The JOURNAL *of* GEOETHICAL NANOTECHNOLOGY

Terasem Movement, Inc.

Volume 2, Issue 4
4th Quarter, 2007

How Neuronanotechnology Will Lead to Melding of Mind and Machine..2

Ray Kurzweil, Ph.D.

An Entrepreneur, author, futurist and producer, Ray Kurzweil, deftly explains a path to understanding the basic principles of human and artificial intelligence, and reverse engineering of the brain in attempting to copy intelligently biological subsystems and brain uploading.

Upgrading Humans - Technical Realities and New Morals..30

Kevin Warwick, Ph.D.

He is the Professor of Cybernetics at the University of Reading, United Kingdom, describes how his 1998 experiment allowed him the title of the world's first, 'Human Cyborg' when he implanted a Radio Frequency Identification Device (of his own design), within his body. Dr. Warwick also explains the present and possible future benefits of the technology of merging man and machine.

Terasem Movement, Inc.
201 Oak Street
Melbourne Beach, Fl 32951

Editor-in-Chief: Martine Rothblatt, Ph.D., J.D.
Managing Editor: Loraine J. Rhodes

The JOURNAL of GEOETHICAL NANOTECHNOLOGY

Volume 2, Issue 4
4th Quarter, 2007

How Neuronanotechnology Will Lead to Melding of Mind and Machine

Ray Kurzweil

This article was adapted from a lecture given by Ray Kurzweil at Terasem's 2nd Annual Workshop on Geoethical Nanotechnology on July 20, 2006 in Lincoln, VT.

Entrepreneur, author, futurist and producer, Ray Kurzweil, deftly explains a path to understanding the basic principles of human and artificial intelligence, and reverse engineering of the brain in attempting to copy intelligently biological subsystems and brain uploading.

Brain uploading is a complex issue because it comes after some of the other developments. I very often talk about GNR: genetics, nanotechnology and robotics. G is Genetics- basically reverse engineering biology and reprogramming. N is Nanotechnology- basically bringing massively parallel self organizing programmability to matter and energy. R is robotics, which really stands for general artificial intelligence, so-called strong AI. This is AI at the human level, which I pointed out earlier, would not be necessarily beyond human levels, because it would combine the strengths of human and machine intelligence. Anytime you use Google, or in many other examples, the advantages of machine intelligence are evident. But brain uploading goes beyond that.

Image 1

GNR will unfold in that order. We are now in the golden age of reprogramming biology. Nanotechnology is not at that level yet of programming, although I will point out many examples of being able to build complex systems at the molecular level that are performing complex functions already. And in robotics -artificial intelligence- we have many impressive examples of narrow AI. [1] The narrowness will gradually give yield to general intelligence, [2] but that's further off.

And brain uploading is even more daunting. We could create an intelligent entity, one that's convincing at the human level, which is a very broad span of abilities and a broad array of capabilities. Consider human musical intelligence: on one end it's a guy whistling, the other end could be Mozart or Beethoven. There's a very broad range of human intelligence in each area. So for an entity to pass the Turing Test, [3] it's a very fuzzy horizon. Brain uploading is recreating a specific person; I think it's a much more demanding capability.

"the thing that we're going to have to understand is the basic principles of human intelligence to do brain uploading."

You could in theory say, well, we don't have to understand how general intelligence works, we could just methodically copy each of the structures and the final entity would then preserve the capabilities of the original even though we have no overall idea of how that system works.

I think that's very unrealistic. You have to understand what the salient features are that we're copying of an individual, because it's not feasible to recreate an individual at the atomic level. You would have to understand which issues are important and what we're looking for. The thing that we're going to have to understand is the basic principles of human intelligence to do brain uploading. But if I have general AI- human level AI- in 2029, I have brain uploading not until the 2040's, although I do believe it will be feasible.

Image 2: Reader [4]

I want to start by demonstrating my recent invention to make a few points about the invention and how this fits into my own thinking. This is a pocket sized reading machine for the blind. As previously mentioned, I've been in this field for 33 years; I started in 1973. In 1976 I introduced the first print-to-speech reading machine for the blind which had three new technologies of omnifont OCR (optical character recognition), text-to-speech synthesis, and the first flat-bed scanner. OCR was able to intelligently deal with patterns it had never seen before. You could present a whole different type style, which could be quite different than anything it'd seen before and it would be able to intelligently deal with it and recognize letters.

It's the same as what you would do if you see a capital A you've never seen before, because it was looking for the actual geometric features that a capital A has. A has a concave region facing south, it has a triangular shaped loop portion, it has certain connections, and even if the letters were serifed and so on it could intelligently figure that out. It could intelligently separate letters if they're touching one another, or if a piece of the letter was broken off it could reattach it.

If all the AI systems in the world suddenly stopped functioning, our economic infrastructure would grind to a halt. Your bank would cease doing business. Most transportation would be crippled. Most communications would be crippled. This was the not the case a decade ago. Of course, our AI systems are not smart enough—yet-- to organize such a conspiracy. Strong AI. If you understand something in only one way, then you really don't understand it at all. This is because if something goes wrong, you get stuck with a thought that just sits in your mind with nowhere to go.

The reading machine is a good example of the Law of Accelerating Returns. It is a thousand times smaller and lighter than the original reading machine, and a lot less expensive. The computer is thousands of times more powerful in terms of speed and capacity, and depending on how you measure this stuff, it's thousands of times more capable.

> *"…the key to success as an inventor is being able to model and anticipate technology trends, or basically timing."*

It's also a good application of applying technology trends. I got involved in modeling trends of technology 30 years ago because I realized that the key to success as an inventor is being able to model and anticipate technology trends, or basically timing.

We get a lot of business plans, and we fund some projects and do some mentoring. I think 95% of those groups would do exactly what they say if given the funding, and 95% of those projects would fail because the timing is wrong: not all the enabling factors are in place when they're needed.

So I began to be an ardent student of technology trends. Being an engineer I gathered a lot of data and built some mathematical models. I now have a team of ten people that collect data in different areas. We built models of technology in different areas in information technology, not anything else.

Each industry goes through pre and post information eras. An interesting field that's now transforming from a pre-information era to a post one is biology. Biology used to be hit or miss. We'd find something out, for example, "here's something that lowers blood pressure, we don't know why this works but it seems to work."

We've automated that drug discovery process now, but I think 99 percent of the drugs on the market today were created with no clear model of these biological processes. Most drug development now is rational drug design. This term has been out for a while but it is only recently that we have had the tools and genetic knowledge to do it. The Human Genome Project [5] has been completed since 1993.

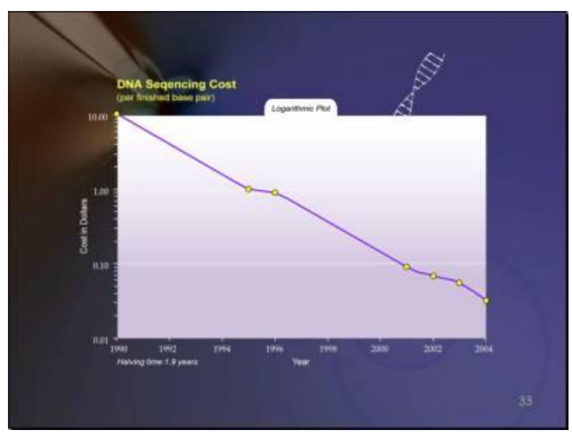

Image 3: DNA Sequencing Cost

And really, we're in the early stages of this, of reprogramming biology. RNA interference can turn genes off: we can send little pieces of RNA that block the messenger RNA expressing the gene.

There are new forms of gene therapy. Gene therapy's been around for a long time but there are new techniques that can put the genetic information in the right place, at the right time. United Therapeutics [6] has an interesting technology that takes the cell in vitro, modifies it in vitro, and then you can inspect it and make sure that it got done correctly. Then you replicate the cell, re-inject those millions of cells back into the organism and they work their way into the right tissues, which in this case are the lungs, and it's cured pulmonary hypertension - a fatal disease. They've done that in animal models. It's now going into human trials.

So there are many examples of being able to reprogram these information processes in biology. You can turn enzymes on and off via the expressions of genes. Pfizer's Torcetrapib [7] turns off one specific enzyme at one specific stage in the development of arteriosclerosis, and the phase 2 trials showed it was very effective. They'll be spending a record $1 billion in phase 3 trials, and there are thousands of developments like this.

And any one of them is unpredictable to address. But the overall impacts of these technologies are remarkably predictable, really. I've been making predictions going forward, I'm not just looking at this data and over-feeding to back data, for a quarter century.

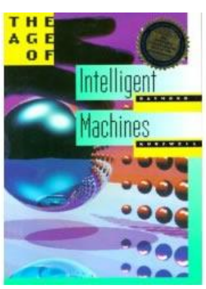

Image 4: "The Age of Intelligent Machines"

The first book I wrote, *The Age of Intelligent Machines* [8], written twenty years ago in the 1980's, had hundreds of predictions about the 1990's and early 2000's, based on these models, which are quite accurate, and I'm not embarrassed about any of them. Some were a few years off. I was generally overly conservative because some outlier- like Novamente [9]- makes some breakthrough and things happened a little bit more quickly, that's what happens.

Image 5

In the mid-80's I predicted the emergence of a worldwide communication network because I saw the ARPANET [10] doubling every year. It was 10,000 nodes going to 20,000 nodes, and it wasn't on anybody's radar map. But doubling every year is multiplying by a thousand in ten years, so I figured in the mid-1990's it would be ten million going to twenty million and then forty to eighty million- it would be a worldwide phenomenon, and that is what we saw. I predicted that for 1995 or 1996, and it happened in 1993 to 1994, and that progression was very predictable.

The genome project was controversial when it was announced. It was not a mainstream project. Mainstream scientists said back in 1989, when we'd collected one-ten-thousandths of the genome, how are you going to do the whole thing in 15 years? And the skeptics were going strong, seven and a half years later, halfway through the 15 year project, saying "I told you this wasn't going to work, and here you are seven and a half years into your 15 year project and you've finished one percent of the project."

But if you double one percent seven more times you get 100 percent and that is exactly what happened. And that has continued since that time. If you apply this to other types of data, to the Proteome Project [11] and so on, you'll see the same kind of exponential progression. When you can measure the information content of a process, we find these very smooth exponential progressions. We find a doubling time of 11, 12, 13 months, depending on what you're measuring.

The same thing's true with the brain. Spatial resolution for 3D volume brain scanning is doubling every year. The amount of data we're collecting on the brain is doubling every year. We're also showing that we can turn this data into working simulations.

> **"The amount of data we're collecting on the brain is doubling every year."**

Let me quickly go through a few dozen examples of just how pervasive and predictable these exponential trends are and how we can use them to make useful predictions. Exponential growth is seductive and surprising. It's smooth, there are no discontinuities. I sometimes refer to the 'knee of the curve.' It's been pointed out that exponential doesn't have a knee of the curve as it's mathematically identical. But it does make a difference in terms of its impact whether you're doubling a thousand nodes on the ARPANET or 100 million nodes on the internet, if it gets to a point where it has real traction and makes a dramatic difference, and I'll show you many examples of that.

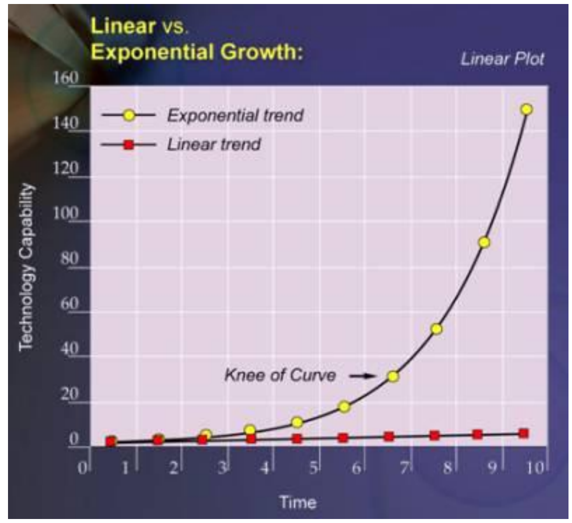

Image 6: Linear vs. Exponential Growth

I have a whole theory as to why that is the case. In fact, the paradigm shift rate is doubling ever year: the rate of technical progress is itself not a constant.

