Building Natural Language Generation Systems

Ehud Reiter and Robert Dale
(University of Aberdeen and Macquarie University)

Reviewed by
Helmut Horacek
University of the Saarland

Reiter and Dale's book is about natural language generation (NLG), and describes methods for building systems of this kind. It offers a comprehensive overview of the field of NLG, from a perspective that has rarely been described elsewhere. It is very valuable for people with little or even no knowledge about NLG particulars who intend to investigate applications where NLG technology might play a certain role, as well as for people who want to obtain an application-oriented overview of the field. In this sense, all communities intended to be served by the book—students, academics in related areas, and software developers—are well addressed. In addition, some sections provide material that is probably new even to researchers with considerable experience in the field, such as the discussions on corpus analysis and domain modeling.

In the introduction, the field of NLG is briefly characterized from research- and application-oriented perspectives and is illustrated by screen shots produced by several systems. Then conditions for effective uses of this technology are elaborated and contrasted with conditions where other techniques are more appropriate, and methods for determining the intended functionality of a system to be built are discussed. The main sections of the book are devoted to the prototypical architecture of application-oriented NLG systems and their major processing phases: document planning, microplanning, and surface realization. Each of these three phases is illustrated by a number of detailed examples, demonstrating the successive refinements of utterance specifications in the course of processing. In the final section, the embedding of natural language processing technology is discussed, featuring typography, combined uses with graphics and hypertext, and integration with speech. The methods are illustrated by a large number of examples—the book contains more than 120 figures on its 248 pages. At the end of each section, a number of useful references for further reading are related to the section topics. In the appendix, a table summarizing the 35 systems referred to in the book is given.

In general, the presentations in the book are beneficial for their high degree of detail in documenting representations. Running examples are illustrated by intermediate representations at various stages of processing. They demonstrate very well the increasing degrees of precision in which utterances are specified and the associated commitments about possible text portions that express these specifications. Particularly valuable are the comparisons of alternative representations, ranging from skeletal propositions over lexicalized case frames to canned text, including a comparison of the consequences for the distribution of work over the system components involved and discussions of the associated pros and cons. These chapters provide a number of valuable hints, given the currently low degree of standardization in the field. Further important aspects addressed, which have been widely ignored in the literature so far,

are the role of corpus analysis in determining the coverage of a system to be built, domain modeling, and embedding NLG in larger environments, although the space devoted to issues of typography is a bit oversized.

However, one can hardly expect that a book on a subject as diverse as NLG will cover all relevant issues in an equally detailed manner. Hence the book by Reiter and Dale is biased to a certain extent towards standardized, formatted document generation, mostly in English. As a consequence, some NLG-related application-oriented topics are underrepresented, prominently issues of multilinguality. Though several multilingual generation systems are mentioned throughout the book—the FOG system, which produces English and French texts, is even used for the running examples—difficult issues in handling differences across languages are not dealt with. For example, problems with phenomena such as head switching, collocations, and rich inflections are briefly mentioned at best, and techniques for solving them are not discussed. But the concentration on certain kinds of applications is fully legitimate and does not affect the usefulness of the book if one is aware of this focus.

The biggest weakness of the book is probably that it requires a reader who is at best vaguely familiar with the concepts of NLG to be able to assess the delicacy of the field. This includes the capability of judging the extent of difficulty of any particular subtask of a system to be built, which depends on depths of understanding of the phenomena addressed. While acquiring the ability to assess this problem in sufficient degree is a matter of long-term experience, novices in the field, especially, should be given some hints about potential problems and limitations of today's techniques. I think that the authors could have addressed this aspect by stressing the fact that the tasks to be accomplished by a generation system are of much greater variety than those needed for natural language understanding:

- On the one hand, many aspects of pragmatics, which is one of the least understood areas in linguistics, are more important in generation than in analysis and, therefore, they have to be captured to a certain extent.

- On the other hand, the initial representations for a generation system are much more distant from the language surface than the results of understanding systems, and they can appear in various, mostly nonlinguistic forms.

These properties have crucial consequences for whether a certain functional aspect of a system to be built can be handled, and how difficult this might be. Although the variety of possible initial representations for NLG is discussed at some length in the book, the associated delicacy does not become clear. In particular, some information in domain model representations may appear only implicitly, and such pieces of information may or may not be required for the functionality of the generation system, depending on subtleties in the intended coverage. Moreover, the many examples used in the book, which by themselves are very valuable to illustrate details of representation, are in a sense simplistic, which they must be in order to allow for a clear and concise presentation. Therefore, some hints about the nature of more complex and more difficult examples would be helpful. Finally, some considerations about typical sizes of applications in terms of lexicon, grammar, and world models would be beneficial, too.

Apart from this drawback, there are a number of further issues that one could wish had been addressed in the book. These range from general topics, such as relating NLG system design to methods in requirements engineering, through application-oriented

issues, such as dependencies between order variations in the standard pipeline architecture and functional coverage of a system, to specific ones, such as lexicalist approaches to surface realization. The latter topics arise mostly in machine translation rather than in standard document generation, the kind of application primarily addressed in the book. As for the other topics, it is not surprising that they have not been addressed in the book, given the currently low degree of understanding of these issues. However, for progress in system-building techniques, a deeper understanding of these questions will be crucial. Hence, all these issues are candidates for being addressed in a complementary book or in a new edition.

In general, the book is well written and easy to read, although it is repetitive in some places, and the section on referring expressions is not too well linked to the remaining parts of the book. Major benefits are in the large number of illustrations and examples, the high degree of detail in which they are presented, and the discussions of the pros and cons of alternative representations. Altogether, I am convinced that the book serves its intended purpose of enabling its readers to build an NLG system on their own, provided the functionality of this system is similar to that of the systems focused on in the book, which particularly means that multilingual issues must not be a major concern.

Helmut Horacek is a senior researcher at the University of the Saarland, Germany. Within the field of natural language generation, he has contributed to lexicalization, referring expression generation, multilingual document generation, and proof presentation. He has served as an organizer or program committee member for several workshops in this area. Horacek's address is University of the Saarland, Department of Computer Science, P.O. Box 1150, D-66041 Saarbrücken, Germany; e-mail: horacek@cs.uni-sb.de.