

The Brain in Silicon: History, and Skepticism

Alessio Plebe, Giorgio Grasso

► **To cite this version:**

Alessio Plebe, Giorgio Grasso. The Brain in Silicon: History, and Skepticism. 3rd International Conference on History and Philosophy of Computing (HaPoC), Oct 2015, Pisa, Italy. pp.273-286, 10.1007/978-3-319-47286-7_19 . hal-01615293

HAL Id: hal-01615293

<https://hal.inria.fr/hal-01615293>

Submitted on 12 Oct 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



The brain in silicon: history, and skepticism

Alessio Plebe and Giorgio Grasso

Department of Cognitive Science, University of Messina, Italy
{aplebe, gmgrasso}@unime.it

Abstract. This paper analyzes the idea of designing computer hardware inspired by the knowledge of how the brain works. This endeavor has lurked around the twists and turns of the computer history since its beginning, and it is still an open challenge today. We briefly review the main steps of this long lasting challenge. Despite obvious progress and changes in the computer technology and in the knowledge of neural mechanisms, along this history there is an impressive similarity in the arguments put forward in support of potential advantages of neural hardware over traditional microprocessor architectures. In fact, almost no results of all that effort reached maturity. We argue that these arguments are theoretically flawed, and therefore the premises for the success of neuromorphic hardware are weak.

Keywords: neurocomputer; neuromorphic hardware; neural networks; canonical neural circuit

1 Introduction

The suggestion of getting inspiration from neural architectures in designing microprocessors, the “brain in silicon” idea, has lurked around the twists and turns of the computer history, almost since its beginning. More precisely, three main periods during which this idea turned into practical projects can be identified. The first neural hardware was designed by Marvin Minsky in 1951 [52], building upon the logical interpretation of neural activity by McCulloch and Pitts [48], and was followed by just a few more attempts. After a long period of an almost complete lack of progress, a renewed interest sparked at the end of the 80’s, driven by the success of neural algorithms in software [64], with several funded projects in Europe, US, and Japan [75]. However, at the beginning of this century, almost no results of all that effort had reached maturity. In the last few years, a new wave of enthusiasm on neural hardware spread around, propelled by a few large projects funded in Europe and US for realistic brain simulations [45]. Again, a revolution in microprocessor design was forecast, in terms that closely remind of those of the previous two periods. This history is described in the next section.

Despite obvious progresses and changes in the technology and knowledge of neural mechanisms, the analysis of those three periods shows a common view of the reasons to believe the brain in silicon should be successful. This view,

discussed in depth in Section §3, essentially hinges on the computational nature of the nervous systems, and the extreme degree of efficiency it offers, thanks to millions of years of evolution. We will first remark the difference between the possibility of a shared computational nature of the brain and of human-made digital computers, as held within the cognitive science and artificial intelligence communities, and the more problematic equivalence of the features, at the implementation level, of brain and silicon computations. Second, even if mimicking the brain at the circuital level might be effective, a further point is that nobody knows exactly what to copy, as what turns the brain into an extremely powerful computer is still obscure. We will review the neuroscientific community efforts in characterizing the key circuital elements of the brain, and their yet inconsistent answers. We will add that it is very likely that the most powerful computational feature of the brain relies on the plasticity mechanisms, which are the less likely to be reproduced in silicon. At the end, even if we are unable to make prediction about the future of neuromorphic hardware, we can say that the principles used to promote this enterprise are theoretically flawed, and therefore the premises for the success of this approach are weak.

2 A Short Historical Account

As a matter of fact, the events identified in the following short historical sketch, can be framed into the wider context of the history of the mind as a machine. An history that one might reasonably suggest began with the thorough reflections on the possibility of mechanizing mind by Descartes [16]. Even if he dismissed a mechanical nature of the human mind, he brought to light this exciting idea, embraced shortly after by LaMettrie in the essay *L'Homme Machine* [39], and put in practice by the the French engineer and inventor Jacques de Vaucanson with the *Tambourine Player* and *Flute Player* [9]. In the period around the advent of digital computers, the history of the mind as a machine evolves into, at least, three main movements: the cybernetic project aimed at explaining living behaviors by principles of the information-feedback machine in the 1940s and 1950s [78, 38]; the advent of artificial intelligence [47, 13]; and finally the enterprise of understanding the mind by modeling its workings within the field of cognitive science [7]. The history traced here is much narrower, focused on the more specific idea that hints gathered from the observation of the structure of the brain may provide guiding principles for the design of computer hardware, which may turn out to be competitive with the standard von Neumann architecture. Even if entrenched with the broader histories just mentioned, there are important differences, for example the goal here is rather practical and does not attempt to tackle an understanding of how the mind works.

