Book review

The Age of Intelligent Machines. Raymond Kurzweil, MIT Press. 1990. ISBN 0-262-11121-7. $39.95.

There is a plethora of books on Artificial Intelligence (AI) and cognitive science available today. Many of these are textbooks or collections of research papers. However, there are not too many books for the average science person, without a background in Computer Science, who wants to know more about AI.

Margaret Boden's *Artificial Intelligence and Natural Man*[1] is a tightly-knit account of early AI. At the time it first came out (1977), it was probably the best available account of AI. However, even this book is not oriented towards the lay science person. The second edition[2] has just one new chapter (on Connectionist ideas). One would expect the book to include more recent examples, and a discussion of current ideas.

Most of the books on cognitive science tend to be highly cerebral, and are definitely not for the casual reader.

What this book offers

If you want a popular but not watered-down account of AI and related areas, which talks both about the theories in the area, as well as practical systems, which describes both commercial products that have come out of AI as well as the genuinely intellectual results, you must look at *The Age of Intelligent Machines* (henceforth abbreviated AIM) by Raymond Kurzweil.

AIM delivers all this and more. It is vividly and generously illustrated. There has been a great deal of effort put into presenting a unified picture of AI and cognitive science, and the author and his research team seem to have delved deep and wide to bring us all sorts of relevant information.

AIM also has a set of delightful interludes, written by the people who have made and continue to make this field—and nothing brings work more to life than a description by one who is immersed in that area. AIM also has photographs of such people, and sample outputs from their creations, creating a lasting impression in the reader's mind.

I guess part of the reason why this book is unusual is that Raymond Kurzweil seems to be unusual. He seems to be a person with several achievements in diverse areas of interest, and a strong believer in technology. To quote from the dust jacket: 'Raymond Kurzweil is the principal developer of the world's first print-to-speech reading machine for the blind, the first electronic keyboard to accurately re-create the sounds of the grand piano and other orchestral instruments, and the first large-vocabulary voice-activated word processor. ··· (he) holds seven honorary doctorates in science, engineering, music and the humanities from leading universities.'

In this review, I want to give the reader a flavour of AIM. I will describe the contents of the first two chapters of the book in detail, but skim over the rest of the book.

AIM's view of AI

The books starts with a prologue on the 'second industrial revolution', which describes the first industrial revolution and its impact, and looks at the potential impact of intelligent machines in this new age.

The first section of the book is about the roots of AI. The first chapter entitled 'What is AI, anyway?' discusses the moving frontier definition of AI: *AI is the study of computer problems that have not yet been solved*. There is then a discussion of intelligence, which concludes: 'The development of machine intelligence helps us to understand natural intelligence by showing us methods that may account for the many skills that comprise intelligence. ... As we begin to comprehend the depth of design inherent in such 'deep' capabilities as intuition and common sense, the awe inherent in our appreciation of intelligence should only be enhanced.' (p. 19). This is essentially the humanist view.

In a section on evolution as an intelligent process, Kurzweil talks about the fact that 'evolution has achieved intelligent work on an extraordinarily high level; yet has taken an extraordinarily long period of time to do so...a strong case can be made that we are far more intelligent than the slow process (of evolution) that created us...' (p. 21). He then goes on to ask 'So what about the intelligence that we are in turn creating? ... our thinking is today significantly superior to that of our machines. Yet the intelligence of our machines is increasing at a very rapid pace... it appears likely that (at some point in the future) human intelligence will be surpassed...' (p. 21). It is interesting to note that Kurzweil has not shied away from the possibility 'that an intelligence may ultimately outperform its creator'; he not only discusses it, he also views it as counter-intuitive but consistent with the observation quoted above.

