

AI in the UK: Past, Present, and Future

Richard Wheeler (rew@pandora.be) Starlab Research Laboratories Engelandstraat 555, Brussels B-1080 BELGIUM http://www.starlab.org/	Professor Donald Michie (D.Michie@ed.ac.uk) University of New South Wales http://www.aii.ed.ac.uk/~dm/dm.html
---	---

Abstract/Introduction

*Expert Update is pleased to include this feature article from Donald Michie and Richard Wheeler responding to some questions about the past, present, and future of AI technology in the UK and Europe. **Donald Michie** is one of the fathers of artificial intelligence, first as an associate of Alan Turing's at Bletchley Park, later as Professor of Machine Intelligence at Edinburgh, and then as Chief Scientist of the Turing Institute in Glasgow. Professor Michie's publications include four books and about 170 papers in experimental biology, AI and computing. He is Editor in Chief of the Machine Intelligence series (seventeen volumes since 1967); also founder of the SGES and of the registered UK charity, the Human Computer Learning Foundation. He is now Adjunct Professor at the University of New South Wales where his work continues. **Richard Wheeler** has spent time in the trenches of applied AI research and application – first with the World Health Organisation in Geneva, then with the Artificial Intelligence Applications Institute in Edinburgh, and finally at Starlab Research in Brussels – one of the flash-points of new technology and new economy research. As 2001 draws to a close we asked some pointed questions about the history (and possible futures) of artificial intelligence.*

Looking back nearly 50 years to the birth of artificial intelligence research in the UK, what do you think were the major founding achievements?

Professor Michie: The following is a top-of-the-head list. I haven't attempted to continue it beyond 1975 except for a selected bit of commercial relevance with which I was directly familiar. My overview stretches to the point at which I mostly lost touch with detailed UK events, owing to spending so much of my time overseas. I should, however, here add Stephen Muggleton's development of Inductive Logic Programming to which I and others attribute great importance, - see the *New Scientist* on applications to biomolecular theory-discovery, approx. 20th February, a must-read for anyone interested in a current European lead in AI relative to the USA. Here's a list.

In the mid-1950s the Experimental Programming Unit, the Metamathematics Unit at Edinburgh, and the Elcock-Foster group at Aberdeen were established. The Machine Intelligence Workshop series was launched in 1965, and in the same year the world's first AI graduate course (Diploma in Machine Intelligence). Collectively the above UK centres grew within 5 years to be on level terms with the four major US centres. Burstall and Popplestone designed the Pop-2 language and the Multi-Pop time-sharing OS, which came to serve a country-wide network of remote interactive users. Pop-2 incorporated

Burstall's immensely powerful "partial application" of functions, and also their "memoisation" (software cacheing, Michie and Popplestone). Chambers and I developed 'BOXES', the first rule-based re-inforcement learner, of which descendants are in use today, and also the first rule-based "learning by imitation", further developed in 1990 with Michael Bain and Jean Hayes-Michie at the Turing Institute, Glasgow, and subsequently by Claude Sammut, myself and others under the name "behavioural cloning".

By 1970 Barrow and Popplestone had developed a teachable vision system for the FREDDY 1 robot, later greatly refined by Barrow and Burstall for FREDDY 2. The distinguished ex-patriate J.A. Robinson was recruited by Bernard Meltzer as a regular and influential visitor to Edinburgh. Over many subsequent years he had repercussions throughout the UK scene. Elcock and Boyer's Absys and Abset languages prefigured the later development of Logic Programming from Kowalski's (Edinburgh) and Colmerauer's (Marseilles) work. In 1967 Pat Hayes, also in Metamathematics, collaborated with Stanford's John McCarthy, on working leave at the in Edinburgh, to produce what has become a classic of the AI literature: "Some philosophical problems from the standpoint of artificial intelligence" (MI-4, 1969) Around this time Alan Bundy was already studying the mechanization of mathematical reasoning under Meltzer, thus sowing the seeds of much of Edinburgh AI's subsequent intellectual history. In 1968 Gregory and Longuet-Higgins had moved from Cambridge to Edinburgh to join with Meltzer and myself to set up the Department of Machine Intelligence and Perception. This ultimately unstable alliance of Bionics (RLG), Computational Logic (BM), Theoretical (HCL-H) and Experimental Programming (DM) was of consequence and great net benefit for UK developments and beyond.

