Message Types in Goal-Oriented Discourse

Seminar in Artificial Intelligence – Fall 1997
James Pustejovsky
Alex Feinman
afeinman@cs.brandeis.edu

Rosencrantz: Do you want to play questions?
Guildenstern: How do you play that?
Rosencrantz: You have to ask questions.
Guildenstern: Statement. 1-love.
(from "Rosencrantz & Guildenstern Are Dead", by Tom Stoppard)

Abstract

In goal-oriented communication, a conversation can be thought of as consisting primarily of transitions between a number of distinguishable states. The utterances that signify these transitions fall into a small number of recognizable categories. The purpose of this paper is to explore how to enumerate, for a chosen goal, these categories and to determine whether it is possible to create a theory whereby they can be consistently recognized, with an eye towards designing a system which will classify utterances for use in Mover's World, a multi-user problem solving environment.

Background

Current research in linguistics, and some common sense, points to the existence of prototypical situations that arise repeatedly in goal-oriented communication. It also points to stereotypical utterances that occur in these situations. From these ideas it seems plausible to construct a theory that explores classifying these predictable utterances into classes. The primary purpose of this paper is to investigate the validity of this idea.

This idea of prototypical situations has been covered before; in slightly varying forms, these scenarios have been referred to as conversational games (Mann88), and even conversation as a form of dance (Winograd88). Certainly it is readily apparent on examination that most conversation appears to follow complex rules that could be likened to a dance, with give and take on the part of the partners in conversation towards, if not some goal, at least the completion of the conversation. Winograd explores this idea at length. This paper, however, concerns itself not so much with the dance itself as with the steps: that is to say, the actual utterances which allow the conversants to move from one state of the conversation to another.

Definitions

Before we can start to classify utterances, we must decide what an utterance is. The term as it is used in this paper is borrowed from (Andernach96), among others. We will consider an utterance to be piece of a communication that communicates a single concept. This is ideal for our purposes; however, breaking up natural language into utterances is not a simple process, and has been the subject of some debate.

Work in (Andernach96) points to the straight-forward approach of using punctuation and conjunctions to break compound sentences into their constituent clauses, and using these clauses as the utterances. This may not prove entirely effective in the sort of communication the author is interested in; however, it is a good start. For the purposes of this paper, it will be assumed that the problem of extracting utterances is either solvable or surmountable, in that it may become unimportant what the exact boundaries of an utterance are if the information contained in it can be extracted anyway.

Another important decision to make is what constitutes a conversation. As suggested by Winograd, to restrict conversation to just the spoken dialogue is to miss much of the important information. A great deal is said by silence or omission, and even more by non-verbal communication. Obviously, without having facial and gesture recognizers, body language interpreters, and so on, it is difficult to get the full meaning of a conversation between people that can see each other. To attempt to avoid these problems, this paper will restrict its discussion to text-only communication, such as that seen in exchange of email or via a UNIX *talk* session. One justification for this is that the system this research will eventually be applied to limits the users to just this sort of communication. In addition to limiting communication to typed text, it restricts the conversants further to taking turns in conversation, not allowing simultaneous communication.

Finally, one more restriction must be added to the form of the conversation. As indicated by the references to goal-oriented communication above, this paper does not pretend to be able to determine a system for classifying all of human speech. In ordinary conversation, the topic of discussion shifts wildly, there are often interruptions and side bars to the discussion, and occasionally the conversants will purposely attempt to sabotage the goals of each other (for example, lying or otherwise restricting information). All of these complications can throw a real wrench into the dynamics of the conversation, and make it very difficult not only for any of the participants to achieve their goals, but also to recognize and categorize the communications that are taking place.

Since the type of communication the author is attempting to classify consists mainly of that between friendly users attempting to work together to solve a goal quickly, it is assumed that not only will the users 'play nice', that is, not give each other false or misleading information, but that they will stick to discussing the topic at hand, and not suddenly veer off into a tangent about Star Trek in the middle of solving a problem set. However unrealistic these restrictions may be in real-world communication, they seem quite valid to impose given the domain of interest.

