THE HUMAN BRAIN PROJECT

Building a vast digital simulation of the brain could transform neuroscience and medicine and reveal new ways of making more powerful computers

By Henry Markram

IT'S TIME TO CHANGE THE WAY WE STUDY THE BRAIN.

Reductionist biology—examining individual brain parts, neural circuits and molecules—has brought us a long way, but it alone cannot explain the workings of the human brain, an information processor within our skull that is perhaps unparalleled anywhere in the universe. We must construct as well as reduce and build as well as dissect. To do that, we need a new paradigm that combines both analysis and synthesis. The father of reductionism, French philosopher René Descartes, wrote about the need to investigate the parts and then reassemble them to re-create the whole.

Putting things together to devise a complete simulation of the human brain is the goal of an undertaking that intends to construct a fantastic new scientific instrument. Nothing quite like it exists yet, but we have begun building it. One way to think of this instrument is as the most powerful flight simulator ever built—only rather than simulating flight through open air, it will simulate a voyage through the brain. This "virtual brain"

IN BRIEF

Computer simulation will introduce ever greater verisimilitude into digital depictions of the workings of the human brain.

By the year 2020 digital brains may be able to represent the inner workings of a single brain cell or even the whole brain.

A sim brain can act as a stand-in for the genuine article, thus fostering a new understanding of autism or permitting virtual drug trials.
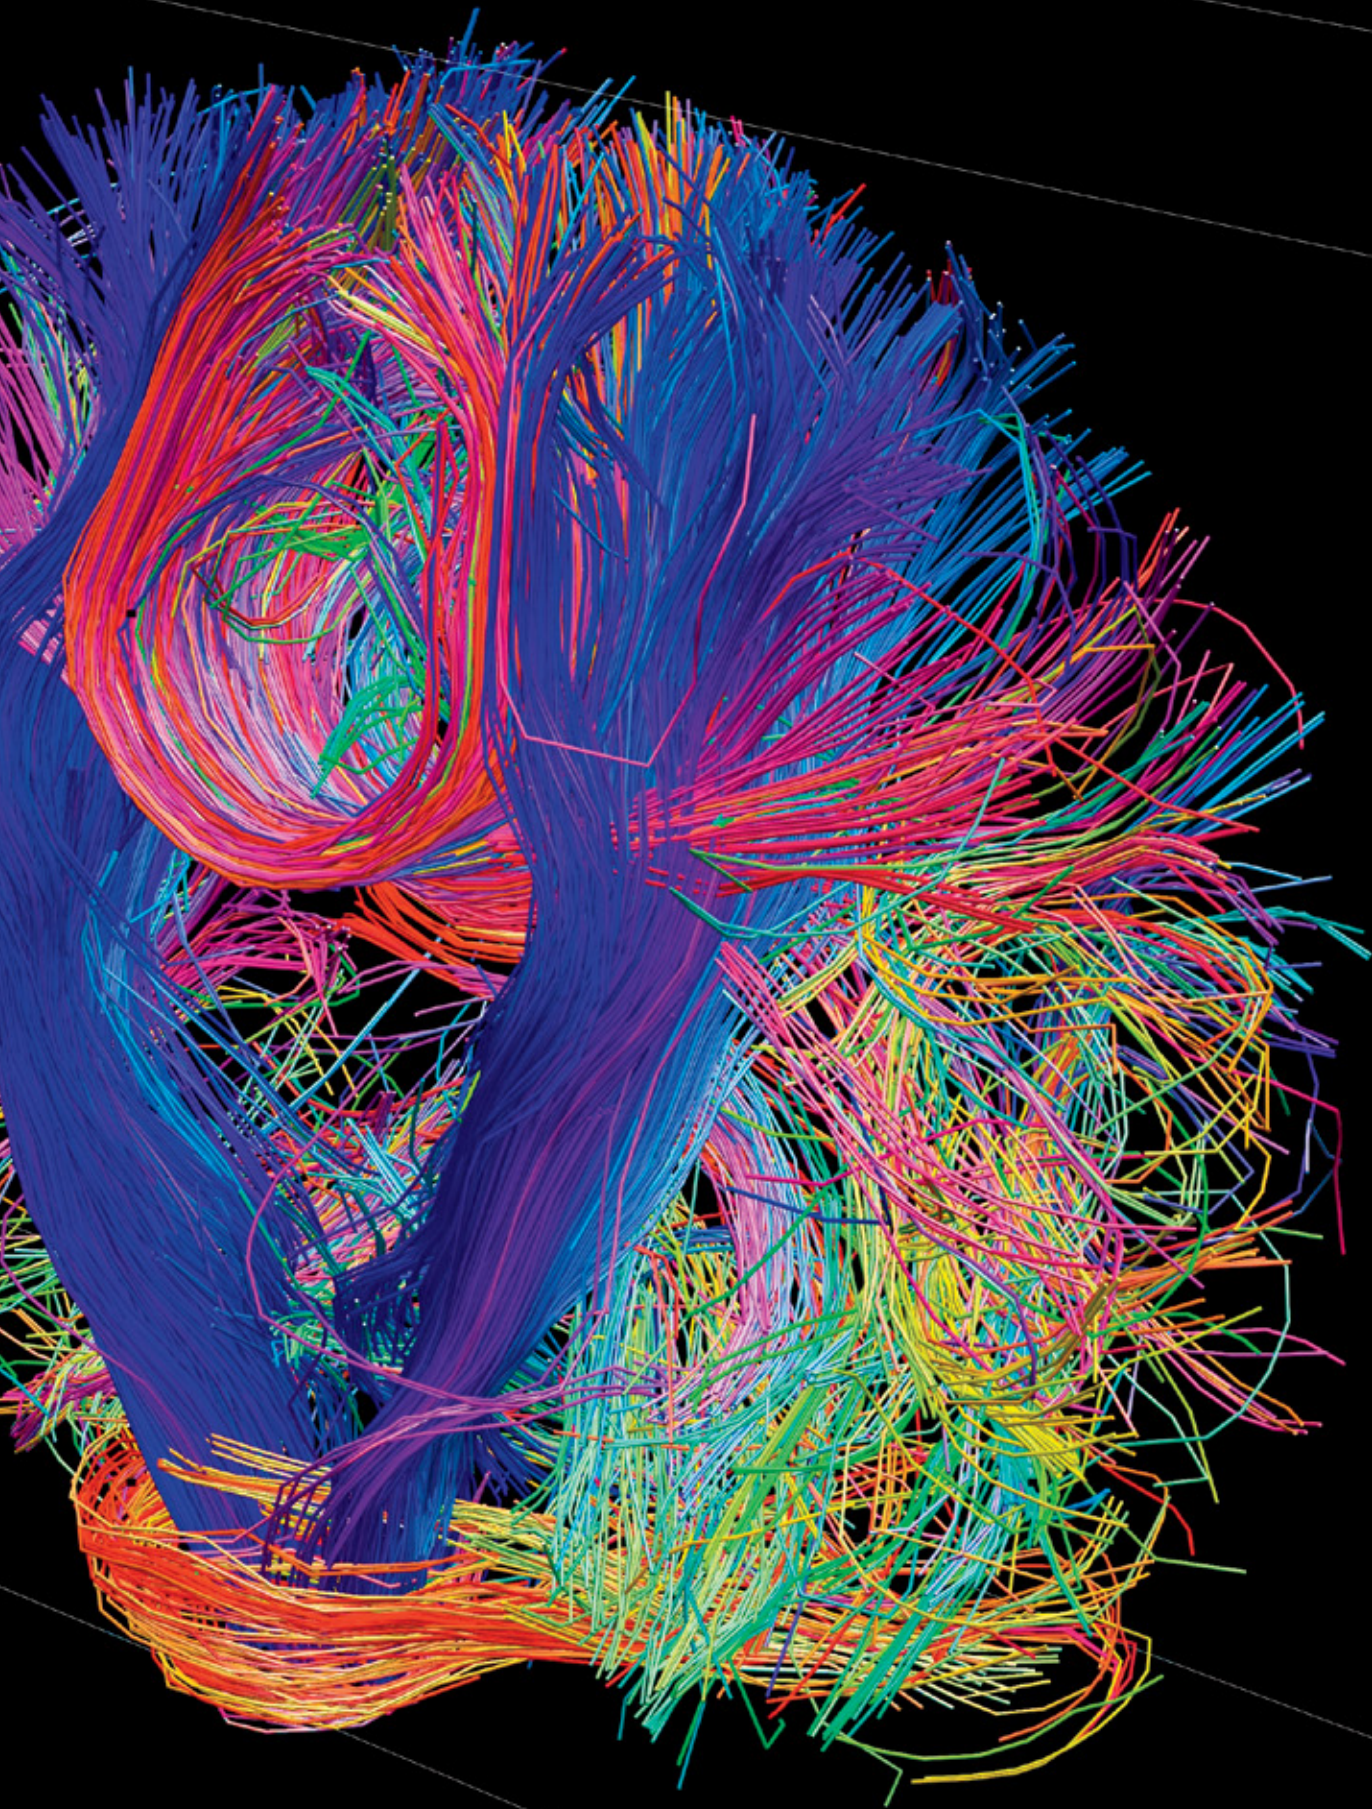
will run on supercomputers and incorporate all the data that neuroscience has generated to date.

