conversations. But surely a basic human ability is that of empathizing, of understanding how another would react and what another would do. The programmer ought to be able to do considerably better with the Block machine than Collins allows.

The other main theses in the book are that not all knowledge can be expressed in language and that machines are able to replace humans in some contexts (e.g., electronic calculators) because we have reduced some human action to machine-like behavior. Collins argues (as have Fred Dretske (1985) and Wallace Matson (1976) long before) that calculators do not really calculate, given that to calculate is to act. Collins distinguishes between behavior and action, denying that machines are capable of the latter. But I could find no argument for this claim. Since the primary distinction Collins makes between behavior and action seems to be the presence of an intention in action, one would expect an argument that machines cannot have intentions, but there is none. In sum, and hopefully no great surprise, although there is much to be learned about the interaction of humans with computers, it appears that there is not much that sociology can tell us about the limits of artificial intelligence.

References

Dretske, Fred (1985), 'Machines and the Mental', *Proceedings and Addresses of the American Philosophical Association* **59**, 23–33.
Matson, Wallace (1976), *Sentience*, Berkeley: University of California Press.

Department of Philosophy,
University of Minnesota,
Duluth, MN 55812, U.S.A.

DAVID J. COLE

Steven L. Tanimoto, *The Elements of Artificial Intelligence Using Common LISP*, New York, Computer Science Press, 1990, 559 pp., approx. $48.00 (cloth), ISBN 0-7167-8230-8.

This review is written from three points of view: those of the researcher, the teacher, and the student. The final summary raises the question whether the task attacked by this book is a solvable task to begin with.

1. Researcher's view

Research publications are often hard to read, because implicit assumptions have to be interpolated by the reader, difficult background material is assumed, steps of explanation are missing, or there are simply not enough examples. Whoever sits down and fills in all the missing links for the first time in the form of a textbook does a praiseworthy job and deserves considerable credit for making the research work accessible to students and to colleagues who might not be exactly in this subfield of specialization. (Compare J. Alan Robinson's (1965) original paper on resolution with the coverage of it in a good text on AI.)

On the other hand, it is much less valuable for material that has already been elucidated by one author, to be presented in a new book, just from a slightly different perspective. From my view point as a researcher, I do not see any major effort by Tanimoto in elucidating new research. For instance, the chapter on machine learning reuses Winston's arch-learning example, which had been sufficiently well described in Winston's own text (1984). Even taking into account the long delays between writing and publishing a book, the age of the cited references in the chapter on machine learning is depressing to this reviewer.

The reader might feel that I am expecting too much from an introductory text, and that might be true. On the other hand, it is necessary for new results to migrate as quickly as possible from research papers to students and peers. After all, we do not teach students Basic in a first course in computer science only to tell them in a graduate course in software engineering that they should use structured programming; rather, we have adopted research results on program design into our lowest level classes.

Quite in line with these comments, it should be noted that Tanimoto's text does not have a section on neural networks (for the record, I am not a connectionist). This seems completely unacceptable, since the connectionist revolution happened five years ago. In summary, from a researcher's point of view, I am not satisfied with Tanimoto's text, but it really needs to be added that this criticism applies equally to many other current texts.

2. Teacher's view

In this category, my impression is mostly positive. Does it ever happen to you that you read an explanation of something you know in a text, and you realize that there is a crucial step missing? You realize that you would not have understood it, if you did not already know? Unfortunately, this happens in too many texts, but in this book the problem occurs only seldomly.

Although explanations vary in their depth, most of them betray the gentle and thoughtful approach of the author. Some of them appear genuinely original and refreshing. Given the background of the author, it is not surprising that computer vision receives deep coverage and is well explained.

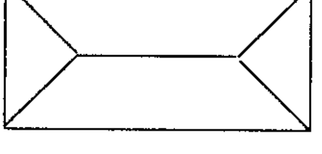

Fig. 1.

