UAB

MODELLING FOR SCIENCE AND ENGINEERING

Is there a Science for AI?
Part II

S. Xambó

IMTech & BSC

23/1/2023

https://web.mat.upc.edu/sebastia.xambo/99/s-uab2.pdf

Tribute to Ramón y Cajal

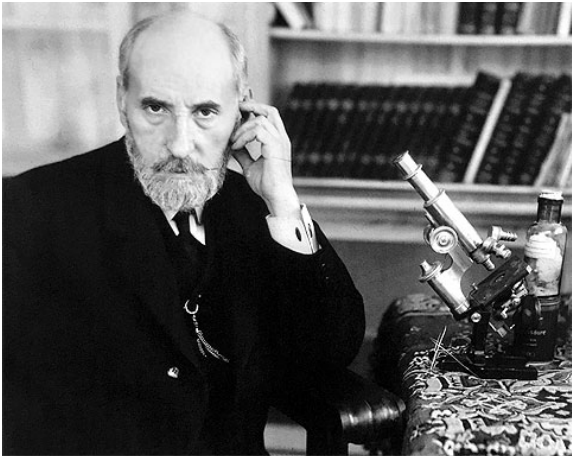


Santiago Ramón y Cajal (1852-1934). Nobel Prize of Physiology or Medicine (1906, shared with Camilo Golgi) for his discoveries about the *structure of the nervous system and the role of the neuron*. The father of modern neuroscience. The (core of) the *Neural Doctrine* was discovered in 1888 at his *home lab* in the Notariado street 7 (Barcelona), 200 m away from the "Reial Acadèmia de Medicina de Barcelona".

[1] (cajal-pasik-pasik-99-00-02)

"His studies on the *architectural organization of the brain*, and his *prophetical predictions* of its functions became the basis of neuroanatomy, neurophysiology, neuropathology, and what he named as *'rational psychology'*. This monumental work justifies his well deserved title of founder of modem Neuroscience."

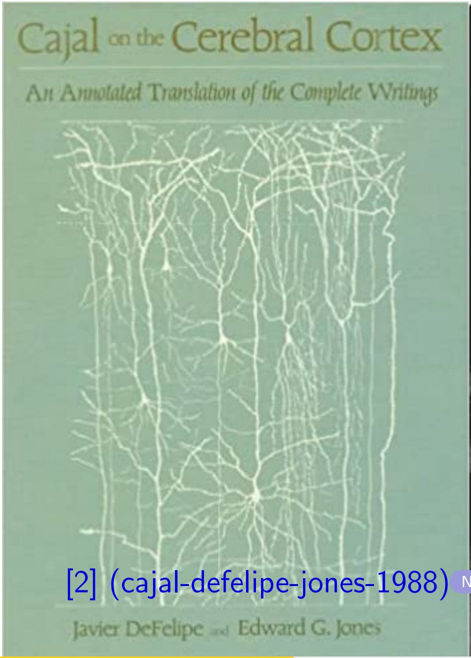

[2] (cajal-defelipe-jones-1988) N

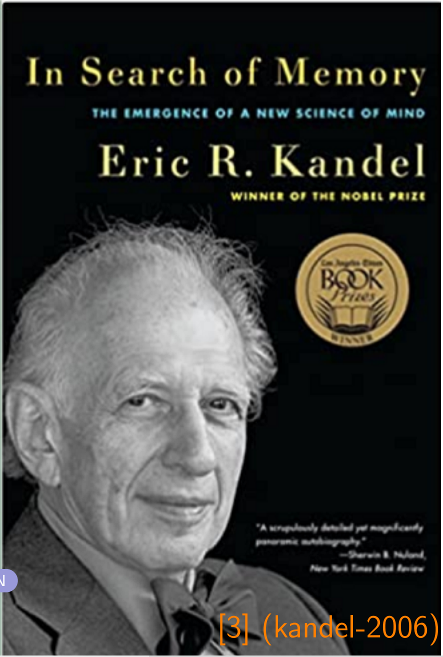

[3] (kandel-2006)

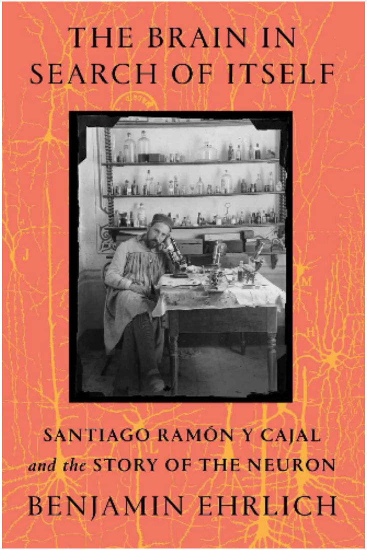

[4] (ehrlich-2022) N

- The highest ideal for a biologist is to clarify the *enigma of the self*.
- The mysterious butterflies of the soul, [the neurons] whose beating of wings may one day reveal to us the *secrets of the mind*.
- Every man can be, if he wants to, a sculptor of his own brain.
- Nothing inspires me more reverence and awe than an old man who is willing to change his mind.
- The car of Spanish culture lacks the wheel of science.

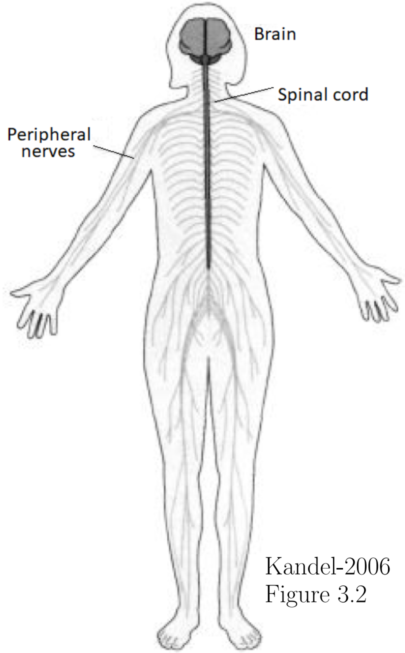

Kandel-2006
Figure 3.2

The central and peripheral nervous systems

- The central nervous system, which consists of the **brain** and the **spinal cord**, is bilaterally symmetrical.
- The spinal cord receives sensory information from the skin through bundles of long axons that innervate the skin. These bundles are called peripheral nerves.
- The spinal cord also sends motor commands to muscles through the axons of the motor neurons.
- These sensory receptors and motor axons are part of the peripheral nervous system.

