



I L L I N O I S

UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN

-

PRODUCTION NOTE

University of Illinois at
Urbana-Champaign Library
Large-scale Digitization Project, 2007.

370.152
T2261
No. 398

Technical Report No. 398

REFERENCE AND REFERENCE FAILURES

Bradley A. Goodman

BBN Laboratories Inc.
Cambridge, Massachusetts 02238

December 1986

Center for the Study of Reading

TECHNICAL REPORTS

UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN
174 Children's Research Center
51 Gerty Drive
Champaign, Illinois 61820

CENTER FOR THE STUDY OF READING

Technical Report No. 398

REFERENCE AND REFERENCE FAILURES

Bradley A. Goodman

BBN Laboratories Inc.
Cambridge, Massachusetts 02238

December 1986

University of Illinois
at Urbana-Champaign
51 Gerty Drive
Champaign, Illinois 61820

Bolt Beranek and Newman Inc.
10 Moulton Street
Cambridge, Massachusetts 02238

The work upon which this publication was based was performed pursuant to Contract No. 400-81-0030 of the National Institute of Education and by the Advanced Research Projects Agency of the Department of Defense under Contract No. N00014-85-C-0079. Special thanks to Candy Sidner for her insightful comments and suggestions during the course of this work. Marie Macaïsa made helpful comments on this paper. Special thanks also to Phil Cohen, Scott Fertig and Kathy Starr for providing the water pump dialogues and for their invaluable observations on them.

Abstract

The goal of this work is the enrichment of human-machine interactions in a natural language environment. Because it cannot be assumed that a speaker and listener have the same beliefs, contexts, perceptions, backgrounds or goals at each point in a conversation, difficulties can occur when a listener interprets a speaker's utterance. These mistakes can lead to various kinds of misunderstandings between speaker and listener, including reference failures or failure to understand the speaker's intention. We call these misunderstandings miscommunication. Such mistakes can slow and possibly break down communication. Our goal is to recognize and isolate such miscommunications and circumvent them. This paper will highlight a particular class of miscommunication--reference problems--by describing a case study and techniques for avoiding failures of reference and will illustrate a framework less restrictive than earlier ones. This allows a speaker leeway in forming an utterance about a task and in determining the conversational vehicle to deliver it and also promotes a new view for extensional reference.

Reference and Reference Failures¹1. Introduction

Reference in the real world differs greatly from the reference processes modelled in current natural language systems. A speaker in the real world is a rational agent who must make a decision about his description in a limited time, with limited resources, knowledge, and abilities. In particular, the speaker's perceptual and communicative skills may be imperfect, he may be sloppy in his description, or his model of the listener may be erroneous or incomplete. Since the speaker's goal in the reference process is to construct a description that "works" for the listener, the listener, from his viewpoint, must take these imperfections into account when trying to interpret the speaker's utterances. Yet, listeners, too, have imperfect perceptual or communicative skills and can be sloppy. Hence, they must be prepared to deal with their own imperfections when identifying references. In real reference, listeners often recover from initial misunderstandings with or without help from the speaker. Natural language understanding systems must do this, too. Therefore, in the reference process, a system should assume and expect problems.

In order to build robust natural language processing systems that can detect and recover from miscommunication we have investigated how people communicate and how they recover from problems in communication (Goodman, 1984, 1985, 1986). This investigation centered on reference problems, which a listener has in determining whom or what a speaker is talking about. A

collection of protocols of a speaker explaining to a listener how to assemble a toy water pump were studied and the common errors in speakers' descriptions were categorized. This led to developing techniques for avoiding failures of reference that were employed in a natural language understanding program.

The traditional approaches to identifying reference in natural language systems were found to be less flexible than people's real behavior. In particular, listeners often find the correct referent even when the speaker's description does not describe any object in the world. To model a listener's behavior, a new component was added to the traditional reference identification mechanism to resolve difficulties in a speaker's description which uses knowledge about linguistic and physical context in a negotiation process that determines the most likely places for speaker error. To repair the speaker's description knowledge sources are used to guide relaxation techniques that delete or replace portions. The algorithm developed more closely approximates people's behavior than reference algorithms designed in the past.

