

Future Technologies, Today's Choices

**Nanotechnology, Artificial Intelligence and Robotics;
A technical, political and institutional map of
emerging technologies.**

Alexander Huw Arnall

Department of Environmental Science and Technology
Environmental Policy and Management Group
Faculty of Life Sciences
Imperial College London
University of London

A report for the Greenpeace Environmental Trust



July 2003

Greenpeace Environmental Trust
Canonbury Villas
London N1 2PN

www.greenpeace.org.uk

ISBN 1-903907-05-5

Published by **Greenpeace Environmental Trust**
Canonbury Villas, London N1 2PN

ISBN 1-903907-05-5

Printed on 100% recycled paper

Contents

List of Tables	2	3. Artificial Intelligence and Robotics	
Abbreviations and Acronyms	3	3.1 Introduction	42
Foreword	4	3.1.1 About AI and robotics	42
<i>Dr Doug Parr, Greenpeace Chief Scientist</i>		3.1.2 Where are we now?	42
Acknowledgements	9	3.2 Aspects of Research	
1. Introduction		3.2.1 Introduction	43
1.1 About nanotechnology, artificial intelligence and robotics	10	3.2.2 Learning	44
1.2 Report structure	10	3.2.3 Reasoning about plans, programs and action	45
1.3 Key references	11	3.2.4 Logical AI	45
2. Nanotechnology		3.2.5 Collaboration	45
2.1 Introduction	12	3.2.6 Perception	46
2.1.1 About nanotechnology	13	3.2.7 Human-computer interaction	46
2.1.2 Where are we now?	13	3.2.8 Public funding	46
2.2 Research and Development		3.3 Applications	
2.2.1 Introduction	14	3.3.1 Introduction	49
2.2.2 Novel materials	14	3.3.2 Intelligent simulation systems	49
2.2.3 Nanotubes	14	3.3.3 Intelligent information resources	50
2.2.4 Tools and fabrication	16	3.3.4 Intelligent project coaches	50
2.2.5 Public funding for research and development	18	3.3.5 Robotics	52
2.3 Applications and Markets		3.3.6 Corporate funding	53
2.3.1 Introduction	21	3.4 Reality and Hype	
2.3.2 Informatics	22	3.4.1 Introduction	54
2.3.3 Pharmaceuticals and medicine	25	3.4.2 Barriers to strong AI	54
2.3.4 Energy	27	3.4.3 A future for strong AI?	55
2.3.5 Defence	29	3.5 Concerns	
2.3.6 Corporate funding	31	3.5.1 Introduction	56
2.4 Reality and Hype		3.5.2 Predictive intelligence	56
2.4.1 Introduction	32	3.5.3 AI and robotic autonomy	57
2.4.2 Molecular nanotechnology	33	3.6 Discussion	59
2.4.3 Fundamental barriers to these visions	35	4. Conclusion	60
2.5 Concerns		Endnotes	62
2.5.1 Introduction	35	References	63
2.5.2 Environmental concerns	36		
2.5.3 Socio-political concerns	37		
2.5.4 Public acceptance of nanotechnology	39		
2.5.5 The regulation debate	40		
2.6 Discussion	41		

List of Tables

Table 1:	Summary of the major nanomaterials currently in research and development and their potential applications.
Table 2:	Applications for new materials and devices resulting from self-assembly and self-organisation.
Table 3:	World-wide government funding for nanotechnology research and development.
Table 4:	Breakdown of spending on the US's National Nanotechnology Initiative from 2001–2003.
Table 5:	Top five government spending on nanotechnology in the Far East in 2002.
Table 6:	Estimated Japanese government nanotechnology research and development expenditures.
Table 7:	Top six European government nanotechnology spending from 1998–2000.
Table 8:	Summary of future estimated global markets in nanotechnology.
Table 9:	Anticipated technological computing developments for 2001–2014.
Table 10:	Maturity of lithography options.
Table 11:	Summary of application areas for informatics.
Table 12:	Summary of application areas for nanoscale pharmaceuticals and medicine.
Table 13:	Summary of applications for energy processing.
Table 14:	US historical funding for technology transitioning into the marketplace.

Abbreviations and Acronyms

AAAI	American Association of Artificial Intelligence	IT	information technology
AI	artificial intelligence	MEMS	micro-electrical-mechanical systems
ANN	artificial neural network	METI	Ministry of Economy, Trade and Industry
ASIMO	Advanced Step in Innovative Mobility	MEXT	Ministry of Education, Culture, Sports, Science and Technology
CBEN	Centre for Biological and Environmental Nanotechnology	MIT	Massachusetts Institute of Technology
CMOS	complementary metal oxide semiconductor	MNT	molecular nanotechnology
CNID	Centre for Nanoscience Innovation for Defence	NASA	National Aeronautics and Space Administration
DARPA	Defence Advanced Research Project Agency	NBIC	nanoscience, biotechnology, information technology and cognitive science
DoD	Department of Defence	NII	National Institute of Informatics
DRAM	dynamic random access memory	NNI	National Nanotechnology Initiative
DTI	Department of Trade and Industry	NSF	National Science Foundation
DDT	dichlorodiphenyltrichloroethane	PC	personal computers
EC	European Commission	PV	photovoltaic
EU	European Union	QIP	quantum information processing
EELD	Evidence Extraction and Link Discovery	RAM	random access memory
EPA	Environmental Protection Agency	RWCP	Real World Computing Project
EPSRC	Engineering and Physical Sciences Research Council	SCI	Scientific Citation Index
FP	Framework Programme	TIA	Total Information Awareness
GM	genetically modified	UCAV	Unmanned Combat Air Vehicle
ISS	Intelligent Simulation System		

Foreword

Dr Doug Parr, Greenpeace Chief Scientist

Why is Greenpeace interested in new technologies? New technologies feature prominently in our ongoing campaigns against genetic modified (GM) crops and nuclear power; however, they are also an integral part of our solutions to environmental problems, including renewable energy technologies, such as solar, wind and wave power, and waste treatment technologies, such as mechanical-biological treatment. So while Greenpeace accepts and relies upon the merits of many new technologies, we campaign against other technologies that have a potentially profound negative impact on the environment.

Greenpeace is in the business of evaluating both future and current threats. Our mission must be to survey upcoming innovations for several reasons. First, we are conscious of unintended (but foreseeable) consequences that impact on the environment. No one intended, for example, that pesticide use in the 1970s and 1980s would have the impact on wildlife that it did. Becoming aware of, and ultimately preventing, the environmental downside of technological developments is clearly a core interest – indeed, the ‘precautionary principle’ has become an important part of international law, such as the Biosafety Protocol on GM organisms. There is also increasing interest in the wider concept of precaution, which is now recognised to include the need for wider participation in the control and direction of technological innovation. This kind of process produces not only a better evidence base, but also more informed decisions. Unintended consequences of a particular new technology cannot always be foreseen; however, if these consequences become a collective problem, it is unreasonable to expect collective responsibility if the decision to proceed with the technology was made by an elite few.

Second, and more subtly, the interests of those who own and control the new technologies

largely determine how a new technology is used. Any technology placed in the hands of those who care little about the possible environmental, health, or social impacts is potentially disastrous. When entire national economies are adapted to take advantage of the economic opportunities offered by new technologies, it is a matter of huge public importance, and the potential environmental and social consequences are clearly of importance to Greenpeace. Global technologies can, particularly in the long term, be of greater significance than Prime Ministers or presidents. Will the power afforded to people and organisations in control of these new technologies be properly controlled? If a single person – a computer-virus writer or a biochemist dealing with anthrax – can cause huge political and financial problems, how much more damage could those with more resources do? Thorough public scrutiny before financial or political commitments to new technologies become irreversible could be hugely beneficial, and surely a matter of democratic rights.

In April and May 2002, Greenpeace and *New Scientist* magazine co-sponsored a series of four debates on the impacts of new technologies, entitled *Science, Technology and the Future*. These debates generated much interest, but the difficulties in locating speakers highlighted the fact that few people could give an overview of either developments in these technologies or their impact in the physical, political and commercial domains. Even more problematic was identifying what the initial technological products would be and their social or environmental consequences.

This prompted Greenpeace to commission a comprehensive review of nanotechnology and artificial intelligence/robotics developments from an organisation with a reputation for technological expertise – Imperial College London. We asked them to document existing applications and to analyse current

research and development (R&D), the main players behind these developments, and the associated incentives and risks.

New Technologies in Context

Beyond the contents of this report, the political and social processes surrounding the introduction of technologies are very important. For example, compare the public response to GM crops in Europe to the wide acceptance of mobile phones. The 'social constitution' of the technology appears key to its acceptability. This social constitution provides the answers to questions such as:

- Who is in control?
- Where can I get information that I trust?
- On what terms is the technology being introduced?
- What risks apply, with what certainty, and to whom?
- Where do the benefits fall?
- Do the risks and benefits fall to the same people (e.g. mobile phones are popular, while mobile phone masts are not)?
- Who takes responsibility for resulting problems?

The evidence presented in this report suggests that, depending on the development pathway, some aspects of nanotechnology might get a rocky ride, as its social constitution is more like that of GM crops than mobile phones. In particular, future disputes surrounding new technology seem certain in the light of globalised, rapid technology transfer. The general public is also increasingly unwilling to accept the word of a company or Government (on the basis of brutal experience), on the risks and benefits of technology, particularly as science and commerce become more closely linked.

At the time of commissioning this report, civil society critiques of the immense R&D and commercial efforts taking place in nanotechnology were quite sparse, but already there are signs that this is changing. In the wake of the furore over genetic modification, the idea of a 'public debate' about new technologies is in vogue, but this has to be meaningful or it will simply promote cynicism.

If public dialogue on science is to mean anything, the approach of nanotechnology is a huge opportunity. Instead of waiting for potential adverse reactions, the scientific community could be proactive. Why not hold a citizens jury to determine scientific priorities on nanotechnology? From each of the agricultural, defence, energy, pharmaceutical, and information technology (IT) sectors (and the numerous cross-overs), the jury could examine current research and its potential. It could suggest which areas need to be highest priority. It would look at the potential short- and long-term applications and the 'blue skies' element necessary for any research programme. Research councils such as the Biotechnology and Biological Sciences Research Council (BBSRC) and the Engineering and Physical Sciences Research Council (EPSRC) in the UK could commit to considering results and utilizing the insights from the findings of such a jury. If dialogue between science and society is to be more than just a sophisticated means of engineering user-acceptance, research councils must adopt this kind of participatory initiative to allow ordinary people to have a say in the types and trajectories of technological innovation.

Nanotechnology

The most common definition of nanotechnology is that of manipulation, observation and measurement at a scale of less than 100 nanometres (one nanometre is one millionth of a millimetre). However, the emergence of a multi-disciplinary field called

'nanotechnology' arises from new instrumentation only recently available, and a flow of public money into a great number of techniques and relevant academic disciplines in what has been described as an 'arms race' between governments. Nanotechnology is really a convenient label for a variety of scientific disciplines which serves as a way of getting money from Government budgets. The figures involved are becoming very large; indeed this report indicates that over US\$2 billion was spent by national governments in 2002, and that these figures will be even larger in 2003. Although the US is said to be the leader, the Japanese government is expected to spend more than the US in 2003. It is also thought that 2002 will prove to be the year when corporate funding matched or exceeded state funds. This is because transnational companies realise that nanotechnology is likely to disrupt their current products and processes, and because the investment community has decided that nanotechnology is the 'next big thing'. Three new business alliances have recently been formed in the US, Europe and Asia, whose sole purpose is to translate research into economically viable products. The UK Government's Department of Trade and Industry estimates that the market for nanotechnology applications will reach over US\$100 billion by 2005. There is now a great deal of momentum behind nanotechnology that has built up into a force which might already struggle to incorporate the outcomes of organised public debate, or meet well-founded public concerns, although by no means will all of the developments be controversial – many will not.

The difficulty in making predictions about the future is that R&D could still take several different directions, and the materials and processes being developed are technology-pushed rather than market-led. After the hype about possible applications, the first real nanotechnology products are starting to appear in the semiconductor

industry – to increase storage densities on microchips – and in the pharmaceutical industry to improve drug targeting and diagnostic aids. Both sectors expect that in the future nanotechnology will provide a dramatic leap forward, but that for now the products seem relatively modest compared to the preceding hype. Other areas of future applications appear to be within the energy sector and defence. With regard to the former, more effective solar cells and highly efficient lighting hold promise on a ten-year time-scale. In the latter, there is no shortage of ideas for military applications and at least two new institutions in the US have been created expressly for the purpose of exploiting nanotechnology for military gain.

Notice that none of these applications deal with the far more distant but highly-publicised prospect of replicator robots or the so-called 'general assembler' – a nano-machine which would produce anything desired given the right raw-materials, and which formed some of the ideas behind Michael Crichton's novel, *Prey*. These applications are currently a long way off due to the difficulties involved in engineering chemical building blocks, information management, and systems design. The challenges are formidable but even so, two US companies are known to be researching molecular assembly. The 'runaway replicator' concerns (also known as the 'grey goo' scenario) raised by Crichton's novel are hideous, but the prospects of it remain way off, and some experts suggest that it would be very difficult to achieve this deliberately, let alone by accident (but see below).

All of this suggests that the development of nanotechnology will go through various different stages, and thus societal debate will need to be an ongoing process rather than a single outcome. There will need to be continual incorporation of the insights from such a debate into policy and product development as the prospects become more

tangible. Already some concerns are becoming evident. Some new materials may constitute new classes of non-biodegradable pollutant about which we have little understanding. Additionally, little work has been done to ascertain the possible effects of nanomaterials on living systems, or the possibility that nanoparticles could slip past the human immune system. Carbon nanotubes are already found in cars and some tennis rackets, but there is virtually no environmental or toxicological data on them. Despite this, of the US\$710 million being spent by the US Government on nanotechnology, only US\$500,000 is being spent on environmental impact assessment, even though a major feature of the product pipeline is that it consists of new materials. Current proposals at EU level on synthetic chemicals regulation are belatedly ensuring that a rule of 'no data, no market' will apply to the basic information about hazardous properties of such chemicals. Knowing the basics about the dangers of new materials is a pre-requisite for effective environmental responsibility. From the Greenpeace perspective, this suggests that whilst 'societal debate' is highly desirable, it is a bit of a luxury if the same old mistakes are being repeated by a new generation of technologists. There is no need for grand, new mechanisms of public involvement to point out the blindingly obvious. With cause for concern, and with the precautionary principle applied, these materials should be considered hazardous until shown otherwise.

Still other concerns are evident in the social arena that revolve around the uses to which the new technology is put – closely linked with ownership and control. One possible dystopian future would be the shift of the control of nanotechnology towards being driven by military needs. This report does not generally support such a prospect at present, although military interest in nanotechnology is considerable. Alternatively, corporate control has been

flagged up by the ETC group, and this implies the pursuit of income streams from those already possessing disposable income. Is the future of nanotechnology then, a plaything of the already-rich? Will the much talked about 'digital-divide' be built upon, exacerbating the inequities present in current society through a 'nano-divide'?

Nanotechnology can only be made available to the poor and to developing countries if the technology remains open to use. Already a company in Toronto has applied for patents on the carbon molecule

Buckminsterfullerene. If ownership of molecules is allowed, the nanotechnology techniques for the precise manipulation of atoms open up a whole new terrain for private ownership. As with genetic engineering where genes have become controlled by patents, things that were once considered universally owned could become controlled by a few.

Artificial Intelligence and Robotics

Unlike the situation for nanotechnology, researchers in artificial intelligence (AI) feel that their work has suffered because of 'public discussion' – hype might be a better term – in the 1960s and 1980s which adversely affected advances in the field after the delivery did not live up to expectations and funding dropped. Many researchers now feel that the goal of mimicking the human ability to solve problems and achieve goals in the real world (the so-called 'strong AI') is neither likely nor desirable because a long series of conceptual breakthroughs is required. Instead the focus is on 'weak AI' – applications that model some, but not all, aspects of human behaviour.

The number of applications for weak AI is growing. AI-related patents in the US increased from 100–1700 between 1989 and 1999, with a total of 3900 patents mentioning related terms. AI systems are generally embedded within larger systems – applications can be found in video games,

speech recognition, and in the ‘data mining’ business sector. Full speech recognition, leading to voice-led Internet access or recognition in security applications, is anticipated relatively soon. However, the ability to extract meaning from natural language recognition remains way off. The data mining market uses software to extract general regularities from online data, dealing in particular with large volumes or patterns humans may not look for. Such systems could be used to predict consumer preferences or extract trends from market data such as patents and news articles. Sales already have reached US\$3.5 billion and are anticipated to be US\$8.8 billion in 2004. Weak AI is already behind systems that detect ‘deviant’ behaviour in credit card use, which has led to improved credit card fraud detection. Potential applications of these techniques to state-security situations are likely to be controversial (see below).

The field of robotics is closely linked to that of AI, although definitional issues abound. ‘Giving AI motor capability’ seems a reasonable definition, but most people would not regard a cruise missile as a robot even though the navigation and control techniques draw heavily on robotics research. After the hype from the 1960s rebounded on investment (as for AI), experts moved away from the idea of complete automation as it was neither desirable nor feasible. Instead, more practical applications have been found, such as in cervical smear screening and, predictably, in the military sphere, where Unmanned Combat Air Vehicles (UCAVs) are being developed, with the hope of fielding them by 2008.

Despite these developments, current AI systems are, it is argued, fundamentally incapable of exhibiting intelligence as we understand it. Current AI is only as smart as the programmer who wrote the code. AI software designers point out that existing computer architecture means that most AI

applications necessarily arise through classical design and programming techniques, rather than new approaches that aim to allow programmes to train and evolve. An example of such an alternative approach may be possible through artificial neural networks, although these systems are so complex that it is not generally possible to follow the reasoning processes that they exhibit.

The funding of AI research is far more difficult to uncover than for nanotechnology as no existing overview seems to exist on the topic, and information on spending is usually placed under a general computer science budget. Industry reportedly leads, with two-thirds of spending on research in computer science, even though public spending has proved an important source of funding in the past, largely because of the field’s high-risk conceptual challenges. Nevertheless it is clear that the US is the leader in spending. It leads, in part, due to military-related institutions, such as the Defence Advanced Research Project Agency (DARPA) and the National Aeronautics and Space Administration (NASA) who used AI systems and robotics for the exploration of Mars. Japan and Europe are also investing (and indeed collaborating) in this field, but are playing catch-up with the US, although Japan remains the leader in using industrial robots.

Far more likely than the tyrannical take-over of society by hyper-intelligent robots (a frequent science fiction theme) or concerns about ‘rights’ for intelligent machines, a more likely issue will be the use of AI systems to spy on people. The US Department of Defense has established a group to look at information gathering and analysis on a huge scale, including government and commercial sources, which would use AI systems to scrutinise the data and extract information about people, relationships, organisations, and activities for counter-terrorism purposes. The concerns about infringing personal privacy or possible

misuse of the data are clear. Furthermore, the use of computer systems for the US National Missile Defence, and possibly for UCAVs, has created a different moral dilemma in that “they will be the first machines given the responsibility for killing human beings without human direction or supervision”.

AI and robotics are likely to continue to creep into our lives without us really noticing. Unfortunately, many of the applications appear to be taking place amongst agencies, particularly the military, that do not readily respond to public concern, however well articulated or thought through.

The Future

Nanotechnology and AI/robotics, together with biotechnology, may well be on a convergent path. In 2001 the National Science Foundation held a large workshop to look at the implications of this convergence and the implications for human abilities and productivity. AI could be boosted by nanotechnology innovations in computing power. Applications of a future nanotechnology general assembler would require some AI and robotics innovations.

Equally, nanotechnology may converge much sooner with biotechnology as it uses the tools and structures of biological systems to generate tiny machines. Although the prospect of general assemblers may be quite distant, self-replicating ‘machines’ that use the tools of biology – and look more like living things than machines – might be closer at hand through the convergence of bio- and nanotechnologies. ‘Grey goo’ might not be a realistic prospect; ‘green goo’ may be closer to the mark – quite how close is difficult to judge on the basis of the evidence in this report. Any creation that posed the prospect of being self-replicating would need to be handled with immense care to ensure environmental protection.

Whether any of the technological futures being scoped out in laboratories are what our general public would like is a question that can only be answered by asking them. If those concerned with the development of new technologies, and nanotechnology in particular, are convinced that the benefits they hope to generate will withstand scrutiny they should have no concerns about engaging and winning public support.

Acknowledgements

Many thanks to my supervisors, Timothy Foxon and Robert Gross, Imperial College London, for their guidance and advice in completing this report; to Douglas Parr, Greenpeace, for commissioning the work; and to Ken Green, University of Manchester Institute of Science and Technology, for his review and commentary.

In addition, I would to thank Gareth Parry, Jenny Nelson, and Murray Shanahan of Imperial College London; Abid Khan of the London Centre for Nanotechnology; and

Olivier Bosch of the International Institute for Strategic Studies (IISS) for allowing me to interview them.

Finally, I am grateful to Jon Glick of the American Association for Artificial Intelligence (AAAI); Andre Gsponer of the Independent Scientific Research Institute (ISRI); Hope Shand of the ETC Group; and Loretta Anania, Ramon Compano, and Jakub Wejcher of the European Community’s Future and Emerging Technologies programme (EC FET) for their assistance.

1. Introduction

1.1 About nanotechnology, artificial intelligence and robotics

The aim of this report is to provide basic, background information of global scope on three emerging technologies: nanotechnology, artificial intelligence (AI) and robotics.

According to the Department of Trade and Industry (DTI), it is important to consider these emerging technologies now because their emergence on the market is anticipated to *'affect almost every aspect of our lives'* during the coming decades (DTI, 2002). Thus, a first major feature of these three disciplines is *product diversity*. In addition, it is possible to characterise them as *disruptive*, *enabling* and *interdisciplinary*.

Disruptive technologies are those that displace older technologies and enable radically new generations of existing products and processes to take over. They can also enable whole new classes of products not previously feasible.

The implications for industry are considerable: companies that do not adapt rapidly face obsolescence and decline, whereas those that do sit up and take notice will be able to do new things in almost every conceivable technological discipline (DTI, 2002).

Nanotechnology is also an enabling technology and, like electricity, the internal combustion engine, or the Internet, its impact on society will be broad and often unanticipated. Unlike these examples, however, nanotechnology is generally considered harder to 'pin down' – it is a general capability that impacts on many scientific disciplines (Holister, 2002). In addition, the interdisciplinary features of these new technologies result in another driving factor for innovation and discovery: they can bring together people from traditionally separate academic groups. For example, the boundaries between physical sciences and life sciences are blurring within these fields.

1.2 Report structure

This report is divided in two main parts: the first examines the field of nanotechnology, and the second looks at AI and robotics.

Furthermore, both parts are divided into six equivalent sections. The Section 1 of each presents an introduction. Following this, the current status of research and development (R&D) is described for both fields in Section 2, with particular attention being paid to the areas of research attracting the most attention. Much of the work described here cuts across traditional academic boundaries and contains a significant technical element. This is because a firm understanding of the nature of the technology itself is essential in understanding its future impact (Holister, 2002). In addition, the perspective taken here is global in scope since governments and corporations world-wide are investing in these areas and research is active on several continents. This suggests that, with international flows of information, technological innovation will be transboundary in nature.

The applications and markets of these emerging technologies are described in Section 3. Specifically, this report aims to highlight the kinds of products which have already been introduced into the global market and those applications due for introduction in the short- and medium-term. In addition, the range of market values that are currently being anticipated are pointed out, although these figures are necessarily highly speculative. Underpinning these R&D and application developments is a wide array of key players. While interest in these technologies is increasing rapidly, particularly in nanotechnology, most of the recent growth of interest comes from those with a strategic interest, such as governments, venture capitalists, large technology-orientated corporations and scientists working in the field (Holister, 2002).

One problem with many of the hundreds of documents written about emerging technologies every year is that they do not distinguish between science and science fiction, let alone the desirable and undesirable in terms of ethics, choice and safety (Ho, 2002b). Thus, Sections 4 and 5 aim to deal with some of these issues: Section 4 separates out some of the hype from the more visionary but solidly placed applications, whereas Section 5 provides an account of the potential environmental and social risks that such uses could pose in the future. Finally, Section 6 highlights some of the key messages of each part.

