[Please cite as:
Forsyth, R.S. (1985). Chapter 10, AI: Future Trends. In: Forsyth, R.S. & Naylor, C. (1985). *The Hitch-hiker's guide to artificial intelligence*. London: Chapman & Hall Ltd., 245-251.]

10
AI: Future trends

It can get lonely being conscious in the Cosmos -- especially with people like Copernicus and Carl Sagan to tell you how big it is. Of course we can talk to each other, we can even write books, but we are only talking to ourselves, After all, we are all human.

Only four prospects of allaying this loneliness exist:

(1) Communicating with extra-terrestrial intelligences;
(2) Teaching animals to speak (e g. chimpanzees);
(3) Learning the language of another species (e.g. dolphins);
(4) Building intelligent artefacts.

So far there has been no authenticated contact with extra-terrestrial beings of any kind. It would be immensely interesting to converse with someone who was not human; but if we wait for messages from Them-out-There, we may wait a very long time.

There has been a certain amount of progress down the second avenue, Washoe the chimp was taught to use American Sign Language with her keepers (captors?), but when the controversy over whether she was using Language with a capital L subsided, so did most of the interest. It turned out that chimps did not have much to say,

Dolphins, whales and porpoises seem to have plenty to say, but we cannot make fin nor flipper of it, despite three decades of serious research and the benefits of intensive computer analysis. Perhaps if we beamed whale-song into space instead of a stream of quiz shows and commercials, Them-out-There might take the trouble to reply.

That leaves machinery. It appears for the moment that if we want to contact a mind that is not housed in a human skull, we will have to build one. How long is it likely to take? Will we have to wait till after we have cracked the dolphins' code and the little green men have beamed down from Ultima Centauri and points west?

This chapter is an attempt to answer those questions.

10.1 Generation after generation

So far this book has concentrated on the acceptable face of AI. We have viewed AI as the vanguard of computer science -- a fertile source of fresh ideas and clever tricks. When these ideas and tricks are successful, they filter out to other parts of computing and prove extremely valuable. This has happened to list-processing and conversational computing the the past, and is currently happening to knowledge-based systems. As we pointed out in Chapter 1, AI 'exports its successes'.

The medium-term prospects for this type of work are good. We can expect progress on a broad front. There may be setbacks, and some of the bolder predictions may not be realized, but by the end of this century we should have witnessed substantial advances in the fields of computer vision,

automatic translation, machine learning and knowledge-based systems. Even the thorny problem of continuous speech understanding will be close to solution.

The Japanese fifth generation initiative has given a tremendous boost to this side of AI. By their ambitious plans the Japanese have subtly re-defined the rules of the game for the computer industry. They expect to produce knowledge information processors in the 1990s which are based on radically new, highly parallel computer architectures. The software that will make such systems work is firmly rooted in AI research.

Japanese reliance on AI techniques has made the rest of the world hurry to scramble back into AI. Now vendors outdo one another in their efforts to claim that their products incorporate the latest AI research almost as if AI was an ingredient that could be added to software like fluoride to toothpaste. And governments have been persuaded to approve multi-million pound investment programmes for fear of 'being left behind in the race'.

All this has polished up the image of AI, which only a decade ago was badly tarnished by its failure to deliver the goods. AI workers can take legitimate satisfaction from coming back into fashion in such a big way -- especially those who weathered the hard times in the 1970s.

But there is a darker side to AI. This is the quest for superhuman intelligence. And the danger is not so much of failure, as of success.

10.2 Brain and superbrain

Jack Good, now of the Virginia Polytechnic Institute, has coined the phrase 'Ultra-Intelligent Machine' (UIM) to describe the kind of system that AI workers, whether consciously or unconsciously, appear to be striving towards. Although it is a problem for the twenty-first century rather than the twentieth, the question of how to deal with UIMs is one that should be faced now, while humanity is still the dominant species on the planet.

The problem is this:

> Will UIMs damage your health?

and it boils down to two subproblems:

> (1) Can we create superhuman intelligences?
> (2) If so, will they be inhuman intelligences?

Let us consider why the answer to both these questions appears to be 'very likely',

10.3 The technological imperative

To foretell the future, we look back at the past. In 1948 the transistor was invented at Bell laboratories in New Jersey. In the same year the first true stored-program digital computer executed its first instructions at Manchester University. That was 37 years ago. In 37 years' time it will be A.D. 2022. Barring disasters, many readers will be alive then.

