

Homo Sapiens Facticus

Peter J Bentley

The purpose of life is to make more life. From the first self-replicating molecules to the complex organisms that make up our ecosystems, all life simply aims to function long enough to reproduce. To achieve this trick, life changes its environment, it changes other life, and it changes itself. Self-replicating molecules join forces in cells, parasites merge into hosts, viruses transfer genes across species, different species co-evolve perfect partnerships with each other, and the rocks and air of our planet are transformed from barren volcanic desolation into blue-green paradise.

Life has always meant change, evolution, adaptation. Human life is no exception. The only difference is that humans have so many capabilities to change life, to change our environments and to change ourselves. We also have the capacity to wonder if we should.

Parameterised People?

Genes. Sections of a long, complicated molecule called DNA that helps produce other molecules, called proteins. Biological nanotechnology, controlling the growth, death, movement and transformation of our cells.

Genes are no blueprint of life, they are the facilitators of development, the transducers of environments into organisms. Genes were written by evolution, and are still being written, although evolution may have lost its sense of direction because of human inventions. Wars, commerce, selective breeding, medicine, and genetic engineering are all taking our species in different genetic directions. But genes are not sacred any more than the complicated patterns of molecules of a snowflake are. Genes are commonly occurring molecules – common because the particular reaction caused by their structures makes them so. That reaction is perhaps the most complex chemical reaction in the universe: an extraordinary event that causes the replication of the molecules involved, as well as the development of hugely complex machines (organisms) to house and protect those molecules, making future self-replication even more likely.

Genes are often misconceived as parameters for people: “a gene for obesity”, or “a gene for intelligence.” Change the parameter and you change the person in that one aspect. This is a gross simplification of the truth, as science is now revealing. Genes work in partnership with a developing organism in its environment. If the environment is abnormal, then so will be the developing organism, regardless of the design of its genes. The total number of genes in a human genome only contains a fraction of the information necessary to create a complete human being. The additional information is provided by the environment. Our visual centres wire themselves up (properly) as we try to see the complexities around us, our bones and muscles form properly through use under normal Earth gravity, the complexities of our social interactions and education are only possible because of our historical and present social environments.

In reality, genes only produce proteins. That's all they ever do. Those proteins may affect a billion cells and interactions in horribly complex patterns, or they may produce fewer, more focussed effects on an organism's development. Currently analysis of the effects of genes is simplistic and assumes the latter. Consequently medicines and therapies based on such knowledge cause "side effects" – effects unforeseen by us, but effects that the genes have always had, if only we knew. As more biological data is generated and computer analysis improves, so our understanding of the true role of genes will improve. Nevertheless, the relationship between genes and organisms will never be simple, because it takes more than genes to make an organism.

The obvious repercussions from this information-theoretic perspective is that any desire to change life for the better can be achieved in two ways. In the first method we can modify genes through gene therapy or direct manipulation in zygotes. For example a special kind of virus known as a retrovirus can be genetically modified so that when it infects our cells, it inserts genes where we wish. But these are technologies still in their infancies. A much more straightforward method to modify life is to modify the environment of the developing organisms. The same genes will produce healthier and longer-lived organisms if their effect is modified by changing the environment of those genes.

To help achieve this, an understanding of *why* the genes perform the functions that they do, i.e., a recognition of our evolutionary history, will provide just as many benefits as an understanding of *how* our genome functions. For example, in times of food shortage, those individuals capable of storing excess food as fat had better chances of survival than those who remained slim. In times of danger, those who were more often agitated by noises and changes in their environments had a better chance of survival compared to those who calmly ignored danger signals.

We are organisms designed for environments that no longer exist. This is hardly new or extraordinary as almost all organisms must constantly adapt to their changing environments (unless they live in the relatively unchanging depths of our oceans). But the design of our environments and similarly the progress of our technology usually ignores our genetic heritage. While this continues, our biological forms may never match our environments. But this mismatch is not unsolvable. Our technology may one day enable us to modify ourselves to suit our new environments, or conversely to modify our environments to suit ourselves. The former may result in "designer people," built for the cities and lifestyles of the future. The latter may result in human beings living in more "biological" cities, and having more "natural" lives.

It could be argued that the beginnings of these trends are already visible. Perhaps "designer people" will naturally evolve from contemporary individuals so dependent on technology that nature is not an important component of their lives. It is not such a leap to move from changing faces and bodies through surgery, and enhancing communication through gadgetry, to changing bodies through gene modification and implanted circuitry. And the trends towards more biological cities are also evident as more individuals increasingly reject the complexities and stresses of modern cities and migrate to simpler, more organic, environmentally friendly lifestyles. Either method will work: changing our genes to suit our environment, or changing our

environments to suit our genes. And since we rarely leave technologies unused, it seems likely that both solutions may be used together.

