

But the way Hirst has souped up his thesis with general expository matter has had not altogether satisfactory results. The textbook-like sections against which the accounts of his own work are set treat some topics, like structural ambiguity, at length, but do not form a very well-balanced or comprehensive whole. They nevertheless dilute Hirst's own work enough to prevent the reader from experiencing that feeling of excitement that good theses provoke, and to leave her disappointed in not getting the fuller and more concentrated account of Hirst's system as such, and its performance, she would have liked. The idea that language interpretation involves quite disparate processes is one deserving investigation, but it is not clear whether Hirst's particular way of combining such different lines as Montague and markers, via frames, is on the right track, and his book does not compel his reader, in a Pied Piper imperative, to follow him.

REFERENCES

- Alshawi, H. 1987 *Memory and Context for Language Interpretation*. Cambridge University Press, Cambridge, England.
 Carter, D.M. 1987 *Interpreting Anaphors in Natural Language Texts*. Ellis Horwood, Chichester, England.

Karen Sparck Jones is co-editor of *Automatic Natural Language Parsing* (Ellis Horwood 1983). Her present research concerns user modeling. Sparck Jones's address is: Computing Laboratory, University of Cambridge, Pembroke Street, Cambridge CB2 3QG, England. E-mail: sparckjones@cl.cam.uk.ac

NOTE: The author of this book is also the book review editor of this journal. Therefore, this review was edited by James Allen, editor of *Computational Linguistics*.

THE FIFTH GENERATION FALLACY: WHY JAPAN IS BETTING ITS FUTURE ON ARTIFICIAL INTELLIGENCE

J. Marshall Unger

(University of Hawaii, Honolulu)

Oxford University Press, 1987, x+230 pp.
 ISBN 0-19-504-939-X; (hb)

Reviewed by
Harold Somers
 UMIST

The main proposal of this book is that the motivation behind the Japanese Fifth Generation project is not a desire to push the barriers of computer technology research, nor even to attain economic superiority through advanced technology, but quite simply to overcome the problems of the Japanese writing system. In advancing this claim, Unger demonstrates that he is at least well read—an impressive array of literature in both English and Japanese is cited—and the book is well written, although sometimes in a journalistic rather than academic style. Despite that, this reviewer, and at least

four other informed colleagues that I have discussed this book with, find the claim quite preposterous. The book is an enjoyable read, in the way that the Letters page in an extremist newspaper is sometimes enjoyable, though I fear that only those who already have a fair background knowledge of the issues will get very far with it: if you don't know the difference between a kanji and a kana, then you might find it rather hard going.

The book is divided into three sections, each of two chapters, dealing with Japanese linguistics and orthography, political and cultural issues, and economics and technology. This is preceded by an Introduction which sets the scene, actually getting the book off to a very promising start, in which Unger virulently attacks the typical reaction to the Fifth Generation, as exemplified by the well-known Feigenbaum and McCorduck (1983) book, which exhibits paranoia, misunderstanding of Japanese attitudes to AI, the weak/strong AI debate, and the reasons for the Fifth Generation:

What the Japanese have in mind when they speak of AI, however, turns out to be something else yet again, and unless one dispels the fog of strong-AI hyperbole that surrounds the Fifth Generation project, it is impossible to appreciate its significance. (p.4)

It never occurs to [Feigenbaum and McCorduck] that, despite Japan's new-found affluence, its intellectual climate continues to suffer from such fundamental conditions as geographical and linguistic isolation, academic factionalism, and the overweening influence of government bureaucracy and giant corporations. (p.5)

It is as an answer to this misunderstanding that the remainder of the book is intended, and it begins with a detailed description of the Japanese writing system, and some of its practical consequences. As I suggested above, this whole section is probably too complex for the reader who does not know about it already, and this complexity is not helped by the unusual approach, starting from a romanization presented in alphabetical order and working towards the kana rather than vice versa. Some minor criticisms of this section would include the fact that the term *mora* (roughly, the consonant-plus-vowel syllable that each kana represents) is not included in the otherwise extensive glossary; the error (pp.21–22) of describing the Hepburn romanization as phonemic, when it is precisely not phonemic (e.g. /tu/ is written *tsu*), whereas the Cabinet (*kunrei*) romanization is; and there is no mention of a third romanization (especially used by Japanese) which is a simple transliteration of kana, e.g. *Toukyou* for *Tōkyō*. While we are on the subject of romanization, we mention also the quaint use by the publisher of the circumflex rather than the macron to indicate long vowels in the Japanese transcription.

