

DICK PASSINGHAM INTERVIEW FOR FIVE BOOKS

1. The Concept of Mind, Gilbert Ryle (1949).
2. Perception and Communication, Donald Broadbent (1958)
3. The Evolution of the Brain and Intelligence, Harry Jerison (1975)
4. Images of Mind, Posner and Raichle (1994)
5. The Organization of Behavior, Donald Hebb (1949)

I wonder if, before we start, you might explain to our readers what we might understand the term ‘cognitive neuroscience’ to encompass? How does it differ from cognitive psychology, or straight neuroscience?

The term, cognitive neuroscience was devised by Mike Gazzaniga while talking to George Miller in a taxi—rather, it was an American taxi, so it was a cab. The idea is to try to understand how human and animal cognition can be supported by the brain.

You have chosen five books that both illuminate advancements in this rapidly evolving field, and mark steps in the development of your own career. But initially, you began your academic life in an entirely different field.

Yes. I did classics at school and also at Oxford for two years, before changing to philosophy and psychology. That was because I went to a public school—private school—and if you were clever the prep school made sure you did classics because the top scholarship at the public school would be in classics. It’s really just that these schools believed at the time, we’re talking about the late 1950s, that ‘gentlemen do classics.’

And I tried, while I was at school to change... But biology was very badly taught, so it would have been very silly of me to do biology. I could have done maths, physics, chemistry, but I wouldn’t have been good at them, I don’t think very well mathematically. So very strangely doing Classics doesn’t seem to have been as disastrous as one might think.

But of course, if you think of the sorts of people who go into psychology, they’re all sorts. Look at the psychologist Daniel Kahneman; he trained in the humanities. Stuart Sutherland, once the professor at Sussex, did Classics and Nick MacIntosh, once the professor at Cambridge, did Classics... In other words, you can do psychology very well, even if you haven’t had a scientific background.

Yes, I studied Experimental Psychology, here at Oxford, and I remember the department underlining that students from both humanities and science backgrounds should apply.

Having done Classics at school, I came to Oxford and started doing the course called ‘Greats’. This combines philosophy with Ancient history. But I wasn’t interested in the Ancient History.

I happened to know Nigel Walker, the reader in criminology at Nuffield College, a lovely man and probably the first person in my life who was a mentor. I was very interested in

crime because I had spent about four months working in the slums of Everton – then there were slums in this country – and I had got interested in why the kids were breaking the law the whole time.

Nigel said that he wished that he had done psychology. So I changed to psychology. At that time you could only do psychology with philosophy, so I went on doing philosophy and I started doing psychology. I went to Gilbert Ryle's lectures; in 1949 Ryle had produced the 'Concept of Mind' and he was still lecturing on it...

This being your first book choice: Gilbert Ryle's *The Concept of Mind* (1949).

I actually remember his lectures almost better than anybody else's, they were in a huge L-shaped room, with him standing on a podium in the middle - very dramatic. And the main burden of his lectures was that we should banish the 'ghost in the machine'—dualism—which most philosophers had gone on believing since the time of Descartes.

This being that the realm of the physical and the realm of the mental being entirely separate.

Exactly, and strangely enough Anthony Kenny, who's a Catholic philosopher, was still a dualist when he taught me and I suspect he still is. I have been told that there are still philosophers around in Oxford who are dualists.

Does this belief have a religious aspect to it?

Yes, of course it does: if you want to believe in an afterlife, and you know perfectly well that the body decays, you are forced to believe that there must be something that can be independent of the body.

The only experiment ever done to find out if that's true was done by Peter Fenwick, a neuro-psychiatrist in London. He had this very good idea: some people, after heart attacks, tell you afterwards that they had out of body experiences and that while they were lying on the table in the operating theatre they were floating above their bodies. So he said, 'I'll test whether that's true.' So what he did—and he's still doing it—is arrange a shelf high up hanging from the ceiling in the operating theatre, and on it was a message; and he tested whether anybody ever read the message.

He published a paper with Parnia in 2014. Of course many people don't have out of body experiences because they die; and of those that don't die, lots don't have out of body experiences; but of the few who report out of body experiences, none have yet read the message!

