



Interview: Richard Walker, Human Brain Project

Ivana Greguric

Scientific Center of Excellence for Integrative Bioethics,
Faculty of Humanities and Social Sciences, University of Zagreb
Zagreb Business School

ibanez_ivana@yahoo.com

Journal of Evolution and Technology - Vol. 27 Issue 2 – September 2017 - pgs 4-11

Decoding the human brain is the most fascinating scientific challenge of the twenty-first century. Neuroscientist Henry Markram from the Human Brain Project (HBP) has claimed that it should be possible within ten years to model the human brain in its entirety. He has already simulated elements of a rat brain, and his team is now working on a detailed, functional artificial human brain that could be built as soon as 2020.

In 2015, during the term of a scholarship at the Brocher Foundation, I visited Campus Bioetech in Geneva, where I interviewed the HBP's official spokesperson, Richard Walker. We discussed how AI will change our evolution, and he warned that emerging information and communications technology (ICT) raises significant questions about losing control over the goals of machines.

Special thanks to Mr. Nenad Bunčić for his help in arranging this interview, and for his hospitality and assistance during my visit to Campus Bioetech.

What is the goal of the Human Brain Project?

The goal of the Human Brain Project is to build a completely new ICT infrastructure for future neuroscience, future medicine, and future computing that will catalyze a global collaborative effort to understand the human brain and its diseases – and ultimately to emulate its computational capabilities. By the emulation of the brain we mean that today's computers are very different from the brain: they can do some things much better than the brain, they consume more energy, but they are less reliable. For example, there are a lot of things that a 5-year-old child can do easily that no computer can do. I can show a 5-year-old child a picture and he recognizes it's a cat, or it's a house, which no computer can do. We would like to take our brain circuits, the ones we look at in rats, mice, and ultimately in humans, and put them on chips which can be a new class of computing system. And so this would be a technological product which you can use to do computing.

How do you plan to create a functional artificial human brain, and will you achieve your goal by 2020?

Most of our efforts have been in making a preparation for the human brain, based on rats and mice. This has helped us to produce a draft digital reconstruction of the microcircuitry of the rat neocortex – it is just 30,000 neurons. If you consider that the average human brain has about 100 billion neurons, it's a huge difference. But this is a significant accomplishment, because it demonstrates to us that it is possible to make a successful digital approximation of brain tissue – and that is a step toward simulation of the whole brain. The key point is, we are developing generic tools that will enable us to simulate any part of the brain of any species – including humans if we have the human data.

I can say that we will achieve our HBP goal by 2023. With basic mapping technologies, we will try to achieve simulation of the human brain. The maps that we will be able to build by 2023 will still be incomplete, and the outcome will not be a faithful representation of every aspect of the human brain. Biologically, lots of things in the brain are duplicated: one neuron, one parameter or cell, is very similar to a parameter and cell elsewhere in the brain – even in different species. The same applies to synapses, some basic circuits, etc. Then, the big constraint we have is that detailed brain simulation requires significant costs of computation power. We think we have enough computing power now, probably, to do a whole branch of the brain – we can't do it, because we don't have the data but computationally we could do it.

What are the proposed simulation requirements for the Human Brain Project supercomputer?

The Human Brain Project supercomputer will be built in stages. Simulation requirements regarding capacity of the hardware and software for the supercomputing and data infrastructure are still in development. We replace our supercomputers about every three years – when they are obsolete. We can currently do a little bit more than a petaflop, which is about 0.1 per cent of what we need. But that's not bad, because we expect to have a machine in 2016/2017 which can do 50 petaflops, and then we only need a twenty-times speed-up. The ultimate goal of the project is to perform cellular-level simulations of the 100 billion neurons of an entire human brain. To that end, we are counting on having an exaflop supercomputer with 100 petabytes of memory before 2023.

Will this artificial brain have the same intelligence and consciousness as humans?

That is a philosophical question. Consciousness is still a huge problem – we don't know what consciousness is. If we take a commonsense definition of consciousness, we know it's associated with certain neural states, so I know that you get synchrony across the brain, across different areas of the brain, and you get this in conscious people, conscious animals – you don't get it in people under anesthesia. Our models will definitely simulate those neural states, they will simulate neural states which in animals and humans we say is consciousness. Whether that will make them conscious or not is such a big philosophical question that I don't even know how we would test it.

