Decoding Consciousness

Thankyou very much for that very kind introduction, [...]. So, what I’m going to talk to you about today is just to give you an idea – [...] – to give you an idea of the kind of work I do. I’m interested in consciousness. It’s the inner life that we all share, the sense of someone being ‘in there’. There might be some zombies in the audience, some people without an inner life who just walk around, the lights are on but nobody is at home. I suspect the majority of you all have a very vivid vibrant inner life, what it is like to be you. I’m interested in that and I’m specifically interested in how brain activity leads to consciousness. So what I’m going to show you is some of the technology we use here at University College London and elsewhere in the world to investigate that, and some of the ways in which clinicians and scientists like myself investigate those problems, and finally show you a couple of examples of my recent work. I’m going to try to keep it non-technical because I know you’re a general audience, and I’m going to try to speak slowly like I’m doing now – I’m a very excitable enthusiastic person, and when I get excited I tend to talk a lot quicker like this, so should I start talking quickly and getting very excited, please stop me and slow me down, tell me to calm down, and get back to something we can all understand. The idea is just to give you a taste of the sorts of things I do, not to give you a kind-of detailed research talk. Good, that seems to be going down well already, okay, that’s the end of the lecture, we can all go. Just kidding. I’m also very happy to take questions as we go along, that’s fine, so just stick up your hand, wave it so I can see you in a big lecture theatre, and we’ll do questions as we go along. There are no stupid questions in science. So don’t be afraid; if you have a question, ask.

So, this is the kind of thing I do. Summarised in one single slide. I use a machine to do brain imaging – it’s a very annoying projector, if it’s going to do that all the talk that’s going to be very annoying. You see the flickering of the projector? I’m sure that’s [...] to be solved. I think it’s a projector problem. – Anyway, I use functional brain imaging devices to measure activity in the brain when people look at particular images, hear sounds or have touch sensations. And from that we derive pictures of activity in their brain. This, for example, is activity in my brain while I listen to someone giving a lecture. There is activity in this particular part of the brain which corresponds to auditory cortex, the part of the brain just inside here, [...]. We also produce rather more exciting and hard to understand pictures, like this picture which I’m sure most of you have no idea what it represents. But that’s a view of my brain from the back where I’ve coloured in some of the early visual areas that respond to the visual information that comes in through our eyes. And this image is even more confusing because this is an image where I’ve taken the back of my brain and, if you like, stretched it out and ironed it flat. So this point marked *f* here corresponds to the very tip of the back of my brain. And it’s called *f* because that bit of the brain represents the fovea which is the most sensitive part of your eyes. It’s the bit that if I ask you to look at me now, you’re pointing the most sensitive bit of your eyes, the fovea. It’s a little bit of your retina that’s about the size of a thumbnail [...]. In fact, we all view the world like that. We don’t look at the world – we all think we see the world as if it’s like a camera, with a big beautiful visual field full of interesting things like your lovely faces looking at me. But if we actually go and measure that scientifically, what people can see, the detailed bit with which you see is just the fovea. If I damage that bit in your eyes with a laser like this, or a more powerful one, you would not be able to read, even though you think you can read text out here, you can’t. And in fact, the way we experience the world scientifically, it’s more like moving something like this around awfully quickly because our eyes move around very rapidly, about three or four times every five or six seconds, that kind of frequency. So the scientists’ view of how we see the world is very different to our actual experience. And yet somehow, I have to connect or explain how we have all these visual experiences and explain it in terms of brain activity.

So I’m going to talk about three different things in my talk. In the first bit of my talk, I’m going to just talk generally about how we measure brain activity in humans. I’m going to focus on new techniques like this functional brain imaging I described. Then I’m going to talk about consciousness, how we might get a grip on it. We all wake up in the morning and start having experiences. They’re rich and they’re vibrant and they’re detailed. How are we going to start studying that? How are we going to start understanding that type of – how that happens in the brain. And I’m going to explain general ways in which we start to address that problem. And then finally, I’m going to talk – to give you just two examples of recent work from my group.

