Making Pedagogical Agents More Socially Intelligent

Lewis Johnson
CARTE, USC / Information Sciences Institute
Johnson@isi.edu
ftp://ftp.isi.edu/isd/johnson/si/

Advances in agent and Web technologies are making it increasingly possible to incorporate guidebots, or animated pedagogical agents, in a wide range of Web-based learning materials. If designed properly these guidebots can promote deeper learning and improve the learner's subjective experience. Guidebots exploit a person's natural tendency to interact socially with computers, as documented by Reeves, Nass, and their colleagues (Reeves & Nass, 1996).

In order for guidebots to work more successfully, they need to respond to learners in socially appropriate ways. This is an important part of any guidebot design activity, and a key theme of research the Center for Advanced Research for Technology in Education (CARTE) at USC / Information Sciences Institute (Johnson, 2001). In particular, the Social Intelligence Project at CARTE is investigating how to improve the social interaction skills of guidebots. It is developing a capability that manages the learner-guidebot interaction to maintain a positive learner experience, promote learner motivation, and avoid and repair communication breakdowns. We see this capability as potentially useful for any guidebot-enhanced learning application, including Web-based learning environments (Johnson, 2002). It should be particularly important for helping novice learners, who have the greatest need for the kind of assistance that guidebots provide, and who may lack self-confidence. Our initial test application is the Virtual Factory Teaching System, an on-line simulation-based training system for teaching factory management concepts (Dessouky et al, 2002). The social intelligence component will be as an augmentation to the Automated Laboratory Instructor agent (D'Souza et al., 2001), a pedagogical agent designed to coach learners as they work with scientific simulations. An effective social intelligence capability will make it possible to offer the VFTS to a wider range of learners, particularly those in allied fields such as business administration who may be intimidated by complex engineering applications and therefore need the kind of tailored guidance that a socially intelligent guidebot can provide.

Background Research

Our work on social intelligence was motivated in part by our experience with Adele, a guidebot designed for support health science education (Shaw et al., 1999). Adele observes students work through simulated medical cases. If the student takes an inappropriate action, e.g., commence a patient examination without having adequately reviewed the patient's medical history, Adele would interrupt and critique the learner's action. This was acceptable in individual cases. However if the student made multiple mistakes Adele would repeatedly interrupt and criticize the learner's actions in the same fashion. This kind of treatment would give learners the impression that Adele had a very stern personality and had low regard for the student's work, none of which was intended either by Adele or her designers. Although we have ameliorated these effects by adjusting the wording of Adele's comments, a basic problem remains: that Adele is unaware of the effects that her comments have on the learner. She does not know whether the learner accepts her comments, is annoyed by them, or is discouraged or even offended by them. These effects cannot be predicted with certainty, since they may depend upon the learner's knowledge and individual characteristics, current activities, and the history of previous interactions. To take these factors into account in communicating with learners clearly requires social intelligence.

To better understand the role of social intelligence in tutorial interaction, we videotaped students and human tutors working together through on-line tutorials and problems using the VFTS. Subjects read though an on-line tutorial in a Web browser, and performed actions on the VFTS system as indicated by the tutorial. Our goal was to draw lessons from such interactions that could be applied to the design of a guidebot that helps students with a system such as VFTS. Although a number of studies (e.g., [Chi et al., 2001]) have been conducted of students working with tutors, in these studies the students and tutors typically use pencil and paper materials, not computer-based materials. We were interested in observing how the tutors respond to the student's actions on the computer, since an automated guidebot would need to respond in a similar way. For example, a tutor needs to be able to judge when is a good time to interrupt the student, based upon what the student is doing on the computer from moment to moment. The tutor needs to be able to fit his tutorial interactions into the context of the student's activities, unlike face-to-face tutoring where the tutor is often in control of the interaction.

We chose as the primary focus of this phase of the study a tutor who was an industrial engineering professor who

had won awards for excellence in teaching. Although his teaching obligations do not require him to do much one-on-one tutoring, our hope was that his general teaching skills would also correlate with tutoring expertise. We plan to expand the study to other tutors as needed, to obtain more examples of socially intelligent tutorial interaction. Although these studies reveal a wide variety of tutorial interactions, and illustrate some principles that I expect apply in a general way to student-guidebot interaction, we can expect that other students and other tutors might in different ways from what is illustrated here.

General Observations

In these sessions we observe a number of techniques that the tutor uses to promote learner interest, engagement, and sense of self-confidence. The tutor emphasized that the simulated factory was the student's factory, and that the factory planning decisions were the student's, or were being performed jointly by the student and the tutor. For example, to get the student started in working through the tutorial the tutor said "OK, why don't you go ahead and read your tutorial factory."

