

Building a Case for Agent-Assisted Learning as a Catalyst for Curriculum Reform in Medical Education

Erin Shaw, Rajaram Ganeshan, and W. Lewis Johnson

Center for Advanced Research in Technology for Education

Information Sciences Institute, University of Southern California

4676 Admiralty Way, Marina del Rey, CA 90292-6695 USA

{shaw, rajaram, johnson}@isi.edu, <http://www.isi.edu/isd/carte/>

Douglas Millar

Graduate School of Education and Psychology, Pepperdine University

24255 Pacific Coast Highway, Malibu, CA 90263 USA

doughhelen@moonlink.net

Abstract

Animated pedagogical agents offer promise as a means of making computer-aided learning more engaging and effective. Realizing its full potential fully will involve integrating agents into curricula on a massive scale, so that they are as pervasive as textbooks currently are. This paper describes progress toward this end: the development of an agent-assisted learning environment designed for widespread use in health science curricula. The system features Adele, an animated pedagogical agent who guides and assesses students as they work through clinical cases. The system is designed for use across a range of health science disciplines, and by students at various levels, from entering medical student through practicing professional. We present the results of a formal evaluation of Adele by twenty-five medical students, and draw several conclusions about computer-based tutoring in a clinical domain. We also describe our experience building the clinical case for the evaluation, from local prototype to Web-based product, and propose a process for large-scale case authoring.

Keywords: Animated pedagogical agent, intelligent tutoring, real-world evaluation, simulation-based training

Introduction

Animated pedagogical agent technology has been proposed as a new approach for making computer-based learning more engaging and effective (Johnson 1998). It builds upon previous work on intelligent tutoring systems (Wenger 1987) and extends it in several important respects. As we view the concept, animated pedagogical agents are a type of autonomous agent (Johnson & Hayes-Roth 1998): they are capable of pursuing goals in complex environments, adapting their behavior as needed in response to unforeseen events. Their environment is typically an educational simulation, together with the learners and other agents that interact with the simulation. A pedagogical agent may seek to achieve pedagogical goals (e.g., to help a learner to learn about a topic), communicative goals (e.g., to acknowledge a learner's action), and task goals (e.g., to help the learner solve some particular problem). Animated pedagogical agents also have life-like, animated personas. They can respond to learners with a combination of verbal communication and non-verbal gestures such as gaze, pointing, body stance, and head nods. They can convey emotions such as surprise, approval, or disappointment. Taken together these capabilities allow animated pedagogical agents to interact with learners in a manner that is closer to face-to-face collaborative learning.

The technical sophistication of animated pedagogical agents has progressed rapidly. Steve, a 3D animated agent, can interact with learners in individual and team scenarios (Rickel and Johnson 1998). PPP Persona is able to generate tutorial presentations of Web-based learning materials (André et al 1998). Cosmo is able to generate critiques and explanations using a combination of speech and emotive gestures (Towns et al 1998). Early empirical results show that these agents can enhance the learning experience and improve its effectiveness (Lester et al 1997).

Adele, the Agent for Distance Education – Light Edition, is designed to help realize the full potential of animated pedagogical agent technology (Shaw et. al 1999). She is designed to be used in a wide range of health science courses, where she monitors students as they solve problems, gives them feedback, points them to relevant on-line reference materials, and evaluates their performance. Adele operates in a Web-based, distributed environment, in which students may interact on line with faculty, other students, and with courseware. This paper focuses on our work in supporting

clinical skills in medicine; however, course materials for dentistry are being developed concurrently, and other courses are planned.

This project, if successful, will result in the adoption of learning technology at a breadth and scale that few if any AI and Education projects have achieved before. It was not clear at the start of the project whether the technology was up to the challenge. Empirical studies of animated pedagogical agents have shown promising results with school children, but it was not known whether university students and practicing professionals would respond positively as well. This paper presents important progress toward widespread adoption of agent-assisted learning techniques. We present the results of a formal evaluation of Adele by twenty-five medical students, and draw several conclusions about computer-based tutoring in a clinical domain. We also describe our experience building the clinical case for the evaluation, from local prototype to Web-based product, and propose a process for large-scale case authoring.

