Readings for this week
(available on the course website)

Today:

- Dreyfus, H.L. and Dreyfus, S.E. *Making a Mind vs. Modeling the Brain: Artificial Intelligence Back at a Branchpoint*. Daedalus, Winter 1988.
- Mitchell, M. *Why AI is Harder Than We Think.*

Next week:

- Sterling & Laughlin chapters 6 (pp. 138-154), and 7
- Dayan & Abbott, Chapter 5.1-5.6
- Mead, C. *Analog VLSI and Neural Systems.* Chapter 1: Introduction and Chapter 4: Neurons.
- Handout on *Linear Neuron Models*
- Carandini M, Heeger D (1994) *Summation and division by neurons in primate visual cortex.*

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

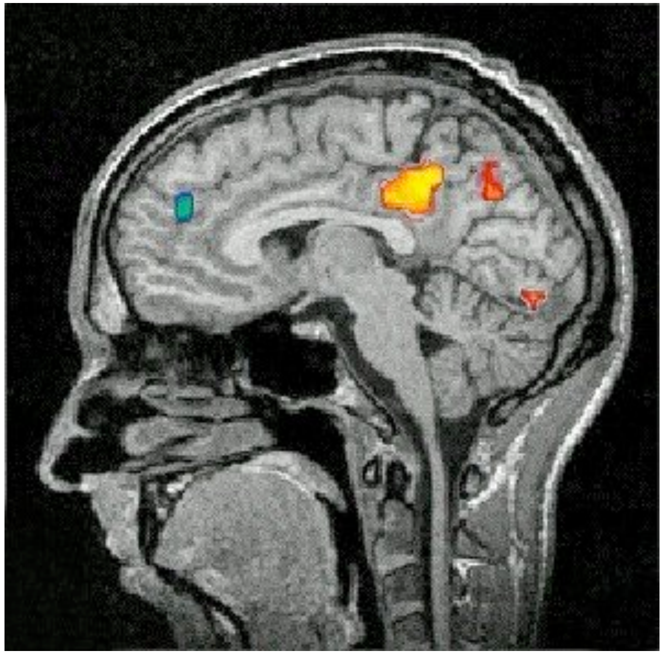

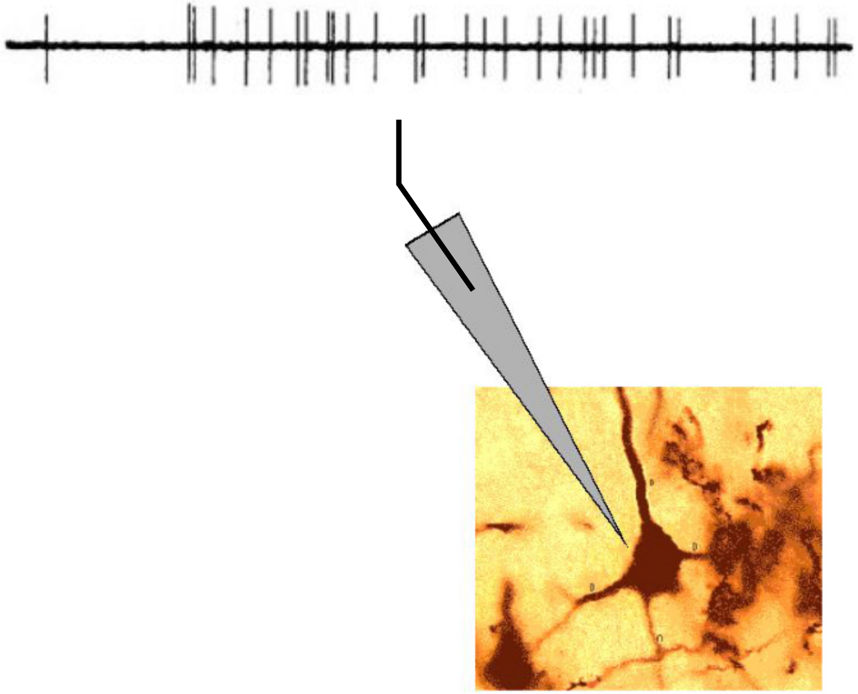

Spontaneous activity of more than 10,000 neurons in visual cortex of awake mice

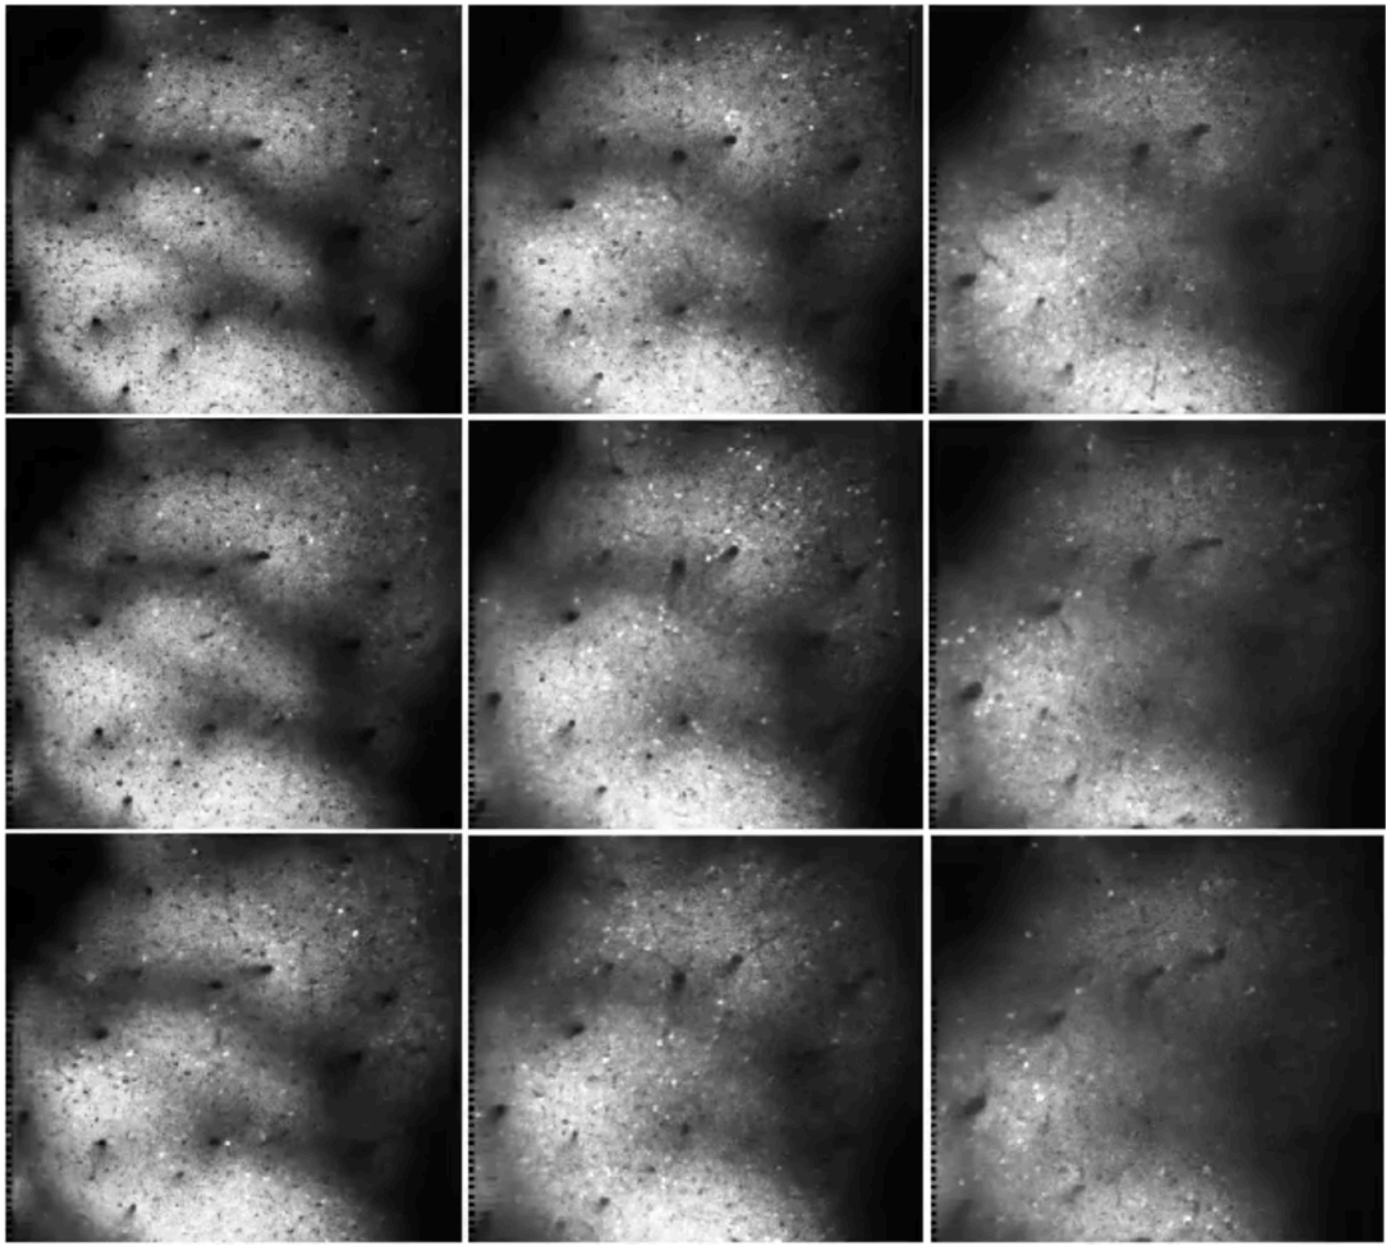

(from Stringer et al., 2019)

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

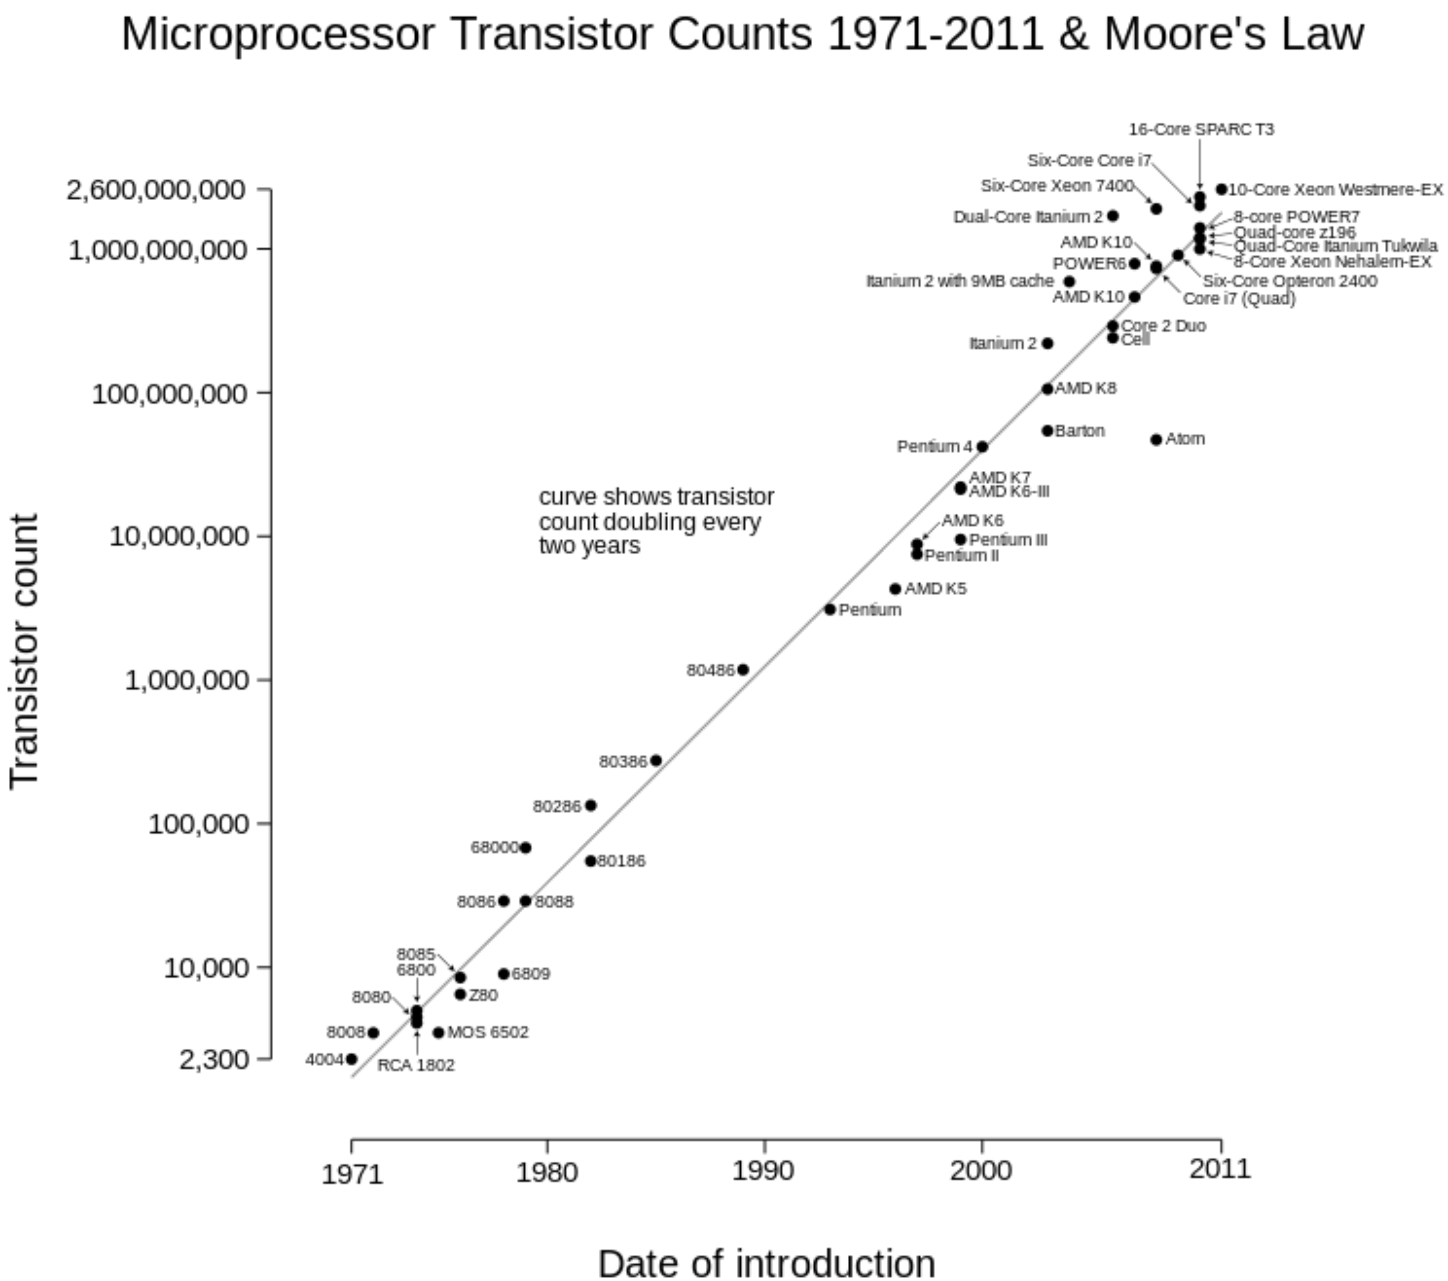

After >**50 years** of concerted research efforts...

