# RECOGNIZING AND CONVEYING ATTITUDE AND ITS UNDERLYING MOTIVATION

**Sandra Carberry** and **Leah Schroeder**
Department of Computer and Information Sciences
University of Delaware
Newark, Delaware 19716
carberry@dewey.udel.edu, schroede@cis.udel.edu

**Abstract** [1]

This paper contends that attitude, emotion, and personality are very different concepts that must be treated differently by an artificial agent. It discusses the importance both of 1) recognizing the attitude of a conversational partner toward a communicated proposition and 2) conveying one's own attitude toward different propositions, in order to enable appropriate responses to be constructed. The paper then focuses on one kind of attitude, doubt regarding a proposition, and presents examples of the variety of ways that this doubt can be expressed in a naturally occurring dialogue. It then summarizes our implemented algorithm for recognizing expressions of doubt, and describes our current research on an algorithm for effectively conveying doubt in a collaborative dialogue.

## Introduction

Previous research on human-computer interfaces has been concerned primarily with conveying 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 address the notions of attitude, emotion, or personality of the computer and human agents.

Our current research is concerned with realizing natural language utterances that effectively convey the attitude of an artificial agent. This paper first argues that emotion, personality, and attitude are very different concepts that must be treated differently by an artificial agent. It then argues that in a collaborative dialogue, communicating attitude requires that the agent also convey (explicitly or implicitly) the beliefs that underly this attitude and the strength of these beliefs. The paper reviews our previous work on recognizing attitude and then gives an overview of our recent research on constructing natural language utterances that effectively convey an agent's attitude during a collaborative dialogue. We conclude with a discussion of our future work that will incorporate intonation.

[1]This work was supported by NSF grant #GER-9354869 and #CDA-9703088.

## 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 manisfestation 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.

The system 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 system recognizes these different attitudes can it address the doubt or confusion, resolve it, and thereby continue the dialogue in an effective manner. Similarly, the system 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.

## 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 agents express doubt in naturally occurring dialogue, we analyzed a number of dialogue corpora, ranging from a corpus of financial planning dialogues that were transcribed from radio talk shows[Tra82] to a corpus of simulated travel planning dialogues[Tra92]; we also noted instances of expressions of doubt appearing in the dialogues of popular novels. Our analysis showed that doubt is expressed in a wide variety of ways, including:

- by contending or suggesting that some conflicting proposition might be true. Such expressions of doubt often take the form of a surface negative question, as in the following example[2] where S2's last utterance expresses doubt at S1's recommendation.

  S1: *I would like to see that into an individual retirement account rollover in a mutual fund group.*
  S2: *At my age?*
  S1: *Yes.*
  S2: *Uh, yeah but isn't there any risk?*

- 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 the dialogue excerpt shown above.

[2]This example is taken from the Harry Gross financial planning dialogues[Tra82].

- By drawing attention to an expectation that is violated by the doubted proposition, such as *"You're kidding, what happened to the seventy-eight dollar fares or those sort of things?"*

- by repeating some facet of the doubted proposition, sometimes with extra information, as in the following two examples.

  S1: *OK, what I'd like you to do is to lend him the 20 thousand.*
  S2: *20 thousand?*

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

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 are not the dominant form for conveying an attitude of doubt.

## Our Previous Work on Recognizing an Attitude of Doubt

In previous work, we developed an algorithm for recognizing when an agent was expressing doubt at a proposition $P_{doubt}$ by contending that some other proposition $P_i$ is true[CL99]. For example, consider the following dialogue segment:

S1: *Who is teaching CS360?*
S2: *Dr. Smith is teaching CS360?*
S1: *Isn't Dr. Smith on sabbatical?*

In the last utterance, S1 (by contending that Dr. Smith is on sabbatical) expresses an attitude of doubt toward the proposition that Dr. Smith is teaching CS360. This expression of doubt relies on the presumed mutual belief that faculty on sabbatical do not teach courses.