So, let’s think about measuring brain activity. Well, first of all, you have to know a bit about the brain. I’m not sure that all of you do know much about the brain. I’m sure you do know that the brain is inside your head, - and you can take it out and put it in a jar – but what actually makes up the brain? Well, like all the other organs in the body, the brain is made up of a large number of specialised cells: those cells are known as neurons. Here’s a very early drawing of neurons from [...] – gross apologies if I mis-pronounced that for those of you that speak Spanish – but he was an anatomist at the turn of the twentieth century. He was very interested in understanding the anatomy of the brain and spinal cord. And he drew these neurons: as you can see, they have a cell body, they have this tree-like structure, these little dendrites, and they have a long process called an axon which extends down here. That’s another picture – a schematic picture – of a neuron, again showing that there are these frilly dendrites, a cell body, and then a long axon. I can't give you a – there's a whole lecture series on how neurons work – but a simple way to think about it is the neurons all connect up to each other by their axons. The dentrites take input from other cells. The neuron then processes that input and sends its output down the axon as an electrical spike. Neurons are very small: a typical neuron is about ten microns. That's very very small, okay, bear that in mind. There are a lot of neurons in your brain: about ten to the power ten, in other words that's ten with ten zeroes after it, which is a very big number. For comparison, this is a rough estimate of the current population of the earth: six point five times ten to the nine. In other words, in each one of your brains there are ten times more neurons than there are people on the earth. And in this lecture theatre of a hundred, there are therefore going to be a thousand times more neurons in your collective brains than the population of the earth. So that kind-of makes it a very hard problem to understand how the brain works. There's a lot a machinery, a lot of neurons that we have to try and understand, build models how they interact. But the first step is to try and measure the activity of these neurons. Thinking about how the neurons are gathered together in the brain, this is a typical human brain, and if we took a cross-section here – a slice – this is what it would look like. The exterior would look kind-of grey; the interior would look kind-of white and shiny. And there would also be some patches of the brain in the middle of the brain. That's a basic distinction between what we call grey matter and white matter. Grey matter is where all the neurons' cell bodies live. So these bits, these bits and these bits. The white matter are these axons. One way of thinking about it is your grey matter is what you think with, it's where the bits of the neuron for thinking live, and the white matter are the connections between neurons: the axons, the wiring that goes between. That's one way of thinking about it. So what we want to do is measure the activity in the grey matter of the cortex. There's quite a lot of it. This is the area of the cortex. If you flattened it out, because this is like sheet that is convoluted or folded up. If you stretched it out, and ironed it, it would be about half the size of an ironing board – the sort of ironing board you'd iron your blouses or shirts on. And it's pretty thin: this is about 1.7 millimatres. But still, I want to remind you even though it's about half the size of an ironing board, I want to remind you just how much stuff there is in it. Each cubic millimetre of this grey matter contains ten to the power five, that's 100,000, neurons. And each cubic millimetre – millimetre remember, just this much – contains three kilometres of that wiring coming out of it or going around it. So we're still not – and these things called synapses are connections between neurons. And there's about four thousand of those per neuron, making about five times ten to the eight – ten with eight zeroes after it. So the numbers of all big, that's the real message. You've got a complex task to explain it. And the techniques we have at the moment are fantastic, but they have limitations. These are the three major techniques that we have for looking at human brain function in normal, healthy human volunteers. In animals, we can put electrodes in the brain: we can drill a hole in the animal's skull, put an electrode in under anaesthetic, and measure the activity of single neurons. That's very good, that's very exciting. The problem with that is, of course, we have ten to the power ten neurons to explain, we can't put ten to the power ten electrodes in nor can we measure that number of neurons in a single recording session. These techniques have the disadvantage they measure much larger populations of neurons. But they can do it non-invasively in humans without needing operations or anything like that.

So how do they work? This is Magnetic Resonance imaging. This is an example of a subject undergoing EEG examination. EEG stands for Electroencephalography. And this is an example of an EEG scanner. The subject sits here in this chair and their head goes in to a sort of bucket there. Imagine you had a toilet, you turned it upside down, that's kind of similar to an EEG device, without the water... How do these devices work? First of all, let's tell you some more interesting information. How much do they cost? How much do you think one of those would cost? Do you have any idea how much that would cost? How much would you pay for it? One thousand pounds. Anyone willing to go higher than a thousand pounds? Higher? One million pounds, okay, that's good. One million pounds, that's a good guess. Ten million – no, lower, no. One point five million pounds would be a rough cost for how much one of those [...]. How much would one of these cost? This looks a bit cheaper, doesn't it. What do you think? Give me a guess. What do you think? How much would you pay for this lovely device here? [...] A thousand. Ten thousand. Okay, a little bit higher. A hundred thousand. Yeah, okay, somewhere between fifty and a hundred thousand. What about this price? This looks very good, doesn't it; very nice. How much would you pay for that? A million? A million. That's about right. I think you've got the hang of it now, you can price these machines very well, I'm impressed. So, they're a combination of quite cheap and very expensive technologies. How do they work and what do they do? Well, I'm going to talk about MRI – Magnetic Resonance Imaging – in a little bit of detail shortly, but briefly, MR machines that are used in research are very similar to the ones that are used in hospitals around the world. This machine is a 1.5 Tesla Siemens [...]. It's made by Siemens, and it will be found in many hospitals throughout the world. And the typical way that MR machines are used in neurology, or neurosurgery, is to produce structural images of the brain, like this one here. It's an image of the brain's structure as it looks like, the anatomy. So physicians or surgeons can diagnose disorders like stroke or brain tumour, which are structural abnormalities. We use it in research to produce functional images: images that reflect not just the structure of the brain, but how it lights up, how it activates, how the neurons change their activity, and I'm going to explain a little bit about how that works. It's a good technique because it has relatively good spatial resolution, you can see the detail in these images is quite fine, and we can see detail in the brain down to one or two millimetres.