2.1 First ideas and realization

The first suggestion to design computers borrowing hints from the brain came from Alan Turing [76], who envisioned a machine based on distributed interconnected elements, called *B-type unorganized machine*. It befell even before the first

general-purpose electronic computers were up and running. Turing’s neurons are simply NAND gates with two inputs, randomly interconnected, and each NAND input can be connected or disconnected, and thus a learning method can “organize” the machine by modifying the connections. His idea of learning generic algorithms by reinforcing successful and useful links and of cutting useless ones in networks was the most farsighted of this report. Not so farsighted was his employer at the National Physical Laboratory, for which the report was produced, who dismissed the work as a “schoolboy essay”. Therefore, this report remained hidden for decades, until upheld by Copeland and Proudfoot in 1996 [12].

On the contrary, an earlier paper by McCulloch and Pitts [48], suggesting that neurons work akin logical ports, had a tremendous impact on the yet unborn field of Artificial Intelligence. They attempted to adapt the logic formalism of Carnap [10] to “neural” elements, that have two possible values only, corresponding to the logical truth values. There are two types of connections, excitatory and inhibitory, and an intrinsic feature of all neurons is a threshold, corresponding to the net number of “true” values in input to produce “true” as output value. Today we all know well that McCulloch and Pitts’ idea was simply wrong, they too became well aware of the different direction the growing neuroscientific evidence was pointing to [59], nevertheless the brain as a logic machine was a fascinating hypothesis, and galvanized more than one scholar at that time. Including Minsky [52], who designed SNARK (*Stochastic Neural Analog Reinforcement Computer*), the first neural computer, assembling 40 “neurons”, each made with six vacuum tubes and a motor to adjust mechanically its connections. The objective of SNARK was to try and find the exit from a maze where the machine would play the part of a rat. Running a mouse through a maze to investigate behavioral patterns was one of the most common paradigms in empirical psychology at that time. The construction of SNARK was an ambitious project, one of the first attempt to reproduce artificial intelligent behavior, but it had no influence on the contemporary progress of digital general-purpose computers. Later on, Minsky himself was one of the most authoritative voices begetting the dismissal of artificial neural research as a whole. The analysis he carried out, together with Papert, of Rosenblatt’s perceptron neural network [61], concluded with these discouraging words [51]:

The perceptron has shown itself worthy of study despite (and even because of!) its severe limitations. It has many features to attract attention: its linearity; its intriguing learning theorem; its clear paradigmatic simplicity as a kind of parallel computation. There is no reason to suppose that any of these virtues carry over to the many-layered version. Nevertheless, we consider it to be an important research problem to elucidate (or reject) our intuitive judgment that the extension is sterile.
[pp.231-232]

2.2 The artificial neural networks period

In the 80’s artificial neural networks became rapidly one of the dominant research fields in artificial intelligence, in part thanks to the impressive development of

Rumelhart and McClelland [64] and their group, who introduced simple and efficient learning algorithms, like backpropagation. In the 90's artificial neural networks were the preferred choice for problems in which the rules governing a system are unknown or difficult to implement, thanks to their ability to learn arbitrary functions [28], but applications were typically implemented in software. The enormous world-wide resurgence of artificial neural networks raised again the interest for building brain-like hardware. The European Commission program ESPRIT from the 90's promoted the development of neural hardware with several research projects: ANNIE, PYGMALION, GALATEA and SPRINT. The Japanese made neuromorphic hardware a key component of their 6th generation computing, and in the US funding for the subject was provided by DARPA, ONR and NFS [75]. In the mid 90's about twenty neuromorphic hardware were commercialized, ranging from Siemens' SYNAPSE-1 (*Synthesis of Neural Algorithms on a Parallel Systolic Engine*), to Philips' L-Neuro, to Adaptive Solutions' CNAPS (*Connected Network of Adaptive Processors*), to Intel's ETANN and Hitachi's MY-NEUPOWER [29]. Despite differences in the technical solutions adopted, the shared approach was essentially to achieve in hardware the maximum performance on array multiplication and summation of the results. This is actually the most intensive computation on a feed-forward network scheme [64], where each layer is defined by a weight matrix, to be applied to the current input vector. The way the elements of a neural chip are connected with one another in paralleling the application of the weight matrix is variable. There are *n-parallel* configurations, where the synapses of one neuron are mapped to the same processing element, and several neurons are computed in parallel, and *s-parallel* configurations, where the synapses of one neuron are mapped to different processing elements, so that several synapses, not belonging to the same neuron, are computed at once. The best configurations depends on the network architecture, and on the number of processing elements of the chip. The CNAPS processor has 64 processing elements, Siemens' MA-16 is designed for 4x4 matrix operation, and in SYNAPSE-1 eight MA-16 chips are cascaded to form systolic arrays.

