WHAT LANGUAGE UNDERSTANDING RESEARCH SUGGESTS ABOUT DISTRITBUTED ARTIFICIAL INTELLIGENCE [1] [2]

Earl D. Sacerdoti
Artificial Intelligence Center
SRI International

Natural language communication requires that dialogue participants know not only the structure of language but also something about the subject matter under discussion and about the processes through which concepts in the subject area may be expressed by means of words, phrases, and sentences in the language. Similarly, a rich dialogue among distributed intelligent processors will require that the processors have a (partial) knowledge base in common and that they employ and understand common ground rules about how to convey portions of that knowledge base within the available communication protocols. Therefore, designers of distributed artificial intelligence (DAI) systems might find useful concepts in the designs of artificial intelligence systems for natural language understanding.

Computer science, as opposed to traditional mathematics, views computation as a process performed with finite resources over time. In an analogous manner, natural language understanding research in artificial intelligence, as opposed to traditional linguistics, is increasingly concerned with communication as a process performed with finite resources over time. DAI systems of any power will certainly also have to view interprocessor communication as a process performed with finite resources over time. The particular perspective of artificial intelligence on the problem of understanding natural language may help us to articulate some issues of importance that must be faced in developing distributed systems that display intelligent behavior.

---

[1] The preparation of this paper was supported by the Advanced Research Projects Agency of the Department of Defense. It was prepared in haste, for which I apologize to the reader. I wish to thank Barbara Grosz and Peter Hart for their helpful suggestions. Please don't expect them to stand firmly behind the papers' conclusions, though. For that matter, don't expect me to!

[2] Originally published in *Proc. Distributed Sensor Nets Workshop*, Carnegie-Mellon University, Pittsburgh, Pennsylvania (1978)

This brief note will attempt to draw an extended analogy between natural language communication and interprocess communication, thereby hopefully identifying in advance some of the gaps that must be filled as the processors that are communicating become increasingly powerful.

A. Introduction

Communications specialists have generally concerned themselves with the format and content of individual classes of interprocess messages, rather than on the ongoing interaction (spanning many instances or many classes of message) through which particular information is communicated. This is analogous to the linguists' traditional concerns with the form and meaning of individual words and sentences. In contrast, AI researchers on language understanding view communication as an <u>activity</u> performed by two or more cooperating parties. This viewpoint on language understanding systems suggests that a crucial issue for DAI is a careful articulation of the processes that underlie participation in dialogue (as opposed to the structuring of the individual interprocess messages and communication protocols).

Below we will characterize three kinds of knowledge that language understanding research suggests must be available to a communicating process. Then we will discuss the activities involved in interpreting and responding to a message.

B. Knowledge about the Subject Domain

Communication between people or processors can be viewed as the incremental sharing and building of their respective knowledge bases. Recent work in computational linguistics has demonstrated the importance to each participant in a dialogue of having a very rich knowledge base that is much more than a static description of the facts that are currently true in the domain of discourse. The knowledge base must encode a dynamic environment consisting of various actors, objects, relationships and events, ordered (or partially ordered) with respect to time. It must describe not only the environment as it truly is, but also the knowledge and beliefs about the environment that are held by each participant in the dialogue.

C. Knowledge about Context

Natural language is used for communication in a dynamically changing context. All utterance in a dialogue cannot typically be interpreted in isolation; it must be analyzed within the context in which it was produced. An interpretation is influenced by the current state of the environment, by a history of the previous states, by the overall structure and content of the dialogue, by knowing who produced the utterance and for whom it was intended, and by the preceding utterances in the dialogue.

This complex collection of required state information renders genuine comprehension of natural dialogue beyond the current state of the art. It is, however, an extremely efficient means of communicating parsimoniously over a noisy medium when there is sufficient processing power available on both sides of the communications link. By each having all this knowledge about the subject domain and the current context of the interaction, processors can communicate using many fewer bits. Furthermore, because the processors are continually engaged in augmenting and checking a shared knowledge base, errorful transmissions are much more likely to be noticed, and a subdialogue requesting confirmation of the suspicious message can be initiated.

Employing this state information, then, may provide significant additional efficiency and reliability in the communications process. On the other hand, maintaining and exploiting this state information imposes a very significant additional computation load. Researchers exploring the distribution of artificial intelligence capabilities will need to evaluate the tradeoffs between the increased processing required and the enriched communication provided by this approach.