It's really hard. In the end, I think we have to make a public judgment and I'm not sure that it's a question that science can really answer with our models today. Some people, such as Christopher Koch who writes books about consciousness, works at the Allen Institute for Brain Science, and works with us, believe that computers can be conscious – any system with enough information complex can become conscious, and that includes computers. Other people say this is nonsense, so it's a very controversial topic. I think the only thing we know is that we will produce those neural states, but everything else we obviously don't know today.

Since innovation can raise questions and dilemmas, it seems there should be a collective responsibility. Is this a responsible innovation on the part of HBP?

We are responsible for our actions in the society. For instance, we do medical things and I think it is very important to involve patients and their relatives. We have certain laws we are required to obey, but that is relatively easy. But the real thing we have to do is build up trust, so the people won't be scared of us. This is not required by law, since there's no law which says I have to discuss ethical and medical issues with patient groups, but if I want patients to contribute data voluntarily I think we have to involve them in frank discussion or otherwise they'll never do it. We have three areas of the project: basic neuroscience, medicine, and computing. They each have specific issues.

For the basic neuroscience there are a lot of people who are scared that we could do things that could harm people, but I think we have to discuss the issues openly so people can understand what we are doing. I personally think the science is very harmless, and it might bring us interesting information about our nature as human beings – but I don't think these are ethical concerns.

How do you achieve responsible research and innovation, and have you had any difficulties to date?

We will have difficulties, but so far no, because we are long way from practical things. The first aspect is legal responsibility – that's the easy one. We do have laws that we must adhere to, this is the *sine qua non*. There are laws about animal experimentation, laws about human experimentation, laws about privacy, laws about informed consent ... we have to respect those laws. So we have an ethics committee which looks at compliance and just makes sure that we obey the rules. Where possible, we leave it to local committees because in each different country in Europe the rules are different, and so where there are rules we allow the local committees to decide and we don't do anything. But in some cases there aren't rules or in some cases the rules are ambiguous – for instance, we use data from China so there we try to make sure they meet European standards, so we will not do research in China if that research would be prohibited in Europe.

The second aspect is to discuss with citizens how we should use our findings, what the rules should be. We are doing things that have long-term consequences, so suppose we have success in developing the science of computing – this could change industry. When we introduced conventional computers this changed industry; if we make better computers, this will change industry again, and it will affect employment, competitiveness, and many other things. But this is a long way away. If I made a new computer today, the economic effects would not appear for 20 years. So we have time.

We have to discuss how to use data: we could make big gains if we could access all the patients' data in hospitals, but patients might not want that, so we must discuss what is the right way of getting medical information while respecting patient rights. We can't succeed if we propose a solution that patients don't like. At the same time, we have to be ethically aware of ourselves because public opinion is important. I do think that working scientists understand some implications well before the public does, and that they are more realistic. They know what is possible, and they are not worried about science fictional scenarios. So, we need people to be informed about ethics within our project. Foresight Lab is a group based at King's College, in London. They talk to people and try to work out the issues in advance. They look at different possibilities – like, suppose it produces side effects that could possibly happen. Responsible innovation is all these things together.

Do you think that the outcome of the Human Brain Project could be used for nontherapeutic goals such as changing neurons or for military applications?

We're producing knowledge and we're producing technologies.

It's completely impossible to stop the use of knowledge. If you were interested in basic physics in 1929/1930 and you understood how the atomic nucleus works, you were preparing for Hiroshima – but you did not stop Hiroshima. We have a commitment in the project that we will not work directly with military organizations, but we don't have illusions that this will prevent the military from adapting our work.

I'm not very scared about neuroscience, because I think the military – what it does there, it will probably get from other places before it gets it from us. I'm more worried about technological applications of the science. Personal weapon systems might make use of our work – there's no doubt that you can use our work in negative ways. But what I think we can do as scientists is tell people about these risks, and then there can be a social, and even international, debate about what we do and what we don't. It is not true that all science is turned to weapons, even when it's possible. We have decided as an international society that we will not build bacterial biological weapons, chemical weapons – so society can decide that some things are just so bad that we don't want them.

What everyone is most scared of are autonomous weapon systems. On the one hand, some technologies could contribute to that, but on the other the military has been doing extremely well with today's computing technology. The real question is not so much the intelligence of neurotechnology because there is already a lot of drone intelligence today. We have to have an ethical discussion in society whether this is acceptable or not, and scientists can help to inform. Science can have a critical role; scientists can get rid of myths. But if the military wants to abuse our technology they can undoubtedly do it.

