Journal of Pragmatics 1 (1977) 171–176

REVIEW

Eugene Charniak and Yorick Wilks, eds., *Computational semantics. An introduction to artificial intelligence and natural language comprehension.* Fundamental studies in computer science, vol. 4. Amsterdam: North-Holland Publishing Company 1976. xiii + 294 pp. US $19.00/Dfl. 55.00.

1.

Since natural language processing within artificial intelligence (AI) is currently the most active branch of computational linguistics, a book like the one reviewed here deserves special interest on the part of linguists. Even more so as this book is the first introduction to AI which is entirely dedicated to language understanding and which explicitly deals with the linguistic aspects of some of the problems involved.

The book is a reader consisting of twelve articles previously used in a tutorial course on computational semantics at the Istituto per gli studi semantici e cognitivi (Lugano) in 1975. By interconnecting the various contributions with numerous helpful cross-references and an extensive subject index, the authors and editors have put the material into a systematical textbook format. The structure of most of the articles is the following: first, a certain problem is presented, then it is shown how this problem is handled in some of the existing or proposed language understanding systems, and finally the authors try to draw some general conclusions. In this way, the reader becomes acquainted with some of the basic theoretical issues, as well as with what has been achieved in the field so far.

The articles are grouped into five sections. In the introductory text of section I ('Inference and knowledge I') Charniak motivates the use of world knowledge and inferences in language analysis and describes some basic techniques. Section II ('Linguistics') contains contributions by Charniak ('Syntax in linguistics'), Philip Hayes ('Semantic markers and selectional restrictions'), Wolfgang Samlowski ('Case grammar'), and Margaret King ('Generative semantics'). Section III discusses some basic issues of the AI approach to natural language analysis in articles by Wilks ('Parsing English I & II'), Greg Scragg ('Semantic nets and memory models') and Charniak ('Inference and knowledge II'). In the next section, headed 'Related fields', we find a brilliant contribution by Walter F. Bischof ('Psychology of language and memory') and a discussion of the relevance of Montague's and Wittgenstein's philosophy for AI by Wilks.

The only article whose function in the book I really cannot see is the one in section V ('Computation') by King and Hayes, entitled 'Programming in LISP'. While there are many things that can be said on the subject of computation, the trouble with a contribution like this one (which is nothing, really, but an introduction to LISP programming) is that its target is unclear. Surely nobody is going to learn programming from a book like the one under review.

The remaining articles contain each some information of probable interest to every category of potential readers, though it is not quite clear here, either, who those intended readers are. Seemingly as little as possible is presupposed in the line of knowledge of the fields touched upon. For instance, we are told what a data base is; the predicate calculus is described; terms such as 'programming language', 'long term memory', and 'data base' are explained. Nevertheless, the book wants, and is, to be taken as a serious contribution to the theory and simulation of language understanding processes and should be reviewed as such.

The rest of this review takes up two topics which may be of special interest to the readers of this journal. We shall first be concerned with what is said about the role of linguistics in AI and then make some remarks on the problem of knowledge about language vs. knowledge about the world.

2.

Linguistics has not had much influence on natural language processing within AI. There has not even been much of a discussion on the relations between the two fields. Since this situation is of course unsatisfactory, every attempt to come to a better mutual understanding deserves attention. The discussions in our book *are* useful in this respect because they show – though mostly involuntarily – some of the reasons for the poor cooperation between linguistics and AI-language processing. By far the most important reason is that linguistics is simply identified with generative linguistics. The unfortunate consequence is that all language phenomena are interpreted against this background. Since some of the basic tenets of generative grammar have never been accepted in AI, this in turn necessarily leads to misunderstanding and a permanent confusion as to what are facts about language and what are facts about generative linguistics.

As an instance, consider Charniak on the usefulness of transformational grammars for language understanding systems. Like many workers in AI, he presupposes that generating all and only English sentences is the main task of a transformational grammar of English; so he suggests its use in order "to have a machine produce grammatical English itself" (p. 37). To give just one further example, King calls a generative grammar "a means for determining which sentences are grammatical" (p. 88). This is a profound misunderstanding. For more than twenty years now, generative linguistics has considered as its most important issue, the question of *how* to generate the sentences of a language, i.e. how to formulate the rules of the grammar

under the assumption that there are infinitely many possibilities for doing so with respect to the grammatical–ungrammatical distinction. Now you cannot buy a transformational grammar without buying the structures it produces. When Charniak says that 'Raising' is a useful device because it tells us that *Jack* is the subject of the embedded sentence in *Jack was believed by the boy to have won the contest*, he should have been aware of the fact that 'Raising' maps one structure onto another structure. If these structures are not the kind he wants for his language understanding system, then he should not talk about 'Raising', but about the traditionally well known fact that *believe* belongs to a class of verbs whose grammatical properties are such and such. This has nothing to do with transformational grammar. If on the other hand transformational grammar *is* useful for AI-language processing, then we should expect him to tie in with the linguistics-dependent language processing of the sixties Charniak does not even mention this tradition with all its theoretical and practical work.

