Pedagogical Agent Research at CARTE

W. Lewis Johnson

■ This article gives an overview of current research on animated pedagogical agents at the Center for Advanced Research in Technology for Education (CARTE) at the University of Southern California/Information Sciences Institute. Animated pedagogical agents, nicknamed *guidebots,* interact with learners to help keep learning activities on track. They combine the pedagogical expertise of intelligent tutoring systems with the interpersonal interaction capabilities of embodied conversational characters. They can support the acquisition of team skills as well as skills performed alone by individuals. At CARTE, we have been developing guidebots that help learners acquire a variety of problem-solving skills in virtual worlds, in multimedia environments, and on the web. We are also developing technologies for creating interactive pedagogical dramas populated with guidebots and other autonomous animated characters.

The purpose of the Center for Advanced Research in Technology for Education (CARTE) at the University of Southern California (USC)/Information Sciences Institute (ISI) is to develop new technologies that promote effective learning and increase learner satisfaction. These technologies are intended to result in interactive learning materials that support the learning process and that complement and enhance existing technologies relevant to learning such as the World Wide Web.

Our work draws significant inspiration from human learning and teaching. We seek to understand and model how people learn from their teachers, peers, and their environment and draw on this knowledge to guide our development of educational technologies. We seek a better understanding of the characteristics that make learning experiences captivating and exciting and attempt to find ways of fostering these characteristics more broadly and systematically. Our work thus builds on research in AI and cognitive science and, at the same time, draws on the experience of specialists in the arts.

A Major Theme: Guidebots

A broad theme of our work has been the creation of *guidebots,* or animated virtual guides for learners. We also refer to these guidebots by the more lengthy term of *animated pedagogical agents* (Johnson, Rickel, and Lester 2000). Guidebots are animated characters that can interact with learners in computer-based interactive learning environments to stimulate and encourage learning. In their complete implementation, they have a number of important features that are relevant to CARTE's goals. They interact naturally with learners, generally in a manner that is inspired by the behavior of human tutors; these interactions include a combination of verbal communication and nonverbal gestures. They express both thoughts and emotions; emotional expression is important to portray characteristics of enthusiasm and empathy that are important for human teachers. They are knowledgeable about the subject matter being learned, of pedagogical strategies, and also have knowledge about how to find and obtain relevant knowledge from available resources such as the World Wide Web.

Figure 1 shows one of the guidebots that we have developed, named STEVE, or SOAR TRAINING EXPERT FOR VIRTUAL ENVIRONMENTS (Rickel and Johnson 1999a). This shot was taken as STEVE explains to the user how to operate a particular piece of equipment called a high-pressure air compressor aboard United States Navy ships. Note how it engages in face-to-face communication, directing its gaze toward the user as it

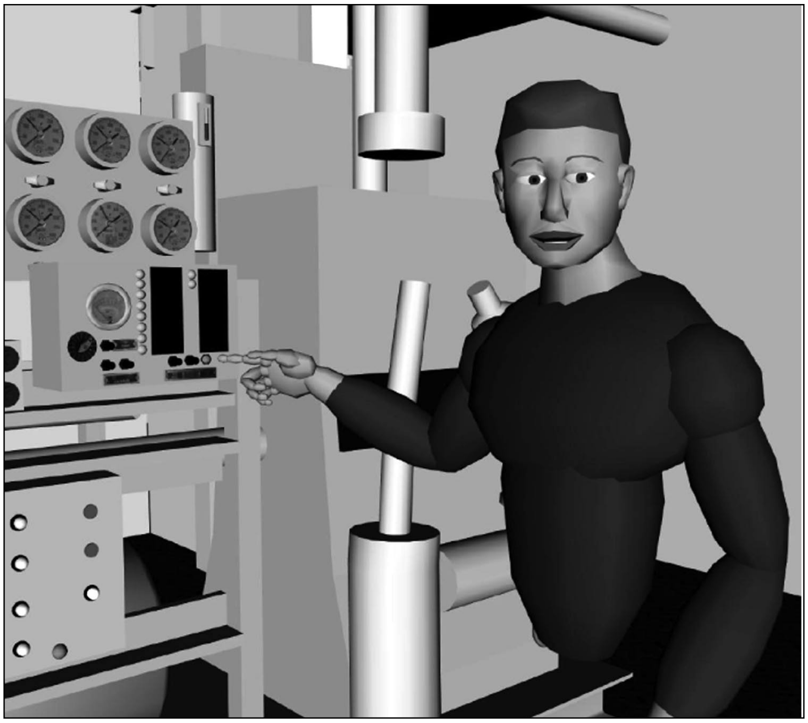

Figure 1. STEVE *in a Virtual Environment.*

explains what to do. It is able to demonstrate to the user how to manipulate the device and then point to device components such as indicator lights that are affected by these manipulations. The resulting interaction is similar in many respects to the interactions between a learner and a human coach.

Ultimately, we see guidebots as elements of a variety of rich interactive experiences. In these experiences, students have significant freedom to explore learning materials and perform activities that reinforce their learning, either individually or in groups. The guidebots serve a number of important functions in these environments.

They help keep the learning on track. People have a tremendous capacity to learn, but they fail to use this capacity effectively if they fail to take advantage of the resources available to them or if they fail to apply the correct metacognitive skills. Guidebots can remind learners of available resources and can reinforce good learning habits. They can offer help and guidance when the students get stuck and provide feedback to the learners on their progress.

They provide an additional communication channel. Guidebots are a user interface component and play important roles within an overall educational user interface design. They are extremely effective at directing the learner's attention to what is important in the learning environment. As learners view instructional materials, guidebots can provide useful commentary on these materials. They can also provide learners with help in navigating complex learning environments. They provide both highly salient and highly nuanced interaction. Human personae have much greater expressive capabilities than conventional user interface elements, and people are well practiced in reading these expressions.

