NEUROSCIENCE IN 2064

A LOOK AT THE LAST CENTURY

As told to

Christof Koch and Gary Marcus

On a warm summer evening, not long ago, while we, the authors, were engaged in a spirited debate about the nature of consciousness, a traveler, going by the name of Lem, appeared, claiming to be from the future; at first, we were skeptical. But his recollections were vivid, and detailed, and more than that, internally consistent. Try as we could, we couldn't break his story; he claimed to be from the year 2064, and his knowledge of neuroscience seemed to be exceptional. Over time, we began to believe that his reports were authentic; in what is below, we have transcribed his story as near as we can recall it.

One hundred years ago, in 1964, the United States and the Soviet Union were jockeying for world supremacy, "computers" still meant human beings, trained to carry out long chains of calculations, and gas guzzlers dominated the highways. Global warming and nanotechnology were not even in the vocabulary, and a British band known as The Beatles had just arrived in America.

What difference one hundred years make! Extreme weather and greatly diminished fossil fuels, the decline of the American and Russian empires, the rise of the Chinese Dragon, and the widespread intrusion of artificial intelligence agents into daily life has transformed the stable, dichotomous Cold War world of 1964 into a more splintered world, vibrant yet at the edge of chaos in its own way. Some of us live longer and more healthfully than our ancestors, as dozens of once-deadly diseases have been cured. Yet the bulk of mankind still lives less than four score and ten years; and the promises of trans-humanists to extend the maximal life span past 120 years have thus far proved illusory.

Molecular biology has finally delivered on the early promises of the Human Genome Project, albeit decades later than forecast. Previous monolithic diseases, such as breast cancer, brain cancer, depression, dementia and autism, have splintered off into a myriad of more specific pathologies, defined not so much by common behavioral phenotypes but by shared mutations, molecular pathways and biochemical mechanisms. In combination with cheap, reliable, and fast genetic tests, the age of personalized medicine, long trumpeted by Leroy Hood, Craig Venter, and other pioneers, arrived in which familial predispositions to behavioral traits, pharmacological interventions, and diseases, permit much more targeted interventions.

Bioterrorism has occasionally struck, but the combination of personal genomics, personal immunizers, and a ubiquitous surveillance state has largely kept the population safe.

Advances in the brain sciences have been in many ways even more impressive; a hundred years ago, humanity knew that the brain—and not the heart or liver—was the seat of the mind, but little about how neural tissue governed perception, comprehension, or consciousness; brain-machine interfaces, now common, did not even figure in the most popular science fiction television program of the day (*Star Trek*). If our understanding of neuroscience is still incomplete, it is shocking how much progress there has been. Yet one also forgets that the seeds for our modern understanding were already in place.

The Romantic Era of Neuroscience: 1964

The first blossoming of the romantic era in neuroscience started almost two centuries ago. It was powered by two technologies, the *optical microscop*e and the refinement of *chemical dyes*, in particular Golgi's staining method of using silver chromate salt. Together, these allowed Santiago Ramón y Cajal to visualize in stunning detail the circuitry of the nervous systems in animals and people, demonstrating in aesthetically pleasing images that brains, like kidneys, hearts and all other biological organs, are composed of a myriad of discrete, cellular units, neurons, and their supporting actors, glial and astrocytes. Neurons, he discovered, came in a dizzying variety of shapes, sizes and geometries.

Later, the *electron microscope* established beyond doubt that nerve cells were linked at discrete specialized junctions, chemical and electrical synapses, and the *microelectrode* recorded the electrical activity of individual nerve cells. In 1963, the Nobel Prize was awarded to John Eccles for discovering the discrete (quantal) nature of synaptic transmission, and to Alan Hodgkin and Andrew Huxley for describing the sodium and potassium membrane currents that power the electrical impulse, the famed action potential or spike, as it travels along the axon. The mathematical formalism they pioneered has proved enduring; the reign of the Hodgkin-Huxley equations describing the biophysics of individual nerve cells would last until they were replaced by molecular dynamics model in the 2020s.