Neural networks

Modeling challenges
Artificial neurons
Artificial Neural networks
Computational resources and techniques

[5] (sporns-2011), [6] (sporns-2018)

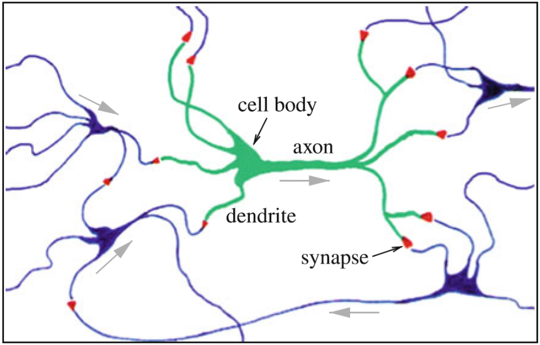

"Network neuroscience is a thriving and rapidly expanding field. Empirical data on brain networks, from *molecular* to *behavioral scales*, are ever increasing in size and complexity. These developments lead to a strong demand for appropriate tools and methods that model and analyze brain network data, such as those provided by graph theory" [6, from the summary].

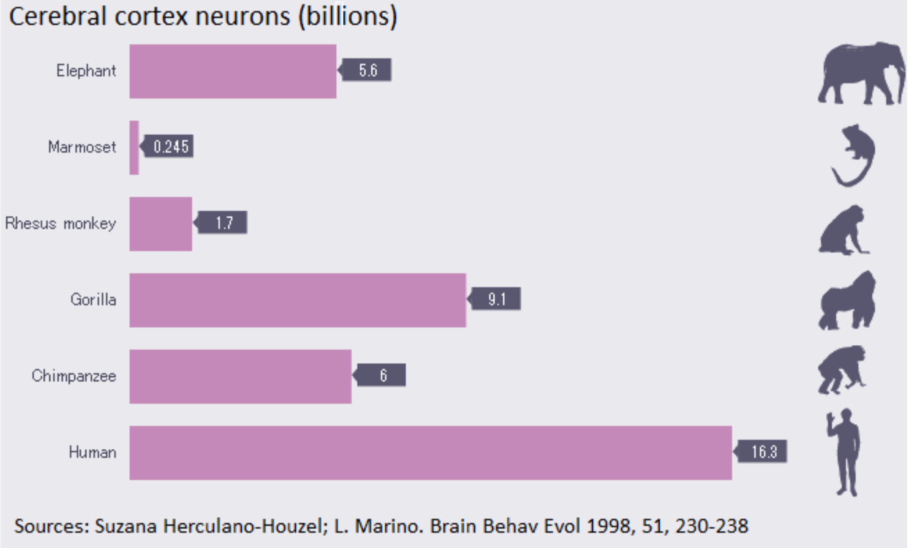
Cerebral cortex neurons (billions)
Elephant
5.6
Marmoset
0.245
Rhesus monkey
1.7
Gorilla
9.1
Chimpanzee
6
Human
16.3
Sources: Suzana Herculano-Houzel; L. Marino. Brain Behav Evol 1998, 51, 230-238

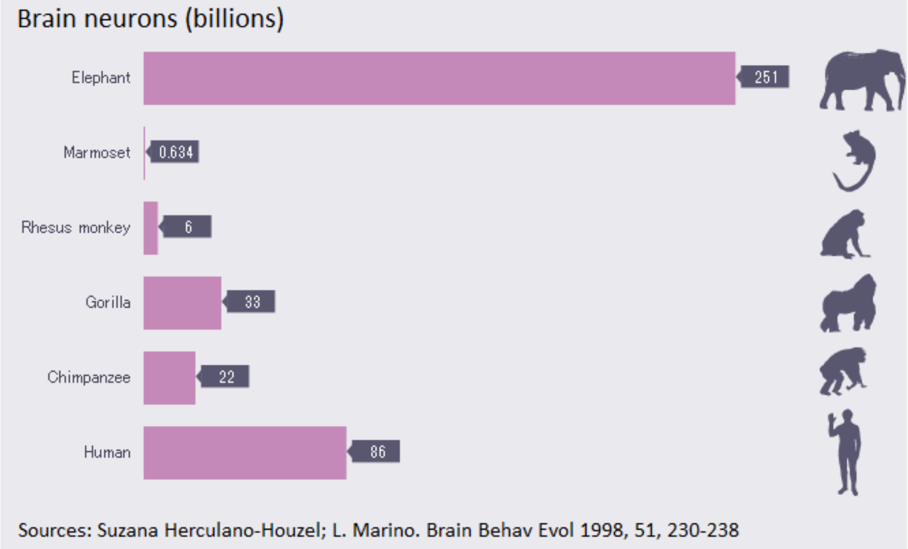
Brain neurons (billions)
Elephant
251
Marmoset
0.634
Rhesus monkey
6
Gorilla
33
Chimpanzee
22
Human
86
Sources: Suzana Herculano-Houzel; L. Marino. Brain Behav Evol 1998, 51, 230-238

In AL, a useful model of a *neuron* is depicted in Fig. 3.1:

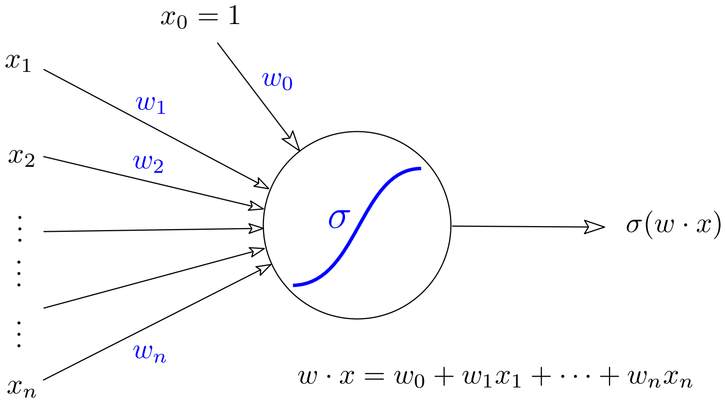

Figure 4.1: Model of a neuron. The neuron's output depends on the inputs x, the *weights* w, and on σ (*activation function*), and this functionality is represented by the decorated circle.

In mathematical terms, a neuron is a function

$$x \mapsto f_w(x) = \sigma(x \cdot w), \tag{1}$$

where $w \in \mathbf{R}^n$ (*weights* or *parameters*) and σ is a *sigmoid* function (*activation function*), like for instance the *logistic function* $\sigma(t) = (1 + e^{-t})^{-1}$, in which case the neuron computes a *logistic regression*. Other widespread choices are $\mathrm{ReLU}(t) = \max(t, 0)$ (Rectified Linear Unit) and $\tanh(t)$.

Augmenting x with $x_0 = 1$ and providing an extra weight w_0 (called the *bias*), the neuron computes $\sigma(w_0 + w_1x_1 + \cdots + x_nw_n)$.

A *neural network* (NN) can be construed as a *composition of neurons* according to a directed graph of connections called the *architecture* of the net.

The standard architecture of a NN is a directed graph structured in *layers* L_j, as illustrated in Figure 4.1, and its *functional signature* can be condensed as a chain:

$$\mathcal{N}: \quad \text{Input} \to L_0 \xrightarrow{f_1} L_1 \xrightarrow{f_2} \cdots \to L_{d-1} \xrightarrow{f_d} L_d \to \text{Output} \tag{2}$$

The integer d is the *depth* of the net. Conventionally, the net is *deep* if $d > 2$, and *shallow* otherwise. The layers $L_1, \ldots, L_{d-1}$ are considered to be *hidden*, while the input and output layers (L_0 and L_d), are *visible*.

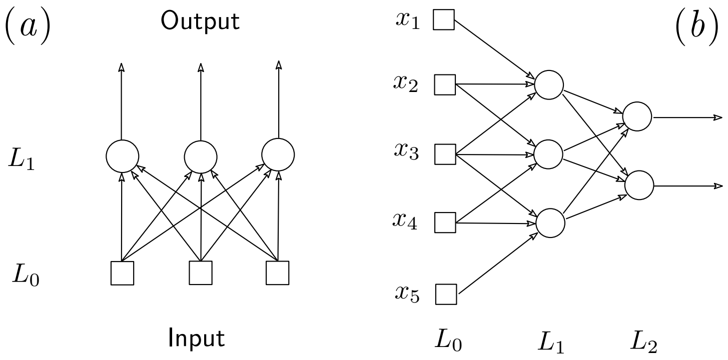

Figure 5.1: (a) Neural network with no hidden neurons and *fully connected*. (b) Network with a hidden layer L_1 of three neurons fully connected to the two output neurons of L_2. The input layer, L_0, is only partially connected to L_1.

Set of data $x^j \in \mathbf{R}^n$, $j \in [m]$ (*dataset*).

Want to predict values y^j provided by a *supervisor* or *expert* in such a way that for objects x not in the dataset the value y corresponding to x is predicted with high probability (*generalization* capacity).

Hypothesis space: the space of functions computed by the network,

$$\mathcal{H} = \{h_w\}_{w \in W}.$$

Problem: to find $w \in W$ such that $h_w(x^j) \approx y^j$.

Method: If the criterion for fitness of $h \in \mathcal{H}$ depends on a function $\ell(h(x), y)$ (*local cost*), the problem amounts to finding w that minimizes $\sum_j \ell(f_w(x^j), y^j)$ (*empirical cost* or *risk*):

$$\mathrm{argmin}_w \sum_j \ell(h_w(x^j), y^j).$$

This is the *empirical risk minimization* rule, ERM.

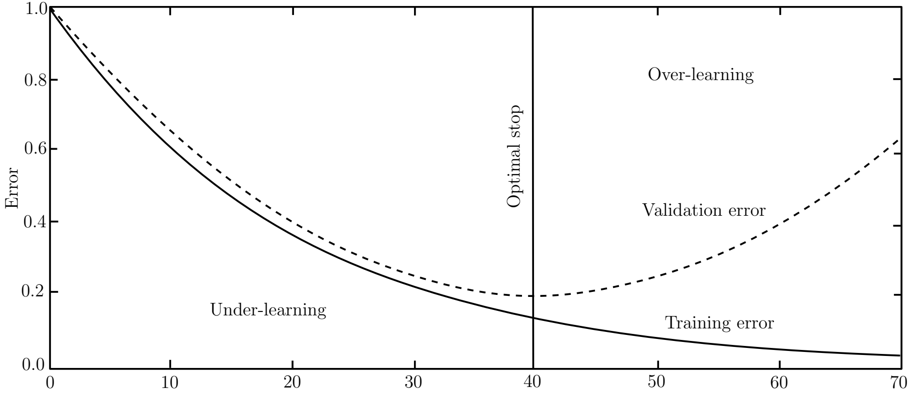

Scheme of a basic training process. In one *epoch* (one step in the training loop), f_w is applied to the dataset. The proportion of errors is the *training error* for that epoch. The learning algorithm determines an update $w = w - \Delta w$ (for example with some variant of gradient descent) and a new epoch is run. The *validation error* is the error incurred by f_w on a *validation dataset*. The training error is decreasing, while the validation error decreases until the *optimal stop* and thereafter it increases.

