Linguistics, Malostranské n.25, 118 00 Praha 1, Czechoslovakia.

Philosophy, Language, and Artificial Intelligence: Resources for Processing Natural Language

Jack Kulas, James H. Fetzer, and Terry L. Rankin (eds.)
(University of Idaho, Moscow; University of Minnesota, Duluth; and IBM AI Support Center, Palo Alto, CA)

Dordrecht: Kluwer Academic Publishers, 1988, xii + 421 pp.
(Studies in Cognitive Systems)
ISBN 1-55608-073-5, $99.00, £58.00, Dfl 195.00 (hb)

Reviewed by
Peter Ludlow
State University of New York at Stony Brook

Sometimes philosophers and linguists say things that are of interest to computer scientists or would be of interest if properly explained. Unfortunately, philosophical work and theoretical linguistics are seldom explained in ways that make them seem relevant to computer scientists. That is, they are seldom explained in ways that suggest directions for future research. Given this communication problem, there is a definite need for a volume that can bring together some of the best work in the philosophy of language, linguistics, and natural language processing, and show how work in the philosophy of language and linguistics has been and can be applied to problems in natural language processing. Unfortunately, despite the title[1] and the promotional literature, this volume does not fulfil the need. Part of the problem is that none of the papers in the volume describe work in which the theoretical insights of philosophers and linguists are *implemented*. There are simply no papers on natural language processing. There is an introduction by Kulas that is supposed to show why these essays are important to AI researchers. But while the introduction occasionally provides reasonable summaries of the essays, it doesn't come close to suggesting applications of the theoretical work to AI. It may well be that computer scientists ought to be familiar with the essays in the book. Most of the essays are, after all, classics. But it is pedagogically naive to think that one can drop, for example, Davidson's "Truth and Meaning" in the lap of a computer scientist and suppose that it suggests anything in the way of a research program.

Another part of the problem is that the volume has been limited to papers that originally appeared in books or journals published by D. Reidel. This is unfortunate, for a number of sections of the book could have been greatly strengthened with the addition of material published elsewhere. Now it may be that the editors intended to provide something like "D. Reidel's Greatest Hits", but such a strategy seems to me misguided. There is nothing intrinsically interesting about the fact that a paper was first published in a Reidel journal. It may be that the move was inspired by the desire to hold down costs,[2] but if the editors were successful in holding down costs, the savings certainly were not passed along to the customer. The price of the 12-essay volume is $99!

Thus far my comments have been strong, but general. I think a section-by-section discussion of the book may be in order, to point out specific concerns and to suggest ways in which the collection might have been improved.

The book begins with a prologue entitled "Modes of Meaning", which contains Grice's "Utterer's Meaning, Sentence Meaning, and Word Meaning"—a great paper, but what it is supposed to suggest to the AI researcher is unclear. Nor is Kulas's introduction much help here, as he offers only a one-sentence summary of the paper.

Part I of the book is entitled "Formal Syntax of Natural Language", but is really about the question of whether natural languages are context-free. The first essay has no place in this volume. It is a short piece that Pullum wrote as a "Topic . . . comment" column for *Natural Language and Linguistic Theory*, which attempts to sort out who was the first to show that natural language is not context-free. The piece is not, nor do I believe it intended to be, a serious piece of academic research. (Pullum has even expressed amusement that "Topic . . . comment" columns have been indexed and abstracted.)[3] This is not a criticism of Pullum; he has other work on the problem of the context-freeness of natural language (e.g., Pullum 1983) that would have made much better additions to the volume. For that matter, Gazdar and Pullum (1985) would also have been a better choice. The second piece in Part 1 is Shieber's "Evidence Against the Context Freeness of Natural Language", which is a solid piece showing how Swiss-German is weakly non–context-free. But Shieber would be the first to admit that the implications of the result for natural language processing are of no great moment.

Part 2 of the collection is entitled "Semantic Aspects of Natural Language" and contains Davidson's "Truth and Meaning" and Hintikka's "Semantics for Propositional Attitudes". I've already noted that the import of Davidson's paper (which makes no concessions to the non-philosopher) is never made clear. As for Hintikka's paper, it isn't really about semantics, but is an application of possible-world semantics to the problem of the attitudes. It more properly belongs with the Perry and Stalnaker papers in Part 5. What is unfortunate about this part of the collection is that there really has been confusion between philosophers and computer scientists about what semantics is—a confusion that could have been straightened out here. Most computer scien-

tists construe the semantics of a term as a data structure that represents conceptual dependencies of the term. But philosophers (e.g., Putnam 1975) take semantics to be about how the terms hook up with the world. For Putnam and others meanings just aren't in the head. So, strictly speaking, you can't *implement* a semantics, for semantics describes the relationship between states of your machine and the world.

