A HIGH-LEVEL INTERPRETATION ALGORITHM FOR POINTS AND RANGES

Elaine A. Rich

Department of Computer Sciences
University of Texas at Austin
Austin, Texas 78712

TR-84-17 June 1984

A High-Level Interpretation Algorithm for Points and Ranges

Elaine A. Rich
Department of Computer Sciences
The University of Texas at Austin

1 Introduction

It is, by now, commonly accepted that many natural language statements "mean", pragmatically, a good deal more than they "mean" literally. The theory of indirect speech acts [Searle 69, Cohen 79] describes ways in which people use language as one form of behavior in the larger context of their goals and techniques for achieving those goals. Given this, we can describe algorithms for computing the additional meanings of statements by referring to the behavioral context in which they occur. The theory of conversational postulates [Grice 75, Gordon 75] describes conventions that are assumed to be shared by all participants in a conversation. Additional intended meanings of statements in a conversation may be computed by analyzing the ways in which the accepted conventions appear to be violated if the statements are interpreted strictly literally. The common thread here is the reliance of natural language communication on a set of *high-level interpretation algorithms*, which each user of the language knows and assumes each other user also to know. These algorithms are called high-level because they rely on the output of the lower-level lexical and syntactic interpretation processes that must also be shared among the users of a language.

The obvious analogy is to encryption systems. If I send an encryped message to someone whom I know to possess the code I am using, then I assume that my message will mean, to its receiver, exactly what I can compute it to mean by applying the decryption algorithm to it. I will design my message so that that computed meaning is the meaning I want to convey. In many modern encryption systems, one additional piece is necessary. The receiver of the message must possess not just the correct algorithm for decryption but also the correct key or keys to be used by the algorithm. In other words, the decryption algorithm must refer to an external database (perhaps a small one) in order to compute the meaning of a statement. If the decryption program does not use the database that the sender of the message expected it to use, then the sender's message will be misinterpreted.

The motivation for the use of encryption systems is simple -- security. The motivation for the use of the corresponding process in natural language is more varied. Sometimes it is the desire to appear polite. Sometimes it is an inability to anticipate in advance all of the situations that may be going on when the statement is interpreted. And sometimes it is simply the desire to be concise. In this paper, we will look at statements whose structure arises from this last motivation. It is simpler to analyze than the others because it does not involve understanding human motivations. Thus it allows us to get a clearer picture of the processes that occur. We will look at how a high-level interpretation algorithm (which is what we are calling the natural analog of a decryption program), coupled with an appropriate database of world knowledge, can be used to define the complete meaning of English statements whose literal meaning is incomplete. The existence of such an interpretation algorithm allows people to communicate in a more concise fashion than would otherwise be possible.

2 Specifying Points and Ranges

Consider the following statements:

(1) free HBO
(2) Tie and jacket required
(3) No stovetop or broiler
(4) No teenagers allowed
(5) No overnight parking
(6) No stopping or standing
(7) Consume 8 glasses of water and 1 hamburger for lunch
(8) No lawyers will be allowed on the jury
(9) A home for troubled teenagers
(10) Pipe smoking not allowed.
Cigarette smoking at the discretion of the individual managers.

Now consider the following questions that might be able to be answered using the information in the statements above. The relevant statement(s) is/are given in parentheses after each question. Following that, the answer appears.

- Is local TV available? [(1) seen on a motel marquis] Yes.
- Do I have to wear shoes? [(2) on restaurant sign] Yes.
- Can I put it in the oven? [(3) stamped on the bottom of a glass casserole dish] Yes.
- Are children allowed? [(4) on a sign next to a merry-go-round with small horses] Yes.
- Are adults allowed? [(4) by the same merry-go-round] No.
- Are children allowed? [(4) on a sign outside a pool hall] No.
- Are adults allowed? [(4) outside the pool hall] Yes.
- Are people in their 20's allowed or is this just for the old folks? [(4) outside the pool hall] Yes, they are allowed.
- Can I park for an hour? [(5) or (6) posted by desired parking place] Yes with (5); no with (6).
- Can I have 9 glasses of water? [(7) in a description of a new fad diet] Yes.
- Can I have 2 hamburgers? [(7) in the same place] No.
- Are physicians also excluded? [(8) in the instructions sent to prospective jurors] No.
- Will someone be kicked out on their 20th birthday? [from the place described by (9)] No, probably not.

The answers that can be derived to these questions suggest several observations about the meaning of the statements involved:

1. When a particular object is mentioned, it is possible that just that object is being referred to (e.g., in (8)).

2. When a particular object is mentioned, it is also possible that a cluster of things near that object on some scale is being referred to (e.g., in (9), where ages near teenage are included, but with some likelihood that drops off with distance from the specified point).

