







September 16, 2002

Controlling Robots with the Mind

People with nerve or limb injuries may one day be able to command wheelchairs, prosthetics and even paralyzed arms and legs by "thinking them through" the motions

By Miguel A. L. Nicolelis and John K. Chapin

Belle, our tiny owl monkey, was seated in her special chair inside a soundproof chamber at our Duke University laboratory. Her right hand grasped a joystick as she watched a horizontal series of lights on a display panel. She knew that if a light suddenly shone and she moved the joystick left or right to correspond to its position, a dispenser would send a drop of fruit juice into her mouth. She loved to play this game. And she was good at it.

Belle wore a cap glued to her head. Under it were four plastic connectors. The connectors fed arrays of microwires--each wire finer than the finest sewing thread--into different regions of Belle's motor cortex, the brain tissue that plans movements and sends instructions for enacting the plans to nerve cells in the spinal cord. Each of the 100 microwires lay beside a single motor neuron. When a neuron produced an electrical discharge--an "action potential"--the adjacent microwire would capture the current and send it up through a small wiring bundle that ran from Belle's cap to a box of electronics on a table next to the booth. The box, in turn, was linked to two computers, one next door and the other half a country away.

In a crowded room across the hall, members of our research team were getting anxious. After months of hard work, we were about to test the idea that we could reliably translate the raw electrical activity in a living being's brain--Belle's mere thoughts--into signals that could direct the actions of a robot. Unknown to Belle on this spring afternoon in 2000, we had assembled a multijointed robot arm in this room, away from her view, that she would control for the first time. As soon as Belle's brain sensed a lit spot on the panel, electronics in the box running two real-time mathematical models would rapidly analyze the tiny action potentials produced by her brain cells. Our lab computer would convert the electrical patterns into instructions that would direct the robot arm. Six hundred miles north, in Cambridge, Mass., a different computer would produce the same actions in another robot arm, built by Mandayam A. Srinivasan, head of the Laboratory for Human and Machine Haptics (the Touch Lab) at the Massachusetts Institute of Technology. At least, that was the plan.

If we had done everything correctly, the two robot arms would behave as Belle's arm did, at exactly the same time. We would have to translate her neuronal activity into robot commands in just 300 milliseconds--the natural delay between the time Belle's motor cortex planned how she should move her limb and the moment it sent the instructions to her muscles. If the brain of a living creature could accurately control two dissimilar robot arms--despite the signal noise and transmission delays inherent in our lab network and the error-prone Internet--perhaps it could someday control a mechanical device or actual limbs in ways that would be truly helpful to people.

Finally the moment came. We randomly switched on lights in front of Belle, and she immediately moved her joystick back and forth to correspond to them. Our robot arm moved similarly to Belle's real arm. So did Srinivasan's. Belle and the robots moved in synchrony, like dancers choreographed by the electrical impulses sparking in Belle's mind. Amid the loud celebration that erupted in Durham, N.C., and Cambridge, we could not help thinking that this was only the beginning of a promising journey.

In the two years since that day, our labs and several others have advanced neuroscience, computer science, microelectronics and robotics to create ways for rats, monkeys and eventually humans to control mechanical and electronic machines purely by "thinking through," or imagining, the motions. Our immediate goal is to help a person who has been paralyzed by a neurological disorder or spinal cord injury, but whose motor cortex is spared, to operate a wheelchair or a robotic limb. Someday the research could also help such a patient regain control over a natural arm or leg, with the aid of wireless communication between implants in the brain and the limb. And it could lead to devices that restore or augment other motor, sensory or cognitive functions.

The big question is, of course, whether we can make a practical, reliable system. Doctors have no means by which to repair spinal cord breaks or damaged brains. In the distant future, neuroscientists may be able to regenerate injured neurons or program stem cells (those capable of differentiating into various cell types) to take their place. But in the near future, brain-machine interfaces (BMIs), or neuroprostheses, are a more viable option for restoring motor function. Success this summer with macaque monkeys that completed different tasks than those we asked of Belle has gotten us even closer to this goal.

From Theory to Practice

Recent advances in brain-machine interfaces are grounded in part on discoveries made about 20 years ago. In the early 1980s Apostolos P. Georgopoulos of Johns Hopkins University recorded the electrical activity of single motor-cortex neurons in macaque monkeys. He found that the nerve cells typically reacted most strongly when a monkey moved its arm in a certain direction. Yet when the arm moved at an angle away from a cell's preferred direction, the neuron's activity didn't cease; it diminished in proportion to the cosine of that angle. The finding showed that motor neurons were broadly tuned to a range of motion and that the brain most likely relied on the collective activity of dispersed populations of single neurons to generate a motor command.