2. Reference

Communication involves a series of utterances from a speaker to a hearer. The hearer uses these utterances to access his own knowledge and the world around him. Some of these utterances are noun phrases that refer to objects, places, ideas and people that exist in the real world or in some imaginary world. They cannot be considered in isolation. For example, consider the utterance "Give me that thing." It can be uttered in many

different situations and can result in different referents of that thing. Understanding such referring expressions requires the hearer to take into account the speaker's intention, overall goal, the beliefs of the speaker and hearer, the linguistic context, the physical context, and the syntax and semantics of the current utterance. The hearer could misinterpret the speaker's information in any one of these parts of communication. Such misunderstandings constitute miscommunication. This research focuses primarily on effects of the linguistic and the physical context.

To explore such reference problems, the following method was devised and followed. First, protocols of subjects communicating about a task were analyzed. Knowledge that people used to recover from reference miscommunications was then isolated. Algorithms were designed to apply a person's knowledge about linguistic and physical context to determine the most likely places for error in the speaker's utterance. Then, computer programs were written: (1) to represent a spatially complex physical world, (2) to manipulate that structure to reflect the changes caused by the listener's interpretation of the speaker's utterances, (3) to identify referent noun phrases, and, when that fails, (4) to search the physical world for reasonable candidates for the referent. These programs, which differ from other research in artificial intelligence, form one component of a natural language system.

2.1 The Domain and Methodology

The task-oriented paradigm of Grosz (1977) was followed since it is easy to study (through videotapes), it places the world in front (a primarily extensional world), and it limits the discussion while still providing a rich environment for complex descriptions. The task chosen as the target is the assembly of a toy water pump. The water pump is reasonably complex, containing four subassemblies that are built from plastic tubes, nozzles, valves, plungers, and caps that can be screwed or pushed together. A large corpus of dialogues containing instructions from an "expert" to an "apprentice" that explain the assembly of the toy water pump was collected by Cohen (see Cohen, 1981; Cohen, Fertig, & Starr, 1982; and Cohen, 1984). Both participants were working to achieve a common goal--the successful assembly of the pump. This domain is rich in perceptual information, allowing for complex descriptions of its elements. The data provide examples of imprecision, confusion, and ambiguity, as well as attempts to correct these problems.

The following exchange exemplifies one such situation. In it, E is instructing A to assemble part of the water pump. Refer to Figure 2-1(a) for a picture of the pump. E and A are communicating verbally but neither can see the other. (The bracketed text in the excerpt tells what was actually occurring while each utterance was spoken.) Notice the complexity of the speaker's descriptions and the resultant processing required by the listener. This dialogue illustrates that (1) listeners repair the speaker's description in order to find a referent,

(2) they repair their initial reference choice once they are given more information, and (3) they can fail to choose a proper referent. In Line 7, E describes the two holes on the BASEVALVE as "the little hole." A must repair the description, realizing that E doesn't really mean "one" hole, but is referring to the "two" holes. A apparently does this since he doesn't complain about E's description and correctly attaches the BASEVALVE to the TUBEBASE. Figure 2-1(b) shows the configuration of the pump after the TUBEBASE is attached to the MAINTUBE in Line 10. In Line 13, A interprets "a red plastic piece" to refer to the NOZZLE. When E adds the relative clause "that has four gizmos on it," A is forced to drop the NOZZLE as the referent and to select the SLIDEVALVE. In Lines 17 and 18, E's description "the other-- the open part of the main tube, the lower valve" is ambiguous, and A selects the wrong site, namely the TUBEBASE, in which to insert the SLIDEVALVE. Since the SLIDEVALVE fits, A doesn't detect any trouble. Lines 20 and 21 keep A from thinking that something is wrong because the part fits loosely. In Lines 27 and 28, A indicates that E has not given him enough information to perform the requested action. In Line 30, A further compounds the error in Line 18 by putting the SPOUT on the TUBEBASE.

Excerpt 1 (Telephone)

- E: 1. Now there's a blue cap
[A grabs the TUBEBASE]
2. that has two little teeth sticking
3. out of the bottom of it.
- A: 4. Yeah.
- E: 5. Okay. On that take the
6. bright shocking pink piece of plastic
[A takes BASEVALVE]

7. and stick the little hole over the teeth.

[A starts to install the BASEVALVE,
backs off, looks at it
again and then goes ahead
and installs it]

A: 8. Okay.

E: 9. Now screw that blue cap onto
10. the bottom of the main tube.

[A screws TUBEBASE onto MAINTUBE]

A: 11. Okay.

E: 12. Now, there's a--

13. a red plastic piece

[A starts for NOZZLE]

14. that has four gizmos on it.