Discussion

To determine likely categories for utterances it is instructive to look at some example communication. The examples below are artificial, and as such may not be representative of real communication, but the author feels they are sufficient to demonstrate the principles of classification.

In this simplistic example, A is inviting B out for dinner. It is assumed that they know each other, recognize each other, and have performed this task before:

1. A: Want to go out for dinner?
2. B: Sure.

Sentence (1) is a simple question, requiring only a yes or no answer. In this case, however, (1) also provides information about what the task is at hand. Another way to communicate this might have been "I am going out for dinner. Do you want to come?", which shows the two utterances separately.

Sentence (2) is a 'yes'. Discussion of how both parties know that replying "sure" means 'yes' harkens to work in (Walker96) and others. Here "sure" is an explicit form of acknowledgment, of the sort that could be listed and merely checked against. However, in addition to enumerating these sorts of acceptance, Walker shows how to extract acceptance from such utterances as "Dinner sounds good." or even "Let's go.", given the context of the task. The gist of that work is that logical consistency is a necessary indicator of acceptance, but not sufficient, and that logical inconsistency is sufficient as an indicator of rejection, but it is not necessary. Also, Walker investigates and explains such conversational gambits as repeating the fact in question to indicate belief in it, or altering some important facet of the new information to indicate rejection of the assertion.

It is important to note that after (2), the conversation is completed; no more communication is needed or expected by either party. This is a clear example of the existence of final states in a conversation. These states are places where it is socially acceptable for the conversation to be over; they do not indicate that the task has been successfully completed, or that either conversant is necessarily happy with the results. Winograd identifies these as states of competition.

This is about the simplest meaningful communication possible, but there are many more possible outcomes given the initial situation. Examination of a few more possibilities will reveal more categories:

1. A: Want to go out for dinner?
2b. B: Sorry, I can't.

This is a 'no', couched in polite terms. The conversation may continue into explanation or negotiation, but this conversation act is basically over, having again reached a final state.

1. A: Want to go out for dinner?
2c. B: I don't know if I'm free. Let me call you back in 10 minutes.

This is a much more complicated response. Here, B has promised A an answer. Winograd refers to this as a "Commit-to-commit". Despite the fact that B has not given A a definite answer one way or the other, the commitment to answer is enough to allow the conversation to end, at least temporarily.

1. A: Want to go out for dinner?
2d. B: Sure. Where are we going?

This is another very common situation. (2d) actually consists of two utterances, a 'yes' and a new question, which B must now deal with. Obviously, at the end of this segment, the conversation is not over. If A ceased communicating, changed topic, or otherwise treated the conversation act as finished, A would be at best confused and perhaps insulted.

1. A: Want to go out for dinner?
2e. B: Only if C isn't going.

This is a slight expansion on the type in (2d), above; in (2e), B has implicitly asked the question "Is C going?", but has also given A his answer to the question in (1), dependent on this new information.

As can be seen, even in the most straightforward of situations many complications can arise; nevertheless, most of these communications consist of sub-conversations which have the same general structures as the original conversation.

More complicated negotiation leads to even more convoluted conversational games. In addition, not all conversations consist of communication between only two conversants. However, superficial inspection points to the appearance of similar categories as those mentioned above, along with new categories for passing the conversational buck. In (Sacks79), examples of explicitly tracking this shift of control in the conversation reveal that categorizing this sort of utterance should at least be no more difficult to perform than recognizing the categories above.

To this point we have been giving a somewhat *ad hoc* system of classification. However, other researchers have done a great deal of work on classifying speech. In particular, (Searle75), expanding upon work in (Austin62), created another level of categorization, termed "illocutionary points", which Winograd uses exclusively in his research. This involves categorizing the utterance by the overall purpose in the communication the utterance has. Searle splits all utterances into five categories:

Directive, where the speaker attempts to get the listener to perform or commit to some action. ("Get the phone!")

