A Model for Reasoning about Interaction with Users in Dynamic, Time Critical Environments for the Application of Hospital Decision Making

Hyunggu Jung and Robin Cohen

Cheriton School of Computer Science
University of Waterloo
{h3jung,rcohen}@uwaterloo.ca

1 Overview

In this paper, we present a model for reasoning about interaction with users in dynamic, time critical environments in a way that is sensitive to the cost of bother. We project the model into the scenario of decision making in hospital emergency room settings, providing a framework for modeling the doctors in that environment, to determine whom to ask to attend to a current patient. A simulation of our model demonstrates that it offers valuable improvements due to its reasoning about bother.

2 A Model for Dynamic, Time Critical Scenarios

Our research aims to develop a model that can be used for scenarios where an agent is reasoning about which human users to enlist to perform decision making, in an environment where decisions need to be made under critical time constraints and where the parameters that serve to model the human users are changing dynamically, to a significant extent.We offer a hybrid transfer of control strategy that takes as its starting point the model of Cheng [1], which includes reasoning about interaction (partial transfers of control or PTOCs) as well as about full transfers of control of the decision making (FTOCs) to another entity.

Each possible transfer control strategy is generated in order to select the strategy that offers the highest expected utility after period of time. A transfer of control chain includes provision to ask a different entity after waiting a determined period of time. Distinct from Cheng's original model, attempts at FTOCs are in framed as PTOCs with the question Q: "Can you take over the decision making?". This then enables either a "yes" response, which results in an

FTOC[1] or a "no" response or silence.

The "no" and silence responses ultimately lead to a new method aimed at coping with the dynamics of the environment and the criticality of time. As in Cheng's model, a TOC strategy will arrange for alternate users to query,

[1] As a simplication, we assume that a "Yes" response results in the user successfully assuming control of the decision.

A. Farzindar and V. Keselj (Eds.): Canadian AI 2010, LNAI 6085, pp. 319–323, 2010.

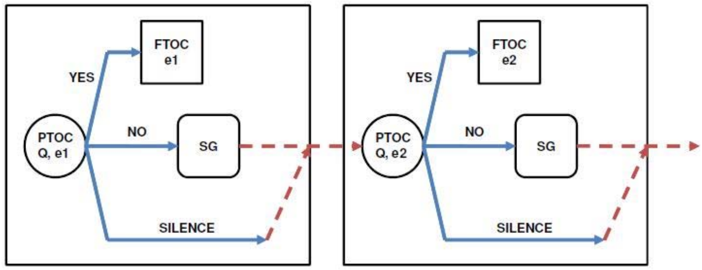

Fig. 1. Visual representation of strategy with the FTOCs and PTOCs; each world occupies one square

once the current respondent has not volunteered to assist. We first make the simplification that the strategies do not ask different entities within the same chain. This is because we are limiting ourself to only one question, that of asking the expert to help. Then, at the end of this chain of attempts, we inject a final decision of strategy regeneration. Strategy regeneration will then allow for an updating of parameter values. This new generation of possible strategies allows us to reason at this point in time regarding the users who are available to help and their expected quality of decision, cost of bother, etc. with information that is no longer stale; this is of particular value in circumstances where choices that are less than optimal can be undesirable, to a dramatic extent.

A diagram outlining the FTOCs and the PTOCs that we envisage is presented in Figure 1 where an arrow with a solid line means the stream of time, but a dotted line means there is no break by the end of the arrow. In addition, we introduce a concept of *world* to facilitate the computation of the utility of any given strategy.

One *world* consists of one PTOC, one FTOC, and one SG (strategy regeneration) node and includes all the parameters currently used to calculate benefits and costs to reason about interaction with entities. Therefore, when the current *world* is moved to the next step, our system asks a new entity. The number of worlds is equivalent to the number of entities that will be asked.

There are n FTOC nodes, n PTOC nodes, n SG nodes, and one virtual node in the overall framework with n worlds. We obtain the overall EU of strategy s by summing up n EU values for FTOC nodes, n EU values for SG nodes and one EU value for the virtual node as follows:

$$EU(s) = EU_n(dfl) + \sum_{j=1}^{n}(EU_j(fn_l) + EU_j(sg)) \tag{1}$$

where dfl[2] reflects a virtual node, n denotes the number of worlds, $EU(fn_l)$ reflects the utility of ending in a FTOC, and $EU(sg)$ reflects the utility of ending in SG node.

[2] The leaf node for the silence response is set to sg.

3 Selection of Experts for Medical Decision Making

We are also interested in determining an appropriate reasoning strategy to find the right person, at the right time, to assist with the care of patients who are arriving at a hospital emergency room. Typically in these settings, patients who appear to require further assistance than can be immediately provided (what we could call "a decision") require soliciting aid form a particular specialist.