1970-1975 Max Clowes founded machine perception studies at Sussex. In Edinburgh under Meltzer's direction Boyer and Moore developed efficient programs for automatic deduction based on Robinson's Resolution Principle, and Kowalski developed the logical and mathematical basis of Logic Programming (LP). In the EPU I talked Maarten van Emden into using the current Boyer-Moore algorithm to show LP's practical feasibility by implementing Hoare's Quicksort. The result was, I believe, the first-ever working logic program, marginally beating programs run on Colmerauer's Prolog implementation. Meanwhile Gordon Plotkin's Ph.D. work on inductive reasoning laid a solid and profound base for the entire subsequent development in Britain and overseas of machine learning for those inductive inference tasks requiring the full expressiveness of first order predicate calculus, i.e. what is known today as Inductive Logic Programming (ILP). With logistic direction from Jim Howe and myself and technical direction from Rod Burstall, a collaboration of Bionics and EPU developed the first versatile and inductively instructable assembly robot. It could operate in randomly disposed starting conditions, and formed and successively updated its own stored model of its task environment. In 1974 with the departure of Longuet-Higgins, following Gregory's in 1970, Edinburgh established under Meltzer's Headship the Department of Artificial Intelligence. This launched the world's first undergraduate courses in AI and the first, and some will say still the best-written and presented, university text for the subject (Ambler, Bundy and Burstall). All this occupied ten years, at the end of which significant and substantial AI

centres had yet to be established in any other country in the world outside the USA and Scotland. The UK's thriving centres have since multiplied and consolidated. Today's result is a manifest challenge to US hegemony, conspicuous in quality rather than sheer dollar power.

Richard Wheeler: Of course I cannot speak about the earliest days of AI, but would add a few words here about my experiences in the UK since about 1995. As Professor Michie notes above, the UK is second only to the US in the developing field of AI, and has produced a disproportionate share of founders, great thinkers, and innovators. What continues to impress me about the AI community in the UK is its cohesiveness, depth of focus, and forward-looking and innovative nature, and as the 90s progressed, a great deal of functional consolidation took place. Now, in 2001, I think the UK is uniquely positioned to forge new relationships throughout Europe and bring its broad expertise in AI research and technology transfer to a much wider community. The Expert Systems Conference in Cambridge every year is a great illustration of this broadening appeal and influence of UK AI – only six years ago I remember it as being a mostly UK affair, while in the recent past it has taken on a distinctly international flavour, with participants submitting papers and attending from all over Europe, the US, and Asia.

I can note a few sea changes I have seen in the field of AI in the last 20 years that many people (even in the field) may not have noticed. Firstly is the move away from symbolic AI (which in many ways has now been assimilated into mainstream computer science) toward more biologically inspired and abstract methods (connectionist and distributed AI). This represents a much more fundamental change than it might at first appear. While there are as many opinions about the nature of human cognition as there are people studying it, it was only recently within our field that the mechanistic shadow of Babbage, Turing, and others began to recede and a timid recognition emerged: we don't really know very well what makes up human thought, but it is unlikely to be rules and statistical order. This goes against the prevailing winds of the twentieth century which has firmly held that science, the human mind, and the laws of the universe were finally within our grasp and knowable. Perhaps Turing understood something we are only now grasping: teaching computers to think like humans may be inherently counter-productive, as they have their own manner in which they are naturally productive (rules and mathematics) as do we (intuition); we are no better at mathematics than machines are at formulating common sense or natural language. Men and machines are, for now anyway, fundamentally different; and while they may not mimic their creators very well, have their own strengths and manner of perceiving their environment - they may even have their own form of machine consciousness which we are unable to perceive. The message of the new millennium may be that we may no longer know what we thought we knew, and are returning to some of the most essential questions in our field.