Commissive, which commits the speaker to a course of action, or declines that course of action, with shades in between. ("I got it.")

Assertive, which states something the speaker thinks is true or false (to some degree). ("It's your mom.")

Declarative, whereby the speaker causes a change in reality by the speech act itself. ("You're in charge till I get off the phone, then.")

Expressive, in which the speaker expresses feelings or other psychological affects regarding the situation. ("I'm tired.")

Declarative and expressive illocutionary points would seem to come up much less often in goal-oriented communication than the other three. In fact, an argument could be made for the validity of a scheme that completely ignored all expressive utterances as irrelevant in such a domain. However, this may end up overlooking utterances whose function is not apparent from their category, and so for now all five illocutionary points must be considered.

It is unclear whether these universal categories should be treated as a basis in all domains, with no other categories, or should merely be used to aid construction of the set of categories for each domain. (Andernach96) uses them as the whole of the classification scheme. It may prove useful to first categorize each utterance into one of these, following work in (Winograd88), and then subcategorize the utterance according to domain.

To apply this classification to our example above, (1) is a directive, attempting to get B to respond with an assertive or commissive. From there we can sub-classify it as a question, and then perhaps what type of question it is. Likewise, (2) is an commissive act, which could then be sub-classed into a positive commitment to perform the task. On the other hand, (2c) could be seen as an assertive or a commissive, depending on the classifiers point of view. Given the fact that it acts as a commitment in the conversation, it must here be classified as a commissive. Likewise, an utterance such as "I don't feel good." seems like an expressive, but as a response to (1) it functions instead as a commissive, declining the dinner invitation.

Determining the Type of an Utterance

Once categories have been set up, the next trick is determining what category (or sub-category) a particular utterance belongs to. This is, of course, the real trick. Studying work by Walker and others in interpreting fragments of natural text in an attempt to extract meaning, it seems likely that a number of different tactics will have to be employed in combination to get any sort of results.

A straight vocabulary and punctuation analysis may yield the most information for the least work. Sentences ending in a question mark, for example, are nearly always questions, and in goal-oriented communication, almost always non-rhetorical ones. Likewise, sentences containing explicit references to a particular object in the problem domain – in the case of the author's research, things like boxes, handtrucks, and pianos – are either requests for information or the information itself. To determine which is the case, we must then turn to work such as that done by Walker, who uses the logical consistancy of the utterance with respect to the domain to determine the category of a statement.

Perhaps most importantly, the categorization of an utterance is often strongly dependent on the history of the conversation. Hence, any analysis of an utterance must be based on the history of the conversation. As with any such situation, we can condense this history to a single state; since the networks of the conversations are not terribly complex, the number of states will be manageably small. In light of this state-based approach, it makes sense to draw charts of typical conversational structures. Winograd calls these commitment networks, and the utterances which this paper discusses form the transitions between states. Hence, in any state, relevant types of communication will fall into one of the outgoing transition categories. This makes the job of deciding the category of an utterance quite a bit easier, provided the network is valid.

This idea of conversational networks is not a new one. The commitment networks idea seems to have grown out of, or at least be closely related to, earlier work by Searle, Austin, and others (see, for example, Searle75), who were among the first to identify the illocutionary points described about. This was folded into research by Sacks, Shegloff, and others (e.g., Sacks78), who proposed a formalization of turn-taking amongst conversational participants, with certain rules to be observed, and designated stopping points, to create the concept of commitment networks, with more complex rules governing who is to speak next and when the conversation is over than accounted for by the simple rules in the original methods proposed by Sack et al.

The aim of this paper is to leverage the deterministic nature of these commitment networks to narrow down the possible number of categories that valid utterances could fall into at any given place in a conversation, and to then use the classification of previous utterances to locate the conversational state within the network. Obviously it can be seen that if the state of the conversation is known, and the conversants are guaranteed (or at least asked nicely) to follow the 'conventional' sort of goal-oriented conversation style described toward the beginning of this paper, it should be possible to narrow down any particular utterance to a handful of valid categories.