[A switches to SLIDEVALVE]

A: 15. Yes.

E: 16. Okay. Put the ungizmoed end in the uh

17. the other--the open

18. part of the main tube, the lower valve.

[A puts SLIDEVALVE into hole in
TUBEBASE, but E meant
OUTLET2 of MAINTUBE]

A: 19. All right.

E: 20. It just fits loosely. It doesn't

21. have to fit right. Okay, then take

22. the clear plastic elbow joint.

[A takes SPOUT]

A: 23. All right.

E: 24. And put it over the bottom opening, too.

[A tries installing SPOUT on
TUBEBASE]

A: 25. Okay.

E: 26. Okay. Now, take the--

A: 27. Which end am I supposed to put it over?

28. Do you know?

E: 29. Put the--put the--the big end--

30. the big end over it.

[A pushes big end of SPOUT on
TUBEBASE, twisting it to
force it on]

The example illustrates the complexity of reference identification in a task-oriented domain. It shows that people do not always give up when a speaker's description isn't perfect (or isn't readily assimilable for them) but that they try to plow ahead anyway.

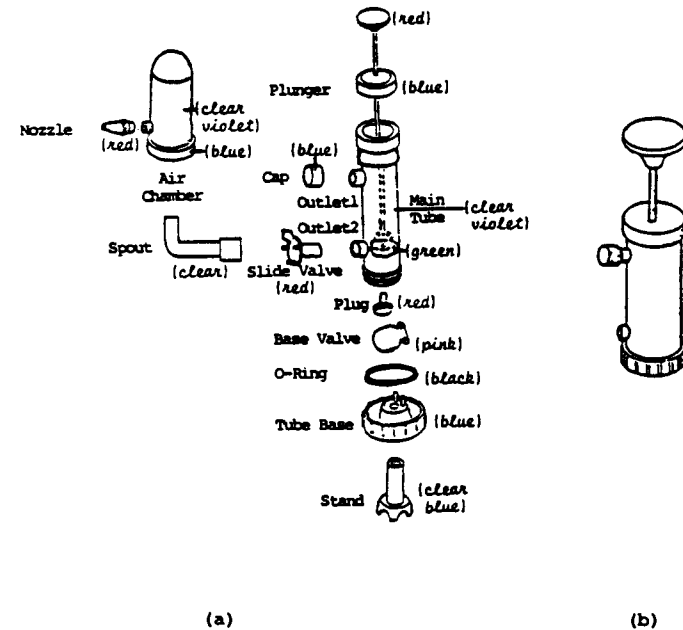


Figure 2-1: The Toy Water Pump

2.2 A New Reference Paradigm from a Computational Viewpoint

Reference identification is a search process where a listener looks for something in the world that satisfies a speaker's uttered description. A computational scheme for performing such reference identifications has evolved from work by other artificial intelligence researchers (Grosz, 1977). That traditional approach succeeds if a referent is found, or fails if no referent is found (see Figure 2-2(a)). However, a reference identification component must be more versatile than those previously constructed. The excerpts provided (Goodman, 1984) show that the traditional approach is inadequate because people's real behavior is much more elaborate. In particular, listeners often find the correct referent even when the speaker's description does not describe any object in the world. For example, a speaker could describe a turquoise block as the "blue block." Most listeners would go ahead and assume that the turquoise block was the one the speaker meant, since turquoise and blue are similar colors.

A key feature to reference identification is "negotiation." Negotiation in reference identification comes in two forms. First, it can occur between the listener and the speaker. The listener can step back, expand greatly on the speaker's description of a plausible referent, and ask for confirmation that he has indeed found the correct referent. For example, a listener could initiate negotiation with "I'm confused. Are you talking about the thing that is kind of flared at the top? Couple inches long. It's kind of blue." Second, negotiation can

be with oneself. This self-negotiation is the one that this research was most concerned with. The listener considers aspects of the speaker's description, the context of the communication, the listener's own abilities, and other relevant sources of knowledge. He then applies that deliberation to determine whether one referent candidate is better than another or, if no candidate is found, what are the most likely places for error or confusion. Such negotiation can result in the listener testing whether or not a particular referent works. For example, linguistic descriptions can influence a listener's perception of the world. The listener must ask himself whether he can perceive one of the objects in the world the way the speaker described it. In some cases, the listener's perception may overrule parts of the description because the listener can't perceive it the way the speaker described it.