The extraordinary acceleration in the development of computer hardware has also brought about another difficult realization within the last ten years: that the machines are getting exponentially more capable while human kind remains, biologically, mired in a rut. To be realistic, many of the advances in AI software and application in the last two decades has been due not to any renaissance in the field (although one has occurred) but

to the prevalence of cheap PC hardware. The common desktop PC now has the kind of horsepower to explore and test AI theory only dreamt of just a decade ago (a good example of this is the victory of Big Blue – after all these years of research it was finally pure horsepower that beat a chess grandmaster – again, IBM was allowing the machine to do what machines do best).

I'd also mention that just in the last few years, even though academic consolidation has collapsed many AI departments and divisions back under the umbrella of informatics, AI as an academic discipline has become ever more distinct. Even ten years ago AI was often regarded as the curious offspring of the fields of engineering, mathematics, and cognitive science, while it is my impression that it is now more generally being viewed as a discipline in its own right. We're getting good at building things that no one else is. This may be due to the fact that AI is becoming ever more transparently integrated into other disciplines and industries (soft computing, machine learning, and expert systems into computer science, NLP and HCI into cognitive science, optimization into engineering and industrial design, planning into military applications, neural systems and fuzzy logic into a dizzying array of complex domains, etc.) while itself becoming more and more unique, advanced, and inspired. As things become more complex, the need for good embedded AI continues to grow. Clearly, AI is becoming a mature field.

How do you see UK AI research being incorporated into products, especially where this had a major commercial impact?

Professor Michie: To get a fix on just this question, in 1984 I moved to Glasgow and with Alty of Strathclyde University founded the Turing Institute, financed by affiliate subscriptions and corporate and Government contracts. From this grew a uniquely instrumented if historically undocumented robotics laboratory funded largely by the US Corporation Westinghouse. There, a FREDDY 3 was developed by Mowforth and Shepherd, culminating in innovative uses of robot-generated and robot-recognized English-language voice signals whereby a pair of robots (e.g. an assembly robot and a find-and-fetch robot) communicated in co-operative tasks. Although this potentially exploitable product did not find a commercial market, it was wholly financed by Westinghouse Corporation in explicit acknowledgement of generic commercial benefits derived from their Turing Institute connection, as touched on below.

Harnessed alongside the Turing Institute, the company Intelligent Terminals Ltd, was dedicated to exploiting a novel programming-by-examples technique. "Structured Induction" (programming by examples) was developed by Alen Shapiro in an Edinburgh Ph.D. Thesis which he completed in Glasgow. One of the Turing Institute's first moves, as a joint venture with ITL, had been to equip Westinghouse at Monroeville, USA, with inductive software and know-how which brought this corporate client immediate, and generously acknowledged, returns in excess of \$10 million per year in their automated uranium fuel refining operation. With Shapiro on board until his departure to the USA, and myself as part-time Technical Director, ITL sold software products and services world-wide, ending with a profitable sale of the company in 1988 to Infolink Ltd. ITL also stimulated the formation of two companies, in the UK and Sweden respectively,

which have continued independently to flourish to this day on the basis of ITL's initial sale to them of software and know-how. The technique of Structured Induction itself, although fully documented in the open literature, remains academically unknown, or at best unused. In conclusion, both the Turing Institute and ITL amassed corporate client lists which in sum must certainly amount to many scores of millions of subsequent quantifiable benefit to those clients.