Our algorithm uses information from several knowledge sources. Linguistic knowledge includes the surface form of the utterance and the presence of cue words. For example, a surface negative question of the form *Isn't* $P_i$*?* can be used to convey doubt at a proposition $P_{doubt}$ whereas a simple interrogative such as *Is* $P_i$*?* cannot, unless the doubt is marked by intonation. Contextual knowledge includes beliefs derived from the discourse, the structure of the discourse, and the salience of different propositions. World knowledge includes stereotypical beliefs that are presumed to be held by speaker and hearer and which indicate that the proposition $P_i$ the speaker is contending is true conflicts with the doubted proposition $P_{doubt}$.

In order to recognize an utterance as an expression of doubt in a collaborative dialogue, our algorithm requires evidence that (1) the speaker believes that the hearer has some belief in $P_{doubt}$ since it is pointless to express doubt at a proposition about which there is no disagreement; (2) the speaker believes that if $P_i$ is true, then $P_{doubt}$ is false; and (3) the speaker has some belief in the proposition $P_i$. Evidence for (1) is provided by the dialogue context in which the other agent conveys his belief in $P_{doubt}$. Evidence for (2) is provided either by world knowledge about stereotypical beliefs suggesting that $P_i$ and $P_{doubt}$ are in conflict or by the clue word *but* which generically suggests an expression of doubt. Evidence for (3) is provided by the surface form of the speaker's utterance, namely a surface negative question. When there is equivalent evidence for interpreting an utterance as expressing doubt at one of several propositions, contextual knowledge in the form of the relative salience of the different propositions is used to arbitrate among the interpretations.

Our implemented algorithm resulted in a system that could recognize a variety of discourse acts and could model embedded negotiation subdialogues[CL99].

## Effectively Communicating Doubt

### Requisite Content of an Expression of Doubt

Our current work is focused on algorithms for realizing appropriate expressions of doubt in a collaborative dialogue. 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 — ie., he has conflicting beliefs that must be resolved before the proposition can be accepted. In order for the agents to resolve their relevant disparate beliefs and continue the collaboration, the doubted agent must know (1) that the other agent has doubt — in the absence of an objection to a posited proposition, the speaker will assume implicit acceptance[CL99], (2) the proposition $P_{doubt}$ that is being doubted, (3) the conflicting beliefs that are blocking acceptance and the strength of these beliefs — as argued in [CCC98], if the doubted agent decides to retain his original beliefs, this information helps him to construct an argument that will be maximally effective and efficient in his attempt to convince the doubting agent.

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* $\neg P_{doubt}$ *to degree D2*, this is not the usual form in which expressions of doubt are realized. Why is this? Grice's maxims[Gri75] 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 the doubted agent needs to know in order to collaborate in resolving the doubt and how much the doubted agent can be expected to infer without being explicitly told. In addition, Clark[Cla96] argues that participants in discourse select utterances that express their communicative intent efficiently, often in sentence fragments. 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 the above criteria can be satisfied while adhering to Grice's maxims and the efficiency of expression noted by Clark.

### Related Work

Our approach to realizing expressions of doubt draws on the work of Vander Linden and Di Eugenio[VLD96] 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 of which of these naturally occurring forms was more or less successful.

**Features Influencing Expressions of Doubt** As argued earlier, an utterance that conveys an attitude of doubt will be most effective if it communicates not only the doubting attitude but also the proposition that is doubted, any support for this doubt, and the strength of this support. In addition, speakers tend to differentiate new supporting information from that which is already part of the common ground and which should already have been considered. These beliefs are often not explicitly stated, but are assumed to be inferrable by the doubted agent based on his knowledge of the preceding dialogue, knowledge of the other agent's beliefs, a model of stereotypical beliefs, linguistic knowledge, and the particular realization of the doubting agent's utterance.

Expectations also play a role in expressions of doubt. In the simplest case, the violated expectation is just that $P_{doubt}$ is false. In other situations, an agent may have an expectation that a proposition will be false if instantiated with some particular subset of its possible instantiations. Responses that conflict with these expectations may provoke an expression of doubt, 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:

S1: *John Smith gets $600 a month in Social Security.*
S2: *$600?*

In this example, S2 conveys her doubt at the proposition posited by S1 as well as the fact that this doubt stems from the instantiation of the amount as $600. In the following, $P_{doubt}$ will refer to the doubted proposition, $P_i$ will refer to a proposition other than $\neg P_{doubt}$ that is the reason for the doubt, and $P_e$ will refer to the violated expectation.