The next major advance came from electrical recordings from anesthetized and, subsequently, from awake and behaving animals with microelectrodes coupled to miniaturized differential amplifiers (and loudspeakers), which made the hitherto silent brain come alive with the staccato sounds of spiking nerve cells. In their classical 1959 and 1962 studies, David Hubel and Torsten Wiesel discovered the selectivity of visual cortical cells to the orientation of lines that the animal looked at. This work in turn launched the bold exploration of the higher order visual cortex that culminated in the late 1960s with the discovery of individual neurons that responded preferentially to faces.

Clinical studies, always a fecund source of knowledge about human nature, had given birth to neurology and to neurosurgery, both of which contributed to neuroscience. The neurologist Paul Broca had first inferred in 1861 from a singular patient that a specific region of the left inferior frontal gyrus is critical to speech. By the 1930s and 1940s, the neurosurgeon Wilder Penfield had stimulated the exposed cortex of epileptic patients with electrodes, thereby triggering simple visual percepts, movements, or vividly recalled memories, again and again. This was a compelling demonstration of the intimate link between the physical brain and the subjective mind.

In mathematical logic, Warren McCulloch and Walter Pitts demonstrated back in 1943 that interconnected networks of very simple neuron-like units could compute any logical expression. In conjunction with the Church-Turing thesis formalizing what is algorithmically computable, theoreticians and engineers established a foothold into the all-important challenge of conceptualizing how the brain could think, reason, and remember. Whereas René Descartes, three hundred years earlier, needed to postulate a vague cognitive substance (*res cogitans*) that did the thinking for people (famously, not for animals), computer scientists such as Frank Rosenblatt, inspired by McCulloch and Pitts, began taking the first tentative steps toward building computer simulations of brain-like circuits. If "Perceptrons," the single-layered neural networks of the 1960s, seem comically simplistic in hindsight, it must be remembered that such simple networks ultimately inspired a revolution. This period of boundless optimism and excitement was cross-fertilized by the launch of Artificial Intelligence in 1955 at Dartmouth College.

Neurophysiologists, computer scientists, and psychologists alike naively imagined that an understanding of the brain was near to hand. Of course, we now know that robust artificial intelligence took a century, not a few decades, to come about, and that neither psychology nor neuroscience was close to having reached the maturity that physics has. But the roots were all there. Nobody really knew remotely how the human brain worked, or how to emulate it, yet the revolution was well underway.

Neuroscience Becomes Big Science: 2014

Fifty years on, studying the brain was no longer a niche field but a full-on movement. The US-based Society for Neuroscience alone had more than forty thousand members, annual funding was well in excess of several billion dollars, and writers, journalists, and an inchoate neuro-industry all thrived on the public interest in the brain.

One major advance was molecular. Scientists had discerned the structure and function of ionic channels and receptors, the miniaturized stochastic switches and modulators embedded in the bilipid membrane that endows neurons with their ability to process information, to shape and guide action potentials along axons, and to release neurotransmitters. Also well understood was the action of sensory receptors that transduce the signals impinging onto the body—photons of light, sound perturbations in the air, or molecules of some odorant—into electrical activity. Indeed, neuroscientists had tracked down how single nucleotide changes in the DNA that encodes one or another photo-pigment protein in the retina impacted the way a subject perceives color.

The molecular revolutions of the day are perhaps best exemplified by the Nobel Prize–winning work of Eric Kandel, which elucidated how the sea slug *Aplysia* learns the gill-withdrawal reflex, the first form of long-term memory to be well understood. It demonstrated the importance of protein synthesis and changes in synaptic connectivity in long-term memory. Kandel's work furthered the growing realization that much of memory is encoded in the specific pattern and strength of connectivity among large ensembles of active neurons (as hypothesized already in 1895 by none other than Sigmund Freud), though the many ways in which memories could be stored *within* an individual neuron were not yet recognized. As Kandel and his contemporaries began to realize, the rules that determine how the influence that one synapse brings to bear on the neuron it is connected to, its *weight*, is up- or downward adjusted depends on the relative timing of the arrival of the pre- and postsynaptic electrical activity. (Cleverly, this gives individual synapses a rudimentary capacity for learning causal relationships, in which event **A** is followed by event **B** but never the other way around.) In 2013, the group of another Nobel laureate, Susumu Tonegawa, became the first to induce a false memory into mice by directly manipulating the underlying neural engram in their hippocampus. A great many molecular details—of the underlying neurotransmitters, second messenger systems, protein kinases, ionic channels, and transcription factors—were all steadily being filled in, even though the overall logic of the brain remained a mystery.

