



GRESHAM COLLEGE
Founded 1597

Future Brain Transcript

Date: Thursday, 17 March 2011 - 6:00PM

Location: Barnard's Inn Hall



17 March 2011

Future Brain

Professor Keith Kendrick

Good evening ladies and gentlemen, and welcome to Gresham College. This is the third in a series of lectures that I have been giving, looking at fundamental and also clinical advances in neuroscience, aiming to bring you as up-to-date as possible. This lecture differs from the other two in that it is much more speculative. I also hope you will find it that little bit more humorous.

“Future Brain” includes an element of crystal ball gazing. I shall be considering what may or may not happen over the next twenty-five years. It is a fairly arbitrary time period, but if you go much beyond that, you would be foolish to make any reasonable prediction as to what may or may not happen in terms of discovery; the field is moving incredibly fast, and that is really what I want to try and give you a flavour of. I do not know how long it will take for some of the things in this lecture to take place, but they will, and we are going to develop a very different understanding of the brain over the next few decades, which will change our lives.

I could talk about many different areas of neuroscience and we would be here for quite a long time. This actually happens to be my 30th Gresham lecture, which suggests that there is an awful lot going on in neuroscience. In this lecture, I am just going to focus on a number of areas – not so much where the next super-magic pills are coming from, but rather to give you a better understanding of how we are going to reach that stage, of how both healthy and diseased brains function, of new ways to modify brain function for therapeutic purposes (particularly non-invasive ones), and the use of the brain to control external devices – brain machine interfaces as they are often called. I will avoid giving you a lot of mathematical modelling because it is likely to get extremely complex, but nevertheless, underpinning a lot of neuroscience advance in the next few decades will be an important interaction between neuroscientists, mathematicians and computational scientists. Indeed, I will also consider an area that is not as “science fictional” as you might think – the design of brain-like computers.

I shall start with recording brain activity changes. Clearly, if we really want to understand more about how the brain works, we have to be able to record its activity and try to make sense of what it is doing, at quite a detailed level. We are now beginning to develop the kinds of technologies that allow us to really interrogate brain activity, and which give us the opportunity to begin to understand how such a complex organ is functioning.

I am sure most of you have heard of brain imaging. This is a magnetic resonance imager, through which we can carry out studies and look at indices of brain activity changes at increasingly high resolution. This slide shows brain activity using a massive magnet – in this case, a 7 tesla magnet, which is becoming pretty routine now in MRI scans, but in fact there are magnets up to 12 tesla. They are not yet used on humans – there are some issues about whether these high magnetic fields might have effects on the human brain. However, the advantage of these very high magnetic field strength machines is that you can get down to millimetre resolution within the brain, or even less, which allows us to tease out changes in quite complex groups of nuclei. This is the case, for example, in the amygdala, where the nuclei themselves can be quite small, but they all do slightly different things.

Another major advance is that, up until recently, these scans have very much been an offline analysis of brain activity. You took the scans and then did averaging, and you worked out, after the person had had their scan, what was going on in their brain. Now, it is actually possible to put somebody in an MRI scanner and get real-time feedback, which of course allows you to instruct the subject to do things, or for them to try and control their own brain activity – which is something I shall return to in a minute.

Magnetoencephalography is another possible method, through which you are again measuring magnetic fields. In this case, however, you are measuring cortical fields at a high density, much higher than with normal electroencephalography. Nonetheless, both methods require the patient to remain totally still within highly expensive pieces of machinery. This is one of the big drawbacks to these kinds of technologies, the restriction to design experiments where a human subject can be sitting inside one of these machines but they have to be totally still, at least with their head. They can do things with their hands and they can speak, but their head has to be still, and that carries a certain amount of limitations.

This has meant that, particularly for the development of brain machine interfaces, a less high-resolution methodology has been developed. I have called this section of my lecture “Shining Light on the Brain” because these technologies use light to look at and even influence brain activity. This is near-infrared spectroscopy, where you are either measuring activity in terms of thermal or haemoglobin changes. You are shining lots of light of infrared wavelength on the brain in order to image general brain activity changes, not anything like as acutely or accurately as you can with fMRI or MEG, but sufficiently enough so that an individual does not have to remain totally still. A Japanese company is now producing special headgear where you can carry out this kind of recording with somebody moving around.