These techniques are rather different. They record electrical activity produced by the brain. So the EEG technique is the oldest, and you see the subject wears this cap, like a swimming cap, but it has little electrodes which make contact with the skull. Those allow you to pick up the tiny electrical currents produced on the surface of the skull by the brain's activity underneath. I'm sure many of you have heard of EKGs or ECGs, Electrocardiograms which are used to measure the electrical activity of the heart, and I'm sure you've seen television programs where that's measured. Just as the heart produces electrical activity [...] the body, the brain produces electrical activity that can be measured on the surface of the skull. It's much smaller and so requires more sophisticated amplifiers to measure, but this also can be used in clinical practise by doctors and physicians, particularly for patients with disorders such as epilepsy, which are disorders of electrical activity in the brain. It produces traces like this: wiggly lines from each electrode, that reflect the moment-by-moment changes in electrical activity in the brain. So what this produces is a very high temporal resolution; a very detailed picture of how things are changing over time, much better than MRI. But it's not very good at telling you where that activity occurred, because the electrical activity at any one electrode comes from a number of different places, because your skull conducts electricity, weakly. Not strongly, that's why when lightening hits you it goes through your body and you die, because it's a weak conductor, not an insulator. If any of you have been struck by lightening and survived, please see me, it would be good to do some experiments. Anyway. Imaging is a more modern technique that works in a very similar way, but now records magnetic fields that are produced. Some of you may have done physics at school. I did, I can't remember any of it. But one of the bits I remember is this relationship between changing electrical and changing magnetic fields. Does anyone remember that? Yeah, nodding faces when I do this. When you change an electrical field, you change the magnetic field at right-angles to it. That's how electrical generators work in power stations [...]. So what I'm saying is the electrical activity in each and every one of your brains causes very tiny changes in the magnetic field around your heads. Very small, because if they were bigger you might attract coins to your head if you thought a lot. Those can be picked up using an MEG devide, which produces a trace very similar to the EEG. But it's a lot snazzier equipment, it's more expensive and it makes us feel like we've got a really nice toy to play with. So it's quite a popular technique right now. We have one of these and one of that one, and two of those in our laboratory, purely dedicated to research. And the combination of those two are used by about 120 researchers at UCL in our square. So although they're expensive pieces of machinery, they're typically used by large numbers of scientists.

Okay, a little bit more about how MRI works. The reason we can use MRI to detect activation in areas of the brain, is that some some scientists discovered in the early 1900s, in the USA and Japan, that haemoglobin behaves differently in a magnetic field when it carries oxygen and when it doesn't. Just to explain that, haemoglobin is a protein that is in our red blood cells, the cells that circulate and we see when we bleed or cut ourselves. Haemoglobin carries oxygen from the lungs to the tissues and oxygen is needed by every cell in our body. As it gives up the oxygen to the tissues and returns in circulation to our lungs its magnetic properties change, and it turns out we can detect that using an MRI scanner. The second exciting thing that happened was people realised that blood flow and oxygenation – oxygenation, the proportion of oxygenated haemoglobin – changed when nerve cells fired. Because nerve cells need oxygen, when they fire more they need more oxygen, so more blood flow goes to the part of the brain that is active. And that meant – these two facts meant that you could combine these two to get a measure at each location in the brain of how the brain blood flow responded to something happening in it: for example, a word, or a flashing light, or a thought, happening. At that point in the bain, over this sort of timescale: this is about five seconds, this is about ten seconds, this is about twenty seconds – there will be a rise in signal – sorry – measured with the MR scanner. And that is known as Blood Oxygenation Level Dependent Contrast – BOLD Contrast. You may read that, sometimes even in newspapers, BOLD contrast, when they're talking about that sort of study. So we use an MRI scanner to collect a lot of images of people's brains, and the signal at each point in the brain goes up and down as they think or are shown images or hear sounds, and we can make pictures that reflect their brain activity. Finishing off the segment, bear in mind that even though that's very cool - it's very cool because it's non-invasive, we don't have to do anything to the subject, and yet we can acquire a detailed picture of activity in our brain – it's still at a very low resolution. So what I'm showing here is the size of one typical voxel we get in one of these brain images. A voxel is like a pixel in a digital camera picture. But a digital camera picture is two-dimensional, where as this is a brain so it's three-dimensional. So a voxel is to a volume as a pixel is to a picture. The size of voxels is about three millimetres on the side. Now that's a lot of brain, isn't it. Three by three by three, you do the mathematics – it's 27 cubic millimetres. And you remember at the start of the lecture, I told you that in every cubic millimetre there was 100,000 neurons. So in 27 there's a lot of neurons. So functional MRI measures the activity from a large population of neurons. Which can differ in important ways, so this fantastic colour image here comes from an experiment in a monkey where they showed different regions of the monkey cortex respond best to different orientations of bars shown to the monkey. All I want you to get from this picture is nothing complicated about how a monkey brain works, but just the idea that different patches of monkey cortex have different colour, so the cells in those regions – the neurons in those regions are doing different things, and yet our functional imaging voxel is quite big, so we're averaging across all those that do different things.

