Toward Recognizing and Conveying An Attitude of Doubt Via Natural Language

Sandra Carberry*, Lynn Lambert, and Leah Schroeder***

*** Department of Computer Science**
University of Delaware
Newark, DE 19716 USA

**** Department of Physics, Computer Science, and Engineering**
Christopher Newport University
Newport News, Virginia 23606 USA

Abstract: This paper addresses the dual problems of recognizing when an agent is expressing an attitude of doubt toward a proposition and of realizing natural language utterances that effectively convey such doubt. It presents naturally occurring examples that illustrate the wide variety of ways in which doubt can be expressed. The paper presents a recognition algorithm that uses linguistic, contextual, and world knowledge to determine that an agent is expressing doubt at a proposition and to identify the proposition that is being doubted. It then presents a realization algorithm that uses these knowledge sources to formulate utterances that will convey doubt when an agent cannot accept a proposition due to conflict between it and some other proposition in the agent's knowledge base.

1 Introduction

Much of the early research in human-computer interaction was concerned with systems that could understand a query, extract information from a database, and return an answer to the user. Subsequent research moved from simple question-answering to intelligent exchange of information between humans and machines, using menus, icons, natural language text or speech, pointing gestures by artificial agents, etc. Only recently has research begun to investigate computational mechanisms for recognizing and exhibiting attitude, emotion, and personality.

Our research addresses the dual problems of recognizing the attitude of an agent toward a proposition and of realizing natural language utterances that effectively convey the attitude of an artificial agent. Section 2 first argues that emotion, personality, and attitude are very different concepts, all of which are important in the design of artificial agents that are natural, realistic, and effective. Section 3 describes some previous work that is related to our research on recognizing and conveying an attitude of doubt. Section 4

presents naturally occurring examples that illustrate the wide variety of ways in which doubt can be expressed, along with other naturally occurring examples showing that the form of the utterance does not by itself determine the attitude of the speaker. Section 5 presents our work on recognizing and conveying an attitude of doubt. It begins by arguing that in a collaborative dialogue, communicating an attitude of doubt toward a proposition requires that the agent also convey (explicitly or implicitly) the beliefs that underly this attitude and the strength of these beliefs. Section 5.1 presents our algorithm for recognizing an agent's doubt, and Section 5.2 presents our rules for constructing natural language utterances that effectively convey an agent's attitude during a collaborative dialogue. Section 6 concludes with a discussion of our future work that will incorporate intonation and gesture.

2 Attitude, Emotion, and Personality

Emotion is any strong overall feeling that is physically evident and affects one's mental processing. Examples of emotion include anger, fear, and joy. *Personality* is the composite of one's behavioral traits or distinguishing characteristics, such as the behavioral trait of being outgoing. On the other hand, we take *attitude* to be one's disposition, or mental state, toward a particular proposition or situation. For example, one may doubt a particular proposition expressed by another agent or be confused by a proposal offered by the agent. Attitudes often have no physical manifestation nor do they interfere with mental processing.

Recognizing the personality and emotions of a human agent can help a computer system exhibit behavior that will establish rapport with the user. Consider, for example, a system that acts as a psychologist to humans with minor affective disorders. The system must recognize the emotions of the user as he or she discusses problems and issues with the system and must adapt its own behavior accordingly.

A collaborative artificial agent must also recognize a user's attitude toward a proposition expressed in a dialogue. For example, the user might fail to accept a proposition due to some conflicting evidence and thus try to convey his doubt to the system. Or the user might fail to fully understand a proposition and thus try to request clarification. Only if the artificial agent recognizes these different attitudes can it address the doubt or confusion, resolve it, and thereby continue the dialogue in an effective manner. Similarly, an artificial agent must communicate to the user its attitude toward different propositions and situations. If it fails to do so, the user will be unable to resolve discrepancies in their beliefs and the collaboration will in all likelihood be unsuccessful or at best inefficient.

In their experimental work, Nass et. al.[Nass et al., 1995] found that "the social rules guiding human-human interaction apply equally to human-computer interaction". We hypothesize that artificial agents are more likely to achieve their full potential as cooperative partners in problem-solving and decision-making if they can interact with users in a natural manner and can recognize and exhibit the same attitudes as human agents. Our work on recognizing and conveying attitude, particularly an attitude of doubt, is directed toward such natural, effective communication.

3 Related Work

Recent research in affective computing has run the gamut from psychological research exploring how humans respond to and interact with embodied agents[Nass et al., 1995, Nass et al., 1996], to computational research investigating issues important in constructing effective embodied artificial agents[Bartneck, 2001, Chopra-Khullar and Badler, 2001], to applied research aimed at actually constructing animated agents that are natural and believable[Cassell et al., 1994, Lester and Stone, 1997, Cassell et al., 2000]. In the area of language, researchers have explored issues such as how cognitive overload may be inferred from an agent's speech[Muller et al., 2001] and how emotion is conveyed by prosodic aspects of speech[Mozziconacci, 2001]. [De Rosis, 2001] provides an excellent overview of work on affective computing.

Our research has been concerned with recognizing and conveying an attitude of doubt, such as occurs during the initiation of negotiation subdialogues. Although a number of researchers[Cohen, 1987, Young and Cohen, 1991, Zukerman et al., 1998, Chu-Carroll and Carberry, 2000, Zukerman, 2001] have investigated argument understanding and generation, they have been primarily concerned with identifying the structure of a discourse, determining intended evidence relations, generating arguments that support a proposition, and identifying the content of a response that will resolve conflicts during negotiation. None have investigated how to recognize that an utterance is expressing doubt or how to generate an utterance that will effectively convey doubt.

Our research has been influenced by previous work on discourse processing. Grosz and Sidner[Grosz and Sidner, 1986] noted three kinds of information that contribute to determining the purpose of a discourse segment and its relation to the previous discourse: linguistic markers, utterance-level intentions, and general knowledge about actions and objects. Other researchers[Reichman, 1985, Hirshberg and Litman, 1993, Knott and Mellish, 1996] have investigated the contribution of cue words such as *But* and *By the way* to

understanding discourse. Our algorithms for recognizing and realizing utterances that convey doubt use a combination of linguistic, contextual, and world knowledge, including cue words and stereotypical beliefs about the world.

4 Communicating Doubt

Our project has the goal both of recognizing a variety of attitudes and of developing algorithms for an artificial agent to effectively convey such attitudes. In our current work, we have focused on instances in which an agent doubts a proposition or proposal communicated by another agent and uses natural language to convey that doubt.