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

• machines are still incapable of solving simple perceptual and 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

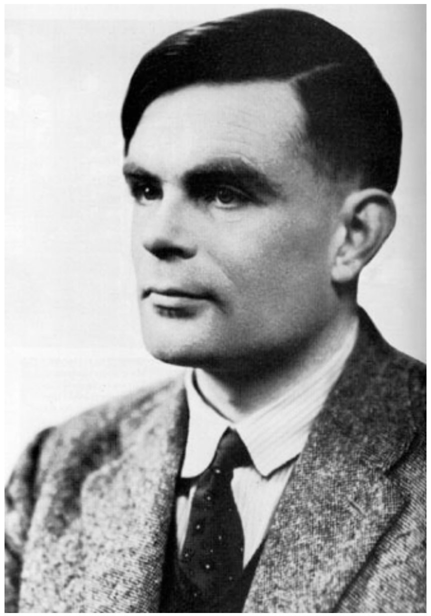

Alan Turing

John von Neumann

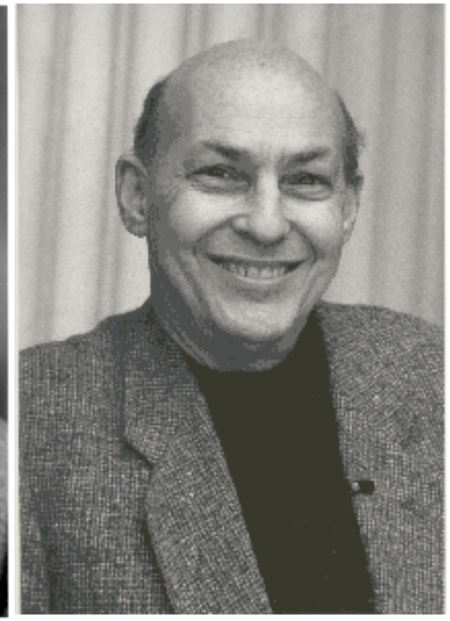

Marvin Minsky

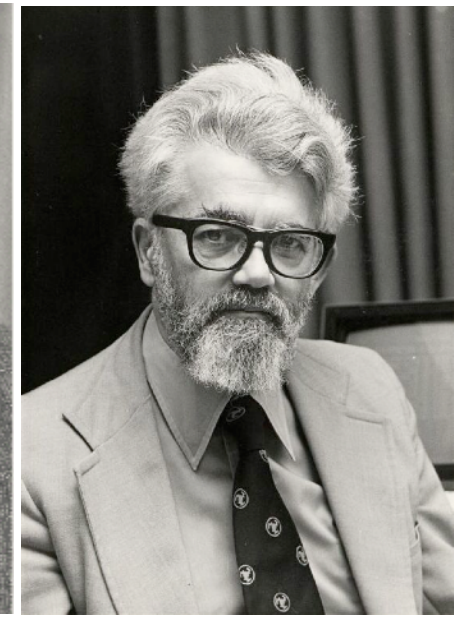

John McCarthy

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

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.

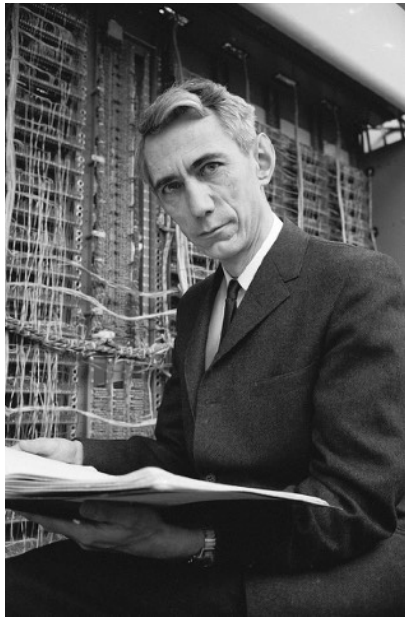

I confidently expect that within a matter of 10 or 15 years, something will emerge from the laboratory which is not too far from the robot of science fiction fame.
— Claude Shannon, 1961

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

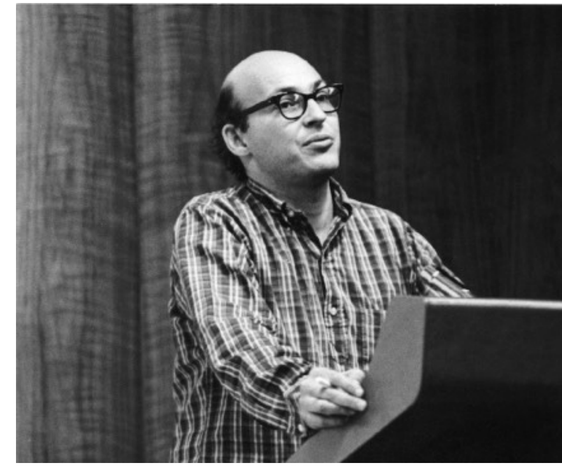

Within a generation...the problem of creating 'artificial intelligence' will be substantially solved.
— Marvin Minsky, 1967

The Lighthill debate (1973)

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

Sir Ja

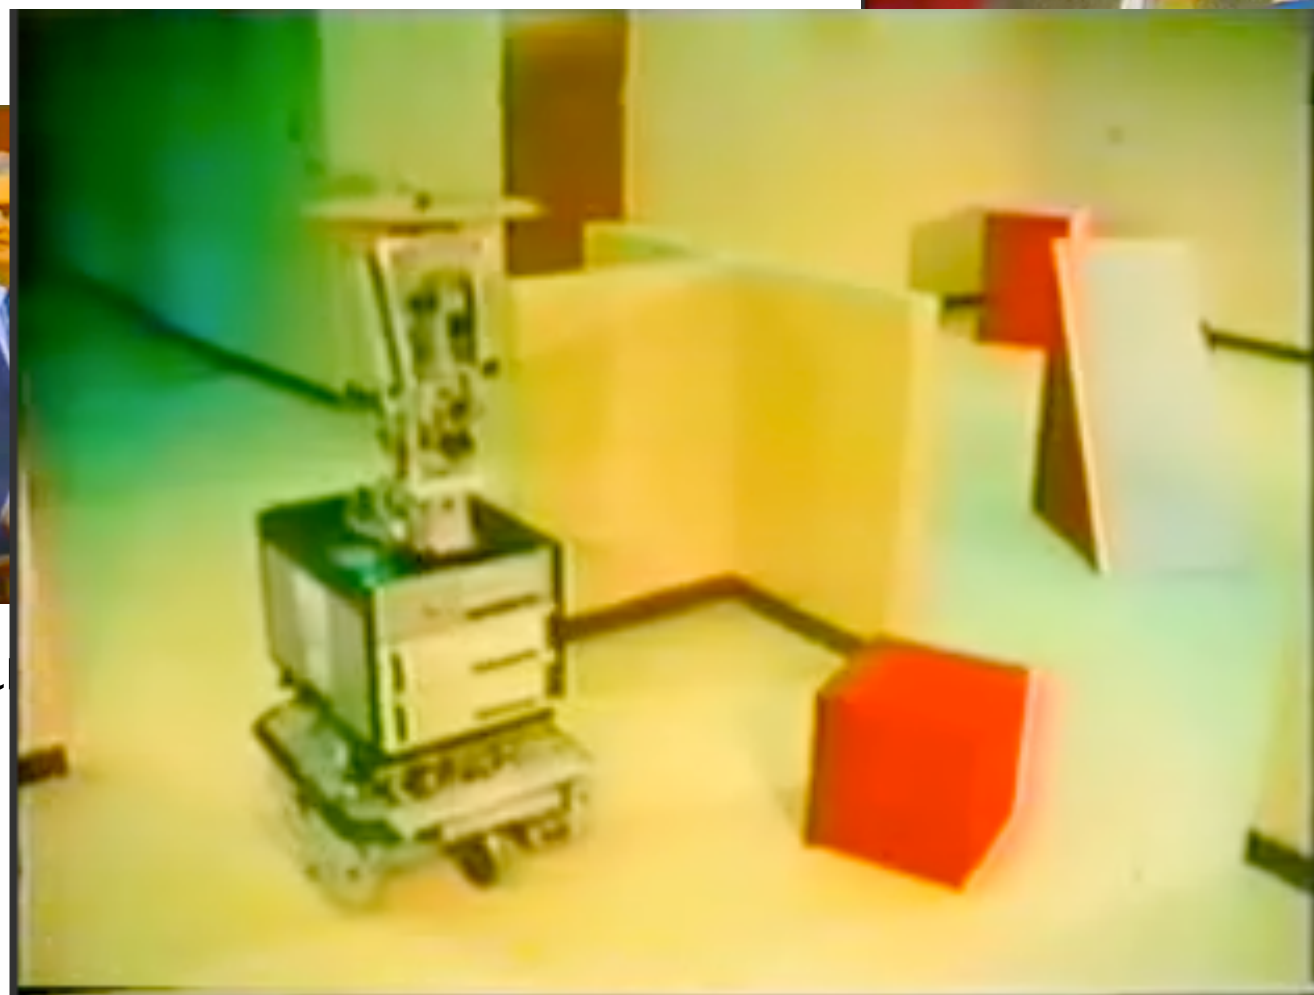

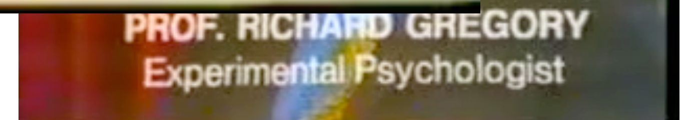

Our first foray into Artificial Intelligence was a program that did a credible job of solving problems in college calculus. Armed with that success, we tackled high school algebra; we found, to our surprise, that it was much harder. Attempts at grade school arithmetic, involving the concept of numbers, 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

Cybernetics/neural networks

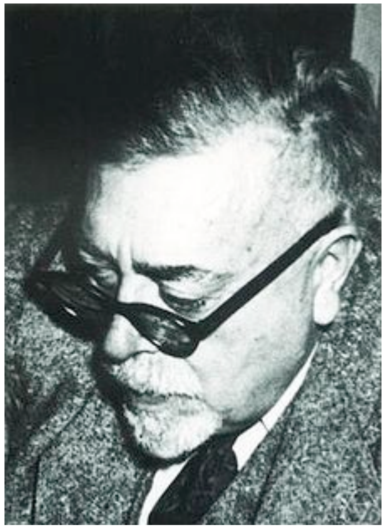

Norbert Wiener

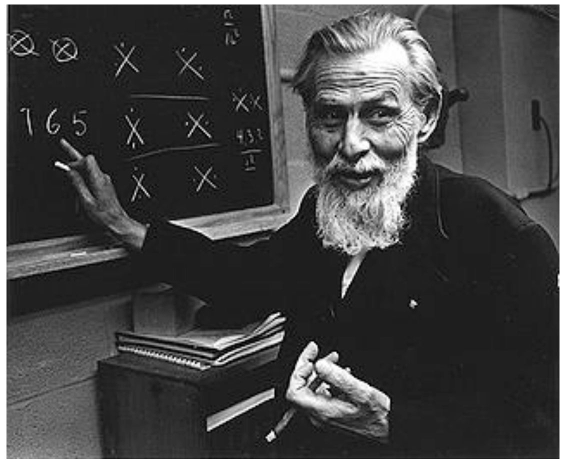

Warren McCulloch & Walter Pitts

Frank Rosenblatt

“The implicit assumption [of the symbol manipulating research program] is that it is relatively easy to specify the behavior that we want the system to perform, and that the challenge is then to design a device or mechanism which will effectively carry out this behavior. . . . It is both easier and more profitable to axiomatize the *physical system* and then investigate this system analytically to determine its behavior, than to axiomatize the *behavior* and then design a physical system by techniques of logical synthesis.”

-- Frank Rosenblatt, 1962

Perceptron model
(Rosenblatt, ca. 1960)

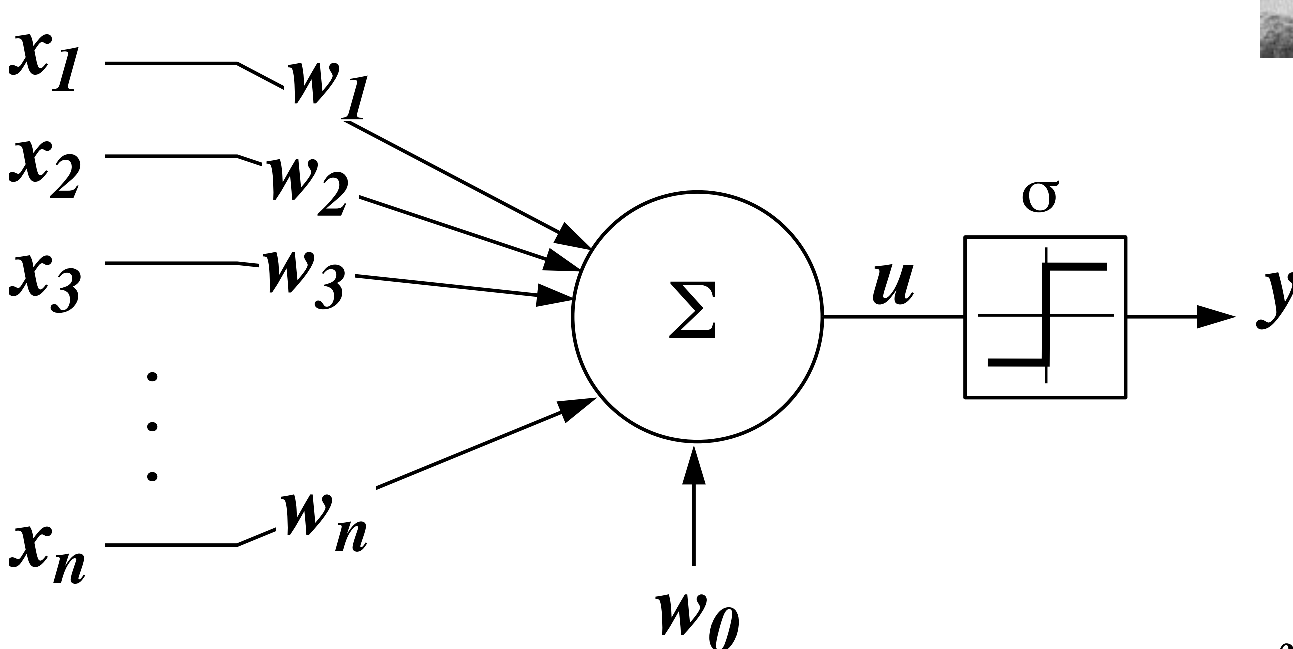

$$u = w_0 + \sum_{i=1}^{n} w_i \, x_i$$

$$y = \sigma(u)$$

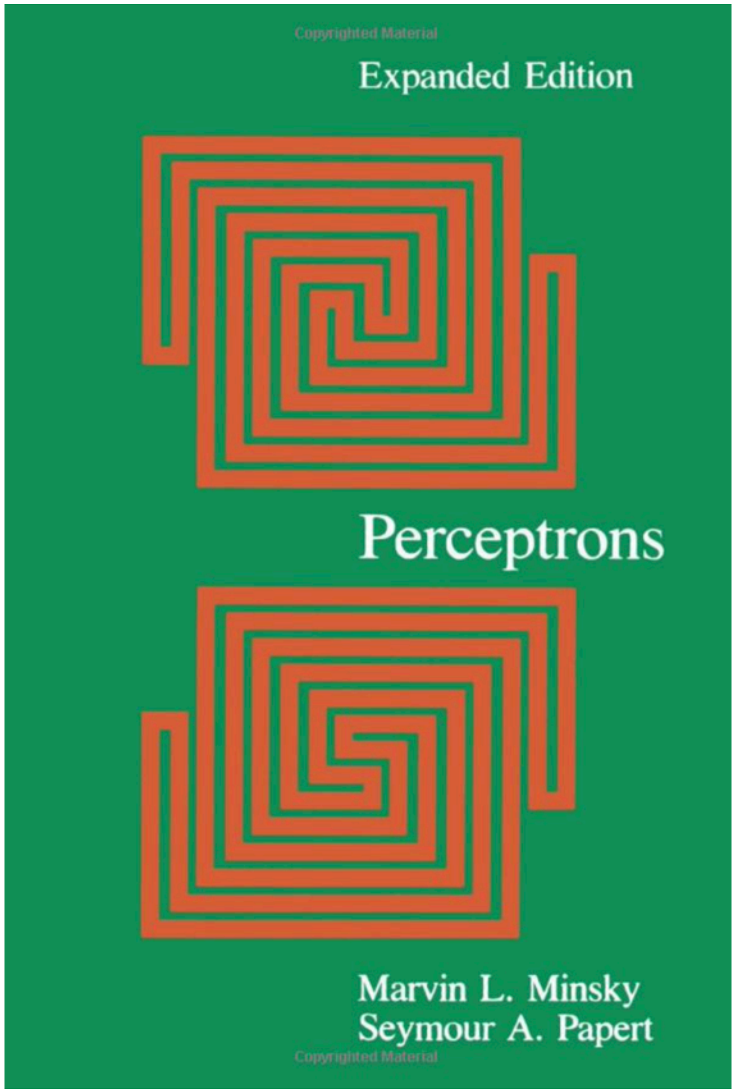

MIT Press 1969

Perceptrons have been widely publicized as "pattern recognition" or "learning" machines and as such have been discussed in a large number of books, journal articles, and voluminous "reports." Most of this writing … is without scientific value.