Adele system overview

Adele's system, shown in Figure 1, consists of four main components: the pedagogical agent, the simulation, the client-server, and the server store. The pedagogical agent consists further of two sub-components, the animated persona and the reasoning engine which monitors the student's actions and provides appropriate feedback to the student. A fifth component, the session manager, is employed when the system is run in multiple-user mode. The central server maintains a database of student progress and when appropriate, provides synchronization for collaborative exercises carried out by multiple students on multiple computers.

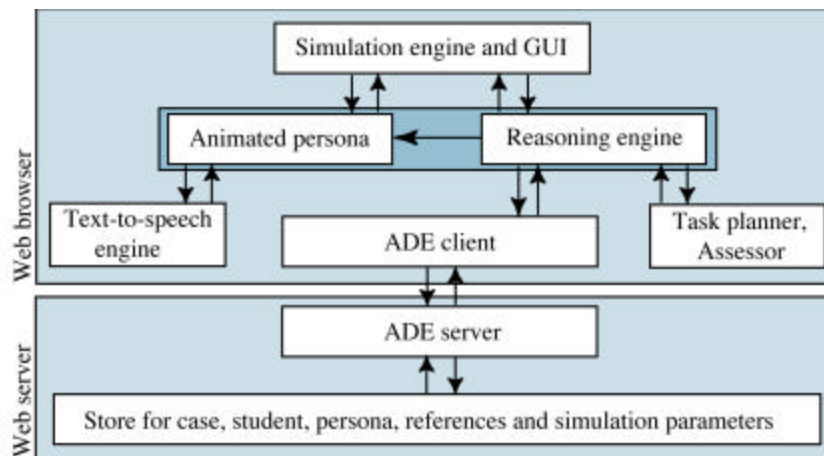


Figure 1. Architectural overview of Adele's system.

The reasoning engine performs all monitoring and decision making. Its decisions are based on a student model, a case task plan, and an initial state, which are downloaded from the server when a case is chosen, and on the agent's current mental state, which is updated as a student works through a case. Upon completion, a record of the student's actions is saved to the server where it can be used to assess the level of the student's expertise and determine how Adele will interact with the student in future cases.

The simulation can be authored using the language or authoring tool of one's choice. All simulations communicate with the agent via a common application programming interface (API) that supports event (e.g., the student orders a lab) and state change (e.g., the lab value is updated) notifications as defined by the simulation logic. The animated persona is simply a Java applet that can be used alone with a JavaScript interface or incorporated into larger applications, such as the simulation-based exercises we describe here. The system runs on the client's side for execution efficiency. This contrasts with the architecture of most other Web-based Intelligent Tutoring Systems where the intelligent tutor sits on the server side, resulting in increased latency in tutor response to student actions (e.g., Brusilovsky et. al 1997).

Task representation and reasoning

Previous efforts in medical intelligent tutoring (e.g., Clancey (1983) or Azevedo et al (1997)) have either been large expert systems incorporating sizeable medical knowledge bases, or have been prototypes focusing on narrow subfields of medicine. Neither approach is suitable for a system like Adele that is designed to apply to variety of health science courses, and yet can be downloaded and run on client computers. Adele employs a different approach: to encode for

each case just the knowledge needed to tutor that case. Decisions of what knowledge to include are made at the time the case is authored. Adele's formal knowledge representation is concerned mainly with the procedural knowledge necessary to work through the case and cope with the contingencies that might arise in the course of the case work-up. Other relevant medical knowledge, such as background information about disease etiologies, is incorporated into textual hints and Web-based reference materials, and thus can be presented to the learner as needed, but is not represented formally within Adele's knowledge base. The run-time knowledge base and reasoning engine is thereby simplified. Adele's knowledge representation focuses on the steps that the student should take to solve the case, the dependencies between them, and their rationales. The task steps and their dependencies are represented using a standard hierarchical plan (Russell & Norvig, 1995) which is similar to the representation used in Steve (Rickel & Johnson, Johnson 1998). A plan hierarchy is comprised of steps, each of which is either a primitive action (e.g. corresponds to a simulation event) or a complex action (e.g. is itself a plan). A step can have preconditions and endconditions, as well as hints, a rationale, a context and a role. Preconditions and end conditions are represented by Boolean expressions in conjunctive normal form. The plan hierarchy is evaluated at each step to account for the dynamic nature of a simulation and the unpredictability of a student's actions. Actions whose goals become "undone" are automatically re-executed while those whose goals are implicitly satisfied are skipped.

