

the book more of a recommended reading is the renewed empiricism in the field, largely promoted by the very practical need to scale up natural language systems, and largely due to the realization that linguistic information about words could be derived from massive on-line text resources.

Whether these resources come in the shape of on-line dictionaries or text corpora is immaterial here. By reading *Looking Up*, one becomes acutely aware of the richness of lexical information available in (and distributed over) millions of words of text. One also understands that careful inspection of a dictionary entry (or a set of related entries) is likely to reveal considerably more in terms of lexical properties of the word (or class of words) than is apparently visible. The nature of lexical information in, and its extractability from, such text resources has been much discussed recently in the computational linguistics and computational lexicography literature; in particular, the statement that there is a wealth of implicit information available in on-line dictionaries and corpora has been made over and over again recently (see, for instance, Atkins et al. 1988; Boguraev and Briscoe 1989; Hindle 1989; Church and Hanks 1989). However, there is a world of difference between "retro-engineering," by whatever means, methods and rules for inferring lexical information from dictionary entries, and being told in advance the kinds of lexical regularities, generalizations, and properties encoded in these entries. For that reason alone, and particularly given that the COBUILD dictionary is available in machine-readable form, *Looking Up* is a book that should not be ignored by researchers interested in computational lexicography, lexical semantics, or simply the nature of word meaning.

REFERENCES

- Atkins, B. T., Kegl, J. and Levin, B. 1988 Anatomy of a Verb Entry. *International Journal of Lexicography* 1(2):84-126.
- Boguraev, B. K., and Briscoe, E. J. 1989 *Computational Lexicography for Natural Language Processing*. Longman, London and New York.
- Church, K. W. and Hanks, P. 1989 Word Association Norms, Mutual Information, and Lexicography. *Proceedings of the 27th Annual Meeting of the Association for Computational Linguistics*, Vancouver, B.C.
- Fillmore, C. J. 1989 Two Dictionaries. *International Journal of Lexicography* 2(1):57-83.
- Hindle, D. 1989 Acquiring Disambiguation Rules from Text. *Proceedings of the 27th Annual Meeting of the Association for Computational Linguistics*, Vancouver, B.C.
- Landau, S. I. 1989 *Dictionaries: The Art and Craft of Lexicography*. Cambridge University Press, Cambridge, U.K.

Branimir Boguraev is a Research Staff Member at the IBM T.J. Watson Research Center, where he is also in charge of the Lexical Systems Project. He holds a Ph.D. degree from the University of Cambridge, where he subsequently worked on a number of projects in natural language processing. His recent work has been in the areas of computational lexicography and lexicology. Boguraev's address is: IBM Research, P.O. Box 704, Yorktown Heights, New York 10598. E-mail: bkb@ibm.com

GENERATING NATURAL LANGUAGE UNDER PRAGMATIC CONSTRAINTS

Eduard H. Hovy

(Information Sciences Institute, University of Southern California)

Hillsdale, NJ: Lawrence Erlbaum Associates, 1988, xiii + 214 pp

Hardbound, ISBN 0-8058-0248-7, \$29.95

Paperbound, ISBN 0-8058-0249-5, \$19.95

Reviewed by

Wolfgang Hoepfner

University of Koblenz

This book is a revised version of the author's dissertation, which was submitted to Yale University in February 1987 and published as a research report (Hovy 1987). One shouldn't be too surprised to learn from the preface that Roger C. Schank, Drew McDermott, and Bob Abelson were the honorable members of the thesis committee. Various well-known AI researchers—not exclusively from Yale—have given a hand while the thesis was on its way and are therefore mentioned in the acknowledgments. The acknowledgments, by the way, give a first example of how stylistic features affect the generation of text: Hovy switches between several styles (formality, verbosity, gratefulness, haste) while expressing his thanks to different classes of persons. The preface stresses that the book is not only useful to computational linguists, but also to theoretical linguists, especially those working in generation. Thus, the last section of every chapter deals with implementation and might be skipped by readers not interested in the computational issues.

The book describes the generation system PAULINE, which was developed and implemented by the author. The system's name is an acronym: Planning And Uttering Language In Natural Environments. (It is also the name of Hovy's sister).