A digital brain will be a resource for the entire scientific community: researchers will reserve time on it, as they do on the biggest telescopes, to conduct their experiments. They will use it to test theories of how the human brain works in health and in disease. They will recruit it to help them develop not only new diagnostic tests for autism or schizophrenia but also new therapies for depression and Alzheimer's disease. The wiring plan for tens of trillions of neural circuits will inspire the design of brainlike computers and intelligent robots. In short, it will transform neuroscience, medicine and information technology.

BRAIN IN A BOX

SCIENTISTS could be running the first simulations of the human brain by the end of this decade, when supercomputers will be powerful enough to support the massive number of calculations needed. The instrument will not require that all mysteries of the brain be unraveled first. Instead it will furnish a framework to accommodate what we do know, while enabling us to make predictions about what we do not. Those predictions will show us where to target our future experiments to prevent wasted effort. The knowledge we generate will be integrated with existing knowledge, and the "holes" in the framework will be filled in with increasingly realistic detail until, eventually, we will have a unified working model of the brain—one that reproduces it accurately from the whole brain down to the level of molecules.

Building this instrument is the goal of the Human Brain Project (HBP), an initiative involving about 130 universities around the world. The HBP is one of six projects competing for a glittering prize, up to €1 billion to be provided over 10 years by the European Union to each of two winners, who will be announced in February 2013.

We need the simulator for at least two reasons. In Europe alone, brain diseases affect 180 million people, or roughly one in three—a number that is set to grow as the population ages. At the same time, pharmaceutical companies are not investing in new treatments for the ailing nervous system. A holistic view of the brain would enable us to reclassify such diseases in biological terms rather than looking at them simply as sets of symptoms. The breadth of this perspective would allow us to move forward to develop a generation of treatments that selectively target the underlying abnormalities.

The second reason is that computing is fast approaching barriers to further development. Computers cannot do many tasks that animal brains do effortlessly, despite the inexorable increase in processing power. For instance, although computer scientists have made huge progress in visual recognition, the machines still struggle to make use of context in a scene or to use arbitrary scraps of information to predict future events in the way the brain can.

