Natural Language Processing: What's Really Involved?

Roger Schank and Alex Kass

Yale University Department of Computer Science
Box 2158 Yale Station, New Haven, CT 06520

Introduction

The question before us is whether or not NLP has gone anywhere since the last TINLAP. The answer depends strongly on what we take NLP to be. If we accept the common assumption that language processing can be isolated from the rest of cognition, and that the goal of NLP research to study language as a formal system independent of such (admittedly difficult) issues as memory, learning, and explanation, then our conclusion is grim. Not only has this approach failed to make significant progress in the eight years since TINLAP2, it is permanently doomed to failure because it misses the point of what language is all about. It is only possible for research on understanding natural language to make progress when the researchers realize that the heart of NLU is the understanding process, not language *per se.*

Language is so fascinating because it can be a vehicle for communicating such fascinating content. The proper focus of NLP research is on the content, and how it is extracted from language. As we all know by now, this is a difficult problem because much of what is communicated by natural language is not explicitly stated. Users of natural languages rely on the fact that they can assume a tremendous amount of shared knowledge to help resolve ambiguities, determine anaphoric reference, fill in ellipsis, etc. The two fundamental problems we must solve in order to get computers to approach human levels of language understanding are, first, to endow them with the kind of knowledge bases that humans have and second, to program them to use that knowledge to guide the understanding process.

NLP researchers must, therefore, address such questions as how we understand and represent the concepts that language can communicate, how we learn new concepts, and how we organize this knowledge in memory so that it will be available

to guide the understanding process in the future. A science of NLP built around questions such as these *can* make progress, and in fact, has been making progress in our lab and elsewhere.

Our Evolving View of the Understanding Process

We have been working on programs that understand natural language text for many years, now in our laboratory at Yale, and at the Stanford AI lab before that. But within that context, our focus has shifted considerably because our conception of what it means to understand has changed drastically. It is this evolution of our notion of what constitutes understanding that represents the real progress we have made.

We started out working on Conceptual Dependency (CD) theory [Schank 75], which was a theory of language-free representation of the content of natural language sentences. This led to the development of programs that could map from language to CD (parsers) [Riesbeck 75] and back to language again, (generators) [Goldman 75]. Understanding meant getting the mapping to CD right, as demonstrated by the ability of the generator to produce a correct paraphrase of the original input.

Of course, much of what someone who is paraphrasing or translating must understand from the input is inferred rather than stated, so we were motivated to develop a theory of inference. Our first theory of inference was quite simple. We attached various types of inferences to the types of CD forms. Each time the parser produced a CD form, the associated inferences would be fired off, producing more CD forms which would fire off more inferences, etc. (See [Rieger 75]).

So now our theory of understanding included a theory of inference, which was, in a sense, a theory of context, but the context didn't really drive the understanding process. While reading a sentence in a text, our programs (unlike people) did not develop expectations about what the following sentences would say. The lack of top-down guidance made the inference process to unconstrained; irrelevant inferences overwhelmed the relevant ones. In response ot this problem we developed a theory of scripts and plans as memory structures that would provide top-down expectations to the parser. Our view of understanding now broadened: rather than simply mapping sentences to an internal representation we were now interested in finding a larger memory structure into which the representation would fit. Script-based understanders developed in our lab included SAM [Cullingford 78] and FRUMP [DeJong 77].

The success of the script-based approach led us to focus on memory issues. Two inter-related issues not adequately addressed by script-applier projects such as SAM

and FRUMP were how scripts are organized in memory and how they are learned dynamically. Since a person cannot truly be said to understand if he doesn't learn from his experiences, it seems odd that so many NLP researchers seem comfortable with the notion that language understanding should be studied independent of issues of memory modification. We began closing this gap by proposing an organizational scheme (called Memory Organization Packets, or MOPs [Schank 81]) by which memory structures could be interconnected. We began to view understanding as a process of becoming reminded of the appropriate prior episodes and building new generalizations that could account for both the old and the new. The CYRUS [Kolodner 80] and IPP [Lebowitz 80] projects represented the first cuts at implementing this theory.

At this point it should be clear that the trend in our research has been steadily away from viewing language understanding as an encapsulated, data-driven process of mapping input strings to internal representations. The more we worked with language-understanding systems in our lab, the clearer it became how active and memory-driven the language understanding process was. The understander's idiosyncratic set of prior experiences, and the way these are organized, have a powerful effect on the meaning of any input to that understander: Different people understand the same thing differently.

This insight led us naturally to our current work, which is on question-asking and explanation. Stories that don't relate to any previously understood experience are cryptic, while stories that contain only things that we expect are boring and teach us nothing. The truly interesting stories — the ones we find ourselves wondering over — are the ones which are close enough to things we understand to be able to interact with our prior experiences, but which cannot quite be explained by any of the ready-made explanations sitting in our memory. The problem with programs that relied completely on applying pre-established knowledge structures is that they ignored anomalies in the stories instead of realizing that the anomalies are the most interesting part.

