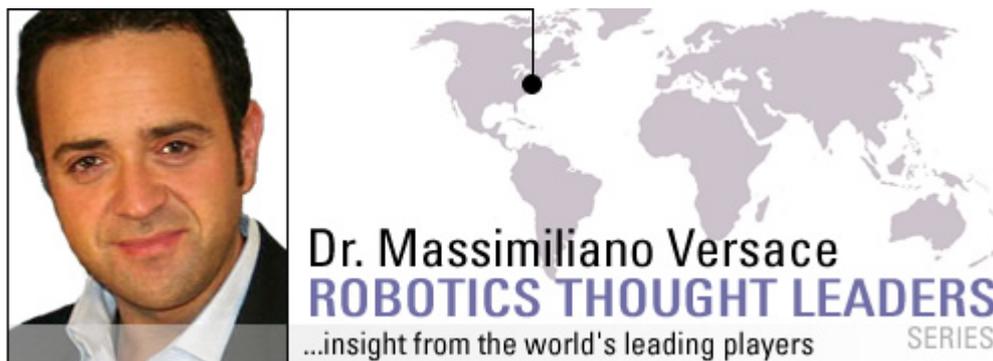


Silicon Brains



By [Dr Massimiliano Versace](#)

[Dr Massimiliano Versace](#), Director, Neuromorphics Lab, Boston University. Corresponding author: maxversace@gmail.com

"We are close to building silicon brains": If you have even a glancing interest in the field of Artificial Intelligence (AI), you would have heard this promise countless times, and seen it broken as often. Even before there were machines, there was the Golem, an animated anthropomorphic being from Jewish folklore, Frankenstein's monster, created by harnessing the power of lightning. These were to be followed by mechanized robots, HAL, and Skynet, among others

In a slightly less theatrical setting, in 1997 Garry Kasparov faced off with the IBM supercomputer, Deep Blue. After a series of favorable games, Kasparov conceded to Deep Blue in the sixth game, prompting the AI community to declare victory. Unfortunately, Deep Blue was far from intelligent: a computer with software ad-hoc designed for rapid calculations of chess moves (I think you agree with me that playing chess requires intelligence, but the opposite is not true). You may think that the lesson was learned, and that AI was not going to be confused again with the notion of "special intelligence". Think again: the dream of an intelligence machine was recently resurrected by the same company, IBM, who engineered a new "human brain", Watson, able to defeat Jeopardy champion Ken Jennings. The examples of special purpose intelligent machines are as fascinating as deceiving: despite their success in specific, very carefully delimited domains, the promise of a general purpose machine with 'human-like' intelligence has yet to be realized.

The reason is simple: building intelligence comparable to that of a human brain in silicon requires an understanding of how the interactions of billions of neurons and synapses give rise to complex behaviors, as well as the ability to replicate those interactions in sophisticated, powerful, and low-power computer hardware. Researchers in fields ranging from neuroscience, psychology, computer science, material science, and robotics have been steadily working to uncover the core design principles underlying neural computation and how this controls behavior, and inventing key technologies needed to build smart machines. To achieve this goal, innovations in fields ranging from computational modeling, hardware, and software design need to converge. This convergence is happening now.

Modeling whole-brain systems

In building whole brain models for virtual animats (artificial animals) and robots, the core challenges include how to autonomously plan, explore, perceive, and understand new environments while both maintaining a sense of the agent's localization and safety within that environment and creating learned memories about the objects and affordances (threat or goal) within the environment. To date, animal capacities in such tasks far exceed those of autonomous machines. Researchers at Boston University (BU)'s [Neuromorphics Lab](#), part of the NSF-sponsored Center of Excellence for Learning in Education, Science and Technology ([CELEST](#)), are translating decades of basic neuroscience research and computational modeling in a whole-brain system that will bring us closer to this goal. Our approach to building the neural models that power the artificial brain is modular. The macro-structure of the brain is initially specified with the goal of being able to "swap in" more refined neural circuits when they become available.