Regarding Ryle's book, it has been cited as the beginning of philosophy of mind—

Well, it's not the *beginning* of philosophy of mind. Russell wrote a book on the mind, and others such as William James. But I think that Ryle was a landmark, because most

psychologists and neuroscientists now believe in physicalism—that is the belief that I am my brain, my body, and my past history. Ryle's book was the start of that.

How did that affect the way you thought about psychology?

The strange thing is that psychologists had independently decided that all there was was behaviour. I read psychology from 1964 to 1966 and behaviourism was still very dominant. I remember B F Skinner coming from the USA and giving a lecture in Oxford. He taught pigeons to do tricks by what is called operant conditioning.

And the book that was given to students of psychology was by Osgood: *Method and Theory in Experimental Psychology*, which was extremely dull, all about rats running in mazes. The reason behaviourism was strong was that you can observe the inputs and the outputs, you can shine a light on a rat's eye and see what it does, or you can present a pigeon with a choice between two lights and see what it does. You can control what goes in and measure what goes out.

Behaviourism at the time had banned words like 'expect,' 'attend,' 'decide,' because the dictum was that there was no objective way of knowing what, if anything, was happening in the head between the input and the output. Therefore, all you could talk about was the inputs and the outputs. So behaviourism ruled, and of course Ryle's lectures were essentially arguing the same: you shouldn't think of this ghostly mind in the machine, all there was what people did and said.

The problem was, if that's psychology, it's deadly dull. And indeed I found the first tutorials in psychology to be deadly dull. I was given, for example, tutorials on what are called taxes—

Taxes?

Yes, it means movement towards or away from something. Worms, for example, move away from light. It didn't seem to me to be very interesting from the point of view of human behaviour. Having gone into psychology because I was interested in crime, it seemed rather arid.

Which is why when I had tutorials with Anne Treisman —suddenly psychology perked up because Anne was interested in attention. Huh! That wicked word! She was doing experiments following up those that Donald Broadbent had done at what was then the MRC Applied Psychology Unit in Cambridge.

The idea was that you put headphones on and play different messages into two ears. The reason that Donald Broadbent had originally done this was that he was working on applied problems, one of which concerns the airport control tower. The controller will be speaking to many pilots so as to guide them in, and so will have to attend to what they say. The question is, how on earth, given the many voices coming in over the headphones, do you attend to one rather than the other?

Donald had the idea that he would play different messages to the two ears, and he and Cherry found that if you got somebody to repeat back what was in one ear, strangely enough they couldn't tell you anything about what was played to the other ear. So it looked to Donald as if, somehow in the brain, what came into the second ear was being filtered out.

This being what we call the 'cocktail party effect'?

When Cherry worked on it, he called it the cocktail party effect. Exactly: when you're in a cocktail party, voices are coming from different directions, and you've got to use the direction of the voice that you're interested in, even though the voices from other directions may be equally loud. So this was a very simple experimental way of looking at that effect.

And Broadbent's theory was that the unattended messages were filtered by physical properties?

Yes, but Anne Treisman found if two messages are played to the two ears you do hear certain things on the unattended ear, like your name. So not everything is filtered out and meaning and familiarity are relevant.

So suddenly psychology was talking about things like attention, which a behaviourist would not allow. And Donald Broadbent in his first book, *Perception and Communication* [1958], produced diagrams of what he thought must be happening in the brain. And these consisted of boxes that were linked by arrows. One such box might be a filter, something that filtered out what was happening on the unattended ear.

What Broadbent was saying was: 'We have no way of visualising it what is happening in the brain. But my experiment tells me that such and such *must* be happening in the brain. I don't know where or how it's happening but there must be something that essentially acts as a filter.' So he's saying yes, we can only study inputs and outputs, but I can still tell you that a particular operation must be happening in between those.

With Broadbent's book, and the work of Treisman, what they were doing was building a model of the internal process.

It's the beginning of saying what must be happening in the brain. Now of course there were other experiments being done at the same time that also made one think that things were happening between the inputs and the outputs. If you think of Pavlov's dogs—the dog hear a metronome and this was followed by meat powder. The dog starts to salivate at the metronome before the food appears. Now, if that happened in your house, you'd say: 'It's expecting dinner.'

