Jesse Gray

Project Exercises

First Project Idea, Week 1 Assignment:

Huge amounts of content (news, information, music, etc.) are being produced and readily available today because of the ease of information access on the internet. Many people use human recommendations for what to view (weblog's, slashdot, etc.), and in some cases automated recommendation agents attempt to show users what they want to see. I think that these automated agents that take into account a user's past choices show promise, but people are frustrated with the one way communication that occurs with these agents. The agent acts as an advisor, but is very sensitive to a user's choices and the user cannot explain the reason behind the choices or correct erroneous impression that the agent has. I think a system that builds on the statistical methods of these agents, but incorporates a feedback channel for the user could be both more effective and less frustrating for the user.

The System's Goal The goal of the system is to combine the statistical recommendation systems employed today with a relational agent interface. Some people express concern that these statistical recommendation systems "judge" them based on a few outlier purchases, and wish to correct the system by giving it more information (i.e., the special circumstances under which they like the outlier purchases, or the actual thread that ties those purchases together)

Advantage of Relational This system should be relational because it is occupying the role of an advisor or salesman, and in its currently non-relational state it frustrates people because it can give them advice, but they cannot talk back and explain more about what they want or don't want, or the context in which they want it. People already sometimes personify these systems, but the personification only makes them more annoyed that it has the wrong idea about their desires.

Relational Aspects The relational aspects of this system are the advice it gives the user, and the system for accepting meta information from the user. The system will offer recommendations whenever the user interacts with it using the knowledge it has gained about the user's tastes. The user will give feedback to the system about the advice, indicating what it agrees with, what it doesn't agree with, and some measure of why.

Evaluation The system would be evaluated based on the quality of the recommendations, and the satisfaction of the user. I separated those out because I think that in many cases the current systems make quality recommendations, but users are still frustrated by their inability to give feedback and fix small errors the systems have made about their tastes.

Week 2 Assignment, Description of the relationship:

Type of relationship I envision the relationship being modeled on the relationship between a human and an assistant or advisor that understands the desires of the human, and conducts research to find new content (music, products, etc., depending on the application) that the user would enjoy. I want the system to be able to understand the context for each interaction, so the system can learn how the user needs different information based on the current context.

Reason for relationship I feel that recommendation agents are becoming more important as the amount of content increases. People already seem to personify these agents, so the existing agents that cannot accept feedback are frustrating to the user. Not only would a relational agent reduce the user's frustration, but would hopefully provide better quality results through the user's feedback.

Establishing relationship The relationship can be established by making contact with the user with some initial recommendations (perhaps based on statistical methods like the recommendation agents in use today). The user will be allowed to give feedback about the recommendations and explain their preferences to some degree, and the agent will refine its recommendations.

Relationship maintenance In order to maintain the relationship, the agent will need to have a satisfactory system for taking feedback from the user, and use that feedback and the user's preferences to make recommendations that satisfy the user. The quality of the recommendations should increase with the relationship.

Week 4 Assignment:

First, I am thinking of changing my project idea at this point. The music recommendation seems like a useful tool, but I'm not confident that the relationship aspects of the system are central enough to the functionality to make it it relevant to this class. I'm currently thinking of a very different idea, one related to my thesis work last summer. That work involved developing a goal and action system to run on the Leonardo robotic platform based on the concept of "simulation theory". In short, the robot uses its own movement repertoire to try to classify movements observed in the human, then it uses its own action repertoire to classify the human's movements as an instance of one of its own actions. Since the robot is classifying the human's actions against its own action structure, it also can determine the goal of the action, and thus infer that it will be the goal of the human. It can further help by completing the action in case the human fails, or by providing assistance in multipart actions. I would like to extend this system with additional functionality, such as being able to make similar simulation theoretic inferences about mental state of the human, both about information concerns (i.e., the robot can determine that the human needs a certain piece of information to complete the action) and about emotional concerns (the robot can determine what the human needs to alter its sad mood).

