Face-to-Face Interaction with Pedagogical Agents, Twenty Years Later

W. Lewis Johnson¹ · James C. Lester²

Published online: 24 September 2015

© International Artificial Intelligence in Education Society 2015

Abstract Johnson et al. (*International Journal of Artificial Intelligence in Education*, 11, 47–78, 2000) introduced and surveyed a new paradigm for interactive learning environments: animated pedagogical agents. The article argued for combining animated interface agent technologies with intelligent learning environments, yielding intelligent systems that can interact with learners in natural, human-like ways to achieve better learning outcomes. We outlined a variety of possible uses for pedagogical agents. But we offered only preliminary evidence that they improve learning, leaving that to future research and development. Twenty years have elapsed since work began on animated pedagogical agents. This article re-examines the concepts and predictions in the 2000 article in the context of the current state of the field. Some of the ideas in the paper have become well established and widely adopted, especially in game-based learning environments. Others are only now being realized, thanks to advances in immersive interfaces and robotics that enable rich face-to-face interaction between learners and agents. Research has confirmed that pedagogical agents can be beneficial, but not equally for all learning problems, applications, and learner populations. Although there is a growing body of research findings about pedagogical agents, many questions remain and much work remains to be done.

Keywords Pedagogical agents · Game-based learning · Virtual tutors · Virtual coaches · Virtual environments · Robotics · Teachable agents

✉ W. Lewis Johnson
ljohnson@alelo.com

¹ Alelo Inc., Los Angeles, CA, USA

² North Carolina State University, Raleigh, NC, USA

Introduction and Retrospective on the Article

In the 1990s our colleagues and we started developing a new approach to the design of interactive learning environments: *animated pedagogical agents*. We were inspired by the early work on interface agents (Laurel 1990; Nagao and Takeuchi 1994; Hayes-Roth and Doyle 1998), and saw potential for applying this approach to the design of intelligent learning environments. The resulting learning environments might be able to interact more naturally with learners, and this might make them more effective as learning tools.

We conducted some preliminary projects with animated pedagogical agents, including Herman the Bug (Lester et al. 1999), STEVE (Johnson et al. 1998), and Adele (Shaw et al. 1999). These efforts gave us further insight into possible uses for animated pedagogical agents in learning environments. Seeing commonalities across different animated pedagogical agent applications, we decided to collaborate and write an article summarizing the range of possible capabilities and benefits of pedagogical agents as we had come to understand them. This article, “Animated pedagogical agents: Face-to-face interaction in interactive learning environments” (Johnson et al. 2000), proved to be highly influential and is commonly cited by researchers interested in pedagogical agents.¹

At the time we wrote the article we were excited by the potential of the technology, but did not yet have a clear understanding of what learning problems it could help solve. Looking at the article again after 15 years we see various claimed benefits such as broadening the bandwidth of tutorial communication or conveying emotional responses to the tutorial situation. But was there evidence that existing intelligent tutors lacked sufficient bandwidth of communication with the learner, or that they suffered from being insufficiently emotional? Early studies provided evidence that they can improve learning (Lester et al. 1997a, b), but we only had a preliminary notion of what features of agents contributed to learning and why. The truth is, we thought pedagogical agents were a clever idea and so we ran with it, with only a general idea of where it might lead us.

Twenty years since the initial work on animated pedagogical agents, and 15 years since the publication of Johnson et al. (2000), we can look back and assess whether it led us in a good direction. Did pedagogical agents live up to their expectations? Of the various capabilities described for agents in the article, which ones proved useful in the long run? What have we learned since then? And are there ideas and concepts in the 2000 article that are worth revisiting, as we look toward the future? This article will explore these questions.

Key Capabilities and Benefits Cited in the Paper

In our original article, we conceptualized pedagogical agents as “autonomous characters” that “cohabit learning environments with students to create rich, face-to-face learning interactions.” We emphasized the multimodal nature of pedagogical agents

¹ According to Google Scholar, the article has been cited more than 1,000 times.

and used the phrasing of “animated pedagogical agents” to highlight their visual embodiment.

The following is the set of potential agent capabilities presented in the original article, illustrated by agents described in the literature up to that point. These capabilities become possible when an agent is provided with an animated persona and is able to interact with learners.