Strictly speaking, not much of the computing architecture design in these projects is inspired by the brain, it is tailored to the abstract interpretation of the computation performed by the brain from a connectionist point of view: summations of inputs weighted by connections strengths. All solutions met a negligible market interest and disappeared shortly.

2.3 The brain “reverse-engineering” challenge

At the beginning of this century a new wave of efforts towards neuron-like hardware mounted, in part driven by the large worldwide enterprise of brain reverse-engineering, taken up by projects like the Blue Brain Project [43] and the Human Brain Project in Europe [45], the DARPA C2S2 (*Cognitive Computing via Synaptronics and Supercomputing*) project at IBM Research [54]. For most of these long-term projects, the ultimate goal is to emulate the entire human brain, and the dominant approach is the emulation of neurons in software, running on

the world top supercomputing systems, like IBM Blue Gene [67]. However, a number of smaller projects started developing new neural hardware too, like FACETs, Neurogrid, SpiNNaker and NeuroDyn [53]. In most of these projects the imitation of the brain reduces, in fact, to a fast parallel implementation of the most time consuming part of mathematical abstractions of neural activity. For example SpiNNaker (*Spiking Neural Network Architecture*, run by Furber and his team at Manchester [32]), is based on conventional digital-based low-power computer ARM9 core, the kind of CPU commonly found in smartphones, programmed to solve a popular simplified formulation of action potential [31]. Currently the SpiNNaker consists of 20,000 chips, each of which emulates 1000 neurons. On a similar path is the new effort taken by Modha at IBM with the TrueNorth chip, based on conventional digital devices running Izhikevich's algorithm, but using an entirely new chip design, consisting of an array of 4096 cores, each with 256 neurons [49]. There are alternatives to digital devices too, that aims at reproducing the analog behavior of neurons. For example, in the FACETS (*Fast Analog Computing with Emergent Transient States*) project [66] an ASIC chip simulates analog neuron waveforms, for as many as 512 neurons with 256 inputs. As soon as an analog neuron reaches the conditions for an action potential, the digital part of the chip generates a spike event with the event time and the address of the spiking neuron.

Even if the most advanced achievements in brain reverse-engineering has been obtained on traditional supercomputers [45], neuromorphic hardware systems may offer a valid alternative for this challenging enterprise. But hopes are high, again, on the potential of neuromorphic computation as a general purpose computer, not just for simulating brain behavior, or executing advanced AI applications, but for ordinary software run by mundane users in their ordinary computers [77].

In order to get a grasp of the motivations for such a periodic impulse toward brain in silicon, it is instructive to compare three overviews of neural hardware and forecasts for the future, spaced each one a decade apart [25, 17, 24]. It is impressive the similarity that they share, in finding an unsatisfactory current impact of neural hardware but expressing the confidence in the potential, in the long run, of this approach. Heemskerk, in 1995, first expressed his concerns [25]:

“Neurocomputer building is expensive in terms of development time and resources, and little is known about the real commercial prospects for working implementations [...] Another reason for not actually building neurocomputers might lie in the fact that the number and variety of (novel) neural network paradigms is still increasing rapidly”

but concluded with:

“If progress advances as rapidly as it has in the past, this implies that neurocomputer performances will increase by about two orders of magnitude [...]. This would offer good opportunities”

Similarly, in 2004 Dias and coworkers [17] stated that:

“A few new neurochips are reported in this survey while the information collected indicates that more neurochips are no longer available commercially. The new solutions that have appeared indicate that this field is still active, but the removal of the market of other solutions does not seem to be good news. [...] there is no clear consensus on how to exploit the currently available [...] technological capabilities for massively parallel neural network hardware implementations.”

nevertheless, they believe in future opportunities:

“These might be the reasons for the slow development of the ANN hardware market in the last years, but the authors believe that this situation will change in the near future with the appearance of new hardware solutions.”

In 2013 Hasler and Marr acknowledge that not much has been achieved so far:

“A primary goal since the early days of neuromorphic hardware research has been to build large-scale systems, although only recently have enough technological breakthroughs been made to allow such visions to be possible.”

but their hopes are even higher:

“Neuromorphic engineering builds artificial systems utilizing basic nervous system operations implemented through bridging fundamental physics of the two mediums, enabling superior synthetic application performance [...] research in this area will accelerate by the pull of commercial ventures that can start utilizing these technologies to competitive commercial advantage.”

3 Computational Secrets of the Brain

There is a noteworthy difference between the neurocomputer project and the AI enterprise, which is grounded on the famous *multiple-realizability* thesis [21]: cognition is characterized as computations independent on their physical implementation. Why then the mechanisms that cause computational power in a biophysical system, like the brain, would cause, in a radically different system, efficient computation in executing generic (including non cognitive) algorithms? The most common argument used to support the belief in the future superiority of neuromorphic hardware is well summarized in these words of Lande [40]:

“One possible answer [to CPU design problems] is to look into what life has invented along half a billion years of evolution [...] Numerous principles found in the brain can provide inspiration to circuit and system designers.”