So, what this means is we have techniques that can non-invasively measure brain activity. And it's got reasonable resolution. The spatial resolution is in the millimetres range and the temporal resolution is in the seconds range. We would like to do better, and the next ten or twenty years might bring new advances and new technologies that go even further. But even this, the availability of MRI scanners in this way, has revolutionised a lot of neuroscience – the study of the human brain – in the last decade. And in every major university in the world, you will now find dedicated functional imaging centres. And, as I say, it's transformed the type of experiments you can do. It was only discovered in 1994 so this is very young technology still.

Okay, so that's the first part of the talk. How do we – so we've got this brilliant technology, how do we use it to study consciousness. We all sort-of know what consciousness is, isn't it. It's – you're thinking stuff right now. You've got thoughts inside your head. But how will we actually try and understand, or measure that? Do we just put someone in the scanner and say, 'think. Do some thinking. We'll measure your brain activity'? Well we might do that but it would be rather hard to interpret the results because different people would be thinking different things. So the first thing is just to sit down and think, well, what do we mean by consciousness? What are the components of consciousness? I think it's useful to think in terms of three different things: level, contents and self. Let me explain that.

Level of consciousness. Right, you all have a particular level of consciousness at the moment. I guess – I think – you're all conscious. You're not unconscious. You might be a little bit sleepy, some of you, because it's quite hot. You might be very excited because you're in the front row and I'm speaking loudly and very excitable. You might be particularly excited by my lecture so your level of consciousness might be up or down, but you're certainly awake rather than asleep. And we know of course that tonight, when we go to sleep, we will become unconscious – I hope. So there's something in the brain that controls that, that controls the level of consciousness. And we might use our brain imaging device to look for a neural correlate for level of consciousness. But when you are awake, throughout the day, the contents of your consciousness change. So, for example, the smell of a rose is a particular quality that is not the same as the rasp of sandpaper, or the taste of a fine wine. Or the feeling of your buttocks on the seat that you're sitting on. I bet, until I just said that, you weren't really thinking about your bottom, unless you're a rather unusual individual. And yet I can direct the contents of your consciousness to think about something else: your bottom on your seat. And it has a particular quality. A lot of us call these things qualia: the quality of experience. And so we might think about the subjective aspect, the particular content of consciousness. That's what I'm trying [...]. And then we might be awake, we might have particular contents of consciousness. But – and that might change, this might change through the day. But there's always someone in there, that's Me or I, having those experiences, self-consciousness. You don't ever mistake any of your thoughts do you? You don't ever think, oh, that was your thought in my head. Or maybe you do – I hope not. Some people do, of course, so a psychiatric illness called schizophrenia is one where one of the problems that individuals with schitzophrenia can have is that they can have difficulty distinguishing the thoughts in their head that are theirs and they think that some of the thoughts in their head have come from other people. So they might say, 'I'm not thinking that – someone out there has a special brain machine which is projecting thoughts into my head, because this isn't my thought'. But that's a very unusual situation, it's very hard for us to imagine what it must be like. But that is a disorder of self-consciousness. And you can have medical disorders of the other things as well, the contents of consciousness: well, imagine a hallucination, many psychiatric patients have hallucinations – 'ooh, who's that over there' – that's an example where I have a content in my consciousness that has no physical stimulus out there. There's nothing happening over there. 'Aah, it's my mum's ghost' – or whatever. That's another mental disorder, of contents of consciousness. And of course we can all think of mental disorders at the level of consciousness. There are famous cases of people in persistent vegitative states and minimally conscious states that have made news headlines around the world where people wake up from their coma. And of course every time someone has surgery in a hospital, that depends on the anaesthesiologist manipulating the level of consciousness, because none of us want to be conscious when we're having surgery. So I'm interested in these neural structures not just for the scientific interest, but also because I'm a physician and I want to understand the physicians' processes that are affected in that way.