Light is also playing a very big role in fundamental neuroscience, perhaps one of the so-called hottest areas of neuroscience at the minute. It is much talked about and very interesting, but I am not going to spend a lot of time on it. The technology is called optogenetics and involves inserting opns or channel opns into the brain,

light-activated molecule genes that you can attach to the membrane of cells. You use a viral vector to get them into the brain, and then when you shine a light onto neurons that are expressing these opcin genes, you can switch them on or off. Down the line, you can also use this same light-activated change in neurons to switch on and off genes that you have put in there.

This is potentially a much more powerful technique than mere electrical stimulation and allows you to stimulate quite wide areas of the brain, which is obviously not possible with a single wire electrode. The only real disadvantage, at the minute, is that it uses fairly standard wavelengths of light and because the light source needs to directly activate these opcins, it is fairly invasive. Perhaps in the future, we will be able to find ways of shining light down into the brain that do not involve directly probing brain tissue.

Here are some examples of where we might be going. In my last lecture, I talked about the increase in deep-brain stimulation research for the treatment of Parkinson's, but also depression. This is clearly a very invasive technique, but we are beginning to understand which parts of the brain you need to stimulate in order to overcome some of the symptoms of these disorders. It would be nice if we were able to do this in a much more non-invasive manner. There are already speculative technologies which target brain areas using infrared pulse laser, focused pulse microwaves or focused pulse ultrasound - ways of stimulating deep-brain structures physically without putting electrodes into the brain.

So, what about being able to read what you are seeing, thinking, feeling from recording the activity of the brain? This is clearly an area that disturbs people. The cartoon in this slide shows a young lady walking down the steps at Wimbledon and the thoughts of those around her. You could argue that we do not really want to know these kinds of things - it would open up a whole can of worms. Nevertheless, we are making sufficient advances in reading brain activity profiles and analysing them in order to start understanding what people are thinking and feeling. In my first lecture, I actually gave an example of a recent study which allowed individuals in a vegetative and apparently unconscious state to communicate with individuals in the outside world. This study used the very simple idea that you get different patterns of activation depending on the things you are imagining - for example, imagining walking through your house requires more spatial imagery than imagining playing tennis. You get completely different patterns of activation. The study found that a few of these patients were actually able to do this, even though they could not communicate in any other way. Based on this discovery, they were then able to ask the patients simple, yes or no questions, in which "yes" would, for example, activate motor imagery, and "no" the spatial imagery. This is obviously a very crude way of interpreting what is going on in the brain, but the patterns of activation that emerge can be very well-demarcated.

In 2005, a much more controversial paper was published, claiming that you could detect people's lies from their fMRI scans. This is a real-life example of Pinocchio's nose, in which one pattern shows that you are guilty, while another pattern shows you to be innocent. This is from the paper by Langleben *et al* in *Human Brain Mapping*, and they claimed to be able to successfully detect deception in people 80-85% of the time.

It has been used several times in court. It has never actually become admissible as evidence, but, as time goes by, it is likely that the technology will improve, and it is not beyond the realms of possibility to imagine that, in ten or twenty years, the lie detector or polygraph, will have been replaced by routine brain scans as a way of determining whether someone is telling the truth or not.

What about knowing what somebody is looking at? Kay *et al* (2008) imaged the visual cortex of brains processing lots of pictures - anything up to 1,000 different pictures in the sets they were giving - and they used various methods to analyse the resulting brain scans with the aim of accurately predicting which picture the subjects were looking at. For sets of up to 200 and 300 pictures, they achieved a 70-80% success rate, and this is sure to improve over time. It is already fairly amazing to be able to detect the correct picture out of quite a large set.

You can certainly envisage this future scenario, where you have forgotten your car keys and you know the location is stored somewhere in your head, and you simply go along for a scan to retrieve this information - "Your car keys are under the sofa." However, you might also expose information that you did not want people to know - as is the case with the gentleman in this cartoon, who seemingly likes to wear women's shoes.