I was at the 50th anniversary of the discovery of DNA and all of us speakers were asked, "Well, what will the next 50 years bring?" And every speaker there, except for Bill Joy and myself, used the last couple of years as a model for the next 50 years, and came up with very tame predictions. I think you could have asked them about 500 years and you would have gotten the same predictions.

Watson [12] said in 50 years we'll have drugs that enable you to eat as much as you want and remain slim. And I said, "Jim, we've done that in animals by using RNA interference to block the fat insulin receptor gene in the fat cells."

These animals ate ravenously and remained slim and it wasn't a fake slimness. They got the benefit of caloric restriction without the restriction, they didn't get heart disease, they didn't get diabetes, and they lived 20 percent longer.

There are a number of pharmaceutical companies rushing to bring fat insulin receptor gene inhibitors to the human market; we'll have that in five to ten years (it would be even less if it wasn't for the FDA) and not 50 years.

We are doubling the paradigm shift rate every decade. That means the next 50 years will have 32 times, approximately, the number of paradigm shifts we've seen in the last 50 years.

These are all log charts. As you go up a chart it represents multiplying some key feature, some key measurements, by factors of ten. A straight line on a log chart is exponential growth.

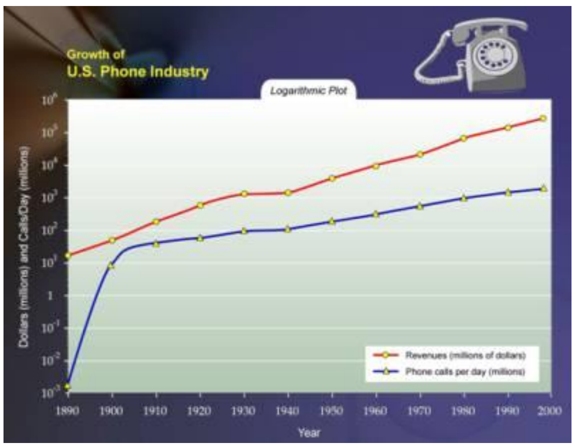

Image 7: Growth of U.S. Phone Industry

Telephones took 50 years to be adopted by a mass audience, a quarter of the U.S. population. Cell phones did that in seven years. These early communication technologies- telephone, radio, and television- took decades to be adopted. The cell phone, the PC, the Web, is measured in a few years time, and that acceleration has continued. Search engines were not used five or six years ago; the word blog wasn't used three years ago; 'social network' wasn't used two years ago. The pace of the paradigm shift continues to accelerate.

Image 8: Mass Use of Inventions

I have a whole theory, which I don't want to spend a lot of time on, that an evolutionary process inherently accelerates because it basically creates a capability and uses that capability to evolve the next stage. You can even see that when we simulate evolution in genetic algorithms on the computer.

This is a double log chart; with the number of years ago that this event took place on the X axis, and how long the paradigm shift took to take place until the next paradigm shift on the Y axis.

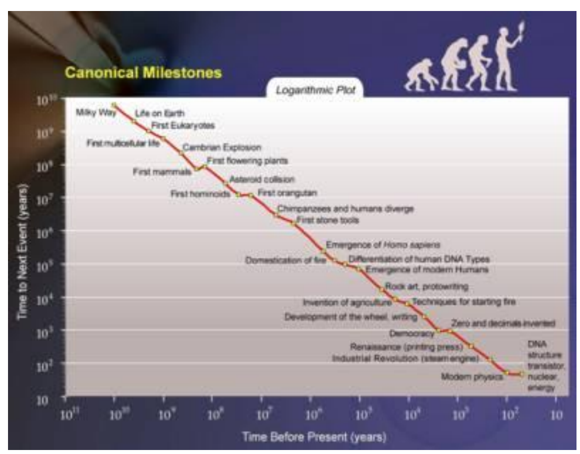

Image 9: Canonical Milestones

The evolution of an information backbone to biology, DNA, took billions of years, although actually RNA came first. And then evolution used DNA to evolve the next stage. It now had an information backbone which it's used ever since, so the Cambrian explosion, when all the body plans of all the animals were evolved, took place 100 times faster, in 10 or 20 million years.

And then those body plans became a mature technology and evolution concentrated on higher cognitive functions and that only took millions of years. And then Homo sapiens evolved in only a few hundred thousand years.

There are only three simple but giant changes that distinguish us from our primate ancestors, involving only a few tens of thousand of bytes of information. One was a larger skull, at the expense of a weaker jaw, so don't get into a biting contest with another primate. Another is our cerebral cortex so we can do "what if" experiments in our mind, abstract reasoning, such as if I took that stone and that stick and tie them together with that twine, I can actually extend the leverage of my arm.

And then we have opposable appendages that work so we could carry out these "what if" experiments and change the environment. It looked like a chimpanzee hand. That is similar but the pivot is down one inch, and basically it just doesn't work very well. They don't have a power grip, they don't have a fine motor coordination, and they're pretty clumsy if you watch them.

So we could effectively change the environment and create tools, and the first stage of that was a little bit faster. We took tens of thousands of years for fire, wheels, stone tools, but then we always use the latest generation of tools to create the next set of tools and so technology evolution has accelerated. And from the straight line on this logarithmic graph, technology evolution is emerging smoothly out of biological evolution.

If we put this on a linear scale it looks like everything has just happened.

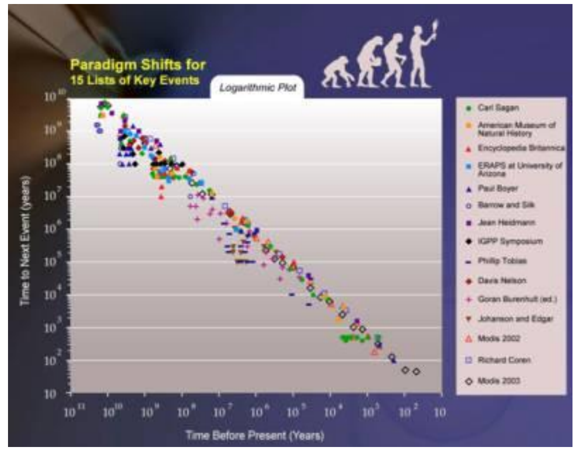

Image 10: Paradigm Shifts

Now some people said Kurzweil would only put points on this graph if they fit on a straight line. To address that I took 14 different lists from 14 different thinkers- Encyclopedia Britannica, American Museum of Natural History, Carl Sagan's cosmic calendar, a dozen other lists. There can be disagreement between the lists: some people think the Cambrian explosion took 25 million years, some people include the ARPANET with the Internet, and there is disagreement when language started.

Therefore there's some spreading of the points. But there's a clear trend line, and clear acceleration. Nobody thinks the Internet took a million years to evolve, nobody thinks the Cambrian explosion happened in ten years. A billion years ago not much happened in a million years.

There's a clear acceleration in this evolutionary process. And it's to the point now where it's very fast, and will continue to accelerate.

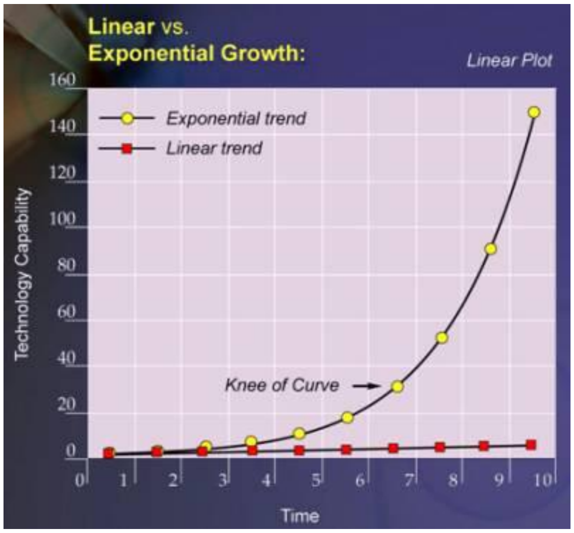

Image 11: **Linear vs. Exponential Growth**

If you compare an exponential to a linear progression, they look the same for a few years. Even if you go to the steep part of the exponential and take a small piece of it, it looks like a straight line. A straight line is a good approximation of an exponential for a short period of time. It's a very bad approximation for a long period of time.

And I think we are hardwired to think in linear terms. I've had countless debates with scientists who take a lineal progression. For example, one was with a neuroscientist who had spent 18 months modeling one ion channel. He's adding up all the other similar nonlinearities in other ion channels and said it would be 100 years.

He's assuming that it's going to take 18 months for every ion channel and every nonlinearity that 50 years from now there's going to be no progress to computers, to scanners, or the ability to model these types of phenomena. He's basically thinking in linear terms.

The ongoing exponential is made up of a series of S-curves. People say exponential growth can't go on forever. Rabbits in Australia ate up all the foliage and then the exponential growth stops. And that's true, any specific paradigm stops when it reaches the limit of that paradigm to provide exponential growth.

Image 12: Growth in Supercomputer Power

What we find in information technology is it leads to research pressure that creates the next paradigm. I'll show you in a moment how this happened five times in computers.

People say there must be some ultimate limit of matter and energy, based on what we know about physics, to support computation and communication, and yes, there are. In Chapter 3 I talk about that. There are ultimate limits but they're not very limited. One cubic inch of nanotube circuitry would be 100 billion times more powerful than the estimate that I used for simulating the several hundred regions of the human brain, and that's not even the ultimate limit.

So, there are limits but they're not very limiting and we will be able to continue the double exponential growth in computation well into this century.

Information technology doubles its power, in terms of price performance, capacity, and bandwidth, every year. When I was at MIT in '67 an eleven million dollar computer took up about five times the size of this room. It was shared by thousands of students and was about one thousand times less powerful than the computer in your cell phone today.

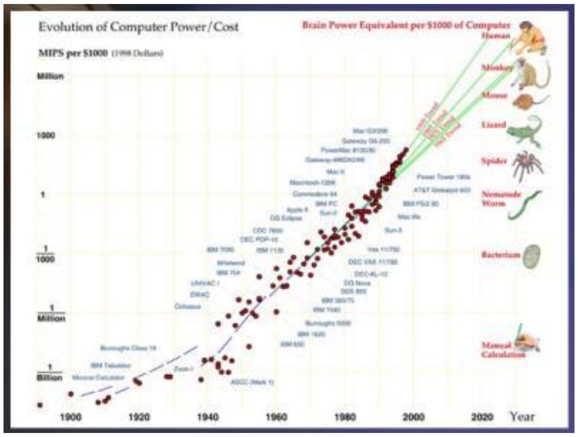

Image 13: Computer Power/Cost

If we look at this chart of over a century of computing, the data processing equipment used in the 1890 census is in the lower left hand corner. Then relays were used to crack the German Enigma code. There's a lot of interesting literature about how the English had to either ignore that information or convince the Germans that they didn't get it through breaking the code but got it from some other way

But they used it without reservation in the Battle of Britain. It enabled the outnumbered RAF to win that battle. In the '50s, they were using vacuum tube based computers. CBS predicted the election of Eisenhower, the first time the networks did that.

They were shrinking vacuums every year, making them smaller and smaller to keep the exponential growth going, and finally, that hit a wall, and that was the end of the shrinking of vacuum tubes. It was not the end of exponential growth of computing. We then went on to transistors, which are not small tubes; it's a whole different approach. Then we've had several decades of integrated circuits and Moore's Law [13], which basically states that the size of transistors on an integrated circuit is shrinking. You can put twice as many in every two years, and they run faster.