Chapter 1 (11 pp.) introduces the specific research area: how do pragmatic and stylistic issues influence the generation of natural language texts? Starting with real-world descriptions of an event at Yale, the destruction of a shantytown by university authorities, it is demonstrated how different viewpoints of the respective authors affect texts. The same event is then described by PAULINE in different pragmatic adjustments. The event itself is represented in a network of about 120 elements (presumably some version of conceptual dependency), and it is claimed that the system produces over 100 different texts. The second example is a description of a fictitious primary election between Carter and Kennedy as Democratic presidential candidates. "Well, so Carter lost the primary to Kennedy by 1335 votes" is one example of a very condensed description. This terseness is beaten only by example number 12, which consists of nothing but blanks: "The program didn't find any topics that it liked and the hearer also liked,

and it didn't have time to search the story representation for other topics or to find ways of mitigating the unpleasant ones" (p. 8). The third event deals with a fight between Mike and Jim, resulting in the death of the latter. All three events are used throughout the book. I didn't count the different versions of the texts generated by the system and included as examples in the book, but one gets a little bored when reading, e.g., the *n*th version of the shantytown story.

The programmatic essence of Hovy's approach can best be characterized by the following quotation:

Current artificial intelligence work in natural language processing places far too little emphasis on the role of pragmatics. This is a mistake. . . . The time is ripe to start examining what kinds of goals are relevant, what strategies achieve them, and what planners and realizers must be like to operate under their control. We must face the fact that, in order to have real, flexible text, we simply cannot do without recourse to the airy world of pragmatics. We need to experiment! (p. 11)

True enough!

Chapter 2 (23 pp.) is briefly called "Pragmatics," and contains the theoretical underpinnings of PAULINE. Starting with a short discussion on what pragmatics is all about and what relation to semantics might be established, Hovy identifies three basic categories that are relevant to his system (p. 17):

- Interlocutor's personal characteristics: factual knowledge, opinions, emotional states, interpersonal relationship, etc.
- The speaker's goals with respect to the hearer: effects on future behavior, opinions, relative status, etc.
- The conversational atmosphere: tone, time, physical setting, etc.

Each of these categories is illustrated with several examples in subsequent sections. This leads to a set of 13 goals, which are a first approximation to the underlying pragmatic features necessary for generation. Some examples of these goals are: increase knowledge; make the topic seem good or bad, contrary to the hearer's opinion; make the hearer feel inferior to, equal with, or dominant over the speaker; be hasty, normal, or effusive. Hovy does not claim that these goals are complete or adequate in a general way. They rather appear to be necessary for the way PAULINE is supposed to produce text. One might wonder how a system would be able to choose rationally among these goals, and this, of course, is beyond the scope of an implemented system. The user has to determine the interpersonal goals and the conversational setting by selecting a value for 23 relevant features, each of these being represented by three possible values. For instance, the emotional state of the hearer can be set to happy, angry, or calm; the conversational atmosphere with respect to time can be set to much time, some time, or little time. Rhetorical goals relate the pragmatic adjustments defined by the user, or rather the experimenter, to decisions for the generator, e.g., word choice, sentence inclusion. The chapter closes with a discus-

sion of rhetorical goals of opinion and of style. The latter has seldom been investigated in computational linguistics (exceptions are the EPISTLE project [Heidorn et al. 1982], or the BOGUE system [Ryan 1989]), and it is certainly a merit of Hovy's work to establish an initial set of 12 style features (e.g., timidity, haste, respect).

In Chapter 3 (17 pp.), Hovy argues for an interpretation component as part of a natural language generator. Interpretation in this sense means that a set of inference rules has to be present to "digest" the conceptual input representation and produce a structure that is suited for the task of verbalization. Examples for such inference rules are the determination of an appropriate level of detail, presentation in a confrontative or conciliatory manner, or the inclusion of reminders. In PAULINE, the interpretation mechanism is realized as patterns that change the original network, when matched with some input configuration.

