Where will the neurosciences lead us?

Włodzisław Duch
Department of Informatics and Neurocognitive Laboratory,
Nicolaus Copernicus University, Toruń, Poland

Where will the neurosciences lead us? What can we expect of the development of research into the way our brains work and cooperate with other brains? The positive results of such research would radically modify our understanding of ourselves and the world we live in. Now, we're not only able to observe our brains better, but also to influence them directly. In the near future, a synergy of IT, neurobiology and cognitive science will lead to the emergence of neurocognitive technology, whose influence on society cannot be yet imagined.

It all begins with quantum physics, which enabled nanotechnology to come to life. Technologies which emerged from this branch became the bases for cognitive robotics, autonomous vehicles, intelligent personal assistants – and they are applied in a number of ways. How can we utilise human developmental potential to the full? We won't know without proper research.

Obviously, there are many other areas whose development will radically change the world, such as the Internet of Things, synthetic biology, lab-grown organs and DNA programming. Even now, DNA programming languages already exist which allow for the modification of DNA chain sequences and the introduction of other fragments, thus creating organisms with particular qualities. Still, I will focus on research which affects our understanding of human nature – above all, of our brains. The 21st century is supposed to be the century of the brain, and it is becoming apparent. Woody Allen called the brain his "second favourite organ", but everything else depends on this organ, including liking things and experiencing pleasure.

A human being is an immensely complex system of cells and bacteria. We consist of 50 trillion (50,000 billion) cells! It's worth realising how great our biological complexity is. Comparing it to the complexity of the software which operates an android such as Bina48 shows how primitive the devices we can build are. Each of the 50 trillion cells contains an approximately two-metre-long DNA chain. Altogether, our bodies contain about 100 billion kilometres of DNA! This entire mechanism depends on the genome, which contains only 20,000 genes – almost as many as *C elegans*, a small nematode with only 302 neurons. Some cells live only for a few days, others – our entire life. Our body constantly changes – it's a continuous process, not a stable structure. The human brain itself contains nearly 100 billion neurons connected by about 1,000 trillion synapses, and so we have about one million billion connections that are created and may be modified throughout the years. These connections emerge so that a baby can understand what he or she sees, touches and hears, includes repetitive elements, can be interpreted, and enables interaction with the world.

What is the intellectual life of such a complex organism like? It's difficult to measure. In a few cases, internal life is very rich, but it can often be schematic. On a cellular level, animals are almost as complex as humans, but their thinking abilities are frequently limited, and their reactions are instinctive (based on embedded patterns). They interact with the environment in a very complex way, are able to survive in difficult conditions, move with agility, find food and process it into energy. These activities don't require conscious mental processes, even though brains direct the internal processes of living organisms and control their interactions with the environment in a complex way. But only humans are capable of creating refined mental models.

In these circumstances, developing artificial intelligence seemed unlikely, because thinking requires a highly complex brain, so a machine wouldn't be capable of it. In 1995, Chinook software beat Dr Tinsley, many-time world champion at a game of draughts, causing a shock. In 1997, chess players could hardly believe Deep Blue software's victory against multiple world champion Garry Kasparov. In 2011, the IBM-developed Watson program won at Jeopardy! against two champions. This success brought a new paradigm to IBM's activity – the company focused on cognitive IT, in which computers are supposed to perform tasks earlier reserved exclusively for people. In certain areas, artificial intelligence is already far ahead of us. In the case of chess, certain programs surpass human skill, so we can talk about computer superintelligence in this field. There are even programs which don't utilise the vast chess knowledge base, but reach the highest levels learning from scratch, merely by observing the positions of figures on the board. It seemed that in Go, the game considered the most difficult, achieving master-level skills would take many more years, but in early 2016 Google AlphaGo, based on deep learning algorithms and neural networks, beat multiple champion Lee Sedol, with the result 4-1. Artificial intelligence has made huge progress and appears in commercial applications, which is especially apparent in the recognition of structures, signals, speech and images. Road sign recognition is a standard in many car models, as is autonomous parking, and soon we can expect full automation of driving. People shouldn't drive – each year, 1.2 million people die in traffic accidents, and another large number become permanently disabled. Why are we afraid of terrorists, but not of cars?