The reasoning engine can be run in three modes. In its most restrictive mode, it will simply block actions whose preconditions are unsatisfied. Adele uses this opportunity to provide unsolicited feedback about what *should* be done to satisfy the desired step's preconditions. The persona displays a *Hint* button so that a student may also ask for hints directly, before guessing or taking an incorrect action. In *practice* mode, the engine does not block – the student can make mistakes – and Adele does not provide unsolicited feedback, but still allows a student to ask for hints. In *exam* mode, Adele is not available. The modes are analogous to those of the SICULE tutor (Alexe & Gecsei, 1996).

Adele also uses the notion of a *situation space* (Marsella and Schmidt 1990) as a means of structuring the space of states associated with a domain so that it can be used to guide planning activity in dynamic situations. A situation is defined by a name, world state, goal expression, priority, and set of transitions which describe possible situations that can result from this situation whenever the associated conditions become true in the world state. Typically, when a situation is entered, a situation-appropriate sub-plan is instantiated to achieve the goal expression. Because the tutoring domain allows us to 'know' all possible situations *a priori*, all situational plans are pre-authored.

Opportunistic learning

Situation-based reasoning can also be applied to the problem of recognizing pedagogical opportunities as a student works through a given task. For example, student requests for diagnostic tests are opportunities for Adele to ask the student questions about the results of the tests. Adele supports two kinds of opportunistic learning: (1) quizzes to assess student's understanding of material; and, (2) introduction of additional reference material (e.g. videos or URL's) on the current subject. By maintaining an awareness of the situation the agent can undertake situationally appropriate interactions with the student, for example, giving quiz on a recently taken step.

Case-based medical education

The School of Medicine at the University of Southern California is beginning to implement an undergraduate (years I and II) curriculum that emphasizes the analysis of clinical cases. A clinical case typically includes the circumstances leading to the presentation of a patient in a doctor's office, the physical examination, the differential diagnoses of the disease, with an emphasis on determining a final diagnosis, and its treatment and management. "Clinical case"-based instruction is, more generally speaking, a case-based approach to instruction [Gragg, 1940], one in which patient cases create a familiar context for learning. The move to a case-based approach to medical instruction is in response to the problem of *knowledge application*: a student may recall learned knowledge but can not apply it in an operational setting [Whitehead 1929]. While not particular to the medical domain, the problem is pervasive within it.

Though the instruction is case-based, it is neither situated nor interactive. Because clinical instruction doesn't begin until year three, cases and their related learning materials, are presented in a classroom setting, and presentation to a class of over a hundred students necessarily places constraints on the amount of interaction that can occur between an instructor and her students. Cases are inherently situated and interactive, however, because they are based on a relationship between a patient and a physician, but have been condensed for presentation. With some effort they can be reconstructed as situated, interactive exercises. In this section we explain how Adele guides students through one of these interactive exercises and discuss the benefits and drawbacks computer-based tutoring in the medical domain.

Using Adele

Figure 2 shows a typical case-based diagnosis exercise in which students are presented with a simulated patient in a clinical setting. In the role of physicians, students are able to perform a variety of actions on the simulated patient; they may ask questions about medical history, perform a physical examination, order diagnostic tests, and make diagnoses. Adele monitors the student's actions and provides feedback accordingly. Depending upon the instructional goals, Adele may highlight aspects of the case, suggest correct actions, provide hints and rationales for particular actions, reference relevant background material, and provide contextual assessment, to test a student's understanding.