So, let's think just about contents of consciousness, what are we going to do? Well, we want to understand how the contents of consciousness: smell of a rose, taste of a fine wine, the colour red – how that subjective quality is encoded in brain activity. But there's a problem. We know a lot of brain activity is just unconscious. So if we think about the colour red, there's something – I'm looking at your red top, I have a subjective view – but there's obviously a lot of unconscious processing in my brain going on before I get to the quality. And then of course if I talk to you about it there's a lot of unconscious motor behaviour. To give you the intuition, I'm sure many of you will have seen Terminator 2. Arnold Schwarzenegger. If not, it's worth looking at. Arnold is a cybernetic organism, a robot, who at the beginning of the movie comes to earth naked, with no clothes. And he needs to go and find some clothes and so he wanders into a bar where there are a large number of bikers, and essentially tries to identify the individual who has the right size clothes for him to take. As part of that movie, you see a sort of visor-view of inside Arnie's brain, where as you're going around is outlining a person and there's a little text saying, calculating size, calculating the right size and the wrong size, calculating [...] the right size. That's, if you like, what happens in our brains every time we identify an object. So if I say, what's that?, you don't have a display inside your head that says, 'outlining object, calculating object, probability overhead projector: 100%' and you've got the answer. You just say, it's an overhead projector. The way you do that is your brain has somehow taken the pattern of light falling on your eyes and done all that calculation for you unconsciously. Similarly, when you say, overhead projector, or when I'm talking to you, I don't think about how to put the tongue in my mouth in the right position and how to open my mouth in the right way. All I do is think I'm going to say the word 'hello' and out it comes. And it's the same for each one of you and has been since you were a child and learned to talk. But there must be some machinery in our brain that allows us to position our tongue correctly, move our mouth correctly, breathe correctly in order to talk. And that, if you like, is an example of unconscious motor activity. So if we just put people in the scanner and measure their response while they look at things, what we see in their brain is a mixture of stuff that's unconscious Arnold Schwarzenegger-style processing, the good conscious bit that we're interested in and the unconscious boring bit to do with them saying, overhead projector or whatever it is they're looking at. So as scientists, to understand the bit of activity associated with this, we have to try and figure out how to dissociate these three things. The way we do it is by studying clever illusions. This is an example of one. So, audience participation. There is a large change – there are two pictures here alternating – there is a large change in one of them, some of you will have already noticed. When you notice it, don't shout out, put your hand up. There's a large change in that picture. Yep, I can see it. A very large change between the two pictures. One person's seen it. Two people have see it. Three people – ah, that's good. The latency to detect this is inversely proportional to your [...]. That's a joke of mine. [...] Okay people, come on, come on, come on. More people please. See, people are getting very stressed down here at the front because lots of people at the back have seen it but hardly anyone at the front has seen it. It's a good predictor of how well you're doing at UCL actually – the shorter time you take... That's a joke as well. Okay, you're getting tired, aren't you. Right, it's an aeroplane, isn't it. If you were going on an airplane flight you really would like the engine to be there. It's good, isn't it. It's good. That's a phenomenon called change blindness which works wonderfully well in large lecture theatres. And the mean time – don't feel bad if you didn't see it – but people who did see it, in a large lecture theatre to see it you have to be pointing your eyes just at the change. So in a large lecture theatre full of boards and pre-sessional students, everyone's eyes are pointing in slightly different places. So some people in the lecture theatre will notice it quickly and other people will have to search around, so don't feel bad at the front. Would you like to see another example, redeem yourself? Yes? No? Yes, I think so. Okay, how about this one? Someone's seen it immediately. The front row are doing much better this time, I knew it. [...]. Okay, those who haven't seen it, behind the couple having dinner the bar is moving up and down rather dramatically. So these are two examples of the phenomenon called change blindness. What's going on, how did I do that? What wizardry? Well, it's quite simple. All I do is put in a grey screen in between the two pictures. So if you look just more gently, you can see it's picture 1; grey screen; picture 2; grey screen; repeat. Picture 1; grey screen; picture 2; grey screen; repeat. That's all it is. You can do it at home, you can make movies like this. Amaze your friends, astound your enemies. How does it work? Well, if I'd just alternated those two pictures without the grey screen – picture 1; picture 2; picture 1; picture 2 – I guarantee all of you would have put you hands up straight away. We think the visual system probably operates like this: it probably has a lot of dumb little bits of your brain, each one looking at a different bit of the scene, and to notice change, the visual detectors shout out, 'look over here! look over here! something's changing here'. And then our attention is directed to that change very quickly and we see it. And we think that's a general explanation for how the visual system notices change. A lot of dumb things – detectors – early in the visual sytem will grab our attention and [...]