So what about learning to control or heal our own brain? Clearly, there is a huge amount of interest in the potential for self-healing, and there are advances occurring very rapidly now which make this into a real future possibility.

A recent study by Caria *et al* (2010) imaged the right insular cortex, which is particularly important for responding to negative emotive stimuli, such as angry faces or disgusted faces. The subjects were shown these faces and then given feedback, or were told to actually control the levels of their activity. This was a real-time fMRI study, for patients to try and boost the activity of their own cingulate cortex, so they actually look at an indicator of how successful they are in increasing the activity of their right insular. As you increase the BOLD activation in the MRI, the patients report even more negative responses towards these emotive pictures. You are actually capable of changing the activity of one of these areas so important in controlling emotional responses, and with feedback you are able to change your emotional response.

This brings us to metacognitive awareness, which is even more complicated. This means an awareness of your

thoughts, attending to them and turning them off and on – an inability to do so is clearly marked in some psychiatric disorders. McCaig *et al* (2011) carried out a study into this, again using real-time fMRI and feedback, causing an increase in the activity of the frontal cortex across training sessions – but there is no increase when feedback is not given. As a result, it seems that even higher brain functions are susceptible to being affected by feedback training, and real-time fMRI has been very useful in revealing this.

It is not the only approach though. This was a widely reported study by Cerf *et al* in *Nature* last year, carried out on patients being treated for epilepsy. In order to understand what was going on in their brains before surgery, an array of recording electrodes was implanted into the affected area for a week or so, which allowed them to take part in neuroscience-type experiments. In the experiment shown here, the temporal lobe is used.

Patients respond particularly well to high-level images, and in this particular case, images of famous people were used: Venus Williams, Josh Brolin, Marilyn Monroe, and Michael Jackson. It does not matter who they are. They recorded responses from single neurons: in the left parahippocampal cortex, this single neuron responds preferentially to Marilyn Monroe but not Josh Brolin, while the neuron in the right hippocampus responds preferentially to the face of Josh Brolin and not to Marilyn Monroe. The subjects are then given a composite image of both Josh Brolin and Marilyn Monroe to view on their computer, and their task is to increase the activity, by their own thoughts, of the neuron that is responding to Marilyn Monroe. They are given feedback as to how well they are doing by the fact that this composite picture becomes deconvolved, becoming progressively like the target image, which is Marilyn Monroe. As you can see, over a period of trials they become quite capable of doing this. So there is one image that is switching from a composite image to Marilyn Monroe, and on this side it is switching into an image of Josh Brolin. When they are successful, you can see that their activity changes. If they do not actually manage to do it, there are no activity changes. It sounds strange that we should be so interested in this, but the ability to change the activity of single, high-order neurons in the temporal cortex by way of feedback suggests future achievements in altering brain activity for therapeutic purposes.

One of the areas of neuroscience that is particularly exciting at the moment is the brain's various electrical, oscillatory rhythms. It is an area that has been known about for a long time, but we are now beginning to get a handle on understanding exactly what these rhythms are doing. They are performing extremely important functions of coordinating activity, either locally or across wide areas of the brain, and there are different rhythms in different states, from very low frequency rhythms like delta (around one or two hertz), up to a high frequency rhythm like gamma (from about 30 to 70 hertz, or even higher). Sometimes, these high frequency rhythms are found coupled with low frequency ones, so they seem to be working together.

A lot of research has shown how these rhythms are changed during learning and under various conditions. Here is some of my own research, carried out on sheep. In this study I looked at a fairly low frequency rhythm, theta, and a high frequency rhythm, gamma. We know that gamma is coupled with theta. The sheep were subjected to a face discrimination trial and these graphs plot the changes that occurred in theta/gamma coupling, the ratio between theta and gamma, the amplitude of gamma, and also of theta. You will notice that, as the animal learns (at 50 %, they do not know what they are doing, whereas by 80%, they are choosing one picture over another to get a food reward), there is a positive correlation between performance and the changes in these aspects of oscillatory rhythms. This is represented by the pseudo colour charts on the right of this slide. As soon as the animal actually learns - as they go from 70% up to 95% - you get sudden increases in theta amplitude and theta/gamma coupling, and you also start to see changes in the ratio of theta to gamma. These correlations strengthen as learning progresses.

