

The  
Economist

JANUARY 6TH-12TH 2018

The message of protests in Iran

Blue-state Republicans: no right left

Education lessons from Pakistan

The world's worst airports

# The next frontier

When thoughts control machines



## The next frontier

# Using thought to control machines

*Brain-computer interfaces may change what it means to be human*



Print edition | Leaders

Jan 4th 2018

TECHNOLOGIES are often billed as transformative. For William Kochevar, the term is justified. Mr Kochevar is paralysed below the shoulders after a cycling accident, yet has managed to feed himself by his own hand. This remarkable feat is partly thanks to electrodes, implanted in his right arm, which stimulate muscles. But the

real magic lies higher up. Mr Kochevar can control his arm using the power of thought. His intention to move is reflected in neural activity in his motor cortex; these signals are detected by implants in his brain and processed into commands to activate the electrodes in his arms.

An ability to decode thought in this way may sound like science fiction. But brain-computer interfaces (BCIs) like the BrainGate system used by Mr Kochevar provide evidence that mind-control can work. Researchers are able to tell what words and images people have heard and seen from neural activity alone. Information can also be encoded and used to stimulate the brain. Over 300,000 people have cochlear implants, which help them to hear by converting sound into electrical signals and sending them into the brain. Scientists have “injected” data into monkeys’ heads, instructing them to perform actions via electrical pulses.

**Latest updates**

Will Donald Trump’s pro-coal policies put miners’ lives at risk?  
GRAPHIC DETAIL

---

Two security flaws in modern chips cause big headaches for the tech business  
SCIENCE AND TECHNOLOGY

---

Are America’s airports the worst in the world?  
GULLIVER

See all updates

As our [Technology Quarterly](https://www.economist.com/technology-quarterly/2018-01-06/thought-experiments) (<https://www.economist.com/technology-quarterly/2018-01-06/thought-experiments>) in this issue explains, the pace of research into BCIs and the scale of its ambition are increasing. Both America’s armed forces and Silicon Valley are starting to focus on the brain. Facebook dreams of thought-to-text typing. Kernel, a startup, has \$100m to spend on neurotechnology.

Elon Musk has formed a firm called Neuralink; he thinks that, if humanity is to survive the advent of artificial intelligence, it needs an upgrade. Entrepreneurs envisage a world in which people can communicate telepathically, with each other and with machines, or acquire superhuman abilities, such as hearing at very high frequencies.

These powers, if they ever materialise, are decades away. But well before then, BCIs

could open the door to remarkable new applications. Imagine stimulating the visual cortex to help the blind, forging new neural connections in stroke victims or monitoring the brain for signs of depression. By turning the firing of neurons into a resource to be harnessed, BCIs may change the idea of what it means to be human.

### **That thinking feeling**

Sceptics scoff. Taking medical BCIs out of the lab into clinical practice has proved very difficult. The BrainGate system used by Mr Kochevar was developed more than ten years ago, but only a handful of people have tried it out. Turning implants into consumer products is even harder to imagine. The path to the mainstream is blocked by three formidable barriers—technological, scientific and commercial.

Start with technology. Non-invasive techniques like an electroencephalogram (EEG) struggle to pick up high-resolution brain signals through intervening layers of skin, bone and membrane. Some advances are being made—on EEG caps that can be used to play virtual-reality games or control industrial robots using thought alone. But for the time being at least, the most ambitious applications require implants that can interact directly with neurons. And existing devices have lots of drawbacks. They involve wires that pass through the skull; they provoke immune responses; they communicate with only a few hundred of the 85bn neurons in the human brain. But that could soon change. Helped by advances in miniaturisation and increased computing power, efforts are under way to make safe, wireless implants that can communicate with hundreds of thousands of neurons. Some of these interpret the brain's electrical signals; others experiment with light, magnetism and ultrasound.

Clear the technological barrier, and another one looms. The brain is still a foreign country. Scientists know little about how exactly it works, especially when it comes to complex functions like memory formation. Research is more advanced in animals, but experiments on humans are hard. Yet, even today, some parts of the brain, like the motor cortex, are better understood. Nor is complete knowledge

always needed. Machine learning can recognise patterns of neural activity; the brain itself gets the hang of controlling BCIS with extraordinary ease. And neurotechnology will reveal more of the brain's secrets.

### **Like a hole in the head**

The third obstacle comprises the practical barriers to commercialisation. It takes time, money and expertise to get medical devices approved. And consumer applications will take off only if they perform a function people find useful. Some of the applications for brain-computer interfaces are unnecessary—a good voice-assistant is a simpler way to type without fingers than a brain implant, for example. The idea of consumers clamouring for craniotomies also seems far-fetched. Yet brain implants are already an established treatment for some conditions. Around 150,000 people receive deep-brain stimulation via electrodes to help them control Parkinson's disease. Elective surgery can become routine, as laser-eye procedures show.

All of which suggests that a route to the future imagined by the neurotech pioneers is arduous but achievable. When human ingenuity is applied to a problem, however hard, it is unwise to bet against it. Within a few years, improved technologies may be opening up new channels of communications with the brain. Many of the first applications hold out unambiguous promise—of movement and senses restored. But as uses move to the augmentation of abilities, whether for military purposes or among consumers, a host of concerns will arise. Privacy is an obvious one: the refuge of an inner voice may disappear. Security is another: if a brain can be reached on the internet, it can also be hacked. Inequality is a third: access to superhuman cognitive abilities could be beyond all except a self-perpetuating elite. Ethicists are already starting to grapple with questions of identity and agency that arise when a machine is in the neural loop.

These questions are not urgent. But the bigger story is that neither are they the realm of pure fantasy. Technology changes the way people live. Beneath the skull

lies the next frontier.

*This article appeared in the Leaders section of the print edition under the headline "The next frontier"*

## You've seen the news, now discover the story

Get incisive analysis on the issues that matter. Whether you read each issue cover to cover, listen to the audio edition, or scan the headlines on your phone, time with *The Economist* is always well spent.

Enjoy 12 weeks' access for \$12

+ receive a free notebook



## Thought experiments

# How brains and machines can be made to work together

*Brain-computer interfaces sound like the stuff of science fiction. Andrew Palmer sorts the reality from the hype*



Print edition | Technology Quarterly

Jan 4th 2018

IN THE gleaming facilities of the Wyss Centre for Bio and Neuroengineering in Geneva, a lab technician takes a well plate out of an incubator. Each well contains a tiny piece of brain tissue derived from human stem cells and sitting on top of an array of electrodes. A screen displays what the electrodes are picking up: the characteristic peak-and-trough wave forms of firing neurons.

To see these signals emanating from disembodied tissue is weird. The firing of a neuron is the basic building block of intelligence. Aggregated and combined, such

“action potentials” retrieve every memory, guide every movement and marshal every thought. As you read this sentence, neurons are firing all over your brain: to make sense of the shapes of the letters on the page; to turn those shapes into phonemes and those phonemes into words; and to confer meaning on those words.