Both of the present authors (first independently and later together) became involved with a somewhat therapeutic compulsion: to dispel what we feared to be the first shadows of a "holistic" or "Gestalt" misconception that would threaten to haunt the fields of engineering and artificial intelligence as it had earlier haunted biology and psychology.

Well, we have considered Gamba machines, which could be described as "two layers of perceptron." We have not found (by thinking or by studying the literature) any other really interesting class of multilayered machine, at least none whose principles seem to have a significant relation to those of the perceptron…. We consider it to be an important research problem to elucidate (or reject) our intuitive judgment that the extension is sterile.

Hubel & Wiesel (1962, 1965)

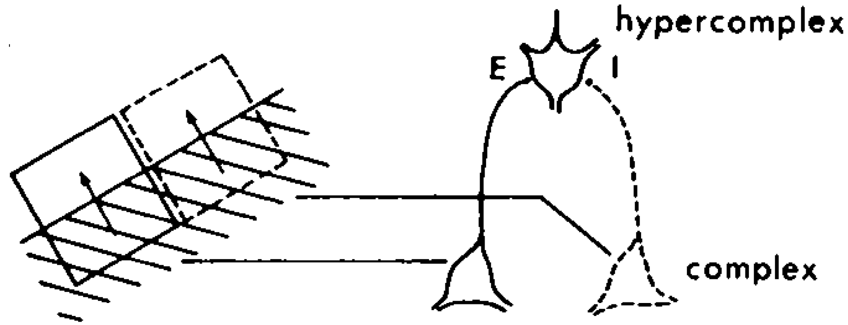

Hypercomplex

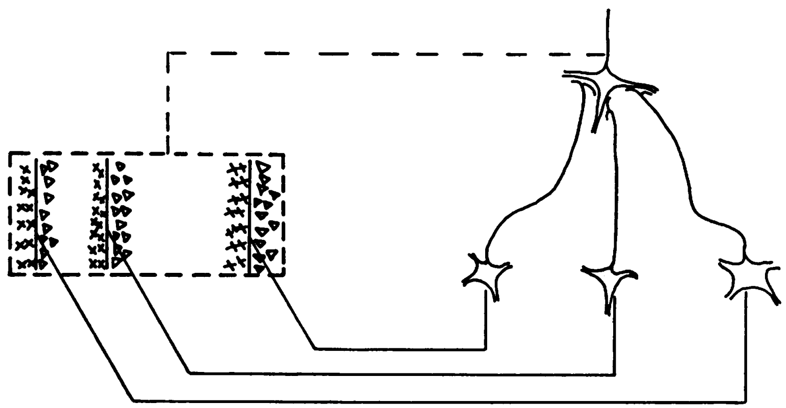

Complex

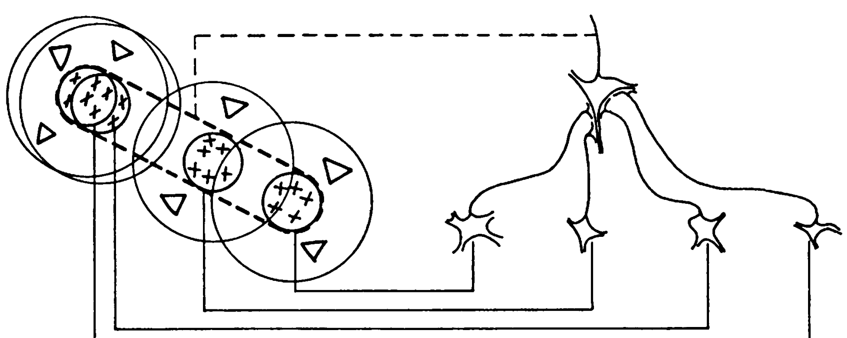

Simple

Neocognitron

(Fukushima 1980)