They act as the teacher's representatives. They are able to interact with the learners at times when a human teacher is not available. At the same time, they collect data about their interactions with the learners that can be valuable to teachers in assessing the learners' skills and planning future computer-based learning experiences.

In the overall environment envisioned here, guidebots interact with a variety of human and automated agents. The following are some important roles for these agents that the environment should support.

Supporting characters: Many learning experiences include additional characters, such as synthetic team members or inhabitants of virtual worlds. These do not act as the learner's guides, but they nevertheless support the learning objectives of the experience through their interactions with the learners and each other.

The director: Borrowing the dramatic metaphor, it is important to have a director who can guide the overall action. In some applications, a human is in the loop to help direct the action, and in other applications, we want the direction to be performed automatically. The director needs to assess whether the current interaction, or "scene," has achieved its intended instructional function and when it is time to move to another scene. The director might influence the way the guidebots and supporting characters interact with the learners depending on how the action unfolds. If the director determines that the current learning objectives have been met, it might guide the interaction toward situations that address new learning objectives. The director thus requires real-time planning and assessment capabilities. In some applications, the director also needs to control the visual presentation of the scene, thus serving as cinematographer.

The author: Guidebot technology will become practical only if it becomes easy to design and create interactive experiences that incorporate guidebots. New kinds of authoring tools are needed that support the unique characteristics of guidebot-enhanced learning experiences and that take advantage of the learning materials that are available on the World Wide Web, in digital libraries, and elsewhere.

Guidebots can be integrated into a variety of interactive media, and the resulting interactive experiences can be delivered in different ways. The following are the media that are of interest to CARTE. They currently require distinct technologies; however, technology advances are blurring the distinctions between them.

Virtual environments: The action takes place within a virtual world. Learners can manipulate objects in the virtual world and, through their manipulations, can practice performing tasks. Each learner has a presence in the virtual world in the sense that guidebots and other participants can observe them and interact with them. Three-dimensional (3D) worlds such as STEVE's immersive mock-ups are prime examples. However, two-dimensional (2D) interactive simulations used by web-based guidebots can also serve as virtual environments, insofar as the guidebot can observe the learner's manipulation of objects on the screen and react to them.

Web-based information spaces: The World Wide Web is a common medium of instruction, and therefore, we design guidebots to operate in conjunction with this medium. In general, the web does not support a strong spatial geometry, and the learners don't have a spatial location within it. Three-dimensional spaces can be found on the web, but they constitute a small fraction of web sites at this time. To assist web-based learning, guidebots receive notifications when learners view pages and click on links on the pages and can respond to these actions.

Interactive pedagogical dramas: The software presents a dramatic story, using sound and images. Learners can influence what takes place in the drama, and the characters can adapt their behavior in response. If the drama is presented in a virtual environment, then the learners can play roles within the story themselves. However, added flexibility can be attained if there is some separation between the learner and the action, that is, where the computer screen frames the action. The learner then shares duties with the automated director in guiding the action. This structure creates opportunities to present back-story material as needed and control the learner's view through cinematographic techniques. In such scenarios, the guidebot typically acts as buddy or adviser to the character or characters that are being directed by the learner.

Embedded training environments: Ultimately, we want to incorporate guidebots into real-world training environments. This incorporation is possible if the training environment is instrumented so that the guidebots can detect what actions the learners are performing. For example, if the learners are operating equipment, then the guidebots might be able to monitor the state of the equipment through a software interface. In such environments, guidebots can be displayed on computer monitors placed in the environment or can be superimposed over the environment using augmented reality goggles. It is even possible to use physical robots as embodiments for the guidebots.

Finally, guidebots can serve useful instructional roles in environments whose purpose is not primarily instructional. For example, guidebots can be embedded in software packages to provide just-in-time training and in interactive games to teach users how to play. Whenever users have particular tasks to perform or problems to solve, guidebots can potentially provide assistance.

A number of CARTE projects are conducted in collaboration with other research groups in computer science and in educational application areas. We are working with USC's new Institute for Creative Technologies to develop more engaging immersive training experiences. Other collaborators include health science research centers' educational psychologists and teaching faculty at USC and elsewhere. These collaborations are essential both for identifying the deficiencies in current teaching practice that guidebots can address and assisting in the evaluation of guidebot applications in instructional settings.

Example Guidebots

CARTE has developed a number of animated pedagogical agents that have served as vehicles for developing and testing key aspects of the underlying technologies. [1]

STEVE

STEVE assists procedural skill training in immersive virtual environments, such as virtual mock-ups of engine rooms, as in figure 1. STEVE supports both individual and team training; it can advise learners as they perform roles within teams, and it can play the role of a missing team member (Rickel and Johnson 1999b). Figure 2 shows two STEVE agents working together, one observing and the other operating a console.

STEVE operates within a networked virtual environment, interoperating with a collection of other components including the VISTA VIEWER

Guidebots can remind learners of available resources and can reinforce good learning habits. They can offer help and guidance when the students get stuck and provide feedback ….