- **Interactive demonstrations.** Agents in virtual environments can demonstrate how to perform tasks such as operating or repairing equipment, as illustrated by STEVE (Johnson et al. 1998). In the process the agent can explain what it is doing and why, and also direct the learner’s attention to important features of the environment, such as by pointing. The learner can also collaborate with and interact with the agent during the course of the demonstration.
- **Navigational guidance.** An agent can lead learners around a complex virtual environment and prevent them from getting lost. Again, STEVE provided a good demonstration of this capability.
- **Gaze and gesture as attentional guides.** Agents can point to and look at objects in the environment to draw the learner’s attention to them. Many of the agents described in the Johnson et al. (2000) article have this capability.
- **Nonverbal feedback.** Agents with animated personas can give feedback nonverbally as well as verbally. These nonverbal cues can take various forms, such as nodding or shaking the head, facial expressions such as smiling or surprise, or even cartwheels across the screen (in the case of Herman the Bug).
- **Conversational signals.** Similarly, agents can use nonverbal signals to regulate conversation with learners. For example, STEVE used head nods as back-channel feedback to indicate that he understands the learner’s spoken utterances.
- **Conveying and eliciting emotion.** Animated agents can express emotion, and can also elicit emotions in learners. We hypothesized that this might influence learner motivation, e.g., by expressing empathy toward the learner.
- **Virtual teammates.** Animated agents can play roles as team members, as part of team training scenarios. STEVE demonstrated this capability (Rickel and Johnson 1999). In a similar vein we hypothesized that animated agents could collaborate with learners as **learning companions**.
- **Adaptive pedagogical interaction.** We argued that the dynamic nature of face-to-face interaction between an agent and a learner made it necessary to make pedagogical interaction highly adaptive, so that the agent can respond to interruptions, turn-taking, and miscellaneous actions that the learner might take during instruction. STEVE is a good demonstration of such adaptive instructional capabilities.

Evolution of the Pedagogical Agent Concept

The field’s conceptualization of pedagogical agents has remained relatively consistent since our original formulation. For example, a recent meta-analytic review of pedagogical agents by Schroeder et al. (2013) defines agents as “on-screen characters that facilitate instruction.” However, as research on pedagogical agents has evolved, it became common to drop “animated” and simply refer to them as “pedagogical agents.”

Departing from the vision presented in the original article, researchers and practitioners have considered a wider space of alternative agent designs. Some researchers such as Veletsianos (2010) have explicitly extended the pedagogical agent concept to include agents portrayed using static images. Some researchers have conducted studies to evaluate the relative benefits of static vs. animated agents. For example, Baylor et al. (2003) and Mayer and DaPra (2012) found that subjects learned more when the persona was animated than when it was static. Meanwhile Moreno et al. (2001) and Baylor et al. (2003) found that the agent's voice contributes significantly to learning, sometimes more than the animated persona. In practice the choice of agent realization depends upon the educational context and the role of the agent in that context.

In the 2000 article we suggested that pedagogical agents could act in a variety of roles. Subsequent research has sought to clarify and investigate these roles (e.g., see Kim and Baylor 2015). In addition to the classic pedagogical agents discussed in the 2000 article, which played a somewhat didactic role, the intervening years have seen the emergence of pedagogical agents that are designed to be taught, *teachable agents* (Biswas et al. 2005), and pedagogical agents that are designed to serve as peers, *learning companions* (Woolf et al. 2010). *Virtual role-players*, an extension of the virtual teammate concept in the 2000 paper, have emerged as an important application area for pedagogical agent technology (Johnson 2015a, b). Each type offers its own potential benefits to learners, and calls for a different set of agent capabilities. Teachable agents often require learners to explain their reasoning to “teach” their students (the teachable agents), thereby triggering the self-explanation effect (Chi et al. 1989) in learners and perhaps enabling deeper learning. Learning companions, in contrast, promote social interactions between students and “near-peers” that may stimulate engagement, which may increase learner motivation. Virtual role-players perform instructional functions through their reactions and responses to learners in educational simulations.

Alelo's VCATs (Virtual Cultural Awareness Trainers) (Johnson et al. 2011) provide a good example of how pedagogical agents can perform multiple roles and employ different capabilities in each role. In VCATs learners acquire knowledge about other cultures and apply their knowledge in simulated encounters with people from that culture. A Virtual Coach (Fig. 1, left) provides guidance and feedback throughout the course, and narrates the instructional material. During the role-play simulations (e.g.,

Fig. 1 VCAT Virtual Coach (left) and virtual role-play (right)

Fig. 1, right) the learner's avatar carries out cross-cultural exchanges with virtual role-players, with advice and feedback from the Virtual Coach.

VCATs have been developed for over 80 countries to date, and over 50,000 trainees have taken VCAT courses. They thus provide useful information regarding what works in agent design. Regarding the question of whether agents should be animated, static images, or disembodied voices, it clearly depends upon the intended role of the agent, and there is no single answer. There are very few instances in which it makes sense for virtual role-players to be disembodied voices. For virtual tutors and coaches on the other hand, a variety of realizations are possible. In the case of the VCAT Virtual Coach there is a combination of realizations. At the beginning of the course the Virtual Coach appears as an animated character. Later on, after the learner is accustomed to working with the Virtual Coach, she fades away and becomes less obtrusive. When she acts as narrator she is a disembodied voice. When she is coaching virtual role-plays she mainly is an off-screen character, and pops up on the screen only when the learner has a question. When she does pop up she appears as a static image. At the end of the training she fades away entirely, and learners are required to demonstrate that they can perform the task unassisted.

