

Actions from thoughts

Miguel A. L. Nicolelis

Real-time direct interfaces between the brain and electronic and mechanical devices could one day be used to restore sensory and motor functions lost through injury or disease. Hybrid brain-machine interfaces also have the potential to enhance our perceptual, motor and cognitive capabilities by revolutionizing the way we use computers and interact with remote environments.

After a clever throw in by Tostão, a simple flick of Rivelino's magic left foot was enough to send the ball soaring into the thin air of the Azteca stadium in Mexico City. As the immaculate white object flew towards the middle of the penalty box on that hot afternoon, the colourful crowd that packed the stands slowly rose in anticipation. They roared, already celebrating, because they had seen that scene a thousand times before: the same graceful black man, dressed in blue shorts and a yellow jersey with the green 10 sewn in the back, defying logic, making fun of physics. The early celebration was warranted. As expected, Pelé floated above all Italian defenders to encounter the ball in mid air, and, with a gentle kiss of a forehead, changed its trajectory towards the net. Brazil had scored the first of its four goals in the final game of the 1970 World Cup and a whole country was about to start dancing in the streets.

The vast range of human abilities and behaviour illustrated here, as well as the gift of remembering the multitude of sensations associated with an instant of joy many decades ago, offer us a glimpse of the awesome repertoire of tasks that the human brain can accomplish. Through mechanisms that still elude our comprehension, the electrical activity of millions of brain cells (neurons) can be translated into precise sequences of skilled movements. Coordinated neuronal activity also provides us with exquisite perceptual and sensorimotor capabilities, illustrated in this example by Pelé's ability to track the ball's trajectory and plan the timing of his jump to hit it head on. But this is not all. Highly distributed patterns of neuronal firing underlie our ability to generate expectations about the outcome of a future event, learn the complex laws of nature and create art. One could argue, therefore, that hidden within the intricate principles that govern the way brain circuits operate lies the key to understanding the very essence of what it is to be human.

Witnessing the relentless growth of the disciplines that define modern neuroscience, one cannot help wondering what kind of insights, clinical applications and technologies may emerge from brain research in the future, and, more important, how they will impact on our lives. Although many of the imagined possibilities may not be feasible at this time, recent work indicates that some current ideas will come to fruition in the not-so-distant future. Here, I focus on one of these — the development of direct interfaces between machines and the human brain.

Brain-machine interfaces

I propose that the introduction of new methods for measuring large-scale brain activity, new techniques for microstimulating neuronal tissue, and emerging developments in microchip design, computer science and robotics have the potential to coalesce into a new technology devoted to creating interfaces between the human brain and artificial devices. One day, it is conceivable that such technology could allow patients to use brain activity to control electronic, mechanical or even virtual devices, leading to new therapeutic alternatives for restoring lost sensory, motor and even cognitive functions. Although many fundamental neurobiological questions and technical difficulties need to be solved, we can be optimistic about the feasibility of implementing this concept in the next few decades. Indeed, one brain-machine interface — the auditory prosthesis known as the cochlear implant — was introduced years ago and has improved the quality of life of many deaf patients^{1,2} (see Box 1).

Neuroscientists have long relished the possibility of using brain signals to control artificial devices³. As a consequence, there are already many terms in the literature⁴ to describe devices that could accomplish this goal (for example, brain-actuated technology, neuroprostheses or neurorobots). Here I will refer to these devices collectively as 'hybrid brain-machine interfaces (HBMI)s'. The word 'hybrid' reflects the fact that these applications rely on continuous interactions between living brain tissue and artificial electronic or mechanical devices.

My definition of HBMI)s incorporates two main types of application. Type 1 devices use

artificially generated electrical signals to stimulate brain tissue in order to transmit some particular type of sensory information or to mimic a particular neurological function. The classic example of this application is an auditory prosthesis. Future applications aimed at restoring other sensory functions, such as vision, by microstimulation of specific brain areas would also belong to this group. In addition, type 1 HBMI)s include methods for direct stimulation of the brain to alleviate pain, to control motor disorders such as Parkinson's disease⁵, and to reduce epileptic activity by stimulation of cranial nerves⁶. These last three applications rely on the observation that direct microstimulation of brain tissue can disrupt pathological patterns of brain activity that underlie some neurological disorders. Type 2 HBMI)s rely on the real-time sampling and processing of large-scale brain activity to control artificial devices. An example of this application would be the use of neural signals derived from the motor cortex to control the