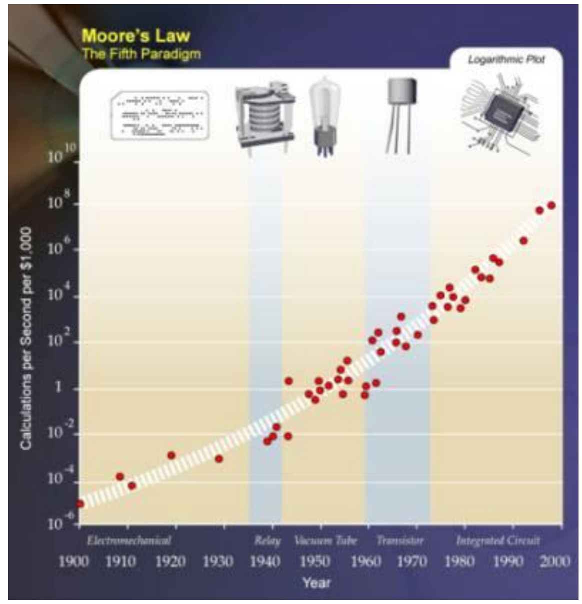

Image 14: Moore's Law

And the end of Moore's law has been predicted for some time; the first predictions were for 2002. Predictions now are for 2022. There's an ITRS roadmap, it's a thick document that's existed for years in the semiconductor industry with tremendous detail of every aspect of every kind of chip, going out fifteen to twenty years in the future; it runs now to 2020. By that time chips will have four nanometer features, twenty carbon atoms, and this has been followed very rigorously for several decades. But by that time, conventional two dimensional chips will be able to establish enough computation to simulate all regions of the brain, and I'll get to that in a moment, for one thousand dollars, based on conventional chips.

If you speak to Intel engineers, they're quite confident of the sixth paradigm, which is three dimensional molecular computing. We live in a three dimensional world, we might as well use the third dimension.

It's very much a mainstream view today. In my last book, *The Age of Spiritual Machines,* which came out in '99 that was a controversial notion. But there's been enough progress on nanotube circuits and other types of molecular circuits, so that it is very much a mainstream view. In fact, there's a nanotube based circuit due to hit the market next year.

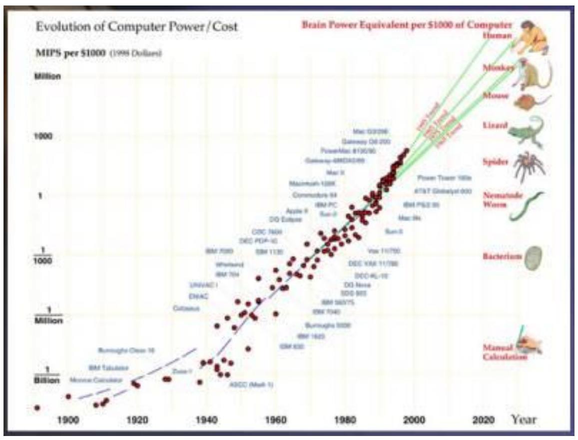

Image 15: Evolution of Computer Power/Cost

These are Hans Moravec's [14] computers, these are different computers than are on the previous chart, but it's the same progression.

Supercomputers are marching along and will hit ten to the sixteenth, which is the amount of computation required to do functional simulation of all regions of the brain, and I have four different derivations of that. The estimates range from ten to the fourteenth to ten to the sixteenth- I use the latter.

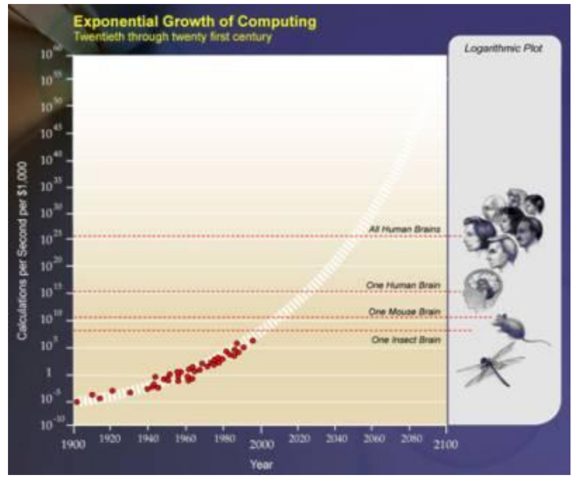

Image 16: Exponential Growth of Computing

And, again, these predictions end up being conservative. Six weeks after the book came out Japan announced two new supercomputer projects to hit that computation level in 2010. And I don't want to dwell on these examples of Moore's law, but look at this chart. This is the price of a transistor. When I was in high school, I could buy a device this big, a relay with support circuitry, equivalent to one transistor, but one thousand times slower, for forty dollars. And I could buy a whole fast transistor for only one dollar in 1968. We could get ten million in 2002 and get 100 million today.

Image 17: Average Transistor Price

We've all heard these very dramatic comparisons, but what I think is interesting, and noteworthy, is how smooth this curve is and how predictable. You think that this would be a very jagged curve. This is the measurement of millions of people's activity in thousands of companies in dozens of countries, with wars and recessions and IPO's and accusations of one country dumping products on another.

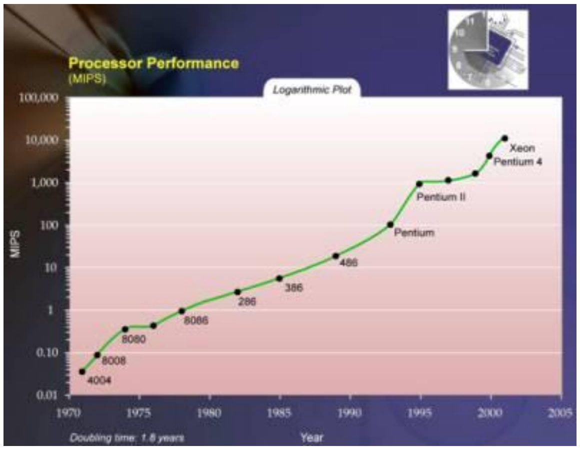

Image 18: Processor Performance

All this unpredictable human behavior and yet you get this very smooth progression: it looks like the outputs of some table top experiment.

You might wonder how can that be, if specific projects like Microsoft and Novamente, (well, Novamente, we're very confident of its success) are unpredictable. How can the overall impact of information technology have such predictable impact?

We see other examples. Thermodynamics [15] was mentioned earlier, where the path of each particle is completely unpredictable, it follows a random walk. In fact that's built into the derivation of the formulas of thermodynamics.

But the overall system of a gas made up of a large number of dynamic, chaotically interacting random particles is very predictable, to a very high degree of precision according to the laws of thermodynamics.

You have a large complex system, and thermodynamics is the first one to be introduced in science. The overall impact can be very predictable despite the fact each element is chaotic and random. We can use that principle in building intelligent systems by having and simulating a lot of processes that are simple and chaotic and unpredictable at the levels of each process. They interact in a dynamic system and create a self-organizing process that shows emergent intelligence.

That's how a lot of these pattern recognition systems work, and there's a system like that in that device there.

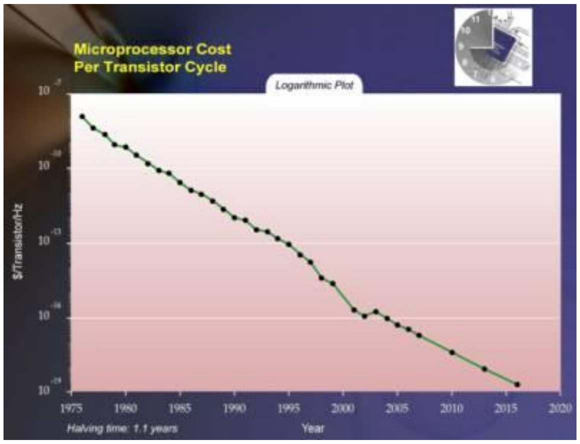

Image 19: Microprocessor Cost

As transistors have gotten smaller, they've gotten faster, pretty smoothly. The cost of a transistor cycle has come down by half, every 1.1 years. If you add all the other forms of innovation- pipe lining, data caching, electronics- it brings it down to one year.

We're doubling price performance in electronics, every year, going back decades. That's 50 percent deflation. It turns out that's also true of lots of other types of information processes, whether you talking about genomic data, proteomic data, brain data, databases of all kinds, all kinds of hardware: 50 percent deflation.

The economists, depending on what week it is, will worry about inflation or deflation. Bernanke [16] and the Fed are back on inflation. They'll say deflation's bad too, although nominally it's a good thing if you buy the same stuff a year later for half the money.

They'll say as information technology ultimately becomes the majority of the economy, which according to my projection it will in the 2020's, it's going to lead to shrinking the economy, because people will increase their consumption but they're not going to really double their consumption every year, year after year. They'll increase it some, but the overall economy as it has to do with information technology will shrink in dollars and that's a bad thing.

That's not what we see. We see people more than double their consumption. There's been 18 percent growth in constant dollars every year for the last 45 years in information technology of every kind despite the fact that you can buy twice as much capability each year for the same amount of money. And the reason for that is as price performance reaches certain capabilities at certain levels whole new applications just explode. People didn't buy iPods for $10,000 fifteen years ago, which is what it would have cost. People didn't buy thousand dollar genomes, which is what it will cost pretty soon.

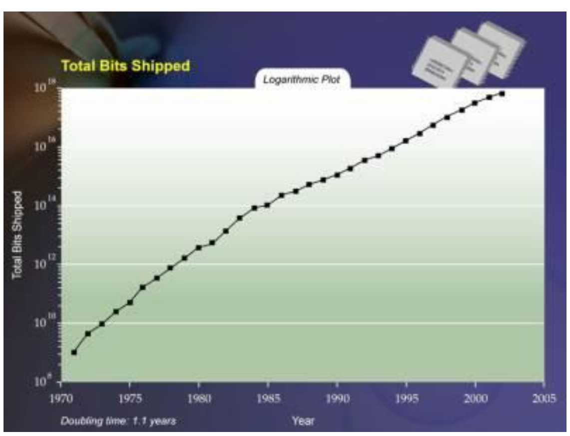

Image 20

As new applications open up, they explode on the horizon and this is the only thing that is growing, that is fueling economic growth. The reason that economic growth is increasing is because more of the economy is shifting to this information economy which has 18 percent growth versus a non-information economy which has negative growth.

Image 21

Magnetic data storage: I just put that up because it has different engineers and different companies but the same expediential progression. We've mentioned biotechnology: sequencing costs were $10 per base pair in 1990, it's a penny today. This slope on this log graph represents a doubling of genetic data every year.

And we see this in many other examples of biology. We're also gaining means to reprogram these biological processes and that will lead to a profound revolution in our ability to reprogram biological processes away from things like disease, cancer, heart disease; and away from aging processes which are themselves information processes which we're beginning to understand. I talk about that in another book, *Fantastic Voyage*. [17]

Communication technology: we won't dwell on this but we see the same progression. Here's that graph of the Internet. I had a little piece about this in the 1980's when it was the

ARPANET, I predicted that it will be a worldwide phenomena.

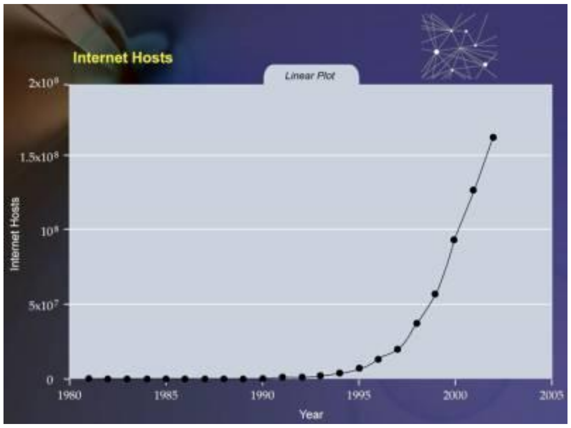

Image 22

This is the same graph on a linear scale. To the casual observer it looks like the Internet just came out of nowhere in the mid 1990's. The first reference to the World Wide Web in the New York Times was in 1993, December of '93. But you could see it coming if you look at the exponential progression, which you can see on a logarithmic graph.

