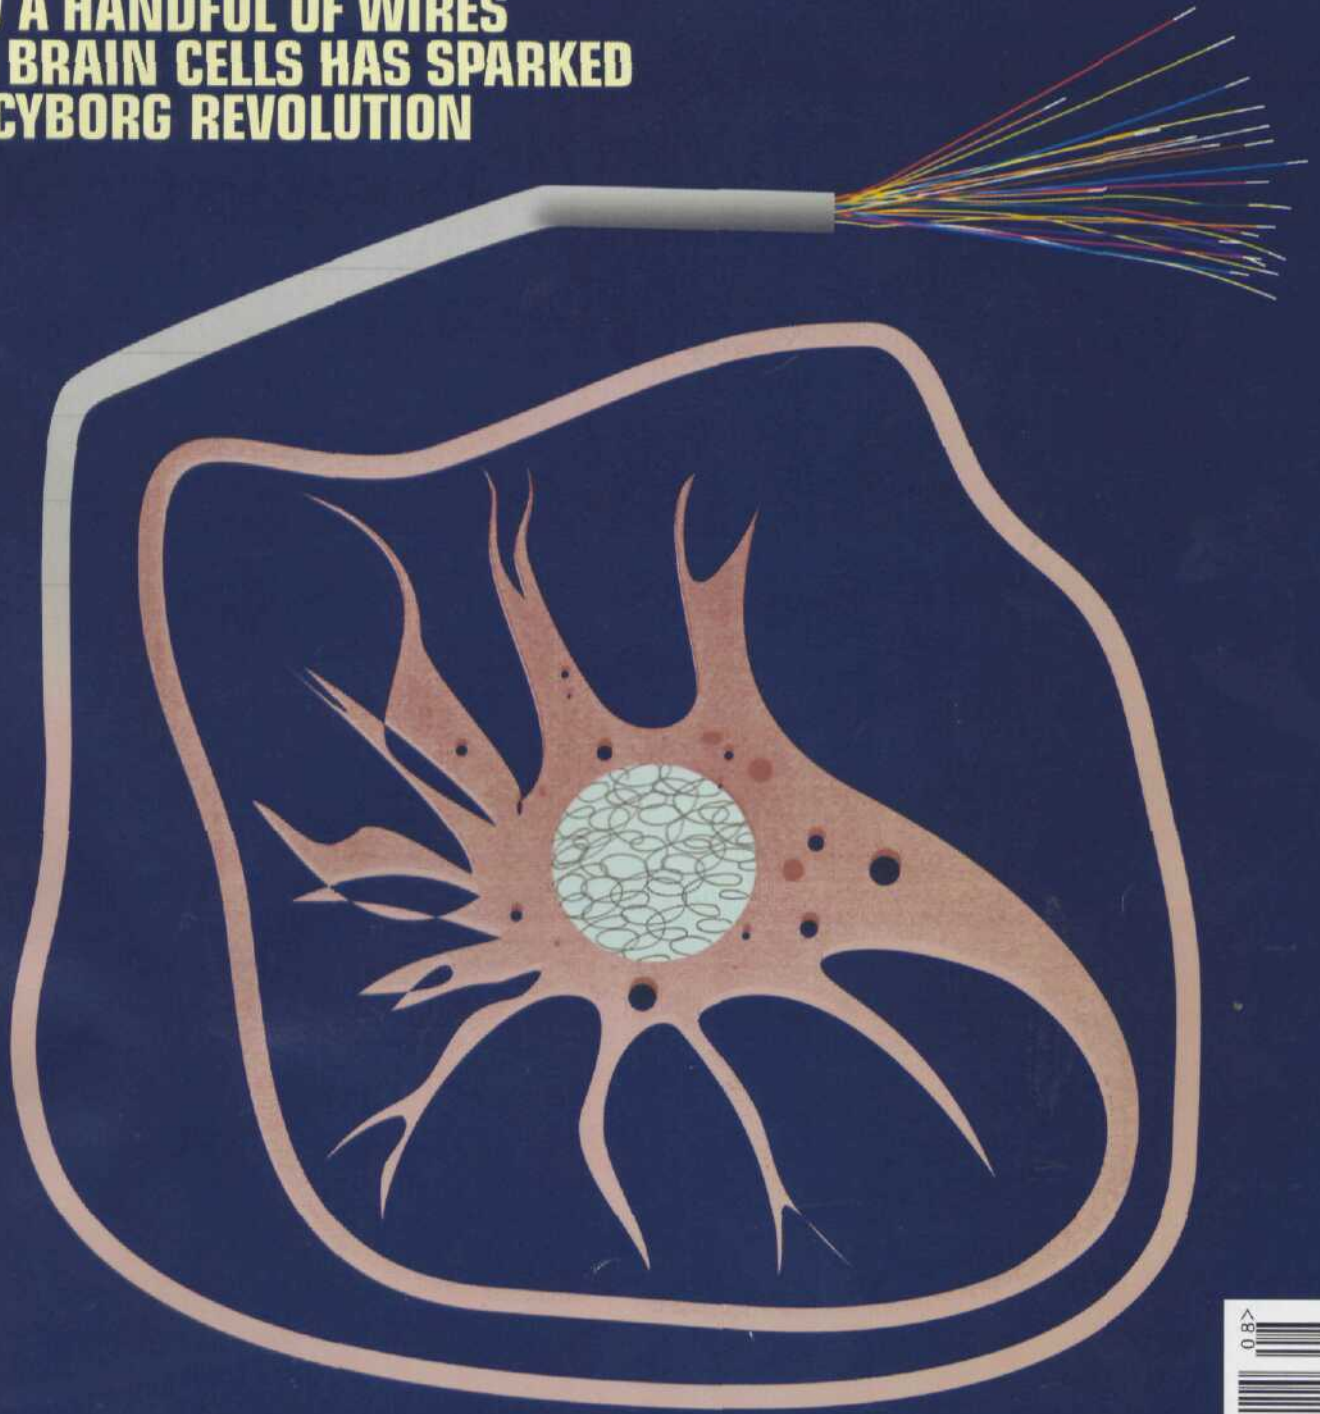