**Latest updates**

Will Donald Trump's pro-coal policies put miners' lives at risk?  
GRAPHIC DETAIL

---

Two security flaws in modern chips cause big headaches for the tech business  
SCIENCE AND TECHNOLOGY

---

Are America's airports the worst in the world?  
GULLIVER

---

See all updates

This symphony of signals is bewilderingly complex. There are as many as 85bn neurons in an adult human brain, and a typical neuron has 10,000 connections to other such cells. The job of mapping these connections is still in its early stages. But as the brain gives up its secrets, remarkable possibilities have opened up: of decoding neural activity and using that code to control external devices.

A channel of communication of this sort requires a brain-computer interface (BCI). Such things are already in use. Since 2004, 13 paralysed people have been implanted with a system called BrainGate, first developed at Brown University (a handful of others have been given a similar device). An array of small electrodes, called a Utah array, is implanted into the motor cortex, a strip of the brain that governs movement. These electrodes detect the neurons that fire when someone intends to move his hands and arms. These signals are sent through wires that poke out of the person's skull to a decoder, where they are translated into a variety of outputs, from moving a cursor to controlling a limb.

The system has allowed a woman paralysed by a stroke to use a robotic arm to take her first sip of coffee without help from a caregiver. It has also been used by a paralysed person to type at a rate of eight words a minute. It has even reanimated useless human limbs. In a study led by Bob Kirsch of Case Western Reserve University, published in the *Lancet* this year, BrainGate was deployed artificially to stimulate muscles in the arms of William Kochevar, who was paralysed in a cycling accident. As a result, he was able to feed himself for the first time in eight years.

Interactions between brains and machines have changed lives in other ways, too. The opening ceremony of the football World Cup in Brazil in 2014 featured a



paraplegic man who used a mind-controlled robotic exoskeleton to kick a ball. A recent study by Ujwal Chaudhary of the University of Tübingen and four co-authors relied on a technique called functional near-infrared spectroscopy (fNIRS), which beams infrared light into the brain, to put yes/no questions to four locked-in patients who had been completely immobilised by Lou Gehrig's disease; the patients' mental responses showed up as identifiable patterns of blood oxygenation.

Neural activity can be stimulated as well as recorded. Cochlear implants convert sound into electrical signals and send them into the brain. Deep-brain stimulation uses electrical pulses, delivered via implanted electrodes, to help control Parkinson's disease. The technique has also been used to treat other movement disorders and mental-health conditions. NeuroPace, a Silicon Valley firm, monitors brain activity for signs of imminent epileptic seizures and delivers electrical stimulation to stop them.

It is easy to see how brain-computer interfaces could be applied to other sensory inputs and outputs. Researchers at the University of California, Berkeley, have deconstructed electrical activity in the temporal lobe when someone is listening to conversation; these patterns can be used to predict what word someone has heard. The brain also produces similar signals when someone imagines hearing spoken words, which may open the door to a speech-processing device for people with conditions such as aphasia (the inability to understand or produce speech).

Researchers at the same university have used changes in blood oxygenation in the brain to reconstruct, fuzzily, film clips that people were watching. Now imagine a device that could work the other way, stimulating the visual cortex of blind people in order to project images into their mind's eye.

If the possibilities of BCIs are enormous, however, so are the problems. The most advanced science is being conducted in animals. Tiny silicon probes called Neuropixels have been developed by researchers at the Howard Hughes Institute, the Allen Institute and University College London to monitor cellular-level activity in multiple brain regions in mice and rats. Scientists at the University of California, San Diego, have built a BCI that can predict from prior neural activity what song a zebra finch will sing. Researchers at the California Institute of Technology have


worked out how cells in the visual cortex of macaque monkeys encoded 50 different aspects of a person’s face, from skin colour to eye spacing. That enabled them to predict the appearance of faces that monkeys were shown from the brain signals they detected, with a spooky degree of accuracy. But conducting scientific research on human brains is harder, for regulatory reasons and because they are larger and more complex.

Even when BCI breakthroughs are made on humans in the lab, they are difficult to translate into clinical practice. *Wired* magazine first reported breathlessly on the then new BrainGate system back in 2005. An early attempt to commercialise the technology, by a company called Cyberkinetics, founded. It took NeuroPace 20 years to develop its technologies and negotiate regulatory approval, and it expects that only 500 people will have its electrodes implanted this year.

Current BCI technologies often require experts to operate them. “It is not much use if you have to have someone with a masters in neural engineering standing next to the patient,” says Leigh Hochberg, a neurologist and professor at Brown University, who is one of the key figures behind BrainGate. Whenever wires pass through the skull and scalp, there is an infection risk. Implants also tend to move slightly within the brain, which can harm the cells it is recording from; and the brain’s immune response to foreign bodies can create scarring around electrodes, making them less effective.

Moreover, existing implants record only a tiny selection of the brain’s signals. The Utah arrays used by the BrainGate consortium, for example, might pick up the firing of just a couple of hundred neurons out of that 85bn total. In a paper published in 2011, Ian Stevenson and Konrad Kording of Northwestern University showed that the number of simultaneously recorded neurons had doubled every seven years since the 1950s (see chart). This falls far short of Moore’s law, which has seen computing power double every two years.

Indeed, the Wyss Centre in Geneva exists because it is so hard to get neurotechnology out of the lab and into clinical practice. John Donoghue, who heads the centre, is another of the pioneers of the BrainGate system. He says it is designed to help



promising ideas cross several “valleys of death”. One is financial: the combination of lengthy payback periods and deep technology scares off most investors. Another is the need for multidisciplinary expertise to get better interfaces built and management skills to keep complex projects on track. Yet another is the state of neuroscience itself. “At its core, this is based on understanding how the brain works, and we just don’t,” says Dr Donoghue.

### **Me, myself and AI**

This odd mixture of extraordinary achievement and halting progress now has a new ingredient: Silicon Valley. In October 2016 Bryan Johnson, an entrepreneur who had made a fortune by selling his payments company, Braintree, announced an investment of \$100m in Kernel, a firm he has founded to “read and write neural code”. Mr Johnson reckons that the rise of artificial intelligence (AI) will demand a concomitant upgrade in human capabilities. “I find it hard to imagine a world by 2050 where we have not intervened to improve ourselves,” he says, picturing an ability to acquire new skills at will or to communicate telepathically with others. Last February Kernel snapped up Kendall Research Systems, a spinoff from the Massachusetts Institute of Technology (MIT) that works on neural interfaces.

Kernel is not alone in seeing BCIs as a way for humans to co-exist with AI rather than be subjugated to it. In 2016 Elon Musk, the boss of SpaceX and Tesla, founded a new company called Neuralink, which is also working to create new forms of implants. He has gathered together an impressive group of co-founders and set a goal of developing a BCI for clinical use in people with disabilities by 2021. Devices for people without such disabilities are about eight to ten years away, by Mr Musk’s reckoning.

