Book Reviews

User Modeling in Text Generation

Cécile L. Paris
(Information Sciences Institute, University of Southern California)

Reviewed by
T. Pattabhiraman
Simon Fraser University

That language varies according to the situation is a truism; however, the details and implications of that truism are far from obvious, whether your enterprise is theory formation or system construction. In the book under review, Cécile Paris identifies the user's level of knowledge about the domain of discourse as a significant factor causing variation in text and presents a computational theory of the mechanisms and processes by which such variation can be caused in the text produced by a natural language generation (NLG) system. Paris also presents details of the TAILOR system, which generates an impressive range of texts describing physical devices, the content and organization of each text being tailored to suit the user's level of knowledge, encoded in a *user model*. As Paris illustrates, a high-school textbook, an adult encyclopedia, or a car manual would each describe a car differently, basing the description significantly on considerations of the domain knowledge of the intended reader of the text.

This book is based on doctoral research carried out by Paris at Columbia University. While previous researchers have observed that a user's domain knowledge affects the *amount* of information in a text (a novice needs more information; an expert needs less), Paris discovered that the user's domain knowledge also affects the *kind* of information that needs to be presented in a text. User modeling and NLG are at present identifiably distinct but significantly interacting research areas, and this book is the first to be published in the area of their overlap. It provides stimulus for further research along several directions of interest to computational linguists in general.

The book is divided into eight chapters. The introductory chapter contains a concise overview of NLG and user modeling in generation and summarizes the author's methodology and major results. The illustrations motivate the author's research clearly, and the chapter should enable even those not specializing in user modeling or NLG to appreciate the author's research goals and results.

Chapter 2 is devoted mainly to a discussion of what a user model might contain and how the information contained in a user model could be used to tailor the output of a generation system. Paris ascertains that a generation system that produces descriptions of physical devices should employ a user model containing a list of spe-

cific objects in the knowledge base that the user knows (e.g., *telephone*), and a list of general, abstract concepts known to the user that underlie the mechanisms of devices (e.g., *electricity*). Paris notes that such explicit representation of the user's knowledge gives considerably greater flexibility to an NLG system to tailor effective responses than do user models that characterize users in terms of a limited number of fixed stereotypes, labeled, for example, as naive, intermediate, or expert.

Through analyses of naturally occurring texts, Paris establishes convincingly in Chapter 3 that descriptions of complex devices do vary according to the assumed level of domain knowledge of the intended readers. Descriptions have been drawn from a wide variety of sources such as junior and adult encyclopedias, high-school textbooks, and car manuals, and analyzed using rhetorical predicates (McKeown 1985). The analyses reveal Paris's chief discovery, which is that descriptions aimed at different types of audiences present different *types* of information and are structured differently.

Texts aimed at expert audiences describe physical devices chiefly in terms of their subcomponents and their attributes. The structure of such texts could be captured in terms of the *constituent schema* (McKeown 1985). However, texts intended for domain novices are process-oriented, in that they primarily express the causal relations in the knowledge base and literally *trace* the processes underlying the function of the device. This discovery leads naturally to the invention of *process trace*, another major result of Paris's research, as a discourse strategy that captures the organization of process-oriented texts aimed at novices. Chapter 4 is devoted mainly to the technical details of the process trace.

Paris represents both constituency schema and process trace strategies as augmented transition networks. Using the same representational formalism for the two strategies is of considerable significance, as it enables the TAILOR system to switch from one strategy to another mid-discourse.

Chapter 5, perhaps the most significant chapter in the book, presents details of how a strategy is initially chosen for generating a description, and when and under what conditions a switch of strategies is possible. The conditions for choosing and switching strategies are articulated primarily in terms of the user's domain knowledge. Thus, the explicit representation of the user's domain knowledge in the user model, combined with the ability to switch strategies, enables TAILOR to generate texts suitable for a wide range of users who are neither naive nor expert.

The implementation details of TAILOR are presented in Chapter 6. This chapter, the longest in the book, highlights the important modules of the system. The presentation is by and large at the appropriate level of detail, which is neither too high-level (sketchy) nor too profuse with program code (tedious). Chapter 7 presents an overview of related work in NLG and user modeling. Finally, Chapter 8 reflects on the research presented thus far in the book and discusses directions of further research that have followed Paris's main work.