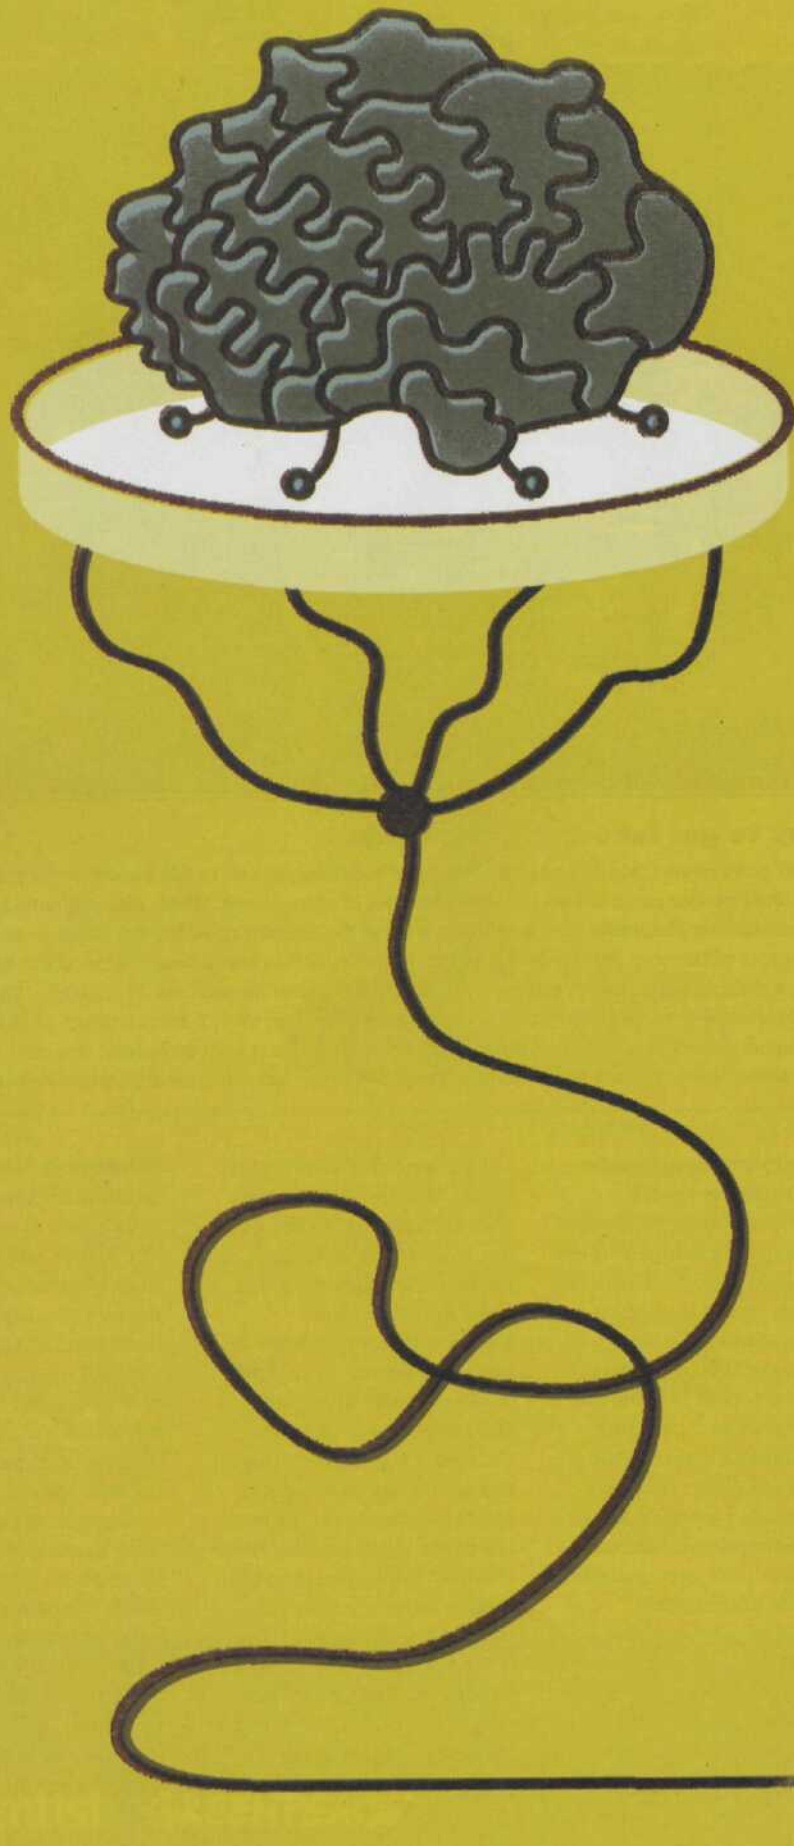23 FEBRUARY 2002 No 2331 WEEKLY £2.20 US $3.95