3. When a particular object is mentioned, it is also possible that an entire range is being referred to and the mentioned object is an endpoint of the range (e.g., (1) - (7)).

4. If a range is specified by a point, then correct interpretation depends on correct selection of the appropriate scale. This is sometimes easy (e.g., (4) refers to the age scale). But sometimes it is difficult. (1) refers to the scale of TV availability as defined by cable companies. (2) refers to a formality of dress scale as defined by our culture. (3) refers to a scale of trauma of heat sources on glass, in which uneven heat such as that of a stovetop is more dangerous than the even heat of an oven.

5. If a range is specified by a point, then correct interpretation depends on deciding which side of the specified point is the desired range. The importance of this can be seen from the meaning of (4), which depends on the context in which it appears. On the merry-go-round, only people younger than teenagers are allowed. At the pool hall, only older people are welcome. (7) contains both an upper bound (on the number of hamburgers) and a lower bound (on the number of glasses of water). Deciding this requires the knowledge that the idea is to consume lots of water so you will feel full, while at the same time consuming only a small amount of food so that you can lose weight.

6. When an object is used to specify a range, it is assumed to specify the strongest possible bound on the range. If the pool hall only wanted retired people, it would not describe its range in terms of teenagers. Likewise, if the glass dish could not stand any heat at all, that stronger statement would have been made.

7. When an object is used, it is assumed to specify a range unless there is no range that makes sense (as in the case of (8)). The fact that range specification is the default gives rise to statements such as (10). If only the first sentence were present, cigarette smoking would be assumed to be allowed since it is below pipe smoking on the allowable-smoking scale. Another sentence is required to block that assumption.

The following high-level interpretation algorithm provides an analysis of a statement ST that appears to assert some property P of some object O:

Does there exist a scale S with the following properties:

- O is on the scale.
- It makes sense to assert P of some objects other than O on the scale and not P of still others.

Yes ⇒ Decide whether O is a boundary in S by asking whether P makes sense of objects on one side of O but not the other.

Yes (Statements 1-7) ⇒ Decide which side of O on S contains the objects for which it makes sense to assert P. Define the meaning of ST to be the assertion of P for all elements of S that are equal to O or are on the side of O just selected, combined with the assertion of not P for all elements of S that are on the other side of O.

No (Statement 9) ⇒ Define the meaning of ST to be the assertion of P for O and for other objects near O on S, where the likelihood of P for a particular object O' declines as the distance between O and O' increases.

No (Statement 8) ⇒ Define the meaning of ST to be simply the assertion of P for O.

This algorithm defines the meaning of statements such as those found in 1-9. Unfortunately, from a computational point of view, it does so only by referring to an external (to the algorithm) measure of the sensibility of a class of assertions. This measure must make use of a large database of world knowledge. But there is no simpler way to analyze statements such as these, as can be seen from the fact that the meaning of (4) cannot be decided without knowing the context in which it occurs.

High-level interpretation algorithms are the recognition counterparts of generative theories such as speech acts and conversational postulates. For example, the conversational postulate, "Be cooperative and give the information that is desired", can be converted into an interpretation algorithm a sketch of which might be, "If the desired information is not present in a response, conclude that the meaning of the response is a relevant piece of information that it would be impolite to assert, or a piece of information that the hearer does not want to hear or a piece of information that the speaker cannot accept, or..." Stating these shared semantic conventions as recognition algorithms, rather than as generative ones, has the practical advantage that they are more easily incorporated into natural language understanding systems. Just as it has proven difficult to exploit generative syntactic theories directly as a basis for understanding systems, so too is it difficult with semantic ones, because information is often lost during generation.

3 Summary

In this paper we have described the role of high-level interpretation algorithms in the understanding of English statements. Such algorithms are the recognition counterparts to generative notions such as conversational postulates. They work only to the extent that both the algorithms and the databases that they reference are shared between the sender and the receiver of a message. Some recognition algorithms, such as the one presented here, are relatively easy to state. But to operate, the algorithms must reference a possibly very large database of world knowledge. It is the construction of this database that is the bottleneck in the automation of the understanding process that we have described.

References

[Cohen 79] Cohen, P. R. & C. R. Perrault.
Elements of a Plan-based Theory of Speech Acts.
Cognitive Science 3, 1979.

[Gordon 75] Gordon, D. & G. Lakoff.
Conversational Postulates.
In P. Cole & J. L. Morgan (editors), *Studies in Syntax, Vol. III*, .
Seminar Press, New York, 1975.

[Grice 75] Grice, H. P.
Logic and Conversation.
In D. Davidson & G. Harman (editors), *The Logic of Grammar*, .
Dickenson, Encino, Calif., 1975.

[Searle 69] Searle, J. R.
Speech Acts.
Cambridge University Press, Cambridge, 1969.