FUTURE PROSPECTS FOR COMPUTATIONAL LINGUISTICS

Gary G. Hendrix
SRI International

Preparation of this paper was supported by the Defense Advance Research Projects Agency under contract N00039-79-C-0118 with the Naval Electronic Systems Command. The views expressed are those of the author.

A. Introduction

For over two decades, researchers in artificial intelligence and computational linguistics have sought to discover principles that would allow computer systems to process natural languages such as English. This work has been pursued both to further the scientific goals of providing a framework for a computational theory of natural-language communication and to further the engineering goals of creating computer-based systems that can communicate with their human users in human terms. Although the goal of fluent machine-based nautral-language understanding remains elusive, considerable progress has been made and future prospects appear bright both for the advancement of the science and for its application to the creation of practical systems.

In particular, after 20 years of nurture in the academic nest, natural-language processing is beginning to test its wings in the commercial world [8]. By the end of the decade, natural-language systems are likely to be in widespread use, bringing computer resources to large numbers of non-computer specialists and bringing new credibility (and hopefully new levels of funding) to the research community.

B. Basis for Optimism

My optimism is based on an extrapolation of three major trends currently affecting the field:

(1) The emergence of an engineering/applications discipline within the computational-linguistics community.

(2) The continuing rapid development of new computing hardware coupled with the beginning of a movement from time-sharing to personal computers.

(3) A shift from syntax and semantics as the principle objects of study to the development of theories that cast language use in terms of a broader theory of goal-motivated behavior and that seek primarily to explain how a speaker's cognitive state motivates him to engage in an act of communication, how a speaker devises utterances with which to perform the act, and how acts of communication affect the cognitive states of hearers.

C. The Impact of Engineering

The emergence of an engineering discipline may strike many researchers in the field as being largely detached from the mainstream of current work. But I believe that, for better or worse, this discipline will have a major and continuing influence on our research community. The public at large tends, often unfairly, to view a science through the products and concrete results it produces, rather than through the mysteries of nature it reveals. Thus, the chemist is seen as the person who produces fertilizer, food coloring and nylon stockings; the biologist finds cures for diseases; and the physicist produces moon rockets, semiconductors, and nuclear power plants. What has computational linguistics produced that has affected the lives of individuals outside the limits of its own close-knit community? As long as the answer remains "virtually nothing," our work will generally be viewed as an ivory tower enterprise. As soon as the answer becomes a set of useful computer systems, we will be viewed as the people who produce such systems and who aspire to produce better ones.

My point here is that the commercial marketplace will tend to judge both our science and our engineering in terms of our existing or potential engineering products. This is, of course, rather unfair to the science; but I believe that it bodes well for our future. After all, most of the current sponsors of research on computational linguistics understand the scientific nature of the enterprise and are likely to continue their support even in the face of minor successes on the engineering front. The impact of an engineering arm can only add to our field's basis of support by bringing in new suport from the commercial sector.

One note of caution is appropriate, however. There is a real possibility that as commercial enterprises enter the natural-language field, they will seek to build in-house groups by attracting researchers from universities and nonprofit institutions. Although this would result in the creation of more jobs for computational linguists, it would also result in proprietary barriers being established between research groups. The net effect in the short term might actually be to retard scientific progress.

D. The State of Applied Work

1. Accessing Databases

Currently, the most commercially viable task for natural-language processing is that of providing access to databases. This is because databases are among the few types of symbolic knowledge representations that are computationally efficient, are in widespread use, and have a semantics that is well understood.

In the last few years, several systems, including LADDER [9], PLANES [29], REL [26], and ROBOT [8], have achieved relatively high levels of proficiency in this area when applied to particular databases. ROBOT has been introduced as a commercial product that runs on large, mainframe computers. A pilot REL product is currently under development that will run on a relatively large personal machine, the HP 9845. This system, or something very much like it, seems likely to reach the marketplace within the next two or three years. Should ROBOT- and REL-like systems prove to be commercial successes, other systems with increasing levels of sophistication are sure to follow.