Figure 2. Adele explains the importance of palpating the patient's abdomen

Multiple-level learning design

A clinical case presents a rich context for learning at many levels. Adele can emphasize the procedure, namely, the *best practice* approach to a particular diagnosis, or she can concentrate on the related learning materials. To address multiple-level learning, we partitioned the instructional capability into three levels: pre-clinical, clinical, and resident, as is appropriate for the domain. The first two years of medical school are pre-clinical; they stress fundamental facts and are examination-oriented; the next two years are clinical and emphasize practice with patients and disease diagnosis; finally, as residents, students will learn to manage the treatment of a disease. One case can support all three levels, although the pedagogical emphasis of the learning task will change for each level.

Computer-based learning design

The pedagogical emphasis will also change with respect to the learning content in order to exploit the benefits, or mitigate the drawbacks, of computer-based learning. For example, for clinical practice, there is no substitute for working with live, or even standardized (i.e., actors playing patients) patients. Discovery by palpation and percussion, two forms of touch, cannot be taught on a computer. Teaching listening skills is problematic, too, because it is currently difficult to record high fidelity lung and heart sounds without special equipment. The best a program can do is to provide a textual description of what something feels like or sounds like.

On the other hand, the visual bias of the computer makes it especially good for teaching observation skills, something that is often under-emphasized in clinical settings, where students are apt to miss important visual cues in their rush to complete an examination. Instructors who evaluated Adele, however, commented that the system's response to *inspection*, was often insufficient: A typical response such as "The neck shows no venous jugular distention.", or "The throat is not inflamed.", leaves open the question of what venous jugular distention, or an inflamed throat, actually look like, so that simply describing what something looks like is often not enough, either, but is easily remedied with the addition of remedial material like images that show both normal and diseased areas.

Case Development

If Adele is to be adopted as part a system-wide curriculum for case-based medical education, the issue of authoring the task representation and simulation logic for each case becomes very important. This section discusses issues in large-scale case authoring and tools to assist in this process.




Figure 3. Adele's pedagogical task model is developed by layering domain, simulation, and pedagogical data.

An interactive case for Adele can be authored in steps as shown in Figure 3. Case data are typically data that have been summarized from a case for presentation in a classroom or seminar setting and typically includes the presenting complaint, the clinical finding, the differential diagnoses, and the treatment plan. The disease data includes information about the differential diagnoses for the case. The best practice model for the diagnosis is built from the case and disease data. This task model specifies an ordering of steps to be taken, places constraints on the steps, includes optional steps, and relates the steps to the diagnosis. Rationales and hints for the steps are added here. How much of the case is used to build the task model depends on the instructional level and objective. Once the best practice model is determined, the author supplies the responses and media assets needed for the simulation. These include image and audio files for the patient and clinical findings. The pedagogical model is then layered over the resulting task model. Contextual comments, reference materials, and quizzes for assessment are added to the appropriate places in the task step sequence.

Building an authoring tool

Currently, case authors enter data for new cases via a spreadsheet, as we have not finished the development of a dedicated authoring system). A translator takes information from the table-based spreadsheet in the form of comma-separated files and parses it into an object-oriented case representation scheme from which both the task representation as well as simulation logic for the case are generated automatically. This translation process performs a variety of syntactic consistency checks on the input provided by the author and also saves the author the burden of matching simulation events with the actions in the plan. The spreadsheet-translator system isn't graphical; authoring a non-sequential path through a case can be confusing; and it does not provide for a persistent repository so that data can be easily reused across cases. We are actively designing an authoring tool that will meet the needs of both instructor and developer that has all of these qualities.

In their paper on authoring via task-specific tools, Bell and Zirkel (1997) argue that many ITS authoring tools sacrifice pedagogical power for content flexibility and that developing an authoring tool for a specialized category of instructional applications has advantages even if it sacrifices flexibility for power. We argue that while targeted authoring tools do indeed sacrifice flexibility, it need not be for power. Power is afforded to authors who understand the types of tasks they are authoring, and the constraints that are inherent within those tasks. From our experiences with one-on-one authoring between physicians and programmers, we propose that a successful authoring system must be: (1) **Domain-intuitive**: Medical professionals understand how clinical cases are constructed and presented, and their knowledge of the domain can be accounted for to help tailor a system for them, rather than one for the programmer, or architecture; (2) **Reusable**: Use standard ontologies and provide a repository-based selection paradigm, where users select previously-authored steps, rationales, etc. from a repository, or create new objects, which are then added to the repository; the repository thus serves as a broader knowledge base from which knowledge for particular cases is drawn as needed; and, (3) **Testable**: Case authors need feedback on the cases they author and familiarizing them with the system they are authoring for will make them better authors. Adele's system may be a good candidate for an authoring tool that utilizes Programming By Demonstration (PBL). Diligent (Angros, et al., 1998), is a good example of a PBL system in which users author tasks within the tutoring environment; it has been shown that instructors more accurately do in context, than explain out of context.