Two techniques proved transformative. First, in the 1980s, the physics of nuclear magnetic resonance was exploited to routinely, reliably, and safely image the static, anatomical structure of the human body by bombarding subjects with radio waves while they were lying inside powerful magnets. Applied to the brain, magnetic resonance imaging (MRI) revolutionized neurology. In the 1990s, MRI was refined to image the *functional* architecture of the active brain with spatiotemporal resolution at the scale of millimeters and seconds. Although the popular images of that time seem laughably crude by contemporary standards, they gave birth to the field of cognitive neuroscience as scientists began to investigate the neural basis of seeing, hearing, feeling, thinking, and remembering. Wars broke out about the "localization" hypothesis when many neuroscientists rejuvenated the old phrenologist program of linking specific mental faculties to specific parts of the brain, identifying more than one hundred brain regions on the basis of functional specializations. By 2014, theories of cognitive neuroscience began to grow in sophistication, as investigators realized that these specific regions formed parts of larger, more complex networks, which at that time eluded understanding. Only a few brain scientists were concerned with the coupling between fMRI signals, reflecting the power consumption of the brain at a sedate pace of seconds, and the switching in the underlying neural lattice at the millisecond scale. Indeed, the elementary spatial unit of brain imaging, *voxels*, at that time about 2 x 2 x 2 mm^3, encompasses about one million highly diverse neurons, glial cells, and astrocytes and ten billion synapses, firing two to twenty times within one MRI scan cycle, way too coarse to infer neuronal mechanism, akin to trying understanding language by listening to a smeared-out recording of the chattering among all the spectators at a sports arena. And few people had any conception of how important glial cells would turn out to be. Techniques like EEG and MEG were better temporally; they recorded electrical and magnetic fields with millisecond precision, but with even less spatial precision. The blurriness of these instruments was mirrored by the primitive and edentate tools used to safely perturb the human brain—electrical stimulation in patients, and extracranial electromagnetic fields and drugs in volunteers.

The other major advance fifty years ago was the birth of opto- and pharmaco-genetics, methods that delicately, transiently, reversibly, and invasively control defined events in defined cell types at defined times, initially in a few model organisms—the worm, the fly, and the mouse. Equipped with these tools for perturbing the brain, scientists systematically moved from correlation to causation, from observing that this circuit is activated whenever the subject is contemplating a decision to inferring that this circuit is necessary for decision making or that those neurons mark a particular memory. By the early 2020s, the complete logic of thalamo-cortical circuits could be manipulated, in hindsight a tipping point in our ability to bridge the gap between cortex and theories of its universal and particular functions.

An enormous amount of work characterized how sensory systems process their information and represent it in the cortical tissue. Silicon microelectrodes and live brain imaging using fluorescent dyes and genetically encoded proxy markers of electrical activity allowed intrepid neuroscientists to track the electrical activity of hundreds of neurons in the behaving animal simultaneously, a significant increase over the previous decades in which the brain was sampled by a single wire. Theoreticians could thereby infer from the firing of neurons the probabilistic manner in which the nervous system represents the visual, auditory, and olfactory environment, as well the animal's physical location, the animal's uncertainty in the face of a perceptual or a subjective decision, and even the presence of familiar individuals such as celebrities.