Technologies are developing very quickly. Surgical robots such as Da Vinci still require the doctor to operate using available tools, but robots capable of performing surgery on their own already exist. STAR (Smart Tissue Autonomous Robot) performs complex surgical operations (so far only on animals), almost entirely without human control. Soon we may discover that a surgeon can't see as sharply or use their hands as precisely as a robot. Machines begin to surpass us in abstract thinking as well as in the skills which require precise hand movements connected with thinking. In 2015, Google received a patent for a method and computer system of robot personality development. In my futurological article from 2000, I predicted that the need for robot personality designers will appear before the end of this decade. A robot needs a personality in order to become more universal and to interact with humans. It's not as difficult as it seems. A more difficult challenge for robotics will be the requirements about materials, power, agility and speed of movement, and navigation, which will allow the robot to survive in a complex world. Creating an artificial rat which would survive in a hostile environment is already a challenge which requires the synergy of many functions. Building a robot which could beat humans at table tennis will be much easier than building a winning football team.

What are humans still better at? We desperately try to cling to something which will allow us to maintain our sense of superiority over machines. According to a common opinion, robots can think logically but can't experience emotions. But neurophysiologists who research the brain know that evoking emotions and manipulating them is a simpler task than processing the information related to language and symbolic thinking. From an evolutionary point of view, language is our latest function. Robots which read our emotions and react to them already exist: Paro the toy seal, Cuddler, Huggler, and other robots used in the therapy of people with depression and autism, or in the care of elderly people who have difficulties with communication. In the early 21st century, a new branch of artificial intelligence emerged – affective computing. An array of robots such as Nexi or the Wrocław-built EMYS can use advanced facial expressions and react emotionally, based on the analysis of a camera image or acoustic signal. By analysing such information, one is able to notice irony or sarcasm in someone's voice, to find out if someone is joking or being serious.

The most difficult task is gaining verbal dexterity. This requires not only the knowledge of notions and the relationships between them, but also imagining a certain model of how the world looks. When we try to translate something, for example using Google Translate, an algorithm, based on a large number of texts in both languages, will try to translate single words, and then to match fragments from the language we translate to, copying longer phrases. The result may be legible, but it won't convey the meaning in areas where we need deeper understanding – the kind of understanding most people have in their heads. This type of technology can be used for developing chatterbots. Equipping them with memory, in which someone's life story is encoded, will enable such bots, androids or humanoid robots to save certain aspects of personality. Such algorithms can create templates, write screenplays for TV series or summarise TV and press news. Associated Press, for example, offers automatons which summarise stock exchange events and various numerical data. This way, an automatic commentary sent by chatterbot can be made. Still, it's not enough to understand speech on the level sufficient to pass the Turing test. Some people are naive and let themselves be taken in, because they've never dealt with a chatterbot capable of a conversation, but the situation is changing now that Siri, Google Now, Cortana and Viv are present in millions of smartphones. Bots try to move discussion towards a direction that is controlled by them, because they have a script with ready answers – it suffices to notice keywords in a conversation in order to appear intelligent or offer a sensible answer. This isn't based on a deepened model of the world. It seems that understanding of the world requires activity in it, manipulating one's own body and external objects, constructing a model of one's environment and 'myself in this world'. Thanks to this, names – symbols which point at objects in the world – gain meaning through association, by understanding the means of interaction with objects and people, and by understanding the effects of one's own activity on the environment.

Text and speech analysis is constantly being perfected, so it's capable of answering an increasing number of questions, as we can observe in the example of IBM Watson's victory. *The Day a Computer Writes a Novel*, a novel written by a computer and submitted to the Hoshi Shinichi Literary Award competition in Japan, was promoted to the second level of the competition. The Todai Robot Project aims to create software which will solve the entrance exams for Japanese universities. So far, Todai has managed to receive 511 of the 950 possible points in the standardised exams in Japan (with a nationwide average of 411). The creators of the project hope for the program to pass the especially difficult exams for the University of Tokyo – including exams on various subjects like history and literature. Already in the 1960s, computer algebra systems appeared which solved mathematical tasks in the entrance exams for MIT, better than the students could.