IS IT ALIVE?

HOW A HANDFUL OF WIRES AND BRAIN CELLS HAS SPARKED THE CYBORG REVOLUTION

WHY SLEEP IS A GAME OF CHANCE   CLONING AND THE COPYCAT KITTEN

MIND OVER METAL

The cyborgs are here and they're helping
us understand how the brain works.
Anil Ananthaswamy investigates

IT'S certainly different, some might even say distasteful, but there's no doubt that Steve Potter has the strangest lab rat around. Its body is a virtual one, running around inside a computer generated world. And its brain is like nothing you've ever seen — an amorphous mass of brain cells from a real rat, living in a shallow glass dish and wired up to the computer. Potter wants to see if the "rat" can learn its way around, which means the neurons must sense and remember the rat's virtual world and control its movements, just as a real brain would.

For decades, neuroscientists have struggled to understand just how tangles of neurons let us carry out everyday wonders such as learning and moving. They've studied the behaviour of single neurons, eavesdropped on the electrical signals from various areas of the brain and even scanned the brains of people carrying out complex tasks. But answers have been hard to come by.

It's difficult to get neurons isolated in a dish to do what they would normally do, and it's tricky to see what's going on at a cellular level when you're studying people inside a scanner. What's needed is a way to bridge the gap, a way to see neurons close up as they interact with the outside world.

Potter, a neuroscientist at Caltech in Pasadena, is one of several researchers who think they've found a way. Some have built artificial machines in which real nervous tissue connects directly to electronics. Others are taking brainstems from fish to control robots or co-opting the brain of a live monkey to move an artificial arm. As the living neural networks adapt to their cybernetic parts, the researchers hope to see the learning process in action.

But it doesn't end there. Some scientists are toying with the idea of tapping directly into the computing power of neurons, to build "neurocomputers" far more powerful than their silicon cousins. They would have to train brain cells to perform tasks that they wouldn't normally do — an ambitious idea, not to mention an ethical minefield — yet the work is well on its way.

William Ditto, co-founder of the Laboratory for Neuroengineering at the Georgia Institute of Technology in Atlanta, argues that the first simple neurocomputers may be built within five years. Eventually, they'll carry out tasks that conventional computers cannot seem to cope with, such as recognising handwriting, speech and faces — things our brains can do in the blink of an eye when only fragments of information are available. Biological systems could also be better at routeing Internet traffic and navigating autonomous vehicles, which are surprisingly complex tasks in computing terms.

Of course, it was hoped that artificial neural networks — silicon versions of the real thing — would solve these problems. But these haven't lived up to their promise. Even a single neuron is a complicated information-processing system, with thousands of inputs and outputs, combining and modulating signals in vastly more complex ways than we understand. We are nowhere near emulating one cell, let alone a network of neurons, says Joel Davis, an expert on adaptive neural systems at the Office of Naval Research in Arlington, Virginia.

For Peter Fromherz, a neuroscientist from the Max Planck Institute for Biochemistry in Munich, neurocomputing is still a while away. He is working on ways to mix biology and silicon so that he can probe individual neurons to study learning and memory.

Last year, his team immobilised two snail neurons between silicon structures that looked like picket fences etched onto a chip. Within two days the neurons grew extensions to connect with each other. They could then exchange electrical signals with one another, or with electrodes on the chip. The connection allows him to see precisely how the cell responds to electrical signals, and with more cells he hopes to study how a whole network changes physically to store a memory. "We have the basic components for integrating digital electronics with neural networks," says Fromherz. "The next step is to have more neurons on the chip. The goal is to make a very small learning network."

But it's too early to speculate where this may lead, says Fromherz. It's true that we could use living neural networks without understanding how they work, but there are other problems. The notion of anchoring neurons to a chip wouldn't work for a neurocomputer of any significance, he says, simply because it would need hundreds of thousands or even millions of neurons. The neurons would have to be free to move around, making and breaking connections as they please.

Potter has gone at least some way towards tackling this problem. While he doesn't have the precise control over individual neurons that Fromherz has, Potter is interested in how networks of neurons solve complex problems, and so has thousands of neurons growing in wells etched in silicon, and 60 electrodes that pick up the constant mass of electrical activity of the network. When hundreds of neurons are connected together, certain behaviours start to emerge. "I'm hoping that some of these emergent properties will be applicable to our artificial neural networks, or perhaps to actual physical hardware," he says.