Currently, there is a wealth of software (frameworks) for deep learning (see `Comparison_of_deep-learning_software` in Wikipedia).

For example, *Tensorflow* (see [7]) provides

> ... an interface for expressing machine learning algorithms, and an implementation for executing such algorithms.

Most of them offer a Python interface and increasingly also a *Julia* interface, as for instance *Tensorflow*. An interesting case is Flux (2017), which is pure Julia (framework and interface).

Critical appraisals. [8] (marcus-2018), [9] (marcus-davis-2019), [10] (lamb-et5-2020), [11] (marcus-2020), [12] (marcus-2022-jul), [13] (marcus-2022-oct)

N

Interactions AIs–Mathematics

Maths applied to ML:

–Generalization bounds

–Recent work of A. Figalli at al.

–The double descent phenomenon

DL applied to Mathematics

A *generalization bound* for the hypothesis $h \in \mathcal{H}$ returned by a *learning algorithm* has the form

$$L(h) \leqslant_\delta L_{\mathcal{D}}(h) + \mathrm{F}(\mathrm{C}, m, \delta),$$

where $L(h)$ is the *expected loss* (or *error*) of h; $L_{\mathcal{D}}(h)$ is the empirical error of h on the data set $\mathcal{D}$; $\mathrm{C} = \mathrm{C}(\mathcal{H})$ is some measure of the *capacity* (or *complexity*) of $\mathcal{H}$; δ is a positive (small) real number; $\mathrm{F}(\mathrm{C}, m, \delta)$ is a mathematical expression involving the parameters C, $m = |\mathcal{D}|$, and δ; and where $\leqslant_\delta$ means that the inequality is true with probability at least $1 - \delta$ relative to the data samples $\mathcal{D}$.

Example [14, Cor. 4.6] Let $\mathcal{H}$ be a finite set of binary hypothesis. Then

$$L(h) \leqslant_\delta L_{\mathcal{D}}(h) + \sqrt{\frac{\ln|\mathcal{H}| + \ln\frac{2}{\delta}}{2m}}.$$

In this case $\mathrm{C} = |\mathcal{H}|$.

N

[15] (abbas-sutter-figalli-woerner-2021), *Effective dimension of machine learning models* (17 pages): "we propose the local effective dimension as a capacity measure which seems to *correlate well with generalization error* on standard data sets. Importantly, we prove that *the local effective dimension bounds the generalization error* and discuss the aptness of this capacity measure for machine learning models". Table 1:

	VC-dimension	Rademacher complexity	Margin-based	Norm-based	Sharpness-based	Local ED
1. Generalization bound	✓	✓	✓	✓	✓	✓
2. Correlation generalization	✗	✗	✗	✗	✗	✓
3. Scale invariant	✗	✓	✗	✗	✗	✓
4. Data dependent	✗	✓	✓	✓	✓	✓
5. Training dependent	✗	✗	✓	✓	✓	✓
6. Finite data	✗	✓	✓	✓	✓	✓
7. Efficient evaluation	✗	✗	✓	✓	✓	✓

The paper [16] (abbas-sutter-zoufal-lucchi-figalli-woerner-2021) establishes that "well-designed *quantum neural networks offer an advantage over classical neural networks* through a higher effective dimension and faster training ability, which we verify on real quantum hardware".

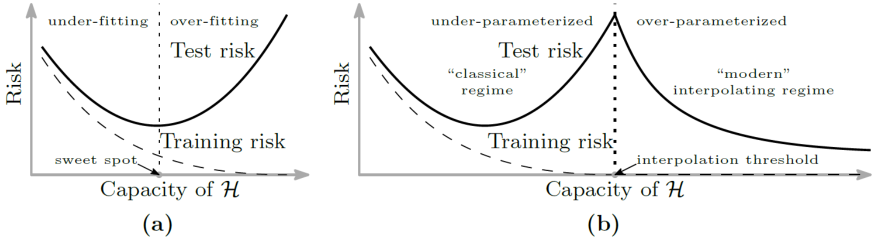

[17] (bach-2022), §11.2, and references therein.

Other applications:

Graph theory: [18] (calvo-tyukin-makarov-2020)

Topology: [19] (gardner-et7-2022).

- [20] (he-2021) (general). Publicity. Y-H He.
- [21] (he-2021-CY).
- [10] (lamb-et5-2020) (graph NNs).
- [22] (wagner-2021) (see NL01, Reviews, page 20). A nice application of reinforcement learning.

N

Brain science

A hint from David Hestenes
Philosophiæ Neuronalis Principia Mathematica?
Further Readings

- [23] (hestenes-1987-brain) (*How the brain works: the next great scientific revolution*):

"In spite of the enormous complexity of the human brain, there are good reasons to believe that only a few basic principles will be needed to understand how it processes sensory input and controls motor output. In fact, the most important principles may be known already! These principles provide the basis for a definite mathematical theory of learning, memory, and behavior." N

Conference AGACSE-2015 (Barcelona). Awarding of the **First Hestenes Prize**. Left to right: D. Hestenes, Lei Huang (winner), Silvia Franchini (finalist), Pierre Dechant (finalist), S. Xambó (organizer), E. Bayro (organizer).