Chapter 4 (25 pp.) discusses the notion of affect in texts, i.e., the opinion a speaker has about the content of his or her utterances. The effect of such a biased generation strategy is demonstrated with different issues of the fighting scene between Jim and Mike, where the differences originate from the agent the system sides with. Hovy develops an "affect rule": "In order to convince the hearer that some topic is GOOD or BAD, combine it with other GOOD or BAD topics using enhancers and mitigators" (p. 61). This rule appeals to both aspects of the generation process. Content is determined by the selection of topics and their linearization. Three plans are provided that have a similar function to the text schemata used in the TEXT system (McKeown 1985). These plans are called DESCRIBE, RELATE, and CONVINCING, the last one being presented in some detail, as this is the one where affect sneaks in. The production of form is influenced by affect-bearing phrases (e.g. "not only X but Y"), by adverbs (e.g. "really," "only," "slightly"), by clause order, and by word choice (e.g. "wimpy" versus "small"). Chapter 4 closes with some observations concerning partiality of the generated text. The rhetorical goals "timidity" and "partiality," or rather their respective values, are calculated on the basis of other features.

Chapter 5 (23 pp.) is about style. Hovy introduces three style notions: formality, haste, and force. For each of these rhetorical goals, some rules are given that allow the system to derive values by combining already established rhetorical goals. "RG:haste", for example, "is set to 'pressured' if the time is marked 'little', the relative social status is marked 'subordinate', and the depth of acquaintance is marked 'acquaintances' or 'strangers'" (p. 93). Such an "algorithmic approach to the creation of style" is what leads to the diversity of texts generated by PAULINE and is motivated by the assumption "that style is the result of following a coherent policy" (p. 104).

In Chapter 6 (24 pp.), Hovy argues for the central role of the lexicon as the one and only device for linguistic knowledge within a generator. This, obviously, is quite contrary to current grammatical theories, as for instance the varie-

ties of unification-based approaches. Hovy briefly sketches issues in transformational, systemic, and functional grammar, and stresses the need for a "phrasal lexicon" as originally introduced by Becker (1975). He then introduces the "syntax specialists," which incorporate the three tasks of inclusion of topics, ordering, and casting (selection of syntactic classes) in a procedural manner. This approach resembles the one earlier developed for analysis and known as "word-expert parsing" (Small and Rieger 1982). This analogy, however, is an inference of the reviewer and is not discussed in the book. The last section of Chapter 6 illustrates the incremental generation of the sentence "The small 23-year-old from New Haven, Sue, was told by Jim that Janet died" as a realization of a conceptual dependency representation of something like "Jim told Sue that Janet died." Appendix B contains the complete phrasal grammar implemented for PAULINE.

Chapter 7 (17 pp.) once more turns to the planning process incorporated in PAULINE. Hovy distinguishes between prescriptive planning and restrictive planning, the first of these methods being the one used in generators up to now. He argues for the incorporation of both techniques, which is called "limited-commitment or interleaved planning." The essence of this method of planning is that there is no fixed and predefined order between the planner (expansion of states) and the realizer (testing whether a planned action has contributed to the ultimate goal). Thus results of the realizer should be able to improve the planner at all times during the planning process. Hovy discusses planning approaches in other generators, e.g., KAMP (Appelt 1985), the system which is best known for its extreme plan-based philosophy. The last section of this chapter describes the interleaved planning employed by PAULINE.

The book's last chapter contains a nine-page section, "Review of Language Generation" a title that is not fully justified as it is with McKeown and Swartout (1988) or McDonald (1987). The main contribution of this chapter is the addition of one fundamental question for generation to the already existing two questions. Hovy adds to 'What should I say?' and 'How should I say it?' the question 'Why should I say it?', thus manifesting his pragmatic approach. A question, but quite a different one, has always puzzled me with American dissertations: why is the state of the art always contained in the last chapter, unlike European dissertations, where this is usually the starting point? For my part, I have developed the habit of starting to read an American dissertation at the last chapter and then continuing at the beginning.

Appendix A contains a short annotated example, which shows the generation of a shantytown text. Appendix B gives a description of PAULINE's phrasal grammar. The book is supplemented with an index of names and a topic index.