But this argument is flawed on several counts.

3.1 Evolution is not design

First, the mechanisms implemented by biological evolution are carved on the specific constraints of the organic system, which imposes roads that in principle are not better or worse than man-made solutions. Both the brain and the digital computers are based on electricity. However, electrical power for digital computation is conducted by metals, like copper, the fastest available conductor, and semiconductors, such as silicon and germanium, which allow control over electron flow at the highest possible speed. Thus, nature had to deal with an enormous disadvantage in dealing with electricity, compared to man made devices, in that metals and semiconductors cannot be used inside living organisms. Nature opted for the only electrical conductors compatible with organic materials: ions. The biophysical breakthrough of exploiting electric power in animals has been the ion channel, a sort of natural electrical device, first appeared as potassium channel about three billion years ago in bacteria, then evolved into the sodium channel 650 million years ago, and currently the most important neural channel [79]. The success of this particular ion channel is very likely due to the abundant availability of sodium in the marine environment during the Paleozoic era. How the neural cell emerged from early ion channels is still uncertain. As with many events in evolution, contingencies may have prevailed to pure adaptation. In the case of neurons, an intriguing theory is that their phylogeny is independent from the history of ion channels. One shared prerequisite of all neurons, from a genomic standpoint, is the capacity to express many more genes and gene products than other cell types, a behavior exhibited by most cells too, as a result of severe stress responses. Neurons might have evolved in ancestral metazoans from other cell types, as the result of development in the adaptive response to localized injury and stress [55]. Similar obscure and contingency dependent histories concern the whole span of neural machinery, from its first appearance in cnidarians, the early central nervous system of echnoderms, in the simple brain in flatworms, to the full brains in polychaetes, insects, cephalopods, and vertebrates [62].

Nature has its troubles but also its vantage points with respect to computer designers, in playing with electricity. The repertoire of organic structures and compounds is quite vast: there are more than 100 types of different known neurotransmitters, and a vast galaxy of neural cell types, diverse morphologically and in their molecular basis [74]. The organization of the neural circuitry span over three dimensions. Even in the new era of three dimensional semiconductors, which could accommodate for the density of a nervous system massive connection pool, technology would not suffice as ground to device the layered structure including the growth of a dendritic tree, which has no foreseeable equivalent in artificial systems.

3.2 What should be copied from the brain?

For the sake of argument, let the evolutionary trajectory of the brain be irrelevant, and the profound physical differences between brain and silicon be somehow

overcome. What would the key elements of the brain structure be to be used for guiding microprocessor design? One may risk to struggle to reproduce in silicon some architectural aspect of the nervous system which is just necessary for some metabolic maintenance, but irrelevant from the computational point of view. The obvious answer should be to take the essential features that make computation in the brain powerful and efficient. The sad point is that nobody has yet been able to identify those features.

The many attempts to relate the complexity of behavior of an organism to macroscopic measures of the brain remain inconclusive. Both weight and size of the brain in vertebrates scale with body size, in a regular way. The relative brain size as a percentage of body size is also an index with scarce relevance, highest in the mouse and the shrew, and average values for primates, human included. Another much-discussed general factor is the encephalization quotient, which indicates the extent to which the brain size of a given species deviates from a sort of standard for the same taxon. This index ranks humans at the top, but is inconsistent with other mammals [62]. The picture is even more complicated when including invertebrates in the comparison [11]. A honeybee's brain has a volume of about 1 cubic millimeter, yet bees display several cognitive and social abilities previously attributed exclusively to higher vertebrates [68]. Even the pure count of neurons lead to puzzling results, for example the elephant brain contains 257 billion neurons, against the 86 billion neurons of the human brain [26].

A suggestion for a designer might be to avoid seeking inspirations from the brain in all its extension and to focus instead on a specific structure that exhibits computational power at the maximum level. A good candidate is the cortex, a milestone in the diversification of the brain throughout evolution, occurred about 200 millions years ago [73]. It is well agreed upon that the mammalian neocortex is the site of the processes enabling higher cognition [50, 23], and it is particularly attractive as a circuit template, for being extremely regular over a large surface.

However, the reason why the particular way neurons are combined in the cortex makes such a difference with respect to the rest of the brain, remains obscure. There are two striking, and at first sight conflicting, features of the cortex:

- the cortex is remarkably uniform in its anatomical structure and in terms of its neuron-level behavior, with respect to the rest of the brain;
- the cortex has an incredibly vast variety of functions, compared to any other brain structure.