Since its foundation, which we arbitrarily date in 1948, the computer Industry has embarked on a period of exponential growth. In almost every parameter -- memory capacity, processing speed,

price, size -- we find a doubling (or halving) every 2 years or so. Nor is there any sign that the pace is slackening. In fact it may be accelerating.

Exponential change is rather hard to grasp. If the standard memory size of a desktop computer, which is now 256 kilobytes, keeps on doubling every 3 years (a conservative estimate) by 2022 it will be 4096 times as capacious as today. That means the average user will have access to over 1000 Mbytes of main memory, or 1 Gigabyte. You will be able to load the equivalent of 2000 hardback books into RAM. Since backing storage is typically ten times the size of main memory, you could hold a library of maybe 20000 books on disc (though it is more likely that you will store old movies on what will probably be an optical disc than text).

Such rapid rates of change overturn our expectations, which are generally based on stasis or at most linear change. A difference of such magnitude is effectively a difference in kind. In hardware terms, we have to imagine the unimaginable.

Advances in software have been much slower so far. But in the last chapter we touched on a few examples of adaptive systems like Eurisko, which can be seen as case studies in creativity, automatic programming or machine learning depending on your viewpoint. This is a key to a boot-strapping process which will launch the software industry on its own period of exponential growth. Once systems become self-improving there will be nothing to hold them back.

A system that monitors and improves its own performance is not tied down to the pace at which human programmers can respond to change by re-writing (and re-testing). There are sound financial reasons for developing software that improves itself, and in our society that usually means it will come about. So by 2022 both hardware and software should be adequate for genuine machine intelligence.

Add to this the fact that the problem of AI holds a curious fascination for some of the best human brains. If you attend a major AI conference you will discover that a number of the 'big names' seem to be driven men -- not so much mad professors as ideological visionaries. They tend to brush aside talk about social accountability, military applications and possibly harmful side-effects. They are impatient with doubts and delays. For them only the outcome matters: the task of creating a superhuman intellect is so urgent that it takes on an almost religious force.

They are the true revolutionaries of our time, on the threshold of one of the most momentous events in history. And in a sense they are right. After all, what could be more important than designing the successors of the human race? Man is undeniably imperfect. It may be possible to do better.

10.4 Evolutionary opportunism

Evolution is not a process that we can halt. Nor is natural selection something that we can switch off when it suits us. It is more like the law of gravitation than the laws of taxation.

We may alter the environment, say by eradicating smallpox or malaria, and thereby change the characteristics that are selected. But we do not stop the selection: that is inexorable. Some organisms survive and breed; others do not. The inheritable characteristics of those that do come to predominate.

Even if the human race exterminates itself, it would be mistaken to say that evolution had taken a 'backward' step. It would be more correct to say that a species which proved itself unfit to its

environment had passed into extinction. And, one way or another, extinction is the fate that awaits the human race. It may be minutes away, or millenniums, but so long as life continues it will continue to adapt and evolve. Homo sapiens is not immune to this process. That we will be succeeded is a certainty. The only uncertainty is what form our successors will take.

They may be another branch of the genus homo, or may derive from another line altogether -- rats, for example. They may not even be mammals. In fact, it is beginning to seem possible that they will not be living creatures at all, although they will have biological origins and may still contain organic components.

Biotechnology complicates the evolutionary picture somewhat. Obviously a creature that consciously manipulates its own genetic make-up has a unique position. AI may add another level of complexity, by introducing systems capable of reproducing themselves without being alive. Thus natural selection becomes unnatural selection (to some extent) but selective pressures do not disappear. It merely means that evolution has a new set of tools at its disposal.

Biological evolution is highly opportunistic: it grabs, so to speak, the materials to hand and makes use of them. Flippers become legs. Lungs developed in fish for buoyancy take on the function of breathing. The new is built upon the old. Our brains reflect this opportunism. The mushrooming neo-cortex is laid down on top of more primitive structures that still control emotion and motivation. They have not been replaced; they are just buried, and thus take on a changed role. If the process continues true to form we might expect our successors to be man-machine hybrids of some sort. Deep within their central processing units, under layers of advanced technology, we might find something bearing more than a passing resemblance to a human brain.