Turning to the question of kanji, it is always difficult to convey accurately to readers used to alphabetic writing systems the pros and cons of a radically different orthography, and like most Western observers, Unger overestimates the problems they cause

(humans), and understates their advantages. When he states that

one can argue from tradition for the use of kana on computers, but any claim that they provide a better or more accurate representation of Japanese speech is completely without merit (p.28)

he is, of course, right, but if we substitute "language" for "speech" we do, in fact, get a fairly convincing argument.

Unger prepares us for the contents of the second chapter dealing with practical consequences (especially vis à vis computers), but it is difficult to know, when he talks of "technical disadvantages for data processing" (p.28) or "considerable cost implications for computer applications" (p.32) if he means for input, output, word processing, OCR, or what, since none of these has been mentioned explicitly up to this point.

The second chapter contains some good sections on the difficulties of ordering in dictionaries or indexes, especially due to individual idiosyncrasies (though he does not mention similar problems with Western dictionaries: are words listed in strict alphabetical order or are compounds listed under head words?). In both cases, humans quickly adapt, even if in Japanese it is a bit more complex.

Unger goes on to discuss (irrelevantly) the book trade, before turning to the much more crucial question of typing and computer input. His description of the different methods of kanji input (pp.64-65) is unclear, and in general this section is very shallow on what is clearly a vital issue. There is a wide range of people in Japan who take a keen interest in issues surrounding the use of word processors, which include the reticence on the part of older senior managers (who still prefer to pass handwritten documents to secretaries for typing, thereby emphasizing the office hierarchy) at one end, and the increasing use of obscure kanji made easily available, at the other, with numerous other issues (such as the demise of writing skills) in between. Apart from input, the storage and output problems are again overstated: more-or-less satisfactory solutions have been found here, and Unger overemphasizes (p.69) the speed and memory overheads. Regarding output, it could be said that the concentration on graphics means that the Japanese have had to make faster strides. If language processing software is so problematic for the Japanese, how is it that it is top of the list of software registered by commercial companies (p.75), and how does this compare with the West? Likewise, Unger seems to shoot himself in the foot when he states that "only 2% of Japanese people can operate a computer" (p.76). Again, how does this compare with the West? Unger's claim is that the Fifth Generation project was started in order to address these problems, but the evidence he presents suggests that the Japanese are well on their way to overcoming them anyway.

The two chapters that make up Part 2 of this book concern political and cultural aspects of the problem.

The first of these has a long section on literacy rates in Japanese followed by a discussion of some of the popular preconceptions about the writing system. The chapter is not engrossing, and its overall relevance is not brought out. The second chapter deals with the alleged conflict of technologies of the writing system on the one hand and computers on the other. Here Unger really shows his prejudices, dismissing the survival of the Japanese writing system over the generations as due to superstition on the part of users, but perpetrating his own narrow-minded superstitions making us wonder what his experience of non-numeric computing might be:

... kanji on computers waste system resources, impede software development, and unnecessarily complicate both computation and I/O functions. It is clearly wasteful to use thousands of intricate symbols when a few simple ones would suffice to convey the same information. The fuzzy definitions of the number, use, and form of kanji demand data-bound programs and guarantee that almost every program that does anything more with kanji than display them is a one-of-a-kind endeavor. Lack of a nonarbitrary collating sequence restricts the use of sorting, a fundamental step in countless applications

Kanji are a problem for computers because, by their very nature, they lack the properties that well-behaved computer data should have. (p.110)

In the third part of the book, Unger returns to the safer ground of hard facts and discusses current language processing technology in Chapter 5 and, finally, the Fifth Generation in the book's last chapter. The fifth chapter mainly concerns Japanese text input, and is quite interesting. Little mention is made of the possibilities of speech input: if some people are to be believed, advances in this area will in any case soon make many of Unger's concerns outdated, and, by the way, it could be said that the phonetic structure of Japanese makes it more amenable to speech processing than some other languages. However, here we are concerned with text input, of which there are three kinds distinguished by Unger: **inscriptive**, **transcriptive**, and **descriptive**.

Inscriptive input is where the user manipulates a stylus, the physical strokes being interpreted by the machine. Although closest to the traditional form of writing, this method is expensive and limited. One thing Unger does not mention is that this form of input is not useful for learners, since stroke order and weight is crucial: it is not sufficient to produce an apparently good copy of a kanji. OCR can also be regarded as a kind of inscriptive input, though OCRs for Japanese are not very robust, and handwriting is as yet out of the question (even for English).