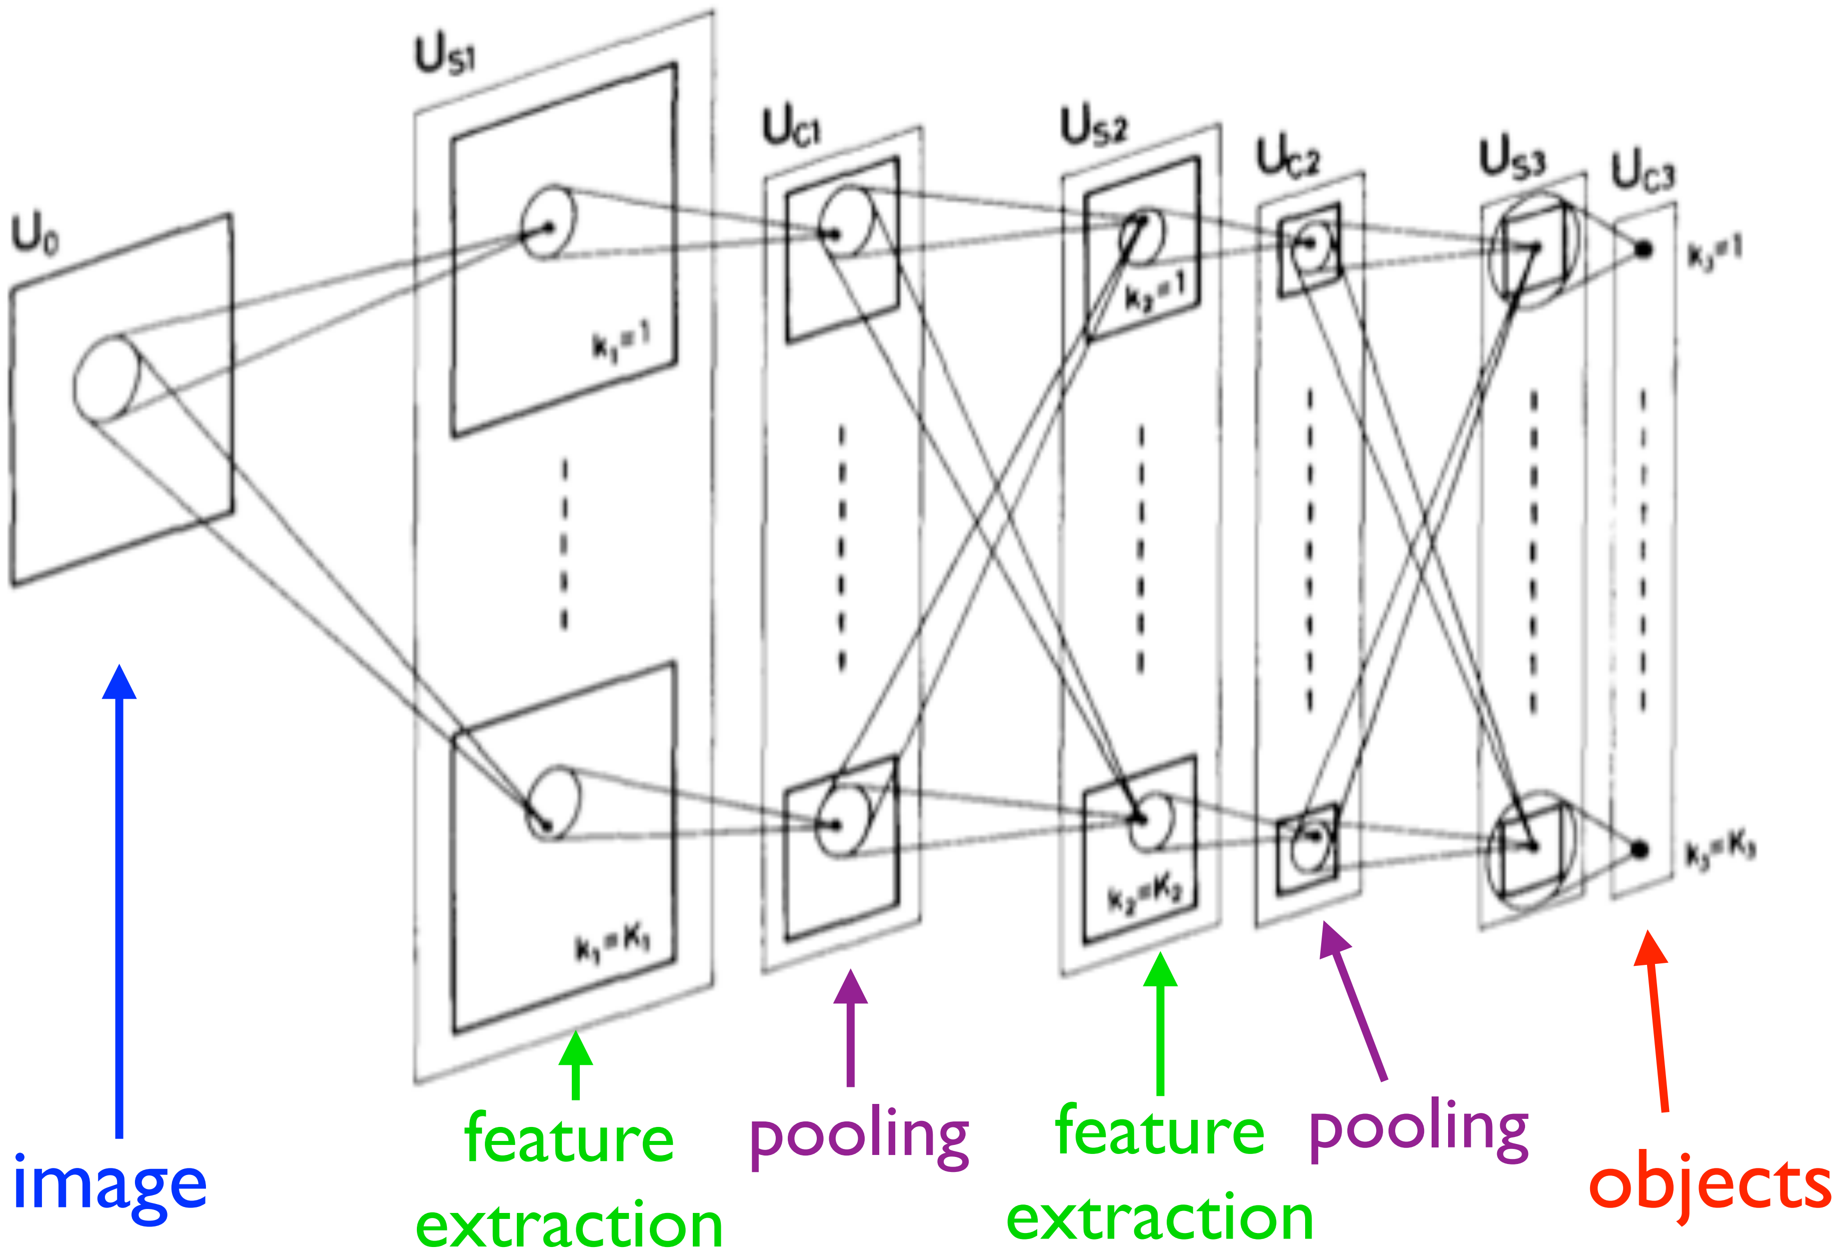

Nervous systems are difficult to observe and manipulate

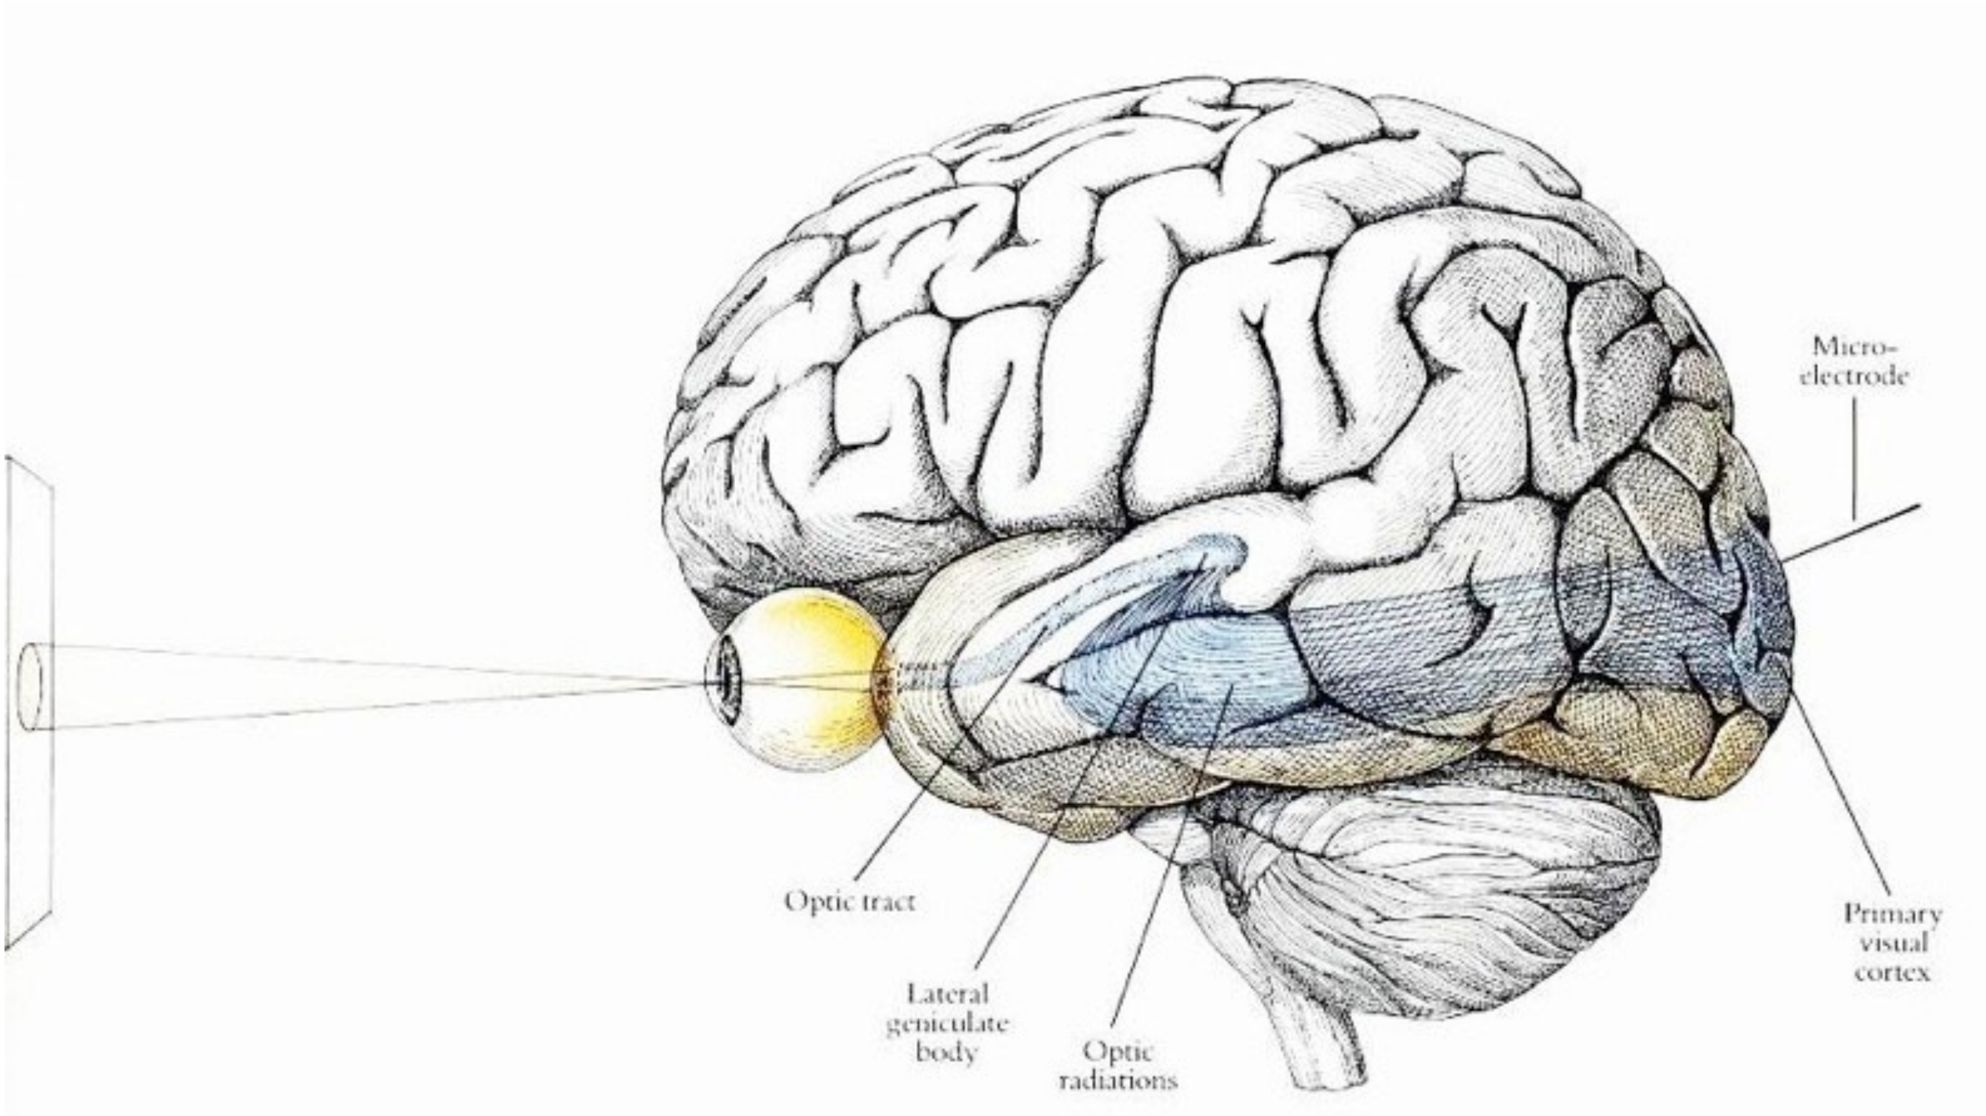

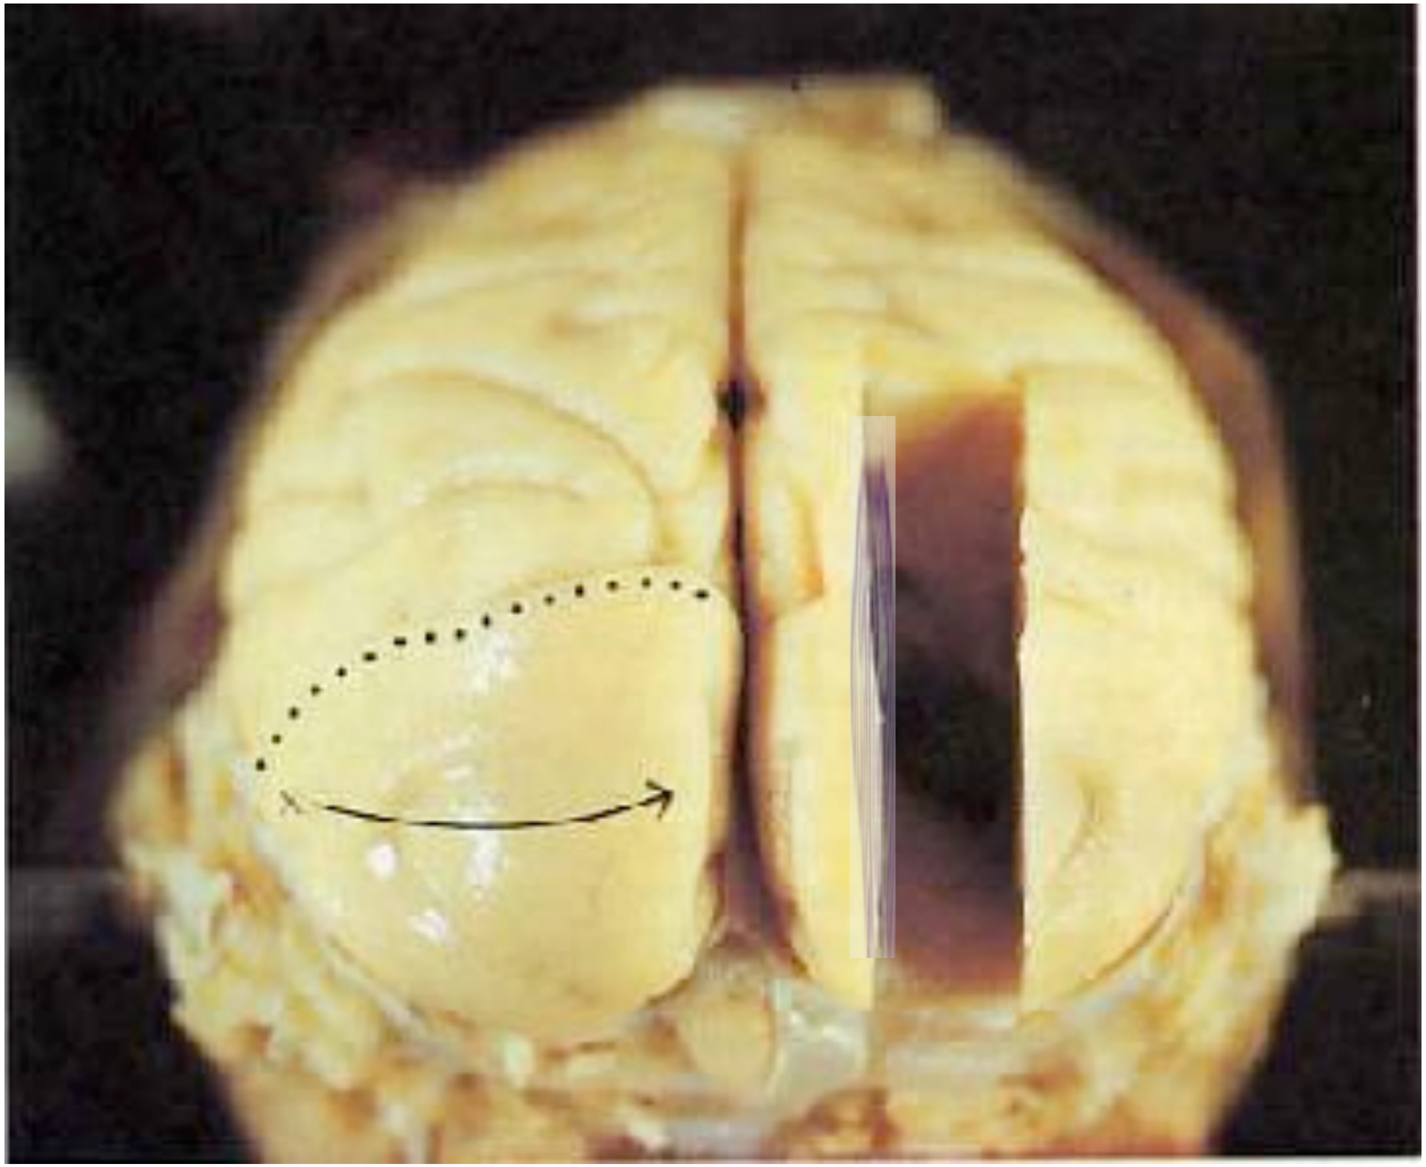

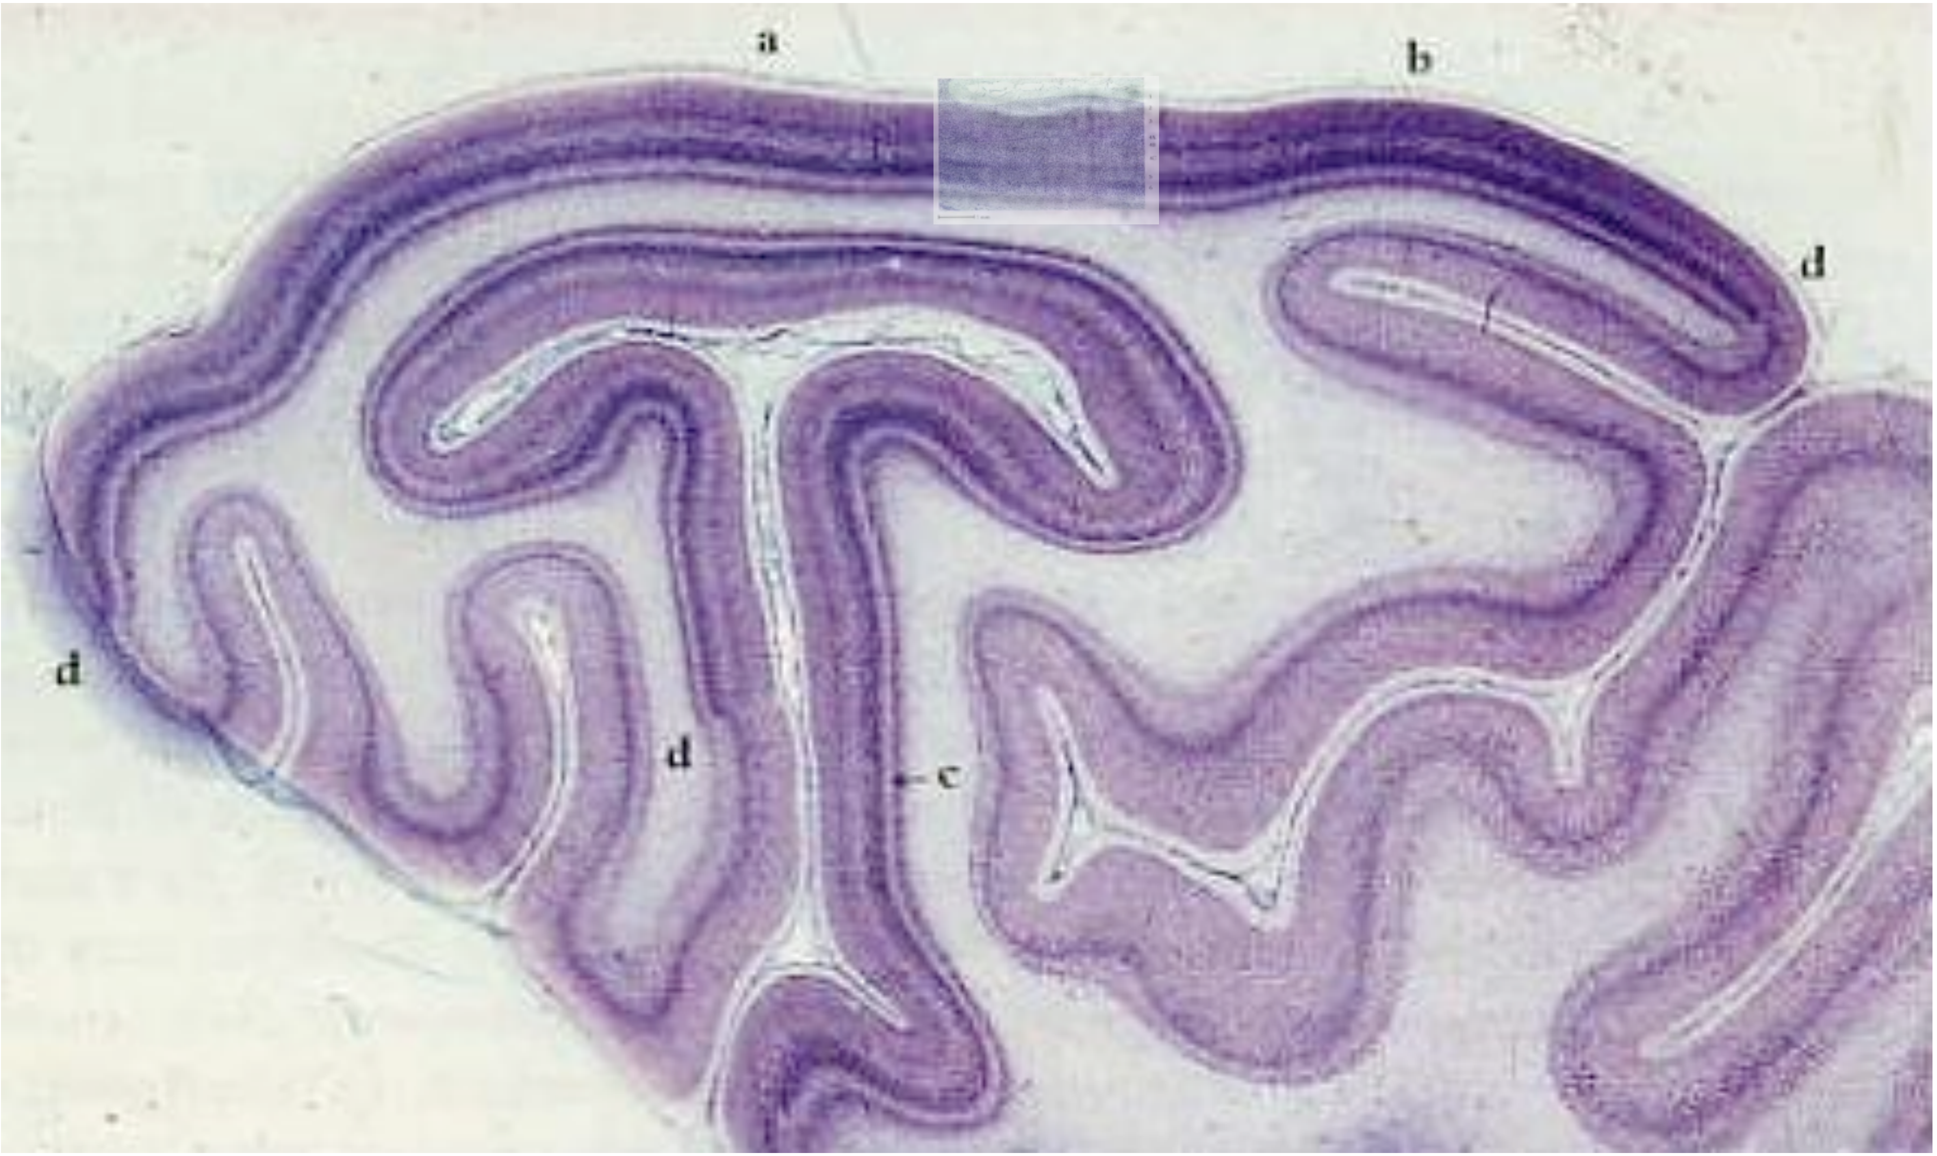

a
b
d
d
d
c

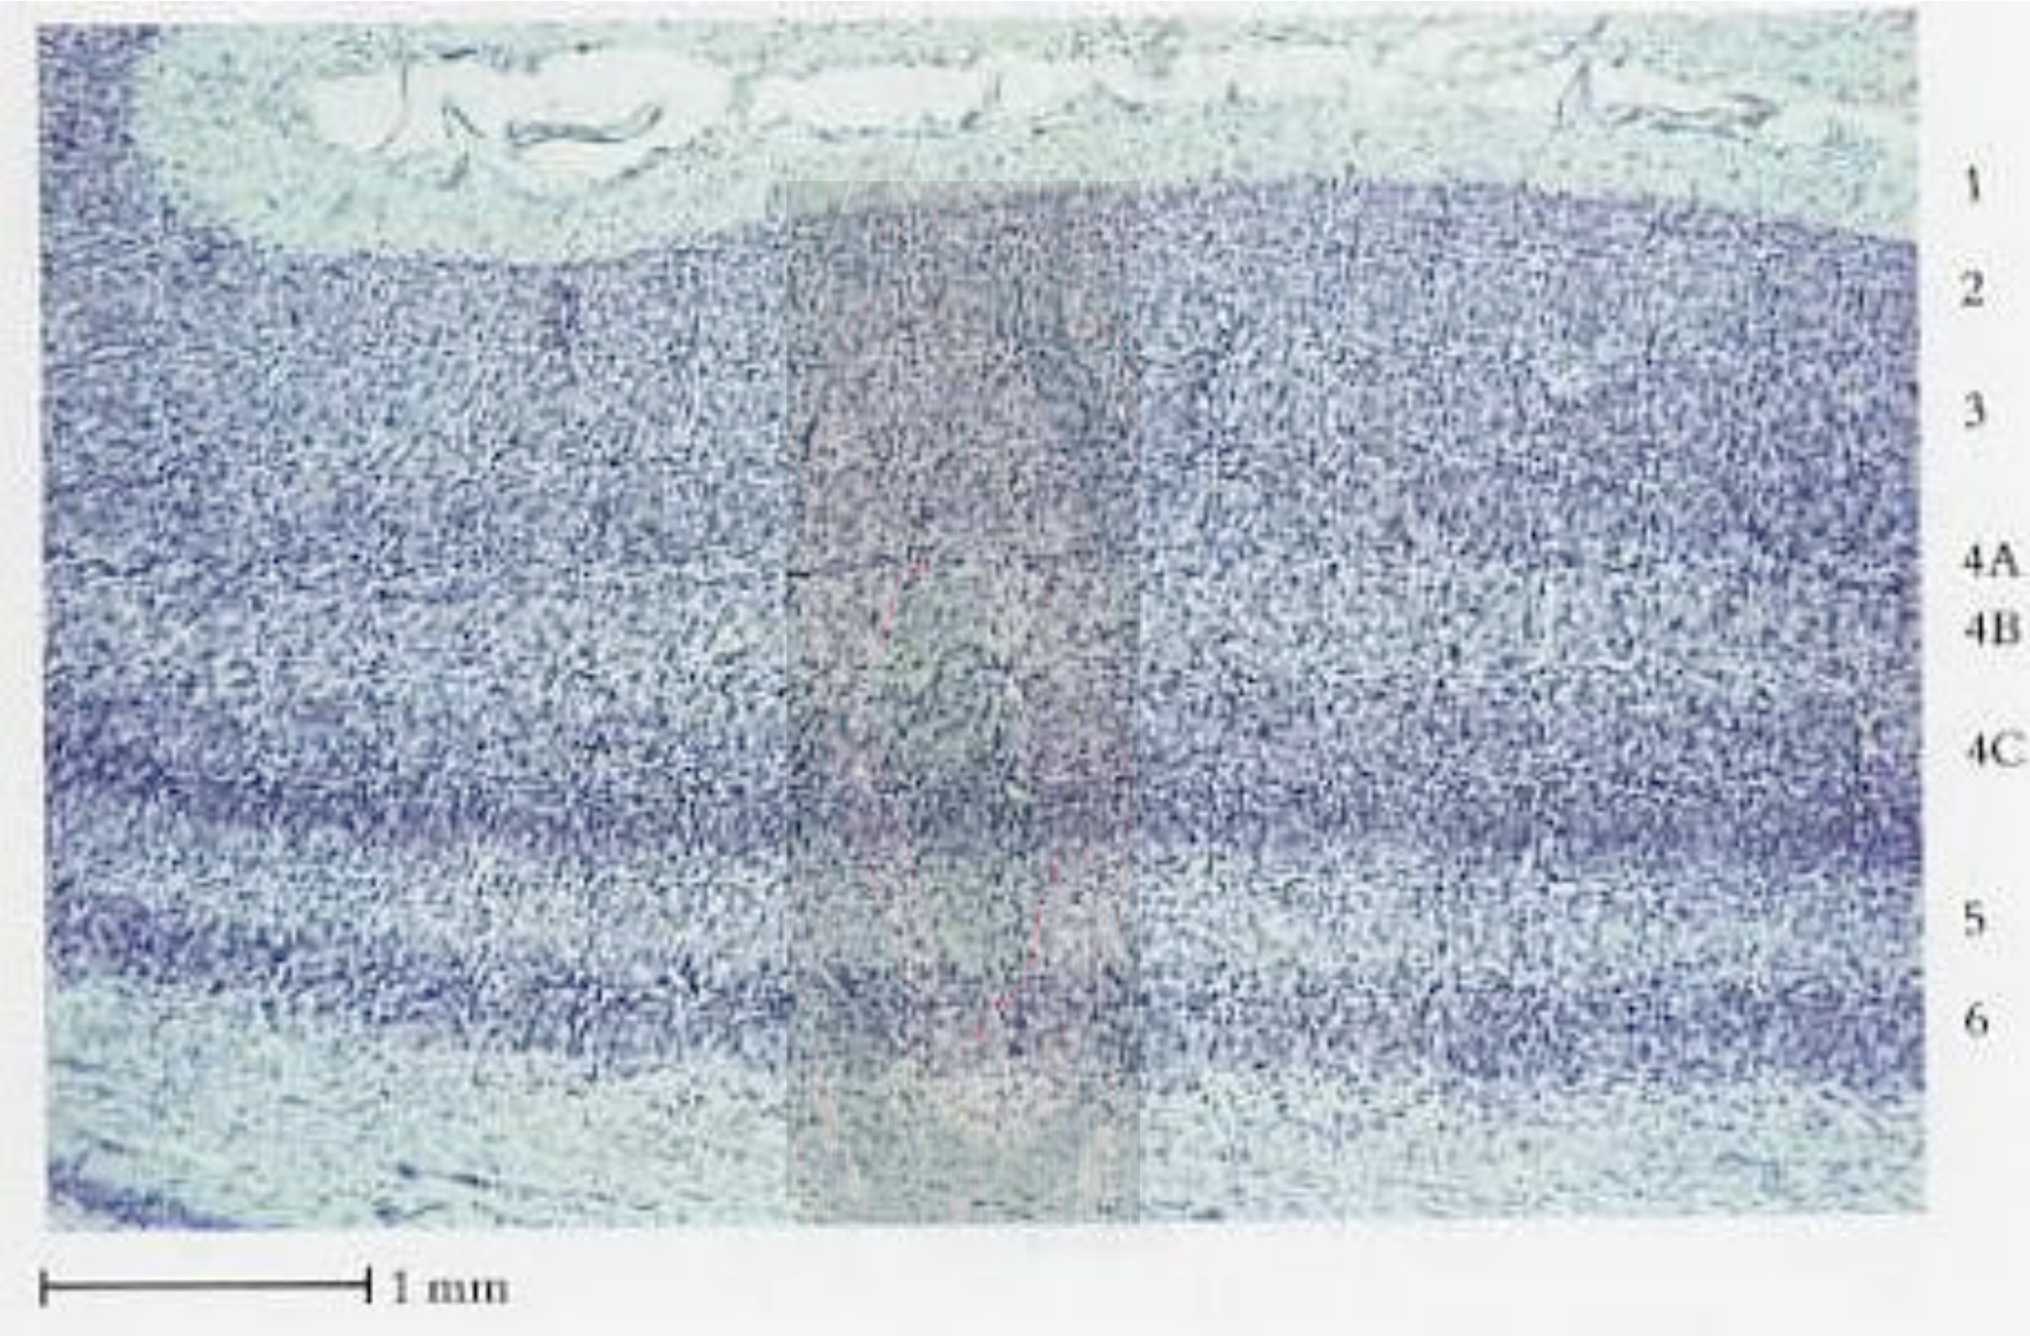
1
2
3
4A
4B
