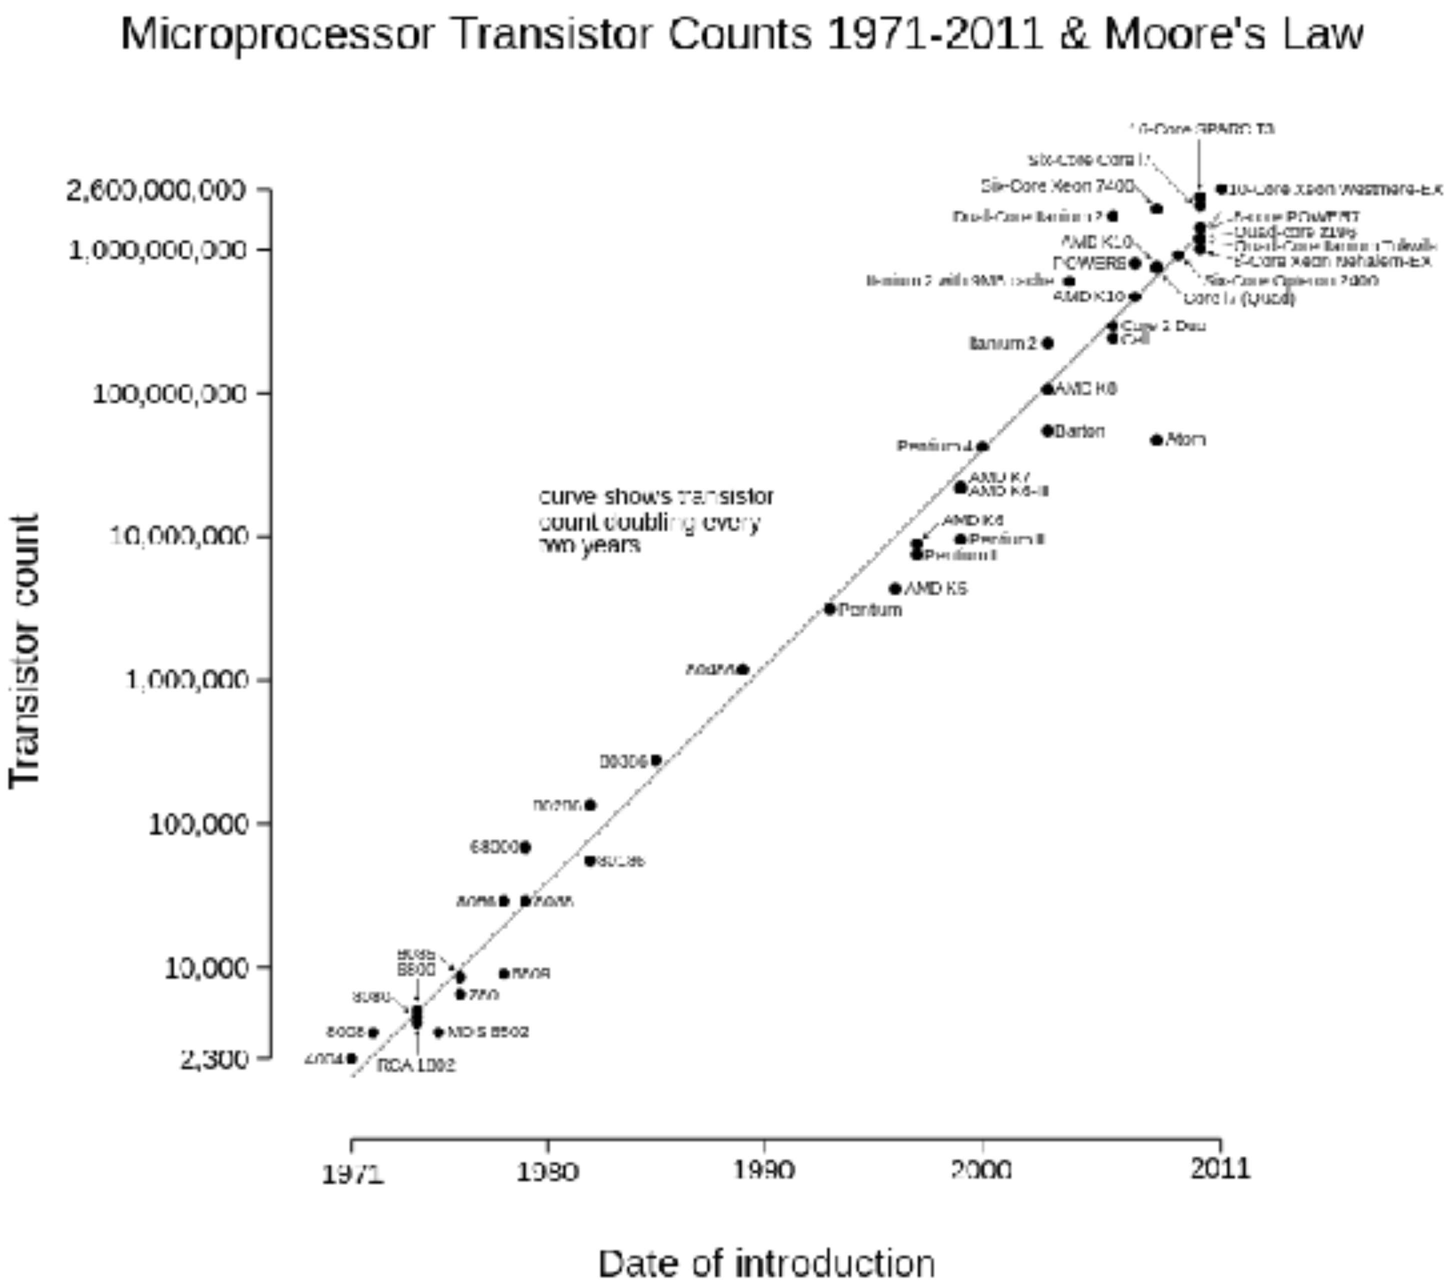

After **50 years** of concerted research efforts...

- 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.*

Why’s it so hard?

Artificial Intelligence



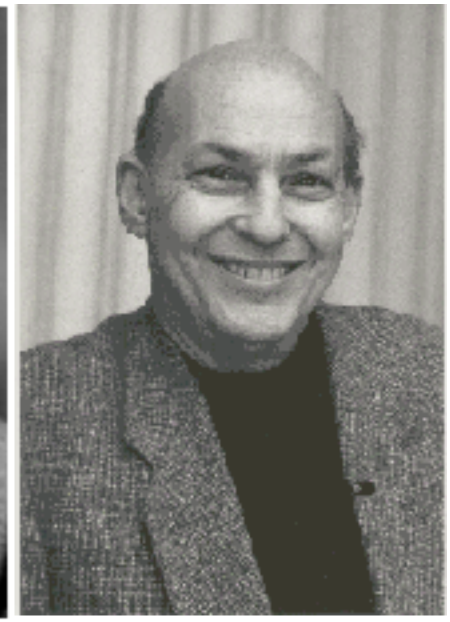


Alan Turing John von Neumann Marvin Minsky John McCarthy

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

Subgoal for July

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.

Extensions for August

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.

The Lighthill debate (1973)

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

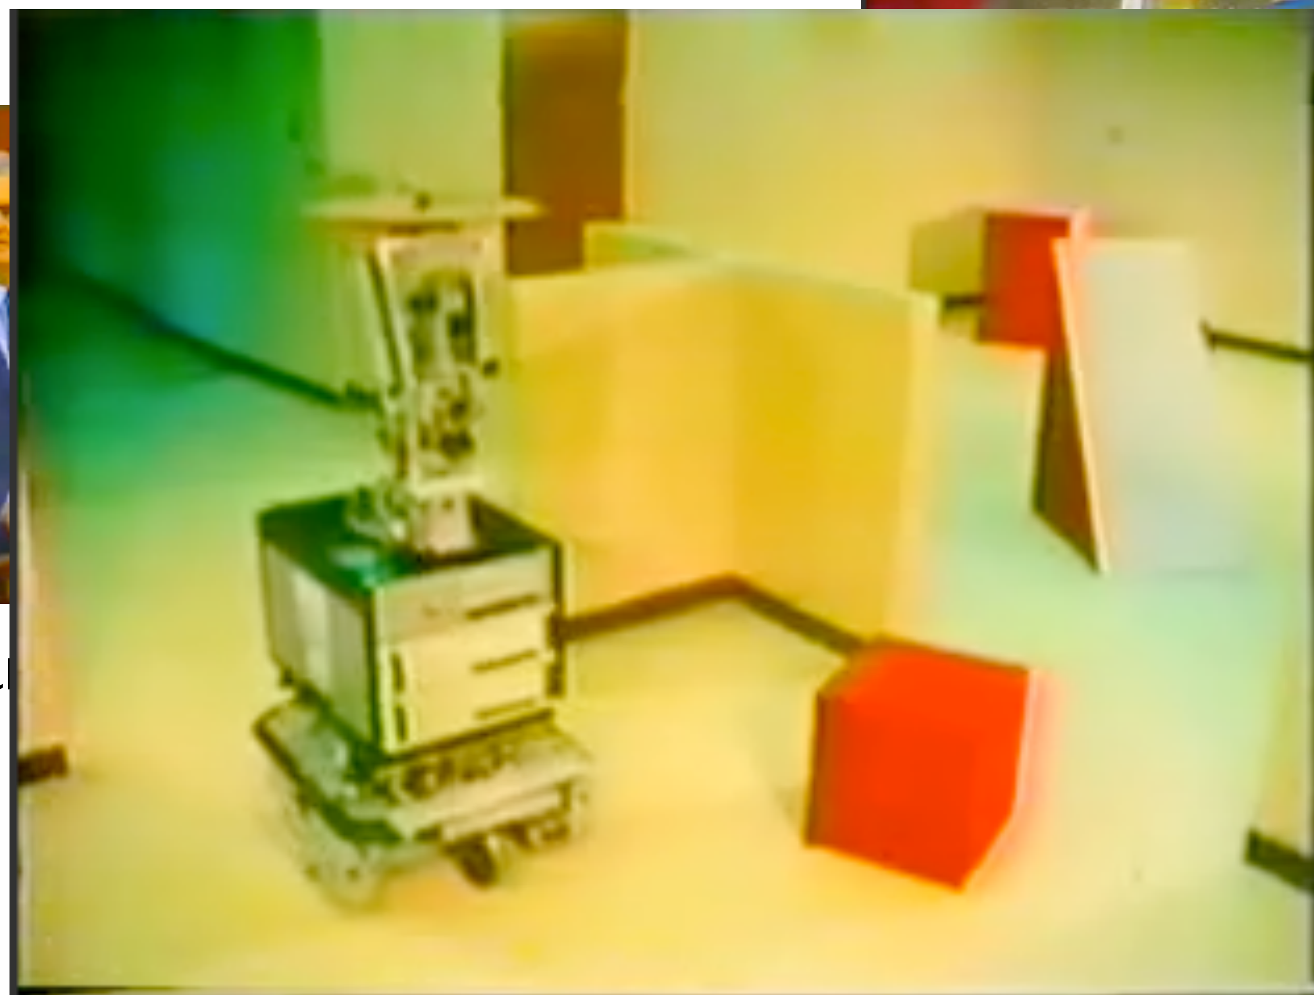

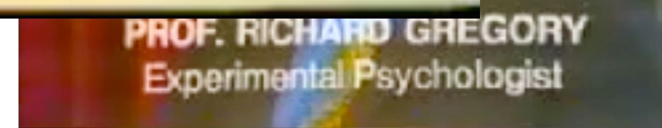

Sir Ja

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

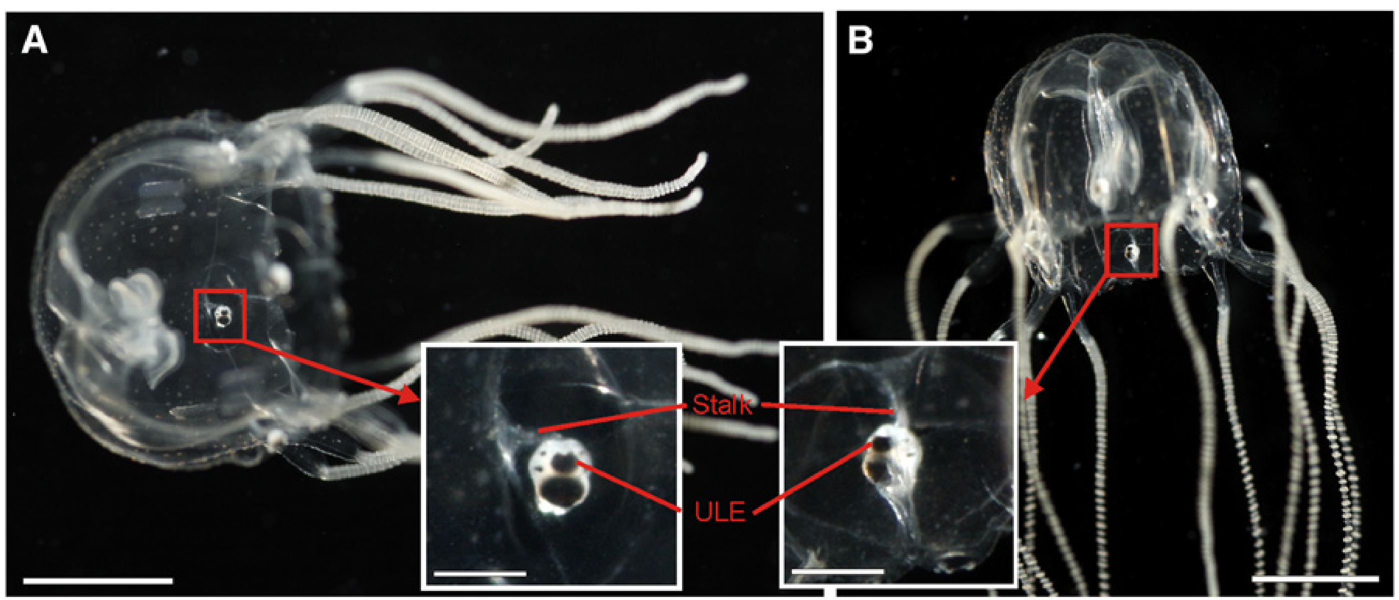

box jellyfish

…problem solving behavior, language, expert knowledge and application, and reason, are all pretty simple once the essence of being and reacting are available. That essence is the ability to move around in a dynamic environment, sensing the surroundings to a degree sufficient to achieve the necessary maintenance of life and reproduction. This part of intelligence is where evolution has concentrated its time--it is much harder.

