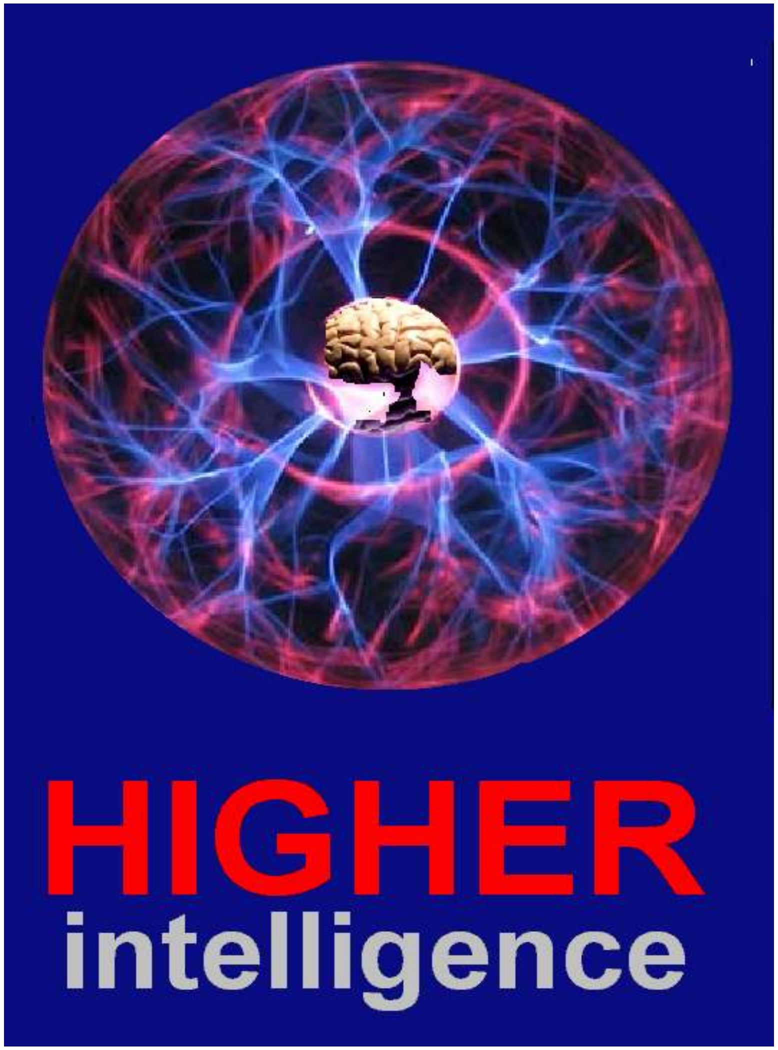

---------------------Page Break---------------------

About the Author

Computers were an almost magical realm in the late 1950's and early 60's. The idea that a machine could be 'smart' fascinated me. It was a boyhood dream of mine that one day I would work with those machines. I had also an eager interest in medicine, and I had to make a decision to pursue a career in either medicine or in computer engineering. Medical students were all too happy to share their horror stories of gore and bodily fluids, so I decided on computer engineering. When I finished my formal education in the early 70's I went to work at Sperry Univac. A few years later I got into design and followed the rules to design products within budget and mostly on time. I developed bits that were buried deep inside the computer; integrated circuits for graphics, add-in boards for speech synthesis, Analog to Digital conversion and video display boards. I wrote programs in Assembler, (a language that is close to machine language), FORTRAN, Basic and C. After a few more years the excitement wore off, and I got a little bored with the computer industry and I looked longingly back at the wonders of the human body.

In the early 90's I was researching a better method to detect computer viruses. I looked at biology for inspiration. In the natural world a DNA signature protein identifies the body's cells. The immune system recognizes this protein and destroys any cells that have an unknown protein signature. To make a long story short, the detection method that I developed for computer viruses works by extracting a program's "DNA". It does this by performing a full analysis of all its functions within a simulated 'software' computer. The resulting sequence of the program's actions is the program's "DNA". An analysis routine then evaluates the behavior and detects if this was a new program, an updated version of a previously analyzed program, or a program containing a computer virus. If the program is 'clean' it is marked with an identifying "protein signature" that can be quickly checked next time without the need to analyze the program.