The question is, can we use words like expect? Well, Donald Broadbent worked at the applied psychology unit, and so did Kenneth Craik, who unfortunately was killed in a bike accident in Cambridge during the war. Craik had a mock-up of an aircraft, and in this the pilot would see enemy aircraft coming in. The question was how does the pilot aim at the enemy aircraft? What Craik found was that you don't aim at where the enemy aircraft is,

you at aim at where it *will be*. And you can't explain that without saying that you're *predicting* where the enemy aircraft will be. The layman's word for that is 'expect.'

If you come to a roundabout or traffic circle, it's the same problem. There's a car coming in from the right, and you judge whether it will be on the roundabout by the time you get there—in which case you have to give way to it. Or, will it not be on the roundabout, in which case you can go. So studying problems like this began to break the ice for words like 'expect', 'attend,' and so on.

At the time, of course, little was known about what was actually happening in the brain during these processes. But when I was a student Hubel and Wiesel in America had just begun recording from individual brain cells in the primary visual cortex in animals, and finding out what the cells responded to. Originally it was thought that they'd respond to spots of light, but they didn't; they responded to bars. Then it turned out that some responded to more complex stimuli. Those experiments were the most exciting thing that we heard about as undergraduates.

Now this was some years after Broadbent produced *Perception and Communication*, and of course it was very far away from looking at issues like attention. These days there are people working on the physiological mechanisms of attention in animals, and you can use brain imaging to do the same thing as I have done. But at that time you wouldn't have been able to.

Do you think that one needs to have modelled it before one can understand what these physiological measurements might mean?

I think Donald would have said that. In other words, I think it's quite a common claim amongst psychologists that you need to have some logical, formal claim of what the operation must be before you look at how it's actually implemented in the brain. And this idea was put forward specifically by David Marr.

Your next book choice, Harry Jerison's *The Evolution of Brain and Intelligence* (1975), will bring us back to the question of animals and their brains.

Yes, I did my PhD in London; I had gone to London to do clinical psychology. The course at the Institute of Psychiatry was outstanding, and I did it because I wanted to go on to do criminology—indeed my MSc thesis was on Eysenck's 'Theory of Criminal Personality,' so I was still passionately interested in crime. But on the course seminars were given by various people and one was by a man called George Ettlenger.

The year before I did the course, David Milner had done the clinical course and had asked a question in Ettlenger's seminar, and Ettlenger had asked him if he'd like to come and work with him. So the year I did the course, I asked a question in George Ettlenger's seminar and he said: 'Would you like to come and work with me?!' So David and I sat in George Ettlenger's laboratory, back-to-back because there wasn't very much room, and we did our theses simultaneously. But now I was not working on crime; we were working on animals...

If I ask now how I came to make that huge leap, the reason is that as an undergraduate I was inspired by the lectures given by Marcel Kinsbourne who described various clinical phenomena. The one I really remember is the rare phenomenon when someone with a lesion in the right parietal cortex says, 'Nurse, somebody is in bed with me.' It turns out that they think that the left side of their body is somebody else. I've never forgotten hearing that.

The malfunctioning brain is fascinating.

It was very far away from rats and how they find their way down mazes. So I think that when I agreed to work with Ettliger on animals, I must have had in mind that, yes, there is something really interesting about the brain, but given that at that time we couldn't look at the human brain during life, the only way of actually looking at the brain is by looking at the brain in an animal.

So I worked on animals. And I had a crazy idea, and it comes back to crime. I thought that these kids that I'd worked with were bad at controlling their impulses. And that the prefrontal cortex, or in common parlance, the 'frontal lobes,' must be involved in controlling your impulses. Crazy idea. Anyway, I did an experiment in which I had two lights—one on the left and one on the right, and the one on the left came on eight times out of ten, and the one on the right came on two times out of ten. I wanted to know if an animal which had a lesion in its frontal lobes would be tempted to go to the more common light when the less common light came on? In other words, would it be bad at controlling its impulses? And that's what happened, and I published it.

But if you're going to work on animal brains, the problem is, what if what you learn from animal brains simply doesn't generalise to people? George Ettliger was very worried about that. All of us working in the lab met regularly in an internal workshop, and we wrote a paper on whether or not what you find in animals generalises to people.