2. Immediate Problems

A major obstacle currently limiting the commercial viability of natural-language access to databases is the problem of telling systems about the vocabulary, concepts and linguistic constructions associated with new databases. The most proficient of the application systems have been hand-tailored with extensive knowledge for accessing just ONE database. Some systems (e.g., ROBOT and REL) have achieved a degree of transportability by using the database itself as a source of knowledge for guiding linguistic processes. However, the knowledge available in the database is generally rather limited. High-performance systems need access to information about the larger enterprise that provides the context in which the database is to be used.

As pointed out by Tennant [27], users who are given natural-language access to a database expect not only to retrieve information directly stored there, but also to compute "reasonable" derivative information. For example, if a database has the location of two ships, users will expect the system to be able to provide the distance between them--an item of information not directly recorded in the database, but easily computed from the existing data. In general, any system thatis to be widely accepted by users must not only provide access to database information, but must also enhance that primary information by providing procedures that calculate secondary attributes from the data actually stored. Data enhancement procedures are currently provided by LADDER and a few other hand-built systems. But work is needed to devise means for allowing system users to specify their own database enhancement functions and to couple their functions with the natural-language component.

Efforts are now underway (e.g. [26] [13]) to simplify the task of acquiring and coding the knowledge needed to transport high-performance systems from one database to another. It appears likely that soon much of this task can be automated or performed by a database administrator, rather than by a computational linquist. When this is achieved, natural-language access to data is likely to move rapidly into widespread use.

E. New Hardware

VLSI (Very Large Scale Integration of computer circuits on single chips) is revolutionizing the computer industry. Within the last year, new personal computer systems have been announced that, at relatively low cost, will provide throughputs rivaling that of the Digital Equipment KA-10, the time-sharing research machine of choice as recently as seven years ago. Although specifications for the new machines differ, a typical configuration will support a very large (32 bit) virtual address space, which is important for knowledge-intensive natural-language processing, and will provide approximately 20 megabytes of local storage, enough for a reasonable-size database.

Such machines will provide a great deal of personal computing power at costs that are initially not much greater than those for a single user's access to a time-shared system, and that are likely to fall rapidly. Hardware costs reductions will be particularly significant for the many small research groups that do not have enough demand to justify the purchase of a large, time-shared machine.

The new generation of machines will have the virtual address space and the speed needed to overcome many of the technical bottlenecks that have hampered research in the past. For example, researchers may be able to spend less time worrying about how to optimize inner loops or how to split large programs into multiple forks. The effort saved can be devoted to the problems of language research itself.

The new machines will also make it economical to bring co siderable computing to people in all sectors of the economy, including government, the military, small business, and to smaller units within large businesses. Detached from the computer wizards that staff the batch processing center or the time-shared facility, users of the new personal machines will need to be more self reliant. Yet, as the use of personal computers spread, these users are likely to be increasingly less sophisticated about computation. Thus, there will be an increasing demand to make personal computers easier to use. As the price of computation drops (and the price of human labor continues to soar), the use of sophisticated means for interacting intelligently with a broad class of computer users will become more and more attractive and demands for natural-language interfaces are likely to mushroom.

F. Future Directions for Basic Research

1. The Research Base

Work on computational linguistics appears to be focusing on a rather different set of issues than those that received attention a few years ago. In particular, mechanisms for dealing with syntax and the literal propositional content of sentences have become fairly well understood, so that now there is increasing interest in the study of language as a component in a broader system of goal-motivated behavior. Within this framework, dialogue participation is not studied as a detached linguistic phenomenon, but as an activity of the total intellect, requiring close coordination between language-specific and general cognitive processing.

Several characteristics of the communicative use of language pose significant problems. Utterances are typically spare, omitting information easily inferred by the hearer from shared knowledge about the domain of discourse. Speakers depend on their hearers to use such knowledge together with the context of the preceding discourse to make partially specified ideas precise. In addition, the literal content of an utterance must be interpreted within the context of the beliefs, goals, and plans of the dialogue participants, so that a hearer can move beyond literal content to the intentions that lie behind the utterance. Furthermore, it is not sufficient to consider an utterance as being addressed to a single purpose; typically it serves multiple purposes: it highlights certain objects and relationships, conveys an attitude toward them, and provides links to previous utterances in addition to communicating some propositional content.