Richard Wheeler: In the rather dim economic climate we now all seem to find ourselves, it seems little consolation to give reminder that AI technology is in widespread use at many companies and on many levels: most credit transactions are reviewed by an expert system of one kind or another, data mining and KDD is becoming a cornerstone of new economy and old economy businesses alike, the military and the aviation industry makes wide use of planning, scheduling, optimization, and a myriad of other critical tools, and so forth. The adoption of these techniques not only indicates the validity and usefulness of AI research through the faith placed in it by the commercial world (which has benefited greatly), but also the state of development and maturity of the field as a whole. What I find more encouraging is the emerging attitude in industry that AI research simply offers better and better plug and play tools with which to optimize commercial value, and while this may leave many researchers feeling a bit underwhelmed and misunderstood, I take this as a comfortable sign of stability in the market for our skills as tool-builders and engineers (and of course, what inventor or researcher does not want to see their creations in fruitful use?).

Even more promising is the movement within our field to embrace open source development of AI platforms, code, and techniques. While at AIAI in Edinburgh I designed a case-based reasoning (CBR) design and diagnostic shell, and in the four years since making it open-source, have received thousands of requests for its commercial and academic use despite having no advertising or real web presence. While I have never made a pence from all this, it proves to me that a market (and a strong one at that) exists for good tools which are well designed and do what they do simply and effectively, and I think this bodes well for AI and the technology industry as a whole. From the commercial requests I have gotten for the CBR shell system (and another open-source system, an artificial immune system shell) most seem to be hopeful that a machine learning system will be capable of sorting out and making sense of genuinely intractable and poorly understood problems, and I think this indicates a general shift in engineering design and analysis (both commercially and industrially) to harness the power of ever more powerful tools for ever more confusing problems. As a planet we are designing systems that are beyond our own abilities to monitor and control, and we are increasingly turning to intelligent machines to sort it all out.

And as with AI more generally, the UK has a unique position and stature within Europe and the wider global community providing high quality and commercially driven AI and technology transfer solutions. As a final thought I would stress the importance of those working in AI not to let their creations sit on the shelf, but to open them up to the larger academic community and industry and to focus on making our technologies ever more useful, valuable, and available.

Looking back, how might you describe machine learning (ML) as a sub-field of AI and do you see it as having biological roots?

Professor Michie: ML is the use by machines of data samples for the purpose of responding more effectively to further data samples from the same source. It is commonly subdivided into (a) rote learning, (b) parameter learning, (c) description learning. A further category is (d) concept learning, in practice misapplied as a general label for (c). It is best restricted to that subset of (c) concerned with descriptions interpretable by human brains.

AI scientists can gain enormously from studying present neuro-cognitive advances. We all take too little note of what is going on in brain science and in the study of human learning. Just taking one's own (possibly under-informed) pet theory of biological learning, and more specifically human learning, and then embodying it in a program, and then showing that the said program learns things, does not in itself prove very much. I am personally inclined to proceed with McCarthy's dictum in mind, namely that before you try making a machine capable of learning in some competence at a non-trivial level, you should see if you can make it capable of being *instructed* at a non-trivial level. In the process you may find that the biological skill which you would like the machine to acquire has a totally different nature from what you had imagined. I am currently finding this in my attempts to instruct and test conversational agents. Existing logic-and-linguistic approaches just don't seem to apply to the task of simulating realistic human chat. More to the point are associatively linked activation networks of pattern-fired rules within frame-like "contexts". Biologically, chat is more like the mutual grooming of primates than it is like Platonic dialogues!

Anyone who really tries to follow McCarthy's dictum will soon recoil at the inadequacy of programming as a means of instruction, meaning by "programming" the most powerful methods and languages available today. *Inductive* programming is another matter. With suitable tools one churn out and test over 100 lines of C or Fortran per day, and code maintenance becomes a snap. A big bonus from programming by examples lies in the slogan: "Don't fix the code, fix the examples". Then just re-induce from the fixed set.