1.3 Key references

This report has been compiled by consulting a wide variety of sources across the entire spectrum of the debate, from industry advocates to environmental and social pressure groups. In doing so, a number of sources have been particularly important. For the section on nanotechnology, the DTI's (2002) *New Dimensions for Manufacturing:*

UK Strategy for Nanotechnology provides a useful introduction to the field. In addition, Ramon Compano (2001) of the European Commission; Professors J.N. Hay and S.J. Shaw (2000) of the University of Surrey and Defence Evaluation and Research Agency (DERA); Paul Holister (2002) of CMP Cientifica; Ian Miles and Duncan Jarvis (2001) of the National Physical Laboratory (NPL); and Otilia Saxl (2000) of the Institute of Nanotechnology have been used extensively for construction of summary tables. Finally, the National Science Foundation (NSF) report, *Societal Implications of Nanoscience and Nanotechnology*, supplies good information on a wide range of issues (Roco and Bainbridge, 2001). For the section on AI and Robotics, Barbara Grosz and Randall Davis – President and President-Elect of the American Association for Artificial Intelligence (AAAI) – and Daniel Weld of the University of Washington provide some useful technical information.

2. Nanotechnology

2.1 Introduction

2.1.1 About nanotechnology

A major difficulty of characterising nanotechnology is that the field does not stem from one established academic discipline (The Economist, 2002). In fact, there are a number of ways in which nanotechnology may be defined. The most common version regards nanoscience as *'the ability to do things – measure, see, predict and make – on the scale of atoms and molecules and exploit the novel properties found at that scale'* (DTI, 2002).

Traditionally, this scale is defined as being between 0.1 and 100 nanometres (nm), 1 nm being one-thousandth of a micron (micrometre; μm), which is, in turn, one-thousandth of a millimetre (mm). However, as will become clear in the later stages of this study, this definition is open to interpretation, and may readily be applied to a number of different technologies that have no obvious common relationship (The Economist, 2002).

Another way to characterise nanotechnology is by distinguishing between the fabrication processes of top-down and bottom-up. Top-down technology refers to the *'fabrication of nanoscale structures by machining and etching techniques'* (Saxl, 2000). However, top-down means more than just miniaturisation: at the nanoscale level different laws of physics come into play, properties of traditional materials change, and the behaviours of surfaces start to dominate the behaviour of bulk materials. On the other hand, bottom-up technology – often referred to as molecular nanotechnology (MNT) – applies to the creation of organic and inorganic structures, atom by atom, or molecule by molecule (Saxl, 2000). It is this area of nanotechnology that has created the most excitement and publicity. In a mature nanotech world, macrostructures would simply be grown from their smallest constituent components: an

'anything box' would take a molecular seed containing instructions for building a product and use tiny nanobots or molecular machines to build it atom by atom (Miller, 2002). Indeed, as Forrest (1989) points out, *'the development of [bottom-up] technology does not depend upon on discovering new scientific principles. The advances required are engineering.'* In short, fully-fledged bottom-up nanotechnology promises nothing less than complete control over the physical structure of matter – the same kind of control over the molecular and structural makeup of physical objects that a word processor provides over the form and content of text (Reynolds, 2002).

2.1.2 Where are we now?

At present it is clear that this bottom-up 'dream' is far from being realised. As Saxl (2000) notes: *'Top-down and bottom-up can be a measure of the level of advancement of nanotechnology, and nanotechnology, as applied today, is still mainly in the top-down stage.'* This state of relative infancy is often compared in the literature to the information technology (IT) sector in the 1960s, or biotechnology in the 1980s. So, with the science fiction aspects of the debate rapidly receding, industry has now necessarily adopted much more realistic expectations (pers. comm., Abid Khan, London Centre for Nanotechnology, 6 Nov 2002.)

This is not to say, however, that we have long to wait before nanotechnology makes its mark in the global market. In fact, current industry jargon would probably describe nanotechnology as 'coming on stream'. For, although the underlying technologies and their applications are still at an early stage of development, there are applications emerging into the market that are likely to be making a significant impact on the industrial scene by 2006 (Miles and Jarvis, 2001). The best evidence of this move into commercialisation concerns the recent emergence of three alliances whose sole purpose is to translate

this underlying research into economically viable products: the US NanoBusiness Alliance, the Europe Nanobusiness Association, and the Asia-Pacific Nanotechnology Forum. In addition to this, laboratories around the world are working on new approaches and on new ways to scale up nanotechnology to industrial levels. For example, the first factories to manufacture carbon nanotubes and fullerenes are under construction in Japan (DTI, 2002).

In spite of these developments, there has been criticism recently over the amount of hype and, consequently, funding that research into nanoscience and nanotechnology has received. For example, the much-heralded US National Nanotechnology Initiative (NNI) has been criticised for using ‘nano’ as a convenient tag to attract funding for a whole range of new science and technologies (e.g. see Roy, 2002). This reinvention is one way of attracting more money because politicians like to feel they are putting money into something new and exciting (pers. comm., Gareth Parry, Imperial College London, 22 Nov 2002). For these reasons, the nanotechnology sector is far broader than you would usually expect to see and the resulting lack of a clear definition is hampering meaningful discussion of its potential costs or benefits. Thus, if we use the standard definition given above, we can say that nanoscience and technology have been around for several decades, particularly in research, development, and manufacturing in IT. Rather, it is the wide availability of tools and information to diverse scientific communities that has generated the current interest in this area (Chaudhari, 2001).

2.2 Research and Development

2.2.1 Introduction

The absence of a universally accepted strict definition of nanotechnology has allowed the research emphasis to broaden, encompassing many areas of work that have traditionally been referred to as chemistry or biology (DTI, 2002). Thus, the first major characteristic of activity grouped under this section is that contemporary R&D cuts across a wide range of industrial sectors. In some cases, major markets are fairly well defined. The food industry serves as a good example here, where there are significant drivers at work (pers. comm., Abid Khan, London Centre for Nanotechnology, 6 Nov 2002). To illustrate, ‘smart’ wrappings for the food industry (that indicate freshness or otherwise) are close to the market (Saxl, 2000). By 2006, beer packaging is anticipated by industry to use the highest weight of nano-strengthened material, at 3 million lbs., followed by meats and carbonated soft drinks. By 2011, meanwhile, the total figure might reach almost 100 million lbs. (nanotechweb.org, 2002). In other cases, important applications are identified but the eventual market impacts are more difficult to predict. For example, nanotechnology is anticipated to yield significant advances in catalyst technology. If these potential applications are realised then the impact on society will be dramatic as catalysts, arguably the most important technology in our modern society, enable the production of a wide range of materials and fuels (Saxl, 2000).

A second characteristic of current work in this area is that the kinds of materials and processes being developed are necessarily ‘technology pushed’: urged on by the potential impacts of nanotechnology, the R&D community is achieving rapid advances in basic science and technology. This level of scientific interest is gauged by Compano and Hullman (2001) who examine the world-

wide number of publications in nanotechnology in the Science Citation Index (SCI) database. They conclude that for the period between 1989 and 1998 the average annual growth rate in the number of publications is an ‘impressive’ 27%. This rise in interest is not confined to a small number of central repositories however (Smith, 1996). Instead, research is spread across more than 30 countries that have developed nanotechnology activities and plans (Holister, 2002). In this way, Compagno and Hullman (2001) also examine the distribution of this interest. Based upon their findings, the most active is the US, with roughly one-quarter of all publications, followed by Japan, China, France, the UK and Russia. These countries alone account for 70% of the world’s scientific papers on nanotechnology. In particular, for China and Russia the shares are outstanding in comparison with their general presence in the SCI database and show the significance of nanoscience in their research systems.

2.2.2 Novel materials

The third major characteristic of activity grouped under this section concerns that fact that nanotechnology is primarily about making things (Holister, 2002). For this reason, most of the existing focus of R&D centres on ‘nanomaterials’: novel materials whose molecular structure has been engineered at the nanometre scale (DTI, 2002). Indeed, Saxl (2000) states that: *‘material science and technology is fundamental to a majority of the applications of nanotechnology.’* Thus, many of the materials that follow (Table 1) involve either bulk production of conventional compounds that are much smaller (and hence exhibit different properties) or new nanomaterials, such as fullerenes and nanotubes (ETC Group, 2002a). The markets range of nanomaterials are considerable. Indeed, it has been estimated that, aided by nanotechnology, novel materials and processes can be expected to have a market

impact of over US\$340 billion within a decade (Holister, 2002).

2.2.3 Nanotubes

Nanotubes provide a good example of how basic R&D can take off into full-scale market application in one specific area. Described as *‘the most important material in nanotechnology today’* (Holister, 2002), nanotubes are a new material with remarkable tensile strength. Indeed, taking current technical barriers into account, nanotube-based material is anticipated to become 50–100 times stronger than steel at one-sixth of the weight (Anton et al., 2001). This development would dwarf the improvements that carbon fibres brought to composites. Harry Kroto, who was awarded the Nobel Prize for the discovery of C60 Buckminsterfullerene, states that such advances will take *‘a long, long time’* to achieve (2010 Nanospace Odyssey lecture, Queen Mary University, 6 Jan 2003), the first applications of nanotubes being in composite development. However, if such technologies do eventually arrive, the results will be awesome: they will *‘be equivalent to James Watt’s invention of the condenser’*, a development that kick-started the industrial revolution. The concept of the space elevator serves as a good illustration of the kind of visionary thinking that recent nanotube development has inspired. The idea of a ‘lift to the stars’ is not itself particularly new: a Russian engineer, Yuri Artutanov, penned the idea of an elevator – perhaps powered by a laser that could quietly transport payloads and people to a space platform – as early as 1960 (cited in Cowen 2002). However, such ideas have always been hampered by the lack of material strength necessary to make the cable attachment. The nanotube may be the key to overcoming this longstanding obstacle, making the space elevator a reality in just 15 years time (Cowen, 2002). This development, though, will rely on the successful incorporation of nanotubes into fibres or ribbons and successfully avoiding

Table 1: Summary of the major nanomaterials currently in research and development and their potential applications.

Material	Properties	Applications	Time-scale (to market launch)
Clusters of atoms			
Quantum wells	Ultra-thin layers – usually a few nanometres thick – of semiconductor material (the well) grown between barrier material by modern crystal growth technologies (Saxl, 2000). The barrier materials trap electrons in the ultra-thin layers, thus producing a number of useful properties. These properties have led, for example, to the development of highly efficient laser devices.	CD players have made use of quantum well lasers for several years. More recent developments promise to make these nanodevices commonplace in low-cost telecommunications and optics.	Current – 5 years
Quantum dots	Fluorescent nanoparticles that are invisible until ‘lit up’ by ultraviolet light. They can be made to exhibit a range of colours, depending on their composition (Miles and Jarvis, 2001).	Telecommunications, optics.	7–8 years
Polymers	Organic-based materials that emit light when an electric current is applied to them and vice versa (pers. comm., Jenny Nelson, Imperial College London, 2 Dec 2002).	Computing, energy conversion.	?
Grains that are less than 100nm in size			
Nanocapsules	Buckminsterfullerenes are the most well known example. Discovered in 1985, these C60 particles are 1nm in width.	Many applications envisaged e.g. nanoparticulate dry lubricant for engineering (Saxl, 2000).	Current – 2 years
Catalytic nanoparticles	In the range of 1–10 nm, such materials were in existence long before it was realised that they belonged to the realms of nanotechnology. However, recent developments are enabling a given mass of catalyst to present more surface area for reaction, hence improving its performance (Hay and Shaw, 2000). Following this, such catalytic nanoparticles can often be regenerated for further use.	Wide range of applications, including materials, fuel and food production, health and agriculture (Hay and Shaw, 2000).	Current – ?
Fibres that are less than 100nm in diameter			
Carbon nanotubes	Two types of nanotube exist: the single-wall carbon nanotubes, the so-called ‘Buckytubes’, and multilayer carbon nanotubes (Hay and Shaw, 2000). Both consist of graphitic carbon and typically have an internal diameter of 5 nm and an external diameter of 10 nm. Described as the ‘most important material in nanotechnology today’ (Holister, 2002), it has been calculated that nanotube-based material has the potential to become 50–100 times stronger than steel at one sixth of the weight.	Many applications are envisaged: space and aircraft manufacture, automobiles and construction. Multi-layered carbon nanotubes are already available in practical commercial quantities. Buckytubes some way off large-scale commercial production (Saxl, 2000).	Current – 5 years
Films that are less than 100nm in thickness			
Self-assembling monolayers (SAMs)	Organic or inorganic substances spontaneously form a layer one molecule thick on a surface. Additional layers can be added, leading to laminates where each layer is just one molecule in depth (Holister, 2002).	A wide range of applications, based on properties ranging from being chemically active to being wear resistant (Saxl, 2000).	2–5 years
Nanoparticulate coatings	Coating technology is now being strongly influenced by nanotechnology. E.g. metallic stainless steel coatings sprayed using nanocrystalline powders have been shown to possess increased hardness when compared with conventional coatings (Saxl, 2000).	Sensors, reaction beds, liquid crystal manufacturing, molecular wires, lubrication and protective layers, anti-corrosion coatings, tougher and harder cutting tools (Holister, 2002).	5–15 years
Nanostructured materials			
Nanocomposites	Composites are combinations of metals, ceramics, polymers and biological materials that allow multi-functional behaviour (Anton et al., 2001). When materials are introduced that exist at the nanolevel, nanocomposites are formed (Hay and Shaw, 2000), and the material’s properties – e.g. hardness, transparency, porosity – are altered.	A number of applications, particularly where purity and electrical conductivity characteristics are important, such as in microelectronics. Commercial exploitation of these materials is currently small, the most ubiquitous of these being carbon black, which finds widespread industrial application, particularly in vehicle tyres (Hay and Shaw, 2000).	Current – 2 years
Textiles	Incorporation of nanoparticles and capsules into clothing leading to increased lightness and durability, and ‘smart’ fabrics (that change their physical properties according to the wearer’s clothing) (Holister, 2002).	Military, lifestyle.	3–5 years

various atmospheric dangers, such as lightning strikes, micrometeors, and human-made space debris.

The market impetus behind such developments, then, is clear: the conventional space industry is anticipated as the first major customer, followed by aircraft manufacturers. However, as production costs drop (currently US\$20–1200/g), nanotubes are expected to find widespread application in such large industries as automobiles and construction. In fact, it is possible to conceive of a market in any area of industry that will benefit from lighter and stronger materials (Holister, 2002). It is expectations such as these that are currently fuelling the race to develop techniques of nanotube mass-production in economic quantities. The ETC Group (2002b) states that there are currently at least 55 companies involved in nanotube fabrication and that production levels will soon be reaching 1 kg/day in some companies. For example, Japan's Mitsui and Co. will start building a facility in April 2003 with an annual production capacity of 120 tons of carbon nanotubes (Fried, 2002). The company plans to market the product to automakers, resin makers and battery makers. In fact, the industry has grown so quickly that Holister (2002) believes that the number of nanotube suppliers already in existence are not likely to be supported by available applications in the years to come. Fried (2002) also supports this contention, stating that the *'carbon nanotube field is already over-saturated'*.

2.2.4 Tools and fabrication

It is a simple statement of fact that in order to make things you must first have the fabrication tools available. Therefore, many of the nanomaterials covered above are co-evolving with a number of enabling technologies and techniques. These tools provide the instrumentation needed to examine and characterize devices and effects during the R&D phase, the manufacturing

techniques that will allow the large-scale economic production of nanotechnology products, and the necessary support for quality control (DTI, 2002). Because of the essential nature of this category, its influence is far greater than is reflected in the size of the economic sectors producing these products. For this reason, the tools and techniques highlighted below have a strong commercial future and the greatest number of established companies (pers. comm., Gareth Parry, Imperial College London, 22 Nov 2002). The following sections cover methods for top-down and bottom-up manufacture, software modelling and nanometrology. However, in the near future, this area will mainly feature extensions of conventional instrumentation and top-down manufacturing. More futuristic molecular scale assembly remains distant (Miles and Jarvis, 2001).

2.2.4.1 Top-down manufacture

Scanning Probe Microscope. This is the general term for a range of instruments with specific functions. Fundamentally, a nanoscopic probe is maintained at a constant height over the bed of atoms. This probe can be positioned so close to individual atoms that the electrons of the probe-tip and atom begin to interact. These interactions can be strong enough to 'lift' the atom and move it to another place (pers. comm., Gareth Parry, Imperial College London, 22 Nov 2002).

Optical Techniques. These techniques can be used to detect movement – obviously important in hi-tech precision engineering. Optical techniques are, in theory, restricted in resolution to half the wavelength of light being used, which keeps them out of the lower nanoscale, but various approaches are now overcoming this limitation (Holister, 2002).

Lithographics. Lithography is the means by which patterns are delineated on silicon chips and micro-electrical-mechanical systems (MEMS). Most significantly, optical lithography is the dominant exposure tool in

use today in the semiconductor industry's Complementary Metal Oxide Semiconductor (CMOS) process

2.2.4.2 Bottom-up manufacture

The tools here support rather more futuristic approaches to large-scale production and nanofabrication based on bottom-up approaches, such as nanomachine production lines (Miles and Jarvis, 2001). This approach is equivalent to building a car engine up from individual components, rather than the less intuitive method of machining a system down from large blocks of material. Indeed, although such techniques are still in their infancy, the DTI (2002) report a recent movement away from top-down techniques towards self-assembly within the international research community. Scientists and engineers are becoming increasingly able to understand, intervene and rearrange the atomic and molecular structure of matter, and control its form in order to achieve specific aims (Saxl, 2000).

Self-assembly and self-organisation. Self-assembly refers to the tendency of some materials to organise themselves into ordered arrays (Anton et al., 2001). This technique potentially offers huge economies, and is considered to have great potential in nanoelectronics. In particular, the study of the self-assembly nature of molecules is proving to be the foundation of rapid growth in applications in science and technology. For example, Saxl (2000) reports that the Stranski–Krastanov methods for growing self-assembly quantum dots has rendered the lithographic approach to semiconductor quantum dot fabrication virtually obsolete. In addition, self-assembly is leading to the fabrication of new materials and devices. The former area of materials consists of new types of nanocomposites or organic/inorganic hybrid structures that are created by depositing or attaching organic molecules to ultra-small particles or ultra-thin manmade layered structures (Hay and Shaw, 2000).

Similarly, the latter area of devices range from the production of new chemical and gas sensors, optical sensors, solar panels and other energy conversion devices, to bio-implants and *in vivo* monitoring. The basis of these technologies is an organic film (the responsive layer) which can be deposited on a hard, active electronic chip substrate. The solid-state chip receives signals from the organic over-layer as it reacts to changes in its environment, and processes them. The applications for these new materials and devices are summarised in Table 2.

Table 2: Applications for new materials and devices resulting from self-assembly and self-organisation.

Name	Technique	Application
<i>New materials</i>		
Sol-gel technology (Miles and Jarvis, 2001)	Inorganic and organic component combination.	The design of different types of materials; functional coatings.
Intercalation of polymers (Miles and Jarvis, 2001)	Intercalation of polymers with other materials (DNA, drugs).	Toxicity testing, drug delivery and drug performance analysis.
Nano-emulsions (Saxl, 2000)	Nanoparticle size and composition selected.	Production of required viscosity and absorption characteristics.
Biomimetics (Anton et al., 2001)	Design of systems, materials and their functionality to mimic nature.	High strength, structural applications, such as artificial bones and teeth.
<i>New devices</i>		
Field-sensing devices (Saxl, 2000)	Combination of molecular films with optical waveguides and resonators.	Biosensing and optical switching.
Material-sensing devices (Saxl, 2000)	Surfaces of liquid crystals or thin membranes and other organic compounds can be used to detect molecules via structural or conductive changes.	Gas and chemical sensing.

2.2.4.3 Software Modelling

Molecular modelling software is another fabrication technique of wide-ranging applicability as it permits the efficient analysis of large molecular structures and substrates (Miles and Jarvis, 2001). Hence, it is much used by molecular

nanotechnologists, where computers can simulate the behaviour of matter at the atomic and molecular level. In addition, computer modelling is anticipated to prove essential in understanding and predicting the behaviour of nanoscale structures because they operate at what is sometimes referred to as the mesoscale, an area where both classical and quantum physics influence behaviour (Holister, 2002).

2.2.4.4 Nanometrology

Fundamental to commercial nanotechnology is repeatability, and fundamental to repeatability is measurement. Nanometrology, then, allows the perfection of the texture at the nanometre and sub-nanometre level to be examined and controlled. This is essential if highly specialised applications of nanotechnology are to operate correctly, for example X-ray optical components and mirrors used in laser technologies (Saxl, 2000).

2.2.5 Public funding for research and development

The main reason for government interest in nanotechnology is strategic: to achieve an advantageous position so that when nanotech applications begin to have a significant effect in the world economy, countries are able to exploit these new opportunities to the full. Harper (2002), who describes the current situation as a global 'arms race', puts these ideas into perspective:

'You only have to look at how IT made a huge difference to both the US economy and US military strength to see how crucial technology is. Nanotechnology is an even more fundamental technology than IT. Not only has it the ability to shift the balance of military power but also affect the global balance of power in the energy markets.'

This emphasis on military power is well founded: Smith (1996) echoes this sentiment when he speculates that it is entirely possible

that much, or even most, US government research in the field is concentrated in the hands of military planners.

Levels of public investment in nanotechnology are reminiscent of a growing strategic interest: this is an area that attracts both large and small countries. Global R&D spending is currently around US\$4 billion (ETC Group, 2002a), with public investment increasing rapidly (503% between 1997 and 2002 across the 'lead' countries²). Table 3 summarises these rises.

Table 3: World-wide government funding for nanotechnology research and development (US\$million).

Area	1997	1998	1999	2000	2001	2002	2003
US*	116	190	255	270	422	604	710
Western Europe	126	151	179	200	225	~400	NA
Japan	120	135	157	245	465	~650	NA
Others**	70	83	96	110	380	~520	NA
Total	432	559	687	825	1502	2174	NA
(% of 1997)	100	129	159	191	348	503	NA

NA: not available.

* Excluding non-federal spending e.g. California.

** 'Others' includes Australia, Canada, China, Eastern Europe, the Former Soviet Union, Singapore, Taiwan and other countries with nanotechnology R&D. For example, in Mexico there are 20 research groups working independently on nanotechnology. Korea, already a world player in electronics, has an ambitious 10-year programme to attain a world-class position in nanotechnology (DTI, 2002).

2.2.5.1 The US

The US is widely considered to be the world-leader in nanoscale science research (Saxl, 2000). Certainly, in terms of leading centres for nanotechnology research, the USA dominates, with eight institutions making the DTI (2002) top list of 13. These centres are University of Santa Barbara, Cornell University, University of California at Los Angeles, Stanford University, IBM Research Laboratories, Northwestern University, Harvard University and the Massachusetts Institute of Technology (MIT). In total, more than 30 universities have announced plans for nanotech research centres since 1997 (Leo, 2001). Further, the US is widely

regarded as the benchmark against which nanotechnology funding should be compared (Roman, 2002). Indeed, Howard (2002) states that, *'while other governments are investing in a range of nanotechnology research, the US effort is by far the most substantial.'* From 1985–1997 the total support for projects related to nanotechnology was estimated at US\$452 million, coming in roughly equal parts from the NSF, various industrial sponsorship, and other government funding. Then in 2000, the much-heralded NNI was launched – a multi-agency programme designed to provide a big funding boost for nanotechnology. There are currently 10 US government partners in the NNI³. These are shown in Table 4.

Table 4: Breakdown of spending on the US's National Nanotechnology Initiative from 2001–2003 (US\$million).

Recipient	2001 actual	2002 estimate	2003 proposed
National Science Foundation	145	199	221
Department of Defence	125	180	201
Department of Energy	78	91	139
National Aeronautics and Space Administration	0	46	49
National Institute of Health	40	41	43
National Institute of Standards and Technology	28	37	44
Environmental Protection Agency	5	5	5
Department of Transportation	0	2	2
US Department of Agriculture	0	2	5
Department of Justice	1	1	1
Total	422	604	710

DTI, 2002.

Table 4 shows that the NSF and Department of Defence (DoD) are the two major recipients of investment in nanoscience and technology R&D. Indeed, the NSF has designated *'nanoscale science and engineering'* as one of its six priority areas, while the DoD has dedicated its funding to

elaborating a *'conceptual template for achieving new levels of war-fighting effectiveness'* (DoD, 2002). This table provides a fairly accurate picture of current research priorities in the US. However, state funding, which can sometimes be substantial, is not included in the estimates. For example, the state of California, which is home to virtually all the work in molecular nanotechnology, has invested US\$100 million in the creation of a California Nanosystems Institute. And neither are the figures static; levels of funding are anticipated to increase rapidly once the economic benefits of US funding begin to be felt, whether in new company start-up activity, or progress towards military or social goals.

