

even inconsistencies(!), among the different authors. Anthologies, in a discipline such as philosophy of language, are better in that they (usually) contain primary sources. My own preference when teaching philosophy of language is not to use a single-author text (except at most to supplement the primary sources and to serve as a guide to the problems and the literature, for the student who prefers such a guide). Rather, I have the students read original sources, while I provide background, connecting material, and explications in lectures.

To sum up, as a "single"-author text, Devitt and Sterelny's book is probably better than Martin's, but the appropriate audience for it (advanced undergraduates at the very least, graduate students (or beyond) at best) could do as well with an anthology. It would certainly serve as an excellent, if somewhat idiosyncratic, supplement to an anthology. Martin's book would be better for (primarily undergraduate) students who need the security of a single-author text, but the instructor would *need* to correct the errors along the way. It could, in any case, be usefully supplemented by an anthology of primary sources.

ACKNOWLEDGMENT

The preparation of this review was supported in part by the National Science Foundation under Grant Nos. IST-8504713 and IRI-8610517.

REFERENCES

- Ashby, R.W. 1967. Verifiability Principle. In Edwards, P. (Ed.), *Encyclopedia of Philosophy* 8. Macmillan and Free Press, New York, NY: 240-247.
- Castañeda, Hector-Neri 1967. Indicators and Quasi-Indicators. *American Philosophical Quarterly* 4: 85-100.
- Castañeda, Hector-Neri 1968. On the Logic of Attributions of Self-Knowledge to Others. *Journal of Philosophy* 54: 439-456.
- Castañeda, Hector-Neri 1972. Thinking and the Structure of the World. *Philosophia* 4: 3-40. Reprinted in *Critica* 6, 43-86.
- Church, Alonzo 1949. Review of Ayer's *Language, Truth and Logic* (2nd ed.). In *Journal of Symbolic Logic* 14.
- Frege, Gottlob 1892. On Sense and Reference. Black, M. (Trans.). In Geach, P. and Black, M. (Eds.). 1960 *Translations from the Philosophical Writings of Gottlob Frege* (2nd ed.). Basil Blackwell, Oxford, England, 56-78.
- Lakoff, George 1987. *Women, Fire, and Dangerous Things: What Categories Reveal about the Mind*. University of Chicago Press, Chicago, IL.
- Maida, Anthony S. and Shapiro, Stuart C. 1982. Intensional Concepts in Propositional Semantic Networks. *Cognitive Science* 6: 291-330.
- Parsons, Terence 1980. *Nonexistent Objects*. Yale University Press, New Haven, CT.
- Rapaport, William J. 1978. Meinongian Theories and a Russellian Paradox. *Noûs* 12: 153-80.
- Rapaport, William J. 1981. How to Make the World Fit Our Language: An Essay in Meinongian Semantics. *Grazer Philosophische Studien* 14: 1-21.
- Rapaport, William J. 1985. Meinongian Semantics for Propositional Semantic Networks. *Proceedings of the 23rd Annual Meeting of the Association for Computational Linguistics*. University of Chicago, Chicago, IL; 43-48.
- Rapaport, William J. 1987. Syntactic Semantics: Foundations of Computational Natural-Language Understanding. In Fetzer J. (Ed.), *Aspects of Artificial Intelligence*. D. Reidel, Dordrecht, Holland.
- Routley, Richard 1979. *Exploring Meinong's Jungle and Beyond*. Department of Philosophy, Research School of Social Sciences, Australian National University, Canberra, Australia.
- Shapiro, Stuart C. and Rapaport, William J. 1987. SNePS Considered as a Fully Intensional Propositional Semantic Network. In McCalla, G. and Cercone, N. (Eds.). *The Knowledge Frontier: Essays in the Representation of Knowledge*. Springer-Verlag, New York, NY; 262-315.
- Zalta, Edward 1983. *Abstract Objects*. D. Reidel, Dordrecht, Holland.

William J. Rapaport holds graduate degrees in both philosophy and computer science. His present research includes intensional knowledge representation and narrative deixis. He has been at SUNY Buffalo since 1984. Rapaport's address is: Department of Computer Science, SUNY Buffalo, Buffalo, NY 14260. E-mail: rapaport@CS.buffalo.edu.

NATURAL LANGUAGE GENERATION: NEW RESULTS IN ARTIFICIAL INTELLIGENCE, PSYCHOLOGY, AND LINGUISTICS

Gerard Kempen (ed.)

(Department of Experimental Psychology, University of Nijmegen)

(NATO Advanced Science Institutes Series E: Applied Sciences, No. 135)

Martinus Nijhoff Publishers: Dordrecht, The Netherlands (distributed by Kluwer Academic Publishers), 1987, xiv+466 pp.