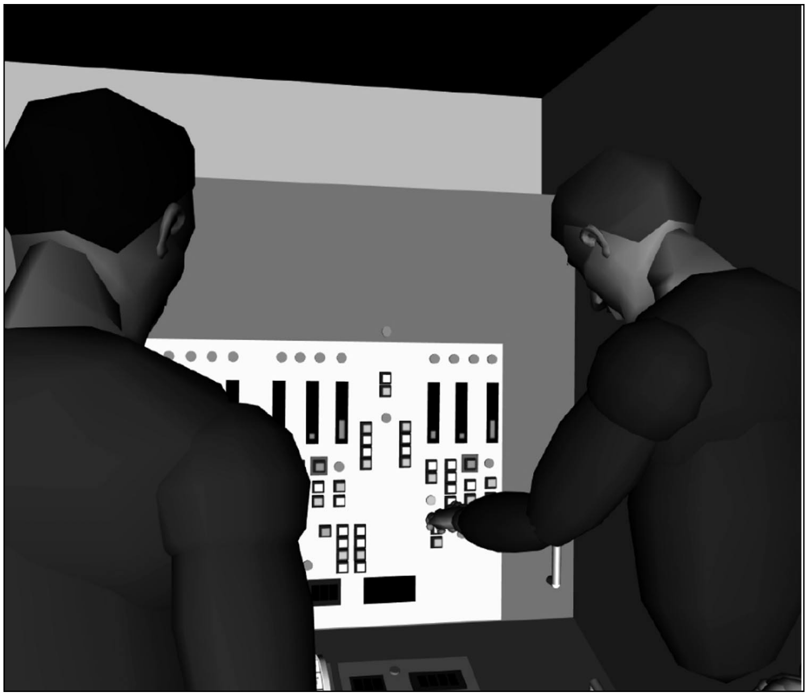

Figure 2. STEVE *Pointing at an Indicator Light.*

visualization system of Lockheed Martin and the VIVIDS simulation authoring tool (Johnson et al. 1998). One set of components—a copy of VISTA VIEWER, a component that plays the sounds of the virtual environment, and a speech recognizer and a speech synthesis engine for communicating with STEVE—is assigned to each student or instructor in the environment. One copy of VIVIDS maintains the simulation state of the virtual environment. One or more STEVE cognitive modules control the actions of the STEVE virtual bodies, which, in turn, are implemented as graphics routines linked into the VISTA VIEWER software. These components exchange messages about the state of the virtual world and the actions of the human and virtual agents in the world, so that all participants have a consistent view of the virtual world.

STEVE uses the SOAR cognitive architecture (Laird, Newell, and Rosenbloom 1987) to model adaptive task execution in dynamic environments. Like other intelligent tutoring systems, it can monitor students' actions, point out errors, and answer questions such as What should I do next? and Why? However, because it has an animated body and cohabits the virtual world with students, it can provide more humanlike assistance than previous automated tutors could (Rickel and Johnson 2000). For example, it can demonstrate actions, use gaze and gestures to direct the student's attention, and guide the student around the virtual world. When playing the role of a student's teammate, STEVE's human body allows students to track its activities as they would track those of a human teammate. Finally, STEVE's agent architecture allows it to robustly handle a dynamic virtual world, potentially populated with people and other agents; it continually monitors the state of the virtual world, always maintaining a plan for completing its current task and revising the plan to handle unexpected events.

STEVE relies on a model of the task being performed, represented as hierarchical nonlinear plans. The plan representation indicates what actions must be performed to carry out the task, what effects each operation is intended to have, and what other actions depend on those actions being performed. An action can either be a primitive operation on the virtual world, such as pressing a button, or another subplan. Each action is annotated to indicate which team role is expected to perform the action, allowing all STEVE agents working the team to track the team's progress in achieving the goals of the task and decide what action is appropriate to perform next. The STEVE agent can then perform the action itself; wait for the student to perform the action; or explain to the student what to do and why, as determined by the STEVE's role and the state of the tutorial interaction.

In addition to STEVE's novel instructional capabilities, it was designed to simplify the development of new training applications. Unlike many computer-aided instruction programs, STEVE's instructional interaction with students need not be scripted. Instead, STEVE includes a variety of domain-independent capabilities for interacting with students, such as explanation generation and student monitoring. STEVE can immediately provide instruction in a new domain given only simple knowledge of the virtual world and a description of the procedural tasks in the domain. This approach results in more robust interactions with students than course designers could easily anticipate in scripts.

Web-Based Guidebots

A number of guidebots for web-based instruction have been developed. They rely on a common distributed architecture and share component modules, which facilitates the creation of new guidebot applications. All focus on helping students acquire particular types of problem-solving skills, such as planning or diagnostic reasoning. Students work through interactive exercises, receiving guidance, feedback, and evaluation from the guidebot as needed.