Formative evaluation

The evaluation we describe here is the first formative evaluation of the Adele system by a medium-sized group of arbitrary and unmonitored college students. The evaluation was given to a class of second year medical students during the third week of November 1998. For this evaluation, physicians in the Department of Family Medicine authored a new case on Lung Cancer that the students had covered in class. The students were asked to work through the case just before their final examination because the case contained questions that would be on the exam. In retrospect, it was not a good idea to require a time-consuming evaluation right before an exam. Over one hundred students worked through the case, but only twenty-five percent of them completed the final evaluation questionnaire. Two face-to-face evaluations were also conducted.

The questions on the evaluation addressed both specific elements of the tutoring system, such as the interface and the rationales, and the general reaction of the students to Adele and the concept of the system. A short pre-evaluation was given in order to assess the students level of computer literacy. The final form contained thirty questions in six categories: system use, system components, rationales, Adele, and learning. Each answer was scored using a Likert scale of 0-4 (0=strong disagree; 1=disagree; 2=neutral; 3=agree; 4=strong agree). We discuss the analysis in the next section.

Analysis of the results

All the quantitative results reported here are based on a sample size of 24-25. The majority of students thought Adele would be a good distance education tool (0-4%; 1- 8%; 2-16%; 3-44%, 4-28%), and useful in class as a preparation tool (0-4%; 1- 17%; 2-8%; 3-38%, 4-33%), but would not suffice as a replacement for a class lecture (0-29%; 1- 33%; 2-21%; 3-8.5%; 4-8.5%). Students found Adele's hints helpful (0-0%; 1- 4%; 2-33%; 3-21%, 4-42%), and liked her rationales (0-0%; 1- 4%; 2-29%; 3-33.5%, 4-33.5%). There were mixed feelings, though, about how they wanted the information. Adele knows why each step should be taken and will explain the reason when the student clicks on her "Why?" button after she suggests taking a step. During a one-on-one evaluation, however, we noticed the student went halfway through the case without ever asking "Why?" A key part of the learning was being ignored. We also observed that when the instructor tested his case, he often asked "Why?" *after* he took the step for which Adele had given a hint. We had three choices: to have Adele give the rationale only when asked, to have her give it automatically after she gave a hint, and to have her give it automatically after the user took a step. We decided to give the students a taste of each and then put the question to them. We split up the especially verbose rationales and gave the both parts automatically: the first after the hint and the second after the action. The shorter rationales were left to discovery. While a small number of students preferred to hear the rationale after, rather than before, an action, even more preferred to hear only when they asked. It is not yet known though, how many of the students who preferred to have to ask for a rationale, actually asked. This is an area that deserves more attention. Analysis of interaction traces with Adele will help answer this question.

These students did not find Adele believable as an attending physician (0-8%; 1- 24%; 2-44%; 3-24%, 4-0%), but actually none of these students had ever worked with an attending physician. A fourth year student who evaluated the first case claimed Adele was just like an attending physician because she gave quizzes as one worked through the case. The quizzes for the Lung Cancer case came only at the end and weren't particularly contextual. If there was a consensus, it was that they wanted more cases to work on. In the words of one participant, "the more practice, the better."