— Rodney Brooks, "Intelligence without representation," *Artificial Intelligence* (1991)

Cybernetics/neural networks

Norbert Wiener

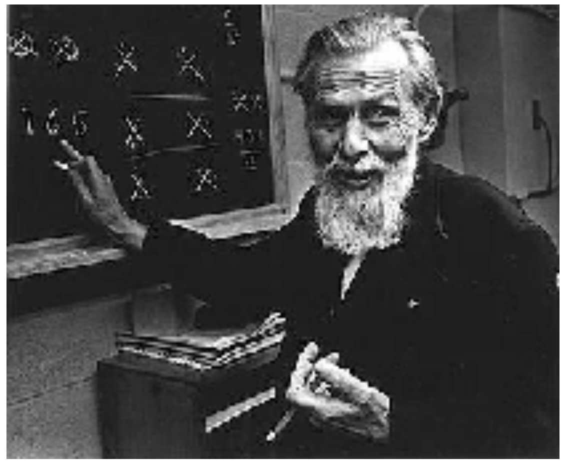

Warren McCulloch & Walter Pitts

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.

Hubel & Wiesel (1962, 1965)

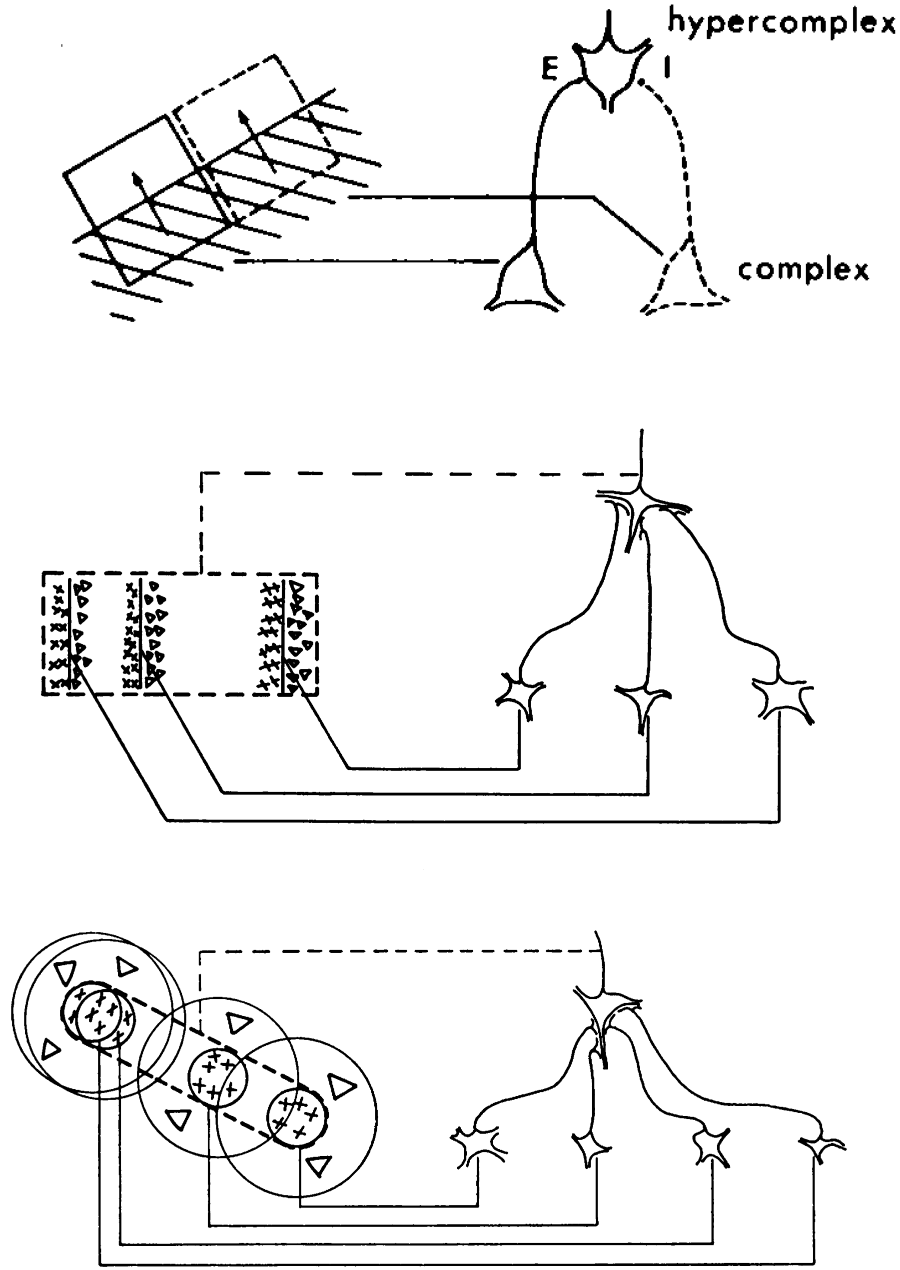

Hypercomplex

↑

Complex

↑

Simple

Neocognitron

(Fukushima 1980)

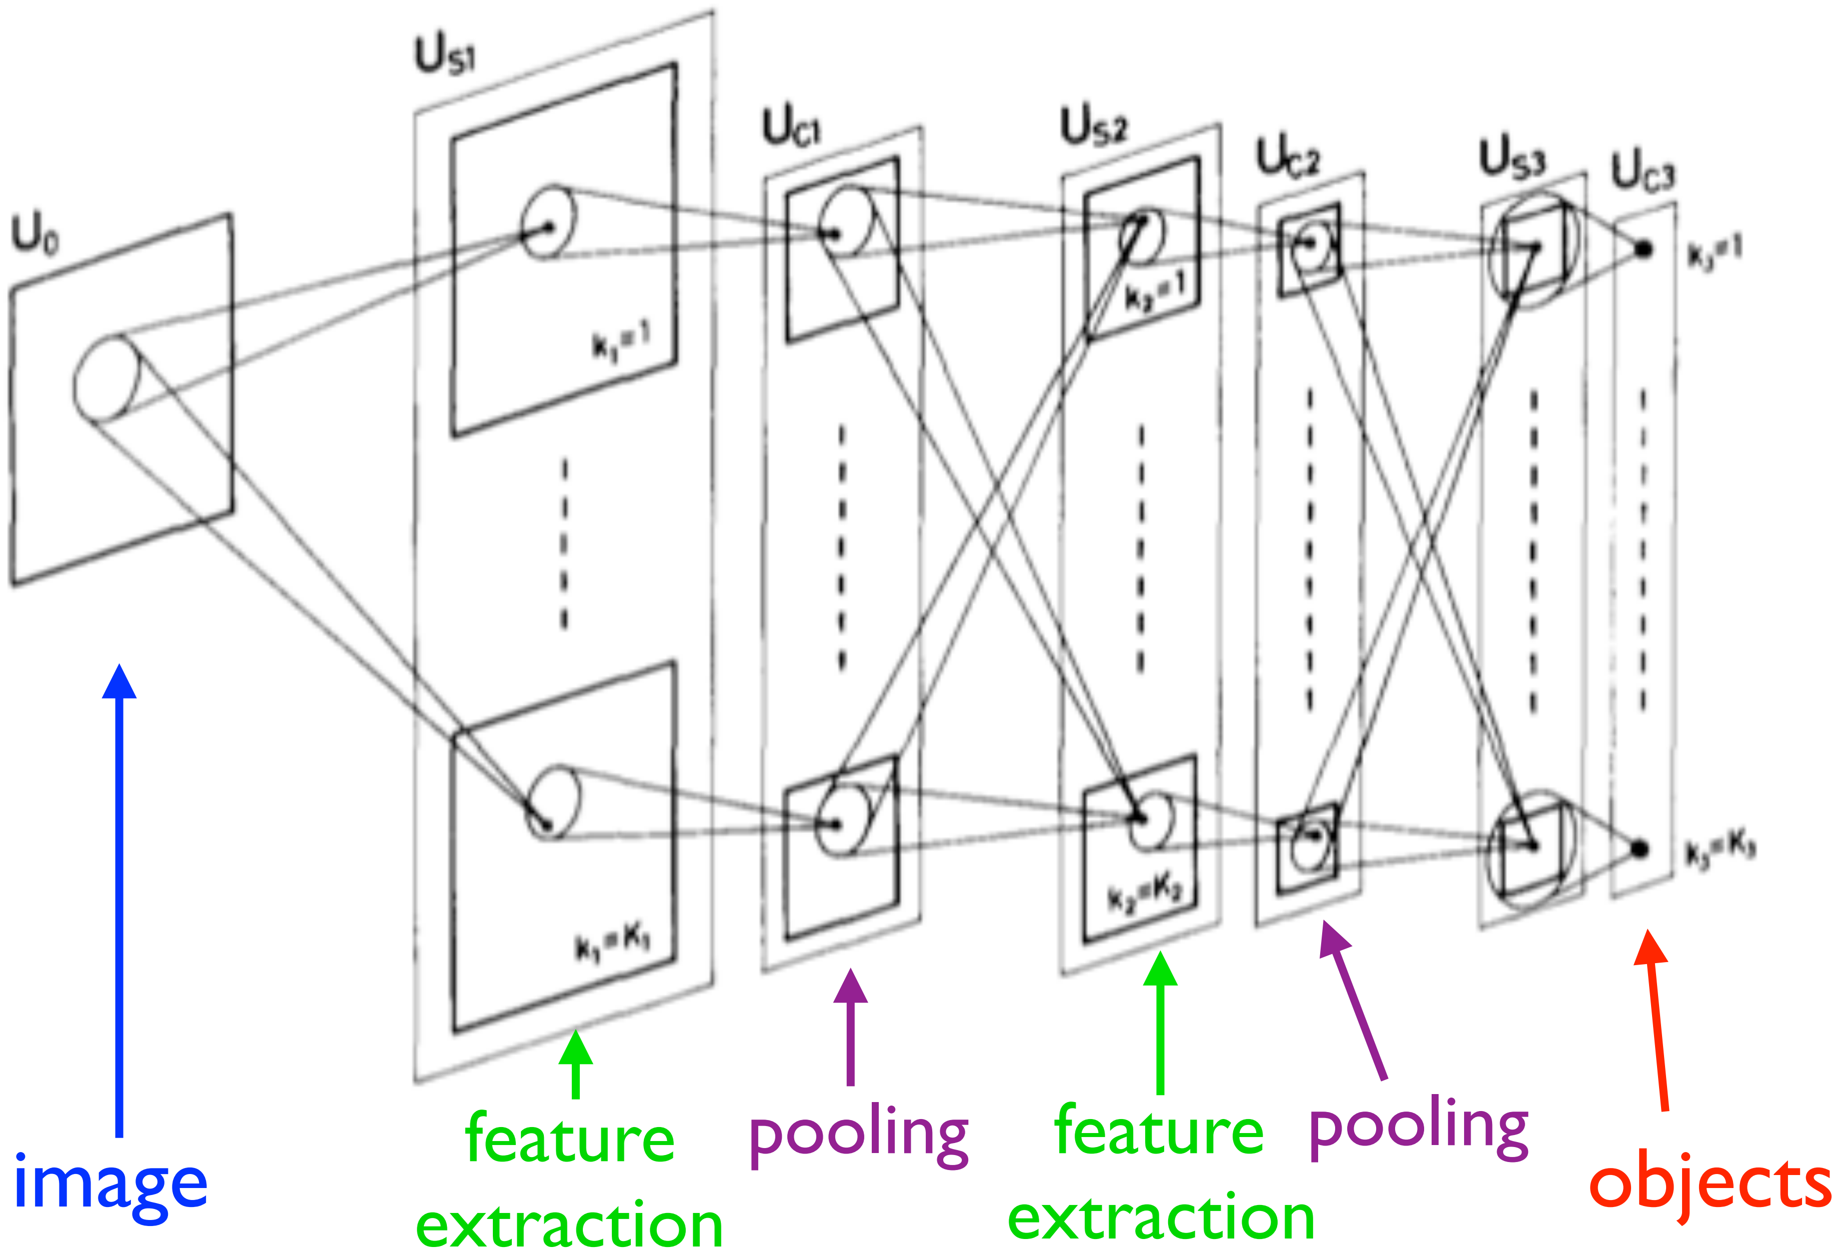

The approach of David Marr (ca. 1980)