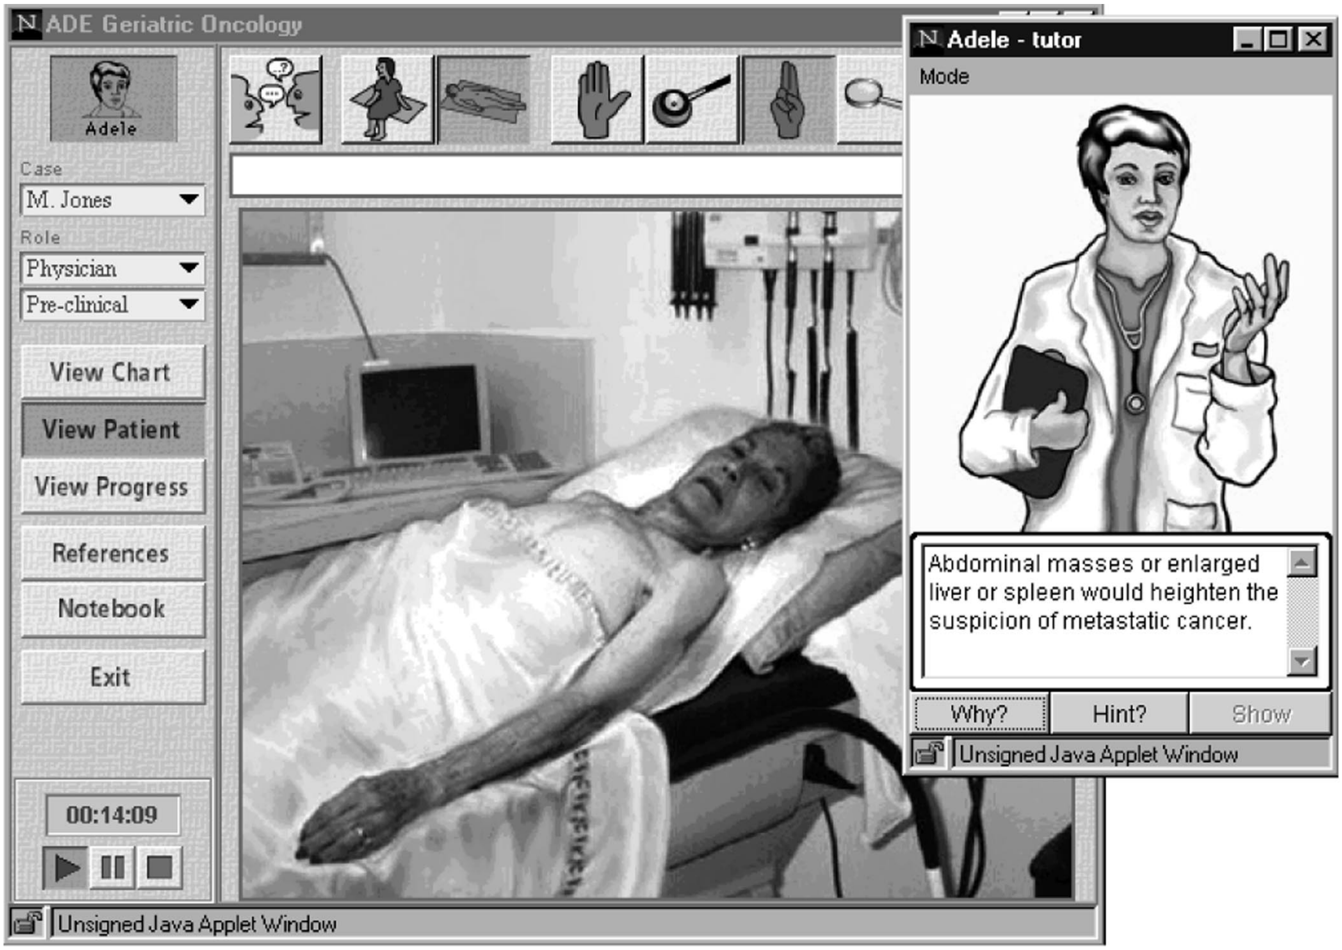

Case: M. Jones

Role: Physician

Pre-clinical

Abdominal masses or enlarged liver or spleen would heighten the suspicion of metastatic cancer.

Figure 3. ADELE *Guiding a Student through Clinical Case.*

Our web-based guidebots are designed to complement, and overcome limitations in, web-based instruction. The web is an excellent source of information for learners. However, learners need more than information; they need opportunities to practice problem-solving skills, and they need feedback on their performance. Instructors can sometimes provide such feedback through e-mail, but learners are then required to wait for the e-mail response, and instructors cannot be expected to be available at all times to answer questions. Guidebots, however, can provide feedback and answer routine questions at any time. In a typical use, guided exercises and web-based information are closely linked together. A learner might start by reviewing web pages on a topic and then completing an exercise assisted by a guidebot. As the student works through the exercise, the guidebot is able to assess the student's mastery of the material that has just been learned and can point the learner back to the lesson material or to supplementary material as appropriate.

ADELE

ADELE (AGENT FOR DISTRIBUTED LEARNING ENVIRONMENTS) supports on-line case-based problem solving, particularly in the health sciences (Shaw, Johnson, and Ganeshan 1999). ADELE monitors the student as he/she works through a simulated case, downloaded to the student's client computer. As the student proceeds, ADELE compares the student's actions against a model of how the task ought to be performed. The task model is in the form of a set of hierarchical partial-order plans, similar to STEVE's representation, with different plans applying to different situations. ADELE can give hints, explain the rationale for recommended actions, point the learner to relevant online references, and intervene if the student is making serious mistakes. Figure 3 shows ADELE explaining the medical rationale for a recommendation to palpate the patient's abdomen, for example. The amount of intervention can be changed depending on the instructional objectives of the module and the needs of the student. During the interac-