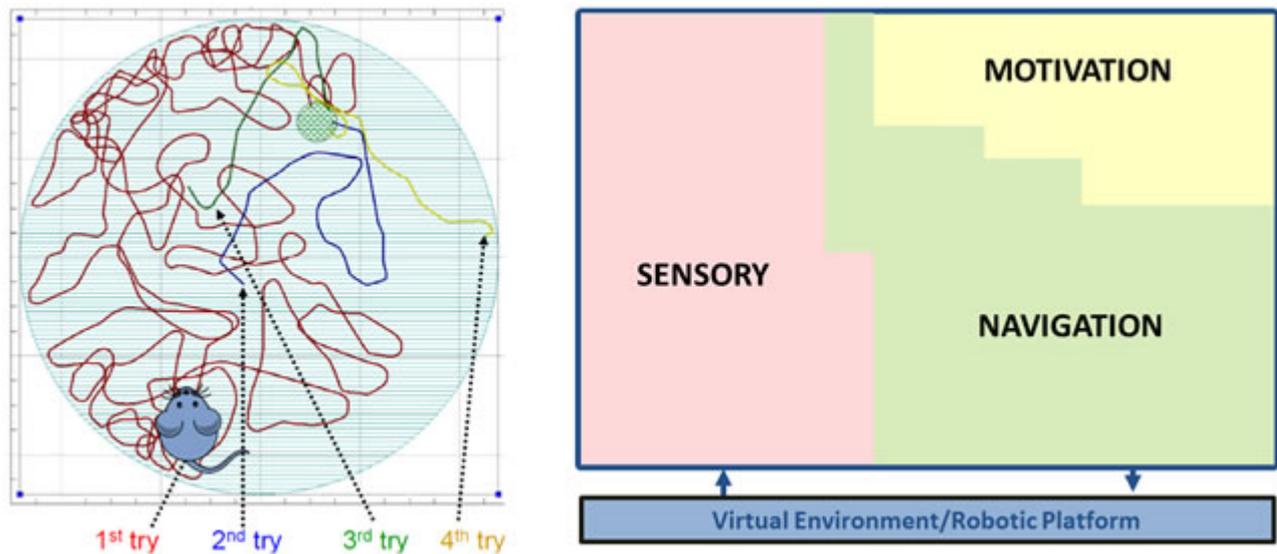


Figure 1. *Left:* MoNETA is a whole-brain-system model consisting of sensory, motivation, and navigation areas interfaced with a virtual environment or robotic platform. *Right:* MoNETA learns to perform the Morris Water Maze task. In the first trial, MoNETA explores its environment driven by lack of comfort and curiosity drives. Once it accidentally swims on top of the submerged platform (green), MoNETA learns its position by using some visual landmark at the border of the pool. As training progresses, MoNETA is able to swim directly from its current position to the platform. MoNETA integrates sensory information into higher order representations of its emerging reality, and is able to react to novel situations not explicitly programmed within the software. It can perceive its surroundings, decide what is useful, and in certain applications even be able to formulate plans that ensure its survival. MoNETA is motivated by analogs of the same drives that underlie the behavior of rats, cats, or humans.

This whole brain system is called MoNETA¹ (Modular Neural Exploring Traveling Agent, Figure 1, left), and its first version has been tested in a virtual Morris Water Maze task (Figure 1, right). The Morris Water Maze is a task used to probe the navigation skills of a rodent. The rat is placed in a water tank and has to use visual cues in order to locate submerged platform and swim to it. The rat is motivated to find the platform to get out of the water tank. Researchers have studied this task at great length, so we know a great deal of the brain areas a rat utilizes in completing the task. Although an apparently simple one, solving the water maze actually requires that the integrated functioning of object recognition and localization, goal selection, and navigation be simulated in several interacting brain areas. To build and simulate MoNETA, new software tools and hardware are needed that implement these models at biological scale.

Software and hardware infrastructure

Building artificial brains in silicon is a huge task, and there is a big divide between computing in neurons and silicon (see Figure 2).

