BOOK REVIEW'S

These reviews first appeared in *Nature* and are republished by kind permission of MacMillan publications.

GIVING 'EM THE BUSINESS

The AI Business: Commercial Uses of Artificial Intelligence.
Edited by Patrick H. Winston and Karen A. Prendergast.
MIT Press: 1984 Pp. 324.

W.F. Clocksin, Cambridge University

The past, it seems, is not always another country. In the 1960s, over-optimistic advocacy by researchers of the nascent field of artificial intelligence (AI) created expectations impossible to satisfy. In the 1980s, spurred by the prospect of serious international competition, we again have enthusiastic promotion of the technology of AI — especially expert systems and "fifth-generation" computers — and it is no less fervent: but this time the promoters move in commercial circles. The passionate propaganda from the marketing men dismays the researcher and perplexes the technical manager; as in the 1960s, the hyperbole obscures a realistic assessment of the true potential of AI technology.

In an attempt to clear the air, a colloquium was called by the MIT Industrial Liaison Program to bring together representatives of the four main groups involved: the academics, who started it all, the investment bankers and venture capitalists; those concerned with industrial research and development; and those in industry who use techniques devised by AI researchers because there is a job to be done. *The AI Business* is a collection of the edited transcripts of the colloquium. The book opens and closes with perspective and summary chapters by the editors. The first chapter sets the scene by placing MIT at the centre of AI research, and by giving a short "history" of the field complete with its own Prehistory, Dawn Age, Renaissance and so forth. The practising researcher may find it difficult to take this chapter seriously, while others may be dangerously misled by such cavalier treatment. Similarly, while the last chapter does a fine job of summing up, and offers some homespun advice on what action should be taken by universities and companies, there is again the tendency to lapse into dottiness, viz., "...strength in Artificial Intelligence ensures strength at everything that Artificial Intelligence can help with. And that may be just about everything". These two chapters are very much the bread of the sandwich though; the meat is proved by the 16 other contributors.

Although Winston (and later, Minsky) admits to regarding the term "expert system" with a certain amount of scorn, the casual reader will not understand why, for the first section of the book contains six chapters describing some of the best examples of work in the field. Here the practitioner will find accounts of his old friends XCON and DIPMETER ADVISOR, true products of AI research which have reached commercial credibility not by hype, but by quietly saving their users substantial amounts of money.

The next section, entitled "Work and Play", contains six chapters on a subject close to the hearts of AI researchers: interactive workstations. Part of AI research is about making computers easier to use, particularly as automated design assistants and as apprentices. This goes much farther than the computer aided draughting found in industry today. The computer should be a participant in the design process, synthesizing candidate designs that conform to sets of constraints with which the designer experiments. There is a fascinating contribution by Alan kay, who rambles through a compelling mixture of reminiscence and advice. Other chapters cover commercially available systems for electronic design automation and English-language communications with databases. The overall theme of this section is the life-cycle from basic research to exploitation in the marketplace.

The next section deals with robotics. Mike Brady gives a good survey of the relationship of AI research areas to their application in industrial robotics, while the other two contributions survey real-life industrial applications, giving examples of sensor-controlled robots and more general issues in factory automation. This is the most straightforward and informative section of the book, perhaps as a result of the authors' application of well-understood techniques to well-defined problems, with clear criteria for the success or failure of an application. Other areas of AI are not yet as much under control.

The final section, "Today and Tomorrow", centres on trends, both technological and economic. Minsky's chapter is a prolonged snort of disapproval at contemporary events in research, politics, business and industry. His basic message, though, is of prime importance. The commercial successes of today rely on AI research carried out ten years ago. Because there is practically no basic work in AI now being done, the future may already be lost. This theme is repeated in the other two contributions by investment analysts. Particularly impressive is the cogent and analytical chapter by the economist, William Janeway. Anyone closely involved with recent Government initiatives in "information technology" in any country will find considerable sympathy with Janeway's view that,

> The packaging for ignorance by artifice to obtain finance distorts not only the returns actually received by investors. Distortion of research activity in order to obtain finance on such terms is both a professional and a social cost that we in this country seem ill equipped to appreciate.

The main question to be asked of this book is, "Who is it for?". As the editors point out in the last chapter, the agreements and differences of opinion expressed at the colloquium have been preserved in the book. The transcripts of the informal discussion sessions at the end of each section are particularly effective in this connection. Great stuff for the practising researcher, who will be intimately familiar with the elusive issues involved, and who will thus be in a position to read between the lines. However, the discussions, together with the editors' own contributions, give the book a tone of superficiality and off-handedness which is dangerous and inappropriate for a more general audience. In spite of this, some of the individual chapters are particularly sound, and deserve wider exposure.

I want to see technical managers and engineers becoming acquainted with the issues covered in this book. But I would tell them not to believe everything they read in it. The simplistic, shoot-from-the-hip flavour may be instinctively appealing, but it endangers credibility, and should not be interpreted as solid worldly-wise experience. Despite Winston's "Six Ages of Artificial Intelligence", 20 years is not a very long time. What we need, above all is a little patience.

KNOWLEDGEABLE MACHINES

Knowledge-based Systems in Artificial Intelligence.
Randall Davis
Douglas Lenat.
McGraw-Hill: 1983. Pp. 490.

John Fox, Imperial Cancer Research Fund

Knowledge engineering is all the rage, just as genetic engineering was a few years ago. There are many similarities between these two new fields. Both have spun off from relatively recent sciences, molecular biology and artificial intelligence. Both attract the attention of the public and of investment managers to an almost embrassing degree. Both have yet to supply practical products on a large scale.

Fashion is fickle, however, and one might be forgiven for thinking that the knowledge engineering bubble will burst, or that the field will mature in the usual way and be rewarded with a respectable but properly modest place in the scientific establishment. Knowledge engineers, of course, do not believe that this will be the fate of their subject. Indeed, devotees of the parent discipline, artificial intelligence (AI), seem to have higher aspirations than merely "a proper place in the scientific establishment" (though the contributor to a television programme who said that "the three most important events in the history of the Universe were The Creation, the appearance of life, and the discovery of artificial intelligence" must have made even the most flamboyant AI scientist wince a little).

Given such high embition, I approached the volume by Davis and Lenat with the questions of scientific and engineering credibility in mind. It seemed promising. Both authors did the work described in the book as PhD dissertations at the Stanford University Heuristic Programming Project, still perhaps the leading centre of knowledge engineering. Furthermore the book consists of two complementary monographs; one with a scientific flavour and one with more of an engineering bias.

Douglas Lenat sees his work on AM (the Automated Mathematician) as scientific research with engineering implications. His efforts centred on the concept of "discovery", and the extent to which interesting discoveries in number theory can be made using programs that can use simple but powerful rules to generate and evaluate large sets of possible conjectures and hypotheses. For example "if $f(x, y)$ is interesting then $f(x, x)$ may also be interesting". If f is ADD, this suggests the concept of doubling, if TIMES then squaring, and so forth. Lenat gives many examples of more sophisticated "heuristics". The hypothesis that led to AM is that discovery (invention, creativity...) should not be assumed to be a mysterious mental faculty, but is the product of mental information processing mechanisms which are subject to logical analysis and thus to computer simulation. AM produced a few ideas which were new to Lenat, a couple that were "new to the world" and several "...additional 'new' concepts were created which both AM [and Lenat] agreed were better forgotten".

Lenat's work does not show that heuristic search is all there is to human discovery, only that there is a prima facie case that computers can be inventive. His hypothesis that discovery could be achieved by a physical device is not new — nothing is — but the demonstration that heuristics can lead to interesting discoveries, now, is dripping with provocative questions.