Our algorithm for expressing doubt takes into account several attributes of $P_{doubt}$, $P_i$, and $P_e$. It uses the authority behind the doubted proposition $P_{doubt}$ since whether it was conveyed by an expert impacts the strength of the hearer's doubt[CCC98]. If there is a conflicting proposition $P_i$, the algorithm uses the strength of the agent's belief in $P_i$, the degree to which $P_i$ is believed to imply $\neg P_{doubt}$, and whether $P_i$ is new information in the dialogue or already part of the established common ground. If the violated expectation is not just a general belief that $\neg P_{doubt}$, then the algorithm takes into account whether it is a term-value violation (ie., a failure to accept $P_{doubt}$ due to an expectation that the instantiated value is not one of the instantiations of the term that would make $P_{doubt}$ true) or a constraint violation (ie., a failure to accept $P_{doubt}$ due to an expectation that $P_{doubt}$ will be false under a specified constraint).

**Rules for Realizing Expressions of Doubt** The following are a few of the rules that we use to realize efficient and natural expressions of doubt. Our algorithms presume a model of the hearer's beliefs and a belief revision system such as that presented in [Gal92, LRC+94] where endorsements are used to evaluate the belief in a proposition.

The first two rules illustrate two situations where an expression of doubt is realized differently: as a surface negative question and as an assertion. When used to express an attitude of doubt, surface negative questions are intended to convey uncertain belief in a proposition $P_i$ and a belief that the doubted agent will be able to infer that the doubting agent believes that $P_i \rightarrow \neg P_{doubt}$. This form is appropriate when $P_i$ is considered more questionable than the implication or if $P_i$ and the implication are endorsed equally strongly. In such cases, the surface negative question suggests to the doubted agent that addressing the truth of $P_i$ is the most effective way to resolve the disparity in beliefs.

Rule-C1: **IF** the endorsement of the doubting agent's belief in $P_i$ is at most *default*, and the strength of the implication $P_i \rightarrow \neg P_{doubt}$ is at least as strong as the doubting agent's belief in $P_i$, **THEN** realize the expression of doubt as a surface negative question that queries the truth of $P_i$.

If the doubting agent is quite certain of $P_i$ (such as his age or a preference about his personal wants), then there is little point in the hearer trying to attack the validity of $P_i$. Instead, the doubted agent must either weaken belief in $P_i \rightarrow \neg P_{doubt}$ or concede that $P_{doubt}$ is false. In such cases, an assertion of $P_i$ conveys the doubting agent's very strong belief in $P_i$ and implicitly invites the doubted agent to address the implication.

Rule-C2: **IF** the endorsement of the doubting agent's belief in $P_i$ is as *first-hand knowledge* (ie., personal preference or direct experience) and endorsement of the implication $P_i \rightarrow \neg P_{doubt}$ is stronger than *default*, **THEN** realize the expression of doubt as a simple declaration of $P_i$.

The above two rules presume that the hearer can infer that the speaker believes that $P_i \rightarrow \neg P_{doubt}$. This relies on a user model that contains stereotypical beliefs that the agents mutually hold. But occasionally the doubting agent will have reason to personally believe that $P_i \rightarrow \neg P_{doubt}$ but may not believe that the doubted agent will recognize the implication (because the doubted agent does not have the same beliefs). In such cases, a cue word of *but* explicitly conveys the implication. For example, consider the following dialogue segment:

S1: *What time does CS360 meet?*
S2: *CS360 meets at 5pm on Tuesday and Thursday.*
S1: *But doesn't the ACM chapter meet on Tuesday evening?*

Without the cue word *but*, S1's second utterance may be interpreted as a simple query verifying when the ACM chapter meets; the cue word *but* explicitly signals that S1's utterance is an expression of doubt, though the doubted agent might not believe that there is a conflict between both the ACM chapter and CS360 meeting on Tuesday evening. This leads to the following rule regarding the cue word *but*:

Rule-Cue-1: **IF** the endorsement of the implication $P_i \rightarrow \neg P_{doubt}$ is first-hand (or more generally is not part of the model of beliefs that the doubted agent holds or believes are held by the doubting agent), **THEN** initiate the expression of doubt with the cue word *but*.

