



Edition 1.0

Against Transhumanism

The delusion of technological
transcendence

Richard A.L. Jones

Preface

About the author

Richard Jones has written extensively on both the technical aspects of nanotechnology and its social and ethical implications; his book “[Soft Machines: nanotechnology and life](#)” is published by OUP. He has a first degree and PhD in physics from the University of Cambridge; after postdoctoral work at Cornell University he has held positions as Lecturer in Physics at Cambridge University and Professor of Physics at Sheffield. His work as an experimental physicist concentrates on the properties of biological and synthetic macromolecules at interfaces; he was elected a Fellow of the Royal Society in 2006 and was awarded the Institute of Physics’s Tabor Medal for Nanoscience in 2009. His blog, on nanotechnology and science policy, can be found at [Soft Machines](#).

About this ebook

This short work brings together some pieces that have previously appeared on my blog Soft Machines (chapters [2,4](#) and [5](#)). Chapter 3 is adapted from an early draft of a piece that, in a much revised form, appeared in a [special issue](#) of the magazine IEEE Spectrum devoted to the Singularity, under the title “[Rupturing the Nanotech Rapture](#)”.

Version 1.0, 15 January 2016

The cover picture is The Ascension, by Benjamin West (1801). Source: [Wikimedia Commons](#)

Transhumanism, technological change, and the Singularity



Rapid technological progress – progress that is obvious on the scale of an individual lifetime - is something we take for granted in the modern world. The world I live in – as a prosperous inhabitant of a developed and wealthy country – is quite different to the world of my grandparents and great-grandparents. My everyday living conditions are comfortable, and I have the necessities of life – food and shelter – in abundance. I can travel – even to the other side of the world – with ease, I have access to devices for education and entertainment undreamt of a few decades ago, and when I fall ill or meet accidental injury, many conditions that would have been a death sentence in former times have cures, often quite straightforward ones. This technological progress has led to economic growth – every year, in developed countries at the technological frontier, people work out how to do things a little bit better. On average this results in a percent or two more economic output each year, which compounds over the years to produce exponential growth.

How will this story unfold in the future? For some, it's a period of growth that won't prove sustainable, and so must come to an end. Perhaps we'll run out of the resources that underpin growth – the easily accessible energy sources we've come to rely on will run out sooner or later, or perhaps supplies of some element we've come to depend on will run short. Maybe we'll cause our environment such irreversible damage, for example

by setting off a runaway climate change event, that it will be no longer compatible with civilization. Less dramatically, we might just run out of good ideas, or the conviction to implement them, and growth and progress might slowly peter out. Or our self-destructive tendencies might finally manifest themselves in a final and culture-destroying war, in which we turn the full destructive force of our technologies against ourselves.

Such pessimism isn't entertained by transhumanists, who regard the technological progress of the modern world as the harbinger of much greater change to come. An industrial revolution has led to an information technology revolution, and this, in their view, has begun to change the very essence of what it means to be human. Our destiny, then, is for the technology we create to transform not just our way of life, but the essence of our existence. Perhaps – and possibly within the lifetimes of those already alive - we will see new forms of human beings in which the biological and technological seamlessly merge. Maybe we will be able to leave our biological forms behind entirely to take up an entirely new form of post-human existence.

What are the new technologies that might enable such a fundamental change? Transhumanists believe that these technologies are already with us, or at least are conceivable on a time

scale of years or decades. From current developments in information technology – and particularly the acceleration of computing power implied by Moore’s law – will come true artificial intelligence, an intelligence which surpasses human intelligence, and perhaps will subsume it.

Meanwhile, they say, the realization of a radical vision of nanotechnology [1] will grant us complete control of the material world, effectively eroding the distinction between software and hardware. This, they anticipate, will end scarcity in all forms, leading to a world of material superabundance, and will lead to medical technologies of such power as to render death essentially voluntary.

One of the first consequences of these new technologies, in their nascent forms, will to accelerate the progress of technological innovation itself. So what can we say about the future that this will lead us to?

Transhumanists look forward to a technological singularity, which we should expect to take place on or around 2045, if Ray Kurzweil is to be relied on [2]. The technological singularity is described as something akin to an event horizon, a date at

which technological growth becomes so rapid that to look beyond it becomes quite unknowable to mere cis-humans. In some versions, this is correlated with the time when, due to the inexorable advance of Moore’s Law, machine intelligence surpasses human intelligence and goes into a recursive cycle of self-improvement.

The original idea of the technological singularity is usually credited to the science fiction writer Vernor Vinge, though as we’ll see later, earlier antecedents can be found, for example in the writing of the British Marxist scientist J.D. Bernal. Even amongst transhumanists and singularitarianists there are different views about what might be meant by the singularity, but I don’t want to explore those here. Instead, I note this - when we talk of the technological singularity we’re using a metaphor, a metaphor borrowed from mathematics and physics. Let’s begin by probing the Singularity as a metaphor.

A real singularity happens in a mathematical function, where for some value of the argument the result of the function is undefined. So a function like $1/(t-t_0)$, as t gets closer and closer to t_0 , takes a larger and larger value until when $t=t_0$, the result is infinite. Kurzweil’s thinking about technological advance revolves around the idea of exponential growth, as exemplified by

Moore's Law, so it's worth making the obvious point that an exponential function doesn't have a singularity. An exponentially growing function - $\exp(t/T)$ - certainly gets larger as t gets larger, and indeed the absolute rate of increase goes up too, but this function never becomes infinite for any finite t .

An exponential function is, of course, what you get when you have a constant fractional growth rate - if you charge your engineers to make your machine or device 20% better every year, for as long as they are successful in meeting their annual target you will get exponential growth. To get a technological singularity from a Moore's law-like acceleration of technology, the fractional rate of technological improvement must itself be increasing in time [3].

It isn't totally implausible that something like this should happen - after all, we use technology to develop more technology. Faster computers should help us design more powerful microprocessors. On the other hand, as the components of our microprocessors shrink, the technical problems we have to overcome to develop the technology themselves grow more intractable. The question is, do our more powerful tools outstrip the greater difficulty of our outstanding tasks? The past has certainly seen periods in which the rate of technological progress has under-

gone periods of acceleration, due to the recursive, self-reinforcing effects of technological and social innovation. This is one way of reading the history of the first industrial revolution, of course - but the industrial revolution wasn't a singularity, because the increase of the rate of change wasn't sustained, it merely settled down at a higher value. What isn't at all clear is whether what is happening now corresponds even to a one-off increase in the rate of change, let alone the sustained and limitless increase in rate of change that is needed to produce a mathematical singularity. The hope or fear of singularitarians is that this is about to change through the development of true artificial intelligence. We shall see.

Singularities occur in physics too. Or, to be more precise, they occur in the theories that physicists use. When we ask physics to calculate the self-energy of an electron, say, or the structure of space-time at the centre of a black hole, we end up with mathematical bad behaviour, singularities in the mathematics of the theories we are using. Does this mathematical bad behaviour correspond to bad behaviour in the physical world, or is it simply alerting us to the shortcomings of our understanding of that physical world? Do we really see infinity in the singularity - or is it just a signal to say we need different physics [4]?

The most notorious singularities in physics are the ones that are predicted to occur in the middle of black holes - here it is the equations of general relativity that predict divergent behaviour in the structure of space-time itself. But like other singularities in physics, what the mathematical singularity is signalling to us is that near the singularity, we have different physics, physics that we don't yet understand. In this case the unknown is the physics of quantum gravity, where quantum mechanics meets general relativity. The singularity at the centre of a black hole is a double mystery; not only do we not understand what the new physics might be, but the phenomena of this physical singularity are literally unobservable, hidden by the event horizon which prevents us from seeing inside the black hole. The new physics beyond the Planck scale is unobservable, too, but for a different, less fundamental reason - the particle accelerators that we'd need to probe it would have to be unfeasibly huge in scale and energy, huge on scales that seem unattainable to humans with our current earth-bound constraints. Is it always a given that physical singularities are unobservable? Naked singularities are difficult to imagine, but don't seem to be completely ruled out.

The biggest singularity in physics of all is the singularity where we think it all began - the Big Bang, a singularity in time which it is unimaginable to see through, just as the end of the universe

in a big crunch provides a singularity in time which we can't conceive of seeing beyond. Now we enter the territory of thinking about the creation of the universe and the ultimate end of the world, which of course have long been rich themes for religious speculation. This connects us back to the conception of a technologically driven singularity in human history, as a discontinuity in the quality of human experience and the character of human nature. As we'll see in the next chapter, this conception of the technological singularity is a metaphor that owes a great deal to these religious forbears.

The strange ideological roots of transhumanism



Transhumanists are surely futurists, if they are nothing else. Excited by the latest developments in nanotechnology, robotics and computer science, they fearlessly look ahead, projecting consequences from technology that are more transformative, more far-reaching, than the pedestrian imaginations of the mainstream. And yet, their ideas, their motivations, do not come from nowhere. They have deep roots, perhaps surprising roots, and following those intellectual trails can give us some important insights into the nature of transhumanism now. From antecedents in the views of the early 20th century British scientific left-wing, and in the early Russian ideologues of space exploration, we're led back, not to rationalism, but to a particular strand of religious apocalyptic thinking that's been a persistent feature of Western thought since the Middle Ages[1].

