

STEVE: A Pedagogical Agent for Virtual Reality

Jeff Rickel and W. Lewis Johnson
Information Sciences Institute & Computer Science Department
University of Southern California
4676 Admiralty Way, Marina del Rey, CA 90292-6695
rickel@isi.edu, johnson@isi.edu
<http://www.isi.edu/isd/VET/vet.html>

Abstract

We are exploring the use of virtual reality for training people how to perform tasks, such as operating and maintaining complex equipment. This video describes Steve, an agent we are developing that assists in the training. Steve is an autonomous, animated agent that cohabits the virtual world with students. Steve continuously monitors the state of the virtual world, periodically manipulating it through virtual motor actions. His objective is to help students learn to perform physical, procedural tasks. He can demonstrate tasks, explaining his actions, as well as monitor students performing tasks, providing help when they need it. In addition to teaching students individual tasks, he can also help them learn to perform multi-person team tasks: he can serve as a tutor for a student learning a particular role in the team, and he can play the role of a teammate when a human teammate is unavailable. By integrating previous work in agent architectures, intelligent tutoring systems, and computer graphics, Steve illustrates a new breed of computer tutor: a human-like agent that can interact with students in a virtual world to help them learn.

1 Introduction

To master complex, real-world tasks, such as operating complicated machinery, people need hands-on experience facing a wide range of situations. They also need a mentor that can demonstrate procedures, answer questions, and monitor their performance, and they may need teammates if their task requires multiple people. Since it is often impractical to provide such training on real equipment, we are exploring the use of virtual reality instead; the training takes place in a three-dimensional, interactive, simulated mock-up of the student's work environment (as shown in the first video clip). Since mentors and teammates are often unavailable when the student needs them, we are developing an autonomous, animated agent that can play these roles. His name is Steve (Soar Training Expert for Virtual Environments).

Steve integrates methods from three research areas: intelligent tutoring systems, computer graphics, and agent architectures. This novel combination results in a unique set of capabilities. Steve has many pedagogical capabilities one

Appears in the Proceedings of the Second International Conference on Autonomous Agents, Minneapolis/St. Paul, ACM Press, May 1998.

©1998 ACM




Figure 1: Steve describing a power light

would expect of an intelligent tutoring system. For example, he can answer questions such as “What should I do next?” and “Why?”. However, because he has an animated body, and cohabits the virtual world with students, he can assist in ways that previous disembodied tutors cannot. For example, he can demonstrate actions, he can use gaze and gestures to direct the student's attention, and he can guide the student around the virtual world. Thus, virtual reality allows more human-like interactions among synthetic agents and humans than desktop interfaces can. Moreover, Steve's agent architecture allows him to robustly handle a dynamic virtual world, potentially populated with people and other agents; he continually monitors the state of the virtual world, always maintaining a plan for completing his current task, and revising the plan to handle unexpected events. All these capabilities are illustrated in the second video clip.

2 Steve's Architecture

Like many other autonomous agents that deal with a real or simulated world, Steve consists of two components: the first, implemented in Soar [5], handles high-level cognitive processing, and the second handles sensorimotor process-

ing. The cognitive component interprets the state of the virtual world, carries out plans to achieve goals, and makes decisions about what actions to take. The sensorimotor component serves as Steve's interface to the virtual world, allowing the cognitive component to perceive the state of the world and cause changes in it. It monitors messages from the simulator describing changes in the state of the world, from the virtual reality software describing actions taken by the student and the student's position and field of view, and from speech recognition software describing the student's requests and questions posed to Steve. The sensorimotor module sends messages to the simulator to take action in the world, to text-to-speech software to generate speech, and to the virtual reality software to control Steve's animated body. Because the sensorimotor component provides a mapping between high-level action commands sent from the cognitive module and their realization in the virtual environment, it controls Steve's appearance in the virtual world. We have experimented with several graphical representations for Steve; in this video, Steve appears as a head and torso, with a hand that can manipulate and point at objects, as shown in Figure 1.

3 Team Training

We have recently extended Steve to understand team tasks, which require the collaboration of multiple people. To do this, we extended Steve's representation of tasks to distinguish the roles of different team members and we generalized all his pedagogical capabilities to use this extended representation. The third video clip shows two Steve agents demonstrating their ability to work together as a team. Their understanding of the team task shows in their anticipation and awareness of each other's actions.