Similar criticisms hold for the following two articles. I don't see any reason for Hayes' extensive presentation of the oldest version of the Katz-Fodor semantic theory. Semantic restrictions on the co-occurrence of words in complex expressions had been noted in structural and traditional linguistics long before Katz and Fodor. It is an almost trivial fact that language understanding systems have to deal with restrictions of this kind if the analysis is not to ignore the level of word semantics altogether. Again, this has nothing to do with generative linguistics. What is specific for the Katz–Fodor theory is its classification of semantic elements and the way these elements are combined to form semantic representations of complex expressions. None of the systems which Hayes mentions (Winograd, Riesbeck, Wilks) makes use of this theory, so the reader does not know what he is supposed to learn from the whole discussion. The same holds for Samlowski's article on case grammar. It is true that in language understanding systems (Simmons, Schank, Wilks) we often find structures which *formally* resemble the deep structures of Fillmore's case grammar. But case grammar is a transformational grammar and you cannot buy the form of its deep structures without buying its transformations. Samlowski is very explicit as to why Fillmore postulated his case grammar (including the transformational part), but he tells us nothing about the reasons why workers in AI did not accept this grammar. Instead he reminds us that case grammar has "an important influence on AI" (p. 55) and that there is a "constant interaction between work in theoretical linguistics and AI" (p. 72), which actually is not the case.

King's contribution on generative semantics is the only one that not even attempts to demonstrate that this theory is relevant for AI (though this has sometimes been claimed, especially because of Lakoff's lexical decomposition rules). The reader does not learn why he should be concerned with this kind of linguistics at all. Yet I think this is better than to be told that there is an influence of linguistics without being told what exactly it consists in.

3.

Workers in AI have always stressed the fact that a machine which is to understand natural language has to have access to what is called knowledge of the world. Nevertheless it seems obvious – and well known in AI – that it is at present not clear where we have to look for the borderline between language and facts. As a typical example, take Scragg's statement that it is "occasionally necessary to represent items in a data base which are not known to be true, or even that are known to be false" (p. 108). In storing a sentence like *Peter said that he went to the store* one has, according to Scragg, to state *as a fact in the data base* that it is unknown whether Peter really went because we "know that Peter said this but not whether he actually went" (p. 108). It seems to me that this kind of approach neglects a clear difference between knowledge and semantics. The fact that we do not know whether Peter *really* went is implicit in the meaning of *say*, and should therefore not be stated as a fact by itself in the data base. Of course we can state elsewhere in the data base whether or not Peter went, and whether or not we know that he went or didn't. But this has nothing whatever to do with what is said in the quotation above.

Consider now what Wilks has to say about Winograd's procedural formulation of word meanings. For Winograd, the meaning of a verb is represented by a series of commands describing what to do when the verb is applied to one or more appropriate objects. Wilks' question is whether something like this can be considered as a meaning representation, or must be taken as the description of a special use of that verb. The consequence of the latter alternative would be to regard a system like Winograd's as being "not *about* natural language at all, but about the other technical question of how goals and subgoals are to be organized in a problem-solving system capable of manipulating simple physical objects" (p. 99).

I think one important reason for the uncertainty just illustrated is that we have two different modes of fact usage AI. These modes are not always seen as distinct; nevertheless, they have completely different consequences as to what can be said about 'language vs. the world'. First, there is the case where knowledge is exclusively used for language analysis in its narrower sense, i.e. for resolving word ambiguities, structural ambiguities, referential ambiguities of pronouns and the like. What world knowledge does in these cases is to make use of presupposed contextual features telling us that an expression will probably be meant one way and not another. The corresponding knowledge-based rules, such as Charniak's 'demons' (cf. pp. 129ff.), are therefore rules which hold only with a certain probability. Hence, Wilks is clearly wrong when he says that in the sentence *The soldiers fired at the women and I saw several fall* "the meaning of the pronoun is perfectly clear" (p. 99). It is easier to construct a context for the reading "several women fell" than to construct one for the reading "several soldiers fell" (or even other readings, maybe), but that is all there is to it. This mode of fact usage thus selects some relatively arbitrary contextual features for resolving ambiguities. Nothing can be said about the reasons for which a certain reading *must* be favored.