Transhumanism is an ideology, a movement, or a belief system, which predicts and looks forward to a future in which an increasing integration of technology with human beings leads to a qualitative, and positive, change in human nature. It sees a trajectory from a current situation in which certain human disabilities and defects can be corrected, through an increasing tendency to use these technologies to enhance the capabilities of humans, to world in which human and machine are integrated to a cyborg existence. Finally, we may leave all traces of our biological past behind, as humans "upload" their intelligence into powerful computers. These ideas are intimately connected with the

idea of a "Singularity", a moment at which accelerating technological change becomes so fast that we pass through an "event horizon" to a radically unknowable future. According to Ray Kurzweil, transhumanism's most visible and well known spokesman, this event will take place in or around 2045[2].

The idea of transhumanism is associated with three predicted technological advances. The first is a vision of a radical nanotechnology as sketched by K. Eric Drexler, in which matter is effectively digitised, with "*matter compilers*" or "*molecular assemblers*" able to build any object with atomic fidelity [3]. This will be the route to the end of scarcity, and complete control over the material world. The second is a conviction - most vocally expounded by Aubrey de Grey [4] - that it will shortly be possible to radically extend human lifespans, in effect eliminating ageing and death. The third is the belief that the exponential growth in computer power implied by Moore's law, to be continued and accelerated through the arrival of advanced nanotechnology, makes the arrival of super-human level artificial intelligence both inevitable and imminent.

One should be sceptical about all three claims on technical grounds, as later chapters will discuss. But here I want to focus, not on technology, but on cultural history. What is the ori-

gin of these ideas, and how do they tap into deeper cultural currents?

We can summarise the position of singularitarians like Kurzweil like this: we're approaching a world where everything is abundant, where we all live for ever, and where a super-intelligent, super-benevolent entity looks after us all. What's more, this is all going to happen in our lifetimes. We've heard this story before, of course. The connection between singularitarian ideas and religious eschatology is brilliantly captured in the phrase attributed to SF writer Ken MacLeod - the singularity is the *"Rapture of the Nerds"*.

The reason this jibe is so devastatingly effective is that it contains a deep truth. Kurzweil himself recognises the religious overtones of his ideas. In his book *The Singularity is Near* [2] he writes *"Evolution moves towards greater complexity, greater elegance, greater knowledge, greater beauty, greater creativity, and greater knowledge of subtler attributes such as love. In every monotheistic tradition God is likewise described as all of these qualities, only without any limitation..."*, concluding, tellingly, *"...we can regard, therefore, the freeing of our thinking from the severe limitations of its biological form to be an essentially spiritual undertaking."*

This line of thought has a long and fascinating pedigree. One can identify at least two distinct routes by which this kind of eschatological thinking developed to contribute to the modern transhumanist movement. For the first, we can look to the origin of the coinage *"transhumanism"* itself, by the British biologist Julian Huxley (not at all coincidentally, the brother of the author of the dystopian novel, *"Brave New World"*, Aldous Huxley). It was among the British scientific left between the wars that many of the themes of transhumanism were first developed. In a remarkable 1929 essay *The World, The Flesh and the Devil* [5] the Marxist scientist Desmond Bernal gives a slogan for transhumanism *"Men will not be content to manufacture life: they will want to improve on it."* Bernal imagines a process of continuous human enhancement, until we arrive at his version of the Singularity: *"Finally, consciousness itself may end or vanish in a humanity that has become completely etherealized, losing the close-knit organism, becoming masses of atoms in space communicating by radiation, and ultimately perhaps resolving itself entirely into light. That may be an end or a beginning, but from here it is out of sight."*

The title of Bernal's essay hints at the influence of his Catholic upbringing - what was the influence of the Marxism? The aspect of Marxism as a project to fundamentally change human

nature by materialist methods is made very clear in a Leon Trotsky pamphlet from 1923 [6], describing life after the revolution: *“Even purely physiologic life will become subject to collective experiments. The human species, the coagulated Homo sapiens, will once more enter into a state of radical transformation, and, in his own hands, will become an object of the most complicated methods of artificial selection and psycho-physical training.”*

The second route to transhumanism also has a Russian dimension. It comes through the pioneer of rocketry and influential ideologue of space travel, Konstantin Tsiolkovsky [7]. Tsiolkovsky was a key proponent of the philosophy of *Cosmism*, and was profoundly influenced by *Cosmism's* founder, the 19th century philosopher and mystic Nikolai Fyodorov [8]. Fyodorov's system of thought blended religion and materialism to create a vision of transcendence not in a spiritual heaven, but in our own material universe. *“God, according to the Copernican system, is the Father, not only doing everything for people, but also through people, demanding, as the God of the fathers, from everyone alive an uniting for the resuscitation of the dead and for the settling by the resurrected generations of worlds for the governing of these lastly”*. It would be through science, and the complete mastery over the material world that this would give humans, that the apocalypse would happen, on earth: *“We pro-*

pose the possibility and the necessity to attain through ultimately all people the learning of and the directing of all the molecules and atoms of the external world, so as to gather the dispersed, to reunite the dissociated, i.e. to reconstitute the bodies of the fathers such as they had been before their end”.

Both routes converge on on the idea of a *Millennium* - a period, believed to be imminent, when mankind would enjoy a sin-free existence of abundance, not on any spiritual plane, but in this world. The origins of these beliefs can be found in readings of the biblical books of Daniel and the Revelation of St John, but these interpretations are not strictly orthodox - to church fathers such as Augustine, events such as the Millennium and the Second Coming were spiritual events in the lives of individual believers. But millennial thinking was widespread in Europe from the middle ages onwards, in a myriad of fissiparous sects led by prophets and revolutionaries of all kinds. But if there was a single inspiration for these movements, it was probably the 12th century abbot Joachim of Fiore, whose prophetic system was described by the historian Norman Cohn as *“the most influential one known to Europe until the appearance of Marxism”*.

One enormously important legacy of Joachim's prophetic writings was a theory of history as unfolding in a predetermined

way through three great ages. The first, the age of the law, was ended by the coming of Jesus, who initiated a second age, the age of the gospel. But a third age was imminent, the age of the spirit, a thousand year reign of the saints. In Cohn's view, it is Joachim's three age theory of history that has led, via Hegel and Marx, to all theories of historical inevitability; bringing the story up to date, we can include in these the transhumanist convictions about the inevitable progress of technology that have such clear precursors in the views of the British scientific Marxists. In the title of one of Kurzweil's earlier books, "*The age of spiritual machines*", one can hear the echoes of Joachite prophecies down the centuries.

Do these colourful antecedents to transhumanism matter? A thoughtful transhumanist might well ask, what is the problem if an idea has origins in religious thought? We can enjoy for a moment the irony that many transhumanists think of themselves as ultra-rational, skeptical atheists. But looking at the history of thought in general, and of science in particular, we see that many very good ideas have come out of religious thinking (and, for that matter, not everything that came out of Marxism was bad, either). The problem is that mixed up with those good ideas were some very bad and pernicious ones, and people who are ignorant of the history of ideas are ill-equipped to distinguish good from bad. One particular vice of some religious pat-

terns of thought that has slipped into transhumanism, for example, is *wishful thinking*.

A transhumanist might well also point out that just because the antecedent to an idea was misguided in the past, that doesn't mean that as it develops it will always be wrong. After all, people have anticipated being able to fly for a long time, and they looked silly to some right up to the moment when it was possible. That's a good argument, and the proper sceptical response to it is to say "*show me*". If you think that a technology for resurrecting dead people is within sight, we need to see the evidence. But we need to judge actually existing technologies rather than dubious extrapolations, particularly those based on readings of historical trends.

This leads me to what I think is the most pernicious consequence of the apocalyptic and millennial origins of transhumanism, which is its association with technological determinism. The idea that history is destiny has proved to be an extremely bad one, and I don't think the idea that technology is destiny will necessarily work out that well either. I do believe in progress, in the sense that I think it's clear that the material conditions are much better now for a majority of people than they were two hundred years ago. But I don't think the continuation of this

trend is inevitable. I don't think the progress we've achieved is irreversible, either, given the problems, like climate change and resource shortages, that we have been storing up for ourselves in the future. I think people who believe that further technological progress is inevitable actually make it less likely - why do the hard work to make the world a better place, if you think that these bigger impersonal forces make your efforts futile?

**Nanotechnology will not lead
to super-abundance**

3

The dream of molecular nanotechnology is, in effect, to reduce all material things to the status of software. Everything, elementary science teaches us, is made out of atoms, and there are only a very limited number of different types of atoms; therefore, it is clear that if one knows the position and type of every atom in an object, and one has a technology which can place atoms in any arbitrary position consistent with the laws of physics and chemistry, then one can in principle reproduce with absolute fidelity any material thing from its constituent atoms. At a stroke, it is predicted, this will end scarcity – any material or artefact, from the most basic commodities to the most precious objects, will be available for virtually no cost. Replacement parts for humans will be simple to make, and will have capabilities that hugely exceed their natural prototypes. Everything – the economy, the environment, even what it is to be human – will be utterly transformed. Yet this vision, despite its wide currency amongst transhumanists and singularitarians, is not shared by nanotechnologists in industry and academia, who are working on a much more disparate set of technologies, many of which may prove to be useful, lucrative, or even transformative, but which are very different in their philosophy and approach to the transcendent technology foreseen by singularitarians.

Nanotechnology now allows us, in some very controlled circumstances, to place individual atoms in prescribed places. The

most famous demonstration of this was the IBM logo picked out in xenon atoms on a surface using a scanning tunnelling microscope, by Don Eigler in 1991 [1]. Molecular nanotechnology – so called by followers of the ideas of K. Eric Drexler, to distinguish it from the group of disparate technologies presently known as nanotechnology in industry and academia – takes this idea of control to the extreme. Rather than controlling the placement of a single atom or molecule with an essentially macroscopic object like the piezo-electric scanner of a scanning tunnelling microscope, MNT imagines arrays of manipulators that are, themselves, on the nanoscale. It is these “nanofactories” that are able to arrange atoms under software control to any pattern consistent with the laws of physics.