. That works very well if we have just a few things changing, like if I alternated the two pictures here. The only thing that would be changing would be that jet engine. So the dumb detectors there would say, 'hey, something's happening over here, have a look' and ti would direct our attention. Now, when I put a grey screen in there what's happening is the whole of the picture is changing, everything is changing from something to grey and then back again. [...] So even a bit like this is going green then grey then green. The bit here is going grey then another grey then another grey. So all over the image, feature detectors are shouting out, 'hey, something's changing here; something's changing here' so we don't know where to look and we have to now search around voluntarily to try and find the bit that's changing and then [...]. So that's change blindness. Why's it useful? It's just a funny little illusion, but one that's useful for the following reason. This stimulus now, everyone in the audience is conscious of a particular type of content – a jet engine, [...]. But nothing has changed about this stimulus since I started talking. It is physically exactly the same pattern of light on your eyes as when you were completely unaware of that change. So if we compare the brain activity before you're aware and after you were aware, the unconscious bits of processing are all exactly the same, so the new thing is the other half, the consciousness bit. So what we've done in terms of this diagram is keep all of this constant and make just this change. And that's one way we can get at the subjective nature of consciousness. Here's another way. This is a more complicated thing called binocular rivalry but I can do a demonstration of this. If you have a piece of paper in front of you, take the piece of paper and roll it up. So - I don't have a piece of paper [..] I love the sound of people rolling bits of paper up. I once did this in Korea [...] it's a nice thing to do. Now, put your hand over the tube like this. So with your left hand, hold the tube in your right hand, put your left hand next to it, now, put the end of the tube up to your right eye, keep both eyes open – like this. Now, look at me. And we're looking at [...]. Look at me, look at me, look at me – you're looking at a hole in your hand, right? Yes, isn't that good? Now those of you who haven't got it right, follow your neighbour. Okay, look around, there's a hole in your hand, yes, it's all very good. Isn't that great. Right, I've lost you for the rest of the lecture now, everyone's going to be looking at the roof through holes in your hand, things like that. Okay, settle down now [...]. Right, I think that's the most successful demo ever. Okay, so what's happened there? What's actually happening? Why do you get a hole in your hand? There isn't a hole in your hand, is there? What was happening is what I – we were – that's a way of getting the two eyes to look at different things, so the same patch of your eye – one eye was looking at this bit in your hand, the other bit was looking down the tube at my head. And what's surprising is that what you saw wasn't a mixture. You didn't see my head on your hand. You saw a hole in your hand through which the head was. So that bit of your palm must have been supressed, must have been made unconscious. There must somehow be competition in that little bit of the visual field. The hand or head? Hand or head? And I won out, you saw me, which is good, I'm the lecturer. But that's an example of what we would call binocular rivalry. There is rivalry between these two alternative interpretations. Head / hand. Head / hand. And in that particular region where there's rivalry, the head won. Now, it turns out, you can produce that experimentally in a laboratory by putting mirrors in front of the eyes and so the vision goes this way and that way. And out here, you can put two different visual stimuli. For example, a face and a grating: a series of horizontal lines [...]. In that situation, what you experience is rather like what you just saw in the demo but, as time goes on, you experience alternations between the two. It flips back and forth. With the hole in the hand it doesn't work like that, and the reason it doesn't work like that is technical. It's to do with distance, the hand being very close and the head being very far away. If we'd got the distances more equal it would rival – back and forth. Now again, that visual illusion's been known for a hundred years by psychologists, but why it's very exciting to people like me is – here's a situation where we can produce exactly the same thing going in the eyes at all times, the physical stimulation is constant, it's not changing at all. But, the contents of your consciousness flips back and forth. So if we record brain activity, things that change when your consciousness flips back and forth must be things to do with the contents of your consciousness, and not to do with the Arnold Schwarzenegger-style analysis of these images. So it's another axample of what we call an experimental paradigm, a thing that we show subjects in the scanner that we use to try and understand the contents of consciousness. And here are the basic sort-of results that have been established by my group and others across the world over the last ten years. There are a whole family of things that rival in fact, here's another example which you might like. This is an example which is known as the Rubin – after the person who invented it, Rubin – face-vase illusion. This is a vase – in fact it's a real three-dimensional object, it's actually owned by Queen Elizabeth II which is kind of nice. If you look carefully – if you think about this, there's not a hole in the vase, but there's two noses looking at each other. You will see that there appears to be a sillhouette of a face in it – here's the nose, here's the mouth, here's the forehead. Here's the nose, forehead [...]. And it's owned by Queen Elizabeth II, see can see that is Queen

