This is a modification of the article as published on NewScientist.com.
This is reproduced with permission from NewScientist magazine.

NewScientist.com

Human life: The next generation

- 24 September 2005
- NewScientist.com news service
- Ray Kurzweil

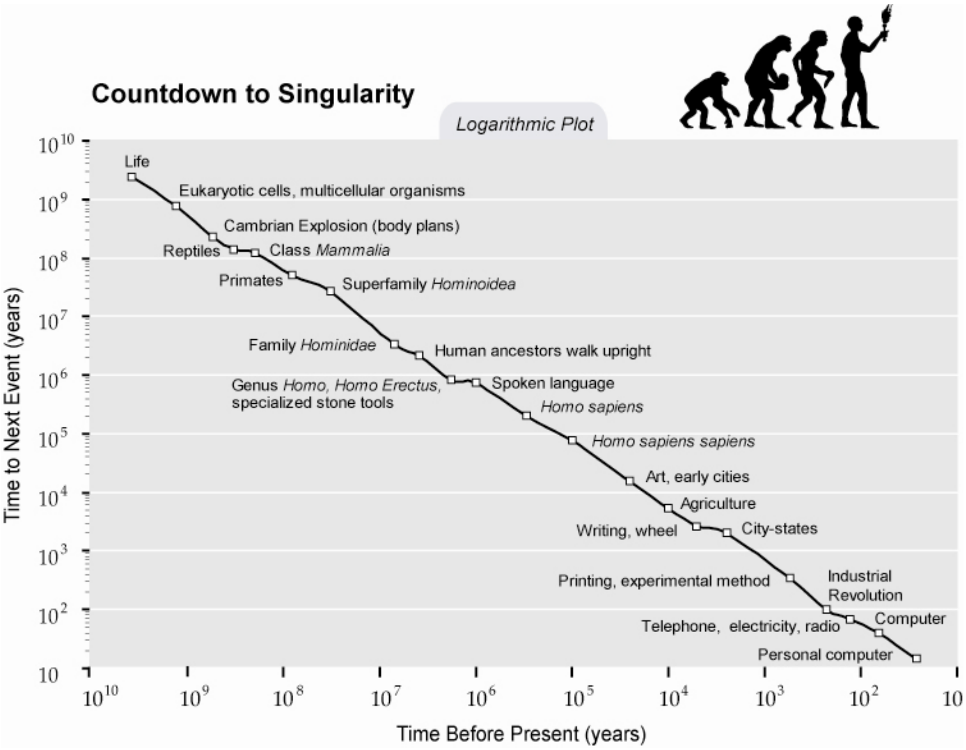

Credit: Ray Kurzweil and KurzweilAI.net

IN 2003, *Time* magazine organised a "Future of Life" conference celebrating the 50th anniversary of Watson and Crick's discovery of the structure of DNA. All the speakers - myself included - were asked what we thought the next 50 years would bring. Most of the predictions were short-sighted.

James Watson's own prediction was that in 50 years, we'll have drugs that allow us to eat as much as we want without gaining weight. "Fifty years?," I replied. In my opinion that's far too pessimistic. We've already demonstrated it in mice, and human drugs using the relevant techniques are in development. We can expect them in five to 10 years, not 50.

The mistake that Watson and virtually every other presenter made was to use the progress of the past 50 years as a model for the next half-century. I describe this way of looking at the future as the "intuitive linear" view: people intuitively assume that the current rate of progress will continue for future periods.

But a serious assessment of the history of technology reveals that technological change is not linear, but exponential. You can examine the data in different ways, on different timescales and for a wide variety of technologies, ranging from electronic to biological. You can analyse the implications, ranging from the sum of human knowledge to the size of the economy. However you measure it, the exponential acceleration of progress and growth applies.