The uniformity of the cortex, with the regular repetition of a six-layered radial profile, has given rise to the search of a unified computational model, able to explain its power and its advantages with respect of other neural circuitry. It is the so called “canonical microcircuit”, proposed in several formulations. Marr proposed the first canonical model in a long and difficult paper that is one of his least known works [46], trying to derive an organization at the level of neural circuits from a general theory of how mathematical models might explain classification of sensory signals. The results were far from empirical reality, both at

the time and subsequently through experimental evidence, so they were almost totally neglected and later abandoned by Marr himself. A few years later, Shepherd [69–71] elaborated a model that was both much simpler and more closely related to the physiology of the cortex, compared to Marr’s work. This circuit has two inputs: one from other areas of the cortex making excitatory synapses on dendrites of a superficial pyramidal neuron, and an afferent input terminating on a spiny stellate cell and dendrites of a deep pyramidal neuron. There are feed-forward inhibitory connections through a superficial inhibitory neuron and feedback connections through a basket cell.

Independently, a similar microcircuit was proposed [19], using a minimal set of neuron-like units, disregarding the spiny stellate cells, whose effect is taken into account in two pyramidal-like cells, and the only non-pyramidal neuron is a generic GABA-receptor inhibitory cell. Several more refined circuits were proposed after these first models [56, 18, 57]. Features of the proposed canonical circuits, such as input amplification by recurrent intracortical connections, have been very influential on researchers in neural computation [15]. Yet these ideas are far from explaining why the cortex is computationally so powerful. Specifically, none of the models provide a direct mapping between elements of the circuits and biological counterparts, and there is no corresponding computational description, with functional variables that can be related to these components, or equations that match dependencies posited among components of the cortical mechanism. Moreover, it is fundamentally quite difficult to see how any of these circuits could account for the diversity of cortical functions. Each circuit would seem to perform a single, fixed operation given its inputs, and cannot explain how the same circuit could do useful information processing across diverse modalities.

3.3 Power from flexibility

While the search for a cortical microcircuit advanced modestly in revealing the computational power of the cortex, there is very strong evidence that a key feature of the cortex is the capability to adapt to perform a wide array of heterogeneous functions, starting from roughly similar structures. This phenomenon is called neural plasticity, it comes in different forms, and it has been investigated from a wide variety of perspectives: the reorganization of the nervous system after injuries and strokes [42, 22], in the early development after birth [6, 63], and as the ordinary, continuous way that the cortex works, such as in memory formation [72, 3, 8].

With respect to the circuitual grain level, plasticity can be classified into:

1. *synaptic plasticity*, addressing changes at single synapse level;
2. *intracortical map plasticity*, addressing internal changes at the level of a single cortical area;
3. *intercortical map plasticity*, addressing changes on a scale larger than a single cortical area.

Synaptic plasticity encompasses long-term potentiation (LTP) [5, 1, 4, 2], where an increase in the synaptic efficiency follows repeated coincidences in the timing of the activation of presynaptic and postsynaptic neurons, long-term depression (LTD), the converse process to LTP [30, 80], and spike-timing-dependent (STDP), inducing synaptic potentiation when action potentials occur in presynaptic cell a few milliseconds before those in the postsynaptic site, whereas the opposite temporal order results in long-term depression [41, 44].

Intracortical map plasticity is best known from somatosensory and visual cortical maps, with behavioral modifications such as perceptual learning [20, 60, 65], occurring at all ages, and the main diversification of cortical functions taking place early, in part before birth [14, 36, 37].

Intercortical map plasticity refers to changes in one or more cortical maps induced by factors on a scale larger than a single map's afferents. A typical case is the abnormal developments in primary cortical areas following the loss of sensory inputs, when neurons become responsive to sensory modalities different from their original one [34]. The most dramatic cortical reorganizations are in consequence of a congenital sensory loss, or very early in development, although even in adults significant modifications have been observed [33]. Most of the research in this field is carried out by inducing sensorial deprivation in animals, resulting in an impressive crossmodal plasticity of primary areas [35].

Thus, the overall picture of how the cortex gains its computational power is rather discouraging from the perspective of taking hints for microprocessor design. Very little can be identified as key solutions in the adult cortex, and the plasticity processes, where its real power is hidden, involve organic transformations, like axon growing, neural connection rearrangement, appearance and disappearance of boutons and dendritic spines [27], programmed cell death [81, 58] which are easy to implement in the organic medium, but alien to silicon.

4 Conclusions

There are fundamental features of the brain that remain at present not well known, which may account for the lack of computational models able to encompass the real key aspects of the biological processing power. We believe that these key aspects could be researched in the very nature and complexity of the plasticity mechanisms rather than in the details of the neural unit processing alone. Until a theoretical framework emerges able to capture the essential aspects of neural plasticity and an appropriate technology able to mimic it is devised, the quest for the "brain in silicon" could be severely impaired. We are agnostic concerning the future of neurocomputers, our point is that the justification put forward for its realizability is scientifically flawed, and it may be the cause of the scarce success met so far.