Elizabeth and her husband Prince Philip looking at each other, which is kind of nice. But the neuroscience significance of that is you can't see both at the same time. If you stare at that, sometimes you'll see a vase, sometimes you'll see two faces. You'll always know there's a vase there but you can't see both at the same time. It's an example of bistable perception. Bistable means there's two different things you can see at any one time. And that's similar to binocular rivalry in terms of there being competition: two different ways of interpreting the world. It turns out that the things – the parts of the brain that change their activity in association with those changes in the contents of consciousness are not just the visual areas that I talked about earlier in the back of the brain, but these areas in what are called parietal and pre-frontal cortex. These are areas outside the normal areas used in human vision that we think are something to do with making things – bringing things into your awareness and making you say, 'aha! it's an overhead projector' or 'aha! it's Gareth giving a lecture'. That kind of thing.

Okay, so final five minutes, we're running a little bit over because of the exciting demonstration. Just a couple of more recent results about how we're going forward with this. Here are two examples of more recent work. What we've become interested in is not just studying how consciousness or mental states are encoded [...], but going the other way: instead of showing someone a stimulus and trying to measure the brain activity and figure out what's happening, we're trying to go the other way, trying to take someone's brain activity and predict what they're thinking. So to decode rather than encode. And we started simply and we've shown two things quite recently that we think are quite exciting. One is that if we show you an object, a simple object, we can predict from your brain activity what its identity is. But we can do that even when you can't tell us about it because it's invisible – I'll explain that in a second. And the second thing we've done is a very simple form of mind-reading, showing that in that binocular rivalry where your consciousness changes back and forth, we can actually measure your brain activity; we can actually predict what you're experiencing over periods of time up to about five minutes.

So the first example, we show you a very simple visual stimulus like this. This is a visual stimulus where you fixate – look – at this red spot here, and this is what's known as a grating. It's a series of lines that have a particular orientation, in this case right-tilting. And we can compare activity in the brain when you look at a grating like this or when you look at a grating like this. We know from all these studies with monkeys that cells very early – neurons very early in your visual system respond differently to these two different things. But you recall this diagram where I showed you the responses in early visual cortex to say right and left tilted gratings: they're very finely intermingled at a resolution too fine for our functional imaging device. What we found though, was that we could measure the pattern of activity across many functional imaging voxels and accurately predict the orientation of the subject's view. More excitingly – I'm rushing through the first bit so I can tell you the more exciting bit – we found you can also predict which of these two orientations was in front of the subject's eyes when even the subject themselves couldn't. What we did was we masked the stimuli. Masking is a way of making stimuli invisible or subliminal. Maybe some of you have thought about advertising sometimes and in many countries in the world subliminal advertising is illegal. Subliminal advertising is where we're watching a movie of something, [...] and suddenly there's a flash, a frame, that says, 'buy Persil washing powder' or something like that, and then the movie goes on. And it's just a single frame, so it's very quick. It's called subliminal because it's below our awareness. If you ask most people at the end of the movie, 'did you see anything unusual?', they'll say, 'no'. They're not aware of it. But of course, the reason why it's illegal in many countries is because a lot of people will be influenced by that advert to go out and buy Persil, they'll say, 'ah, what shall I buy, what washing powder, Persil or Ariel? Oh, I'll take Persil.' Any you'll ask them why and they'll say, 'I don't know why, I just prefer Persil'. Now actually, the evidence for that is quite poor, but that's the idea in people's minds why unconscious stimuli can influence our behaviour. So here's a very simple scientific way to study it. So what we do is we modify – remember these are gratings that have orientation – we modify them a little bit by making them dashed rather than straight lines. And this is a – the blown up bits you can see there are dashes. And we put those dashes in so we can alternate the target with a mask. The mask is a criss-crossy pattern, it's the opposite polarity, that means it's white on black, not black on white, and it fits neatly into the gaps in this one. Now if we alternate target and mask very quickly: mask; target; mask; target; mask; target; mask; target; mask; target – like that – what the subject experiences is just a flashing criss-crossy pattern and they don't see this target. We know they don't see it because we can ask them after fifteen seconds. We say, 'well, we were actually presenting a target which either had a right tilt or a left tilt. Which one was it?' And when they press a button to indicate that, the responses are chance – fifty percent, just as if they were guessing. So we know they're not aware of the orientation. But what's exciting is our brain imaging device, using activity from the visual cortex, can do better than they can. So remember, they've got information coming into their eyes, it produces activity in their brain. And we asked them, using that activity in your brain, 'can you guess whether it was right or left?' And they can't. However, we have the same brain activity, now measured with our MRI machine, and the question is can we guess which was being presented? Yes, we can. We can predict significantly above chance. This is showing the size of the pattern in their brain – number of voxels – and how well we do at predicting. And we do, about sixty or seventy percent of the time, get it right. Where they're getting it wrong all the time. So put in simple language, we can tell more about what's in front of people's eyes from measuring their brain activity than they themselves can. And what we can conclude is there must be an unconscious subliminal representation of that stimulus in their brain. So, we can say that the information about the features of invisible visual stimuli must be present in these very early stages of visual processing. And we've gone on - we are going on to use that to study unconscious representations much more generally in the brain. So more complicated representations of more complicated things like faces or houses, or emotions or complicated things like that.

