Stefanie Tellex

Project Exercises

Week 3

I will still consider a voice controlled wheelchair. I would like to study aspects of the human-horse relationship, since that also involves a sentient vehicle. In addition, I would like to capture aspects of the friend relationship, because the machine is going to make mistakes, and I want the human to understand its limitations and work together with it to work around them.

Both the user and system will benefit from establishing a relationship. The system will benefit most from the relationship when it makes a mistake. It will rely on its partner to get it out of difficult situations. The system gives the user increased mobility but will not be able to always deliver. Since the wheelchair will understand and use language, it will automatically trigger high user expectations of both social and linguistic intelligence, even if the user tries to consciously suppress them. Given that people become attached to Aibos and Rhoombas, the human will probably form some form of relationship with the chair on their own terms. For both ethical and practical reasons, it makes sense to try to make this experience as positive as possible. First, it will avoid, as much as possible, raising the human's expectations through language use and not fulfilling them. Even if we cannot completely fulfill the user's expectations, thinking about the problem and evaluating the system will identify ways the system is failing, and there is an opportunity to build work-arounds (or at least to write a disclaimer!). I don't think there is much cost to the system to establish a relationship. There is cost to implement and evaluate it (since this will be a long-term relationship), but given the level of user interaction with and user dependence on this system, it is worth considerable effort to make it as smooth and natural as possible.

One way to establish the relationship is through a tutorial session, where the user "gets to know" the chair, and learns how to use it. In this phase, the chair will teach the user how to use it, but also teach the user about its limitations. This might take the form of navigating through a closed-course obstacle course. An environment and task could be designed that takes the user through common tasks. Although in general the chair would not know what to do when, in a closed environment it should be able to offer more guided advice. In this setup it would be important to make it clear that the chair will not typically have so much information about its environment.

In order to maintain the relationship, the chair has to be benefiting the user enough to be worth using. If this isn't the case, it is to the user's benefit to revert to whatever solution they were using before for mobility. For the rest of this analysis, I will assume that the chair is working well enough to be worth using. In that case, the major danger is that the user will take the chair's capabilities for granted, and become less and less forgiving of its weaknesses. I think there are two types of measures to take to avoid this. First, the chair should fail gracefully. For example, when it is confused it could ask the user to help by

using the joystick. It could apologize when it does something wrong, or ask a clarifying question. If possible, it should explain the underlying problem (eg, a noisy environment, a low battery, a motor stall.) Second, the chair should try to remind the user when it performs well. This will step a fine line between being obnoxious and establishing a relationship. For example, after obeying a command, it could ask "Is this right?", even it is fairly sure it is correct, in order to get the user to acknowledge its good performance. (This would probably have to be a pretty low frequency occurrence of course, and it should also be able to respond to corrections if it turns out it was in the wrong place.) Or it could somehow convey its emotional state to the user - maybe through lights, or a monitor? It could then show its happiness when it performs well, and show sadness when it makes a mistake.

Week 4

First, I want to measure the effectiveness of the language understanding. This evaluation doesn't directly measure the relational component directly, but it is important because one of the reasons I'm building this system is to understand spatial language. In addition, it will serve as a baseline for trust measures; it will be a way of quantifying how trustworthy the system is in fact, which can then be compared to how people perceive it.

To evaluate the relationship, I want to measure how much the human trusts the wheelchair. Although it would be most appropriate to measure this in a long term relationship, this is probably impractical given time and subject constraints. To properly establish a relationship, a physical chair would need to exist and be used for an extended period of time, while for this project I plan to implement it in a simulator. So probably the right thing is to evaluate the first impressions of the chair, after a tutorial session and completing some tasks together.

To do this, subjects will come in and interact with the simulated vehicle. They will complete the tutorial session, and then use it to navigate through the simulated world to solve problems. I am currently running a user study where the chair is replaced by a human driver in order to collect instances of spatial language use. I am using two tasks. First, a search task where the team is searching in a maze for treasure items. Second, a navigation task, where one member of the team is given a map of the maze and instructs the driver on where to go. I will use these same two tasks for the evaluation, replacing the human driver with the program.

After subjects complete the evaluation, they will complete a questionnaire to determine how much they trusted the chair. I would compare this data to data from human-human pairs, to see how the relationship compares. It might make sense to also compare it to human-human pairs where the human driver has been handicapped in some way (perhaps showing a more limited environment or putting them in separate rooms with a poor quality speech feed.) The program would be successful based on how similar the human-machine questionnaire data was to the human-human data.

Week 5

I started a thread on a wheelchair forum asking for feedback on the idea. I've had some interesting responses, although not as many as I'd hoped for.

I would like to interview people although I haven't actually set anything up yet. I think it makes sense to talk to people in wheelchairs. I want to find out whether they think the technology would be useful, and whether they would feel a stigma talking to their machine. I also want to record the kind of mobility aid they're using - how they are currently controlling the powerchair. It would be good to talk to people who use a variety of interfaces.

Interview Questions

- I'm designing a voice controlled wheelchair. Imagine you could say anything to your chair that you could to a person who was pushing it, and it would obey your commands. What do you think of the idea?
- Would you feel self conscious talking to the chair?
- Would you like that interface better or worse than your existing interface?
- In practice, such an interface would make mistakes. It would misunderstand your command, and just do the wrong thing. How would you feel when it made a mistake?
- What if it asked you a clarifying question if it didn't understand your command? For example, it might say "Please repeat that?" Would you feel self conscious if your chair talked back?