We're shrinking technology at a predictable rate. That's obvious in electronics, that's the whole force behind Moore's Law. It's also true of mechanical systems: it's a factor of 100 per 3D volume per decade. A doubling every year is multiplying by 1000 in ten years of price performance, capacity and bandwidth of information technology. That's a billion in thirty years, and there's also a slow second level of exponential growth, so it's really 25 years. In 25 years we'll also be shrinking the size of electronic and mechanical technologies by a factor of 100,000.

So, if you look at the impressive things we can do today and factor in this billion fold increase in capacity and one hundred-thousand fold decrease in size, it gives some idea about what will be feasible.

These illustrations, which have now been simulated, which are in Eric Drexler's [18] book, "Engine's of Creation" of 20 years ago (which founded nanotechnology [19]) are now being built.

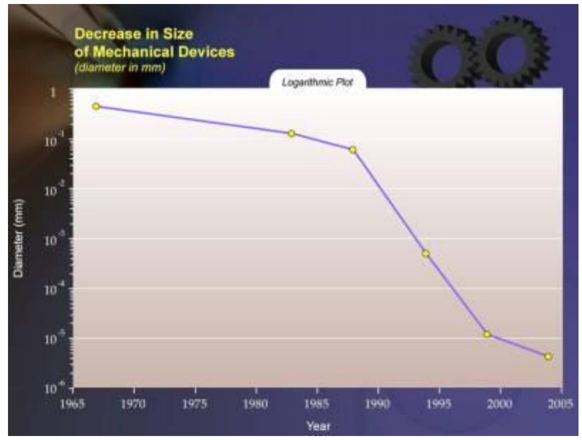

Image 23

I give several dozen examples in my book of molecular systems, in particular those being used in biology. For example, one scientist actually cured Type 1 diabetes in rats with a blood cell sized device. It's quite intricate. It lets insulin out in a controlled fashion. It blocks antibodies and can stop Type 1 diabetes as an autoimmune disorder. It works, and they're gearing up for human trials.

At MIT and at the University of Rochester they have a little device that's blood cell sized that can detect antigens, specifically cancer cells. Once it detects them it latches onto the cell, burrows inside the cells, detects it's inside the cells, and releases toxins. So, it's a fairly elaborate procedure. It works quite reliably in-vitro and they're now gearing up for animal trials.

So, this is what's feasible today. There are four major conferences on this kind of stuff, bio-MEMS.

People say 'nanobot' sounds very futuristic, but blood cell sized devices can actually do things inside the human body.

Image 24: Respirocyte

They're not putting it in the human body yet, although we have pea size devices and bb sized devices we are putting in the human body, but we are putting them in animals and they're performing therapeutic functions today.

If you apply the billion-fold increase I've spoken about, then very sophisticated nanobots will be feasible in the decades ahead. This is a robotic red blood cell that was designed by Robert Freitas [20], and we do understand how red blood cells work.

It brings up an interesting observation about biology. It's very intricate but it's also very suboptimal in terms of its capability, because biology got stuck with some earlier substance like building everything out of proteins, which are interesting substances but a limited class of materials.

We find out we can re-engineer these systems to be far more powerful. If you replace ten percent of your red blood cells with these respirocytes, you can do an Olympic sprint for 15 minutes without taking a breath, or sit at the bottom of your pool for four hours.

And a very dramatic example is in fact the brain, which transmits information through signals that are electrochemical and travel at a few hundred feet per second. That's a million times slower than electronics.

For another example, the computation that takes place in the interneuronal connections such as dendrites are computed at 200 calculations per second. They're not exactly calculations, they're digitally controlled analog transactions, but there are about 200 of them typically, as the reset time is about five to ten milliseconds. That's a million times slower than electronics.

Ultimately we can reengineer these things to become far more powerful. If we follow this doubly exponential progression, $1,000 of computation will equal ten to the fourteen calculations per second, which I maintain is enough to do functional emulation-simulation of all regions of the brain- for $1,000. You don't even have to believe in three dimensional molecular computing, although there are relatively few people who are knowledgeable and express skepticism about that anymore because they're now working in labs.

But even if you are skeptical, two dimensional conventional photolithograph [21] chips on the ITRS [22] roadmap by around 2018 or 2019 will provide this level of computation for $1,000.

Image 25

I pose reverse engineering the brain as more of an existence proof, that we'll ultimately understand the principles of the operation of the human brain. I agree with Ben Goertzel [23] that we can achieve human level AI even if we never looked at the brain at all. Largely that's been the case in AI (Artificial Intelligence) because up until recently we couldn't see in the brain, other than with these fuzzy fMRI [24] images, which aren't sufficient to reverse engineer anything.

That's changing. I talk about five or six emerging scanning technologies where we can see, in-vivo individual interneural connections, see them firing in real time and begin to develop models of how a real brain works, combining them with models of individual neurons that can be tested in-vitro.

We're building very sophisticated model and simulations of the brain like the IBM Project on the cerebral cortex.

I make the case, although I don't have time to go into details, in chapter 4 of the book, that at the progress we're making (and you can measure certain aspects of this, like the spatial resolution of brain scanning, the amount of data, the number of simulations we're doing of specific regions) it's really conservative to say within 20 years we'll have detailed models and simulations of the different regions of the brain.

I do think that human level intelligence will come from AI research as it has, which is just trying to emulate human levels of intelligence. I believe narrow AI is getting less narrow, and combining multiple methodologies, and that will continue through hundreds of steps until it's as broad as human intelligence, which is not completely general either.

We will get some tips on reverse engineering the human brain. I just saw Tomaso Poggio [25] at one of these celebrations at of the 50th anniversary of AI- he's head of vision research at MIT. He's been a skeptic of finding out anything from the brain because we haven't been, because we haven't been able to see in the brain. He just came up to me and said, "You know, you were right, we've just gotten these fantastic insights from brain sciences into the coding of the early stages of the visual cortex and we applied that to our computer simulations. And at first we didn't understand why it works and we thought about it some more and now we do."

It was a real giant leap in performance. We had the same experience in speech recognition in that we got early stages of the auditory cortex which were reverse engineered. And those were also counter intuitive, but we applied them to the transformations in our speech recognition program and got a big leap in performance.

So there is some useful information to be found in reverse engineering the human brain. It will expand the AI tool kit. And all of this is gearing up exponentially, the spatial resolution and the amount of data.

We're also showing that we can simulate regions of the human brain and these simulations are a completely different level than trying to simulate at the molecular level of neurons, for example, although that's also very valuable.

You do need to understand one level of science before you can go up to the next. So you have to understand physics before you can understand chemistry. But chemistry has it own rules. You need to understand chemistry before you go up to levels of biology, but biology has its own rules.

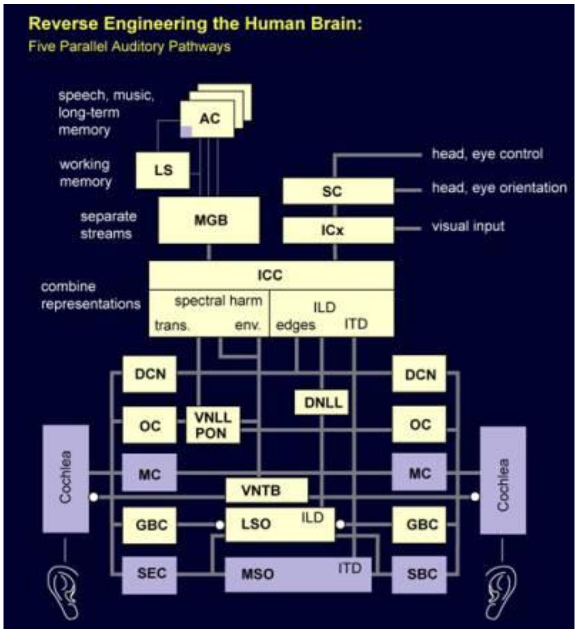

Image 26

This is a block diagram of 12 regions of the auditory cortex, and as mentioned earlier, they've applied sophisticated psychoacoustic tests to the simulation and get similar results as applying the same tests to human auditory perception.

This doesn't prove the model and simulation is perfect but it does show it's pretty good.

There's a simulation of the cerebellum, which is comprised of more than half the neurons in the brain. The simulation is not of that many neurons, but it can easily be scaled up. And then they apply skill formation tests- the cerebellum is responsible for skill formation- and get similar results.

This brings up an important issue which I think has come up a number of times this morning, which is how complicated is the human brain? We need to distinguish between the human brain as it's manifested versus the design of the human brain.

Because if you look at the human brain as it's manifested in a mature person, it's vastly complex. We can argue about the amount of complexity, but there are a hundred billion neurons and thousands of connections per neuron, so that's a hundred trillion connections (these are all order of magnitude), and thousands of ion channels for a connection and so forth.

That's a vast amount of complexity. I've estimated it would take thousands of trillions of bytes of information to characterize the state of the human brain: that might be low.

But the design of the human brain is a billion times simpler than that. How do we know that? Well, the design of the human brain and body is in the genome, and the genome doesn't have very much information in it, relatively speaking. There are three billion rungs, or six billion bits, or a hundred billion bytes, but it's replete with redundancies.

One system called ALU [26] is a lengthy sequence is repeated 300,000 times, and there are many other sequences like that. I show in the book how you can, using lossless compression and knowledge of the structure of the genome, compress the genome, including the junk- the so called junk which isn't junk because it effects gene expression.

We can compress the whole genome into 50 to 100 million bytes. This is not simple either, but it is less than Microsoft Word and it's a level of complexity that we can manage.

Image 27

They might say how can that be? How could something that is 50 to 100 million bytes then describe an entity that's a billion or more times complex than that? Well, the relationship of the genome to the brain is the relationship of a description of a recursive probabilistic fractal to the expression of that fractal.

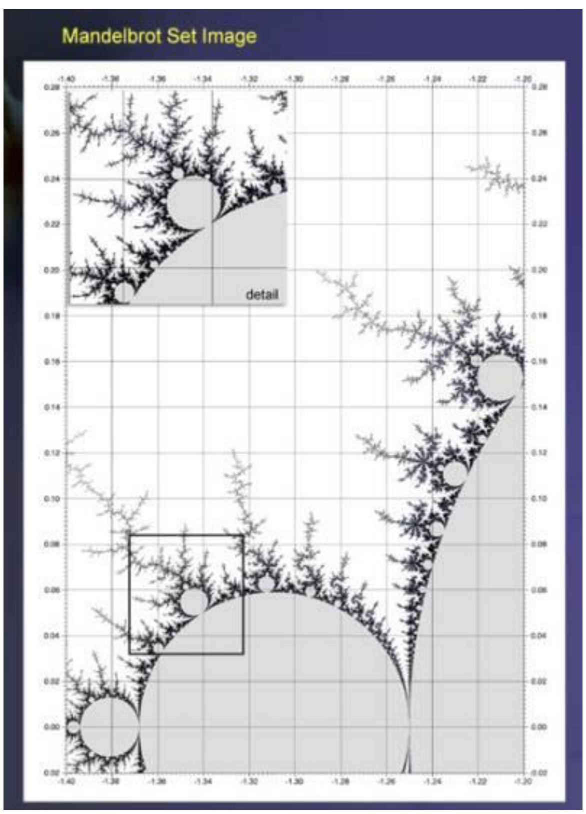

Image 28: Mandelbrot Set Image

You've all seen this image. In fact it's on the cover of a book called *Complexity* [27]. It's a very complex image: it's the Mandelbrot set. And how complex is it? Well, the manifestation of it could be billions of bytes depending on what resolution you present it in. As you zoom in on small regions, you get more complexity. There's complexity within complexity.

But how complex is the design of the Mandelbrot set? It is six letters long. And that is a good example of the relationship of the genome to the brain.