Having these conversational networks be the basis for determining utterance type just adds another level to the problem; to classify an utterance, it is necessary first to build a commitment network for the task at hand. This brings up yet another level: deciding what task the users are attempting to perform. Plan recognition is yet another hot area of research; it may be beyond the scope of this paper to discuss it at length, however.

At this point it is important to again bring up some more complications. As Winograd points out, most conversations do not go through every state in a commitment network explicitly, or is every transition marked by a clear communication. Often, silence indicates that the default action should be taken; sometimes, rather than respond by communication, one conversant or the other simply takes the action discussed. Classifying the meaning of silences, or of actions other than communication, is still a problem.

For the system this research is aimed at, all forms of interaction between users are recorded, and so there is at least the hope of analyzing their non-written communication in hopes of picking up on alternate methods that the users communicate. Taking silence to indicate the default action is the easiest route, but then trouble arises deciding what the default action is for a given scenario.

The varying number of categories complicates matters if the domain being analyzed is unfavorable; however, confining our area of investigation to goal-oriented communication helps eliminate many spurious types of communication, and makes classification more feasible. But there may still be too many difficulties to get any sort of consistent classification.

Structured Communication

Given all these complications, it appears very difficult to classify utterances automatically without helping the system somewhat. In further research it may be possible to work around these problems by letting the communicator provides more cues as to the type of utterance he or she is expressing. Obviously it is not feasible to have the user of some system be required to check the "Assertive" box every time he submits an assertive measure; not only is this horribly intrusive, but the user may not even be aware of the correct type for a message he is sending.

One possible alteration of the playing field that may make it easier to extract information about the correct category for an utterance while simultaneously reducing the user's workload is to require all communication be structured or semi-structured. In this case structured means allowing communication only from a small subset of commands and nouns in a particular order, allowing only a communication chosen from a preset list, or some such restriction. This is probably not feasible for most domains. Interestingly, even such a restricted domain does not reduce the problem to the trivial one of enumerating the category for each possible communication. This is due to the fact that most communications could serve more than one purpose, given the context. Humans are exceedingly clever at finding ways to communicate more than a restricted system would ordinarily allow; for example, using timing of a message to denote importance or even reverse the meaning of an utterance. So restricting the users to fully structured communication would harm much more than help.

Semi-structured communication is a much less constrictive style of communication where certain portions of the communication may be tagged or otherwise marked, but the actual body of the message is unrestricted. Malone uses E-mail as his example; in a standard email message, as defined by RFC-822 (Crocker82), a message has a number of set headers, each of which has a range of valid values that vary from completely restricted (the date header, restricted to a valid date in a specified format) to completely free (the subject header, which is only restricted by a length limitation). In addition, email has a completely unstructured message body, allowing any amount of information in any form.

Work by (Malone88) and others indicates that requiring a certain amount of structure in messages not only allows much easier machine parsing of the message, but actually improves human performance even when the structuring information is not acted upon by some external agent. This result is not entirely surprising (as Malone would have it) given that in many cases natural communication acquires a tightly defined structure in situations where performance is critical. Conversation in such situations as a cockpit or crowded kitchen often follow very tightly defined protocols, with explicit checks on everything to avoid costly errors. To avoid excess mental energy being expended trying to figure out what the speaker has said, the conversants tend to use formulaic forms of conversation, including complete repeated phrases with only the subject or object changed. This allows the listener to pick out the important piece of information immediately, and then (if necessary) figure out the rest of the message to get context.