4C
5
6
1 mm

1 mm^2 of cortex contains 100,000 neurons

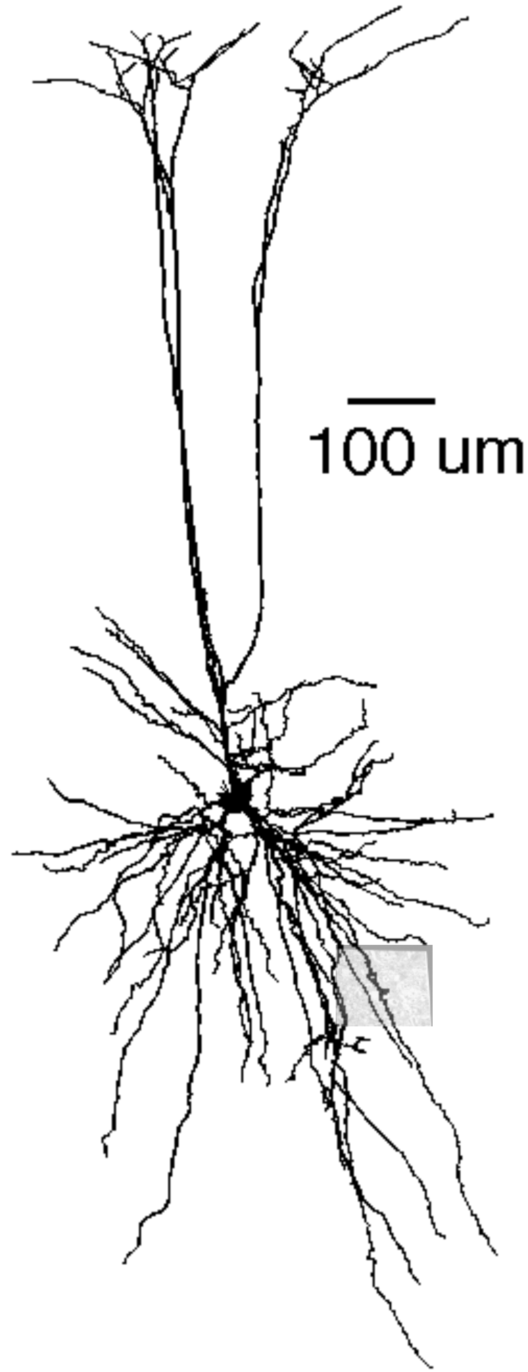
100 um

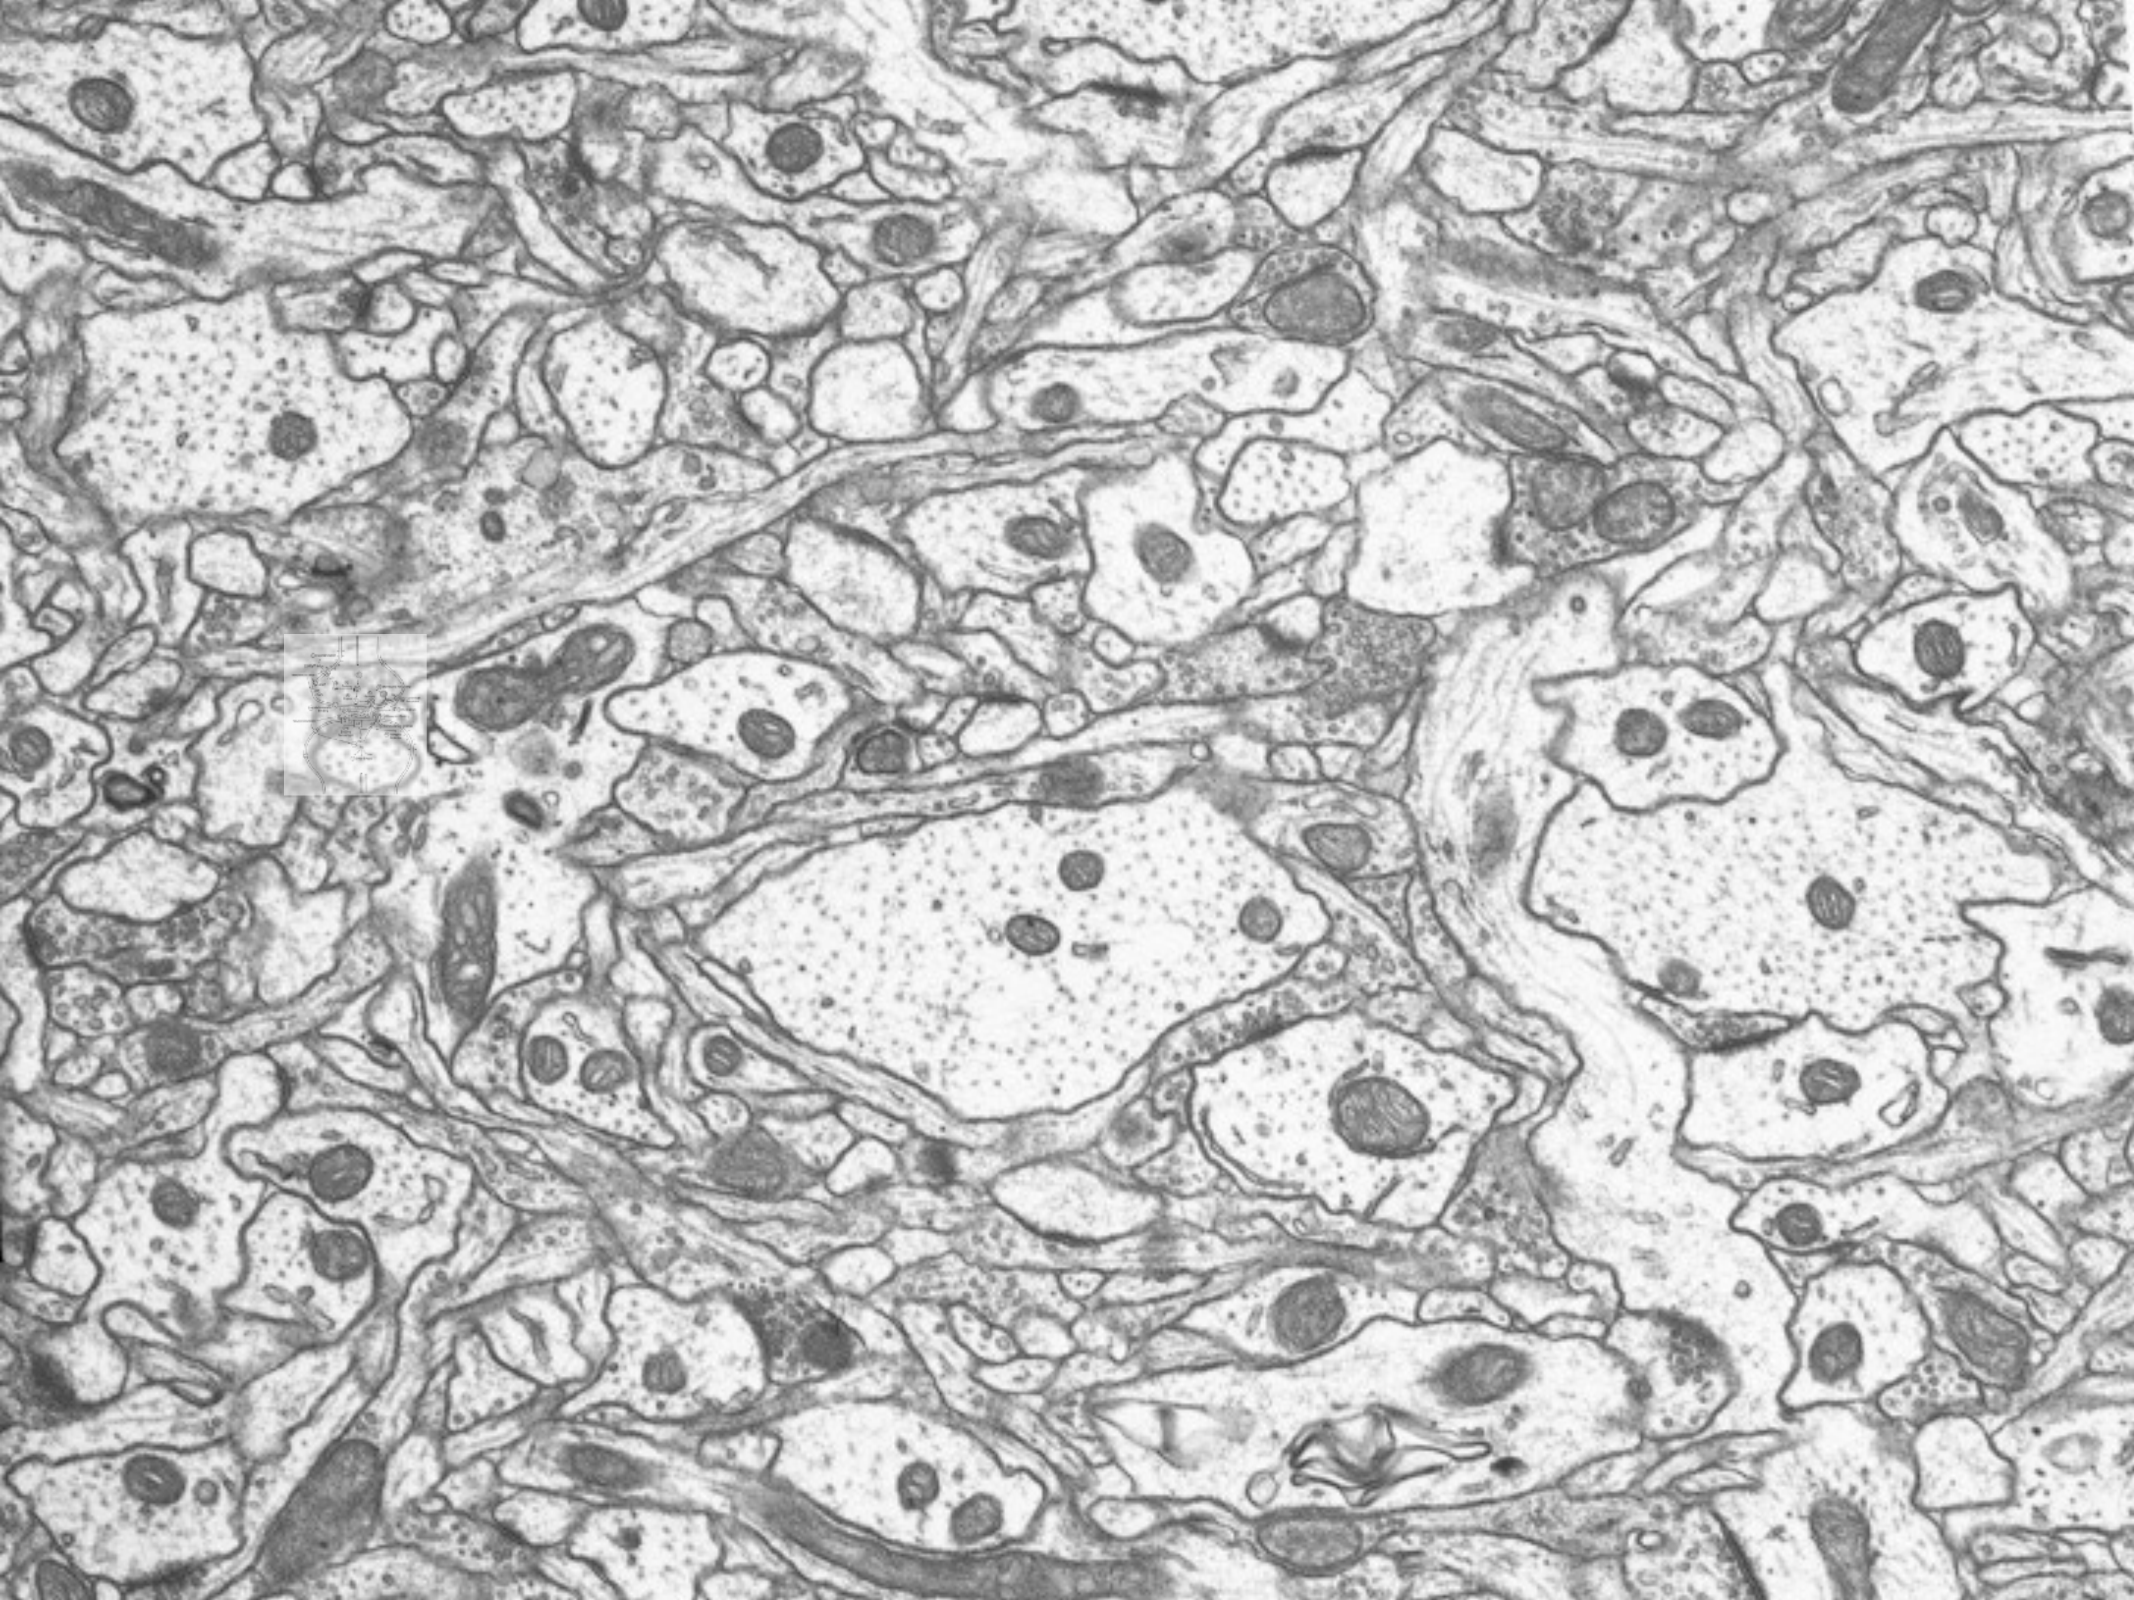

Anatomy of a synapse

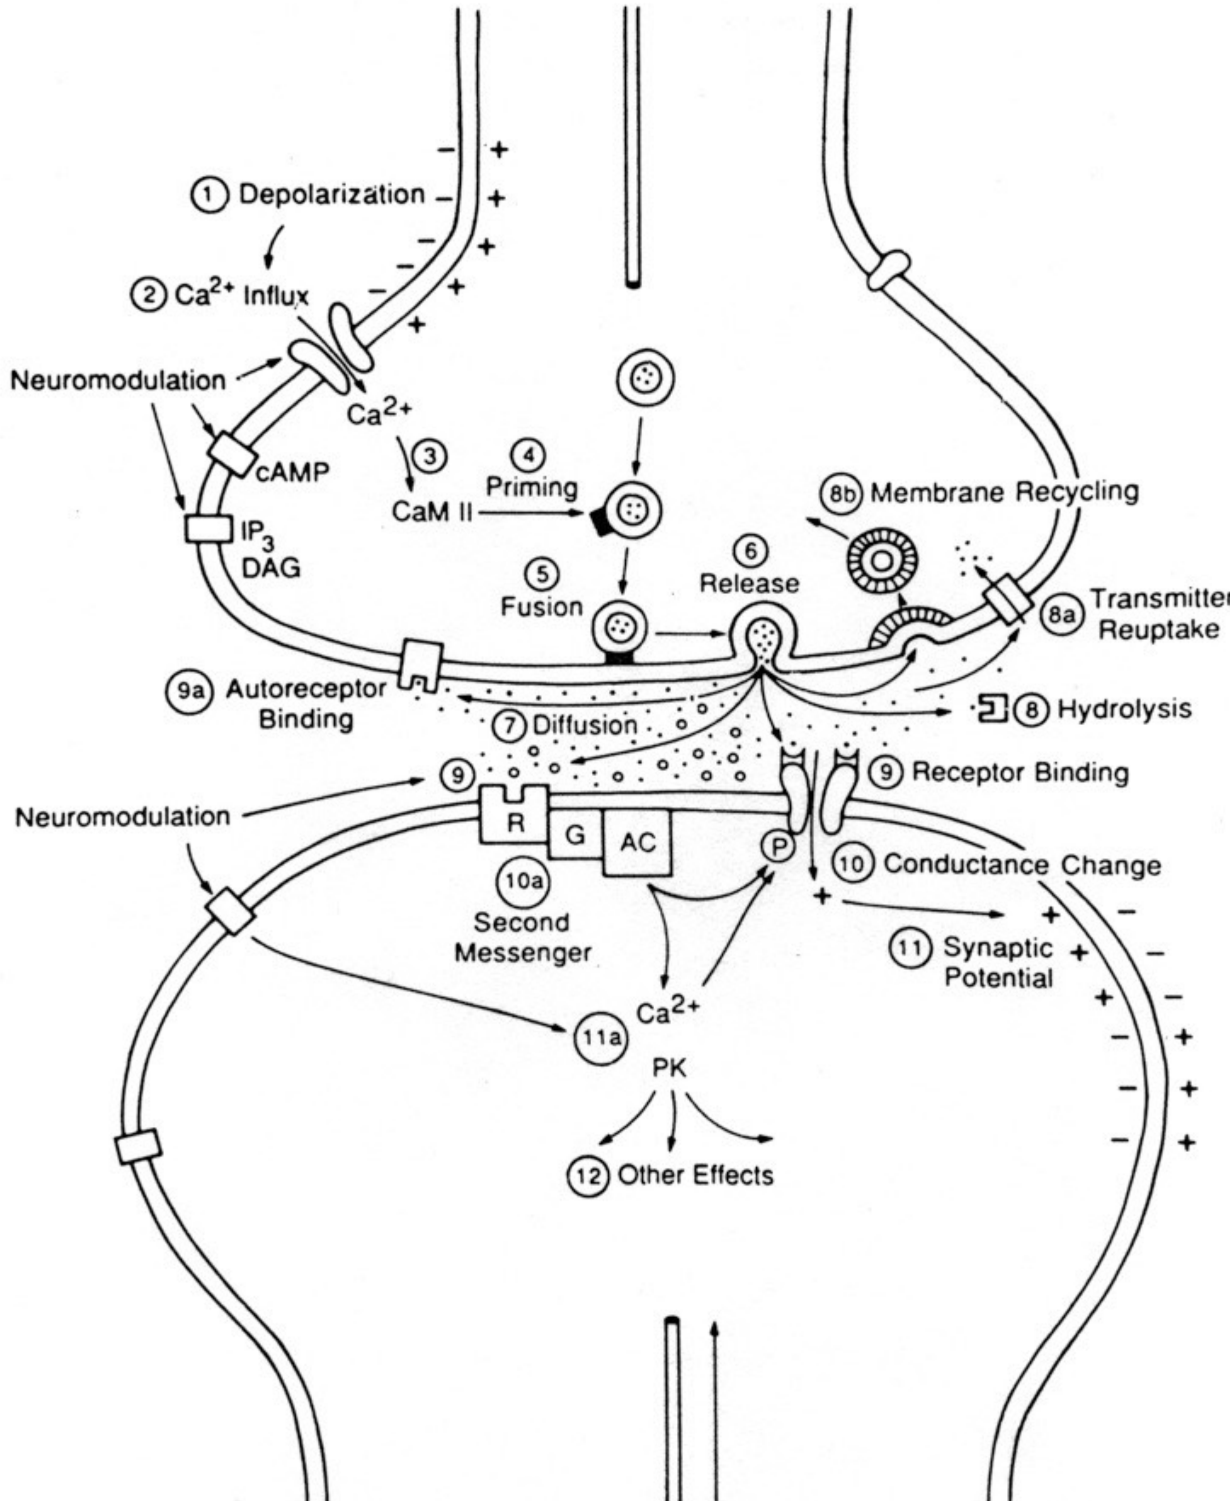

Are there principles?

"God is a hacker"
– Francis Crick

