

PARSING

Ralph Grishman
Dept. of Computer Science
New York University
New York, N. Y.

One reason for the wide variety of views on many subjects in computational linguistics (such as parsing) is the diversity of objectives which lead people to do research in this area. Some researchers are motivated primarily by potential applications - the development of natural language interfaces for computer systems. Others are primarily concerned with the psychological processes which underlie human language, and view the computer as a tool for modeling and thus improving our understanding of these processes. Since, as is often observed, man is our best example of a natural language processor, these two groups do have a strong commonality of research interest. Nonetheless, their divergence of objective must lead to differences in the way they regard the component processes of natural language understanding. (If - when human processing is better understood - it is recognized that the simulation of human processes is not the most effective way of constructing a natural language interface, there may even be a deliberate divergence in the processes themselves.) My work, and this position paper, reflect an applications orientation; those with different research objectives will come to quite different conclusions.

WHY PARSE?

One of the tasks of computer science in general, and of artificial intelligence in particular, is that of coping in a systematic fashion with systems of high complexity. Natural language interfaces certainly fit that characterization.

A natural language interface must analyze input sequences, communicate with some underlying system (data base, robot, etc.), and generate responses. In the transition from the natural language input to the language of the underlying system there is in principle no need to make explicit reference to any intermediate structures; we could write our interface as a (huge) set of rules which map directly from input sequences into our target language. We know full well, however, that such a system would be nearly impossible to write, and certainly impossible to understand or modify. By introducing intermediate structures, we are able to divide the task into more manageable components.

Specific intermediate structures are of value insofar as they facilitate the expression of relationships which must be captured in the system - relationships which would be more cumbersome to express using other representations. For example, the representations at the level of logical form (such as predicate calculus) are chosen to facilitate the computation of logical inferences.

In the same way, a representation of constituent structure (a parse tree), if properly chosen, will facilitate the statement of many linguistic constraints and relationships. Grammatical constraints will enable the system to identify the pertinent syntactic category for many multiply classified words. Some constraints on anaphora (such as the notion of command) and on quantifier structure are also best stated in terms of surface structure.

Equally important, many sentence relationships which must be captured at some point in the analysis (such as the relation between active and passive sentences or between reduced and expanded conjoinings) are most easily stated as transformations between constituent structures. By using syntactic transformations to regularize the

constituent structure, we can substantially simplify the specification of the subsequent stages of analysis.

SPECIFICATION VS. PROCEDURE

The arguments just given for parse trees (and other intermediate structures) are arguments for how best to specify the transformations which a natural language input must undergo. They are not arguments for a particular language analysis procedure. A direct implementation of the simplest specifications does not necessarily yield the most efficient procedure; as our systems become more sophisticated, the distance from specification to implementation structure may increase. We should therefore favor formalisms which (because of their simple structure) can be automatically adapted to a variety of procedures. Among these variations are:

PARALLEL PROCESSING. Phrase structure grammars and augmented phrase structure grammars lend themselves naturally to parallel parsing procedures - either top-down (following alternative expansions in parallel), bottom-up (trying alternative reductions in parallel), or a combination of the two. In particular, some of the parsing algorithms developed as part of the speech recognition research of the past decade are readily adaptable to parallel processing. To minimize parallelism, however, the grammatical constraints must be organized to minimize or at least postpone the interactions among the analyses of the various parts of a sentence.

ANALYSIS AND GENERATION. In the same way that sentence analysis involves a translation to a "deep structure," an increasing number of systems now include a generation component to translate from deep structure to sentences. If the mapping from sentence to deep structure is direct (without reference to a parse tree), the generation component may require a separate design effort. On the other hand, if the mapping is specified in terms of incremental transformations of the constituent structure, producing an inverse mapping may be relatively straightforward (and the greater the non-procedural content of the transformations, the easier it should be to reverse them).

AVOIDING THE PARSE TREE. To emphasize the distinction between specification and procedure, let me mention a possibility for an "optimizing" analyser of the future: one whose specifications are given in terms of transformations of the constituent structure followed by interpretation of the regularized ("deep") structure, but whose implementation avoids actually constructing a parse tree. Instead, the transformations would be applied to the deep structure interpretation rules, producing a (much larger) set of rules for interpreting the input sequences directly. Some small experiments have been done in this direction (K. Konolige, "Capturing Linguistic Generalizations with Grammar Metarules," Proc. 18th Ann'l Meeting ACL, 1979). By avoiding explicit construction of a parse tree, we could accelerate the analysis procedure while retaining the descriptive advantages of independent, incremental transformations of constituent structure. While development of any such automatic grammar restructuring procedure would certainly be a difficult task, it does indicate the possibilities which open up when specification and implementation are separated.