To identify how human agents express doubt, we analyzed a number of naturally occurring dialogue corpora, ranging from a corpus of financial planning dialogues that were transcribed from radio talk shows[Transcripts, 1982] to a corpus of simulated travel planning dialogues[Transcripts, 1992]; we also noted instances of expressions of doubt appearing in the dialogues of popular novels and natural conversations. Our analysis showed that doubt is expressed in a wide variety of ways, including[1]:

- by contending or suggesting that some conflicting proposition is true. Such expressions of doubt often take the form of a surface-negative question, as in the first two examples below where the surface-negative questions in utterances (4) and (9) express doubt at S1's recommendation, or as a statement posed as a question such as *Even though it's 4 more years?* The expression of doubt can also take the form of a simple declaration of the conflicting proposition, as in utterance (15) of Example-3, though such utterances are often closer to outright rejections.

Example-1:

(1) S1: *I would like to see that into an individual retirement account rollover in a mutual fund group.*
(2) S2: *At my age?*
(3) S1: *Yes.*
(4) S2: *Uh, yeah but isn't there any risk?*
(5) S1: *Oh hey, there's risk in anything–There's risk in those eight stocks you have.*
(6) S2: *Yeah-*
(7) S1: *There's risk in your money market fund. If you are concerned about risk, then put it into a money market fund that invests in only United States securities.*

[1] All of the examples in this section are taken from our dialogue transcripts.

Example-2:

(8) S1: *Leave a maximum of a 1000 in there and put 9000 into a money market fund.*

(9) S2: *well can i do that through - but that's not insured is it? ...*

(10) S1: *Well that's ok. don't worry about it. if you're worried about it being insured what you do is put it in a money market fund that's involved in United States securities.*

Example-3:

(In this example, the name of the computer lab is *Linc.*)

(11) S1: *Do you know the name of the printer?*

(12) S2: *Linc.*

(13) S1: *Just Linc?*

(14) S2: *Yes.*

(15) S1: *But there are two printers.*

(16) S2: *I know. But the one is not hooked up to the network.*

(17) S1: *Oh. Thanks.*

- by drawing attention to a feature that is inconsistent with the doubted proposition or proposal. Such utterances were often realized by elliptical fragments, such as *"At my age?"* in utterance (2) of Example-1 above.

- by repeating some facet of the doubted proposition, sometimes with extra information, as in utterances (20) and (22) of the following two examples.

Example-4:

(19) S1: *OK, what I'd like you to do is to lend him the 20 thousand.*

(20) S2: *Lend it — 20 thousand?*

Example-5:

(21) S1: *The maximum amount of credit that you will be able to get will be $400 on their tax return.*

(22) S2: *$400 for the whole year?*

- by a general query of the doubted proposition, as in *"You have 40 thou(sand) in a mm fund?"*

- by an explicit statement of disbelief, as in *"I'm not so sure of that."* or *"That doesn't sound right."*

However, the form of the utterance (such as a surface-negative question or an elliptical fragment) does not by itself indicate that the utterance is expressing doubt. Consider the following dialogue excerpts, where the surface-negative questions in utterances (26) and (30) merely seek confirmation of a proposition, rather than expressing doubt:

Example-6:

(23) S1: *And if you have more money left after you pay the taxes, what difference does it make if you pay a few bucks more in taxes?*

(24) S2: *I'm telling my wife but she won't listen.*

(25) S1: *Well maybe she'll listen to me.*

(26) *If you get 200 bucks – isn't it better to have 200 bucks and have 200 left than to have nothing at all?*

Example-7:

(27) S1: *We have an all savers certificate that's due next october ...*
(28) S2: *how much is that for?*
(29) S1: *15-8.*
(30) S2: *Ok - and you got in right on the ground floor, didn't ya?*
(31) S1: *First day.*
(32) S2: *Yup – 12.61% – that was a goodie.*

Similarly, elliptical fragments can repeat some facet of the preceding proposition or introduce a new feature, yet serve other purposes than to express doubt. In the following two dialogue excerpts, the elliptical fragment in utterance (35) seeks clarification by repeating a term from the proposition expressed in utterance (34), and the elliptical fragment in utterance (39) attempts clarification by introducing a new feature into the proposition in utterance (38). But in neither case is S1 expressing doubt.

Example-8:

(33) S1: *Okay, so that's all that form is going to ask me – how much I paid for it and how much I sold it for?*
(34) S2: *Gross receipts less cost of goods sold.*
(35) S1: *Cost of goods sold? To me?*
(36) S2: *Right. Your cost of the goods that you sold.*
(37) S1: *So in other words, the actual worth of the product, the gross receipts?*

Example-9:

(38) S2: *Have you made any deposits at all to the bank?*
(39) S1: *These FTD deposits?*

5 Recognizing and Conveying Doubt in Collaborative Dialogue

As illustrated by the examples in the preceding section, there are many ways in which a speaker can convey an attitude of doubt. Moreover, the surface form of the utterance is insufficient to distinguish expressions of doubt from other actions. Thus not only is recognizing doubt a complex problem, but also an artificial agent that is attempting to convey doubt must be careful to select a mechanism that will be correctly interpreted by the hearer.

In a collaborative dialogue, it is most effective if a speaker not only communicates his doubt at a proposition but also the underlying basis for this doubt and the strength of these beliefs. This enables the listener to respond by addressing the reason for the doubt and facilitates an efficient negotiation subdialogue in which the agents resolve the discrepancies in their beliefs. We conjecture that this is why simple statements of disbelief, such as *"I don't accept what you are saying"*, are not the dominant form for conveying an attitude of doubt.

The next two sections present our work on recognizing and conveying doubt. Our recognition al-

gorithm addresses the problems of determining when an agent is expressing doubt and of identifying the proposition P_{doubt} that is being doubted. Our realization algorithm addresses the problem of generating natural language utterances that effectively convey an attitude of doubt toward a proposition P_{doubt} when there is a conflicting proposition P_i. Of course, when spoken language is involved and the conversants can see one another, prosody and gesture play a significant role in conveying an attitude of doubt. Our future research will take prosody and gesture into account.

5.1 Recognizing an Attitude of Doubt

Our research has produced an algorithm for recognizing when an agent is expressing doubt at a proposition P_{doubt} by querying some other proposition P_i that he believes is true. Consider the following dialogue segment and the alternative responses by S1:

(40) S1: *How is Tel-Micro expected to perform this year?*
(41) S2: *Tel-Micro is projected to have a $200 million profit this year.*

(42a) S1: *Didn't Tel-Micro have a large loss this past year?*
(42b) S1: *Did Tel-Micro have a large loss this past year?*
(42c) S1: *Isn't Tel-Micro expanding its research division?*
(42d) S1: *But isn't Tel-Micro expanding its research division?*

Utterance (42a) is clearly expressing doubt at the proposition that Tel-Micro will have a $200 million profit this year, and the doubt is motivated by the belief that Tel-Micro had a large loss for the past year. On the other hand, utterance (42b) seems to be merely requesting further information about Tel-Micro's performance. Thus one might hypothesize that surface-negative questions constitute expressions of doubt. However, this is not always the case, as illustrated by utterance (42c) which only requests information about Tel-Micro's research activity. The difference between (42a) and (42c) seems to be that companies with large losses for the past year would not generally be expected to have huge profits the following year, whereas expanding one's research division would not seem to generally contradict huge profits. On the other hand, if we precede utterance (42c) with the cue word *But*, as in utterance (42d), then it does seem to be conveying an attitude of doubt, but the reason for the doubt is unclear — perhaps S1 believes that expanding a research division is a costly enterprise and that such expansions prevent huge profits. Thus recognizing expressions of doubt is not an easy task.