Advice and feedback were phrased as suggestions, comments, and questions, rather than explicit criticisms or evaluation of the student's action as right or wrong. These were phrased in a number of different ways. Examples of different interactions are illustrated in the following section, "Characteristics of Tutor Interaction Tactics."

The tutor was able to decide how to interact with the learners in part by observing the learners and comparing their actions against their expectations of what the learner should be doing next. The tutor could watch the students work through the tutorial, and when they failed to do what the tutorial suggested the tutor would intervene.

Some tutorial exchanges were contingent on the overall problem solving session context. For example, when the students began follow-up problem solving sessions, the tutor would summarize what was accomplished in the previous session and what needs to be achieved in the following session.

Many tutorial exchanges were triggered by the students' questions. The student would ask a question, and the tutor would answer the question. Answering the question might require the tutor to mention a number of individual points, each of which might be acknowledged or questioned by the learner.

Looking at these sessions, we see big differences in how each student interacts with the tutor. One student reported enjoying interacting with the tutor, and felt that the VFTS system would not have been as fun and interesting if another person were not present. This student communicated freely with the tutor, and frequently asked for confirmation that she was solving the problem correctly. She also asked many questions, which led the tutor to communicate more in response. She initially had very little confidence in her ability to perform the task, and seemed to have quite a bit uncertainty as to whether she was performing the task correctly. By the end of the second session, though, she was much more confident and self-motivated, and furthermore was able to answer the post-test questions correctly. The other student had more confidence all along, initiated communication with the tutor much less frequently, and responded much more briefly, sometimes with just a nod of the head. The tutor needed to prompt this student on occasion to find out whether she had any questions. Yet this student performed more poorly on the post-test. As it turned out, these students also scored very differently on the Myers-Briggs Type Indicator, suggesting that they had very distinct personality types. Yet although the students were very different in personalities and work styles, we did not see a big difference in the style in which the tutor interacted with each tutor; the main differences were that the tutor was more communicative and initiated more exchanges with the less talkative student. This raises a number of questions that will be investigated in future studies.

- It turns out in this case that the less confident student was more talkative. How would the tutor have interacted with an unconfident, untalkative student?
- Would the less confident student have responded better to a tutor that was more reassuring? If so I expect that the difference in performance would not have been great, since the student gained confidence over time in any case and became very self-motivated.
- Would the students have performed better if the tutor had adapted his interaction style more according to the personality characteristics of the learner? The learners' performance suggests that this might be so. The tutor needs to be willing to initiate conversation even with students who are not very talkative. However, there is a risk that too much conversation could be off-putting for some students. Post-session interviews with the tutor, observations of the tutorial interactions, indicate that the tutor recognized when the student wanted to be left alone.

Interaction Tactics

Tutorial dialog exchanges in these sessions can be characterized as a series of interaction tactics, where each tactic is intended to communicate particular information or have a particular effect on the listener. When a dialog exchange is initiated the speaker looks to see that the

intended effect has been achieved, and if not rephrases the comment accordingly. We see this in the following student-tutor exchange. The tutor is watching the student read the tutorial, notes that the tutorial calls for the student to perform a regression analysis using the VFTS, waits for the student to do it, and when the student fails to carry out the step follows up with more specific advice about what to do.

Tutor: So it's asking for regression
Student: Right, that wasn't an option
(*Student clicks a few times on the wrong button on the screen, then stops*)
Student: There's no place...
Tutor: You want to click on regression here and make sure that it matches the button that's up there.

Interaction tactics frequently fail to have the intended result, as in this example. People seem to be quite good at recovering from such failures, both to make communication more effective and to maintain a fluid social interaction between the dialog participants. Likewise a guidebot needs to be prepared for failures in interaction tactics, and adapt responses accordingly.

Some of the interaction exchanges involve multiple interactions regardless of the responses of the learner. For example, at the beginning of the second lesson the tutor reviewed what was accomplished in the previous lesson, and multiple utterances were required to do this. Thus we sometimes need to organize individual interaction tactics into higher-level dialog structures. We can model such interactions as complex plans composed of multiple individual interaction tactics. Each interaction tactic may or may not achieved the intended effect, requiring further adaptive responses from the tutor. Or, an interaction exchange may be a mixture of utterances and task steps. For example, when the tutor guides a student through the process of running the factory simulation the tutor guides the student through the sequence of steps involved in running the simulation, commenting and answering questions along the way.

Characteristics of Tutor Interaction Tactics

The following are some important characteristics of the interaction tactics initiated by the tutor in these sessions.

Unsurprisingly, many tutor-initiated interactions are offers of hints. What was surprising, though, was that these hints were expressed in a variety of different ways, and relatively few were explicit instructions to perform a particular action. Some were expressed as questions, as in the following examples:

Tutor: Want to look at your capacity?
Tutor: Do you want to move that over so that way you'll be able to see this while you're doing it at the same time?