Let's look a bit further into the future. It's the third year of the Human Brain Project, whose goal is to gather information about how brains work – and to create a computer supersimulator which would integrate all the collected data into one system, on which further research would be conducted. This would be of great significance for medicine, but also for artificial intelligence. Within the framework of the Brain Initiative in the US, many interesting projects for mapping brain activity were created. Large-scale projects have also been announced in China, South Korea, Japan and Brazil. Poland seems completely invisible in this company – there's no large-scale programme devoted to such subjects here at all.

Within these large-scale projects, new algorithms for artificial intelligence are being planned. The MICrONS (*Machine Intelligence from Cortical Networks*) project is supposed to create a new generation of machine learning algorithms, working on the level of human competency. The cerebral cortex is only 2-4 millimetres thin and is composed of cortical columns of less than a millimetre in diameter. These columns contain microcircuits. Understanding the way they function will facilitate transferring such functions to computers. We have a detailed model of a cortical column, created in the Blue Brain project. The general idea is to apply reverse brain engineering to extract information about the organisation of neural data processing, and transfer it to a smartphone user's personality model. The brain-computer connection will be strong enough to create an avatar with its user's personality qualities, familiar with his or her preferences and ways of thinking. Equipped with a vast knowledge of the world and quick access to information, such an avatar will become an extension of its user in a way: it will constantly provide advice and recommendations, making the user increasingly self-dependent. The goal of MICrONS is also to fix damaged brains. With knowledge about what was damaged in the brain and where, it will be possible to implant new neurons there – or rather, stem cells developing into neurons. In turn, new connections will emerge and fix missing functions – perhaps even adding new abilities to the brain.

So far, our connection with computers is not that strong. We often use GPS or Google Maps to find directions, but not for advice with complex problems. Still, all our choices will slowly become more and more algorithm-driven. Right now, most of the information we receive is based on data collected about us. Google, Amazon and Netflix work this way: search engines know our preferences and suggest advertising based on them. When we rent a film from Netflix, our previous choices are analysed, so we receive suggestions of potentially interesting films. It's also a natural function of intelligent TV. Banks do it when they grant loans, the stock exchange does it too – there, even microseconds matter when it concerns access to the central server. Social networks, increasingly frequented by bots, do it too. Sellers hire a bot to advertise their products and to comment about them online. Because this is common, we are increasingly becoming managed by algorithms. Once we're in our niche, we're pulled further into it and isolated from other interests, opinions and social groups.

The SyNAPSE project, developed since 2008 by a large, IBM-coordinated consortium, has currently reached its final phase. For the first time, there's the chance to construct chips which can function like synapses in the brain. When a series of impulses flow through synapses, their conductivity changes, and new information physically impacts the brain. If nothing changed, no trace in memory would remain. Constructing chips which could modify their structure physically under the influence of electricity flow makes electronics akin to biological neural networks. We could call it *neuronics*. This became possible thanks to memory storage, nanotechnology and understanding the way the cortical columns and microcircuits in the cortex work. One artificial neuron in the SyNAPSE-made TrueNorth processor requires about 500 transistors, so a model of a neuron is not merely a primitive logical switch. Quite the contrary – it's a fairly complex impulse-transmitting structure. Thanks to nanotechnology, it is possible to produce neuroprocessors consisting of a million artificial neurons. The complexity of one such neuroprocessor is equal to 5.5 billion transistors, but it requires 10,000 times less energy than conventional processors. Supercomputers require megawatts of power, while such a module uses no more energy than a tiny lightbulb. For the first time in history, we have an artificial system whose complexity nearly matches the human brain. Still, we have to learn how to utilise such chips. They will soon appear in our phones and, after deep learning, algorithm-based training, will be able to recognise complex structures, voices, images, video and abstract relations. In short, they will be able to see, hear and understand.

In order for this to happen, we have to know which regions connect with each other in the brain and how various subnetworks process information. In 2010, Dharmendra Modha, head of the SyNAPSE project, described the connection between 383 hierarchically organised areas of a macaque's brain. This is a modular structure in which not everything is interconnected. From a functional point of view, when we hear, understand or use language, very complex communication occurs in the brain – almost all areas transfer information between each other, adding previously remembered associations in order to recognise and interpret the signals. In experiments examining perception, memory or attention, functional connections aren't as vast, and active subnetworks are smaller. The human brain has a similar structure, only more complex. Most complex activities are possible due to cooperation between multiple areas of the brain.