5.1.1 Conditions for Recognizing Doubt

Expressions of doubt often refer implicitly to some proposition that is part of the dialogue context, namely the proposition P_{doubt} that is the object of the speaker's doubt. In order for such complex discourse acts to be recognized as intended, there must be evidence available that the hearer can use in deducing the speaker's intention. In particular, to recognize that an utterance is expressing doubt at a proposition P_{doubt} by querying some other proposition P_i that the speaker believes is true, our algorithm requires evidence that:

1. the speaker has some belief in the proposition P_i
2. the speaker believes that the hearer believes P_{doubt}
3. the speaker believes that if P_i is true, then P_{doubt} is false

First, there must be evidence that the speaker believes that P_i is true; otherwise the hearer could not be expected to believe that the speaker was using a conflict between P_i and P_{doubt} to express an attitude of doubt toward P_{doubt}. Second, there must be evidence that the speaker believes that the hearer believes P_{doubt}, since it is pointless in a collaborative dialogue to express doubt toward a proposition about which there is no disagreement. And third, there must be evidence that the speaker believes that if P_i is true, then P_{doubt} is false; otherwise the hearer could not be expected to infer that the speaker believes that the truth of P_i raises questions about the truth of P_{doubt}.

In addition, certain salience requirements must be met. P_{doubt} must be salient or prominent in the dialogue and must not yet have been accepted by the speaker. Otherwise, the hearer has no basis for associating the speaker's doubting attitude with the proposition P_{doubt}. In addition, the speaker must have made the conflicting evidence P_i salient in the dialogue. Otherwise, the hearer has no reason to use P_i in trying to discern the speaker's attitude.

5.1.2 Sources of Evidence

Our algorithm uses information from several knowledge sources as evidence in recognizing expressions of doubt:

- Linguistic knowledge: including the surface form of the utterance (such as a surface-negative question) and the presence of cue words

- Contextual knowledge: including beliefs derived from the dialogue, the structure of the dialogue, and the relative salience of different propositions in the dialogue
- World knowledge: including stereotypical beliefs that are presumed to be held by speaker and hearer

These knowledge sources provide evidence for the three requisite beliefs noted in Section 5.1.1. Linguistic knowledge about the surface form of the utterance provides evidence for the first belief, that the speaker has some belief in P_i. For example, a surface-negative question of the form *Isn't* P_i*?* conveys some degree of belief in P_i, whereas a simple interrogative such as *Is* P_i*?* does not. Contextual knowledge provides evidence for the second belief, that the speaker believes the hearer believes P_{doubt}. In a collaborative dialogue, an agent's statements are presumed to be felicitous; thus if the hearer has previously stated P_{doubt}, then one is justified in ascribing to the speaker the belief that the hearer believes P_{doubt}. World knowledge, in the form of stereotypical beliefs, provides evidence that the speaker believes that P_i and P_{doubt} are in conflict. Consider the stereotypical beliefs shown in Figure 1 and the dialogue segment shown in Figure 2. After utterance (47), there are three open propositions that have been posited by S2 but not yet accepted by S1: 1) the proposition that ARG-Growth is a good purchase, 2) the proposition that all growth mutual funds have fallen in value, and 3) the proposition that ARG-Growth does not have any bad investments. Each of these propositions might motivate an expression of doubt by S1. The challenge for S2 is twofold: to recognize any expression of doubt and, when one occurs, to identify the proposition that is being doubted.

Let us examine three possible continuations of the dialogue in Figure 2:

(48a) S1: *Doesn't ARG-Growth own stock in companies that have filed for bankruptcy?*
(48b) S1: *Doesn't ARG-Growth have high management fees?*
(48c) S1: *Isn't ARG-Growth part of the Horizon group?*

While (48a) and (48b) seem to be expressing doubt, (48c) is merely seeking further information about ARG-Growth. The reason for this difference in interpretation is that evidence from world knowledge (the stereotypical beliefs) suggests, in the case of (48a) and (48b), that S1 believes that the proposition he is contending is true implies that one of the three open propositions is false, namely the proposition that he is doubting. No such evidence exists in the case of (48c). Let us consider (48a) and (48b) in more detail. Since it is stereotypically believed that owning stock in companies that have filed for bankruptcy implies that one has bad investments, utterance (48a) should be interpreted as expressing doubt at the proposition conveyed by utterance (47), namely that ARG-Growth does not have any bad investments. On the other hand, since

it is stereotypically believed that mutual funds with high management fees are not good purchases, utterance (48b) should be interpreted as expressing doubt at the proposition conveyed by utterance (44), namely that ARG-Growth is a good purchase. Note that in this case, the expression of doubt is not aimed at the most recent proposition posited by S2 but instead reverts to expressing doubt at an earlier proposition — in this case, one that utterance (45) had previously doubted. Thus world knowledge in the form of stereotypical beliefs provides evidence both for recognizing an utterance as an expression of doubt and for identifying the proposition P_{doubt} that is being doubted.

In addition to world knowledge, linguistic knowledge in the form of cue words also provides evidence for the third requisite belief, namely that the speaker believes that P_i is in conflict with some proposition P_{doubt} — but it says nothing about how to instantiate P_{doubt}. If there is also world knowledge (as described above) that suggests the proposition P_{doubt} that P_i calls into question, then it determines P_{doubt}. Otherwise, the most salient proposition that meets the salience conditions is used to instantiate P_{doubt}. Before we give an example, let us discuss the salience conditions.

Contextual knowledge captures the salience of individual propositions and thus indicates whether the salience conditions noted in Section 5.1.1 hold. Contextual knowledge includes a stack of propositions ordered according to salience or degree of prominence or focus[Lambert and Carberry, 1992, Carberry, 1990].[2] Thus contextual knowledge provides evidence for two of the salience conditions, namely that the doubting agent has made P_i salient (in our case, via the current utterance) and that P_{doubt} is salient in the dialogue.