A few years ago there was a discussion in *Computational Linguistics* initiated by Ballard (1983) about standards for

reports on natural language systems. This discussion was motivated by the observation that descriptions of NL systems very often lack the desirable details of implementation, examples, or theoretical underpinning, as well as precision. Probably everybody who has been engaged in teaching courses on NL processing has been frustrated in trying to get a careful description of how some natural language system *really* works. Now, let's look at Hovy's book from this angle. Unlike some other published work from the well-known Yale Department, Hovy is very modest in his claims of generality or coverage. It is often pointed out that the specific approach of PAULINE is in no way a solution of all pragmatic issues in NL generation, which, by the way, would be difficult to prove. So here we have a pleasant side of the book.

There are, however, other properties that make the reader feel a little uneasy. Hovy never states precisely what the input representations of the texts look like. One can easily infer that there is some version of conceptual dependency (Schank 1975) representing the contents of the stories. There is even a linearized ATRANS example for a prototypical sentence, such as "John gave Mary the book and she gave him the money" (p. 26). And at the end of the book (p. 140), Hovy informs us that "PAULINE's input is represented in a standard case-frame-type language based on conceptual dependency . . . and is embedded in a property inheritance network." But still, one wonders how complex events, as the shantytown example, might be represented in conceptual dependency networks, which types of primitive actions are employed, and how a basic network is enhanced with pragmatic features.

A second shortcoming is the vagueness in how conceptual dependency constructs are verbalized. One might guess that parts of older generators (e.g., Goldman 1975; Meehan 1976; or even Simmons and Slocum 1972) are employed for, e.g., the verbalization of primitive acts, but how this is done in PAULINE is not stated explicitly.

The interaction of the various pragmatic attributes and their respective values is described in several different parts of the book. The reader gets the impression that this is a very complicated and tricky business, but it is very hard to state in a somewhat complete way just how this is achieved. The impact of the features is often demonstrated with generated texts, and these examples are quite convincing. However, they are not sufficient for invoking a detailed picture of how things interact with each other. Apart from this, there is only a marginal remark on the programming language, which is T, and on the size of the program (over 12,000 lines, a rather dubious indication). Neither hardware nor processing time is mentioned.

Hovy refers to most of the important work done in NL generation so far (226 entries in the bibliography). He could have been more detailed in using these entries. One often sees a heap of references at the end of a statement, especially in the theoretical sections. A more specific use of other people's work would have been a benefit for the book.

The final critical remark is related to the degree of redundancy. There are many statements that occur time and again throughout the book. The reader often wonders why one is being told things one already knows. It is not as bad as in another recent dissertation from Yale (Hammond 1988), which was recently reviewed in a German AI journal (Hertzberg and Horz 1990). A translated quotation from this review: "If the reader really wants to read the complete book, she is punished and lulled to sleep by the recurrent repetitions. . . . The sentence 'Plans are indexed by the goals they satisfy and by the problems they avoid' is learned by heart just by reading any two chapters. If Hammond is an efficient LaTeX user, he has probably written a macro; at least, it would have been worthwhile." Once more: it is not as bad with Hovy, but it is annoying nevertheless.

My general opinion about the book is that it is necessary reading for anybody working in NL generation. It is profitable reading for anybody engaged in natural language processing. And it is worthwhile reading for linguists specializing in pragmatics, style, or language production. I have stated above those properties that I missed or didn't like. What I did like was Hovy's scientific rigor in advocating the pragmatic basis of text generation. This attitude is significant for an AI-oriented approach to language processing, but this does not imply that it is a scruffy one.