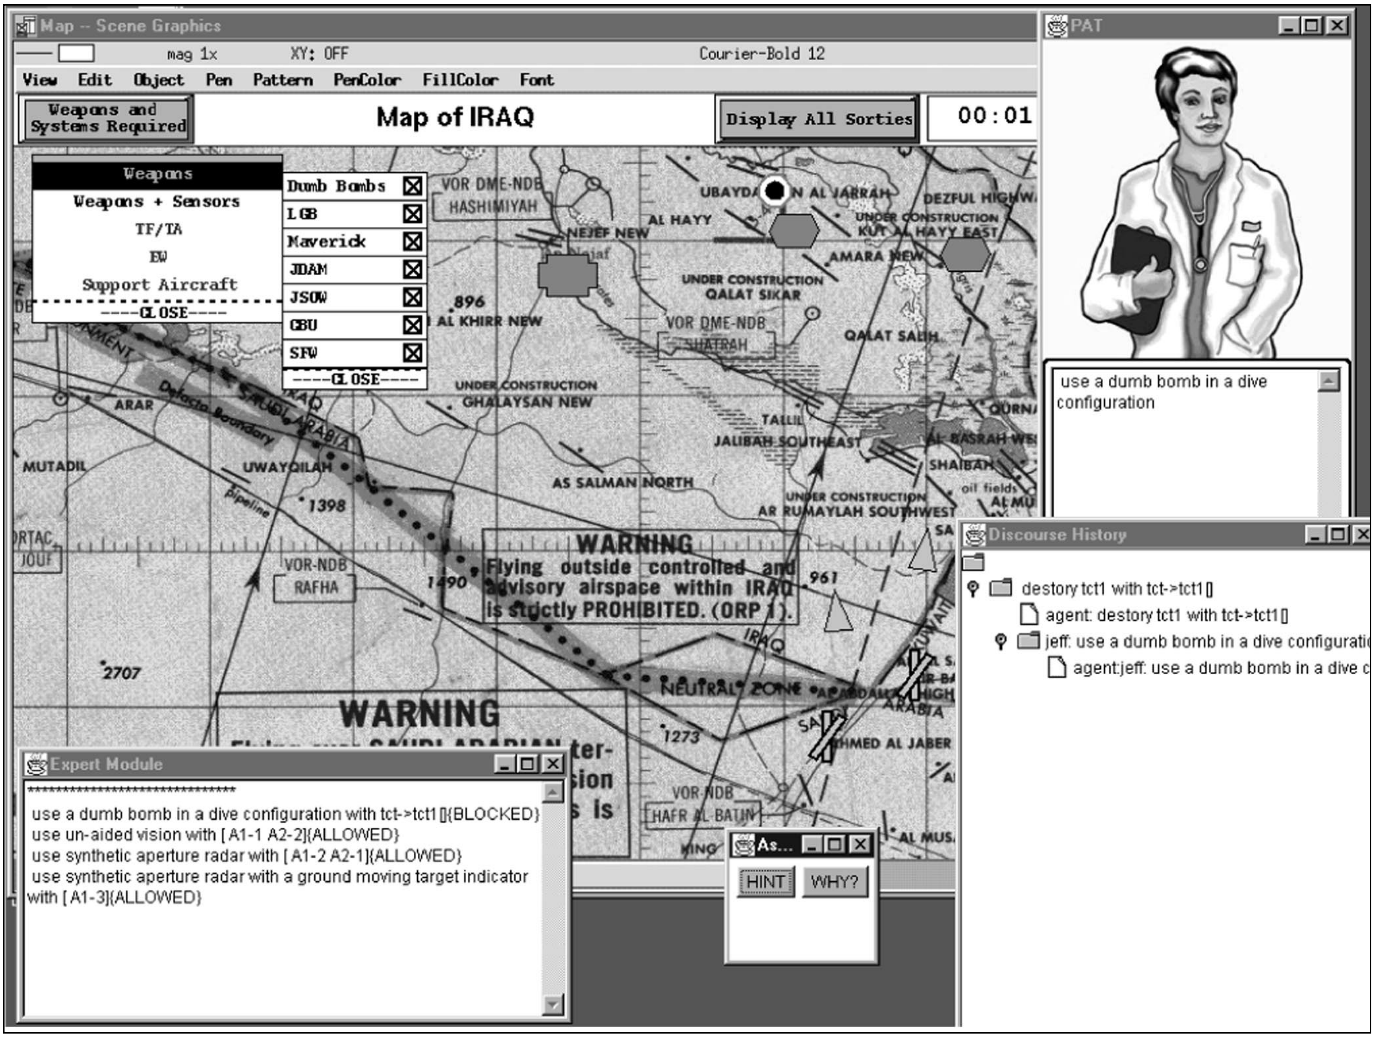

use a dumb bomb in a dive configuration

destory tct1 with tct->tct1 []
agent: destory tct1 with tct->tct1 []
jeff: use a dumb bomb in a dive configurati
agent:jeff: use a dumb bomb in a dive c

```
******************************
use a dumb bomb in a dive configuration with tct->tct1 []{BLOCKED}
use un-aided vision with [ A1-1 A2-2]{ALLOWED}
use synthetic aperture radar with [ A1-2 A2-1]{ALLOWED}
use synthetic aperture radar with a ground moving target indicator
with [ A1-3]{ALLOWED}
```

Figure 4. PAT *Applied to Time-Critical Targeting.*

tion, ADELE currently responds to the learner using a combination of verbal and nonverbal gestures.

The ADELE architecture supports alternative student monitoring engines for different types of problem solving. The original engine, developed around hierarchical partial-order plans, is well suited for training students to carry out clinical procedures, which have a substantial procedural element. However, it is of limited effectiveness in critiquing the student's diagnostic problem-solving skills. ADELE could explain to the learner what clinical tests to perform and why these tests were important, but it only had a limited ability to critique the student's diagnostic conclusions from the evidence gathered during the case.

To address this concern, Rajaram Ganeshan in our group developed an alternative engine that focuses on modeling diagnostic reasoning (Ganeshan et al. 2000). In this model, findings and hypotheses are linked together in a Bayesian network. The overall network, which could be quite large, is simplified and abstracted to include just those findings and hypotheses that are relevant to the particular case, resulting in a relatively small Bayesian network that can be downloaded quickly to the student's computer and executed. A diagnostic strategy module monitors the student's evidence-gathering actions and tries to infer the hypothesis that the student is trying to evaluate; if this hypothesis is not clear, ADELE will require the student to explain what hypothesis he/she is currently following. The student is free to pursue any hypothesis, as long as the student pursues each hypothesis thoroughly enough to determine with high confidence whether it applies to the current case.