Our system also includes in contextual knowledge an indication of which propositions have been accepted by the doubting agent. Acceptance can be explicit, such as by saying *"I accept that"*, but implicit acceptance is much more common. In a collaborative dialogue, if a listener does not believe a communicated proposition, he must convey this disagreement as soon as possible[Walker and Whittaker, 1990]; otherwise the speaker of the proposition is justified in believing that the posited proposition has been accepted. This is similar to Clark and Schaefer's contention[Clark and Schaefer, 1989] that by passing up the opportunity to ask for a repair, a listener implicitly conveys that he has understood an utterance. Our system models implicit acceptance as passing up the opportunity to challenge a proposition by instead addressing a less

[2] Our system actually recognizes a wider variety of discourse acts than just expressions of doubt and employs a discourse model that stacks the discourse acts with their constituent propositions. Discourse acts can contribute to accomplishing other discourse acts; when a new discourse act A_i is initiated that contributes to a discourse act A_j on the stack, all discourse acts above A_j are popped from the stack and A_i is pushed onto the stack (along with possibly a chain of discourse acts that lead to A_i being part of A_j). A full description of the discourse model is beyond the scope of this paper and can be found in [Lambert and Carberry, 1991]. Salience of a proposition is equated with salience of the discourse act that references it.

salient proposition in the dialogue. Thus contextual knowledge also provides evidence for which propositions have not been accepted and therefore remain open for rejection.

Let us now consider an expression of doubt that relies on the cue word *But*. Suppose that S1 had continued the dialogue in Figure 2 with

(48d) S1: *But isn't ARG-Growth part of the Horizon Group?*

Stereotypical beliefs do not provide evidence that ARG-Growth being part of the Horizon group conflicts with any posited proposition in the dialogue. However, the presence of the cue word *But* at the start of the utterance *does* provide such evidence, but it does not help in determining the object of the doubt. Thus the most salient open proposition (the most salient proposition that has not been accepted by S1) is used to instantiate P_{doubt}. Consequently, utterance (48d) is interpreted as expressing doubt at the proposition that ARG-Growth does not have bad investments, but S2 has no indication of why S1 believes that being part of the Horizon Group suggests that ARG-Growth has bad investments — perhaps S1 believes that managers of Horizon Group funds typically make bad investment decisions.

Our methodology for recognizing expressions of doubt uses linguistic, contextual, and world knowledge, as described above, to determine that the necessary salience conditions are satisfied and that there is evidence for the three requisite beliefs listed in Section 5.1.1. Section 5.3 discusses implementation and evaluation of the work.

5.2 Realizing Expressions of Doubt

Besides recognizing doubt expressed by another agent, a collaborative artificial agent must be able to express its own doubt at claims made by the other agent, especially in situations where it has incomplete or uncertain information. For example, an artificial agent's knowledge base might not be up-to-date due to rapidly changing data and only periodic updating — a knowledge base of mutual fund information might only be updated at the end of each day. In addition, information might be added to the knowledge base from sources of different reliability. Thus an agent might have varying degrees of confidence regarding the propositions in its knowledge base. Furthermore, a communicated proposition P_{doubt} might conflict with a proposition P_i in a knowledge base, and both propositions might in fact be true due to some exceptional condition. Thus as artificial agents mature into full collaborative partners, they must be provided with the

ability to respond in a natural and effective manner that will help resolve discrepancies in the beliefs of the participants.

There is a significant difference between recognition and realization. In the case of recognition, the system is presented with evidence in the form of an utterance, the current context, presumed speaker beliefs, etc. and must use this evidence to discern the speaker's intentions. Realization, on the other hand, must incorporate into the communication medium sufficient evidence for the attitude of the agent to be recognized. Too much evidence can make the dialogue appear unnatural, and too little evidence will result in either an ambiguous or a misinterpreted utterance. Moreover, in the case of realization, there may be several alternative ways of conveying the same attitude. Not only might the best choice depend on other factors such as the personality of the agent, but these factors might affect the realization algorithm. For example, an aggressive agent might precede every utterance with a cue word of *But*, and an insecure agent might use an utterance form that conveys a less certain belief in the conflicting evidence than he actually holds. Our research thus far has not taken personality into account. It has focused on realizing an expression of doubt in a form that effectively conveys the attitude of an agent and the agents underlying beliefs that motivate this attitude, so that the collaborative partner can respond effectively to the agent's doubt. In the future we will extend our research to take other factors such as personality into account and will investigate how these factors affect the appropriate realization.

5.2.1 Requisite Content of an Expression of Doubt

An agent expresses doubt at a proposition P_{doubt} not because he will *never* believe P_{doubt} but because he is uncertain about accepting it — perhaps he has conflicting beliefs that must be resolved before the proposition can be accepted or perhaps he had an expectation of a response different from P_{doubt}. This paper will focus on situations where the agent believes a proposition P_i that conflicts with a communicated proposition P_{doubt}; we will briefly discuss violated expectations in Section 6.

In order for the agents to resolve their relevant disparate beliefs and continue the collaboration, a collaborative partner must know:

1. that the other agent has doubt — in the absence of an objection to a posited proposition, a conversational partner will assume implicit acceptance
2. the proposition P_{doubt} that is being doubted

3. the conflicting beliefs that are blocking acceptance

4. some estimation of the relative strength of these conflicting beliefs.

As argued in [Chu-Carroll and Carberry, 1998], knowledge about the doubting agent's conflicting beliefs and how strongly they are held will help the doubted agent construct a response that most effectively resolves the agents' disparate beliefs and allows the collaborative problem-solving to continue.

Although an agent could conceivably satisfy the above criteria by saying *I doubt X because I have belief Y which I hold to degree D1 and Y implies* $\neg X$ *to degree D2*, this is not the usual form in which expressions of doubt are realized. Why is this? Grice's maxims[Grice, 1975] state that an agent should be as informative as needed, without including extraneous or irrelevant information. Thus, in formulating an expression of doubt, the speaker must consider how much his conversational partner needs to know in order to collaborate in resolving the doubt and how much he can be expected to infer without being explicitly told. In addition, Clark[Clark, 1996] argues that participants in discourse select utterances that express their communicative intent efficiently. Generation is a problem of choices, and the choices made by the speaker, as well as the choices that are discarded, say something about the speaker's intent. Since *efficiency of expression* is the expected natural form of discourse, a hearer is likely to derive unintended implications from significantly less economical realizations. Thus we must identify how to provide the above four requisite pieces of information to the agent whose communicated proposition is being doubted, while adhering to Grice's maxims and the efficiency of expression noted by Clark.

5.2.2 Our Approach

Our approach to realizing expressions of doubt draws on the work of Vander Linden and Di Eugenio[Vander Linden and DiEugenio, 1996] on negative imperatives. They posited a relationship between the form in which a negative imperative should be expressed (*Don't X*, *Never X*, or *Take care not to X*) and features of the action X's relationship to the reader in terms of attention, awareness, and safety. In their work, they used machine learning to identify correlations between the two. Our work differs from theirs in that we must deal with an agent's beliefs motivating his doubt and we consider a wider range of different realizations. In addition, Vander Linden and DiEugenio had a set of instructional texts that constituted a *gold standard* of examples to which machine learning could be applied. Many of the expressions of doubt in our corpus are non-ideal, because they were not recognized as doubt or because information that was not included in the

utterance, but could have been, was ultimately needed to resolve the doubt. Thus it was not possible to use machine learning, and our rules are based on our examination of naturally occurring dialogues as well as our judgements about which of these naturally occurring forms was more or less successful.