Genetic determinism is a commonly known idea. A person born with microcephaly or congenital hydrocephalus won't become a brilliant student due to serious brain damage. Neural determinism is a lesser-known phenomenon. The results of various experiences in life, 'brainwashing', or our entire personal history shape the functional construction of the brain – in other words, we think in the way we're able to. When someone grows up in a specific environment, his or her behaviour and worldview will be difficult to change. One can't think differently than brain processes allow them to. In order to change the way of thinking, physical changes in the brain are necessary. These can happen slowly, or can be caused by electric shocks or strong impulses from a magnetic field. In such cases, neuroplasticity grows and the brain can change, but such procedures are not in use because of ethical reasons.

In order to understand the way the brain works, a metaphor is useful: the mind is a shadow of the brain's activity. We want to research this shadow in various ways, watching the way neurons connect in the brain, how neural energy spreads. Events – especially traumatic ones – form new connective paths between neurons, spreading out in all directions. Neurons have 10,000 inputs – they create an incredibly complex system. Why do they connect in this way and not in another one? Where do our habits and memories come from? In the dialogues between the Indian king Milinda and the sage Nagasena, written down 1600 years ago, habits are explained using a water analogy: water usually flows the same way it did many times before, down familiar paths, only sometimes creating new channels. We know beautiful examples of rock erosion. A similar erosion process occurs in brain matter: in an analogous way, paths of functional connections emerge in order to carry electric impulses. The brain needs to create 100,000 billion connections – every second, millions of new ones emerge, and pre-existing ones are reinforced. As a result, a 2-year-old child is able to recognise and understand words, walk, and control his or her movements. The structure of connections in the brain makes it possible. In order to understand this better, we need to conduct a detailed analysis – divide the brain into, for example, 1,000 areas and find out how activation spreads due to connections between those areas. This mechanism is used in analysis of the connectome – a collection of multiple connections between areas of the brain.

We will be able to connect neurochips following the connectome pattern, so that the entire system can work in a way analogous to our brains. When we imagine something visual, the activity of the visual cortex increases. The more vivid someone's visual imagination is, the stronger the activity of their primary visual cortex. It's the brain's 'neural space' in which images are created. By analysing activity in this space, we're able to reconstruct measured signals and to transfer them into sounds and images. We're able to see what visions or intentions appear in our minds. When we try to watch thoughts in a non-invasive way (without opening the skull), we're looking through a thick wall. The EEG signal is dispersed. Sometimes, in the case of severely ill people, we're able to look inside the brain and to put a net of electrodes onto, for example, the auditory cortex. On the basis of the local activity of groups of neurons, we're able to replicate the reactions of the auditory cortex to the sound of speech or imaginary speech – the impression of an internal voice. The frequency of impulses, along with time, space and energy, is information which we want to transfer into sounds and images. We want to predict how the brain will be activated by various words.

Neuroimaging, as a method of 'watching the mind', is becoming increasingly successful. Recently, Jack Gallant's laboratory at the University of Berkeley managed to build a semantic atlas that shows which patterns of brain activation refer to various concepts, for example the names of devices, animals, roads, places and many other categories. Every time we think about a particular notion, a specific brain activation emerges, equipping such notions with meaning, bringing to mind associations which help interpret them, and allowing for a particular response. Neuroimaging allows us to observe processes in the visual cortex and the entire brain – and, as a result, we can recreate the images we have in our heads. The world of the mind ceases to be closed to an external viewer. To an extent, we can also estimate what someone dreams about. It requires recording the activity of the brain and waking up the person sleeping in the scanner in order to ask about his or her dream. After a longer series of experiments, we can – on the basis of brain activity alone – choose the most likely of 20 dream themes. Because the skull blurs the image, it's not very precise. More detailed information can be received via a net of electrodes placed directly on the cerebral cortex. The device, called BrainGate: Turning Thought into Action, allows for the collection of data from the cerebral cortex and for more precise operation than by EEG electrodes placed on the head. A really good signal can be achieved only through electrodes located on the cortex or in the brain. Currently such operations are performed only for medical purposes, but in the future we may increase the scope of our perception and steer our thoughts this way – as such technologies become safer. What is happening in our brains when we read or watch a film? We recognise characters, their actions, motivations, time and place. We can follow information processing step by step, using functional resonance. The brain segments the stream of events, as if editing a film. With every change in situation, a quick change in brain activation configuration follows. Interestingly, when we try to present the facts using logical formulas, brain activations are different than in the case of verbal interpretation with the same logical structure. It means that we can learn formal logic for a long time, but that won't help to understand complex situations in daily life any better.