An existence proof for such a nanofactory is provided by the ribosome – the biological machine that makes proteins according to a precise sequence dictated by the genetic code, embodied by the sequence of bases in a stretch of messenger DNA. The existence of the ribosome, according to proponents of MNT, demonstrate the principle of a software controlled nanofactory; the power of MNT will come from taking this principle and re-engineering it to a new level of perfection. In particular, synthetic nanotechnology will not be constrained to use the weak and flexible materials that biology uses, nor will it be limited by the random and contingent design processes of evolution. In-

stead it will be able to use the strongest and stiffest materials available – such as diamond – and will use the rational design principles of mechanical engineering. In this way, we can expect the capabilities of the products of molecular nanotechnology to exceed those of biology “*as a 747 exceeds the capabilities of a sparrow*”.

What is the transhumanist wish list for molecular nanotechnology? At the beginning is material plenty. If one has nanofactories, able to assemble anything from its component atoms, given a supply of a few common chemical feedstocks, the only thing limiting the availability of any object or material whatsoever is the availability of its software representation. Just as any piece of music can be reproduced simply by downloading an MP3 file, a perfect copy of the most intricate piece of engineering, or of the most precious artwork, could be made by simply downloading the file specifying its atomic structure.

But with the ability to place atoms in arbitrary arrangements comes considerable more power than simply the ability to copy existing materials. It should be possible to design materials stronger than anything currently known, and to integrate within these materials almost unlimited functionality, in terms of sens-

ing and information processing, and to provide them with motors of unparalleled power density.

Implicit in these visions is the assumption that we have available very great computing power in tiny packages. This is desired both for its own sake – to make possible macroscopic computers of power and speed orders of magnitude greater than is available today – and to enable the other ambitions of nanotechnology to be fulfilled. Nanofactories will need to be controlled by computers of immense power, while autonomous nano- or micro- scale vehicles, for use in medicine or for military purposes, will need to have substantial on-board computing capability, together with sensors and communications devices.

It is in the area of medicine that the connection between nanotechnology and transhumanist aspirations becomes clearest. Given the molecular origins of many diseases, it is natural to imagine tiny nanoscale robots (nanobots) which can identify this damage and repair it. Infectious diseases present a simpler problem – we can simply imagine nanobots with the ability to detect, chase and destroy undesirable viruses and bacteria.

We can go further – we can imagine nanobots that can supplement the natural functions of the body, or even replace them. For example, it has been proposed to replace the oxygen carrying function of red blood cells by “*respirocytes*” – specialised nanobots designed to concentrate oxygen, transport it and release it on demand. On a larger scale, it should be possible to replace organs or limbs with more robust and effective synthetic replacements. Meanwhile, implants into the brain should allow a direct interface between our biological “*wetware*” and powerful computers. Areas of the brain damaged by accident or degenerative illness could be by-passed by neural prostheses, and new connections established with the senses and with other parts of the body to replace damaged nerves.

We can imagine a situation in which more and more of the body is replaced by more durable and functional synthetic replacements. If a neural prosthesis is possible, why not use implants to give the brain access to computers able to carry out complex calculations for it and to look up data in vast databases? If damaged senses can be repaired, to cure blindness and deafness, why not add additional senses, allowing us to see in the infra-red or directly detect radio waves? In this way, we can imagine replacing more and more of our frail bodies and brains by robustly engineered replacements of vastly more power.

What is to stop us leaving our bodies entirely? Only the need to preserve the contents of our memories and consciousness, our mental identities – and maybe those nanobots will be able to swim through the capillaries of our brains to make that final readout.

How is it envisaged that these dramatic advances will be achieved? The basic outlines were laid out K. Eric Drexler’s 1992 book “*Nanosystems*”, and they have been further developed by Drexler and some coworkers [2]. The underlying idea is summed up in the phrase “*the principles of mechanical engineering applied to chemistry*”.

The basic constructional principle is positionally controlled mechanosynthesis – a tool is used to grasp a reactive molecular fragment, which is then mechanically brought into contact with an appropriate surface, where it will react in the desired way. It is recognised that a lot of care will need to go into designing working chemistries; the surface to which the fragment will be attached will probably need to be passivated, so there will probably need to be quite an intricate sequence of deprotection and reaction steps. For this, and other reasons, it is envisaged that before the general aim of building structures from any elements in any arrangement permitted by the basic laws of chem-

istry can be achieved, structures will be built from a much more limited palette of elements. The favourite candidate for an early working system is sp³ bonded carbon (i.e. diamond), with surfaces passivated, where necessary, with hydrogen. The reasons for this choice include the fact that diamond is very strong and stiff, the friction of diamond interfaces is very low, and some planar hydrogen terminated diamond surfaces are stable against surface reconstructions.

Diamond, then, is the building material of choice. The design philosophy for building the structures and devices that MNT needs is essentially the adoption of the methods of mechanical engineering at the atomic scale. The basic components are the cogs and gears that are such a prominent feature of the imagery of Drexlerian nanotechnology. The purpose of these mechanisms include the distribution of power from nanoscale electrical motors, sorting devices that extract the feedstocks needed for the mechanosynthesis, and even devices for processing information. Drawing inspiration from a previous generation of mechanical computing devices, Drexler envisages ultra-small mechanical computers exploiting “*rod logic*”.

One major focus of the Drexler vision of nanotechnology is manufacturing – to achieve the full benefits of reducing matter

to software, one must be able to use these principles to manufacture usable, macroscopic artefacts. This, of course, presents problems of scale. Currently, one can imagine using a scanning tunnelling microscope to position individual atoms with some precision, but to make a macroscopic object we will need to scale up this operation by factors of order Avogadro’s number. The key conceptual tool introduced to meet this difficulty is the idea of “*exponential manufacturing*”. In its simplest form, one imagines a nanoscale “*assembler*” which can make arbitrary objects on its own scale. Of course, one of the things such an assembler could make would be another assembler, and this assembler itself could go on to make further assemblers, until the resulting exponential growth led to the production of enough of them that they could combine forces to make a macroscopic object.

The vision of autonomous, self-replicating devices multiplying exponentially does, of course, bring to mind the old story of the sorcerer’s apprentice. The idea of out-of-control replicators voraciously consuming the resources of the biosphere goes under the graphic description of the “*grey goo*” problem. To neutralise this threat, Drexler and his coworkers have recently been emphasising alternative manufacturing methods which avoid the use of free living replicators. In this new vision of the “*nanofactory*”, the rather organic picture of reproducing replicators is re-

placed by a Fordian vision of nanoscale mass production, with endlessly repeated elementary operations on countless production lines.

One place where the idea of free-living nanodevices (albeit ones which do not necessarily have the capability to self-replicate) remains prominent is in the projected application of Drexlerian nanotechnology to medicine. Once again, the design principles envisaged here are entirely mechanical, with specialised nanobots produced for the detection and destruction of pathogens, for the repair of damage to cells and for the replacement of underperforming or damaged cells entirely. One application of particular importance for the idea of the singularity is the idea of nanobots as a way of scanning the state of the brain.

The picture of nanoscale mechanical systems and devices outlined in “*Nanosystems*” is supported by a number of detailed calculations, which I will discuss below. But first, it’s worth looking at the general issue of the way this kind of nanotechnology is imagined to relate to biology. In a general way, the idea of sophisticated nanoscale machines owes a great deal to cell biology, which of course offers us a number of remarkable models, in the form of molecular machines such as atp-synthase. Cell

biology also gives an example of software controlled synthesis, in the form of protein synthesis. The ribosome is a remarkable molecular machine which is able to read information from a strand of messenger RNA, and convert the code transferred by the RNA molecule from the archival copy on the DNA molecule into a sequence of amino acids in a protein, which sequence in turn defines the three dimensional structure of the protein and its function. It’s clear that a ribosome, then, fulfils many of the functions envisaged for the assembler.

Cell biology, then, offers us an existence proof that a sophisticated nanotechnology is possible, involving many of the functions imagined for an artificial nanotechnology. Molecular motors convert chemical energy into mechanical energy, active ion channels in membranes effectively sort molecules, and above all, the ribosome carries out atomically accurate software directed synthesis. The strongest argument for the possibility of a radical nanotechnology of the kind Drexler proposes, then, is the fact that biology exists.

The followers of Drexler take this argument further. If biology can produce a sophisticated nanotechnology based on soft materials like proteins and lipids, which seem to an engineering eye to be transparently unsuitable, then how much more power-

ful would our synthetic nanotechnology would be if we could use strong, stiff materials like diamond? And if biology can produce working motors and assemblers using only the random design methods of Darwinian evolution, how much more powerful could the devices be if they were rationally designed using all the insights we've learnt from macroscopic engineering.

In this view, cell biology offers us an existence proof that shows that a radical nanotechnology is possible, but we should expect our artificial nanotechnology hugely to surpass the naturally occurring prototype in power, just as macroscopic technologies like cars and aeroplanes exceed the power of horses and birds.

But there is another point of view [3]. This starts with the recognition that the physical environment in which cell biology takes place is very different from the familiar world for which the assumptions and approximations of mechanical engineering have been developed.

