Belief ascription in mixed initiative dialogue

Mark Lee
Department of Computer Science
University of Sheffield
Regent Court, 211 Portobello Street
Sheffield S1 4DP, UK
M.Lee@dcs.shef.ac.uk

Abstract

Real dialogues often feature conflicting beliefs and goals between the participants. In such cases, it is often necessary for one participant to reason about the nature of the conflict and respond appropriately. This paper presents a brief outline of a computational theory of dialogue understanding which can deal with such conflicts. Our theory is capable of distinguishing between cases of pragmatic communication, deception and mistaken belief and using such distinctions to infer the goals of the speaker. We will also discuss some initial work on dialogue control to resolve such conflicts.

1 Introduction

The majority of work in computational pragmatics has been involved in modelling and understanding cooperative, mutual goal-driven dialogues. However, dialogues featuring private goals and motives are common in reality. For example, consider the medical counselling dialogue in (1):

(1). Doctor: What are the symptoms?
Patient: I have a pain in my stomach.
Doctor: Is the pain below your navel?
Patient: Yes, an inch further down and to the right.
Doctor: Your medical record reports you've had an inflamed appendix previously. Is it a similar type of pain?
Patient: Yes. But I was told that this is never serious
Doctor: Often it isn't, but it can become serious if left untreated. How long have you felt pain?

In the above dialogue, the doctor and patient have both mutual and private goals. For example, both wish to discover the cause of the symptoms; but in addition, the doctor would typically have additional goals of patient education; and risk assessment and communication. Similarly, the patient might have private goals such as giving informed consent and, potentially, goals of misdirection and deception which could be detrimental to his or her own health. In such a domain, the doctor must be able to reason about the implicit goals of the patient and be capable of guiding the dialogue by controlling who has the initiative at certain stages.

In addition, agents in a dialogue may hold different beliefs and use different referring expressions to convey certain concepts. For example in (1) the patient uses the term "stomach" to refer to his or her entire abdomen rather than using it as a precise medical term. The doctor must be able to reason about such reference uses and decide whether or not to interrupt the dialogue to correct the patient.

Any practical dialogue understanding system must be able to make the same decisions. This involves being able to distinguish genuine cases of pragmatic meaning such as metaphor from cases of deception and mistaken belief. Such understanding is only possible, if the system can accurately predict and ascribe goals to the speaker which may or may not be explicitly communicated.

In this paper, we will outline some of the work implemented in ViewGen, a dialogue understanding system which is capable of dealing with conflicting belief and goal driven dialogues in such a medical counselling domain.

2 ViewGen

ViewGen [Ballim and Wilks, 1991],[Wilks et al., 1991] is a dialogue understanding system which reasons about the attitudes of other agents in nested belief structures called *environments*. It does this by *ascription* - i.e. assuming that attitudes held in one attitude box can be ascribed to others. There are two main methods of ascription - default ascription and stereotypical ascription. Default ascription applies to common attitudes which ViewGen assumes that any agent will hold and also ascribe to any other agent unless there is contrary evidence. Stereotypical ascription applies to stereotypical attitudes which ViewGen assumes apply to instances of a particular stereotype e.g.: ViewGen ascribes expert medical knowledge to doctors; goals to agents in a particular situations e.g. knowledge elicitation goals to agents involved in information seeking dialogues; or intentions to stages in a dialogue e.g. introduction acts in the introduction of a new topic.

The nested structure of the environments allows ViewGen to deal with the resolution of intensional references. For example, in (1) the patient uses the term "stomach" to refer to the general area of the abdomen. However, there is nothing in the context at that point in the dialogue to suggest that the patient is not using "stomach" in its precise sense. ViewGen is capable of maintaining a stereotype for patients where the concept of stomach is typically used to refer to its common meaning and when interpreting the patient's utterance, this stereotype would be preferred to the precise medical meaning of stomach.

ViewGen plans and understands utterances using is a set of task related speech acts and dialogue control acts. The speech acts are specified as plan operators which are used to plan what is communicated by the System and also used to recognise plans by the user. For example, informing is specified as:

Inform(Speaker,Hearer,Proposition)
Preconditions: goal(Speaker,bel(Hearer,Proposition)
bel(Speaker,Proposition)

where the predicates goal and bel refer to goals and beliefs.

