

Current Research in Natural Language Generation

Robert Dale, Chris Mellish, and Michael Zock (editors)
(University of Edinburgh and CNRS)

London: Academic Press, 1990,
ix + 356 pp.
(Cognitive science series) Hardbound
ISBN 0-12-200735-2, \$45.00, £23.50

Reviewed by
Ingrid Zukerman
Monash University

Current Research in Natural Language Generation is derived from the Second European Natural Language Generation Workshop, which was held in Edinburgh in April 1989. The papers included in this volume were selected from revised versions of some of the papers presented at the workshop. The book provides a snapshot of the current research in NLG, with particular emphasis on the work conducted by the European research community. It is aimed at an audience already familiar with NLG. Even though the different papers provide introductory technical material where necessary, in general, this material alone seemed insufficient to enable the uninitiated to understand completely the different arguments. However, adequate references are provided.

The book is divided into four main sections: text planning (four papers), linguistic realization (two papers), building descriptions (three papers), and connectionist approaches (two papers).

The first section contains contributions by Hovy, Scott and de Souza, Cawsey, and McKeown et al. The first paper, entitled "Unresolved issues in paragraph planning," by Hovy, is a position paper that raises seven unresolved problems in discourse planning, mainly from the perspective of Rhetorical Structure Theory (RST) (Mann and Thompson 1988). These problems are divided into two groups: problems concerning the theory and representation of coherence relations, and algorithmic problems. The importance of this paper is that it focuses the discussion on NLG on crucial issues to which researchers have to address themselves.

The second paper, "Getting the message across in RST-based generation," by Scott and de Souza, addresses the problem of generating text that achieves a communicative goal effectively. To this end, the authors look upon style in terms of a reader's ease of processing a text, rather than in terms of aesthetics. The main contribution of the paper is that it presents explicit heuristics grounded in psycholinguistic evidence to control the realization of RST discourse relations. I found Section 3.2, "Making the text sensitive to the communicative setting," where the authors link the generation of textual markers to context sensitivity, somewhat problematic. In addition, since the heuristics presented in the paper were not implemented, some analysis of how they interact would have been useful.

The third paper, "Generating explanatory discourse," by Cawsey, discusses an interactive content and discourse planner whose output is tailored to the changing capabilities of the user, and which can also handle interruptions and remedial discourse. The paper offers a novel approach that addresses specific criteria in discourse planning. Its contributions include: the use of an approach similar to Litman's (1985) to separate content and discourse planning; the hierarchical decomposition of schemata; and the use of an agenda for content planning. This paper provides a clear exposition

on the integration of several ingredients that are required in discourse planning.

The last paper in this section, by McKeown et al., "Natural language generation in COMET," describes the approach taken in the nongraphic modules of COMET (Co-Ordinated Multimedia Explanation Testbed). The thrust of the paper is two-fold: the derivation of new schemata from old ones for content planning; and the use of the Functional Unification Formalism (FUF) to allow for interaction across modules. The modules implemented in FUF are: media coordinator, lexical interface, and grammatical realization. The first of these points would have benefited from a comparison between the time and effort expended in developing schemas from scratch versus verifying the derived schemas against the texts in the domain at hand. With respect to the second point, I found the use of a uniform formalism to carry out such diverse tasks particularly appealing. A minor stylistic shortcoming of this paper is a certain lack of uniformity in the presentation.

The section on linguistic realization contains contributions by van Noord and De Smedt. Van Noord's paper, entitled "An overview of head-driven bottom-up generation," describes and analyzes the technique in the title, beginning with a "vanilla" algorithm, and then proposing extensions to address different problems that crop up. The main message of the paper was both understandable and convincing. However, following the details was more difficult. This was partly due to the fact that some nontrivial statements were left to the reader to figure out, and was exacerbated by some inconsistency in notation and a sample Dutch grammar (the author's Figure 6.4) from which it is unclear how the sample sentences can be derived.

The paper authored by De Smedt, "IPF: An incremental parallel formulator," presents a sentence generator based on Segment Grammars (Kempen 1987), which is designed to operate in parallel with a content planner. The main feature of the system is that it handles conceptual input that can be provided by the content planner in any order.

The third section, "Building descriptions," features papers by Horacek, Dale, and Reiter. Horacek's paper, entitled "The architecture of a generation component in a complete natural language dialogue system," provides a description of an entire generation system, with particular emphasis on two novel components: a component that performs meaning-preserving transformations at a conceptual level; and a component that maps predicates used at this level onto lexemes and grammatical functions. In addition, it addresses the requirements from a generation system by embedding it in an advisory system.