The world of cell biology is the world of water at very low Reynolds numbers, in which water behaves more like the most viscous molasses than the free-flowing liquid that we are familiar with on macroscopic scales.

It is a world dominated by the fluctuations of constant Brownian motion, in which components are ceaselessly bombarded by fast moving water molecules, in which devices will flex and stretch randomly.

Unfamiliar forces of little importance at the macroscopic scale – such as the van der Waals force – dominate, resulting in overwhelming tendencies for components to stick together as soon as they approach. Stickiest of all, in biological environments of the kind that will be important in nanomedicine, we find a plethora of protein molecules, whose tendency to stick underlies a number of undesirable phenomenon, like the rejection of medical implants.

Looked at this way, it becomes difficult at first sight to see how biology works at all, so hostile does the watery nanoscale environment seem to be for engineering. However, biology does work, and it works very well on these scales. The reason for this is, of course, that this is the environment for which evolution has optimised it.

Where human engineering, founded on assumptions appropriate for the macroscopic world, sees features like lack of rigidity, excessive stickiness, and constant random motion as difficulties to be designed around, biology has evolved design principles which exploit these very same features.

The principle of self-assembly, for example, exploits the combination of strong surface forces and random Brownian motion to make the sophisticated structures used by cell biology, for example the assemblies of intricately folded protein molecules, sometimes in association with lipid membranes.

It is a combination of lack of stiffness and the bombardment of Brownian motion that is used in molecular motors, where it is the change in shape of a protein molecule that provides the power stroke converting chemical energy to mechanical energy.

It is the fascinating insights of single molecule biophysics, combined with the atomic-resolution structures of biological machinery that are coming from structural biology, that are allowing us to unravel in detail how the marvellous machinery of the cell actually works.

What is increasingly clear is how much the operating principles differ from those that we are familiar with in macroscopic engineering, and how much these operating principles are optimised for the unfamiliar environment of the nanoscale.

What, then, of the specific feasibility of the proposals for a nanotechnology as “*the principles of mechanical engineering applied to chemistry*”? Even if it is less obvious than it first seems that the approach used by biology can be hugely improved upon by using stiff materials and rational engineering-based design approaches, is there any reason to suppose that the mechanical engineering approach might not in fact work? Although there is no proof of this negative, there are a number of potentially serious issues whose potential impact on the viability of the MNT approach have, in my view, been seriously underestimated by its proponents [4].

The first difficulty relates to the question of whether the “*machine parts*” of molecular nanotechnology - the cogs and gears so familiar from MNT illustrations – are actually stable. These are essentially molecular clusters with odd and special shapes. They have been designed using molecular modelling software,

which works on the principle that if valencies are satisfied and bonds aren't distorted too much from their normal values then the structures formed will be chemically stable.

But this is an assumption - and two features of MNT machine parts make this assumption questionable. These structures typically are envisaged as having substantially strained bonds. And, almost by definition, they have a lot of surface. In fact, we know that the stable structure of clean surfaces is very rarely what you would predict on the basis of simple molecular modelling - they "*reconstruct*". One highly relevant finding is that the stable form of some small diamond clusters actually have surfaces coated with graphite-like carbon.

The second problem relates to the importance of thermal noise and Brownian motion on the nanoscale at room temperature. The issue is that the mechanical engineering paradigm that underlies MNT depends on close dimensional tolerances. But at the nanoscale, at room temperature, Brownian motion and thermal noise mean that parts are constantly flexing and fluctuating in size, making the effective "*thermal tolerance*" much worse than the mechanical tolerances that we rely on in macroscopic engineering. Clearly one answer is to use very stiff materials like diamond, but even diamond may not be stiff enough. Will it

be possible to engineer complex mechanisms in the face of this lack of dimensional tolerance?

The high surface area in highly structured nanoscale systems also has implications for friction and energy dissipation, because of the importance of surface forces. As people attempt to shrink micro-electromechanical systems (MEMS) towards the nanoscale the combination of friction and irreversible sticking (called in the field "*stiction*") causes many devices to fail. MNT systems will have very large internal areas, and as they are envisaged as operating at very high power densities; thus even rather low values of friction (as we can expect between diamond surfaces, especially if the sliding surfaces are not crystallographically related) may in practise compromise the operations of the devices by generating high levels of local heating which in turn will make any chemical stability issues much more serious.

In fact, it is perhaps questionable how useful friction is as a concept at all. What we are talking about is the leakage of energy from the driving modes of the machines into the random, higher frequency vibrational modes that constitute heat. This mode coupling will always occur whenever the chemical bonds are

stretched beyond the range over which they are well approximated by a harmonic potential (i.e. they obey Hooke's law).

One situation in which friction will certainly cause irreversible damage is if uncontrolled, reactive species (such as water or oxygen) get caught up in the mechanisms. The presence of uncontrolled, foreign chemical species will almost certainly lead to molecular adsorption on any exposed surfaces followed by uncontrolled mechanochemistry leading to irreversible chemical damage to the mechanisms. MNT will need an extreme ultra-high vacuum to work, so it is envisaged that the operations of MNT will take place in a completely controlled environment sealed from the outside world - the so-called "*eutatic*" environment.

But, to be useful, MNT devices will need to interact with the outside world. A medical MNT device will need to exist in bodily fluids - amongst the most heterogeneous media its possible to imagine - and a MNT manufacturing device will need to take in raw materials from the environment and deliver the product. In pretty much any application of MNT molecules will need to be exchanged with the surroundings. As anyone who's tried to do an experiment in a vacuum system knows, it's the interfaces between the vacuum system and the outside world - the feed-

throughs - that cause all the problems. Nanosystems includes a design for a "*molecular mill*" to admit selected molecules into the eutactic environment, but again it is at the level of a rough sketch.

The main argument about the feasibility of such selective pumps and valves is the existence of membrane pumps in biology. But, although a calcium pump is fairly effective at discriminating between calcium ions and sodium ions, its operation is statistical - its selectivity doesn't need to be anything like 100%. To maintain a eutactic environment common small molecules like water and oxygen will need to be excluded with very high efficiency.

Of course, none of these issues constitutes a definitive proof that the MNT route will not work. But they certainly imply that the difficulties of implementing this program are going to be substantially greater than implied by proponents of the mechanical approach, and that, if it does prove possible to implement these ideas, the range of environments in which such devices could operate may well be quite limited. If, for example, it should only prove possible to make devices like this to operate in conditions of low temperature and ultra-high vacuum, this would dramatically reduce their impact and economic importance. The hu-

man body, in which these devices would have to work if they were to be important for nanomedicine, is probably one of the most hostile environments one could imagine for this approach to nanotechnology.

The only real proof that MNT will work, of course, will come from an experimental demonstration. Since Drexler's "*Nanosystems*" was published, in 1992, there has been an explosion of work on nanotechnology in academic, industrial and government laboratories around the world. What is striking, though, is how little of this is directly relevant to or inspired by the MNT vision of a mechanically inspired nanotechnology.

Amongst adherents of this vision, the explanation of this is that this view of nanotechnology was deliberately suppressed, and the resources of the National Nanotechnology Initiative were re-directed at the much more incremental problems in chemistry and material science. There may be some truth in the view that there was a conscious effort by some, particularly in the US nano-business community, to distance the idea of nanotechnology from the grand visions of Drexler, with what might be perceived as their disreputable associations with transhumanism, science fiction and the spectre of grey goo[5].

However, this can't be a complete explanation of the lack of experimental progress towards the mechanically inspired MNT vision, for two reasons.

Firstly, it's a view that neglects the fact that the US government funds only a minority of the research in nanotechnology in the world; countries in Europe and Asia may look at the science policy of the USA with some interest, but they certainly don't feel bound to follow it. It is my impression that Drexler has been a much less polarising figure outside the USA than within that country, and while his ideas may still be controversial, I don't sense the imperative to write him out of the history of nanotechnology that I suspect does exist in parts of the US nanotechnology community in business and academia.

Secondly, this view overestimates the extent to which the science enterprise is steered from above. Any scientists able to see a way of making significant progress towards the goals of MNT would try to do it, given the fame and rewards that would undoubtedly follow such an achievement.

Of course, it's possible to conceive of other ways of achieving the grand goals of radical nanotechnology, without appealing to mechanical engineering analogies. Drexler himself has always stressed that early progress in nanotechnology is most likely to follow a biologically inspired path.

DNA itself can be used as a constructional material, exploiting its remarkable properties of self-assembly to make programmed structures and devices, while chemists have made nanoscale shuttles and motors. Self-assembly has become an important paradigm for making nanoscale structures, with exciting applications in drug delivery and tissue engineering.

The interface between nanotechnology and biotechnology is very important. We will see significant medical applications from this, but the timescales to reach the clinic are likely to be long, and unravelling the complexity of the human organism remains a huge barrier to the simple visions of cell-by-cell repair. Brain implants are with us now. But these are crude tools, helpful though they are beginning to be for the severely disabled. A measure of the huge gulf between what's achieved now and the visions of large scale interfaces between the brain and computers, let alone a complete reading of a mental state, is the difference between the 100 billion or so neurons of the human brain

with the 128 or so that a state of the art brain interface can read now.

What are the prospects for truly disruptive breakthroughs coming from what's happening in university labs now? What is possible is that our efforts to copy the design philosophy of biology – the soft nanotechnology of self-assembly and responsive molecules – may allow us to make systems that do, in some very crude way, share some of the characteristics of biological nanosystems. This would be a synthetic biology, in which we have gone further than current efforts in that field to re-engineer radically existing micro-organisms.