In enunciating his dictum, McCarthy had symbolic skills in mind. What do we do then about subsymbolic skills? How about the task of instructing a team of dogs to play three-a-side football, as is done today using Sony's robot dogs? The UNSW AI Lab are the reigning world champions in this section of the annual RoboCup. The programming task proves to be a backbreaking, brain-busting labour of fearsome tediousness. Can we get any help from looking at real dogs? No-one knows how to program (in the McCarthy sense) a real dog. And no dogs as far as I am aware know anything about football. But from what I glean about dog training (mainly from Stanley Coren's masterpiece "The Intelligence of Dogs") I would bet that top dog-trainers could in a few months turn a team of bright dogs, like poodles, into competent players, even able to observe the off-side rule etc. Why? Because dogs are born pre-programmed to be *instructable by example*. So trainers know very well that "inductive programming" is the way to go. If Sony had

equipped their latest Aibo robot dogs with the kind of programming-by-example tools I remember from IITL days, RoboCup programmers might be delivered from much unwelcome grind and bind!

Inductive Logic Programming (ILP) of course is a further step in the direction of inductive instructability. But no-one yet uses it in *structured* mode. The expressivity of individual clauses is so great as to let the user off the hook of having to structure the code to retain transparency. As the complexity of tasks increases, I will predict that ILP-ers will in the end begin to use their theory-discovery tools in just the way that Alen Shapiro pioneered at propositional level all those years ago.

Richard Wheeler: AI has undergone such a transformation in the last 10 to 15 years that it is hard to separate machine learning from broader AI, and ML (like AI) has both biologically inspired and non-biologically inspired adherents, but I think machine learning as a field has come to be most strongly regarded as a symbolic and non-biological sub-field (for example, induction, theorem proving, planning, ILP). My definition of ML (and AI) is generally “making things which get better and learn as they go along”. And as we have only just begun on our journey to understand the workings of our own biology and of the web of life around us so too is new inspiration coming for ever more effective and powerful machine learning.

Do you consider AI and ML to be a form of advanced engineering or tool building?

Professor Michie: R.A. Fisher built many useful tools. He would have been astonished to be told that what he was doing was engineering! So would Abraham Wald, whose "sequential analysis" was anticipated by Turing. Ditto Robinson *re* Resolution. The relation to science is that AI is abstract instrumentation. But so is arithmetic, or any kind of applied maths. What am I saying then? Simply that AI people are instrumentation engineers, just as are the designers and builders of telescopes. What in AI we design and build could be called “epistemoscopes”.

Richard Wheeler: Tool making has had a very pivotal role in the advancement of human kind as a species, and this sort of technical evolution, especially in human terms, is a very difficult topic to consider and address. It is commonly heard that human evolution has shifted from the organic (nature making better toolmakers) to the purely inorganic (human kind making better tools) - that technology is an evolutionary "extension" to humankind, and as such, has taken over from Darwin. The more technologically fit among us, it would seem, are prospering at an ever-increasing rate. While I would not argue with this concept, I see real challenges growing out of the fundamental unfitness of humankind to control the artifacts it is beginning to create; surely the day will come when technology takes over as the dominant form of life on the planet - it has been widely proposed that this may even be nature's plan as human life becomes increasingly untenable. I do not necessarily view this as a bad thing, and believe it may be reasonable to assume that the mechanisms nature uses to keep organisms in check, and from outgrowing or destroying their niche, may apply to humanity as well. However, I believe that the great promise of technology and AI is still as tools to advance and improve the

human condition. AI is not unique in this perspective. We created the hammer because our hands are terrible at pounding in nails. No one was afraid of the hammer until someone sharpened one end. We created the calculator and modern computing devices for the same reasons - because our minds are terrible at manipulating numbers. Many people are looking to AI in a similar fashion - to create tools to monitor and manipulate systems that we are unable to understand ourselves, and like all tools, AI is undoubtedly dangerous. The real problem is the power a hammer, calculator, or AI device gives its creator - as before: the tools are getting smarter, but we are not.

