Modeling User Perception of Interaction Opportunities in Collaborative Human-Computer Settings

Ece Kamar, Barbara J. Grosz and David Sarne
School of Engineering and Applied Sciences
Harvard University, Cambridge, MA 02138 USA
{kamar, grosz, sarned} @eecs.harvard.edu

Introduction

Interruptions are important for effective collaborative work, because agents often possess information required by others on their team. This need to get information from another agent arises in mixed human-computer teams as well as in homogeneous computer-agent environments. For example, a (human) driver may see changes in weather conditions that affect route selection while an automated navigation system without sensors does not. The navigation system may need this information to identify the best route. It is crucial to time interruptions, which are inherently disruptive, appropriately. Efficient interruption timing improves task performance as well as emotional state and awareness of the user, and decreases the negative effects of interruption (Adamczyk & Bailey 2004).

A key aspect of reasoning about interruptions in collaborative settings is the ability to accurately estimate the costs and benefits of the interruption so that the outcome of the interruption positively affects group task outcomes. Cost estimation has been investigated in prior work on interruption management (Horvitz & Apacible 2003), but this work presumes a benefit to the user of having information the computer system can provide. The benefits of interruption have been studied in the adjustable-autonomy literature, but that work focuses on when to turn control over to a person (Tambe *et al.* 2006). Few models have combined these two aspects into an integrated decision making mechanism (Fleming & Cohen 2001), and none have done so in the kinds of fast-paced domains we consider, i.e., domains in which agents are distributed, conditions may be rapidly changing, actions occur at a fast pace, and decisions must be made within tightly constrained time frames. Furthermore, almost no attention has been paid to the possible discrepancy between a computer agent's calculation of the utility of the interruption and a person's estimation of the usefulness of the interruption. The failure to estimate accurately may lead to a person rejecting the interruption, and thus to a missed opportunity to improve team performance, turning the interruption into an unnecessary disturbance.

Our research proposes a new model for interruption management. This model aims to help maximize the efficiency of collaboration between an agent and a person by better estimating interruption outcomes and by taking into account the possible mismatch between the computer's calculation of utility and the person's perception of it. It focuses on determining the factors that influence people's perception of interruptions, and thus their overall tendency to accept or reject them, when they are generated by a computer system. The results will enable the design of more efficient interfaces, ones for which the likelihood that valuable interruptions will be accepted by the user is higher.

To investigate the interruption management problem empirically, we developed a new, abstract game using the Colored Trails (CT) infrastructure which has been used previously as a research test-bed for a variety of decision-making problems (Grosz *et al.* 2004). This framework enables us to focus on investigating the interruption problem without the specification overhead of real world domains, but is sufficiently interesting for human participants to play; thus it provides a good test environment.

Modeling Interruptions

The CT game we defined involves two players, one controlled by a computer agent and the other by a person. Players are allowed to move one step at a turn in one of four directions on a board. They have individual goals which they aim to reach as quickly as possible. The dynamically changing nature of the real world is mimicked by having the goals move stochastically with a probability determined by a Gaussian function with a center at the current position of the goal. When players reach a goal, they and the goal are randomly relocated on the board, and another round of play starts. This new round is the analogue of being assigned a new task. The game continues until a certain number of turns are played. This model can be generalized to real world applications such as email assistance by replacing the simple cost and utility functions with application specific functions.

The remainder of this paper describes an initial investigation which considers a collaborative setting in which the person has complete information—including information from which the agent can benefit —and the agent has incomplete information[1]. In particular, the person has complete

[1]This initial setting provides a baseline for assessing the influences on a user's determination of interruption usefulness. Future

information about both players and their goals, whereas the agent lacks current information about its goal's position (but knows its position, the position of the person and the person's goal). The agent must initiate interactions with the user to learn its actual goal position. We model the cost of interruption as the loss of the opportunity to move for one turn. At the beginning of each turn, the agent decides whether to interrupt the user or not. The person is free to accept or reject an interruption request. If the person responds positively to an interruption request, the agent is told the current position of its goal and both players are prohibited from moving for this turn.

The players share a common scoring function S,

$$S = S_P + S_A \text{ and } S_i = \sum_k (s - h_k)$$

where S_A and S_P denote the agent's and the person's accumulated (individual) points, respectively, s is the number of points given for reaching a goal, h_k is the number of moves it takes for the player to get to the kth goal, and the sum is over all goals that have been reached by player i. The objective of players is to maximize S. The person has an incentive to accept interruption requests, because overall success depends on the agent's ability to reach its goal.