Second, there is the mode of fact usage, where we do not just have some world *knowledge* as such, but where indeed we do have a *world* from which our knowledge is drawn, as is the case in Winograd's system. Here we are in the situation of knowing two things: one is the relation between facts in the world and the machine's knowledge of the world. This is the identity relation, since the worlds of blocks at given time is nothing but a special form of part of the machine's memory. Secondly, we know for every fact of the memory how it can be verbalized by the machine. We can now communicate with the machine *if* we have access to the same world as the machine has, *if* we regard the same entities as facts in this world as the machine does, and *if* we verbalize these facts in the same way as the machine does. In my opinion, the crucial point if Winograd's system is in this common access to world facts and their common verbalization. Wilks reminds us of later Wittgensteinian philosophy's dictum: one should not take a mini-language and assume its properties (that may be appropriate to a mini-world) to be properties of natural language as a whole, *precisely because* we do not know how this language would function in a more complicated world. If we are to take this admonition seriously, then Winograd's system is neither about facts nor about language as such, but about some well-defined language with respect to some well-defined world; "with respect to" meaning what has been said above about the two's relation (cf. also p. 224). Of course there is nothing new in this statement. The reason for its reiteration here is that, in my opinion, many workers in AI do not seem to be aware of its necessary consequences for their practical work. Thus, we are often confronted with statements to the effect that there exists something like an abstract of language understanding which nobody possesses, but which can be investigated relatively independently from world knowledge. World knowledge would then come in when, and where and to the extend that, it is needed to support the process of understanding. In Wilks' words, "The point of view...has been that these representations [of knowledge, P.E.] must be justified in terms of some concrete problems that they solve, such as word-sense or pronoun reference ambiguity, otherwise they may not be essentially connected with the understanding of natural language" (p. 181). If the machine is supposed to communicate in natural language (i.e. also to act by using natural language), then this view does not hold. The machine would not only have to *know* everything in the world it wants to talk about, but it would have to *be* in the world. However, the expression "to know everything" does not imply the old Bloomfieldian view according to which we are only able to say something about meanings if we know how the world 'really' is. But it does mean that we know how the natural language user understands the world and hence organizes his language: we are not interested in knowing that a whale is a mammal but in knowing that it is a fish.

What this means in practice is that we will learn less (in the longer run) about natural language understanding from systems which translate (such as Wilks'), or paraphrase and draw consequences (such as Schank's), than from systems which analyze language within a specified context (such as Charniak's children's stories).

And we will learn even more from systems which talk about, and act in, a world which can be perceived from outside the machine, and which is accessible to the machine by all sensory channels. I really think that much of the discussion on what language understanding machines should be able to do would be cut short, if one realized that the only machine that can understand language is a robot.

4.

Let me conclude this review with a critical remark on *computational semantics*. We are told that it is new, that it draws elements from different fields and that there is no consensus on many questions *within* computational semantics. But what *is* computational semantics? Sometimes it looks like a theory: "A chapter on programming clearly has a place in a textbook on computational semantics. If one claims to be constructing theories of natural language comprehension..." (p. 235). Sometimes it looks like a discipline: "Computational Semantics, the name we have given to the study of language based upon Artificial Intelligence methods... assumes that 'language is as language does', an idea not unknown to the older disciplines of linguistics, psychology, etc..." (p. 1). And sometimes it looks like a name for certain methods: "Computational Semantics is not so much a new subject as a new way of looking at old questions...". "Computational Semantics is the best way we have of solving the difficult problem of language comprehension" (Editors' Preface).

Peter Eisenberg
Born 1940. Studied computer science and linguistics in Berlin and Cambridge, Mass.; presently teaching linguistics at the Technische Universität, Hannover.
Publications related to the topic of this review:
1976. Oberflächenstruktur und logische Struktur. Tübingen: Niemeyer.
1976. Die Bedeutung semantischer Theorien für die künstliche Intelligenz. Studium Linguistik 2: 1–23.
1977. Zum Begriff der syntaktischen Mehrdeutigkeit. Linguistische Berichte 48: 28–43.
Editor of:
Beiträge zur automatischen Sprachbearbeitung. Berlin, New york: de Gruyter. (Vol. 1: Maschinelle Sprachanalyse, 1976; vol. 2: Semantik und künstliche Intelligenz, 1977.)