An examination of the current state of the art in natural-language processing systems reveals several deficiencies in the combination and coordination of language-specific and general-purpose reasoning capabilities. Although there are some systems that coordinate different kinds of language-specific capabilities [3] [12] [20] [16] [30] [17], and some that reason about limited action scenarios [21] [15] [19] [25] to arrive at an interpretation of what has been said, and others that attempt to account for some of the ways in which context affects meaning [7] [10] [18] [14], one or more of the following crucial limitations is evident in every natural-language processing system constructed to date:

- Interpretation is literal (only propositional content is determined).
- The user's knowledge and beliefs are assumed to be identical with the system's.
- The user's plans and goals (especially as distinct from those of the system) are ignored.

Initial progress has been made in overcoming some of these limitations. Wilensky [28] has investigated the use of goals and plans in a computer system that interprets stories (see also [22] [4]). Allen and Perrault [1] and Cohen [6] have examined the interaction between beliefs and plans in task-oriented dialogues and have implemented a system that uses information about what its "hearer" knows in order to plan and to recognize a limited set of speech acts (Searle [23] [24]). These efforts have demonstrated the viability of incorporating planning capabilities in a natural-language processing system, but more robust reasoning and planning capabilities are needed to approach the smooth integration of language-specific and general reasoning capabilities required for fluent communication in natural language.

2. Some Predictions

Basic research provides a leading indicator with which to predict new directions in applied science and engineering; but I know of no leading indicator for basic research itself. About the best we can do is to consider the current state of the art, seek to identify central problems, and predict that those problems will be the ones receiving the most attention.

The view of language use as an activity of the total intellect makes it clear that advances in computational linguistics will be closely tied to advances in research on general-purpose common-sense reasoning. Hobbs [11], for example, has argued that 10 seemingly different and fundamental problems of computational linguistics may all be reduced to problems of common-sense deduction, and Cohen's work clearly ties language to planning.

The problems of planning and reasoning are, of course, central problems for the whole of AI. But computational linguistics brings to these problems its own special requirements, such as the need to consider the beliefs, goals, and possible actions of multiple agents, and the need to precipitate the achievement of multiple goals through the performance of actions with multiple-faceted primary effects. There are similar needs in other applications, but nowhere do they arise more naturally than in human language.

In addition to a growing emphasis on general-purpose reasoning capabilities, I believe that the next few years will see an increased interest in natural-language generation, language acquisition, information-science applications, multimedia communication, and speech.

Generation: In comparison with interpretation, generation has received relatively little attention as a subject of study. One explanation is that computer systems have more control over output than input, and therefore have been able to rely on canned phrases for output. Whatever the reason for past neglect, it is clear that generation deserves increased attention. As computer systems acquire more complex knowledge bases, they will require better means of communicating their knowledge. More importantly, for a system to carry on a reasonable dialogue with a user, it must not only interpret inputs but also respond appropriately in context, generating responses that are custom tailored to the (assumed) needs and mental state of the user.

Hopefully, much of the same research that is needed on planning and reasoning to move beyond literal content in interpretation will provide a basis for sophisticated generation.

Acquisition: Another generally neglected area, at least computationally, is that of language acquisition. Berwick [2] has made an interesting start in this area with his work on the acquisition of grammar rules. Equally important is work on acquisition of new vocabulary, either through reasoning by analogy [5] or simply by being told new words [13]. Because language acquisition (particularly vocabulary acquisition) is essential for moving natural-language systems to new domains, I believe considerable resources are likely to be devoted to this problem and that therefore rapid progress will ensue.

Information Science: One of the greatest resources of our society is the wealth of knowledge recorded in natural-language texts; but there are major obstacles to placing relevant texts in the hands of those who need them. Even when texts are made available in machine-readable form, documents relevant to the solution of particular problems are notoriously difficult to locate. Although computational linguistics has no ready solution to the problems of information science, I believe that it is the only real source of hope, and that the future is likely to bring increased cooperation between workers in the two fields.

Multimedia Communication: The use of natural language is, of course, only one of several means of communication available to humans. In viewing language use from a broader framework of goal-directed activity, the use of other media and their possible interactions with language, with one another, and with general-purpose problem-solving facilities becomes increasingly important as a subject of study.

Many of the most central problems of computational linguistics come up in the use of any medium of communication. For example, one can easily imagine something like speech acts being performed through the use of pictures and gestures rather than through utterances in language. In fact, these types of communicative acts are what people use to communicate when they share no verbal language in common.