I think that the primary factor in evaluating the effectiveness of this system is the ease with which the human can complete tasks with the robot's help. I envision a simple task, but one which the human cannot complete on their own. This may be because the human

was not provided with the complete information necessary to complete the task, or because it is a task that cannot be completed by one person (for example, it requires three hands). Another type of task might be one that is possible to do on one's own, but is easier with the help of the robot (for example, there are too many conditions to easily keep track of).

The ease of task completion could be measured by the time it takes to complete the task. To determine the potential benefit of the goal inference mechanism to the task, there could be a control task presented to a second group of participants. While in the goal inference group the robot will attempt to infer the goals of the human and proactively provide assistance, in the control group the robot will have no goal inference ability. Participants in the control group will have to explicitly communicate what they are doing and what they want the robot to do.

The task times could be compared between goal inference case and the control case to determine if the goal inference mechanism allowed the task to be completed more quickly.

Another factor that I would like to measure (and compare between the groups) is the level of satisfaction or frustration that the user felt while completing these tasks, as well as a measure of how helpful they thought the robot was to the task completion. While some type of physiological mechanism could potentially be used to measure this reaction, it is likely that I will need to rely on a self-report questionnaire for this data.

Week 6, more task info and pre-interviews

As the agent I have described here is intended to infer certain cognitive information about a human interacting with it (immediate action, goal, and knowledge about the local environment), I have been consider two possible ways of evaluating it, either by using the information about the human to attempt to provide timely information and help in a cooperative task, or by using that same information about the human to thwart the activities of the human in a competitive setting.

In either case I would like to design a system that I can get a few subjects to interact with, and measure both their performance at the task and a subjective sense of how they felt about the interaction (how helpful/frustrating was the agent). I will probably have to conduct these experiments with a virtual on-screen character instead of a robot in order to create (in a timely manner) a sophisticated enough task for the experiment. However, I hope to be able to use some extra modalities, such as speech recognition and gaze detection. It would also be interesting, if feasible, to measure physiological signs of frustration such as pressure on the mouse or galvanic skin response.

Another interesting pilot study to perform would be to set up the same task but with two humans (no agent) and see what kind of helpful or competitive techniques they use. This could possibly provide some insight into the cues the agent will respond to.

Helpful Case: I have been thinking of a task where the robot and human have slightly different information (some barriers, exist, perhaps, so neither can see the entire task surface). In the helpful case, the agent will be set up to try to help the human, once it has determined what the human is attempting, so it will provide relevant information about the task based on what the human needs to know to proceed. As a control, the same task could be attempted without any proactive abilities in the robot. In both cases the robot will respond to commands from the user about the task; however, hopefully in the helpful case the user will need to use these commands less frequently. In addition to the success metrics above, another interesting piece of data could be the number of direct commands the user had to issue to accomplish the task.

As a pre-interview, I have been thinking of talking to people who work with Robonaut, since it is a robot designed to work in complicated tasks with humans. I would like to talk to some of the designers/operators of Robonaut, who have thought about many robot-human cooperation issues and likely encountered relevant problems as they perform experiments in simulated Robonaut/astronaut cooperative tasks. I would also like to talk to people who have participated in the other half of the interaction, i.e. an astronaut who has performed cooperative tasks with Robonaut as well as with other astronauts, and can provide information about successes and difficulties in both of those setting.

Competitive Case: In the competitive case, I am envisioning a situation where the user is playing some sort of competitive game (not necessarily a symmetric game) against the agent. A good control for this experiment is not obvious to me; perhaps based on the type of game they are playing, there could be a version of the agent that tries to win ignoring the human and a version that actively attempts to thwart the human's plans (but the game would have to be chosen so that both those strategies are relevant).

As a pre-interview, I think it would be good to talk to both designers of game agent AI, as well as people who have some experience playing games. It will be interesting to see kind modeling the human is done by antagonistic game agents, as well as what people expect and find lacking in current game AI.