What do you find exciting in AI in the new millennium and what are you currently working on? How do you see the future of AI in the UK and abroad?

Professor Michie: With Prof Claude Sammut at the University of New South Wales, and with help from Dr Zuhair Bandar's group at Manchester Metropolitan University, UK, we are developing a natural-language conversational agent. Our aim with "Sophie", is to automate human-computer chat in something like the style of the "Turing Test". In 1950 he outlined this as follows:

“I believe that in about fifty years' time it will be possible to programme computers, with a storage capacity of about 10^9 (ten to the power of 9), to make them play the imitation game so well that an average interrogator will not have more than 70 per cent chance of making the right identification [as between human and machine] after five minutes of questioning.”

A few commercial niches for such "conversational agents", if they can be achieved, are listed below:

1. Personal guides to trade shows, conferences, exhibitions, museums, galleries, theme parks, palaces, archaeological sites, festivals and the like.
2. Web-based guides for e-commerce. E-businesses like Amazon Books have Web-sites that actively broker direct interaction between buyers and warehouses. They build incrementally assembled profiles of the individual tastes of each customer. There is now a need to do more than personalize, namely to humanize. Interaction should flow through a virtual person to which the customer relates as to a human catalogue guide and advisor.
3. Coaches for English as a second language. Growing numbers are today displaced across national frontiers into new lives in which they face language barriers. Improvements are required to what can be delivered by current Computer-Aided Language Learning packages, including distance learning. There is a need to enable learners to practise conversational skills. Chat companions endowed with virtual personalities may make this possible.
4. In advanced countries the proportion of the population past retirement age is growing. To counter isolation of the elderly, it may be possible to supplement

books, pets and conventional televised entertainment with personal chat companions as Web TV add-ons.

5. Every country has a growing underclass that, if left unemployed and untrained, overflows into streets and prisons. Inexpensive autotutors with conversational interfaces could support and expand current skill-training and re-skilling programmes.

6. Electronic books, could have their footnotes extended by chat agents. When mouse-clicked, each such point could evoke a knowledgeable pseudo-human source able to elaborate it in discussion with the reader.

Current chat performance is still superficial and incoherent. Approximation even to the standard of Turing's 1950 "imitation game" still lies in the future, even though it is undemanding. Suppose that as many as 40% of the judges see through the disguise and correctly pronounce which of the two remote conversants is the machine. The remaining 60%, being unable to say which is which, must assign the label "machine" at random to one of the two. On average therefore one half of these 60% will make the correct identification by chance. So the expected percent of correct identifications comes to 40% + 30%, equal to Turing's criterion. Since this presupposes that the agent fails to fool the judges almost half the time (40%), and since Turing only allows the judge five minutes to penetrate the machine's disguise, his formulation of the Test sounds very permissive.

I believe that it may prove possible to "tune" even today's breed of conversational agent, popularly termed chatbots, to this relatively undemanding level. But transition from chat to genuinely rational discourse confronts an apparently profound difference between chat and discussion. My aim is to demonstrate over the next few years, that even discussion can be simulated by the simple-minded methods that we are applying. These methods are based on associative retrieval from huge cross-indexed dictionaries of stock responses. My hunch is that the common idea that discussion proceeds by logical reasoning is mistaken. The time that the brain requires to reason "on the fly" is too long for the tempo of chat. Rather, people have already constructed reasoned justifications for their opinions, probably at leisure in countless earlier readings and ruminations. In this way they accumulate the canned fruits of prior reasoning, and accumulated experience of stock counter-arguments. Later they repeatedly retrieve and deploy "on the fly" the canned materials.