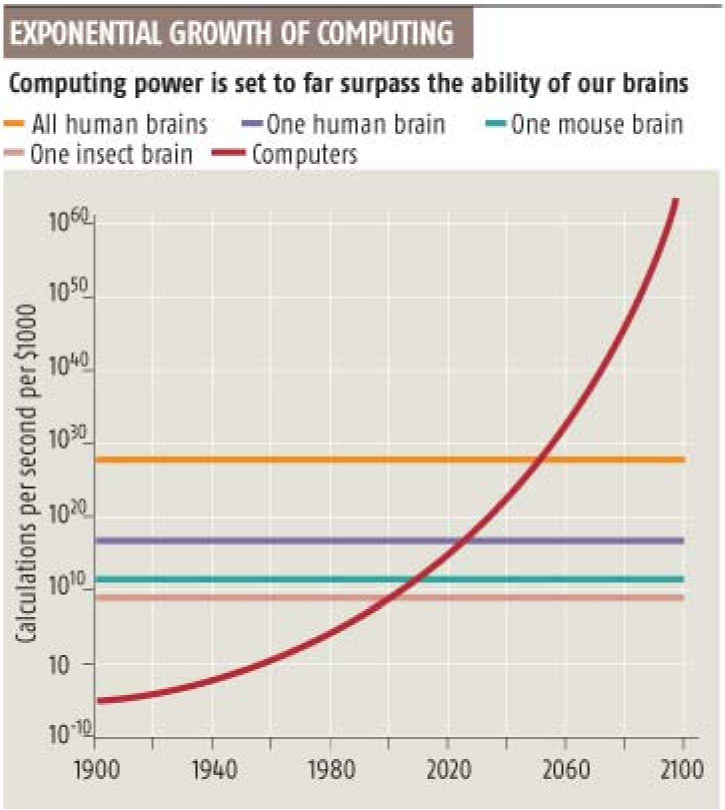
EXPONENTIAL GROWTH OF COMPUTING

Computing power is set to far surpass the ability of our brains

Understanding exponential progress is key to understanding future trends. Over the long term, exponential growth produces change on a scale dramatically different from linear growth. Consider that in 1990, the human genome project was widely regarded as controversial. In 1989, we sequenced only one-thousandth of the genome. But from 1990 onwards the amount of genetic data sequenced doubled every year - a rate of growth that continues today - and the transcription of the human genome was completed in 2003.

We are making exponential progress in every type of information technology. Moreover, virtually all technologies are becoming information technologies. If we combine all of these trends, we can reliably predict that, in the not too distant future, we will reach what is known as The Singularity. This is a time when the pace of technological change will be so rapid and its impact so deep that human life will be irreversibly transformed. We will be able to reprogram our biology, and ultimately transcend it. The result will be an intimate merger between ourselves and the technology we are creating.

The evidence for this ubiquitous exponential growth is abundant. In my new book, *The Singularity is Near*, I have more than 40 graphs from a broad variety of fields, including communications, the internet, brain scanning and biological technologies, that reveal exponential progress. Broadly speaking, my models show that we are doubling the paradigm-shift rate (roughly, the rate of technical innovation) every decade. Throughout the 20th century, the rate of progress gradually picked up speed. By the end of the century the rate was such that the sum total of the century's achievements was equivalent to about 20 years of progress at the 2000 rate.