References

1. Artola, A., Singer, W.: Long term potentiation and NMDA receptors in rat visual cortex. *Nature* 330, 649–652 (1987)

2. Bear, M., Kirkwood, A.: Neocortical long term potentiation. *Current Opinion in Neurobiology* 3, 197–202 (1993)
3. Bermúdez-Rattoni, F. (ed.): *Neural plasticity and memory : from genes to brain imaging*. CRC Press, Boca Raton (FL) (2007)
4. Bliss, T., Collingridge, G.: A synaptic model of memory: long-term potentiation in the hippocampus. *Nature* 361, 31–39 (1993)
5. Bliss, T., Lømo, T.: Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path. *Journal of Physiology* 232, 331–356 (1973)
6. Blumberg, M.S., Freeman, J.H., Robinson, S. (eds.): *Oxford Handbook of Developmental Behavioral Neuroscience*. Oxford University Press, Oxford (UK) (2010)
7. Boden, M.: *Mind as Machine: A History of Cognitive Science*. Oxford University Press, Oxford (UK) (2008)
8. Bontempi, B., Silva, A., Christen, Y. (eds.): *Memories: Molecules and Circuits*. Springer-Verlag, Berlin (2007)
9. Brown, P.: The mechanization of art. In: Husbands, P., Holland, O., Wheeler, M. (eds.) *The Mechanical Mind in History*, pp. 259–281. The Guilford Press, New York (2008)
10. Carnap, R.: *Der logische Aufbau der Welt*. Weltkreis Verlag, Berlin-Schlactensee (1928)
11. Chittka, L., Niven, J.: Are bigger brains better? *Current Biology* 19, R995–R1008 (2009)
12. Copeland, J., Proudfoot, D.: On Alan Turing’s anticipation of connectionism. *Synthese* 108, 361–377 (1996)
13. Cordeschi, R.: *The discovery of the artificial – Behavior, Mind and Machines Before and Beyond Cybernetics*. Springer-Verlag, Berlin (2002)
14. Crair, M.C.: Neuronal activity during development: permissive or instructive? *Current Opinion in Neurobiology* 9, 88–93 (1999)
15. Dayan, P., Abbott, L.F.: *Theoretical Neuroscience*. MIT Press, Cambridge (MA) (2001)
16. Descartes, R.: *Discours de la méthode*. Ian Maire, Leyde (1637)
17. Dias, F.M., Antunes, A., Mota, A.M.: Artificial neural networks: a review of commercial hardware. *Engineering Applications of Artificial Intelligence* 17, 945–952 (2004)
18. Douglas, R.J., Martin, K.A.: Neuronal circuits of the neocortex. *Annual Review of Neuroscience* 27, 419–451 (2004)
19. Douglas, R.J., Martin, K.A., Whitteridge, D.: A canonical microcircuit for neocortex. *Neural Computation* 1, 480–488 (1989)
20. Fahle, M., Poggio, T. (eds.): *Perceptual Learning*. MIT Press, Cambridge (MA) (2002)
21. Fodor, J.: Special sciences (or: The disunity of science as a working hypothesis). *Synthese* 28, 77–115 (1974)
22. Fuchs, E., Flügge, G.: Adult neuroplasticity: More than 40 years of research. *Neural Plasticity* 2014, ID541870 (2014)
23. Fuster, J.M.: *The Prefrontal Cortex*. Academic Press, New York (2008), fourth edition
24. Hasler, J., Marr, B.: Finding a roadmap to achieve large neuromorphic hardware systems. *Frontiers in Neuroscience – Neuromorphic Engineering* 7, 118 (2013)
25. Heemskerk, J.N.H.: *Overview of Neural Hardware – Neurocomputers for Brain-Style Processin – Design, Implementation and Application*. Ph.D. thesis, Unit of Experimental and Theoretical Psychology, Leiden University (1995)