In Potter's experiment, a program analyses the neurons' electrical outputs, looking for repeatable patterns of activity. The neurons seem to produce regular patterns, each of which Potter arbitrarily assigns to control a specific movement of the virtual rat. So
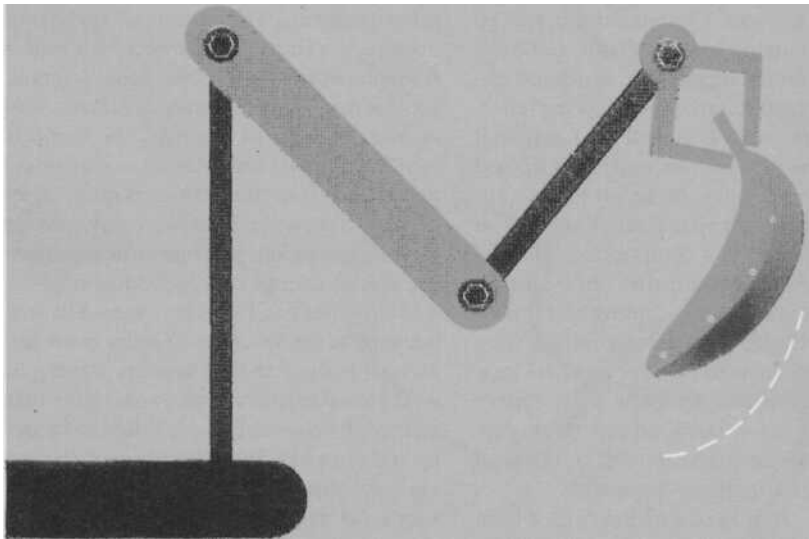
pattern one might make it move to the right, pattern two to the left, say. When the neurons are active, the rat moves about according to the sequence of patterns they generate.

Potter also provides the neurons with feedback. Information about the rat's movements and "sensory" signals—from a kind of artificial touch and vision system—is sent back to the neurons via the electrodes. The neurons respond by changing their firing patterns and connections—making the rat move differently (see Diagram, p 29). These changes should reveal just how adaptable the brain is. Given the right feedback, Potter hopes that his artificial animal—he calls it an animat—will become familiar with its virtual surroundings and that this learning will be reflected in changing signals and connections within the neural network. Ultimately, if the animat bumps into a virtual wall, it should learn to avoid the wall next time round.

Whether or not this will actually happen, it's too early to tell. But research in Chicago suggests he might be on the right track. Neuroscientist Sandro Mussa-Ivaldi of Northwestern University Medical School believes he's already seen learning in action. Instead of using dissociated brain cells, Mussa-Ivaldi's team is using whole brainstems extracted from lampreys—eel-like fish older than dinosaurs—to control a modern-day robot. Like Potter, he's hoping to develop an experimental system that can be used to understand the properties of biological neural networks.

The robot has light sensors, and the electrical signal they generate is fed into one of the neural circuits in the brainstem. The circuit generates signals that are fed back to the robot's wheels. This odd combination of organic brain and inorganic body creates a robot that may move towards a light or away from it.

Mussa-Ivaldi took a robot that had an affinity for light, and held it motionless for five minutes with an intense light shining on its left side. When he freed the robot into normal light, the researchers got a big surprise. It ignored lights on the left and moved only towards lights on the right. It seemed to have learned to override its natural affinity.

While the robots are giving Mussa-Ivaldi the opportunity to design experiments to study that learning process, he's also interested in how we might imitate it, even how we might tap into the neural tissue to develop new methods of computation. The ability to adapt, as his lamprey robot did, is what will set apart living neural computers from conventional ones, he believes. The challenge will be to learn enough about such behaviour to control it so that robots learn to do useful tasks.

But there are moral issues to be resolved before we move too far down this road.

The death of animals is accepted by society because there's no other way of finding out the details of how neurons keep us thinking and moving, and because it could lead eventually to treatments for brain disorders such as Alzheimer's and Parkinson's diseases, perhaps even paralysis. But creating a new type of computer or robot is clearly not reason enough to sacrifice animals. It might be more acceptable if we could use neurons grown entirely in culture, from stem cells perhaps. Indeed Potter has kept alive neuronal cultures for 18 months now.