2.2.5.2 Far East

Table 5 shows the levels of 2002 government spending on nanotechnology within five countries in the Far East. On average, these figures are lower than in the US although, given the increased purchasing power in countries such as China, they may be considered as relatively high (Roman, 2002). However, while the figures given are up-to-date, the time-scales over which they operate are ambiguous.

Table 5: Top five government spending on nanotechnology in the Far East in 2002 (US\$million).

Country	Spending
Japan	750
China	200
Korea	150
Taiwan	111
Singapore	40
Total	1251

Roman, 2002.

Of all the countries shown in Table 5, Japan's nanotech investments are by far the greatest. Indeed, it is universally agreed that Japan has the only fully co-ordinated and funded national policy of nanotechnology research. The most prominent product of this

national policy has been the Ministry of Economy, Trade and Industry (METI) programme on atomic manipulation, 1991–2001, entitled *Research and Development of Ultimate Manipulation of Molecules* (Tam, 2001). The programme was funded at the ¥25 billion level (approximately US\$210 million). Of the total, US\$167 million has been allocated for the development of microbots (Saxl, 2000). Nowadays, the Japanese government views the successful development of nanotechnology as key to restoration of its economy: nanotechnology is one of the four strategic platforms of Japan's second basic plan for science and technology. For example, the Japanese government has founded the Expert Group on Nanotechnology under the Japan Federation of Economic Organisations Committee on Industrial Technology. In another initiative, which it calls its 'e-Japan strategy', the Japanese government aims to become '*the world's most advanced IT nation within five years*' (IT Strategic Headquarters, 2001). Japan's government nanotechnology expenditures are given in Table 6.

Table 6: Estimated Japanese government nanotechnology research and development expenditures (US\$million).

1997	1998	1999	2000	2001	2002	2003
120	135	157	245	465	~750	~1000

Roman, 2002.

Although the figures given in Table 6 are impressive, Roman (2002) believes that the annual 50% increase does cast some doubt over their accuracy. For while there is no doubt that funding will continue to increase, increasing the number of researchers available to absorb this extra funding does not seem possible on an annual basis.

2.2.5.3 European Union

All European Union (EU) member states, except Luxembourg where no universities are located, have research programmes. For some countries, such as Germany, Ireland or Sweden, where nanotechnology is considered of strategic importance, nanotechnology programmes have been established for several

years. On the other hand, many countries have no specifically focused nanotechnology initiatives, but this research is covered within more general R&D programmes (Compano, 2001). Table 7 summarises the situation for the top six countries.

Table 7: Top six European government nanotechnology spending from 1998–2000 (€million).

Country/institution	1998	1999	2000
Germany	49.0	58.0	63.0
UK	32.0	35.0	39.0
European Commission	26.0	27.0	29.0
France	12.0	18.0	19.0
Netherlands	4.7	6.2	6.9
Sweden	3.4	5.6	5.8
European total	139.8	164.7	184.0

Compano, 2001.

The European Commission (EC) funds nanoscience through its so-called Framework Programme (FP). The aim of the FP6 is to produce breakthrough technologies that directly benefit the EU, either economically or socially. Under this, €1.3 billion are earmarked for '*nanotechnologies and nanosciences, knowledge-based multifunctional materials and new production processes and devices*' in the 2002–2006 FP out of a total budget of €11.3 billion. This thematic priority is only partly dedicated to nanoscience, while other thematic priorities also have a nanotechnology component. At first glance this may seem a small figure compared to the 2003 NNI budget of US\$710 million (€0.72 billion). However, it does not take into account the substantial contributions made by individual member states (Compano, 2001). The UK serves as a good example of this, where public spending on nanotechnology R&D was around £30 million in 2001 (DTI, 2002), 70–80% of it from the Engineering and Physical Sciences Research Council (EPSRC). However, this is set to rise quite rapidly in 2002–2003 as the new interdisciplinary research collaborations and university technology centres start to spread.

2.3 Applications and Markets

2.3.1 Introduction

The applications of nanotechnology are extremely diverse, mainly because the field is interdisciplinary (Miles and Jarvis, 2001). In addition, the effect that nanotechnology will have during the next decade is difficult to estimate because of potentially new and unanticipated applications. For example, if simply reducing the microstructure in existing materials can make a big market impact, then this may, in turn, lead to a whole new set of applications. However, it seems reasonable to assume that during the next two to three years most activity in nanotechnology will still be in the area of research, rather than completed projects or products. Holister (2002) estimates that there are currently 455 public and private companies, 95 investors, and 271 academic institutions and government entities that are involved in the near-term applications of nanotechnology world-wide. The ability of such institutions to transfer research results into industrial applications can be indicated by the number of filed patents. Compano and Hullman (2001) provide an analysis of this, using the number of nanopatents filed at the European Patent Office (EPO) in Munich. Over the whole 1981–1998 period, the number of nanopatents rises from 28–180 patents, with an average growth rate in the 1990s amounting to 7%.

One important characteristic of activity grouped within this section is that much of the work in near-term applications of nanotechnologies is ‘market-pulled’: in each case, a particular and potentially profitable use within industry and/or the consumer market has been identified. However, as with the difficulty in predicting the future applications of nanotechnology, many market analysts believe that it is too soon to produce reliable figures for the global market – it is simply too early to say where and when markets and applications will come

(DTI, 2002). In spite of these difficulties, some forecasts exist that do hint at the kind of growth we might expect.

Table 8: Summary of future estimated global markets in nanotechnology.

Year	Estimated global market
2001	£31–55 billion
2005	£105 billion
2008	£500 billion
2010	£700 billion
2011–2015	Exceeds US\$1 trillion (£0.6 trillion)

DTI, 2002.

Most strikingly, the NSF predicts that the total market for nanotech products and services will reach US\$1 trillion by 2015 (Roco and Bainbridge, 2001). The accuracy of this claim is difficult to assess, given the doubts expressed above. Compano and Hullman (2001) approach the problem through the comparison of publication (representing basic science or R&D) and patent (representing technology applications) nanotechnology data with Grupp’s (1993) theory of Stylised Technological Development. As a result, they conclude that the peak of scientific activity is still to come, possibly in three to five years from now, and large-scale exploitation of nanotechnological results might arise ten years from now.

Considering the above comments about nanotechnological development and market-pull, it is instructive to examine which areas of industry will be affected first. Mihail Roco, the NSF senior advisor for nanotechnology, believes that *‘early payoffs will come in computing and pharmaceuticals’* (quoted in Leo, 2001), whereas Holister (2002) points out that medicine is a huge market, thereby implying that revenue for nanotechnology in this area could be substantial. On the other hand, the NSF believe that, due to the high initial costs involved, *‘nanotechnology-based goods and services will probably be introduced earlier in those markets where performance*

characteristics are especially important and price is a secondary consideration' (Roco and Bainbridge, 2001). Examples of these are medical applications and space exploration. The experience gained will then reduce technical and production uncertainties and prepare these technologies for deployment into the market place.

A good indication of the areas of current and near-future commercial nanotech activity is the type of patents made. Compano and Hullman (2001) state that one-quarter of all patents filed are focused on instrumentation. This supports the view that nanotechnology is at the beginning of the development phase of an enabling technology where the first focus is to develop suitable tools and fabrication techniques. The most important industrial sectors are informatics (information science), and pharmaceuticals and chemicals. For the first sector, *'massive storage devices, flat panel displays, or electronic paper are prominent IT patenting areas. In addition to this, extended semiconductor approaches and alternative nanoscale information processing, transmission or storage devices are dominant.'* In the case of chemistry and pharmaceuticals, a large number of patents are directed towards *'finding new approaches for drug delivery, medical diagnosis, and cancer treatments which are supposed to have huge future markets. Nanotechnology patenting for other sectors (e.g. aerospace, construction industries and food processing) show yearly increasing values, but their absolute numbers are relatively small.'* In summary then, IT and medicine look set to have an impact on the market first. The next two sections deal with both these areas in more detail. Following this, the widely cited potential impacts of nanotechnology on the energy and defence sectors are examined.

2.3.2 Informatics

Informatics, or information science, can be thought of as consisting of three interrelated

areas: electronics, magnetics and optics. This section primarily concentrates on electronics, acknowledged by Compano (2001) as one of the major drivers of the world-wide economy. In fact, the current market for miniaturised systems is estimated at US\$40 billion and the market for IT peripherals to be more than US\$20 billion, although semiconductor products have a dominant role and their turnover grows at a higher rate than the overall electronics market. The field is dominated by the US and Japan. In fact, apart from a few niche markets where Western European companies are able to compete, recent technological breakthroughs have been largely due to major manufacturers in these countries (Miles and Jarvis, 2001). Japan has a particularly strong commercial basis in this area, although Japanese R&D tends to be organised through lines determined by the government (via the MicroMachine Centre): the METI funds much of the work (US\$100 million in the last five years). In the US too, government is very involved in applied research. Here, the activities of military funding agencies are of note – such institutions tend to be generous in their company funding in this field, even when there is a clear commercial benefit for the companies involved.

In general, it is much harder to predict the commercially successful technologies in the world of electronics than in the world of materials (Holister, 2002). However, if one considers that the major driving force in nanoscience for the last decade has been microelectronics (Glinos, 1999), then it makes sense that nanotechnology will play an important role in the future of this industry. The ETC Group (2002a) provide a notable statistic here, stating that by 2012 the entire market will be dependent on nanotech. For, although there are few nanotechnology products in the market place at present, future growth is expected to be strong, with a predicted composite annual growth rate of 30–40%, with emerging

markets around 70% (DTI, 2002). A number of recent forecasts, although varying greatly, reflect this market confidence. For example, Miles and Jarvis (2001) put the market for nanotechnology-based IT and electronics devices at around US\$70 billion by 2010. A second estimate states that nanotechnology will yield an annual production of about US\$300 billion for the semiconductor industry and about the same amount again for global integrated circuits sales within 10–15 years (NSF, 2001). Similarly, for micro- and nanotechnology systems in the telecommunications sector, the market is presently estimated as being in the order of US\$35 billion with an anticipated compound annual growth rate of around 70%.

2.3.2.1 Moore's Law

The microelectronics industry had looked ahead and seen serious challenges for its basic CMOS process. CMOS technology has been refined for over 20 years, driving the 'line width'-the width of the smallest feature in an integrated circuit (IC)-from 10 mm down to 0.25 μm (Doering, 2001). This is

the force behind Moore's law, which predicts that the processing power of ICs will double every 18 months (Glinos, 1999). Based on Moore's law, industry predictions are summarised in Table 9.

Semiconductor industry associations assume that they will be close to introducing 100 nm ground-rule technology by 2004 (Compano, 2001). The significance of this lies in the fact that 100 nm is widely viewed as a kind of 'turning point', where many radically new technologies will have to be developed. To begin with, optical lithography will become obsolete somewhere around 100 nm. As a result, 'next generation lithography' options are currently being investigated. These are summarised in Table 10.

Table 9: Anticipated technological computing developments for 2001–2014.

Feature	Year					
	2001	2003	2005	2008	2011	2014
Memory						
Minimum feature size DRAM (1/2 pitch in nm)	150	120	100	70	50	35
Gbits/chip	2	4	8	24	68	194
Density (Gbits/cm ²)	0.49	0.89	1.63	4.03	9.94	24.50
Logic (processing power)						
Minimum feature size (gate length in nm)	100	80	65	45	30–32	20–22
Density (million transistors per cm ²)	13	24	44	109	269	664
Logic clock (GHz)	1.7	2.5	3.5	6.0	10.0	13.5

DRAM: Dynamic Random Access Memory, a type of memory used in most personal computers.

Adapted from Compano, 2001

Table 10: Maturity of lithography options.

Year of introduction	2001	2003	2006	2009
Minimum feature size	150	120	90	65
Optical 193 nm	X*	X		
Optical 157 nm		X	X	
Extreme UV			X	X
X-rays				X
Electron beam			X	X
Ion beam			X	X
Printing				X

*An 'X' designates the date at which the respective fabrication technology is expected to become economically viable for mass production.

Adapted from Compano, 2001.

Excluding the printing process, each fabrication technique essentially works on the same principle where a reactive silicon-based agent is exposed to increasingly focused electromagnetic radiation: optical to X-rays representing a successive reduction in photon wavelength; E-beam and ion beam projection technologies using focused electron and ion beams respectively. All of these techniques are currently under active evaluation-the aim is to have the appropriate equipment for the corresponding time-frame. To date, X-ray and ion beam projection have

received the greatest research investment (Compano, 2001). Printing technologies, however, are the ultimate goal, where sheets of circuits can be rolled off the production line like a printing press.

2.3.2.2 Beyond Moore's law

Moore's law cannot continue indefinitely. In the years following 2015, additional difficulties are likely to be encountered, some of which may pose serious challenges to traditional semiconductor manufacturing techniques. In particular, limits to the degree that interconnections or wires between transistors may be scaled could in turn limit the effective computation speed of devices because of the properties and compatibility of particular materials, despite incremental present-day advances in these areas (Anton et al., 2001). Thermal dissipation in chips with extremely high device-densities will also pose a serious challenge. This issue is not so much a fundamental limitation as it is an economic consideration, in that heat dissipation mechanisms and cooling technology may be required that add to the total system cost, thereby adversely affecting the marginal cost per computational function for these devices. Eventually, however, CMOS technology may hit a more crucial barrier, the quantum world, where the laws of physics operate in a very different paradigm to that experienced in everyday life. For example, futuristic circuits operating on a quantum scale would have to take Heisenberg's Uncertainty Principle into account. Overcoming this barrier is a different matter altogether, where the problems are no longer merely technological (Glinos, 1999), and industry has already begun to investigate the problem in a number of ways. Three of the most commonly cited approaches-molecular nanoelectronics and quantum information processing (QIP)-are expanded upon below. In addition, computational self-assembly is acknowledged as a potentially key fabrication technique of the future.

Molecular nanoelectronics. Organic molecules have been shown to have the necessary properties to be used in electronics. Devices made of molecular components would be much smaller than those made by existing silicon technologies and ultimately offer the smallest electronics theoretically possible without moving into the realm of subatomic particles (Holister, 2002). Molecular electronic devices could operate as logic switches through chemical means, using synthesised organic compounds. These devices can be assembled chemically in large numbers and organised to form a computer. The main advantage of this approach is significantly lower power consumption by individual devices. Several approaches for such devices have been devised, and experiments have shown evidence of switching behaviour for individual devices.

For example, in 'DNA computing', the similarities between mathematical operations and biological reactions are used to perform calculations. The key idea is to find the parallelism between DNA-the basic genetic information-and well-known digital computers. This is because a string of DNA can be used to solve combination problems if it can be put together in the right sequence (Compano, 2001). One issue is that molecular memories must be able to maintain their state, just as in a digital electronic computer. Also, given that the manufacturing and assembly process for these devices will lead to device defects, a defect-tolerant computer architecture needs to be developed. Fabricating reliable interconnections between devices using carbon nanotubes (or some other technology) is an additional challenge. A significant amount of work is ongoing in each of these areas. Even though experimental progress to date in this area has been substantial, it seems unlikely that molecular computers could be developed within the next 15 years that would be relatively attractive (from a price and performance standpoint) compared

with conventional electronic computers (Anton et al., 2001).

QIP. This crosses the disciplines of quantum physics, computer science, information theory and engineering with the aim of harnessing the fundamental laws of quantum physics to '*dramatically improve the acquisition, transmission and processing of information*' (Miles and Jarvis, 2001). QIP represents computing at the smallest possible scale, in which one atom is equivalent to one byte of information. Other aspects of quantum computing also considered attractive relate to their massive parallelism in computation (i.e. the ability to perform simultaneous calculations) (Holister, 2002). These concepts are qualitatively different from those employed in traditional computers and will hence require new computer architectures. A preliminary survey of work in this area by Anton et al., (2001) indicates that quantum switches are unlikely to overcome major technical obstacles, such as '*error correction, de-coherence and signal input/output*' within the next 15 years. If this proves to be the case, QIP-based computing, as with molecular nanoelectronics, does not appear to be competitive with traditional digital electronic computers for some time.

Computational self-assembly. A major barrier to the introduction of nanoelectronics is that there are no established mass production techniques for creating devices on a commercial basis (Saxl, 2000). In the short to medium term, Table 10 covers the most promising fabrication approaches. In the long-term, however, more ambitious bottom-up methods based on self-assembly techniques are proposed. Bottom-up approaches are elegant, cheap and possibly enormously powerful techniques for future mass replication. The relatively straightforward architecture of molecular memory means that self-assembly techniques in this area may bear fruit in a few years (Table 11). Tackling processors is another

matter, however, because of the greater complexity involved-their applicability remains limited until total control over the emerging structures in terms of wiring and their interconnections can be obtained. These are formidable obstacles. Self-assembly, therefore, will likely be combined initially with some more traditional top-down approaches. For example, many believe that inducing molecular components to self-assemble on a patterned substrate in some sort of hybrid system will represent the first commercialisation of nanoelectronics (Holister, 2002).

2.3.2.3 Summary of applications

The main drivers for current applied microelectronics research are computing, telecommunications, consumer electronics and military applications. It is not evident how long personal computing will act as a driver. On the one hand, personal computers (PCs) already offer sufficiently good performance for a large number of users; on the other, new applications, such as automatic voice recognition or PC wireless communications, may give further impulses for further progress (Compano, 2001). Military applications have a restricted volume but are of strategic importance. Thus, it is generally anticipated that most new technologies within this area will emerge in (US) military use first before eventually finding their way into the civil sector (Saxl, 2000). These and other ideas are summarised in Table 11.

2.3.3 Pharmaceuticals and medicine

Nanotechnology, combined with biotechnology, are the underpinning technologies pushing the rapid advances in '*genomics, combinatorial chemistry, high throughput robotic screening, drug discovery, gene sequencing and bioinformatics and their applications*' (Saxl, 2000). In medicine, advances can take place at the nanoscale where, for example, either passive or active nano-engineered systems can be used that

Table 11: Summary of application areas for informatics.

Material/technique	Applications	Time-scale (to market launch)
Pre 2015		
Quantum well structures (pers.comm., Gareth Barry, Imperial College London, 22 Nov 2002).	Telecommunications/optics industry. Potentially very important applications in laser development for the data communications sector. The aim is to use fibre optic communications in building and computers. The problems are cost and high temperature operating conditions. Quantum well/dot structures can potentially solve this problem	Quantum well lasers already utilised in CD players. Not yet optimised for the communications market (i.e. fibre optics): 4–5 years.
Quantum dot structures (source as above).		Quantum dots still in research stage: 7–8 years.
Photonic crystal technologies (Miles and Jarvis, 2001).	Optical communication sector, i.e. fibre optics. Photonic integrated circuits can be nearly a million times denser than electronic ones. Their tighter confinement and novel dispersion properties also open up opportunities for very low power devices	Still in basic R&D, but very strong commercial interest emerging.
Carbon nanotubes in nanoelectronics. These hold promise as basic components for nanoelectronics – they can act as conductors, semiconductors and insulators (Holister, 2002).	Memory and storage; commercial prototype nanotube-based (non-volatile); RAM; display technologies; E-paper.	Commercial prototype nanotube-based RAM predicted in 1–2 years. Consumer flat screen by the end of 2003. Limited commercialisation of E-paper in 1–2 years.
Spintronics – the utilisation of electron spin for significantly enhanced or fundamentally new device functionality (Science Blog, 2002).	Ultra-high capacity disk drives and computer memories.	A read head has been demonstrated that can deal with storage densities of a terabit per square inch. In 2001 Fuji announced a new magnetic coating promising 3-gigabyte floppy disk.
Polymers (Compano, 2001).	Display technologies – this sector is driven by the electronics consumer market.	Some commercialisation, e.g. Cambridge Display Technologies has been formed specifically to exploit this technology.
Post 2015		
Molecular nanoelectronics (including DNA computing) (Compano, 2001).	Circuits based on single molecule and single electron transistors will appear, initially in special applications.	Single atom transistor demonstrated recently. Still immature, but huge potential (Miles and Jarvis, 2001).
Quantum information processing (QIP) (Compano, 2001).	Several researchers have devised algorithms for problems that are very computationally intensive (and thus time-consuming) for existing digital computers, which could be made much faster using the physics of quantum computers. E.g. factoring large numbers (essential for cryptographic applications), searching large databases, pattern matching, simulation of molecular and quantum phenomena (Anton et al., 2001).	Still in pure research phase, although some US defence money has been made available (Holister, 2002).

enable the required dose of drug to be delivered at the correct time to the target area, or at the macro-level, such as induced tissue growth. This reduces unwanted side effects, improves patient compliance, leads to lower doses and opens up new possibilities (that would be impossible without nanotechnology approaches) (Miles and Jarvis, 2001). The size of this market is the main driving force behind such innovation. LaVan and Langer (2001) predict that: *'fundamental changes in drug production and delivery are expected to affect about half of the [US]\$380 billion world-wide drug production in the next decade.'* At present, nanotechnology is estimated to have a 1% stake in this, but whole sectors will continue to grow and this contribution is expected to increase rapidly (Ho, 2002a). The US is widely recognised as a leader in this area, with a company market share of about 40%, and many applications close to the market.

2.3.3.1 Drug delivery

The most promising aspect of pharmaceuticals and medicine as it relates to nanotechnology is currently drug delivery. In the words of LaVan and Langer (2001): *'It is likely that the pharmaceutical industry will transition from a paradigm of drug discovery by screening compounds to the purposeful engineering of targeted molecules.'* At present, there are 30 main drug delivery products on the market. The total annual income for all of these is approximately US\$33 billion with an annual growth of 15% (based on global product revenue) (Miles and Jarvis, 2001). Two major drivers are primarily responsible for this increase in the market. First, present advances in diagnostic technology appear to be outpacing advances in new therapeutic agents. Highly detailed information from a patient is becoming available, thus promoting much more specific use of pharmaceuticals (LaVan and Langer, 2001). Second, the acceptance of new drug formulations is expensive and slow, taking up to 15 years to obtain accreditation of new drug formulas

with no guarantee of success. In response, some companies are trying to hurry the long clinical phase required in Western medicine (Ho, 2002a). However, powerful incentives remain to investigate new techniques that can more effectively deliver or target existing drugs (Saxl, 2000). In addition, many of these new tools will have foundation in current techniques: a targeted molecule may simply add spatial or temporal resolution to an existing assay. Thus, although many potential applications are envisaged, the actual near-future products are not much more than better research tools or aids to diagnosis (Ho, 2002a). These are summarised in Table 12.

2.3.4 Energy

The global energy sector is likely to be particularly affected by coming advances in nanotechnology. To illustrate, significant changes in lighting technologies are expected in the next 10–15 years. Semiconductors used in the preparation of light-emitting diodes can increasingly be sculpted on nanoscale dimensions. Projections indicate that such nanotechnology-based advances have the potential to reduce world-wide consumption of energy by more than 10% (NSF, 2001). The various applications showing most promise are summarised in Table 13.

Most current photovoltaic (PV) production is based upon crystalline and amorphous silicon technologies. However, as Table 13 shows, research is now focusing upon new technologies which may result in significant reductions in PV costs and/or improvements in efficiency. Nanotechnology is anticipated to play an important part in this future. Although total PV power output remains relatively low, the industry is growing rapidly-the production of PV modules expanded by 40% in 1997 (Saxl, 2000). This increase is largely due to the building and construction industry, the largest and fastest growing sector at present. In addition, developing countries represent a potentially vast market (pers. comm., Jenny Nelson,

Table 12: Summary of application areas for nanoscale pharmaceuticals and medicine.