We have been able to model this kind of change in a formal way, involving very simple circuits. We have also been able to model how these are affected by learning, using the molecular participants in learning - particularly the glutamate receptor, the NMDA receptor. This is the kind of approach that is very much going to be important in the future - looking at biological information, discovering effects, and then generating models in order to formally describe what is going on, as well as predicting things that you cannot necessarily do with recording experiments. These models will be the building blocks used to train artificial devices to work like brains.

This is another study showing learning. This is recording either high frequency or low frequency oscillatory activity, but just plotting it in terms of levels of activation, shown by the pseudo colour scale here. This is a very simple motor or movement learning task, and I am sure many of you will know, from training in sport, that actually imagining yourself doing activities plays a very important part of training. If you look at the activity patterns you get in motor areas of the brain when you are imagining a particular set of movements, they are not usually as strong as the real movements, but they nonetheless follow the same pattern and with almost the same level of activation. However, if you give the individuals feedback as well, you can actually improve the level of activity in both the high frequency and the low frequency bands, beyond that which you would get with normal movement. So you can improve, with feedback, these kinds of changes in oscillations as well.

Uhlhaas and Singer (2010) report some interesting findings, which pose more questions than answers but also show why looking at these kinds of oscillations is going to be important for us down the line. They have been looking at the changes in gamma (high frequency) power and also beta (slightly lower frequency) power, and they have found that during development there are changes in gamma power such that, particularly when you reach adulthood, there is an increase in gamma power and an increase in the synchronisation of beta. In

general, what they are showing is that, as you go through adolescence, there is an increased synchronicity in these brain oscillations. As I am sure you know, a lot of psychiatric disorders tend to start off in adolescence, and it is notable that a lot of these disorders are characterised by changes in oscillatory rhythms.

This is an example from the same group. You can see that there is a reduction in gamma power in patients with schizophrenia and a change in the beta frequency as well. There are also examples of this from a number of other disorders, like depression.

This whole field has led to the emergence of a kind of fringe medicine, although those who practise it would obviously not label it as such. A number of papers have been published – unfortunately not often in major journals – about this field, focussing attention on attention deficit hyperactivity disorder as well as autism. For example, studies have observed measurable changes in high and low frequency rhythms in the activity of the brains of children with ADHD. They have shown how by simply recording activity from a very few sites on the scalp, you can pick up these rhythms. You can then give these children a feedback task where they learn to increase or decrease the activities of these particular rhythms. Swingle (2008) used a DVD of *Toy Story* as a feedback task: when the children get it right, the DVD plays, when they don't get it right, it stops. Over about twenty-five, twenty-six sessions, it is claimed that they can actually control these brain rhythms and improvement in behaviour has been reported. It is emphasised, however, by all of the people practising this, that you need to do other things as well; it is not enough to rely completely on this feedback task. Perhaps this kind of bio-feedback task may well prove very useful in treating various disorders down the line, but we really do need to understand a bit more about what is going on, and maybe this is a rather simple version of doing it.

Many eminent people who are experienced in treating ADHD claim that this is all a placebo effect, so it is quite controversial. However I hope, from the kinds of evidence that I have shown you already, that you can see how it might not be as controversial as it first appears. In the future, once we understand more about what is altering these brain rhythms, then the basis for altering them, and hopefully affecting behaviour, will be somewhat more accepted.

I show this slide to convince you how amazingly adaptive the brain is. This is a CAT scan in which the person's skull is almost empty. This individual has virtually no brain – the brain is right on the periphery here, just a couple of centimetres under the skull. This individual is actually a French civil servant and does have a low IQ, but still manages to survive. This was caused by hydrocephalus, but because it happened over such a long period of time, the brain was capable of adapting and allowing them to lead a relatively normal life. In general, if you give the brain enough time to adapt to various insults to it, then it can overcome dysfunction quite remarkably. Unfortunately, in most instances of brain damage or deterioration, the time span is too quick for us to be able to adapt properly.

