GENERATION - A NEW FRONTIER OF NATURAL LANGUAGE PROCESSING?

Aravind K. Joshi

Department of Computer and Information Science
University of Pennsylvania

Comprehension and generation are the two complementary aspects of natural language processing (NLP). However, much of the research in NLP until recently has focussed on comprehension. Some of the reasons for this almost exclusive emphasis on comprehension are (1) the belief that comprehension is harder than generation, (2) problems in comprehension could be formulated in the AI paradigm developed for problems in perception, (3) the potential areas of applications seemed to call for comprehension more than generation, e.g., question-answer systems, where the answers can be presented in some fixed format or even in some non-linguistic fashion (such as tables), etc. Now there is a flurry of activity in generation, and we are definitely going to see a significant part of future NLP research devoted to generation. A key motivation for this interest in generation is the realization that many applications of NLP require that the response produced by a system must be flexible (i.e., not produced by filling in a fixed set of templates) and must often consist of a sequence of sentences (i.e., a text) which must have a textual structure (and not just an arbitrary sequence of sentences containing the necessary information). As the research in generation is taking roots, a number of interesting theoretical issues have become very important, and these are likely to determine the paradigm of research in this "new" area.

Based on the input from several researchers in NLP, I prepared a set of questions that the panel on Generation was invited to address in their position papers. These questions were as follows:

- What is the relationship between NL comprehension and generation?
Is there inherently an asymmetry between comprehension and generation?
Is comprehension more heuristic than generation?
- Will the demands of language generation bring AI and linguistics closer together than the demands of comprehension did in the past. Is there something special about generation?

- Does generation constrain the problem differently from comprehension in that it would not matter if some high-powered machine could comprehend things no human could say, but would matter if the same machine generated them.
- How should the generation and comprehension capabilities of a system be matched. By looking at the sentences or texts a system generates, the user may ascribe comprehension capabilities to the system, which the system may or may not have. In other words how will generation affect user's behavior with respect to the input he/she provides to the system?
- Are knowledge structures of the world as much as language, the same or different for comprehension and generation?
- How does one control for syntactic choice and lexical choice?
- What is the status of different grammatical formalisms with respect to generation? Should the formalism be the same for generation as for comprehension?

The panelists have chosen to focus on some of these questions. They have, of course, raised some additional questions. Some of the key issues discussed by the panelists are as follows. Appelt has explored the notion of bidirectional grammars, i.e., grammars that can be used by processors of approximately equal computational complexity to parse and generate sentences of language. In this sense, he wants to treat comprehension and generation as strict inverses of each other. He suggests that by using bidirectional grammars the problems of maintaining consistency between comprehension and generation components when one of them changes can be eliminated. Kroch is concerned with the limits on the capacity of the human language generation mechanism, which translates preverbal messages into sentences of a natural language. His main point is that there are limits to the competence the generation mechanism is trying to model. He suggests some theoretical characterizations of these limits that should help in circumscribing the problem of generation. McDonald points out that although one could have a common representation of linguistic knowledge, the processes that draw on this knowledge for comprehension and generation cannot be the same because of the radical differences in information flow. He also points out that in generation it is difficult to ignore syntax and control

of variation of linguistic form. Mann considers various aspects of lexicon, grammar, and discourse from the point of view of comprehension and generation. Although both comprehension and generation have to deal with all these problems, there are differences with respect to particular problems addressed in generation. He suggests that these differences arise because the technical problems that limit the quality of generated text are very different from the corresponding set of problems that limits the quality of comprehension. Marcus focusses on the problem of lexical choice, which has not received much attention in the work on generation so far. He suggests that if the generation systems are to be both fluent and portable, they must know about both words and meanings. He is concerned about the fact that much of the current research on generation has focussed on subtle and difficult matters as responding appropriately to the user's intentions, correctly utilizing rhetorical structures etc., but it has avoided the issue of what would make such systems mean the literal content of the words they use.

Comprehension and generation, when viewed as functions mapping from utterances to meanings and intentions and vice versa, can certainly be regarded as inverses of each other. However, these functions are enormously complex and therefore, although at the global level they are inverses of each other, the inverse transformation (i.e, computation of one function from the other) is not likely to be so direct. So, in this sense, there may be an asymmetry between comprehension and generation even at the theoretical level. There is an asymmetry certainly at the practical level. In comprehension, under certain circumstances, some of the linguistic knowledge may be ignored (of course, at some cost) by utilizing some higher levels of knowledge, which is required in any case. However, under the same circumstances, one cannot avoid the use of the very same linguistic knowledge in generation, the quality of the output becomes quite unacceptable to a human user very rapidly, otherwise. It is this asymmetry that, I think, will force us to examine in detail the relationship between grammar, lexicon, and message planning and may elucidate the relationship between linguistic knowledge and conceptual knowledge. All these questions are equally relevant to comprehension. However, work on generation seems to require us to be more sensitive to these relationships than we may have been in the past, when the focus was on comprehension only.

Comprehension and generation are not just inverses, they are related to each other also in another manner. The human generation mechanism also involves some monitoring of the output,

presumably by the comprehension mechanism. Computer generation systems so far have not been concerned with this issue (as far as I know). The generation and comprehension components work independently, even if they share some procedures and data structures, they have no knowledge of each other. Whether or not comprehension and generation should be related to each other in this sense in a computer system is an open question and needs considerable attention. The panelists have not paid much attention to this question (one of them has declared it as a non-problem). Perhaps, the audience will make some contributions here.

Acknowledgements

This work is partially supported by DARPA grants NOOO14-85-K-0018 and NOOO14-85-K-0807, NSF grants MCS8219196-CER, MCS-82-07294, 1 RO1-HL-29985-01, U.S. Army grants DAA6-29-84-K-0061, DAAB07-84-K-F077, U.S. Air Force grant 82-NM-299, AI Center grants NSF-MCS-83-05221.