Material/technique	Property	Applications	Time-scale (to market launch)
Diagnostics			
Nanosized markers i.e. the attachment of nanoparticles to molecules of interest (Holister, 2002).	Minute quantities of a substance can be detected, down to individual molecules	E.g. detection of cancer cells to allow early treatment.	?
'Lab-on-a-chip' technologies (Saxl, 2000).	Miniaturisation and speeding up of the analytical process.	The creation of miniature, portable diagnostic laboratories for uses in the food, pharmaceutical and chemical industries; in disease prevention and control; and in environmental monitoring.	Although chips currently cost over £1250 (US\$2085) each to make, within three years the costs should fall dramatically, making these tools widely available.
Quantum dots (pers. comm., Gareth Barry, Imperial College London, 22 Nov 2002).	Quantum dots can be tracked very precisely when molecules are 'bar coded' by their unique light spectrum.	Diagnosis	In early stage of development, but there is enough interest here for some commercialisation (e.g. Q-dot Inc.).
Drug delivery			
Nanoparticles in the range of 50–100 nm (Miles and Jarvis, 2001).	Larger particles cannot enter tumour pores while nanoparticles can easily move into a tumour.	Cancer treatment.	?
Nanosizing in the range of 100–200 nm (Miles and Jarvis, 2001).	Low solubility.	More effective treatment with existing drugs.	?
Polymers (Holister, 2002).	These molecules can be engineered to a high degree of accuracy.	Nanobiological drug carrying devices.	?
Ligands on a nanoparticle surface (Holister, 2002).	These molecules can be engineered to a high degree of accuracy.	The ligand target receptors can recognise damaged tissue, attach to it and release a therapeutic drug.	?
Nanocapsules (Holister, 2002).	Evading body's immune system whilst directing a therapeutic agent to the desired site.	A Buckyball-based AIDS treatment is just about to enter clinical trials (Ho, 2002a).	Early clinical.
Increased particle adhesion (Holister, 2002).	Degree of localised drug retention increased.	Slow drug release.	?
Nanoporous materials (Holister, 2002).	Evading body's immune system whilst directing a therapeutic agent to the desired site.	When coupled to sensors, drug-delivering implants could be developed.	Pre-clinical: an insulin-delivery system is being tested in mice.
'Pharmacy-on-a-chip' (Saxl, 2000).	Monitor conditions and act as an artificial means of regulating and maintaining the body's own hormonal balance.	E.g. Diabetes treatment.	More distant than 'lab-on-a-chip' technologies.
Sorting biomolecules (Holister, 2002).	Nanopores and membranes are capable of sorting, for example, left- and right-handed versions of molecules.	Gene analysis and sequencing.	Current – ?
Tissue regeneration, growth and repair			
Nanoengineered prosthetics (Miles and Jarvis, 2001).	Increased miniaturisation; increased prosthetic strength and weight reduction; improved biocompatibility.	Retinal, auditory, spinal and cranial implants.	Most immediate will be external tissue grafts; dental and bone replacements; internal tissue implants (Miles and Jarvis, 2001).
Cellular manipulation (Miles and Jarvis, 2001).	Manipulation and coercion of cellular systems.	Persuasion of lost nerve tissue to grow; growth of body parts.	More distant: 5–7 years.

Table 13: Summary of applications for energy processing.

Material/technique	Applications	Time-scale (to market launch)
Power generation (PV technology)		
Polymer materials (organic).	Solar cells (pers. comm., Jenny Nelson, Imperial College London, 2 Dec 2002). Current developments aim to balance moderate efficiency with low cost. Another big advantage is that these layers can easily be incorporated into appliances. Current problems stem from the material's instability.	The research stage has advanced much more quickly than expected. As a result, polymer-based PV cells should enter the market in 5 years.
Combinations of organic and inorganic molecules.	Dye-sensitised solar cells made from a thin hybrid layer (pers. comm., Jenny Nelson, Imperial College London, 2 Dec 2002). These cells are potentially very cheap because fabrication is from cheap, low purity materials by simple and low cost procedures (Saxl, 2000). Photocatalytic water treatment.	Low power applications will enter market first. Limited commercialisation already occurring (e.g. by Sustainable Technologies International).
Quantum wells (inorganic).	Quantum-well solar cells (pers. comm., Jenny Nelson, Imperial College London, 2 Dec 2002). Current research is taking place in high-efficiency applications because the infrared part of the solar spectrum may be absorbed.	Pure research.
Nanorods.	These structures can be tuned to respond to different wavelengths of light forming cheap and efficient solar cells (Holister, 2002).	Long-term.
Fuel conversion/storage		
Improved fuel catalysts through nanostructuring.	Fuel conversion (Saxl, 2000).	Current – 3 years.
Nanotubes.	Fuel storage. E.g. a methane-based fuel cell for powering mobile phones and laptops is currently being developed. (Holister, 2002).	2 years.
Nanoparticles.	Vastly increased (e.g. x10) charge and discharge battery rate (Holister, 2002).	Distant.

Imperial College London, 2 Dec 2002).

In spite of these developments, however, nanotechnology, as a new and radical technology, still faces an uncertain future in this area as a number of alternative technologies are also competing for attention (e.g. inorganic silicon). Indeed, it may be 20 years before nanotech-based PV begins competing as a viable energy source with this example. As reminiscent of so many of the aforementioned applications in this section, there is much hype but no one really knows how to achieve these things yet (pers. comm., Jenny Nelson, Imperial College London, 2 Dec 2002).

2.3.5 Defence

Nanoscale informatics, pharmaceuticals and medicine remain the most high-profile areas of near-term market application. However, Gsponer (2002) contends that the most significant near-term applications of nanotechnology will be in the military domain. This is because micromechanical and MEMS engineering is historically connected to nuclear weapons laboratories: it was within this domain that the field of nanotechnology was born a few decades ago. Today, it is not difficult to understand why nanotechnology might appeal to military planners. Through technologies such as

steam navigation, repeating firearms, and high explosives, Western powers have enjoyed virtually unchallenged military supremacy throughout the 19th Century (Reynolds, 2002). It is not absurd, then, to imagine that nanotechnology could play a similar role in the 21st Century. Indeed, new technologies, notably IT, are playing an increasingly important part in modern warfare-as reflected by recent investments in the US DoD (see Table 4). Trends such as these have led leading strategic commentators, such as David Jeremiah (1995), to conclude that military applications of nanotechnology have an even greater potential than nuclear weapons to radically change the balance of global power in the future. Fundamentally, this potential lies in a greater range of military options when deciding how to respond to aggression. Scott Pace (1989) of RAND expands upon this:

'How might nanotechnology contribute to US military power? In peacetime or crisis, nanocomputers may allow more capable surveillance of potential aggressors. The flood of data from world-wide sensors could be culled more efficiently to look for truly threatening activities. In low-intensity warfare, intelligent sensors and barrier systems could isolate or channel guerrilla movements depending on the local terrain. In conventional theatre war, nanotechnology may lead to small, cheap, highly lethal anti-tank weapons. Such weapons could allow relatively small numbers of infantry to defeat assaults by large armoured forces. At nuclear conflict levels, accurate nanocomputer guidance and low nanomachine production costs would accelerate current trends in the proliferation of 'smart' munitions. Rather than requiring nuclear weapons to attack massive conventional forces or distant, hard targets, nanotechnology enhancements to cruise missiles and ballistic missiles could allow them to destroy their targets with conventional explosives. Conventional explosives themselves might be replaced by

molecular disassemblers that would be rapidly effective, but with less unintended destruction to surrounding buildings and populations.'

Other stated applications include (NSF, 2001):

- information dominance through advanced nanoelectronics
- more sophisticated virtual reality systems
- increased use of enhanced automation and robotics
- required improvements in chemical/biological/nuclear sensing
- design improvements in systems used for nuclear non-proliferation monitoring and management
- combined nanomechanical and micromechanical devices for control of nuclear defence systems.

In addition, such nanotechnologies might be 'cleaner' and 'safer' and less likely to cause collateral damage than the technologies they replace, making them especially appealing to military planners (Reynolds, 2002). For example, MEMS have many potential uses in the battlefield, largely due to their built-in mechanical functions that allow them to act as sensors and actuators (RAND, 2002). Actuators in particular extend the functionality of sensors by allowing them to respond to the environment with the usage of force. Applications of MEMS in military systems include ammunition, petroleum, food, as well as enabling a host of other smarter, more efficient logistics operations.

The infantry soldier too is anticipated to receive a nanotech-based 'makeover': a new Institute for Soldier Nanotechnology (ISN) has been created at MIT, with a US Army grant of US\$50 million over five years. The

goal of this research centre is to greatly enhance the protection and survival of the infantry soldier using nanoscience (New Scientist, 2002). For example, US army planners are hoping to lighten the load that soldiers carry into battle (currently around 64 kg) by redesigning the equipment from the atomic scale up. Current signs indicate that progress towards these objectives may soon begin to bear fruit: a Centre for Nanoscience Innovation for Defence (CNID) was created in January 2003 to facilitate the rapid transition of research innovation in the nanosciences into applications for the defence sector (Science Blog, 2002). CNID is sponsored by two federal agencies – the Defence Advanced Research Project Agency (DARPA) and Defence MicroElectronics Activity (DMEA) – to the tune of US\$20 million over three years.

2.3.6 Corporate funding

The difficulties involved in drawing upon accurate corporate data from within the public domain are far more substantial than those encountered with regard to public investment. Thus, a detailed analysis of corporate activity is mainly beyond the scope of this report. However, it is important to recognise that, urged on by the growing interest (and hype) currently surrounding nanotechnology, spending by big firms in 2002 is anticipated to match or even exceed government spending (Holister, 2002). Furthermore, this private investment is very often at the forefront of application development in the marketplace. Helsel (2002) demonstrates this by showing how historical funding for technology transitioning into the US market place is led by corporate sources.

In total, there are an estimated 470 nanotech companies distributed across North America, Europe and Asia (ETC Group, 2002a). Of these, about 230 are based in the US, about 130 in Europe, and about 75 in the Asia-Pacific.

Table 14: US historical funding for technology transitioning into the marketplace.

Source	Percentage
Corporate	34%
Federal government	29%
Angels*	25%
State and local government	5%
Venture capital institutions	4%
University endowments	3%

* Angels are individuals who provide capital to one or more start-up companies. An angel is usually affluent or has a personal stake in the success of the venture. Such investments are characterised by high levels of risk and a potentially large return on investment.

Adapted from Helsel, 2002.

2.3.6.1 Transnational companies.

Transnational companies often carry out their own nanotech-related R&D. This is because they understand that nanotechnology is likely to disrupt their current products and processes, and therefore recognise the need to understand and control the pace of such implications (DTI, 2002). In this way, some of the world's largest companies, including IBM, Motorola, Hewlett Packard, Lucent, Hitachi, Mitsubishi, NEC, Corning, Dow Chemical and 3M have launched significant nanotech initiatives through their own venture capital funds or as a direct result of their own R&D (Holister, 2002). In the US and Switzerland for example, IBM is providing some US\$100 million nanotech-related funding for its hi-tech research laboratories (DTI, 2002). In Japan too, many of the nation's largest players have now entered the nanotech field, including Fuji, Hewlett-Packard Japan, Hitachi, Mitsubishi, NEC and Sony. For example, Toray Industries, a global maker of synthetic fibre, textiles and chemicals, is establishing a US\$40 million centre specialising in nanotechnology and biotechnology near Tokyo. The building is expected to be finished by March 2003 (Fried, 2002).

2.3.6.2 Start-up companies

At present there are about 100 business start-ups – new business ventures in their earliest stages of development – in operation today, about half of which are located in the US (Thibodeau, 2002). Such companies rely on their understanding of where new opportunities and markets may lie and thus play an important role in commercialising research. This increase in activity amongst start-ups is mirrored by the investment community, who, according to Abid Khan (pers. comm., London Centre for Nanotechnology, 6 Nov 2002), have decided that nanotechnology is the ‘next big thing’ – the new computing or biotechnology. Indeed, some large investment groups now have specialists who follow developments in the subject. Although such activity tends to produce little in the way of a coherent picture, business investment in nanotechnology start-ups is on the rise. There were over 20 nanotechnology investments in the first half of 2002 in the US and Europe, and more than US\$100 million invested in the US in the first half of 2002 (Holister, 2002). According to Thibodeau (2002), this level of funding is projected to increase to US\$1 billion by 2003.

2.4 Reality and Hype

2.4.1 Introduction

Nanotechnology advocates have been criticised within recent years for hyping the potential impact that nanoscale science and technology will have upon the economy and society. For example, in response to the NSF claim that the size of the nanotechnology market will reach US\$1 trillion in 10 years time (Roco and Bainbridge, 2001), The Economist (2002) points to ‘*nano-enthusiasts*’ being responsible for ‘*recklessly setting impossibly high expectations for the economic benefits.*’ This sentiment is even echoed by some material scientists themselves: Roy (2002) describes the term ‘nano’ as a ‘halo regime’ – a term that is sold

to budget managers in order to increase funding. He concludes that: ‘*the [term] should be new, different, euphonious, and connected somehow, however tenuously, to science.*’ It is not difficult to identify the kind of claims that can fuel such reaction. For example, Pergamit and Peterson (1993) state that: ‘*Humanity will be faced with a powerful, accelerated social revolution as a result of nanotechnology. In the near future, a team of scientists will succeed in constructing the first nano-sized robot capable of self-replication. Within a few short years, and five billion trillion nano-robots later, virtually all present industrial processes will be obsolete as well as our contemporary concept of labor.*’

Regardless of the accuracy of these claims, however, there can be no doubt that the language in which they are framed has helped to attract large amounts of investment. The pinnacle of this came in 1997 when the US NNI was launched by President Bill Clinton to ‘*an extraordinary amount of hype*’ (ETC Group, 2002a). Amongst the various documents produced by the White House about the subject was one entitled: *National Nanotechnology Initiative: Leading to the Next Industrial Revolution* (White House Fact Sheet, 2000). The fact sheet lists seven ‘*potential breakthroughs*’ anticipated over the next quarter-century. These include ‘*making materials and products from the bottom-up*’ (i.e. by building them up from atoms and molecules) and ‘*improving the computer speed and efficiency of minuscule transistors and memory chips by factors of millions.*’ However, these ambitious claims were accompanied with very little serious investigation of their feasibility, or indeed whether nanotechnology – rather than some other competing technology – will even deliver within the allotted time-frame.

At present, there is a general understanding amongst industry that the level of hype

surrounding nanotechnology has, to some extent, damaged investment potential (DTI, 2002). For example, Schultz (2002) advocates the need for nanotech researchers and supporters to dampen unquestioning enthusiasm for nanotechnology. This is because, without discussion of the potential pitfalls, future nanotechnology research could be subjected to such extreme pressure that funding is jeopardised and research progress is slowed, perhaps halted altogether in some cases. This realisation has quickly led to an attempt by industry to diffuse some of the wilder claims surrounding the field. Glenn Fishbine's *Investor's Guide to Nanotech and Micromachines* (2002) provides one such example: it is through this type of work that the science fiction aspects of the debate are now receding (pers. comm., Abid Khan, London Centre for Nanotechnology, 6 Nov 2002). However, in spite of these developments, it is clear that the distinction between near-future applications of nanotechnology (see Section 2.3) and some of the more visionary aspects of the debate has become blurred. An attempt to distinguish between the two areas here will help draw out some of the more legitimate concerns currently being voiced about nanotechnology in the following section.

2.4.2 Molecular nanotechnology

The more hyped aspects of nanotechnology have generally revolved around MNT. Proponents of this approach suggest that environmentally clean, inexpensive, and efficient manufacturing of structures, devices, and 'smart' products, based on the flexible control of architectures and processes at an atomic or molecular scale of precision, may be feasible in the near future (i.e. 10–20 years from the present). The ambitious goal is to produce complex products on demand using simple raw materials, such as by inserting the basic chemical elements in a molecular assembly factory to yield a common household appliance (Nelson and Shipbaugh, 1995). These visions have

attracted a great deal of public interest, and impressive demonstrations have been made of microscopic devices. For example, in August 2001 scientists from Osaka University built the smallest micromechanical system ever, a spring whose arm is only 0.3 μm wide (Ho, 2002a). However, although almost qualifying as a nanodevice, the question of whether it is possible to attain extreme capability and, if so, how to develop the field, is a point of contention in both scientific and policy circles (Nelson and Shipbaugh, 1995).

In spite of the above controversies, it remains clear that bottom-up technologies, while having the potential to be immensely important in the longer term, are not likely in the near future (DTI, 2002). However, some products benefiting from research into molecular manufacturing may be developed in the near term. As initial nanomachining, novel chemistry and protein engineering (or other biotechnologies) are refined, initial products will likely focus on those that substitute for existing high-cost, lower-efficiency products. Likely candidates for these technologies include a wide variety of sensor applications, tailored biomedical products (including diagnostics and therapeutics), extremely capable computing and storage products, and unique, tailored materials (i.e. smart materials using nanoscale sensors, actuators, and perhaps controller elements) for aerospace or similar high-capability needs (Nelson and Shipbaugh, 1995). Predictions of when bottom-up processes will begin to become available on a widespread basis vary across the literature. In general, the hyped aspects of the industry are operating around a 20-year time-scale, with estimates for economically viable self-assembly techniques tending to convene around 2015 (Ho, 2002a). However, to reach a fully mature nanotechnology society – where it is possible to manipulate objects on all scales from atom to macroscopic – is expected to take at least

35 years (Nelson and Shipbaugh, 1995). This is partly due to the economic advantages of competing technologies. For example, with regard to advanced computing, Anton et al., (2001) state that: *'the odds-on favourite for the next 15 years remains traditional digital electronic computers based on semiconductor technology. Given the virtual certainty of continued progress in this area, it is hard to imagine a scenario in which...quantum-switch-based computing, molecular computers, or something else could offer a significant performance advantage at a competitive price.'* The major technical obstacles to development in other areas of MNT – namely molecular manufacturing, general assembly and nanobots – are expanded upon below.

2.4.2.1 Molecular manufacturing

To realise molecular manufacturing, a number of technical accomplishments are necessary (Nelson and Shipbaugh, 1995). First, suitable molecular building blocks must be found. These building blocks must be physically durable, chemically stable, easily manipulated, and (to a certain extent) functionally versatile. The second major area for development is in the ability to assemble complex structures based on a particular design. A number of researchers have been working on different approaches to this issue. One uses atomic-force or molecular microscopes with very small nanoprobe to move atoms or molecules around with the aid of physical or chemical forces. An alternative approach uses lasers to place molecules in a desired location. Chemical assembly techniques are also being addressed, including an approach to building structures one molecular layer at a time. A third major area for development within molecular manufacturing is systems design and engineering. Extremely complex molecular systems at the macroscale will require substantial subsystem design, overall system design, and systems integration, much like complex manufactured systems of the present day. Although the design issues are

likely to be largely separable at a subsystems level, the amount of computation required for design and validation is likely to be quite substantial. Performing checks on engineering constraints, such as defect tolerance, physical integrity, and chemical stability, will be required as well (Anton et al., 2001).

2.4.2.2 Nanobots and other nanoscale devices

This area can be accredited with receiving the most severe hype, where headline-grabbing predictions include curing cancer, eliminating infections, enhancing our intelligence, and even making us immortal. In fact, according to Saxl (2000), it will take 25 years at least before tiny machines circulate in the bloodstream cleaning out fat deposits from our arteries. Indeed, although the implications of such revolutionary technologies are awesome, developments that appear achievable in the short and medium term are not particularly dramatic. Perhaps the most advanced work in this area concerns MIT's Bioinstrumentation Laboratory where an autonomous miniature robot, dubbed the 'NanoWalker', is being designed (MIT, 2002). Measuring approximately 25 mm², the name NanoWalker stems from its ability to take thousands of steps per second in the nanometre range. The ultimate goal of this type of robotic machine, generically referred to as an assembler, is the construction of materials an atom or molecule at a time by precisely placing reactive groups. This is called 'positional assembly' (Holister, 2002).

2.4.2.3 General assembly

The General Assembler is considered to be the 'Holy Grail' of nanotechnology and represents the ultimate utility of atom-manipulating nanobots. In general, such an assembling device is regarded as extremely distant (e.g. more than 25 years). However, there are presently two US companies known to be going after molecular assembly, in addition to engineering several 'magical' assembler dependent products. One of these

companies, Zyvex (2002), aims *'to become the leading world-wide supplier of tools, products, and services that enable adaptable, affordable, and molecularly precise manufacturing' and offers a 'variety of products, services, and licensing opportunities,'* including a number of nanomanipulation devices. Such nano-advocates claim the first major breakthrough in this area might occur as early as 2007.

2.4.3 Fundamental barriers to these visions

This report does not intend to refute that significant progress has been made in constructing macroscale objects using MNT techniques. Although the building blocks for these systems currently exist only in isolation at the research stage, it is certainly reasonable to expect that an integrated capability could be developed over the next 15 years. Such a system might be able to assemble structures with between 100 and 10,000 components and total dimensions of perhaps tens of microns (Anton et al., 2001). In particular, a series of important breakthroughs would certainly cause progress in this area to develop much more rapidly, especially if research continues to accelerate at today's rate. However, particularly in light of some of the wilder claims regarding nanotechnology-enabled futures, it must also be stressed that, although molecular manufacturing holds significant promise, it remains the least concrete of all the technologies discussed in this report. Certainly, there are a number of major technical obstacles to be overcome, some of which might be virtually insurmountable. Indeed, in the most carefully considered dismissal to date, Professor Smalley upholds the notion of nanobot replicators as fundamentally problematic (Smalley, 2001). First, the fingers of such atomically sized manipulators are too 'fat' to allow sufficient control of the reaction chemistry; second, they are too 'sticky' – the atoms of the manipulator hands would be adhered to the atom that is being moved.

Furthermore, other commentators such as Ho (2002c), point to major problems concerning energy sources and dissipation, or just the sheer complexity of the task at hand. For example, diamond assemblies might be relatively easy to assemble; other structures, such as biological configurations, are infinitely more complicated.

2.5 Concerns

2.5.1 Introduction

Given the difficulty in foreseeing nanotechnology outcomes and estimating likelihood, it is difficult to extrapolate predictions of specific threats and risks from current trends (Anton et al., 2001). And yet, in spite of this, recent discussions of the possible dangers posed by future technologies (such as AI, genetic engineering and MNT) have made it clear that analysis of the major classes of risks of nanotechnology is warranted. Perhaps the greatest difficulty in predicting the impacts of new technologies has to do with the fact that, once the technical and commercial feasibility of an innovation is demonstrated, subsequent developments may be as much in the hands of users as in those of the innovators (NSF, 2001). As a result, new technologies can affect society in ways that were not intended by those who initiated them. Sometimes these unintended consequences are beneficial, such as spin-offs with valuable applications in fields remote from the original innovation. A good example of this concerns the early days of the Internet – the subject is covered in Part 2 of this report. Other times, intended benefits may also have unintended or 'second-order' consequences. Interestingly, while a few far-sighted scientists are focusing on potentially negative second-order impacts of future nanotech applications, virtually no one has been tracking the potentially negative impacts of nanotechnology's present-day products (ETC Group, 2002a). This section, therefore, will attempt to distinguish between these two time-frames, as well as introducing

the main environmental and socio-political concerns. For the purposes of this report, 'long-term' refers to a hazard that, due to challenges associated with technological development, is unlikely to manifest itself within a 10–15 year time-frame.

2.5.2 Environmental concerns

The potential impact of nanostructured particles and devices on the environment is perhaps the most high profile of contemporary concerns. Quantum dots, nanoparticles, and other throwaway nanodevices may constitute whole new classes of non-biodegradable pollutants that scientists have very little understanding of. Essentially, most nanoparticles produced today are mini-versions of particles that have been produced for a long time. Thus, the larger (micro) versions have undergone testing, while their smaller (nano) counterparts have not (ETC Group, 2002a). For example, Vicki Colvin, Executive Director of Rice University's Centre for Biological and Environmental Nanotechnology (CBEN) has recently postulated that nanomaterials provide a large and active surface for adsorbing smaller contaminants, such as cadmium and organics. Thus, like naturally occurring colloids, they could provide an avenue for rapid and long-range transport of waste in underground water (cited in Colvin, 2002).