And the tantalising possibility remains that we will truly learn to harness the unfamiliar quantum effects of the nanoscale to implement true quantum computing and information processing. The interaction of light, electrons and matter in complex nanostructured materials is leading to exciting new discoveries in plasmonics, optical metamaterials, spintronics and optoelectronics.

Yet the methods available for structuring materials to achieve these effects are still crude. Here, I suspect, is the true killer application for the idea of “*software control of matter*”; devices that

integrate electronics and optics, fully exploiting their quantum character, in truly novel ways. This is a long way from the mechanical paradigm of the molecular nanotechnologists.

We're left with some questions, and a few tentative answers. Are the predictions of some singularitarians, that molecular nanotechnology could arrive within 15 or 20 years, and hasten the arrival of a technological singularity before 2050, plausible? I don't think so; we're already 17 years since the publication of "Nanosystems", and, far from experiencing exponentially accelerating technological progress towards the goals set out in that book, not a lot has been achieved. Is the vision of molecular nanotechnology impossible in principle, or do we simply need more time to get there? It's not possible to say for certain, but the obstacles in the way of the vision seem to be growing.

There's another possibility, though, which does remain interesting. Will there be progress towards some, at least, of the more radical goals of nanotechnology, by routes quite different from those foreseen by the proponents of molecular nanotechnology? I think the answer to this is quite possibly yes; developments in synthetic biology (understood in its broadest sense) and in making systems in which quantum computing is possible

may well have far reaching consequences, even if it isn't at all clear what these are yet.

What are beginning to take shape are new paradigms for radical nanotechnologies; in place of a mechanical paradigm, inspired by macroscopic engineering, we are seeing the development of biological paradigms and quantum paradigms, which acknowledge the different physics that dominates the nanoscale world and makes the best of the opportunities this offers. Perhaps we should applaud Drexler for alerting us to the exciting general possibilities of nanotechnology, while recognising that the trajectories of new technologies rarely run smoothly along the paths foreseen by their pioneers.

**Your mind will not be
uploaded**

4

The recent movie “*Transcendence*” will not be troubling the sci-fi canon of classics, if the reviews are anything to go by. But its central plot device - “*uploading*” a human consciousness to a computer - remains both a central aspiration of transhumanists, and a source of queasy fascination to the rest of us. The idea is that someone's mind is simply a computer programme, that in the future could be run on a much more powerful computer than a brain, just as one might run an old arcade game on a modern PC in emulation mode. “*Mind uploading*” has a clear appeal for people who wish to escape the constraints of our flesh and blood existence, notably the constraint of our inevitable mortality.

In this chapter I want to consider two questions about mind uploading, from a scientific perspective. I’m going to use as an operational definition of “*uploading a mind*” the requirement that we can carry out a computer simulation of the activity of the brain in question that is indistinguishable in its outputs from the brain itself. For this, we would need to be able to determine the state of an individual’s brain to sufficient accuracy that it would be possible to run a simulation that accurately predicted the future behaviour of that individual and would convince an external observer that it faithfully captured the individual’s identity. I’m entirely aware that this operational definition already glosses over some deep conceptual questions, but it’s a good concrete starting point.

My first question is whether it will be possible to upload the mind of anyone reading this now. My answer to this is no, with a high degree of probability, given what we know now about how the brain works, what we can do now technologically, and what technological advances are likely in our lifetimes.

My second question is whether it will ever be possible to upload a mind, or whether there is some point of principle that will always make this impossible. I’m obviously much less certain about this, but I remain sceptical.

This will be a long chapter, going into some technical detail. To summarise my argument, I start by asking whether or when it will be possible to map out the “*wiring diagram*” of an individual’s brain - the map of all the connections between its 100 billion or so neurons. We’ll probably be able to achieve this mapping in the coming decades, but only for a dead and sectioned brain; the challenges for mapping out a living brain at sub-micron scales look very hard.

Then we'll ask some fundamental questions about what it means to simulate a brain. Simulating brains at the levels of neurons and synapses requires the input of phenomenological equations, whose parameters vary across the components of the brain and change with time, and are inaccessible to in-vivo experiment.

Unlike artificial computers, there is no clean digital abstraction layer in the brain; given the biological history of nervous systems as evolved, rather than designed, systems, there's no reason to expect one.

The fundamental unit of biological information processing is the molecule, rather than any higher level structure like a neuron or a synapse; molecular level information processing evolved very early in the history of life. Living organisms sense their environment, they react to what they are sensing by changing the way they behave, and if they are able to, by changing the environment too. This kind of information processing, unsurprisingly, remains central to all organisms, humans included, and this means that a true simulation of the brain would need to be carried out at the molecular scale, rather than the cellular scale. The scale of the necessary simulation is out of reach of any currently foreseeable advance in computing power.

Finally I will conclude with some much more speculative thoughts about the central role of randomness in biological information processing. I'll ask where this randomness comes from, finding an ultimate origin in quantum mechanical fluctuations, and speculate about what in-principle implications that might have on the simulation of consciousness.

Why would people think mind uploading will be possible in our lifetimes, given the scientific implausibility of this suggestion? I ascribe this to a combination of over-literal interpretation of some prevalent metaphors about the brain, over-optimistic projections of the speed of technological advance, a lack of clear thinking about the difference between evolved and designed systems, and above all wishful thinking arising from people's obvious aversion to death and oblivion.

On science and metaphors

I need to make a couple of preliminary comments to begin with. First, while I'm sure there's a great deal more biology to learn about how the brain works, I don't see yet that there's any cause to suppose we need fundamentally new physics to understand it. Of course, new discoveries may change everything, but it seems to me that the physics we've got is quite compli-

cated enough, and this discussion will be couched entirely in currently known, fundamentally physicalist, principles.

The second point is that, to get anywhere in this discussion, we're going to need to immunise ourselves against the way in which almost all popular discussion of neuroscience is carried out in metaphorical language. Metaphors used clearly and well are powerful aids to understanding, but when we take them too literally they can be badly misleading. It's an interesting historical reflection that when computers were new and unfamiliar, the metaphorical traffic led from biological brains to electronic computers. Since computers were popularly described as "*electronic brains*", it's not surprising that biological metaphors like "*memory*" were quickly naturalised in the way computers were described.

But now the metaphors go the other way, and we think about the brain as if it were a computer (I think the brain **is** a computer, by the way, but it's a computer that's so different to man-made ones, so plastic and mutable, so much immersed in and responsive to its environment, that comparisons with the computers we know about are bound to be misleading). So if what we are discussing is how easy or possible it will be to emulate the brain with a man-made computer, the fact that we are so ac-

customed to metaphorical descriptions of brains in terms of man-made computers will naturally bias us to positive answers. It's too easy to move from saying a neuron is analogous to a simple combination of logic gates in a computer, say, to thinking that it can be replaced by one.

A further problem is that many of these metaphors are now so stale and worn out that they have lost all force, and the substance of the original comparison has been forgotten. We often hear, for example, the assertion that some characteristic or other is "*hard-wired*" in the brain, but if one stops to think what an animal's brain looks and feels like there's nothing much hard about it. It's a *soft machine*.

Mapping the brain's "wiring diagram"

One metaphor that is important is the idea that the brain has a "*wiring diagram*". The human brain has about 100 billion neurons, each of which is connected to many others by thin fibres - the axons and dendrites - along which electrical signals pass. There's about 100,000 miles of axon in a brain, connecting at between a hundred to a thousand trillion synaptic connections. It's this pattern of connectivity between the neurons through the axons and dendrites that constitutes the "*wiring diagram*" of the brain. I'll argue below that knowing this "*wiring diagram*" is not

yet a sufficient condition for simulating the operation of a brain - it must surely, however, be a necessary one.

So far, scientists have successfully mapped out the “wiring diagram” of one organism’s nervous system - the microscopic worm *C. elegans*, which has a total of 300 neurons. This achievement was itself a technical tour-de-force, which illustrates what would need to be done to determine the immeasurably more complex “*wiring diagram*” of the human brain. The issue is that these fibres are thin (hundreds of nanometers, for the thinnest of them), very densely packed, and the fibres from a single neuron can pervade a very large volume [1].

Currently electron microscopy is required to resolve the finest connections, and this can only be done on thin sections. Although new high resolution imaging techniques may well be developed, it’s difficult to see how this requirement to image section by section will go away. Magnetic resonance imaging, on the other hand, can image an intact brain, but at much lower resolution - more like millimetres than nanometers. The resolution of MRI derives from the strength of the magnetic field gradient you can sustain. You can have a large gradient over a small volume but if you’re constrained to keep the brain intact that provides quite a hard limit.

Proponents of mind uploading who recognise these difficulties at this point resort to the idea of nanobots crawling through the brain, reading it from the inside. The previous chapter discussed why I think it will be very much more difficult than people think to create such nanobots.

Mapping out all the neural connections of a human brain, then, will be difficult. It probably will be done, on a timescale perhaps of decades. The big but, though, is that this mapping will be destructive, and the brain it is done on will be definitively dead before the process starts. And massive job though it will be to map out this “*micro-scale connectome*”, there's something very important it doesn't tell you - the difference between a live brain and a dead lump of meat - that is what the initial electrical state of the brain is, where the ion gradients are, what the molecules are doing. But more on molecules later...

Modelling, simulation, emulation: why mind uploading might make sense if you believed in intelligent design

If you did have a map of all the neural connections of a human brain, dead or alive, is that enough to simulate it? You could combine the map with known equations for the propagation of electrical signals along axons (the Hodgkin-Huxley equations), models of neurons and models for the behaviour of synapses. This is the level of simulation, for example, carried out in the “*Blue Brain*” project [2].