Steve Grossberg received 2015 Lifetime Achievement Award from the Society of Experimental Psychologists. Photo by Eric Levin. N

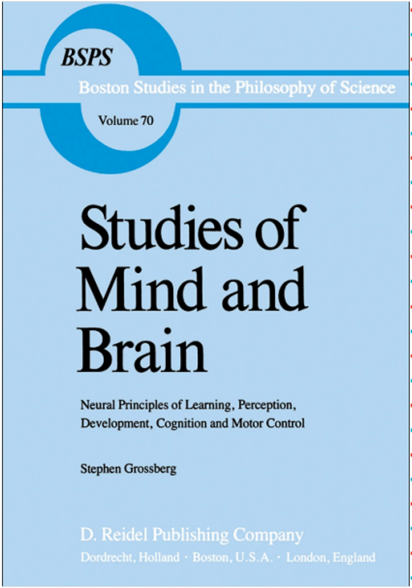

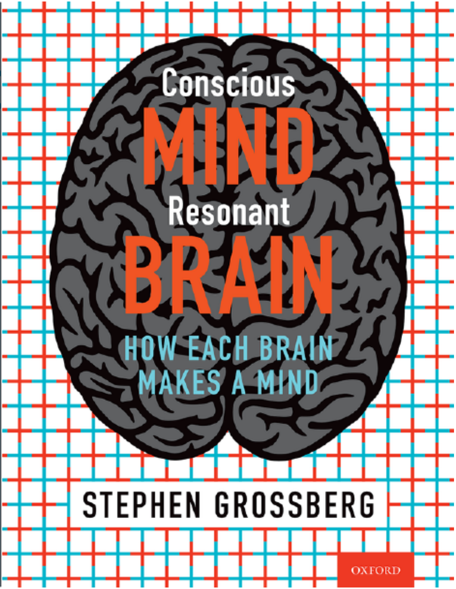

[24] (grossberg-1982) [25] (grossberg-2021) OUP flyer

See also [26] (grossberg-2017) and [27] (wunsch-2019).

- Newton's Principia: *New intuitions and new mathematics*.

"I hope that this brief summary [split between physical sciences and psychology and neuroscience at the beginning od the XX century] helps to put into historical perspective why it has taken so long to begin to theoretically understand how a brain can give rise to a mind. Such progress required the introduction of a new scientific paradigm that simultaneously discovered new conceptual intuitions and new mathematics with which to understand the nonlinear, nonlocal, and nonstationary laws that link brain to mind." N

- Einstein's dark hole: *cognitive impenetrability*.

Einstein: In a letter to Queen Elizabeth of Belgium, 1933, Einstein wrote: "Most of us prefer to look outside rather than inside ourselves; for in the latter case we see but a dark hole, which means: nothing at all."

"The neural computations that regulate our daily experiences are inaccessible to us due to the property of *cognitive impenetrability*. Because of this impenetrability, we can behave in a world of interacting percepts, feelings, and ideas without having to worry about, or to be distracted by, our individual nerve cells, electrical potentials, and chemical transmitters. We experience the apparent simplicity of our behavioral mastery rather than the actual complexity of its generative neural machinery."

Feynman and human vision

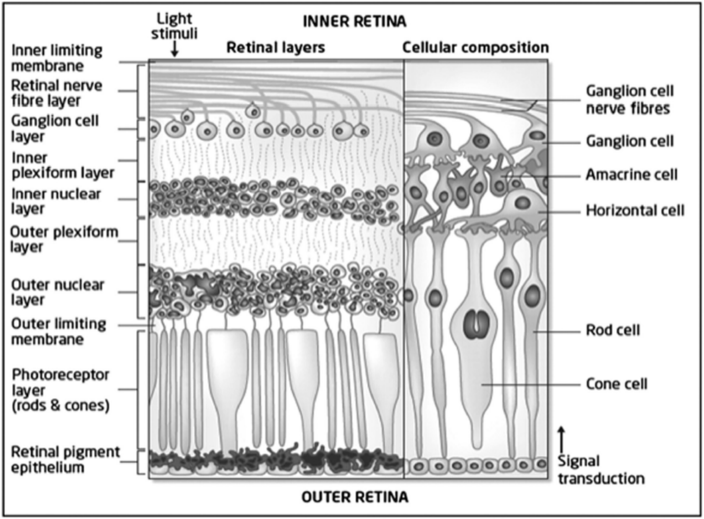

http://brain.oxfordjournals.org/content/early/2011/01/20/brain.awq346

- **von Neumann**: *The Computer and the Brain*.

The **McCulloch–Pitts model** [of neural networks, 1943] had an influence far beyond the field of neural networks. One of its most enduring influences was upon **John Von Neumann**, one of the greatest mathematicians of the twentieth century, when he was at the Institute for Advanced Study in Princeton, developing the digital computer that has totally revolutionized our lives.

This influence is illustrated by a famous series of Silliman Lectures that Von Neumann gave at Yale University in 1956, which he called *The Computer and the Brain*.

Shun-Ichi Amari: *Back propagation algorithm*

Variants of the back propagation model were introduced independently by several authors, including **Shun-Ichi Amari** in 1972 in Japan (Amari, 1972), **Paul Werbos** in 1974 as part of his PhD thesis at Harvard University (Werbos, 1974, 1994), and **David Parker** in 1982 in California (Parker, 1982, 1985, 1986, 1987). **Paul Werbos** seems to be the first person to have published the algorithm in its modern form and use it successfully in applications.

Back propagation finally became popular in response to an oft-cited article by **David Rumelhart**, **Geoffrey Hinton**, and **Ronald Williams** that was published in 1986 (Rumelhart, Hinton, and Williams, 1986). Their 1986 article has often been incorrectly cited as the source of the algorithm by people who are unaware of its history.

Towards a synthesis of biological and artificial intelligence. “As I noted , the same canonical design for neocortical circuits can be specialized to accomplish many different perceptual and cognitive tasks. This claim has been partially realized by the 3D LAMINART model for 3D vision and figure-ground perception; the cARTWORD model for conscious speech perception, learning, and recognition; and the LIST PARSE model for cognitive working memory and planning. *These models provide a foundation for a future unified theory of biological intelligence.* They also illustrate *how software and hardware with a similar canonical circuit design can be specialized to carry out different types of intelligence in future engineering and technological applications*. Because of their compatible design, these software and hardware circuits can be assembled into the ‘brain’ of increasingly autonomous adaptive agents. *If and when this happens, biological and artificial intelligence will be able to seamlessly interact, and complement each others’ strengths and weaknesses*.

