THE NEW YORKER

ELEMENTS

A COMPUTER TO RIVAL THE BRAIN

By Kelly Clancy February 15, 2017

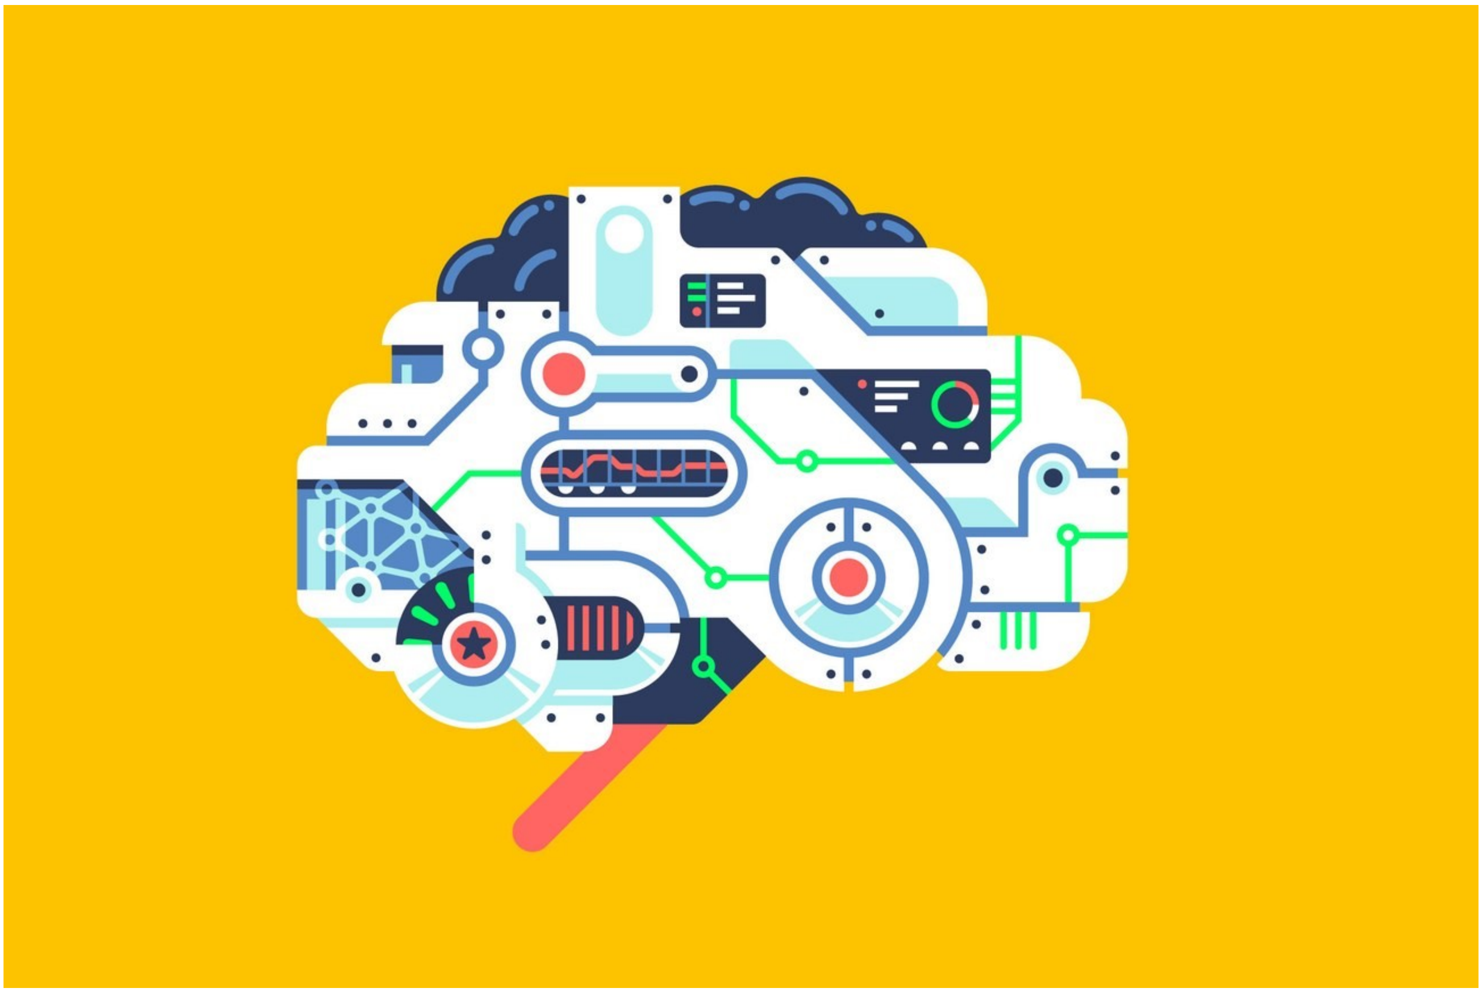

Artificial intelligence has achieved much of its recent success by mimicking biology. Now it must go further.

ILLUSTRATION BY GORAN FACTORY

More than two hundred years ago, a French weaver named Joseph Jacquard invented a mechanism that greatly simplified textile production. His design replaced the lowly draw boy—the young apprentice who meticulously chose which threads to feed into the loom to create a particular pattern—with a series of paper punch cards, which had holes dictating the lay of each stitch. The device was so successful that it was repurposed in the first interfaces between humans and computers; for much of the twentieth century, programmers laid out their code like weavers, using a lattice of punched holes. The cards themselves were fussy and fragile. Ethereal information was at the mercy of its paper substrate, coded in a language only experts could understand. But successive computer interfaces became more natural, more flexible. Immutable program instructions were softened to "If *x*, then *y*. When *a*, try *b*." Now, long after Jacquard's invention, we simply ask Amazon's Echo to start a pot of coffee, or Apple's Siri to find the closest car wash. In order to make our interactions with machines more natural, we've learned to model them after ourselves.

Early in the history of artificial intelligence, researchers came up against what is referred to as Moravec's paradox: tasks that seem laborious to us (arithmetic, for example) are easy for a computer, whereas those that seem easy to us (like picking out a friend's voice in a noisy bar) have been the hardest for A.I. to master. It is not profoundly challenging to design a computer that can beat a human at a rule-based game like chess; a logical machine does logic well. But engineers have yet to build a robot that can hopscotch. The Austrian roboticist Hans Moravec theorized that this might have something to do with evolution. Since higher reasoning has only recently evolved—perhaps within the last hundred thousand years—it hasn't had time to become optimized in humans the way that locomotion or vision has. The things we do best are largely unconscious, coded in circuits so ancient that their calculations don't percolate up to our experience. But because logic was the first form of biological reasoning that we could perceive, our thinking machines were, by necessity, logic-based.