This was the vCIS anti-virus immune system. It was sold to IBM-Internet Security Systems in Atlanta in 2002 and it is now used in IBM's 'Proventia' ™ security products. After working with the company as Chief Scientist for behavior analysis technology, I left at the end of the following year. That gave me the time and the finances to return to an old passion.

The human brain has fascinated me for some time. In 1986 I had started work on a cell computer to simulate brain function, when I interrupted that work to spend my time on the more pressing Computer Immune System. I have a 1986 patent paper that was submitted for the cell processing system, but I did not follow it through.

I returned to the study of the human brain in 2003 and by 2008 I had developed a chip that contained just ten neurons which showed promise by learning to recognize simple sounds. The device learns by updating values in synapses, which contain a register that is updated every time the neuron triggers.

In the course of this work I also considered existing Artificial Neural Networks, and found them to be as lacking as a music box compared to a symphony orchestra. Actually, this is not a bad comparison. In a music box the tones are fixed; it plays the same primitive 'pling / plong' melody over and over. In the same manner, Artificial Neural Networks (ANNs) work with static weights that assign a value to each input. The ANN (to avoid confusion with a functional neuron) is triggered, e.g. produces an output, when the summed value of all input values reaches a certain threshold. Therefore, the same input results in the same output, unless the weights are changed by some external means. Rather like a music box playing the same tune over and over.

In contrast, the brain is far more complex, more like a very large symphony orchestra. Many instruments are all doing their own thing. But they all play of the same sheet music and the performance comes together as one. No two performances are ever exactly alike, although they may be very similar. Like the brain, an orchestra also has the ability to learn an unlimited number of new tunes. The tunes are played haltingly at first but improve over time with practice. In the same way, when our brain learns new things they feel awkward at first. With practice, we feel more comfortable until the task becomes routine and we perform it with almost no thought. Driving a car is a good example of this. All the regions of our brain are constantly active playing their own tunes, learning all the time. It all adds up to one unique symphony that makes us who we are. My passion is to decipher that symphony to build machines that learn like humans learn, evolve intelligence as they learn, and serve man with their acquired intelligence. I believe that the answer is not to build clever machines, but machines that have the ability to learn to get clever.

--.o0o.--

~~~~~~~~~~~~~~~~~~~~Page Break~~~~~~~~~~~~~~~~~~~~
~~~~~~~~~~~~~~~~~~~~

Introduction

Some well respected individuals would have you believe that intelligent systems are just around the corner. They even speculate on machines that are smarter than we are. That we can program a computer not just to mimic the actions of the brain, but to surpass it. Intelligent robots have appeared in science fiction movies as far back as the 1950's.

The reality is very different. Huge supercomputers are needed, using many hundreds of thousands of processors, to mimic just a small part of the human brain. Your brain is truly amazing. Most of us take human abilities for granted and we don't realize what an amazing brain we have. I see the brain as the best example that we have of the existence of a Higher Power, a creator. No other organism has ever had a brain anywhere near the complexity of our brain. Some animals have much larger brains. A whale brain for instance is roughly five times the size of a human brain, but they do not have the intellectual abilities that we have. It is assumed that most of the whale brain is used in sensory processing due to their large body size. On the other hand, African grey parrots have tiny brains but have amazing abilities in understanding and producing human speech, to the point where a trainer can carry on a conversation with the bird. Alex was a famous African grey parrot whose videos, demonstrating his amazing abilities, can still be viewed on the internet. Parrots brains have particularly well developed areas for speech and communications. It is not brain size, but brain structure that appears to define our capabilities.

Each generation has attempted to explain the brain from the perspective of the level of technology that was available at the time. We are now in the age of computers, so we try to explain the brain from that perspective. But the brain is not a computer; it does not 'compute'. It works by completely different principles.

The earliest evidence of successful brain surgery is found around 7000 years ago in Ensisheim, France. A skull was found there that had been opened up while its owner was still alive. Judging by bone growth around the edges he lived for some time following the surgery. The Edwin Smith papyrus (named after its discoverer) originates in Egypt approximately 1700 B.C. and contains 48 case histories of brain injuries, treatment and symptoms. It is copied from an earlier papyrus that is estimated to be a 1000 years older. Hippocrates (460-370 BC) wrote a paper "On the injuries of the Head" in which he describes cases in which the skull is to be opened to treat the brain. Aristotle (400 BC) thought the brain was nothing more than a cooling system for the blood and the heart. By the middle ages the consensus was that the brain harbored thought, memory and fantasy.