Neuralink is not saying what exactly it is doing, but Mr Musk’s thinking is outlined in a lengthy post on *Wait But Why*, a website. In it, he describes the need for humans to communicate far more quickly with each other, and with computers, if they are not to be left in the dust by AI. The post raises some extraordinary possibilities: being able to access and absorb knowledge instantly from the cloud or to pump images from one person’s retina straight into the visual cortex of another; creating entirely new sensory abilities, from infrared eyesight to high-frequency hearing; and ultimately, melding together human and artificial intelligence.

In April it was Facebook's turn to boggle minds as it revealed plans to create a "silent speech" interface that would allow people to type at 100 words a minute straight from their brain. A group of more than 60 researchers, some inside Facebook and some outside, are working on the project. A separate startup, Openwater, is also working on a non-invasive neural-imaging system; its founder, Mary Lou Jepsen, says that her technology will eventually allow minds to be read.

Many BCI experts react to the arrival of the Valley visionaries by rolling their eyes. Neuroscience is a work in progress, they say. An effective BCI requires the involvement of many disciplines: materials science, neuroscience, machine learning, engineering, design and others. There are no shortcuts to clinical trials and regulatory approval.

In all this, the sceptics are right. Many of the ambitions being aired look fantastical. Still, this is a critical moment for BCIs. Vast amounts of money are pouring into the field. Researchers are trying multiple approaches. Mr Musk in particular has a track record of combining grandiose aspirations (colonising Mars) and practical success (recovering and relaunching rockets via SpaceX).

To be clear, "The Matrix" is not imminent. But BCIs may be about to take a big leap forward. For that to happen, the most important thing is to find a better way of connecting with the brain.

*This article appeared in the Technology Quarterly section of the print edition under the headline "Thought experiments"*

## Headache

# Reading the brain from the outside

*Can brain activity be deciphered without opening up the skull?*



Print edition | Technology Quarterly

Jan 4th 2018

PATRICK KAIFOSH'S left hand lies flat on the table in front of him. Occasionally his fingers twitch or his palm rises up slightly from the surface. There is nothing obvious to connect these movements with what is happening on the tablet in front of him, where a game of asteroids is being played. Yet he is controlling the spaceship on the screen as it spins, thrusts and fires.

What enables him to do so is a sweatband studded with small gold bars that sits halfway up his left forearm. Each bar contains a handful of electrodes designed to pick up the signals of motor units (the combination of a motor neuron, a cell that

projects from the spinal cord, and the muscle fibres it controls). These data are processed by machine-learning algorithms and translated into the actions in the game. Dr Kaifosh, a co-founder of CTRL-Labs, the startup behind the device, has learned to exercise impressive control over these signals with hardly any obvious movement.

**Latest updates**

Will Donald Trump's pro-coal policies put miners' lives at risk?  
GRAPHIC DETAIL

---

Two security flaws in modern chips cause big headaches for the tech business  
SCIENCE AND TECHNOLOGY

---

Are America's airports the worst in the world?  
GULLIVER

---

See all updates

Some say that the claims of Dr Kaifosh and Thomas Reardon, his co-founder, that CTRL-Labs has created a brain-machine interface are nonsense. The sweatband is nowhere near the brain, and the signals it is picking up are generated not just by the firing of a motor neuron but by the electrical activity of muscles. "If this is a BCI, then the movement of my fingers when I type on a keyboard is also a brain

output," sniffs one researcher. Krishna Shenoy, who directs the neural prosthetics systems lab at Stanford University and acts as an adviser to the firm, thinks it is on the right side of the divide. "Measuring the movement of the hand is motion capture. They are picking up neural activity amplified by the muscles."

Whatever the semantics, it is instructive to hear the logic behind the firm's decision to record the activity of the peripheral nervous system, rather than looking directly inside the head. The startup wants to create a consumer product (its potential uses include being an interface for interactions in virtual reality and augmented reality). It is not reasonable to expect consumers to undergo brain surgery, say the founders, and current non-invasive options for reading the brain provide noisy, hard-to-read signals. "For machine-learning folk, there is no question which data set—cortical neurons or motor neurons—you would prefer," says Dr Reardon.

This trade-off between the degree of invasiveness and the fidelity of brain signals is a big problem in the search for improved BCIs. But plenty of people are trying to find a better way to read neural code from outside the skull.

The simplest way to read electrical activity from outside is to conduct an electroencephalogram (EEG). And it is not all that simple. Conventionally, it has involved wearing a cap containing lots of electrodes that are pressed against the surface of the scalp. To improve the signal quality, a conductive gel is often applied. That requires a hairwash afterwards. Sometimes the skin of the scalp is roughened up to get a better connection. As a consumer experience it beats going to the dentist, but not by much.

Once on, each electrode picks up currents generated by the firing of thousands of neurons, but only in the area covered by that electrode. Neurons that fire deep in the brain are not detected either. The signal is distorted by the layers of skin, bone and membrane that separate the brain from the electrode. And muscle activity (of the sort that CTRL-Labs looks for) from eye and neck movements or clenched jaws can overwhelm the neural data.

Even so, some EEG signals are strong enough to be picked up pretty reliably. An “event-related potential”, for example, is an electrical signal that the brain reliably gives off in response to an external stimulus of some sort. One such, called an error-related potential (Errp), occurs when a user spots a mistake. Researchers at MIT have connected a human observer wearing an EEG cap to an industrial robot called Baxter as it carried out a sorting task. If Baxter made a mistake, an Errp signal in the observer’s brain alerted the robot to its error; helpfully, if Baxter still did not react, the human brain generated an even stronger Errp signal.

### **If the cap fits**

Neurable, a consumer startup, has developed an EEG headset with just seven dry electrodes which uses a signal called the P300 to enable users to play a virtual-reality (VR) escape game. This signal is a marker of surprise or recognition. Think of the word “brain” and then watch a series of letters flash up randomly on a screen; when the letter “b” comes up, you will almost certainly be giving off a P300 signal. In Neurable’s game, all you have to do is concentrate on an object (a ball, say) for it to come towards you or be hurled at an object. Ramses Alcaide, Neurable’s boss, sees the potential for entertainment companies like Disney (owner of the Star Wars and Marvel franchises) to license the software in theme parks and arcade games.

Thorsten Zander of the Technische Universität in Berlin thinks that “passive” EEG signals (those that are not evoked by an external stimulus) can be put to good use too. Research has shown that brainwave activity changes depending on how alert, drowsy or focused a person is. If an EEG can reliably pick this up, perhaps surgeons, pilots or truck drivers who are becoming dangerously tired can be identified. Studies have shown strong correlations between people’s mental states as shown by an EEG and their ability to spot weapons in X-rays of luggage.

Yet the uses of EEGs remain limited. In a real-world environment like a cockpit, a car or an airport, muscle activity and ambient electricity are likely to confound any neural signals. As for Neurable’s game, it relies not solely on brain activity but also deploys eye-tracking technology to see where a player is looking. Dr Alcaide says the system can work with brain signals alone, but it is hard for a user to disentangle the two.

Other non-invasive options also have flaws. Magnetoencephalography measures magnetic fields generated by electrical activity in the brain, but it requires a special room to shield the machinery from Earth’s magnetic field. Functional magnetic resonance imaging (fMRI) can spot changes in blood oxygenation, a proxy for neural activity, and can zero in on a small area of the brain. But it involves a large, expensive machine, and there is a lag between neural activity and blood flow.