To repair the traditional approach, an algorithm was developed that captures for certain cases the listener's ability to negotiate with himself for a referent. It can search for a referent and, if it doesn't find one, it can try to find possible referent candidates that might work, and then loosen the speaker's description using knowledge about the speaker, the conversation, and the listener himself. Thus, the reference process becomes multi-step and resumable. This computational model, which is called "FWIM" for "Find What I Mean," is more faithful to the data than the traditional model (see Figure 2-2(b)).

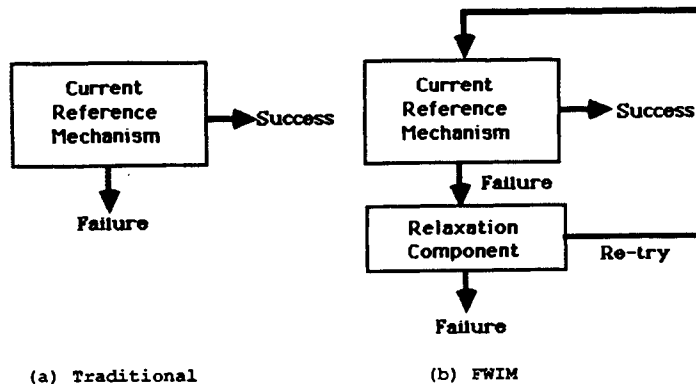


Figure 2-2: Approaches to reference identification

One means of making sense of a failed description is to delete or replace the portions that cause it not to match objects in the hearer's world. This program uses "relaxation" techniques to capture this behavior. The reference identification module treats descriptions as approximate and relaxes a description in order to find a referent when the literal content fails to provide the needed information. Relaxation, however, is not done blindly but is modelled on a person's behavior. A computational model was developed that can relax aspects of a description using many of the sources of knowledge used by people. Relaxation then becomes a form of communication repair (in the style of the work on repair theory found in Brown & VanLehn, 1980). A goal in the model is to use the knowledge sources to reduce the number of referent candidates that must be considered while making sure that a particular relaxation makes sense.

The component works by first selecting a set of reasonable referent candidates for the speaker's description (see also Joshi, 1978) by searching a knowledge base that represents objects in the world, scoring partial matches of each candidate to the speaker's description, and selecting those with higher scores. The component then generates, using information from the knowledge sources, a relaxation ordering graph that describes the order-to-relax features in the speaker's description. Finally, it combines the candidates with the ordering to yield the most likely referent. An ordered relaxation of parts of the speaker's description can be provided by consulting knowledge known about linguistics (the actual form of the speaker's utterance), perception (physical aspects of the world and the listener's ability to distinguish different feature values), specificity (hierarchical knowledge to judge how vague or specific a particular feature value is), etc.

The algorithm attempts to show how a listener might judge the importance of the features specified in a speaker's description using knowledge about linguistic and physical context. Figure 2-3 illustrates this process. The speaker's description is represented at the top of the figure. The set of specified features and their assigned feature value (e.g., the pair Color: Maroon) are also shown there. A set of objects in the real world are selected by the partial matcher as potential candidates for the referent. These candidates are shown near the top of the figure (C_1, C_2, \dots, C_n). Inside each box is a set of features and feature values that describe that object. A set

of partial orderings are generated that suggest which features in the speaker's description should be relaxed first--one ordering for each knowledge source (shown as "Linguistic," "Perceptual," and "Hierarchical" in the figure). For example, linguistic knowledge recommends relaxing Color or Shape before Function, and relaxing Function before Size. This rule was developed when it was noticed that speakers typically add more important information at the end of a description. Since the syntactic constituents at the end often are either relative clauses or predicate complements, a relaxation rule was created that recommends that the features in a speaker's description are relaxed in the order: adjectives, then prepositional phrases, and finally relative clauses and predicate complements.² A control structure was designed that takes the speaker's description, puts all the (partial) orders together, and then attempts to satisfy them as best it can. This is illustrated at the bottom of the diagram by the reordered referent candidates.

2.3. Related Work in Reference and Miscommunication

There are two major pieces of work in AI literature that laid the foundation for this research: those in reference and those in miscommunication.