There were caveats, however. Georgopoulos had recorded the activity of single neurons one at a time and from only one motor area. This approach left unproved the underlying hypothesis that some kind of coding scheme emerges from the simultaneous activity of many neurons distributed across multiple cortical areas. Scientists knew that the frontal and parietal lobes--in the forward and rear parts of the brain, respectively--interacted to plan and generate motor commands. But technological bottlenecks prevented neurophysiologists from making widespread recordings at once. Furthermore, most scientists believed that by cataloguing the properties of neurons one at a time, they could build a comprehensive map of how the brain works--as if charting the properties of individual trees could unveil the ecological structure of an entire forest!

Fortunately, not everyone agreed. When the two of us met 14 years ago at Hahnemann University, we discussed the challenge of simultaneously recording many single neurons. By 1993 technological breakthroughs we had made allowed us to record 48 neurons spread across five structures that form a rat's sensorimotor system--the brain regions that perceive and use sensory information to direct movements.

Crucial to our success back then--and since--were new electrode arrays containing Teflon-coated stainless-steel microwires that could be implanted in an animal's brain. Neurophysiologists had used standard electrodes that resemble rigid needles to record single neurons. These classic electrodes worked well but only for a few hours, because cellular compounds collected around the electrodes' tips and eventually insulated them from the current. Furthermore, as the subject's brain moved slightly during normal activity, the stiff pins damaged neurons. The microwires we devised in our lab (later produced by NBLabs in Denison, Tex.) had blunter tips, about 50 microns in diameter, and were much more flexible. Cellular substances did not seal off the ends, and the flexibility greatly reduced neuron damage. These properties enabled us to produce recordings for months on end, and having tools for reliable recording allowed us to begin developing systems for translating brain signals into commands that could control a mechanical device.

With electrical engineer Harvey Wiggins, now president of Plexon in Dallas, and with Donald J. Woodward and Samuel A. Deadwyler of Wake Forest University School of Medicine, we devised a small "Harvey box" of custom electronics, like the one next to Belle's booth. It was the first hardware that could properly sample, filter and amplify neural signals from many electrodes. Special software allowed us to discriminate electrical activity from up to four single neurons per microwire by identifying unique features of each cell's electrical discharge.

A Rat's Brain Controls a Lever

In our next experiments at Hahnemann in the mid-1990s, we taught a rat in a cage to control a lever with its mind. First we trained it to press a bar with its forelimb. The bar was electronically connected to a lever outside the cage. When the rat pressed the bar, the outside lever tipped down to a chute and delivered a drop of water it could drink.

We fitted the rat's head with a small version of the brain-machine interface Belle would later use. Every time the rat commanded its forelimb to press the bar, we simultaneously recorded the action potentials produced by 46 neurons. We had programmed resistors in a so-called integrator, which weighted and processed data from the neurons to generate a single analog output that predicted very well the trajectory of the rat's forelimb. We linked this integrator to the robot lever's controller so that it could command the lever.

Once the rat had gotten used to pressing the bar for water, we disconnected the bar from the lever. The rat pressed the bar, but the lever remained still. Frustrated, it began to press the bar repeatedly, to no avail. But one time, the lever tipped and delivered the water. The rat didn't know it, but its 46 neurons had expressed the same firing pattern they had in earlier trials when the bar still worked. That pattern prompted the integrator to put the lever in motion.

After several hours the rat realized it no longer needed to press the bar. If it just looked at the bar and imagined its forelimb pressing it, its neurons could still express the firing pattern that our brain-machine interface would interpret as motor commands to move the lever. Over time, four of six rats succeeded in this task. They learned that they had to "think through" the motion of pressing the bar. This is not as mystical as it might sound; right now you can imagine reaching out to grasp an object near you--without doing so. In similar fashion, a person with an injured or severed limb might learn to control a robot arm joined to a shoulder.

