SOME ISSUES IN PARSING AND NATURAL LANGUAGE UNDERSTANDING

Robert J. Bobrow
Bolt Beranek and Newman Inc.

Bonnie L. Webber
Department of Computer & Information Science
University of Pennsylvania

Language is a system for encoding and transmitting ideas. A theory that seeks to explain linguistic phenomena in terms of this fact is a functional theory. One that does not misses the point. [10]

PREAMBLE

Our response to the questions posed to this panel is influenced by a number of beliefs (or biases!) which we have developed in the course of building and analyzing the operation of several natural language understanding (NLU) systems. [1, 2, 3, 12] While the emphasis of the panel is on parsing, we feel that the recovery of the syntactic structure of a natural language utterance must be viewed as part of a larger process of recovering the meaning, intentions and goals underlying its generation. Hence it is inappropriate to consider designing or evaluating natural language parsers or grammars without taking into account the architecture of the whole NLU system of which they're a part.¹ This is the premise from which our beliefs arise, beliefs which concern two things:

- o the distribution of various types of knowledge, in particular syntactic knowledge, among the modules of an NLU system
- o the information and control flow among those modules.

As to the first belief, in the NLU systems we have worked on, most syntactic information is localized in a "syntactic module", although that module does not produce a reified data structure representing the syntactic description of an utterance. Thus, if "parsing" is taken as requiring the production of such a reified structure, then we do not believe in its necessity. However we do believe in the existence of a module which provides syntactic information to those other parts of the system whose decisions ride on it.

As to the second belief, we feel that syntax, semantics and pragmatics effectively constitute parallel but interacting processors, and that information such as local syntactic relations is determined by joint decisions among them. Our experience shows that with minimal loss of efficiency, one can design these processors to interface cleanly with one another, so as to allow independent design, implementation and modification. We spell out these beliefs in slightly more detail below, and at greater length in [4].

¹We are not claiming that the only factors shaping a parser or a grammar, beyond syntactic considerations, are things like meaning, intention, etc. There are clearly mechanical and memory factors, as well as laziness - a speaker's penchant for trying to get away with the minimal level of effort needed to accomplish the task!

RESPONSES

The Computational Perspective

The first set of questions to this panel concern the computational perspective, and the useful purposes served by distinguishing parsing from interpretation. We believe that syntactic knowledge plays an important role in NLU. In particular, we believe that there is a significant type of utterance description that can be determined on purely syntactic grounds², albeit not necessarily uniquely. This description can be used to guide semantic and discourse level structure recovery processes such as interpretation, anaphoric resolution, focus tracking, given/new distinctions, ellipsis resolution, etc. in a manner that is independent of the lexical and conceptual content of the utterance. There are several advantages to factoring out such knowledge from the remainder of the NLU system and providing a "syntactic module" whose interactions with the rest of the system provide information on the syntactic structure of an utterance. The first advantage is to simplify system building, as we know from experience [1, 2, 3, 4, 5, 12]. Once the pattern of communication between processors is settled, it is easier to attach a new semantics to the hooks already provided in the grammar than to build a new semantic processor. In addition, because each module has only to consider a portion of the constraints implicit in the data (e.g. syntactic constraints, semantic constraints and discourse context), each module can be designed to optimize its own processing and provide an efficient system.

The panel has also been charged with considering parallel processing as a challenge to its views on parsing. This touches on our beliefs about the interaction among the modules that comprise the NLU system. To respond to this issue, we first want to distinguish between two types of parallelism: one, in which many instances of the same thing are done at once (as in an array of parallel adders) and another, in which the many things done simultaneously can be different. Supporting this latter type of parallelism doesn't change our view of parsing, but rather underlies it. We believe that the interconnected processes involved in NLU must support a basic operating principle that Norman and Bobrow [14] have called "The Principle of Continually Available Output" (CAO). This states that the interacting processes must begin to provide output over a wide range of resource allocations, even before their analyses are complete, and even before all input data is available. We take this position for two reasons: one, it facilitates computational efficiency, and two, it seems to be closer to human parsing processes (a point which we will get to in answering the next question).

The requirement that syntactic analysis, semantic interpretation and discourse processing must be able to operate in (pseudo-)parallel, obeying the CAO

²that is, solely on the basis of syntactic categories/features and ordering information

principle, has sparked our interest in the design of pairs of processes which can pass forward and backward useful information/advice/questions as soon as possible. The added potential for interaction of such processors can increase the capability and efficiency of the overall NLU process. Thus, for example, if the syntactic module makes its intermediate decisions available to semantics and/or pragmatics, then those processors can evaluate those decisions, guide syntax's future behavior and, in addition, develop in parallel their own analyses. Having sent on its latest assertion/advice/question, whether syntax then decides to continue on with something else or wait for a response will depend on the particular kind of message sent. Thus, the parsers and grammars that concern us are ones able to work with other appropriately designed components to support CAO. While the equipment we are using to implement and test our ideas is serial, we take very seriously the notion of parallelism.

Finally under the heading of "Computational Perspective", we are asked about what might motivate our trying to make parsing procedures simulate what we suspect human parsing processes to be like. One motivation for us is the belief that natural language is so tuned to the part extraordinary, part banal cognitive capabilities of human beings that only by simulating human parsing processes can we cover all and only the language phenomena that we are called upon to process. A particular (extraordinary) aspect of human cognitive (and hence, parsing) behavior that we want to explore and eventually simulate is people's ability to respond even under degraded data or resource limitations. There are examples of listeners initiating reasonable responses to an utterance even before the utterance is complete, and in some case even before a complete syntactic unit has been heard. Simultaneous translation is one notable example [8], and another is provided by the performance of subjects in a verbally guided assembly task reported by P. Cohen [6]. Such an ability to produce output before all input data is available (or before enough processing resources have been made available to produce the best possible response) is what led Norman and Bobrow to formulate their CAO Principle. Our interest is in architectures for NLU systems which support CAO and in search strategies through such architectures for an optimal interpretation.

