

# Building Natural Language Generation Systems

Ehud Reiter  
Department of Computing Science  
University of Aberdeen  
King's College  
Aberdeen AB9 2UE, BRITAIN  
email: ereiter@csd.abdn.ac.uk

## 1 Introduction

Natural Language Generation (NLG) systems generate texts in English (or other human languages, such as French) from computer-accessible data. NLG systems are (currently) most often used to help human authors write routine documents, including business letters [SBW91] and weather reports [GDK94]. They also have been used as interactive explanation tools which communicate information in an understandable way to non-expert users, especially in software engineering (eg, [RK92]) and medical (eg, [BMF<sup>+</sup>94]) contexts.

From a technical perspective, almost all applied NLG systems perform the following three tasks [Rei94]:

**Content Determination and Text Planning:** Decide what information should be communicated to the user (content determination), and how this information should be rhetorically structured (text planning). These tasks are usually done simultaneously.

**Sentence Planning:** Decide how the information will be split among individual sentences and paragraphs, and what cohesion devices (eg, pronouns, discourse markers) should be added to make the text flow smoothly.

**Realization:** Generate the individual sentences in a grammatically correct manner.

In the rest of this paper, I shall briefly discuss different ways of performing each of these tasks

## 2 Content Determination and Text Planning

Content determination (deciding what information to communicate in the text) and text-planning (organizing the information into a rhetorically coherent structure) are done simultaneously in most applied NLG systems [Rei94]. These tasks can be done at many different levels of sophistication. One of the simplest (and most common) approaches is simply to write a 'hard-coded' content/text-planner in a standard programming language (C++, Lisp, etc). The resultant system may lack flexibility, but if the texts being produced have a standardized content and structure (which is true in many technical domains), then this can be the most effective way to perform these tasks.

On the other end of the sophistication spectrum, many standard AI techniques have been adopted for content determination and text planning, including rule-based systems [RML95] and planning [Hov88]. Systems built in this way are in principal very flexible and powerful, although in practice they have sometimes not been robust enough for real-world use.

An intermediate approach which has been quite popular is to use a special 'schema' or 'text-planning' language [McK85, KKR91]. Such languages typically allow the developer to represent

text plans as transition networks of one sort or another, with the nodes giving the information content and the arcs giving the rhetorical structure. In many cases text-planning languages are implemented as macro packages, which gives the developer access to the full power of the underlying programming language whenever necessary.

### 3 Sentence Planning

Sentence planning includes

- Conjunction and other aggregation. For example, transforming (1) into (2):
  - 1) Sam has high blood pressure. Sam has low blood sugar.
  - 2) Sam has high blood pressure *and* low blood sugar.
- Pronominalization and other reference. For example, transforming (3) into (4):
  - 3) I just saw Mrs. Black. Mrs Black has a high temperature.
  - 4) I just saw Mrs. Black. *She* has a high temperature.
- Introducing discourse markers. For example, transforming (5) into (6):
  - 5) If Sam goes to the hospital, he should go to the store.
  - 6) If Sam goes to the hospital, he should *also* go to the store.

The common theme behind these operations is they do not change the information content of the text, but they do make it more fluent and easily readable.

Sentence planning is important if the text needs to read fluently and, in particular, if it should look like it was written by a human (which is usually the case for business letters, for example). If it doesn't matter if the text sounds stilted and was obviously produced by a computer, then it may be possible to de-emphasize sentence planning, and perform minimal aggregation, use no pronouns, etc.

If the text does need to look fluent, then a good job of sentence planning is essential. There are formal models of all of the operations mentioned above, and some applied NLG systems have incorporated them, eg, [MKS94]. It is also possible to do effective sentence-planning in an ad-hoc manner, at least in a limited domain; Knowledge Point's *Performance Now* system is a good example of this.