The second chapter, on 'Philosophical Roots', takes us on a whistle-stop tour of philosophy, covering Plato, Descartes, Newton, Kant and Wittgenstein. This chapter also (impartially) discusses the arguments of Hubert and Stuart Dreyfus, who have been strongly critical of AI, and what AI can achieve. The next section is a Platonic dialogue (reminiscent of Hofstadter's interludes in '*Godel, Escher, Bach: The Eternal Golden Braid*') on the nature of human thought. After this comes a piece by a Margaret Litvin who apparently wrote the article as a seventh grade student. She prefers to live in an age of intelligent people, and argues that computers can be made more intelligent by making people more stupid, and perhaps that is already happening!

Daniel Dennett, a philosopher from Tufts University, then comes in with an excellent piece 'Can Machines Think'. He argues in his usual erudite and articulate fashion that the Turing Test has been misunderstood and misapplied.

In 1950, Alan Turing proposed what is now known as the **Turing Test**[3] as an operational definition of AI. He proposed (to quote Dennett, p. 48) 'a parlour game of sorts, the imitation game, to be played by a man, a woman and a judge.... The man and woman are hidden from the judge's view but are able to communicate with the judge by teletype; the judge's task is to guess, after a period of questioning each contestant, which interlocuter is the man and which the woman. The man tries to convince the judge that he is the woman, and the woman tries to convince the judge of the truth. The man wins if the judge makes the wrong identification....

Now suppose, Turing said, we replace the man or woman with a computer

and give the judge the task of determining which is the human being and which is the computer. Turing proposed that any computer that can regularly or often fool a discerning judge in this game would be intelligent, a computer that thinks, *beyond any reasonable doubt*.'

Dennett considers various alternatives to the Turing Test, and shows how all of them are inadequate. In the process he talks of computer chess, and AI systems such as the natural language interface LUNAR, the computer simulation of a paranoid patient PARRY, and the CYRUS program which modelled the memory of former Secretary of State of the USA. Dennett warns that 'the problem of overestimating cognitive prowess, comprehension, and intelligence is not, then, just a philosophical problem. It is a real social problem, and we should alert ourselves to it, and take steps to avert it' (p. 59).

Next, Mitchel Waldrop, whom many of us know from his writings in *Science*, has an article on *Can Computers Think?*. The article starts with a nice overview of AI. Waldrop then describes the Chinese Room problem of the philosopher John Searle, and provides counter-arguments to Searle's parody of AI. He discusses science as a message of despair—how Copernicus, Darwin and Freud brought us from the centre of the universe to a small point in an infinite universe, with man as another animal, at the mercy of demons in our own unconscious minds. He also talks of science as a message of hope, saying: 'the symbol-processing model (of AI) ... shows us how feeling, purpose, thought and awareness can be part of the physical brain and yet transcend the brain ... The individual enzymes, lipids, and DNA molecules that go to make up a cell are comparatively simple things ... And yet when all those molecules are brought together in an exquisitely ordered pattern, they *are* life... In the same way our minds are perhaps nothing more than machines ... Perhaps we are just processors of neuronal symbols. ... Perhaps *the Magic Flute* is only a sequence of sound waves... And perhaps, in illuminating the nature of mind and intelligence, AI is only reaffirming how unique and precious the mind is.' (p. 67)

Sherry Turkle from MIT then talks about children's perception of the world, about Piagetian ideas, and how modern children are accommodating computers into their model of the world. She warns against 'the too easy acceptance of the idea that computers closely resemble people in their thinking and differ only in their lack of feelings ...' (p. 73). The next interlude, by Blaine Mathieu, founder of a Canadian computer firm, is a conversation between 'two "people", a philosopher and a computer, (handling) some of the physical, emotional, and moral issues of machine intelligence. Seymour Papert from MIT comes up with a (possibly fictitious but definitely humourous) anecdote about the way Weizenbaum's famous ELIZA program (which sought to imitate a Rogerian psychotherapist) 'passed' the Turing Test.