Moreover, because more powerful computers require more energy, supplying their needs will one day no longer be feasible. The performance of today's supercomputers is measured in petaflops—quadrillions of logic operations per second. The next generation, due around 2020, will be 1,000 times faster and will be measured in exaflops—quintillions of operations per second. By itself, the first exa-scale machine will probably consume around 20 megawatts, roughly the energy requirement of a small town in winter. To create increasingly powerful computers that perform some of the simple but useful things that humans are capable of in an energy-efficient way, we need a radically new strategy.

We could do worse than take inspiration from the human brain, which performs a range of intelligent functions on a mere 20 or so watts—a million times fewer than an exa-scale machine and equivalent to a weak lightbulb. For that, we need to understand the multilevel organization of the brain, from genes to behavior. The knowledge is out there, but we need to bring it together—and our instrument will provide the platform on which to do that.

Critics say that the goal of modeling the human brain is unachievable. One of their principal objections is that it is impossible to reproduce the connectivity among the brain's 100 trillion synapses because we cannot measure it. They are correct that we cannot measure the web of connections, which is why we are not going to—at least, not all of it. We intend to reproduce the myriad connections among brain cells by different means.

The key to our approach is to craft the basic blueprint according to which the brain is built: the set of rules that has guided its construction over evolution and does so anew in each developing fetus. In theory, those rules are all the information we need to start building a brain. The skeptics are right: the complexity they generate is daunting—hence our need for supercomputers to capture it. But unraveling the rules themselves is a far more tractable problem. If we pull it off, there is no logical reason why we cannot apply the blueprint in the same way that biology does and build an "in silico" brain.

The kind of rules we are talking about are ones that govern the genes that lead to the types of cells there are in the brain and the underlying plan for the way those cells are distributed and how they are connected. We know that such rules exist because we discovered some of them while laying the groundwork for the HBP. We started doing that almost 20 years ago by measuring the characteristics of individual neurons. We collected vast amounts of data about the geometric properties of different neuronal types and digitally reconstructed hundreds of them in three dimensions. Using a painstaking method called patch clamping, which involves placing the tip of a microscopic glass pipette up against a cell membrane to measure the voltage across its ion channels, we also recorded the neurons' electrical properties.

In 2005 modeling a single neuron took a powerful computer and a three-year Ph.D. project. It was clear that more ambitious goals would soon become achievable, however, and that we could model larger elements of brain circuitry even if our knowledge of those elements was incomplete. At the Brain Mind Institute at the Swiss Federal Institute of Technology in Lausanne, we launched one of the HBP's predecessors, the Blue Brain Project. We would build what we call "unifying computer models"—models that integrate all existing data and hypotheses about a

Henry Markram directs the Blue Brain Project at the Swiss Federal Institute of Technology in Lausanne. He has done extensive work on how neurons interconnect, communicate and learn. He also discovered fundamental principles of brain plasticity and is co-discoverer of the intense world theory of autism and the theory of how the brain computes as a liquid that is constantly perturbed.

given brain circuit, while reconciling conflicts in that information and highlighting where knowledge is lacking.

SYNTHESIS BIOLOGY

AS A TEST CASE, we set out to build a unifying model of a brain structure called the cortical column. The column is the equivalent of a processor in your laptop. To use a crude metaphor, if you were to put a miniature apple corer through the cortex and pull out a cylinder of tissue about half a millimeter in diameter and 1.5 mm in height, that would be a column. Within that tissue core, you would find a very dense network consisting of a few tens of thousands of cells. The column is such an efficient design for an information-processing element that once evolution had hit on the formula, it kept applying this recipe again and again until no more space was left in the skull and the cortex had to fold in on itself to create more room—hence, your convoluted brain.

The column penetrates the six vertical layers of the neocortex, the cortex's outer layer, and the neural connections between it and the rest of the brain are organized differently in each layer. The organization of these connections resembles the way telephone calls are assigned a numerical address and routed through an exchange. A few hundred neuron types reside in a column, and using our IBM Blue Gene supercomputer, we integrated all the available information about how those types mix in each layer until we had a "recipe" for a column in a newborn rat. We also instructed the computer to allow the virtual neurons to connect in all the ways that real neurons do—but only in those ways. It took us three years to build the software facility that, in turn, allowed us to construct this first unifying model of a column. And with it we had our proof of concept of what we call synthesis biology—a simulation of the brain from the full diversity of biological knowledge—and how it can serve as both a feasible and an inventive new way of doing research.