The expected outcome of interruption (EOI) is the difference between the expected outcome (EO) of the game when there is an interruption and the EO when no such interruption takes place. The agent interrupts the user when the EOI is estimated to be positive. This calculation requires deriving joint policies for the collaborative group of agent and user, a problem which may be modeled as a Decentralized POMDP (Dec-POMDP). However, the complexity of the solution is *NEXP-complete*. [2] As a result, our approach is to estimate the EOI by combining agent and person-sided estimates of interruption outcomes using the equations below, where I indicates Interruption, NI indicates No Interruption, and EU indicates Expected Utility:

$$EOI = EOI_P + EOI_A$$
$$EOI_P = EU_P^I - EU_P^{NI} \text{ and } EOI_A = EU_A^I - EU_A^{NI}$$

The person-side estimate may be modeled by a Markov Decision Process (MDP), because the person is able to observe the complete state of the world. ExpectiMax is used to calculate EU_P for the current state of the world, which is represented by person-player position p, goal position g and current turn h. Given that B indicates the set of possible board positions and $MP(g', p, g)$ is the probability of a goal move from position g to g' with player position p,

$$EU_P^{NI} = EU_P(p, g, h)$$
$$EU_P^I = \sum_{g' \in B} MP(g', p, g) \times EU_P(p, g', h+1).$$

The modeling of the agent-side requires a Partially Observable MDP (POMDP), because the agent does not have complete information about the state of the world. The current state of the game is represented by agent-player position p; belief state b, which is a probability distribution over the possible goal positions; and current turn h. After each turn, b is updated to b' with State Estimator (SE).

$$b'(c' \in B) = SE(c') = \sum_{c \in B} b(c) \times MP(c', p, c)$$
$$EU_A^{NI} = EU_A(p, b, h) \text{ and } EU_A^I = EU_A(p, b', h+1)$$

Experimental Setting

We are currently running experiments of participants playing this CT game in a lab setting. To evaluate the agent's interruption decision making performance, we simulate three homogeneous agent settings and compare the outcomes with the results of the human-agent experiments. The agent settings comprise two computer agents with different capabilities: (1) both agents have complete information; (2) one agent has incomplete information, but is able to interrupt the other agent; and (3) one agent has incomplete information and is not able to interrupt the other.

A second set of computer-human experiments uses handcrafted interruptions generated by the computer agent, which vary such parameters as the complexity of calculating costs and gains, the magnitude of the EOI, and the type of the collaborator. The experiments aim to identify the subset of factors an agent needs to focus on for learning or profiling to better predict a user's tendency to accept interruptions.

Conclusion and Future Work

This paper describes an experimental design and a computational framework for exploring an integrated interruption model in collaborative settings to determine how people perceive the effectiveness of interruption.

References

Adamczyk, P., and Bailey, B. 2004. If not now, when?: the effects of interruption at different moments within task execution. In *CHI '04*, 271–278.

Fleming, M., and Cohen, R. 2001. A user modeling approach to determining system initiative in mixed-initiative ai systems. In *UM '01*, 54–63.

Grosz, B.; Kraus, S.; Talman, S.; Stossel, B.; and Havlin, M. 2004. The influence of social dependencies on decision-making: Initial investigations with a new game. *AAMAS'04* 02:782–789.

Horvitz, E., and Apacible, J. 2003. Learning and reasoning about interruption. In *ICMI '03*, 20–27.

Kamar, E., and Grosz, B. J. 2007. Applying MDP approaches for estimating outcome of interaction in collaborative human-computer settings. In *MSDM 2007*, 25–32.

Tambe, M.; Bowring, E.; Pearce, J.; Varakantham, P.; Scerri, P.; and Pynadath, F. 2006. Electric elves: What went wrong and why. In *AAAI'06*.

work will vary the allocation of information.

[2]Comprehensive analyses of the problem, MDP, POMDP and Dec-MDP approaches may be found in a longer paper (Kamar & Grosz 2007).

The research reported in this paper was supported in part by NSF grants IIS-0222892 and CNS-0453923 and in part by contract number 55-000720, a subcontract to SRI International's DARPA Contract No. FA8750-05-C-0033. Any opinions, findings and conclusions, or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of NSF, DARPA or the U.S. Government.