D. Knowledge about Communicating

To enable the rich interactions we have been describing, a third kind of knowledge is needed. This is knowledge about the "rules of the game" of communicating. In typical systems that perform interprocess communication this consists of no more than the encoding of routines that can either create or interpret instances of particular classes of messages. By analogy with natural language understanding systems, communicating processes may also have to know enough about the activity of communication itself to determine when the focus of the communication is shifting, what the current goals of the sender are that lie behind his current transmissions, and when (and why) the current sequence of transmissions is coherent as a whole. The cues for these kinds of information are often encoded in subtle ways in natural language dialogues. Human participants seem able to perform the deductions required to pick up these cues easily. Language understanding systems have been rather poor at this to date. By encoding the cues more explicitly for interprocess communication, the requirement for the individual processors to perform complex deductions can be greatly reduced. This aspect of communication, then, appears to be one that can be incorporated rather easily into DAI systems.

E. Interpreting a Message

In analyzing and interpreting messages expressed in natural language, a variety of kinds of information must be brought to bear. Firstly, there is knowledge about the syntax, how individual words are combined into phrases and how phrases are combined into sentences. Secondly, there is lexical knowledge, about the meanings of individual words and the roles they can take on within larger phrases. Thirdly, there is knowledge abut the mapping between the phrases of the input and descriptions of objects in the internal representation of the subject domain. (Strictly speaking, this is what philosophers mean when they refer to semantics although the term is almost always used in a much wider sense in the literature of artificial intelligence.) Finally, knowledge about the current state is needed to refine the descriptions of objects into designations of particular entities in the "real world."

While these types of knowledge are listed in order of increasing complexity and difficulty of use, they are not employed in a strictly linear order. Since the analysis at each level affects the confidence in the conclusions drawn at other levels, the overall interpretation is usually built up incrementally with many partial contributions from all levels.

The designer of a system that involves interprocess communication typically builds in to the communicating processes knowledge of the first two sorts described above. He worries about the structure of each message and about the values and meanings of the fields within those messages. The other types of knowledge are “hard wired" into his programs, and are typically extremely simple. Each message typically has an unambiguous meaning independent of its ordering within the overall dialogue.

This design of a communication protocol is appropriate for situations where the processing cost of sending or receiving each message must be kept low. The use of the semantically oriented kinds of knowledge makes the processing for message transmission very much more expensive. However, if processing at the source and destination is relatively cheap (and we expect this to be the case for DAI systems), the overall cost of the system might be minimized by trading off higher processing requirements for lower bandwidths and higher noise levels.

F. Responding to a Message

We have just sketched the (rather complex) process by which a message might be interpreted by a processor in a DAI system. Once the message is interpreted, it must then be responded to. As might be expected, this is also a rather complex process for a natural language understanding system, and will probably be complex for a DAI system as well.

The complexity of the response derives mostly from the need to update and maintain the complex knowledge about state information described in Section C above. The response must include:

- a check for the validity of the current assumptions about the state of the other communicating process;
- a check for cues that the sender of the message is shifting the focus of the communication;
- a determination of the overt actions to be taken in response to the message;
- a determination of whether (and under what conditions and when) a return message is requested or required;
- the generation, if needed, of a return message (which will involve encoding sufficient information for the other process to perform the same set of tasks)

As was the case with the task of interpreting messages, responding to messages will be easier for computer-to-computer interactions than for natural language communication. By requiring explicit indications of shifts of focus and, perhaps, an explicit indication of nonstandard or unexpected assumptions, the response to messages can probably be performed with tolerable efficiency.

G. Conclusions

We have speculated on the possible relevance of the AI approach to natural language understanding to the problem of communication between processors in a distributed artificial intelligence system. While the processing required to interpret a natural language is almost certainly far more than a DAI system needs, a simplified form of this processing that performs all or most of the tasks performed in understanding natural language may well be required. This suggests that designers of DAI systems might best design their communications protocols by narrowing down the capabilities required for natural language communication rather than by building up from the traditional, individual-message-oriented approach. The key point I have tried to make is that, to develop a robust distributed system that is able to communicate over a noisy, relatively narrow-bandwidth channel, one must consider not just the format and content of the individual messages that are used to communicate, but also the role that sequences of messages play in the overall process of communication.

H. Relevant References

1. Cohen, P.R., "On Knowing What to Say: Planning Speech Acts," Ph.D. Thesis, University of Toronto, 1978.

2. Grosz, B.J, "Focusing in Dialog," Proc. *TINLAP-2*, Urbana, Ill., 1978.

3. Lesser, V.R. and Erman, L.D., "A Retrospective View of the HEARSAY-II Architecture," *Proc. 5th IJCAI*, Cambridge, Mass., 1977.