Yet despite these advances combined with the exponential increase in relevant data and the efforts of the brightest minds on the planet, comprehension of the brain's circuits in health and disease increased sublinearly. Even the smallest of all multicellular "model organisms," the roundworm *C. elegans*, whose nervous system contains a mere 302 neurons, was scarcely understood as a whole. Hundreds of worm specialists focused on isolated reductionist accounts of one function or another. Yet no one attempted to integrate all this knowledge into a single, coherent, comprehensive, holistic, and explanatory framework. Nor had any brain disease yet been cured. Many in the rapidly growing elderly population faced symptoms of dementia, yet little could be done to slow down the ravages of the disease; it must have been heartbreaking to witness. When the once dominant *Diagnostic and Statistical Manual of Mental Disorders*—at the time the psychiatrist's bible for treating patients with mental afflictions—appeared in its fifth edition in 2013, it did not list a *single* biomarker nor a single fMRI diagnostic criterion. If you were depressed, heard voices, or felt persecuted in the early twenty-first century, your only options were to talk to a therapist, fill out questionnaires, and take little-understood drugs that swamped your brain and had untold side effects.

In fairness, such slow progress was inevitable. Historically, science had been most successful when studying isolated systems with reduced degrees of freedom that tamed their complexity: a marble rolling down an inclined plane, a planet that plows its orbit around its center star, a lone electron in a magnetic field, a double strand of DNA. Even though it was obvious that living systems were characterized by large numbers of highly heterogeneous components, be they proteins, genes, or nerve cells, it was far from obvious how to deal with that complexity. A fundamental problem in the brain sciences has always been the numerous ways in which components interact causally across a large spectrum of space-time, from nanometers to meters and from microseconds to years. A complete understanding demands that a large fraction of these interactions be experimentally or computationally probed. This is fiendishly difficult. Bioinformaticians had few clues about how to integrate computations that spanned so many scales of time and space, and they lacked the relevant hardware, as cloud computers were primitive.

It was already becoming clear just how hard the problem was; even today no single human understands how the brain works at anything but an abstract and highly simplified level. Nature provides few shortcuts; a complete understanding of the brain comes not from any one experiment but from the integration of thousands of experiments that bridge many levels. Engineered systems such as spacecraft or computers that contain billions (then) or trillions (now) of discrete components are quite different. They are purposefully built to *limit* the interactions among the parts to a small number. Thus design rules for the layout of integrated electronic circuits impose a minimum distance between wires and other components to eliminate coupling, and the power supply is kept separate from computing, with computing separate from memory. Yet nervous systems interdigitate practically everything, from power supply to computation to memory. Nature couldn't have made herself more difficult to understand if she had tried. Early twenty-first-century scientists had begun to recognize this complexity but were unprepared and unable to deal with its consequences.

The next major revolution was not technological, but organizational. A private American initiative, the Allen Institute for Brain Science, taking cues from the biotechnology industry, was the first to approach neuroscience as "Big Science," moving from a model oriented around autonomous "star" investigators toward a team-based approach in which several hundred scientists from molecular biology, anatomy, physiology, genomics, optics, physics, and informatics worked together on industrial-scale projects, the first several of which had been launched by 2014 (see the chapter by Koch and colleagues, this volume). One generated the complete ontology of cortical cell types—the shape of their dendritic tree, the near- and far-flung target zones of their axons, the genes they express, their electrical behavior, and the rules governing their connectivities in the mouse and the human brain. The other was the construction of brain observatories—cerebroscopes—to record, make publicly accessible, analyze, and model the cellular events in the cortico-thalamic system underlying visual information processing in behaving mice. Other, even larger enterprises were spawned in the 2020s, as China and India became scientific world powers.

Also notable from that time was the publicly funded European Human Brain Project, which built a series of ever-larger supercomputer facilities to simulate, at the cellular, and, ultimately, at the subcellular level, the biophysics of neurons and their supporting cellular actors, in brains of increasing size, from the mouse to the human brain. Early on, their combination of morphological, anatomical, and physiological knowledge yielded an electrical model of a cortical column in rodents, a proof-of-principle that the electrodynamics of a chunk of brain matter could be understood by combing detailed biological knowledge with sufficient computational resources. The vision of a gigantic computer model of the human brain with the promise to comprehend its functioning, eliminate brain diseases, and ultimately upload ourselves, excited the public imagination with its near-religious imagery. As those initial simulations proved to be computationally underpowered and inaccurate, this promise backfired, leading to the withdrawal of public support for some time in the 2020s. Much was learned, but the public was disappointed.