Computers are often likened to brains, but they work in a manner foreign to biology. The computing architecture still in use today was first described by the mathematician John von Neumann and his colleagues in 1945. A modern laptop is conceptually identical to the punch-card behemoths of the past, although engineers have traded paper for a purely electric stream of on-off signals. In a von Neumann machine, all data-crunching happens in the central processing unit (C.P.U.). Program instructions, then data, flow from the computer's memory to its C.P.U. in an orderly series of zeroes and ones, much like a stack of punch cards shuffling through. Although multicore computers allow some processing to occur in parallel, their efficacy is limited: software engineers must painstakingly choreograph these streams of information to avoid catastrophic system errors. In the brain, by contrast, data run simultaneously through billions of parallel processors—that is, our neurons. Like computers, they communicate in a binary language of electrical spikes. The difference is that each neuron is pre-programmed, whether through genetic patterning or learned associations, to share its computations directly with the proper targets. Processing unfolds organically, without the need for a C.P.U.

Consider vision. We sense the world with an array of millions of photoreceptors, each of which plays a small and specific role in representing an image with neural activity. These cells shuttle the representation through a hierarchy of brain areas, progressively forming the conscious percept of sight. A von Neumann computer would have to stream that same amount of data, plus the instructions to process it, through a single logical core. And though a computer's circuits move data much faster than the brain's synapses, they consume a large amount of energy in doing so. In 1990, the legendary Caltech engineer Carver Mead correctly predicted that our present-day computers would use ten million times more energy for a single instruction than the brain uses for a synaptic activation.