Students certainly preferred the persona to a text-only tutor (0-12%; 1- 32%; 2-24%; 3-16%, 4-16%). Though these numbers are not conclusive, the students' provided very favorable impressions of Adele in their general textual comments. Students also preferred a real voice to a synthesized one (0-0%; 1- 12%; 2-32%; 3-32%, 4-24%). Surprisingly, however, they did not mind that Adele's lips and voice were not synchronized (0-4%; 1- 22%; 2-48%; 3-22%, 4-4%). Adele's lips and voice move at the same time but are not yet phonetically synchronized and some people find this phenomenon disturbing. It may not be as important as we once thought, although we are changing it anyway.

In general, they thought that the system was easy to access (0-0%; 1- 8%; 2-8%; 3-42%; 4-42%), the interface was easy to use (0-0%; 1- 4%; 2-24%; 3-52%; 4-20%), and that it was easy to navigate between different views - chart, patient, progress - (0-0%; 1- 12%; 2-8%; 3-60%; 4-20%), but they found it difficult to figure out what to do (0-12%; 1- 16%; 2-24%; 3-40%; 4-8%), and did not think two of the system components were terribly useful. The latter comments indicate a need for more guidance than the system currently provides. We are looking at Adele to fill this role; to suggest actions at a *system* level, as opposed to a *task* level, when a user becomes confused, and to point out the importance of interface elements if they are not utilized. Individual comments ran the gamut, however;

from frustrated students who were unable to use the system to enthusiastic users whose comments sound more like testimonials. As researchers and initial users of the systems we develop, we are inclined to bias the usability of the system in favor of users like ourselves, in other words, sophisticated users. Before any system can be successfully deployed on a large scale, this bias must be remedied.

Conclusion

Adele and other pedagogical agents like her have lessons to offer regarding autonomous agent design. As indicated in the previous section, the evaluation confirmed the usefulness of Adele as a distance education tool that supplements class lectures. The animated persona was clearly preferable to a text-only tutor even though the students did not find Adele as believable as an attending physician. Adele's proactive interventions with quizzes and references seem to make her more believable. We should also note that the Adele persona made a positive impression on the medical faculty, whose cooperation and support was crucial. We are looking into conversational interfaces to improve the interaction with the student. Many of the students also complained that the user interface was not very intuitive and hence found it difficult to use. We are modifying the system so that Adele can provide assistance at the system level.

Another area of concern is the use of rationales for actions, which are a key part of the learning process. The evaluation indicated that students would prefer to explicitly ask for rationale, but don't always do so. Adele's task representation needs to be enhanced to explicitly reason with hypotheses and their likelihoods as in GUIDON (Clancey, 1983). This will allow Adele to present rationales in a more flexible manner and also automatically generate quizzes to verify the student's knowledge of the hypotheses underlying their actions. From an authoring perspective, the explicit representation of hypotheses and their relationships to findings could allow for automated generation of rationale without requiring the author to provide the rationale for every diagnostic step in a case. These limitations notwithstanding, the approach of orienting Adele's run-time knowledge base to individual cases appears to be workable, and will continue to be followed as the library of cases is developed.

Finally, we believe that the issue of case authoring is critical for the widespread adoption of Adele as part of the medical curriculum. Any authoring environment should be intuitive to the medical domain, support re-use of case material via a standard ontology and be integrated with the Adele system to allow testing of cases.

Acknowledgments

CARTE staff members Kate Labore and Dr. Jeff Rickel, 'A' Team members Ami Adler, Andrew Marshal, Anna Romero, graduate student Chon Yi all contributed to the work presented here. Dr. Allan Abbott led the evaluation at the medical school and he and medical student Michael Hasler authored the case. Our collaborators, Drs. Demetrio Demetriatis, William La, Wesley Naritoku, Sidney Ontai, and Beverly Wood, and Angela Atencio and Leah Flodin at the USC School of Medicine provided indispensable assistance. This work was supported by an internal research and development grant from the USC Information Sciences Institute.