26. Herculano-Houzel, S., de Souza, K.A., Neves, K., Porfirio, J., Messeder, D., Feijó, L.M., Maldonado, J., Manger, P.R.: The elephant brain in numbers. *Frontiers in Neuroanatomy* 8, Article 46 (2014)
27. Holtmaat, A., Svoboda, K.: Experience-dependent structural synaptic plasticity in the mammalian brain. *Nature Reviews Neuroscience* 10, 647–658 (2009)
28. Hornik, K., Stinchcombe, M., White, H.: Multilayer feedforward networks are universal approximators. *Neural Networks* 2, 359–366 (1989)
29. Ienne, P.: Digital connectionist hardware: Current problems and future challenges. In: Mira, J., Moreno-Díaz, R., Cabestany, J. (eds.) *Biological and Artificial Computation: From Neuroscience to Technology*, pp. 688–713. Springer-Verlag, Berlin (1997)
30. Ito, M.: Long-term depression. *Annual Review of Neuroscience* 12, 85–102 (1989)
31. Izhikevich, E.M.: Simple model of spiking neurons. *IEEE Transactions on Neural Networks* 14, 1569–1572 (2003)
32. Jin, X., Lujan, M., Plana, L.A., Davies, S., Temple, S., Furber, S.: Modeling spiking neural networks on SpinNaker. *Computing in Science & Engineering* 12, 91–97 (2010)
33. Kaas, J.H.: Plasticity of sensory and motor maps in adult mammals. *Annual Review of Neuroscience* 14, 137–167 (1997)
34. Karlen, S.J., Hunt, D.L., Krubitzer, L.: Cross-modal plasticity in the mammalian neocortex. In: Blumberg et al. [6], pp. 357–374
35. Karlen, S.J., Kahn, D., Krubitzer, L.: Early blindness results in abnormal corticocortical and thalamo cortical connections. *Neuroscience* 142, 843–858 (2006)
36. Khazipov, R., Buzsáki, G.: Early patterns of electrical activity in the developing cortex. In: Blumberg et al. [6], pp. 161–177
37. Khazipov, R., Colonnese, M., Minlebaev, M.: Neonatal cortical rhythms. In: Rubenstein and Rakic [63], pp. 131–153
38. Kline, R.R.: *The Cybernetics Moment – Or Why We Call Our Age the Information Age*. Johns Hopkins University Press, Baltimore (2015)
39. de La Mettrie, J.O.: *L’Homme Machine*. Elie Luzac, Leyden (1748)
40. Lande, T.S.: *Neuromorphic Systems Engineering - Neural Networks in Silicon*. Kluwer, Dordrecht (NL) (1998)
41. Levy, W., Steward, O.: Temporal contiguity requirements for long-term associative potentiation/depression in the hippocampus. *Neuroscience* 8, 791–797 (1983)
42. Lövdén, M., Bäckman, L., Lindenberger, U., Schaefer, S., Schmiedek, F.: A theoretical framework for the study of adult cognitive plasticity. *Psychological Bulletin* 136, 659 – 676 (2010)
43. Markram, H.: The blue brain project. *Nature Reviews Neuroscience* 7, 153–160 (2006)
44. Markram, H., Lübke, J., Frotscher, M., Sakmann, B.: Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science* 275, 213–215 (1997)
45. Markram, H., Muller, E., Ramaswamy, S., et al., M.W.R.: Reconstruction and simulation of neocortical microcircuitry. *Cell* 163, 456–492 (2015)
46. Marr, D.: A theory for cerebral neocortex. *Proceedings of the Royal Society of London B* 176, 161–234 (1970)
47. McCorduck, P.: *Machines Who Think: A Personal Inquiry into the History and Prospect of Artificial Intelligence*. Freeman, San Francisco (1979)
48. McCulloch, W., Pitts, W.: A logical calculus of the ideas immanent in nervous activity. *Bulletin of Mathematical Biophysics* 5, 115–133 (1943)