A major starting point of this research was Cohen's (1981; Cohen, Fertig, & Starr, 1982; Cohen, 1984) detailed analysis of the pragmatics of reference and the effects of different modalities of communication. He showed that it was reasonable to

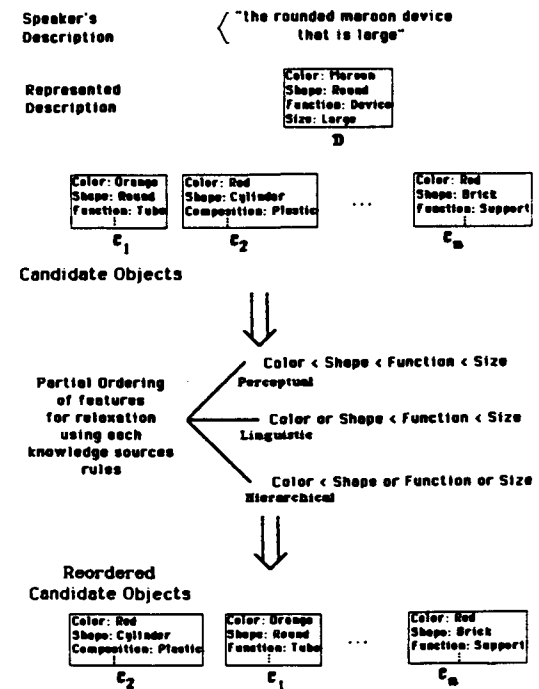


Figure 2-3: Reordering referent candidates

consider reference identification, instead of being intimately tangled, as separate from the whole process of language understanding. Cohen presented evidence (1981, 1984) that a speaker attempts as a separate step in his overall plan of communication to get a hearer to identify a referent. He provided grounds for an IDENTIFY action by illustrating particular requests to identify from his water pump protocols. For example, utterances like "Notice the two side outlets on the tube end" or "Find the rubber ring shaped like an O" showed that the speaker wanted the hearer to perform some kind of action. That action is the IDENTIFY act, which is to search the world for a referent for the speaker's description (and thus identify it). Cohen also showed that the hearer's response to a request to identify provided further evidence. He pointed out excerpts in the protocols where hearers responded to a request to identify with a confirmation that the identification had actually occurred (e.g., "Got it."). Cohen went on to show how referent fits into a plan-based theory of communication.

The reference paradigm we followed was closest to that developed by Grosz (1977), which was similar to others in the past (Winograd, 1971; Woods, Kaplan, & Nash-Webber, 1972): Put the speaker's description into a searchable form (i.e., parse and semantically interpret it) and then use that form as a pattern that can be compared against objects. A referent is found when the pattern matches one or more of the objects. The pattern and a target referent match each other if all the attributes specified in the pattern exactly fit the corresponding attributes

in the target. There is variability in each of the past reference schemes in what pattern is generated, how the world is represented, and how the actual search progresses, but the general scheme remains the same. Success occurs if, and only if, a perfect match exists between all the pattern's attributes and the corresponding attributes on a target. Grosz's reference mechanism departed from past works by introducing the notion of focus which provides a better way to resolve referents by constraining the search space. For definite noun phrases, the choice of possible referent candidates is guided by the focus mechanism. The information provided in the definite noun phrase itself (i.e., by the head noun and modifiers) is used to distinguish the referent from other objects in focus. Grosz showed how both the surrounding nonlinguistic environment and the global context of preceding discourse are part of focus and how it is used to resolve definite noun phrases. Grosz also proposed the need for inexact matching in the reference process should something go wrong:

The retrieval component can fail to find such a match even though for most people the noun phrase suffices to identify an object. . . . Alternatively, more than one object may match, but the ambiguity may not matter for the purposes of the utterance. The problem in either case is to determine the nature of the mismatch and whether it matters. . . . The focus mechanism provides one crucial element for deciding about inexact matches. It separates those items that are in the focus of attention from all other known

items. If an exact match cannot be found in focus, it is reasonable to ask if any of the items in focus come close to matching the description of the noun phrase (the question of what is close is the other crucial element in such decisions) and if so which is closest (Grosz, 1977, p. 161).