Box 1 Cochlear implants: the first HBMI

Auditory prostheses work by converting features of acoustic signals, such as speech, into patterns of electrical stimuli that are then delivered through an array of chronically implanted electrodes to auditory nerve fibres lying on the basilar membrane of the cochlea. As the basilar membrane contains a representation of sound frequencies, known as a tonotopic map, auditory prostheses deliver high-frequency information to the basal region of the cochlea, and low-frequency signals to the apical region, to mimic normal auditory processing. More than 30,000 deaf patients, ranging in age from 12 months to 80 years, have had such devices successfully implanted². Although results vary from case to case, even slight improvements in auditory performance have helped people to communicate better and to become more aware of their surrounding environment.

insight feature

movements of a prosthetic robotic arm in real time. Obviously, clinical applications that require reciprocal interaction between the brain and artificial devices will combine both type 1 and 2 HBMI.

The design and implementation of HBMI will involve the combined efforts of many areas of research, such as neuroscience, computer science, biomedical engineering, very large scale integration (VLSI) design and robotics. I have selected a few current developments in these fields to illustrate below some of the conceptual advances and technologies that will be required to design and implement useful HBMI. I will then describe two potential clinical applications of such technology that should emerge in the near future: a system to monitor and treat epileptic seizures and a device to control a robotic prosthetic arm.

Building a HBMI

The first of the many challenges associated with the development of any HBMI is the need to understand better the principles by which neural ensembles encode sensory, motor and cognitive information. This is rapidly becoming one of the main goals of modern neuroscience, but our present knowledge is elementary at best. In the case of motor control, for instance, the areas of the primate brain involved are well known, and considerable information is available on the physiological properties of individual neurons located in each of them. But we know little about how the brain makes use of information from these neurons to generate movements. To design a type 2 HBMI that uses brain-derived signals to control a prosthetic robotic arm, we will need to learn how to sample and decode the motor signals generated by neurons and how to feed them into an artificial device to mimic the intended movement.

Recording brain activity

It is clear that neurobiological principles will be central in devising a strategy to overcome these hurdles. For example, classic experiments in primates have demonstrated that fundamental parameters of motor control emerge by the collective activation of large distributed populations of neurons in the primary motor cortex (M1). Single M1 neurons are broadly tuned to the direction of force required to generate a reaching arm movement⁷. In other words, even though these neurons fire maximally before the execution of a movement in one direction, they also fire significantly before the onset of arm movements in a broad range of other directions. Therefore, to compute a precise direction of arm movement, the brain may have to perform the equivalent of a neuronal 'vote' or, in mathematical terms, a vector summation of the activity of these broadly tuned neurons⁷. This implies that to obtain the motor signals required to control an artificial