Paraphrasing the twentieth-century British war leader Winston Churchill, neuroscience was at the end of the beginning of the quest to understand the brain and the mind. Neuroscientists had not yet figured out how to bridge the many levels of neurophysiology, from molecules to cells to circuits to behavior, but they had discerned enough to make the mission clear, and many critical tools were in place.

The Modern Era: 2064

Today, by identifying hierarchies of modules and submodules in the cortical sheet, we've largely tamed the sheer diversity and the vast extent of the neocortex. The basic organization of the cortical six-layered sheet is now known to schoolchildren, and if the overall interconnectivity is far too hard for any individual to understand, the nervous system of laboratory organisms like flies can now be emulated—successfully—with computers; human brains, too, have been simulated with some fidelity, although in time frames—about one-hundredth of real time—that make them less useful than was originally anticipated.

The retina was the first piece of neural tissue to be understood, in the sense that its output—action potentials along the optic nerve—can be quite accurately predicted from its input—patterns of light. One reason the retina led the way is its (relative) simplicity; unlike other nervous matter, the retina has primarily feed-forward connections—without any significant connections from the brain proper back to the retina. Most of its cellular elements had been recognized in the late twentieth century. By 2020, a Big-Science consortium of anatomists, physiologists, biophysics modelers, and machine learning specialists had arrived at a nearly complete description of retinal input-output, and the firing rates of the two dozen ganglion cell types, whose axons make up the optic nerve, could be reliably predicted, in response to arbitrary visual stimuli. That understanding (in combination with advanced optogenetics and implantable ocular electronics) led to effective treatments for macular degeneration, diabetic retinopathy, and retinitis pigmentosa.

Similar techniques helped crack the codes used in the visual thalamus and early visual cortical areas, as the onion layers of the brain began to be peeled back, one by one. A complete cellular-based working model of how the mouse moves through a maze in response to what it sees, together with the ontology of the approximately one thousand different cell types that make up the brain, was achieved in the mid 2020s. The senses of touch, hearing, and smell were decrypted a few years later.

This success fed the hope that understanding the entire mouse brain could not be far behind. Mechanistic explanations for what happens when the brain goes to sleep, dreams, wakes up, decides to run, remembers a location for another day, and develops across its lifespan, from birth to senescence, seemed close at hand. But these hopes were dashed. Yes, plenty of individual stories were told, but they could not be assembled into a coherent whole.

Funding for brain research slowed down because of the inability to translate these insights to people and their pathologies. Not that anybody seriously argued that the human brain was fundamentally different from that of the mouse. Of course, the two differ dramatically in size and accessibility. The human brain is more than a thousand times bigger than the mouse brain—1.4 kg versus 0.4 g in mass; a papaya versus a sugar cube in volume; eighty-six billion nerve cells versus seventy-one million for the entire brain and sixteen billion versus fourteen million nerve cells for the neocortex. Even more importantly was the ethical constraint: the living human brain could only be probed at the required cellular level under rare conditions, primarily during neurosurgery. fMRI, EEG, MEG, and other noninvasive techniques that peered at the brain from the outside were blind to genes, proteins, and cell types. While a rice- or corn-sized chunk of human gray matter is by and large similar to that of the mouse, there are many, many minute differences. Given the divergent ways in which *Mus musculus* and *Homo sapiens* evolved over the last seventy-five million years since their last common ancestor, their genes and gene regulatory mechanisms, proteins, synapses, neurons, and circuits differ in a multitude of small ways. Yet these trivial but elusive differences made generalizations from the mouse to humans difficult. Indeed, pharmaceutical companies had realized this earlier on and had discontinued much of their mouse research already in the early 2010s. After the animal rights movement managed to shut down almost all invasive research on nonhuman primates worldwide by the end of the 2020s, neuroscience entered what is now known as the lost decade. This was marked by low funding and pessimism that neuroscience could ever truly ameliorate the staggering toll that brain diseases took on the aging population, estimated to be 10 percent of world GDP.

