NATE KITCH

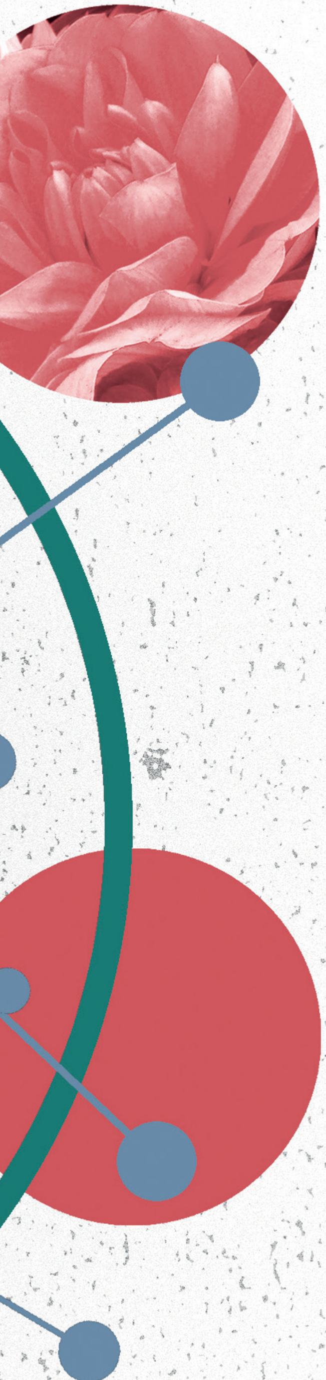

SPECIAL ISSUE

CONSCIOUSNESS

THERE are a lot of hard problems in the world, but only one of them gets to call itself "the hard problem". And that is the problem of consciousness – how a kilogram or so of nerve cells conjures up the seamless kaleidoscope of sensations, thoughts, memories and emotions that occupy every waking moment.

The intractability of this problem prompted British psychologist Stuart Sutherland's notorious 1989 observation: "Consciousness is a fascinating but elusive phenomenon… Nothing worth reading has been written on it."

The hard problem remains unsolved. Yet neuroscientists have still made incredible progress understanding consciousness, from the reasons it exists to the problems we have when it doesn't work properly.

Is consciousness still fascinating? Yes. Elusive? Absolutely. But Sutherland's final point no longer stands. Read on…

We're no longer forced to speculate about the mind, we can watch it in action, says **Daniel Bor**

This is your brain on consciousness

STEPHAN ZAUBITZER/PICTURETANK

THE first time I saw my father in hospital after his stroke, I was disturbed to find that my strong and confident dad had been replaced by someone confused and childlike. Besides being concerned about whether or not he would recover, I was struck by the profound metaphysical implications of what had just happened.

At the time I was a few weeks away from my final university exams in philosophy and neuroscience, both of which addressed consciousness. In my philosophy lectures I had heard elegant arguments that consciousness is not a physical phenomenon and must be somehow independent of our material, corporeal brains. This idea, most famously articulated by Descartes as dualism, nearly 400 years ago, seemed in stark contrast to the neuroscientific evidence in front of me: my father's consciousness had been maimed by a small blood clot in his brain.

Soon after, I abandoned plans for a PhD in the philosophy of the mind, opting for one on the neuroscience of consciousness instead. There are certainly questions about our minds that seem more in the realms of philosophy. What is it like to be a bat? Is your experience of seeing the colour red the same as mine? In fact, how do we know for certain that other people are conscious at all? But I would argue that it is neuroscience, not philosophy, that has the best chance of answering even these most difficult questions.

One area in which we have made great progress is in discovering the physical or neural correlates of consciousness – what consciousness in the brain "looks like", you might say. One way to investigate this question is to see what changes when consciousness is reduced or absent, as happens when people are in a vegetative state, with no sign of awareness.

Brain scans show that such people usually have damage to the thalamus, a relay centre located smack-bang in the middle of the brain (see diagram, right). Another common finding is damage to the connections between the thalamus and the prefrontal cortex, a region at the front of the brain, generally responsible for high-level complex thought.

The prefrontal cortex has also been implicated using another technique – scanning the brain while people lose consciousness under general anaesthesia. As awareness fades, a discrete set of regions are deactivated, with the lateral prefrontal cortex the most notable absentee.

Seeing red

Those kinds of investigations have been invaluable for narrowing down the search for the parts of the brain involved in us being awake and aware, but they still don't tell us what happens in the brain when we see the colour red, for example.

Simply getting someone to lie in a brain scanner while they stare at something red won't work, because we know that there is lots of unconscious brain activity caused by visual stimuli – indeed, any sensory stimuli. How can we get round this problem?

One solution is to use stimuli that are just at the threshold of awareness, so they are only sometimes perceived – playing a faint burst of noise, for instance, or flashing a word on a screen almost too briefly to be noticed. If the person does not consciously notice the word flashing up, the only part of the brain that is activated is that which is directly connected to the sense organs concerned, in this case the visual cortex. But if the subject becomes aware of the words or sounds, another set of areas kick into action. These are the lateral prefrontal cortex and the posterior parietal cortex, another region heavily involved in complex, high-level thought, this time at the top of the brain, to the rear.