Rather than specify the effects of an act, there are separate ascription rules for the speaker and hearer of any act:

Update on the Speaker's belief set

For every condition C in a dialogue act performed:
ascribe(Speaker, Hearer, bel(C))

Update on the Hearer's belief set

For every condition C in a dialogue act performed:
ascribe(Hearer, Speaker, C)

This account differs from standard accounts of computational speech acts (e.g. [Allen, 1983]) in that the underlying operation of ascription takes over implicitly many of the conditions and effects that are more extensively specified elsewhere. Further discussion of ViewGen's speech act representation and ascription are given in Lee and Wilks [Lee and Wilks, 1996a,Lee and Wilks, 1996b].

ViewGen plans its communicative goals using a partial order clausal link planner [McAllester and Rosenblatt, 1991]. The planner is also used to understand utterances by a process of plan recognition. Given a communicative act performed by the user, ViewGen attempts to generate a plan involving the act which results in one or more ascribable goals which it can then ascribe to the speaker. Ascribable goals are either common goals which can be ascribed by default or stereotypical goals for what is known about the user or context of the utterance. Understanding can, therefore, be seen as a form of ascription where ViewGen has to reason which beliefs, goals and intentions are ascribed to the modelled agent.

However, it is rare for there to be only one goal to be satisfied at any turn of a dialogue. Instead, any utterance choice is usually informed by the multiple satisfaction of both task related and dialogue control based goals. As a planning heuristic, ViewGen selects and plans to satisfy as many communicative goals as possible in any plan.

This heuristic is assumed to be also used by the other (human) dialogue participant. This complicates the use of plan recognition in understanding utterances but allows ViewGen to draw further pragmatic inferences based on its model of the speaker's belief state. In the next two subsections, we will briefly describe how ViewGen understands utterances by ascribing additional goals based on the perceived plan of the speaker.

2 Inferring beliefs and goals based on felicity conditions

As noted above, any utterance is represented is represented by a speech act with attendant felicity conditions which are communicated by ascription. As in (1), often the communicated conditions contradict the beliefs held by the hearer. Such cases complicate plan recognition since recognising plans based on untruths is difficult. In particular, ViewGen has to distinguish between cases of mistaken belief, deception, and genuine cases of intentionally flouting truth conditions to communicate pragmatically [Grice, 1975].

Mistaken beliefs occur when the speaker performs a speech act where one or more of the felicity conditions are known by the System to be untrue. In such a case, the false condition can still be ascribed to the speaker as a mistaken belief i.e. if P represents a proposition representing the propositional content of the felicity condition, then the following beliefs are true:

bel(System, not(P))
bel(System, bel(Speaker, P))

However, the latter belief may be blocked by contrary evidence that the speaker does, in fact, believe the contrary to the communicated proposition. This is the paradigm case of deception, i.e.

bel(System, not(P))
bel(System, bel(Speaker, not(P)))

An essential feature of deception is that the speaker attempts to convince the System that he or she actually believes the fraudulent proposition. In such a case, the following attitude can be ascribed by the System:

bel(System, goal(Speaker, bel(System, bel(Speaker, P))))

A third possibility is that the Speaker is attempting to satisfy an implicit goal pragmatically. Such an interpretation is licensed if the sets of conditions for mistaken belief and deception can be eliminated. Such an elimination typically requires the ascription of further beliefs to the speaker. We claim that the additional ascriptions can be regarded as part of the context sensitive meaning of the utterance and can be used in plan recognition to trigger further goal ascriptions. This work is described further in [Lee, 1996].

3 Inferring goals using plan evaluation

In addition to being based on false premises, plans can be inefficient, i.e. a plan could achieve the ascribed goal but use a larger set of actions than necessary. Given an apparently inefficient dialogue plan, ViewGen attempts to reason why the speaker chose such a plan, assuming as it does so that the speaker is attempting to implicitly satisfy additional goals. The System attempts to ascribe either conjunctive goals to the speaker which the speaker is attempting to achieve or avoidance goals which the speaker is avoiding and intending the hearer to realise this avoidance. A typical example of the latter would be a patient initiating an unnecessary topic change to avoid answering a difficult question.