A young philosopher, Robert French, published a persuasive case some years ago which essentially argued that Turing-Test intelligence cannot be mechanically simulated unless the machine has opportunity to directly experience the real world in a wide range of ordinary physical interactions, which experience he believed formed an indispensable substratum. It was this kind of question that motivated the Edinburgh experimental robotics project, FREDDY. The had long-term goal I had in mind was to see whether conjectures of the Robert French sort are true. The real world of scientific politics proved so dismissive of the approach that it had to be abandoned at an early stage. An enquiry along similar lines is I believe beginning to emerge from the work of the leading robo-

soccer implementers, whose latest creations we saw at the 2001 Seattle IJCAI. They have not yet reached the levels of inductive instructability, nor world-model maintenance and updating, that we were beginning to achieve (*see* Ambler et al., AI Journal 1975). But they are moving in that direction and beyond, with facilities for interagent intercommunication already in place. Can “Turing-Test intelligence” ultimately be evolved bottom-up, so that the players, or their robot coach, can summarize a game and points of play in subsequent QA sessions? The answer lies in the mists of the future.

Richard Wheeler: I am terrifically excited to be working in the field of artificial intelligence right now, and at Starlab have been exploring three directions as broadly as possible: new AI paradigms (focusing on machine learning through behavioural observation, artificial immune systems, and methods for chaotic modeling of complex systems), applied AI (multi-modal teaching methods, collaborative agent environments, child aware technology and advanced methods for anomaly detection), and new computing methodologies (immunological and chaotic computing). I also find the explosion of applied research in alife, cellular automata, inductive logic programming, behavioural cloning, hybrid optimization techniques, and AI applied to future computing methods such as chaotic, immunological, and parasitic computing very important and exciting – to play a role in the ongoing birth of a new field and new technologies is a wonderful thing indeed and I am glad to play some small part.

Some general comments about the present state of AI and what I look forward to in the years to come in our field. AI, like science in general, is in its infancy - of course, this gives us the opportunity to contribute to the very root of the field; be there at its birth, as it were. Again, I find this very exciting. The 20th century has left us with an impressive legacy of human myopia, greed, and hubris, which is strongly reflected in our artifacts and technologies. Despite how it may seem, we really know very little about ourselves, our minds, our planet, or our universe - and these deficiencies have been passed on to (and have held back) AI since the days of Turing. As above, it has not been conceptual or theoretical advances that have fuelled growing research and innovation in AI (even today's cutting-edge AI has existed in theoretical form for 50+ years), but the availability of fast inexpensive PC hardware. Similar future advances in computational hardware (especially evolutionary and inherently parallel hardware) will surely bring about similar effects. The rise of chaotic, quantum, optical or other as yet unharnessed computing methodologies will undoubtedly spark another vast revolution in what we now call AI.

One of the essential conflicts in the public's mind about AI seems to be that of mankind versus machine, but many people fail to grasp one simple fact which is at the root of the controversy: that the machines are getting smarter, but we are not. This, almost inevitably, will cause some Copernican re-appraisal of human kind's place in the universe within our lifetime. Considering the current rate of progress in the physical sciences, it may not be unreasonable to assume that we will see the dawn of "real" AI (the possibility of human-level cognition) within the next 20 years or so - guessing about the future of AI systems beyond that point is unfruitful. There are a few things we can guess at, however. The first is the rise of evolutionary reasoning devices rising out of the context of present-day genetic programming and artificial life methods. Around the turn

of the century (the 19th century) we began to build artifacts and create technology which we are unable to understand, properly monitor, or control; systems of such complexity that we as a species may lack the intellectual band-width ever to fully comprehend. A jet engine is one such complex and chaotic device (another common example is the internet) - despite following a very simple design principle and being made of fairly well understood components, once it is assembled and put into use, it defies our abilities to monitor, control, and predict its behaviour. This reflects a number of fundamental failures: our lack of advanced sensing equipment to properly monitor the device's components, our lack of understanding of chaotic physical systems, and our willingness to build and use things which we do not understand and cannot properly control. While all AI methodologies will play a part, evolutionary methods are the most likely way forward for real AI - we cannot describe and model systems which we ourselves do not have the "wetware" capacity to understand. It may be that you cannot design a brain (nature didn't), but must evolve it over time - so too may one day our machines develop and grow.