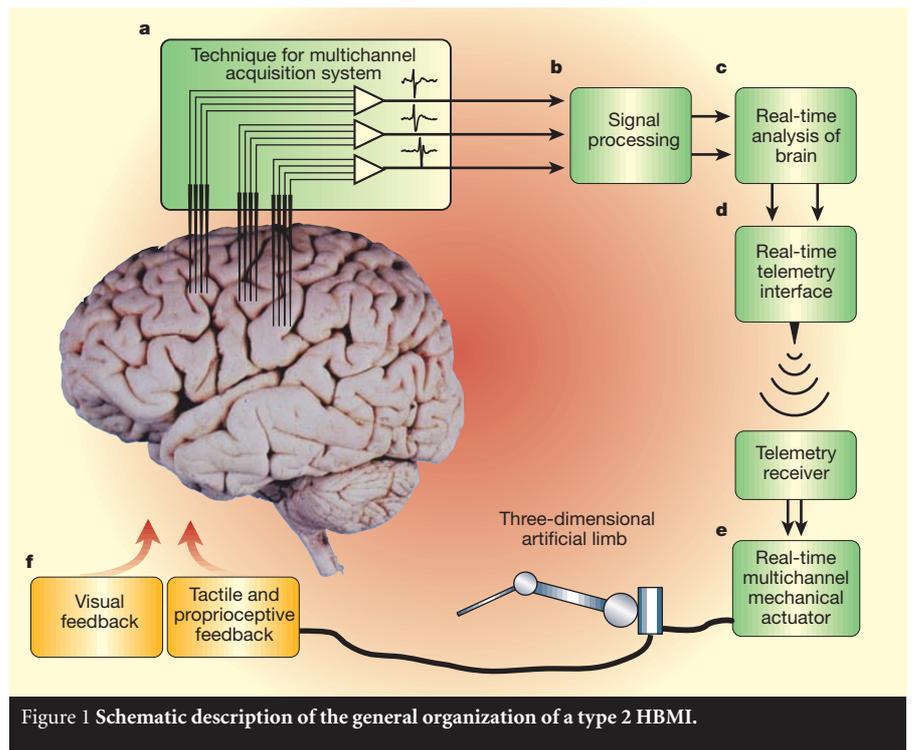


Figure 1 Schematic description of the general organization of a type 2 HBMI.

device we will need to sample the activity of many neurons simultaneously and design algorithms capable of extracting motor control signals from these ensembles. Moreover, it will be crucial to investigate how these neural ensembles interact under more complex and 'real-world' experimental conditions⁸ to generate different motor behaviours. These data will be vital to answering basic questions in regard to the development of type 2 HBMI. For example, what is the minimum neuronal sample required to generate reliable brain-derived control signals? Should these samples be obtained from one or multiple brain areas? Does the same population of neurons code for single or multiple control parameters? Finally, how might neural encoding mechanisms change with time, experience and learning?

Figure 1 illustrates the general organization of a type 2 HBMI and depicts some of the technological challenges involved in designing such devices. The first design step involves the selection of a technique (Fig. 1a) that yields reliable, stable and long-term recordings of brain activity that can be used as control signals to drive an artificial device. From recent studies in animals^{9,10}, clinical applications of HBMI will probably require sampling of large numbers of neurons (in the order of hundreds or thousands) with a temporal resolution of 10–100 ms, depending on the application.

Although neuroscientists have long recognized the need to investigate the properties of large neural ensembles¹¹, it is very difficult to obtain reliable, long-term measurements of neural ensemble activity with high spatial and temporal resolution. Starting in the

1940s and 50s with multichannel recordings of scalp electroencephalographic (EEG) activity and of the general electrical activity evoked by movement or sensory stimulation, a variety of metabolic, optical and electrophysiological methods have been introduced for monitoring large-scale brain activity. Modern multichannel electrophysiological recordings are made from arrays of microelectrodes surgically implanted in the brain. They currently allow neurophysiologists to record simultaneously, with a resolution of milliseconds, the extracellular activity of up to 100 individual neurons, distributed across multiple brain structures, in animals carrying out some task or other¹². Although future improvements might allow long-term and non-invasive sampling of human neural activity with the same temporal resolution as intracranial recordings, the first generation of HBMI will probably rely on improved versions of electrophysiological methods, such as multichannel EEG or multielectrode intracranial recordings. Indeed, preliminary studies in paralysed patients have shown that EEG signals can be used to trigger the movement of computer cursors¹³ or offer a way for patients to communicate¹⁴.

Unfortunately, less invasive electrophysiological methods, such as scalp EEG recordings that reflect the common electrical activity of millions of neurons in widespread areas of the cortex, lack the resolution to provide the kind of time-varying motor signals needed to control a robotic arm in real time⁴. Multichannel intracranial recordings of brain activity, obtained by surgical implantation of arrays of microwires within one or more cortical motor areas, will

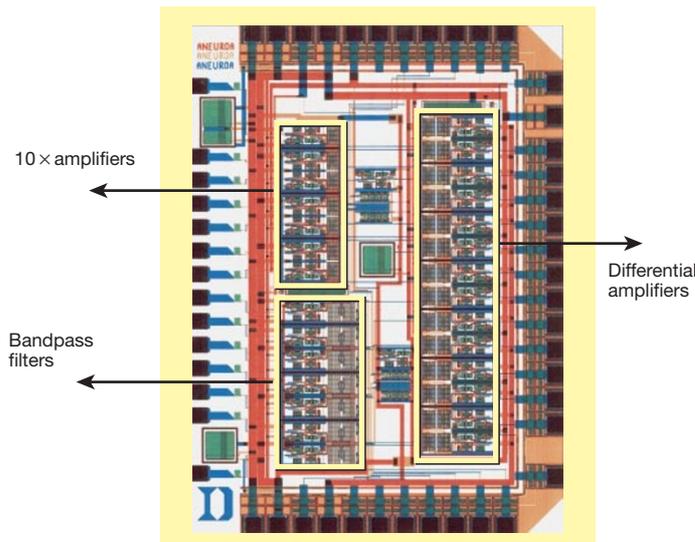


Figure 2 A prototype of an instrumentation neurochip for processing brain-derived signals. This chip, containing a portion of the analog signal processing for 16 neural channels, was designed by I. Obeid, H. Aurora, J. Morizio and P. Wolf in the Departments of Biomedical and Electrical & Computer Engineering at the Pratt School of Engineering, Duke University. The mixed-signal CMOS (complementary metal-oxide semiconductor) process used in the design supports digital signal processing modules which will be included in future generations of this device.