In evaluating the diagnostic strategy version of ADELE, we found that students learned effectively when they worked in pairs with ADELE. The students would work together to solve the case and discuss among themselves which

Figure 5. Carmen in Gina's Office.

hypotheses to consider and how best to respond to ADELE's questions. In this configuration, students learn both from their interactions with ADELE and their interactions with each other, and ADELE helps to ensure that the learning episode stays on track.

Other Web-Based Guidebots

PAT (PEDAGOGICAL AGENT FOR TRAINING) (Rickel et al. 2000) is an implementation of STEVE's procedural skill-training representation compatible with the ADELE architecture. PAT incorporates a model of tutorial dialog, enabling it to place its interactions with the learner in the context of a coherent dialog. Figure 4 shows a prototype application of PAT in the domain of time-critical target planning. This initial version simply reused the existing body artwork for ADELE; of course, a complete application would use new artwork suitable for a military training context.

ALI (AUTOMATED LABORATORY INSTRUCTOR) (D'Souza et al. 2001) applies the guidebot approach to online science experiments. ALI monitors students as they conduct experiments on simulated physical systems. As the student manipulates simulation variables, ALI interprets the simulation results using a qualitative model of the system, defined using qualitative process theory. ALI can then quiz the student to see whether he/she interpreted the simulation results in the same manner. ALI can also critique the student's experimental technique to make sure that he/she is experimenting with the system thoroughly and systematically.

Gina and Carmen

CARTE is also integrating guidebots into interactive pedagogical dramas. An example of an application using this approach is CARMEN's BRIGHT IDEAS, an interactive multimedia training course designed to help mothers of pediatric cancer patients develop problem-solving skills (Marsella, Johnson, and LaBore 2000). Learners are presented with a multimedia dramatic story about a mother, Carmen, who faces a variety of problems at home and at work relating to her daughter's illness. Gina, a clinical trainer discusses Carmen's problems with her, as shown in figure 5, and helps her to learn how to better address problems and try to solve them. In each run through the story, the user can direct the thoughts and moods of either Carmen or Gina, who then behaves in a manner consistent with those thoughts and moods. The character that the user is not controlling acts autonomously in response to the other character. The unfolding of the story line thus is different for each run through the story.

Any interactive pedagogical drama in our conception requires believable characters capable of portraying roles in the story, one or more guidebots to reinforce the pedagogical goals of the story, and a story that is believable and flexible enough to accommodate the learner's range of possible actions. All these are present at least to some degree in CARMEN'S BRIGHT IDEAS. To create a believable, flexible story, we started with a linear story written by a professional scriptwriter. We then deconstructed the story, identifying the overall organization of the story into scenes, inferring the motivations of the characters in the story, and proposing alternative actions that they might perform instead if their emotional state and attitudes were different. Thus, a story structure resulted that consisted of a sequence of relatively fixed back-story scenes that set the stage for the drama, followed by a highly nonlinear scene in which Carmen and Gina discuss Carmen's problems and consider possible solutions and a set of scenes in which Carmen carries out the proposed solutions. The linear scenes were then created using scripted animation, and we developed the nonlinear scenes using autonomous agent characters.

To implement the nonlinear scenes, we developed an infrastructure for autonomous virtual actors using agent technology. Each onscreen character consists of an animated puppet, capable of speaking a set of prerecorded lines and performing a combination of body and facial gestures, and an agent control that determines what lines to say and what gestures to perform. The spoken lines were recorded by professional voice actors to make sure that the emotional portrayal of the characters was as realistic as possible. The agent control maintains a dynamic model of the emotional state of the character, incorporating factors relating to stress and coping, based on the clinically motivated work of Smith and Lazarus (1990). This model determines, for example, whether Carmen responds emotionally to its situation, for example, by blaming herself or others, or in a problem-directed fashion, by attempting to develop solutions to its problems.

Gina serves as a guidebot in CARMEN'S BRIGHT IDEAS, responding to what Carmen says and attempting to guide the interchange toward successful resolution of Carmen's problems. The Gina agent has a repertoire of possible responses that it can give to Carmen that can help to keep Carmen's problem solving on track. Gina also has a model of the overall structure of the story in terms of scenes and decides when Carmen has made enough progress to warrant a transition from one scene to the next. Thus, in this story, Gina is acting as a virtual director, sharing responsibility for directing the story with the learner, who is influencing how Carmen responds to the unfolding situation. Meanwhile, an off-screen cinematographer agent makes decisions about how to frame the unfolding action, deciding which characters should be included in the shot and how closely to zoom in.

Current Research

The following are some guidebot-related research topics that CARTE is currently investigating.

Increasing Realism of Appearance and Behavior

Currently, our guidebots have a distinctly stylized and artificial appearance. In some applications, such as CARMEN'S BRIGHT IDEAS, this is a deliberate stylistic choice, which can be appropriate when realism is not necessary to help the guidebots achieve the intended instructional objectives. However in immersive environments such as the engine room that STEVE inhabits in figure 1, a realistic appearance is important both to increase believability and to give the guidebot a greater ability to demonstrate skills. We are currently working with USC's Institute for Creative Technologies to improve a new version of STEVE with highly realistic appearance and behavior. STEVE's agent "brain" is being used to control a realistic human figure, developed by Boston Dynamics, coupled with a realistic face model developed by Haptek, Inc. The new STEVE will be able to model complex behaviors and exhibit believable facial expressions. Thus, it will enable STEVE to be employed in new types of training applications, including those where social interaction skills are crucial, such as peacekeeping operations.