Our realization algorithm assumes that the agent has decided to express doubt at a proposition P_{doubt} posited by the other agent, that P_{doubt} is currently the most salient proposition in the dialogue, and that the doubting agent's knowledge base includes a proposition P_i that conflicts with P_{doubt}.[3] Our rules use linguistic knowledge about the impact of the surface form of the utterance and the impact of cue words, contextual knowledge consisting of the common ground[Clark, 1996] of the dialogue participants, and world knowledge consisting of the beliefs held by the doubting agent, the strength of those beliefs, and stereotypical beliefs that can be ascribed to the dialogue participants. We assume appropriate mechanisms for detecting conflict, for determining when to initiate a negotiation subdialogue by expressing doubt, and for determining the conflicting evidence that should be used in a response[Chu-Carroll and Carberry, 1998]. We also assume an appropriate belief revision mechanism (such as a Bayesian reasoning system[Zukerman et al., 1998], or a belief system such as that described in [Galliers, 1992, Logan et al., 1994] in which endorsements are used to evaluate the relative support for a proposition, or a belief logic such as that presented in [Driankov, 1988, Bonarini et al., 1990] in which belief/disbelief pairs capture how strongly a proposition is believed), and we assume that the beliefs can be classified according to the strength of their support into the following categories:

- *First-hand*: the belief is a personal preference or something directly experienced and essentially certain
- *Very-strong*: the belief is supported by testimony from an expert or generally reliable source
- *Strong*: the belief is strongly held
- *Weak*: the belief has only weak support

The classification of a proposition P_i into one of the above categories captures the strength of the evidence supporting it, and the classification of an implication $P_i \rightarrow \neg P_{doubt}$ captures the strength of evidence supporting the belief that P_i being true implies that P_{doubt} is not.[4] We also assume that the doubting agent has

[3] Currently we assume that there is only one conflicting proposition P_i. Future work will investigate both how to handle multiple conflicting propositions and how to express doubt at a less salient proposition in the dialogue.

[4] For implications we assume an additional category which we refer to as *Logical* to capture instances where a proposition P_{doubt} is a logical consequence of a proposition P_i.

world knowledge consisting of a set of stereotypical beliefs that can be ascribed to the dialogue participants and contextual knowledge that includes a set of beliefs that are part of their common ground[Clark, 1996].

5.2.3 Rules for Realizing Expressions of Doubt

A surface-negative question conveys uncertain belief in a proposition P_i. If the hearer recognizes that the speaker believes that $P_i \rightarrow \neg P_{doubt}$, then the hearer will recognize the conflict between the speaker's belief in P_i and the proposition P_{doubt}, and will thus recognize that the speaker is expressing doubt at P_{doubt} by contending that P_i is true. Therefore, a surface-negative question is appropriate if the speaker believes P_i, that $P_i \rightarrow \neg P_{doubt}$, that the hearer will recognize the implication, and that P_i is at least as questionable as $P_i \rightarrow \neg P_{doubt}$ (since the surface-negative question draws P_i into focus and thus invites the doubted agent to address it, thereby suggesting that refuting P_i is the easiest way to get P_{doubt} accepted). This leads to our first rule:

Rule-R1: Surface-Negative Question

IF the agent has some belief in P_i, the endorsement of the agent's belief in P_i is at most *strong*, and the strength of its belief in $P_i \rightarrow \neg P_{doubt}$ is at least as strong as its belief in P_i, **THEN** use a surface-negative question that queries the truth of P_i.

Since there are other forms that convey stronger belief in P_i, our rules select the surface-negative question only if the agent's belief in P_i is less than *very-strong*.

Suppose that S1 has a weak belief that the stock Net-Tel has a number of patents that are being challenged, and a strong belief that companies whose patents are being challenged are not generally highly rated. Then Rule-R1 would lead to utterance (50) in the following dialogue excerpt (with one caveat noted below):

(48) S1: *What telecommunications stock has a strong buy rating from a majority of your research analysts?*
(49) S2: *Net-Tel has a strong buy rating.*
(50) S3: *Aren't Net-Tel's patents being challenged?*

However, Rule-R1 does not stipulate that the implication $P_i \rightarrow \neg P_{doubt}$ be part of the beliefs held by the listener. If the listener does not believe that P_i and P_{doubt} are generally in conflict, then, without additional cues, the listener will not be able to recognize that the speaker's utterance is expressing doubt. Rule-R2 addresses such situations by inserting the clue word *But*.

Rule-R2: Clue word *But*

IF $P_i \rightarrow \neg P_{doubt}$ is not part of the set of stereotypical beliefs that can be ascribed to the dialogue participants, then initiate the expression of doubt with the cue word *But*.

Suppose, for example, that S1 is a financial agent that has been asked to secure a mortgage for a client. Suppose further that S1 has the strong belief that corporate banks do not supply low-income mortgages, but does not know whether this belief is shared with the bank information agent (S2) that he is getting information from. Then Rule-R2 would lead to utterance (53) in the following dialogue segment:

(51) S1: *What nearby bank gives low income mortgages?*
(52) S2: *Alma Trust gives low income mortgages.*
(53) S1: *But isn't Alma Trust a corporate bank?*

Rule-R2 is not limited to just surface-negative questions, but is applicable whenever the agent suspects that the implication he wants to convey might not be recognized.

Propositions embedded in belief clauses (such as *I thought that* P_i) appear to convey stronger belief in P_i than do those embedded in surface-negative questions. In an informal survey, graduate students were given dialogues which ended with several alternative forms of expressions of doubt. They were asked to rate the strength of belief in the underlying proposition conveyed by each form. For the majority of instances, belief clauses with embedded propositions were judged to convey stronger beliefs than did surface-negative questions. Thus since embedding the conflicting evidence in a belief clause conveys stronger belief than a surface-negative question but still invites the doubted agent to address the truth of P_i, it is appropriate when the agent's belief in the conflicting evidence is very strong but the agent's belief in the implication $P_i \rightarrow \neg P_{doubt}$ is at least as strong as its belief in P_i. This leads to Rule-R3.

Rule-R3: Belief-Clause

IF the agent's belief in P_i is very strong and the agent's belief in the implication $P_i \rightarrow \neg P_{doubt}$ is at least very strong, **THEN** use a statement that embeds P_i in a belief clause of the form *I thought that* P_i.

For example, suppose that S1 has the very strong beliefs that Net-Tel had its bonds downgraded and that companies with downgraded bonds are not good buys. Then Rule-R3 would lead to utterance (56) in the following dialogue excerpt:

(54) S1: *What telecommunications stock has a strong buy rating from a majority of your research analysts?*
(55) S2: *Net-Tel has a strong buy rating.*
(56) S1: *I thought that Net-Tel had its bonds downgraded.*

Alternatively, if the agent believes P_i because it has been conveyed by a reliable or expert source, the agent might choose to cite the source of the belief. This leads to Rule-R4.