We can also read people’s intentions by observing the activity of their brains. What did I want to do? I went to the kitchen, and I forgot why; perhaps I had a specific intention. I usually learn about it when I'm already acting. Thanks to neuroimaging, we can see this intention in the medial part of the cerebral cortex. It's well hidden inside the brain. Activations of this area are related to spontaneous self-reflection. When I get two numbers which I'm supposed to add, subtract or multiply, an external viewer can see activation in the frontal part of the cingulate gyrus cortex before I make the decision. This activation makes it possible to predict what I'm going to do even before I know about it. As various experiments have shown (as early as in the 1960s), by observing the activation of the motor cortex we can predict if the subject will press a button or not about half a second earlier than his or her decision. If we observed the activity of the prefrontal cortex, in which the primary plans for action emerge, we would be able to make the prediction even 10 seconds earlier. Sometimes, various alternative plans are made, and one of them is chosen at the very last moment. I become aware of what I want to do (what my brain has planned), and I feel the will to act only in the moment in which the prefrontal cortex, which transfers information to the muscles, is sufficiently stimulated. 'Self' is one of the processes implemented by the brain; it is commonly thought that the brain follows the 'self'. But the bulk of research conducted until now shows clearly that it's not the 'self' which 'has' the brain, but rather the brain creates the 'self'.

When I discuss the issue of free will with my students and I explain these mechanisms to them, they usually say: "In that case, it's not me, it's my brain". It's a good excuse when we do something stupid. I don't always pay attention to it, but I've just scratched my head. I haven't thought 'now I want to lift my hand and scratch my head' – I did it spontaneously. Was it me or my brain? I guess it was me, even though I didn't act consciously. What 'my' brain does are my actions. Our problem with understanding the mind-body relationship is that we imagine the 'self' as some kind of idealised image of ourselves, not what our brain, or indeed our entire body does. 'Self' appears to be an abstract image, a model, which usually turns out to be false. After all, I sometimes do things which I don't accept. For example, I don't stick to my New Year's resolutions. I'm not like this – if I did something nasty, it wasn't me, but 'something told me to'. Or perhaps I am like this, and I fail to change, even though I try. When we understand this problem, it can become the starting point on the way to actual change.

It appears that using EEG, signals retrieved directly from the cortex, or those from electrodes implanted deeper in the brain, we can direct our own selves and change our behaviours in a more conscious way. In order to abstract our intentions from brain activity, it has to be processed in a very complex way, using machine learning algorithms. It allows us to direct not only our own actions but, for example, could also direct the actions of a distant robot. The robot would behave like me because it would be directed using information derived from my brain. That's the mechanism of a brain-computer interface. It became a popular technology recently. A thought-operated car travelled across the United States in 1997, and later through Siberia. It's an achievement in sport rather than in science. It requires stronger attention than during standard driving. Such interfaces are still very primitive – they can only recognise a few commands such as 'left' and 'right', and sometimes 'go forward', 'go back', 'stop'.

We would like to read the states of the brain better and to connect them to mental states – in other words, to shift from objectively measured data to subjective experience. Neurodynamics describes quite well what is happening in our brains. The number of interesting methods which can be used to research brain activity grows larger – we have NIRS, PET, fMRI and others. But the problem is that we aren't able to describe our inner mental states well. Phenomenologists made some attempts at this in the 1920s, but they failed. Recent works by philosophers of the mind, such as Eric Schwitzgebel's book *Perplexities of Consciousness*, describe situations in which people can't tell what they're actually feeling. They're unable to describe their 'psychological space'. I've been trying to connect neurodynamics with events in psychological space for 20 years. In certain simple cases of colour and shape perception, or simple decisions, it's possible – but it's very difficult on a larger scale. We have problems with accessing and describing the processes which take place in our brains.

