

Press Release



When a computer is not a computing machine...

Peter AJ van der Made, vWISP Pty Ltd.

After 50 years of research and billions of dollars poured into Artificial Intelligence, the only ‘intelligent’ machines have so far appeared in fictional Hollywood movies. They vary from cute little animated robots that would not hurt a fly to laser-beam shooting monsters that wipe out the human race and take over the world. Several movie producers have done their homework and understood that the human brain has awesome capabilities. Their movies show massive computers with petaflop processors and multi-terabytes of memory. Others seem to think that a computer with such impressive capabilities will fit into the size of a can of sardines. Whether in deep space, on another planet or Earth, the machines eventually turn murderous and wipe out everyone who comes near, or develop stunning new intelligence and compassion by rewiring themselves. So far for fiction and wishful thinking, for the average guy it can be difficult to make a distinction between fact and fiction in a world where the boundaries are increasingly fuzzy.

You sure have nothing to fear from your wiz-bang quad-core desktop PC. No matter what the salesman and the marketing gurus have told you, it is just too slow and dumb to make any attempt against your intelligence. You are amazing – let me rephrase that before anyone gets the wrong idea; I mean your brain is amazing. Constructed from at least 100 billion processing nerve cells and somewhere in the order of 100 trillion synaptic “memory” junctions, the human brain leaves any computer for dead. To put this into perspective, if both your PC and your brain were competing race cars in which processing speed is represented as velocity, your brain racer would have gone up and down to the moon before your desktop PC racer had crawled 6.5 meters across the floor. To approach your capabilities the PC would have to speed up 116 million times. How do I know this? Elementary my dear Watson, let me explain. A brain nerve cell, generally known as a neuron, can be simulated with remarkable realism by a computer program. The program runs a mathematical model that was derived from the operation of a real nerve cell. It takes about 0.0011 seconds to execute this model for a single cell on a 3

GHz state-of-the-art desktop PC. To execute the model for the entire brain would therefore take 100 billion times the execution time of a single cell, which comes to 30556 hours or 1273 days. In contrast, the response time of our brain is less than a second or at least 116 million times faster. When a little girl crossed the street in front of my car I responded a lot faster. In a fraction of a second I pulled the handbrake, hit the brakes and turned the steering wheel to avoid her. In a panic situation our brain can respond 440 million times faster than your state-of-the-art PC. The best supercomputers, machines the size of a rather large fridge and a lot more power-hungry run only 100,000 times faster than your PC and cost in the range of a billion dollars. That is still a long way from the 116 million-increase that is needed to simulate the entire human brain in real time. The best supercomputer still takes 19 minutes to do what you do in less than a second. It is good to know that a billion machine is nearly 1200 times slower than we are, unless your happen to be making billion dollar calculators.

In your near infinite wisdom, which I just worked out to be better than the fastest supercomputer, you can probably pick several holes in this theory. For one, would we have to simulate all the neurons all the time? Or could we do away with the idea of simulating neurons altogether and simulate higher brain functions instead? Maybe, but no one has yet been able to formulate a unifying brain algorithm, one that works for all the things our brain does all the time; vision, audio perception, olfactory perception, tactile perception, spatial perception, association and thought, that works and does not depend on a similar or far more complex algorithm. Mathematical models that take into consideration everything that is going on in the brain become too complex to comprehend even for mathematicians. The results of AI research that ignored the neuron model have been meager even after many years of research and development. Computers are good at computing, after all, they are huge programmable calculators. Our brains are not calculators, they are association engines genetically programmed to associate events at birth, which form the basis to learn new associations from real-world events over our lifetime. We store information by repetition and recall information by tags.

Brains do not have a central processor, or a separate memory module. Memory is an intrinsic part of each processing junction. Brains do not run programs but respond to real-world events that comply with certain rules. The rules remain the same. A word is always composed from a similar frequency pattern – however not the same frequencies. The structural differences between the brain and a computer are so immense that I wonder if it makes sense to simulate the brain on a computer. Aside from the ethical questions, would it be possible to contrive a biological brain that executes Boolean logic and simulates a computer? If this question seems absurd then perhaps the previous one is absurd too, unless we build a computer that is not a computing machine but consists out of association machines that are similar to nerve cells. By repeating this structure in a hierarchy in the same way it is done in the brain we may be able to simulate the entire brain without the need for unattainable processing speeds.

We can certainly build simple systems consisting from a few million artificial nerve cells, interconnected through artificial synapses. Such hierarchical models are complex, but when we build them we learn what we are doing wrong, and improve on the design in

each consecutive model. One thing appears clear; the science of Artificial Intelligence is about to take a leap forward after too many years of stagnation. Artificial Intelligence (AI) systems are not a replacement for people or for computers. Computers, as calculating machines are irreplaceable in a lot of current applications. Association machines or “Artificial Brains” if you like, can offer improvements in areas where computers are not very good, such as face recognition, and other cognitive tasks. A whole world of new applications comes to mind; for instance, they could be used to protect teenage drivers from killing themselves on our roads. To do that we would need to teach an AI system to drive safely and replicate that knowledge to other AI systems. Such a system would then be built into a car to monitor the actions of the driver and say something like “You can’t do that, Dave” when the driver tries to do something stupid. They might even think it’s cool. We could design AI devices that connect straight into the human nervous system to control prosthesis. We could teach AI devices to understand spoken words and sentences in context with a 100% success rate, or teach AI systems to explore deep under the earth for minerals, systems that know what to look for without human supervision. They could be used in educational toys, intelligent alarm systems, games, and face recognition security systems.

What happens when such machines start to approach the complexity of the human brain? At this point we enter the realm of speculation. Will the system develop awareness? How much of the limbic system do we need to build into our ‘cyborg’? What will the system learn and how do we control what it learns? Asimov’s “three laws of robotics” can not be put into practice, because our machines don’t run programs. Will machines take over the world? Was Hollywood right after all? Fear has always surrounded new technology.

A single simulated nerve cell with just 16 synapses requires about 3500 logic gates. The processor in a PC consists of about 400 million logic gates. A theoretical neural chip with the same transistor count would contain 114,000 artificial neurons and just over 1.8 million synapses. To simulate the human brain we would need nearly 900,000 of such chips. It may be more viable to start with something simple like a cockroach. It has a million neurons and we could build one using just 10 chips – but it is still a big cockroach. Eventually we may learn new techniques that make it possible to make these things even smaller. For now, there are physical limits on the size of transistors in a chip. At the point where a transistor is just a few molecules across it stops working like a binary switch.

Realistically, trainable machines that learn to perform a specific task are a concrete possibility. The vWISP NeoCortex processor is a teachable association engine that works in the same way as a neural cell and its hierarchical structure is based on the human brain.

vWISP Pty. Ltd. is located at the Innovation Centre, USC, Sippy Downs Drive, Sippy Downs, Queensland. For more information visit our web site: www.vwisp.net.au