Although there is a huge amount of talk about the therapeutic use of neuronal stem cells in treating a large number of disorders, most of these studies involve inserting stem cells into various regions of the brain to cause regeneration - for example, to cause regeneration of dopaminergic neurons in patients with Parkinson's. But, in fact, the brain has its own neural stem cell population, and there are a couple of structures that are involved in various aspects of learning (such as the dentate gyrus) and odour recognition (the olfactory bulb) that do actually show routine evidence for what we call neuro-genesis, the production of new neurons. That seems to be very important within these structures, for the processes of learning. But there are also large numbers of neural stem cells in what is called the sub-ventricular zone, which is the lining of the cerebral ventricles, and when the brain is damaged these actively track into that area of the brain and attempt to repair it. Of course, their actions are in most cases insufficient when you have large amounts of brain damage, but there is huge interest in the possibility of boosting the activity of these self-repair systems within the brain, so that they could react to damage to various structures in a more effective way. So this is likely to be another big area of advance in the next twenty-five years, trying to control the activity of self-repair mechanisms within the brain to help deal with progressive damage particularly.

What about using the brain to control external devices? We are seeing increasing amounts of evidence of being able to record the kind of brain activity that I have been talking about in order to control prosthetic limbs, for example, but also completely other devices, like robots.

In this slide, the woman has been fitted with a robotic arm, which she controls through links to nerves in her chest, which she can obviously control with her thoughts. This kind of advance is picking up speed rapidly now, particularly with the treatment of amputees.

More controversially, Guenther *et al* (2009) have published a study on one particular individual who is mute. They have implanted electrodes into the speech areas of the brain, recorded the activity of these and passed it through various analytical systems to drive a speech synthesiser. The subject was given feedback from the synthesiser, hearing what they were producing. This individual was therefore able to produce recognisable speech through changing the activity of the speech centres of their brain, which is quite remarkable. *PLoS ONE* is perhaps not the best journal in the world, but this study almost beggars belief and it is certainly the kind of thing that I expect to see happening with neuro-prosthesis over the next twenty-five years.

Indeed, there was a very big review in the annual reviews of Psychology just last year, looking at the different things we might be able to do with individuals such as paraplegics in order to use recordings from their brains

(from neural infrared spectroscopy as well as from a standard EEG) to control a number of external devices, including the TV.

Of course, there is a lot of work going on in developing brain machine interfaces which are not just therapeutic. For example, individual brain waves are recorded by neural infrared spectroscopy in order to control the Honda robot, and quite successfully. This is very much an area where a lot of effort is being invested. It is still relatively early days, but nevertheless, it is possible to control very complicated robots using input from human brain waves.

In fact, there is an app for it already! iPhone is selling something called "X-Wave", where you can control the iPhone by recording electrical activity from your brain. I have no idea how successful it is, but if you want to flash out \$100, you can find out for yourself. Obviously, this is a rather simplistic approach, but clearly it has attracted a lot of attention already.

So, can we construct new brain parts or even an artificial brain? Now we are very much approaching the "science fictional." Could we really insert a sense of humour into those lacking it, or help someone with tone-deafness? Could give someone's brain some "spare parts"? These ideas are often called neural-morphic approaches, and there are essentially three basic approaches that you can use. The first is to make existing computer technology simulate brains, and there is quite a lot of work being done towards that. Alternatively, you could say that existing technology is not sufficient to emulate a brain, so we need to develop new technologies using more brain-like components. This is the first kind of approach.

The second approach is very important. We will need to be able to integrate these kind of chips into the brain, so hybrid systems using neuronal and *in silico* components are being created in order to interface between brain and computer. This is already being implemented and I have already shown you an example. You can see this approach at work in cochlea and retina implants, where you are inserting electrode arrays in order to record the activity of nerves, and then translating those into digital devices.

The third – and perhaps most disturbing – approach is to make an organic-type brain, either based on current biological principles – neurons, synapses and so forth – or perhaps other novel approaches, and the potential for this lies in polymer chemistry and molecular physics.

Henry Markram is Director of the Blue Brain Project in Lausanne, Switzerland. Some five or so years ago, he set up a unit out there with access to a massive IBM super-computer. Their aim is to model and simulate very basic building blocks of the cortex of the brain. So far, they are making progress; they have got past first principles and there is certainly a huge amount of modelling going on. We shall have to wait and see what the project produces.