The chapter ends with a lengthy but well-written piece by Doug Hofstadter entitled 'A Coffeehouse Conversation on the Turing Test'. After an initial dialogue between three students on the Turing Test, Hofstadter describes a program called *NICOLAI*, shown to him by students of the University of Kansas in Lawrence, which could converse in natural language. There is a long transcript of his amazing dialogue with *NICOLAI*. Hofstadter then uses this example to talk of the efficacy and nature of probing questions in the Turing Test, and how the way people use the interactive communication tool called 'talk' is revealing. Hofstadter concludes that various modifications notwithstanding, the original Turing Test admirably filled the purpose of converting a philosophical question into an operational question.

Hopefully, this description gives an idea of the inter-disciplinary nature of the book, the extent of the topics discussed and the depth to which issues are examined. Clearly, articles such as Dennett's and Hofstadter's vastly enhance the quality of the book, since the authors are good communicators and technically sound.

About the rest of the book

The rest of the book is assembled in a similar fashion, with discussions of major issues interspersed with interludes from various contributors.

The third chapter is on mathematical roots of AI, and the fourth deals with modelling intelligence. The fifth chapter deals with mechanical computers while the sixth is about early electronic computers.

The second part of the book on the *Moving Frontier*, starts with a discussion of the areas of Pattern Recognition and Vision. Kurzweil adds a personal postscript to this chapter, describing work by him and his companies on Optical Character Recognition, Automatic Speech Recognition and Electronic Music. In each area, he has brought out products such as, for example, reading machines for the blind, and voice-controlled word processors.

The next chapter, on the search for knowledge, deals with the idea of using vast amounts of domain-specific knowledge to solve problems. There are pieces from Ed Feigenbaum, one of the early proponents of Expert Systems, Kazuhiro Fuchi, the head of the Japanese Fifth Generation Computer Systems Project, and Brian Oakley who headed the Alvey Project of UK.

The ninth chapter, titled 'The Science of Art' explores the way AI has interacted with art, in particular with painting, story-telling, music and creativity. Several pictures created by Harold Cohen's program AARON and Cohen's own article on art and AARON add to this chapter. Cohen concludes his article thus: '... remarkably little of (AARON) has to do with art; it constitutes a cognitive model of a reasonably general kind, and I even suspect it could be adapted to other modes without too much distortion. But the lack of art specificity is not as puzzling as it may seem at first glance. The principal difference between artists and nonartists is not a cognitive difference. It is simply that artists make art and nonartists do not.' (p. 385).

Kurzweil then describes an experiment conducted with the Kurzweil Cybernetic Poet (!) which takes in original poems written by humans, creates a model of the poet(s), and generates original stanzas of poetry using such models. Kurzweil put together some stanzas written by actual poets along with some created by this program, and asked people to decide which were created by humans and which by the program, in a kind of a Turing test The program did manage to confuse the judges!

The tenth chapter deals with Visions

of the future, and lists potential breakthroughs, including a translating telephone, an intelligent assistant, an intelligent telephone-answering machine, a cybernetic chauffeur, invisible credit cards and artificial people! Allen Newell, one of the founders of the field of AI has a short piece on *Fairy Tales*, where he says: '... the aim of technology, when properly applied is to build a land of Faerie' (p. 421), and further that 'computer science and technology are the stuff out of which the future fairy land can be built' (p. 423).

The last chapter discusses the impact of AI on employment and the economy, education, communications, etc. This is followed by a piece by Margaret Boden on the *Social Impact of AI*, where she looks at both the positive and negative implications of AI.

Extensive notes, a good bibliography, a glossary and an index round off the book.

The only complaint I have about the book is that, in spite of the large size and the excellent printing, the publishers have chosen to use a dark typeface that is hard on the eyes.

Summary

To many of us, AI is about bettering the quality of life, using what have been termed 'intellectual steam-hammers', just as the steam-hammers of the first industrial revolution eased the physical burden of people.

AI is an amalgam of sciences and of technology. AI involves cross-disciplinary work covering a myriad of disciplines, understanding human processes, and human thought, and applying this knowledge to make more 'intelligent' systems. Clearly, AI is not a discipline insulated by *Keep Away* signs.