Another rule addresses the use of the cue phrase *even though* which occurred in our corpus when the proposition $P_i$ was already part of the common ground established by the preceding dialogue.

A proposition can be doubted without the presence of a conflicting proposition $P_{doubt}$. Several rules address such situations. For example, the doubting agent may fail to accept $P_{doubt}$ due to an expectation that $P_{doubt}$ will be false in situations in which *_constraint* holds. This constraint is not a term in $P_{doubt}$, but the doubting agent believes that the doubted agent intends that the constraint holds. For example, consider the following dialogue segment:

S1: *The most you will get back on your taxes is $400.*
S2: *$400 for the whole year?*

In this example, S2 presumably had the expectation that he would receive more than $400 for the whole year and this expectation is violated. The following rule addresses such cases:

Rule-E1: **IF** there is no $P_i$ as evidence for $\neg P_{doubt}$ yet the doubting agent believes that $P_{doubt}$ is false when *_term* in $P_{doubt}$ is instantiated with *_value* in situations where *_constraint* holds, **THEN** realize the expression of doubt as an elliptical fragment that repeats *_value* followed by *_constraint*.

## Preliminary Evaluation

We performed a preliminary evaluation of our rules by asking six subjects to judge dialogue segments which ended with a response that was in some cases an expression of doubt and in other cases was not. The expressions of doubt in the dialogue segments corresponded to the forms that our rules would generate. All responses that were intended as expressions of doubt were recognized as such. In addition, the subjects were asked to identify the $P_{doubt}$, $P_e$, and $P_i$ when present, estimate the strength of the beliefs, and rate the strength of the doubt. The results corresponded well with the basis for our rules.

## Future Work

Prosody is an important feature of expressions of doubt that are part of spoken dialogue. For example, consider the following dialogue segment:

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

S2'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 are necessary in order to communicate the correct attitude of the speaker. Thus we intend to explore the contribution of prosody to our generation methodology. We will also be extending our work to realize other speaker attitudes besides doubt.

## Summary

This paper has distinguished emotion, personality, and attitude. It argued that a collaborative dialogue system must be able both 1) to recognize an attitude of doubt and its underlying motivation and 2) to convey an attitude of doubt when it does not accept a communicated proposition or proposal. The paper reviewed our previous work on recognizing expressions of doubt and then presented a brief overview of our current research on realizing natural language utterances that communicate an attitude of doubt. In future work, we will explore the use of prosody and will investigate the communication of other attitudes besides doubt.

## References

[CCC98] Jennifer Chu-Carroll and Sandra Carberry. Collaborative response generation in planning dialogues. *Computational Linguistics*, 24(3):355–400, 1998.

[CL99] Sandra Carberry and Lynn Lambert. A process model for recognizing communicative acts and modeling negotiation subdialogues. *Computational Linguistics*, 25(1):1–53, 1999.

[Cla96] Herbert Clark. *Using Language*. Cambridge University Press, 1996.

[Gal92] Julia Rose Galliers. Autonomous belief revision and communication. In P. Gardenfors, editor, *Belief Revision*, Cambridge tracts in theoretical computer science. Cambridge University Press, Cambridge, England, 1992.

[Gri75] H. Paul Grice. Logic and Conversation. In P. Cole and J. L. Morgan, editors, *Syntax and Semantics III: Speech Acts*, pages 41–58, N.Y., 1975. Academic Press.

[LRC+94] Brian Logan, Steven Reece, Allison Cawsey, Julia Galliers, and Karen Sparck Jones. Belief revision and dialogue management in information retrieval. Technical report, University of Cambridge Computer Laboratory, 1994.

[Tra82] Harry Gross Transcripts. Transcripts derived from tapes of the radio talk show harry gross: Speaking of your money, 1982. Provided by the Dept. of Computer Science at the University of Pennsylvania.

[Tra92] SRI Transcripts. Transcripts derived from audiotape conversations made at sri international, menlo park, california, 1992. Prepared by Jacqueline Kowtko under the direction of Patti Price.

[VLD96] Keith Vander Linden and Barbara DiEugenio. A corpus study of negative imperatives in natural language instructions. In *Proceedings of the 16th International Conference on Computational Linguistics*, 1996.