Many of you will probably know far more than I do about computers, but one of the big differences between computers and neuronal based systems is that computers process information first and then physically shift it off to separate registers. This slows things down and also increases the amount of space and power that you have to use, and is actually a limitation to current computer technology. Neural networks both process and store information in the same place, or they can do. So that is a big difference between brain-based systems and computers, which allows, computationally, the brain to do, in a relatively small size, as much as a super-computer that would fill two of these rooms and would actually consume probably a million times more power.

In America, the Defence Advance Research Projects Agency is currently spending millions of dollars on this particular scheme. It is called Systems of Neuromorphic Adaptive Plastic Scalable Electronics, or SYNAPSE for short – they obviously chose the acronym very carefully! This programme is aimed to develop brain inspired electronic chips that mimic the function, size, and power consumption of a biological cortex, and multimillion dollar grants have already been awarded to IBM, Hughes Research Labs, and Hewlett-Packard. They are working together with academic partners, but these are the companies who are spearheading this. As you can imagine, this is potentially a multimillion dollar commercial gain if they are successful.

Advances in this field have been helped by the discovery of a new electronic component. We have known about resistors, capacitors and inductors for a long time, but in 1971, Leon Chua made a theoretical prediction that there must be another form of component – what he describes as a resistor with memory, a memristor. It is a device that can actually change its length according to the amount of current that is being passed through it, and therefore change its resistance. It is also sensitive to the direction the current goes. The memristor has the amazing ability to remember what charge has been passed through it and alter its resistance, and the important thing is that it remembers that resistance even when the power is turned off. So it is effectively a non-volatile component, and in theory this kind of device could improve flash drives and so forth. Chua stated that memristors must exist in principle, but it was a group at Hewlett-Packard who published a letter in *Nature* in 2008 with the first evidence that these things could be constructed. About a year after that, they showed that you could build chips with memristors that could process and also memorise information, and they were four to eight times smaller than current chips.

This is interesting enough from the point of view of the computer industry, but the idea of an electronic component with memory is of course exactly what a neuron and synapse are. Research as recent as last year has constructed these switch-point apparatus chips using memristors, and shown that they mimic the activities

of neurons during basic aspects of learning, (for example, neuroplasticity changes). It is still very early days, but in principle these would be the sorts of components that you would use in order to build a chip that functions just like a brain.

Incidentally, there is a huge amount of concern, particularly in the US, that the CIA is going to be developing mind control chips. Addressing a recent conference of the All American United National Association of Foil Hat Wearers Against CIA Mind Control Rays (AAUNAFHWACMR), the Public Relations Director of the CIA, assured those gathered, "Although it is true that only small electrical charges are needed to stimulate the brain, that is not the whole story. To actually change the average American mind on anything at all would require at least the voltage of a car battery, and because of the size and weight issues involved, that is just not practical." At least he had a sense of humour!

It does not stop people like the Foil Hat Wearers worrying about the development of mind-control chips. The technology, in theory, is not completely out of reach, but it would be a long time before we got there.

So what about organic brains? Could we actually create anything that functions similar to our current biology? Yes, we already can. You can obviously culture the neuronal stem cells, and you can make them into neural circuits, but, in a sense, it is going to take a long time to develop something that will be performing any very complex functions. On the other hand, you could conceive of developing organic-based, but nevertheless nanotechnology-based, brains. To give you an example: the basic unit is a nerve cell, with dendrites and axon, and you can generate that with a gold nano-particle bound to conducting polymer strands. The axon, with different sizes, can be conducting polymer strands. Receiving dendrites conducting polymer strands, ending in a gold particle, for conduction, and the synapse can be the interception between polymers and electrolytes. You can effectively duplicate (at the nano-scale) something that is constructed like a neural network with existing technologies in molecular physics and polymer chemistry. That is the kind of thing that you could come up with. One of the advantages of these kinds of networks is that you can actually have bio-directional communication, which, of course, in neurons you cannot - it is only one direction. The axon will only normally propagate in one direction when it is functioning.