ISBN 90-247-3558-0, \$79.50 (hb), Dfl 195., £53.50

Reviewed by
Marie Bienkowski
SRI International!

Natural Language Generation is a collection of papers that were presented at the Third International Workshop on Natural Language Generation in Nijmegen, The Netherlands, on August 19-23, 1986. Instead of a soft-cover proceedings, the workshop contents are captured in this hardcover book containing edited versions of the papers. The contributions are from computational linguistics, linguistics, and psychology. In the preface, Kempen, the editor, states that the interactions among workshop participants demonstrated how much these different disciplines share. Unfortunately, the interactions do not appear to be reflected in the edited versions of the papers, even though they might have been of interest to non-attendees.

Language generation research has been viewed as the poorer cousin of work on language understanding. This has been true of computational work as well as psychological research. People sometimes claim that until computers have something to talk about, language generation is not worth studying. Or, they assert that language understanding is much 'harder', so is more deserving of attention. This book presents the work of researchers who have ignored such pronouncements,

and who have been worrying about a variety of hard problems concerning 'what to say and how to say it' for some time. Those concerned with computational modeling have looked for any information that could be expressed in language. And they've found it, in data bases, tic-tac-toe games, stock market reports, primary election reports, visual scenes, newspaper reports, encyclopedic knowledge, and many other places.

Language generation is becoming especially popular because of the increase in interest in explanation for intelligent systems. People are creating knowledge bases that can support explanations of reasoning and examining how explanations should proceed and get translated into natural language. But work on natural language generation is not focused exclusively on explanation, as this book shows. The papers here are organized under six headings: pragmatic aspects, generation of connected discourse, generator design, grammars and grammatical formalisms, stages of human sentence production, and aspects of lexicalization.

The book contains descriptions of some state-of-the-art work on language generation. It is probably not suitable for someone unfamiliar with problems in natural language processing or linguistics, because the contributions are quite varied. Better sources for introductory material are Cullingford (1986), McKeown (1985), and Grishman (1986). A few of the contributions are too specific and detailed, a few very general and vague. Some report computational work, others do not. Some address very general issues that are of importance and interest to all natural language researchers (such as the contributions on pragmatic representations, levels of processing, and systemic grammar formalizations), but others address very narrow issues (such as generating answers from linguistically coded data bases). However, the book is a good source of some important work, including contributions from European researchers that might be hard to find elsewhere.

Because of the length of the book (29 papers on 29 different topics), I will mention only a few from each section. My selections are not representative of the book as a whole, of course, but will give a flavor of the type of contributions. I chose well-written papers that present contributions of general interest. My overall impression of the book is that publishing it as a soft-cover proceedings would have been more suitable, because of the wide variations in content and quality.

The first section, on pragmatic aspects, begins with Eduard Hovy's excellent work on introducing pragmatic information, such as affect and speaker-hearer relationship, into a computational model of generation. His idea is to link pragmatic goals with syntactic realizations through rhetorical goals, such as verbosity, formality, and haste. An important result of his work is his description of how rhetorical goals are used to affect five points in the generation process: topic collection, topic organization, sentence organization, clause con-

tent, and lexicalization. This paper gives a good overview and several examples of his work.

Kathleen McCoy's paper on responding to misconceptions presents a succinct overview of her system and its treatment of the classification and response to a user's misconceptions. She describes how the decision of whether to respond, the general way to respond, and the specifics of what to say in the response, are affected by a model of the user's misconceptions.

Doug Appelt's contribution is an addition to his work on the generation of referring expressions. He presents a description of referring represented in modal logic, and precisely defines concepts like sincerity and competence. He also presents a schema for determining the beliefs that a speaker and hearer hold after a referring act, this being the ultimate goal of referring. The paper following Appelt's, by Norbert Reithinger, examines the production of referring expressions in a full dialog system, complete with a graphical system for generating actual pointing actions. His work is important in expanding the discourse context of a generator to include a shared visual situation.

Cecile Paris and Kathy McKeown's paper (in the next section on connected discourse generation) describes an extension of McKeown's familiar discourse strategies. The extension adds process strategies for traversing causal nets to produce narrative-like descriptions of complex physical objects. The implementation they describe has a facility for switching between the new procedural strategies and the old declarative ones.

Hans-Joachim Novak's paper describes discourse strategies for generating object movement descriptions. Interestingly, he uses a method that tries to model how a description might be visualized by a hearer. He also proposes some practical methods for generating referring expressions. His work is significant in tying generation to real-world object representations.

The third section is on generator design. It begins with an excellent paper by Dave McDonald, Marie Vaughan, and James Pustejovsky. They propose an abstract reference model for generation, to enable comparison among generators. The reference model, simply stated, is to identify the speaker's situation, map it onto an utterance, and then read out the utterance. This model is used to structure a discussion of factors in generator design which contribute to efficiency. The factors they discuss are precomputation of structure, the size of steps between representational levels, exploitation of regularities in natural languages, efficient control, and delayed evaluation of information.