At that point, we had a static model—the equivalent of a column in a comatose brain. We wanted to know whether it would start to behave like a real column, albeit one isolated from the rest of the brain in a slice of living brain tissue, so we gave it a jolt—some external stimulation. In 2008 we applied a simulated electrical pulse to our virtual column. As we watched, the neurons began to speak to one another. "Spikes," or action potentials—the language of the brain—spread through the column as it began to work as an integrated circuit. The spikes flowed between the layers and oscillated back and forth, just as they do in living brain slices. This was behavior we had not programmed into the model; it emerged spontaneously because of the way the circuit was built. And the circuit stayed active even after the stimulation had stopped and briefly developed its own internal dynamics, its own way of representing information.

Since then, we have been gradually integrating more of the information generated by laboratories around the world into this unifying model of the column. The software we have developed is also being refined continuously so that each week we rebuild the column, we do so with more data, more rules and more accuracy. The next step is to integrate data for an entire brain region and then for an entire brain—to begin with, a rodent brain.

Our effort will depend heavily on a discipline called neuroinformatics. Vast quantities of brain-related data from all over the world need to be brought together in a coherent way, then mined for patterns or rules that describe how the brain is organized. We need to capture the biological processes those rules describe in sets of mathematical equations, while developing the software that will enable us to solve the equations on supercomputers. We also need to create software that will construct a brain that conforms to the inherent biology. We call it the "brain builder."

The predictions of how the brain operates offered up by neuroinformatics—and refined by new data—will accelerate our understanding of brain function without measuring every aspect of it. We can make predictions based on the rules we are uncovering and then test those predictions against reality. One of our current goals is to use knowledge of genes that give rise to the proteins for certain types of neurons to predict the structure and behavior of those cells. The link between genes and actual neurons constitutes what we call an "informatics bridge," the kind of shortcut that synthesis biology offers us.
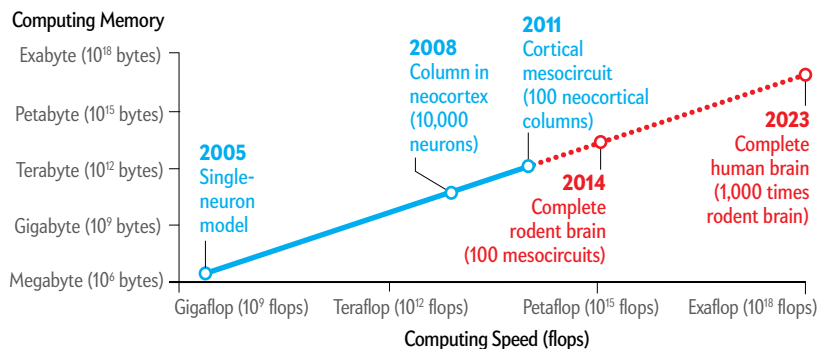
Another kind of informatics bridge that scientists have made use of for years has to do with genetic mutations and their link to disease: specifically, how mutation changes the proteins that cells manufacture, which in turn affect the geometry and electrical characteristics of neurons, the synapses they form and the electrical activity that emerges locally, in microcircuits, before spreading in a wide swath across whole brain regions.

In theory, for example, we could program a certain mutation into the model and then observe how that mutation affects it at each step along the biological chain. If the resulting symptom, or con-


POWER OF THE EXAFLOP

More Computer = More Brain

The ability to simulate the brain in enough detail to carry out vital scientific research will grow with computer power. A digital facsimile of a cylindrical piece of tissue in the rat cortex became a reality in 2008, when speed was clocked in teraflops. As computers climb to the peta and exa scales, the Human Brain Project envisages full-brain simulations of a mouse and of the same species that conceived *Hamlet* and Einstein's general theory of relativity.