Dale's paper, "Generating recipes: An overview of Epicure," also presents an entire generation system, but in this case, the system stands by itself. Like Horacek's system, Dale's mechanism relies on processes that map the information between different levels of representation. The focus of Dale's work is the construction of referring expressions for entities that are undergoing constant change. The main contribution of the paper is a representation language that supports this task. Most of the paper is described at a good level of detail. However, the section on the generation of pseudo-partitive noun phrases and the section on the unification grammar are minimal. In particular, the presentation of a fragment of the unification grammar, without some explanation that traces the realization process, makes it extremely difficult to understand how the discourse is realized.

The last paper in this section, by Ehud Reiter, is entitled "Generating descriptions that exploit a user's domain knowledge." It provides a detailed algorithm for the generation of object descriptions that are accurate (truthful), valid (fulfill the speaker's communicative goal), and free of false implicatures for a particular user model, and presents a complexity analysis of this algorithm. The focus of this work is on the

selection of content words, rather than rhetorical structures. The main contribution of the paper lies in its formal and detailed treatment of the subject. At the definitional level, there wasn't a clear distinction between the author's interpretation of Grice's maxims of brevity and of quantity, but this shortcoming was overcome in the examples.

The last section of the book, "Connectionist approaches," contains papers by Houghton and by Kitano. Houghton's paper, "The problem of serial order: A neural network model for sequence learning and recall," differs from the rest of the papers in the book in that it addresses a low-level NLG task, namely the task of learning and recalling phonemic sequences in monosyllabic English words. Houghton treats this phenomenon as a temporal one, proposing the notion of competitive queuing, where the order in which phonemes surface depends on their level of activation. The form of sequence recall produced by this model is in line with psycholinguistic evidence.

The paper by Kitano, "Parallel incremental sentence production for a model of simultaneous interpretation," describes a concurrent parsing and generation process used in a simultaneous English-Japanese translation system. Like human translators, this system formulates hypotheses about forthcoming utterances, and supports the commencement of generation before parsing is completed. The architecture of the system combines a parallel marker-passing scheme with a connectionist network. In addition, case-based processing is integrated with constraint-based processing to provide flexibility in generation and maintain the ability to generate specific expressions. The system described in this paper exhibits impressive capabilities. However, this paper would have been enhanced by the presentation of some statistics of the implementation, e.g., how large is the memory? what types of concept sequences are included in the memory? how many concepts and concept sequences typically get activated? and how long does it take to translate a sentence of average length?

The criticisms above pertain mainly to presentation and to the omission and/or lack of clarity of some explanations. However, in general, the papers in this book are well written and contain substantial contributions. Overall, the book provides a good coverage of the NLG field, and it is useful for researchers in NLG and related fields. It offers a good balance of position papers, papers that provide a complete system perspective, and papers that focus on particular capabilities of a system, e.g., realization, content planning, user modeling, and lexical planning. In particular, papers that describe complete NLG systems, such as those by McKeown et al., Horacek, and Dale, are essential in such a collection, in order to give the reader a view of how the different pieces fit together. At the same time, the papers that deal with particular components of such systems provide details regarding the different aspects that have to be considered.

The discourse planning systems described in the book were implemented in technical domains, such as electrical circuits, equipment repair, investment advice, and recipes. Nontechnical domains, such as story generation, were conspicuously absent. On the realization front, Systemic Grammars were an unhappy omission. Nevertheless, a variety of research issues and approaches were presented in the book. This gives a sense both of the youth and the depth of the field. At the same time, there is an emerging consensus with respect to approaches to the more defined issues. For example, schemas and RST for discourse content planning, and unification-based grammars for the realization component.

References

- Kempen, Gerard (1987). "A framework for incremental syntactic tree formation." *Proceedings, 10th International Joint Conference on Artificial Intelligence (IJCAI-87)*, Milan, Italy, 655–660.
- Litman, Diane Judith (1985). *Plan recognition and discourse analysis: an integrated approach for understanding dialogues*. Doctoral dissertation, published as technical report 170, Department of Computer Science, University of Rochester.
- Mann, William C.; and Thompson, Sandra (1988). "Rhetorical Structure Theory: toward a functional theory of text organization." *Text*, 8(3), 243–281.

Ingrid Zukerman is a Senior Lecturer in Computer Science at Monash University, conducting research in discourse planning and user modeling. She received her B.Sc. in industrial engineering and management and her M.Sc. in operations research, both from the Technion—Israel Institute of Technology. In 1986, she completed her Ph.D. in Computer Science from UCLA on the generation of meta-comments. Zukerman's address is: Department of Computer Science, Monash University, Clayton, Victoria 3168, Australia; e-mail: ingrid@bruce.cs.monash.oz.au