What are your strategies and methods to prevent this negative usage?

You can't prevent it.

What we can do is raise public awareness and discuss it with the public, and then the public collectively can prevent things. We can inform the public at large, journalists, politicians, non-governmental organizations, etc. The military can use technology without inventors even knowing it, and that is a concern. But I think it's naive to say people will give up any generic technology – brain technology is extremely generic. It is as if I say, "I will make an atomic bomb that is only used for killing people," and then you can say "I don't want you to make it." But if you say, "I will make a computer chip" ... well, computer chips are used everywhere from mobile phones, to washing machines, to bombs. Is it realistic to think people will give up the chip? I think it's not. It's like a hammer. You can use hammers to make weapons, but hammers have lots of other uses as well. Are we really going to give up powers that become available to us? It's not going to happen. I think we should encourage discussion and see what society wants and doesn't want – society wouldn't want to ban hammers, and I don't think anyone wants to ban computer chips. What we want to ban is *some uses* of hammers and chips.

The Human Brain Project is collating clinical data that will be available to researchers across the world as part of its open data policy. This could raise many issues: among them, obligations of informed consent, privacy, and confidentiality. What do you do to protect the data of patients?

This is an interesting issue. Our goal is to make sure that no one can ever identify individual data. The first thing we do is we take away all the personal identifiers so you'll never find obvious things that tell date of birth, etc. Second, we aggregate the data, so what actually happens is that data stays inside the hospital. For example, you are a researcher and you want to know what is the size of a certain part of the brain – the hippocampus – in patients who are having difficulty walking and have some sort of paralysis. You will make that query and you want to know those sizes. Our center splits that query into a separate query to each hospital. Then the hospital looks at it, using its own data, and identifies all the patients who can't

walk or who have that sort of problem. Then it looks at this brain measurement and gives back the statistics of sizes, so you don't get the individual patient. You will get the data from each hospital, and then we further aggregate this data and put it all together.

If a hospital has only one patient, or a small number of patients, with a certain condition, it refuses to answer and it doesn't give any data. So you can't get it. But security researchers will tell you that, under very special conditions, with a lot of effort, you might still be able to get the individual patient's data. That raises questions about different ethical regimes. In some countries, they demand it should be technically impossible to get the data – and that we can't do. In other countries, they say “impossible with reasonable efforts” – and that we can do. We can make it incredibly difficult to get any personal data. Ethically, I think that's enough.

Impossibility in science doesn't exist; in science nothing is impossible. So that depends also: even in these countries where they say “totally impossible,” a court might not interpret it that way. This is where trustworthiness comes in, and so we can argue about interpretation of the law. But you also have other kinds of cases, for instance in the UK where some hospitals did things that were completely legal but still caused much confusion. They had a project, called “Cared in the UK,” that was a bit like ours. They wanted to take hospital data, aggregate it, and use it for research. By British law this was completely legal. But there was a huge scandal, patients said no, and they had to cancel a project that cost them a hundred million. So, I think the real thing to understand is that being legal is necessary – but it is not enough. The real thing is that we must win the trust of patients and their representatives.

Which could be the most serious potential danger of the project, and what are the ethical issues that your Steering Committee of the Society and Ethics Programme (SEP) has discussed from the beginning?

I'm not actually particular scared. We have discussed two big, crucial ethics issues so far, and a third is coming up. The first one is medical informatics and the issue of privacy, informed consent, and confidentiality. This is the real issue. As you may know, Europe is just formulating a new data protection regulation, and this will affect us. And we don't know what the content of this regulation is because it is highly debated. Researchers throughout Europe are pressing for what's called a research exemption – that means we can use anonymous data without informed consent. But privacy advocates are pushing that there should not be such an exemption, and the Parliament has not yet decided so we don't know. We obviously would find it very convenient to have the exemption, but I still think we have these trust issues. If we have the exemption what we do is definitely legal, but that isn't enough for trust.

The second issue is dual use. This is not imminent for five to ten years, so we have time. Privacy and data protection is the immediate issue – it's our priority – but our ethics committee has formed a special work group to work on dual use. We must not have any illusions that we can stop dual use, but we should at least be aware of it.