Transcriptive input ("phonetic" typing with kanji conversion) is the most popular of the three, and here the issues are whether to use Roman or kana keyboards, the burden of shift-, control-, and alt-keying, and the efficiency of kanji conversion. Unger is rather negative on transcriptive input, though he overestimates the awkwardness of double keystrokes (at least to judge

from my own observations of native Japanese users of transcriptive input word processors). His discussion is principally concerned with alternative keyboard layouts, with rather less attention to the probably more important and interesting problem of conversion to kanji. Here Unger reveals his lack of familiarity with computational linguistics, belittling the role of partial syntactic analysis in kanji conversion. For us this is an important factor, especially if one thinks of machine translation, since transcriptive input with interactive kanji conversion (the system offers a choice of—or often a single—possible kanji for a given input sequence, which the user then confirms) can be seen as a type of pre-editing: choice of kanji effectively reduces homograph ambiguities that would otherwise further complicate the analysis.

The third type of input is called descriptive, since the user has to indicate the desired character via some non-linguistic device. This includes using a stylus to touch the appropriate cell on a tablet listing all the available characters; assigning (hopefully) mnemonic three-character codes to the kanji and then accessing them via the standard 30-key keyboard (30 keys give 27,000 codes, not 2,700 as Unger states (p.157)); and finally a technique that Unger calls “patternless keyboards”, but which, after three readings, this reviewer was unable to understand (halfway through there is a section heading “Romanization”, which seems misplaced). There are other possibilities of descriptive input that Unger does not mention, notably systems requiring the user to characterize the desired kanji in terms of its components—this might be using mnemonics, as with a system I saw demonstrated at Tokyo University, or on a workstation involving graphics, as in the system being developed here at UMIST.

The chapter ends with an attempt to look at the costs of kanji input, which Unger first measures by comparing the costs of PCs that support kanji input with those that do not: after a bit of juggling of which any “creative accountant” would be proud and which includes disregarding cheap word processors (loss leaders) and expensive PCs (no reason given) and recalculating median prices, Unger is obviously pleased to report that “the overhead of look-up dictionaries and transcriptive input software is substantial” (p.168).

And finally, the author deals directly with the topic featured in the book’s title. This chapter, the shortest of the six, does contain some interesting background information about how the project was set up, but mostly Unger is concerned once again to demonstrate his thesis concerning the connection between the Fifth Generation project and the Japanese writing system. But in all the citations of commentaries (by Japanese authors), which he offers as proof, there is barely one which really makes his point. He suggests that “the desire to use kanji with complete freedom on computers was the one technical dream that everyone involved took for granted” (p.172), but in doing so seems to ignore the

large percentage of work in the project totally or largely unrelated to kanji handling—especially work on vision, speech, and knowledge representation. He presents a series of quotes about the need for naive-user-friendly human-machine interfaces over the next few pages, consistently implying that these make his point: again not one of them explicitly mentions kanji, and can all be assumed to be referring to a wide range of human-machine interface devices. Unger goes on to talk about Japanese misunderstandings of Western AI and related fields, perhaps with some justification, but then perpetuates another misunderstanding in the other direction concerning the apparent secrecy of Japanese researchers: noting that results of ICOT research are rarely published, or aired at conferences, he misses the point that this is largely because of insufficient English: if one can read the Japanese language conference proceedings and journals, there is a mine of information to be had.

In conclusion, Unger talks a little about the aims of the Fifth Generation and the dangers of its failure. Like many commentators, he seems to suppose that the Japanese are simply very naive about these things: why else would they set up a 15-year project with fantastically ambitious goals? A different view is shared by some observers, including myself: of course the Japanese know that their goals are overly ambitious, and they quite readily admit this on occasion. However, even if they do not achieve them, they are going to achieve a lot of things incidentally on the way, and indeed are already doing so. And by the way, the same is true of a similarly apparently over-ambitious project in the machine translation field, namely the ATR telephone dialogue translation research project. How many Western governments would fund such long-term basic research programs? A colleague commented to me recently that he would not even know how to go about formulating such a long-term program. Unger cites another observer, commenting that

in one sense, it does not matter how the project turns out. As an exercise in the politics of Japanese R&D, an effort to enlist the cooperation of government bureaucrats, university faculty, graduate students, and the big corporations, it has already succeeded. (pp.191–192)

As a final word about this book, I note that its publication has received wide coverage in the lay press: I have seen it reviewed already in the *New York Times* and in *New Scientist*. But these reviews, like the book, do no more than propagate the myths: Japanese is hard to read and write (but English isn’t?), and the Japanese collude in this myth propagation because they like to think foreigners can’t learn their language. In this respect, Unger’s book is interesting. But the connection between this and the Fifth Generation is tenuous, and while the title is eye-catching and may promote sales, it is not appropriate to the content of the book. I have to borrow from the Scottish legal system to hand down the verdict: not proven.