Just fifty years ago, John von Neumann described the brain as a large collection of switches with some analog properties, according to the state of computer technology at the time. Today, many still view the brain as some sort of computer and attempt to simulate it using an increasingly complex mathematical model. While we may attain new understanding of the brain from these efforts, it is not likely to lead to intelligent machines. How then do we build a brain? We need a new vision.

The brain's function is to interpret the world, learn from it, and to contribute to the survival of the species. What the brain learns is perceived through the senses. The brain constantly learns from the moment it is formed. Much of what we learn takes place at a non-conscious level. In our early youth we had to learn how to walk, how to speak and how to catch a ball. These are physical abilities that require precise control of motor skills. Once learned, we perform them without thought. We remember the places we have been to, street directions, books we have read, movies we have seen, and millions of other facts. No memory is ever recalled with the same exact details.

In contrast, a computer is a programmable calculator, related to a distant relative; the mechanical tabulating machine. Like these early mechanical machines, it runs programs that govern its function in exact, repeatable steps. It is not aware of what it does. Without a program – A human programmer's recorded intelligence – it does nothing. It repeats the same exact steps every time. Data is stored in an exact order and with perfect recall. It can search lists very quickly, but no one has ever observed a search routine operating in the brain. We recall information instantly.

The different principles by which computers and brains operate are the cause of increasing complications in simulating the brain on a computer. We need a new perspective on the brain. One in which we separate the information carrier part of the brain from its 'software', e.g. its accumulated knowledge that determines its behavior as much as a programs determines what a computer does.

Intelligence is not governed by the rules of logic; it evolves as a function of learning. A baby is not born with complete knowledge. The brain is a huge three dimensional information carrier in which the stored information itself determines the processing algorithm. The 'Mind' consists of knowledge that overlays this physical information carrier. Intelligence is not programmed, but evolves as the information carrier gets filled with data. The information itself forms the selection criteria in the processing path and facilitates further learning.

You may think that we do not just store information. We store impressions, feelings, emotions, tastes, sounds, tactile senses, and etcetera. However, in the brain all this information is treated in the same way. Whether it is a taste sensation, a beautiful sunset, or a symphony, the senses convert it all to sequences of pulses. There is no difference in the perception of a Rolling Stones record or a symphony in the auditory part of the brain.

The brain is a dark place enclosed in the skull. A sunset is converted by the eyes to millions of pulse streams that express the location of horizontal, diagonal and vertical lines, color information and intensity. Similarly, the sound from a symphony is converted into many pulse streams in the cochlea. Each of these pulse streams consists of events that denote the number of audio vibrations per second.

It is the same with tactile information, the location of our limbs, smells and tastes. Massive channels containing pulse streams from all the senses enter the brain every second of every day. At this point there is no 'sound' or 'vision' or 'feeling' – there are only pulse streams.

We know how the cells in the brain communicate. We know how they are organized. The number of papers that are published on brain science has exploded in recent times as new machines and methods of research have become available. Continuing research on animals still reveals new information. With new functional MRI machines we can observe the brain of a living human being, and see which parts activate while certain tasks are performed. It is still unclear what the exact function is of Glial cells in the brain. There is uncertainty about the exact function of many neurotransmitters and neuromodulators, especially in combination. Although our model may not be complete, we are now in a position to build machines that work by the principles that know about. Building these models is useful to get to a more complete understanding of the brain.

The Blue Brain project is an attempt to simulate a complex model of the entire brain on an IBM Blue Gene ™ supercomputer. The proof is still in the pudding, so to speak.

Blue Gene/P Supercomputer © 2009 IBM. 213,000 parallel PowerPC processors

Dr. Eugene Izhikevich of the Neuroscience Institute in San Diego performed a simulation of a model of the brain using normal PCs in a shared processing network. He used a neural model that was much simpler than the model used in the Blue Brain project. In this simulation he used 27 computers. It took 50 days to simulate 1 second of brain time. A quick calculation shows us that our brain is at least 117 million times faster than a 3 GHz PC. If he were to use an accurate neural model, incorporating all the information that we have, this figure would be many times higher. Hence the need for a supercomputer.