Rule-R4: Source-Clause

IF the agent's very strong belief in P_i is the result of testimony from an expert or reliable source X and the agent's belief in the implication $P_i \rightarrow \neg P_{doubt}$ is at least very strong, **THEN** use a statement that embeds P_i in a source clause of the form *X said that* P_i.

Thus if S1 had read about the downgrading of Net-Tel's bonds in the Wall Street Journal, Rule-R4 would produce (56a) as an alternative to (56) in the dialogue segment shown above:

(56a) S1: *Yesterday's Wall Street Journal reported that Net-Tel had its bonds downgraded.*

Rule-R3 and Rule-R4 both convey very strong beliefs by the doubting agent. We have considered distinguishing between the two rules on the basis of whether the agent positing P_{doubt} is itself an expert, but further research is needed to determine whether this distinction is appropriate.

An utterance of the form *Even though* P_i? (i.e., a statement posed as a question preceded by the cue phrase *even though*) not only conveys a very strong belief in P_i, but seems to suggest that the belief is so strong that the agent who posited P_{doubt} will have more success in refuting the implication $P_i \rightarrow \neg P_{doubt}$ which is not quite as strongly held. This leads to Rule-R5:

Rule-R5: Statement-as-Question

IF the agent's belief in P_i is at least very strong or P_i is part of the common ground of the two agents, and the agent's belief in the implication $P_i \rightarrow \neg P_{doubt}$ is strong, **THEN** use a statement as question of P_i preceded by the cue phrase *Even though*.

For example, suppose that S1 is again a financial consultant who is interacting with a bank information agent regarding a client. Suppose further that S1 knows that his client can only invest $2000, and has a strong belief that high interest rates are available only for larger investments. Then Rule-R5 would lead to utterance (58) below:

(57) S2: *State Savings Bank will give a very high interest rate.*
(58) S1: *Even though he only has $2000 to invest?*

A simple declaration of P_i conveys an extremely strong, essentially certain, belief in P_i; it also tends to convey that the speaker does not think that the hearer will be able to successfully defend P_{doubt}. Thus such a form is appropriate if the speaker has first-hand knowledge of P_i and a very strong belief that P_i implies P_{doubt}. This leads to Rule-R6, which produces utterances that are very close to outright rejections:

Rule-R6: Declare

IF the agent's belief in P_i is first-hand knowledge and the agent's belief in $P_i \rightarrow \neg P_{doubt}$ is at least very strong, **THEN** use a simple declaration of P_i.

For example, suppose that S1 is the financial agent for a large non-profit agency and is interacting with the Internal Revenue Service (the tax authority in the United States). Suppose further that S1 has a very strong belief that non-profit agencies do not have to file quarterly tax returns. Then Rule-R6 will produce utterance (60) below:

(59) S2: *You need to file a quarterly tax return.*
(60) S1: *We're a non-profit organization.*

Lastly, we must consider how to respond when the agent has at least a strong belief in P_i but its belief in the implication $P_i \rightarrow P_{doubt}$ is weak. In this case, the utterance should invite the other agent to attempt to refute the implication since the doubting agent's belief in the implication is much weaker than its belief in P_i. Thus an appropriate utterance is one that specifically queries the possibility of P_{doubt} while noting belief in P_i. This leads to Rule-R7:

Rule-R7: Query-Possibility

IF the agent's belief in the implication $P_i \rightarrow \neg P_{doubt}$ is weak but its belief in P_i is at least strong,

THEN use a simple interrogative about whether P_{doubt} is possible when P_i is true.

In our corpus, agents often use a question like *Can I join the IRA when I'm 65?* to emphasize the implication. Although this form might be interpreted as a simple request for information, this is not problematic since the agent has so little belief in the implication $P_i \rightarrow \neg P_{doubt}$ that he really wants to verify that it does not hold.

5.3 Evaluation

Our algorithm for recognizing a subclass of expressions of doubt (namely where a speaker is expressing doubt at a proposition P_{doubt} by querying some other proposition P_i that he believes is true) has been implemented as part of a larger system for modeling negotiation subdialogues and has been tested in a University advisement domain. The system is given the semantic representation of a new utterance, along with contextual knowledge in the form of a discourse model of the preceding dialogue[Lambert and Carberry, 1991] and world knowledge in the form of stereotypical beliefs that can be ascribed to the dialogue participants. Linguistic knowledge about discourse acts (such as that an agent uttering a surface-negative question has some degree of belief in the proposition contained in the question) is encoded in the recipes for performing the discourse acts; other linguistic knowledge, such as that the cue word *but* provides evidence for an expression of doubt, is encoded in the system. Further discussion of discourse acts and the dialogue model is beyond the scope of this paper and can be found in [Lambert and Carberry, 1991]. The system was evaluated by comparing the system's interpretation of utterances with those of eight human subjects. The subjects were given a set of stereotypical beliefs for the domain, such as that faculty on sabbatical do not normally teach courses. The subjects were presented with a set of dialogues and asked to analyze several utterances from each dialogue. For each utterance, the subjects were given a suggested interpretation and asked whether the suggested interpretation was reasonable and whether they could identify a better interpretation.

For 15 of 20 utterances, the subjects unanimously believed that the system's interpretation was best. All of the instances where some of the subjects either disagreed with the system's interpretation or suggested what they felt was a better interpretation involved surface-negative questions where a clue word was not present and where the stereotypical domain knowledge did not provide a conflict with the queried proposition P_i. These instances could be divided into two categories, as described below:

- One error category involved the use of world knowledge that was not encoded in the system. In two of five instances, some subjects went outside the default beliefs encoded in the system and used their own experience to propose a belief that might suggest a conflict. While these beliefs were not typical and thus should not be represented as stereotypical beliefs, they did represent beliefs that might be encoded in an individual user model and used along with stereotypical beliefs to recognize expressions of doubt. Nonetheless, the subjects were split on how these two utterances should be

interpreted, agreeing with the system slightly more than half the time. In one other surface-negative question that was not interpreted by the system as an expression of doubt, one subject accepted the system's interpretation as reasonable but argued that it would be better to interpret the utterance as an expression of doubt. This subject attributed to the speaker a feature that was neither evident from the dialogue nor stereotypically true. The other subjects agreed with the system that the utterance was not an expression of doubt.

- The other error category involved a more fine-grained interpretation of utterances than was provided by our system. The system interpreted surface-negative questions that did not express doubt as *seeking information about whether the queried proposition was true*; when the subjects did not interpret the utterance as expressing doubt, five of them contended that *seeking verification of the queried proposition* would be a better interpretation than *seeking information*. Nonetheless, they agreed with the system that the utterance was not expressing doubt.