Education changes the structure of students' brains. One can call it, metaphorically, 'brain sculpture'. Teachers obviously think they don't do anything like this – they perceive their students on a mental level, from the angle of symbolic communication. They talk about shaping personality and good habits. But effective learning has to affect the structure of the brain, of memory, of associations and of functioning. Two books on the relationship between education and the 'development of the central nervous system' were published in the late 19th century, before anyone discussed the notion of neuroplasticity. The question is: can we modify the brain directly, with the omission of the senses? Could I teach someone by forming a connection in their brain, without passing information on verbally? Would I be able to affect the associations forming in their brain? We're not always sure if the things we do are our own actions. The subject of agency is widely researched. I may be convinced that something is not my action – for example, when a dowsing rod in my hand wobbles, I may believe that it's not because of my influence, but that it moves on its own. But research shows that my muscles shrink and it's me who is inadvertently causing the wobbling, even though I have the impression that I'm not doing anything. I may also be convinced that I see the results of my actions, even though it's not true. Stimulating the frontal cortex with a magnetic field on one side of the brain may skew test results: the subject will point objects at this side in 80% of their replies, without any awareness that the result was influenced by stimulation. My brain may be affected by something invisible – but I still believe in my own conscious will. So far, placing a coil which would generate a strong magnetic field in someone's hat is not possible, but in the near future, it may well be.

In order to focus on something, effort is necessary. What does it mean to concentrate on sensory cues or on abstract thinking? Particular areas of the brain have to be strongly activated, and have to work in cooperation. Frontal lobes continuously have to send activating signals like 'don't sleep!' to the auditory or visual cortex. A cat waits tensely outside the mouse hole, and the neurons in its brain are working at high speed (about 20 impulses per second) to allow it to notice things and act quickly. When any movement or sound appears, the neurons emit 40 impulses per second – the visual cortex works in sync with the motor cortex, which makes perception and action nearly immediate. The cat won't blank out – there's not enough space in its brain for too many simultaneous processes. But when we observe something for a longer while – if we analyse images, or travel by car – our attention disperses, and we may daydream or fall into microsleep. Direct stimulation of the cortex using alternating current can help maintain concentration without mental effort. Such smartphone-operated stimulators have become popular recently – they're advertised as tools for increasing productivity, gaining mental energy before work, and also for relaxing.

In short, we may soon be able to perform fairly complex tasks consciously, by programming our brains – in fact, programming ourselves. When we grow up and learn about ourselves, we describe our reactions in order to define what we're like. We observe our reactions and the reactions of others to our behaviour. We get to know ourselves by observing the flow of information inside the brain, but also using the information we receive from outside.

Technologies of external control are obviously a subject of interest for the army. In the American army, the Engagement Skills Trainer (EST) has been expanded by Neuro-EST. It's worth analysing the processes in the brains of well-trained experts in order to stimulate similar activations in the student's brain. The transfer of skills between the expert and the student requires identification of states of the brain and stimulation through the skull. If we can read information from one brain and transfer it into another, then perhaps the long-distance transfer of thoughts and impressions may be possible as well. Such attempts have been already made. A message was already sent from one brain to another using Morse code. So far, we can't call that a genuine telepathic transfer – rather a curio of little practical significance. Direct intervention into the brain is applied only in severe medical conditions. Deep brain stimulation is used in the treatment of Parkinson's disease, compulsive-obsessive disorder, severe depression and addiction. People with an implanted stimulator can regulate its influence on their brains when they feel that their hands are too shaky or they begin to feel an internal compulsion driving them to act obsessively. Of course, instead of implanting electrodes, it would be better to use non-invasive methods like neurofeedback. There are many processes taking place in the brain which we don't yet know well, but using EEG we can at least read some of this information and transfer it into images and sounds. It can make us realise that our brains are unnecessarily overstimulated and too prone to distraction when we want to focus – for example on dancing or music improvisation. The technology is still quite primitive – it hasn't evolved much since my 1978 article *Electronics and Stress* in the *Przekrój* weekly – but I hope that our own research will allow for the development of new, more efficient neurofeedback methods.

Certain people with mental disabilities can display extraordinary skills – for example astonishing memory or numeracy. There are few people like this – only about 100 cases of the savant syndrome have been described so far. People with such qualities are often incapable of self-sufficiency and often have an IQ of between 40 and 70, which makes them dependent on permanent care. Half of them are diagnosed with the autistic spectrum disorder. It's possible that no one would like to become a savant forever, but could an able-bodied person be turned into one temporarily? Allan Snyder from the University of Sydney conducted such an experiment. Since savants' brains are impaired, perhaps we could temporarily turn off a part of the brain using a magnetic field or alternate current and find out if the remaining part will work better. The assumption was that without unnecessary activations diverting energy, the brain should function more efficiently. To an extent this proved true. Some people who took part in the experiment slightly improved their drawing or computer gaming skills. Still, the improvement didn't work on everyone, and the difference wasn't striking. The applied method couldn't inhibit specific brain processes; it impacts quite a large area of the brain.