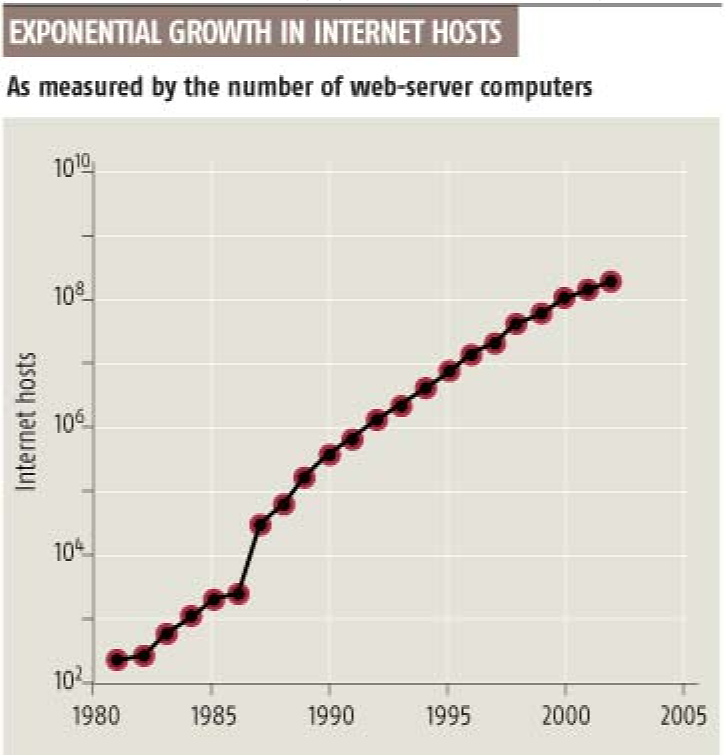
EXPONENTIAL GROWTH IN INTERNET HOSTS

As measured by the number of web-server computers

Growth in information technology is particularly rapid: we're doubling its power, as measured by price-performance, bandwidth, capacity and many other measures, every year or so. That's a factor of a thousand in 10 years, a million in 20 years, and a billion in 30 years, although a slow, second level of exponential growth means that a billion-fold improvement takes only about a quarter of a century.

The exponential growth of computing goes back over a century and covers five major paradigms: electromechanical computing as used in the 1890 US census, relay-based computing as used to crack Nazi cryptography in the early 1940s, vacuum-tube-based computing as used by CBS to predict the election of Dwight Eisenhower in 1952, discrete-transistor-based computing as used in the first space launches in the early 1960s, and finally computing based on integrated circuits, invented in 1958 and applied to mainstream computing from the late 1960s. Each time it became apparent that one paradigm was about to run out of steam, this realisation resulted in research pressure to create the next paradigm.

Today we have over a decade left in the paradigm of shrinking transistors on an integrated circuit, but there has already been enormous progress in creating the sixth major computing paradigm of three-dimensional molecular computing, using carbon nanotubes for example. And electronics is just one example of many. As another, it took us 14 years to sequence the genome of HIV; SARS took only 31 days.

Accelerating returns

The result is that we can reliably predict such measures as price-performance and capacity of a broad variety of information technologies. There are, of course, many things that we cannot dependably anticipate. In