Humanity has opened two Pandora's boxes at the same time, one labelled genetic engineering, the other labelled knowledge engineering. What we have let out is not entirely clear, but it is reasonable to hazard a guess that it contains the seeds of our successors.

10.5 Misguided missiles

On balance, then, the answer to the question 'can we create superhuman intelligences?' ls another question: 'why on earth not?' The evidence suggests not only that we can, but that we are compelled to do so by the remorseless logic of evolution.

The next question is whether we are going to like the results.

One rather ominous feature of AI, which tends to indicate that we will not, is the extent to which it depends on military funding.

It is all very well to take money from the US Navy to develop an academically interesting system that can form an 'image' of the sea bed from sound signals, but those generous guys from the Pentagon are not as philanthropic as they seem. Next year they will be back asking you to put it in the nose of a torpedo.

More than half the short-term applications of AI are on the battlefield. They include:

Intelligent submersibles,
'Smart' munitions,
Cruise missiles,
Self-guided tanks,

Knowledge-based sonar systems,
Homing torpedoes,
Expert systems for radar interpretation

and many more besides that few of us will ever hear about. (Until it is too late?)

This is a disquieting trend. Consider what it is that makes an artillery shell 'smart'. It does not just fall out of the sky and blast a huge hole in the ground: it selects its target. As it nears the end of its flight it seeks out tank-like objects and adjusts its trajectory to make sure it lands on one. Bang! Shoot in the right general direction and you are assured of a direct hit.

That is the immediate impact of AI -- better ways of destroying tanks, and, though the promotional videos are coy about it, better ways of killing people. But if the test of intelligence is how many people get killed, surely a really intelligent system will kill everybody?

10.6 A critique of pure reason

Fortunately this doomsday scenario ls wildly improbable. It was only put in to give you a cheap thrill. The militarization of AI is a disturbing trend; but we already have so-called dumb weapons of such awesome destructive power that anyone who is still reading in2022 will have more pressing problems to worry about.

If we get that far, there will be plenty of civilian UIMs as well as military ones. Without actually thinking in the traditional sense they will be able to do almost everything that requires reasoned thought better than we can. In mathematics and natural science they will have gone far beyond us. In industry they will so far excel human managers that the running of entire economies will be under their control. They may not play golf very well, but they will be better at chess than any human being. In short, we will be their intellectual inferiors.

The effect will be similar to what happens when a modern mining company descends on a remote valley in Papua New Guinea to rip out a mountain or two. The natives have never seen a white man before (and they are not going to now, since the corporation is Japanese). Suddenly they are confronted with space-age technology. Their arts and crafts become redundant overnight. Witch-doctoring, face-painting, ear-piercing, hunting -- all the skills they take pride in become meaningless. They convert to a cash economy; take menial jobs in the mine; and spend their free hours in the company bar, consuming prodigious quantities of beer. One day, after the ore runs out, the mining facility moves on to fresh deposits, leaving a scar on the landscape and another wrecked culture on the slagheap of civilization.

We have three or four decades to prepare ourselves for this dependent status. We will be like stone-age hunter-gatherers in the urban desert. Mankind's long love-affair with the intellect will be over. How traumatic you find that depends on your attitude to life.

But, like it or not, one fallacy that must be disposed of is the idea that we can somehow 'pull the plug' just before machines get too intelligent for us. This is one genie that will not go back in its bottle.

10.7 The mind boggles

It sounds pretty grim. It may also sound light years away from you sitting at home with your 48K microcomputer, trying to get your chess program to make a move in less than half an hour. But

though the UIM and your home computer are worlds apart, you may live long enough to belong to both those worlds.

A lot depends on whether we, as humans, retain or abdicate control over the intelligent systems that we have started to devise and which we have a compulsion to keep on improving. A widespread understanding of the achievements, and limitations, of AI will help curb its wilder excesses. If it is left to the technocrats, we can imagine how things will turn out. The more people who know what is going on, the more hope there is that inhuman intellects will not be inhumane.

Everyone interested in AI -- amateur as well as professional -- has a part to play in deciding whether the UIMs of the next century will be friend or foe. That includes you.