In order for the human first clinical assistants (FCAs) to make the best decisions about which specialists to bring in to assist the patients that are arriving, the proposal is to have our multiagent reasoning system running in the background, operating with current parameter values to suggest to the medical professionals who exactly they should contact to assist the current patient. These experts are then the entities $\{e_1,\ e_2, \ldots,\ e_n\}$ that are considered in our reasoning about interaction, with the FCA contacting the experts according to the optimal strategy our model generates (who first, waiting for how long, before contacting who next, etc.)

We propose the addition of one new parameter as part of the user modeling for the bother cost, a $Lack\_of\_ExpertiseFactor$. This parameter is used to help to record the general level of expertise of each doctor (i.e. medical specialist), with respect to the kind of medical problem that the patient is exhibiting.

We also adjust the calculations proposed by Cheng for estimating bother cost, in order to reduce the number of parameter values that need to be acquired or solicited (for our time-critical scenarios).

We assume that a user's willingness is simply determined by their attentional state factor and their expertise level. We also assume that a user's probability of response is determined by their willingness. Some factors which affect bother cost in hospital settings are thus as follows.

- *User_Unwillingness_Factor=Attention_State_Factor+Lack_of_Expertise_Factor*
- $Init = User\_Unwillingness\_Factor \times Attention\_State\_Factor \times TOC\_Base\_Bother\_Cost$
- $BSF\ (Bother\,So\,Far) = \sum_{toc \in PastTOC} TOC\_Base\_Bother\_Cost(toc) \times \beta^{t(toc)}$
- $BotherCost\ (BC) = Init + BC\_Inc\_Fn(BSF, User\_Unwillingness\_Factor)$

We introduce another new parameter, *task criticality (TC)*, to affect the reasoning about interaction. TC is used to enable the expected quality of a decision to be weighted more heavily in the overall calculation of expected utility, when the case at hand is very critical. This parameter may also be adjusted, dynamically. When a patient has high task criticality, strong expertise is required because the patient's problem may become much more serious if not treated intensively.

With bother to an expert being modeled in detail, this leads to strategies that are less likely to ask an entity who will simply fail to respond. Our detailed bother modeling for time critical environments is an advance on other bother models such as [2], which focuses on modeling annoyance.

4 Validation

Our validation measures performance of our model reflecting dynamic and time critical aspects by comparing with that of a model missing the calculation of bother cost. In the setting of our validation simulating hospital emergency scenarios, there are four entities on the entity list and five patients on the waiting list. Every patient has a task criticality for the specific medical problem and the task criticality of each patient is changed dynamically as time passes. Our simulation first selects the patient whose task criticality is highest among those of patients.

We then obtain a strategy chain by calculating formulae reflecting our model with information of each patient. After choosing a entity in the chain, we ask him/her to treat the current patient and update the task criticality of patients who have been treated by entities, as well as those remaining on the waiting list. When there is no more patients on the waiting list, we finally count the number of dead patients[3]. By comparing the number of dead patients simulated by our model with bother cost and without bother cost, we can validate whether our model reflects dynamic and time critical domains effectively. Our current results demonstrate valuable improvements with our model.

5 Discussion

Our work contrasts with those of user modeling researchers such as [2], who focus on modeling a user's plans in order to determine whether to interact, whereas we focus on deriving the best expected utility from the interaction. The user modeling approach to our research coincides well with those of others such as [3], who advocates that values of variables in user models be determined as a combination of modeling the specific user, the class to which that user belongs and the traits that are typical of all users. For future work, it would be valuable to explore a counterpart to Fleming's stereotypical classes, in order to gain greater insights into how to set the values of the user modeling parameters. For future work, we should also lift the various default parameters suggested in Section 3 and explore methods for learning about our users, over time. One possible starting point for this work is the research of [4], which advocates the use of active learning, to involve the user in the process of progressively determining appropriate parameter values for interruptions.

References

1. Cheng, M., Cohen, R.: A hybrid transfer of control model for adjustable autonomy multiagent systems. In: Proceedings of AAMAS 2005 (2005)
2. Raskutti, B., Zukerman, I.: Generating queries and replies during information seeking interactions. Int. J. of Human Computer Studies 47(6), 689–734 (1997)

[3] This is a rather drastic statistic. We are exploring other measures of system performance.

3. Fleming, M.: The use of increasingly specific user models in the design of mixed-initiative systems. In: Proceedings of the 17th Canadian Conference on AI, pp. 434–438 (2004)
4. Kapoor, A., Horvitz, E.: Experience sampling for building predictive user models: a comparative study. In: CHI 2008: Proceeding of the 26th annual SIGCHI conference on Human factors in computing systems, pp. 657–666. ACM, New York (2008)