REFERENCE

Feigenbaum, Edward A. and Pamela McCorduck 1983 *The Fifth Generation: Artificial Intelligence and Japan's Computer Challenge to the World*. Addison-Wesley, Reading, MA.

Harold Somers is a lecturer at UMIST, and a member of the machine translation research group there, with a special interest in Japanese MT. He recently spent 15 months in Japan as a visiting research fellow, first at NTT's Basic Research Laboratory, and then at Toshiba Corporation's R&D Center, during which time he was able to closely investigate Japanese NLP activities. Somers's address is: Centre for Computational Linguistics, UMIST, P.O. Box 88, Manchester M60 1QD, England. E-mail: hls@ccl.umist.ac.uk.

NATURAL LANGUAGE UNDERSTANDING

James Allen

(University of Rochester)

Menlo Park, CA: Benjamin/Cummings, 1987, xvi+574 pp.

ISBN 0-8053-0330-8; (hb)

Reviewed by

Michael Kac

University of Minnesota

Since the appearance of Tennant (1981), there have been a number of books intended to serve as introductions to computational linguistics, including Harris (1985), Grishman (1986), and Winograd (1983). The book under review is directed at much the same audience and is, in my opinion, the best so far. This is not to say that it is flawless; but its weaknesses are due more to the nature of the subject than to any failings of the author. (I should note, however, that I have not used it in a course, and so cannot report on student reactions to it.) The book is in four parts, entitled "Syntactic Processing", "Semantic Interpretation", "Context and World Knowledge", and "Response Generation", supplemented by two appendices, one on logic and one on symbolic computation; each chapter ends with a set of problems. It is broad enough in its coverage of a variety of approaches to give the reader a sense of the field of computational linguistics in its present state while managing to avoid a mere "Cook's Tour" approach (Pullum 1984).

The author is conscientious about regularly going into certain problems at a level of detail that makes for more than just a superficial overview; at the same time, the range of topics covered is sufficiently broad that anyone wanting to base a course on a cover-to-cover reading of the book had better have a full year in which to do it. Allen does not discuss how to actually formulate algorithms in some selected language, though many Lisp-like data structures appear throughout and the overall perspective is quite clearly that of contemporary AI. (The second appendix covers the essential ideas underlying Lisp and Prolog, though without taking the

form of a primer for either language.) A nice theoretical balance is maintained, and if there is any axe-grinding it is exceedingly subtle.¹

The treatment of syntax takes context-free grammar as the starting point and then considers both top-down and bottom-up approaches to parsing. Two different formalisms are pursued in some detail, namely transition networks and logic grammars, both with and without augmentations that take them outside the class of CFG-equivalent devices. One whole chapter is devoted to handling movement, and includes discussion of how (at the cost of increasing the size of the grammar) use of a hold facility can be avoided by using Gazdar-style slash categories. The problem of non-determinism is also discussed, and there is an overview of Marcus's parser.²

The section on semantics deals partly with representational issues (including what logical forms should look like), partly with the specifics of designing algorithms for semantic interpretation, and partly with a variety of open questions (including the especially thorny problem of correctly interpreting nominal compounds). A notion of "merging", closely akin to unification, is proposed as a basic operation for semantic interpretation and a detailed example showing the interaction of a syntactic parser and semantic interpreter is presented.

The section on context and world knowledge deals first with knowledge representation (a topic that is, of course, of more than just linguistic relevance), with problems in establishing reference, and with inferences based on knowledge. Anaphora is dealt with in two places, once in the chapter on reference (where the focus is on intrasentential relations), and once in the chapter on discourse structure (where the focus is on "long-distance" anaphoric relations, such as the use of a pronoun to refer to something mentioned some time earlier in a dialog). Finally, the last part contains a chapter on question-answering systems and one on "generation",³ the latter of note in part because of the attention it gives to systemic grammar.

I very much like the overall plan of the book (as outlined above) and think that the level of detail into which it goes is about right for a text of this kind. From the organization, I infer that the presumed audience is computer scientists with little background regarding natural language rather than linguists with an interest in applying what they know in a computational domain, which seems reasonable enough: in my own introductory CL course, the former group vastly outnumbers the latter (but then computer scientists as a group must outnumber linguists by an astounding factor anyway). I wonder, however, if the time is not fast approaching when we are going to have to make some linguistics beyond the introductory level a prerequisite for courses in CL, which would eliminate the need for a lot of background that Allen (and others, such as Harris) are forced to try to cram in as best they can. My own prejudices and disciplinary allegiances may be showing