2.5.2.1 Infiltrating humans

The concern that nanomaterials could bind to certain common but harmful substances in the environment, such as pesticides or PCBs, leads to the short-term worry of such materials infiltrating humans. According to the ETC Group (2002a), at a recent fact-finding meeting at the US Environmental Protection Agency (EPA), researchers reported that nanoparticles can penetrate living cells and accumulate in animal organs. In particular, the possibility of toxic elements attaching themselves to otherwise benign nanomaterials inside bacteria and

finding a way into the bloodstream was acknowledged. In addition, very little work has been done in order to ascertain the possible effects of nanomaterials on living systems. One possibility is that proteins in the bloodstream will attach to the surface of nanoparticles, thus changing their shape and function, and triggering dangerous unintended consequences, such as blood clotting. A second possibility relates to the ability of nanoparticles to slip past the human immune system unnoticed, a property desirable for drug delivery, but worrying if potentially harmful substances can attach to otherwise benign nanomaterials and reside in the body in a similar manner. According to Colvin (2002), *'it is possible to speculate that nanoscale inorganic matter is generally biologically inert. However, without hard data that specifically address the issue of synthetic nanomaterials, it is impossible to know what physiological effects will occur, and, more critically, what exposure levels to recommend.'* To illustrate, this report shows how nanotubes, should industry predictions be realised, are set to become relatively ubiquitous within the coming decades – such materials are already finding their way into a number of products. But it has not yet been determined what happens if, for example, large quantities of nanotubes are absorbed by the human body. One prominent concern relates to the structural similarities between nanotubes and asbestos fibres: like the latter, nanotubes fibres are long, extremely durable, and have the potential to reside in the lungs for lengthy periods of time. One recent study, conducted by the National Aeronautics and Space Administration (NASA), has shown that breathing in large quantities of nanotubes can cause damage to lungs. However, as nanotubes are essentially similar to soot, then this is not particularly surprising (The Economist, 2002). On the whole, far more experiments are required before the issue can be resolved.

2.5.2.2 Self-replication

Self-replication is probably the earliest-recognised and best-known long-term danger of MNT. This centres upon the idea that self-replicating nanorobots capable of functioning autonomously in the natural environment could quickly convert that natural environment (i.e. ‘biomass’) into replicas of themselves (i.e. ‘nanomass’) on a global basis. Such a scenario is usually referred to as the ‘grey goo’ problem but perhaps more properly termed ‘global ecophagy’ (Freitas, 2000). The main feature that distinguishes runaway replication as a long-term environmental concern is the extreme difficulty involved in constructing machines with the adaptability of living organisms. As Freitas (2000) notes:

‘The replicators easiest to build will be inflexible machines, like automobiles or industrial robots...To build a runaway replicator that could operate in the wild would be like building a car that could go off-road and fuel itself from tree sap. With enough work, this should be possible, but it will hardly happen by accident. Without replication, accidents would be like those of industry today: locally harmful, but not catastrophic to the biosphere. Catastrophic problems seem more likely to arise though deliberate misuse, such as the use of nanotechnology for military aggression’ (see below).

This is not to imply, however, that the risk that molecular machines designed for economic purposes might replicate unchecked and destroy the world should be written off altogether: while the danger seems slight, even a slight risk of such a catastrophe is best avoided (Zyvex, 2002). To this end, David Forrest (1989) has produced a set of guidelines to assure that molecular machines and their products are developed in a safe and responsible manner.

2.5.3 Socio-political concerns

Clearly, if scientists are successful in

developing nanofabrication techniques for manufacturing nanoelectronic devices in huge volumes at very low cost, then the impact on society will be enormous. The potentially disruptive nature of nanotechnology has already been highlighted in earlier sections through its ability to generate major new paradigm shifts in how things are generated, such as a shift from top-down to bottom-up manufacturing techniques. This section further elaborates upon this and similar concerns.

2.5.3.1 Medical ethics

The ethical questions that have been raised in recent years following the advancement of such technologies as gene therapy are similar to in scope and philosophy to nanotechnology. For example, the emergence of highly specific drug therapies, a nanobased technique that features prominently in earlier sections of this report, may result in genetic discrimination. That is, discrimination directed against an individual or family based solely on an apparent or perceived genetic variation from the ‘normal’ human genotype (LaVan and Langer, 2001). The major concern here lies in the end result of going down such a road: that the de-selection of characteristics judged unwanted by societies (referred to as negative eugenics) will be viewed as the right, responsible, moral thing to do, as will cures and enhancements (Wolbring, 2002). Similarly, on a longer time-scale, concerns over nanotech applications for enhancing the performance of the human body might also arise. A major question here is whether such enhancements can be forced upon people, either when in a position to make a decision for themselves or, more controversially, against their will.

2.5.3.2 The nano-divide

If Moore’s law holds and the miniaturisation of PCs continues unchecked well into the 21st Century, then it seems likely that, in the long-term, society will get to a point where people can carry computers 24 hours a day. As Chaudhari (2001) states: *‘We are evolving*

to the point where every human being will be connected to any other human or to the vast network of information sources throughout the world by a communication system comprised of wireless and optical fibre communication links.' A world in which information is abundant and cheap may well have serious privacy implications for those who can afford to connect. However, little consideration seems to have been given to those who will clearly not be able to afford to participate. Indeed, many nations are already witnessing an IT divide, particularly in reference to Internet usage, that correlates with inequality in the distribution of wealth. This gap is likely to be exacerbated by any impending nanotechnological revolution, forming a so-called 'nano-divide.' It is important not to underestimate the potential scale of this: the transition from a pre-nano to post-nano world could be very traumatic and could exacerbate the problem of haves vs. have-nots. Such differences are likely to be striking (Smith, 2001).

A quick glance at demographics provides some insight into what such a post-nano world might look like. According to the World Bank, the Western industrial democracies will shrink from 12.7% of today's population to 8.6% by 2025. At the same time in the developing world the population will double (cited in Jeremiah, 1995). The kinds of nanotech-inspired wonders alluded to throughout this report may only be feasible for the 8.6% of the 2025 population who live in Western industrial democracies, and the upper layer of society in the developing and non-developing world, not for the rural poor and the underside of all urban populations. In other words, *'the differences in the quality of life will be even starker than today between these two worlds'* (Jeremiah, 1995). The NSF supports these sentiments: *'Those who participate in the nano revolution stand to become very wealthy. Those who do not may find it increasingly difficult to afford the*

technological wonders that it engenders.' (Roco and Bainbridge, 2001). One near-term example will be in medical care, as nanotech-based treatments may be initially expensive and hence only accessible to the very rich.

In the longer-term, campaign groups such as the ETC Group point to what they describe as the '*corporate concentration*' of '*material building blocks and processes that make everything from dams to DNA.*' This concern arises irrespective of the general doctrine in patent law that products of nature cannot be patented because the atomically-engineered elements of today are able to side-step the issue. For example, C Sixty Inc., a Toronto, Canada-based start-up exercise, has filed a series of patents, five of which have been granted, for Buckminsterfullerene. The aim of C Sixty Inc. is to corner the market with respect to this remarkable molecule and its vast potential in drug delivery. A big concern of the ETC Group (2002c) is that patenting offices (such as the US Patent and Trademark Office) understand nanotechnology, so that when approached by industry, examiners understand what are reasonable boundaries to intellectual property rights.

2.5.3.3 Destructive uses

The potentially catastrophic but long-term danger that the deliberate misuse of nanotechnology for military aggression poses has already been sketched out above. Indeed, Howard (2002) concedes that *'once the basic technology is available, it would not be difficult to adapt it as an instrument of war or terror.'* Gsponer (2002), on the other hand, draws attention to the existing potential of nanotechnology to affect dangerous and destabilising 'refinements' of existing nuclear weapons designs – such fourth generation nuclear weapons are new types of explosives that can be developed in full compliance with the Comprehensive Test Ban Treaty (CTBT). Such developments hint at the worrying possibility of a nanotechnology arms race. Zyvex (2002)

sketch out the underlying rationale for such an occurrence:

'It is clear that offensive weapons made using advanced nanotechnology can only be stopped by defensive systems made using advanced nanotechnology as well. If one side has such weapons and the other doesn't, the outcome will be swift and very lopsided. This is just a specific instance of the general rule that technological superiority plays an important and often critical role in determining the victor in battle. Clearly, we will need much further research into defensive systems as this technology becomes more mature.'

2.5.4 Public acceptance of nanotechnology

In spite of the concerns highlighted above, both precautionary principle and industry advocates agree that there is time to create dialogue and consensus that could prevent the kind of confrontations occurring that plagued the development of biotechnology. In this way, the objective of industry is to launch pre-emptive strikes against any problems with public acceptance of nanotechnology that might arise down the line (Gorman, 2002). The earliest example of this is the Foresight Institute, a think-tank founded in 1986 primarily to facilitate public understanding and discussion of the policy issues surrounding the development and deployment of nanotechnology. More recently, nanotech researchers have been urged to build on the example of the Ethical and Social Implications (ELSI) project (an interdisciplinary effort within the Human Genome Project). That is, to *'take a hard look at potential ethical and cultural issues, but follow through much more carefully and get out ahead of the public'* (Paul Thompson, Professor of Ethics at Purdue University, quoted in Leo, 2001). Indeed, the NNI has long acknowledged a need to integrate societal studies and dialogues concerning the perceived dangers of nanotechnology with its investment strategy, and the resulting White

House Fact Sheet (2000) promised that the impact nanotechnology has on society from legal, ethical, social, economic, and workforce preparation perspectives would be studied. These aims have already been realised to some extent. For example, the 2001 NSF report entitled *Implications of Nanoscience and Nanotechnology* takes a long, hard look at a range of hypothetical social ramifications (Roco and Bainbridge, 2001).

This industry strategy has been received with mixed reaction. Some commentators, such as Ho (2002a), have praised scientists for informing the public with *'clarity and candour.'* Others, on the other hand have not been nearly so generous in their assessment. Herrera (2002), for example, sums up the present state of the nanotech industry as being comparable to a *'sitting duck'*, just as biotech was during the 1990s, because it is not taking the issue of public acceptance seriously. Herrera continues: *'Ask members of the nanotechnology community if there are any obvious or potential controversies that they should be watching for, and they will say 'no'... Scientists think about ethics but they don't let it interfere with their work.'*

At present, the majority of controversy in this area surrounds the interaction of nanomaterials with the environment and their implications for human health. Vicki Colvin of CBEN believes that *'scientists' experience with other particulate matter argues for a thorough examination of how nanoparticles might react in mammalian systems when they are inhaled or when there is skin exposure'* (cited in Schultz, 2002). In addition, nanotech manufacturing processes need to be examined for potential health impacts, for example the solvents used in the gases produced in the manufacture of carbon nanotubes. Outside of manufacturing, researchers should investigate the possible consequences of nanoparticles entering and accumulating in the food chain. Indeed, some of the ongoing work by CBEN, and other

organisations such as NASA and the EPA, has been alluded to above. However, it is becoming increasingly clear that this work alone is not sufficient for the scope of these issues. As Colvin (2002) notes:

'It is critical that more organisations and people devote time and money to these questions. This requires a change in the current climate: of the [US\$710 million in funding for the NNI in the fiscal year 2003, less than [US\$500,000 is devoted to the study of environmental impact. It is difficult to convince scientists, or funding managers, to support environmental impact studies. The immediate payback for research that demonstrates ways of using nanomaterials to cure disease, for example, is greater than the reward for uncovering that a nanomaterial may cause disease.'

One way in which prevailing industry attitudes may be influenced is through the idea that information about unintended effects (whatever its conclusions), rather than alarming investors, in fact reassures, thus increasing the likelihood that viable nanotechnology products are developed. Most importantly, hard data on the environmental effects of nanomaterials could go a long way to building the public's trust (Colvin, 2002). This is in contrast to, for example, the controversy that surrounded the pesticide DDT in the 1960s and early 1970s: by refusing to acknowledge the demonstrable environmental harm caused by DDT, the US chemical industry lost a controversial but effective product, particularly for control of mosquitoes and mosquito-borne diseases.

2.5.5 The regulation debate

The precautionary approach upholds that regulatory action may be taken, based on the possibility of significant environmental damage, even before there is conclusive, scientific evidence that the damage will occur (European Environment Agency, 2003). Perhaps the most vigorous example of this

concerns the ETC Group, who have called for a global moratorium on the manufacture of nanomaterials until such a time when their interactions with living systems are more fully understood (McCullagh, 2002). Such an appeal is well-placed within this precautionary worldview, and nano-advocates have had to respond quickly with a number of forceful counter-arguments. Many of these claims stem from the diversity of envisaged nanotech applications and products (i.e. essentially a vast array of very small components), the difficulties of defining nanotechnology, and its broad interdisciplinary scope. Indeed, the convergence of a wide number of scientific disciplines within the field of nanotechnology certainly complicates the practicalities of enforcing such a ban, especially when one considers that pushing research underground may increase either the danger of deliberate misuse, or at least the difficulty of ensuring that usage remains within responsible boundaries.

As an alternative, nano-enthusiasts advocate a more modest regulation structure combined with robust civilian research. Such an approach would focus work on the potential risks and benefits of nanotechnology, whilst ensuring that safe practices are exported to developing countries. (Indeed, it is in the interests of developing countries to adopt good practice, otherwise investment will flop). Thus, such a regime should be based on the monitoring of the sale of such technologies, rather than control. This situation is analogous to biotechnology: the DNA experience, for example, suggests that a combination of self-regulation and government co-ordination can answer legitimate safety concerns while allowing scientific research to flourish (Reynolds, 2002). Thus, while there is no way of knowing, *a priori*, the unintended and higher order consequences of nanotechnology, the participation of environmental and social scientists in the field may allow for

important issues to be identified earlier, the right questions to be raised, and necessary corrective actions to be taken. It does seem likely that some form of regulatory control will be necessary to assure that nanotechnology is developed safely – *‘safe designs, safe procedures and methods to test for potentially hazardous assemblers can be incorporated into standards by consensus of interested parties’* (Forrest, 1989). The greatest danger, however, appears to be intentional abuse of the technology, so certain aspects of development should be performed in a secure environment.

2.6 Discussion

While many of the nanotechnologies covered in this part of the report might appear advanced, it is fair to conclude that most contemporary experimental capabilities in this area are still in their infancy. This means that it is extremely difficult to foresee many outcomes that developments in this field will bring over the next 10 years, let alone assess their likelihood. Initially, it is probable that the impact of nanotechnology will be limited to a few specific products and services, where consumers are willing (or able) to pay a premium for new or improved performance. Looking further ahead, controversy surrounds the possibility of realising some of the wilder visions of a nanotech-enabled future. This is in spite of the fact that many of these ideas stem from quite straightforward concepts founded in solid science (Holister, 2002);

we are unlikely to witness any radical developments during the next 15 years unless a series of fundamental breakthroughs occur between now and then. However, as the range of associated tool and fabrication techniques begin to mature, the field is set to become increasingly commonplace in the coming decades. Ultimately, then, the longer-term structural impact of nanotechnology on a whole range of sectors – in manufacturing, transport, services and domestic practice – could be substantial in 30–50 years. These changes are likely to be gradual as, on the whole, the displacement of an old technology by a new one tends to be both slow and incomplete (NSF, 2001).

In the meantime, a number of well-founded short-term concerns remain, many of which revolve around issues of human health. Considering past experiences of industry and government mismanagement in this area (notably through GM-related controversy), nano-advocates would do well to sit up and take note. For, although an externally imposed nanotech moratorium seems both unpractical and probably damaging at present, industry may find such a fate virtually self-imposed if they do not take the issue of public acceptance seriously. This report has shown some nano-advocate awareness of environmentally-sound practice. Industry must demonstrate a commitment to this by funding the relevant research on a far greater scale than currently witnessed.

3. Artificial Intelligence and Robotics

3.1 Introduction

3.1.1 About AI and robotics

AI has been one of the most controversial domains of inquiry in computer science since it was first proposed in the 1950s. Defined as the part of computer science concerned with designing systems that exhibit the characteristics associated with human intelligence, the field has attracted researchers because of its ambitious goals and enormous underlying intellectual challenges (National Research Council [NRC], 1999). The ultimate aim is to make computer programmes that are capable of solving problems and achieving goals in the world as well as humans – the pursuit of so-called ‘strong AI’. This goal has caught the attention of the media, but by no means do all AI researchers view strong AI as worth investigating – excessive optimism in the 1950s and 1960s concerning strong AI has given way to an appreciation of the extreme difficulty of the problem (Copeland, 2000). To date, progress in this direction has been meagre. Because 50 years of failure eventually starts to affect funding, the AI field has diversified and experts have established themselves in other areas where they can be said to have had some success. These new areas are less concerned with the business of making computers think, focusing instead on what can be referred to as ‘weak AI’ – the development of practical technology for modelling aspects of human behaviour (Goodwins, 2001). In this way, AI research has produced an extensive body of principles, representations, and algorithms. Today, successful AI applications range from custom-built expert systems to mass-produced software and consumer electronics.

Robotics, on the other hand, may be thought of as *‘the science of extending human motor capabilities with machines’* (Trevelyan, 1999). However, a closer look at this definition creates a more complicated picture. For example, a cruise missile,

although not intuitively referred to as a robot, nevertheless incorporates many of the navigation and control techniques explored in the context of mobile-robotics research. Furthermore, robots are not necessarily dependent on hardware for their operation. It is possible, for instance, to conceive of intelligent entities that operate purely within information systems – the so-called ‘softbots’ or ‘software agents’ – as robots (Doyle and Dean, 1996). It is noteworthy, however, that such distinctions between ‘hard’ and ‘soft’ are bound to fade in importance in the future as physical agents enter into electronic communication with each other and with online information sources, and as informational agents exploit perceptual and motor mechanisms. It is difficult, then, to state categorically exactly what constitutes a robot. This report, however, considers robotics research as the attempt to instil intelligent software with some degree of motor capability. Since many of the major areas of AI research play an essential role in work on robots, robotics will be considered here as a sub-section of AI.

3.1.2 Where are we now?

As alluded to above, the field of AI has not moved along as quickly as innovators have predicted. One reason for this has been the damaging cycle of hype and disappointment within the industry, and the accompanying rise and fall in research investment⁴ (pers. comm., Murray Shanahan, Imperial College London, 17 Jan 2003). This began in the 1960s when general enthusiasm surrounding the prospects of AI moved in parallel with the exciting developments of the computer. However, this optimism resulted in downfall during the 1970s when the work failed to produce, climaxing in the UK with the highly damaging 1973 Lighthill Report – a government commissioned paper from the Science and Engineering Research Council which damned AI and recommended withdrawal of research funding. In addition, the same kind of official doubts which the

Lighthill Report made explicit in the UK lay, less explicitly, behind a similar slow down in research funding in the US (Malcolm, 2001). The next big rise in AI funding occurred in the 1980s, mainly in reaction to Japanese enthusiasm for the field. In Japan, the 5th Generation project was born; the UK reacted through the Alvey initiative, which now focused on ‘knowledge based systems’ so as to avoid any awkward parallels between current research and the previously condemned AI. Again, both projects were characterised by a lack of progress and AI research failed to make it into the mainstream. Most recently, the early to mid 1990s has seen the emergence of software agents, and the resulting excitement has once again sparked a rise in investment. In addition, the field of robotics has become much more influential of late, particularly through the entertainment industry. Many of these developments are described in more detail later on in this report.

Today, AI is about at the same place the PC industry was in 1978 (Brooks, 2001) – the waves of enthusiasm that accompanied the developments of computers have long gone and researchers are beginning to come to terms with how hard the problems of AI really are. However, technological know-how is not the only obstacle that the AI industry faces – another is the purported ‘AI effect’ whereby the existence of AI in modern software products go largely unnoticed despite the widespread use of such applications (Stottler Henke, 2002). Indeed, AI is considered by some researchers to be an unimplementable technology: as soon as the technology advances, the perspective shifts, and the quality of intelligence passes to those activities that are still only in the human domain (Joseph, 2001). For example, many of those in industry do not use the term ‘artificial intelligence’ even when their company’s products rely on some AI techniques (Stottler Henke, 2002). The exact reasons for the AI effect are uncertain, but

it is likely that the phenomenon developed in reaction to the kind of historical tendencies to oversell the industry alluded to earlier.

3.2 Aspects of Research

3.2.1 Introduction

The above section has described how researchers have re-evaluated their expectations with regard to achieving strong AI. Associated with this reality check is the recognition that classical attempts at modelling AI, based upon the capabilities of digital computers to manipulate symbols, are probably not sufficient to achieve anything resembling true intelligence. This is because symbolic AI systems, as they are known, are designed and programmed rather than trained or evolved. As a consequence, they function under rules and, as such, tend to be very fragile, rarely proving effective outside of their assigned domain (Hsuing, 2002). In other words, symbolic AI is proving to be only as smart as the programmer who has written the programmes in the first place.

In realisation of this, scientists are beginning to look much more closely at the mechanisms of the brain and the way it learns, evolves and develops intelligence from a sense of being conscious (Aleksander, 2002). For example, AI software designers are beginning to team up with cognitive psychologists and use cognitive science concepts. Another example centres upon the work of the ‘connectionists’ who draw attention to computer architecture, arguing that the arrangement of most symbolic AI programmes is fundamentally incapable of exhibiting the essential characteristics of intelligence to any useful degree. As an alternative, connectionists aim to develop AI through artificial neural networks (ANNs). Based on the structure of the nervous system, these ‘computational-cognitive models’ are designed to exhibit some form of learning and ‘common-sense’ by drawing links between meanings (Hsiung, 2002). ANNs,

then, work in a similar fashion to the brain: as information comes in, connections among processing nodes are either strengthened (if the new evidence is consistent) or weakened (if the link seems false) (Khan, 2002).

The emergence of ANNs reflects an underlying paradigm change within the AI research community and, as a result, such systems have undeniably received much attention of late. However, regardless of their success in creating interest, the fact remains that ANNs have not nearly been able to replace symbolic AI. As Grosz and Davis (1994) remark: *'[Symbolic AI has] produced the technology that underlies the few thousand knowledge-based expert systems used in industry today.'* A major challenge for the next decade, then, is to significantly extend this foundation to make possible new kinds of high-impact application systems. A second major challenge will be to ensure that AI continues to integrate with related areas of computing research and other fields (Doyle and Dean, 1996). For example, the kinds of developments described in Section 2 for nanotechnology may go some way to accelerating progress in AI, particularly through the sensor interface. For these reasons, the list of main research areas that follows should be regarded as neither exhaustive nor clear-cut. Indeed, future categorisations will again change as the field solves problems and identifies new ones.

3.2.2 Learning

According to Daniel Weld (1995) of the University of Washington, machine learning addresses two interrelated problems: *'the development of software that improves automatically through experience and the extraction of rules from a large volume of specific data.'* Systems capable of exhibiting such characteristics are important because they have the potential to reach higher levels of performance than systems that must be modified manually to deal with situations their designers did not anticipate (Grosz and

Davis, 1994). This, in turn, allows software to automatically adapt to new or changing users and runtime environments, and to accommodate for the rapidly increasing quantities of diverse data available today. When designing programmes to tackle these problems, AI researchers have a variety of learning methods at their disposal. However, as alluded to above, ANNs represent one of the most promising of these.

3.2.2.1 Artificial neural networks

There are many advantages of ANNs and advances in this field will increase their popularity. Their main value over symbolic AI systems lies in the fact that they are trained rather than programmed: they learn to evolve to their environment, beyond the care and attention of their creator (Hsuing, 2002). Other major advantages of ANNs lie in their ability to classify and recognise patterns and to handle abnormal input data, a characteristic very important for systems that handle a wide range of data. Furthermore, many neural networks are biologically plausible, which means they may provide clues as to how the brain works as they progress. Like the brain, the power of ANNs lies in their ability to process information in a parallel fashion (that is, process multiple chunks of data simultaneously). This, however, is where the limitations of such systems begin to arise: unfortunately, machines today are serial – they only execute one instruction at a time. As a consequence, modelling parallel processing on serial machines can be a very time-consuming process (Matthews, 2000a). A second problem relates to the fact that it is very difficult to understand their internal reasoning processes and therefore to obtain an explanation for any particular conclusion. As a result, they are best used when the results of a model are more important than understanding how the model works. To this end, these systems are often used in stock market analysis, fingerprint identification, character recognition, speech recognition,

and scientific analysis of data (Stottler Henke, 2002).

3.2.3 Reasoning about plans, programs and action

Intelligent systems must be able to plan – to determine appropriate actions for their perceived situation, and then execute them and monitor the results. However, in spite of the fact that this area has been under active research since the 1950s, AI planning applications are furthest from human-level (Grosz and Davis, 1994). Ordinary people, for example, manage to accomplish an extraordinary number of complex tasks just using simple, informal thought processes based on a large amount of common knowledge. AI, on the other hand, is far behind humans in using such reasoning except for limited jobs, and tasks that rely heavily on common-sense reasoning are usually poor candidates for AI applications (Stottler Henke, 2002). In the past, researchers have mainly had to rely on the development of algorithms that *‘automatically construct and execute sequences of primitive commands in order to achieve high-level goals’* (Weld, 1995). More recently, the field of plausible reasoning has demonstrated its feasibility in tackling the problem of representing, understanding, and controlling the behaviour of agents or other systems in the context of incomplete or incorrect information (Weld, 1995). Another development that may lead to significant advances in the area of artificial reasoning is fuzzy logic. Traditional Western logic systems assume that things are either in one category or another. Yet in everyday life, we know this is often not precisely so. Fuzzy logic, then, provides a way of taking into account our common-sense knowledge that most things are a matter of degree when a computer is automatically making a decision (Stottler Henke, 2002). Thus, in spite of the difficulties inherent in this field of AI, planning systems have been successfully developed for several tasks to date, including

factory automation, military transportation scheduling, and medical treatment planning. These will be covered in more detail below.

