

# Decision Theoretic Dialogue Planning for Initiative Problems

Bryan McEleney and Gregory O'Hare

Department of Computer Science, University College Dublin, Dublin 4, Ireland  
{bryan.mceleney, gregory.o'hare}@ucd.ie

**Abstract.** The taking of initiative has significance in spoken language dialogue systems and in human-computer interaction. A system that takes no initiative may fail to seize opportunities that are important, but a system that always takes the initiative may not allow the user to take the actions he favours. We have implemented a mixed-initiative planning system that adapts its strategy to a nested belief model. In simulation, the planner's performance was compared to two fixed strategies of always taking the initiative and always declining it, and it performed significantly better than both.

## 1 Introduction

In cooperative dialogue one agent will consider a goal that cannot be satisfied through his own action. He may choose to act towards it and in doing so both commits to it and reveals his intention to a second agent. Then on the second agent's turn, a cooperative response is considered to satisfy the intention. Symmetrically, the second agent can act towards an intention, again obtaining recognition of his choice and his commitment. Often, the plan is so structured that there is only one option for a cooperative response, but there are also many times when the agent has two or more options. Among these, there are different levels of initiative - the agent might do nothing, he might add the next child act to the focussed goal, he might choose between different children to add, he might move the focus point, or he might introduce a new goal. An agent acting alone may have several options to choose amongst and chooses the one of maximum expected utility. Where there are two agents and a single goal, he must decide whether to act now, or wait for the other agent to decide whether the goal is important enough. Where there are two agents each with different goals, each agent must consider the coordination of their initiative with the other agent's goal. For these decisions, the agent must use a nested belief model to determine how the second agent is expected to respond to his initiative, what initiative will be taken on the third turn as a result of that, and so on. Typically, the agents will have different beliefs, so one agent may believe he should take the initiative while the other may believe he should not, or neither agent will take the initiative even though a goal was worth pursuing. With belief revision at each turn, initiative may surface at different times, even for the same goal. For example,

one agent might state that it is opportune to fix the car, which would promote the other agent's latent initiative to go to the shops and buy some tools. As an example of deeply nested initiative, the agent with the idea to fix the car might foresee the second agent's initiative to buy some milk while he is at the shops, so that as the last act in the plan, that they would have a cup of tea. For these problems, we will describe a general purpose dialogue planner that uses deeply nested, probabilistic belief models.

## 2 Example Problem

The planner has been tested with several initiative problems, where one agent decides upon his own goals, where two agents compete for the same goal, and where two agents decide on a sequence of different goals. Here we present an instance of the second type, a flight booking problem. The system decides whether to offer a passenger a window seat, but is wary that the passenger may not intend to have one, and would be inconvenienced by the conversation. The system's other option is to not offer the seat, and wait for the passenger to take the initiative. The passenger is also unsure whether to take the initiative since his effort will have been wasted if window seats are not available. The plan for this subdialogue is three steps deep, with two sources of variation - the travel agent's belief about the passenger's intention to have a seat, and the travel agent's belief about the passenger's belief about the travel agent having available seats. The performance of the system can be described as a function of these two variables.

A plan library was constructed for use with the planner, from which each belief set in a nested model was populated. A reward of 100 is obtained by a happy passenger, and 65 by passengers who wanted a window seat but ended up without one. A cost of 5 was uniformly assigned to each dialogue act, except for a special act chat, of cost -1, which allows the agents to decline the initiative by chatting about the weather. Lacking empirical data in the domain, these are estimated costs.

## 3 The Dialogue Planner

The nested belief model of the agent is a quantitative form of the BDI (belief-desire-intention) model [6]. Nested beliefs are probabilistic. At each turn, the belief revision module revises the belief model based on the preconditions and effects of the dialogue act of the last turn. On the agent's turn, the planner decides a dialogue act, using a decision tree. In the tree, there are chance nodes that represent uncertainty about the mental state of the next actor, and choice nodes that represent his options. Figure 1 illustrates the decision tree for our example problem.

Each node of the tree is grown by the plan recogniser. Plan recognition is done recursively, by recognising the tail of the history list, performing a planning step at the next level of nesting (in agreement with Pollack [4]), and then filtering the hypotheses that are consistent with the current history item. For example to

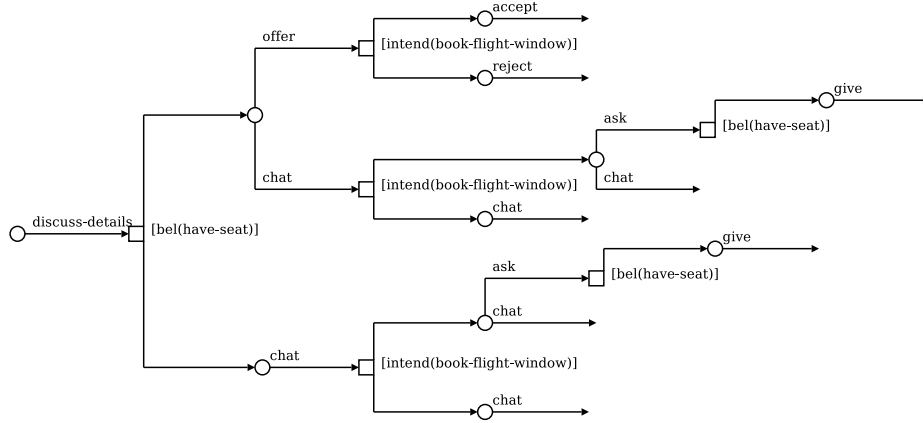


Fig. 1. Decision Tree


continue an “ask” act performed at level 3, the planner at level 4 would call the recogniser at level 5 which produces a hypothesis for the first agent’s intention. Then, to continue the plan, a forward planning step at level 4 adds a child node to the hypothesis. Planning can be done by both decomposition chaining and goal-effect chaining.