The Linguistic Perspective

We have been asked to comment on legitimate inferences about human linguistic competence and performance that we can draw from our experiences with mechanical parsing of formal grammar. Our response is that whatever parsing is for natural languages, it is still only part of a larger process. Just because we know what parsing is in formal language systems, we do not necessarily know what role it plays in the context of total communication. Simply put, formal notions of parsing underconstrain the goals of the syntactic component of an NLU system. Efficiency measures, based on the resources required for generation of one or all complete parses for a sentence, without semantic or pragmatic interaction, do not necessarily specify desirable properties of a natural language syntactic analysis component.

As for whether the efficiency of parsing algorithms for CF or regular grammars suggest that the core of NL grammars is CF or regular, we want to distinguish that part of perception (and hence, syntactic analysis) which groups the stimulus into recognizable units from that part which fills in gaps in information (inferentially) on the basis of such groups. Results in CF grammar theory says that grouping is not best done purely bottom-up, that there are advantages to

using predictive mechanisms as well [9, 7]. This suggests two things for parsing natural language:

1. There is a level of evidence and a process for using it that is working to suggest groups.
2. There is another filtering, inferencing mechanism that makes predictions and diagnoses on the basis of those groups.

It is possible that the grouping mechanism may make use of strategies applicable to CF parsing, such as well-formed substring tables or charts, without requiring the overall language specification be CF. In our current RUS/PSI-KLONE system, grouping is a function of the syntactic module: its output consists of suggested groupings. These suggestions may be at abstract, specific or disjunctive. For example, an abstract description might be "this is the head of an NP, everything to its left is a pre-modifier". Here there is no comment about exactly how these pre-modifiers group. A disjunctive description would consist of an explicit enumeration of all the possibilities at some point (e.g., "this is either a time prepositional phrase (PP) or an agentive PP or a locative PP, etc."). Disjunctive descriptions allow us to prune possibilities via case analysis.

In short, we believe in using as much evidence from formal systems as seems understandable and reasonable, to constrain what the system should be doing.

The Interactions

Finally, we have been asked about the nature of the relationship between a grammar and a procedure for applying it. On the systems building side, our feeling is that while one should be able to take a grammar and convert it to a recognition or generation procedure [10], it is likely that such procedures will embody a whole set of principles that are control structure related, and not part of the grammar. For example, a grammar need not specify in what order to look for things or in what order decisions should be made. Thus, one may not be able to reconstruct the grammar uniquely from a procedure for applying it.

On the other hand, on the human parsing side, we definitely feel that natural language is strongly tuned to both people's means of production and their means of recognition, and that principles like McDonalds' Indelibility Principle [13] or Marcus' Determinism Hypothesis [11] shape what are (and are not) seen as sentences of the language.

REFERENCES

1. Bobrow, R. J. The RUS System. BBN Report 3878, Bolt Beranek and Newman Inc., 1978.
2. Bobrow, R. J. & Webber, B. L. PSI-KLONE - Parsing and Semantic Interpretation in the BBN Natural Language Understanding System. CSCSI/CSEIO Annual Conference, CSCSI/CSEIO, 1980.
3. Bobrow, R. J. & Webber, B. L. Knowledge Representation for Syntactic/Semantic Processing. Proceedings of The First Annual National Conference on Artificial Intelligence, American Association for Artificial Intelligence, 1980.

4. Bobrow, R.J. & Webber, B.L. Parsing and Semantic Interpretation as an Incremental Recognition Process. Proceedings of a Symposium on Modelling Human Parsing Strategies, Center for Cognitive Science, University of Texas, Austin TX, 1981.
5. Bobrow, R.J. & Webber, B.L. Systems Considerations for Search by Cooperating Processes: Providing Continually Available Output. Proceedings of the Sixth IJCAI, International Joint Conference on Artificial Intelligence, 1981.
6. Cohen, P. personal communication. videotape of experimental task
7. Earley, J. An efficient context-free parsing algorithm. Communications of the ACM 13 (February 1970), 94-102.
8. Goldman-Eisler, F. Psychological Mechanisms of Speech Production as Studied through the Analysis of Simultaneous Translation. In B. Butterworth, Ed., Language Production, Academic Press, 1980.
9. Graham, S., Harrison, M. and Ruzzo, W. An Improved Context-Free Recognizer. ACM Transactions on Programming Languages and Systems (July 1980), 416-463.
10. Kay, M. An Algorithm for Compiling Parsing Tables from a Grammar. Proceedings of a Symposium on Modelling Human Parsing Strategies, Center for Cognitive Science, University of Texas, Austin TX, 1981.
11. Marcus, M. A Theory of Syntactic Recognition for Natural Language. MIT Press, 1980.
12. Mark, W. S. & Barton, G. E. The RUSGrammar Parsing System. GMR 3243, General Motors Research Laboratories, 1980.
13. McDonald, D. ????. Ph.D. Th., Massachusetts Institute of Technology, 1980.
14. Norman, D. & Bobrow, D. On Data-limited and Resource-limited Processes. CSL74-2, Xerox PARC, May, 1974.