The darkest hour is often just before the dawn. Help came from a very distant relative of humans, *C. elegans*, and from the triumphant marriage of artificial and biological molecular machines.

To be sure, it took over thirty-five years from when the connectome of two worms were mapped (in 1986) for an accurate, predictive, comprehensive, and fully testable model of its nervous system to be formulated. The key insight—the role of neuromodulators in switching pathways and circuits dynamically—was already faintly recognized fifty years ago by such pioneers as Cornelia Bargmann and Eve Marder, in the worm and other non-vertebrate species, but because worms lack action potentials, the importance of Bargmann and Marder's work for vertebrate creatures was initially overlooked. We now know that principles of dynamic routing are critical in all creatures.

The conquest of the living human brain was finally achieved with nanobotic neural implants, colloquially known as brainbots. These are molecular machines for imaging and manipulating the brain that can be safely injected by the millions into the bloodstream. The first generation of brainbots were designed to sample and measure their local environment, such as the electrical potential, or the concentration of a particular neurotransmitter or small molecule, and could be queried from the outside. More advanced probes read the transcriptional signature of individual neurons, monitor their electrical activity, arrest or trigger spikes, and, most recently, control synaptic release at individual synapses. They intervene at any point in the body by delivering missing or eliminating miss-formed neurotransmitters or proteins, or trigger electrical activity. Some operate transiently while others act as modified viruses that find a permanent home inside nerve and glial cells to arrest and ultimately repair the damages degenerative diseases such as Alzheimer's or Parkinson's cause; by the mid 2050s, almost all medicine, and all neuroscience, had moved to nanobotic platforms; even optogenetics, the workhorse of the early twenty-first century, eventually was displaced. Because of their high spatial specificity—guided by an externally imposed 3-D radio field—properly designed nanobots can target individual cells anywhere in the brain with enormous precision.

Many once-common mental diseases can now be delayed or, in a few cases, cured. To be sure, progress in reducing morbidity and mortality of brain-based pathologies—tumors, traumatic-brain injury, epilepsy, schizophrenia, Parkinson's, Alzheimer's and other forms of dementia—took much longer to realize than anyone conceived of in the early years of the new millennium. (An instructive parallel is the War on Cancer, announced by President Nixon in 1971, when America was flush with the success of the lunar landing; it was nearly five decades before there was a significant decline in the actual death rates for cancer, while death rates for respiratory, infectious, and cardiovascular diseases had plummeted much earlier.) Reducing the collective impact of brain-based pathologies turned out to be *more* difficult than curing the diverse set of pathologies known as cancer; both are highly heterogeneous diseases with an inexhaustible multiplicity of genetic, epigenetic, and environmental causes, but because of mosaicity, the complexity was even greater for the brain.

Brainbot treatment is expensive. And like most medical procedures, it has side effects, restricting it to appropriate patient populations. Yet although traditionalists and religious people object, nanobotic enhancement in healthy subjects is immensely attractive to those who believe in the infinite betterment of the human condition. Its proven ability to boost athletic agility and speed, learning and recall, has given rise to an underground market in brain enhancements. Those able to pay and willing to live with the short- and long-term morbidity and mortality risks are threatening to turn into trans-humans, a cognitive elite that easily outcompetes nonenhanced normals in the marketplace and in warfare.

In academic circles, the ongoing debate is about the growing raft of whole-brain simulations and what they mean both ethically and scientifically. For one thing, the question—first raised over fifty years ago—about the relevant level for brain simulation lingers. The intellectual tension arises between bottom-up simulators, who hold a form of extreme biological chauvinism—the need to consider every ionic channel, synapse, and action potential to fully do justice to the baroque complexity of the brain's circuits—and top-down simulators, who are motivated by the austerity of a purely algorithmic approach of replicating the mind in software (the mind is not wet, after all) and start with behavior or with computation.