Computing Memory

- Exabyte (10^{18} bytes)
- Petabyte (10^{15} bytes)
- Terabyte (10^{12} bytes)
- Gigabyte (10^{9} bytes)
- Megabyte (10^{6} bytes)

2005 Single-neuron model

2008 Column in neocortex (10,000 neurons)

2011 Cortical mesocircuit (100 neocortical columns)

2014 Complete rodent brain (100 mesocircuits)

2023 Complete human brain (1,000 times rodent brain)

Computing Speed (flops): Gigaflop (10^{9} flops) — Teraflop (10^{12} flops) — Petaflop (10^{15} flops) — Exaflop (10^{18} flops)

Deconstructing the Brain

The Human Brain Project intends to create a computer simulation of the 89 billion neurons inside our skull and the 100 trillion connections that wire those cells together. A meticulous virtual copy of the human brain would potentially enable basic research on brain cells and circuits or computer-based drug trials. The project, which is seeking €1 billion in funding from the European Union, would model each level of brain function, from chemical and electrical signaling up to the cognitive traits that underlie intelligent behaviors.
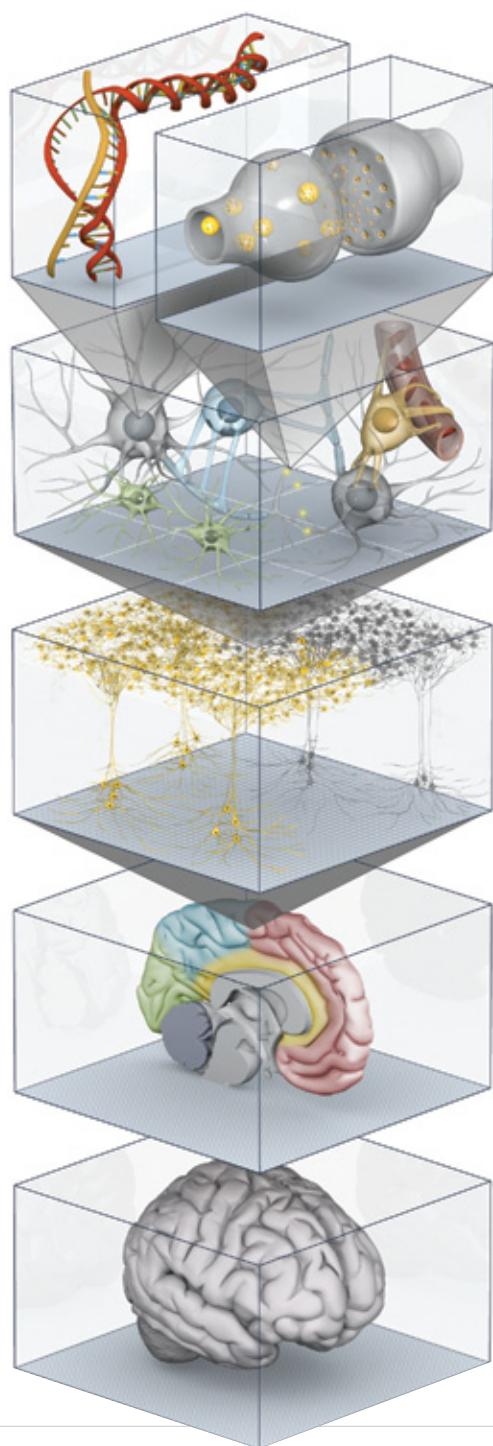
Molecular

A century of research, beginning with the first inspection of a brain cell under a microscope, would translate into a digital facsimile that combines component molecular parts to assemble a cell that demonstrates the essential properties of a neuron—the transmission of electrical and chemical signals.