Creative thinking is required in order to develop explanations that can be used to understand novel stories. But we believe that this kind of creativity can be accomplished by a fairly simple, mechanical process. Understanding these stories involves retrieving explanations that have worked in the past from memory, revising those explanations to make them applicable to the current story, and finally storing the modifications back in memory so that the understander will have learned a new explanation by the time the story has been understood. The creativity stems from finding an interesting, relevant explanation in memory (creative indexing) and knowing how to modifiy the explanation to make it fit the problem at hand (creative tweaking).

For example, suppose you read that Swale, a star 3-year-old racehorse, was found dead in his stall a few days after a big race. If you are not an avid racing fan you are unlikely to have ready-made explanations for racehorse deaths. However, if you think

creatively about the problem you may be reminded of any number of relevant cases you know about and you can probably modify the associated explanations to make them fit Swale. For example, if you think of Swale as an athlete who died despite being in good condition you might be reminded of the jogger, Jim Fixx, who had a hidden heart defect which combined with the stress of jogging to kill him. Perhaps Swale had a hidden defect, and perhaps running in races was for Swale what jogging was for Fixx. If you think of Swale as a star performer who died before his time you may be reminded of Janis Joplin or John Belushi. While Swale probably wasn't taking recreational drugs, this explanation can lead to drug-related explanations that are more plausible. Perhaps Swale was poisoned by a competitor, or maybe his owner was giving him drugs to make him run faster and accidentally gave him too much. The two tricks to generating hypotheses in this way are, first, to characterise the problem in a way that makes the old explanations come to mind, and second, to tweak the old explanations into something that fits the new situation.

A full presentation of our current work is clearly impossible in this paper (see [Schank 86] for a theoretical discussion and [Kass 86] and [Leake and Owens 86] for brief discussions of a program built around these principles); the goal here is simply to point out how our interest in natural language processing has led us naturally, and indeed inevitably, to develop theories of explanation and creativity. Some may say that we have strayed from the core issues of NLP, but our point is that these *are* the core issues. The drive to explain what might be going on in the story guides the entire understanding process.

Conclusion

So, our answer to the original question is that the study of NLP *in terms of the overall understanding process* is making very good progress. It's not so much that we have developed solutions to all our problems, although we certainly have developed some solutions. The point is that we are starting to understand what the problem is, and this is much more important.

Acknowledgements

Thanks to Larry Birnbaum, David Leake, Chris Owens, and Ashwin Ram for their helpful comments. Our work is supported in part by the Air Force Office of Scientific Research under grant 85-0343.

References

[Cullingford 78] Cullingford, R., *Script Application: Computer Understanding of Newspaper Stories,* Ph.D. Thesis, Yale University, 1978. Research Report #116.

[DeJong 77] DeJong, G.F., *Skimming newspaper stories by computer,* Technical Report 104, Yale University Department of Computer Science, 1977.

[Goldman 75] Goldman, N., Conceptual Generation, *Conceptual Information Processing,* North-Holland, Amsterdam, 1975.

[Kass 86] Kass, A., Modifying Explanations to Understand Stories, *Proceedings of the Eighth Annual Conference of the Cognitive Science Society,* Cognitive Science Society, Amherst, MA, August 1986.

[Kolodner 80] Kolodner, J.L., *Retrieval and Organizational Strategies in Conceptual Memory: AComputer Model,* Ph.D. Thesis, Yale University, November 1980.

[Leake and Owens 86] Leake, D. B. and Owens, C. C., Oragnizing Memory for Explanation, *Proceedings of the Eighth Annual Conference of the Cognitive Science Society,* Cognitive Science Society, Lawrence Erlbaum Associates, 1986.

[Lebowitz 80] Lebowitz, M., *Generalization and Memory in an Integrated Understanding System,* Ph.D. Thesis, Yale University, October 1980.

[Rieger 75] Rieger, C., Conceptual Memory and Inference, *Conceptual Information Processing,* North-Holland, Amsterdam, 1975.

[Riesbeck 75] Riesbeck, C., Conceptual Analysis, *Conceptual Information Processing,* North-Holland, Amsterdam, 1975.

[Schank 75] Schank, R.C., *Fundamental Studies in Computer Science,* Volume 3: *Conceptual Information Processing,* North-Holland, Amsterdam, 1975.

[Schank 81] Schank, R.C., *Reading and Understanding: Teaching from the Perspective of Artificial Intelligence,* Lawrence Erlbaum Associates, Hillsdale, NJ, 1981.

[Schank 86] Schank, R.C., *Explanation Patterns: Understanding Mechanically and Creatively,* 1986. Book in press.