Part 3 is entitled "Connecting Syntax with Semantics", and begins with Montague's "The Proper Treatment of Quantification in Ordinary English". The paper does, of course, speak to the relationship between syntax and semantics, though it could perhaps be included in a section called "Doing Away with Syntax". Computer scientists are generally aware of Montague's work, and they generally regard it as computationally intractable (e.g., see Hirst 1987, p. 32). There is no mention of this in the introduction, however, nor is there a discussion of attempts to implement versions of Montague grammar. Friedman, Moran, and Warren (1978a, 1978b) should be discussed if not included, but the work doesn't even make the bibliography.

The second piece in this section is Gazdar's "Phrase Structure Grammar". While there is a standard line about applying Montague's semantics to GPSG, this particular paper does not provide the best statement of the relationship. Moreover, Gazdar's position on the relationship between syntax and semantics is essentially the same as Montague's. Both think that there is an isomorphism between syntax and semantics, the chief difference being that Montague employs a categorial grammar syntax while Gazdar employs GPSG. The back-to-back appearance of these papers gives the impression that this view is uncontroversial, but nothing could be further from the truth. Higginbotham, to name one philosopher with reservations about Montague's program, is discussed in the introduction, but Higginbotham (1985) really ought to be included in a collection of this nature.

In my opinion the Gazdar paper (which is fine as an introduction to GPSG) really belongs in the section entitled "Formal Syntax of Natural Language", but that heading apparently meant something else to the editors. One also has to wonder why, if a paper on GPSG is to be included, there is nothing on LFG, which has been at least as important as GPSG in natural language processing. Moreover, since GB is the most widely employed linguistic theory and has driven NLP research by Berwick and others in MIT's AI lab, one wonders why it is not represented in the collection.

Part 4, "Natural Language and Logical Form", contains two important essays; Barwise and Cooper's "Generalized Quantifiers and Natural Language", and Hintikka's "Quantifiers in Natural Language: Some Logical Problems I". They are important, but important in pointing out the limits of first-order theories of quantification in accounting for the semantics of natural language. The implications of this for AI are never made clear. Moreover, one might suppose that computer scientists will not be concerned if their parsers cannot handle Hintikka branching constructions like *Some product of some subdivision of every company of every conglomerate is advertised in some page of some number of every magazine of every newspaper chain.* The nature of the logical form of natural language is a much broader problem than these two essays alone suggest. This is attested to by work on logical form done within the GB framework and by extensive work in the Davidsonian tradition. Again, the essays included here are excellent, but they belong in a section on nonelementary quantification or on the limits of first-order quantified logic.

Part 5, "Possible Worlds and Situation Semantics", contains essays by Stalnaker and Perry in which they debate (among other things) whether possible worlds or situations provide objects of the appropriate granularity for the purposes of semantics. Together they supply some much-needed clarification to the debate between possible-world semantics and situation semantics. However, the consequences of the debate for AI are far from obvious, and go unexplained.

The epilogue, "From Semantics to Pragmatics", contains Kamp's "Semantics *vs.* Pragmatics", an essay that Kulas summarizes in one sentence as being about "the ebb and flow of ordinary conversations in ordinary language". I, on the other hand, took the essay to be providing a formal pragmatic theory of the illocutionary force of commands and permission statements.

The upshot is that the collection would more appropriately have been titled *Formal Philosophy of Language*, but even as a collection in the philosophy of language it is poorly structured and contains some significant gaps. This is not to criticize the individual papers. As I noted earlier, they are almost all classics. However, they are also readily available in other places. Why is "artificial intelligence" in the title? Probably for marketing reasons. Few philosophers or linguists can afford to lay out $99.00 for a book. One final note: The advertisements for the book and the note on the back cover suggest that work on discourse representation theory is covered, but this is simply not the case.

Acknowledgments

Thanks are due to Stuart Shieber and David S. Warren for helpful discussion. Omissions and errors are my own, and the opinions expressed here are not necessarily theirs, but not necessarily not theirs.

REFERENCES

Friedman, Joyce; Moran, Douglas; and Warren, David S. 1978a Evaluating English Sentences in a Logical Model: A Process Version of Montague Grammar. In *Proceedings of the 7th International Conference on Computational Linguistics*, Bergen, Norway.

Friedman, Joyce; Moran, Douglas; and Warren, David S. 1978b Explicit Finite Intensional Models for PTQ. *American Journal of Computational Linguistics* (microfiche) 74; 3–22.

Gazdar, Gerald and Pullum, Geoffrey 1985 Computationally Relevant Properties of Natural Languages and Their Grammars. *New Generation Computing* 3; 273–306.

Higginbotham, James 1985 On Semantics. *Linguistic Inquiry* 16; 547–593.