therefore be required, with mathematical analysis of the extracellular activity of smaller populations (100–1,000) of neurons providing the raw brain signals for use in most HBMI¹⁰. Despite some degree of recording degradation over time, present technology allows simultaneous sampling of 50–100 neurons, distributed across multiple cortical areas of small primates, to remain viable for several years^{10,12}. Technological advances in multielectrode array design and neural signal instrumentation in the next decade alone are expected to increase the number of neurons that can be recorded simultaneously by at least one order of magnitude.

The precise placement of the electrode arrays for intracranial recording may not be as critical to the ability to control an artificial device as was first conjectured. As motor control signals emerge from the distributed activation of large populations of neurons, and as cortical and subcortical neurons are capable of considerable plastic reorganization during adulthood¹⁵, electrode arrays targeted to brain areas of interest may suffice in most cases. As subjects learn to interact with artificial devices through HBMI, it is likely that sampled neurons that were not originally involved in the type of motor control to be mimicked may be recruited into generating the signals required to control artificial devices.

Generating the output

After selecting a method for acquiring the necessary brain signals, the next challenge is to design the instrumentation (Fig. 1b–d) required for recording and processing these

signals in real time. Currently, this requires specialized, sizeable and expensive electronic equipment, which can amplify and filter the original signal as well as perform analog-to-digital conversion to facilitate further processing and storage of data. To make HBMI viable, new technologies for portable, wireless-based, multichannel neural signal instrumentation are needed.

The central issue of signal conditioning and instrumentation may be solved in the near future by the application of mixed-signal VLSI into the design of neurophysiological instrumentation chips. This technology allows analog and digital signals to coexist in the same microchip, and has the potential to provide the multichannel, programmable and low-noise package required for conditioning brain-derived signals for clinical implementation of HBMI. Moreover, the resulting microchip would be small enough to be chronically implanted in patients and could be powered by replaceable batteries. Such microchips could rely on wireless communication protocols based on a radio frequency link to broadcast neural signals to other components of the HBMI (Fig. 1d–e).

Prototypes of dedicated ‘instrumentation neurochips’ (Fig. 2) are currently being developed, although many complex issues must be solved before they can become operational¹⁶. For instance, efficient solutions will have to be found to provide enough power for performing analog and digital processing, while still ensuring that signals can be transmitted by telemetry. Thus, battery technology, device packing and the bandwidth of the neural signals, among

other factors, will certainly be important in the design of HBMI¹⁶.

Having selected a method for sampling and conditioning brain signals, the next step—and one of the most difficult challenges—is to define a strategy for extracting meaningful control information from neural ensemble activity in real time. Currently, neuroscientists rely on a variety of linear and nonlinear multivariate algorithms, such as discriminant analysis, multiple linear regression and artificial neural networks, to carry out real-time and off-line analysis of neural ensemble data. Preliminary results from animal studies that use these different methods are encouraging, but considerably more experience is needed to apply these techniques in clinical HBMI. The challenge is to produce algorithms that can combine the activity of large numbers of neurons, which convey different amounts of information, and extract stable control signals, even when the firing patterns of these neurons change significantly across different timescales. Research on areas ranging from automatic sorting algorithms for unsupervised isolation of single neuron action potentials, to the design of real-time pattern recognition algorithms that can handle data from thousands of simultaneously recorded neurons will certainly be required. In the same context, clinical applications of HBMI will require considerable computational resources.

In the not too distant future, new developments in the design of brain-inspired VLSI¹⁷, an exciting area of research aimed at modelling neuronal systems in silicon¹⁸, may provide the means for achieving the type of efficient real-time neural signal analysis required for HBMI. This technology may allow pattern recognition algorithms, such as artificial neural networks or realistic models of neural circuits, to be implemented directly in silicon circuits. Among many other technical hurdles, significant work will be required to make these silicon circuits adaptive, perhaps by incorporating learning rules derived from the study of biological neural circuits. This will allow ‘training’ of algorithms as well as ensuring the robustness of the control system. From an implementation point of view, ‘analytical neurochips’ are ideal as they could be interfaced with the instrumentation neurochip and be chronically implanted in the subject.