A final thing about mind-reading then and then we'll finish because we've run out of time. Mind reading, I thought the best thing about this decoding approach was perhaps you could read people's minds. And I was was very excited and so I went to the – as often when I'm very excited, I went to the Oxford English Dictionary. The Oxford English Dictionary, if you've ever seen it, it's very very very big. And contains every single word in the English language and when it dates to. And I thought mind-reading would date back hundreds of years as a word. It doesn't, it dates to 1863. And it dates to this person: does anyone know who that person is? Fantastic, a star, he said Edgar Allan Poe, and that's exactly right. Edgar Allan Poe wrote a book called The Murders in the Rue Morgue, this is an extract here from The Murders in the Rue Morgue. And the interesting thing about this book also is he introduced the idea – he introduced the world's first fictional detective. So he was the first to write a detective story and this was about a Frenchman called C. Auguste Dupin. And this text depicts Dupin, being a detective, he's reading the mind of the person next to him. But he's doing that because he knows something about the context: who that person is; what they're saying. He makes inferences, guesses, hunches about that and because he's a good detective he does it very well and appears to be mind-reading. The question for us though is whether we can mind-read in a much more direct way by just looking at brain activity. Not by [...] or hunches, but can we mind-read? We chose to do this to start off with in a very simple situation. Wouldn't it be great to do an experiment where I put you in a scanner, I say, 'think of a card. Any card.' And I process your brain activity and I say, 'ah, you were thinking of the seven of clubs, weren't you', or the ten of hearts or whatever. That has a lot of different possibilities so we can't do that right now. What we can do, though, is mind-read a very simple two-stage system, and that's binocular rivalry like I showed you earlier. We show you either this red thing or this blue thing, the red thing to the left eye, the blue thing to the right eye, and rotating as well. If we ask you to tell us what you see, red or blue, you'll see red for a bit then blue then red then blue and so on. And over about three minutes you see alterations between the two. What we can do it track very accurately from patterns of activity in your brain, without knowing what you're pressing – red or blue – without knowing that, we can decode and track these alterations over time. So here again over three minutes are displayed from three subjects are decoder – this kind-of solid line – this is our prediction. The dotted line is what subjects reported they were experiencing. We didn't have that information. We try to decode it and we overlay the two to see how well we did. And as you can see from those graphs, we're doing pretty well. These dotted and solid lines are overlapping almost all the time for this subject. There are a few errors here and there are a few errors here, but overall the accuracy is about 80 or 90% accurate. Which we're excited about because that represents a great start to the idea of directly reading minds. And that's of relevance to a whole load of areas. Not just understanding consciousness, which is my interest, but it's also an area that's very hot scientifically, in terms of decoding people's thoughts about intentions and movement. For example, the journal Nature, last week, had two articles about implanting chips in the motor cortex of paralysed patients, which was trying to decode their movement intentions help them move. What we're doing is at the other end of the system, not movement, but perception. But we're fundamentally in nature the same kind of enterprise, we're trying to decode people's thoughts and actions from their brain activity. So that's why this might be important, this kind of decoding is a new approach to identifying how information is encoded in the brain. And it has practical value potentially, as well as the scientific value.

So, that's really it as far as what I want to say to you. If you're interested in finding out more, here are three ways to find out more. You can go to the website of the functional imaging laboratory at UCL. You can go to my group's webpages and read stuff, and we've written a recent review for those of you who are technically minded, which is available in Nature Neuroscience [...]. I'll pop that side up again at the end of the lecture and take a few questions, if anyone wants to dash off. But just to remind you as well, although I told you about the work of my group, science is a collaborative activity, it takes a lot of people, and these are some of the individuals, the phd students, clinical fellows and postdoctoral fellows in my group who are collaborators. And the Wellcome Trust fund it. Just to remind you, it takes a lot of people to do science, and we all work together and also have a good time, and enjoy what we do, 'cause I think that's fundamentally why we're here. Life is short – you've only got one of them and you've got to enjoy it and do the best you can, that's my view anyway. And have a nice work-life balance and enjoy the science. That's why I like being at UCL, that's why I like being here and I hope you all have a wonderful time here too. Thank you very much.

Yeah, questions anyone? Anyone got any questions? Or do you just want to go out and get an ice-cream. Or a cigarette. You shouldn't smoke, it's bad for you. Yes, there's a question up there.

[question]

Yep, the question is, is there anything called a subconscious mind? I think the problem here is terminology. People use different words to refer to different – either subconscious, unconscious, they're kind of the same thing. I think the big mistake is actually believing that Freud's correct. So, Freud was very significant at the start of the twentieth century by introducing the idea that there could be an unconscious at all. But he went the whole hog, he said each one of us has kind-of an unconscious id, a primeval, subconscious thing that just wants to have sex with lots of people and eat huge amounts of ice-cream and become very fat, and do all of those kind of uncontrolled things that are completely inappropriate in a lecture, or indeed in areas of professional life, and that we have to keep that under control with our conscious mind, and our whole life is doing that. That's a rather different idea of the unconscious than the one be believe nowadays scientifically, where the unconscious isn't like a sort of beast living within us, with its own mind and ideas. It's more just a collection of stuff: how we process objects or sometimes unconscious feelings, I guess, but nothing as complicated as beliefs and desires and thoughts [...] – I don't think – I hope not. Any more for any more? Right, well thank you very much indeed, and I wish you all the best for [...].