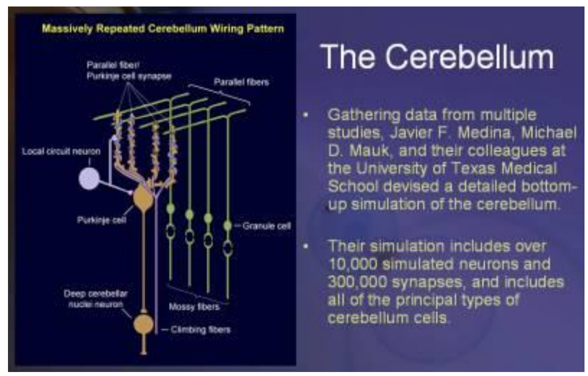

Image 29

Take the cerebellum, for example. If I asked you to reverse engineer the cerebellum and I show it to you and you see the trillions, literally, of these deeply intertwined connections you'd say this is a vast amount of complexity. But it has been reverse engineered. There are only a few genes that control it and there are only a few tens of thousands of bytes in the design.

The genome says there are four different types of neurons, that one module is organized like this, and now repeat ten billion times, and add a certain amount of random variations within the following constraints with each repetition. That's a summary of what the genome says about the cerebellum.

There's tens of thousands of bytes and it creates this system. It looks very complex.

It's stochastic, which means large elements of it are random, but it's self-organizing, and that's the key to its brilliant design. As a child learns to walk and to talk and to catch a fly ball, it gets filled up with meaningful information as it self-organizes.

While the design is not simple it's a level of complexity we can and have figured out, and the design, as it's manifested, is much more complex.

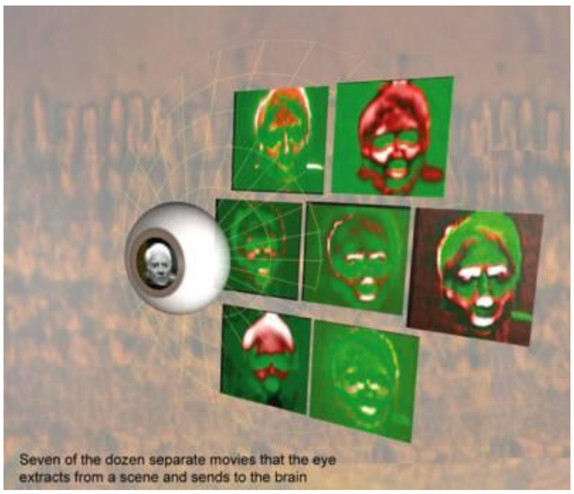

Image 30

This is interesting. This is the early visual cortex that's been reverse engineered. It's interesting in that we receive a lot of data into our eyes, it's a high bandwidth of information going into the early part of the optical system.

But in the early visual cortex, it gets reduced to seven very low resolutions, very simple movies, where each movie is looking at certain types of features. It's sending these very sparsely coded seven different movies. And that's all that the rest of the brain ever gets.

We essentially hallucinate the world, because we actually don't see it. We only see these very low resolution movies, but we're able to have the illusion that we see a high resolution image.

But these insights have enabled simulation of the visual system to perform with much greater performance because there's a very clever coding system.

And to another point, modeling systems at the right level is very useful to simulate one neuron, which is vastly complex. And as we go up to higher levels, and look at the function of the whole region, like the cerebellum, or the hippocampus, those are not simple either.

But we find that they're simpler than a single neuron, just as the pancreas is simpler than the pancreatic islet cell. You might say the pancreas is different because the pancreas is just a glob of cells whereas there's intricate circuitry in the brain. Still the same principles apply. As we find out what these regions are doing to transform information, the models are not simple but they're much simpler than the lower level description of the individual neurons.

All this is driving economic growth. I don't want to dwell on this, but this is an interesting chart on e-commerce growing smoothly and exponentially, and it's now a trillion dollars. You might say wait a second, wasn't there a boom and a bust in e-commerce? How come we don't see that anywhere here?

That turned out to be just a capital market phenomenon. Wall Street looked at the internet and said "wow, this is going to change everything. This is going to turn every business model on its head," and all these valued soared, in 1999 and 2000.

And that is true. It is going to do that, and it's going to do that exponentially, but exponentially isn't instantly. So Wall Street came back a couple of years later and said, "it hasn't changed every business model; I guess we were wrong," and all the values went the other way.

Meanwhile, it has progressed exponentially. Three years ago people were saying you can't make money in internet advertising. Now you've got one company with 99% of its revenues from internet advertising and it's worth $100 billion.

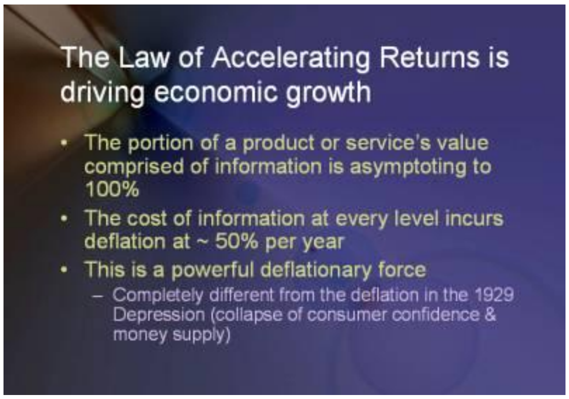

Image 31

There's been smooth exponential progression in this technology, and that boom/bust phenomena is an accurate harbinger of what ultimately is a true revolution.

We had ours in AI in the '80s. There was one in telecommunications, and Martine Rothblatt [28] with her exquisite timing built up a very successful telecommunications company. She got out at the boom then went into biotech, which again was very good timing. We had it with the railroads, this boom bust phenomena.

Let me show you one other example of self-organizing in our own technology. We took some of our speech recognition which we worked on for many years. We developed the first large vocabulary speech recognition system. There was also speech synthesis. We took contemporary versions of that and put them together with language translation. That's an area that's progressed a great deal: it's not so narrow. It's using pattern recognition. The early systems used rules, as it's been pointed out, and now they're using self organizing pattern recognition.

When I was at Google they showed me a system that can translate from English to Arabic and Arabic to English. They have these very large corpuses of translated text, Rosetta stone texts [29], and the system created its own rules. Nobody on the team spoke a word of Arabic, yet the system compared equally to human professional translators and won the DARPA competition.

I put together a translating telephone. I've used this to communicate with people who don't speak English, and we were able to communicate just fine. Although people misunderstand each other even if they're speaking the same language. This will be a routine service of your cell phone early in the next decade.

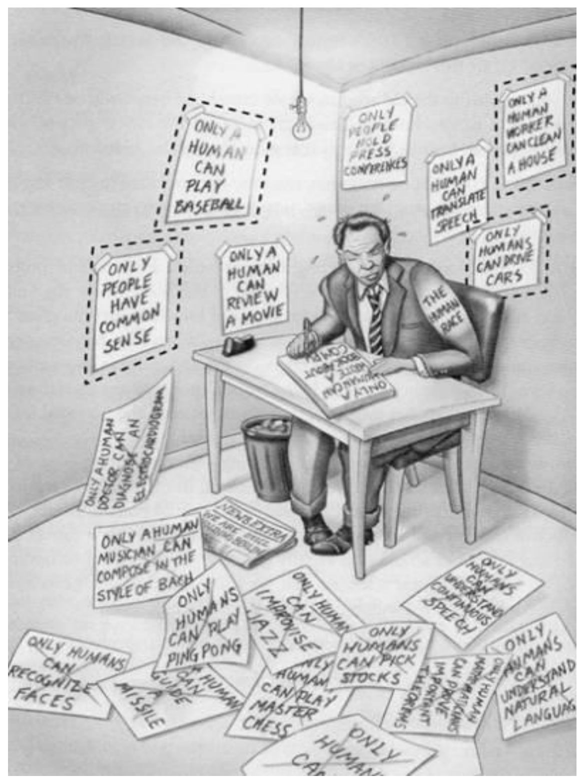

Image 32

I don't want to dwell on this, but I will make a point about this picture, which was in *The Age of Spiritual Machines* in 1999 without the dotted lines. I put it in this book, *The*

Singularity Is Near. This man, symbolizing the human race, is pointing out that computers can't do something or other, and a lot of failed examples of that are on the floor, and there are still a few examples on the wall. I put dotted lines around the things I think will fall off the wall within the next decade.

I was at George Gilder's [30] conference this fall, and Gilder chided me on the automobile driving example. He pointed out that they just had this competition where all the cars ran off the road very quickly. I said because exponential growth progress tends to zoom before the very end, in the next year or two I thought the cars will make that. Five weeks later, Sebastian Thrun's [31] car and four other cars made this very difficult, 300 mile route, a circuitous route.

Image 33

In 2029, $1,000 of computation will be a thousand times more powerful than this figure, ten to the sixteenth, which I believe is sufficient for functional emulations. This is not simulating every neuron, but I don't think that's necessary. We'll have completed the reverse engineering of the human brain in terms of the salient principles of operations. We'll have simulations of all the different regions.

And I believe a computer will pass the Turing Test and I believe the progress will come from general AI and from reverse engineering the human brain working together. All of them are going to provide a handsome expansion of our AI tool kit.

That will necessarily be more powerful than human intelligence because it will combine the natural advantages of the machine with the powers of human pattern recognition. But it is not an alien invasion. I think an important issue with uploading is to understand that we're going to merge with this technology. You're not going to be able to walk into a room and say "humans on the left side of the room and machines on the right."

We're getting closer to our technology. We put it in our pockets now. Martine [Rothblatt] and I are working on a project to put it into your clothing to monitor your health. They're already beginning to go into our bodies and brains in this very early stage, but that's going to accelerate. We have nanobots [32] that are quite sophisticated already in terms of being able to perform therapeutic functions.

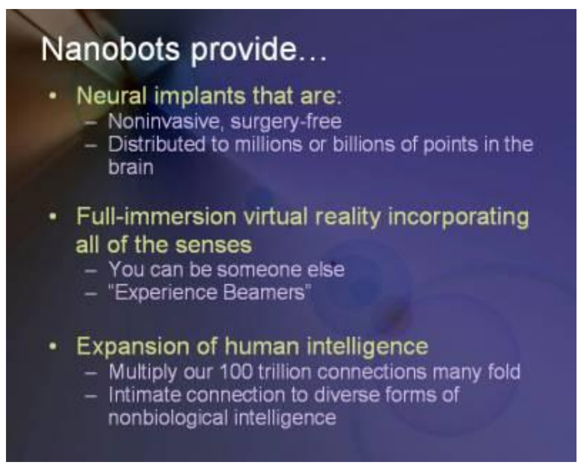

Image 34

If you add a billion fold increase of capability and a hundred thousand decrease in size to what we can do today, then by 2029 these will be quite sophisticated and they'll be going inside our human body, keeping us healthy from inside. They'll also be going inside our human brain, interacting with our biological neurons, creating full immersive virtual reality from within the nervous system, expanding our memory and our cognitive capabilities.

You're going to be hard pressed to find a human who isn't enhanced in some way by technology. And once it gets a foothold in our bodies and brains it's going to be subject to this law of Accelerating Returns, which is just doubling our capabilities, not because it's self-replicated, but just because that's the nature of information technology.

We have ten to the twenty-sixth power calculations per second in the whole human race today, by my estimate. Fifty years from now that number's going to be ten to the twenty sixth power. Biological capability is pretty much fixed. The population might increase but that's not going to change that figure.

The non-biological intelligence is going to intersect with that in the 2020s, and ultimately will become far more powerful.

Most of the action in our brains at some point in the 2030's or 2040's is going to be non-biological. And obviously non-biological intelligence is much more subject to uploading.

It's not going to be the case that in 2045 we will have a fully biological human that we just try to upload. People will already be different by that time.

Average Life Expectancy (Years)

Cro Magnon	18
Ancient Egypt	25
1400 Europe	30
1800 Europe & U.S.	37
1900 U.S.	48
2002 U.S.	78

Image 35

Human life expectancy is increasing, and I don't want to dwell on that, but one of the implications is that it'll go into hyper drive once we get into the mature process of reprogramming our biology. Then we'll also be able to expand beyond the limitations of biology with this intimate merger with nanotechnology.