As computer systems with high-quality graphics displays, voice synthesizers, and other types of output devices come into widespread use, an interesting practical problem will be that of deciding what medium or mixture of media is most appropriate for presenting information to users under a given set of circumstances. I believe we can look forward to rapid progress on the use of multimedia communication, especially in mixtures of text and graphics (e.g., as in the use of a natural-language text to help explain a graphics display).

Spoken Input: In the long term, the greatest promise for a broad range of practical applications lies in accessing computers through (continuous) spoken language, rather than through typed input. Given its tremendous economic importance, I believe a major new attack on this problem is likely to be mounted before the end of the decade, but I would be uncomfortable predicting its outcome.

Although continuous speech input may be some years away, excellent possibilities currently exist for the creation of systems that combine discrete word recognition with practical natural-language processing. Such systems are well worth pursuing as an important interim step toward providing machines with fully natural communications abilities.

G. Problems of Technology Transfer

The expected progress in basic research over the next few years will, of course, eventually have considerable impact on the development of practical systems. Even in the near term, basic research is certain to produce many spinoffs that, in simplified form, will provide practical benefits for applied systems. But the problems of transferring scientific progress from the laboratory to the marketplace must not be underestimated. In particular, techniques that work well on carefully selected laboratory problems are often difficult to use on a large-scale basis. (Perhaps this is because of the standard scientific practice of selecting as a subject for experimentation the simplest problem exhibiting the phenomena of interest.)

As an example of this difficulty, consider knowledge representation. Currently, conventional database management systems (DBMSs) are the only systems in widespread use for storing symbolic information. The AI community, of course, has a number of methods for maintaining more sophisticated knowledge bases of, say, formulas in first-order logic. But their complexity and requirements for great amounts of computer resources (both memory and time) have prevented any such systems from becoming a commercially viable alternative to standard DBMSs.

I believe that systems that maintain models of the ongoing dialogue and the changing physical context (as in, for example, Grosz [7] and Robinson [19]) or that reason about the mental states of users will eventually become important in practical applications. But the computational requirements for such systems are so much greater than those of current applied systems that they will have little commercial viability for some time.

Fortunately, the linguistic coverage of several current systems appears to be adequate for many practical purposes, so commercialization need not wait for more advanced techniques to be transferred. On the other hand, applied systems currently are only barely up to their tasks, and therefore there is a need for an ongoing examination of basic research results to find ways of repackaging advanced techniques in cost-effective forms.

In general, the basic science and the application of computational linguistics should be pursued in parallel, with each aiding the other. Engineering can aid the science by anchoring it to actual needs and by pointing out new problems. Basic science can provide engineering with techniques that provide new opportunities for practical application.

REFERENCES

1. Allen, J. & C. Perrault. 1978. Participating in Dialogues: Understanding via plan deduction. Proceedings, Second National Conference, Canadian Society for Computational Studies of Intelligence, Toronto, Canada.

2. Berwick, R. C., 1980. Computational Analogues of Constraints on Grammars: A Model of Syntactic Acquisition. The 18th Annual Meeting of the Association for Computational Linguistics, Philadelphia, Pennsylvania, June 1980.

3. Bobrow, D. G., et al. 1977. GUS, A Frame Driven Dialog System. Artificial Intelligence, 8, 155-173.

4. Carbonell, J. G. 1978. Computer Models of Social and Political Reasoning. Ph.D. Thesis, Yale University, New Haven, Connecticut.

5. Carbonell, J. G. 1980. Metaphor--A Key to Extensible Semantic Analysis. The 18th Annual Meeting of the Association for Computational Linguistics, Philadelphia, Pennsylvania, June 1980.

6. Cohen, P. 1978. On knowing what to say: planning speech acts. Technical Report No. 118, Department of Computer Science, University of Toronto. January 1978.

7. Grosz, B. J., 1978. Focusing in Dialog. Proceedings of TINLAP-2, Urbana, Illinois, 24-26 July, 1978.

8. L. R. Harris, 1977. User Oriented Data Base Query with the ROBOT Natural Language Query System. Proc. Third International Conference on Very Large Data Bases, Tokyo (October 1977).