We also performed a preliminary evaluation of our methodology for realizing expressions of doubt. In Section 5.2.1 we argued that an expression of doubt must allow the hearer to determine 1) that the agent is expressing doubt, 2) the proposition that is the object of the doubt, 3) the conflicting evidence that motivates the doubt, and 4) some estimate of the relative strength of the conflicting beliefs. In order to determine how well the forms we generate accomplish this, we had six subjects analyze 14 dialogue segments, each of which ended with a highlighted utterance that in 10 instances was intended to be an expression of doubt. The expressions of doubt in the questionnaire corresponded to the forms that our rules would generate[5] and also included elliptical fragments that are part of our further work on realizing expressions of doubt.

To determine whether the forms produced by our realization rules would satisfy criteria 1) and 2) above, the subjects were asked to determine whether the highlighted utterance was expressing doubt and, if so, to identify the utterance that was being doubted. For the ten highlighted utterances that were intended to be expressions of doubt, the subjects unanimously recognized them as expressing doubt. For the four highlighted utterances that were not intended to be expressions of doubt, one utterance was misjudged by two of the subjects. For every expression of doubt, the subjects all identified the intended doubted proposition P_{doubt} in the dialogue segment. Thus we concluded that our realization rules would specify utterance forms that effectively convey doubt and the proposition that is the object of the doubt.

[5] Except for Rule-R5 which was added later.

Subjects were also asked to rate the strength of the doubt. They almost always rated instances in which doubt was expressed by contending P_i as conveying stronger doubt than those instances that did not supply a conflicting proposition. (An example of the latter is an elliptical fragment that repeats some aspect of P_{doubt} such as *$2000 every year?* following the utterance *You must pay $2000 every year.*) Although strength of doubt doesn't correspond directly to one of the features used in our algorithm, this is consistent with the idea that supporting evidence may not be attributed if it is not expressed.

To determine whether the forms produced by our realization rules would satisfy criteria 3) and 4) above, the subjects were asked to select from a list of propositions (or write in additional propositions) those that motivated the expression of doubt and were communicated by it. They were also asked to rate the strength of the doubting agent's belief in the selected propositions (both P_i and $P_i \rightarrow \neg P_{doubt}$). Only five times out of 60 instances (ten expressions of doubt times six subjects) did the subjects specify a belief that our algorithm did not use in selecting the form of the utterance, and only three times did the subjects fail to recognize a belief that was used and intended to be conveyed. We analyzed the strengths of beliefs identified by the subjects to determine how they correlated with the strengths that would have produced this form by our rules. Since subjects varied in the ranges that they used in rating the strengths of the beliefs, we looked at the scores relative to each subject's ratings of the other beliefs. Most of the strength ratings were consistent with our realization rules. The most frequent inconsistency resulted from situations where our rules would have generated a form based on slightly different strengths of beliefs for P_i and the implication but our subjects rated them equivalent strengths. Nonetheless the strength ratings were close. The only significant inconsistency with our rules was one utterance in which doubt was expressed via an utterance in which the speaker's husband was cited as the reliable source of the conflicting evidence P_i. In this instance, the agent's belief in her husband's expertise was estimated by the subjects as relatively weak whereas it was coded as very strong in our belief model. It appears that the subjects were applying their own views about a spouse's likely expertise in the domain in making their judgements rather than solely basing their judgements on the form of the utterance. These experiments led us to conclude that our realization rules would specify utterance forms that effectively convey the conflicting evidence P_i, the implication $P_i \rightarrow \neg P_{doubt}$, and an estimate of the strength of the agent's belief in each. The resulting utterances would then enable a collaborative partner to formulate a response that will address the doubt and help the agents resolve their disparate beliefs[Chu-Carroll and Carberry, 1998].

6 Summary and Future Work

This paper has argued that the concept of attitude is different from those of personality and emotion but important in the design of artificial agents. In particular, it contended that artificial agents must be able both 1) to recognize an attitude of doubt and its underlying motivation and 2) to convey an attitude of doubt when the agent does not accept a communicated proposition or proposal. The paper presented our algorithms that use linguistic, contextual, and world knowledge to recognize a subclass of expressions of doubt and to realize utterances that express doubt when evidence conflicts with a proposition posited by another agent.

However, expressions of doubt are also appropriate when an expectation is violated by a response, even though the doubting agent may have little or no support for the expectation. Such violated expectations are often signaled by elliptical fragments, as in the following dialogue segment:

(61) S2: *John Smith gets $600 in Social Security.*
(62) S1: *$600?*

In this example, S1 conveys doubt at the proposition posited by S2 as well as the fact that this doubt stems from the instantiation of the amount as $600. Expressions of doubt resulting from such violated expectations, without supporting evidence, generally take the form of elliptical fragments. Our research is now examining such expressions of doubt.

Prosody and gesture are important features of expressions of doubt that are part of spoken dialogue and face-to-face interactions. For example, consider the following dialogue segment:

S2: *I just invested $100,000 that I inherited.*
S1: *In the stock market?*

S1's utterance in this example could either be an expression of doubt or a request for clarification. Intonational differences in the realization of the utterance or gestural evidence in the form of facial expression are necessary in order to communicate the correct attitude of the speaker. Thus we intend to explore the contribution of prosody and gesture both to our recognition algorithm and to our generation methodology. In addition, we will be exploring other attitudes besides doubt and will be investigating the contribution of personality to our realization algorithm.

References

[Bartneck, 2001] Bartneck, C. 2001. How Convincing is Mr. Data's Smile: Affective Expressions of Machines. *User Modeling and User-Adapted Interaction* **11**(4): 279–295.

[Bonarini et al., 1990] Bonarini, A., E. Cappelletti, and A. Corrao. 1990. Network-based Management of Subjective Judgements: A Proposal Accepting Cyclic Dependencies. Technical Report 90-067, Dipartimento di Elettronica, Politecnico di Milano, Milano, Italy.

[Carberry, 1990] Carberry, S. 1990. *Plan Recognition in Natural Language Dialogue*, ACL-MIT Press Series on Natural Language Processing. Cambridge, Massachusetts: MIT Press.

[Cassell et al., 1994] Cassell, J., C. Pelachaud, N. Badler, M. Steedman, B. Achorn, T. Becket, B. Douville, S. Prevost, and M. Stone. 1994. Animated Conversation: Rule-Based Generation of Facial Expression, Gesture, and Spoken Intonation for Multiple Conversational Agents. In: *Proceedings of SIGGRAPH '94*: 413–420.

[Cassell et al., 2000] Cassell, J., J. Sullivan, S. Prevost, and E. Churchill (eds.). 2000. *Embodied Conversational Agents.* Cambridge, MA: MIT Press.

[Chopra-Khullar and Badler, 2001] Chopra-Khullar, S. and N. Badler. 2001. Where to look? Automating Attending Behaviors of Virtual Human Characters. In: *Autonomous Agents and Multi-agent Systems*: 9–23.

[Chu-Carroll and Carberry, 1998] Chu-Carroll, J. and S. Carberry. 1998. Collaborative Response Generation in Planning Dialogues. *Computational Linguistics* **24**(3): 355–400.