In addition, these conversations tend to show every step in the commitment networks described above, which allows even casual inspectors to see the flow of the conversations through the sort of states shown above. As opposed to the use of silence as default action described above, in a critical situation almost every step is explicitly acknowledged before any commitment is taken. This is very beneficial for analytical purposes and for verifying the validity of the commitment network idea. Here is an example the author surreptitiously collected while working in a busy kitchen preparing food with about a dozen other people:

1. Cook: Do we need these scraps?
2. Head cook: Just a minute.
3. Cook: Ok. (waits)
4. Head cook: What was that?
5. Cook: Do we need these scraps?
6. Head cook: No.
7. Cook: So I should throw them out?
8. Head cook: Yes.
9. Cook: Ok, thanks.
10. Head cook: (turns away)

In a less hectic environment, at the very least sentences (1) through (4) would be unnecessary. (1) is actually a request for attention, and (2) is an explicit commit-to-commit, with the task they refer to being the actual start of a conversation. Because the attention of the head cook is fiercely competed for, such extra handshaking is needed to assure that the cook is talking to correct person. Of course, this has a cost, as seen by the fact that no new information about the task is exchanged until (6), where the head cook finally responds to the original question.

Sentence (7) is a double-check that the default action is the correct one. Again, this occurs because the tasks are fateful; that is, they have irreversible consequences. Likewise, (9) is an explicit acknowledgment that the conversation is over, and that the attention of the head cook is no longer needed. The head cook's response is one of the rare non-verbal, non-explicit utterances in the domain, allowed because the conversation is more or less over, and leaving the other cook in a wedged state is much less disastrous than holding the attention of the head cook.

For all of this added complication, the conversation still fits nicely into the commitment network idea. There are actually three tasks to be performed here; first, getting the head cook's attention. This task is completed successfully at (4), which allows the main task (asking what to do with the scraps) to continue (5-6). During this we move through the third task (7-8), which is just a clarification of the second, and then to the end of the conversation (9-10).

This explicitness makes extraction of the state trivial by hand, but may not aid automatic classification nearly as much. Certainly the lack of silent transitions will help. But the actual classification may be made even harder by the chopped language and oblique references.

Related Work

There are a number of related topics being investigated by other researchers. (May97) has constructed a system which will automatically classify incoming email from a specific mailing list given a number of key phrases to search the message from. These phrases were selected by hand, and allow the system to categorize the messages into four categories ("Administration", "Question", "Discussion", "Announcement") with at least statistically significant accuracy. This is similar to the idea of classification this paper is concerned with, although the scheme May describes categorized the messages according to a sort of meta-content rather than their direct role in a conversation.

The disparity is not as great as it seems, however; he describes how structures of question and reply often evolved into discussion of a topic, which points to at least one commitment network. This would involve one user asking a "Question", leaving the conversation in a sort of waiting-for-answer state; if someone answered the question, with "Discussion", the conversation would transition to a 'discussion' state, which could spawn new questions (and start the cycle again), or more discussion, with no end in sight. The other two categories, however, would not participate in group discussion, being somewhat one-sided announcements ("Administration" messages generally taking the form of an announcement, but from a figure of authority with respect to the mailing list) which require no response.

The work here is also related to a program written by the author last year (Feinman96), which used a clustering algorithm to classify email messages into folders. The author's work, based on the well-known Cobweb clustering system (Fisher87) was able to achieve high accuracy with only a limited corpus. The system, which automatically determined the phrases to key on, seems to achieve better accuracy than May's system, which uses manually selected phrases. Obviously, this may be due to poorly-chosen key phrases, but puts in another vote for the of completely automating the classification process.

Future Work

The next task is building a computer program that would automatically classify utterances in a very narrowly-defined environment. Given tight constraints on the subject of conversation, it should be possible to use various cues in the utterances to classify them, especially if the conversation follows the networks described above. The constraint of requiring only goal-oriented communication seems legitimate in a system where the conversants are at 'work', as is the case with the author's research. This also holds for requiring

There are still a number of major hurdles to be tackled; separating utterances from the actual communication, building the commitment networks, and so on. But this paper has served to prove the possibility of useful classification of utterances, if not the feasibility of such classification. As the field moves towards more sophisticated analysis of plain text, the classification problem will get easier.