49. Merolla, P.A., Arthur, J.V., Alvarez-Icaza, R., Cassidy, A.S., et al., J.S.: A million spiking-neuron integrated circuit with a scalable communication network and interface. *Science* 345, 668–673 (2014)
50. Miller, E.K., Freedman, D.J., Wallis, J.D.: The prefrontal cortex: Categories, concepts and cognition. *Philosophical Transactions: Biological Sciences* 357, 1123–1136 (2002)
51. Minsky, M., Papert, S.: *Perceptrons*. MIT Press, Cambridge (MA) (1969)
52. Minsky, M.L.: *Neural nets and the brain-model problem*. Ph.D. thesis, Princeton University (1954)
53. Misra, J., Saha, I.: Artificial neural networks in hardware: A survey of two decades of progress. *Neurocomputing* 74, 239–255 (2010)
54. Modha, D.S., Ananthanarayanan, R., Esser, S.K., Ndirango, A., Sherbondy, A.J., Singh, R.: Cognitive computing: unite neuroscience, supercomputing, and nanotechnology to discover, demonstrate, and deliver the brain’s core algorithms. *Communications of the Association for Computing Machinery* 54, 62–71 (2011)
55. Moroz, L.L.: On the independent origins of complex brains and neurons. *Brain, Behavior and Evolution* 74, 177–190 (2009)
56. Nieuwenhuys, R.: The neocortex. *Anatomy and Embryology* 190, 307–337 (1994)
57. Nieuwenhuys, R., Voogd, J., van Huijzen, C.: *The Human Central Nervous System*. Springer-Verlag, Berlin (2008)
58. Oppenheim, R.W., Milligan, C., Sun, W.: Programmed cell death during nervous system development: Mechanisms, regulation, functions, and implications for neurobehavioral ontogeny. In: Blumberg et al. [6], pp. 76–107
59. Pitts, W., McCulloch, W.: How we know universals: The perception of auditory and visual forms. *Bulletin of Mathematical Biophysics* 9, 115–133 (1947)
60. Roelfsema, P.R., van Ooyen, A., Watanabe, T.: Perceptual learning rules based on reinforcers and attention. *Trends in Cognitive Sciences* 14, 64–71 (2009)
61. Rosenblatt, F.: The perceptron: a probabilistic model for information storage and organisation in the brain. *Psychological Review* 65, 386–408 (1958)
62. Roth, G., Dicke, U.: Evolution of nervous systems and brains. In: Galizia, G., Lledo, P.M. (eds.) *Neurosciences – From Molecule to Behavior*, pp. 19–45. Springer-Verlag, Berlin (2013)
63. Rubenstein, J.L.R., Rakic, P. (eds.): *Comprehensive developmental neuroscience: neural circuit development and function in the healthy and diseased brain*. Academic Press, New York (2013)
64. Rumelhart, D.E., McClelland, J.L. (eds.): *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*. MIT Press, Cambridge (MA) (1986)
65. Sasaki, Y., Nanez, J.E., Watanabe, T.: Advances in visual perceptual learning and plasticity. *Nature Reviews Neuroscience* 11, 53–60 (2010)
66. Schemmel, J., Brüderle, D., Grübl, A., Hock, M., Meier, K., Millner, S.: A wafer-scale neuromorphic hardware system for large-scale neural modeling. In: *Proceedings of IEEE International Symposium on Circuits and Systems*. pp. 1947–1950 (2010)
67. Schürmann, F., Delalondre, F., Kumbhar, P.S., Biddiscombe, J., Miguel Gila, e.a.: Rebasement I/O for Scientific Computing: Leveraging Storage Class Memory in an IBM BlueGene/Q Supercomputer. No. 8488 in *Lecture Notes in Computer Science*, Springer-Verlag, Berlin (2014)
68. Seeley, T.D.: What studies of communication have revealed about the minds of worker honey bees. In: Kikuchi, T., Azuma, N., Higashi, S. (eds.) *Genes, Behaviors and Evolution of Social Insects*, pp. 21–33. Hokkaido University Press, Sapporo (2003)

69. Shepherd, G.M.: *The Synaptic Organization of the Brain*. Oxford University Press, Oxford (UK) (1974)
70. Shepherd, G.M.: *The Synaptic Organization of the Brain*. Oxford University Press, Oxford (UK) (1979), second Edition
71. Shepherd, G.M.: A basic circuit for cortical organization. In: Gazzaniga, M.S. (ed.) *Perspectives on Memory Research*, pp. 93–134. MIT Press, Cambridge (MA) (1988)
72. Squire, L., Kandel, E.: *Memory: from Mind to Molecules*. Scientific American Library, New York (1999)
73. Striedter, G.F.: *Principles of Brain Evolution*. Sinauer Associated, Sunderland, MA (2003)
74. Sugino, K., Hempel, C.M., Miller, M.N., Hattox, A.M., Shapiro, P., Wu, C., Huang, Z.J., Nelson, S.B.: Molecular taxonomy of major neuronal classes in the adult mouse forebrain. *Journal of Cognition and Culture* 6, 181–189 (2006)
75. Taylor, J.: *The Promise of Neural Networks*. Springer-Verlag, Berlin (1993)
76. Turing, A.: *Intelligent machinery*. Tech. rep., National Physical Laboratory, London (1948), raccolto in Ince, D. C. (ed.) *Collected Works of A. M. Turing: Mechanical Intelligence*, Edinburgh University Press, 1969
77. Versace, M., Chandler, B.: Meet MoNETA the brain-inspired chip that will outsmart us all. *IEEE Spectrum* 12, 30–37 (2010)
78. Wiener, N.: *Cybernetics, or Control and Communication in the Animal and the Machine*. MIT Press, Cambridge (MA) (1948)
79. Zakon, H.H.: Adaptive evolution of voltage-gated sodium channels: The first 800 million years. *Proceedings of the Natural Academy of Science USA* 109, 10619–10625 (2012)
80. Zhuo, M., Hawkins, R.D.: Long-term depression: A learning-related type of synaptic plasticity in the mammalian central nervous system. *Reviews in the Neurosciences* 6, 259–277 (1995)
81. Zou, D., Feinstein, P., Rivers, A., Mathews, G., Kim, A., Greer, C.: Postnatal refinement of peripheral olfactory projections. *Science* 304, 1976–1979 (2004)