Ringle and Bruce (1982), in a survey of numerous types of miscommunication in conversation, point out problems across a wide spectrum of dialogue types and situations. They describe two primary ways that conversations fail. The first one, input failure, occurs when the listener is unable to form a complete or at least coherent, interpretation for an utterance. Input failure can occur due to such causes as misinterpretation of a single word, incorrect resolution of a referential term, or misplacement of a negation. Such failures cause the listener to misunderstand without weakening the listener's comprehension of the overall context of the communication (making the failures local in nature). The second way that Ringle and Bruce say that people fail, model failure, happens when the listener cannot incorporate the inputs into a coherent belief model as intended by the speaker. The problem can be due to an input failure when information is lost that is needed to assimilate the speaker's utterances into the belief model. It can also occur when a listener does not have sufficient background knowledge, has a different thematic emphasis than the speaker, or does not make the proper inference (or any at all) from the speaker's input. Ringle and Bruce describe repair techniques for failure that often occur between the listener and speaker. Such repairs are

usually initiated by the listener providing a failure cue (e.g., recapitulating the speaker's important points) to the speaker to indicate possible trouble and usually require action by both the listener and the speaker. Sometimes the dialogue situation affects the ability of the listener to provide such cues. For example, in a teacher-student relationship, it is hard for the student to interrupt the teacher's lecture/conversation to initiate a repair due to a mistake the student feels has occurred. In other conversational settings, such interruptions are easier.

McCoy (1985a, 1985b) focuses on a particular class of communication problems and misconceptions about the objects modelled by a system in its knowledge base. She is concerned with discrepancies between the beliefs of the system and that of the user as seen in their dialogue. Her work concentrates on two kinds of misconceptions about the properties of an object: misclassification and misattribution.

Misclassification occurs when one classifies an object incorrectly. For example, a person may think that whales are fish, when, in fact, they are mammals (McCoy, 1985b, p. 17). McCoy called the way to correct this problem a like-super strategy since an expert may believe that the user misclassified the misconception object (whale) because it is similar to the posited superordinate (fish). She defines two other kinds of misclassifications that her system can detect--Like-Some-Super and No-Support. Like-Some-Super occurs when the expert believes a user wrongly classified an object because it is like some

subclass of the posited superordinate. For example, a whale may be viewed by someone as a fish because they think that a whale is like a shark, and a shark is a fish (McCoy, 1985b, p. 24). No-Support occurs when the system can find no support in the user model for the misclassification. McCoy's system simply denies the incorrect information in that case and provides the correct information.

Misattribution is the second class of misconceptions that McCoy deals with. They occur when the user wrongly attributes a property to an object that the object doesn't have. Either the user has confused the object with one he thinks is a similar object or he has made a bad analogy from a similar object (the Wrong Object strategy). McCoy presents an example where the user attributes the "high liquidity" property of a money market fund to a money market certificate. Another reason that misattribution can occur is that the user attributes a related property to an object instead of the actual one (the Wrong Attribute strategy). An example that McCoy presents occurs when the user talks about the "interest" on the stock but really meant the "dividend." The correction in that case is the substitution of the proper property for the incorrect one. The last case of misattribution that McCoy considers is No Support, which occurs when the expert can find no support for the misattribution. In that case, McCoy's system denies the incorrect information and asserts the correct information.

McCoy's work demonstrates the power of representing objects using a taxonomic knowledge base that indicates an object's

superordinates and subtypes, and its attributes and their values. That paradigm allows her to notice several classes of users' misconceptions and to correct them.

3. Summary

The goal of this work is to build robust natural language understanding systems, allowing them to detect and avoid miscommunication. The goal is not to make a perfect listener but a more tolerant one who can avoid many mistakes, though it may still be wrong on occasion. During communication, problems can occur such as reference mistakes which are one kind of obstacle to robust communication. The succeed/fail paradigm followed by previous natural language researchers can be extended to tackle reference errors.

Real world objects are represented hierarchically in a knowledge base using a representation language, KL-One (Brachman, 1977), that follows in the tradition of semantic networks and frames. In this, the reference identification task looks for a referent by comparing the representation of the speaker's input to elements in the knowledge base by using a matching procedure. Failure to find a referent in previous reference identification systems resulted in the unsuccessful termination of the reference task.

A theory of relaxation developed to recover from reference failures provides a much better model for human performance. When people are asked to identify objects, they appear to behave in a particular way: They find candidates, adjust as necessary, re-try, and, if necessary, give up and ask for help. Relaxation

is an integral part of this process. The particular parameters of relaxation differ from task to task and person to person. This work models the relaxation process and provides a computational model for experimenting with the different parameters and incorporates the same language and physical knowledge that people use in identifying referents to guide the relaxation process. Knowledge is represented as a set of rules and as data in a hierarchical knowledge base. Rule-based relaxation provides a methodical way to use knowledge about language and the world to find a referent. The hierarchical representation makes it possible to tackle issues of a speaker's imprecision or over-specification by checking the position of a description and to use that to suggest possible repairs.