References

1. Alexe, C. and Gecsei, J. (1996). A learning environment for the surgical intensive care unit. Frasson, C., Gauthier, G. and Lesgold, A. (Eds.), *Proc. of the Third Int'l Conf. on Intelligent Tutoring Systems*, pp. 439-447. Springer Verlag, Berlin.
2. Angros, R., Johnson, W.L., & Rickel, J. (1997). Agents that Learn to Instruct,. AAAI 1997 Fall Symposium Series: Intelligent Tutoring Systems Authoring Tools, Technical Report FS-97-01, , AAAI Press.
3. André, E., Rist, T., and Müller, J. (1998). Integrating reactive and scripted behaviors in life-like presentation agents. In K.P. Sycara and M. Wooldridge (Eds.), *Proc. of the Second Int'l Conf. on Autonomous Agents*, pp. 261-268, ACM Press, New York.
4. Azevedo, R., Lajoie, S.P., Desaulniers, M., Fleiszer, D.M., & Bret, P.M. (1997). RadTutor: The theoretical and empirical basis for the design of a mammography interpretation tutor. Proceedings of the 8th World Conference on Artificial Intelligence and Education, IOS Press, Amsterdam..
5. Brusilovsky, P., Schwartz, E., and Weber, G. (1997). ELM-ART: An intelligent tutoring system on world wide web.
6. Clancey, W. (1987). "Knowledge-based Tutoring: The GUIDON Program," MIT Press.
7. Gragg, C.I. (1940). Because wisdom can't be told, *Harvard Alumni Bulletin*, October 19, pp.78-85.

8. Johnson, W.L. (1998). Pedagogical agents. Proceedings of ICCE'98, vol. 1, 13-22. China Higher Education Press and Springer-Verlag, Beijing.
9. Johnson, W.L. and Hayes-Roth, B. (1998). The First Autonomous Agents Conference. *The Knowledge Engineering Review* 13(2), pp. 1-6.
10. Lester, J.C., Converse, S., Stone, B., Kahler, S., and Barlow, T. (1997). Animated pedagogical agents and problem-solving effectiveness: A large-scale empirical evaluation. In *Proceedings of the Eighth World Conference on Artificial Intelligence in Education*, pp. 23-30. IOS Press, Amsterdam.
11. Marsella, S.C. and Schmidt, C.F. (1990). Reactive planning using a situation space, *Proc. of the 1990 AAAI Spring Symposium Workshop on Planning*.
12. Munro, A., Johnson, M.C., Pizzini, Q.A., Surmon, D.S., and Wogulis, J.L.(1996). A tool for building simulation-based learning environments. In *Simulation-Based Learning Technology Workshop Proceedings, ITS'96*.
13. Reilly, W.S.N. (1997). A methodology for building believable social agents. In W.L. Johnson and B. Hayes-Roth (Eds.), *Proc. of the First Int'l Conference on Autonomous Agents*, pp. 114-121. ACM Press.
14. Rickel, J. and Johnson, W.L. (1998). Animated agents for procedural training in virtual reality: perception, cognition, and motor control. Accepted for publication in *Applied Artificial Intelligence Journal*.
15. Rickel, J. and Johnson, W.L. (1997). Intelligent Tutoring in Virtual Reality: A Preliminary Report, *Proc. of the Int'l Conf. on Artificial Intelligence in Education*, pp. 294-301 IOS Press, Amsterdam.
16. Russell, S., and Norvig, P. (1995) *Artificial Intelligence: A Modern Approach*. Prentice Hall, Englewood Cliffs.
17. Sengers, P., (1998). Do the Thing Right: An architecture for action-expression. In K.P. Sycara and M. Wooldridge (Eds.), *Proc. of the Second Int'l Conf. on Autonomous Agents*, pp. 24-31. ACM Press, New York.
18. Shaw, E., Johnson, W.L., and Ganeshan, R. (1999). Pedagogical agents on the web. Submitted to *Third Int'l Conference on Autonomous Agents*.
19. Towns, S.G., Callaway, C.B., Voerman, J.L., and Lester, J.C. (1998) Coherent gestures, locomotion, and speech in life-like pedagogical agents. *IUI '98: 1998 International Conference on Intelligent User Interfaces*, pp. 13-20, ACM Press, New York.
20. Wenger, E. (1987). *Artificial intelligence and tutoring systems: Computational and cognitive approaches to the communication of knowledge*. Los Altos, CA: Morgan Kaufmann Publishers, Inc.
21. Whitehead, A.N. (1929) *The Aims of Education*. Cambridge University.