In the problem the author will be applying this research to, initial analysis of data sets indicate the sub-categorization approach described above, where messages are first categorized according to (Searle75), may yield favorable results. This will most likely be the starting point, along with the incorporation of semi-structuring into the system to give it a bit more information to work off of.

This leads to the next challenge; determining the sort of structuring which is appropriate to the problem. Obviously, the sort of structuring should be advantageous to the user, but in addition to providing extra information to aid classification. For example, a simple choice of the general type of communication, perhaps from a list such as "question, answer, other" or some such would give a starting point for classification. Another option is to use the existing information of who is talking to whom, which already must be explicitly indicated by the user, and stitch it to past communications to decipher where in a commitment the agents are.

In other work, the author intends to resurrect his mail-munging code and run it on the corpus used by May's system to compare results directly. This should determine if the improvement in classification from (May97) to the author's system described above was in fact due to better selection of key phrases, as the actual clustering algorithm is very similar. This will have direct bearing on the approach taken to implementing the main classification system; if the clustering system appears advantageous it will be used rather than depending on manual classification.

Many applications would benefit from even such a constrained system of classification as has been described here. The research the author is participating in deals with a system where many users must communicate effectively to coordinate a planning task; the aim is to have the system recognize their plans by classifying their communication and determining where in the planning conversation they are. The system could then provide this communication to the users in any number of ways, including providing default responses to communication, or prompting novice users with possible actions.

Other possible applications include prioritizing or even filtering of Usenet news messages, as defined in RFC-1036 (Horton87). As a public forum, the signal-to-noise ratio of Usenet is often quite low, requiring the paging through of literally thousands of worthless messages to . This is similar to the systems described above, but could make use of the continuance of a conversation to aid in classification of each message. As the Usenet message format tracks the thread of conversation by recording the previous message in its headers (specifically, the References: header, which gives the Message-IDs of the messages this one refers to), it should be quite easy to use this information intelligently. Already, most fancy news readers will diagram the network of messages and replies, but the intent here would be to create a filtering algorithm which actually flags or removes uninteresting or repetitive messages, which abound in Usenet communication.

For example, most Usenet conversations start with a request for information, often followed by a number of "me too"'s, people who also desire the information but are too impatient or clueless to wait for an actual answer to be posted. This is then followed by an answer, followed by a number of other people answering the same question again, either because they think they know better than the original answerer, or because they, too, clueless and haven't checked whether someone else already answered the question. Finally, if the topic seemed interesting, it will spin off other conversations, and often a personal diatribe or two from people complaining about each others' grammar, views, or personal hygiene. Finally, someone will mention Hitler or Nazis[1], at which point the thread is declared dead and a new one is resurrected.

From this morass, usually the reader only desires the original question and the 'best' answer, or perhaps all the answers if they differ significantly. By tracking the cycle shown above, it should be possible to aid the user in weeding out unimportant messages. The cycle shown above does not easily fit into a single commitment network on the face of it, but does fit if some null transitions are allowed. It should be possible to leverage this additional information into existing filtering systems to extract information about each message. Whether it will make that much of a difference in the classification is uncertain, but it is certainly another thing to look into.

Many other applications abound. Anywhere where the state information of a conversation is important, a system that could consistently recognize that state would be useful. For example, researchers at MIT are doing work on wearable computers that are always on and act as auxiliary processors for the human user, supplying side bar information, scheduling, and communication with the network in real time (see, for

---

[1] This particular behavior of Usenet is referred to as Godwin's Law. From (Raymond97):
Godwin's Law, prov. [USENET]
"As a USENET discussion grows longer, the probability of a comparison involving Nazis or Hitler approaches one." There is a tradition in many groups that, once this occurs, that thread is over, and whoever mentioned the Nazis has automatically lost whatever argument was in progress. Godwin's Law thus guarantees the existence of an upper bound on thread length in those groups.