HISTORY OF THE FUTURE

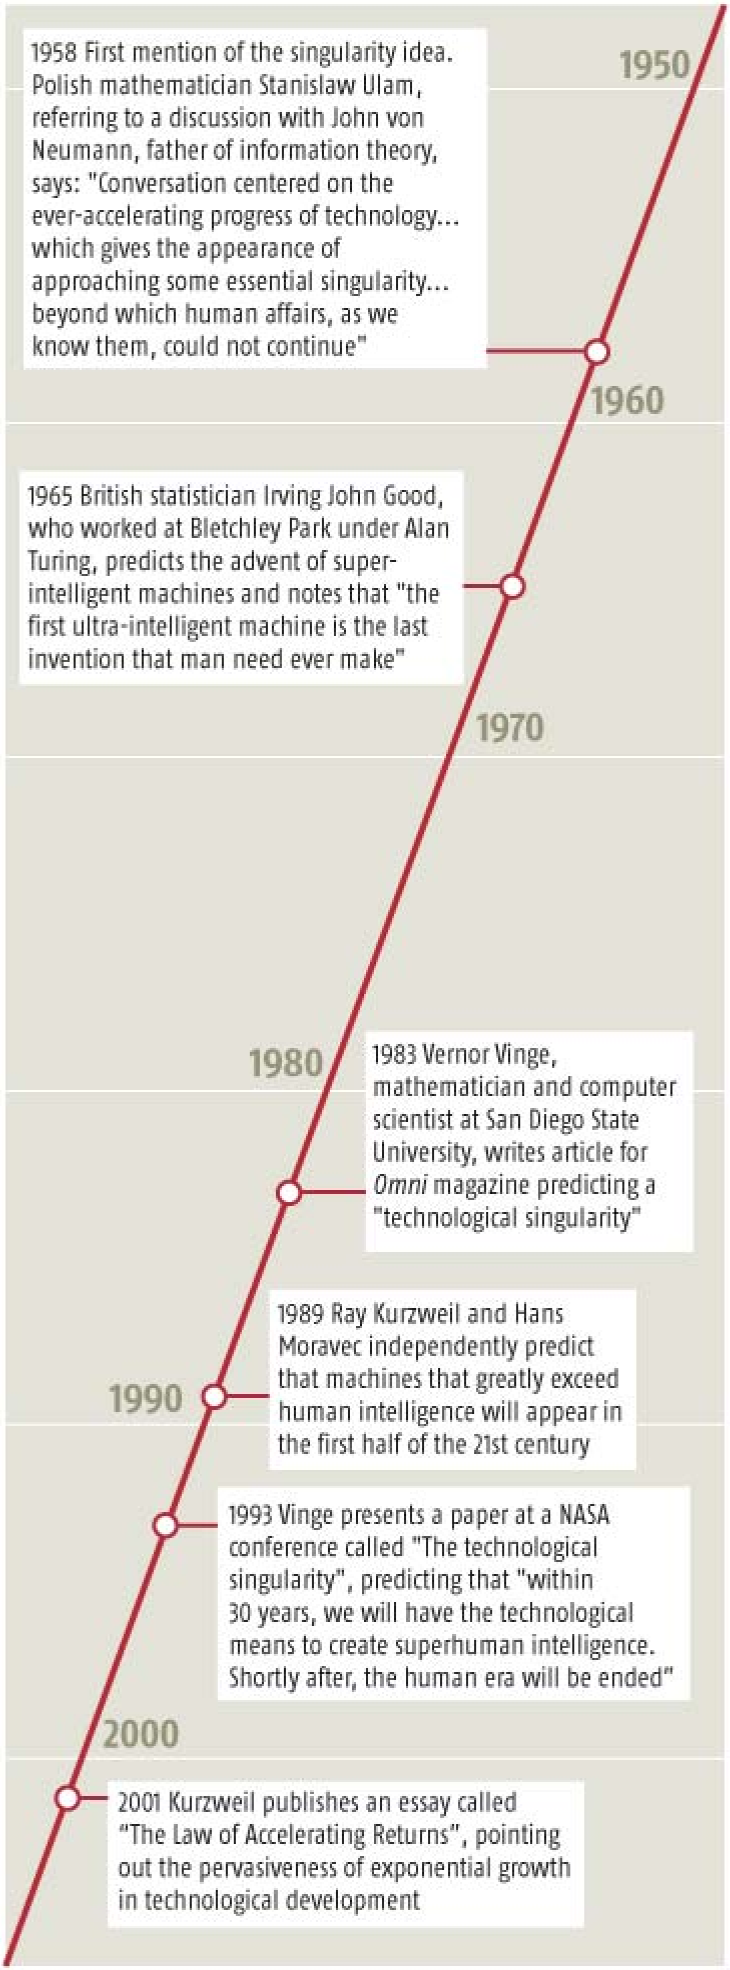

1958 First mention of the singularity idea. Polish mathematician Stanislaw Ulam, referring to a discussion with John von Neumann, father of information theory, says: "Conversation centered on the ever-accelerating progress of technology... which gives the appearance of approaching some essential singularity... beyond which human affairs, as we know them, could not continue"