Both sides have made major advances, but neither has been fully successful. Biophysicists accurately simulate the biochemical and neural activities of worms and flies with near full verisimilitude. Yet for mammals, deviations appear. And these differences between actual and simulated behaviors become more pronounced when moving from rodent brains, via those of monkeys and apes, to the human brain. Thus the spoken language such simulations produce is garbled, and most simulations remain at the kindergarten level on many tasks. What are we missing today? Do we have to simulate every ionic channel and every neurotransmitter molecule? Must we treat the brain as a quantum mechanical system? The brain is, after all, a physical object like any other one, subject to the iron law of quantum mechanics. Yet the vast majority of brain scientists assume that the nervous system, a hot (by QM standards) and wet organ closely coupled to its environment, can be approximated very well as a classical system.

Even considered as a classical system, biophysical brain simulations are dreadfully slow, working at one-hundredth the speed of real human brains; now that Moore's law has run out, and quantum computation proved to be of limited real-world use, it's not clear where the next advance will come from. Top-down modelers, meanwhile, capture some of the essence of human cognition, but with comparatively little fidelity to biological reality. Until the two approaches can be bridged, the thought-reading prosthetics that seemed so near a decade ago will continue to remain elusive. (In part, once again, the problem stems from complexity. Mathematicians and engineers imagined that there would be one true brain algorithm to rule them all, but because of the arbitrary accidents of nature's evolutionary opportunism, that simply hasn't proven to be the case; indeed, there seem to be almost as many algorithms as there are brain circuits, which has left little opportunity for shortcuts along the way.)

Meanwhile, on the cognitive side, processes such as language, planning, social cognition, and higher-level reasoning still resist explanation, especially in the intricate forms they take in people. Nanobotics may bridge this gap in our knowledge eventually, but for now, knowledge of uniquely human faculties still lags. We still don't know how the brain encodes sentences, and only a tiny bit is known about word meanings; complex concepts, like "the sort of person who reads fictitious narratives," remain entirely out of our grasp. If the neural basis of association has been entirely unraveled, the neural basis of higher-level cognition has not.

Ethically, as full-scale human brain emulations have neared, the political battles have been heated. Some see modeled rodents as ethically equal to real rodents and argue that complete human-brain emulations merit rights equal to human beings. Some scholars see emotional distress in the rudimentary human brain simulants. Yet most (chose to) believe that a simulation is an imitation rather than the real thing, just like a computer simulating the aerodynamics of flight will never actually lift off. Politicians avoid the issue, but time is clearly running out. Will it be legal to employ a whole-brain emulation for intellectual work, much as one might employ a human? Would it be ethical? Does all income accrue to the owner of the simulation, or might those whose brains contributed to the simulation also deserve royalty fees, in addition to the hourly fees they were paid for their original participation in extended brain scans?

The final challenge, indubitably, will be how *subjective* feelings, how consciousness itself, emerges from the physical brain. Even today, there remains an explanatory gap between neural activity and subjective feelings, between the brain and the conscious mind. One belongs to the realm of physics, to space and time, energy and mass. The other one belongs to a still poorly understood magisterium of experience. Spearheaded by the molecular biologist turned neuroscientist Francis Crick, cognitive neuroscientists have been tracking down the neuronal correlates of consciousness, but the vast complexity involved has kept us from a full solution. If we by now have a clear understanding of the dynamics by which information passes into awareness, we still don't fully know why experiences feel the way they do. The expectation is that the "hard problem" of consciousness will eventually be dissolved, and even disappear, much in the same way that the problem of "what is life" has disappeared from view, replaced by a host of more tractable problems about the details of reproduction and metabolism. As the behavior of computer artifacts begins to approach, and often to exceed, human capability, more and more people believe that consciousness arises from a privileged form of information associated with highly organized matter, such as brains or artificial intelligence agents, as argued already half a century earlier by Giulio Tononi. But if Descartes's famous conclusions four centuries ago might be paraphrased "I am conscious, therefore I am," the issues of consciousness still haven't been fully resolved. It is to be hoped that the next hundred years will finally bring resolution to the ancient mind-body riddle.

Acknowledgment: We wish to thank Ramez Naam, author of *Nexus*, for very thoughtful comments.