And this book, by Jerison, informed your work?

Yes. Harry Jerison was mainly interested in evolution and in particular in the size of the brain. Of course you don't have the brains of ancestral animals, but if you have skulls or partial skulls, you can work out the size of the brain. You can tell very little from the shape of the inside of the skull, but you can at least measure the size of the brain.

So he plotted the size of the brain in ancestral animals and looked at changes over time. And without going into the technicalities of how you compare the size of the brains, it's obvious that one of the problems is that one of the factors that determines the size of the brain is how big *you* are. There's a relation such that an elephant's got a bigger brain than a mouse. Harry had ideas about how you could get rid of the effect of body size, and look at what he called the 'extra neurons' that might contribute to intelligence. I was very interested in that.

My problem was the animal experiments that I did when I came back to Oxford were very boring to run. Science can often be very dull, collecting the data, and it was. So to keep the mind alive, I started doing some calculations about whether the human brain or different

parts of it were bigger than you'd expect, given our size. So, inspired by Harry's book that came out in 1975, I wrote series of papers on these issues for the next five years.

Then Desmond Morris, the zoologist, produced a book...

—*The Naked Ape*.

That's it. And I thought it was naive.

Ha!

So I thought I should write a professional version. The advantage would be that I could include all these calculations that I'd done about the human brain. So I wrote a book in 1982 called *The Human Primate*, and I hoped it would make me famous like Richard Dawkins, but it didn't... Still it was a worthy attempt to try to ask the question as to how people differ from other primates in their brain and behaviour. In other words, what's special about the human primate?

And so I was really influenced by the ideas and questions that George Ettliger had asked, and by Harry Jerison's book. It led me to write a book, more recently, called '*What is Special about the Human Brain?*'.

Is this based on the idea that there *should* be something special about the human brain?

Well, I have changed my mind. In *The Human Primate*, I suggested that the trends that you can see if you compare a monkey with a chimpanzee are continued if you look from chimpanzee to human. So what I was stressing was the similarities, that we were following trends. Of course the analyses are based on modern species, not the actual ancestors.

But when I came to write my later book I had already done some work using brain imaging. I was beginning to get cold feet because it seemed to me there were some things that might be special, that I should try to investigate. So I went back on some of what I'd said earlier. You're no good as a scientist if you haven't ever been wrong.

There are some things that you and I can do that other animals can't do. One of them is what you might call 'mental trial and error'. We can think: 'If I do A what would happen? If I do B what would happen?' and do this before we act. This means that we don't just rush in. There's a selective advantage in being able to think before you act.

Of course, animals can plan, but the experiments I know of are ones where, let's say, there's a maze on a screen and the animal moves a cursor through the maze so as to find a goal in the maze. There are cells in the brain which specify the end location long before the animal has moved the cursor there. The activity of these cells reflects the planning.

But of course the maze is visible. Yet I can think about whether I'm going to have cornflakes or cauliflower for breakfast tomorrow, and these are not visible. It's not clear to me that a chimpanzee can do this. So this idea of mental trial and error seems to me an important way

in which people differ, one that confers a major selective advantage. Steve Wise and I wrote a book called *'The Neurobiology of the Prefrontal Cortex'*, and we gave it the subtitle, *'The Origins of Insight'*. We were suggesting that the ability to think about the problem before you act depends on the prefrontal cortex.

So, work by Jerison and by yourself in finding the similarities and dissimilarities has been a major step in psychology in as much as you can show—

Now wait, Jerison's book is one of the classics, an example of someone going off and doing something totally new. It's a very major bit of work involving the analysis of a huge number of fossil skulls. So I don't think it's fair to compare the weight of what Harry did with what I did.

This marks a major step in as much as it demonstrates that animal experiments, which have been done for decades, were valid?

Yes, work of this sort looks at those respects in which those experiments are valid but also at the limitations of those experiments.

Perhaps we might move on to your fourth book choice, which takes us into the 1990s and the advent of brain imaging.

The problem is that we don't just want information about the size of different areas of the human brain: we can get this post-mortem. We need information about the living human brain, that is *while* we're doing things.

