

Anticipatory planning for decision-theoretic grounding and task advancement in mixed-initiative dialogue systems

Rossana Damiano

Dipartimento di Informatica
Universita' di Torino
C.so Svizzera 185
10149, Torino (ITALY)
rossana@di.unito.it

David Traum

University of Southern California
Institute for Creative Technologies
13274 Fiji Way
Marina del Rey, CA 90292
traum@ict.usc.edu

A very important issue for dialogue systems is the formulation of a response incorporating both grounding functions (feedback about the state of mutual understanding) as well as task-advancement. For example, consider the system's move after the following sub-dialogue:

- 1 **Sys:** Where do you want to go?
- 2 **User:** Boston.
- 3a **Sys:** When would you like to go?
- 3b Tell me more about your travel plans.
- 3c When would you like to go to Boston?
- 3d Do you want to go to Boston?
- 3e Did you say Boston?
- 3f Boston?
- 3g Boston or Austin?
- 3h Where?
- 3i Please Repeat.

The system could simply try to advance the task, either by prompting for user initiative (3a), or continuing system initiative (3b). The system could also focus more narrowly on grounding, providing one of (3d-3i) or similar utterances, depending on the degree of confidence in system understanding, and anticipation of likely source of the problem. An intermediate solution is also possible, combining a grounding verification move with a task advancement move (3c). Several factors affect the decision of which next utterance to produce, including the degree of confidence in understanding, likelihood of other candidate understandings, importance of mutual understanding, cost of extra turns, likelihood of user performing a subsequent repair, cost of subsequent repair, etc. A tricky point is that lots of these factors vary with individual circumstances (users, physical environment, specific task), and are difficult to estimate precisely.

Previous work (Paek and Horvitz, 1999; Ndiaye and Jameson, 1996) has stressed the importance of evaluating alternate actions in dialogue by adopting a decision-theoretic approach (Haddawy and Hanks, 1999), i.e. based on a trade-off between expected cost and benefit. Decision-theoretic approaches al-

low one to account for the uncertainty about the outcomes of each alternative when evaluating their expected utility (given the uncertainty about the current state). For example, the utility of confirmations depends on the improvement they are likely to cause in mutual understanding, given the cost of performing them.

(Ndiaye and Jameson, 1996) propose an explicit representation of the user's and system's moves in order to evaluate the most promising move on the part of the system. The key idea of an *anticipation feedback loop* ((Kobsa and Wahlster, 1989)) is the following: given a particular move of the system, anticipate the potential responses of the user, and evaluate their utilities with respect to the system's goals. Uncertainty about the user's responses is modelled as a probability distribution on the user's responses (Boella and Damiano, 2000).

We argue that the integration of decision-theory and probabilistic reasoning within a framework that allows explicit representation of the user's subsequent behavior will be advantageous for improving the robustness and the flexibility of dialogue systems. Uncertainty exists not only about whether the system has correctly understood the user's contributions, but also about the nature of user initiative, since, in a mixed-initiative system, the user's behavior is not entirely determined by the system's moves even given the assumption of perfect understanding by the system.

Consider a future utterance by the user, following the dialogue example above

User: I want to go to Austin at 3pm.

There are many possible interpretations for such an utterance. It can be a **repair** (I said *Austin*, not *Boston*), a mis-recognized *confirmation* (user actually said *Boston*), but also a user **change of plan** (now I want to go to Austin rather than Boston), a **qualification** (I'm going to Austin at 3, which means that you can't send me to Boston then), or a **new plan** (after we work out the Boston trip, I want to go to Austin). While some of these are beyond the capabilities of currently implemented systems, they

all exist in human-human travel dialogues, and will need to be handled as users grow more demanding about the flexibility and utility of dialogue systems, especially when planning complex constraint-based trips.

Given this framework, we can illustrate how *anticipatory planning* can help guide the decision-making process. We focus on the choice between the alternatives constituted by performing an explicit verification ((3d, 3e, 3f) or an implicit verification (3c). For simplicity, we assume that the evaluation of the utility of a communicative move is based on the *cost* of the action (in terms of time, sentence planning, etc.), on the resulting *degree of understanding*, and on the overall amount of *repairs* during the interaction (which should be as low as possible); the effect of communicative moves on these parameters, in turn, is affected by the degree of confidence in understanding in the current state. Performing an explicit verification has higher costs (especially when the system understood correctly), but increases the degree of (mutual) understanding, while implicit verifications, on the other hand, are less costly but don't induce as much confidence in mutual understanding (since a user might address only the explicit question, leaving more doubt about whether the verified information was implicitly confirmed or ignored). However, an evaluation which accounts only for the immediate effects would not be fully satisfactory: if the system tries to anticipate the user's reasoning in the state resulting from the execution of each alternative, it becomes apparent that when an explicit confirmation has been performed, the subsequent utility for the user to initiate a repair will be lower, due to the increased degree of mutual understanding established by the explicit confirmation. Such a move can thus be warranted not just by current considerations of degree of confidence, but also by a desire to simplify potential future disambiguation problems and change the likelihoods of user actions.

Another advantage of anticipatory planning may be in the design process for new dialogue systems. In particular, when no accurate observations have been made about the behavior of the real users, the intuitions of human developers about the intention-driven activity of the users in response to the system's communicative acts can be embodied in a natural way. This feature also enhances the transportability of a dialogue system to new application domains, where the sets of user's and system's moves change, together with their goals. Moreover, the use of anticipation feedback allows modelling different degrees of cooperativity between the dialogue participants, since it separately represents the user's utility evaluation of alternative moves, which can be based on the satisfaction of different goals and preferences from the system's. Finally, if the sys-

tem's prediction of the user's responses is explicitly represented, the comparison of the expectations underlying the system's design with the user's actual responses will be more direct with respect to an implicit representation, and reconfiguration made easier.

References

- G. Boella and R. Damiano. 2000. Communication and cooperation among agents. In *Proc. of Gotalog 2000 Workshop on the Semantics and the Pragmatics of Dialogue.*, Gothenburg (Sweden).
- P. Haddawy and S. Hanks. 1999. Utility models for goal-directed, decision-theoretic planners. *Computational Intelligence*, 14(3).
- W. Kobsa and A. Wahlster. 1989. User models in dialog systems. In *User models in dialog systems. Kobsa, A., and Wahlster, W., eds.*, pages 4–34. Springer, Berlin.
- Alassane Ndiaye and Anthony Jameson. 1996. Predictive role taking in dialog: Global anticipation feedback based on transmutability. In *5th International Conference on User Modeling*, pages 137–144.
- Tim Paek and Eric Horvitz. 1999. Uncertainty, utility, and misunderstanding: A decision-theoretic perspective on grounding in conversational systems. In *Working Notes AAAI Fall Symposium on Psychological Models of Communication in Collaborative Systems*.