"While most mind and brain experimentalists ignored theory, and most theorists looked for more hospitable frontiers, there arose the widespread tendency to fill this theoretical vacuum by interpreting brain function in terms of whatever technology happened to be current. The ever-expanding list of technological metaphors to which the brain has been compared includes telegraph circuits, hydraulic systems, information processing channels, digital computers, linear control systems, catastrophes (in the mathematical sense of René Thom (1977)), *holograms*, *spin glasses*, *Bayesian networks*, and *Deep Learning*. *All of these metaphors have been unable to explain substantial databases that link mind and brain*, since none of them arose from a sustained analysis of mind and brain data."

Francis Crick, "most famous for his classical Nobel Prize–winning work on understanding DNA, published a book in 1994, with the arresting title *The Astonishing Hypothesis*, in which he discussed some aspects of consciousness (Crick, 1994). The astonishing hypothesis refers to the fact that *all of our mental phenomena arise from activities of the nerve cells, or neurons, of our brains*." Crick claimed that "this hypothesis is so alien to the ideas of most people alive today that it can truly be called astonishing" [yes, but oblivous of the views of preceding scientists, particularly **Ramón y Cajal**]

Cajal's *Neuron Doctrine* is now taken for granted as gospel. However, on the day that Cajal and Golgi shared the Nobel Prize, Golgi used his own Nobel lecture to deny that individual neurons exist, even as Cajal asserted that they do. Controversy existed then in psychology and neuroscience, and it continues to exist today. How could it be otherwise, given the complexity of the subject matter, and its importance to human identity and society?

- Philosophers, D. Dennett in particular: *Consciousness Explained* (?) (1991)

N

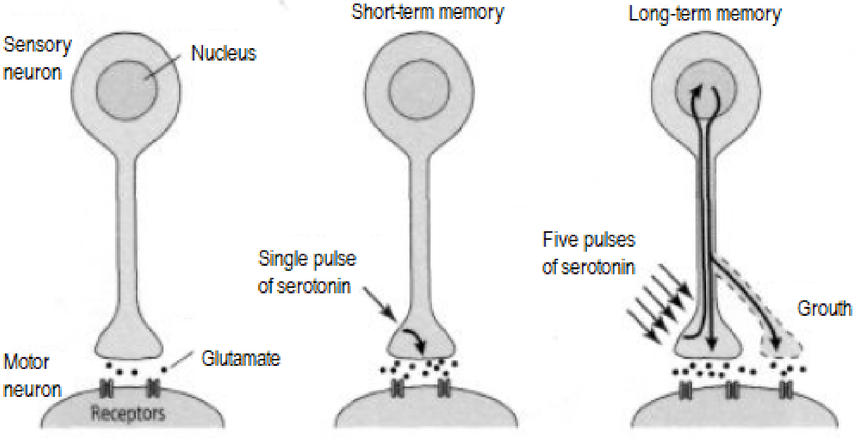

18-4 Changes underlying short- and long-term memory in a single sensory and motor neuron (p. 256)

Functional change: Synapse strenghthened via enhanced release of glutamate. The nucleus is not involved

Anatomical change: Synthesis of proteins in the nucleus and growth of new synaptic connections, as well as enhanced release of glutamate.

[3] (kandel-2006)

For a thorough treatise on Neural Science, see

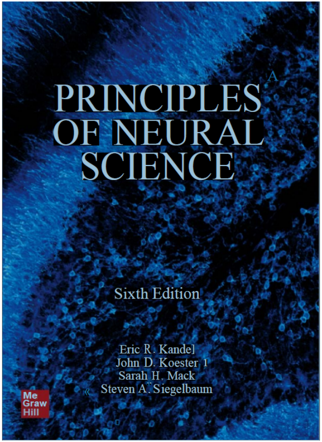

[28] (kandel-koester-mack-siegelbaum-2021)

Some Nobel Prizes in Physiology or Medicine: N

We can only see a short distance ahead,
but we can see plenty there that needs to be done
(Alan Turing)

For example:
(contribute) to produce a comprehensive

Philosophiæ Neuronalis
Principia Mathematica!

Thank you!

Notes

- Cajal formulated the *neuron doctrine*, the basis for all modern thinking about the nervous system. He is arguably the most important brain scientist who ever lived (Kandel 2006) P

- Warmflash-2016 (In "Visionlearning: your insight into science"): By 1891, Ramón y Cajal decided that *the expansion of nerve cells occurred through growth of dendrites*, which he referred to as "protoplasmic branches." Within three years [...] he *speculated that intelligence was related to the number and efficiency of connections between different types of cells with the part of the brain known as the cerebral cortex*. It was a remarkable idea for the time, and it dovetails with what neuroscientists have learned in recent years.

- Historians have ranked Cajal alongside Darwin and Pasteur as one of the greatest biologists [...] among Copernicus, Galileo and Newton as one of the greatest scientists of all time. His masterpiece, *Texture of the Nervous System of Man and the Vertebrates*, is a foundational text for neuroscience, comparable to *On the Origin of Species* for evolutionary biology (SciAM, April 2022).

P

"If one aims to build richer AI systems, i.e. semantically sound, explainable, and reliable, one has to add a sound reasoning layer to deep learning" (from [10] (lamb-et5-2020)).

"Recent research in artificial intelligence and machine learning has largely emphasized general purpose learning and ever larger training sets and more and more compute. In contrast, I propose a *hybrid, knowledge driven, reasoning based approach*, centered around *cognitive models*, that could provide the substrate for a richer, more robust AI than is currently possible" ([11, Abstract]).

P

A learning algorithm has good *generalization rate* if the hypothesis $h \in \mathcal{H}$ it supplies is close to the expert who produced the examples, that is, if the expected loss $L(h)$ is small. How to guarantee this condition if the LA only knows the examples $\mathcal{D}$ and somehow the space $\mathcal{H}$? The problem may seem impossible if we consider that $\mathcal{D}$ is always finite, and that $\mathcal{H}$, as a rule, is infinite.