If any area is likely to yield a big breakthrough in non-invasive recording of the brain, it is a variation on fNIRS, the infrared technique used in the experiment to allow locked-in patients to communicate. In essence, light sent through the skull is either absorbed or reflected back to detectors, providing a picture of what is going on in the brain. This technique does not require bulky equipment, and unlike EEG it does not measure electrical activity, so it is not confused by muscle activity. Both Facebook and Openwater are focusing their efforts on this area.

The obstacles to a breakthrough are formidable, however. Current infrared techniques measure an epiphenomenon, blood oxygenation (the degree of which affects the absorption of light), rather than the actual firing of neurons. The light usually penetrates only a few millimetres into the cortex. And because light scatters in tissue (think of how your whole fingertip glows red when you press a pen-torch against it), the precise source of reflected signals is hard to identify.



Facebook is not saying much about what it is doing. Its efforts are being led by Mark Chevillet, who joined the social-media giant's Building 8 consumer-hardware team from Johns Hopkins University. To cope with the problem of light scattering as it passes through the brain, the team hopes to be able to pick up on both ballistic photons, which pass through tissue in a straight line, and what it terms "quasi-ballistic photons", which deviate slightly but can still be traced to a specific source. The clock is ticking. Dr Chevillet has about a year of a two-year programme left to demonstrate that the firm's goal of brain-controlled typing at 100 words a minute is achievable using current invasive cell-recording techniques, and to produce a road map for replicating that level of performance non-invasively.

Openwater is much less tight-lipped. Ms Jepsen says that her San Francisco-based startup uses holography to reconstruct how light scatters in the body, so it can neutralise this effect. Openwater, she suggests, has already created technology that has a billion times the resolution of an fMRI machine, can penetrate the cortex to a depth of 10cm, and can sample data in milliseconds.

Openwater has yet to demonstrate its technology, so these claims are impossible to verify. Most BCI experts are sceptical. But Ms Jepsen has an impressive background in consumer electronics and display technologies, and breakthroughs by their nature upend conventional wisdom. Developer kits are due out in 2018.

In the meantime, other efforts to decipher the language of the brain are under way. Some involve heading downstream into the peripheral nervous system. One example of that approach is CTRL-Labs; another is provided by Qi Wang, at Columbia University, who researches the role of the *locus coeruleus*, a nucleus deep in the brain stem that plays a role in modulating anxiety and stress. Dr Wang is looking at ways of stimulating the vagus nerve, which runs from the brain into the abdomen, through the skin to see if he can affect the *locus coeruleus*.

Others are looking at invasive approaches that do not involve drilling through the skull. One idea, from a firm called SmartStent, using technology partly developed with the University of Melbourne, is to use a stent-like device called a "stentrode" that is studded with electrodes. It is inserted via a small incision in the neck and then guided up through blood vessels to overlie the brain. Once the device is in the right location, it expands from about the size of a matchstick to the size of the

vessel and tissue grows into its scaffolding, keeping it in place. Human trials of the stentrodre are due to start next year.

Another approach is to put electrodes under the scalp but not under the skull. Maxime Baud, a neurologist attached to the Wyss Centre, wants to do just that in order to monitor the long-term seizure patterns of epileptics. He hopes that once these patterns are revealed, they can be used to provide accurate forecasts of when a seizure is likely to occur.

Yet others think they need to go directly to the source of action potentials. And that means heading inside the brain itself.

*This article appeared in the Technology Quarterly section of the print edition under the headline "Headache"*

## Inside intelligence

# The best way of looking at the brain is from within

*The hunt for smaller, safer and smarter brain implants*



---

Print edition | Technology Quarterly

Jan 6th 2018

---

TALK to neuroscientists about brain-computer interfaces (BCIs) for long enough, and the stadium analogy is almost bound to come up. This compares the neural activity of the brain to the noise made by a crowd at a football game. From outside the ground, you might hear background noise and be able to tell from the roars whether a team has scored. In a blimp above the stadium you can tell who has scored and perhaps which players were involved. Only inside it can you ask the fan in row 72 how things unfolded in detail.

Similarly, with the brain it is only by getting closer to the action that you can really understand what is going on. To get high-resolution signals, for now there is no alternative to opening up the skull. One option is to place electrodes onto the surface of the brain in what is known as electrocorticography. Another is to push them right into the tissue of the brain, for example by using a grid of microelectrodes like BrainGate's Utah array.

**Latest updates**

Will Donald Trump's pro-coal policies put miners' lives at risk?  
GRAPHIC DETAIL

---

Two security flaws in modern chips cause big headaches for the tech business  
SCIENCE AND TECHNOLOGY

---

Are America's airports the worst in the world?  
GULLIVER

---

See all updates

Just how close you have to come to individual neurons to operate BCIs is a matter of debate. In people who suffer from movement disorders such as Parkinson's disease, spaghetti-like leads and big electrodes are used to carry out deep-brain stimulation over a fairly large area of tissue. Such treatment is generally regarded as effective. Andrew Jackson of the University of Newcastle thinks that

recording activity by ensembles of neurons, of the sort that gets picked up by electrocorticography arrays, can be used to decode relatively simple movement signals, like an intention to grasp something or to extend the elbow.

But to generate fine-grained control signals, such as the movement of individual fingers, more precision is needed. "These are very small signals, and there are many neurons packed closely together, all firing together," says Andrew Schwartz of the University of Pittsburgh. Aggregating them inevitably means sacrificing detail. After all, individual cells can have very specific functions, from navigation to facial recognition. The 2014 Nobel prize for medicine was awarded for work on place and grid cells, which fire when animals reach a specific location; the idea of the "Jennifer Aniston neuron" stems from research showing that single neurons can fire in response to pictures of a specific celebrity.

Companies like Neuralink and Kernel are betting that the most ambitious visions of BCIs, in which thoughts, images and movements are seamlessly encoded and decoded, will require high-resolution implants. So, too, is America's Defense Advanced Research Projects Agency (DARPA), an arm of the Pentagon, which this year distributed \$65m among six organisations to create a high-resolution

implantable interface. BrainGate and others continue to work on systems of their own.

But the challenges that these researchers face are truly daunting. The ideal implant would be safe, small, wireless and long-lasting. It would be capable of transmitting huge amounts of data at high speed. It would interact with many more neurons than current technology allows (the DARPA programme sets its grant recipients a target of 1m neurons, along with a deadline of 2021 for a pilot trial to get under way in humans). It would also have to navigate an environment that Claude Clément of the Wyss Centre likens to a jungle by the sea: humid, hot and salty. “The brain is not the right place to do technology,” he says. As the chief technology officer, he should know.

### **Da neuron, ron, ron**

That is not stopping people from trying. The efforts now being made to create better implants can be divided into two broad categories. The first reimagines the current technology of small wire electrodes. The second heads off in new, non-electrical directions.