A Monkey's Brain Controls a Robot Arm

We were thrilled with our rats' success. It inspired us to move forward, to try to reproduce in a robotic limb the three-dimensional arm movements made by monkeys--animals with brains far more similar to those of humans. As a first step, we had to



▶ Archive with
160+ issues &
3,600+ articles

▶ 12 new digital
issues

▶ Unlimited
access 24 x 7

**Subscribe
Today
& Save**

Free Trial Issue +
Free Gift

Low Introductory
Rate

**SCIENTIFIC
AMERICAN** | **DIGITAL**

SUBSCRIBE TODAY ▶

SUBSCRIBE NOW



**SCIENTIFIC
AMERICAN**

▶ Archive with
160+ issues &
3,600+ articles

▶ 12 new digital
issues

▶ Unlimited
access 24 x 7

**Get More
Science.
Get
Digital.**

**SCIENTIFIC
AMERICAN** | **DIGITAL**

SUBSCRIBE TODAY ▶

devise technology for predicting how the monkeys intended to move their natural arms.

At this time, one of us (Nicolletis) moved to Duke and established a neurophysiology laboratory there. Together we built an interface to simultaneously monitor close to 100 neurons, distributed across the frontal and parietal lobes. We proceeded to try it with several owl monkeys. We chose owl monkeys because their motor cortical areas are located on the surface of their smooth brain, a configuration that minimizes the surgical difficulty of implanting microwire arrays. The microwire arrays allowed us to record the action potentials in each creature's brain for several months.

In our first experiments, we required owl monkeys, including Belle, to move a joystick left or right after seeing a light appear on the left or right side of a video screen. We later sat them in a chair facing an opaque barrier. When we lifted the barrier they saw a piece of fruit on a tray. The monkeys had to reach out and grab the fruit, bring it to their mouth and place their hand back down. We measured the position of each monkey's wrist by attaching fiber-optic sensors to it, which defined the wrist's trajectory.

Further analysis revealed that a simple linear summation of the electrical activity of cortical motor neurons predicted very well the position of an animal's hand a few hundred milliseconds ahead of time. This discovery was made by Johan Wessberg of Duke, now at the Gothenburg University in Sweden. The main trick was for the computer to continuously combine neuronal activity produced as far back in time as one second to best predict movements in real time.

As our scientific work proceeded, we acquired a more advanced Harvey box from Plexon. Using it and some custom, real-time algorithms, our computer sampled and integrated the action potentials every 50 to 100 milliseconds. Software translated the output into instructions that could direct the actions of a robot arm in three-dimensional space. Only then did we try to use a BMI to control a robotic device. As we watched our multijointed robot arm accurately mimic Belle's arm movements on that inspiring afternoon in 2000, it was difficult not to ponder the implausibility of it all. Only 50 to 100 neurons randomly sampled from tens of millions were doing the needed work.

Later mathematical analyses revealed that the accuracy of the robot movements was roughly proportional to the number of neurons recorded, but this linear relation began to taper off as the number increased. By sampling 100 neurons we could create robot hand trajectories that were about 70 percent similar to those the monkeys produced. Further analysis estimated that to achieve 95 percent accuracy in the prediction of one-dimensional hand movements, as few as 500 to 700 neurons would suffice, depending on which brain regions we sampled. We are now calculating the number of neurons that would be needed for highly accurate three-dimensional movements. We suspect the total will again be in the hundreds, not thousands.

These results suggest that within each cortical area, the "message" defining a given hand movement is widely disseminated. This decentralization is extremely beneficial to the animal: in case of injury, the animal can fall back on a huge reservoir of redundancy. For us researchers, it means that a BMI neuroprosthesis for severely paralyzed patients may require sampling smaller populations of neurons than was once anticipated.

We continued working with Belle and our other monkeys after Belle's successful experiment. We found that as the animals perfected their tasks, the properties of their neurons changed--over several days or even within a daily two-hour recording session. The contribution of individual neurons varied over time. To cope with this "motor learning," we added a simple routine that enabled our model to reassess periodically the contribution of each neuron. Brain cells that ceased to influence the predictions significantly were dropped from the model, and those that became better predictors were added. In essence, we designed a way to extract from the brain a neural output for hand trajectory. This coding, plus our ability to measure neurons reliably over time, allowed our BMI to represent Belle's intended movements accurately for several months. We could have continued, but we had the data we needed.

It is important to note that the gradual changing of neuronal electrical activity helps to give the brain its plasticity. The number of action potentials a neuron generates before a given movement changes as the animal undergoes more experiences. Yet the dynamic revision of neuronal properties does not represent an impediment for practical BMIs. The beauty of a distributed neural output is that it does not rely on a small group of neurons. If a BMI can maintain viable recordings from hundreds to thousands of single neurons for months to years and utilize models that can learn, it can handle evolving neurons, neuronal death and even degradation in electrode-recording capabilities.