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

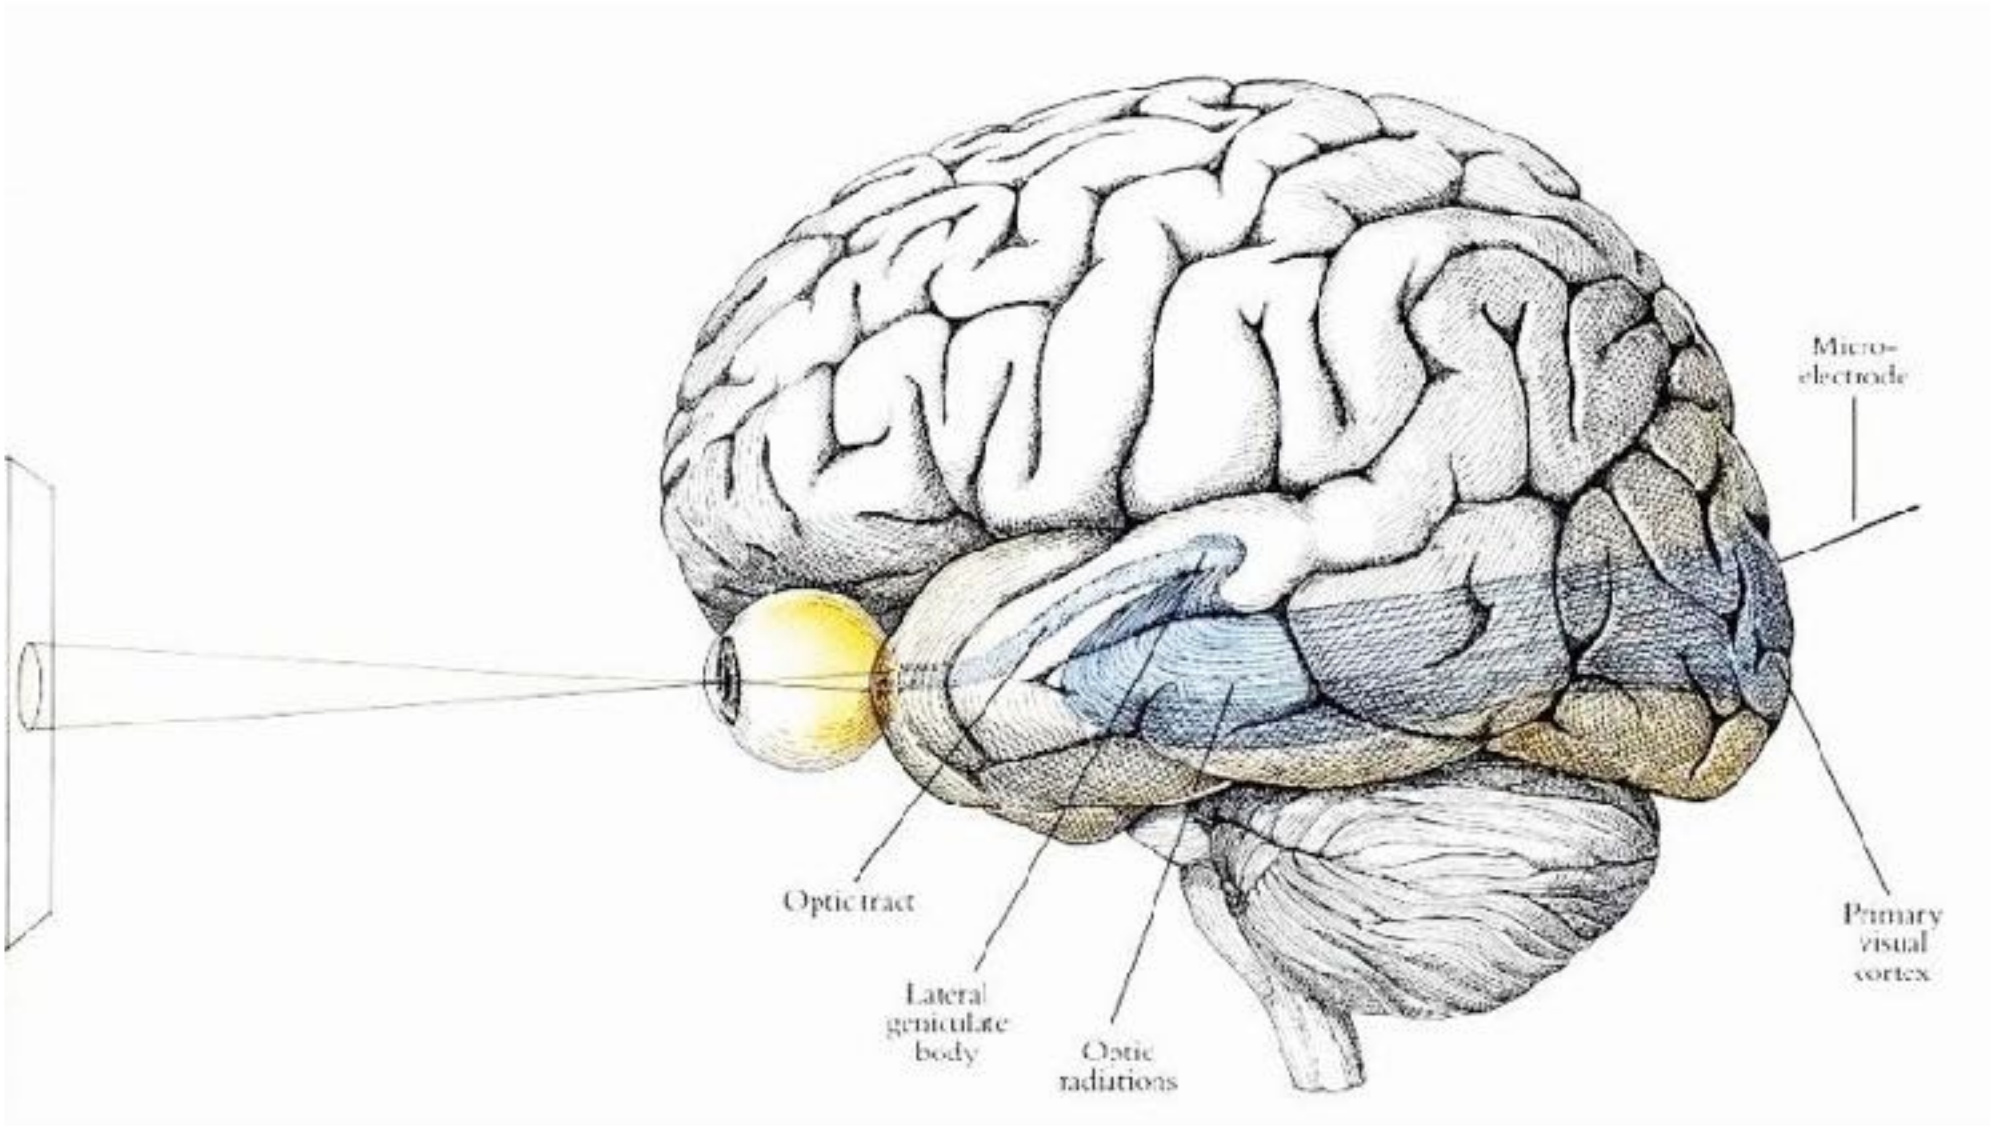

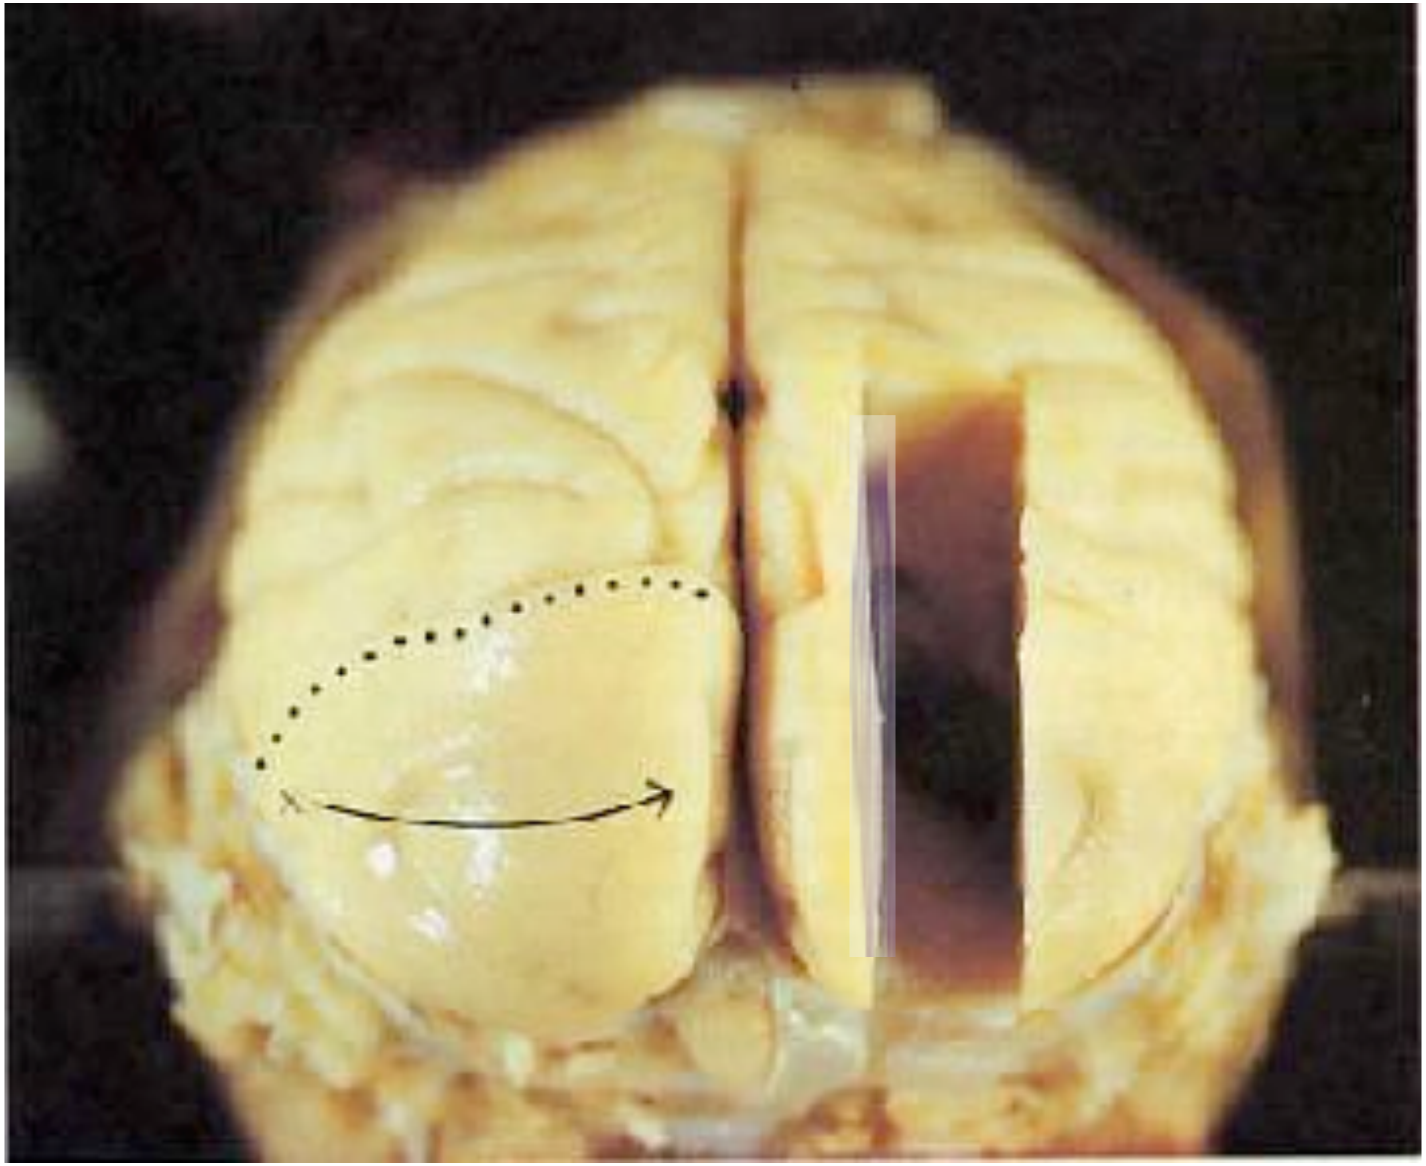

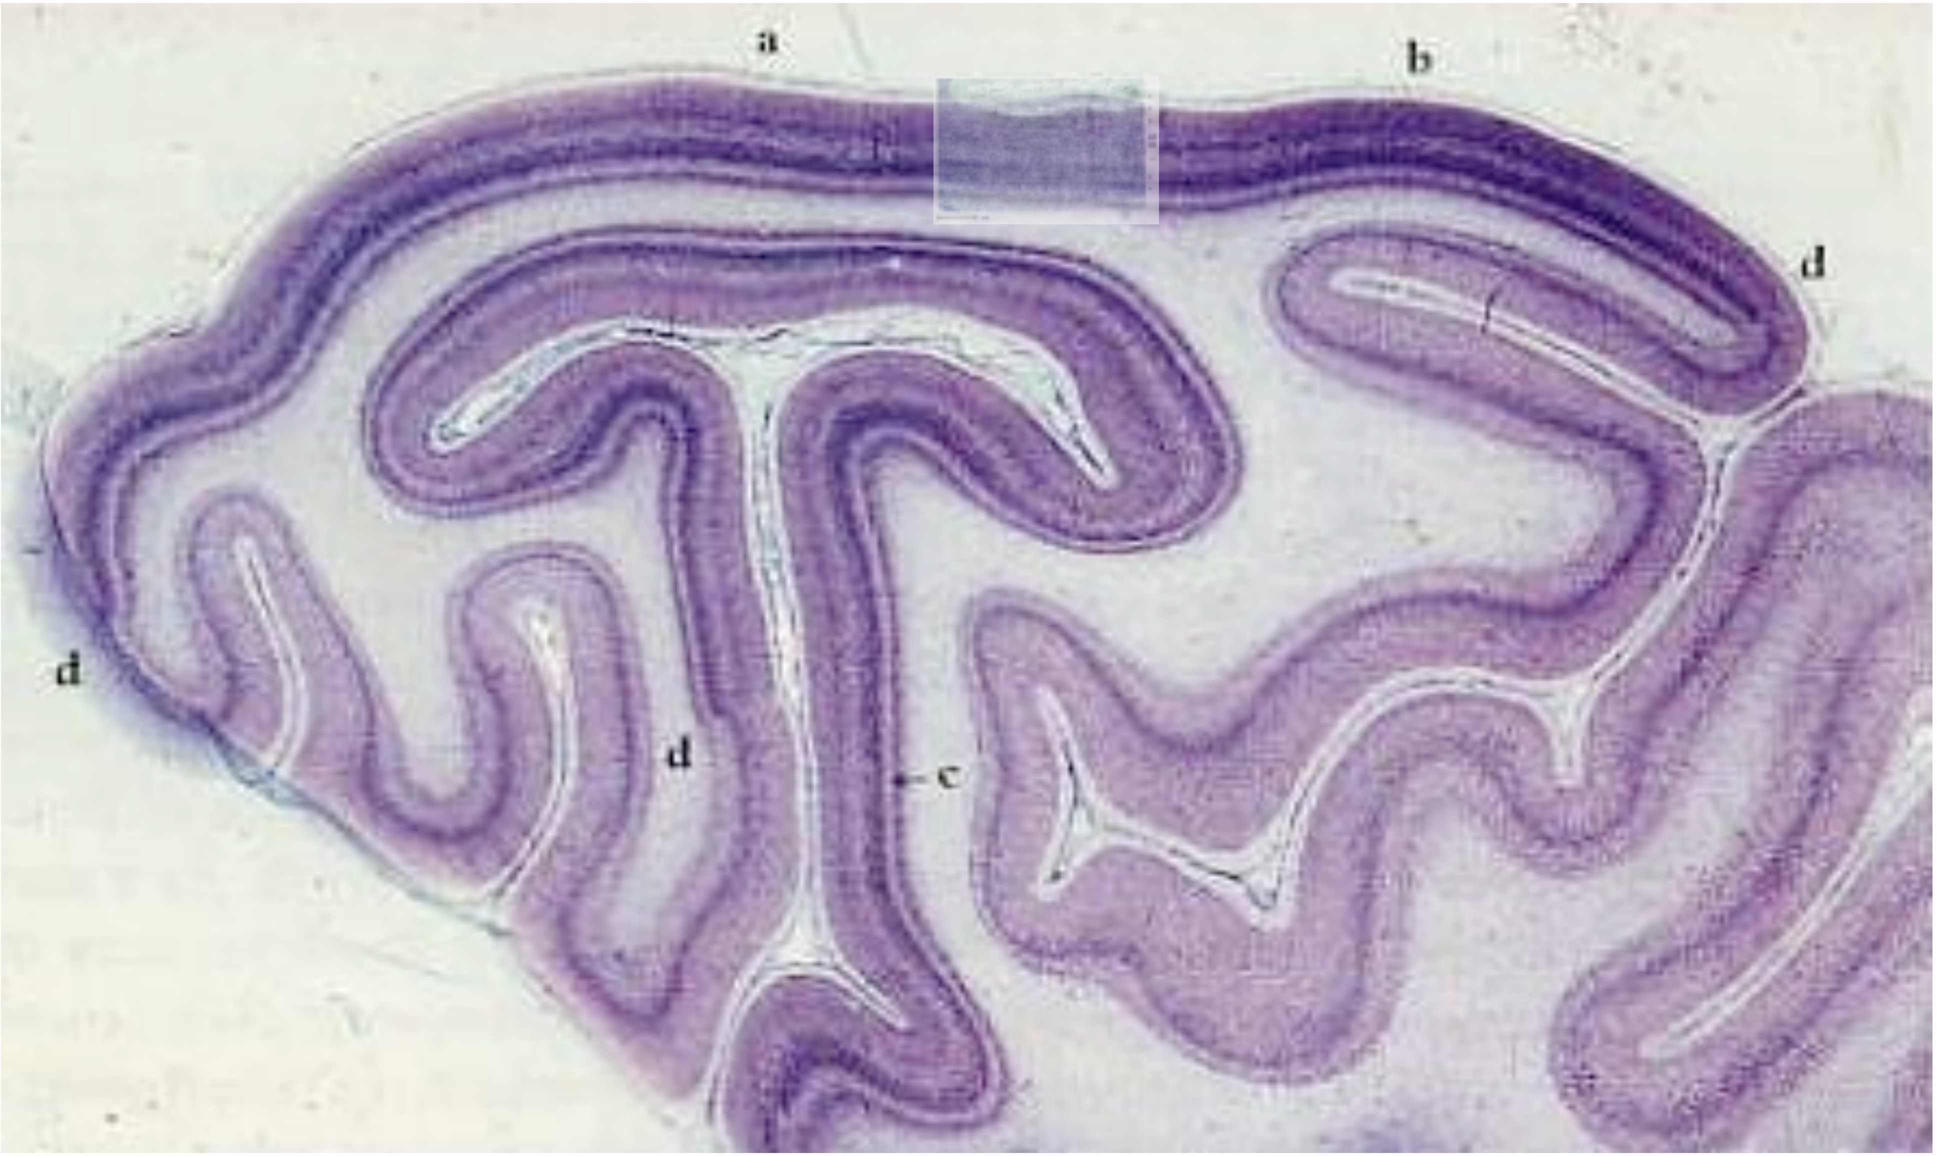