1965 British statistician Irving John Good, who worked at Bletchley Park under Alan Turing, predicts the advent of super-intelligent machines and notes that "the first ultra-intelligent machine is the last invention that man need ever make"

1983 Vernor Vinge, mathematician and computer scientist at San Diego State University, writes article for *Omni* magazine predicting a "technological singularity"

1989 Ray Kurzweil and Hans Moravec independently predict that machines that greatly exceed human intelligence will appear in the first half of the 21st century

1993 Vinge presents a paper at a NASA conference called "The technological singularity", predicting that "within 30 years, we will have the technological means to create superhuman intelligence. Shortly after, the human era will be ended"

2001 Kurzweil publishes an essay called "The Law of Accelerating Returns", pointing out the pervasiveness of exponential growth in technological development fact, our inability to make reliable predictions applies to any specific project. But the overall capabilities of information technology in each field can be projected. And I say this not just with hindsight; I have been making forward-looking predictions of this type for more than 20 years.

We see examples in other areas of science of very smooth and reliable outcomes resulting from the interaction of a great many unpredictable events. Consider that predicting the path of a single molecule in a gas is essentially impossible, but predicting the properties of the entire gas - comprised of a great many chaotically interacting molecules - can be done very reliably through the laws of thermodynamics. Analogously, it is not possible to reliably predict the results of a specific project or company, but the overall capabilities of information technology, comprised of many chaotic activities, can nonetheless be dependably anticipated through what I call "the law of accelerating returns".

So what does the law of accelerating returns tell us about the future? In terms of the aforementioned paradigm-shift rate, between 2000 and 2014 we'll make 20 years of progress at 2000 rates, equivalent to the entire 20th century. And then we'll do the same again in only seven years. To express this another way, we won't experience 100 years of technological advance in the 21st century; we will witness in the order of 20,000 years of progress when measured by the rate of progress in 2000, or about 1000 times that achieved in the 20th century.

Above all, information technologies will grow at an explosive rate. And information technology is *the* technology that we need to consider. Ultimately everything of value will become an information technology: our biology, our thoughts and thinking processes, manufacturing and many other fields. As one example, nanotechnology-based manufacturing will enable us to apply computerised techniques to automatically assemble complex products at the molecular level. This will mean that by the mid-2020s we will be able to meet our energy needs using very inexpensive nanotechnology-based solar panels that will capture the energy in 0.03 per cent of the sunlight that falls on the Earth, which is all we need to meet our projected energy needs in 2030.

A common objection is that there must be limits to exponential growth, as in the example of rabbits in Australia. The answer is that there are, but they're not very limiting. By 2020, $1000 will purchase 10^{16} calculations per second (cps) of computing (compared with about 10^{9} cps today), which is the level I estimate is required to functionally simulate the human brain. Another few decades on, and we will be able to build more optimal computing systems. For example, one cubic inch of nanotube circuitry would be about 100 million times more powerful than the human brain. The ultimate 1-kilogram computer - about the weight of a laptop today - which I envision late in this century, could provide 10^{42} cps, about 10 quadrillion (10^{16}) times more powerful than all human brains put together today. And that's if we restrict the computer to functioning at a cold temperature. If we find a way to let it get hot, we could improve that by a factor of another 100 million. And of course, we'll devote more than 1 kilogram of matter to computing. Ultimately, we'll use a significant portion of the matter and energy in our vicinity as a computing substrate.

Our growing mastery of information processes means that the 21st century will be characterised by three great technology revolutions. We are in the early stages of the "G" revolution (genetics, or biotechnology) right now. Biotechnology is providing the means to actually change your genes: not just designer babies but designer baby boomers.

One technology that is already here is RNA interference (RNAi), which is used to turn genes off by blocking messenger RNA from expressing specific genes. Each human gene is just one of 23,000 little software programs we have inherited that represent the design of our biology. It is not very often that we use software programs that are not upgraded and modified for several years, let alone thousands of years. Yet these genetic programs evolved tens of thousands of years ago when conditions were very different. For one thing, it was not in the interest of the species for people to live very long. But since viral diseases, cancer and many other diseases depend on gene expression at some crucial point in their life cycle, RNAi promises to be a breakthrough technology.