Satisfyingly, while many animals have a thalamus, the two cortical brain areas implicated in consciousness are nothing like as large and well developed in other

"IN THEORY WE COULD CALCULATE HOW CONSCIOUS ANYTHING IS, BE IT A HUMAN, RAT OR COMPUTER"

Seats of consciousness

Brain scanning reveals that three areas of the brain play a pivotal role

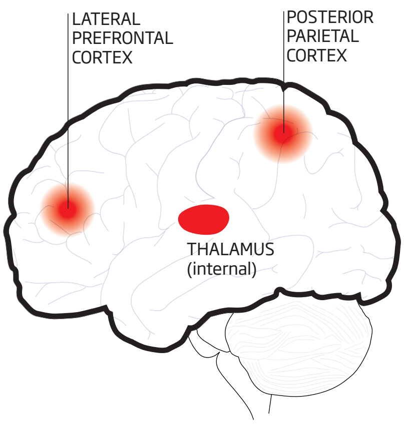

species as they are in humans. This fits with the common intuition that, while there may be a spectrum of consciousness across the animal kingdom, there is something very special about our own form of it.

In humans the three brain areas implicated in consciousness – the thalamus, lateral prefrontal cortex and posterior parietal cortex – share a distinctive feature: they have more connections to each other and to elsewhere in the brain, than any other region. With such dense connections, these three regions are best placed to receive, combine and analyse information from the rest of the brain. Many neuroscientists suspect that it is this drawing together of information that is a hallmark of consciousness. When I talk to a friend in the pub, for instance, I don't experience him as a series of disjointed features, but as a unified whole, combining his appearance with the sound of his voice, knowledge of his name, favourite beer and so on – all amalgamated into a single person-object.

How does the brain knit together all these disparate strands of information from a variety of brain locations? The leading hypothesis is that the relevant neurons start firing in synchrony many times a second, a phenomenon we can see as brainwaves on an electroencephalogram (EEG), whereby electrodes are placed on the scalp. The signature of consciousness seems to be an ultrafast form of these brainwaves originating in the thalamus and spreading across the cortex.

One of the most prominent attempts to turn this experimental data into a theory of consciousness is known as the "global neuronal workspace" model. This suggests that input from our eyes, ears and so on, is first processed unconsciously, primarily in sensory brain regions. It emerges into our conscious awareness only if it ignites activity in the prefrontal and parietal cortices, with these regions connecting through ultrafast brainwaves.

This model links consciousness with difficult tasks, which often require a drawing together of multiple strands of knowledge. This view fits nicely with the fact that there is high activity in our lateral prefrontal and posterior parietal cortices when we carry out new or complex tasks, while activity in these areas dips when we do repetitive tasks on autopilot, like driving a familiar route.

The main rival to global workspace as a theory of consciousness is a mathematical model called the "information integration theory", which says consciousness is simply combining data together so that it is more than the sum of its parts. This idea is said to explain why my experience of meeting a friend in the pub, with all senses and knowledge about him wrapped together, feels so much more › than the raw sensory information that makes it up.

But the model could be applied equally well to the internet as to a human: its creators make the audacious claim that we should be able to calculate how conscious any particular information-processing network is – be it in the brain of a human, rat or computer. All we need to know is the network's structure, in particular how many nodes it contains and how they are connected together.

Fiendishly hard

Unfortunately the maths involves so many fiendish calculations, which grow exponentially as the number of nodes increases, that our most advanced supercomputers could not perform them in a realistic time frame for even a simple nematode worm with about 300 neurons. The sums may well be simplified in future, however, to make them more practical.

This mathematical theory may seem very different from global neuronal workspace – it ignores the brain's anatomy, for a start – yet encouragingly, both models say consciousness is about combining information, and both focus on the most densely connected parts of the information-processing network. I feel this common ground reflects the significant progress the field is making.

We may not yet have solved the so-called hard problem of consciousness – how a bunch of neurons can generate the experience of seeing the colour red. Yet to me, worrying about the hard problem is just another version of dualism – seeing consciousness as something that is so mysterious it cannot be explained by studying the brain scientifically.

Every time in history we thought there had to be some supernatural cause for a mysterious phenomenon – such as mental illness or even the rising of bread dough – we eventually found the scientific explanation. It seems plausible to me that if we continue to chip away at the "easy problems" we will eventually find there is no hard problem left at all. ■

Daniel Bor is a cognitive neuroscientist at the Sackler Centre for Consciousness Science, based at the University of Sussex, UK. His book on consciousness is called *The Ravenous Brain* (Basic Books, 2012)

ALTERED STATES

Sometimes strange states of consciousness can reveal a lot about how we normally construct reality

Liar, liar

Why did you pick that outfit to wear this morning? What made you do your to-do list in that order today? In fact, how did you even end up in that job? You may think you know the reasons, but they could be a work of fiction.

That bizarre conclusion has emerged from studies of people who have had an extreme form of brain surgery - the complete severing of the thick bundle of nerves connecting the two hemispheres of the brain - in an attempt to cure their epilepsy.

Such people usually seem fine, but tasks that test cross-talk between the hemispheres can catch them out. In one test, people had different images shown to each eye, and had to point to a similar image with the hand on the same side as the eye.

When one person saw a snow scene with his left eye, he chose a picture of a snow shovel with his left hand. But when asked to explain his choice, he had a problem. His left eye and hand's actions were under the control of his right brain, as each brain hemisphere controls the opposite side of the body. But language is controlled by the left brain, which could not access the snowy image "seen" by the right brain. So the subject invented a reason that had nothing to do with snow: the shovel was for cleaning out a chicken coup, he said, as a chicken was the last image seen by his left brain.

Such findings have led to the "interpretive brain" theory, which says that the brain makes up narratives about our actions to help us make sense of the world. Any of us can be tempted into this sort of confabulation. In one study, people who have never had brain surgery were told to choose a picture from a selection, then tricked into thinking they had picked another. When asked for their reasoning, their explanations were convincing - and yet had to be entirely imaginary. Who knows how often our consciousness plays these sorts of tricks on us? Clare Wilson