a
b
d
d
d
c

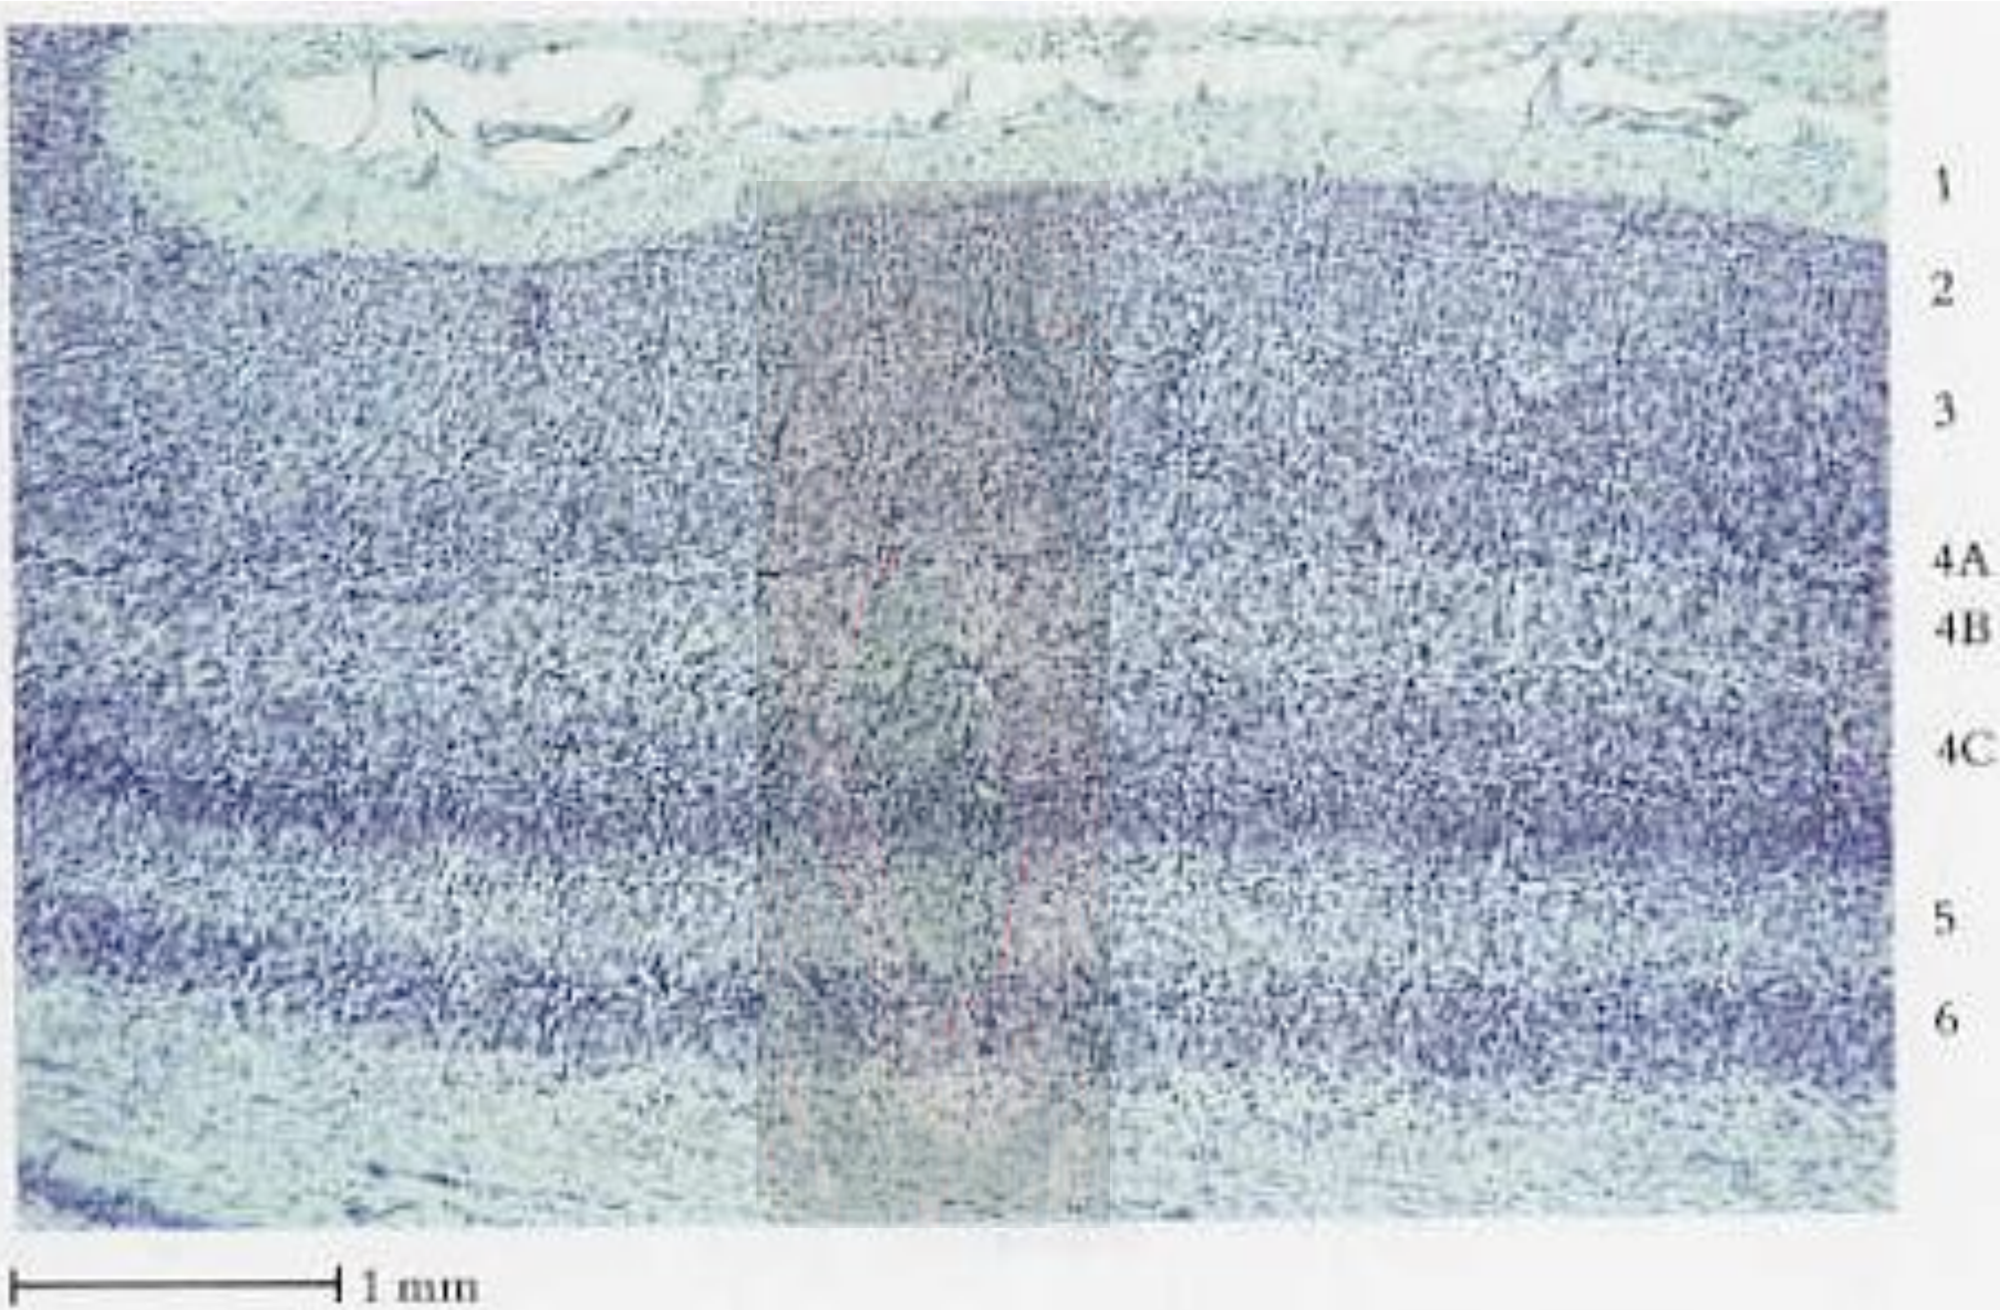

1
2
3
4A
4B
4C
5
6
1 mm

1 mm^2 of cortex contains 100,000 neurons