[Chu-Carroll and Carberry, 2000] Chu-Carroll, J. and S. Carberry. 2000. Conflict Resolution in Collaborative Planning Dialogs. *International Journal of Human-Computer Studies* **53**(6): 969–1015.

[Clark, 1996] Clark, H. 1996. *Using Language.* Cambridge, England: Cambridge University Press.

[Clark and Schaefer, 1989] Clark, H. and E. Schaefer. 1989. Contributing to Discourse. *Cognitive Science*: **13**:259–294.

[Cohen, 1987] Cohen, R. 1987. Analyzing the Structure of Argumentative Discourse. *Computational Linguistics* **13**(1-2): 11–24.

[De Rosis, 2001] De Rosis, F. 2001. Preface: Towards Adaptation of Interaction to Affective Factors. *User Modeling and User-Adapted Interaction* **11**(4): 267–278.

[Driankov, 1988] Driankov, D. 1988. Towards a Many-Valued Logic of Quantified Belief. Ph.D. thesis, Linkoping University, Department of Computer and Information Science, Linkoping, Sweden.

[Galliers, 1992] Galliers, J. R. 1992. Autonomous belief revision and communication. In: P. Gardenfors (ed.): *Belief Revision*, Cambridge tracts in theoretical computer science. Cambridge, England: Cambridge University Press.

[Grice, 1975] Grice, H. P. 1975. Logic and Conversation. In: P. Cole and J. L. Morgan (eds.): *Syntax and Semantics III: Speech Acts*: 41–58, New York: Academic Press.

[Grosz and Sidner, 1986] Grosz, B. and C. Sidner. 1986. Attention, Intentions, and the Structure of Discourse. *Computational Linguistics* **12**(3): 175–204.

[Hirshberg and Litman, 1993] Hirshberg, J. and D. Litman. 1993. Empirical Studies on the Disambiguation of Cue Phrases. *Computational Linguistics* **19**(3): 501–530.

[Knott and Mellish, 1996] Knott, A. and C. Mellish. 1996. A Feature-Based Account of the Relations Signalled by Sentence and Clause Connectives. *Language and Speech* **39**(2-3): 143–183.

[Lambert and Carberry, 1991] Lambert, L. and S. Carberry. 1991. A Tripartite Plan-based Model of Dialogue. In: *Proceedings of the 29th Annual Meeting of the ACL*: 47–54.

[Lambert and Carberry, 1992] Lambert, L. and S. Carberry. 1992. Modeling Negotiation Subdialogues. In: *Proceedings of the 30th Annual Meeting of the ACL*: 193–200.

[Lester and Stone, 1997] Lester, J. C. and B. Stone. 1997. Increasing Believability in Animated Pedagogical Agents. In: *Proceedings of the First International Conference on Autonomous Agents*: 16–21.

[Logan et al., 1994] Logan, B., S. Reece, A. Cawsey, J. Galliers, and K. Sparck Jones. 1994. Belief Revision and Dialogue Management in Information Retrieval. Technical report, University of Cambridge Computer Laboratory.

[Mozziconacci, 2001] Mozziconacci, S. 2001. Modeling Emotion and Attitude in Speech by Means of Perceptually Based Parameter Values. *User Modeling and User-Adapted Interaction* **11**(4): 297–326.

[Muller et al., 2001] Muller, C., B. Groszmann-Hutter, A. Jameson, R. Rummer, and F. Wittig. 2001. Recognizing Time Pressure and Cognitive Load on the Basis of Speech: An Experimental Study. In: *Proceedings of the 8th International Conference on User Modeling*: 24–33.

[Nass et al., 1996] Nass, C., B. J. Fogg, and Y. Moon. 1996. Can Computers be Teammates?. *International Journal of Human-Computer Studies* **45**(6): 669–678.

[Nass et al., 1995] Nass, C., Y. Moon, B. J. Fogg, B. Reeves, and C. Dryer. 1995. Can Computer Personalities be Human Personalities?. *International Journal of Human-Computer Studies* **43**(2): 223–239.

[Reichman, 1985] Reichman, R. 1985. *Getting Computers to Talk Like You and Me.* Cambridge, Massachusetts: MIT Press.

[Transcripts, 1982] Transcripts, H. G. 1982. Transcripts derived from tapes of the radio talk show Harry Gross: Speaking of Your Money. Provided by the Dept. of Computer Science at the University of Pennsylvania.

[Transcripts, 1992] Transcripts, S. 1992. Transcripts derived from audiotape conversations made at SRI International, Menlo Park, California. Prepared by Jacqueline Kowtko under the direction of Patti Price.

[Vander Linden and DiEugenio, 1996] Vander Linden, K. and B. DiEugenio. 1996. A Corpus Study of Negative Imperatives in Natural Language Instructions. In: *Proceedings of the 16th International Conference on Computational Linguistics.*

[Walker and Whittaker, 1990] Walker, M. and S. Whittaker. 1990. Mixed Initiative in Dialogue: An Investigation into Discourse Segmentation. In: *Proceedings of the 28th Annual Meeting of the Association for Computational Linguistics*: 70–78.

[Young and Cohen, 1991] Young, M. A. and R. Cohen. 1991. Determining Intended Evidence Relations in Natural Language Arguments. *Computational Intelligence* **7**: 110–118.

[Zukerman, 2001] Zukerman, I. 2001. An Integrated Approach for Generating Arguments and Rebuttals. In: *Proceedings of the 8th International Conference on User Modeling*: 84–94.

[Zukerman et al., 1998] Zukerman, I., R. McConachy, and K. Korb. 1998. Bayesian Reasoning in an Abductive Mechanism for Argument Generation and Analysis. In: *Proceedings of the Fifteenth National Conference on Artificial Intelligence*: 833–838.

Bankruptcy-Rule: If one holds stock in companies that have filed for bankruptcy, then one has bad investments.
$\forall x\ \forall y$ (Owns-stock(x,y) $\wedge$ Filed-bankruptcy(y) $\rightarrow$ Has-bad-investments(x))

Manage-Fees-Rule: If a mutual fund has high management fees, then it is not a good purchase.
$\forall x$(Mutual-fund(x) $\wedge$ Management-Fees(x,HIGH) $\rightarrow \neg$Good-Purchase(x))

Figure 1: A Set of Stereotypical Beliefs

(43) S1: *Which mutual fund should I purchase?*
(44) S2: *ARG-Growth Fund is a very good purchase.*
(45) S1: *Hasn't the share value of ARG-Growth fallen during each of the last two years?*
(46) S2: *All of the growth mutual funds have fallen.*
(47) *ARG-Growth doesn't have any bad investments.*

Figure 2: A Dialogue with Three Open Propositions

Captions for Figures

Figure 1. A Set of Stereotypical Beliefs

Figure 2. A Dialogue with Three Open Propositions