4. Levy, D.M., "Communicative Goals and Strategies Between Discourse and Syntax," *Proc. Symposium on Discourse and Syntax*," UCLA, Nov. 1977.

5. Robinson, A.E., "Investigating the Process of Natural Language Communication," SRI AI Center Tech. Note 165, 1978.

I. Viewfoils (*remember viewfoils?*) from the Talk

Figure 1: Two processors each have large, complex knowledge bases. The first processor believes that W1 is a box-end wrench. The second believes only that W1 is a wrench.

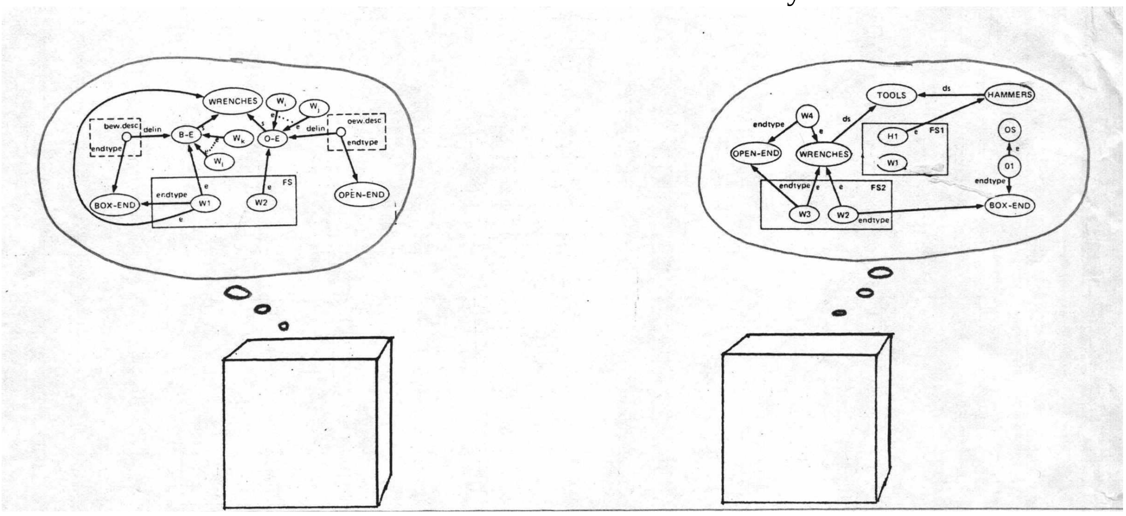

Figure 2: The sending processor need transmit only the information it believes the receiving processor needs in order for its knowledge base to be updated ("The Endtype of wrench W1 is 'box-end.'").

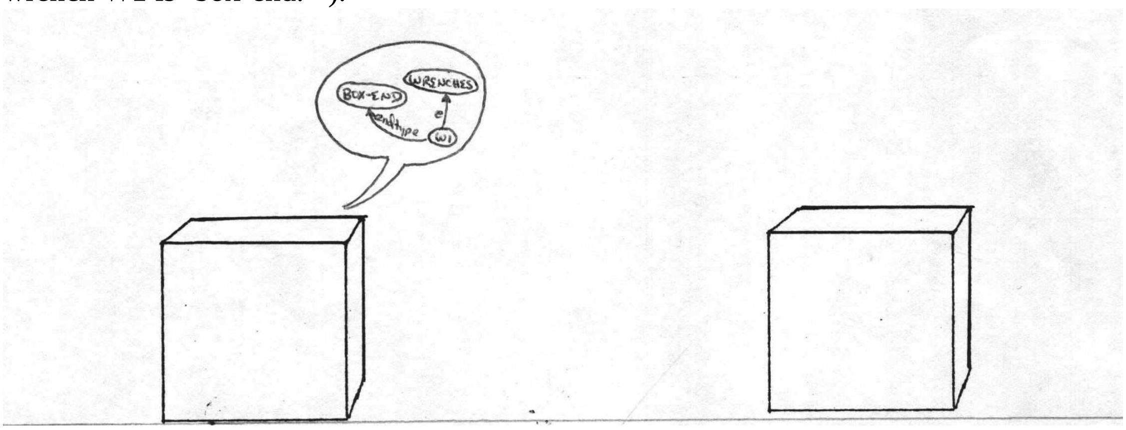

Figure 3: The receiving processor updates its knowledge base.