When I did my PhD, the only way of seeing whether somebody had a brain tumour was to pump air into the spinal cord; it went into the fluid filled cavities, the ventricles in the brain, and you could see those in an X-ray. If there was a tumour, the ventricles were distorted. And that was the only way that you could see the brain. It gave the patient a dreadful headache.

Since then there have been major advances, first of all CT scans in the early 1970s. You take a series of X-rays from different angles and you can then produce a picture of the brain. Doing this involves computed tomography, so called because a computer is used to reconstruct the whole brain from slices— *tomos* being Greek for a cut or section.

Then later in the 1970 MRI was developed for scanning human tissue. Paul Lauterbur and Peter Mansfield got the Nobel Prize for this development. MRI gives exquisite pictures of the structure of the human brain.

But in the 1980s, a new method was invented, which enabled you to look at the brain at work: this positron emission tomography [PET]. The idea is that when an area of the brain is active, it needs oxygen and glucose and these are brought by the arterial blood. So if you can measure the passage of the arterial blood, you will be able to see which areas are active

when somebody is in the scanner. And this particular method introduces a radioactive tracer into the blood so that you can detect the blood flow.

As it happens, I heard Richard Frackowiak lecture in Oxford in the late 1980s, and I went to see him at the end and asked if I could collaborate. He was working at the MRC Cyclotron Unit at the Hammersmith Hospital, a pure research unit. He said yes, so I started going down to the Hammersmith and doing experiments.

Then, yet another method was invented in the early 1990s: it was found that you could measure the ratio of the oxygenated blood and the deoxygenated blood once the oxygen had been removed, and that you could do this using an MRI scanner. The measurement is called the BOLD-contrast and the technique is called fMRI or functional magnetic resonance imaging.

However, experiments using PET and later fMRI only took off after the psychologist Posner worked out a way of analysing the data.

This is Michael Posner, who along with Marcus Reichle, wrote your fourth book, *Images of Mind* (1994).

Yes, Posner is one of the great psychologists. He was interested in things like reaction times. Let's suppose I ask you to press a button on the left if a light comes up on the left, and on the right if the light comes up on the right, and to do so as quickly as possible. That's called a choice reaction. And it takes me roughly 500 milliseconds. But, how long did it take me to make the choice itself? Well, in the 19th century Donders showed that you can compare that time with the time it takes you simply to press a button if a single light comes on, roughly 200 milliseconds. Now subtract the simple reaction time from the choice reaction time and you get an estimate of 300 milliseconds for the time it took you to make your mind up.

So you have two conditions, and you subtract what you find for one from what you find for the other. It was on this basis that Posner, working with Marcus Raichle and Steve Petersen, devised a classic experiment using PET. They were interested in the language system. So in one condition they showed a word and the person simply looked at it. In another condition, they showed a word and the person had to repeat it. But they were not interested in vision; so, though they scanned the subjects in both conditions, they then removed everything there was in the scan for the looking condition. What there were left with was a scan that only showed the areas that are involved in repeating.

Subtraction imaging, I think this method is called.

Yes. Then they had a third condition. They showed a noun but rather than repeating it the subject had to say a verb that's relevant. So, 'cake,' or, say, 'drop'—

'Eat,' 'slice...'

Yes. It's up to the person. So now they were interested in how you generate a verb that is associated with a noun. Of course, the subject had said something, but they were not

interested in speaking since they already had a scan from when the subject repeated the noun. So they removed everything that was in that scan and what they were left with was a scan that only showed the areas that involved in generation.

And that method's described in the book by Posner and Raichle that came out in 1994, as a Scientific American publication. It's a very simple book, but the subtraction method has become fundamental to the analysis of data from functional brain imaging.

So you can scan the brain at work and the subjects don't end up with a dreadful headache and you don't have to kill anyone to do it! But you can also look at what the different bits of the brain do, and compare those results with what you find when you record the activity of brain cells in animals, or—as is now being done—what you see when you scan animals. People are now also looking at the anatomical connections between brain areas because you can visualise these using scanning methods and compare them in human and animal brains.