The result of the example tells us that the generalization rate increases if the empirical risk minimization ensures that $L_{\mathcal{D}}(h)$ is small and that m is sufficient large so that the second sum of the expression is also small.

P

he-2021: We review, for a general audience, a variety of recent experiments on extracting structure from machine-learning mathematical data that have been compiled over the years. Focusing on supervised machine-learning on labeled data from different fields ranging from geometry to representation theory, from combinatorics to number theory, we present a comparative study of the accuracies on different problems. The paradigm should be useful for conjecture formulation, finding more efficient methods of computation, as well as probing into certain hierarchy of structures in mathematics. Based on various colloquia, seminars and conference talks in 2020, this is a contribution to the launch of the journal *Data Science in the Mathematical Sciences*.

P

Hestenes is famously known for his life-long endeavor to show the superiority of geometric calculi to phrase and explore a wide variety of geometrical and physical theories, and their expediency for a deep reform of physics curricula. Since his 1966 PhD and to this day, he has shaped those realms in what I called the *Hestenes era* in a lecture at the ICACGA 2022 Conference (4 Oct 2022: *Spinning spinors with geometric algebra for one century and beyond*). In that lecture I took the opportunity to refer also to a few of his works that lie outside of his main stream, with emphasis on the one you see cited and quoted in the screen.

This 33-page paper would deserve, despite having appeared 35 years ago, a thorough study in a seminar in conjunction with the materials presented in next screen.

Which produced, quote, "a mathematical theory that bridges the gap between neurophysiology and psychology, providing penetrating insights into brain mechanisms for learning, memory, motivation, and the organization of behavior"? P

- Born in 1939, STEPHEN GROSSBERG is Wang Professor of Cognitive and Neural Systems at Boston University, Professor Emeritus of Mathematics & Statistics, Psychological & Brain Sciences, and Biomedical Engineering. Google scholar: 1034 items; 80464 citations (8439 since 2018); h-index 131 (40 since 2018); i10-index 389 (182 since 2018).

P

The book on the left is a collection of 13 papers, and it is the work cosidered by Hestenes five years after its publication. The first paper in the collection is *How does a brain build a cognitive code?* (1982, 42 pages, 1927 citations), "a self-contained introduction to my work from a recent perspective"; the last, *A theory of human memory: self-organization and performance of sensory-motor codes, maps, and plans* (1978, 144 pages, 481 citations), "The theory's main contribution is to show how a temporal stream of data patterns can drive the formation of globally consistent cognitive representations and purposive behavioral plans despite the abysmal ignorance of individual cells. To build this theory, I needed all the conceptual and mathematical machinery that I had been accumulating over the past twenty years. The lesson of the article is that all the pieces fit together".

Two main points about the book on the right (extracted from the OUP flyer):

- Explores how your mind works, notably how you learn to consciously see, hear, feel, and know things.

- Creates a computational foundation for the next generation of autonomous, adaptive, and intelligent algorithms, devices, and mobile agents in engineering, technology, and AI. P

"In many scientific revolutions, such as relativity theory and quantum mechanics, once the new physical intuitions were discovered, relevant mathematics was available with which to convert them into rigorous theoretical science. In the mind and brain sciences, we have not been so lucky, since both new intuitions and new mathematics needed to be developed. The Newtonian revolution was also of this kind, since Newton had to both derive the laws of celestial mechanics and to discover calculus with which to mathematically analyze them. My own scientific work since 1957, when I was a Freshman in Dartmouth College, has been devoted to introducing and developing foundational intuitions, mathematics, and the behavioral and neural models built upon them into the mind-brain sciences. I could never have imagined then how much I would be able to discover and understand with the help of many gifted PhD student, postdoctoral fellow, and faculty collaborators."

P

Feynman: "This inversion of the retinal layers may make sense from the viewpoint of how the retina develops, but it seems very peculiar from the viewpoint of how the brain sees. Richard Feynman, one of the most brilliant physicists of the twentieth century, got very interested in how the brain sees until he became aware of these facts. Feynman was a famously honest man who realized that he could not think of a sensible explanation for how the brain sees so well despite the retina's peculiarities of having a blind spot as well as retinal veins and retinal layers in front of the photoreceptors. *Every scientist needs heuristics, or intuitive design principles, on which to base a principled scientific theory*. These heuristics need to make sense, and in the best cases are even parsimonious and beautiful, for a top theorist to be comfortable with them. Feynman realized that he could not find intuitively plausible principles for how the brain copes with its noisy retinas, so he got out of vision. Of course, a man with Feynman's abundant intellectual and personal resources could not be long deterred by this disappointment. He rapidly moved on to successfully pioneer the fields of quantum computing and nanotechnology." P

The visual illusion of neon color spreading. Neither the square nor the blue color that are perceived within it are in the image that defines a neon color display. The display consists only of black and blue arcs.

P

1906: Camillo Golgi and Santiago Ramón y Cajal: *structure of the nervous system*.

1932: Sir Charles Scott Sherrington and Edgar Douglas Adrian: *functions of neurons*.

1936: Henry Hallett Dale and Otto Loewi: *chemical transmission of nerve impulses* (neurotransmitters)

1944: Joseph Erlanger and Herbert Spencer Gasser: functions of single nerve fibers (*action potentials*)

1949: Walter Rudolf Hess (interbrain) and António Caetano Egas Moniz (lobotomy)

1962: Francis Harry Compton Crick, James Dewey Watson, and Maurice Hugh Frederick Wilkins; *structure nucleic acids*.

1963: Sir John Carew Eccles (*synapse*), Sir Alan Lloyd Hodgkin, and Sir Andrew Fielding Huxley: *ionic mechanisms involved in excitation and inhibition in the peripheral and central portions of the nerve cell membrane*.