3.2.4 Logical AI

This type of reasoning concerns what a programme knows about the world in general, the facts of the specific situation in which it must act, and the goals that it must accomplish (Grosz and Davis, 1994). Such concepts are held within the programme in the form of sentences of some mathematical logical language. The most successful example of this is an expert system, created when a ‘knowledge engineer’ interviews experts in a certain domain and tries to embody their knowledge in a computer programme for carrying out some task, such as diagnosis. However, the usefulness of current expert systems also depends on their users demonstrating a certain level of common-sense too.

3.2.4.1 Algorithms and genetic programming

An algorithm is defined as a *‘detailed sequence of actions to perform to accomplish some task’* (FOLDOC, 2003). One branch of algorithm theory, genetic programming, is currently receiving much attention. This is a technique for getting software to solve a task by ‘mating’ random programmes and selecting the fittest in millions of generations. Khan (2002) elaborates: *‘Genetic algorithms use natural selection, mutating and crossbreeding within a pool of sub-optimal scenarios. Better solutions live and worse ones die – allowing the programme to discover the best option without trying every possible combination along the way.’*

3.2.5 Collaboration

The ubiquity of computers, networks and distributed information resources means that collaboration between these entities is important. The field of multiagent coordination concerns itself with the problem of endowing agents with the ability to communicate with each other to reach

mutually beneficial agreements (Grosz and Davis, 1994). In addition, specialised techniques must also be developed that enable an agent to represent and reason about the capabilities of other agents (Weld, 1995). These types of systems are dealt with by the EU Disappearing Computer project and are expanded upon later (Section 3.3.5.1) due to their focus on spatially distributed artefacts.

3.2.6 Perception

Many AI systems require an ability to handle several different types of perceptual information (Grosz and Davis, 1994). The most important of these are expanded upon below.

3.2.6.1 Pattern recognition

The speed with which people extract information from images makes vision the preferred perceptual modality for most people in the majority of tasks, thus implying that easy-to-use computers should be capable of both understanding and synthesising images. One of the goals of computer-vision research is image understanding and classification. Depending on the application, the imagery to be understood might include a scanned document page, a mug shot, an aerial photograph, or a video of a home or office scene (Weld, 1995). Typical state-of-the-art tasks include facial recognition; object recognition and reconstruction; hand tracking and gesture recognition; and document analysis and recognition. However, while current computer-vision techniques are capable of impressive feats under controlled conditions, such techniques often prove to be brittle and non-robust under real-world conditions (Grosz and Davis, 1994).

3.2.6.2 Understanding natural language

The ultimate goal of natural language-processing research is to create systems able to communicate with people in natural languages. Such communication requires an ability to understand the meaning and

purpose of communicative actions, such as spoken utterances, written texts, and the gestures that accompany them and an ability to produce such communicative actions appropriately. These abilities, in their most general form, are ‘far beyond current scientific understanding and computing technology’ (Weld, 1995). However, the potential relevance of natural language processing to industry is immense, as such systems could be central to the next generation of intelligent interface.

3.2.7 Human-computer interaction

This area of AI follows on from perception in that people use a number of different media to communicate, including: spoken, signed and written languages; gestures; sounds; drawings; diagrams; and maps (Grosz and Davis, 1994). In particular, knowledge representation is important due to its powerful effect on the prospects for a computer or person to draw conclusions or make inferences from that information (Stottler Henke, 2002). Consequently, work in this area seeks to discover expressive, convenient, efficient, and appropriate methods for representing information about all aspects of the world.

3.2.8 Public funding

According to hi-tech consultancy, Gartner Dataquest (cited in BBC, 2002), one billion PCs have been sold across the world, with numbers anticipated to rise rapidly in the next few years, reaching the two billion mark in by 2008. The level of interconnectedness between such machines is also set to rise: in this decade, half a billion human-operated machines and countless computers – in the form of appliances, sensors, controllers, and the like – will be linked (Dertouzos, 1999). This, in turn, will lead to an explosion in the Internet economy. Today, some US\$50 billion changes hands over this system, but by 2030 this flow will amount to US\$4 trillion of today’s dollars, or one quarter of

the world's economy. Obviously, in such a future scenario, the extraordinarily sophisticated systems used to control communications, power, stock exchanges, and monetary assets can break down and might come under attack. Today, given the relative complexity and unreliability of the Internet, it is not surprising that commentators view this scenario with increasing trepidation. AI, then, is seen by many as having an essential role in a future where commercial and military information warfare is a major, perhaps dominant, characteristic. In addition, AI is touted by many (e.g. see Dertouzos, 1999) as having the potential to greatly improve human productivity and ease of use within this prospective network.

As computer science, and AI in particular, is considered to be of strategic importance, it is worth here briefly examining government funding in this area. In general, computer science receives a relatively small proportion of the research funding in many countries, even if anecdotal evidence suggests that the fractions are increasing (Schneider and Robb, 2001). This is in spite of the fact that public funding has played an important part in AI research in the past, largely because of the field's high-risk conceptual challenges. However, the picture is complicated by the fact that international comparisons of research funding in this area are difficult to make, since different countries use different funding methods. France and Japan, for example, rely heavily on national research institutes and laboratories, rather than expecting most research to be done in university departments. In addition, funding for AI research is reported far less thoroughly than it is for nanotechnology. Consequently, relevant information has often been obscure, and it has been necessary instead to report funding in computer science in general. As a rule, AI budgets will represent a small proportion of these figures.

3.2.8.1 The US

Historically in the US, the concept of AI originated in the private sector, but the growth of the field has depended largely on public investments. Today, computer science research in the US is funded by a number of governmental agencies. Total US government computer science research expenditures in 1998 were US\$1399 million, with approximately one third devoted to what was described as 'basic research' (Schneider and Robb, 2001). Three agencies (NSF, DARPA, and the Department of Energy [DOE]) in the US together support US\$365 million of this work, and the NSF is responsible for funding the lion's share (Schneider and Robb, 2001). In addition, a number of other agencies are of note. These include the National Institutes of Health (NIH) and NASA which have also pursued AI applications of particular relevance to their own separate agendas (NRC, 1999).

Of these institutions, DARPA is credited with considerable advancement of the field from the 1960s onwards. This hardly comes as surprising when one considers the close link that exists between the military and computer science – in fact, the early development of computing was virtually exclusively limited to military purposes (Matthews, 2000b). The most famous example of this concerns the development of the Internet, in which DARPA played a central role in the 1970s and 1980s⁵. More recently, less visible but arguably equally significant developments have come to the fore. For example, a 1994 report by the AAAI paraphrased a former director of DARPA, saying that DART (the intelligent system used for troop and material deployment for Operation Desert Shield and Operation Desert Storm in 1990 and 1991) '*justified DARPA's entire investment in AI technology*' (cited in NRC, 1999). One consequence of this is that modern-day battle relies heavily on data networks. This has been stressed by A. Michael Andrews, the US

Army's Deputy Assistant Secretary for Research and Technology, saying: *'Everything relies on a reliable and secure network. Without it, our vulnerability is exposed'* (quoted in Machan, 2002).

Today, DARPA's funding for AI research is spread among a number of programme areas, each with a specific application focus. For example, funding for AI is included in the Intelligent Systems and Software programme, which received roughly US\$60 million in 1995. This applied research programme is intended to leverage work in intelligent systems and software that supports military objectives, enabling information systems to assist in decision-making tasks in stressful, time-sensitive situations. Additional DARPA funding for AI is contained in the Intelligent Integration of Information programme, which is intended to improve commanders' awareness of battlefield conditions. DARPA continues to fund some of the more basic research in AI as well. Such funding is included in its information sciences budget, which declined from US\$35 million to US\$22 million annually between 1991 and 1996. The AI funding supports work in software technology development, human-computer interfaces, microelectronics, and speech recognition and understanding (NRC, 1999).

In addition to DARPA, NASA has also built up a reputation for high-risk, high-impact AI research. One obvious development is the Pathfinder robot, which used a number of modern AI and robotics techniques to explore the surface of Mars. Another is the Deep Space One mission in which an 'autonomous' controller (i.e. without human intervention) was able to fly a spacecraft for part of a mission and exceeded all performance goals. NASA is expected to continue exploring this technology heavily into the future and, based on these earlier successes, can be considered as a major innovative player within this field (Hendler, 2000).

3.2.8.2 Japan and Europe

Although the US has played a central role in developing the AI research agenda, other countries and regions have also played their part. One of the most notable examples of this occurred in the early 1980s when both Japan and Europe dramatically increased their funding of AI research, partly as a reaction to the newly emerged expert systems industry. One of the most ambitious projects undertaken was the 5th Generation Computer Systems Project, an attempt to combine European ingenuity with Japanese industrial skill in order to develop a new sort of AI that might rival the US's domination in the field. However, 5th Generation project technology never really made it into the mainstream, largely because its inflexible theoretical basis was found to be inferior to the less elaborate, rough-and-ready approaches to AI development pursued by the US (Joseph, 2001). The latest collaborative attempt by these two parties to break US hegemony in AI is the Real World Computing Project (RWCP), or the 6th Generation Computing Project, a 10-year programme that started in 1992 (around the end of the 5th Generation project). This time the RWCP has a much broader remit: to focus on a variety of different 'softer' technologies that use neural or fuzzy techniques. Thus, the research components are much more spread out, and appear to have been selected with an eye for more practical applications of the latest technology (Joseph, 2001).

In addition to the combined effort above, both parties have also more recently established research programmes of their own. In Japan, for example, the National Institute of Informatics (NII), an inter-university research institute under the Ministry of Education, Culture, Sports, Science and Technology (MEXT), is pursuing a programme to expand the field of IT. Established in April 2000, intelligent systems, which form one component of MEXT's seven-sided agenda, aim to develop advanced

technology for next-generation symbiotic robots and systems, and new models for information sharing and exchanging (NII, 2002). This programme is closely linked to the Japanese government's seven-year plan to develop humanoid robots. In fact, Japan is the clear leader in using industrial robots as it accounts for over half of all units in the world (The Economist, 2001).

The EU, too, is engaged in AI-related research. Perhaps the most ambitious of this is related to the EU Framework VI proposal for spending €16.29 billion over 2002–2006, of which 27% is destined for IT (Schneider and Robb, 2001). In addition to central EU funding, individual European states are also developing their own research agendas. In the UK, the Information Technology/Computer Science (IT/CS) Programme in the Engineering and Physical Sciences Research Council (EPSRC) budget for 2000/2001 is £70.3 million; investment in computer science research is about 45% of this (EPSRC, 2003). Other EU countries are also of interest, particularly Germany, France and Scandinavia, where the latter is particularly well advanced in the use of computing and IT. However, other smaller countries are not making much in the way of substantial commitments to computer science research.

3.3 Applications

3.3.1 Introduction

The above section has demonstrated the diverse and multifaceted nature of AI research, and this work has resulted in an extensive body of principles, representations, algorithms, and spin-off technologies (Weld, 1995). The relative state of infancy of research into strong AI means that this field can be put aside for the time being. Rather, this section will attempt to elaborate upon weak applications of AI, where, it is fair to say, considerable effort in this area has resulted in some real-world product success. In fact, the actual and potential uses of weak

AI are virtually endless: one measure of the growth of practical applications is the number of patents mentioning the term AI and related terms. According to the US Patent Office, only about 100 patents specifically mentioned AI a decade ago; in contrast, in 1999 about 1,700 patents mentioned AI with another 3,900 or so mentioning related terms (Buchanan and Uthurusamy, 1999). However, it is worth bearing in mind that the actual prevalence of emerging AI technology may be greater than this due to classification-related difficulties and the fact that such products are more likely to be embedded in some larger system than a stand-alone machine. In general, such applications are used to increase the productivity of knowledge workers by intelligently automating their tasks, or to make technical products of all kinds easier to use for both workers and consumers through intelligent automation of their complex functions (Stottler Henke, 2002). It is possible now to identify four families of intelligent systems that have broad applicability across a wide range of sectors (Grosz and Davis, 1994). These are intelligent simulation systems; intelligent information resources; intelligent project coaches; and robotics.

3.3.2 Intelligent simulation systems

These applications are commonly used in a number of different scenarios. First, an Intelligent Simulation System (ISS) may be generated to learn more about the behaviour of an original system, when the original system is not available for manipulation. The modelling of climate systems is a good example. Second, the original system may not be available because of cost or safety reasons, or it may not be built yet and the purpose of learning about it is to design it better (Stottler Henke, 2002). Third, an ISS might be employed for training purposes in anticipation of dangerous situations, when the cost of real-world training is prohibitive. Such technologies are particularly well-

advanced in military applications through the simulation of war ‘games’. Another very big business in the realm of ISSs is the video-game market, comparable to the film business in size. AI systems have become fundamental to this industry because, unlike in film, it is often up to a computer or game console to create a sense of reality for the game-player. Such standards of realism are going up all the time (Broersma, 2001).

3.3.3 Intelligent information resources

Intelligent systems must be able to provide access to a wide variety of information, including visual and audio data, in addition to commonplace structured databases (Grosz and Davis, 1994). One development in this area that is receiving much attention is ‘data mining’, the extraction of general regularities from online data (Weld, 1995). This area is becoming increasingly important due to the fact that all types of commercial and government institutions are now logging huge volumes of data and require the means to optimise the use of these vast resources (Stottler Henke, 2002). Indeed, according to the market research firm IDC (cited in Dalesio, 2002), revenue from sales of all types of data mining software are anticipated to grow from about US\$540 million this year to about US\$1.5 billion in 2005.

Looking beyond data mining, other technologies are also appearing on the horizon. For example, SilverEgg Technologies, a Japanese venture company, have developed Aigent, a system that observes which product categories a customer clicks on, and then makes intelligent guesses about that customer’s preferences (Joseph, 2001). Another development in this area concerns the ‘heuristic’: *‘A rule of thumb, simplification, or educated guess that reduces or limits the search for solutions in domains that are difficult and poorly understood’* (FOLDOC, 2003). Thus, in terms of AI, heuristics is a way of trying to discover something or an

idea embedded in a programme. By 2006, it is anticipated that companies will be able to use this kind of software to analyse customer feedback, whether it comes from the Internet, call centres, or sidewalk surveys. Market research divisions, too, will be able to better track competitors, sales trends and research extracted from huge volumes of patents, scientific articles and news reports (Dalesio, 2002). These developments hint of the ‘next big thing’ in industry – ‘business intelligence’. These systems, already in limited application today, improve on data mining services by presenting their findings in more useful formats – using advanced visualisation tools – and by deploying AI to look for patterns that human users might not look for. The potential value of such technology to business has already created fierce competition: established software companies like IBM, Microsoft and Oracle, along with younger competitors like Business Objects, MicroStrategy and Moreover.com, are vying for their share of a market that is expected to grow from US\$3.5 billion in 2002 to US\$8.8 billion in 2004 (Miller, 2001).

In general, the above examples carry out tasks for one Web site or organisation. However, some innovators envisage the technology going a lot further than this. For instance, it is not hard to imagine a future world of semi-autonomous agents, roaming the Web and carrying out various tasks for their owners. Such agents could be given a rough idea of what we want, do some comparison-shopping, and order the best deal, just like a real personal assistant. Ultimately, virtual organisations composed of autonomous agents, which could form spontaneously to carry out a specific task and then disband again, might be possible (Broersma, 2001).

3.3.4 Intelligent project coaches

This section represents the most diverse range of applications: intelligent project coaches can function as co-workers, assisting

and collaborating in a wide range of design or operations teams for complex systems. For basic personal use, 'interface agents' are computer programs that employ AI techniques to provide active assistance to a user during computer-based tasks. These agents acquire their competence by learning from the user as well as from agents assisting other users. To date, several prototype agents have been built using this technique (Maes, 1994). For example, US start-ups, such as Saffron Technology and Manna, are marketing software tools that learn the individual user's buying patterns and make personalised recommendations accordingly.

In addition to interface agents, the next 10 years are likely to see rapid AI development occurring in speech recognition (Hendler, 2000). Indeed, computer speech input has already arrived and is commercially available – many telephone services use speech recognition at present. In addition, cell phones without keypads are likely to reach the market as early as next year. These devices are anticipated to enhance the use and appeal of the mobile Internet by allowing users to call up any Web page from a mobile device just by speaking its address. Voice recognition also has security applications: in a demo at the 3GSM World Congress in February 2002, Mitsubishi Electric demonstrated a SIM card featuring voice validation software developed by Domain Dynamics Ltd. The software provides a 'biometric template' that can recognise a person's voice to properly identify a user – *'a necessity when providing access to corporate or private databases over the Internet'* (Mokhoff, 2002). With regard to speech recognition's natural successor, natural language processing, such technology is, to date, poorly developed and computers are not yet able to even approach the ability of humans to extract meaning from natural languages (Stottler Henke, 2002). However, due to the many potentially

valuable practical applications of this technology, developments in this area are expected to advance quickly. For example, automated language translation also looks set to mature sometime between 10–15 years from present.

Perhaps the most ambitious examples of AI development that are currently occurring in this area relate to computer learning. One example is the ANN, Falcon. Designed by San Diego-based HNC Software, Falcon maintains a profile of how, when, and where customers use their credit cards and, from this, develops an ability to discern 'deviant' behaviour. To date, this system is used by nine of the ten leading US credit card companies: they claim it has improved fraud detection rates from 30–70% (Khan, 2002). Another example – and one that is probably the most challenging in ANN development today – is being undertaken by DARPA, who have launched an initiative to develop a cognitive (i.e. thinking) system. The aim of this system is to reason in a variety of ways, learn from experience, and adapt to surprises. In the words of Melymuka (2002): *'It will be aware of its behaviour and explain itself...It will be able to anticipate different scenarios and predict and plan for novel futures.'* The ultimate aim is to develop cognitive systems capable of assisting or replacing soldiers on hazardous duty or civilians responding to toxic spills or disasters.

In addition to AIs that focus on novel ways of learning, other programmes exist which can be said to primarily reason. Perhaps the most successful example in operation today is the SmartAirport Operations Centre, a logistics programme created by Ascent Technology. This AI uses genetic algorithms to plan airport timetables by calculating how to optimise complicated scenarios. Other reasoning programmes are based on heuristic classification – a form of expert system – and are generally considered the

most feasible given the present knowledge of AI. These AIs have found their way into cockpits of fighter-pilots, where their main role is to reduce the workload on the pilot by providing advice in certain stressful situations (Matthews, 2000b).

3.3.5 Robotics

A distinction has already been drawn above (Section 3.1.1) between robots working in informational environments and robots with physical abilities. One advantage of the former is that there is little need for investment in additional expensive or unreliable robotic hardware as existing computer systems and networks provide adequate sensor and effector environments. On the other hand, the kinds of robotics systems elaborated on here, physical robots, require mechanisation of various physical sensory and motor abilities (Doyle and Dean, 1996). The challenges involved in providing such a latter environment are considerable, especially when complete automation is sought, as in Honda's humanoid ASIMO project⁶. Thus, rather than focus on the ambitious and distant goal of relative autonomy, this report picks up on Trevelyan (1999) who points out that complete automation is often unfeasible, impossible, or simply unwanted. Indeed, much of today's robotics research focuses instead on far humbler goals, such as simplicity, force control, calibration and accuracy. Thus, we can see that, to some extent, the field of robotics has followed similar lines as that of AI, attempting to rebound from the overly optimistic predictions of the 1950s and 1960s, and coming up against more contemporary problems not dissimilar to the AI effect. Indeed, while few of the innovations that emerge from the work of robotics researchers ever appear in the form of robots, or even parts of robots, their results are widely applied in industrial machines not defined as so (Trevelyan, 1999).

In spite of these significant challenges, there are some good examples of AI-controlled robotic systems. For instance, TriPath Imaging has built FocalPoint, a diagnosis expert system that examines Pap smears for signs of cervical cancer. FocalPoint screens five million slides each year, or about 10% of all slides taken in the US and, like human lab technicians in training, teaches itself by practising on slides that pathologists have already diagnosed. Thus, one big advantage of such a system is that, if implemented properly, FocalPoint allows you to replicate your very best people (Khan, 2002).

A second example and, again, perhaps the most ambitious of all, concerns DARPA, who are in the process of developing an Unmanned Combat Air Vehicle (UCAV). According to Boeing (2002), the UCAV system is designed to *'prove the technical feasibility of multiple UCAVs autonomously performing extremely dangerous and high-priority combat missions.'* In a typical mission scenario, *'multiple UCAVs will be equipped with pre-programmed objectives and preliminary targeting information from ground-based mission planners. Operations can then be carried out autonomously, but can also be revised en route by UCAV controllers should new objectives dictate.'* If the program is a success, the US DoD expects to begin fielding UCAV weapon systems in the 2008 time-frame.

3.3.5.1 Robot teams

Expanding upon the concept of collaboration highlighted above, one area of AI that is showing much promise is 'ubiquitous computing' using information artefacts: future forms of everyday objects that represent a merging of current everyday objects with the capabilities of information processing and exchange. For example, the EU-funded initiative of the Information Society Technologies (IST) research programme aims to show how such artefacts can be made to work together, and in particular how they provide behaviour or

functionality that exceeds the sum of their parts (The Disappearing Computer, 2003). It is from these ideas that the concept of robot teams begins to emerge. Robot teams potentially have applications in a wide range of areas. This is because robots working in teams *'allow for solutions in which knowledge, expertise, and motor capability may be distributed in time and space'* (Maes, 1994). Thus, while individual robots may only have limited capacity, robots working together in groups might be able to perform complex tasks. These include military surveillance, mine removal, automated household tasks, large scale laboratory projects (such as those used in the Human Genome Project) and assembly. In this way, most military planners believe that robots and remote-controlled sensors represent the future of information collection on the battle-field (Jeremiah, 1995).

3.3.6 Corporate funding

While Section 3.3.5 has demonstrated the significant commercial interest in AI, the picture for corporate investment in this area is a far less coherent. To date, unlike the field of nanotechnology, no significant overview of AI funding seems to exist in the literature. Having said this, however, the level of corporate support for AI application development is, in all likelihood, considerable: according to Henry McDonald, Director of the NASA Ames Research Centre, one-third of computer-science funding comes from government and two-thirds from industry (cited in Krill, 2002). This is not to say, though, that the interests of the scientific and business worlds necessarily concur; while AI may pose many fascinating questions for the former, such technology has to be commercially viable in the latter (Broersma, 2001). For this reason, no industry has yet identified a strong motive for developing strong AI and it is unlikely that scientists and business people will get any closer together in the future. The central focus here, then, must be on the utility of

products, rather than their degree of intelligence.

As alluded to above, one of the most commercially valuable frontiers of AI is e-commerce, where technicians are hoping to make the online world simpler and more capable at the same time. Robots, too, are potentially big business for the hi-tech companies prepared to invest in them: investment in robots world-wide increased markedly during 2000, with almost 100,000 new units being installed, raising the total stock of robots to 750,000 at the end of 2000 (The Economist, 2001).

3.3.6.1 The US and Japan

In general, the US is more widely regarded for its private software development than it is for its hardware, for which Japan is most highly thought of (Shim, 2002). Indeed, as noted earlier, the number of US AI-related patents in existence increased from 100 in 1989, to 1,700 in 1999. Private firms, including large manufacturers of electronics and computers, as well as major users of IT, hold a vast majority of these patents. The top three of these are IBM (297 patents), Hitachi (192) and Motorola (114), although another 17 companies make an appearance on the list⁷. Similarly, many of Japan's major companies have plans for AI. According to Shim (2002), the trick for major companies *'is to time things right so as to be on the cutting edge of the next big thing.'* For example, one area in which Sony – one of the most successful companies in the history of consumer electronics – has invested in heavily is the home-robot market, through its Entertainment Robot America division. Sony's latest development in this area concerns Aibo Recognition, a mechanical dog granted with the ability to recognise its owner's name, voice and face, as well as automatically recharge itself. By infusing Aibo with increased AI, such as voice and face recognition, the hope is to give Aibo owners the ability to interact with a robot at an unprecedented level (Spooner, 2002).