The final component of the idealized HBMI (Fig. 1e–f) is a real-time control interface which uses processed brain signals to control an artificial device. The types of devices used are likely to vary considerably in each application, ranging from elaborate electrical pattern generators to control muscles, to complex robotic and computational devices designed to augment motor skills¹⁹.

HBMI for epilepsy control

Estimates indicate that about 0.5–2.0% of the population has epilepsy²⁰. About 10–50% of these patients do not respond well to current

antiepileptic medications and may not be candidates for surgery. Throughout this century, neuroscientists have used multi-channel recording from scalp, brain surface and even chronically implanted intracranial electrodes to investigate the electrophysiological activity that characterizes different types of seizure in humans. By doing so, scientists have not only identified different types of epilepsy, but they have also learned that there are distinct patterns of neurophysiological activity associated with the initiation and establishment of a seizure attack.

Several exciting new developments in epilepsy research indicate that the development of an unsupervised HBMI for monitoring, detecting and treating seizure activity may be possible in the next decade (Fig. 3a). First, for certain types of seizure, there seems to be a particular spatiotemporal pattern of cortical activity that appears seconds or even minutes before the full epileptic attack starts²¹. Recently, a few laboratories have introduced automatic seizure-prediction algorithms that can be applied to intracranial and scalp recordings to forecast the occurrence of a seizure^{21,22}. These and future seizure-prediction algorithms might provide sufficient time (2–5 minutes) to warn the patient of an imminent attack, and to trigger automatic therapeutic intervention before convulsion or loss of consciousness.

But what kind of therapy could be triggered that would work in patients who are refractory to epilepsy medication? The answer may lie in another recent development in epilepsy research. Studies in both animals²³ and human subjects⁶ have revealed that electrical stimulation of peripheral cranial nerves, such as the vagus²³ and trigeminal²⁴ nerves, can substantially reduce cortical epileptic activity. Moreover, if this peripheral nerve stimulation is applied before the initiation of seizure or during its initial stages, significantly higher reduction of seizure activity can be achieved.

From this I believe that a device containing a combination of both type 1 and 2 HBMI could be designed to function somewhat like a modern heart pacemaker (Fig. 3a). This ‘brain pacemaker’ would rely on arrays of chronically implanted electrodes to search continuously for spatiotemporal patterns of cortical activity indicating an imminent epileptic attack. Instrumentation neurochips would be responsible for all the basic signal-processing operations. They would also provide signals to one or more seizure-prediction algorithms, implemented into analytical neurochips, which would carry out real-time analysis of cortical activity. Once pre-seizure activity patterns were detected, the analytical neurochip could trigger electrical stimulation of one or multiple cranial nerves. In patients who respond to pharmacological therapy, the same stimulator could be used to activate a mini-pump to deliver one or more anti-epileptic