Speech Communication

A number of our guidebots use text-to-speech synthesis (TTS) to communicate with learners. STEVE furthermore supports speech recognition. TTS affords flexibility for interactive learning applications because the guidebot can generate utterances dynamically and is not limited to particular utterances that were recorded beforehand. Unfortunately, shortcomings in TTS technology limit its usefulness; guidebot voices can sound monotonous and artificial and can be difficult to understand at times. We are investigating ways of making more effective use of speech-processing technology. We are developing models of how to use prosody effectively in instructional interactions, which involves observing interactions between human teachers and students and recording professional actors as they deliver lines appropriate for instructional interaction. We then are modeling the prosody contours using available TTS packages.

Synthetic Face-to-Face Interaction

STEVE and ADELE both have the ability to engage in face-to-face communication with learners, combining speech and gesture. Face-to-face communication provides them with both direct and subtle means to communicate with learners and provide them with feedback. Gaze, head position, and body position can be used to indicate the guidebots' focus of attention and indicate when they are communicating and to whom. They help make these guidebots appear more aware of their environment and, hence, more believable as intelligent guides. We are now interested in improving the naturalness of our guidebots' face-to-face interaction, using nonverbal gestures to support and regulate face-to-face communication. We also see potential for using body language to depict the mental state of supporting characters and encourage empathy on the part of the learner. This use of body language is exhibited to some extent by Carmen, as illustrated in figure 4, and is something that we continue to investigate. Meanwhile, we continue to develop models of tutorial dialog that can be incorporated into such face-to-face interaction (Rickel et al. 2000).

Student Monitoring

All our work hinges on our agents being able to perform plan recognition in a general sense, to be able to understand what strategies students are following as they solve problems. Our hypothesis is that students can effectively be monitored across a range of applications using a small number of domain-independent monitoring engines focused on particular skills, such as diagnostic problem solving or experimentation. We continue to develop and test this hypothesis. For example, we are currently working with science faculty at California State University at Los Angeles to determine how the ALI experimentation model can be applied to new types of science experiments. In the process, we are testing the limits of qualitative process theory as a method for modeling physical systems and identifying new types of problem-solving skills that are relevant to experimental science.

Team Performance Analysis

The area of team training is important and not well supported by intelligent tutoring systems. We have developed two systems that exhibit important technologies for team performance analysis. One is the PROBES system, which monitors the performance of tank platoons in a simulated armor training exercise (Marsella and Johnson 1998). The other is ISAAC, which analyzes team performance and identifies factors that contribute to team success or failure (Raines, Tambe, and Marsella 2000). Both have a role to play in the team analysis capabilities that we envision. PROBES is effective where a well-defined model of team behavior and team performance objectives exists; this model is encoded in the situation space model that PROBES uses to track the learner. ISAAC is more appropriate when such a model is lacking because it helps to acquire a model from team performance data. We are interested in integrating and developing these team analysis capabilities and applying them in team training exercises using the new realistic version of STEVE.

Models of Emotion

People naturally read emotions into animated characters. Furthermore, emotional response is an essential aspect of face-to-face tutorial interaction. We must therefore design our guidebots so that they impress the learner as having the right emotional responses at the right time. We therefore have been developing dynamic models of emotion and testing various combinations of facial expression and body language to determine how best to convey emotion to learners. We would like to incorporate emotion models into a variety of guidebot characters. The way the guidebot expresses its emotions depends on the instructional context, but it is clear that some emotional response is necessary; otherwise, the learner might read emotions into the character that were not intended by the character's developers.

Models of Personality and Character

Models of emotion are important for believable characters in guidebot-enhanced learning applications, but they are not sufficient, particularly in interactive pedagogical dramas. Pedagogical dramas require characters with personalities, which help to determine how characters tend to react cognitively and emotionally to different situations. Personality characteristics can also influence how characters express their emotions. We also require models of character development, so that the characters' personality attributes can change in appropriate ways over time. We are investigating these issues in the context of a new pedagogical drama project called IMPACT, intended to foster good eating and exercise habits in fourth and fifth graders. The IMPACT game has several main characters, each of which has different personality characteristics that change over time.

Automated Direction

CARMEN'S BRIGHT IDEAS supports nonlinear story interaction by determining what takes place within a scene and when scene transitions take place. It also provides automated cinematography, character direction, and shared directorial control. We are now generalizing these mechanisms so that they can be applied to a range of interactive pedagogical dramas.

Authoring Technology

Current authoring tool research is exemplified by DILIGENT, a system developed by Richard Angros (2000) that learns procedures from a combination of observing expert demonstrations and performing experiments. DILIGENT demonstrates that interactive authoring tools can effectively exploit machine learning techniques. Advanced authoring techniques such as these are important to ensure the broad applicability of guidebot technology; conventional authoring tools designed for scripted instruction are not appropriate for guidebot design. We are engaged in a number of authoring research and development activities. Andrew Scholer is building on Angros's model to support authoring through *interactive tutorial dialog*, tutoring the agent so that the agent can then tutor students (Scholer et al. 2000a, 2000b). A graphic interface to ADELE's plan representation has been developed so that domain specialists can modify ADELE's task knowledge. Finally, we are developing methods for reusing elements of ADELE's cases so that new case-based exercises can be constructed rapidly.

Acknowledgments

CARTE includes the following faculty and staff members: Jeff Rickel, Stacy Marsella, Richard Whitney, Erin Shaw, Kate LaBore, Frederica Stassi, Marcus Thiebaux, and Fanny Mak. Students and interns include Rogelio Adobbati, Shish Aikat, Paul Cho, Teresa Dey, Aaron D'Souza, Daniel Mellet D'Huart, Carlos Fernandez, Stephen Kim, Girish Srinivasan, Brett Rutland, Narayanan Sadagopan, Girish Srinivasan, and Daisy Zhang.