3.4 Reality and Hype

3.4.1 Introduction

The kinds of applications outlined above necessarily rely to some degree on weak AI. It might seem paradoxical, then, when one considers that it is the area of strong AI that features more prominently in the public imagination. To begin with, it is a well known fact that many revered members of the academic community deem the achievement of machine intelligence reaching, or even surpassing our own, as an inevitability (Barry, 2001). Most famously, this category includes Ray Kurzweil, inventor of the first reading machine for the blind, who believes that *'within 30 years, we will have an understanding of how the human brain works that will give us templates of intelligence for developing strong AI'* (cited in Anderson, 2001). In fact, as this section makes clear, the future of strong AI is highly uncertain, with considerable controversy present within the literature concerning whether it is even possible or not. The primary aim here, then, is to consider the technological and philosophical constraints within the field. From this, it should be clear that the issues raised by the possibility of strong AI are so fundamental that they cross many academic boundaries, including philosophy, sociology and psychology.

3.4.2 Barriers to strong AI

The standard test against which the possibility of strong AI is often judged concerns Alan Turing's 1950 article, *Computing Machinery and Intelligence*, in which the author discusses the conditions for considering a machine to be intelligent (Turing, 1950). He argues that if a machine could successfully pretend to be human to a knowledgeable observer then you certainly should consider it intelligent (McCarthy, 2003). This test would satisfy most people but not all philosophers, some of which have challenged the 'inevitable' achievement of strong AI based upon the assertion that the hypothesis of strong AI is itself false.

One famous sceptic of AI is Hubert Dreyfus, who says that a computer will never be intelligent unless it can display a good command of common-sense (Dreyfus, 1992). Dreyfus then follows up by saying that computers will never be able to fully grasp common-sense, since much of our common-sense is on a 'know-how' basis. For example, the notion that one solid cannot easily penetrate another is common-sense, yet the knowledge required to ride a bicycle is not something you can gain from a book, or from someone telling you. You can only learn through experience. Thus, since current computers can only really 'represent' things, the possibility of taking a skill, emotion, or something else equally abstract, and changing it into a series of zeros and ones is, according to Dreyfus, close to impossible (Matthews, 1999). A second famous doubter is John Searle, who, with his Chinese Room analogy, has responded directly to Turing (cited in Goodwins, 2001):

'Take a room with two slots in the wall, an English-speaking man inside and a rulebook. The rulebook tells him how to deal with Chinese sentences that are pushed through the slot – how to choose characters with which to reply, and what order to send them back out through the second slot. The responses may be perfect Chinese, but it does not logically follow on that the man is actually understanding the language as a native speaker would, rather than merely processing it.'

Although a number of convincing rebuttals to the kinds of philosophical arguments presented above exist, there can be no doubt that such positions present intellectually powerful barriers to the ultimate goal of AI research. Following on from this, it might appear that opinion in this area is neatly polarised. However, the picture is significantly complicated by the fact that many researchers consider strong AI as neither particularly likely nor even desirable. In fact, many of the present obstacles to

strong AI research are far more mundane, having been developed as a result of new scientific interest in the mechanisms of the brain and the way they learn, evolve and develop intelligence from a sense of being conscious. To begin with, although computers are certainly becoming faster, such achievements do not necessarily correspond with computers becoming more intelligent. For, as described by Jaron Lanier (cited in Ho, 2002c) as the ‘great shame’ of computer science, Moore’s law in hardware development must be starkly contrasted with the fact that computer engineers do not seem to be able to write software much better as computers get more advanced.

So far, this report has largely focused on ways in which scientists model part of what we know about our capabilities as sentient beings, rather than attempting to provide true sentience. However, even if the ability to programme software advances rapidly within the next few decades, it seems likely that the AI laboratories of the day will be incapable of providing the kind of environment necessary for generating anything resembling well-rounded intelligence. This idea stems largely from the work of Rodney Brooks of the MIT who has worked hard in recent years to challenge prevailing attitudes towards AI research⁸. Humphrys (1997) builds on these ideas by asserting that you can’t expect to build a single, isolated AI alone in a laboratory and expect to simulate much intelligence. This is because, unless AIs are provided with space in which to evolve a rich culture, with repeated social interaction with things that are like them, you cannot really expect to get beyond a certain stage.

In addition to software development, significant challenges also exist in the development of more artificially intelligent robots. For example, while computer vision is good at certain tasks, there also are many things it is not particularly good at, such as general object recognition. According to

Brooks (2002), computer vision systems can do a few things with great skill, but still after 40 years of effort they are not good at the things humans and many animals do effortlessly. Secondly, robots lack the dexterity of the human hand, a primary ingredient in the types of manufacturing that have moved to low-cost locations. According to Brooks, ‘*low-cost dextrous manipulation*’ is essential if progress is to be made. At present, however, even high-cost dextrous manipulation is beyond researchers. Furthermore, such challenges are unlikely to be met in the next few years, possibly requiring 30–40 years before such technologies are refined.

3.4.3 A future for strong AI?

In spite of the many fundamental barriers highlighted above, the fields of AI and robotics are replete with many wonderfully inventive predictions, a domain where reality and science fiction often meet. Indeed, it is likely that in the next two decades ‘*we’ll see more and better capabilities that we tend to attribute as awareness*’ (Hendler, 2000). However, it is unlikely that machines will ever have human awareness in the philosophical sense of the term, although they may come close in the long term. Rather, we can expect to see classical AI going on to produce more and more sophisticated applications in restricted domains, such as expert systems, chess programs and Internet agents. At the same time, the next 30 years will produce new types of animal-inspired machines that are more ‘messy’ and unpredictable than any we have seen before – less rationally intelligent but more rounded and whole (Humphrys, 1997).

One potentially far-reaching development involves side-stepping the seemingly polarised weak/strong AI debate through the development of cyborg technology, the applications of which could lead to humans having certain physiological processes aided or controlled by mechanical or electronic

devices. The most high-profile demonstration in this area concerns ‘robo-rat’, which, through the implantation of electrodes into the parts of the brain responsible for sensing reward and for stimulation from the left and right whiskers, has been successfully guided by a human controller (Graham-Rowe, 2002). A similar experiment has also been demonstrated by Steve Potter, Professor of Biomedical Engineering at the Georgia Institute of Technology, who has developed a ‘rat-controlled robot’ (Cameron, 2002). This device results from placing a droplet of solution containing thousands of rat neuron cells onto a silicon chip and then relaying the resulting electrical activity to a robot. The robot then manifests these signals with physical motion, each of its movements a direct result of neurons communicating with neurons. Such examples of merging computer chips with living tissue may seem crude, but are described by scientists as ‘momentous’ – an event comparable to the first organ transplant or cloned animal (Philipson, 2001). This is because such experiments open up the possibility of using computer technology to supplement human intelligence, rather than replace it.

In conclusion, then, we will not see full AI in our lives. The reason is that there is no obvious way of getting from here to there – from the rather useless robots and brittle software programs in existence nowadays to human-level intelligence. A long series of conceptual breakthroughs are needed, and this kind of thinking is very difficult to timetable.

3.5 Concerns

3.5.1 Introduction

The fields of strong AI and robotics are generally regarded as controversial because of their far-reaching social, ethical, and philosophical implications. Research managers are in no doubt that such controversy has affected the funding environment for AI and the objectives of

many research programmes (NRC, 1999). However, in general, less attention is paid to the implications of weak AI, even though many of the applications of this field, as demonstrated above, are in operation today. In other words, it should be recognised that many of the concerns described below do not rely on the long-term development of strong AI as popularly imagined. As for Section 2.5 on nanotechnology then, this section, as well as considering the connotations of AI, will attempt to distinguish between short- and long-term concerns that advancements in this area will surely bring.

3.5.2 Predictive intelligence

According to Kirsner (2002), the technology world’s big debate for 2003 will centre on predictive intelligence. This aspect of AI, already touched upon above, concerns the ability to use software running on powerful computers to analyse information about one’s prior behaviour. In the private sector, companies are already using predictive intelligence to analyse data profiles and solve more mundane business problems. These include Epsilon – a database marketing company based in the US, which have been combing through transactional data since the 1980s to help its customers market more effectively – along with other projects designed to identify which customers are more likely to spend the most money (Kirsner, 2002).

The most dramatic example of this is provided by the US DoD, which has established a research group to develop technology for information gathering and analysis on a huge scale. Its goal is to mine data sources all over the world – including government and commercial stores of personal information – to look for terrorists and terrorist threats (Anthes, 2002). This programme includes the recently-established controversial Total Information Awareness (TIA) office which aims to ‘revolutionise the ability of the US to detect, classify and identify foreign terrorists, decipher their plans, and take timely action to

pre-empt and defeat terrorist acts.’ The tools which the TIA intends to develop to achieve this rely to a large extent on new AI technologies. These include ‘entity extraction from natural language text’ and ‘biologically inspired algorithms for agent control.’ Furthermore, one of the TIA’s 13 subdivisions, the Human Identification at a Distance (HumanID) programme, is releasing contracts for face, iris and gait recognition. Another of the subdivisions, FutureMap, will concentrate on market-based techniques for avoiding surprise and predicting future events (Hertzberg, 2002).

A second programme, called Evidence Extraction and Link Discovery (EELD), aims to develop technology for *‘automated discovery, extraction and linking of sparse evidence contained in large amounts of classified and unclassified data sources’* (Anthes, 2002). In order to achieve this, EELD will have to develop detection capabilities to extract relevant data and relationships about people, organisations and activities from huge volumes of data.

Apart from the sheer ambitiousness of the programmes, TIA and EELD have generated concern mainly in relation to their implications for infringing individual and group privacy, and the possibility of such information being handled carelessly or even leading to malevolence. Indeed, it only takes a moment of reflection to consider that nearly everyone in modern society has at least one fact about themselves to hide. And yet, in spite of these well-founded concerns, both the TIA and EELD are already in active development; in response, Hertzberg (2002) recommends that, at a minimum, a temporary shutdown of the EELD system pending some sort of congressional review and the creation of safeguards is highly desirable.

3.5.3 AI and robotic autonomy

Many of the major ethical issues surrounding AI-related development hinge upon the

potential for software and robot autonomy. In the short term, some commentators question whether people will really want to cede control over our affairs to an artificially intelligent piece of software, which might even have its own legal powers. Broersma (2001) believes that, while some autonomy is beneficial, absolute autonomy is frightening. For one thing, it is clear that legal systems are not yet prepared for high autonomy systems, even in scenarios that are relatively simple to envisage, such as the possession of personal information. In the longer-term, however, in which it is possible to envisage extremely advanced applications of hard AI, serious questions arise concerning military conflict, and robot ‘take-overs’ and machine rights. Each of these is dealt with in turn below.

3.5.3.1 AI and military conflict

This report shows that the military interest in AI is significant. However, as pointed out above, the difficulties involved in achieving anything resembling hard AI surely mean that any such system will be subject to reliability concerns. This idea is not new; the issue is picked up by Thompson as early as 1977, who sets out his concerns regarding existing and planned uses of computer technology as part of nuclear weapons systems. More generally, it is his belief that no computer system has the capacity to reliably make decisions of the required kind and in the required circumstances, nor can one ever be constructed. This is because the complexity and sensitivity of such systems makes exhaustive characterisation extremely difficult, and any resulting mistakes cannot be corrected via the usual process of use, failure and modification. More recently, the controversial US National Missile Defence programme, which is being designed using the latest AI technology, provides a second example. The system is supposed to dispense ‘kill power’ based on an ability to recognise incoming missiles in a matter of seconds and then decide whether to destroy, intercept or ignore them (Newquist, 1987). However, serious concerns are already

being voiced based upon the workability of such a system. This is because, while testing may be possible for an autonomous tank and other weapons of the electronic battlefield, it is not feasible for National Missile Defence. Such a system can only be realistically evaluated in actual combat (Augarten, 1986). More fundamentally, significant moral difficulties arise out of human distaste for autonomous weapons. Gary Chapman (2000) summarises this concern well:

'[Such arms] are a revolution in warfare in that they will be the first machines given the responsibility for killing human beings without human direction or supervision. To make this more accurate, these weapons will be the first killing machines that are actually predatory, that are designed to hunt human beings and destroy them.'

Indeed, the UCAV example provided above demonstrates that potentially, in battle, humans may be taken out of the decision-making loop and still be on the receiving end – where the 'kill power' goes.

3.5.3.2 Robot 'take-over' and machine rights

Such issues of predatory machines are bound to raise concern over the scenario of AIs overtaking humankind and thus somehow competing with him. This idea has often been popularised by classic science fiction works and populist academics, such as Professor Kevin Warwick, Professor of Cybernetics at the University of Reading, UK, who has repeated this beliefs concerning robot 'take-over' on many occasions in the press, in his books, and on television and radio. Consider the following letter from Nicholas Albery (1999) of the Institute of Social Inventions. Published in *New Scientist* and entitled *Robot Terror*, Albery seeks support for the following petition:

'In view of the likelihood that early in the next millennium computers and robots will be developed with a capacity and complexity

greater than that of the human brain, and with the potential to act malevolently towards humans, we, the undersigned, call on politicians and scientific associations to establish an international commission to monitor and control the development of artificial intelligence systems.'

It is this kind of claim that seems to infuriate many in the AI scientific community. Chris Malcolm (2001) of the School of Artificial Intelligence at Edinburgh University, for example, describes belief in the robot take-over scenario as 'dangerous' and 'misleading'. He points out that public overreaction to AI stems from an assumption that something which displays some of the attributes of creaturehood must possess all the attributes of creaturehood. In his words:

'Intelligence is no more enough to make a real creature than is fur and beady eyes. No matter how much intelligence is added to your word processor it is not going to sulk and refuse to edit any more letters if you don't improve your spelling...Our problem is that while we have got used to the idea that teddy bears are not real even though we may be in the habit of talking to them at length, we are not used to contraptions being intelligent enough to talk back, and are willing to credit them with possession of the full orchestra of creaturehood on hearing a few flute-like notes.'

Perhaps the most measured assessment of the possibility of tyrannical take-over to date stems from the work of Whitby and Oliver (2001), who, in addition, to the classic worst case scenario, focus on the more subtle ideas of 'cultural reliance' and 'co-evolution'. With regard to the former, the authors conclude that: *'although not obviously misguided or incoherent, predictions of tyrannical take-over are wrong. This is due to a number of possible failsafe methods, such as buddy systems, ethical systems programming, and perhaps most importantly, humans as final*

arbitrators in decision making.' In any case, it is not clear in the first place why intelligence should necessarily be regarded as synonymous with aggression. On the other hand, cultural reliance, in which humans somehow allow a position of dependency on AI and robotics to develop, and co-evolution, in which human and machine become inextricably intertwined, are regarded as more probable.

The strong public reaction to machine take-over appears, then, not to be well founded. However, if it is possible to agree, for argument's sake, that humankind will be able to create a truly intelligent machine, a much deeper issue arises: how will a sentient artificial being be received by humankind and by society? Barry (2001) asks pertinent questions: *'Would it be forced to exist like its automaton predecessors who have effectively been our slaves, or would it enjoy the same rights as the humans who created it, simply because of its intellect?'* This is an enormous question that touches religion, politics and law, but to date little serious discussion has been given to the possibility of a new intelligent species and to the rights an autonomous sentient might claim.

3.6 Discussion

The short-term concerns surrounding AI and robotics are mainly ethical in nature. This is in contrast to nanotechnology, the potential dangers of which cover a much larger spectrum and one that includes environmental risk. As shown above, weak AI tends to create concern with respect to its role as a tool for human interaction, throwing up issues of responsibility, privacy and trust. Applications in this area are emerging all the time, making 2003 the right time to begin public debate over these concerns. This is important for three main reasons. First, there might be a tendency for AI technology to creep into our lives largely unnoticed. This is because of the well-documented AI effect, due to which the

major applications of AI research are mostly hidden from view because they are embedded in larger software systems. Second, many of these applications are morally ambiguous – a grey area of ethics that stands in stark contrast to Isaac Asimov's famously clear-cut three laws of robotics⁹. Third, presuming that a public debate over AI can be initiated, there is little evidence to date that this discussion will affect military and commercial interests. Having said that, there is evidence of some attempt to flesh out a code of professionalism for AI. For example, in reference to AI and responsibility, Whitby (1984) writes:

'Where an AI system is introduced into any human system it shall be the responsibility of the AI professional to ensure that a human or group of humans within the system shall take moral and/or legal responsibility for the human consequences of any malfunction of the AI system.'

However, there is little sign in the literature that suggests these ideas have been followed up on.

Strong AI, on the other hand, asks much more fundamental questions as the field necessarily deals with human/machine relationships per se. As a consequence, the kinds of tools that might be necessary to begin debate over strong AI are not even here yet, so great are the implications. However, it is likely that this technology will not occur in our lifetimes; regardless of how often Professor Warwick is presented as an AI expert, the fact remains that his opinions are far removed from the majority view of the AI community (Colton, 2001). On the other hand, this report is by no means intended to downplay such potentially revolutionary developments as 'mere' science fiction. For, if the long-term potential of AI was to be realised, then it would surely have a demonstrable impact in a whole range of industrial and, in particular, service sectors.

4. Conclusion

This report began by stressing the need to provide background information on nanotechnology and AI. In doing so, it was hoped that the prospects of these emerging technologies to affect quality of life in the coming decades could be realistically assessed. One consequence of providing such an overview is that there can be no decisive conclusions as such; the industries characterised here are too dynamic and uncertain to generate any real sense of resolution. However, it is possible to highlight a number of important differences and similarities between nanotechnology and AI which go some way to shedding more light on their character.

Perhaps the greatest contrast between the two industries concerns public interest. Indeed, as this report has demonstrated, nanotechnology is widely regarded as a ‘new’ and exciting branch of science and technology. This belief has contributed to the massive period of growth that this high-profile and wide-ranging field is currently enjoying. AI, on the other hand, is viewed by many as an highly specialised and unproven discipline. One reason for this concerns the gross over-optimism that characterised the industry in the 1960s and 1980s. Another reason reflects the AI community’s seemingly insurmountable difficulty in publicising its own achievements without whipping up general anxiety over machine superiority. The upshot of all this has been the field’s struggle to attract funding in the past and it is likely that this trend will continue for sometime into the foreseeable future.

Revealing similarities also exist between nanotechnology and AI. There has been much talk recently regarding the convergence of traditionally separate scientific fields, in particular the blurring of the boundaries between the physical sciences and life sciences – perhaps even the first step towards the long sought after unification of physics, chemistry and biology (Howard, 2002). For

example, the concourse of nanoscience, biotechnology, IT, and cognitive science (‘NBIC’) was discussed during a December 2001 NSF workshop. NBIC, it was agreed ‘*could achieve a tremendous improvement in human abilities, societal outcomes, the nation’s productivity and the quality of life*’ (Roco and Bainbridge, 2003). In some ways, the above conclusion is hardly surprising given the ambitious and broad scope of the technologies discussed in this report. As pointed out above, ‘convergence’ largely arises from the wide availability of techniques and tools on offer today – the real innovation stems from the process of bringing individuals from traditionally separate disciplines together.

Most importantly for convergence here, it is possible that developments in nanotechnology could lead to advances in AI through improvements in computer miniaturisation, performance, or architecture (but see Section 3.4.2 on barriers to strong AI), or through the sensor interface. In addition, it seems fair to assume that any futuristic nanobots would have to be imbued with a reasonable degree of AI. A second, more contentious similarity concerns reinvention. As demonstrated in this report, the ‘rediscovery’ of AI has been a virtual necessity for the survival of the industry; for nanotechnology the phenomena is less obvious but is arguably there all the same. That is, as a natural extension of the micromechanical and MEMS research that begun in the 1960s, nanoscience is hardly ‘new’ as such; rather, ‘nano’ can be viewed as a useful tag with which to boost funding. Just what the consequences of this strategy will be, it is hard to tell. Ironically, AI provides an excellent example of a promising scientific discipline that has often resulted in disappointment. Whether the same happens to nanotechnology remains to be seen.

The second consequence of providing an overview is that certain elements of

nanotechnology and AI development are bound to be overlooked. First, the difficulties of drawing out accurate statistics for corporate R&D have already been alluded to earlier. Second, there are wide ranging applications across the economy for sensors that can support industrial processes and be incorporated into new or existing products (Miles and Jarvis, 2001). The application of nanotechnology to this area should allow for improvements in functionality and much decreased size. Third, a more in-depth analysis of environmental concerns is warranted. This is because public acceptability of such risk is likely to vary considerably in relation to the application being considered. For example, the application of nanotechnology to computerisation is less likely to cause concern than those practices which might lead to the release of nanoparticles into the environment, such as the disposal of nano-based composites. Fourth, it is possible to conceive of a number of environmental goods that may arise. For example, the potential for gains in energy generation and efficiency have already pointed out above (Section 2.3.4.), and it is conceivable that dramatic improvements in environmental sensing and modelling could also be achieved. However, any pervasive diffusion of nano- and AI-based technologies in the coming decades is bound to have a significant effect on the demand for resources

by industry, transport and the domestic sector. The way in which these more fundamental changes might impact on the environment would have to form the basis of a much larger technology assessment, in which long-term structural changes to global industry and commerce were considered.

Finally, it is easy to overlook the lessons that attitudes towards technological development teach us about human nature. This report has largely relied upon the technique of looking ahead, identifying technological possibilities, and assessing the likelihood of successfully moving towards their realisation. Significantly, this process mirrors that of technological innovators, a kind of thinking that often translates into the belief that technological development is autonomous – the ultimate self-fulfilling prophecy. To some extent we are already on this road. Most technologies covered in this report are within the bounds of current scientific possibility and it is just a matter of time, effort and expenditure before they are realised. However, the contrasting fortunes of the nanotechnology and AI industries remind us that much of this progress hinges on public approval. Ultimately, a 21st-Century acceptance model calls for technological innovations to be received on a voluntary basis where the perceptible usefulness of new technology products are balanced against associated risks that are shown to be manageable.

¹ Grove-White, R., Macnaghton, P. and Wynne B. (2000). Wising Up: The Public and New Technologies, Lancaster, UK: IEPPP, Lancaster University.

² It is worth bearing in mind when consulting this type of information that, given the difficulty of even agreeing on what constitutes nanotechnology, many of the numbers presented below should be treated with caution (Roman, 2002).

³ For a detailed breakdown, see the NNI's own report at:
<http://www.nano.gov/2003budget.html>

⁴ For a detailed description of the history of AI, see the University of Edinburgh's Division of Informatics Website at:
http://www.dai.ed.ac.uk/AI_at_Edinburgh_perspective.html

⁵ For more information, see Ruthfield, 1995.

⁶ See Honda's Website for more information at: <http://world.honda.com/robot/>

⁷ For details, see National Research Council, 1999.

⁸ For a full description of this paradigm shift, see Humphrys, 1997.

⁹ The Three Laws of Robotics are:

- a. A robot may not injure a human being, or, through inaction, allow a human being to come to harm;
- b. A robot must obey the orders given to it by human beings except where such orders would conflict with the First Law;
- c. A robot must protect its own existence as long as such protection does not conflict with the First or Second Law.