Using brain tissue within machines raises some unique questions, however, says Paul Root Wolpe of the Center for Bioethics at the University of Pennsylvania, and NASA's chief bioethicist. "You are creating an organism that by its very definition could not exist in nature," says Wolpe. But he insists that as long as the tissue is collected in a way that's ethically acceptable, without causing pain or distress, it doesn't mean we should never build such organisms. "What I'm saying is then we need to start to have a moral conversation about the implications—about how far we should take them, and what we shouldn't do."

That conversation is not going to be easy. For instance, how do we know if the "creature" we have built is capable of feeling pain? What are its rights? "What really matters here is what the creature is capable of sensing and doing," says philosopher and cognitive scientist Andy Clark of the University of Sussex. The complexity of the robot's behaviour should help us judge it. "That's roughly how we judge the animals we find in the world," says Clark. Something with the complex behaviour of a dog or a chimp would raise greater ethical concerns something that behaved more like a worm.

Miguel Nicolelis and his colleagues at Duke University in Durham, North Carolina, are raising ethical issues of a different type. They are probing the brain of a living animal. It sounds gruesome, but the aim is clear. Nicolelis wants to understand how large populations of cells interact and process information so that he can develop interfaces that control prosthetics for people who are paralysed. At the moment, his team is learning to control a robot arm using signals from an owl monkey's brain. As the monkey reaches for food, surgically implanted microelectrodes in its cortex read the neural signals. A computer analyses these signals, recognises patterns of activity in the brain, and predicts the motion of the monkey's arm. These predictions guide the robotic arm so that when the monkey moves its arm the robotic arm moves too — in eerie coordination.

Does this mean the monkey's brain is being used as a computer? In a very real sense, yes, says Ditto. And by doing these kinds of experiments, we get to understand the brain better, which could eventually help us to control it for other purposes.

But getting the brain to do our bidding would require feedback. For instance, if the robotic arm was reaching for an object, sending feedback to the monkey's brain could help the brain learn how to do it right. Nicolelis's team has just finished developing a system for delivering such feedback. One is visual — the monkey can watch the position of the robot arm displayed as a dot on a screen. The other is tactile — the monkey's skin can be stimulated to provide feedback. It might even be possible to directly stimulate sensory regions of the brain.

The monkey's brain will have to figure out what the feedback means—just as people with cochlear implants must learn to interpret the new spectrum of sounds they hear. What's crucial is that the brain will be learning about an external environment, while the environment, represented by the computer, is learning about the monkey's brain signals. "It's a two-way street. In essence you are generating an interaction between a brain and an artificial device," says Nicolelis.

Mandayam Srinivasan from the Massachusetts Institute of Technology, who worked with Nicolelis on this project, conjures up an even more fantastic future. One day human brains might be wired up to transmit signals and receive feedback. Astronauts could be trained using a virtual reality system where objects bounce as they would on Mars. Before leaving Earth the astronaut's brain could learn how to handle gravity on the Red Planet.

For Wolpe, the idea of tapping into brains—human or animal—raises the biggest issues. Should we be allowed to control other brains? "If we implant things in people's brains that allow impulses to move from the brain out to the world, it isn't that much of a conceptual leap to think of impulses moving from the world back into the brain in such a way as to affect its functioning or its cognitive ability," he says. "We begin to think about remotely controlling brains, including human brains. Then who has the right or the ability to do that?"

The obvious answers—that we should never use this kind of technology to control people's thought processes or to turn animals into slaves—are the easy ones, says Wolpe. But what about the grey areas, where things are not so easily defined? "The real ethical struggle will happen over the more subtle issues, and we don't know what those subtle issues will be."

'YOU WOULD BE CREATING AN ORGANISM THAT BY ITS VERY DEFINITION COULD NOT EXIST IN NATURE'

Tiny electrodes pick up activity from brain cells

Electrical patterns are decoded by computer and used to move the rat's virtual body

Rat brain cells

Sensory feedback from virtual eyes and whiskers

Neurally controlled animat

Digital feedback turned into electrical signals and passed back to brain cells

Anil Ananthaswamy is a science writer based in Berkeley, California