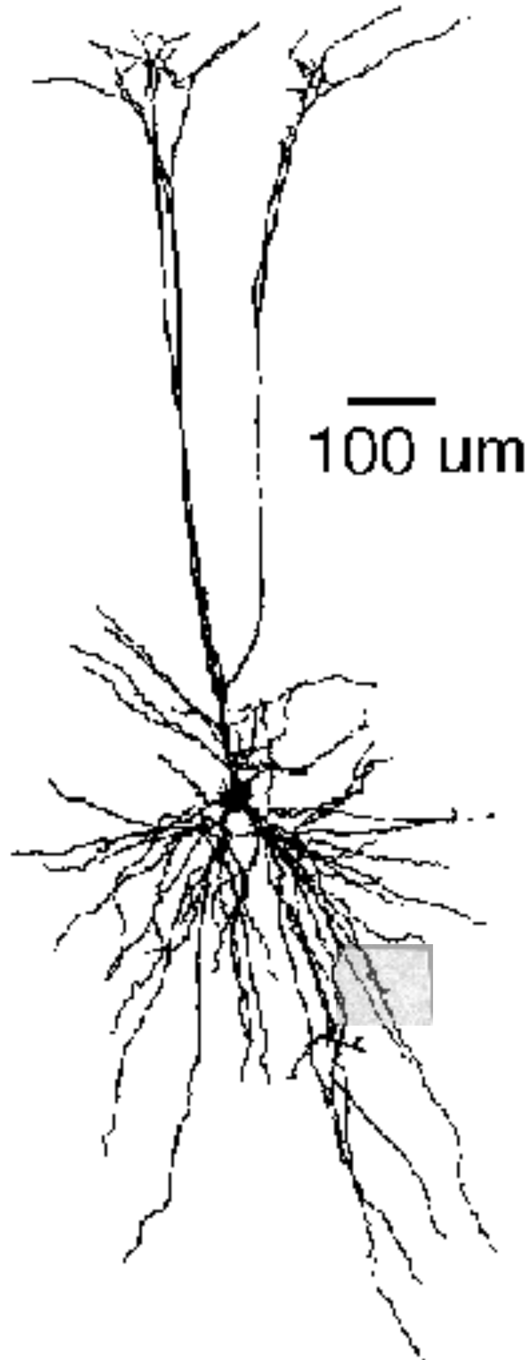
100 um

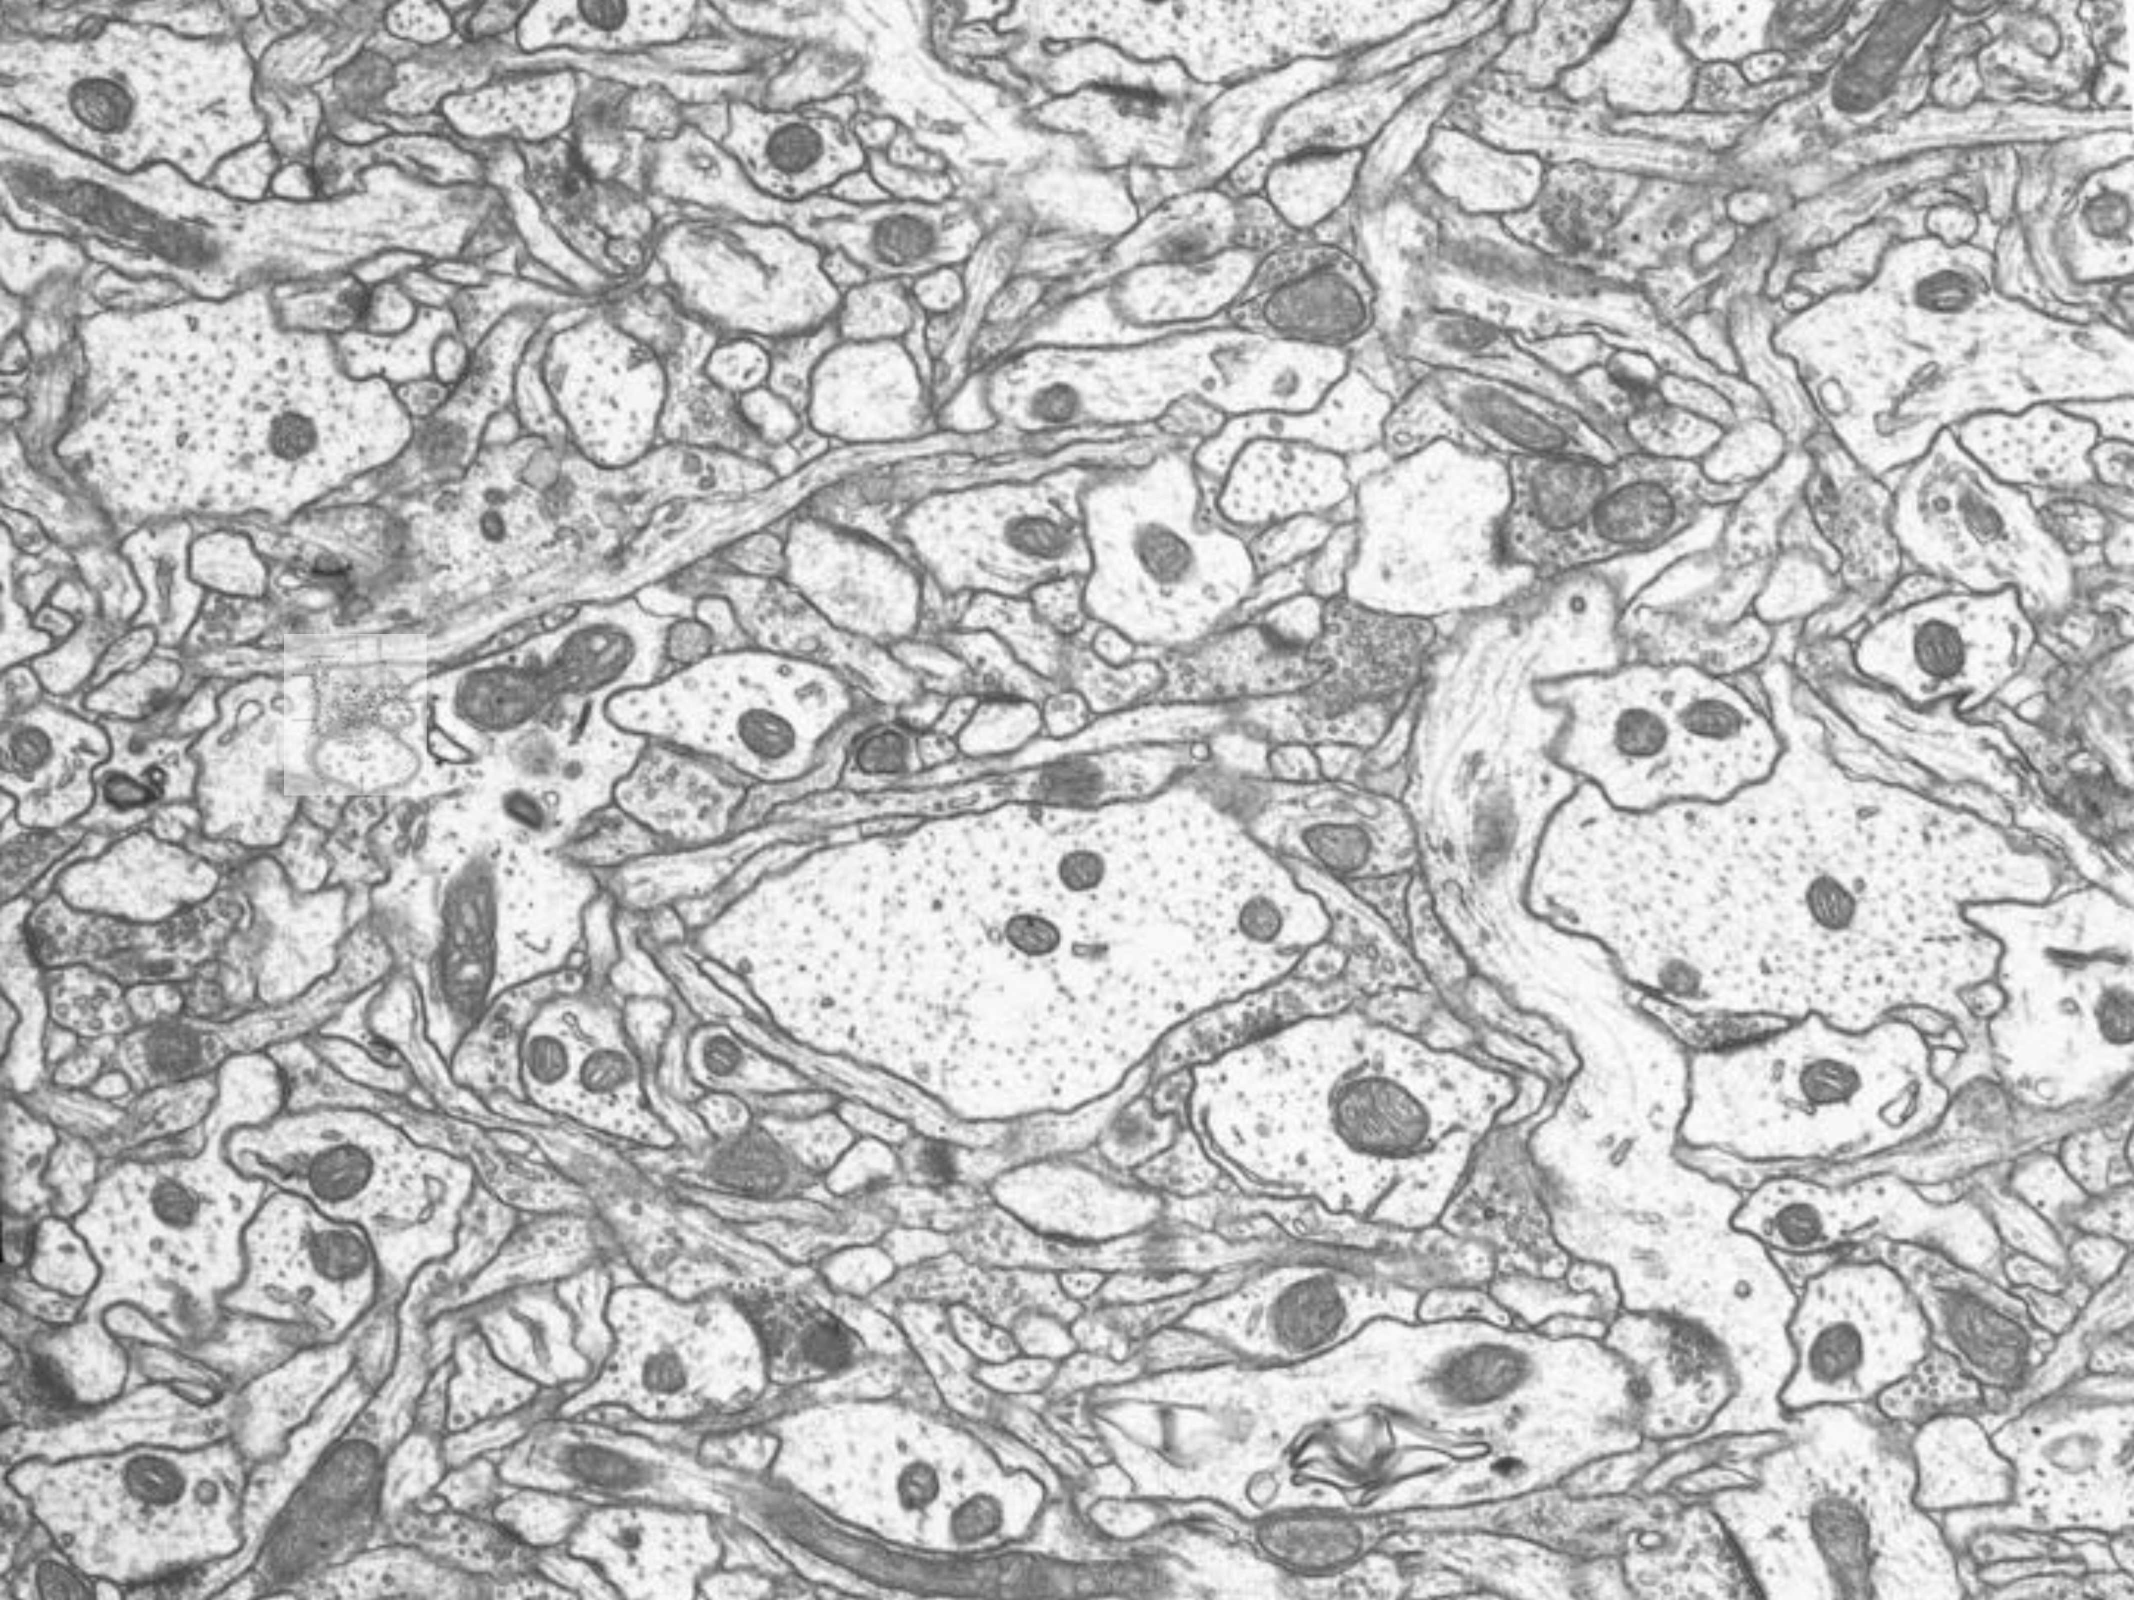

Anatomy of a synapse

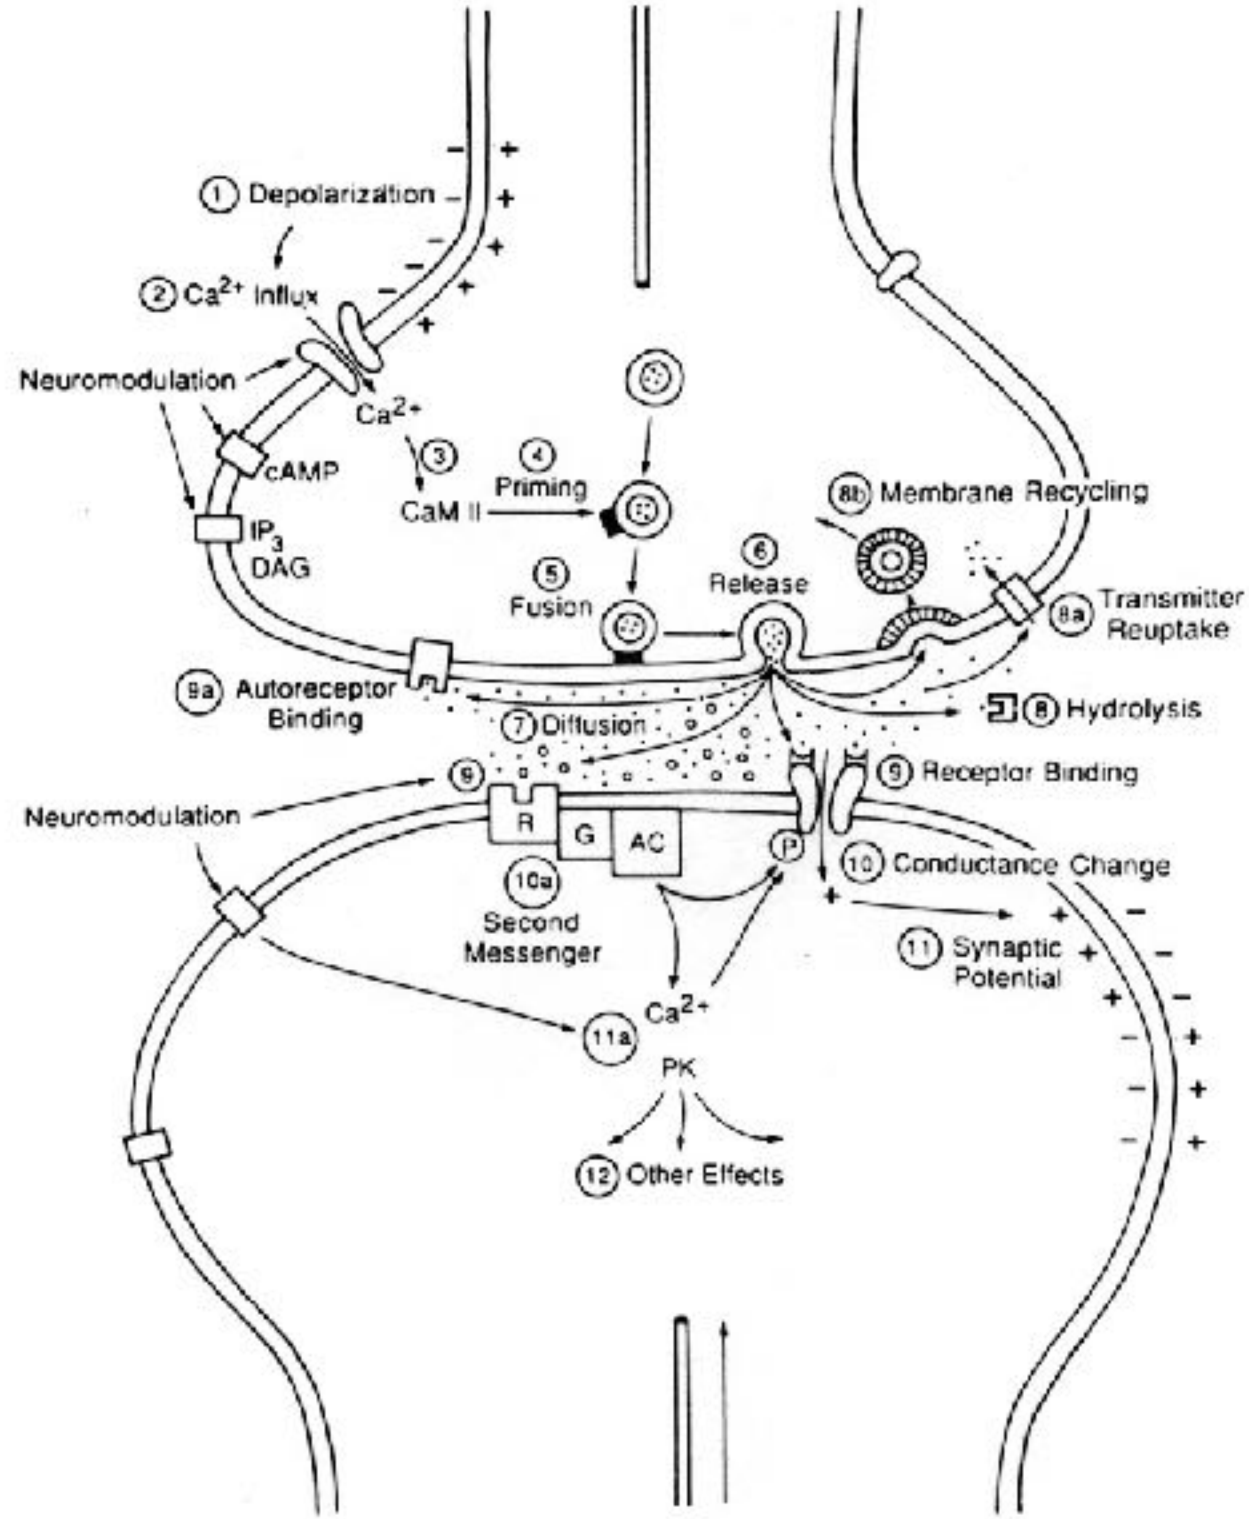

Are there principles?