For a fuller explanation see:
http://whatis.techtarget.com/definition/0,,sid9_gci520366,00.html

- Albery, N. (1999). Robot Terror. *New Scientist*; 161 (2171): 49. [letter]
- Aleksander, I. (2002). The Next Big Thing. *BBCi & The Open University* [online]. http://www.open2.net/nextbigthing/ai/hear_the_arguments/argument3.htm
- Anderson, K. (2001). Predicting AI's Future. *BBC News*; 21 Sep 2001 [online]. http://news.bbc.co.uk/1/hi/in_dept_h/sci_tech/2001/artificial_intelligence/1555742.stm
- Antón, P.S. Silbergliitt, R. and Schneider, J. (2001). *The Global Technology Revolution: Bio/Nano/Materials Trends and Their Synergies with Information Technology by 2015*. Santa Monica, CA: RAND. Document Number: MR-1307-NIC.<http://www.rand.org/publications/MR/MR1307/>
- Anthes, G.H. (2002). Global Surveillance: The Government's Plan. *Computerworld*; Security: Privacy, 25 Nov 2002 [online]. <http://www.computerworld.com/securitytopics/security/privacy/story/0,10801,76117,00.html>
- Augarten, S. (1986). Can Computers be Failsafe? *PC Magazine*; 5 (1): 97.
- Barry, R. (2001). Sentience: The Next Moral Dilemma. *ZDNet UK*; News: 24 Jan 2001 [online]. <http://news.zdnet.co.uk/story/0,,s2083942,00.html>
- .BBC (2002). Computers Reach One Million Mark. *BBC News*; 1 Jul 2002 [online]. <http://news.bbc.co.uk/1/hi/sci/tech/2077986.stm>
- Boeing (2002). Unmanned Combat Air Vehicle (X-45). *Boeing* [online]. <http://www.boeing.com/phantom/ucav.html>
- Broersma, M. (2001). AI Gets Down to Business. *ZDNet UK*; News: 23 Jan 2001 [online]. <http://news.zdnet.co.uk/story/0,,s2083916,00.html>
- Brooks, R. (2001). Interviewed in: Anderson, K. Predicting AI's future. *BBC News*; 21 Sep 2001 [online]. http://news.bbc.co.uk/1/hi/in_dept_h/sci_tech/2001/artificial_intelligence/1555742.stm
- Brooks, R. (2002). Interviewed in: Jusko, J. The Robot Evolution. *Industry Week*; 1 Nov 2002 [online]. <http://www.industryweek.com/CurrentArticles/Asp/articles.asp?ArticleId=1356>
- Buchanan, B. and Uthurusamy, S. (1999). The Innovative Applications of Artificial Intelligence Conference. *AI Magazine*; 20 (1): 11–12.
- Cameron, D. (2002). Rat-Brained Robot. *Technology Review*; 18 Dec 2002 [online]. http://www.technologyreview.com/articles/print_version/wo_cameron121802.asp
- Chapman, G. (2000). Quoted in: Matthews, J. (2000b). Military Applications of AI. *Generation5 Essays* [online]. http://www.generation5.org/app_military.shtml
- Chaudhari, P. (2001). Future Implications of Nanoscale Science and Technology: Wired Human, Quantum Legos, and an Ocean of Information. In: M.C. Roco, W.S. Bainbridge (editors). *Societal Implications of Nanoscience and Nanotechnology*. Arlington, VA, USA: National Science Foundation; pp75–9. <http://www.wtec.org/loyola/nano/NSET.Societal.Implications/nanosipdf>
- Colton, S. (2001). Cyborg Off His Christmas Tree. *Times Higher Education Supplement*; 21 Dec 2001: p14. <http://www.aisb.org.uk/articles/computology.html>
- Colvin, V. (2002). Responsible Nanotechnology: Looking Beyond the Good News. *EurekAlert!*; Nanotechnology In Context: Nov 2002 [online]. <http://www.eurekalert.org/context.php?context=nano&cshow=essays&essaydate=1102>
- Compano, R. editor (2001). *Technology Roadmap for Nanoelectronics*. Luxembourg: Office for Official Publications of the European Communities.

- Compano, R. and Hullman, A. (2002). Forecasting the Development of Nanotechnology with the Help of Science and Technology Indicators. *Nanotechnology*; 13 (3): 243–7.
- Copeland, B.J. (2000). What is Artificial Intelligence? *AlanTuring.net*; Reference Articles: May 2000 [online]. http://www.alanturing.net/turing_archive/pages/Reference%20Articles/what_is_AI/What%20is%20AI02.html
- Cowen, R. (2002). Ribbon to the Stars. *Science News*; 162 (14): 218. <http://www.sciencenews.org/20021005/bob9.asp>
- Dalesio, E.P. (2002). ‘Text mining’ Software Business Grows. *Sun-Sentinel*; 4 Mar 2002 [online]. <http://www.sun-sentinel.com/business/local/sfl-0304textmining.story?coll=sfla-business-headlines>
- Department of Defense (2002). Nanotechnology Research Awards Announced. *DefenseLINK*; News: 23 Feb 2001 [online]. http://www.defenselink.mil/news/Feb2001/b02232001_br079-01.html
- Department of Trade and Industry (2002). *New Dimensions for Manufacturing: UK Strategy for Nanotechnology*. Report of the UK Advisory Group on Nanotechnology Applications; June 2002. UK: Department of Trade and Industry. <http://www.dti.gov.uk/innovation/nanotechnologyreport.pdf>
- Dertouzos, M.L. (1999). The Future of Computing. *Scientific American*; 281 (2): 52–5. <http://www.sciam.com/issue.cfm?suedate=Aug-99>
- Disappearing Computer, The (2003). Mission Statement. *The Disappearing Computer* [online]. <http://www.disappearing-computer.net/>
- Doering, R. (2001). Societal Implications of Scaling to Nanoelectronics. In: M.C. Roco, W.S.Bainbridge (editors). *Societal Implications of Nanoscience and Nanotechnology*. Arlington, VA, USA: National Science Foundation; pp68–75. <http://www.wtec.org/loyola/nano/NSET.Societal.Implications/nanosi.pdf>
- Doyle, J. and Dean, T. (1996). Strategic Directions in Artificial Intelligence. *ACM Computing Surveys*; 28 (4): 653–70. <http://medg.lcs.mit.edu/ftp/doyle/sdai96.html>
- Dreyfus, H. (1992). *What Computers Still Can’t Do: A Critique of Artificial Reason*. Cambridge, MA, USA: Massachusetts Institute of Technology.
- Economist, The (2001). Robots. *The Economist.com*; 29 Nov 2001 [online]. http://www.economist.com/printedition/displayStory.cfm?Story_ID=887290
- Economist, The (2002). Trouble in Nanoland. *The Economist.com*; 5 Dec 2002 [online]. http://www.economist.com/science/displaystory.cfm?story_id=1477445
- Engineering and Physical Sciences Research Council (2003). *EPSRC Web Site* [online]. <http://www.epsrc.ac.uk>
- ETC Group (2002a). No Small Matter! Nanotech Particles Penetrate Living Cells and Accumulate in Animal Organs. *ETC Group Communique*; (76) May/June 2002 [online]. http://www.etcgroup.org/documents/Comm_NanoMat_July02.pdf
- ETC Group (2002b). Nanotech Takes a Giant Step Down! *ETC Group News Release*; 6 Mar 2002 [online]. <http://www.etcgroup.org/documents/nr2002march6.pdf>
- ETC Group (2002c). Patenting Elements of Nature. *ETC Group Geno-type*; 25 Mar 2002 [online]. <http://www.etcgroup.org/documents/nanopatentsgeno.rtf.pdf>
- European Environment Agency (2003). EEA Multilingual Environmental Glossary. *European Environment Agency* [online]. http://glossary.eea.eu.int/EEAGlossary/P/precautionary_approach
- Fishbine, G. (2002). *Investor’s Guide to Nanotech and Micromachines*. Chichester, UK: John Wiley.

- FOLDOC (2003). *Free On-Line Dictionary of Computing* [Online].
<http://wombat.doc.ic.ac.uk/foldoc/>
- Forrest, D. (1989). Regulating Nanotechnology Development. *Foresight Institute*; 23 Mar 1989 [online].
<http://www.foresight.org/NanoRev/Forrest1989.html>
- Freitas, R.A. (2000). Some Limits to Global Ecophagy by Biovorous Nanoreplicators, with Public Policy Recommendations. *Foresight Institute*; Mar 2000 [online].
<http://www.foresight.org/NanoRev/Ecophagy.html>
- Fried, J. (2002). Japan Sees Nanotech as Key to Rebuilding its Economy. *Small Times*; 7 Jan 2002 [online].
http://www.smalltimes.com/document_display.cfm?document_id=2843
- Glinos, K. (1999). Nanotechnology: Blurring the Boundaries at the Atomic Scale. *RTD info*; (21) Feb 1999 [Online].
<http://www.europa.eu.int/comm/research/rtdinf21/en/dossier1.html>
- Goodwins, R. (2001). The Machine that Wanted to be a Mind. *ZDNet UK*; News: 23 Jan 2001 [online].
<http://news.zdnet.co.uk/story/0,,s2083911,00.html>
- Gorman, J. (2002). Taming High-Tech Particles: Cautious Steps into the Nanotech Future. *Science News*; 161 (13): 200.
<http://www.sciencenews.org/20020330/bob8.asp>
- Graham-Rowe, D. (2002). “Robo-rat” Controlled by Brain Electrodes. *New Scientist*; 417: 37.
<http://www.wireheading.com/robobots/ratbot.html>
- Grosz, B. and Davis, R. editors (1994). *A Report to ARPA on Twenty-First Century Intelligent Systems*. Menlo Park, CA, USA: American Association for Artificial Intelligence.
<http://www.aaai.org/Resources/Policy/arpa-report.html>
- Grupp, H. editor (1993). *Technologie am Beginn des 21. Jahrhunderts*. Heidelberg: Physica Verlag. Cited in: Compano, R. and Hullman, A. (2002).
- Gsponer, A. (2002). From the Lab to the Battlefield? Nanotechnology and Fourth-Generation Nuclear Weapons. *Disarmament Diplomacy*; (67) Oct–Nov 2002 [online].
<http://www.acronym.org.uk/dd/dd67/67op1.htm>
- Harper, T. (2002). The Nanotechnology Arms Race: Why Nobody Wants to be Left Behind. *nanotechweb.org*; 14 Nov 2002 [online].
<http://nanotechweb.org/articles/column/1/11/1/1>
- Hay, J.N. and Shaw, S.J. (2000). A Review of Nanocomposites 2000. *Proceedings of the Nanocomposites 2000 Conference*; 6–7 Nov 2000, Brussels, Belgium. Farnborough, UK: Defence Evaluation and Research Agency.
http://www.nano.org.uk/nanocomposites_review.pdf
- Helsel, S. (2002). Are Corporate and Government Funding Overlooked? *NanoelectronicsPlanet*; 26 Jul 2002. [online].
http://www.nanoelectronicsplanet.com/features/article/0,4028,6571_1433691,00.html
- Herrera, S. (2002). No small matter. *Red Herring*; 16 Dec 2002 [online].
<http://www.redherring.com/insider/2002/12/10trends-nanotech121602.html>
- Hertzberg, H. (2002). Too Much Information. *The New Yorker*; 9 Dec 2002. [online]
http://www.newyorker.com/talk/content/?021209ta_talk_hertzberg
- Hendler, J. (2000). A Chat About the Future of Artificial Intelligence. *CNN*; 1 Jan 2000 [online].
<http://www.cnn.com/COMMUNITY/transcripts/1999/12/hendler/>
- Ho, M.W. (2002a). Nanotechnology: A Hard Pill to Swallow. *Science in Society*; (16) Autumn 2002 [online].
<http://www.i-sis.org.uk/nanotechnology.php>

H – M

- Ho, M.W. (2002b). The Brave New World Quartet. *Science in Society*; (16) Autumn 2002 [online]. <http://www.i-sis.org.uk/bravenewworld.php>
- Ho, M. (2002c). Will Computers Become Super-Human? *Science in Society*; (16) Autumn 2002 [online]. <http://www.i-sis.org.uk/computersvshumans.php>
- Holister, P. (2002). Nanotech: The Tiny Revolution. *CMP Cientifica*; July 2002 [online]. http://www.cientifica.info/html/docs/NOR_White_Paper.pdf
- Howard, S. (2002). Nanotechnology and Mass Destruction: The Need for an Inner Space Treaty. *Disarmament Diplomacy*; (65) Jul–Aug 2002 [online]. <http://www.acronym.org.uk/dd/dd/65/65op1.htm>
- Hsiung, S. (2002) An Introduction to Artificial Intelligence. *Generation5* [online]. <http://www.generation5.org/aiintro.shtml>
- Humphrys, M. (1997). AI is Possible...But it Won't Happen: The Future of Artificial Intelligence. *Proceedings of the Next Generation Symposium*; Aug 1997, Jesus College, Cambridge, UK. <http://www.compapp.dcu.ie/~humphrys/newsci.html>
- IT Strategic Headquarters (2001). *E-Japan Strategy*. Policy Initiative, 22 Jan 2001 Tokyo, Japan: IT Policy Office. http://www.kantei.go.jp/foreign/it/network/0122full_e.html
- Jeremiah, D.E. (1995). Nanotechnology and Global Security. *Proceedings of the Fourth Foresight Conference on Molecular Nanotechnology*; 9–11 Nov 1995, Palo Alto, CA, USA. [online]. <http://www.zyvex.com/nanotech/nano4/jeremiahPaper.html>
- Joseph, S. (2001). AI. *J@pan Inc Magazine*; Nov 2001 [online]. <http://www.japaninc.com/article.php?articleID=514>
- Khan, J. (2002). It's Alive! *Wired*; 10 (3) [online]. <http://www.wired.com/wired/archive/10.03/everywhere.html>
- Kirsner, S. (2002). Getting Smart About Predictive Intelligence. *Boston Globe*; 30 Dec 2002 [online]. <http://www.boston.com/globe>
- Krill, P. (2002). CTO Forum: Will Artificial Intelligence Surpass the Human Variety? *InfoWorld*; 11 Apr 2002 [online]. <http://www.infoworld.com/articles/hn/xml/02/04/11/020411hnartificial.xml>
- LaVan, D.A. and Langer, R. (2001). Implications of Nanotechnology in the Pharmaceuticals and Medical Fields. In: M.C. Roco and W.S. Bainbridge (editors). *Societal Implications of Nanoscience and Nanotechnology*. Arlington, VA, USA: National Science Foundation; pp79–83. <http://www.wtec.org/loyola/nano/NSET.Societal.Implications/nanosi.pdf>
- Leo, A. (2001). Get Ready for Your Nano Future. *Technology Review*; May 2001 [online]. http://www.technologyreview.com/articles/wo_leo050401.asp
- Lighthill, J. (1973). Artificial Intelligence: A General Survey. In: *Artificial Intelligence: A Paper Symposium*. London, UK: Science Research Council.
- Machan, D. (2002). A Few Good Toys. *Forbes Magazine*; (12) 9 Dec 2002 [online]. <http://www.forbes.com/forbes/2002/1209/120.html>
- Maes, P. (1994). Agents That Reduce Work and Information Overload. *Communications of the ACM*; 37 (7): 31–40. <http://pattie.www.media.mit.edu/people/pattie/CACM-94/CACM-94.p1.html>
- Malcolm, C. (2001). Robots Won't Rule. *School of Artificial Intelligence, Division of Informatics, Edinburgh University*; 3 Aug 2001 [online]. http://www.dai.ed.ac.uk/homes/cam/Robots_Wont_Rule.shtml
- Martel, S. (2000). Introduction to the NanoWalker: A Miniature Autonomous Robot Capable of Various Tasks at the Molecular and Atomic Scales. *Bio-Instrumentation Laboratory*; Projects [online]. <http://bioinstrumentation.mit.edu>

- Matthews, J. (1999). Philosophical Arguments For and Against AI. *Generation5*; 13 Dec 1999 [online].
http://www.generation5.org/ai_phil.html
- Matthews, J. (2000a). Introduction to Neural Networks. *Generation5*; 20 Mar 2000 [online].
<http://www.generation5.org/nnintro.shtml>
- Matthews, J. (2000b). Military Applications of AI. *Generation5*; 8 May 2000 [online].
http://www.generation5.org/app_military.shtml
- McCarthy, J. (2003). What is Artificial Intelligence? *Computer Science Department, Stanford University*; 29 Mar 2003 [online].
<http://www-formal.stanford.edu/jmc/whatisai/whatisai.html>
- McCullagh, D. (2002). Report Calls for Nanotech Laissez-Faire. *CNET News.com*; 21 Nov 2002 [online].
<http://www.news.com.com/2100-1023-966766.html>
- Melymuka, K. (2002). Good Morning, Dave... *Computerworld*; 11 Nov 2002 [online].
<http://www.computerworld.com/softwaretopics/software/appdev/story/0,10801,75728,00.html>
- Miles, I. and Jarvis, D. (2001). *Nanotechnology – A Scenario for Success in 2006*. Teddington, UK: HMSO. National Physical Laboratory Report Number: CBTL16.
http://libsvr.npl.co.uk/npl_web/pdf/cbtlm16.pdf
- Miller, D. (2001). Five Technologies You Need to Know. *The Industry Standard*; 21 May 2001 [online].
<http://www.thestandard.com/article/0,1902,24308,00.html>
- Miller, P. (2002). Big, Big Plans for Very Small Things. *Green Futures*; Jul/Aug 2002: 50–1.
<http://www.paulmiller.org/Nanotechnology.pdf>
- Tam, A. (2001). Japan's Nanotech Players: Thinking Big. *J@pan Inc Magazine*; May 2001 [online].
http://www.japaninc.net/mag/com/p/2001/05/print/may01p_investor_nanotech.html
- Mokhoff, N. (2002). Speech Technology Looses its Kooky Luster. *EETimes*; 1 Nov 2002 [online].
<http://www.eetimes.com/story/OEG20021031S0046>
- nanotechweb.org (2002). Nanocomposites Set to Wrap up the Packaging Market. *nanotechweb.org*; 23 Aug 2002 [online].
<http://nanotechweb.org/articles/news/1/8/18/1>
- National Institute of Informatics (2002). *National Institute of Informatics 2002*. Tokyo, Japan: National Institute of Informatics.
<http://www.nii.ac.jp/publications/nii-yoran/yoran2002-e.pdf>
- National Research Council (1999). *Funding a Revolution: Government Support for Computing Research*. Washington, DC, USA: National Academy Press.
<http://www.nap.edu/readingroom/books/far/notice.html>
- Nelson, M. and Shipbaugh, C. (1995). *The Potential of Nanotechnology for Molecular Manufacturing*. Santa Monica, CA, USA: RAND. Document Number: MR-615-RC.
<http://www.rand.org/publications/MR/MR615/>
- Newquist, H. (1987). Star Wars: When Imperfect Man Strives to Make Perfect Machines. *Computerworld*; 30 Nov 1987: 17. <http://www-cse.stanford.edu/classes/cs201/Projects/autonomous-weapons/articles/starwars.txt>
- New Scientist (2002). US Army Seeks Nanotech Suits. *New Scientist.com*; News Service, 14 Mar 2002 [online].
<http://www.newscientist.com/news/news.jsp?id=ns99992043>
- Pace, S. (1989). Military Implications of Nanotechnology. *Foresight Update* 6; 1 Aug 1989 [online].
<http://www.foresight.org/Updates/Update06/Update06.2.html>

- Pergamit, G. and Peterson, C. (1993). Nanotechnology. In: Ringo, T. editor. *On The Cutting Edge of Technology*. Indianapolis, IN, USA: Sams Publishing.
- Philipson, G. (2001). Human and Computer Merger a Chip of the Old Block. *The Age*; IT News, 18 Sep 2001, p1. For a summary see: <http://www.gyre.org/news/related/Metacomputing/Biological+Computing>
- RAND (2002). Revolutionising Logistics Support by “Thinking Small”. RAND; National Security [online]. http://www.rand.org/natsec_area/products/thinkingsmall.html
- Reynolds, G.H. (2002). *Forward to the Future: Nanotechnology and Regulatory Policy*. San Francisco, CA, USA: Pacific Research Institute. http://www.pacificresearch.org/pub/sab/techno/forward_to_nanotech.pdf
- Roco, M.C., and Bainbridge, W.S., editors (2001). *Societal Implications of Nanoscience and Nanotechnology*. Arlington, VA, USA: National Science Foundation. <http://www.wtec.org/loyola/nano/NSET.Societal.Implications/nanosipdf>
- Roco, M.C., and Bainbridge, W.S., editors (2002). *Converging Technologies for Improving Human Performance. Nanotechnology, Biotechnology, Information Technology and Cognitive Science*. Arlington, VA, USA: National Science Foundation. http://wtec.org/ConvergingTechnologies/Report/NBIC_pre_publication.pdf
- Roman, C. (2002). *It’s Ours to Lose: An Analysis of EU Nanotechnology Funding and the Sixth Framework Programme*. Brussels, Belgium: European NanoBusiness Association. <http://www.nanoeurope.org/docs/European%20Nanotech%20Funding.pdf>
- Roy, R. (2002). Giga Science and Society. *Materials Today*; 5 (12): 72. http://www.materialstoday.com/pdfs_5_12/opinion.pdf
- Ruthfield, S. (1995). The Internet’s History and Development. *AMC Crossroads*; 2 (1) [online]. <http://www.acm.org/crossroads/xrds2-1/inet-history.html>
- Saxl, O. (2000). *Opportunities for Industry in the Application of Nanotechnology*. London, UK: Office of Science and Technology. [A report for The Institute of Nanotechnology, April 2000]. <http://www.nano.org.uk/contents.htm>
- Schneider, F.B. and Rodd, M., editors (2001). *International Review of UK Research in Computer Science*. Swindon, UK: Engineering and Physical Sciences Research Council. http://www.iee.org/policy/csreport/cs_report.pdf
- Schulz, W.G. (2002). Nanotechnology Under the Scope. *Chemical and Engineering News*; 4 Dec 2002 [online]. <http://pubs.acs.org/cen/today/dec4.html>
- Science Blog (2002). New Centre to Transfer Nano Innovations to Defence Industry. *Science Blog*; 10 Dec 2002 [online]. <http://www.scienceblog.com/community/article553.html>
- Shim, R. (2002). Sony’s Ando: PCs to Function Like a Brain. *ZDNet; News*; 5 Dec 2002 [online]. <http://zdnet.com.com/2100-1105-976269.html>
- Smalley, R.E. (2001). Of Chemistry, Love and Nanobots. *Scientific American*; 285: 76–7. <http://www.ruf.rice.edu/~smalley/rick%27s%20publications/SA285-76.pdf>
- Smith, R.H. (1996). Molecular Nanotechnology: Research Funding Sources. *Nanotechnology Magazine*; 2 (6).

- Smith, R.H. (2001). Social, Ethical, and Legal Implications of Nanotechnology. In: Roco, M.C., and Bainbridge, W.S., editors. *Societal Implications of Nanoscience and Nanotechnology*. Arlington, VA, USA: National Science Foundation. pp203–11. <http://www.wtec.org/loyola/nano/NSET.Societal.Implications/nanosi.pdf>
- Spooner, J.G. (2002). Sony Teaches Abio New Tricks. *ZDNet*; News: 10 Oct 2002 [online]. <http://zdnet.com.com/2100-1106-961536.html>
- Stottler Henke (2002). Glossary of AI Terms. *Stottler Henke Associates, Inc.* [online]. http://www.shai.com/ai_general/glossary.htm
- Thibodeau, P. (2002). Senate Scrutinizes US Nanotech Investment. *Computer World*; 18 Sep 2002 [online]. <http://computerworld.com/managementtopics/management/itspending/story/0,10801,74341,00.html>
- Thompson, H. (1977). There Will Always Be Another Moonrise: Computer Technology and Nuclear Weapons. *AISB Quarterly*; (53/54) [online]. <http://www.aisb.org.uk/articles/moonrise.html>
- Trevelyan, J. (1999). Redefining Robotics for the New Millennium. *International Journal of Robotics Research*; 18 (12): 1211–23.
- Turing, A. (1950). Computing Machinery and Intelligence. *Mind*; 9: 433–60. <http://cogprints.ecs.soton.ac.uk/archive/00000499/00/turing.html>
- Weld, D.S., editor. (1995). The Role of Intelligent Systems in the National Information Infrastructure. *American Association for Artificial Intelligence* [online]. <http://www.aaai.org/Resources/Policy/nii.html>
- Whitby, B. (1984) A Code of Professionalism for AI. *AISB Quarterly*; (64) [online]. <http://www.aisb.org.uk/articles/professionalism.html>
- Whitby, B. and Oliver, K. (2001). How to Avoid a Robot Takeover: Political and Ethical Choices in the Design and Introduction of Intelligent Artifacts. *AISB Quarterly*; (104) [online]. http://www.aisb.org.uk/articles/avoid_takeover.html
- White House Fact Sheet (2000). *National Nanotechnology Initiative: Leading to the Next Industrial Revolution*. Washington, DC, USA: The White House. http://clinton4.nara.gov/textonly/WH/New/html/20000121_4.html
- Wolbring, G. (2002). The Silenced Targets. In: *Institute of Science in Society Special Series – Inside Human Genetics and Genomics*; Part II: 28 Jan 2002 [online]. <http://www.i-sis.org.uk/SilencedTargets.php>
- Zyvex (2002). *Zyvex Web Site* [online]. <http://www.zyvex.com/Products/tools.html>