When we do identify gene-environment combinations that are harmful to individuals, and when we have the capabilities to alter the effects, what should our response be? Do we have a right to change genes? Imagine the situation in a few decades, where gene-modification becomes a straightforward choice for adults. Do we have a right to prevent people from choosing their genes? Maybe not, if that choice harms no-one else. But any choice made by a parent might be inherited by their children. Should we let a fashion-conscious consumer choose the destiny of our species? And what of a gene that produces the same high as heroin? If an individual chooses to have this stimulation encoded in his or her genes (one that will almost certainly also cause harmful side-effects), does that individual still have the right to choose it? Once the genetic change has been made, once a human-made retrovirus has infected the brain cells and modified the genes, what should our response be? A forced cure? A lifetime of expensive care by the state? In a similar vein, do parents have a right to choose new genes for their children? Does a government have a right to impose restrictions on any genes that may cause cancer or premature death from heart failure? Would a government have the right to use gene therapy to reduce the likelihood of an individual harming others?

If your answer is no to any of these questions, then remember that genes and environments are the duality that make up life. Governments already impose restrictions on the environments of individuals – smoking is increasingly being banned, eating excessive quantities of some foods is actively discouraged, imbibing certain types of drugs is forbidden, the liberty of individuals is removed if they harm others. Modifying genes will have a similar effect on individuals as modifying environments does. So following the moral standards of today, certain genes or certain forms of personal gene-modification should also be banned, while governments should use behaviour-correcting gene therapy. Decisions will have to be made about which freedoms should be retained by individuals, and which should be subjugated for the good of a society and a species, with respect to genetics. Such decisions may be the responsibility of a society to make as a whole, once clear and unambiguous scientific data has been established.

The Robot Egg

Genes and cells are nature's nanotechnology. Designed by evolution, they demonstrate the power of using devices designed at the molecular scale to build organisms at macro scales. Our own version of the technology is a long way behind biology.

There are two types of nanotechnology today. Top-down nanotech – or the use of large machines to make very tiny devices, and bottom-up nanotech – the use of tiny devices to make larger artefacts. The former is a reality, as nanoscale computer processors, micro-structured materials and DNA computers demonstrate. Ethically, these are not hugely threatening, for they require conventional design and manufacture, and so represent just one further step towards miniaturisation and optimisation of technology. However, bottom-up nanotechnology is slightly different. Still more a dream than reality, this technology promises much: self-building devices,

self-repairing, organic technology. Perhaps a technology that could repair internal damage within us, or rebuild faulty conventional technologies, or enable our waste to be recycled effectively.

Bottom-up nanotechnology also implies several other capabilities. For such tiny machines to be able to do anything useful at all, we require billions of them. This immediately presents two problems: how to make them and how to control them. It is not feasible to use conventional top-down manufacturing methods to make so many tiny devices, so the solution is to borrow from nature: our nano-machines must self-replicate. They must make copies of themselves. It is also not feasible to program and monitor every single nano-machine, so once again the solution is to borrow from nature: we must place the same program in every machine, and ensure that different parts of the program are activated depending on the context for that machine. This is exactly what genes do within cells as organisms develop and grow.

The scientific and technological challenges to achieve these capabilities are immense and unsolved. We do not have any form of self-replicating machine at any scale (except for examples of biological cells or viruses modified by us). We also are currently incapable of producing a program that will enable billions of nano-machines to behave in the way we desire. But research into both areas is progressing, and in several decades we will have better solutions.

Current state-of-the-art includes the PACE, or “Programmable Artificial Cell Evolution” project, based at the Center for Living Technology in Venice. This research uses custom-built micro-fluidic field programmable gate array devices, or “wet labs on a chip.” Small enough to fit on an ordinary microscope slide, they enable the evolution (using a genetic algorithm) of self-replicating vesicles – the first stage towards the creation of a fully programmable self-replicating artificial cell. Other examples include the creation of DNA computers: the use of millions of custom-built DNA molecules to perform hugely parallel computations. Several other projects are slowly miniaturising robots to produce so-called “smart dust” – tiny self-powered computers capable of sensing, communicating wirelessly and moving. Such technologies are intended to provide new forms of intelligent medicine (for example, the intelligent drug that only produces its effect when in contact with cancer cells), or new forms of information processing and computation (for example, sensors in all devices to enable automatic monitoring for security or fire safety). In the long term, should self-replicating nanorobots become a reality, then an entirely new way of constructing technology would be enabled.