Given the high level of integration, and the complexity of the Blue Gene supercomputer, it is very unlikely that this system will be placed on a single chip anytime soon. To build suitable artificial brains we need a different approach, one that is more closely related to the methods used in the brain.

Would we have found the laws of aerodynamics if we had simulated the flight of a bird on a computer? Would we have to simulate all the tendons in the wing, the airflow over the feathers, the force excerpted by the wing on the downward stroke and the resistance of the air? Would we have discovered from our computer model that it is the shape of the wing that gives a bird lift? A 747 surpasses the capabilities of a bird by millions of times, all thanks to the feeble beginnings of a physical model that was built by the Wright brothers early in the 20th century, which revealed the laws of aerodynamics. In the same way, I believe we need to build functional models of the brain, not on a computer, but operating by the principles of the brain.

A program on a computer simulating the human brain, which in turn is acting on accumulated knowledge, is two levels of abstraction removed from the real world. The first level is the mathematical model of the brain. That model is very complex and likely to contain errors or omissions in its description. The second level is the computer itself. Computers consist of binary logic gates. Why not use the same technology to build a brain according to the mathematical model? This reduces the abstraction level back to one, resulting in much lower clocking speeds to produce the same results, a large reduction in the number of parts and a much lower power consumption. The result is smaller, compact artificial brains that can be used to learn the shortcomings of the model, and eventually in applications ranging from intelligent, self-driving cars to space exploration machines and brain prosthesis.

Chapter 1. What is Intelligence?

If the Aborigine drafted an I.Q. test, all of Western civilization would presumably flunk it. ~ Stanley Garn

Browsing the New York Times archives is an interesting journey into the past, giving us a contemporary view of the dawn of modern computers. Computers were referred to as 'Brains' in the post-WW2 era. In reality these 'Brains' were nothing but simple programmable calculators. Perhaps inspired by these 'brain' stories, many fictive tales and movies have emerged about robots and smart, but often murderous machines. Movies like 'Starwars' and 'A Space Odyssey' have conditioned the public into believing that intelligent machines are not just possible, but likely. The reality is very different; to put it bluntly – an intelligent, aware and thinking machine is not even a remote technical possibility anytime soon. A whole new technology will need to emerge to build such a machine. An accurate simulation of the entire brain takes a very large computer using several megawatts of power. Before we can create an intelligent machine we need to know how we define 'intelligence'. Will an accurate simulation need to sleep? Will it forget?

Let us consider what it takes today for a computer program to pass an I.Q. test. The test questions have to be known to the programmer. Often such questions include sequences of numbers, figures and word combinations. Once the program has determined that one particular test requires the next number in a sequence of random numbers it can calculate the next number. The same is possible for word combinations. Figures become a lot more difficult to analyze.

Varying the questions would cause the program to score much lower or not at all, depending on the level of variation. A finite set of varying questions would increase the complexity of the program by a factor that is directly related to the level of variation.

Hence, an open set of I.Q. test questions would increase the complexity of the program by a factor approaching infinity. Since this is not possible the program would score either by pure coincidence or not at all. A human being faced with the same set of I.Q. test questions will score by spatial recognition, sequence recognition, numeric association, word association, objective analysis and by coincidence. Computer programs can not do these things very well and the complexity of programs quickly becomes a stumbling block.

Intelligence tests are also cultural background sensitive. Depending on the test design, the score will largely depend on the subject's background. Many I.Q. tests presume a western education. An I.Q. test devised for one cultural group will most likely not be suitable for another cultural group. Does this mean that one group of people is more intelligent than the other groups? Not at all. It simply shows that intelligence, more specifically the type of intelligence, is a function of environment.

The first time I went to a village in Papua I was astonished at the villagers' ability to find food everywhere. Their brains have been trained from an early age to survive in that setting and their intelligence evolved to match. A smart business executive, or an academic, will do very poorly in this setting.

When we spent some time in a new setting, we adapt through learning. We develop new skills. Someone from a Papuan village who receives a western education will score very well on the above mentioned I.Q. test. The business man and the academic learn to survive if they receive help and live long enough.

A dictionary definition of intelligence is the ability to learn from the environment and to interact with it. Learning sets up new pathways in the brain that change our behavior. So far we have seen no robots that can adapt to their environment through learning. Intelligent machines do not exist.

In his 1950 paper 'Computing machinery and Intelligence", Alan Turing defined a machine as intelligent when its responses in a