This is a very interesting thing to do from the point of neuroscience, but it is not a simulation of a human brain, and certainly not of any individual’s brain. It’s a model, which aggregates phenomenological descriptions of the collective behaviours and interactions of components like the many varieties of voltage gated ion channels and the synaptic vesicles. The equations you’d use to model an individual synapse, for example, would have different parameters for different synapses, and these parameters change with time (and in response to the information being processed). Without an understanding of what’s going on in the neuron at the molecular level, these are parameters you would need to measure experimentally for each synapse.

An analogy might make this clearer. Let me ask this question: is it possible to simulate the CPU in your mobile phone? At first sight this seems a stupid question - of course one can predict with a very high degree of certainty what the outputs of the CPU would be for any given set of inputs. After all, the engineers at ARM will have done just such simulations before any of the designs had even been manufactured, using well-understood and reliable design software. But a sceptical physicist might point out that every CPU is different at the atomic level, due to the inherent finite tolerances of manufacturing, and in any case the scale of the system is much too large to be able to simulate at the quantum mechanical level that would be needed to capture the electronic characteristics of the device.

In this case, of course, the engineers are right, for all practical purposes. This is because the phenomenology that predicts the behaviour of individual circuit elements is well-understood in terms of the physics, and the way these elements behave is simple, reliable and robust - robust in the sense that quite a lot of variation in the atomic configuration produces the same outcomes.

We can think of the system as having three distinct levels of description. There is the detailed level of what the electrons and

ions are doing, which would account for the basic electrical properties of the component semiconductors and insulators, and the junctions and interfaces between them. Then there is the behaviour of the circuit elements that are built from these materials - the current-voltage characteristics of the field effect transistors, and the way these components are built up into circuits. And finally, there is a description at a digital level, in which logical operations are implemented.

Once one has designed circuit elements with clear thresholds and strongly non-linear behaviour, one can rely on there being a clean separation between the digital and physical levels. It's this clean separation between the physical and the digital that makes the job of emulating the behaviour of one type of CPU on another one relatively uncomplicated.

But this separation between the physical and the digital in an integrated circuit isn't an accident or something pre-ordained - it happens *because we've designed it to be that way*. For those of us who don't accept the idea of intelligent design in biology, that's not true for brains. There is no clean "*digital abstraction layer*" in a brain - why should there be, unless someone designed it that way? In a brain, for example, the digital is continually remodelling the physical - we see changes in connectivity

and changes in synaptic strength as a consequence of the information being processed, changes, that as we see, are the manifestation of substantial physical changes, at the molecular level, in the neurons and synapses.

The unit of biological information processing is the molecule

Is there any general principle that underlies biological information processing, in the brain and elsewhere, that would help us understand what ionic conduction, synaptic response, learning and so on have in common? I believe there is - underlying all these phenomena are processes of macromolecular shape change in response to a changing local environment.

Ion channel proteins change shape in response to the electric field across the membrane, opening or closing pores; at the synapse shape-changing proteins respond to electrical changes to trigger the bursting open of synaptic vesicles to release the neurotransmitters, which themselves bind to protein receptors to transmit their signal, and complicated sequences of protein shape changes underlie the signalling networks that strengthen and weaken synaptic responses to make memory, remodelling the connections between neurons.

This emphasises that the fundamental unit of biological information processing is not the neuron or the synapse, it's the *molecule*. Dennis Bray, in an important 1995 paper^[3], pointed out that a protein molecule can act as a logic gate through the process of allostery - its catalytic activity is modified by the presence or absence of bound chemicals. In this chemical version of logic, the inputs are the presence or absence of certain small molecules, and the outputs are the molecules that the protein produces, in the presence of the right input chemicals, by catalysis. As these output chemicals can themselves be the inputs to other protein logic gates, complex computational networks linking the inputs and outputs of many different logic gates can be built up. The ultimate inputs of these circuits will be environmental cues - the presence or absence of chemicals or other environmental triggers detected by molecular sensors at the surface of the cells. The ultimate outputs can be short-term - to activate a molecular motor so that a cell swims towards a food source or away from a toxin. Or they can be long term, in activating and deactivating different genes so that the cell builds different structures for itself, or even changes the entire direction of its development.

This is how a single celled organism like an amoeba can exhibit behaviour that is in effect purposeful, that is adaptive to the

clues it detects from the environment around it. All living cells process information this way. In the collective alliance of cells that makes up a multi-cellular organism like a human, all our cells have the ability to process information. The particular cells that specialise in doing information processing and long-ranged communication - the neurons - start out with the general capability for computation that all cells have, but through evolution have developed this capability to a higher degree and added to it some new tricks.

The most important of these new tricks is an ability to control the flow of ions across a membrane in a way that modifies the membrane potential, allowing information to be carried over long distances by the passage of shock waves of membrane potential, and communications to be made between neurons in response to these rapid changes in membrane potential through the release of chemicals at synapses. But, as always happens in evolved systems, these are new tricks built on the old hardware and old design principles - molecules whose shape changes in response to changes in their environment, this shape change producing functional effects (such as the opening of an ion channel in response to a change in membrane potential).

The molecular basis of biological information processing emphasises the limitations of the “*wiring*” metaphor. Determining the location and connectivity of individual neurons, or the “*connectome*” as it’s begun to be called in neuroscience - is necessary, but far from sufficient condition for specifying the informational state of the brain; to do that completely requires us to know where the relevant molecules are, how many of them are present, and what state they’re in.

The brain, randomness, and quantum mechanics

The molecular basis of biological computation means that it isn't deterministic, it's stochastic, it's random. This randomness isn't an accidental add-on, it's intrinsic to the way molecular information processing works. Any molecule in a warm, wet watery environment like the cell is constantly bombarded by its neighbouring water molecules, and this bombardment leads to the constant jiggling we call Brownian motion. But it's exactly the same bombardment that drives the molecule to change shape when its environment changes. So if we simulate, at the molecular level, the key parts of the information processing system of the brain, like the ion channels or the synaptic vesicles, or the broader cell signalling mechanisms by which the neurons remodel themselves in response to the information they carry, we need to explicitly include that randomness.

I want to speculate here about what the implications are of this inherently random character of biological information processing. A great deal has been written about randomness, determinism and the possibility of free will, and I'm largely going to avoid these tricky issues. I will make one important point, though. It seems to me that all the agonising about whether the idea of free will is compatible with a brain that operates through deterministic physics is completely misplaced, because the brain just doesn't operate through deterministic physics.

In a computer simulation, we'd build in the randomness by calls to a pseudo-random number generator, as we compute the noise term in the Langevin equation that would describe, for example, the internal motions of an receptor protein docking with a neurotransmitter molecule. In the real world, the question we have to answer is whether this randomness is simply a reflection of our lack of knowledge? Does it simply arise from a decision we make not to keep track of every detail of each molecular motion in a very complex systems? Or is it “real” randomness, that is intrinsic to the fundamental physics, and in particular from the quantum mechanical character of reality? I think it is real randomness, whose origins can be traced back to quantum fluctuations.

To be clear, I'm not claiming here that the brain is a quantum computer, in the sense that it exploits quantum coherence in the way suggested by Roger Penrose. It seems to me difficult to understand how sufficient coherence could be maintained in the warm and wet environment of the cell. Instead, I want to focus on the origin of the forces between atoms and molecules.

Attractions between uncharged molecules arise from the van der Waals force, which is most fundamentally understood as a fluctuation force, a force that arises from the way randomly fluctuating fields are modified by atoms and molecules. The fluctuating fields in question are the zero-point and thermal fluctuations of the electromagnetic field of the vacuum. Because the van der Waals force arises from quantum fluctuations, the force itself is fluctuating [4], and these random fluctuations, of quantum origin, are sufficient to account for the randomness of the warm, wet nanoscale world.

The complexity theorist Scott Aaronson has recently written an interesting, but highly speculative essay that touches on these issues [5]. Aaronson argues that there is a type of unpredictability about the universe today that arises from the quantum un-

knowability of the initial conditions of the universe. He evokes the quantum no-cloning principle to argue that quantum state functions that have evolved unitarily, without decoherence, from the beginning of the universe - he calls these "*freebits*" - have a different character of uncertainty to the normal types of randomness we deal with using probability distributions. The question then is whether the fundamental unpredictability of "*freebits*" could be connected to some fundamental unpredictability of the decisions made by a human mind. Aaronson suggests it could, if there were a way in which the randomness inherent in the molecular processes underlying the operation of the brain - such as the opening and closing of ion channels - could be traced back to quantum uncertainty. My own suggestion is that the origin of van der Waals forces, as a fluctuation force, in the quantum fluctuations of the vacuum electromagnetic field, offers the connection that Aaronson is looking for.

If Aaronson is correct that his "*freebit*" picture shows how the fundamental unknowability of the quantum initial conditions of the universe translate into a fundamental unpredictability of certain physical processes now, and I am correct in my suggestion that the origins of the van der Waals force in the quantum fluctuations of fields provide a route through which such unpredictability translates into the outcomes of physical processes in the brain, then this provides an argument for mind uploading being

impossible in principle. This is a conclusion I suggest only very tentatively.

Your mind will not be uploaded: dealing with it

But there's nothing tentative about my conclusion that if you are alive now, your mind will not be uploaded. What comforts does this leave for those fearing oblivion and the void, but reluctant to engage with the traditional consolations of religion and philosophy? Transhumanists have two cards left to play.