Two other papers in this section describe specific generation systems. Laurence Danlos is concerned with determining what linguistic data should be used in what part of the generation process. She gives a very detailed treatment of various types of clause production, agreement rules, deletion of repeated elements in coordinates, and more in a description of her French and English generator. Paul Jacobs is also concerned with

representation of linguistic knowledge, although his main concern is extensibility. His English generator, KING, uses a simple control scheme that exploits the rich linguistic representations in a separate, frame-based, hierarchical system. The section on grammars and grammatical formalisms present detailed papers on everything from the relevance of Tree Adjoining Grammars to generation (by Aravind Joshi) to a formal model of systemic grammar (by Terry Patten and Graeme Ritchie). There is also a detailed description of a generator, by Harry Bunt, that uses pragmatic information.

The final sections primarily contain the contributions of the psychologists. Koenraad De Smedt and Gerard Kempen propose what is surely the first true computational model of sentence production that mimics the incremental nature of human production. Their model, which includes a monitoring component, captures various phenomena, such as hesitations, syntactic deadlock, and self-corrections, including modifications to conceptual structures.

Another contribution by Willem Levelt and Herbert Schriefers, explores stages of activation of lexical properties such as sound form and conceptual conditions. One of the more interesting aspects that they address is how a lexical item checks if its conceptual conditions are satisfied. They extend the earlier idea of matching a core sense to include checks on specificity of meaning.

One final paper is worth mention. Karen Kukich's paper presents a connectionist implementation of part of her stock market report generator. This is important, but preliminary, work on architectures that could liberate generators from serial processing schemes. However, we should not throw out our previous serial schemes just yet.

On the whole, this book provides an important source for research on many aspects of natural language generation. Although the contributions are not of uniform quality and level of detail, most are very good.

REFERENCES

- Cullingford, R.E. 1986 *Natural Language Processing: A Knowledge-Engineering Approach*, Rowman and Littlefield, Inc., Totowa, NJ.
 McKeown, K. 1985 *Text Generation*, Cambridge University Press, Cambridge, England.
 Grishman, R. 1986 *Computational Linguistics* Cambridge University Press, Cambridge, England.

NOTE

1. The views expressed in this review are the author's and do not necessarily reflect the opinions of SRI International.

Marie Bienkowski is a member of the Applied AI Technology Program at SRI International. Her Ph.D. research dealt with explanation production and language generation for problem-solving systems; her current research is on methods for explanation production in training systems. Bienkowski's address is: SRI International, 333 Ravenswood Ave., Menlo Park, CA 94025. E-Mail: bienk@istc.sri.com.

THE LINGUISTIC BASIS OF TEXT GENERATION

Laurence Danlos

(Centre National de la Recherche Scientifique, Paris, France)

(Studies in Natural Language Processing)

Cambridge University Press: Cambridge, England, 1987, x+222 pp.

ISBN 0-521-32398-8, \$39.50 (hb) (20% discount to ACL members)

Reviewed by

Kathleen McCoy

University of Delaware

In this book Laurence Danlos has been able to achieve a nice balance between straight linguistics and straight computer science (artificial intelligence). She uses a detailed linguistic analysis as the basis for a text generation system. In doing so, she has managed to come up with ideas of interest to both fields.

The book describes the methodology behind a generation system whose aim is to produce "good" texts from semantic representations of what is to be conveyed. Danlos says that there are two kinds of decisions that must be made to do this:

- Conceptual decisions (e.g., what order should the information be presented in, what should be made explicit and what implicit?); and
- Linguistic decisions (e.g., where should sentence boundaries be made, what words should be used, what syntactic constructions?).

Danlos rather convincingly defends a claim that all of these decisions are dependent on each other. For instance, a decision to order the information in one way will limit the choice of syntactic constructions available (which in turn will limit lexical choice) and vice versa. In addition, there is no a priori reason why priority should be given to one of these decisions over the others. The priority decision concerning a particular semantic relation can only be made within a particular domain after detailed linguistic analysis. In order for the generation system to work, it must capture the available conceptual and linguistic choices. Danlos advocates encoding the choices in two structures, and illustrates how the choices are determined and resulting structures used for texts concerning direct causal relationships (between an ACT and RESULTing state) within the terrorist domain. The two structures she advocates are:

1. A lexicon grammar that is specific to a particular domain and semantic relationship and encodes the possible simple sentences (lexical items) that can be used to express concepts in the domain (e.g., the act and result in a direct causal relationship such as a murder attempt); and
2. A discourse grammar that is specific to a particular semantic relation and encodes the remaining choices.