Interestingly, one would expect that "closest" match would suffice to solve the problem of finding a referent. However, this does not usually provide the correct referent. Closest match isn't sufficient because there are many features associated with an object and, thus, determining which of those features to keep and which to drop is a difficult problem due to the combinatorics and the effects of context. The relaxation method described circumvents the problem by using the knowledge that people have about language and the physical world to prune down the search space.

4. Future Directions

The FWIM reference identification system developed models the reference process by the classification operation of KL-One (Lipkis, 1982). A more complicated model for reference with a

complete identification plan that requires making inferences beyond those provided by classification is needed. The model could also require that the listener execute a physical action before determining the proper referent. Cohen gives two excellent examples of such reference plans (Cohen, 1984, p. 101). The first, "the magnetic screwdriver, please," requires the listener to place various screwdrivers against metal to determine which is magnetic. The second, "the three two-inch long salted green noodles," requires the listener to count, examine, measure and taste.

References

- Brachman, R. J. (1977). A structural paradigm for representing knowledge. Doctoral dissertation, Harvard University (Tech. Rep. No. 3605). Cambridge, MA: Bolt Beranek and Newman Inc.
- Brown, J. S., & VanLehn, K. (1980). Repair theory: A generative theory of bugs in procedural skills. Cognitive Science, 4(4), 379-426.
- Cohen, P. R. (1981). The need for referent identification as a planned action. In Proceedings of IJCAI-81 (pp. 31-35). Vancouver, B.C., Canada.
- Cohen, P. R., Fertig, S., & Starr, K. (1982). Dependencies of discourse structure on the modality of communication: Telephone vs. teletype. In Proceedings of ACL (pp. 28-35). Toronto, Ontario, Canada.
- Cohen, P. R. (1984). The pragmatics of referring and the modality of communication. Computational Linguistics, 10(2), 97-146.
- Goodman, B. A. (1984). Communication and miscommunication. Doctoral dissertation, University of Illinois, Urbana, IL.
- Goodman, B. A. (1985). Repairing reference identification failures by relaxation. In Proceedings of the 23rd annual meeting of the association for computational linguistics (pp. 204-217). Chicago, IL.
- Goodman, B. A. (1986). Reference identification and reference identification failures. Computational Linguistics. Accepted for publication.

- Grosz, B. J. (1977). The representation and use of focus in dialogue understanding. Doctoral dissertation, University of California (Tech. Note 151). Menlo Park: Stanford Research Institute.
- Joshi, A. K. (1978). A note on partial match of descriptions: Can one simultaneously question (Retrieve) and inform (Update)? Theoretical issues in natural language processing-2 (pp. 184-186). Urbana, IL.
- Lipkis, T. (1982). A KL-ONE classifier. In Proceedings of the 1981 KL-One workshop (pp. 128-145). Jackson, NH. Also, Report No. 4842, Bolt Beranek and Newman Inc., Cambridge, MA. Also, Consul Note #5, USC/Information Sciences Institute, October 1981.
- McCoy, K. F. (1985a). The role of perspective in responding to property misconceptions. In Proceedings of IJCAI-85, Los Angeles, CA. (pp. 791-793).
- McCoy, K. F. (1985). Correcting object-related misconceptions. Philadelphia: University of Pennsylvania, Department of Computer and Information Science, Doctoral dissertation.
- Ringle, M., & Bruce, B. (1982). Conversation failure. In W. Lehnart & M. Ringle (Eds.), Strategies for natural language processing (pp. 203-221). Hillsdale, NJ: Erlbaum.
- Winograd, T. (1971). Procedures as a representation for data in a computer program for understanding natural language. Doctoral dissertation, Massachusetts Institute of Technology, Cambridge (Report No. TR-84, Project MAC). Cambridge, MA: MIT.

Woods, W. A., Kaplan, R. M., & Nash-Webber, B. L. (1972). The lunar sciences natural language information system (BBN Report 2378). Cambridge, MA: Bolt Beranek and Newman Inc.

Footnotes

¹A version of this paper will be presented at the third Theoretical Issues in Natural Language Processing conference (TINLAP3) in Las Cruces, New Mexico in January 1987.

²A more general and more applicable rule is that information presented at the end of a description is usually more prominent. And therefore, more prominent features should be relaxed last.