It seems likely that we will one day achieve such bottom-up nanotechnology. Biology demonstrates extraordinary prowess in nanotechnology already, with genes, proteins and cells already behaving as nano-machines, able to build and self-repair hugely complex organisms. Biology thus presents us with a vision of what this form of nanotechnology may look like. A device constructed by nanotechnology will necessarily have to grow much like an organism does. The creation of a new device, for example, a robot, would require one or more parents to produce an initial seed, from which a new device would grow. The robot would have to lay an egg, which would grow into a new robot.

The speed of such nano-growth will be limited by chemistry and physics – it is highly likely that our initial attempts will grow exceedingly slowly; much slower than the optimised designs of nature. It is also likely that early nanotech will be very sensitive to its environment and be very simple. Visions of run-away “grey goo” growing out of control are scientifically flawed – nanotech would be much more likely to die prematurely than grow wildly. Even if such a situation arose, it would be no more alarming than mould growing slowly on a wall. Since the fabrication would be our creation, we would find it exceedingly simple to kill it – much simpler than killing biological life, and we are already talented at that.

The true dilemmas of this form of future nanotech have nothing to do with grey goo, and much more to do with self-replication. If humans one day create self-replicating devices that pass their programs onto their children, then evolution will inevitably occur. Evolution will progress at the speed of that replication (a species of virus mutates and evolves considerably faster than elephants do). We are then presented with the problem that evolution will be modifying our nanotech programs. Whatever the selective pressure, i.e., whatever makes that self-replicating device more likely to have successful offspring, will modify the program accordingly. In a very real sense, bottom-up nanotech will be a form of simple life, that will subsequently try and survive and adapt to its environment. This life will not be threatening to us (unless it is designed to be), for it will change too slowly and will be understood by us far better than any other form of life. The laws of physics will prevent it from transforming into monstrous forms as seen in some movies. On its own, evolution would take millions of years to modify our creation. Even after all that time, we would be able to destroy it just as we can destroy all biological life using similar poisons. But it will exist because of us, and its only purpose would be to serve us. Successful bottom-up nanotechnology would make us gods, with correspondingly impressive responsibilities. Breeding your own robots might be as easy to regulate as breeding livestock. But a society in possession of the machinery required to construct bottom-up nanotechnology would surely be as powerful as a society in possession of nuclear technology, except that such power would be constructive rather than destructive.

You, Robot?

Nanorobots may still be little more than dreams, but macro-sized robots are one of the most common devices we use today. Comprising some form of autonomous control (typically a computer), often using sensors and motors, many of our familiar household appliances are properly termed robots. Washing machines, video recorders, refrigerators, central heating systems, microwave ovens, automobiles, even children’s toys – all are examples of how seamlessly robotics has entered our worlds. In our factories, robots have been used for decades, whether they are the impressive-looking robot arms that weld cars in their choreographed dances, or they are the big clunky machines that process our food, build packaging, clean bottles, make electronic circuit boards, sort products, and perform a million other functions.

These robots are not frightening, although most caused controversy when introduced. Ever since the invention of factories and the industrial revolution, increasing automation has always taken away some jobs and created new ones in their places. The robots in today’s factories are no different. Likewise, ever since the

creation of the first labour-saving devices for the home, newer and more impressive “automatic” products have been introduced on a regular basis. While the more traditional person might have been terrified at the idea of a machine that cleans clothes, or a machine that washes dishes, these items become normal in a decade and old-fashioned in three.

The kinds of robots that make people nervous are not every day robots that we take for granted. They are the ones that “might go out of control and kill us all!” The robots that look like *Cybermen* from *Doctor Who*, or even worse, the *Borg* from *Star Trek*. In short, people are nervous of the science fiction image of robots, and are largely oblivious to the true nature of robots around them.

This is perhaps a shame, for in recent years improvements in technology have finally enabled us to produce humanoid robots with similar levels of articulation as humans. Improvements in batteries now mean that these robots can stand on their own two feet and do not need to be plugged into the wall. Advances in computer software mean that they can see, hear, and – to a limited extent – understand their senses enough to move around without hitting too many things. Today, being frightened of these robots makes no more sense than being frightened of a washing machine.

In the near future, advances in mechatronics will result in genuine improvements to quality of life. We now have the capability to scan 3D forms and print 3D shapes in many materials including metals, meaning that it is possible to construct artificial limbs with the exact bone and joint structures of the originals. Improvements in neural and muscle interfacing will mean that those involved in accidents will one day have perfectly matching robotic replacement limbs.