I just touched for 60 seconds on the criticism that exponential trends can't go on forever. I'll also address a big criticism, that hardware is definitely increasing exponentially but software is not. I've got many examples, not only individual anecdotal examples on how software is improving, but on that there are many different ways to measure it. I don't have much time left but I'll mention one that's very good.

Take chess machines. People think that's just hardware driven - just computers getting more powerful therefore they can deal with the exponential combinatorial explosion of move-countermove sequences to greater effectiveness.

But there is a pattern recognition decision that needs to be made, because you can't expand a tree of move-countermove possibilities forever. At some point you have to decide to cut off one of those trees of expansion. And, in fact, you have to do that constantly.

That's a pattern recognition task. Now, you can do that simply by counting up piece values, and that's what was done in the early stages, but the chess machines have applied more and more sophisticated pattern recognition to that decision.

Chess machines got better and better. I predicted that a computer would take the World Chess Championship in 1998. I did that in the 80's: it happened in '97. Yet that was basically from hardware improvement, because these chess supercomputers were getting more powerful each year, adding 40 points to chess score.

Chess score is a very wonderful quantitative logarithmic measure of intelligence within that domain. IBM cancelled the supercomputer in '97 because there was much less interest in chess, so then we had this big drop in computation as we switched to using PC's.

Recently, Deep Fritz [33] running out on a few PC's had the same amount of computation as Deep Thought [34], a chess supercomputer in 1988. They both could do four million board positions per second- they had equal hardware.

The difference between them was not because one had more hardware capability than the other, the PC's today are equal to the supercomputer of 18 years ago.

Yet Deep Fritz has 400 points more than Deep Thought. Another way to think about it is that it matched Deep Blue [35], despite the fact that it has only one percent as much hardware capability.

So, if you follow that, it's a quantitative demonstration of smooth exponential progression in the quality of the software in this pattern recognition domain.

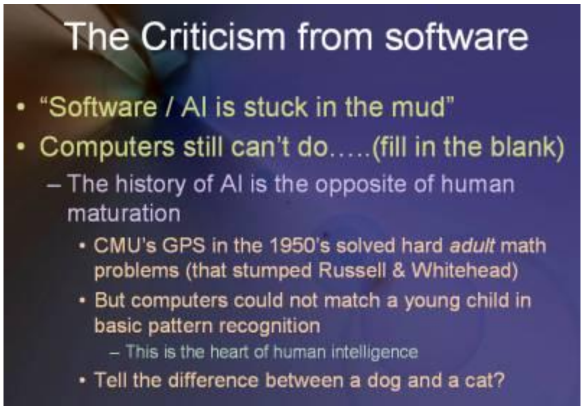

age 36

I have many other examples like that in the book. Software is not stuck in the mud; there's many ways to measure it; and you can also look at things like ability to simulate the brain and reverse engineer aspects of it which is also scaling up exponentially.

As I mentioned earlier, we will have to understand how the brain works. We could in theory say you don't have to do that in order to do brain uploading; you could just blindly copy every little module at a micro level.

I don't think that's feasible. You need to know what is going on to know what to look for. We're finding in some neurons it's the concentration of certain neurotransmitters that matters, in other ones it's the location of certain neurotransmitters- whether they're near a certain ion channel or not.

We need to understand how these mechanisms are working to be able to know what to reverse engineer, to the extent that we're trying to copy intelligently biological subsystems.

Endnotes

1. Narrow AI [Artificial Intelligence] - a kind of AI system that displays Specialized Intelligence in some domain. Most AI software in history, and under development, fits into this category. However, this Mind Ontology

was created to help move toward the development of Narrow AI's antonym, Artificial General Intelligence.
Many Narrow AI systems are based on the use of individual Learning Algorithms, customized for particular domains.
agiri.org/wiki/Narrow_AI November 13, 2007 1:39PM EST

2. General Intelligence - controversial construct used in the field of psychology (*see also psychometrics*) to quantify what is common to the scores of all intelligence tests.
en.wikipedia.org/wiki/General_intelligence_factor November 13, 2007 1:46PM EST

3. Turing Test – a proposal for a test of a machine's capability to demonstrate intelligence. Described by Alan Turing in the 1950 paper "Computing machinery and intelligence," it proceeds as follows: a human judge engages in a natural language conversation
with one human and one machine, each of which try to appear human; if the judge cannot reliably tell which is which, then the machine is said to pass the test.
en.wikipedia.org/wiki/Turing_test November 13, 2007 1:50PM EST

4. K-NFB Reading Technology, Inc. presents the ultimate solution to your everyday reading needs—the Kurzweil-National Federation of the Blind Reader. This portable device will revolutionize the way you deal with the avalanche of print you confront each day.
www.knfbreader.com/ November 7, 2007 2:50PM EST

5. Human Genome Project - Completed in 2003, the Human Genome Project (HGP) was a 13-year project coordinated by the U.S. Department of Energy and the National Institutes of Health. During the early years of the HGP, the Wellcome Trust (U.K.) became a major partner; additional contributions came from Japan, France, Germany, China, and others. See our history page for more information.
www.ornl.gov/sci/techresources/Human_Genome/home.shtml November 14, 2007 3:49PM EST

6. United Therapeutics - a biotechnology company focused on the development and commercialization of unique products for patients with chronic and life-threatening cardiovascular, cancer and infectious diseases. In these segments, United Therapeutics is actively developing four technology platforms: Prostacyclin Analogs, Immunotherapeutic Monoclonal Antibodies, Glycobiology, and Telemedicine.
http://www.unither.com/utcabout.asp
November 13, 2007 3:12PM EST

7. Torcetrapib - (CP-529414, Pfizer) was a drug being developed to treat Hypercholesterolemia (elevated cholesterol levels) and prevent cardiovascular disease. Its development was halted in 2006 when phase III studies showed excessive mortality in the treatment group receiving a combination of atorvastatin and the study group.
http://en.wikipedia.org/wiki/Torcetrapib
November 13, 2007 3:15PM EST

8. The Age of Intelligent Machines - inventor and visionary computer scientist Raymond Kurzweil probes the past, present, and future of artificial intelligence, from its earliest philosophical and mathematical roots to tantalizing glimpses of 21st-century machines with superior intelligence and truly prodigious speed and memory. Generously illustrated and easily accessible to the non-specialist, this book provides the background needed for
a full understanding of the enormous scientific potential represented by intelligent machines as well as their equally profound philosophic, economic, and social implications.

www.amazon.ca
November 13, 2007 3:27PM EST

9. Novamente - Novamente supplies software products and services to power intelligent virtual agents for computer games and simulations.
http://www.novamente.net/

10. ARPANET - developed by DARPA of the United States Department of Defense, was the world's first operational packet switching network, and the predecessor of the global Internet.
http://en.wikipedia.org/wiki/ARPANET
November 13, 2007 3:33PM EST

11. Proteome Project - The Human Proteome Folding project will provide scientists with data that predicts the shape of a very large number of human proteins. These predictions will give scientists the clues they need to identify the biological functions of individual proteins within the human body.
www.worldcommunitygrid.org/projects_showcase/
November 13, 2007 3:36PM EST

12. James Dewey Watson - (born April 6, 1928) is an American molecular biologist, best known as one of the co-discoverers of the structure of DNA. Watson, Francis Crick, and Maurice Wilkins were awarded the 1962 Nobel Prize in Physiology or Medicine "for their discoveries concerning the molecular structure of nucleic acids and its significance for information transfer in living material."
http://en.wikipedia.org/wiki/James_D._Watson
November 13, 2007 3:39PM EST

13. Moore's Law - describes an important trend in the history of computer hardware: that the number of transistors that can be inexpensively placed on an integrated circuit is increasing exponentially, doubling approximately every two years.
http://en.wikipedia.org/wiki/Moore's_law
November 13, 2007 3:41PM EST

14. Hans Moravec (born November 30, 1948 in Austria) is a research professor at the Robotics Institute (Carnegie Mellon) of Carnegie Mellon University. He is known for his work on robotics, artificial intelligence, and writings on the impact of technology. Moravec also is a futurist with many of his publications and predictions focusing on transhumanism. Moravec developed techniques in machine vision for determining the region of interest (ROI) in a scene.
http://en.wikipedia.org/wiki/Hans_Moravec
November 13, 2007 3:46PM EST

15. Thermodynamics - (from the Greek θερμη, *therme,* meaning "heat" and δυναμις,*dynamis,* meaning "power") is a branch of physics that studies the effects of changes in
temperature, pressure, and volume on physical systems at the macroscopic scale by analyzing the collective motion of their particles using statistics.
http://en.wikipedia.org/wiki/Thermodynamics
November 13, 2007 3 :54PM EST

16. Ben Shalom Bernanke - (born December 13, 1953) is an American economist and current Chairman of the Board of Governors of the United States Federal Reserve.
http://en.wikipedia.org/wiki/Ben_Bernanke
November 13, 2007 3:58PM EST

17. Fantastic Voyage: Live Long Enough to Live Forever - In Fantastic Voyage, high-tech visionary Ray Kurzweil teams up with life-extension expert Terry Grossman, M.D., to consider the awesome benefits to human health and longevity promised by the leading edge of medical science--and what you can do today to take full advantage of these startling advances.

http://www.fantastic-voyage.net/ November 13, 2007 4:28PM EST

18. K. Eric Drexler - a researcher and author whose work focuses on advanced nanotechnologies and directions for current research. His 1981 paper in the *Proceedings of the National Academy of Sciences* established fundamental principles of molecular design, protein engineering, and productive nanosystems. Drexler's research in this field has been the basis for numerous journal articles and for books including *Engines ofCreation: The Coming Era of Nanotechnology* (written for a general audience) and*Nanosystems: Molecular Machinery, Manufacturing, and Computation* (a quantitative, physics-based analysis). He helped lead development of the 2007 *Technology Roadmapfor Productive Nanosystems,* a project managed by Battelle and hosted by several of the U.S. National Laboratories.
http://www.e-drexler.com/ November 13, 2007
4:36PM EST

19. Nanotechnology - refers broadly to a field of applied science and technology whose unifying theme is the control of matter on the atomic and molecular scale, normally 1 to 100 nanometers, and the fabrication of devices within that size range.
http://en.wikipedia.org/wiki/Nanotechnology
November 13, 2007 4:40PM EST

20. Robert A. Freitas Jr., J.D. - published the first detailed technical design study of a medical nanorobot ever published in a peer-reviewed mainstream biomedical journal and is the author of Nanomedicine, the first book-length technical discussion of the medical applications of nanotechnology and medical nanorobotics.
http://www.rfreitas.com/ November 13, 2007 4:43PM EST

21. Photolithography (also **optical lithography**) is a process used in microfabrication to selectively remove parts of a thin film (or the bulk of a substrate). It uses light to transfer a geometric pattern from a photomask to a light-sensitive chemical (photoresist, or simply "resist") on the substrate. A series of chemical treatments then engraves the exposure pattern into the material underneath the photoresist. In a complex integrated circuit (for example, modern CMOS), a wafer will go through the photolithographic cycle up to 50 times.
http://en.wikipedia.org/wiki/Photolithography
November 13, 2007 4:49PM EST

22. ITRS - the fifteen-year assessment of the semiconductor industry's future technology requirements. These future needs drive present-day strategies for world-wide research and development among manufacturers' research facilities, universities, and national labs.
http://www.itrs.net/ November 13, 2007 4:51PM EST

23. Ben Goertzel - 20+ yrs in AI R&D and commercialization. Former CTO of 120+ employee, thinking machine company, Webmind. PhD in mathematics from Temple University. Held several university positions in mathematics, computer science, and psychology, in the US, New Zealand and Australia. Author of 70+ research papers, journalistic articles and 8 scholarly books dealing with topics in cognitive sciences and futurism. Principle architect of the Novamente Cognition Engine.