A third issue is employment and the economy. But again, I'm not sure how far these are general issues for computer technology and how far they are specific to our work. Either way, we're certainly contributing to the issue. There's absolutely no guarantee that jobs that are destroyable will be replaced, but I think that is inevitable. So the real question is, How do we build our society under those conditions? That's a general social-political discussion, not just a discussion for HBP, and it will happen even if our project fails completely.

The Ethics, Legal and Social Aspects (ELSA) Committee includes the influential Oxford philosopher Julian Savulescu, director of the Oxford Uehiro Centre for Practical Ethics, which defends cognitive enhancement. How do his arguments about human enhancement fit into your

Society and Ethics Programme and the methods you have developed during investigations of the emerging technologies?

Of course, if you are an ethicist you will have points of view but any individual has his own point of view. We hope we have a committee that contains different people with different points of view. If you don't have varied points of view, you don't think, and we try to make sure we have different points of view. I don't see this as a big problem. We have non-religious people, religious people, scientists, non-scientists, lawyers – a big mix.

When you mentioned making a chip of the human brain, it reminded me of the Swedish hi-tech company Epicenter. It implanted a computer chip under the skin of its workers enabling them to interact with technology (for example, for access through security doors and to pay for lunch). Previously, the chip was used for therapeutic treatment (specifically, to treat Parkinson's disease), but it could also be used for location tracking and to include employee ID. How do you think this concern can be solved?

At the moment, you have a little bit of protection because such chips don't have a lot of range, which makes surveillance hard. This question is a new one for me, but I think it has to involve rules about acceptable and non-acceptable use. This is another case about generic technology.

Radio-Frequency Identification technology is now fundamental. I use it in my badge to get into this building, for contactless payment, and to identify luggage at the airport. Supermarkets use it for tracking stock – so if you think we're going to abolish RFID's this is a myth, it's not going to happen. They're just going to make them better and more powerful. So I don't think the solution lies in banning a technology whenever you can say that certain uses are illegal. For instance, it's like saying NCA reads my email, are we going to ban email because of that? No! But we can make mass email surveillance illegal. I would say, some things control workers too closely, and I would make that illegal. From my personal view, if the Human Brain Project Department accesses your Facebook account, or demands you give them a password, I would send them to jail – but I wouldn't ban Facebook. That explains the cyborg path. This is a moderately harmless technology; we can imagine worse scenarios.

I advocate for the establishment of a new ethical discipline, "Cyborgoethics," which would establish interdisciplinary dialogue between philosophers and scientists who are open to public judgment about the limits of the implementation of artificial techniques in a human being. Are the scientists in the engineering sciences ready for such a dialogue?

I'm not an ethical realist, so I don't believe there are ethical goods we can discover with that sort of ethics-ometer. I think we decide as a collective what sort of society we want to live in, whereas I can't decide it on my own because I have to take account of your needs and we have to discuss this collectively. On some things, we want to impose our moral rules, and there are some standards that we want to impose by legislation. If you want to implant something weird in your head, well I don't really want to stop you from doing it, but I certainly want to make sure that HBP doesn't force you to do it.

Do you consider it would be important to make a framework of ethical norms on how to control the ubiquitous cyborgization of human beings that could allow remote control of the person's health or behavior, or development of communicative information viruses inside the human body by interest groups? How you think this can be prevented in the future?

I think, for me, the big ethical issue with technology today is automation whether applied to the whole economy in the broad or to services. It's highly predictable that in ten years, if you go to McDonald's you'll be served by a machine. You'll just say, "I want a Big Mac" and a Big Mac will get out of the box.

Well, it can't be prevented. The first question is, Do we want to prevent it? I said earlier, it is not something that an ethicist can decide; it's something we as a population can decide. So we as citizens, do we want to keep low-level jobs that are easy to automate, low-level jobs that are dangerous? For instance, today, if you go to a car factory you no longer have workers in the paint shop, instead there are robots – that's probably a good thing because those workers used to die from cancer. We think that one is good. But if we replace bus drivers, what do you think about that? We have autopilots in our planes – we approved that because they are much safer than pilots. But what don't we approve? In predictive medical technology, I probably want to talk to a real doctor, because I want to discuss with someone who is a human being, who can tell me about side effects and try work out my decision with me. Well that's not just a science question. That's the sort of decision we have to make about automation.

Robot evolution will change our world. Do you think intelligent robots will overtake humans in the future?