Cellular

A brain-in-a-box simulation will have to capture every detail of neurons and nonneuronal glial cells, including the exact geometric shapes of the dendrites and axons that receive and send information.

Circuits

A model of the neural connections between different brain areas and among neighboring cells may furnish clues to the origins of complex brain diseases such as autism and schizophrenia.

Regions

Major neural substructures—the amygdala (emotions), the hippocampus (memory), the frontal lobes (executive control)—can be inspected alone or as they interact with one another.

Whole Organ

An in silico brain might substitute for the actual organ. By removing the computer code for a "gene," the virtual system can, for instance, mimic the effects of a mutation, as scientists do today by "knocking out" a gene in mice. The tool would avoid the lengthy breeding process and could simulate a multitude of experimental conditions.

stellation of symptoms, matches what we see in real life, that virtual chain of events becomes a candidate for a disease mechanism, and we can even begin to look for potential therapeutic targets along it.

This process is intensely iterative. We integrate all the data we can find and program the model to obey certain biological rules, then run a simulation and compare the "output," or resulting behavior of proteins, cells and circuits, with relevant experimental data. If they do not match, we go back and check the accuracy of the data and refine the biological rules. If they do match, we bring in more data, adding ever more detail while expanding our model to a larger portion of the brain. As the software improves, data integration becomes faster and automatic, and the model behaves more like the actual biology. Modeling the whole brain, when our knowledge of cells and synapses is still incomplete, no longer seems an impossible dream.

To feed this enterprise, we need data and lots of them. Ethical concerns restrict the experiments that neuroscientists can perform on the human brain, but fortunately the brains of all mammals are built according to common rules, with species-specific variations. Most of what we know about the genetics of the mammalian brain comes from mice, while monkeys have given us valuable insights into cognition. We can therefore begin by building a unifying model of a rodent brain and then using it as a starting template from which to develop our human brain model—gradually integrating detail after detail. Thus, the models of mouse, rat and human brains will develop in parallel.

The data that neuroscientists generate will help us identify the rules that govern brain organization and verify experimentally that our extrapolations—those predicted chains of causation—match the biological truth. At the level of cognition, we know that very young babies have some grasp of the numerical concepts 1, 2 and 3 but not of higher numbers. When we are finally able to model the brain of a newborn, that model must recapitulate both what the baby can do and what it cannot.

A great deal of the data we need already exist, but they are not easily accessible. One major challenge for the HBP will be to pool and organize them. Take the medical arena: those data are going to be immensely valuable to us not only because dysfunction tells us about normal

Illustration by Emily Cooper

function but also because any model we produce must behave like a healthy brain and later get sick in the same way that a real brain does. Patients' brain scans will therefore be a rich source of information.

Currently every time a patient has a scan, that scan resides in a digital archive. The world's hospitals stock millions of scans, and although they are already used for research purposes, that research happens in such a piecemeal way they remain a largely untapped resource. If we could bring together those scans on Internet-accessible "clouds," collecting them with patients' records and biochemical and genetic information, doctors could look across vast populations of patients for patterns that define disease. The power of this strategy will come from being able to mathematically pinpoint the differences and similarities among all diseases. A multiuniversity endeavor called the Alzheimer's Disease Neuroimaging Initiative is trying to do just that by collecting neuroimaging, cerebrospinal fluid and blood records from large numbers of dementia patients and healthy control subjects.

THE FUTURE OF COMPUTING

LAST BUT NOT LEAST, there is the computing issue. The latest generation of Blue Gene is a peta-scale beast consisting of close to 300,000 processors packed into the space of 72 fridges. Petaflops are sufficient to model a rat brain of 200 million neurons at a cellular level of detail but not a human brain of 89 billion neurons. For that achievement, we need an exa-scale supercomputer, and even then a molecular-level simulation of the human brain will be beyond our reach.