Some hints were phrased as suggestions, expressed conditionally so that the student could decide whether or not to follow the tutor's suggestions, e.g:

Tutor: So you could move down and do the basic parameters.

Consistent with the view of the tutor's comments as interaction tactics, nearly all tutor comments elicited a verbal response or acknowledgment from the learner or a nonverbal action to follow the tutor's suggestion. The only observed exceptions to this were as follows:

- The tutor was articulating a general principle that applied to the current situation but not specifically to it, e.g., "You want to save the factory every time you change it."
- The tutor's question was a rhetorical lead-in to subsequent comment, e.g., "Can I give you a hint? On where…"
- The tutor's comment provided context or motivation for suggestion that followed immediately after, e.g., "And you can get individual statistics on products… click on here and get it."

Some hints were phrased as comments of what the tutor would do, rather than suggestions of what the student should do, as in the following:

Tutor: I'd go to the very top.

In other cases the hint was a suggestion of what the student and tutor might do together, e.g.:

Tutor: So why don't we go back to the tutorial factory and work on the planning there.

In yet other cases, the hint was expressed as if a comment about what the learner's goal should be, as if she already had that goal, e.g.:

Tutor: You wanna do one at a time, go back and change it, …

These tactics have the effect of involving the learner in the decision making process, as well as making the activity as shared student-tutor activity. In multi-turn interactions the tutor would typically switch agent stances between turns, enhancing the effect. Here is an example of switching from student goal to tutor goal:

Tutor: … you can tweak it as you go.
Student: Yeah.
Tutor: That's what I'd do.

Here is an example of switching from student goal to joint goal:

Tutor: You want to also read the tutorial. It gives you some—
Student: The what?
Tutor: The tutorial. Let's see what they say.

Towards a Socially Intelligent Interface

We are now drawing from these lessons in order to build a social intelligence model and a guidebot interface that is

able to interact with learners in a more socially appropriate way.

Figure 1 illustrates how the SI model will integrate with the learning environment and other agent modules. As the learner interacts with the learning environment (i.e., simulations and Web pages), learner actions are recorded and passed to an action analysis module that assesses both the correctness of learner actions and the cognitive demands of the learner, based on the difficulty of the activity and the amount of attention required by the learning environment. The SI model uses this information to assess whether pedagogical interventions are appropriate in the current context. The SI model also receives input in the form of visual information and conversational inputs. Visual information is processed using a gesture understanding module that receives image data, recognizes facial gestures, head pose, and point of gaze, and correlates point of gaze against GUI layout to determine focus of attention on the screen. The visual tracking technology is being developed by USC's Laboratory for Computational and Biological Vision, under the direction of Christoph von der Malsburg.

Conversational inputs are user questions and requests, and answers to questions posed by the agent, entered via menus and limited natural language. The SI model controls the presentation of the pedagogical agent by issuing body gesture commands to the agent's animated persona, and sending marked up text to a text-to-speech synthesizer for presentation as a combination of text and speech.

Although this architecture is the ultimate goal, we need to approach this goal incrementally. Although we expect that the above analysis of face-to-face tutorial interaction has important implications for student-guidebot interaction, we need to determine to what extent we can emulate face-to-face interaction with a guidebot, and to what extent students will respond to socially intelligent guidebots as they do to people. The following is a description of the steps that have been taken so far, and the steps that will be underway during this symposium.

The first priority is developing the input and output interfaces, so that we test the interaction model. We used the Digital Puppets / WebTutor architecture of Rizzo, Johnson, and Shaw (2002) as the basis for the interface. The WebTutor is used to annotate Web pages with JavaScript controls that initiate interactions with an animated character, or "digital puppet." In the new SI configuration, the tutorial Web page is enhanced with controls that track which paragraphs that student is currently viewing, and allows the learner to click on phrases in the text and ask questions about them, either chosen from a menu of questions relevant to that topic or typed in directly. These questions will be passed to the guidebot control mechanism which will then use the digital puppet to generate a response. The Virtual Factory Teaching System was also instrumented, so that each interface action (menu selection, button click, etc.) is made available to the SI system.

Based upon the analysis of tutorial interaction described above, Wauter Bosma in our group developed a natural language generator described to produce a range of verbal interaction tactics. To use the NL generator one selects the type of tactic and the desired characteristics of the tactic. The tactics themselves are represented internally as collections of features in the DISCOUNT scheme for marking up tutorial dialogs (Pilkington, 1999), extended with additional features to capture the distinctions observed our transcripts. Given a set of content descriptions and a desired interaction tactic, the generator generates appropriate text; features that are not specified n the tactic description are chosen randomly over a probability distribution. The text is then passed to the digital puppet, which utters the text using a text-to-speech synthesizer.