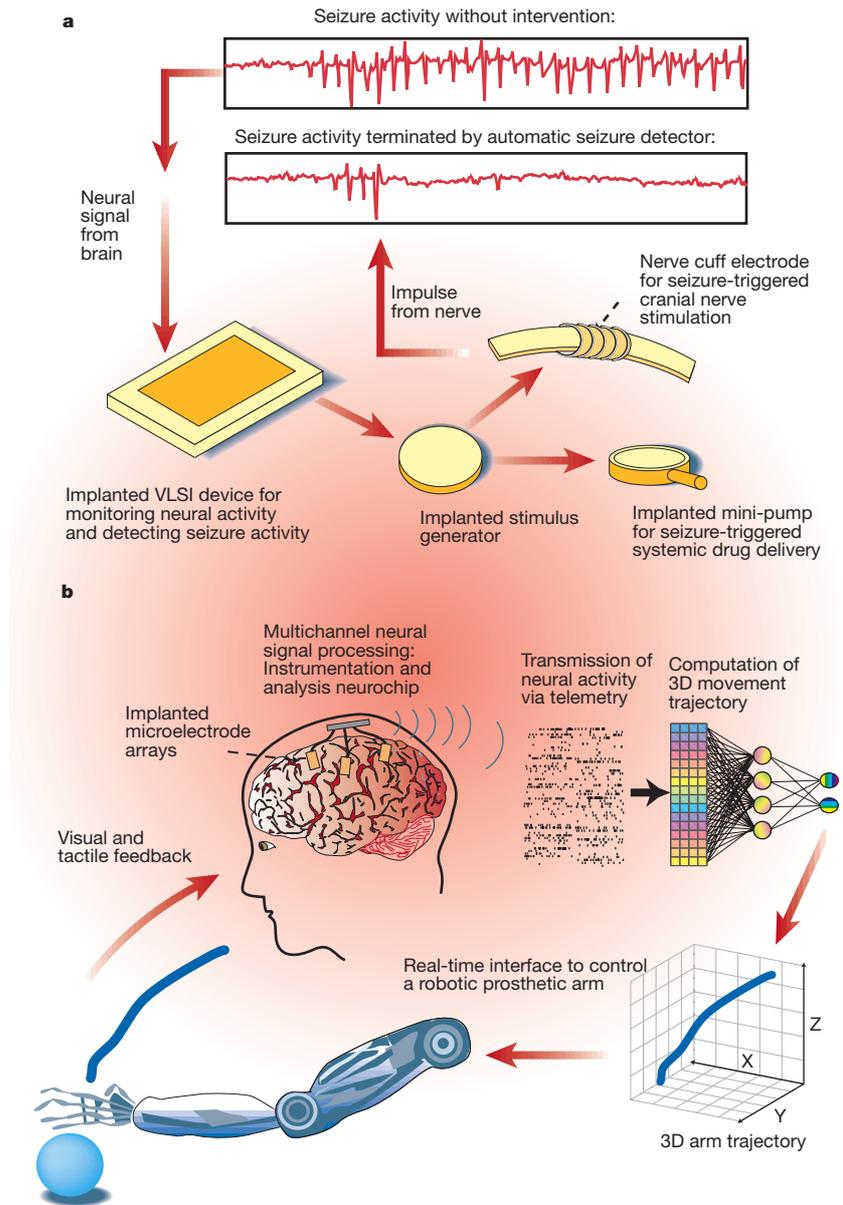


Figure 3 Schematic description of two potential applications of type 2 HBMI. **a**, Design of a ‘brain pacemaker’ that monitors neural activity using a VLSI chip designed to detect seizure activity. When seizure activity is detected, the VLSI chip sends a signal to an implanted stimulus generator that drives either a nerve cuff electrode or a mini-pump for drug delivery, either of which can stop the seizure activity. **b**, HBMI for controlling a robotic prosthetic arm using brain-derived signals. Multiple, chronically implanted, intracranial microelectrode arrays would be used to sample the activity of large populations of single cortical neurons simultaneously. The combined activity of these neural ensembles would then be transformed by a mathematical algorithm into continuous three-dimensional arm-trajectory signals that would be used to control the movements of a robotic prosthetic arm. A closed control loop would be established by providing the subject with both visual and tactile feedback signals generated by movement of the robotic arm.

drugs directly into the blood stream. Recently, a simplified implementation of this concept has been used successfully in rats²⁴, giving hope that a brain pacemaker for seizure monitoring and control in humans may not be far ahead.

HBMI to restore motor function

Another clinical application of HBMI that could emerge in the near future aims at

restoring different aspects of motor function in patients with severe body paralysis, caused primarily by strokes, spinal cord lesions or peripheral degenerative disorders (Fig. 3b). Advances in this rapidly growing field of research indicate that neural signals from healthy regions of the brain could be used to control the movements of artificial prosthetic devices, such as a robotic arm. Preliminary

findings also demonstrate that paralysed patients can learn to use brain signals obtained from their motor cortex to interact with computers²⁵.

Extensive electrophysiological work in primates and imaging studies in humans have shown that multiple interconnected cortical areas in the frontal and parietal lobes are involved in the selection of motor commands that control the production of voluntary arm movements²⁶. Although each of these areas has different degrees of functional specialization, in theory, each of them could be selected as the source of brain signals for controlling the movements of an artificial device. Within each of these cortical areas, different motor parameters, such as a force and direction of movement, are coded by the distributed activity of populations of neurons, each of which is typically broadly tuned to one (or more) of these parameters. This indicates that implementations of HBMs for robotic arm control need to rely on intracranial recordings from large populations of single neurons to derive motor control signals.