***Images of Mind* brought together a cognitive psychologist and a neuroscientist. Is that significant?**

Yes, this happens so often in science. Crick was a physicist, Watson was a biologist. It's true that Kahneman and his colleague Tversky were both psychologists, but Tversky was really a mathematician: he published a book of 1000 pages with dense maths. Anyone who's done science knows that most of the ideas generated are actually generated in discussion with other people, often younger people, and it's never clear who actually thought of a particular idea in the first place.

Take the book that I wrote with Steve Wise on *The Neurobiology of the Prefrontal Cortex*. We Skyped once a week—he's in America, I'm here—so we had regular discussions over two years. He knows more about some things than I do and vice versa, I've no notion where most of the ideas in the book came from. So I think it's very relevant that advances are often made when two people with very different backgrounds come together, such as Posner who is a psychologist and Raichle who is a neurologist.

Your final book choice moves us very much into the physical realm, the neural basis of cognition...

Yes, there's a limitation, you see, of imaging. I ask you to make a decision in the scanner, and using the subtraction technique, I find, let's say, activity in the prefrontal cortex when you made that decision. The image shows a patch in which the cells are active. The patch is in the order of several millimetres. But there are millions of brain cells in that patch!

Though the patch tells me *where* something's happening, it doesn't tell me *how* the brain cells do it. But the fundamental aim of neuroscience is not to ask where are things happening but how they are happening, that is what the mechanisms are.

In recent years methods have been developed to record from individual brain cells in the human brain during surgery. You can do this while the patients are awake because there are no pain fibres in the brain. But there's a real problem: you can record from cells, or groups of cells, but there are an estimated 86 billion cells in the human brain. So if you can only record from 20 cells, or 200 cells, you're in real trouble.

How are you going to work out how 86 billion cells work? You might think that what you need to do is get a computer, and try to teach it how to do the sorts of things that people do. People at Deep Mind in London are doing just that, for example teaching a computer how to play the game Go. But we want to know how the *actual* brain works. And Donald Hebb, in 1949, published a book that is fundamental to this enterprise, my fifth book.

This is *The Organisation of Behaviour*.

Yes, it was a theoretical book, because at that time we knew very little about how the brain worked. But he had two ideas that have become absolutely central to our understanding of how must do so. The first idea was a suggestion as to how the brain learns.

When Donald Hebb wrote his book, electrophysiologists and anatomists had shown the following. Brain cell A has a cell body and a long process or axon; and so-called 'action potentials' are propagated along this axon. But the axon of cell A doesn't actually touch cell B; instead there's a gap between them. We call this the synaptic gap. The terminal of the axon of cell A influences cell B by releasing packets of chemicals which are taken up at so-called post-synaptic sites on cell B.

But what happens during learning? Well, Hebb suggested that there must be changes at the synapse. He further proposed that the more frequently cell B fires at the same time as cell A, the more likely it will fire in future when cell A is active. This idea led to the term the 'Hebbian synapse'.

That idea was taken up by Giles Brindley, a physiologist from Cambridge who then went to the Institute of Psychiatry in London. And he had a PhD student called David Marr, a mathematician, who joined him there. In his thesis, David Marr produced theories as to how three structures of the brain might work: the cerebellum, the hippocampus, and the neocortex. These papers have had a phenomenal influence. David was the only genius of my generation that I have known. He was also a mentor. He once told me at tea that I should read less and think more.

David Marr then went to Cambridge to work with Sydney Brenner and Francis Crick, but he was then poached by Marvin Minsky to go to MIT. While at MIT, David Marr helped found the field of computational neuroscience. Tragically, he died when he was 35.

If you take Marr's theory of the cerebellum, the fundamental idea is that it supports motor learning and that it can do so because of Hebbian synapses. But he died before anybody could succeed in testing the basic prediction which was that there were modifiable synapses in the cerebellum. The Japanese neuroscientist Masao Ito was able to confirm this prediction, but only after David Marr had died. This was ten years after the publication of

Marr's original paper on the cerebellum in 1969. And this was 20 years after the publication of Hebb's book.

Marr then went on to argue that they must be also be modifiable synapses in the hippocampus and the cortex. And it was around that time that a mechanism was found called 'long term potentiation', Understanding the chemical mechanisms for this has become fundamental to understanding how learning occurs.