The decision tree is evaluated using an expected utility calculation at each chance node. At choice nodes, each agent chooses the maximum expected utility branch. Both of these calculations are done in the context of the belief model of the choosing agent. For example, the passenger may have one value for a subtree because he believes a seat is available, but the travel agent, believing none is available, has a different value for the same subtree. The planner can calculate utility in both cooperative and self-interested fashions.

#### 4 Results


The system was evaluated by simulation rather than with a human counterpart, since evaluation over a range of belief model states is time-consuming. A similar approach to evaluating initiative strategies is advocated in [2] and [3]. The configuration of the belief model was varied for the evaluation. The first result is the distribution of initiative, shown in figure 2. The initiative rate is high when either the passenger is expected to intend a window seat, or when the passenger believes that a window seat is available. The left-hand plot represents the initiative distribution of the travel agent, whereas the right hand plot represents that of the passenger. Notice how the travel agent backs off when he believes that the passenger will pick up the initiative in the next turn, and that nobody takes the initiative when both variables have low values.

The second result for the system is its performance. We plotted the difference between the utility of the system and the two fixed initiative strategies of always



**Fig. 2.** Distribution of Initiative: (a) Travel Agent (b) Passenger

declining the initiative (fixed-chat) and always taking it (fixed-offer). To do this, we modified the planner to prune all but the desired strategy from the relevant choice node. Figure 3 shows the results. On the left, the planner always beats the fixed-chat strategy, and by as much as 4.60 units at  $[1,0.2]$ . This is a good fraction of the length of the longest dialogue in the decision tree, which is  $[\text{chat}, \text{ask}, \text{give}]$ , of length 15 units. The gain is similar for the fixed-offer strategy with a margin of as much as 12.0 units at  $[0,1]$ . Similar margins have been found across our battery of sample problems.



**Fig. 3.** Utility Gain over Fixed Strategies (a) Always Decline (b) Always Take

## 5 Discussion

An early approach to initiative planning was that of Smith and Hipp [5], in which a theorem prover constructs plans, employing a user model to plan subdialogues for obstacle subgoals. The planner could select between initiative modes, but it was not quantitative, and did not use a nested model. Recently, Fleming and

Cohen [1] discussed a decision-theoretic approach to initiative that is similar to our own, but do not provide an implemented general purpose planner, nor do they use a nested model. There have also been two notable computer simulation studies in dialogue initiative. Guinn [2] used a quantitative negotiation system with a user model to decide between continuations to a plan, but did not use this approach to plan beyond the current move. Ishizaki et al [3] examined a map-following task and found that mixed initiative dialogue can reduce utility. However their system used an initiative policy rather than adapting to the plan and user at hand.

The results for our system are encouraging but presuppose a wide variation of belief model configurations in the dialogue system's lifetime. Otherwise, a fixed strategy is just as good. The system must have a dynamic model of the user available. This could be provided, or, using the belief revision module, the system could automatically update stereotypes as dialogue evidence accumulates. The computational demands of the system are light, with just a handful of simple calculations and comparisons performed in evaluating the decision tree. It is easy to compile the tree evaluation function to a procedural programming language. It takes little time to configure a plan library for use with a new problem. We are extending our battery of problems to widen the scope of our results.

## 6 Conclusion

We have evaluated a dialogue planner that decides initiative by adapting to a probabilistic nested belief model. In simulation, the planner performed better than otherwise equivalent fixed initiative planners.

## References

1. Fleming, M., Cohen, R.: A User Modeling Approach to Determining System Initiative in Mixed-Initiative AI Systems. Proceedings of the Eighth International Conference of User Modelling (2001) 54–63
2. Guinn, C.: An Analysis of Initiative Selection in Collaborative Task-Oriented Dialogue. User Modelling and User Adapted Interaction **8(3-4)** (1998) 255–314
3. Ishizaki, M., Crocker, M., Mellish, C.: Exploring Mixed-Initiative Dialogue Using Computer Dialogue Simulation. User Modelling and User Adapted Interaction **9** (1999) 79–91
4. Pollack, M.: A Model of Plan Inference that Distinguishes between the Belief of Actors and Observers. Proceedings of the 24th Conference of the Association of Computational Linguistics (1986) 207–214
5. Smith, R., Hipp, R., Biermann A.: A Dialogue Control Algorithm and its Performance. Proceedings of the Third Conference on Applied Natural Language Processing (1992) 9–16
6. Rao, A., Georgeff, M.: Modelling Rational Agents within a BDI architecture. Proceedings of the Second Conference on Knowledge Representation and Reasoning (1991) 473–484