Exploiting Sensory Feedback

Belle proved that a bmi can work for a primate brain. But could we adapt the interface to more complex brains? In May 2001 we began studies with three macaque monkeys at Duke. Their brains contain deep furrows and convolutions that resemble those of the human brain.

We employed the same BMI used for Belle, with one fundamental addition: now the monkeys could exploit visual feedback to judge for themselves how well the BMI could mimic their hand movements. We let the macaques move a joystick in random directions, driving a cursor across a computer screen. Suddenly a round target would appear somewhere on the screen. To receive a sip of fruit juice, the monkey had to position the cursor quickly inside the target--within 0.5 second--by rapidly manipulating the joystick.

The first macaque to master this task was Aurora, an elegant female who clearly enjoyed showing off that she could hit the target more than 90 percent of the time. For a year, our postdoctoral fellows Roy Crist and José Carmena recorded the activity of up to 92 neurons in five frontal and parietal areas of Aurora's cortex.

Once Aurora commanded the game, we started playing a trick on her. In about 30 percent of the trials we disabled the connection between the joystick and the cursor. To move the cursor quickly within the target, Aurora had to rely solely on her brain activity, processed by our BMI. After being puzzled, Aurora gradually altered her strategy. Although she continued to make hand movements, after a few days she learned she could control the cursor 100 percent of the time with her brain alone. In a few trials each day during the ensuing weeks Aurora didn't even bother to move her hand; she moved the cursor by just thinking about the trajectory it should take.

That was not all. Because Aurora could see her performance on the screen, the BMI made better and better predictions even though it was recording the same neurons. Although much more analysis is required to understand this result, one explanation is that the visual feedback helped Aurora to maximize the BMI's reaction to both brain and machine learning. If this proves true, visual or other sensory feedback could allow people to improve the performance of their own BMIs.

We observed another encouraging result. At this writing, it has been a year since we implanted the microwires in Aurora's brain, and we continue to record 60 to 70 neurons daily. This extended success indicates that even in a primate with a convoluted brain, our microwire arrays can provide long-term, high-quality, multichannel signals. Although this sample is down from the original 92 neurons, Aurora's performance with the BMI remains at the highest levels she has achieved.

We will make Aurora's tasks more challenging. In May we began modifying the BMI to give her tactile feedback for new experiments that are now beginning. The BMI will control a nearby robot arm fitted with a gripper that simulates a grasping hand. Force sensors will indicate when the gripper encounters an object and how much force is required to hold it. Tactile feedback--is the object heavy or light, slick or sticky?--will be delivered to a patch on Aurora's skin embedded with small vibrators. Variations in the vibration frequencies should help Aurora figure out how much force the robot arm should apply to, say, pick up a piece of fruit, and to hold it as the robot brings it back to her. This experiment might give us the most concrete evidence yet that a person suffering from severe paralysis could regain basic arm movements through an implant in the brain that communicated over wires, or wirelessly, with signal generators embedded in a limb.

If visual and tactile sensations mimic the information that usually flows between Aurora's own arm and brain, long-term interaction with a BMI could possibly stimulate her brain to incorporate the robot into its representations of her body--schema known to exist in most brain regions. In other words, Aurora's brain might represent this artificial device as another part of her body. Neuronal tissue in her brain might even dedicate itself to operating the robot arm and interpreting its feedback.

To test whether this hypothesis has merit, we plan to conduct experiments like those done with Aurora, except that an animal's arm will be temporarily anesthetized, thereby removing any natural feedback information. We predict that after a transition period, the primate will be able to interact with the BMI just fine. If the animal's brain does meld the robot arm into its body representations, it is reasonable to expect that a paraplegic's brain would do the same, rededicating neurons that once served a natural limb to the operation of an artificial one.

Each advance shows how plastic the brain is. Yet there will always be limits. It is unlikely, for example, that a stroke victim could gain full control over a robot limb. Stroke damage is usually widespread and involves so much of the brain's white matter--the fibers that allow brain regions to communicate--that the destruction overwhelms the brain's plastic capabilities. This is why stroke victims who lose control of uninjured limbs rarely regain it.