24. fMRI - Functional magnetic resonance imaging is the use of MRI to measure the haemodynamic response related to neural activity in the brain or spinal cord of humans or other animals. It is one of the most recently developed forms of neuroimaging.

http://en.wikipedia.org/ November 13, 2007 5:03PM EST

25. Tomaso A. Poggio, Ph.D. - the Eugene McDermott Professor at the Department of Brain & Cognitive Sciences; Director, Center for Biological & Computational Learning; Member for the last 20 years of CSAIL/the Artificial Intelligence Laboratory at MIT; and, since 2000, member of the faculty of the McGovern Institute for Brain Research. http://www.csail.mit.edu/biographies/ November 14, 2007 3:19PM EST

26. ALU - An **Alu sequence** is a short stretch of DNA originally characterized by the action of the Alu restriction endonuclease. Alu sequences of different kinds occur in large numbers in primate genomes. In fact, Alu sequences are the most abundant mobile elements in the human genome. http://en.wikipedia.org/wiki/Alu_sequence November 14, 2007 4:26PM EST

27. *Complexity* – "Chaos and Life: Complexity and Order in Evolution and Thought" http://www.columbia.edu/ November 14, 2007 4:40PM EST

28. Martine Rothblatt - Martine Aliana Rothblatt Ph.D, MBA, J.D. is a lawyer, author, and entrepreneur. Dr. Rothblatt graduated from UCLA with a combined law and MBA degree in 1981, then began work in Washington DC, first in the field of communication satellite law, and eventually in life sciences projects like the Human Genome Project. http://en.wikipedia.org/wiki/Martine_Rothblatt November 14, 2007 4:45PM EST

29. The Rosetta stone text - The Rosetta Stone is a black basalt slab bearing an inscription dating from the year 196 BC. It was the crucial key to the deciphering of Egyptian hieroglyphs, and the foundation of modern Egyptology. The stone was discovered in 1799 by the French troops in Napoleon's military expedition. http://www.egyptologyonline.com/rosetta_stone.htm November 14, 2007 4:48PM EST

30. George F. Gilder - (born November 29, 1939, in New York City) is an American writer, techno-utopian intellectual and co-founder of the Discovery Institute. His 1981 bestseller *Wealth and Poverty* advanced a practical and moral case for capitalism during the early months of the Reagan Administration. http://en.wikipedia.org/wiki/George_Gilder November 14, 2007 4:50PM EST

31. Sebastian Thrun - (born 1967 in Solingen, Germany) is a Professor of Computer Science at Stanford University and director of the Stanford Artificial Intelligence Laboratory (SAIL). He led the development of the robotic vehicle Stanley, which won the DARPA Grand Challenge in 2005 and brought Stanford a two-million dollar prize. http://en.wikipedia.org/wiki/Sebastian_Thrun November 14, 2007 5:04PM EST

32. Nanobots – (Nanorobotics) the technology of creating machines or robots at or close to the scale of a nanometres (10-9 metres). More specifically, nanorobotics refers to the still largely hypothetical nanotechnology engineering discipline of designing and building nanorobots. http://en.wikipedia.org/wiki/Nanorobotics November 14, 2007 5:10PM EST

33. Deep Fritz - **Fritz** is a German chess program developed by Frans Morsch and Mathias Feist and published by Chessbase. There is also a version called *Deep Fritz* that is designed for multi-processing. http://en.wikipedia.org/wiki/Deep_Fritz November 14, 2007 5:28PM EST

34. Deep Thought - a computer designed to play chess. It was second in the line of chess computers developed by Feng-hsiung Hsu, starting with ChipTest and culminating in **Deep Blue**. Deep Thought was easily defeated in both games of a 2-game match with Kasparov in 1989 as well as in a play by email match with Michael Valvo.
It was named after **Deep Thought**, a fictional computer in Douglas Adams' series, *The Hitchhiker's Guide to the Galaxy*.
http://en.wikipedia.org/wiki/Deep_Thought
November 14, 2007 5:31PM EST

35. Deep Blue - a chess-playing computer developed by IBM. On 11 May 1997, the machine won a short 6 game exhibition match (not a world title match) by two wins to one with 3 draws against world champion Garry Kasparov after Kasparov made a remarkable blunder (for a world chess champion) in the opening of the last game.
http://en.wikipedia.org/wiki/IBM_Deep_Blue
November 14, 2007 5:34PM EST

Bio

Ray Kurzweil, Ph.D.

Ray Kurzweil is the best-selling author, inventor, entrepreneur, futurist, and movie producer. In 1999, he received the National Medal of Technology for pioneering and innovative achievements in computer science serving to enrich the lives of many disabled persons.

For additional information, please visit www.kurzweilai.net and www.singularity.com.

The JOURNAL *of* GEOETHICAL NANOTECHNOLOGY

Volume 2, Issue 4
4th Quarter, 2007

Upgrading Humans - Technical Realities and New Morals

Professor Kevin Warwick

This article was adapted from a lecture given by Professor Kevin Warwick, Ph.D. during the 3rd Annual Workshop Webinar on Geoethical Nanotechnology, on July 20th, 2007.

Dr. Warwick, Professor of Cybernetics at the University of Reading, England, describes how his 1998 experiment allowed him the title of the world's first, 'Human Cyborg' when he implanted a Radio Frequency Identification Device (of his own design), within his body. Dr. Warwick also explains the present and possible future benefits of the technology of merging man and machine.

From the last millennium, as far back as 1998, I became the first human to have one of these things that you see on *Image 1*, implanted. Not the coin, because the coin is just there to show what a typical salary of a professor is in the UK. On the right-hand side is an RFID, Radio Frequency Identification Device. [1]

Image 1: RFID

The implant in my left arm actually identified me to the computer in the building at Reading University. We got the computer to open doors for me, to switch on lights, even to say "hello" when I came through the front door.

Since we're using the same sort of technology, in hand somewhat, with different tracking devices, either by the cell phone network or by GPS, to track and monitor individuals. That opens all sorts of ethical questions, not just with regard to the identity of a person, which is one you can get by the radio frequency ID devices, maybe information on the person or even their body state, but also where the person is at any time. I get embroiled now in all sorts of questions as to whether we should use the technology for tracking and locating old people perhaps suffering from dementia, or even abducted children. I am keeping a technological low profile on that.

The article overall, though, is looking at the merger between humans and technology and what that can offer. There are two aspects to that. One is the possibility of enhancement, upgrading, actually giving all humans extra abilities. The other is to look at problems where technology can be used for therapeutic purposes. From an ethical point of view, there are more immediate possibilities; fewer people are concerned about that.

"One is the possibility of enhancement, upgrading, actually giving all humans extra abilities."

One thing we are working on with surgeons is the deep brain stimulator used with Parkinson's disease. [2] We are trying to take that a stage further.

Image 2 shows an actual incident of tremors coming on in a time sequence. This is data taken from an actual patient. On the bottom of the four graphs what we can see, if you go along the graph, is EMG movement signals.[3] There is nothing until about 45 seconds, just over halfway through.

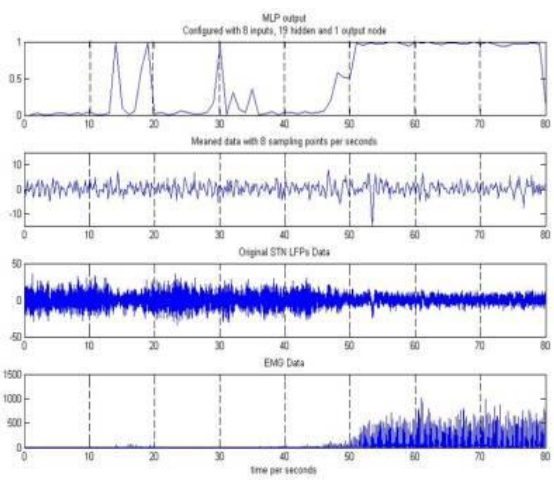

Image 2: MLP Output

Then the sequence is showing movement of a patient at a particular time point. This is at about the 45-second time point. This is when the tremors actually start.

If we look at the third trace then what we can see is the data taken from the actual implant in the individual's brain that has Parkinson's disease. This is local field potential data. Up until about the 45-second mark, the variations of the signal are much greater and then they deteriorate.

We are putting in an artificial brain, a computer brain, an artificial neural network, to learn to recognize the tremors, but also to try to predict when tremors are going to occur before they occur.

We can see on the first/top trace the output of the network operating accurately -- well, it's stating that tremors are occurring -- but more importantly, at an earlier point, particularly this point at 30 seconds, we are seeing the network predicting accurately that Parkinsonian tremors are going to occur, before they occur. In fact, it is about 15 to 20 seconds before they occur.

We are developing an intelligent stimulator that only stimulates when our computer brain predicts that it needs to stimulate. The overall device is really a computer brain sitting within the human brain predicting what the human brain is going to do, and then stopping it from doing it before it actually does it.

"...a computer brain sitting within the human brain predicting what the human brain is going to do, and then stopping it from doing it before it actually does it."

Of course, from an ethical point of view there are all sorts of spillovers here, because the same sort of stimulator can be used to stimulate feelings of happiness, feelings of sadness and, therefore, there is potential for its use with depression or just generally as an electronic drug.

If you're feeling down, this stimulator can predict that you're going to feel down before you feel down, and make you feel happy before you have even thought about it.

Image 3: Articulated Robot Arm

Campbell Aird lost his right arm due to cancer. He has received this articulated robot arm which he can control by muscular movements. It is actually connected up to muscles that he can move, and the signals are translated in order to get the robot arm to move and the hand to operate.

We are working on an interface directly from the nervous system to the robot hand, which is bidirectional such that it can pick up sensory signals and send them back the other way.

From a therapeutic point of view, it's great. We can help people here that have had their hands and arms amputated, so that they can control the hand and arm as though the brain thinks it's theirs.

It is not translating muscular signals as others have done, for example, the University of Chicago, into signals to drive the arm. It is also picking up signals to give an indication directly to the person's brain of how much force the hand is applying, and so on and so forth.

From an ethical point of view, it is great to help people from a therapeutic point of view. However, once we get signals, we are putting brain signals onto the wires, neural signals that appear on the wires in the arm. From a technological point of view, once you get signals on wires, you can send them wherever you like.

We are looking here at the possibility, not just of somebody controlling an arm, but somebody, directly from their neural signals, controlling technology that is connected to the internet, wherever, on another continent, possibly even on another planet.

If we are looking at space travel in the future, it may not be necessary for the human to actually travel, but their extremities, if you like, their body parts can travel. And they can get a feeling and a sense -- maybe even smell and get different sensory input from the distant planet without actually moving from the comfort of their beautiful surroundings.

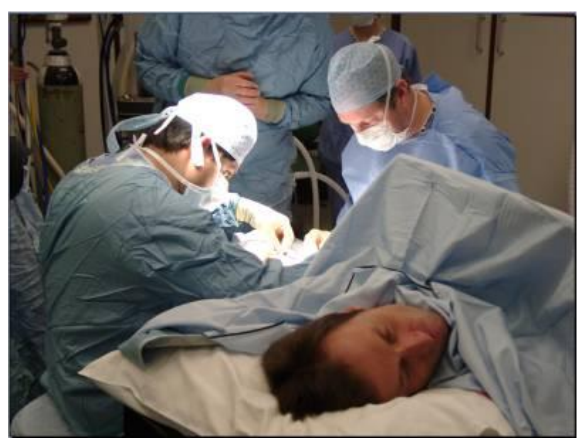

Image 4: Implant Operation

Now, we will now move on to my own last implant experiment. *Image 4* shows me having a wonderful time on the operating theater bed at the Radcliffe Infirmary in Oxford.

The implant took two hours to put in place. This is actually the implant that John Donoghue later called the BrainGate System [4]. In fact we did this a couple of years before John Donoghue did, and created a bidirectional link. This for us was vitally important.

Now we are working at taking the advantages of machine intelligence, when you compare it to human intelligence, and trying to upgrade human intelligence by linking it to the machine directly, in a bidirectional way.

There are other advantages; well, memories ultimately. Can we enhance our memory by linking our brains directly to computer brains, and can we enhance our sensory input?

Humans have a very limited range of senses. We have five that we know about, can we enhance that, by including infrared, ultrasonic, or ultraviolet, et cetera, et cetera?

Can we increase the dimensionality with which we understand things? The human brain understands things in three dimensions. Can we look at understanding things as a computer can, in many, many more dimensions?