The work described in this article was conducted with support from the Office of Naval Research, Air Force Research Laboratory, United States Army STRICOM, National Cancer Institute, U.S. Department of Agriculture, National Institute of Child Health and Human Development, the National Science Foundation, and internal grants from USC and USC/ISI.

Note

1. Papers describing these projects are listed in the references and are available on the web at www.isi.edu/isd/carte.

References

Angros, R. 2000. An Agent That Learns to Instruct: Exploiting Demonstrations through Experimentation. Ph.D. dissertation, University of Southern California.

D'Souza, A.; Rickel, J.; Herreros, B.; and Johnson, W. L. 2001. An Automated Lab Instructor for Simulated Science Experiments. In *Artificial Intelligence in Education: AI-ED in the Wired and Wireless Future,* eds. J. D. Moore, C. L. Radfield, and W. L. Johnson, 65–76. Amsterdam: IOS.

Ganeshan, R.; Johnson, W. L.; Shaw, E.; and Wood, B. 2000. Tutoring Diagnostic Problem Solving. In *Proceedings of Intelligent Tutoring Systems: Fifth International Conference (ITS 2000),* eds. G. Gauthier and K. VanLehn, 33–42. Berlin: Springer-Verlag.

Johnson, W. L.; Rickel, J. W.; and Lester, J. C. 2000. Animated Pedagogical Agents: Face-to-Face Interaction in Interactive Learning Environments. *International Journal of Artificial Intelligence in Education* 11:47–78.

Johnson, W. L.; Rickel, J.; Stiles, R.; and Munro, A. 1998. Integrating Pedagogical Agents into Virtual Environments. *Presence: Teleoperators and Virtual Environments* 7(6): 523–546.

Laird, J. E.; Newell, A.; and Rosenbloom, P. 1987. SOAR: An Architecture for General Intelligence. *Artificial Intelligence* 33(1): 1–64.

Marsella, S., and Johnson, W. L. 1998. An Intelligent Assistant for Team Training in Dynamic Multi-Agent Virtual Worlds. In *Proceedings of Intelligent Tutoring Systems: Fourth International Conference (ITS 1998),* eds. B. P. Goettl, H. M. Halff, C. L. Radfield, and V. J. Shute, 464–473. Berlin: Springer-Verlag.

Marsella, S.; Johnson, W. L.; and LaBore, C. 2000. Interactive Pedagogical Drama. In Proceedings of the Fourth International Conference on Autonomous Agents, 301–308. New York: Association of Computing Machinery.

Raines, T.; Tambe, M.; and Marsella, S. 2000. Automated Assistants to Aid Humans in Understanding Team Behaviors. In Proceedings of the Fourth International Conference on Autonomous Agents, 419–426. New York: Association of Computing Machinery.

Rickel, J., and Johnson, W. L. 2000. Task-Oriented Collaboration with Embodied Agents in Virtual Worlds. In *Embodied Conversational Agents,* eds. J. Cassell, J. Sullivan, S. Prevost, and E. Churchill, 95–122. Cambridge, Mass.: MIT Press.

Rickel, J., and Johnson, W. L. 1999a. Animated Agents for Procedural Training in Virtual Reality: Perception, Cognition, and Motor Control. *Applied Artificial Intelligence* 13(4): 343–382.

Rickel, J., and Johnson, W. L. 1999b. Virtual Humans for Team Training in Virtual Reality. In *Proceedings of the Ninth World Conference on Artificial Intelligence in Education,* 578–585. Amsterdam: IOS.

Rickel, J.; Ganeshan, R.; Rich, C.; Sidner, C. L.; and Lesh, N. 2000. Task-Oriented Tutorial Dialogue: Issues and Agents. Paper presented at the AAAI Fall Symposium on Building Dialogue Systems for Tutorial Applications, 3–5 November, North Falmouth, Massachusetts.

Scholer, A.; Angros, R., Jr.; Rickel, J.; and Johnson, W. L. 2000a. Teaching Animated Agents in Virtual Worlds. Paper presented at the AAAI Spring Symposium on Smart Graphics, 21–23 March, Stanford, California.

Scholer, A.; Rickel, J.; Angros, R., Jr.; and Johnson, W. L. 2000b. Acquiring and Using Task Knowledge for Training in Virtual Reality. Paper presented at the IUI-2000 Workshop on Using Plans in Intelligent User Interfaces, 9 January, New Orleans, Louisiana.

Shaw, E.; Johnson, W. L.; and Ganeshan, R. 1999. Pedagogical Agents on the Web. In Proceedings of the Third International Conference on Autonomous Agents, 283–290. New York: Association of Computing Machinery.

Smith, C. A., and Lazarus, R. S. 1990. Emotion and Adaptation. In *Handbook of Personality: Theory and Research,* eds. L. A. Pervin and O. P. John, 609–637. New York: Guilford.

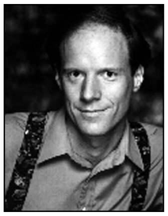

W. Lewis Johnson is research associate professor in computer science and director of the Center for Advanced Research in Technology for Education at the University of Southern California/Information Sciences Institute. He received his Ph.D. in computer science at Yale University and his undergraduate degree in linguistics at Princeton University. His research interests center on advanced technologies for education and training, with a particular focus on applications of animated agents and virtual humans. He is past chair of the ACM Special Interest Group for Artificial Intelligence, past president of the International Artificial Intelligence in Education Society, and program cochair of the 2002 International Joint Conference on Artificial Intelligence and Multi-Agent Systems. His e-mail address is johnson@isi.edu.