In the future then, when you want a new brain, it might just be possible. However, I believe that the key advances in neuroscience will not lie in this direction. They will be in linking up brain-like chips to our existing brain circuit, so the interface between brain-like chips and our brains will be the key area of advance, because you can obviously programme chips with a huge amount of information. If you could actually get that information read out by the brain, then you could almost imagine having a sort of internal Google, information that you just plug in.

Clearly, in light of the issues that I have been raising in this lecture, it will not surprise you to know that there has been quite an increase in the area of neural ethics. After all, we are starting to consider the implications of understanding and controlling human organisms that are unprecedented. When reading and controlling the thoughts of others, are you guilty or not? Furthermore, although most of the examples I have given have been within a therapeutic context, they could also be used to improve normal and non-impaired functions. How are we going to deal with or even police this kind of thing if we aim towards such a *Robocop*-like scenario?

Moreover, if you believe - as we must do really - that consciousness evolves naturally from complex neuronal based-systems, then in the future we may well have to ask ourselves - will we be generating machines capable of some degree of consciousness and therefore feelings?

Of course, this is not twenty-five years away, but the principles are going to be there, and it will be very interesting to see what happens when we do have the technology to generate large-scale neuronal-type networks, and we are faced with their ramifications.

In summary, advances in brain recording technology and computational approaches will increasingly allow interpretation of mental activity. Look out for the guy on the street corner who is going to be offering budget brain scans! There will be an awful lot of fraud and criminal activity going hand in hand with these developments, but increasingly, we are going to be able to see what people are thinking and feeling, just by recording the activity of their brains.

Of course, our thoughts may be used against us. Perhaps every courtroom will have its witness box installed with a real-time fMRI machine or something similar? Or what about the button that connects to the inner ear of all workers, just to see what they are thinking and feeling?

As I have tried to intimate, I personally believe that the whole area of bio-feedback is going to be extremely important in helping us help ourselves treat damage in dysfunctional brains. It is still something of a "quack" area in many ways, but when the hard science emerges to back it up and we begin to understand more of what we are doing, then it is a very real possibility. After all, the brain is extremely adaptable, and we seem to be able to control it ourselves, which suggests the potential to use that capability in a beneficial way. Of course, not only could you use it for beneficial purposes, but you could also use it for less constructive ones as well - a two-edged sword.

We will be able to control external devices using brain machine interfaces. Perhaps we will have robots that we

are able to control with our own thoughts to do things around the house? I suspect that the first real applications of these brain machine interfaces are likely to be in the gaming area – it is such big business – and the idea that you can control games with your thought patterns carries a certain degree of appeal, as you have already seen from the iPhone application.

We already experience enough domestic strife arguing about who controls the television remote, but imagine if both partners could control devices around the house. Who will exert the dominant control? I say this flippantly, but it raises an important point. If everyone can control machines in their local environment, what happens when two people with differing views and motivations come into conflict? In this case, developments in neuroscience may not be such a step forwards but may actually make things worse.

Computers will become more like brains – will we end up with computer psychiatrists? I certainly hope not, but it is a possibility.

Neural-prosthetics, as I have already intimated, will become far more routinely available and not just for the replacement of artificial limbs, but also for replacing functions that we, for one reason or another, are unable to do. So it is not out of the question that we will have interfaces with complex chips that may indeed function like brains – you can imagine a module for acquiring languages or something to improve on what the brain has, or does not have, including a Google implant for those of us who want to know more than we perhaps really ought. Alternatively, I am sure some of us would prefer an implant that just switched it all off!

For some people, the issues raised in this lecture might be enough to make you want to run away – a case of neuroscience gone crazy. It is not that. It is going to be a really exciting period, both in neuroscience but also in computing, and they will go hand-in-hand. We can only hope that this technology will be used for the benefit of mankind. I can certainly assure you that, for the person giving this lecture in twenty years' time, a lot of these things will actually be fairly routine. Our brains will no longer exist as a private place that we carry around with us – the “future brain” may no longer be safe from intruders, looking to know what we are thinking and feeling.

Thank you very much.

©Professor Keith Kendrick, Gresham College 2011