The exposition in this book is felicitous, and the cross-referencing between the chapters judicious. The book is essentially self-contained; however, familiarity with (or access to) McKeown's book (1985) will be helpful, especially to complement the brief presentation of the constituency schema in Chapter 4 and the overview of the ATN implementation of the strategies in Section 6.4.3. Full appreciation of discussions in Section 8.3 requires familiarity with some of the major problems of current interest to NLG researchers. Paris, with clear explication of her methodology (e.g., in Chapter 3), has delivered a book that is instructive in methodology as well.

The book contains a good set of appropriately placed citations and an adequate bibliography, which, together with Chapter 7, serve as a useful resource for researchers both seasoned and raw. A minor deficiency, however, may be noted here: most refer-

ence entries for journal articles, book chapters, and conference proceedings lack page numbers of the items in the respective volumes.[1]

It may be pointed out that generator decisions depending on the user's domain knowledge are conditional not only on *what the user knows*, which is explicitly represented in the user model, but also on:

- *how much the user knows*. For example, "the user knows about most parts of the device": this involves knowing the size of the concept lists in the user model relative to the size of a portion of the knowledge base.

- *what the user does not know*. For example, "check that the process information does not involve basic concepts the user does not know, and use process trace" (Figure 6.15, page 125). There is an implicit closed-world assumption in the user model. However, it is not clear how this check is performed in TAILOR. Does the knowledge base contain a separate list or taxonomy of all basic concepts? As far as I could see, apart from basic concept names that might appear in the *function* slots of object representations, there is no inventory of basic concepts in the knowledge base. How does TAILOR determine that process information does involve basic concepts?

Terminology from the register theory of Systemic-Functional Linguistics (SFL) is introduced in the opening chapter of the book (Section 1.1.1). However, no use is made of it in the expositions in the main chapters of the book (Chapters 3–6). The reader, especially one insufficiently acquainted with SFL, is apt to lose sight of this terminology and may need re-introduction when it is used again much later, in Section 8.3.3. It may have been better to postpone introduction of register-theoretic terms to Section 8.3.3 itself. On a positive note, it should be mentioned that Paris provides tantalizing connections between contextual elements in NLG (especially, user models) and register theory. On the one hand, it will be fruitful to employ register theory as a theoretical framework for conveying the influence of situational variables in NLG. Reciprocally, it will be equally fruitful to enrich linguistic theory (in particular, SFL) with Paris's insights on how discourse structure and content vary according to tenor.

The results presented in this book will be useful to researchers and system developers in intelligent tutoring systems and computer-assisted instruction. The usefulness, however, is limited by the lack of integration of TAILOR's main processes into a dialog system. This limitation is not particular to this book, but is characteristic of much work on monological multisentential text generation. The state of the art in NLG does not as yet permit systems like TAILOR to be integrated easily into those engaging in flexible dialog with users. More positively, we may regard this as a major research direction that merits further development. An important desideratum for dialog capabilities seems to be an explicit representation of the speaker's (generator's) intentions. Paris discusses current work on this problem in Section 8.3.2 of the book.

In sum, this book constitutes a significant step forward, provides impetus for further research along several directions, and is useful to researchers and students

---

1 Some minor errata: The English output for the process trace of Figures 4.21 and 4.22 in the book is shown in Figure 4.23, and not, as mentioned on page 83, in Figure 4.5. On page 102, we read that in the text shown in Figure 5.9, the same user model as in Figure 5.5 is used except for a minor change in the value of *most*. However, Figures 5.9 and 5.5 show user models with *different* basic concepts, viz., *nil* and *loudspeaker*, respectively. The few other editorial oversights that remain are inconsequential, and imperceptible unless you look for them with a vigilant pair of eyes.

professing diverse interests in the study of communication in context. Perhaps the strongest selling point of Paris's book is that it contains material of interest and value to a wide variety of users ranging from the naive to the expert.

Reference

McKeown, Kathleen R. (1985). *Text Generation: Using Discourse Strategies and Focus Constraints to Generate Natural Language Text*. Cambridge University Press.

T. Pattabhiraman is a research scientist in the Natural Language Laboratory at Simon Fraser University. His Ph.D. work was on the role of salience in natural language generation. He guest-edited two special issues of *Computational Intelligence* on generation. His address is: Centre for Systems Science, Simon Fraser University, Burnaby, B.C., Canada V5A 1S6; e-mail: patta@cs.sfu.ca.