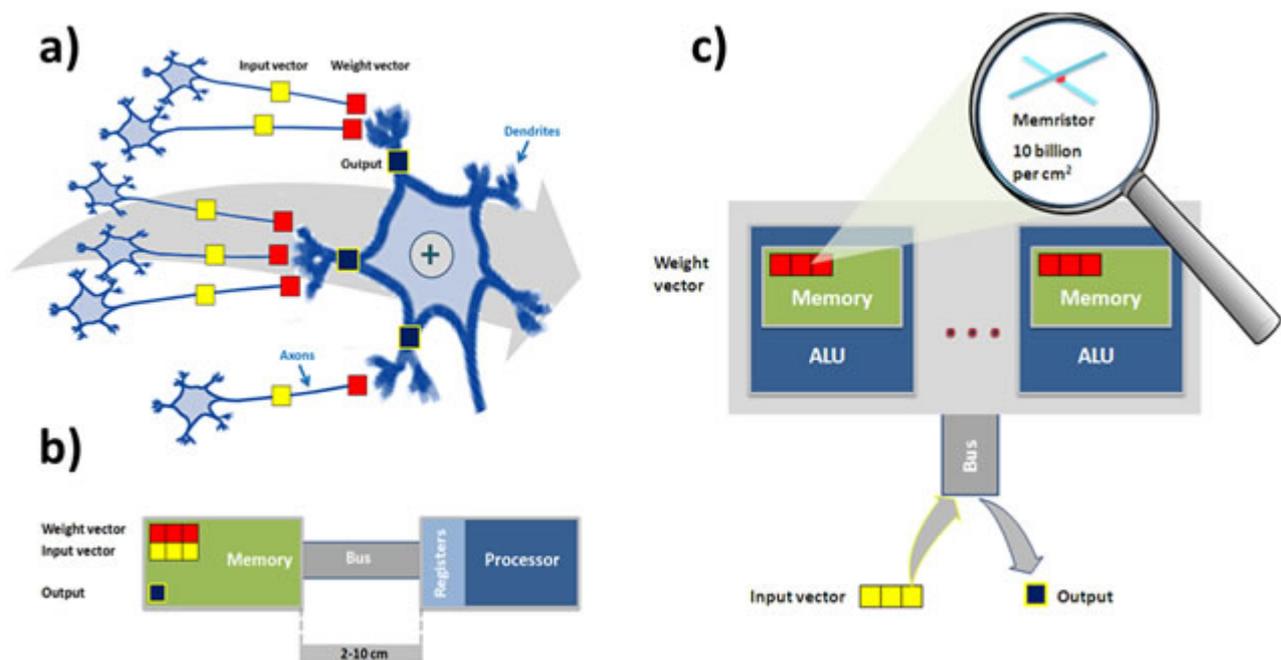


Figure 2. Computing in biology and in silicon. **a)** Biological computation is deeply distributed, with data and computation tightly interwoven. Axons carry the input vector from neighboring neurons to the synapses of the target neuron. Each synapse stores its state locally and computes the product of its input with its state. These individual products flow towards the body of the neuron over the dendrites where the signal is implicitly summed. For an idealized neuron, by the time the combined signals reach the body of the neuron the synapses and dendrites have produced a single scalar value. **b)** In a conventional computing scheme, the memory and processor are physically separated by a fixed-capacity bus. The processor only contains a small number of slots for storing data during computation, called registers. The processor keeps the running total in a register, starting from zero. At each step, the processor loads the next element from the weight vector and the next element from the input vector and adds the product to the running total. After running through all the elements, the processor writes the result back to memory. For computing intelligent algorithms, accessing the bus is highly inefficient primarily because of the distance, both physical and electronic. **c)** Future memristive-based processors will have several architectural features that allow it to compute a dot product much more efficiently than a conventional processor. Most critically, the weights will be stored in a high-density memristive memory directly on top of the multiprocessor cores, decreasing the distance data must travel and increases memory bandwidth. Compared to a conventional processor, these devices will contain many simple cores instead of a few complex cores.

Consider, for example, one dimension of the problem: how to translate synapses into their electronic equivalent. Biological synapses are dense—the cortex needs roughly 10^{10} synapses per square centimeter, and they consume miniscule power; have complex, nonlinear dynamics; and, in some cases, can maintain their memory for decades. Until recently, these characteristics translated to one more unreachable goal for those aspiring to build electronic brains, particularly large models. In the past few years, however, work on memristive devices² has gained momentum which could bring designers closer to an electronic brain architecture that can adaptively interact with the world in real time.

Memristive devices are nanoscale electrical components that have several attractive features: nonlinear resistance that can be altered electrically, they can be packed into extremely dense crossbars, and they are compatible with CMOS processes. The last characteristic allows designers to integrate dense, memristive memories with conventional circuit and thus place memory and computation closer together, decreasing power dissipation and allowing to implement dense synaptic-like memories in conventional CMOS processors.

The problem remains on how to program neural algorithms in these massively parallel processors with dense local memories. To bootstrap this process of building intelligent hardware, Hewlett-Packard and the Boston University

Neuromorphics Lab are jointly developing the Cog Ex Machina³. Along with the memristive-based hardware, the Cog software provide a low-cost, flexible all-digital platform for building large brain models that can interact with a simulated or real environment in real time. This is one of the seven key steps needed to build silicon brains.

Bringing it all together: building silicon brains in seven steps

Evolution has taken millions of years, and billions of organisms to engineer the 'wetware' that powers our perception: decisions and emotions. While biological brains have this unfair head start over their silicon counterparts, progress in basic research and emerging computing technologies can help to substantially close this gap, as shown in Figure 3.