Finally, it's worth asking – can we transfer personality from one brain to another? People try to do it on a minor scale, building a model of a person's memory and specific behaviours on the basis of conversations and observations. A utopian initiative named Project 2045 is aimed at enabling a complete mind transfer from a human brain to a neurocomputer. We already have a neurocomputer that matches the brain in complexity, we will be able to create a convincingly brain-like structure in a few years, and we can extract information from the brain in a number of ways – why not try uploading a human personality into a computer? The Electronic Immortality Corporation believes that this should be possible by around 2045.

How does the development of neurocognitive technologies and digital access to information influence people, their understanding of the world, and their relationships with others? Marc Pronsky in his well-known article *Digital Natives, Digital Immigrants*, claims that children brought up in a multimedia-rich environment won't be interested in traditional, boring way of learning. The problem with multimedia presentations is that they're easily memorized but offer little more than an episodic collection of facts and memories. In order to use them for systematic thinking and drawing conclusions, it's necessary to turn them into elements of semantic memory. That's not an easy task; it requires multiple repetitions. We learn the multiplication tables slowly. It seems that there's a danger from the point of view of episodic and semantic memory. A number of articles have raised the issue of 'dumbing down' due to easy access to information. In *The Shallows* book Nicholas Carry describes the damage to our brains caused by the internet, which makes us intellectually shallow. The ability to concentrate diminishes, we're less able to contemplate the materials we study, our cognitive skills worsen, and the brain is forced into large energetic efforts when we jump from one subject to another. Switching various areas of the brain on and off is very energy-consuming. How can we prepare our children's brains to function efficiently in today's world? We cannot self-regulate ourselves well – and that turns out to be a key issue. We don't understand what can satisfy us in the long run, what we should pay attention to. We have problems with memory and the rational assessment of emotional situations, we develop bad habits, and we have limited capacity of our senses. But at the same time, new possibilities of brain enhancement appear, and people sooner or later will begin to use them on a larger scale. We can imagine sense-augmenting implants, like lenses or artificial eyes which allow for long-distance viewing, will be turned into a microscope, or will see in UV and notice bacteria. We can be sure that all these things will soon be possible.

In this situation, we should ask if we're witnessing the birth of a new human species, *Homo sapiens digital* – a transhumanoid with enhanced senses and modified deep thinking mechanisms. Ted Berger at the University of California has already begun to place electronic implants in the hippocampus in order to enhance memory. Perhaps we will replace parts of our brains one by one, in order to become fully-fledged cyborgs in the end. In such a situation, our brains would never malfunction, just become increasingly perfect. This is one of the future's possibilities. We have to remember that the state of the brain is dependent not only on what is happening inside it, but above all on the results of interaction with the environment. It is an 'event in the world'. Changing the possibilities of perception through hearing, vision and other implants will impact the functioning of the brain and the way it perceives the world.

Systems which nearly match the human brain in complexity already exist. I have serious concerns that we have great opportunities, but also great dangers ahead. Everything may head in the right direction and help us answer currently unanswerable questions. Which factors shape human nature? What builds our images of the world, our memes? How is brain development directed by the culture we exist in, by literature or music? How do we develop human potential to the fullest extent, beginning from the earliest age, supporting development at every stage? We don't understand these issues fully yet. At the same time, we teach facts at school, not how to better understand ourselves. It's the wrong approach, far from the ancient Greek ideal of *paideia*.

When we begin to consider technological development seriously, perhaps it can be utilised for good purposes. But it can also be the other way round, because brainwashing, manipulating public opinion, and rearing fanatics are also within the reach of brain-influencing technologies. Obligatory mind-controlling hats, dispersing and directing our thinking processes known from science-fiction stories, are not merely a fantasy anymore, but a very real danger.

This is a transcript of my talk at the conference "Interfaces, codes and symbols",
Wrocław, May 2016