A.I. owes much of its recent success to biological metaphors. Deep learning, for example, which underlies technologies from Siri to Google Translate, uses several interconnected processing layers, modelled after the neuronal strata that compose the cortex. Still, given that even the most advanced neural networks are run on von Neumann machines, they are computationally intensive and energy-greedy. Last March, AlphaGo, a program created by Google DeepMind, was able to beat a world-champion human player of Go, but only after it had trained on a database of thirty million moves, running on approximately a million watts. (Its opponent's brain, by contrast, would have been about fifty thousand times more energy-thrifty, consuming twenty watts.) Likewise, several years ago, Google's brain simulator taught itself to identify cats in YouTube videos using sixteen thousand core processors and all the wattage that came with them. Now companies want to endow our personal devices with intelligence, to let our smartphones recognize our family members, anticipate our moods, and suggest adjustments to our medications. To do so, A.I. will need to move beyond algorithms run on supercomputers and become embodied *in silico*.

Building on decades of work by Mead and others, engineers have been racing to roll out the first so-called neuromorphic chips for consumer use. Kwabena Boahen's research group at Stanford unveiled its low-power Neurogrid chip in 2014, and Qualcomm has announced that its brain-inspired Zeroth processor will reach the market in 2018. Another model, I.B.M.'s TrueNorth, only recently moved from digital prototype to usable product. It consists of a million silicon neurons, tiny cores that communicate directly with one another using synapse-like connections. Here, the medium is the message; each neuron is both program and processing unit. The sensory data that the chip receives, rather than marching along single file, fan out through its synaptic networks. TrueNorth ultimately arrives at a decision—say, classifying the emotional timbre of its user's voice—by group vote, as a choir of individual singers might strike on a harmony. I.B.M. claims the chip is useful in real-time pattern recognition, as for speech processing or image classification. But the biggest advance is its energy efficiency: it uses twenty milliwatts per square centimetre, more than a thousand times less than a traditional chip.

TrueNorth was also designed to emulate some of the brain's messiness. For the past several billion years, life has had to learn to make do with its own imperfect corporeity —fuzzy eyesight, limited hearing, and so on. Despite sensing the world through a scrim of unpredictable molecular interactions, though, organisms tend to get around with remarkable accuracy. What seems like a bug may be, mathematically speaking, a feature. Randomness turns out to add a great deal of computational power to probabilistic algorithms like the ones underlying modern A.I.; input noise can shake up their output, preventing them from getting stuck on bad solutions. TrueNorth creates its own sort of fuzziness by including a random-number generator with each neuron. I.B.M. is developing another chip that achieves the same goal more elegantly, using a material that changes phase from amorphous to crystalline with a certain degree of randomness. And this is the crux of the conceptual shift that is taking place in computing: increasingly, engineers will exploit the computational properties of matter rather than guarding against its inherent fallibility, as they had to do with the punch cards. Matter will not execute a computation; it will be the computation.

Given the utter lack of consensus on how the brain actually works, these designs are more or less cartoons of what neuroscientists think might be happening. But, even if they don't reflect absolute biological reality, the recent success of A.I. suggests that they are useful cartoons. Indeed, they may eventually confirm or challenge our understanding of the brain; as the physicist Richard Feynman put it, "What I cannot create, I do not understand." Or perhaps their power lies in their simplicity. Eve Marder, a neuroscientist at Brandeis University, has argued that the more details we include in our models, the more wrong we may make them—such is the complexity of neurobiology and the depth of our ignorance. Strict fidelity may not be necessary in designing practical A.I. TrueNorth, for instance, can't learn on its own. The chip has to be optimized for a particular task using A.I. run on a conventional computer. So, though TrueNorth maintains one part of the biological metaphor, it does so at the cost of another. And perhaps there's nothing wrong with that. Who is to say that every feature of the brain is worth mimicking? Our own human algorithms are not necessarily ideal. As Darwin demonstrated, evolution is not an unremitting race toward perfection. It is a haphazard wander around good enough.

Kelly Clancy is a neuroscientist based in Basel, Switzerland. *Read more »*