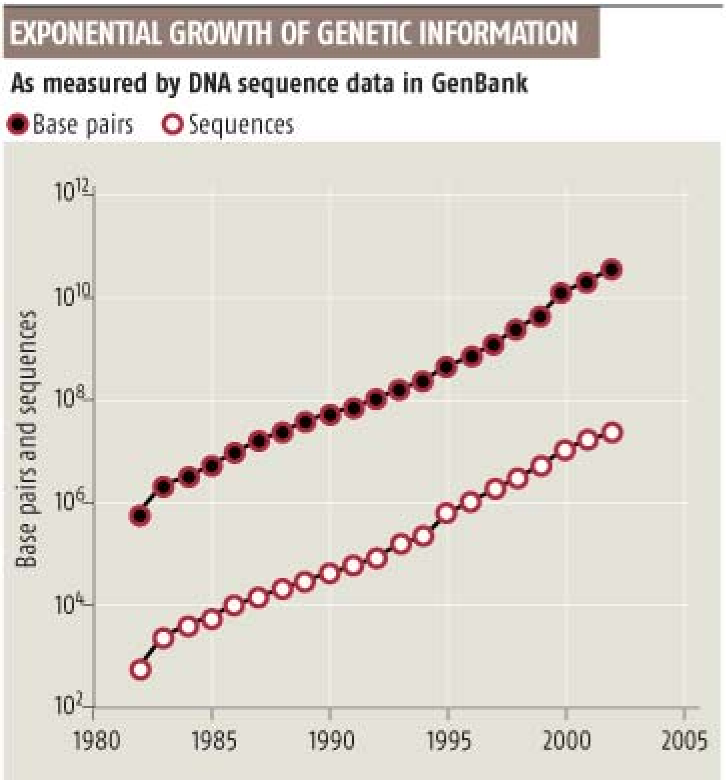
EXPONENTIAL GROWTH OF GENETIC INFORMATION

As measured by DNA sequence data in GenBank

Grow your own

New means of adding new genes are also emerging that have overcome the problem of placing genetic information precisely. One successful technique is to add the genetic information in vitro, making it possible to ensure the genetic information is inserted in the proper place. Once verified, the modified cell can be reproduced in vitro and large numbers of modified cells introduced into the patient's bloodstream, where they will travel to and become embedded in the correct tissues. This approach to gene therapy has successfully cured pulmonary hypertension in rats and has been approved for human trials.

Another important line of attack is to regrow our own cells, tissues and even whole organs, and introduce them into our bodies. One major benefit of this "therapeutic cloning" technique is that we will be able to create these new tissues and organs from versions of our cells that have also been made younger - the emerging field of rejuvenation medicine. For example, we will be able to create new heart cells from your skin cells and introduce them into your system through the bloodstream. Over time, your heart cells will all be replaced, resulting in a rejuvenated "young" heart with your own DNA.

Drug discovery was once a matter of finding substances that produced some beneficial effect without excessive side effects. This process was similar to early humans' tool discovery, which was limited to simply finding rocks and natural implements that could be used for helpful purposes. Today, we are learning the precise biochemical pathways that underlie both disease and ageing processes, and are able to design drugs to carry out precise missions at the molecular level. The scope and scale of these efforts are vast.