"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)

Ramón y Cajal on theorists

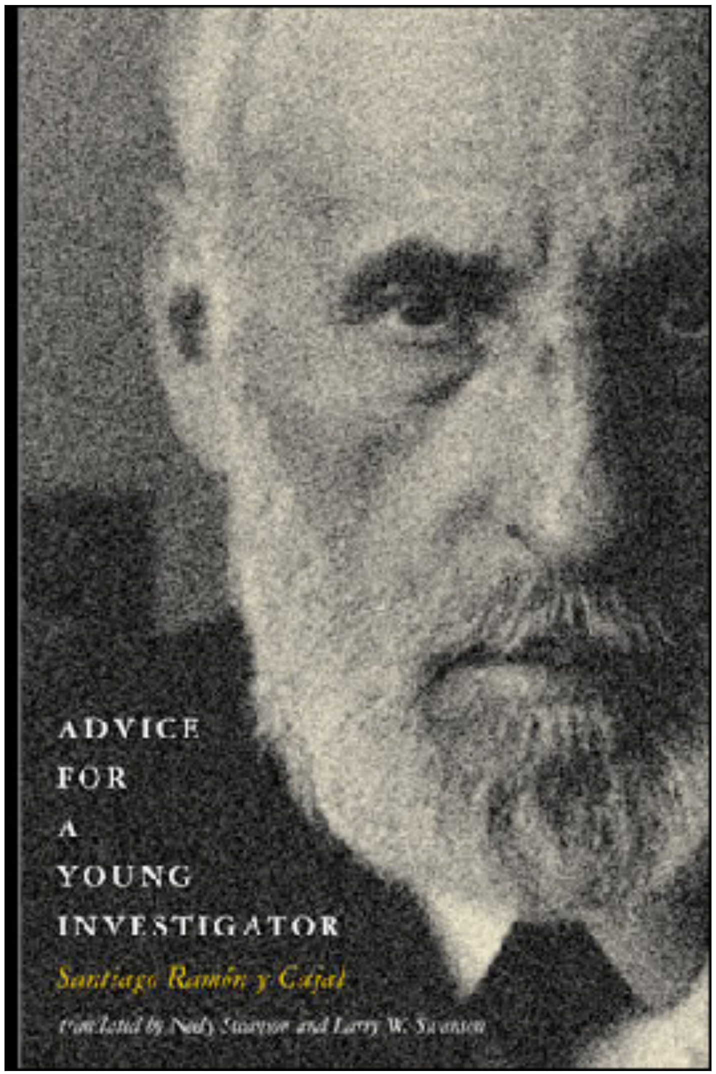

Basically, the theorist is a lazy person masquerading as a diligent one. He unconsciously obeys the law of minimum effort because it is easier to fashion a theory than to discover a phenomenon.

-- From the chapter "Diseases of the Will"

Principles of optics govern the design of eyes

THE EVOLUTION OF EYES

Michael F. Land *Russell D. Fernald*

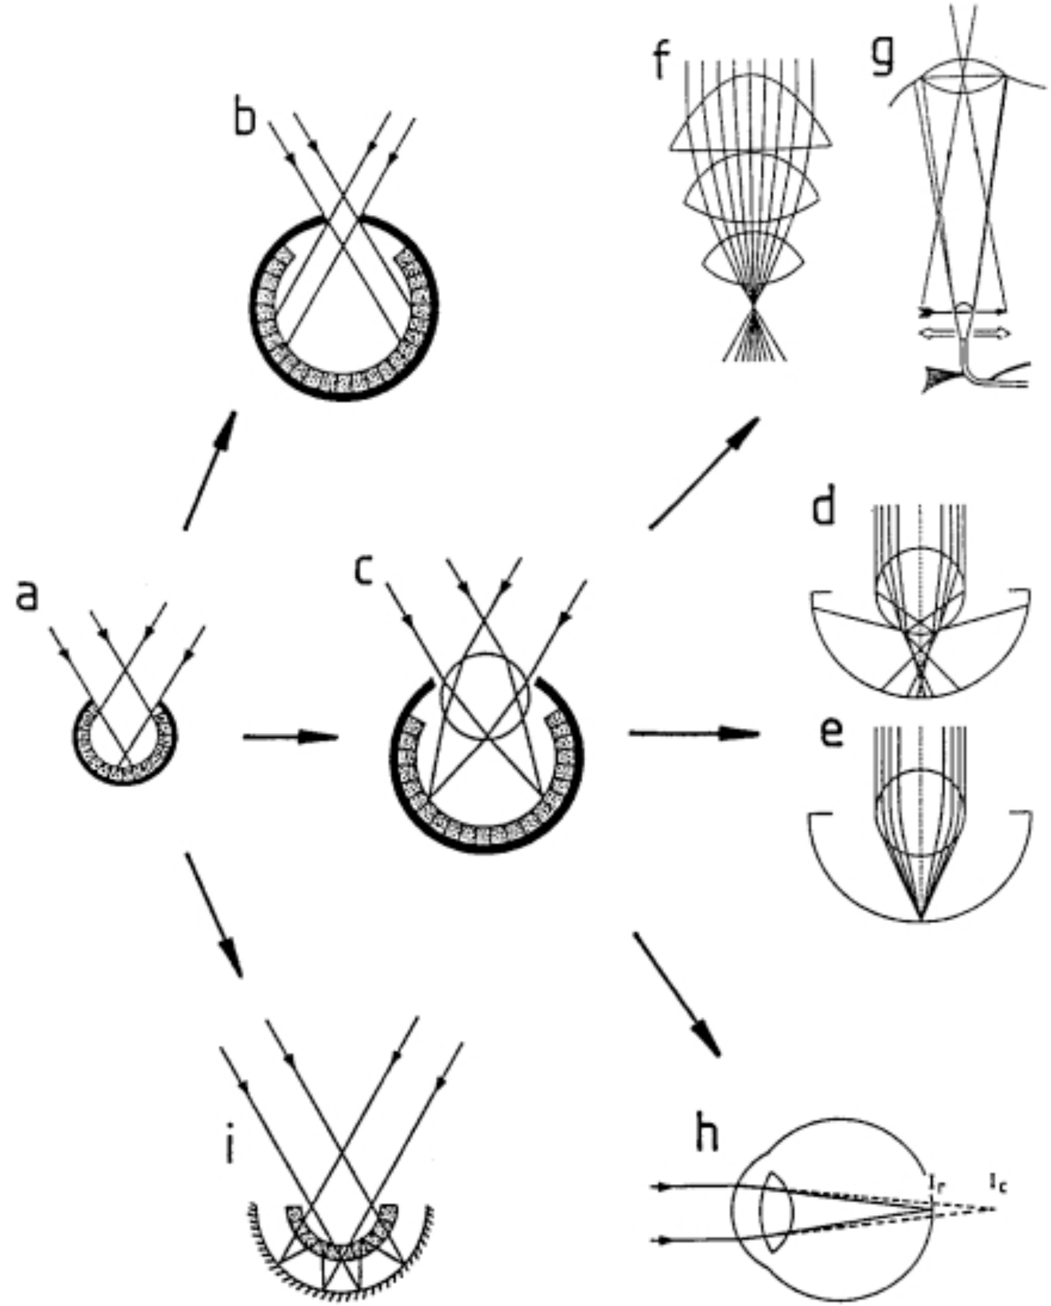

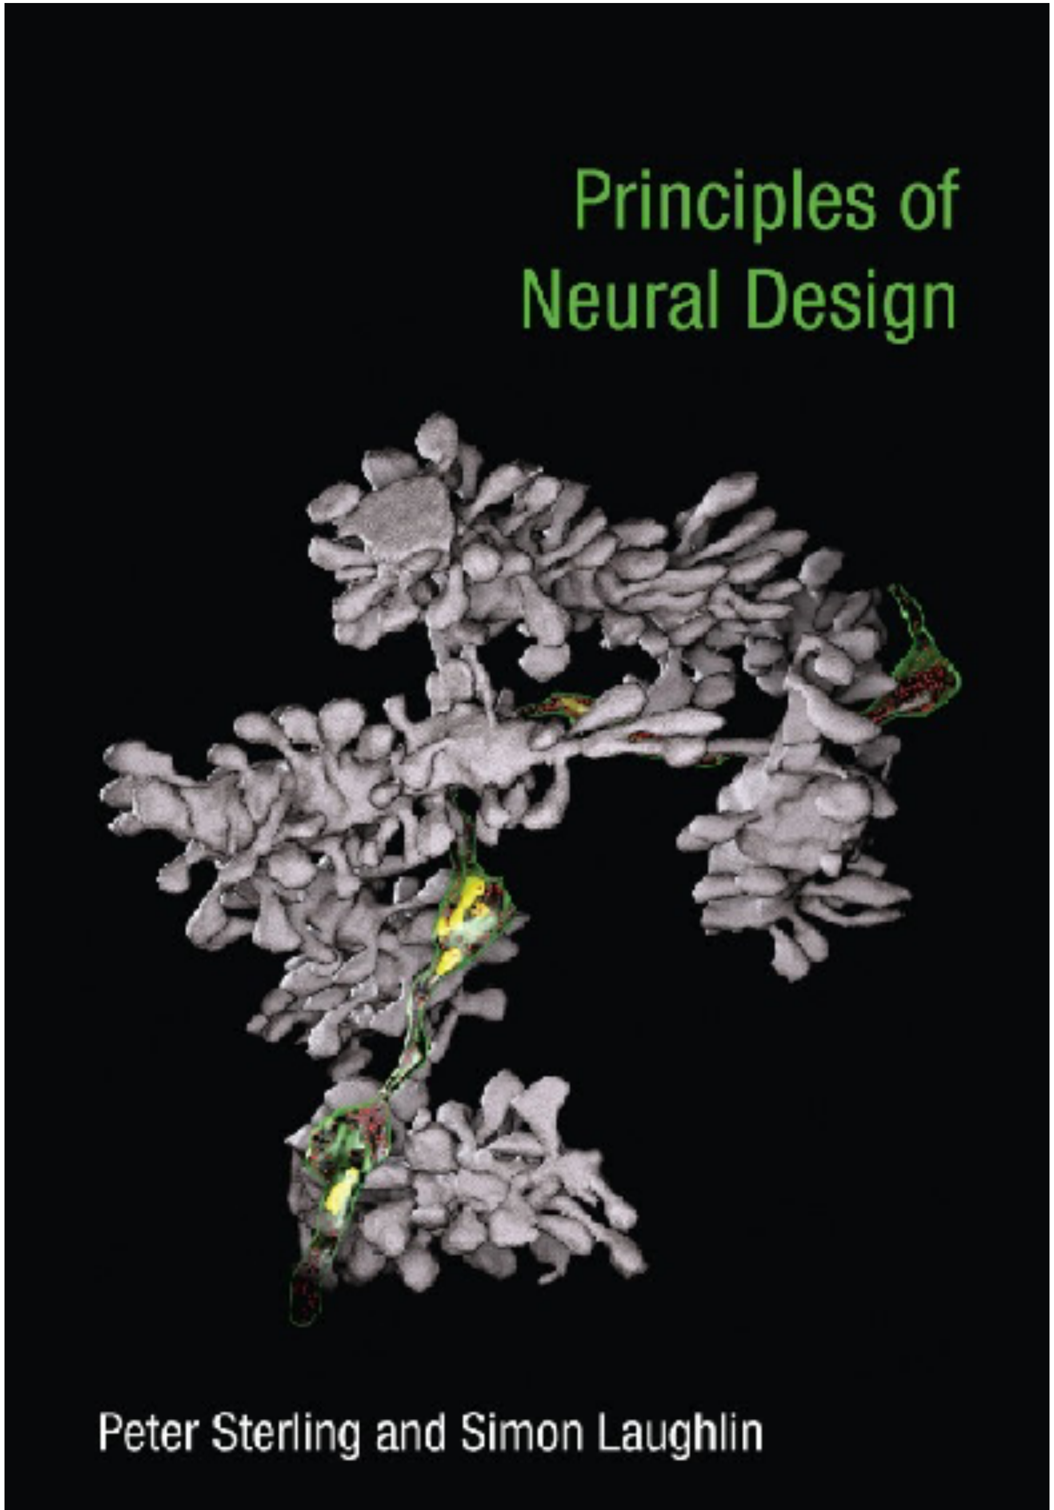

Principles

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

Zhang & Sejnowski (2000)