Hirst, Graeme 1987 *Semantic Interpretation and the Resolution of Ambiguity*. Cambridge University Press, Cambridge, England.

Pullum, Geoffrey 1983 Context Freeness and the Computer Processing of Human Languages. In *Proceedings of the 21st Meeting of the Association for Computational Linguistics*, Cambridge, MA: 1–6.

Pullum, Geoffrey 1986 A Guest of the State. *Natural Language and Linguistic Theory* 4; 284–289.

Putnam, Hilary 1975 The Meaning of 'Meaning'. In *Mind, Language, and Reality: Philosophical Papers 2*, Cambridge University Press, Cambridge, England.

Notes

1. The title shown above is that given on the title page and back cover, but its listing with the Library of Congress is the slightly more appropriate *Philosophy, Language, and Artificial Intelligence: Philosophical Resources for Natural Language Processing*.
2. Not only was the need to pay royalties avoided, but typesetting expenses seem to have been spared as well. For the most part the essays appear as they were originally typeset, pagination being the only change. This policy is unfortunate, for mistakes have surfaced in some of these writings which ought to be noted in editorial footnotes, if not corrected. For example, there is a slip in Montague's paper which is reprinted without mention here (the first meaning postulate on p. 157 contains a biconditional instead of a conditional—fine in the case of intransitive verbs, but not for common nouns). This "bug" has infested some implementations of Montague grammar.
3. Perhaps, under the circumstances, I can be forgiven for citing a "Topic . . . comment" column myself.

   > This oasis, this last respite from the seriousness of professionalized linguistics, this space reserved at the end of each *NLLT* for a piece of writing flippant and inconsequential enough to give an exhausted assistant professor the strength to get up and do what needs to be done, had been mistaken for mere research, and was being indexed and abstracted (Pullum 1986, p. 288).

Peter Ludlow is an assistant professor of philosophy. His current research interests are in the philosophy of language and of cognitive science. He has previously worked on projects in natural language processing for Honeywell's Corporate Systems Development Division. Ludlow's address is: Department of Philosophy, SUNY at Stony Brook, Stony Brook, NY 11794. E-mail: PLUDLOW@SBCCVM.bitnet

---

Systemic Text Generation as Problem Solving

Terry Patten
(Department of Computer and Information Science, The Ohio State University)

Cambridge University Press, 1988, xii + 214 pp.
(Studies in Natural Language Processing)
ISBN 0-521-35076-X, $34.50 (20% discount to ACL members) (hb)

Reviewed by
Eduard Hovy
USC Information Sciences Institute

In thinking about the problem of generating English sentences, it is very reasonable to wonder whether the natural paradigm of the solution is not perhaps that of planning or problem solving. (I use the terms here interchangeably; the distinctions between the two enterprises are small enough to be ignored in this context.) That is to say, given that you start with some specific state, namely the input, and want to end with some other state, the corresponding sentence, and you have a set of rules, namely the grammar, that prescribe how to bring about the change, then why not simply use a planner or problem solver to do the job?

This is exactly what Terry Patten's book is about. It does not address all the issues involved, but it shows convincingly that you can implement a sentence generator in the problem-solving paradigm and it provides as well some very useful information on systemic grammar: both a formal definition of systemic networks and a description of a particular implementation of one.

The book consists of three parts: introductory material, the core idea, and the rest. The introductory material contains a chapter describing AI problem solving and a chapter describing systemic grammar. The core material first describes how one can interpret a grammar as information with which one can perform problem solving, and then provides a formal model of systemic grammar. In the remainder, a particular implementation of a problem solver using systemic grammar, called SLANG, is described and compared with other generators, and parts of SLANG's grammar and samples of its output are provided.

Chapter 3, the background chapter describing systemic grammar, is a simple introduction to a body of thought that often has been called impenetrable. Given the breadth of application of systemic linguistics, an understanding of it is required for anyone who wishes to venture beyond the narrow view of language as taken by the various generative paradigms. Since it focuses on implementational issues, this chapter offers, specially tailored for the computational linguistics community, one of the most readable descriptions of the systemic view of language I have yet encountered. Unfortunately, this fact makes the chapter less suitable as a general overview of the field, since it gives no indication of the depth of linguistic research that underlies the ideas it describes.

Chapter 4 leads the reader through the correspondences between problem solving, which is a search process through a space of alternative states, and systemic language generation, which is a search process through a network of alternative meanings (i.e., aspects of sentences). The network's grammatical choice points, called **systems**, are implemented by production rules. In the way illustrated, AI techniques for handling huge spaces of interdependent alternatives efficiently can be applied to fine-grained grammatical distinctions identified by systemic (or other functional) linguists.

Chapter 5 contains a somewhat exploratory formal