Start with ways to make electrodes smaller and better. Ken Shepard is a professor of electrical and biomedical engineering at Columbia University; his lab is a recipient of DARPA funds, and is aiming to build a device that could eventually help blind people with an intact visual cortex to see by stimulating precisely the right neurons in order to produce images inside their brains. He thinks he can do so by using state-of-the-art CMOS (complementary metal-oxide semiconductor) electronics.

Dr Shepard is aware that any kind of penetrating electrode can cause cell damage, so he wants to build “the mother of all surface recording devices” which will sit on top of the cortex and under the membranes that surround the brain. He has already created a prototype of a first-generation CMOS chip, which measures about 1cm by 1cm and contains 65,000 electrodes; a slightly larger, second-generation version will house 1m sensors. But like everyone else trying to make implants work, Dr Shepard is not just cramming sensors onto the chip. He also has to add the same number of amplifiers, a converter to turn the analogue signals of action potentials into the digital 0s and 1s of machine learning, and a wireless link to send (or

receive) data to a relay station that will sit on the scalp. That, in turn, will send (or receive) the data wirelessly to external processors for decoding.

The device also has to be powered, another huge part of the implantables puzzle. No one in the field puts faith in batteries as a source of power. They are too bulky, and the risk of battery fluid leaking into the brain is too high. Like many of his peers, Dr Shepard uses inductive coupling, whereby currents passing through a coiled wire create a magnetic field that can induce a current in a second coil (the way that an electric toothbrush gets recharged). That job is done by coils on the chip and on the relay station.




Photo: The Leap

Over on America's west coast, a startup called Paradromics is also using inductive coupling to power its implantable. But Matt Angle, its boss, does not think that souped-up surface recordings will deliver sufficiently high resolution. Instead, he is working on creating tiny bundles of glass and metal microwires that can be pushed into brain tissue, a bit like a Utah array but with many more sensors. To stop the wires clumping together, thereby reducing the number of neurons they engage with, the firm uses a sacrificial polymer to splay them apart; the polymer dissolves but the wires remain separated. They are then bonded onto a high-speed CMOS circuit. A version of the device, with 65,000 electrodes, will be released next year for use in animal research.

That still leaves lots to do before Paradromics can meet its DARPA-funded goal of creating a 1m wire device that can be used in people. Chief among them is coping with the amount of data coming out of the head. Dr Angle reckons that the initial device produces 24 gigabits of data every second (streaming an ultra-high-definition movie on Netflix uses up to 7GB an hour). In animals, these data can be transmitted through a cable to a bulky aluminium head-mounted processor. That is a hard look to pull off in humans; besides, such quantities of data would generate far too much heat to be handled inside the skull or transmitted wirelessly out of it.

So Paradromics, along with everyone else trying to create a high-bandwidth signal into and out of the brain, has to find a way to compress the data rate without compromising the speed and quality of information sent. Dr Angle reckons he can

do this in two ways: first, by ignoring the moments of silence in between action potentials, rather than laboriously encoding them as a string of zeros; and second, by concentrating on the wave forms of specific action potentials rather than recording each point along their curves. Indeed, he sees data compression as being the company's big selling-point, and expects others that want to create specific BCI applications or prostheses simply to plug into its feed. "We see ourselves as the neural data backbone, like a Qualcomm or Intel," he says.

### **Meshy business**

Some researchers are trying to get away from the idea of wire implants altogether. At Brown University, for example, Arto Nurmikko is leading a multidisciplinary team to create "neurograins", each the size of a grain of sugar, that could be sprinkled on top of the cortex or implanted within it. Each grain would have to have built-in amplifiers, analogue-to-digital converters and the ability to send data to a relay station which could power the grains inductively and pass the information to an external processor. Dr Nurmikko is testing elements of the system in rodents; he hopes eventually to put tens of thousands of grains inside the head.

Meanwhile, in a lab at Harvard University, Guosong Hong is demonstrating another innovative interface. He dips a syringe into a beaker of water and injects into it a small, billowing and glinting mesh. It is strangely beautiful to watch. Dr Hong is a postdoctoral fellow in the lab of Charles Lieber, a professor of chemistry; they are both working to create a neural interface that blurs the distinction between biology and electronics. Their solution is a porous net made of a flexible polymer called SU-8, studded with sensors and conductive metal.

The mesh is designed to solve a number of problems. One has to do with the brain's immune response to foreign bodies. By replicating the flexibility and softness of neural tissue, and allowing neurons and other types of cells to grow within it, it should avoid the scarring that stiffer, solid probes can sometimes cause. It also takes up much less space: less than 1% of the volume of a Utah array. Animal trials have gone well; the next stage will be to insert the mesh into the brains of epilepsy patients who have not responded to other forms of treatment and are waiting to have bits of tissue removed.

A mile away, at MIT, members of Polina Anikeeva's lab are also trying to build devices that match the physical properties of neural tissue. Dr Anikeeva is a materials scientist who first dived into neuroscience at the lab of Karl Deisseroth at Stanford University, who pioneered the use of optogenetics, a way of genetically engineering cells so that they turn on and off in response to light. Her reaction upon seeing a (mouse) brain up close for the first time was amazement at how squishy it was. "It is problematic to have something with the elastic properties of a knife inside something with the elastic properties of a chocolate pudding," she says.

One way she is dealing with that is to borrow from the world of telecoms by creating a multichannel fibre with a width of 100 microns (one micron is a millionth of a metre), about the same as a human hair. That is denser than some of the devices being worked on elsewhere, but the main thing that distinguishes it is that it can do multiple things. "Electronics with just current and voltage is not going to do the trick," she says, pointing out that the brain communicates not just electrically but chemically, too.

Dr Anikeeva's sensor has one channel for recording using electrodes, but it is also able to take advantage of optogenetics. A second channel is designed to deliver channelrhodopsin, an algal protein that can be smuggled into neurons to make them sensitive to light, and a third to shine a light so that these modified neurons can be activated.

It is too early to know if optogenetics can be used safely in humans: channelrhodopsin has to be incorporated into cells using a virus, and there are question-marks about how much light can safely be shone into the brain. But human clinical trials are under way to make retinal ganglion cells light-sensitive in people whose photoreceptor cells are damaged; another of the recipients of DARPA funds, Fondation Voir et Entendre in Paris, aims to use the technique to transfer images from special goggles directly into the visual cortex of completely blind people. In principle, other senses could also be restored: optogenetic stimulation of cells in the inner ear of mice has been used to control hearing.

Dr Anikeeva is also toying with another way of stimulating the brain. She thinks that a weak magnetic field could be used to penetrate deep into neural tissue and



heat up magnetic nanoparticles that have been injected into the brain. If heat-sensitive capsaicin receptors were triggered in modified neurons nearby, the increased temperature would cause the neurons to fire.

Another candidate for recording and activating neurons, beyond voltage, light and magnets, is ultrasound. Jose Carmena and Michel Maharbiz at the University of California, Berkeley, are the main proponents of this approach, which again involves the insertion of tiny particles (which they call “neural dust”) into tissue. Passing ultrasound through the body affects a crystal in these motes which vibrates like a tuning fork; that produces voltage to power a transistor. Electrical activity in adjacent tissue, whether muscles or neurons, can change the nature of the ultrasonic echo given off by the particle, so this activity can be recorded.