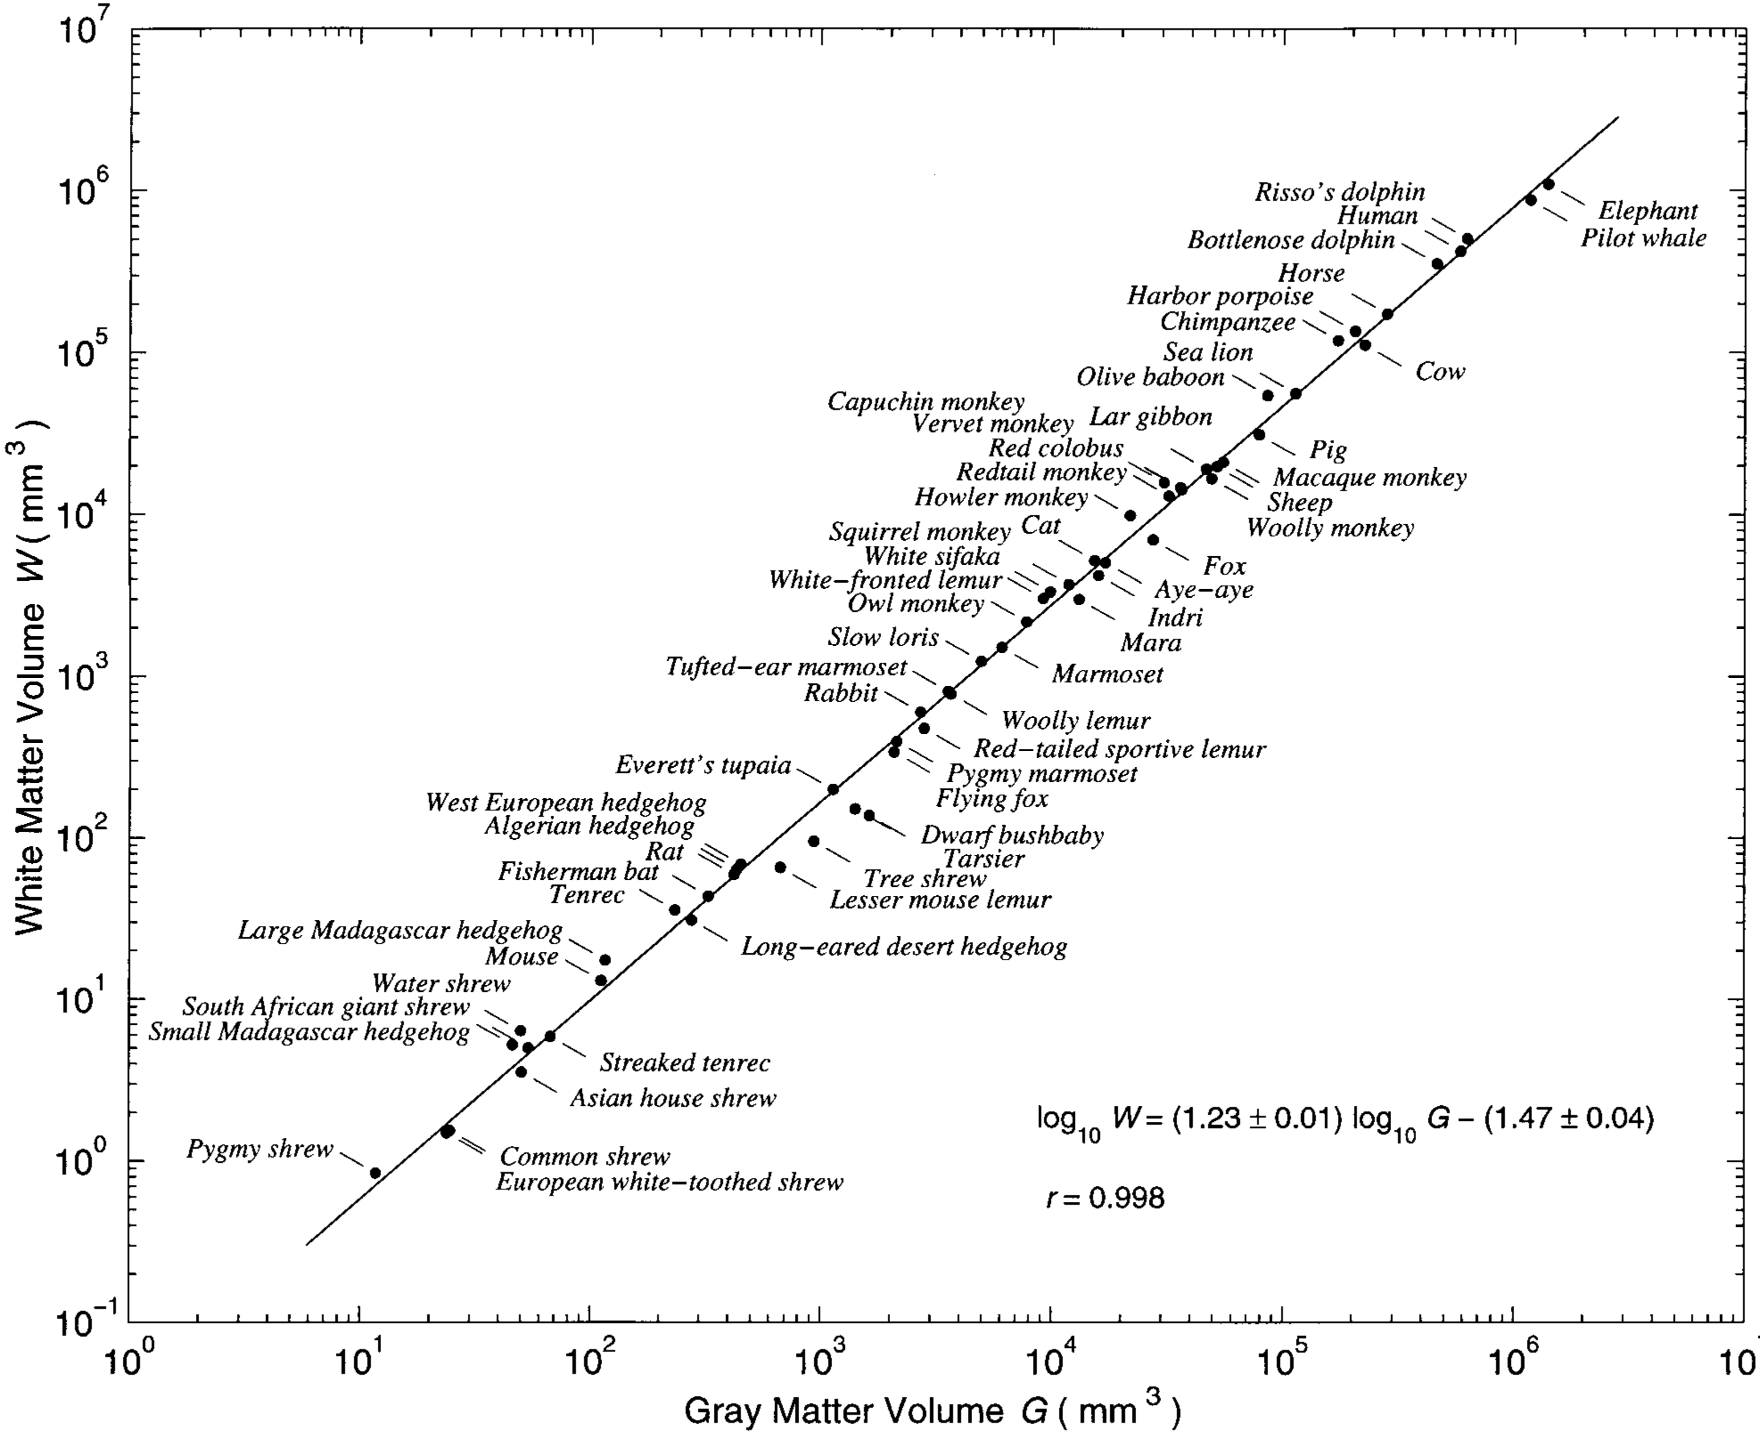

Wallisch & Movshon (2008)

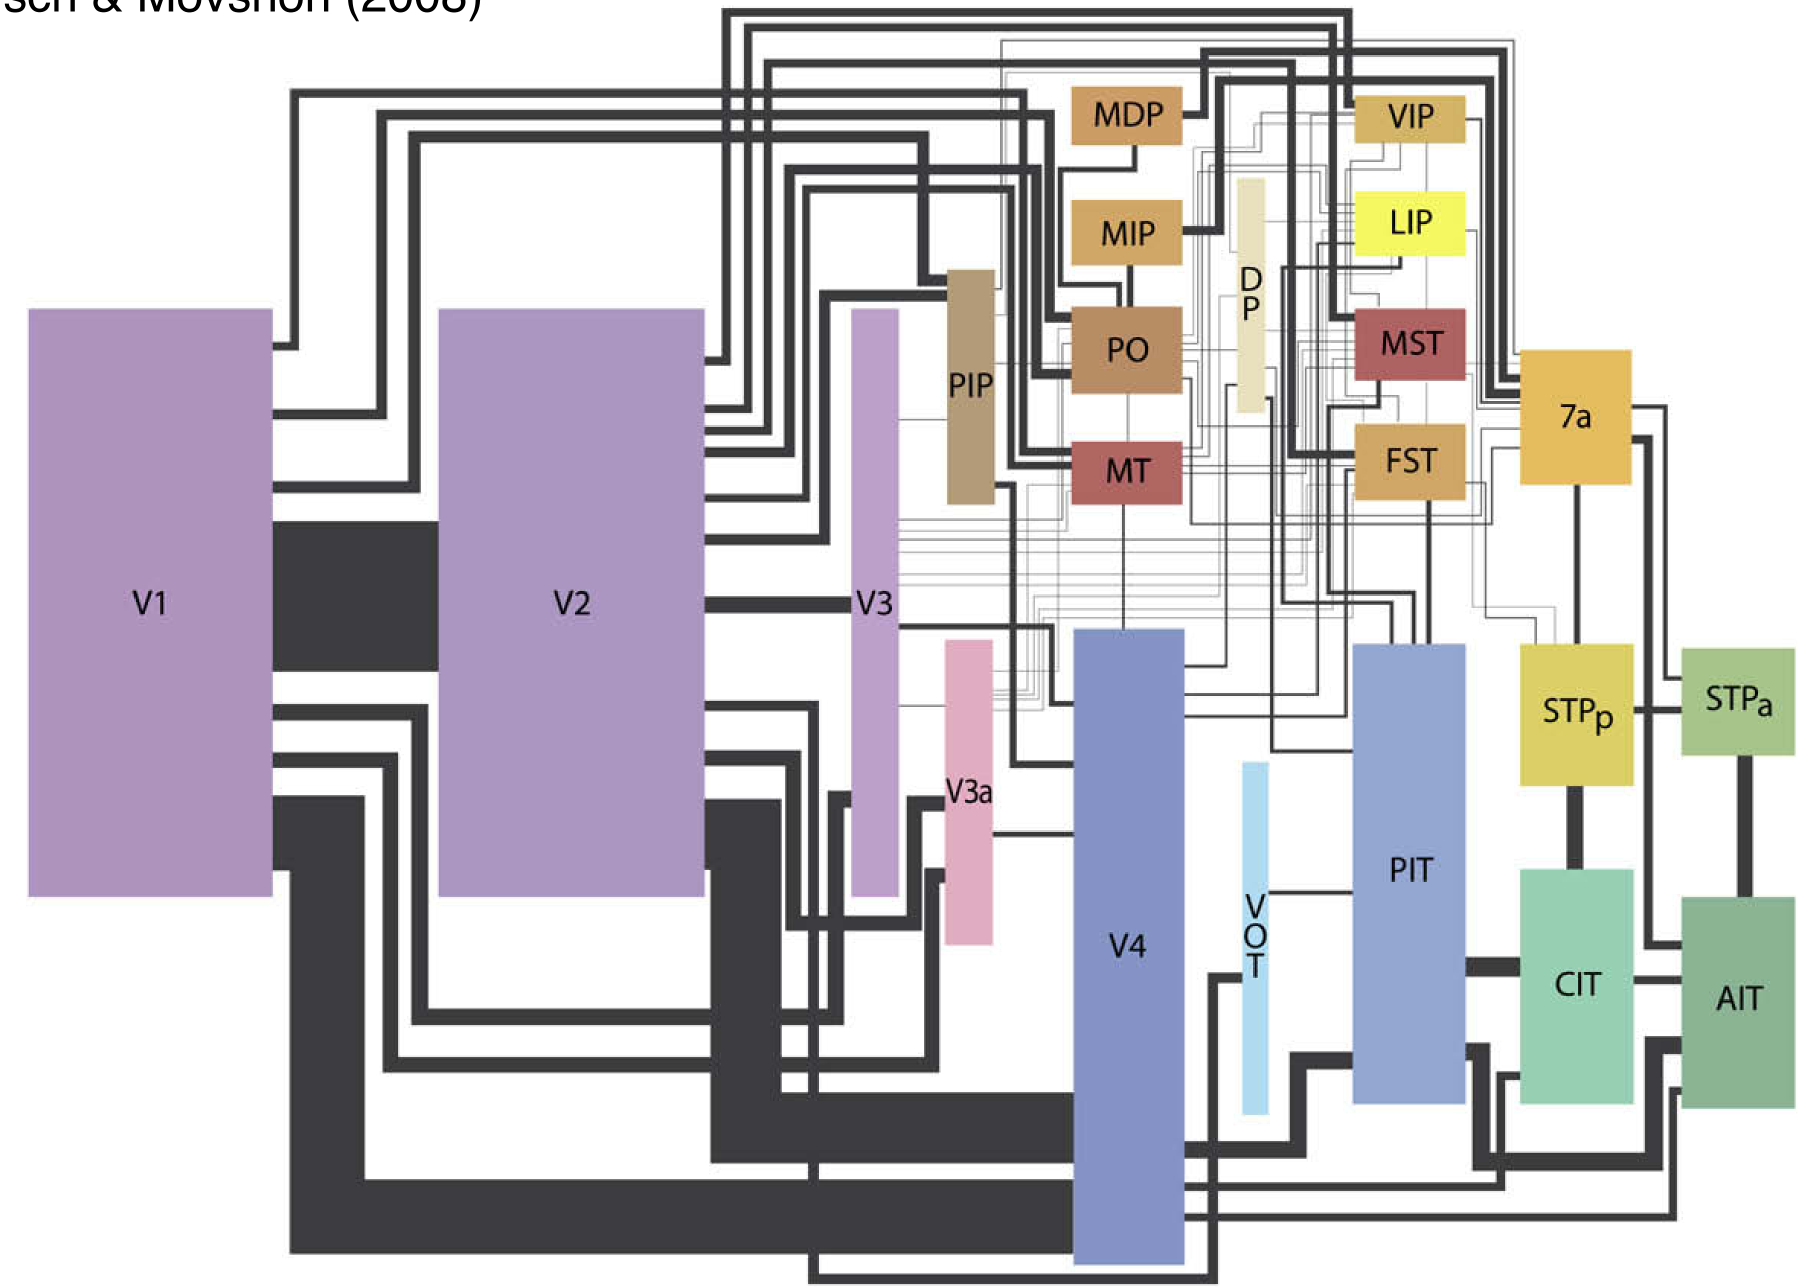

‘Gabor filters’ · · ? · · objects · faces

Cortical neurons

- have elaborate *dendritic* and *axonal* arbors
- are highly organized by layer
- are interconnected in a 'canonical microcircuit'