In such cases, plan recognition proceeds as follows: given an utterance, its minimal meaning is derived from its surface form. An utterance's minimal meaning is the set of communicated attitudes associated with the speech act of the utterance. The planner then attempts to generate a plan which connects the utterance with one or more ascribable goals. If the plan derived is inefficient i.e fails to achieve the ascribed goals in the most direct manner possible, ViewGen reassesses the speaker's plan by re-planning from the initial context to the ascribed goal. It then ascribes additional goals to the speaker to explain the recognised plan's divergence from the "optimal plan".

4 Dialogue control

In the previous two sections, we have outlined a theory of dialogue understanding which can deal with non-cooperative, personal goal-orientated communication. We have argued that real dialogues are not strictly cooperative and often feature personal goals which must be taken into account when understanding. Such reasoning is also essential in understanding the pragmatic aspects of dialogue.

Our current work is involved in applying this work to initiative management. We have added a set of discourse control acts based on Bunt's context change theory [Bunt, 1995] to model purely discourse based goals and actions. Bunt presents a taxonomy of three different classes of dialogue control act: feedback acts, interaction management acts and social obligation acts. Of particular interest to our current purposes, are the interaction management acts which control turn taking and topic shifting. Because ViewGen is able to recognise and respond to implicit, pragmatic and possibly non-cooperative goals, it is important for the System to be able to take and relinquish initiative to control the topic focus.

One additional concern is the correction of mistaken beliefs on the part of the user. As discussed above, ViewGen is able to reason about the mistaken beliefs. For example in Section 1, we discussed how ViewGen deals with common and stereotypical mistakes based on the identification of concepts such as "stomach". Such mistaken references can be safely ignored in medical counselling and interrupting the patient description of symptoms is often not desirable. However, at the end of the dialogue, the patient communicates a mistaken belief that such symptoms can be safely ignored. Such a belief is dangerous and so must be immediately corrected by the use of a correction speech act. Rather than continue with the diagnosis, the system must attempt to satisfy a goal of educating the patient about the symptom.

This is achieved by the use of an *interruption* plan operator. This interrupts the current dialogue plan being pursued by the System, until the mistaken belief can be corrected. In (1), this is achieved by the System directly stating that such a condition can be serious. Once, the mistaken belief is believed to have been corrected, the original dialogue plan to achieve a diagnosis of the illness is continued.

5 Conclusions

We have briefly presented a theory of dialogue understanding based on belief and goal ascription. Such a theory is able to deal with non-cooperative dialogues. We have argued that such a theory must be extended to handle the control and management of initiative in such dialogues and briefly discussed a recent extension of our theory using explicit dialogue control acts.

References

Allen, J. (1983). Recognising intentions from natural language utterances. In Brady, M. and Berwick, R., editors, *Computational Models of Discourse*. MIT Press, Cambridge, MA.

Ballim, A. and Wilks, Y. (1991). *Artificial Believers*. Lawrence Erlbaum Associates, Hillsdale, New Jersey.

Bunt, H. (1995). Dynamic interpretation in text and dialogue. In M.M. Taylor, F. N. and Bouwhuis, D., editors, *The Structure of Multimodal Dialogue 2*. John Benjamins Publishing Company, Amsterdam.

Grice, H. (1975). Logic and conversation. In Cole, P. and Morgan, J., editors, *Syntax and Semantics*, volume 3: Speech Acts. Academic Press, New York.

Lee, M. (1996). Planning ironical utterances. In *Proceedings of the 16th workshop of the UK Special Interest Group on Planning and Scheduling*, Liverpool John Moores University.

Lee, M. and Wilks, Y. (1996a). An ascription-based approach to speech acts. In *Proceedings of the 16th Conference on Computational Linguistics (COLING-96)*, Copenhagen.

Lee, M. and Wilks, Y. (1996b). Speech acts on demand. In *Proceedings of the 5th International Pragmatics Conference*, Mexico City.

McAllester, D. and Rosenblatt, D. (1991). Systematic nonlinear planning. In *Proceedings of AAAI90*, pages 634–639.

Wilks, Y., Barnden, J., and Ballim, A. (1991). Belief ascription, metaphor and intensional identity. *Cognitive Science*, 15(1).