1970: Julius Axelrod, Ulf von Euler, and Sir Bernard Katz: *transmittors in the nerve terminals and the mechanism for their storage, release and inactivation*.

1981: Roger W. Sperry (*functional specialization of cerebral hemispheres*), and David H. Hubel – Torsten N. Wiesel (*information processing in the visual system*).

2000: Arvid Carlsson (*dopamine*), Paul Greengard (*molecular and cellular functions of neurons*), and Eric R. Kandel (*physiological basis of memory storage in neurons*): *signal transduction in the nervous system*.

2014: John O'Keefe, May-Britt Moser, and Edvard I. Moser: *discoveries of cells that constitute a positioning system in the brain*.

P

References I

[1] S. Ramón y Cajal, *Texture of the Nervous System of Man and the Vertebrates (Volumes I, II, III)*.

Springer-Verlag, 1999, 2000, 2002.

An annotated and edited translation of the original Spanish text (1904), with the additions in the French version, by Pedro Pasik and Tauba Pasik.

[2] S. Ramón y Cajal, J. DeFelipe, and E. G. Jones, *Cajal on the cerebral cortex: An annotated translation of the complete writings*.

Oxford University Press, 1988.

XVII+654p.

References II

[3] E. R. Kandel, *In search of memory: The emergence of a new science of mind.*

W. W. Norton & Company, 2006.

[4] B. Ehrlich, *The Brain in Search of Itself: Santiago Ramón y Cajal and the story of the neuron.*

Farrar, Straus and Giroux, 2022.

[5] O. Sporns, *Networks of the Brain.*

MIT press, 2011.

[6] O. Sporns, "Graph theory methods: applications in brain networks," *Dialogues in clinical neuroscience*, vol. 20, no. 2, 2018.

References III

[7] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, *et al.*, "Tensorflow: Large-scale machine learning on heterogeneous distributed systems," *arXiv preprint arXiv:1603.04467*, 2016.

[8] G. Marcus, "Deep learning: A critical appraisal," 2018.
`https://arxiv.org/ftp/arxiv/papers/1801/1801.00631.pdf`.

[9] G. Marcus and E. Davis, *Rebooting AI: Building artificial intelligence we can trust*.
Vintage, 2019.

References IV

[10] L. Lamb, A. Garcez, M. Gori, M. Prates, P. Avelar, and M. Vardi, "Graph neural networks meet neural-symbolic computing: A survey and perspective," 2020.

https://arxiv.org/pdf/2003.00330.pdf.

[11] G. Marcus, "The next decade in AI: four steps towards robust artificial intelligence," 2020.

https://arxiv.org/vc/arxiv/papers/2002/2002.06177v2.pdf.

[12] G. Marcus, "Artificial general intelligence is not as imminent as you might think," *Scientific American*, vol. July 2022, 2022.

References V

[13] G. Marcus, "Artificial confidence," *Scientific American*, vol. October 2022, pp. 44–45, 2022.

[14] S. Shalev-Shwartz and S. Ben-David, *Understanding machine learning: From theory to algorithms*.

Cambridge University Press, 2014.

440 pp.

[15] A. Abbas, D. Sutter, A. Figalli, and S. Woerner, "Effective dimension of machine learning models," 2021.

`https://arxiv.org/pdf/2112.04807`.

References VI

[16] A. Abbas, D. Sutter, C. Zoufal, A. Lucchi, A. Figalli, and S. Woerner, "The power of quantum neural networks," *Nature Computational Science*, vol. 1, no. 6, pp. 403–409, 2021.

https://arxiv.org/pdf/2011.00027.

[17] F. Bach, "Learning theory from first principles, draft *(13-Oct-22)*," 2022.

https://www.di.ens.fr/~fbach/ltfp_book.pdf. xiv+360 p.

[18] C. Calvo Tapia, I. Tyukin, and V. A. Makarov, "Universal principles justify the existence of concept cells," *Scientific reports*, vol. 10, no. 1, pp. 1–9, 2020.

References VII

[19] R. J. Gardner, E. Hermansen, M. Pachitariu, Y. Burak, N. A. Baas, B. A. Dunn, M.-B. Moser, and E. I. Moser, "Toroidal topology of population activity in grid cells," *Nature*, vol. 602, no. 7895, pp. 123–128, 2022.

[20] Y.-H. He, "Machine-learning mathematical structures," 2021. `https://arxiv.org/pdf/2101.06317.pdf`.

[21] Y.-H. He, *The Calabi–Yau Landscape: From Geometry, to Physics, to Machine Learning*, vol. 2293 of *Lecture Notes in Mathematics*. Springer, 2021.

References VIII

[22] A. Z. Wagner, "Constructions in combinatorics via neural networks," 2021.

https://arxiv.org/pdf/2104.14516.pdf.

[23] D. Hestenes, "How the brain works: the next great scientific revolution," in *Maximum-entropy and Bayesian Spectral Analysis and Estimation Problems*, pp. 173–205, Reidel, 1987.

[24] S. Grossberg, *Studies of Mind and Brain. Neural Principles of Learning, Perception, Development, Cognition and Motor Control*, vol. 70 of *Boston Studies in the Philosophy of Science*.

Reidel, 1982.

References IX

[25] S. Grossberg, *Conscious mind, resonant brain: How each brain makes a mind*.

Oxford University Press, 2021.

[26] S. Grossberg, "Towards solving the hard problem of consciousness: The varieties of brain resonances and the conscious experiences that they support," *Neural Networks*, vol. 87, pp. 38–95, 2017.

[27] D. Wunsch II, "Admiring the Great Mountain: A Celebration Special Issue in Honor of Stephen Grossberg's 80th Birthday," 2019.

[28] E. R. Kandel, J. D. Koester, S. H. Mack, and S. A. Siegelbaum (editors), *Principles of Neural Science (sixth edition)*.

McGraw-Hill, 2021.