The other idea that Hebb had was what he called 'cell assemblies'. The idea was that if I see object A, then *this* assembly of cells fire whereas If I see object B *this* assembly of cells fire. Many cells in assembly A will also be part of assembly B. So cell assemblies are arranged in a series of Venn diagram.

When it became possible to record from individual cells, we started thinking that we can understand how the brain works in terms of the firing of individual cells. But, of course, It's not that one cell fires when you see something: whole groups fire. So now we realize that the brain works in terms of so-called 'population coding'. It is whole populations of cells are coding for something.

So the reason why Hebb is so critical, is that—though he proved neither—these two ideas have turned out to be incredibly powerful in understanding how the brain *actually* works. This year the neuroscientist Lucia Vaina—David Marr's wife—and I have edited a book in which a group of computational neuroscientists take up the original ideas of David Marr. They ask how we think the various operations the brain performs are actually implemented, given the cells that are there and the connections that are there? That *must* be the final aim of neuroscience. Neuroscience is coming of age.

What you were describing, about patterns of activity in the brain being the basis of thought, and the demise of dualism—it reminded me of when I was an undergraduate, telling my friends about what I was learning. Many had the same response: 'surely that's very unromantic, to think about everything we experience as being merely a bunch of electrical signals in the brain?' But I remember a module on religious experience, how that actually manifests in the brain, and being struck by the beauty of it—being able to visualise that process in action.

I think that science is phenomenally romantic. Or... I don't know if 'romantic' is the right word. I mean *exciting*, because any day you might be looking in your data wherever the data is, and you might find things that totally change your mind. You might find something that you didn't expect in the slightest. People sometimes think that all scientists do is have hypotheses and test them. But a huge amount of it is seeing things that you never expected to see.

Astrophysicists are constantly being horrified by new phenomena that they didn't know of! It's so exciting. To paraphrase Shakespeare, there are things in heaven and earth we never dreamt of. Poets may dream of them, but scientists have the excitement of finding them.

I suppose looking into the brain is another way of looking into the great abyss. It's almost as unknown.

Yes, it's true that we know very little about the brain. It's the most complicated thing there is to understand. But at least things are on the move. For many years neuroscience was dull because nothing much was happening. And then suddenly, in the nineties, it took off. Because of brain imaging in particular it's suddenly become a fast moving field.

If I think about philosophy—philosophers have been wondering about dualism ever since Descartes, and they've got nowhere. But if you talk to a young graduate student or post-doc in neuroscience today, they're not interested in what happened a few years ago because things are moving so fast. And then they can find something new tomorrow. Exciting.

Do you think modern-day work in cognitive neuroscience is making all these years of philosophical questioning irrelevant?

Yes, I do. Everybody knows that if we use terms, we need to be clear about what they mean. But philosophers worry about things like the fact that we can't prove that there's an outside world—scepticism. They worry that because something has happened regularly in the past doesn't mean that it is bound to continue to do so in the future – the problem of induction. But there are no solutions.

Of course I shouldn't say that philosophers have made no advances. They've made some technical advances. But if I compare that with the rate at which science can advance, then there's no comparison. If you want to study consciousness, philosophers have got nowhere. They ask questions like 'how do I know that what it looks like when I see green is the same as what it looks like when you see green?'. And they've gone on worrying about that for an awfully long time, and they probably won't be able to answer it.

But the problem of consciousness is one that actually *is* open to empirical investigation, and people are studying it in many ways, including people who've worked in my lab. For example, you can give propofol, an anaesthetic, and study what's happening in the brain as people lose awareness of pain, sounds, and so on. You can look at what happens when people are or are not aware of things that they see. In other words the *empirical* study of consciousness has moved fairly quickly. The *philosophical* study of consciousness is static.

Are these the people who really want to know the answer, rather than enjoying the process of questioning?

Peter Medawar, a great biologist, wrote a short book called *The Art of the Soluble*. His point was that there is an art in finding problems in science that are tractable. I'm not interested in the question of whether I can prove whether the outside world exists. Nobody can. The great thing about science is that it takes all sorts and there are a multitude of problems out there - and they can be solved.