example, Rhodes97). This agent is intended to aid the user in remembering details about the person he or she is speaking to, and tracking old conversations. By using conversational state as part of the world context the agent bases its information gathering off of, the user could carry on extended conversations with that person, and avoid repeating information. Similarly, one could imagine a crowded room full of people the user has talked to or needs to talk to, with markings overlaid on the user's visual field color-coded according to where in each conversation the user is. If the user sees a big knot of green people, say, who all need his attention, he could wander over to them to answer their questions or steer away from them as necessary. This is a somewhat fanciful, but is not entirely impossible, as seen by demonstrations of equally fanciful tools at the International Symposium on Wearable Computing.

Once the major hurdles of plain text recognition, utterance extraction, plan recognition, and non-verbal interpretation are conquered, this technique should prove very successful.

References:

T. Andernach (1996) A Machine Learning Approach to the Classification of Dialogue Utterances. To appear in *Proceedings of NeMLaP-2*

E. Auramäki, et al. (1988) A Speech-Act-Based Office Modeling Approach. In *ACM Transactions on Office Information Systems,* Vol. 6 issue 2

J. Austin (1962) *How to Do Things With Words.* Harvard University Press, Cambridge, Massachusetts

G. Brown and G. Yule (1983) *Discourse Analysis.* Cambridge University Press, London

D. Crocker (1982), RFC #822: Standard for the Format of ARPA Internet Text Messages. Among other sources, this is available from: `http://sunsite.doc.ic.ac.uk/packages/rfc/rfc822.txt`

A. Feinman (1996) A Semi-Autonomous Mail Agent for Filing Electronic Mail. Unpublished.

D. Fisher (1987) Knowledge Acquisition via Incremental Conceptual Clustering. In *ML 2*

M. Horton (1987) RFC #1036: Standard for Interchange of USENET Messages. Available from: `http://sunsite.doc.ic.ac.uk/packages/rfc/rfc1036.txt`

T. Malone, et al. (1988) Semi-structured Messages are Surprisingly Useful for Computer-Supported Coordination. In I. Greif, ed., *Computer-Supported Cooperative Work: A Book of Readings,* Morgan Kaufmann Publishers, Inc.

W. Mann (1988) Dialogue Games: Conventions of Human Interaction. In *Argumentation,* Vol. 2

A. May (1997) Automatic Classification of E-mail Messages by Message Type. In *Journal of the American Society for Information Science,* Vol. 48 issue 1

E. Raymond (1996) The Jargon File, v. 4.0.0. The definitive copy is available from: `http://www.ccil.org/jargon/`

B. Rhodes (1997) The Wearable Remembrance Agent: A system for augmented memory. In *The Proceedings of The First International Symposium on Wearable Computers (ISWC '97),* Cambridge, Massachusetts

H. Sacks, E. Schegloff & G. Jefferson (1978) A Simplest Systematics for the Organization of Turn-Taking in Conversation. In J. Schenkein, ed., *Studies in the Organization of Conversational Interaction*, Academic Press, New York

J. Searle (1975) A Taxonomy of Illocutionary Acts. In K. Gunderson (ed.), *Language, Mind and Knowledge,* Minneapolis: University of Minneapolis Press

D. Traum and J. Allen (1994) Discourse Obligations in Dialogue Processing. In *ACL 94*

D. Traum and E. Hinkelman (1992) Conversation Acts in Task-oriented Spoken Dialogue. In *Computational Intelligence*, Vol. 8 number 3

M. Walker (1996) Inferring Acceptance and Rejection in Dialogue by Default Rules of Inference. In *Language and Speech*, 39-2

T. Winograd (1988) A Language/Action Perspective on the Design of Cooperative Work. In I. Greif, ed., *Computer-supported Cooperative Work: a Book of Readings,* Morgan Kaufmann Publishers, Inc.

Corpus from problem sets for Mover's World, a collaborative problem-solving domain written by S. Landsman, et al.