"A wing would be a most mystifying structure if one did not know that birds flew….."
– H.B. Barlow (1961)

Otto Lilienthal experiments with flight (1890’s)

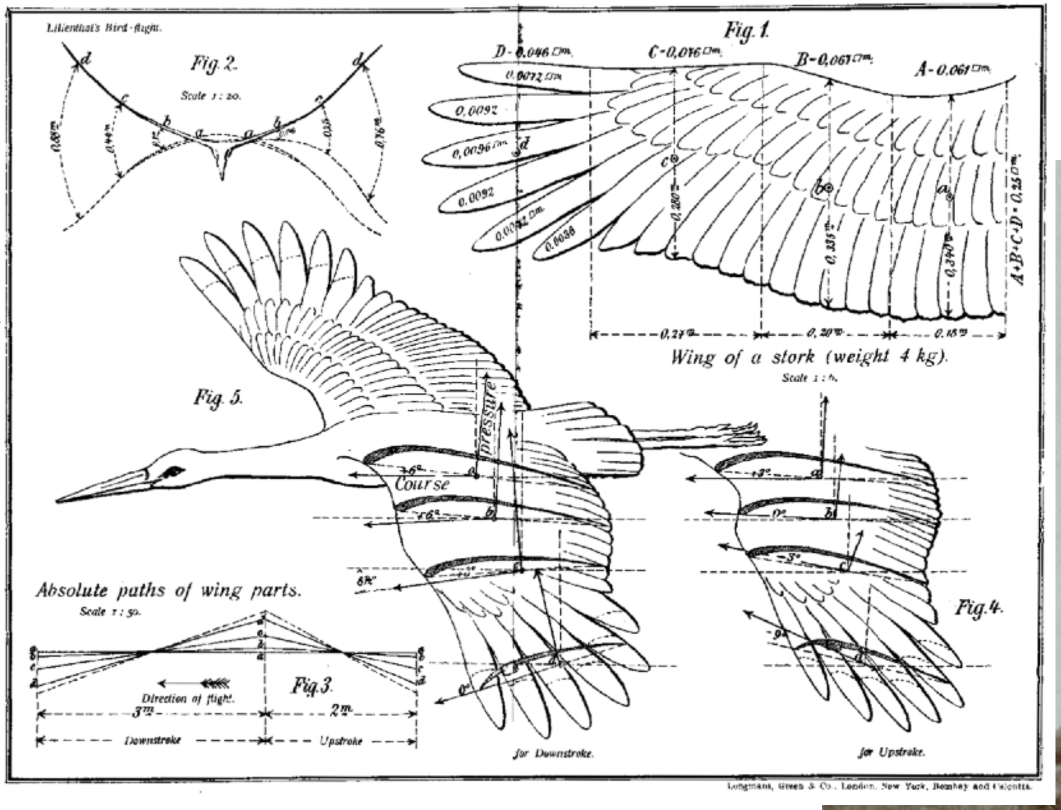

Der Vogelflug als Grundlage der Fliegekunst (1889)

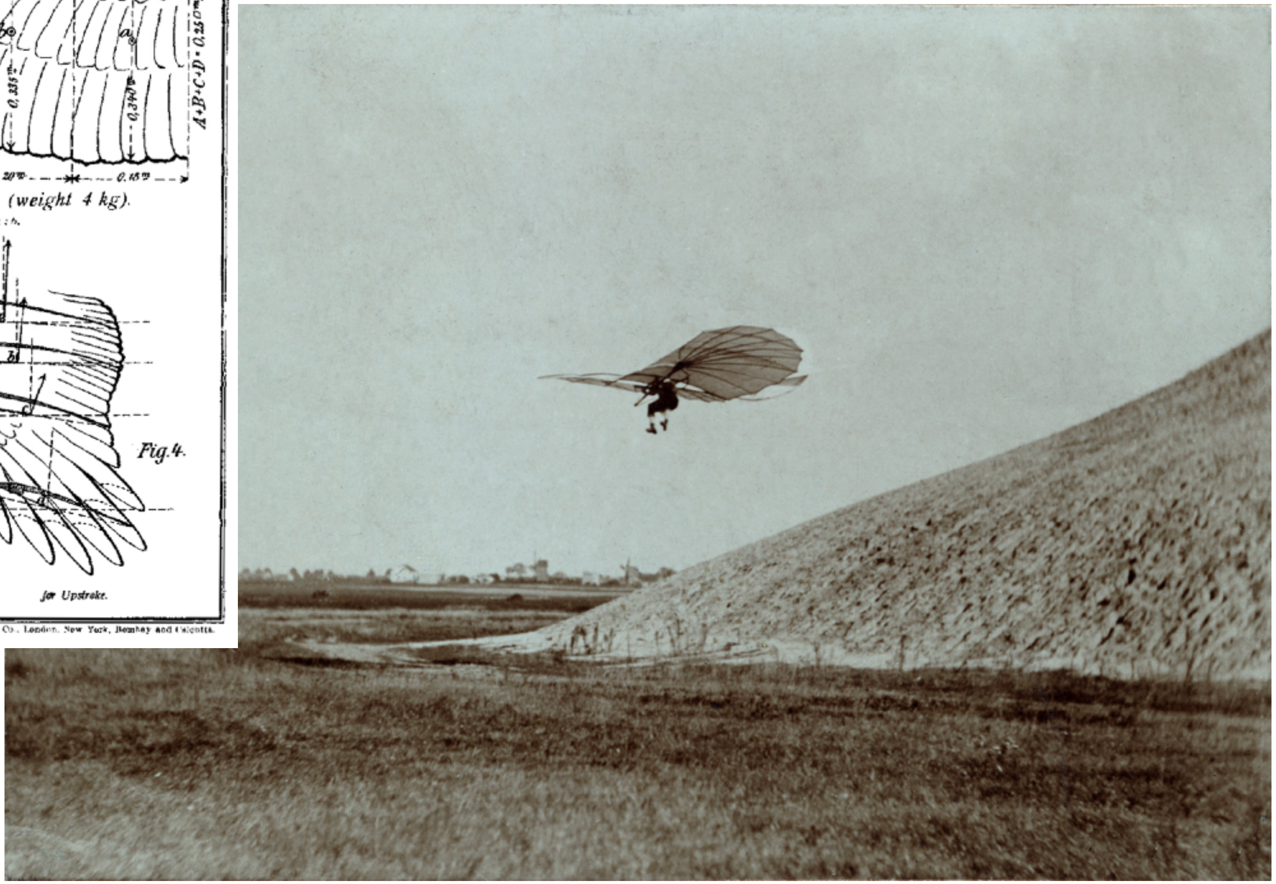

Wright Flyer (1903)