Many of these new efforts raise even more questions. If the ambition is to create a “whole-brain interface” that covers multiple regions of the brain, there must be a physical limit to how much additional material, be it wires, grains or motes, can be introduced into a human brain. If such particles can be made sufficiently small to mitigate that problem, another uncertainty arises: would they float around in the brain, and with what effects? And how can large numbers of implants be put into different parts of the brain in a single procedure, particularly if the use of tiny, flexible materials creates a “wet noodle” problem whereby implants are too floppy to make their way into tissue? (Rumour has it that Neuralink may be pursuing the idea of an automated “sewing machine” designed to get around this issue.)

All this underlines how hard it will be to engineer a new neural interface that works both safely and well. But the range of efforts to create such a device also prompts optimism. “We are approaching an inflection-point that will enable at-scale recording and stimulation,” says Andreas Schaefer, a neuroscientist at the Crick Institute in London.

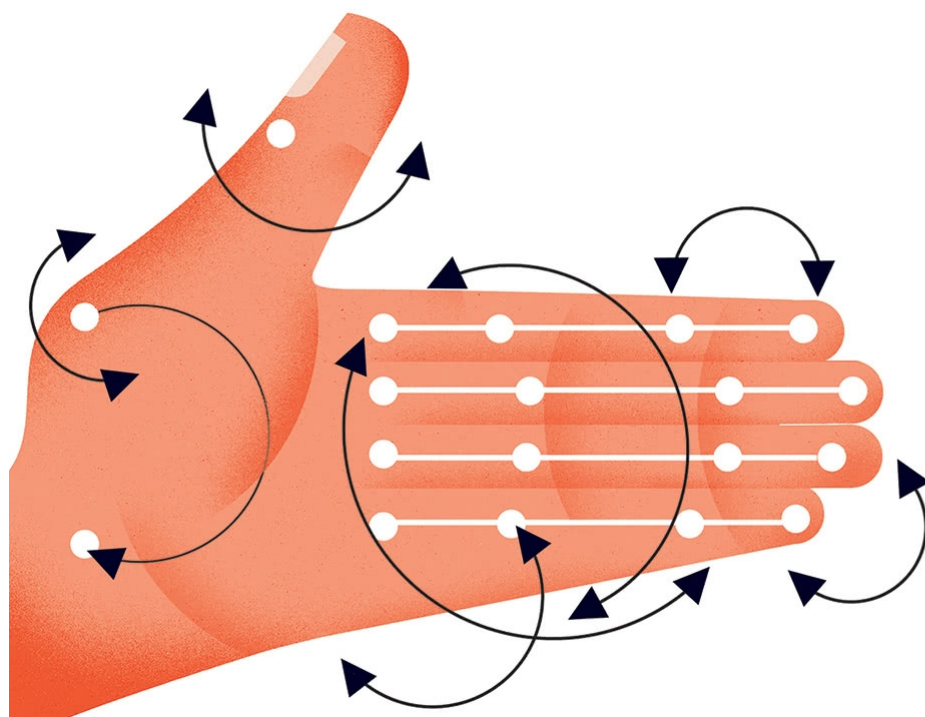
Even so, being able to get the data out of the brain, or into it, is only the first step. The next thing is processing them.

*This article appeared in the Technology Quarterly section of the print edition under the headline "Inside intelligence"*

Translation required

## Turning brain signals into useful information

*Once data have been extracted from the brain, how can they be employed to best effect?*



Nathalie Lees

---

Print edition | Technology Quarterly

Jan 4th 2018

---

FOR those who reckon that brain-computer interfaces will never catch on, there is a simple answer: they already have. Well over 300,000 people worldwide have had cochlear implants fitted in their ears. Strictly speaking, this hearing device does not interact directly with neural tissue, but the effect is not dissimilar. A processor captures sound, which is converted into electrical signals and sent to an electrode in the inner ear, stimulating the cochlear nerve so that sound is heard in the brain. Michael Merzenich, a neuroscientist who helped develop them, explains that the implants provide only a crude representation of speech, “like playing Chopin with your fist”. But given a little time, the brain works out the signals.

That offers a clue to another part of the BCI equation: what to do once you have gained access to the brain. As cochlear implants show, one option is to let the world's most powerful learning machine do its stuff. In a famous mid-20th-century experiment, two Austrian researchers showed that the brain could quickly adapt to a pair of glasses that turned the image they projected onto the retina upside down. More recently, researchers at Colorado State University have come up with a device that converts sounds into electrical impulses. When pressed against the tongue, it produces different kinds of tingle which the brain learns to associate with specific sounds.

ADVERTISING



inRead invented by Teads

### Latest updates

Will Donald Trump's pro-coal policies put miners' lives at risk?

GRAPHIC DETAIL

Two security flaws in modern chips cause big headaches for the tech business

SCIENCE AND TECHNOLOGY

Are America's airports the worst in the world?

GULLIVER

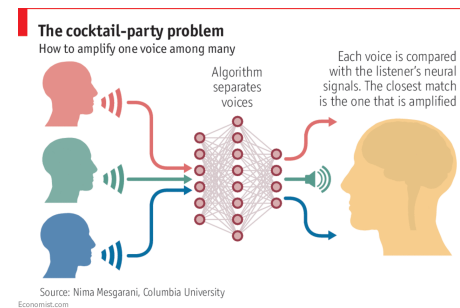
[See all updates](#)

The brain, then, is remarkably good at working things out. Then again, so are computers. One problem with a hearing aid, for example, is that it amplifies every sound that is coming in; when you want to focus on one person in a noisy environment, such as a party, that is not much help. Nima Mesgarani of Columbia University is working on a way to separate out the specific person you want to listen

to. The idea is that an algorithm will distinguish between different voices talking at the same time, creating a spectrogram, or visual representation of sound

frequencies, of each person's speech. It then looks at neural activity in the brain as the wearer of the hearing aid concentrates on a specific interlocutor. This activity can also be reconstructed into a spectrogram, and the ones that match up will get amplified (see diagram).

Algorithms have done better than brain plasticity at enabling paralysed people to send a cursor to a target using thought alone. In research published earlier this year, for example, Dr Shenoy and his collaborators at Stanford University recorded a big improvement in brain-controlled typing. This stemmed not from new signals or whizzier interfaces but from better maths.



One contribution came from Dr Shenoy's use of data generated during the testing phase of his algorithm. In the training phase a user is repeatedly told to move a cursor to a particular target; machine-learning programs identify patterns in neural activity that correlate with this movement. In the testing phase the user is shown a grid of letters and told to move the cursor wherever he wants; that tests the algorithm's ability to predict the user's wishes. The user's intention to hit a specific target also shows up in the data; by refitting the algorithm to include that information too, the cursor can be made to move to its target more quickly.