Teams worldwide are racing to build such computers. When they arrive, like previous generations of supercomputers, they are likely to be adapted to simulating physical processes, such as those used in nuclear physics. Biological simulations have different requirements, and in collaboration with large computer manufacturers and other industrial partners, our consortium of high-performance-computing experts will configure one such machine for the task of simulating a brain. They will also develop the software that will allow us to build unifying models from the lowest to the highest resolution so that it will be possible, within our simulator, to zoom in and out among molecules, cells and the entire brain.

Once our brain simulator has been built, researchers will be able to set up in silico experiments using the software specimen much as they would a biological specimen, with certain key differences. To give you an idea of what these might be, think about how scientists currently search for the roots of disease by using mice in which a gene has been "knocked out." They have to breed the mice, which takes time, is expensive and is not always possible—for example, if the knockout is lethal to the embryo—even if one lays aside ethical concerns surrounding animal experimentation.

With the in silico brain, they will be able to knock out a virtual gene and see the results in "human" brains that are different ages and that function in distinctive ways. They will be able to repeat the experiment under as many different conditions as they like, using the same model, thus ensuring a thoroughness that is not obtainable in animals. Not only could this accelerate the process by which pharmaceutical researchers identify potential drug targets, it will also change the way clinical trials are conducted. It will be much easier to select a target population, and drugs that do not work or that have unacceptable side effects will be filtered out more quickly, with the result that the entire R&D pipeline will be accelerated and made more efficient.

What we learn from such simulations will also feed back into the design of computers by revealing how evolution produced a brain that is resilient, that performs multiple tasks rapidly and simultaneously on a massive scale—while consuming the same amount of energy as a lightbulb—and that has a huge memory capacity.

Brainlike computer chips will be used to build so-called neuromorphic computers. The HBP will print brain circuits on silicon chips, building on technology developed in the European Union projects BrainScaleS and SpiNNaker.

The first whole-brain simulations we run on our instrument will lack a fundamental feature of the human brain: they will not develop as a child does. From birth onward, the cortex forms as a result of the proliferation, migration and pruning of neurons and of a process we call plasticity that is highly dependent on experience. Our models will instead begin at any arbitrary age, leapfrogging years of development, and continue from there to capture experiences. We will need to build the machinery to allow the model to change in response to input from the environment.

The litmus test of the virtual brain will come when we connect it up to a virtual software representation of a body and place it in a realistic virtual environment. Then the in silico brain will be capable of receiving information from the environment and acting on it. Only after this achievement will we be able to teach it skills and judge if it is truly intelligent. Because we know there is redundancy in the human brain—that is, one neural system can compensate for another—we can begin to find which aspects of brain function are essential to intelligent behavior.

The HBP raises important ethical issues. Even if a tool that simulates the human brain is a long way off, it is legitimate to ask whether it would be responsible to build a virtual brain that possessed more cortical columns than a human brain or that combined humanlike intelligence with a capacity for number crunching a million times greater than that of IBM's Deep Blue, its chess-playing computer.

We are not the only ones setting the bar high in attempting to reverse the fragmentation of brain research. In May 2010 the Seattle-based Allen Institute for Brain Science launched its Allen Human Brain Atlas, with the goal of mapping all the genes that are active in the human brain.

Funding is likely to be the main limiting factor for any group making an attempt of this kind. In our case, the goal will be achievable only if we obtain the support we need. Supercomputers are expensive, and the final cost of the HBP is likely to match or exceed that of the Human Genome Project. In February 2013 we will know if we have the green light. Meanwhile we press ahead with an enterprise we believe will give us unparalleled insight into our own identities as creatures capable of contemplating the chiaroscuro of a Caravaggio painting or the paradoxes of quantum physics. ▪

| MORE TO EXPLORE |

Links to a few Human Brain Project sites:
Human Brain Project: www.humanbrainproject.eu
BrainScaleS: http://brainscales.kip.uni-heidelberg.de
SpiNNaker:http://apt.cs.man.ac.uk/projects/SpiNNaker