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

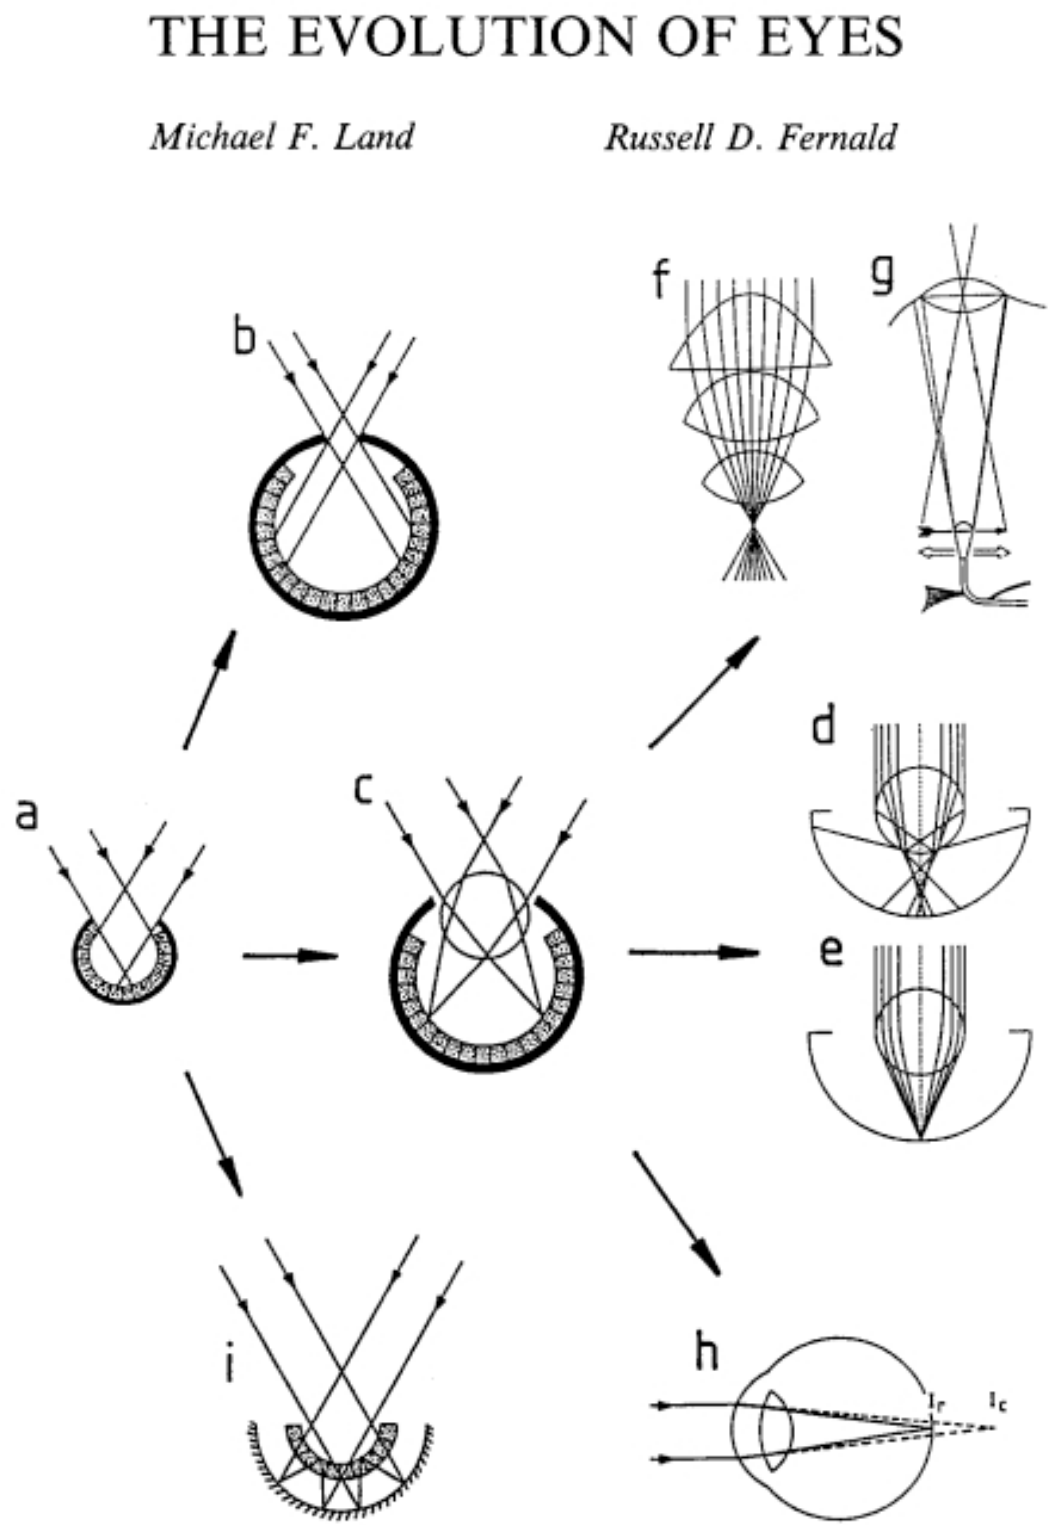

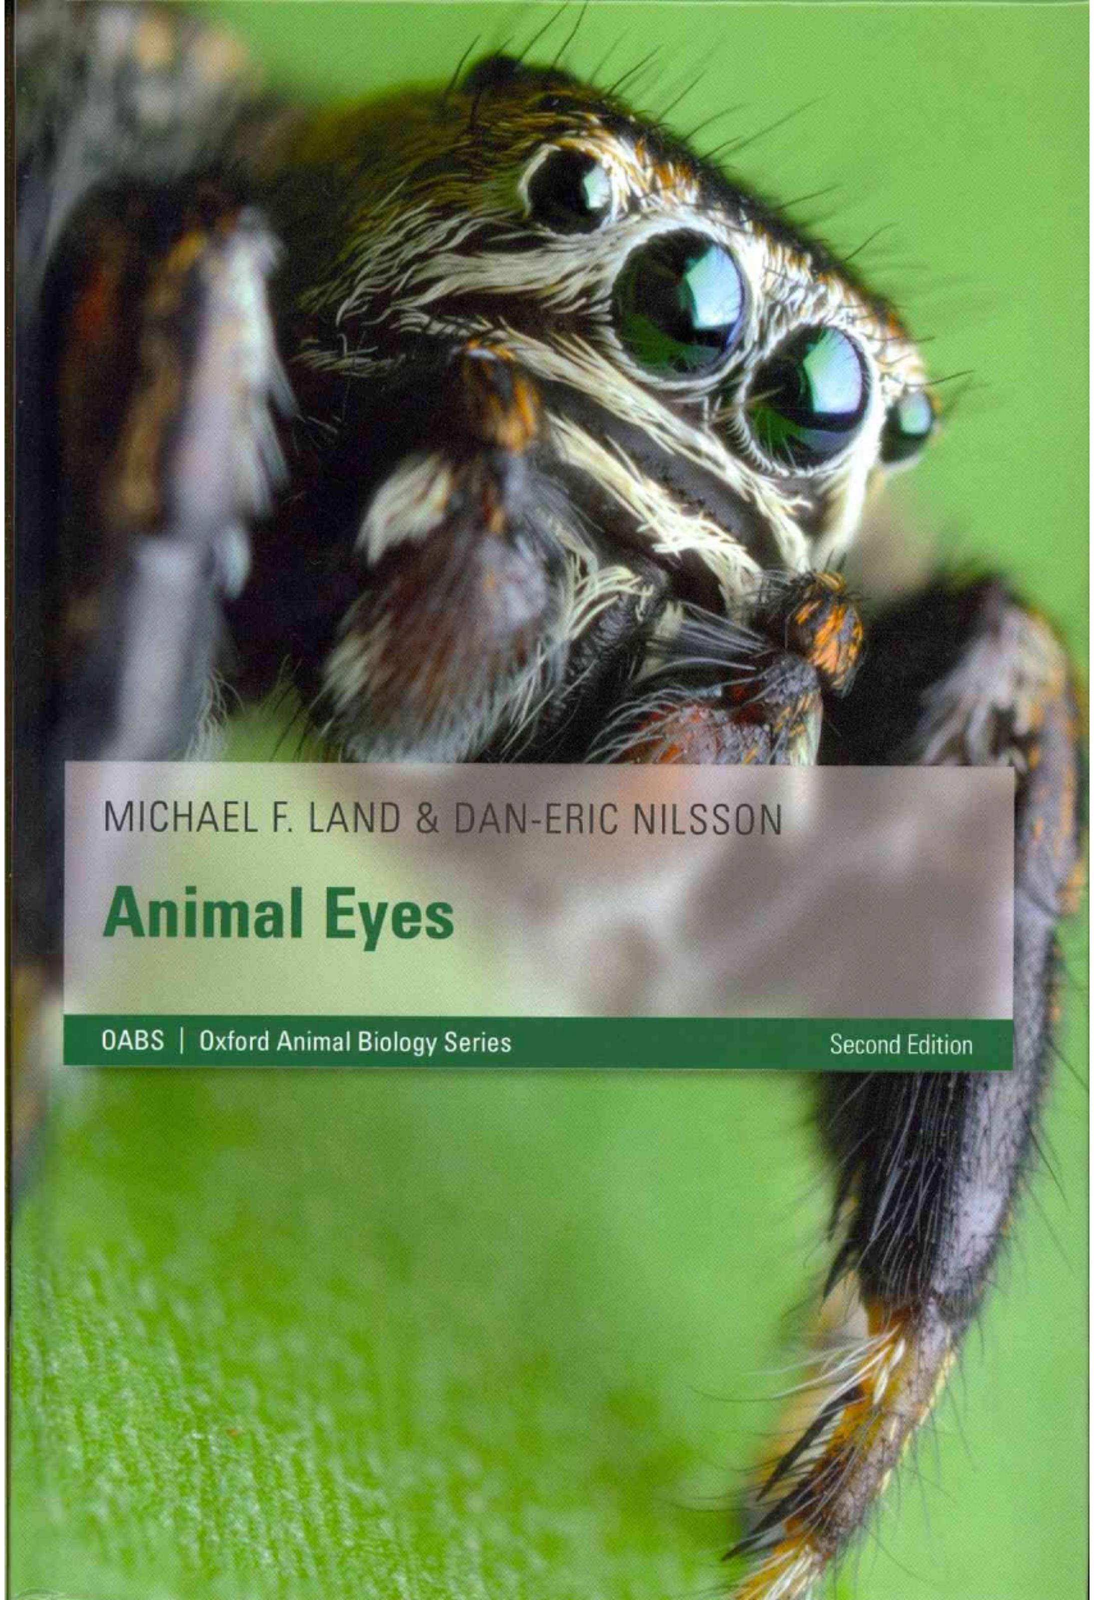

MICHAEL F. LAND & DAN-ERIC NILSSON

Animal Eyes

OABS | Oxford Animal Biology Series

Second Edition

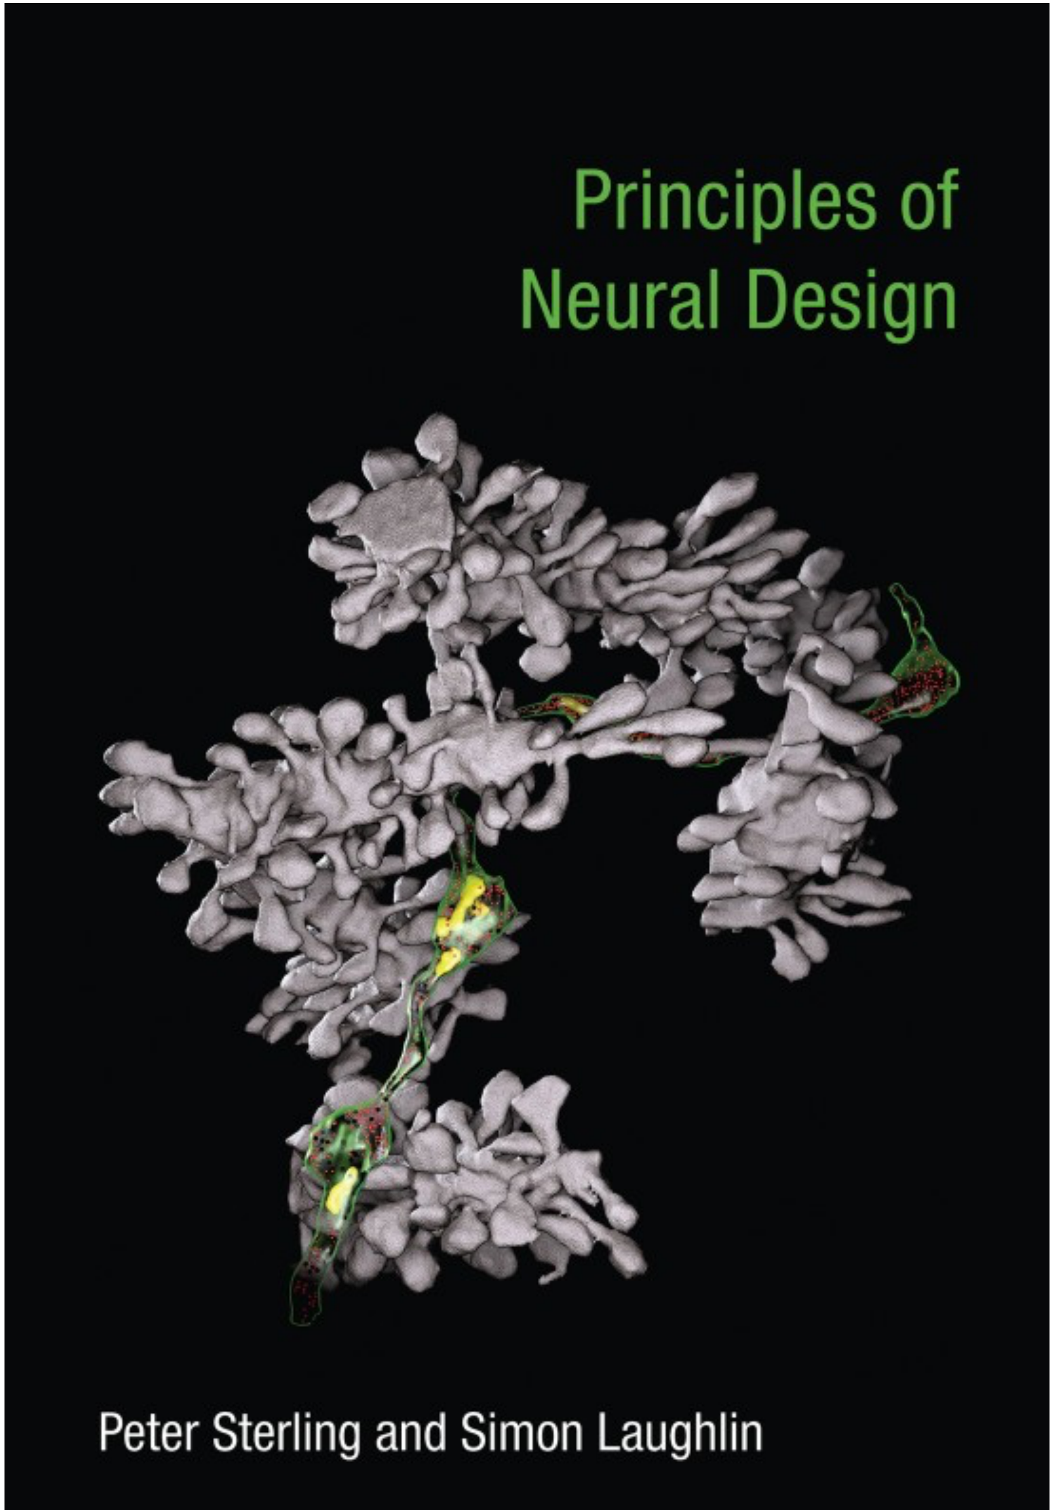

Principles

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

Why AI is Harder Than We Think

Melanie Mitchell

Springs and winters

Four fallacies:

1. Narrow intelligence is on a continuum with general intelligence.
2. Easy things are easy and hard things are hard.
3. The lure of wishful mnemonics.
4. Intelligence is all in the brain.


jumping spider


sand wasp

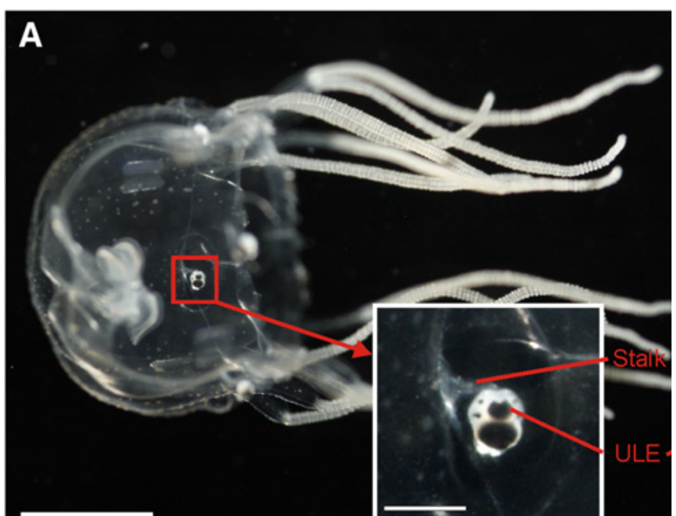

box jellyfish

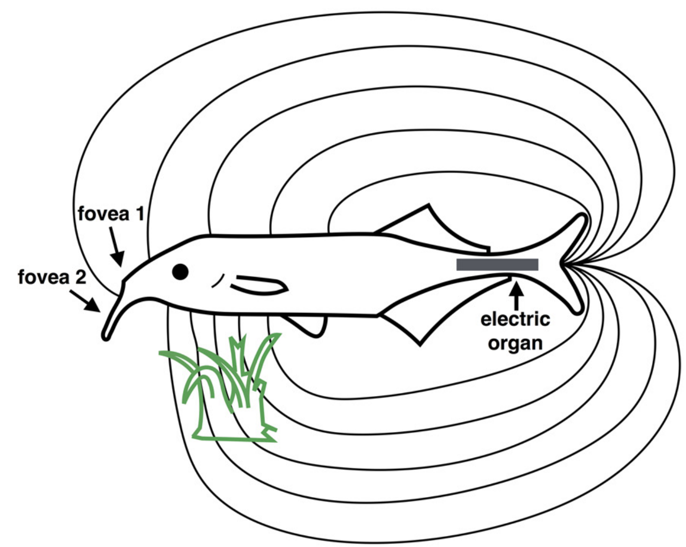

weakly electric fish

NERSC (Lawrence Berkeley Lab) ~ 5 MW

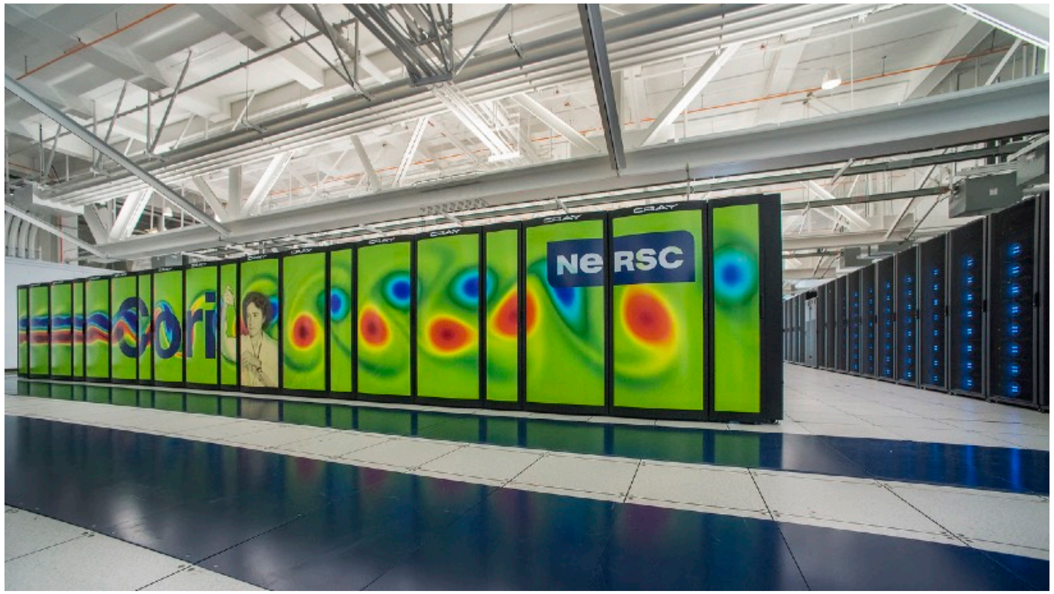

Jumping spider ~ 1 fly/day

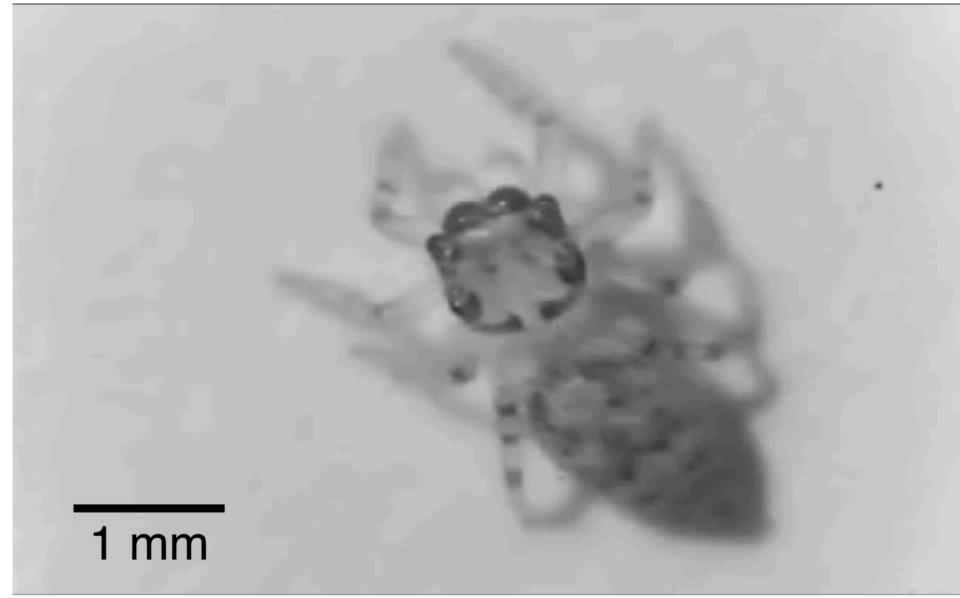

(Bair & Olshausen, 1991)

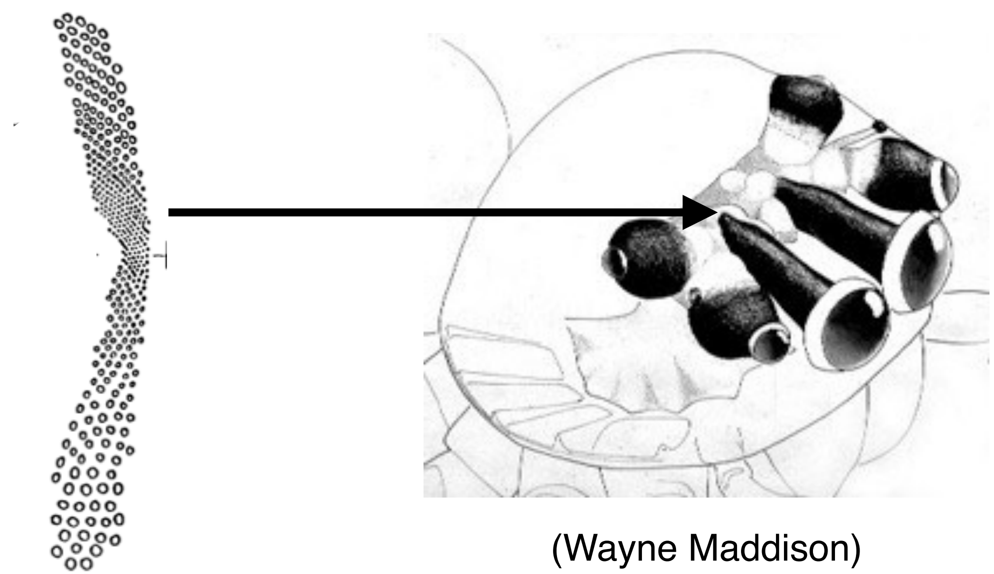

(Wayne Maddison)

…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)

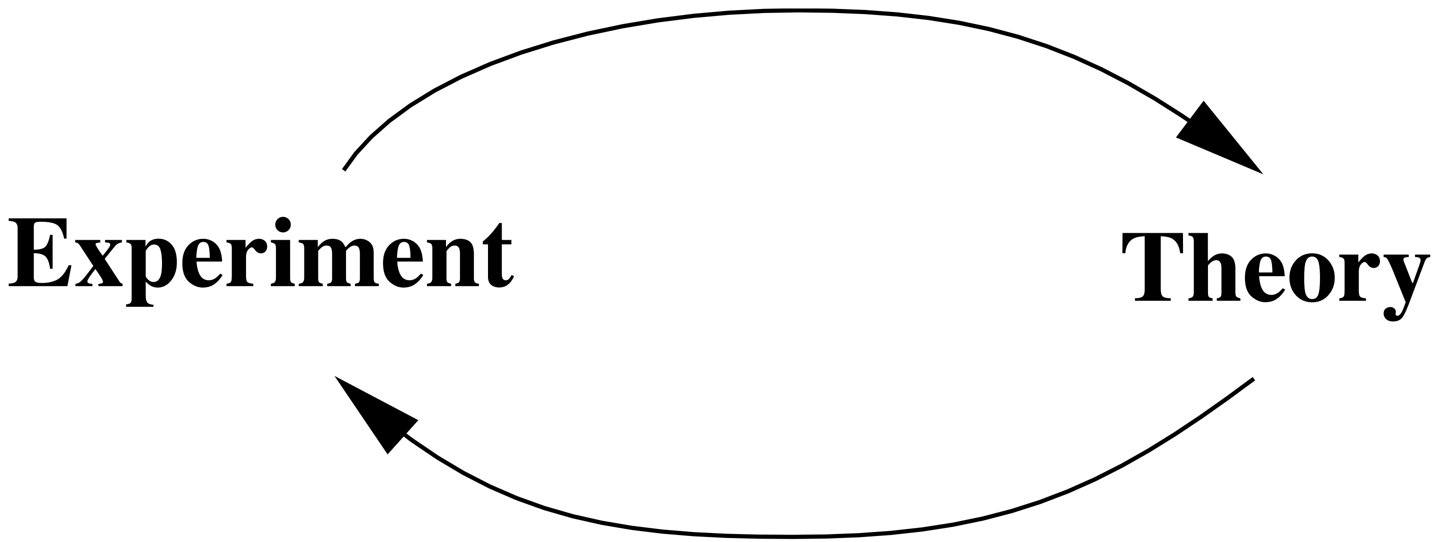
Experiment
Theory

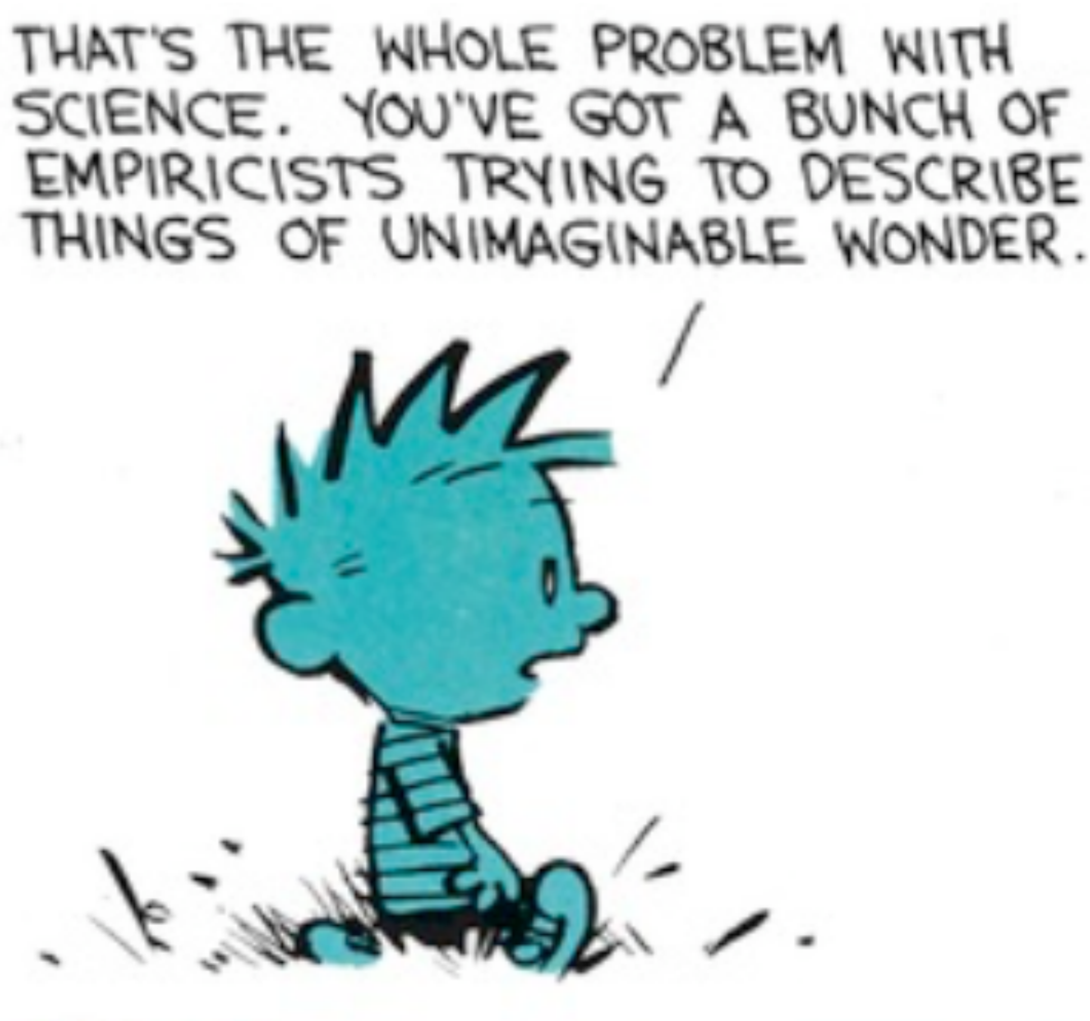
THAT'S THE WHOLE PROBLEM WITH SCIENCE. YOU'VE GOT A BUNCH OF EMPIRICISTS TRYING TO DESCRIBE THINGS OF UNIMAGINABLE WONDER.