9. G. G. Hendrix, E. D. Sacerdoti, D. Sagalowicz, and J. Slocum, 1978. Developing a Natural Language Interface to Complex Data. ACM Transactions on Database Systems, Vol. 3, No. 2 (June 1978).

10. Hobbs, J. 1979. Coherence and coreference. Cognitive Science. Vol. 3, No. 1, 67-90.

11. Hobbs, J. 1980. Selective inferencing. Third National Conference of Canadian Society for Computational Studies of Intelligence. Victoria, British Columbia. May 1980.

12. Landsbergen, S. P. J., 1976. Syntax and Formal Semantics of English in PHLIQA1. In Coling 76, Preprints of the 6th International Conference on Computational Linguistics, Ottawa, Ontario, Canada, 28 June - 2 July 1976. No. 21.

13. Lewis, w. H., and Hendrix, G. G., 1979. Machine Intelligence: Research and Applications -- First Semiannual Report. SRI International, Menlo Park, California, October 8, 1979.

14. Mann, W., J. Moore, & J. Levin 1977. A comprehension model for human dialogue. Proceedings, International Joint Conference on Artificial Intelligence, 77-87, Cambridge, Mass. August 1977.

15. Novak, G. 1977. Representations of knowledge in a program for solving physics problems. Proceedings, International Joint Conference on Artificial Intelligence, 286-291, Cambridge, Mass. August 1977.

16. Petrick, S. R. 1978. Automatic Syntactic and Semantic Analysis. In Proceedings of the Interdsciplainary Conference on Automated Text Processing (Bielefeld, German Federal Republic, 8-12 November 1976). Edited by J. Petofi and S. Allen. Reidel, Dordrecht, Holland.

17. Reddy, D. R., et al. 1977. Speech Understanding Systems: A Summary of Results of the Five-Year Research Effort. Department of Computer Science. Carnegie-Mellon University, Pittsburgh, Pennsylvania, August, 1977.

18. Rieger, C. 1975. Conceptual Overlays: A Mechanism for the Interpretation of Sentence Meaning in Context. Technical Report TR-354. Computer Science Department, University of Maryland, College Park, Maryland. February 1975.

19. Robinson, Ann E. The Interpretation of Verb Phrases in Dialogues. Technical Note 206, Artificial Intelligence Center, SRI International, Menlo Park, Ca., January 1980.

20. Sager, N. and R. Grishman. 1975. The Restriction Language for Computer Grammars. Communications of the ACM, 1975, 18, 390-400.

21. Schank, R. C., and Yale A.I. 1975. SAM--A Story Understander. Yale University, Department of Computer Science Research Report.

22. Schank, R. and R. Abelson. 1977. Scripts, plans, goals, and understanding. Hillsdale N.J.: Laurence Erlbaum Associates.

23. Searle, J. 1969. Speech acts: An essay in the philosophy of language. Cambridge, England: Cambridge University Press.

24. Searle, J 1975. Indirect speech acts. In P. Cole and J. Morgan (Eds.), Syntax and semantics, Vol. 3, 59-82. New York: Academic Press.

25. Sidner, C. L. 1979. A Computational Model of Co-Reference Comprehension in English. Ph.D. Thesis, Massachusetts Institute of Technology, Cambridge, Massachusetts.

26. F. B. Thompson and B. H. Thompson, 1975. Practical Natural Language Processing: The REL System as Prototype. In M. Rubinoff and M. C. Yovits, eds., Advances in Computers 13 (Academic Press, New York, 1975).

27. H. Tennant, "Experience with the Evaluation of Natural Language Question Answerers," &Proc. Sixth International Joint Conference on Artificial Intelligeneå, Tokyo, Japan (August 1979).

28. Wilensky, R. 1978. "Understanding Goal-Based Stories." Yale University, New Haven, Connecticut. Ph.D. Thesis.

29. D. Waltz, "Natural Language Access to a Large Data Base: an Engineering Approach," Proc. 4th Internatioal Joint Conference on Artificial Intelligence, Tbilisi, USSR, pp. 868-872 (September 1975).

30. Woods, W. A., et al. 1976. Speech Understanding Systems: Final Report. BBN Report No. 3438, Bolt Beranek and Newman, Cambridge, Massachusetts.