AI involves learning from humans, who have evolved over several millennia, about how to do things, what to do, and even how not to do things, and what not to do, and using this knowledge to create useful intelligent artifacts.

There have been many expectations of AI, some bordering on the absurd. But AI is not about replicating a poet's ability with words, or a mischievous child's smile.

There is a lot to be said about AI, and about the ideas and the technologies that are used in AI. It takes considerable skill to communicate this effectively And Raymond Kurzweil's book seems to have done just this, in a very comprehensible and attractive fashion.

I would recommend this book to anyone who wants a good introduction to Artificial Intelligence and related technologies. This is a book to be gently savoured and relished, much like an old friendship, or a mellow after-dinner liqueur.

1. Boden, M. A., *Artificial Intelligence and Natural Man*, Basic Books, New York, 1977.
2. Boden, M. A., *Artificial Intelligence and Natural Man*, 2nd Edition, MIT Press, 1987.
3. Turing, A., Computing Machinery and Intelligence, *Mind*, 1950, vol. 59.

Videotape review

The Age of Intelligent Machines: Machines That Think. Written and narrated by Raymond Kurzweil, Videotape, MIT Press, Duration: 28:50 min., $ 39.95.

The book *The Age of Intelligent Machines* has a (companion) videotape associated with it. The blurb on the tape describes this videotape as 'a survey of Artificial Intelligence (AI) showing AI at work and under development'. The videotape, written and narrated by Raymond Kurzweil, has several nice features.

First of all, it is well put-together. It is a good mixture of basic AI ideas with interesting examples of practical systems in the lab and in regular use. The multidisciplinary nature of AI is clearly brought out. It gives a balanced view of AI achievements and potential, without hyperbole or exaggeration. A newcomer to the area would probably get interested in this area with just one viewing of this videotape.

As in the book, prominent AI researchers feature in the videotape, talking about the 'paradoxes, promise and challenges of AI', making the field come alive.

One underlying theme in the tape is that technology in general, and intelligent machines in particular, can help people reach beyond their limits and overcome barriers of disabilities.

There are two applications that catch the viewer's attention. Stevie Wonder, the blind musician describes how speech synthesizers help him use a music synthesizer, manipulating the myriads of controls on such equipment. He also demonstrates how Optical Character Recognition (OCR) equipment using Pattern Recognition technology helps in reading out books for him. 'With technology', he says, 'I can do things. Technology is like a brother, mother ... a friend ... It is another sunshine in my life ...'

The other application is a robotic aid for the disabled, being developed at Stanford University by Stephen Michelowski. He is trying to integrate several technologies (such as robot motion, machine vision, speech recognition etc) in a mobile robot which can respond to spoken commands. The videotape shows clips of how a quadraplegic, paralysed by a fall on concrete, is able to instruct the robot to open a microwave oven, pick up a mug from the oven, and bring it to him, so that he can sip from the mug. When you first hear him talk of his dependence on people for everyday activities—'just to scratch my nose'—you realize what this technology means to him; it gives him the ability to do things for himself.

Roger Schank (then at Yale University) talks about what he likes best—cooking and eating—and about a program called CHEF, which learns how to cook new dishes, using a library of recipes. He talks of figuring out how we do things—composing a sentence, understanding what someone is saying—and about the need to understand a task in such a way that we can 'tell' computers how to do it. Marvin Minsky of MIT, one of the founding fathers of AI, describes how tasks that are (relatively) simple for humans to do, such as tying shoelaces, are quite complex for machines. He later describes his notion of a *Society of Minds*, a set of distributed intelligent agents working together to achieve intelligent behaviour. A machine architecture to test such ideas, the Connection Machine, is briefly discussed. Edward Feigenbaum of Stanford University, the messiah of Expert Systems, talks of using expert systems to simulate biochemistry. You then see the AARON program which creates art. Harold Cohen, of the University of California at San Diego, who