But although algorithms are getting better, there is still a lot of room for improvement, not least because data remain thin on the ground. Despite claims that smart algorithms can make up for bad signals, they can do only so much. "Machine learning does nearly magical things, but it cannot do magic," says Dr Shenoy. Consider the use of functional near-infrared spectroscopy to identify simple yes/no answers given by locked-in patients to true-or-false statements; they were right 70% of the time, a huge advance on not being able to communicate at all, but nowhere near enough to have confidence in their responses to an end-of-life discussion, say. More and cleaner data are required to build better algorithms.

It does not help that knowledge of how the brain works is still so incomplete. Even with better interfaces, the organ's extraordinary complexities will not be quickly unravelled. The movement of a cursor has two degrees of freedom, for example; a

human hand has 27. Visual-cortex researchers often work with static images, whereas humans in real life have to cope with continuously moving images. Work on the sensory feedback that humans experience when they grip an object has barely begun.

And although computational neuroscientists can piggyback on broader advances in the field of machine learning, from facial recognition to autonomous cars, the noisiness of neural data presents a particular challenge. A neuron in the motor cortex may fire at a rate of 100 action potentials a second when someone thinks about moving his right arm on one occasion, but at a rate of 115 on another. To make matters worse, neurons' jobs overlap. So if a neuron has an average firing rate of 100 to the right and 70 to the left, what does a rate of 85 signify?

At least the activities of the motor cortex have a visible output in the form of movement, showing up correlations with neural data from which predictions can be made. But other cognitive processes lack obvious outputs. Take the area that Facebook is interested in: silent, or imagined, speech. It is not certain that the brain's representation of imagined speech is similar enough to actual (spoken or heard) speech to be used as a reference point. Progress is hampered by another factor: "We have a century's worth of data on how movement is generated by neural activity," says BrainGate's Dr Hochberg dryly. "We know less about animal speech."

Higher-level functions, such as decision-making, present an even greater challenge. BCI algorithms require a model that explicitly defines the relationship between neural activity and the parameter in question. "The problem begins with defining the parameter itself," says Dr Schwartz of Pittsburgh University. "Exactly what is cognition? How do you write an equation for it?"

Such difficulties suggest two things. One is that a set of algorithms for whole-brain activity is a very long way off. Another is that the best route forward for signal processing in a brain-computer interface is likely to be some combination of machine learning and brain plasticity. The trick will be to develop a system in which the two co-operate, not just for the sake of efficiency but also for reasons of ethics.

*This article appeared in the Technology Quarterly section of the print edition under the headline "Translation required"*



**Brain scan****Eberhard Fetz**

*The man who helped launch brain-computer interfaces in 1969 has not finished yet*



---

Print edition | Technology Quarterly

Jan 4th 2018

---

BCIs have deep roots. In the 18th century Luigi Galvani discovered the role of electricity in nerve activity when he found that applying voltage could cause a dead frog's legs to twitch. In the 1920s Hans Berger used electroencephalography to record human brain waves. In the 1960s José Delgado theatrically used a brain implant to stop a charging bull in its tracks. One of the field's father figures is still hard at work in the lab.

Eberhard Fetz was a post-doctoral researcher at the University of Washington in Seattle when he decided to test whether a monkey could control the needle of a

meter using only its mind. A paper based on that research, published in 1969, showed that it could. Dr Fetz tracked down the movement of the needle to the firing rate of a single neuron in the monkey's brain. The animal learned to control the activity of that single cell within two minutes, and was also able to switch to control a different neuron.

**Latest updates**

Will Donald Trump's pro-coal policies put miners' lives at risk?  
GRAPHIC DETAIL

---

Two security flaws in modern chips cause big headaches for the tech business  
SCIENCE AND TECHNOLOGY

---

Are America's airports the worst in the world?  
GULLIVER

---

See all updates

Dr Fetz disclaims any great insights in setting up the experiment. "I was just curious, and did not make the association with potential uses of robotic arms or the like," he says. But the effect of his paper was profound. It showed both that volitional control of a BCI was possible, and that the brain was capable of learning how to operate one without any help.

Some 48 years later, Dr Fetz is still at the University of Washington, still fizzing with energy and still enthralled by the brain's plasticity. He is particularly interested in the possibility of artificially strengthening connections between cells, and perhaps forging entirely new ones.

As an example, he points to research in which the recording of an action potential in the brain prompts not only the normal firing of a motor neuron in the spinal cord but also a parallel stimulus from a BCI delivered to the same site. The idea is to take advantage of a relationship made famous in an aphorism by a Canadian psychologist, Donald Hebb: "Neurons that fire together, wire together." This reinforced stimulus strengthens the connection between the original action potential and the motor neuron, which could help recovery from spinal-cord injuries. Such stimulation might also encourage stronger bonds in the brain itself—between the speech-processing area of a stroke victim's brain, say, and the region that controls movements of the lips and mouth.

Asked to explain the slow rate of progress since his breakthrough paper, Dr Fetz points to the technical difficulties of recording from single cells and the high hurdle to doing this sort of work in people. But he does not think there is a need to



find out much more about the brain in order to make further advances: “I have had a lot of progress by just jumping in.”

Grey matter, red tape

## In search of serendipity

*How obstacles to workable brain-computer interfaces may be overcome*




Print edition | Technology Quarterly

Jan 4th 2018

NEUROTECHNOLOGY has long been a favourite of science-fiction writers. In “Neuromancer”, a wildly inventive book by William Gibson written in 1984, people can use neural implants to jack into the sensory experiences of others. The idea of a neural lace, a mesh that grows into the brain, was conceived by Iain M. Banks in his “Culture” series of novels. “The Terminal Man” by Michael Crichton, published in 1972, imagines the effects of a brain implant on someone who is convinced that machines are taking over from humans. (Spoiler: not good.)

Where the sci-fi genre led, philosophers are now starting to follow. In Howard Chizeck’s lab at the University of Washington, researchers are working on an implanted device to administer deep-brain stimulation (DBS) in order to treat a common movement disorder called essential tremor. Conventionally, DBS stimulation is always on, wasting energy and depriving the patient of a sense of control. The lab’s ethicist, Tim Brown, a doctoral student of philosophy, says that some DBS patients suffer a sense of alienation and complain of feeling like a robot.

ADVERTISING



inRead invented by Teads

### Latest updates

Will Donald Trump’s pro-coal policies put miners’ lives at risk?

GRAPHIC DETAIL

Two security flaws in modern chips cause big headaches for the tech business

SCIENCE AND TECHNOLOGY

Are America’s airports the worst in the world?

GULLIVER

[See all updates](#)

To change that, the team at the University of Washington is using neuronal activity associated with intentional movements as a trigger for turning the device on. But the researchers also want to enable patients to use a conscious thought process to override these settings. That is more useful than it might sound: stimulation currents for essential tremor can cause side-effects like distorted speech, so someone about to

give a presentation, say, might wish to shake rather than slur his words.