REFERENCES

- Appelt, D. E. 1985 *Planning English Sentences*. Cambridge University Press, Cambridge, U.K.
- Ballard, B. W. 1983 On the Need for Careful Description of NL Prototypes. *Computational Linguistics* 9(1):23-24.
- Becker, J. D. 1975 The Phrasal Lexicon. In: Schank, R. C. and Nash-Webber, B. L., eds. *Theoretical Issues in Natural Language Processing*. Cambridge, MA, 70-73.
- Goldman, N. M. 1975 Conceptual Generation. In: Schank R. C. *Conceptual Information Processing*. North Holland, Amsterdam, The Netherlands, 289-371.
- Hammond, K. 1988. *Case-Based Planning: Viewing Planning as a Memory Task*. Academic Press, San Diego, CA.
- Heidorn, G. E., Jensen, K., Miller, L. A., Byrd, R. J., and Chodorow, M. S. 1982 The EPISTLE Text-Critiquing System. *IBM Systems Journal* 21(3):305-326.
- Hertzberg, J., and Horz, A. 1990 Review of Hammond 1988. *Künstliche Intelligenz: Forschung und Entwicklung* 1:61-62.
- Hovy, E. H. 1987 *Generating Natural Language under Pragmatic Constraints*. Ph.D. thesis, research report 521, Department of Computer Sciences, Yale University, New Haven, CT.
- McDonald, D. D. 1987. Natural Language Generation. In: Shapiro, S., ed. *Encyclopedia of Artificial Intelligence*. John Wiley, New York, NY, 642-655.
- McKeown, K. R. 1985 *Text Generation*. Cambridge University Press, Cambridge, U.K.
- McKeown, K. R. and Swartout, W. R. 1988. Language Generation and Explanation. In: Zock, M. and Sabah, G., eds. *Advances in Natural Language Generation: An Interdisciplinary Perspective*. Pinter Publishers, London, U.K., 1-51.
- Meehan, J. 1976. *The Metanovel: Writing Stories by Computer*. Garland Press, New York, NY.
- Ryan, M. B. P. 1989 *The Computational Codification of the Semantic Aspects of Style*. Technical report CSRI-231, Computer Systems Research Institute, University of Toronto, Toronto, Canada.
- Schank, R. C. 1975 *Conceptual Information Processing*. North-Holland, Amsterdam, The Netherlands.
- Simmons, R. F. and Slocum, J. 1972 Generating English Discourse from Semantic Networks. *Communications of the ACM* 15(10):891-905.
- Small, S. and Rieger, C. 1982 Parsing and Comprehending with Word Experts (A Theory and its Realization). In: Lehnert, W. G. and Ringle, M. H., eds. *Strategies for Natural Language Processing*. Lawrence Erlbaum Associates, Hillsdale, N.J., 89-148.

Wolfgang Hoepfner is a professor of Computational Linguistics and Artificial Intelligence at the University of Koblenz, West Germany. He holds a doctorate from the University of Hamburg in Computational Linguistics. Formerly engaged in various AI projects (HAM-ANS, WISBER) at the University of Hamburg, his recent research has been in knowledge representation, especially temporal and spatial reasoning, and in NL generation. Hoepfner's address is: University of Koblenz, FB Informatik, Computerlinguistik, Rheinau 3-4, D-5400 Koblenz, West Germany. E-mail: hoepfner@infko.uucp

BRIEFLY NOTED

LEARNABILITY AND LINGUISTIC THEORY

Robert J. Matthews and William Demopoulos, eds.

(Rutgers University and University of Western Ontario)

Dordrecht: Kluwer, 1989, vii + 217 pp.

(Studies in Theoretical Psycholinguistics 9)

Hardbound, ISBN 0-7923-0247-8, Dfl 130.-, \$64.00, £42.00

The ninth volume of Kluwer's Studies in Theoretical Psycholinguistics is, like the previous ones, devoted to what any modern linguistic theory has inevitably to face, namely the logical problem of language acquisition. As before, in most contributed papers the problem is discussed within the parametrized Government-Binding framework. This time, however, issues from formal learnability theory serve as a starting point and as a basis for the subsequent reformulation of the rationalist-empiricist debate on language acquisition. Less general questions are also addressed; as, for example, how children eventually succeed in avoiding overgeneralizations—an intriguing puzzle, given the widely accepted view that not enough negative evidence is directly accessible in the course of first language learning.

No doubt the book may be of interest to cognitive scientists and those computational linguists who deal with modeling natural language acquisition. The idea of having a computer system that gradually learns a language from examples, very much like people do, is an exciting one, and it will surely receive much attention in the foreseeable future. Though the book under consideration does not provide the would-be designers of a computer learning system with algorithms they might immediately employ, it presents a wide selection of topics characteristic of the current literature on modeling natural language acquisition and, consequently, can be used by computational linguists as an important source of information relevant to their research.—Miroslaw Bańko, Institute of Polish Language, Warsaw University