In the long term, will we use these technologies to turn ourselves into cyborgs (a merging of flesh and technology) voluntarily? Will we become reliant on technology to communicate, to eat, to move? Will we have electronic devices attached to our faces and ears?

In fact, we have already done all of these things. We are already a society of cyborgs, linked to the Internet like *Borgs* to each other. But we have the choice to behave in this way. We can choose the freedoms (and costs) of the Internet Age, or we can disconnect ourselves. The choice is slowly being eroded as governments increasingly use the newer, cheaper technologies to communicate with their citizens. But technology has always changed societies, and always will. No doubt hunters were horrified at the reliance of farmers on the new technology of crop growing – what if the crops failed? Our new reliance on information technology is just another step along the road of the changing nature of humankind.

In the future we will no doubt have more choices: perhaps to upgrade our limbs or internal organs for improved models – taking pace-makers, hip or knee replacements to a new level. Or upgrade our eyes. Or even improve our brains by adding extra computer support. We will have the choice to mechanise as much of ourselves as we can afford to (or as much as our health service can afford to). As long as these enhancements or replacements do not damage our brains, they should be seen as no more frightening as a hip-replacement or artificial limb. Perhaps one day it will become normal to update your limbs every 25 years (or more often if you are a

sportsperson). There may be specialised versions for different jobs; there might even be fashions for strange limbs just as there are fashions for strange shoes today.

Such changes cannot affect the core nature of our humanity for they are voluntary and not inherited. Every human child will still be born as their genes and environments dictate. How that child chooses to modify him or herself should be a personal choice, as long as that choice harms no-one else. Perhaps the most likely cause of harm will be the conflicts caused by the haves and have-nots, and the subsequent division of societies and cultures into those who embrace the technologies and those who reject them. While forced uptake might be one solution, it seems more probable that the natural tendency of society will result in the majority taking up the more useful technologies. No-one forces people to wear hearing aides or glasses, but life is so much better for those that do, that the results are inevitable.

In parallel, humanoid robots will eventually become as lifelike as ourselves. It is not beyond our technology to make perfect mechanical copies of our bodies, and thus it is likely that such an achievement will be made. These robots really would be able to help in the home and perhaps perform simple repetitive tasks when taught by us. They will not, however, be capable of the intellectual ability of more than a mouse for a long, long time. So while we might have perfect humanoid robots in three or four decades, their brains will take much, much longer to perfect. They might look like us, but they will not be taking over the world or murdering anyone – they will just be complicated human-shaped machines.

Mind Your Head

There is one important technology we have yet to understand or duplicate. It is grown with the help of genes and built using biological nanotechnology of neurons and glial cells. Your mind is in your head, and we are a long way from understanding it.

Any decent humanoid robot, or indeed, any good computer intelligence that might be used to help us run our machines, finances and society, needs a good brain. Attempts at understanding what brains are and how they work have been made for centuries. Today we know there are regions of the brain that seem to be specialised for different functions, yet we know the whole thing is made from the same stuff: neurons, and we know that damage to the brain can be partially overcome as different regions take over new functions. We also know that if there is such a thing as a soul, it has nothing to do with consciousness, personality, memory, emotions, desires, understanding, planning, learning, prediction, perception or movement, for all of these functions are dramatically altered or removed as a result of brain injury.

We can model neurons and networks of neurons quite successfully and give computers the ability to learn and predict simple things. Artificial Neural Networks can process images and detect faces, or control the path of robots past simple obstacles. We can also evolve neural networks within computers and improve their abilities to perform different functions.

This may prove to be the best way to create a true artificial intelligence. The network of neurons in our skulls is simply too complex to untangle – even if we know the connections, the behaviour of each individual neuron is slightly different, as biological studies have shown. Our brains evolved the structure and plasticity they have because of countless challenges we faced in our evolutionary history. Unlike obvious structures such as limbs or blood vessels, the structures of the brain are much harder to assign functions to. So the solution is to evolve new neural networks within computers, set them countless virtual challenges, and evolve ever-more complex brains.

The result (which may be many decades away) will be a different kind of intelligence, evolved to suit the challenges it faced in its own evolutionary history. It will be up to us to guide it through its potential awakening into consciousness (should that ever happen). If it has known nothing except oppression and warfare, clearly it may not be friendly towards us, so the logical route is to ensure that artificial intelligences survive better when they peacefully co-exist with us. There is no reason why a peaceful ecological balance between human and artificial intelligence should not arise, just as ecological balances arise in the natural world as a result of evolution.