Steve's ability to understand team tasks is important because it allows him to help people learn team tasks. In such training, Steve agents can play two valuable roles: they can serve as a tutor for an individual human team member, and they can play the role of missing team members, allowing students to practice team tasks without requiring all their human team members. The fourth video segment demonstrates this sort of training. A student practices his role in a two-person task while one Steve agent serves as his teammate and another Steve agent serves as his tutor.

4 Status and Future Work

Steve is fully implemented and integrated with the other software components on which he relies (i.e., virtual reality software, a simulator, and commercial speech recognition and text-to-speech products). We have tested him on a variety of Naval operating procedures; he can operate several consoles that control the engines aboard Naval ships, and he can perform an inspection of the air compressors on these engines. (The video shows simplified versions of some of these tasks.) We are continuing to extend his knowledge of these and related tasks. However, he is not limited to this domain; he can provide instruction in a new domain given only the appropriate declarative domain knowledge.

We have a variety of plans for future work. As an alternative to manual entry of domain knowledge, we are extending Steve to learn from demonstrations [1, 3]. To increase Steve's ability to motivate students, we are extending him to include emotions [2]. To allow him to express emotions, and to extend his range of nonverbal communication, we are

giving him control over his facial expressions. We are adapting Steve's architecture for use in distance learning over the World Wide Web [4]. Finally, Steve is under ongoing informal evaluation, both within our group and by external colleagues (e.g., the Air Force Armstrong Laboratory), and we are planning formal evaluations later this year.

5 Conclusion

Steve illustrates the enormous potential in combining work in agent architectures, intelligent tutoring, and graphics. When combined, these technologies result in a new breed of computer tutor: a human-like agent that can interact with students in a virtual world to help them learn.

For more technical details on Steve, as well as a discussion on related work, see [3], [6], [7], and [8].

6 Acknowledgments

This work is a collaboration among many people. We especially thank Marcus Thiebaut and Erin Shaw for creating Steve's body, Ben Moore for his work on speech recognition, Randy Stiles and his team at Lockheed Martin for their virtual reality software and models, and Allen Munro and his team at the USC Behavioral Technology Laboratory for their simulation software and models. This work is funded by the Office of Naval Research, grant N00014-95-C-0179.

References

- [1] R. Angros, Jr., W.L. Johnson, and J. Rickel. Agents that learn to instruct. In *AAAI Fall Symposium on Intelligent Tutoring System Authoring Tools*, Menlo Park, CA, 1997. AAAI Press. AAAI Technical Report FS-97-01.
- [2] C. Elliott, J. Rickel, and J.C. Lester. Integrating affective computing into animated tutoring agents. In *Proceedings of the IJCAI Workshop on Animated Interface Agents: Making Them Intelligent*, pages 113-121, Nagoya, Japan, 1997.
- [3] W.L. Johnson, J. Rickel, R. Stiles, and A. Munro. Integrating pedagogical agents into virtual environments. *Presence*, 1998. Forthcoming.
- [4] W.L. Johnson and E. Shaw. Using agents to overcome deficiencies in web-based courseware. In J. Rickel and J. Lester, editors, *Proceedings of the AI-ED Workshop on Pedagogical Agents*, pages 48-55, Kobe, Japan, 1997.
- [5] J.E. Laird, A. Newell, and P.S. Rosenbloom. Soar: An architecture for general intelligence. *Artificial Intelligence*, 33(1):1-64, 1987.
- [6] J. Rickel and W.L. Johnson. Integrating pedagogical capabilities in a virtual environment agent. In *Proceedings of the First International Conference on Autonomous Agents*. ACM Press, February 1997.
- [7] J. Rickel and W.L. Johnson. Intelligent tutoring in virtual reality: A preliminary report. In B. du Boulay and R. Mizoguchi, editors, *Proceedings of the Eighth World Conference on Artificial Intelligence in Education*, pages 294-301. IOS Press, 1997.
- [8] J. Rickel and W.L. Johnson. Animated agents for procedural training in virtual reality: Perception, cognition, and motor control. 1998. Submitted for publication.