### 4 Realization

A Realizer generates individual sentences (typically from a 'deep syntactic' representation [Rei94]). The realizer needs to make sure that the rules of English are obeyed, including

- Point absorption and other punctuation rules. For example, the sentence *I saw Helen Jones, my sister-in-law* should end in ".", not ".,"
- Morphology. For example, the plural of *box* is *boxes*, not *boxs*.
- Agreement. For example, *I am here* instead of *I is here*.
- Reflexives. For example, *John saw himself*, instead of *John saw John*.

There are numerous linguistic formalisms and theories which can be incorporated into an NLG Realizer, far too many to describe here. There are also some general-purpose 'engines' which can be programmed with various linguistic rules, such as FUF [Elh92] and PENMAN [Pen89].

In many cases, acceptable performance can be achieved without using complex linguistic modules. In particular, if only a few different types of sentences are being generated, then it may be simpler and cheaper to use fill-in-the-blank templates for realization, instead of ‘proper’ syntactic processing.

## 5 Conclusion

Many different techniques are available for performing the three NLG tasks of content determination (and text planning), sentence planning, and realization. These techniques range from the simplistic to the extremely sophisticated, and it is impossible to say that one is always ‘better’ than another. It all depends on the characteristics of the application, such as whether extensive filtering and summarization of information is needed, whether texts need to ‘look like they were written by a person’, and how much syntactic variety is expected to occur in the generated texts. A good NLG engineer will choose the most appropriate set of techniques, given the needs of the application [RM93] and the available resources.

## References

- [BMF<sup>+</sup>94] B. Buchanan, J. Moore, D. Forsythe, G. Carenini, and S. Ohlsson. Using medical informatics for explanation in a clinical setting. Technical Report 93-16, Intelligent Systems Laboratory, University of Pittsburgh, 1994.
- [Elh92] Michael Elhadad. *Using Argumentation to Control Lexical Choice: A Functional Unification Implementation*. PhD thesis, Columbia University, 1992.
- [GDK94] Eli Goldberg, Norbert Driedger, and Richard Kittredge. Using natural-language processing to produce weather forecasts. *IEEE Expert*, 9(2):45–53, 1994.
- [Hov88] Eduard Hovy. Planning coherent multisentential text. In *Proceedings of 26th Annual Meeting of the Association for Computational Linguistics (ACL-88)*, pages 163–169, 1988.
- [KKR91] Richard Kittredge, Tanya Korelsky, and Owen Rambow. On the need for domain communication language. *Computational Intelligence*, 7(4):305–314, 1991.
- [McK85] Kathleen McKeown. *Text Generation*. Cambridge University Press, 1985.
- [MKS94] Kathleen McKeown, Karen Kukich, and James Shaw. Practical issues in automatic document generation. In *Proceedings of the Fourth Conference on Applied Natural-Language Processing (ANLP-1994)*, pages 7–14, 1994.
- [Pen89] Penman Natural Language Group. The Penman user guide. Technical report, Information Sciences Institute, Marina del Rey, CA 90292, 1989.
- [RK92] Owen Rambow and Tanya Korelsky. Applied text generation. In *Proceedings of the Third Conference on Applied Natural Language Processing (ANLP-1992)*, pages 40–47, 1992.
- [Rei94] Ehud Reiter. Has a consensus NL Generation architecture appeared, and is it psycholinguistically plausible? In *Proceedings of the Seventh International Workshop on Natural Language Generation (INLGW-1994)*, pages 163–170, 1994.
- [RM93] Ehud Reiter and Chris Mellish. Optimising the costs and benefits of natural language generation. In *Proceedings of the 13th International Joint Conference on Artificial Intelligence (IJCAI-1993)*, volume 2, pages 1164–1169, 1993.
- [RML95] Ehud Reiter, Chris Mellish, and John Levine. Automatic generation of technical documentation. *Applied Artificial Intelligence*, 9(3):259–287, 1995.
- [SBW91] Stephen Springer, Paul Buta, and Thomas Wolf. Automatic letter composition for customer service. In Reid Smith and Carlisle Scott, editors, *Innovative Applications of Artificial Intelligence 3 (Proceedings of CAIA-1991)*. AAAI Press, 1991.