Another sure element in the rise of AI in the next 20 years will be the internet, or what the internet will become. The internet is about enablement and efficiency. Imagine that you are an infant living in a world where you cannot see, hear, smell, touch, or speak. In this world, you can only manipulate and create using the tools and constructs, which exist within a very narrow presentational and representational "bandwidth" - that is the state of AI now. The systems we create are invariably run and tested in toy domains with little or no recourse to the wider information world, but the internet is set to change all that, by providing a single protocol or access channel for AI to use. Of course information capacity (like complexity) does not make an object intelligent, only "well read", as in the case of the well-known AI system "CYC", but the web is sure to spark off an ever increasing deluge of better-informed devices. The future of the internet is not just to facilitate information transfer, but to enable representational form, function, and reasoning as well; something the printed page (the internet's parent technology) has long been incapable of. These developmental goals overlap heavily with real AI.

A word about robots - most people assume that AI is somehow about building robots, which I suppose used to be true - the idea being to build a "thinking engine" or machine which had human characteristics (classic cybernetics); in time, no aspect of human experience (physical, psychological, emotional, spiritual) went unexplored. Perhaps one of the most dramatic realisations in the field of AI is that we no longer want to mimic the unstable mind of man, but to build the mind of God. Many people, myself included, have little or no interest in recreating the fragile, unlikely, primitive, incoherent, and unreasonable minds of this tiny planet's latest inhabitants, and instead attempt to pursue the root of reasoning back as far as it can go.

AI has already taught us many crucial things about the nature of human thinking, perception, cognition, and reasoning - in the future it may teach us an even more fundamental point: that the universe is information rich and awareness poor. The future of AI may lie in the ability to integrate and reflect upon ever-increasing stores of available information and compress, reference, and recombine it in unique ways.

Humankind calls this “creativity”. My hope is that in the future technology (and AI) will have advanced to the point where humanity will be freed from the tyranny of bad weather, bad genetics, and bad decision making which, up until now, have singularly characterised the planet. In short, we will be freed to pursue those things that best represent humanity and its unique place in the universe: creativity, exploration, discovery, and compassion.

Professor Michie: On the place of AI in the turbulently evolving big picture which Richard has the spirit to tackle above, I support the general thrust. But the mind of God, seen as AI's construction goal, can be of little use without an access channel that mere humans can use, and can enjoy using. Otherwise AI professionals become a new priesthood (as has often occurred with previous keepers of sacred sources) mediating the laity's access to their arcane stores of wisdom. Hence my stress on virtual-personality interfaces.

My personal prognostication is a good deal darker than Richard's. But until I know of a way in which the human species can absorb facts which they find *prima facie* unwelcome, there seems little to gain from opening one's mouth. My limited contacts with aid workers, anti-war protesters, and other objectors to the pace of the Gadarene rush, give me the impression that they are as keen as the other side to see change in terms of a war between Good and Evil, rather than concentrating on needed calculations and simulations -- about things like the worsening exhaustion of world water resources or the accelerating transnational mobility of capital relative to that of people.

In my own qualitative mental simulations, I don't like what I see. From some laboratory reports that I have read, *homo sapiens* may in general prefer to be poorer provided that others of similar socio-economic class are guaranteed to be *even* poorer, rather than that all should be equally enriched. Getting dependable answers to these sorts of unknowns concerning our psychological constitution would seem to be a precondition of rational long-term plans -- e.g. how much should we spend on Kyoto goals, *versus* research into possibilities of motivational re-engineering, *versus* into the interconnections between governments, multinationals and organized crime syndicates, etc.? Little of this seems to figure in public statements by governments, NGO's, media, or other public institutions.

I am mindful of the opportunity to equip the scripts of conversational agents' with as much unobtrusive factology on topics such as these as we can.