At a first glance, a random sample of 100–1,000 cortical motor neurons, which represents a reasonable expectation for the yield of multielectrode intracranial recordings in the near future, may look too small to unveil any useful information. But recent neurophysiological experiments dispute this view. For instance, currently one can obtain precise off-line reconstructions of complex three-dimensional arm trajectories by using simple multiple regression techniques to transform the activity of 300–400 serially recorded cortical motor neurons into a neural population vector²⁷. Moreover, experiments in rats⁹ and primates¹⁰ have shown that simple, real-time algorithms, applied to samples of 50–100 simultaneously recorded cortical neurons, can be used to control robotic devices in real time and mimic the type of three-dimensional arm reaching movements produced by primates.

Another important issue is that, to achieve seamless interactions with prosthetic devices, patients will have to receive sensory feedback information (for example visual or tactile signals) from the prosthetic limbs. These feedback signals will establish a closed control loop between the brain and artificial devices and will probably help patients learn how to operate HBMs. Studies in rats have revealed that, if subjects receive visual feedback information as they learn to use brain activity to interact with a robotic arm, and are rewarded for the successful completion of these movements, they progressively cease to produce overt limb movements⁹. In other words, even though the rats continued to exhibit the patterns of cortical activity required to control the movements of the robotic arm, this motor activity did not result in any significant limb movement. This indicates that motor control signals can be generated by cortical neurons

without any muscle activity, and hence that paralysed patients might be capable of learning to operate a robotic arm even though they cannot move their own limbs.

These observations also raise the intriguing hypothesis that, by establishing a closed control loop with an artificial device (Fig. 3b), the brain could incorporate electronic, mechanical or even virtual objects into its somatic and motor representations, and operate upon them as if they were simple extensions of our own bodies. The fact that the adult cortex is capable of significant functional reorganization (or plasticity) after peripheral and central injuries¹⁵, changes in sensory experience²⁸ and learning of new motor skills²⁹ supports this possibility. Indeed, the notion that adult plasticity can dynamically alter the perception of the limits of our own body is corroborated by studies on patients who have undergone limb amputations. Immediately after the amputation, most of these patients experience the sensation that their amputated limb is still present and moving. These 'phantom limb' sensations are paralleled by a significant plasticity of body maps in the somatosensory cortex³⁰, the part of the brain that receives and interprets sensory signals from areas such as the skin surface. Instead of remaining silent, the areas in these brain maps that used to represent the amputated limb progressively start to respond to stimulation of neighbouring body regions spared by the amputation. Thus, it is conceivable that tactile feedback signals, generated by the movements of a brain-controlled robotic arm and delivered to the patient's skin, could be used to incorporate the representation of such an artificial device into cortical and subcortical somatotopic maps.

Undoubtedly, years of research will be required and many fundamental technological breakthroughs needed before this comes close to reality. Nevertheless, it seems reasonable to predict that a definitive demonstration of such a phenomenon could trigger a revolution in the way future generations interact with computers, virtual objects and remote environments, by allowing never-before-experienced augmentation of perceptual, motor and cognitive capabilities. Such applications, however, will require the introduction of new, non-invasive methods for sampling brain activity.

A final thought

Although developing expectations of a distant future is a risky business and may raise unjustified hope that solutions are just around the corner, I cannot avoid ending this brief overview on an optimistic tone. Despite many significant conceptual and technological obstacles, the possibility of developing clinical applications of HBMs is real and worth pursuing, especially given the potential benefits that they may bring to people afflicted by neurological disorders. At the very least,

research on HBMs will yield powerful new tools to investigate hypotheses of how large populations of neurons process information and adapt according to changes in experience. Some may argue that one could achieve this goal just by building theoretical models and running computational simulations. Perhaps that is true. But as my good friend Idan Segev, a leading computational neuroscientist, always tells me, there is a subtle but fundamental difference between simulating reality and building it. Those of us who saw Pelé scoring that magic goal on that hot Mexican afternoon in 1970 and dreamed about doing the same thing would certainly agree.