Cryonics offers the promise of putting your brain in a deep freeze to wait for technology to catch up with the challenges of uploading. It's clear that a piece of biological tissue that has formed a glass at -192 C will, if kept at that temperature, remain in that state indefinitely without significant molecular rearrangements. The question is how much information is lost in the interval between clinical death and achieving that uniform low temperature, as a consequence both of the inevitable return to equilibrium once living systems fail, and of the physical effects of rapid cooling.

Physiological structures may survive, but as we've seen, it's at the molecular level that the fundamentals of biological information processing take place, and current procedures will undoubtedly be highly perturbing at this level. All this leaves aside, of course, the sociological questions about why a future society, even if it has succeeded in overcoming the massive technical obstacles to characterising the brain at the molecular level, would wish to expend resources in reanimating the consciousnesses of the particular individuals who now choose this method of corporeal preservation.

The second possibility that appeals to transhumanists is that we are on the verge of a revolution in *radical life extension*. One favoured route to this involves a revolution in medicine, achieved through radical medical nanotechnology. In this vision, autonomous nanobots patrolling the body are able to identify and deal with disease at the molecular level. We discussed in the previous chapter the practical difficulties standing in the way of this vision.

What is unquestionably true, of course, is that improvements in public health, typical lifestyles and medical techniques have led to year-on-year increases in life expectancy, but this is driven mostly by reducing premature death. The increasingly preva-

lent diseases of old age - particularly neurodegenerative diseases like Alzheimer's - seem as intractable as ever; we don't even have a firm understanding of their causes, let alone working therapies. While substantial fractions of our older people are suffering from cruel and incurable dementias, the idea of radical life extension seems to me to be a hollow joke.

Why should I worry about what transhumanists, or any else, believes in? I don't think the consequences of transhumanist thinking are entirely benign, and I'll expand on that in the next chapter. But there is a very specific concern about science policy that I would like to conclude with here.

Radical ideas like mind uploading are not part of the scientific mainstream, but there is a danger that they can still end up distorting scientific priorities. Popular science books, TED talks and the like flirt around such ideas and give them currency, if not credibility, adding fuel to the *Economy of Promises* [6] that influences - and distorts - the way resources are allocated between different scientific fields.

Scientists doing computational neuroscience don't themselves have to claim that their work will lead to mind uploading to bene-

fit from an environment in which such claims are entertained by people like Ray Kurzweil, with a wide readership and some technical credibility. I think computational neuroscience will lead to some fascinating new science, but you could certainly question the proportionality of the resource it will receive compared to, say, more experimental work to understand the causes of neurodegenerative diseases.

Why transhumanism matters

5

The political scientist Francis Fukuyama once identified transhumanism as “*the world’s most dangerous idea*” [1]. Perhaps a handful of bioconservatives share this view, but I suspect few others do. After all, transhumanism is hardly part of the mainstream. It has a few high profile spokesmen, and it has its vociferous adherents on the internet, but that’s not unusual. The wealth, prominence, and technical credibility of some of its sympathisers - drawn from the elite of Silicon Valley - does, though, differentiate transhumanism from the general run of fringe movements. My own criticisms of transhumanism, as summarized in the previous two chapters, have focused on the technical shortcomings of some of the key elements of the belief package - especially molecular nanotechnology, and the idea of mind uploading. I fear that my critique hasn’t achieved much purchase. To many observers with some sort of scientific background, even those who share some of my scepticism of the specifics, the worst one might say about transhumanism is that it is mostly harmless, perhaps over-exuberant in its claims and ambitions, but beneficial in that it promotes a positive image of science and technology.

But there is another critique of transhumanism, which emphasises not the distance between transhumanism’s claims and what is technologically plausible, as I have done, but the continuity between the way transhumanists talk about technology and the future and the way these issues are talked about in the

mainstream. In this view, transhumanism matters, not so much for its strange ideological roots and shaky technical foundations, but because it illuminates some much more widely held, but pathological, beliefs about technology. The most persistent proponent of this critique is Dale Carrico, whose arguments are summarised in a recent article [2]. Although Carrico looks at transhumanism from a different perspective from me, the perspective of a rhetorician rather than an experimental scientist, I find his critique deserving of serious attention. For Carrico, transhumanism distorts the way we think about technology, it contaminates the way we consider possible futures, and rather than being radical it is actually profoundly conservative in the way in which it buttresses existing power structures.

Carrico’s starting point is to emphasise that there is no such thing as technology, and as such it makes no sense to talk about whether one is “*for*” or “*against*” technology. On this point, he is surely correct; as I’ve already stressed, technology is not a single thing that is advancing at a single rate [3]. There are many technologies, some are advancing fast, some are neglected and stagnating, some are going backwards. Nor does it make sense to say that technology is by itself good or bad; of the many technologies that exist or are possible, some are useful, some not. Or to be more precise, some technologies may be useful to some groups of people, they may be unhelpful to

other groups of people, or their potential to be helpful to some people may not be realised because of the political and social circumstances we find ourselves in.

Transhumanists have a particular tendency to reify technology, since for them it is technology that is the vehicle for redemption and transfiguration. But the urge to reify technology and even to assign agency to it goes much wider - there is even, after all, an influential book called "*What Technology Wants*"[\[4\]](#). As Carrico stresses, the agency belongs to the people who make the technology and the people who use it. Technology doesn't want anything, people do (but they may not always get what they want, by technology or any other means).

Why would you want to think of technology, not as something that is shaped by human choices, but as an autonomous force with a logic and direction of its own? Although people who think this way may like to think of themselves as progressive and futuristic, it's actually a rather conservative position, which finds it easy to assume that the way things will be in the future is inevitable and always for the best. It's a view common among people associated with what we now call "*the technology sector*" - a name which itself speaks to a strangely narrow view of technology, in which the only thing that counts as a technology is a

wireless connection to a database. Serious damage is being done by the assumption that the rapid recent progress we've seen in one particular group of technologies, to do with information and communication, means that we can be confident that other areas of technology in which we urgently need to see faster progress - for example in healthcare and sustainable energy - are proceeding as fast [\[5\]](#).

But can we even talk in an uncomplicated way about "*progress*"? Carrico thinks not - for him, there can be no single direction of enhancement - the most one can say is that things may get better for one particular group of people in one particular set of circumstances: "*There is no general optimization for every outcome, there is no universal training for every profession, but always only enablements freighted with disablements. To say the least, every pursuit has among its costs the other pursuits we might have tried instead*". Here, Carrico is following in a long tradition of critics of utilitarianism - compare, for example, William Blake: "*He who would do good to another must do it in Minute Particulars: general Good is the plea of the scoundrel, hypocrite, and flatterer, for Art and Science cannot exist but in minutely organized Particulars.*"

It's certainly true to say that previous promises of technological progress have not been universally redeemed. Nuclear power turned out not to be "*too cheap to meter*", but instead led to accidents and intractable waste problems. The internet, rather than empowering the masses, seems to be enabling a universal surveillance state. And productivity gains and improvements in manufacturing technology seem to be leading, not to universal leisure and prosperity, but to increasingly unequal concentrations of wealth and power. Perhaps the promise we should be most fearful of now is the framing of climate change as an "*engineering problem with an engineering solution*", with geoengineering a redemptive technology that relieves us of any obligation to develop a more sustainable energy economy.

One can certainly construct lists like this, lists of regrets for previous technologies didn't live up to their promises, and one should certainly try and learn from them. I would want to sound more optimistic, and point out that what this list illustrates is not that we shouldn't have set out to develop those technologies, but that we should have steered them down more congenial roads, and perhaps that we could have done so had we created better political and economic circumstances.

Ultimately, I think I do believe that there has been progress. To speak personally, my own life is much better than the lives of my grandparents and great-grandparents, and while this experience isn't universal, the same could be said by many billions of people across the world. But we must accept that damage has been done in the name of progress, and above all, we must recognise that progress in the future is not inevitable - it needs to be worked for.

So what will the future look like? One could ask a futurologist, but that's not exactly a solid discipline. We have people writing for think-tanks and management consultancies, spotting trends and weaving scenarios. Then we have the transhumanists, projecting technological futures as destiny. At the pinnacle of futurology, we have Ray Kurzweil, a successful inventor, best-selling writer, and Google Director of Engineering, perhaps the world's most high profile transhumanist. To Carrico, there is a continuity between the mainstream futurologists - "*the quintessential intellectuals propping up the neoliberal order*" - and the "*superlative*" futurology of the transhumanists, with its promises of material abundance through nanotechnology, perfect wisdom through artificial intelligence, and eternal life through radical life extension.

The respect with which these transhumanist claims are treated by the super-rich elite of Silicon Valley provides the link. One can make a good living telling rich and powerful people what they want to hear, which is generally that it's right that they're rich and powerful, and that in the future they will become more so (and perhaps will live for ever into the bargain). And in our society the approval of the rich and powerful itself serves to validate the messages that they like to listen to.

The continuities between mainstream futurology and the superlative futurism of the transhumanists come across in some common themes. There's a persistent strand of greedy reductionism, which in talking about economics manifests itself as market fundamentalism, and in social sciences, the just-so stories of evolutionary psychology. Hyperbole is prevalent, and we see overuse and misuse of metaphor (nanotechnology and synthetic biology providing some classic examples [6]).