For me the big ethical issue is machine decision-making. Pilots who occasionally make decisions are fairly comfortable. The autopilot sets the route of the plane, and we become comfortable with that, but the pilot has an override and can take control. But how much control are we prepared to give to the machines? I don't think that has anything specifically to do with brain technology, because we can give control to machines with our current technologies. And sometimes we like it. But there are other things where we would be very uncomfortable, with drones that decide when to shoot, or with robot policemen that decide when to shoot. I don't think we would want medical decision-making that was completely automatic.

Will robots be capable of evolving? How we can control them?

A robot is a complete machine, so there is another philosophically interesting question. Who decides the goals of machines? As long as we control the goals of machines, we still have control. Take the autopilot – it has a very clear objective, which is to take the plane from place A to place B, following a predefined route, to get it there safely, and to deal with any emergency that might happen on the way. I think most of us are comfortable with that. Suppose that the plane had the right to decide that I don't actually want to go to Athens, I want to go to Venice – will you be deeply disturbed by that if you lose control over the goals? We could make technology which did that today. We don't do it, because no one can see any point to it, but there is no technological reason why we can't do it if we want.

Computers can modify their own code. We don't usually do it because engineers like machines that are under control and predictable. If you allow a machine to change its own goals and change its own code, it becomes almost impossible to predict what it's going to do and engineers don't like that. So, we have good engineering reasons not to do it. But if we make machines that evolve or we make artificial organisms that evolve – evolution is now outside human control.

If I make a genetically modified organism and release it in the environment, I would claim it is mathematically impossible to predict where it's going to go. Unless you make it that it would die in 30 minutes – that's pretty safe. Even that mechanism, evolution could just get rid of it. If we lose control over the goals of the machine, for me we are crossing the barrier. As long as we have machines as tools, which help us, we are safe. I also think even though it's technologically possible to make machines that set their own goals, I can't really see a scenario in which that would happen and who would have an intent to do that.

Recently, there has been a lot of warning from Stephen Hawking that future developments in artificial intelligence have the potential to eradicate mankind. Would you consider this fear serious?

It is not intelligence that worries me, it's intelligence's goals. As long as we control the goals of our machines, I'm happy and relaxed. If we lose control over the goals of the machines, then I'm not happy at all. Our risk is not a risk specific to our technology or artificial intelligence, it's a risk of computers in general. We can lose control with a nuclear reactor and this has nothing to do with clever technology. So I'm worried about losing control, but I think this is a general technological problem.

How we can put technology under the control?

Engineered systems are under control. There are two kinds of control. I'm talking about control of goals. I can make a Google self-driving car that is badly programmed, and in some situation it goes crazy, or I can make an autopilot for a plane that in some seconds doesn't work properly. So, I'm not talking about *that* control. That's just bad engineering.

What I'm scared about is losing control over the goals of the machines, where machines set their own goals, and that would be scary. But it doesn't happen. We don't have it! We could do it! I can make a machine tomorrow which would do that!

How can you make it?

Well it's easy! You set up a function so the machine, let's say take an example – it's a little robot and it has to navigate to a target, not to do anything nasty, it just goes to a target and says, "Hello." When it reaches a target it lights up.

To do that it has in its memory the definition of the target, which I write, and as long as I write that definition the robot is under my control. If I change it, and add a little function that says choose a random location and make that your target – now the robot is no longer under my control. I've lost it. That I can do in five minutes on this computer. That is changing control without any intelligence. There is no intelligence there.

Are there ethical implications of human brain simulation, and what kind of consequences could it have for humanity?

The Human Brain Project has five streams in its *Society and Ethics Programme* (SEP): the *Foresight Lab*, which produces scenarios of potential developments; *Conceptual and philosophical analysis*, with initial focus on simulation; *Public dialogue with stakeholders*; *Research awareness* that encourages ethical reflection among the researchers of the HBP; and *Governance and regulation*, an independent ELSA committee that works with the National Ethics Committees in Europe.

The HBP is committed to civil research only; nonetheless, the HBP also has an open data policy. All partners have undertaken not to accept funding from, or use data or knowledge acquired for, military applications. There are some ethical issues regarding military and "dual use," such as: memory modulation; monitoring, augmenting and enhancing brain capacities; mind reading in counter terrorism; mind control, and so on.

Thank you for taking the time to speak with me today.

You're very welcome.