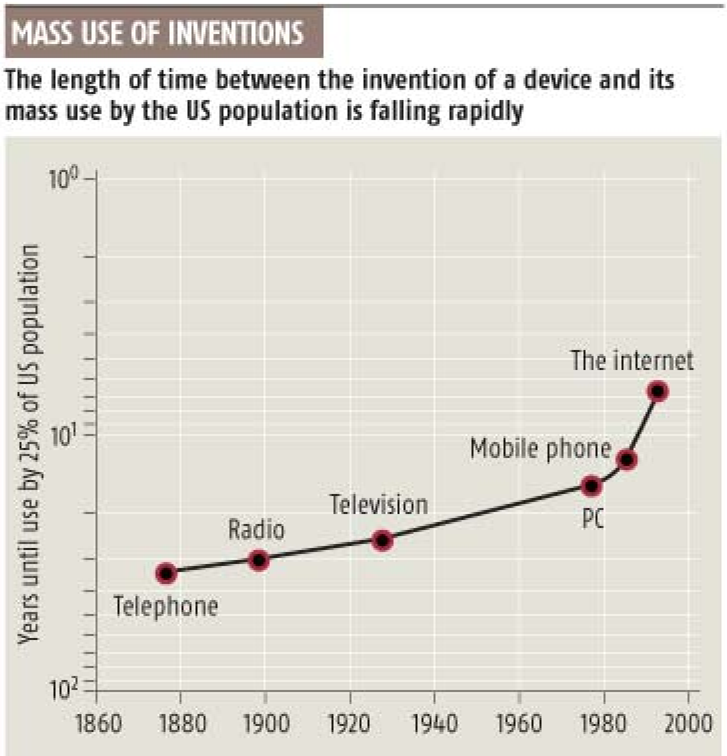
MASS USE OF INVENTIONS

The length of time between the invention of a device and its mass use by the US population is falling rapidly

But perfecting our biology will only get us so far. The reality is that biology will never be able to match what we will be capable of engineering, now that we are gaining a deep understanding of biology's principles of operation.

That will bring us to the "N" or nanotechnology revolution, which will achieve maturity in the 2020s. There are already early impressive experiments. A biped nanorobot created by Nadrian Seeman and William Sherman of New York University can walk on legs just 10 nanometres long, demonstrating the ability of nanoscale machines to execute precise manoeuvres. MicroCHIPS of Bedford, Massachusetts, has developed a computerised device that is implanted under the skin and delivers precise mixtures of medicines from hundreds of nanoscale wells inside it. There are many other examples.

Version 2.0

By the 2020s, nanotechnology will enable us to create almost any physical product we want from inexpensive materials, using information processes. We will be able to go beyond the limits of biology, and replace your current "human body version 1.0" with a dramatically upgraded version 2.0, providing radical life extension. The "killer app" of nanotechnology is "nanobots", blood-cell sized robots that can travel in the bloodstream destroying pathogens, removing debris, correcting errors in DNA and reversing ageing processes.

We're already in the early stages of augmenting and replacing each of our organs, even portions of our brains with neural implants, the most recent versions of which allow patients to download new software to their implants from outside their bodies. Each of our organs will ultimately be replaced. For example, nanobots could deliver to our bloodstream an optimal set of all the nutrients, hormones and other substances we need, as well as remove toxins and waste products. The gastrointestinal tract could then be reserved for culinary pleasures rather than the tedious biological function of providing nutrients. After all, we've already in some ways separated the communication and pleasurable aspects of sex from its biological function.

The most profound transformation will be "R" for the robotics revolution, which really refers to "strong" AI, or artificial intelligence at the human level (see "Reverse engineering the human brain"). Hundreds of applications of "narrow AI" - machine intelligence that equals or exceeds human intelligence for specific tasks - already permeate our modern infrastructure. Every time you send an email or make a cellphone call, intelligent algorithms route the information. AI programs diagnose electrocardiograms with an accuracy rivalling doctors, evaluate medical images, fly and land aircraft, guide intelligent autonomous weapons, make automated investment decisions for over a trillion dollars of funds, and guide industrial processes. A couple of decades ago these were all research projects.

With regard to strong AI, we'll have both the hardware and software to recreate human intelligence by the end of the 2020s. We'll be able to improve these methods and harness the speed, memory capabilities and knowledge-sharing ability of machines.