Our next task is to evaluate the adequacy and coverage of the text synthesizer with human tutors, to see whether the repertoire of interaction tactics is sufficient and whether the text realizations of the tactics are adequate. We then plan some Wizard-of-Oz studies in which the learners use the interface as if it were controlled by a guidebot, but the choice of tactics and responses is actually controlled by an experimenter at another console. This will help us to compare the style of interaction as well as the learners' subjective experience against face-to-face interaction. One question whether visual input is essential, either to determine when to initiate a tutorial exchange or to monitor the learner's responses to those exchanges. Another question is to what extent the learners perceive the tutor as being responsive to the leraner. Based upon the results of these studies we can then proceed to automate the selection of interaction tactics by the guidebot.

Acknowledgements

The author wishes to acknowledge the contributions of following people to the SI project: Wauter Bosma, Cynthia Chan, Heather Collins, Maged Dessouky, Jason Finley, Lawrence Kite, Kate LaBore, Christoph von der Malsburge, Richard Mayer, Kazunori Okada, Helen Pain, Lei Qu, Jeff Rickel, and Erin Shaw, and Ning Wang. This material is based on work that was supported in part by the National Science Foundation under Grant No. 0121330, and in part by a grant from Microsoft Resarch. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author and do not necessarily reflect the views of the National Science Foundation

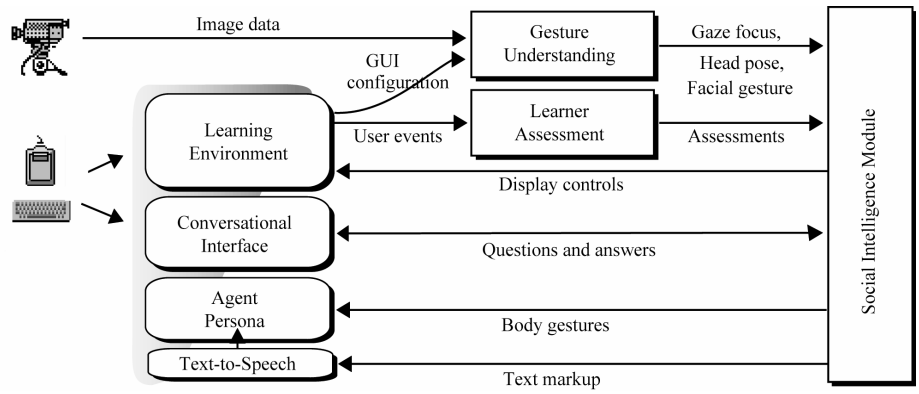

Figure 1. Architecture of the socially intelligent learning system

References

Chi, M.T.H., Siler, S.A., Jeong, J., Yamauchi, T., Hausmann, R.G. (2001). Learning from human tutoring. *Cognitive Science* 25, 471-533.

Dessouky, M.M., Verma, S., Bailey, D., & Rickel, J. (2001). A methodology for developing a Web-based factory simulator for manufacturing education. *IEE Transactions* 33, 167-180.

D'Souza, A., Rickel, J., Herreros, B., & Johnson, W.L. (2001). An automated lab instructor for simulated science experiments. In J.D. Moore, C.L. Redfield, and W.L. Johnson (Eds.), Artificial Intelligence in Education: AI-ED in the Wired and Wireless Future, 65-76. IOS Press, Amsterdam.

Johnson, W.L. (2001). Pedagogical agent research at CARTE. *AI Magazine* 22 (4), 85-95.

Johnson, W.L. (2002). Using agent technology to improve the quality of Web-based education. In N. Zhong and J. Liu (Eds.), *Web Intelligence*, forthcoming. Berlin: Springer. Also available at ftp://ftp.isi.edu/isd/johnson/si/.

Pilkington, R.M. (1999). Analysing educational discourse: the DISCOUNT scheme. Technical report No. 99/2, Computer Based Learning Unit, Univ. of Leeds.

Reeves, B. and Nass, C. (1996). *The Media Equation*. New York: Cambridge University Press.

Rizzo, P., Shaw, E., & Johnson, W.L. (2002). An agent that helps children to author rhetorically structured digital puppet presentations. In S.A. Cerri, G. Gouardères, & Fábio Paraguaçu (Eds.), Proceedings fo the 6th Intelligent Tutoring Systems Comference, 903-912. Berlin: Springer.

Shaw, E.; Johnson, W. L.; and Ganeshan, R. (1999). Pedagogical agents on the Web. In *Proceedings of the Third International Conference on Autonomous Agents*. New York: ACM Press.