There's a very interesting sensibility which manifests itself as a hostility to the actual materiality of the world. This begins with the familiar downgrading of the importance of making things compared to processing information, but ends with an actual desire to upload oneself to a disembodied life as a "cyber-angel". It's this that makes clear the essentially religious charac-

ter of the transhumanist quest. In this view, we're soon entering a world where there is no scarcity, everyone lives for ever, and we're watched over by a benevolent super-intelligence and its going to happen in our lifetimes! We've seen this story before, of course, as we saw in chapter 2. In Carrico's words, transhumanists are "*infantile wish-fulfillment fantasists who fancy that they will quite literally arrive at a personally techno-transcendentalizing destination denominated The Future.*" One could argue that transhumanism/singularitarianism constitutes the state religion of Californian techno-neoliberalism, and like all state religions its purpose is to justify the power of the incumbents.

There is, of course, a powerful counter-argument to this kind of scepticism - the reality and scale of the technical and scientific changes in recent years, and the promise of changes yet to come. It's difficult to write critically about technological change in a way that doesn't lay you open to charges of ignorance of this reality. The counterargument to this is that, in the superlative version of futurology, real technological advances and real promise - for better manufacturing, better healthcare, digital access to information, network security and user-friendly software - are co-opted into an essentially crypto-religious project. It's in this sense that the speculative superlative futurology of the tran-

shumanists contaminates the discussion we ought to be having about technology and society's relationship to it.

Another prominent critique of transhumanism comes from the conservative, often religious, strand of thought sometimes labelled "*bioconservatives*". Carrico strongly dissociates himself from this point of view, and indeed regards these two apparently contending points of view, not as polar opposites, but as "*a longstanding clash of reactionary eugenic parochialisms*". Bioconservatives regard the "*natural*" as a moral category, and look back to an ideal past which never existed, just as the ideal future that the transhumanists look forward to will never exist either. Carrico sees a eugenic streak in both mindsets, as well as an intolerance of diversity and an unwillingness to allow people to choose what they actually want.

It's this diversity that Carrico wants to keep hold of, as we talk, not of **The Future**, but of the many possible futures that could emerge from the proper way democracy should balance the different desires and wishes of many different people.

The Californian tech culture in which transhumanism finds its natural home is characterised by a conspicuous conformism

and lack of diversity. This by itself should make us suspicious of a movement that imposes its own parochial vision of an ideal future. But what can we know about the future, except that it is literally unknowable? There's a fundamental unknowability of the consequences of human actions, and this in itself is a fundamental limit, on humanity's knowledge of what it is capable of.

What we need from technology



Transhumanism is wrong about many things, but there's one thing it gets right – the human condition has been qualitatively and irreversibly changed by the technologies we have developed up to now.

In fact, one can go further – our collective existence, as a world community of seven and a half billion people, is existentially dependent on the technologies we have developed up to now. To give just one single example of this, it is the Haber-Bosch process, which uses fossil fuel energy to fix atmospheric nitrogen for use as a fertilizer, that underlies the so-called “green revolution” that has transformed agricultural yields. Without this, between a third and a half of the world's population would starve. But to do this, we have completely re-engineered the earth's natural nitrogen cycle.

And in doing this, we have inadvertently reengineered the earth's atmosphere and climate systems. In our reliance on the Haber-Bosch process – and in many other ways – we have come to rely on the cheap energy provided by fossil fuels. So we have come to rely on technologies for our collective existence, but we know the technologies we rely on are not sustainable and must be replaced.

It's clear what the shape of the new technologies we need should take – renewable energy technologies like solar photovoltaics, together with the necessary energy storage technologies, need to be made cheaper and more scalable, for example. But, contrary to the technological determinism espoused by the transhumanists, technologies don't develop themselves.

Technologies are developed by the focused collective efforts of organized groups of people. The mobilization of these organizations, and choice of which technological problems they direct their efforts towards, is a matter of politics and social organization.

So technologies will advance, and it's essential that they do. But it's our choices that determine how fast, and in what direction, technologies go forward. We need those choices to be driven, not by the delusional dreams of transhumanism, but by the all too real problems we face.

Notes



Notes on chapter 1

[1] This vision was set out in a series of books by K. Eric Drexler, notably [*Engines of Creation*](#) (Anchor, 1986) and the more technical [*Nanosystems: Molecular Machinery, Manufacturing and Computation*](#).(Wiley, 1992)

[2] Ray Kurzweil has been by the most visible and widely publicized proponent of the Singularity, and his book [*The Singularity is Near*](#) (Penguin, 2006), which contains this claim, remains the most coherent narrative of the Singularitarian position.

[3] This argument tacitly assumes that technology is a single thing, and that there's some simple scalar variable that can be used to describe "technological progress" in general. It isn't – there are many different technologies, and at any given time some may be accelerating, some may be stagnating, and some may indeed be regressing. We pick up this point later, and I've discussed it in more detail in my blogpost [*Accelerating change or innovation stagnation?*](#)

[4] One argument for the latter is the close mathematical similarity between some field theories that are used in condensed matter physics and the quantum field theories used in high en-

ergy physics. For field theories in condensed matter, singularities arise because of the neglect of the atomic nature of matter, which makes the breakdown of theories based on an assumption of continua unsurprising.

Notes on chapter 2

- [1] The classic treatment of millennial and apocalyptic thinking in the middle ages is Norman Cohn's [*The Pursuit of the Millennium*](#). (OUP, 1992)
- [2] Ray Kurzweil, [*The Singularity is Near*](#) (Penguin, 2006),
- [3] K. Eric Drexler, [*Engines of Creation*](#) (Anchor, 1986) and [*Nanosystems: Molecular Machinery, Manufacturing and Computation*](#).(Wiley, 1992)
- [4] Aubrey de Grey, [*Ending Ageing*](#), (St Martin's Press, 2008)
- [5] Desmond Bernal, [*The World, The Flesh and the Devil*](#), 1929
- [6] The connections to modern political movements are brought up to date by John Gray in [*Black Mass: apocalyptic religion and the death of Utopia*](#) (Penguin, 2008), which argues that the science based Utopian movements of the twentieth century should be viewed as perverted versions of religious visions of the apocalypse. The Trotsky quotation is from this book.
- [7] Patrick McCray's [*The Visioneers*](#) (Princeton University Press, 2012) (reviewed by me in this blogpost - [*New Dawn Fades*](#)) is a sympathetic account of the connections between the 1960's and 70's space colonies movement, and their inspirations from the thought of Tsiolkovsky, and K. Eric Drexler.

[8] The Fyodorov quotation is from his *The Philosophy of the Common Task*, as quoted by [N. Berdyaev](#).

[9] The correspondence between the three “superlative technologies” of transhumanism and traditional religious superlatives was made by Dale Carrico, whose cogent critique of transhumanism - [Futurological Discourses and Posthuman Terrains](#), *Existenz* 8 47 (2013), is discussed in detail in chapter 5.

Notes on chapter 3

[1] D.M. Eigler and E.K. Schweizer, [*Positioning single atoms with a scanning tunneling microscope*](#). Nature 344 524 (1990).

[2] K.E. Drexler, [*Nanosystems: Molecular Machinery, Manufacturing and Computation*](#).(Wiley, 1992)

[3] This is the argument of my own book [*Soft Machines: nanotechnology and life*](#), R.A.L. Jones, OUP (2004)

[4] For more details of these difficulties for MNT, see my blog-post [*Six challenges for Molecular Nanotechnology*](#).

Ralph Merkle and Rob Freitas responded to my criticisms of MNT in: [*Research challenges for the diamondoid mechanosynthesis path to advanced nanotechnology*](#), to which I replied in: [*Nanobots, nanomedicine, Kurzweil, Freitas and Merkle*](#).

[5] See my 2014 article [*What has nanotechnology taught us about contemporary technoscience?*](#) for my perspective on the development of nanotechnology as a category of academic research.

Notes on chapter 4

[1] For an excellent overview of what's possible now in tracing the connectivity of nervous systems and what the challenges are, see [*The Big and the Small: Challenges of Imaging the Brain's Circuits*](#), J.W. Lichtman & W. Denk, *Science* 334 618 (2011)

[2] For a semi-technical overview of the Blue Brain project, see [*On the Blue Brain Project*](#), H. Markram, *Nature Reviews Neuroscience*, 7 153 (2006)

[3] D. Bray, [*Protein molecules as computational elements in living cells*](#), *Nature* 365 307 (1995). See also Dennis Bray's book [*Wetware*](#) (Yale University Press, 2011)

[4] For the random quality of van der Waals forces, see my blogpost [*Where the randomness comes from*](#).

[5] S. Aaronson, [*The Ghost in the Quantum Turing Machine*](#), 2013

[6] For a critique of this tendency in science policy, see my essay, [*The Economy of Promises*](#), *Nature Nanotechnology*, 3 65 (2008)

Notes on chapter 5

1. Francis Fukuyama, [*Transhumanism*](#). Foreign Affairs, October 23 2009
2. D. Carrico, [*Futurological Discourses and Posthuman Terrains*](#), *Existenz* 8 47 (2013). See also Dale Carrico's blog [*Amor Mundi*](#).
3. I elaborate on this, for example, in R.A.L. Jones, [*Accelerating change or innovation stagnation?*](#), Soft Machines blog March 25 2011
4. K. Kelly, [*What Technology Wants*](#), Viking 2010
5. See my blog post [*The economics of innovation stagnation*](#), Soft Machines blog May 3 2014. Vaclav Smil makes similar points in [*Moore's Curse*](#), *Spectrum* 19 March 2015.
6. See my blog post [*Three things that Synthetic Biology should learn from Nanotechnology*](#), Soft Machines blog April 15 2011.