Miguel A. L. Nicolelis is in the Departments of Neurobiology, Experimental Psychology, and Biomedical Engineering, Duke University, Durham, North Carolina 27710, USA

(e-mail: nicoleli@neuro.duke.edu)

- Merzenich, M. M., Schindler, D. N. & White, M. W. *Laryngoscope* **84**, 1887–1893 (1974).
- Pfingst, B. E. in *Neural Prostheses for Restoration of Sensory and Motor Function*. (eds Chapin, J. K. & Moxon, K. A.) 3–43 (CRC, Boca Raton, 2000).
- Schmidt, E. M. *Ann. Biomed. Eng.* **8**, 339–349 (1980).
- Chapin, J. K. & Moxon, K. A. (eds) *Neural Prostheses for Restoration of Sensory and Motor Function*. (CRC, Boca Raton, 2000).
- Benabid, A. L. *et al. Lancet* **337**, 403–406 (1991).
- Uthman, B. M., Wilder, B. J., Hammond, E. J. & Reid, S. A. *Epilepsia* **31**(Suppl. 2), S44–S50 (1990).
- Georgopoulos, A. P., Swartz, A. B. & Ketter, R. E. *Science* **233**, 1416–1419 (1986).
- Ghazanfar, A. A. & Hauser, M. D. *Trends Cog. Sci.* **3**, 377–384 (1999).
- Chapin, J. K., Moxon, K. A., Markowitz, R. S. & Nicolelis, M. A. L. *Nature Neurosci.* **2**, 664–670 (1999).
- Wessberg, J. *et al. Nature* **408**, 361–365 (2000).
- Hebb, D. O. *The Organization of Behaviour: A Neuropsychological Theory* (Wiley, New York, 1949).
- Nicolelis, M. A. L. *et al. Nature Neurosci.* **1**, 621–630 (1998).
- Wolpaw, J. R., McFarland, D. J., Neat, G. W. & Forneris, C. A. *Electroencephalogr. Clin. Neurophysiol.* **78**, 252–259 (1991).
- Schutz, S. *et al. Nature* **398**, 297–298 (1999).
- Wu, C. W. & Kaas, J. H. *J. Neurosci.* **19**, 7679–7697 (1999).
- Moxon, K. A., Morizio, J., Chapin, J. K., Nicolelis, M. A. L. & Wolf, P. D. in *Neural Prostheses for Restoration of Sensory and Motor Function*. (eds Chapin, J. K. & Moxon, K. A.) (CRC, Boca Raton, 2000).
- Hahnloser, R. H., Sarpeshkar, R., Mahowald, M. A., Douglas, R. J. & Seung, H. S. *Nature* **405**, 947–951 (2000).
- Mead, C. *Analog VLSI and Neural Systems* (Addison-Wesley, Reading, MA, 1989).
- Srinivasan, M. A. in *In Virtual Reality: Scientific and Technical Challenges* (eds Durlach, N. I. & Mavour, A. S.) 161–187 (National Academy Press, 1994).
- McNamara, J. O. *Nature* **399**(Suppl.), A15–A22 (1999).
- Martinerie, J. *et al. Nature Med.* **4**, 1173–1176 (1998).
- Webber, W. R. S., Lesser, R. P., Richardson, R. T. & Wilson, K. *Electroencephalogr. Clin. Neurophysiol.* **98**, 250–272 (1996).
- Zabara, J. *Epilepsia* **33**, 1005–1012 (1992).
- Fanselow, E. E., Reid, A. P. & Nicolelis, M. A. L. *J. Neurosci.* **20**, 8160–8168 (2000).
- Kennedy, P. R. & Bakay, R. A. *NeuroReport* **9**, 1707–1711 (1998).
- Wise, S. P., Boussaoud, D., Johnson, P. B. & Caminiti, R. *Annu. Rev. Neurosci.* **20**, 25–42 (1997).
- Schwartz, A. *Science* **265**, 540–542 (1994).
- Polley, D. B., Chen-Bee, C. H. & Frostig, R. D. *Neuron* **24**, 623–637 (1999).
- Laubach, M., Wessberg, J. & Nicolelis, M. A. L. *Nature* **405**, 567–571 (2000).
- Ramachandran, V. S. *Proc. Natl Acad. Sci. USA* **90**, 10413–10420 (1993).

Acknowledgements. I thank J. Chapin, F. Ebner, E. Fanselow, A. Ghazanfar, C. Henriquez, J. Kaas, J. Kralk, D. Krupa, S. Ribeiro, V. de Sa, I. Segev, M. Shuler, S. Simon, J. Wessberg and P. Wolf for comments and suggestions, and E. Fanselow, D. Krupa, P. Beck and J. Wessberg for their help in creating the illustrations for this article. The author's research on HBM is funded by NIH, DARPA and ONR grants.