The few cases where explanations are not clear seem to be due to omitted figures or to the understandable effort to present a one-page explanation in half a page. To demonstrate this point, can you draw the medial-axis transform of a rectangle from the following description, if you don't know it already?

The medial-axis transform of a rectangle consists of diagonal line segments, open at one end, and meeting (if the rectangle is not square) on a fifth segment, also part of the medial-axis transform. The segments are open at the corners of the rectangle, because each of the corners, being a member of the rectangular set of points, has only one member in its $\mu(A, p)$. (p. 460)

If you cannot visualize this description, there is no figure in this book to help you. You have to refer back to Ballard and Brown's vision text (1982: 253) for the figure (see Figure 1).

3. Student's view

In order to judge the student's appreciation of Tänimoto's text, I asked my students at the end of an open-book final exam in my class that used it: "What is your opinion of the Tanimoto textbook?" (Alternatively, the students could have discussed the Lisp book used.) A credit of 5% of the exam was given, to ensure that students would be sufficiently "interested" in the question. One 8.5×11 page was available for comments.

The responses given to this question were considerably more often positive than negative. Probably the most negative responses were these:

> Also the examples given are not very clearly explained. Look[s] like the author[s] themselves were not very clear.

> I did come across some explanations that were not understandable. And could not get an understanding to what was trying to be explained.

A few positive quotes:

> . . . His explanation is simple and easy to understand and goes into just the right amount of mathematical analysis.

> Tanimoto's textbook provides good explanation of principles and actual programs which provides an ideal learning combination.

In general, the impression garnered from this informal survey is that the use of Common Lisp to make abstract principles concrete is appreciated. The sections vary in quality, with praise for the Computer Vision section and criticism for the coverage of logical reasoning.

4. Summary

Is it still possible to write a single 500+ page AI text that is sufficiently detailed to be understandable, that covers AI programming, and that covers all subareas at an up-to-date level? The sad answer is "probably no." The 55 pages that the author expends on an introduction to Common Lisp should have been replaced by an introduction to neural networks.

The approach used in a number of schools to teach Lisp first in a separate course and then concentrate on the real AI ideas and their implementation should prevail over the long run. Given an almost unsolvable task, the author has delivered a graceful and considerate attempt that ranks high among current texts.

References

1. Robinson, J. Alan (1965), 'A Machine Oriented Logic Based on the Resolution Principle', *Journal of the Association for Computing Machinery* **12**, pp. 23–41.
2. Winston, Patrick H. (1984), *Artificial Intelligence*, 2nd ed. (Reading, MA: Addison-Wesley).
3. Ballard, Dana H. and Brown, Christopher M. (1982), *Computer Vision* (Englewood Cliffs, NJ: Prentice-Hall).

Department of Computer and Information Science, JAMES GELLER
New Jersey Institute of Technology,
Newark, NJ 07102, U.S.A.

K.A. Mohyeldin Said, W.H. Newton-Smith, R. Viale, and K.V. Wilkes (eds.), *Modelling the Mind*, Oxford, Clarendon Press, 1990, viii + 216 pp., $49.95 (cloth), ISBN 0-19-824973-X.

Many people, on opening this book to an article on the Turing Test, will be torn between the feeling of "oh no, not again" and the interest of seeing what Donald Davidson could find to say about it. I will devote some of the little space available for detailed comment at the end of this review to that chapter.

This book is a collection of essays from roughly equal numbers of philosophers and what the blurb calls "scientists" (actually four psychologists and two physiologists). In so far as you can have a cognitive science collection without any artificial intelligence/computer science contribution (though Tomaso Poggio might be taken to provide it, I suppose, from within computer vision), then this is one, and its title certainly directs the reader to the central question of cognitive science, that of "modelling the mind". Artificial Intelligence is well represented in the bibliography, if not in the contributors' list. Conversely, Davidson, represented by two articles ("Turing's Test" and "Representation and Interpretation"), is modestly absent from the bibliography.

Some of its distinguished contributors have set out their views elsewhere on