If we also look, perhaps most importantly, at communication. The way we are communicating here in terms of translating signals from our brains, electrochemical signals, into what are these trivial coded messages that we call speech. Then try to translate them back again with our ears to understand what someone else is thinking, it is a very poor way of communicating. The possibilities of communicating directly, ultimately brain-to-brain, well, that's very much what I want to look at.

"The possibilities of communicating directly, ultimately brain-to-brain, well, that's very much what I want to look at."

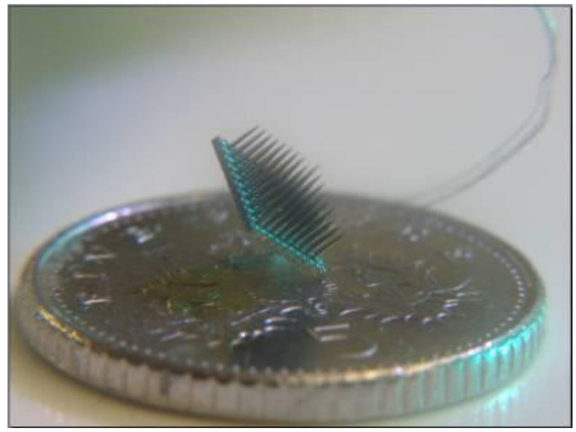

Image 5: Implant

Image 5 shows what I actually had implanted. This is 100 silicone electrodes. We see the wires moving off from them. This was surgically implanted into the median nerve fibers of my left arm. It gave a bidirectional access to the internet, to the computer. When I move my hand the neural signals could be translated by the computer into signals that could operate a robot hand, or operate any piece of technology.

I was able to drive a wheelchair around directly from neural signals. We are looking at this in terms of: could this be used directly in a person's brain? For us it was a little bit safer, I guess, to use it in my nervous system rather than my brain for experimental purposes. But also we wanted to feed in signals to see what we could actually do to stimulate my brain to give me different sensory input.

Image 6

I'll give you some examples of what we actually achieved -- this is my wife (*Image 6*) and she is wearing some jewelry that was put together by a student of the Royal College of Art.

> *"...when I was calm and relaxed, my wife's jewelry was a cool blue, and when I got excited the jewelry flashed red."*

The jewelry changes color from red to blue, so when this was connected to my nervous system, by way of the implant, we could pick up either movement signals or motor signals of one type or another, and pick out different indications, for example, excitement. One way of operating it was when I was calm and relaxed, my wife's jewelry was a cool blue, and when I got excited the jewelry flashed red.

Now if you can imagine - Irena, my wife, works in a different place than me, so she's walking around with the jewelry on and it is a cool blue. Fine, no problems at all, he is just nice and relaxed.

Then it starts flashing red: What is he doing? And more importantly, who is he doing it with? So I don't know whether it is that good of an idea.

Image 7

This is a picture of me in New York City, Columbia University and the real-time computing lab. What the guys there helped me do was to put my nervous system onto the internet live in real-time. I am not sure whether anyone else has actually done that even now.

I moved my hand in New York and my neural signals were transmitted across the internet to Reading, to the UK, in order to move a robot hand.

When the robot hand moved, it mimicked my hand movements. It then gripped an object, and we got signals fed back from fingertip sensors in the hand to -- back across the internet -- to stimulate my nervous system, literally to stimulate my brain. What I was getting was a change in current frequency stimulations. The more the hand gripped, the more the current pulses I received personally by my nervous system to my brain.

I was trying to get the robot hand on a different continent to grip an object, to apply a particular force to an object. It extended my nervous system across the internet. I have to say that worked very, very well. I have to thank the guys at Columbia for allowing us to work with them on that one.

Image 8

In the next image, we see the extrasensory input. I have a blindfold on and a Computer Associates baseball cap. They were one of our sponsors. Also on the cap are ultrasonic sensors.

The output from the ultrasonic sensors was fed down to that little sci-fi type of arrangement I've got on which was a radio transmitter receiver. And it was able to transmit signals from external devices to stimulate my nervous system, and hence to stimulate my brain, sending electro-chemical signals up through my nervous system.

> *"The closer I got to an object, the more current pulses my brain was receiving, the further away from an object, the current pulses died away."*

It took about six weeks for my brain to recognize the signals we were sending in, but when we did this particular experiment, as I moved closer to an object and no longer could see it, I was receiving ultrasonic signals that altered the current frequency and was stimulating my brain. The closer I got to an object, the more current pulses my brain was receiving, the further away from an object, the current pulses died away.

I was very accurately able to decide how far away objects were. With the blindfold on, I could move around in the laboratory without bumping into things.

When Ian, my researcher, moved a board towards me, the signals increased in frequency. But what I was trying to do in the experiment was to keep the same distance away from the board to follow the board ultrasonically.

Actually it's not that difficult to do. This is like a sense that a bat has, and it's extending our range of senses. I was able to use that, and rather more easily than I thought was going to be possible, I guess.

It is the sort of thing that a blind person could use not to repair their blindness, but to give them an alternate sense, or a different type of sensory input. For sighted people, of course, it means they can have extra senses.

Certainly ultrasonic is fine, it works. We have robots that have it, but here we know it's humans. If anybody wants an ultrasonic sense, fine, we have the technology. It will cost a little bit of money, and involve a little bit of surgery, but you could have it now.

I cannot see why we could not have infrared. After all, your television set has an infrared sense, surely you feel a little bit jealous that you don't have that.

Why not have for yourself an ultraviolet sense? Well, go for it. X-ray is the one perhaps you remember from an old Ray Milland film, "The Man with X-Ray Eyes". Well, before too long we could all have X-ray senses. Why not?

Years ago I worked in the telecommunication industry, and communication has always been very, very important to me. I think it is important to what we are as humans. I really feel terrible that we are so pathetic at communicating, and can we improve? Well, why not?

I have the implant in my nervous system. Irena, my wife, had electrodes pushed into her nervous system, a process called microneurography [5], and we linked our nervous systems together.

Image 9

When I moved my hand, my motor neural signals were transmitted across into Irena's nervous system. When she moved her hand, the other way, my brain actually received her neural pulses. When she went, one, two, three, in terms of moving her hand, my brain received, one, two, three pulses.

What we actually achieved was a telegraphic [6] form of communication, directly from nervous system to nervous system. Which is like a telephone system, if you like, but missing out the interface of giving the signals by your mouth into the microphone, converting them from electro-chemical signals to mechanical pressure waves, and then back to electrical signals, which seems such a laborious process, why not stay with electrical signals? That is what we were doing. It is the first step really on a form of communicating ultimately directly from brain-to-brain. That is where my main focus of research is now, looking at the possibilities of communicating directly from brain-to-brain, both human-to-machine and then back the other way, but also human-to-human.

I think it will take us into a completely new world with lots of exciting possibilities for me. Not just from a therapeutic point of view but from an enhancement point of view. A little bit matrixy, if you like, when we look at machines sending signals directly into the brain of a human, but from a communications point of view, tremendously exciting that in the future we will be able to send signals, graphical pictures maybe, who knows where we can stop. It is a very exciting future with this type of technology linking human brains and machine brains together.

Follow-up Issues from the Workshop

On interference from other wireless electronic devices

Interference is perhaps a funny way of looking at it. When we were monitoring my own nervous system, we were able to pick up signals that everybody's nervous system gets, such as particularly text messages on mobile phones. Once, when I was trying to move my hand, we were looking at the motor neural signals directly in real-time and suddenly they started shooting around all over the place. For a moment I was worried, I thought there was some major problem with my nervous system.

It turned out that one of the researchers had a text message coming through directly close by. It was just giving an indication of what my nervous system was doing. The signals were appearing on my nervous system, but essentially my brain was not doing anything with the signals.

After all the nervous system is – it is not a wonderful aerial, but it is quite a good antenna itself, a resistor/capacitor network. We were able to pick up a local radio station that has a big mast about ten miles away. We had a sound card attached to my nervous system and were listening. It was not good reception in the sense, it was very poor, but it was in fact Barry Manilow.

On bandwidth limitations

We have been trying to make everything as simple as possible. The communication has just been telegraphic, a bit like a Morse Code communication, because we wanted to see if this implant be used for very simple control devices, such as the work of John Donoghue.

The signals are very simple: yes, no, left, right, and using a cursor and selecting from things on the menu, so that a very simple implant could be more widely used for people who are paralyzed. Potentially, in this simple form, a person could use it to drive a wheelchair; a person could drive a car around with a very simple implant of this version.

When we looked to richer forms of communication it would need much greater bandwidth. When we communicate, we are not using all of our neurons, we are using some select neurons. If we were looking at transmitting signals, maybe it would need a lot more electrodes, but not a lot, lot more. Maybe a thousand would be sufficient for quite a rich form of communication.

On misuse

Rather than a chemical form of drug, an electrical, electronic form of drug could ultimately be quite dangerous. But it could also be very, very positive from a medical point of view, overcoming pain with electronic signals, you could be very specific and overcome particular pain signals without the need to affect the whole of the body, as with chemicals.
There are also some potential misuses where one group would like to control another group and this technology, whether we like it or not, is opening up that possibility.
It is coming up with very much of a Matrix sort of scenario. There are good points and bad points for that, therapy and maybe enhancing electronic signals, but also potential negatives of individual control. I am really a technological person when I look at what is possible and what is not. Society, though, must contend with these enormous questions.

Endnotes

1. Radio-frequency identification (**RFID**) - an automatic identification method, relying on storing and remotely retrieving data using devices called RFID tags or transponders. An RFID tag is an object that can be applied to or incorporated into a product, animal, or person for the purpose of identification using radiowaves. Some tags can be read from several meters away and beyond the line of sight of the reader.
http://en.wikipedia.org/wiki/RFID October 12, 2007 3:49PM EST

2. Parkinson's disease - a movement disorder that is chronic and progressive and occurs when a group of cells in an area of the brain called the substantia nigra begin to malfunction and die.
http://www.pdf.org/AboutPD/ October 12, 2007 4:10PM EST

3. EMG - Electromyography, or EMG, involves testing the electrical activity of muscles. Often, EMG testing is performed with another test that measures the conducting function of

nerves.
www.emedicinehealth.com October 12, 2007 4:15PM EST

4. The BrainGate Neural Interface System - is currently the subject of a pilot clinical trial being conducted under an Investigational Device Exemption (IDE) from the FDA. The system is designed to restore functionality for a limited, immobile group of severely motor-impaired individuals. It is expected that people using the BrainGate System will employ a personal computer as the gateway to a range of self-directed activities.
www.cyberkineticsinc.com/content/medicalproducts/
October 12, 2007 5:19PM EST

5. John Donoghue, Ph.D., MS – Is a Henry Merritt Wriston Professor with the Department of Neuroscience and Director of the Brain Science Program at Brown University in Providence, RI.
ttp://research.brown.edu/myresearch/John_Donoghue%20 October 12, 2007 5:25PM EST

6. Microneurography is a method to record the traffic of impulses in human nerves with percutaneously inserted needle electrodes.
www.neurofys.uu.se/v2/index.php?lnk=research/micron October 12, 2007 5:41PM EST

7. Telegraphy is the long-distance transmission of written messages without physical transport of letters, originally by changing something that could be observed from a distance (*optical telegraphy*).

8. Radiotelegraphy or **wireless telegraphy** transmits messages using radio. Telegraphy includes recent forms of data transmission such as fax, email, and computer networks in general.
http://en.wikipedia.org/wiki/Telegraphy
October 12, 2007 5:44PM EST

BIO

Kevin Warwick

Dr. Warwick Is the Professor of Cybernetics at the University of Reading, England, where he carries out research in artificial intelligence, control, robotics and biomedical engineering. He is also Director of the University KTP Centre, which links the University with Small to Medium Enterprises. At 22 he took his first degree at Aston University, followed by a PhD and a research post at Imperial College, London. He subsequently held positions at Oxford, Newcastle and Warwick universities before being offered the Chair at Reading, at the age of 33.