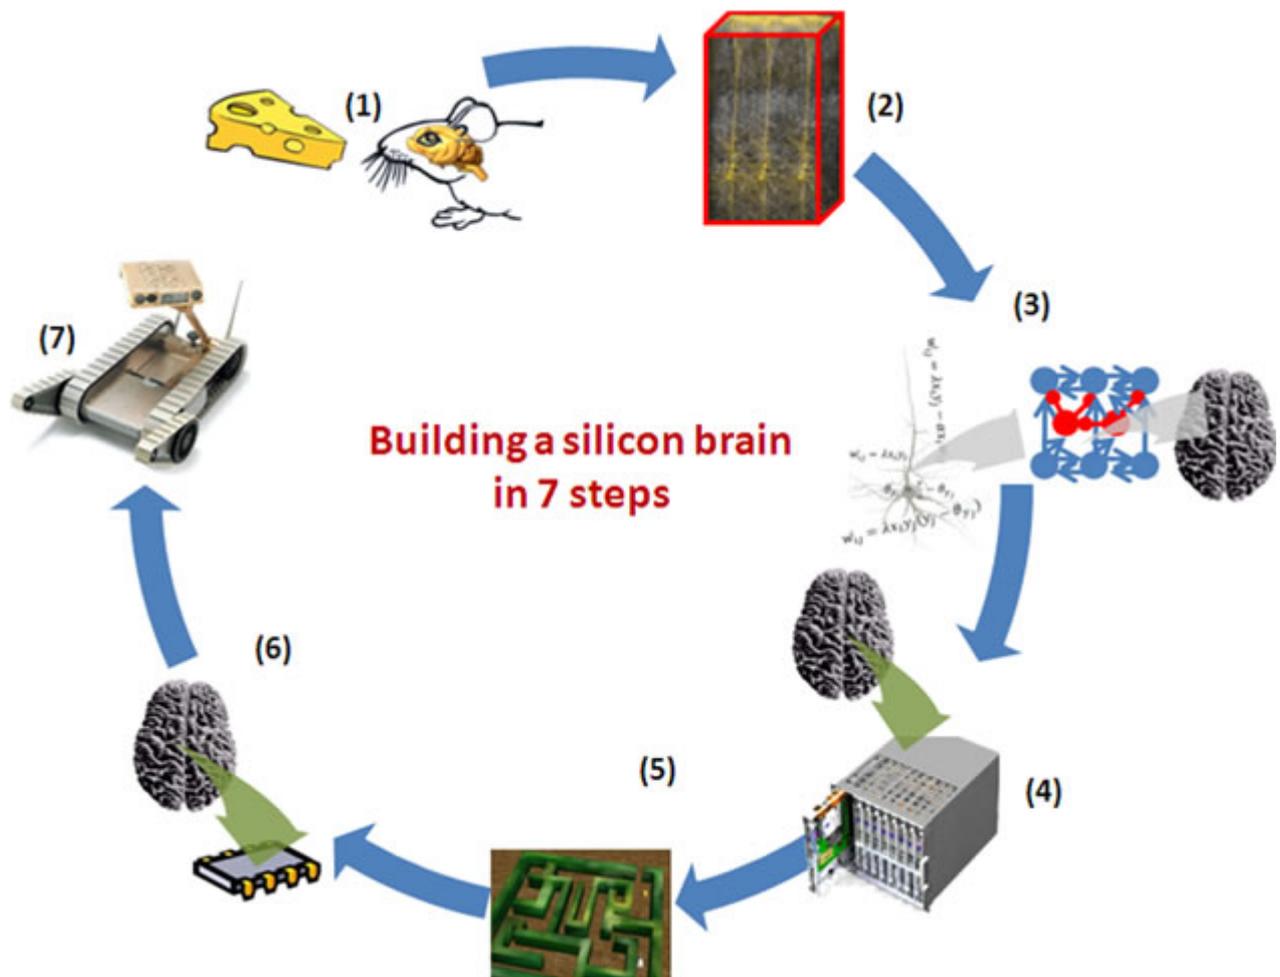


Figure 3. Building a silicon brain in in seven steps.

Basic neuroscience research studies how brain circuits give rise to behavior (1), and isolates the main computations of these circuits (2). These are translated into mathematical modeling of neurons, networks, and whole brain systems (3). This step paves the way to implementation of these models in software, which can then be run on large clusters of GPUs (4) and simulate virtual agents behaving in virtual environments (5). These models can be then run on a much smaller, denser, portable, and low power hardware (6) that can then power mobile robotic platforms (7).

Building silicon brains requires coordinated advances in several fields to enable the development of biological-grade intelligence in machines. One can hope that the next time a chess grandmaster or a Jeopardy player sits down to a game with an intelligent computer, the computer may win or it may lose, but its behavior, emotions, and decisions will be indistinguishable from its human counterpart. And that will be a true victory for AI!

References

1. Versace M. and Chandler B. (2010) MoNETA: A Mind Made from Memristors. Cover feature, *IEEE Spectrum*, December 2011.
2. Strukov D.B., Snider G.S, Stewart D.R., and Williams R.S. (2008) The missing memristor found. *Nature* 453, 80-8
3. Snider G., Amerson R., Carter D., Abdalla H., Qureshi S., Leveille J., Versace M., Ames H., Patrick S., Chandler B., Gorchetchnikov A., and Mingolla E. (2011) Adaptive Computation with Memristive Memory. □ Cover feature, *IEEE Computer* 44(2), 21-28.
Date Added: Jul 18, 2011 | Updated: Jul 18, 2011