Reality Check

Good news notwithstanding, we researchers must be very cautious about offering false hope to people with serious disabilities. We must still overcome many hurdles before BMIs can be considered safe, reliable and efficient therapeutic options. We have to demonstrate in clinical trials that a proposed BMI will offer much greater well-being while posing no risk of added neurological damage.

Surgical implantation of electrode arrays will always be of medical concern, for instance. Investigators need to evaluate whether highly dense microwire arrays can provide viable recordings without causing tissue damage or infection in humans. Progress toward dense arrays is already under way. Duke electronics technician Gary Lehw has designed ways to increase significantly the number of microwires mounted in an array that is light and easy to implant. We can now implant multiple arrays, each of which has up to 160 microwires and measures five by eight millimeters, smaller than a pinky fingernail. We recently implanted 704 microwires across eight cortical areas in a macaque and recorded 318 neurons simultaneously.

In addition, considerable miniaturization of electronics and batteries must occur. We have begun collaborating with José Carlos Principe of the University of Florida to craft implantable microelectronics that will embed in hardware the neuronal

pattern recognition we now do with software, thereby eventually freeing the BMI from a computer. These microchips will thus have to send wireless control data to robotic actuators. Working with Patrick D. Wolf's lab at Duke, we have built the first wireless "neurochip" and beta-tested it with Aurora. Seeing streams of neural activity flash on a laptop many meters away from Aurora--broadcast via the first wireless connection between a primate's brain and a computer--was a delight.



More and more scientists are embracing the vision that BMIs can help people in need. In the past year, several traditional neurological laboratories have begun to pursue neuroprosthetic devices. Preliminary results from Arizona State University, Brown University and the California Institute of Technology have recently appeared. Some of the studies provide independent confirmation of the rat and monkey studies we have done. Researchers at Arizona State basically reproduced our 3-D approach in owl monkeys and showed that it can work in rhesus monkeys too. Scientists at Brown enabled a rhesus macaque monkey to move a cursor around a computer screen. Both groups recorded 10 to 20 neurons or so per animal. Their success further demonstrates that this new field is progressing nicely.

The most useful BMIs will exploit hundreds to a few thousand single neurons distributed over multiple motor regions in the frontal and parietal lobes. Those that record only a small number of neurons (say, 30 or fewer) from a single cortical area would never provide clinical help, because they would lack the excess capacity required to adapt to neuronal loss or changes in neuronal responsiveness. The other extreme--recording millions of neurons using large electrodes--would most likely not work either, because it might be too invasive.

Noninvasive methods, though promising for some therapies, will probably be of limited use for controlling prostheses with thoughts. Scalp recording, called electroencephalography (EEG), is a noninvasive technique that can drive a different kind of brain-machine interface, however. Niels Birbaumer of the University of Tübingen in Germany has successfully used EEG recordings and a computer interface to help patients paralyzed by severe neurological disorders learn how to modulate their EEG activity to select letters on a computer screen, so they can write messages. The process is time-consuming but offers the only way for these people to communicate with the world. Yet EEG signals cannot be used directly for limb prostheses, because they depict the average electrical activity of broad populations of neurons; it is difficult to extract from them the fine variations needed to encode precise arm and hand movements.

Despite the remaining hurdles, we have plenty of reasons to be optimistic. Although it may be a decade before we witness the operation of the first human neuroprosthesis, all the amazing possibilities crossed our minds that afternoon in Durham as we watched the activity of Belle's neurons flashing on a computer monitor. We will always remember our sense of awe as we eavesdropped on the processes by which the primate brain generates a thought. Belle's thought to receive her juice was a simple one, but a thought it was, and it commanded the outside world to achieve her very real goal.

Miguel A. L. Nicolelis and John K. Chapin have collaborated for more than a decade. Nicolelis, a native of Brazil, received his M.D. and Ph.D. in neurophysiology from the University of São Paulo. After postdoctoral work at Hahnemann University, he joined Duke University, where he now co-directs the Center for Neuroengineering and is professor of neurobiology, biomedical engineering, and psychological and brain sciences. Chapin received his Ph.D. in neurophysiology from the University of Rochester and has held faculty positions at the University of Texas and the MCP Hahnemann University School of Medicine (now Drexel University College of Medicine). He is currently professor of physiology and pharmacology at the State University of New York Downstate Medical Center.

© 1996-2005 Scientific American, Inc. All rights reserved.
Reproduction in whole or in part without permission is prohibited.