Perhaps in 50 or 100 years we may have a true artificial intelligence. A brain that evolved and developed under our guidance (but not according to our preconceived design of what a brain “should” do). Such an artificial mind would be made from neural networks like our own. It could be taught, but not compelled to think in certain ways. When that brain is given a voice, and perhaps a robot body, do we also give it any rights in our society? It will be a threatening vision for many – for this intelligence would have the potential to be more than us. It could live for as long as it chose to, or as long as there was circuitry to support its mind. Regardless of the length of our lives then, this would seem like an immortal being.

There might be some who would claim that such a mind would be nothing more than a soulless machine. But if that mind was a hugely complex network of neurons, evolved over thousands of millions of generations, having faced countless challenges in complex virtual environments, would it matter that it was running on silicon and electricity, not flesh and electricity? If that mind expressed a desire not to be killed, or not to be treated in a way that would harm it, could we ignore that request? Even if it did not make such a request, would we have a right to destroy it?

Such questions are easier to answer if one imagines a mind that is as well-developed as a human brain. But our research will be very slow and incremental. Artificial minds will improve in capability, each one better than the next. When do we decide that a mind deserves any rights? When do we decide that a mind is conscious? What would the punishment be for destroying such a mind? If the requirement is that it should possess a soul or that it should be conscious, then how do you tell? Is a mouse conscious? Does it have a soul? Yet it deserves some rights as a living creature, does it not? Should a mind evolved in a computer not deserve the same rights as its biological counterparts?

There will be some who also try to reverse-engineer the brain. If in two or three centuries from now, a human brain is successfully translated into a neural network running on a computer, will that mind not deserve the right to live as long as it

chooses to, if it is harming no-one else? And what of all the incremental attempts along the way; the partial successes at representing human thought and memory in computers? Which of these should be preserved from destruction? If you happened to be a neural network running on a computer, what would you want? And what would you think if you were copied so that there was more than one of you, just so that the original you could be vivisected?

This is a vision of the future sufficiently far ahead that it is not possible to be accurate. Yet we already evolve artificial life inside computers with the mental capacity of insects. There are enough scientists around the world dedicating their careers to this subject, that it is inevitable that we will create ever more complex artificial minds.

Closing Thoughts

Humans are the products of millions of years of evolution. We have a long history of changing our environments, which includes changing our ideas, beliefs and cultures. But we have a much longer history of more fundamental behaviours, programmed into us by evolution: finding a mate, tribal and territorial behaviours, social interaction, hierarchical leadership and follower roles, nurturing, rearing and nest-making. These deeper instincts link us inexorably to our ancestors, despite differences in thoughts, ideas or languages. The same instincts are likely to make us want to keep such aspects of ourselves constant. Yet for the first time, we will have the choice to modify even these deeper parts of ourselves. By modifying our genes, we will be able to sever the links to our ancestors should we choose to. By using nanotechnology, we would have the ability to duplicate many mechanisms found only in nature, and potentially replace them. Robotics would give us another route to change our physical forms, or duplicate them. And a successful artificial intelligence would show us how intelligence can arise from a different evolutionary path to our own.

Our technology will enable us to transform ourselves in ways never before imagined. It will also enable us to stay the same, despite a need to change. Our dilemma is to choose which parts of us must be altered, and which must be preserved.

Technology never stands still, and where it exists, someone will use it. With our current rate of progress, almost certainly using a combination of genetics, nanotechnology, robotics and artificial intelligence, one day we may well create a new kind of human. Perhaps our grandchildren's grandchildren might have friends belonging to the species *Homo Sapiens Facticius*.

Bibliography

Peter J. Bentley (2002). *Digital Biology. How nature is transforming our technology and our lives*. Simon & Schuster. ISBN: 0743204476

Rodney Brooks (2002) *Flesh and Machines: How Robots Will Change Us*. Pantheon; 1st edition (February 12, 2002) ISBN: 0375420797

Cris Calude, Gheorghe Paun (2000) *Computing with Cells and Atoms: An Introduction to Quantum, DNA and Membrane Computing*. CRC (October 26, 2000) ISBN: 0748408991

John McCaskill (2006) "Using IT to help synthesis artificial living cells." PACE IST Press Release, 16 October 2006.

Douglas Mulhall (2002) *Our Molecular Future: How Nanotechnology, Robotics, Genetics and Artificial Intelligence Will Transform Our World*. Prometheus Books (July 2002) ISBN10: 1573929921

Scientific American (2002) *Understanding Nanotechnology*. Warner Books (December 1, 2002) ISBN: 0446679569

Steve Jones. (2005) *Introducing Genetics*. Totem Books; New Ed edition (December 25, 2005) ISBN: 1840466367