Giving humans more options of this sort will be essential if some of the bolder visions for brain-computer interfaces are to be realised. Hannah Maslen from the

University of Oxford is another ethicist who works on a BCI project, in this case a neural speech prosthesis being developed by a consortium of European researchers. One of her jobs is to think through the distinctions between inner speech and public speech: people need a dependable mechanism for separating out what they want to say from what they think.

That is only one of many ethical questions that the sci-fi versions of brain-computer interfaces bring up. What protection will BCIs offer against neural hacking? Who owns neural data, including information that is gathered for research purposes now but may be decipherable in detail at some point in the future? Where does accountability lie if a user does something wrong? And if brain implants are performed not for therapeutic purposes but to augment people's abilities, will that make the world an even more unequal place?

### **From potential to action**

For some, these sorts of questions cannot be asked too early: more than any other new technology, BCIs may redefine what it means to be human. For others, they are premature. "The societal-justice problem of who gets access to enhanced memory or vision is a question for the next decades, not years," says Thomas Cochrane, a neurologist and director of neuroethics at the Centre for Bioethics at Harvard Medical School.

In truth, both arguments are right. It is hard to find anyone who argues that visions of whole-brain implants and AI-human symbiosis are impossible to realise; but harder still to find anyone who thinks something so revolutionary will happen in the near future. This report has looked at some of the technological difficulties associated with taking BCIs out of the lab and into the mainstream. But these are not the only obstacles in the way of "brain mice" and telekinesis.

The development path to the eventual, otherworldly destination envisaged by organisations like Neuralink and Kernel is extremely long and uncertain. The money and patience of rich individuals like Elon Musk and Bryan Johnson can help, but in reality each leg of the journey needs a commercial pathway.

Companies such as CTRL-Labs and Neurable may well open the door to consumer applications fairly quickly. But for invasive technologies, commercialisation will

initially depend on therapeutic applications. That means overcoming a host of hurdles, from managing clinical trials to changing doctors' attitudes. Frank Fischer, the boss of NeuroPace, has successfully negotiated regulatory approval for his company's epilepsy treatment, but it has been a long, hard road. "If we tried to raise money today knowing the results ahead of time, it would have been impossible to get funded," he says.

Start with regulation. Neural interfaces are not drugs but medical devices, which means that clinical trials can be completed with just a handful of patients for proof-of-principle trials, and just a couple of hundred for the trials that come after that. Even so, ensuring a supply of patients for experiments with invasive interfaces presents practical difficulties. There is only one good supply of these human guinea pigs: epilepsy patients who have proved unresponsive to drugs and need surgery. These patients have already had craniotomies and electrodes implanted so that doctors can monitor them and pinpoint the focal points of their seizures; while these patients are in hospital waiting for seizures to happen, researchers swoop in with requests of their own. But the supply of volunteers is limited. Where exactly the electrodes are placed depends on clinical needs, not researchers' wishes. And because patients are often deliberately sleep-deprived in order to hasten seizures, their capacity to carry out anything but simple cognitive tasks is limited.

When it comes to safety, new technologies entail lengthier approval processes. Harvard's Dr Lieber says that his neural mesh requires a new sterilisation protocol to be agreed with America's Food and Drug Administration. Researchers have to deal with the question of how well devices will last in the brain over very long periods. The Wyss Centre has an accelerated-ageing facility that exposes electrodes to hydrogen peroxide, in a process that mimics the brain's immune response to foreign objects; seven days' exposure in the lab is equivalent to seven years in the brain.

The regulators are not the only people who have to be won over. Health insurers (or other gatekeepers in single-payer systems) need to be persuaded that the devices offer value for money. The Wyss Centre, which aims to bow out of projects before devices are certified for manufacturing, plans with this in mind. One of the applications it is working on is for tinnitus, a persistent internal noise in the ears

of sufferers which is often caused by overactivity in the auditory cortex. The idea is to provide an implant which gives users feedback on their cortical activity so that they can learn to suppress any excess. Looking ahead to negotiations with insurers, the Wyss is trying to demonstrate the effectiveness of its implant by including a control group of people whose tinnitus is being treated with cognitive behavioural therapy.

That still leaves two other groups to persuade. Doctors need to be convinced that the risks of opening up the skull are justified. Mr Fischer says that educating physicians proved harder than expected. “The neurology community does not find it natural to think about device therapy,” he says.

Most important, patients will have to want the devices. This is partly a question of whether they are prepared to have brain surgery. The precedents of once-rare, now-routine procedures such as laser eye and cosmetic surgery suggest that invasiveness alone need not stop brain implants from catching on. More than 150,000 people have had electrodes implanted for deep-brain stimulation to help them control Parkinson’s disease. But it is also a matter of functionality: plenty of amputees, for example, prefer simple metal hooks to prosthetic arms because they are more reliable.


### **Waiting for Neuromancer**

These are all good reasons to be cautious about the prospects for BCIs. But there are also reasons to think that the field is poised for a great leap forward. Ed Boyden, a neuroscientist at MIT who made his name as one of the people behind optogenetics, points out that innovations are often serendipitous—from Alexander Fleming’s chance discovery of penicillin to the role of yogurt-makers in the development of CRISPR, a gene-editing technique. The trick, he says, is to engineer the chances that serendipity will occur, which means pursuing lots of paths at once.

That is exactly what is now being done with BCIs. Scientific efforts to understand and map the brain are shedding ever more light on how its activity can be harnessed by a BCI and providing ever more data for algorithms to learn from. Firms like CTRL-Labs and Neurable are already listening to some of the more accessible neural signals, be it from the peripheral nervous system or from outside

the skull. NeuroPace's closed-loop epilepsy system creates a regulatory precedent that others can follow.

Above all, researchers are working hard on a wide range of new implants for sending and receiving signals to and from the brain. That is where outfits like Kernel and Neuralink are focused in the short term. Mr Musk's four-year schedule for creating a BCI for clinical use is too ambitious for full clinical trials to be concluded, but it is much more realistic for pilot trials. This is also the rough timeframe to which DARPA is working with its implantables programme. With these and other efforts running concurrently, serendipity has become more likely.



Once a really good, portable, patient-friendly BCI is available, it is not hard to think of medical conditions that affect a large number of people and could potentially justify surgery. More than 50m people worldwide suffer from epilepsy, and 40% of those do not respond to medication. Depression affects more than 300m people worldwide; many of them might benefit from a BCI that monitored the brain for biomarkers of such mental disorders and delivered appropriate stimulation. The quality of life of many older people suffering from dysphagia (difficulty in swallowing) could be improved by a device that helped them swallow whenever they wanted to. "A closed-loop system for recording from a brain and responding in a medically useful way is not a small market," says Dr Hochberg.

That may still bring to mind the aphorism of Peter Thiel, a Silicon Valley grandee, about having been promised flying cars and getting 140 characters. There is a large gap between dreamy talk of symbiosis with AI, or infrared eyesight, and taking years to build a better brain implant for medical purposes. But if a device to deliver a real-time, high-resolution, long-lasting picture of neural activity can be engineered, that gap will shrink spectacularly.

*This article appeared in the Technology Quarterly section of the print edition under the headline "Grey matter, red tape"*