Ultimately, we will merge with our technology. This will begin with nanobots in our bodies and brains. The nanobots will keep us healthy, provide full-immersion virtual reality from within the nervous system, provide direct brain-to-brain communication over the internet and greatly expand human intelligence. But keep in mind that non-biological intelligence is doubling in capability each year, whereas our biological intelligence is essentially fixed. As we get to the 2030s, the non-biological portion of our intelligence will predominate. By the mid 2040s, the non-biological portion of our intelligence will be billions of times more capable than the biological portion. Non-biological intelligence will have access to its own design and will be able to improve itself in an increasingly rapid redesign cycle.

This is not a utopian vision: the GNR technologies each have perils to match their promise. The danger of a bioengineered pathological virus is already with us. Self-replication will ultimately be feasible in non-biological nanotechnology-based systems as well, which will introduce its own dangers. This is a topic for another essay, but in short the answer is not relinquishment. Any attempt to proscribe such technologies will not only deprive human society of profound benefits, but will drive these technologies underground, which would make the dangers worse.

Some commentators have questioned whether we would still be human after such dramatic changes. These observers may define the concept of human as being based on our limitations, but I prefer to define us as the species that seeks - and succeeds - in going beyond our limitations. Because our ability to increase our horizons is expanding exponentially rather than linearly, we can anticipate a dramatic century of accelerating change ahead.

Reverse engineering the human brain

The most profound transformation will be in "strong" AI, that is, artificial intelligence at the human level. To recreate the capabilities of the human brain, we need to meet both the hardware and software requirements. Achieving the hardware requirement was controversial five years ago, but is now largely a mainstream view among informed observers. Supercomputers are already at 100 trillion (10^{14}) calculations per second (cps), and will hit 10^{16} cps around the end of this decade, which is the level I estimate is required to functionally simulate the human brain. Several supercomputers with 10^{15} cps are already on the drawing board, with two Japanese efforts targeting 10^{16} cps around the end of the decade. By 2020, 10^{16} cps will be available for around $1000. So now the controversy is focused on the algorithms.

To understand the principles of human intelligence we need to reverse-engineer the human brain. Here, progress is far greater than most people realise. The spatial and temporal resolution of brain scanning is progressing at an exponential rate, roughly doubling each year. Scanning tools, such as a new system from the University of Pennsylvania, can now see individual interneuronal connections, and watch them fire in real time. Already, we have mathematical models of a couple of dozen regions of the brain, including the cerebellum, which comprises more than half the neurons in the brain. IBM is creating a highly detailed simulation of about 10,000 cortical neurons, including tens of millions of connections. The first version will simulate electrical activity, and a future version will also simulate chemical activity. By the mid 2020s, it is conservative to conclude that we will have effective models of the whole brain.

There are a number of key ways in which the organisation of the brain differs from a conventional computer. The brain's circuits, for example, transmit information as chemical gradients travelling at only a few hundred metres per second, which is millions of times slower than electronic circuits. The brain is massively parallel: there are about 100 trillion interneuronal connections all computing simultaneously. The brain combines analogue and digital phenomena. The brain rewires itself, and it uses emergent properties, with intelligent behaviour emerging from the brain's chaotic and complex activity. But as we gain sufficient data to model neurons and regions of neurons in detail, we find that we can express the coding of information in the brain and how this information is transformed in mathematical terms. We are then able to simulate these transformations on conventional parallel computing platforms, even though the underlying hardware architecture is quite different.

One benefit of a full understanding of the human brain will be a deep understanding of ourselves, but the key implication is that it will expand the tool kit of techniques we can apply to create artificial intelligence. We will then be able to create non-biological systems that match human intelligence. These superintelligent computers will be able to do things we are not able to do, such as share knowledge and skills at electronic speeds.

Close this window Return to article