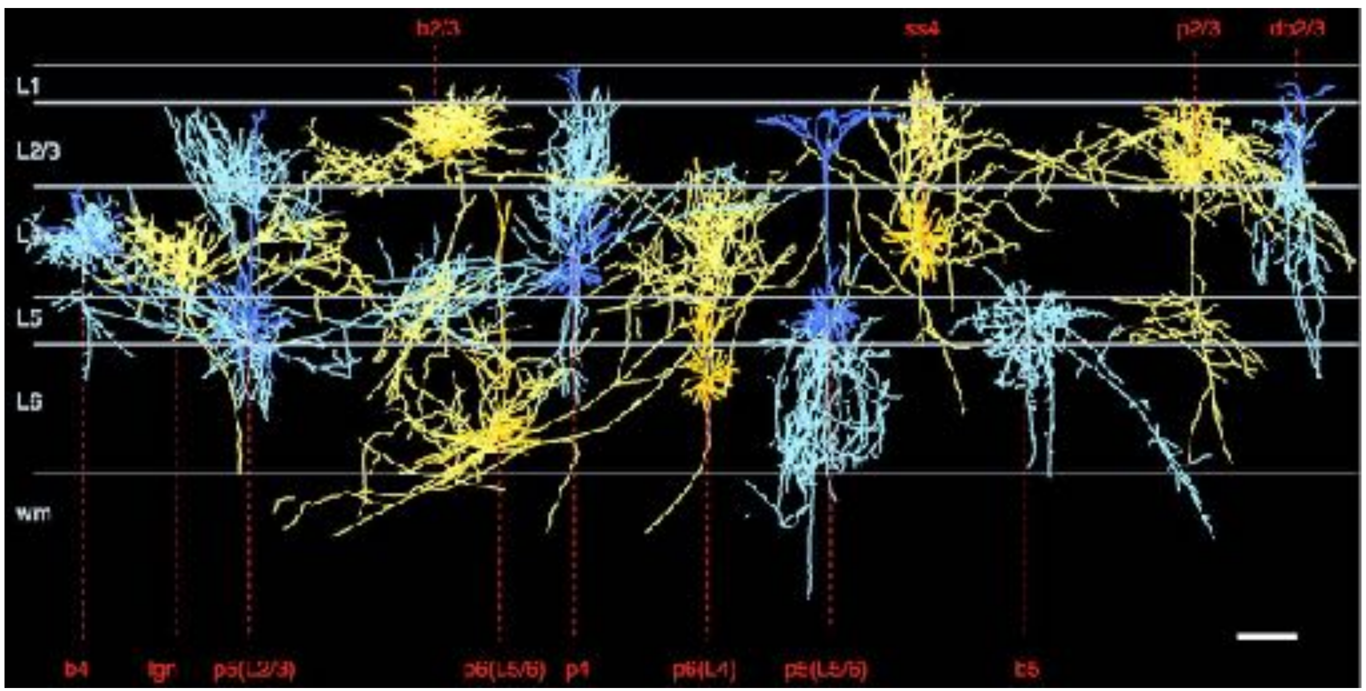

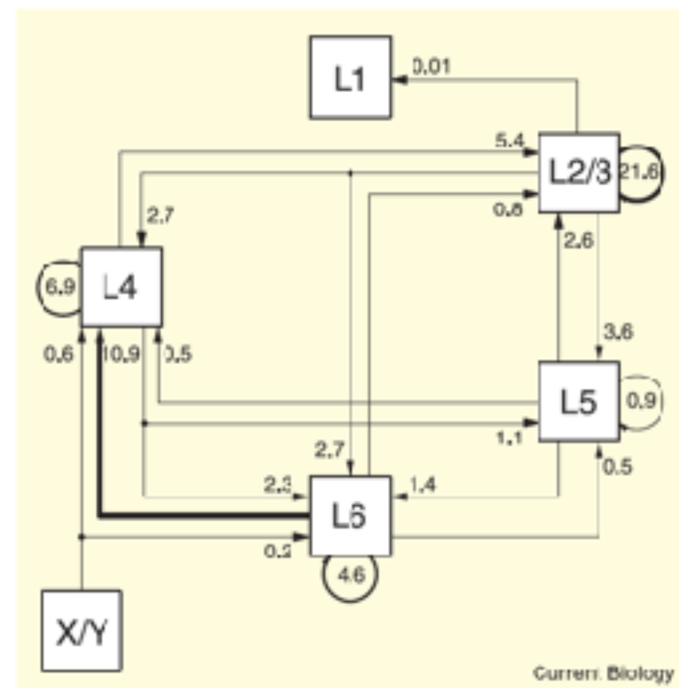

(Douglas and Martin, 2007)

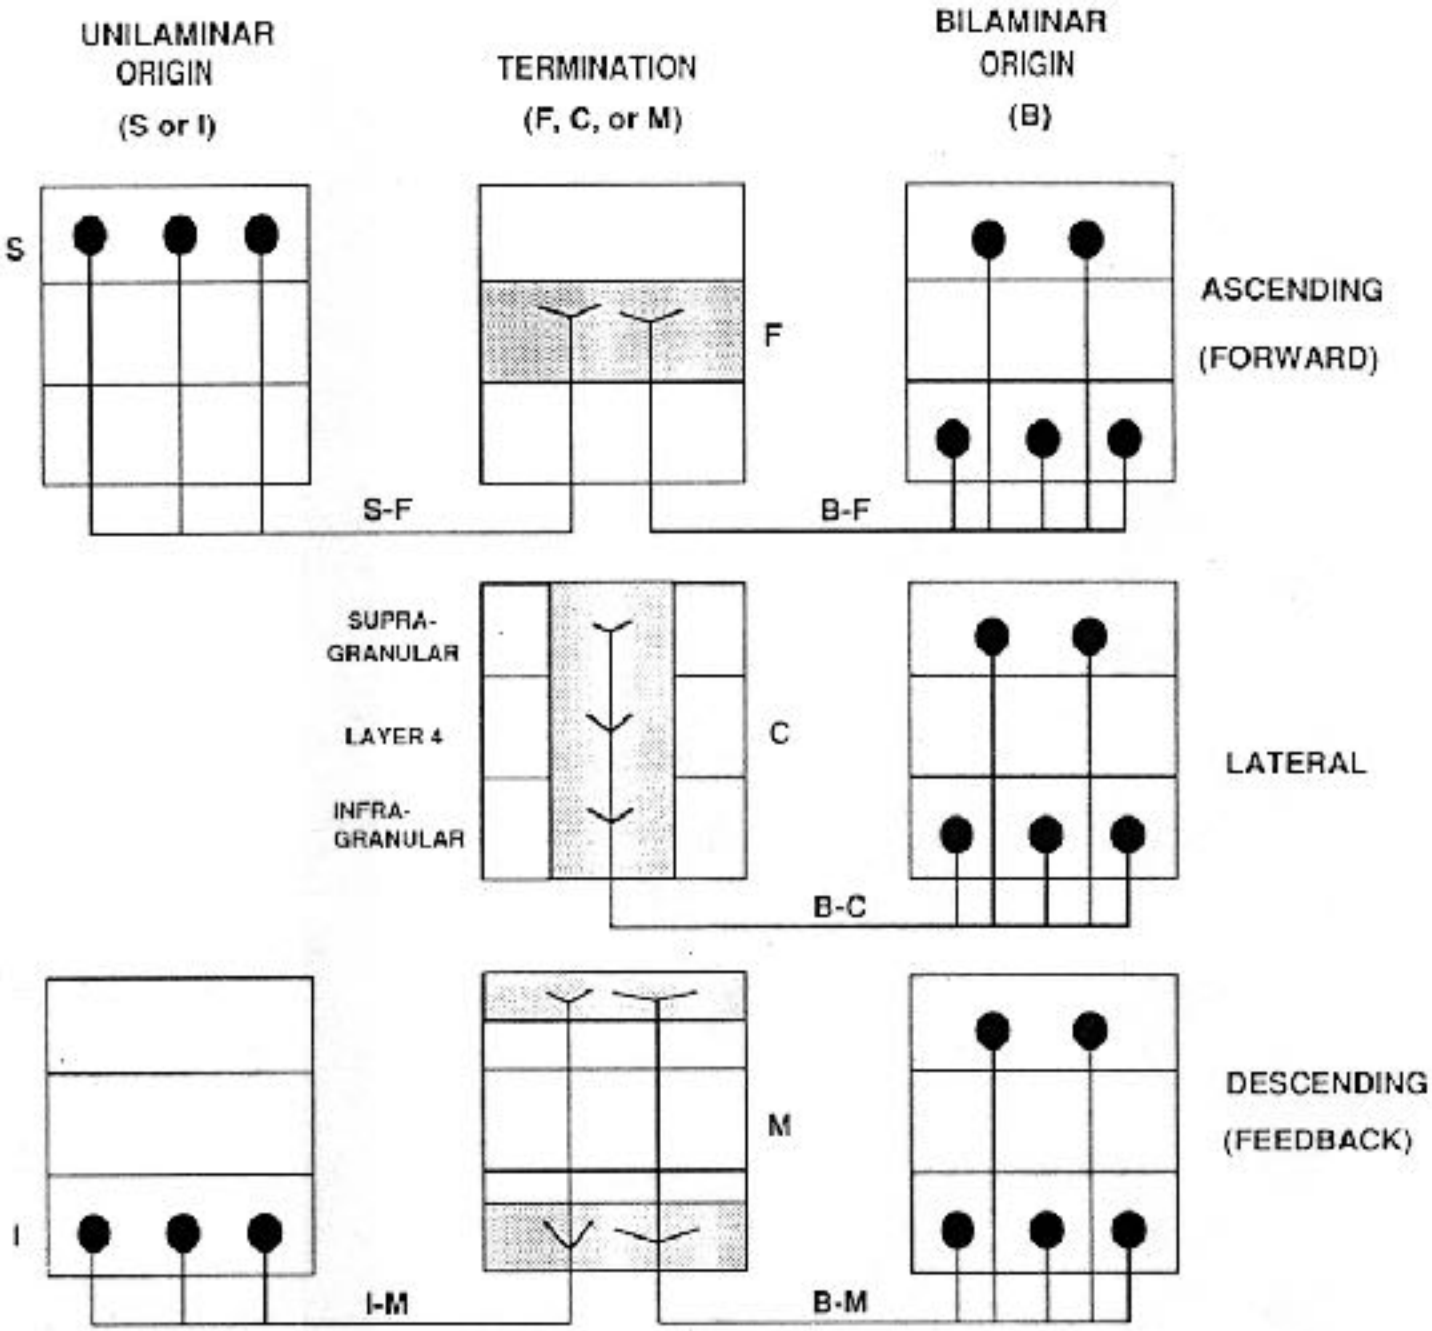

Felleman & Van Essen 1991

Computational principles

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

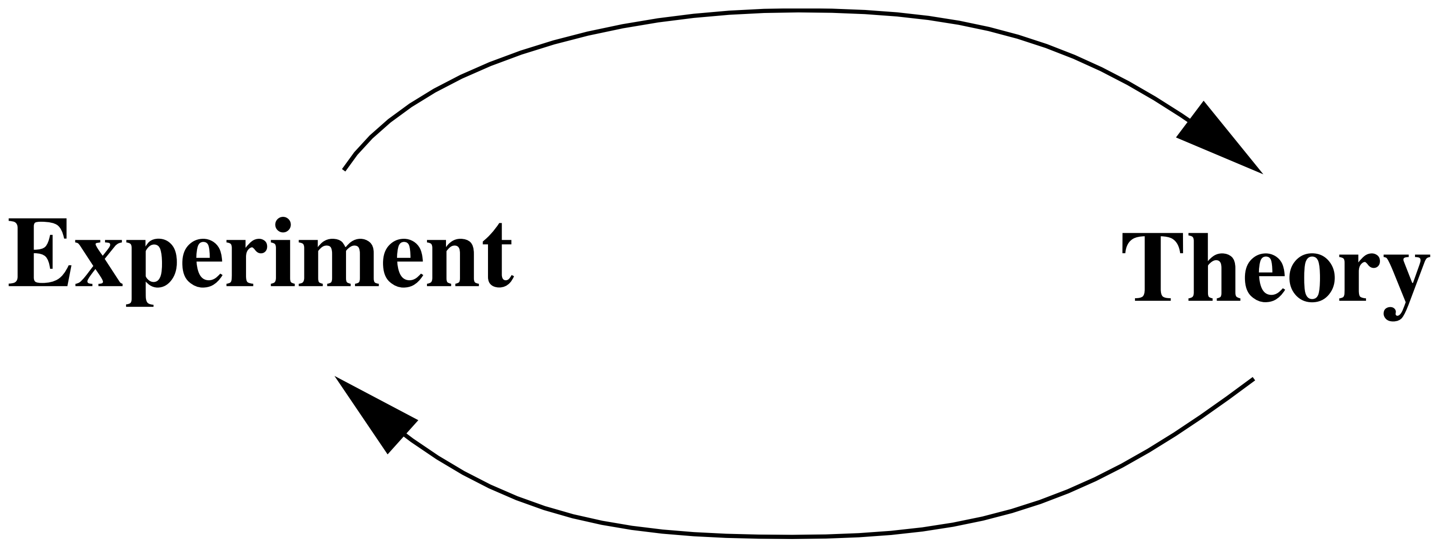
Experiment
Theory

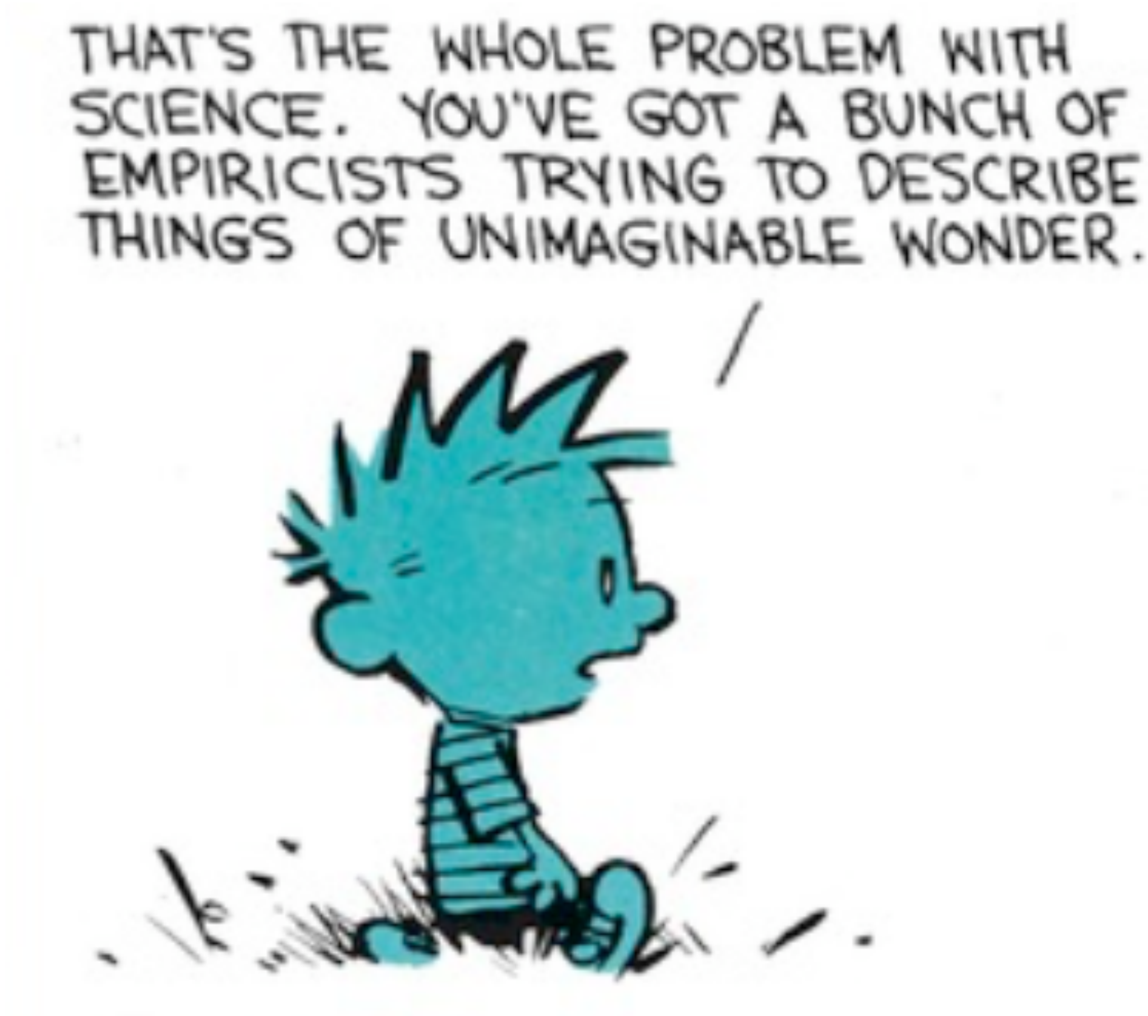
THAT'S THE WHOLE PROBLEM WITH SCIENCE. YOU'VE GOT A BUNCH OF EMPIRICISTS TRYING TO DESCRIBE THINGS OF UNIMAGINABLE WONDER.