VCATs also illustrate how video stories of real people can complement interaction with pedagogical agents. Research in the use of video stories has developed in parallel with work with pedagogical agents (e.g., Schank 2010). Each serves distinct pedagogical functions, at different phases at the learning process (Johnson 2015a, b). Videos of real people, whom learners admire and respect, can motivate a learner to want to learn a new skill. Animated pedagogical agents are not real and so do not engender respect, but they respond to learners in ways that prerecorded videos do not. They are more useful when the learner is motivated and actively engaged in learning a new skill.

Since 2000 some roles for pedagogical agents have become much more common than others, and this has affected their capabilities as well as their realizations. For example there has been relatively little work since 2000 on agents that can engage in interactive demonstrations or navigational guidance. Most agents do not cohabit virtual worlds with learners, and so do not have opportunities to interact with the learner and the virtual world at the same time, as in an interactive demonstration. However this may now be changing as agents are taking on robotic form, as we will discuss below.

We have found that when the role of the agent is unclear it can have a negative effect on learning effectiveness. For example in the Tactical Iraqi game (Johnson 2010a) a single agent, Samia Faris, played the role teammate as well as coach. Because Samia was a teammate she was present in the simulation at all times. As a result learners tended to rely heavily on her to tell them what to do in the simulation. When we separated the functions of role-player and coach, as in VCATs, we could make the coach fade away while continuing to provide the role-playing functions.

New Technical Developments

Since the 2000 article there have been significant advances in three communities that contribute to agent functionalities: virtual humans, affective computing, and natural language processing. This has made it possible to create pedagogical agents that are, by the standards of the year 2000, phenomenally responsive. Perhaps we did not directly

anticipate that these technologies would become as sophisticated as they have as quickly as they did, but in the intervening 15 years it is clear that the kinds of pedagogical agents we envisioned are quickly becoming a reality. Together, the virtual human, affect, and natural language capabilities support precisely the kinds of multi-modal communication capabilities noted in the 2000 article.

The International Conference on Intelligent Virtual Agents has emerged as a showcase for virtual human technologies. Each year brings new advances in agents' ability to navigate through complex virtual environments, generate expressive gestures, direct gaze, communicate with speech, and synthesize full-body movements that integrate each of these functionalities to create seamless behaviors in real time. Current virtual human technologies, while still requiring much work to be done in both the supporting graphics and the AI behind the scenes, have arrived at a state that makes the interactive embodiment of agents possible.

The affective computing community has developed sophisticated computational models of affect recognition, affect understanding, and affect synthesis (D'Mello and Graesser 2010; McQuiggan, et al. 2007; Sabourin et al. 2011). Affect-aware computing systems are growing increasingly aware of their users through multiple modalities that draw on an ever widening array of sensors, which are themselves becoming less cumbersome to use and more cost effective (Arroyo et al. 2009; Grafsgaard et al. 2014; Kapoor and Picard 2005). These in turn are used in conjunction with environments populated with virtual humans who can respond, again in real time, to users' unfolding affective states. Much of this work has been conducted with an eye toward education and training applications, and pedagogical agents have thereby acquired increasingly greater abilities to recognize when users (learners) exhibit *frustration*, *boredom*, *confusion*, and states of *flow*, among others. Thus pedagogical agents have developed to the point where they do more than convey and elicit emotion, as suggested in the 2000 paper, but can also reason about and respond to emotion. They thus have become more emotionally intelligent. Moreover agents utilize emotion as a means to an end—to comfort, support, and motivate learners (Swartout et al. 2013; Johnson et al. 2011). This, together with social interaction tactics such as politeness tactics (Wang et al. 2008) is making agents more socially intelligent and thereby more instructionally effective.

The natural language community has made advances in human language technology that have significantly improved pedagogical agents' communicative capabilities. In particular, advances in natural language understanding and dialogue management have created the possibility of highly interactive natural language dialogue. Coupled with advances in automated speech recognition and text-to-speech technologies, state-of-the-art pedagogical agents can now engage in spoken language dialogue. There are limitations, of course, but the possibility of having tutorial conversations with pedagogical agents is now a reality.

These advances in capability enrich the interactive user experience, but they also raise new questions about how users will respond to and use the technology. For example, it is now possible to record video clips of real people and combine them with natural language technology to create interactive video-based agents that tell stories and respond to learner questions (Artstein et al. 2014). Will learners regard these agents as “real”, and therefore worthy of trust and respect, or as synthetic and fictional, and therefore less trustworthy than a real person? Also sophisticated and complex

technologies can sometimes be difficult to users to understand, control, and author. This could stand in the way of acceptance and adoption of these technologies by teachers and other stakeholders. These considerations are likely to become more significant as pedagogical agents transition into the educational mainstream.

What have we Learned Since then?

The 2000 article cited some preliminary studies. Since then, there have been many research studies, and we have a much better understanding of what features contribute to the effectiveness of pedagogical agents. Recently, a meta-analysis was conducted to determine what effect, if any, pedagogical agents have on learning (Schroeder et al. 2013). The meta-analysis considered 43 studies, which collectively involved more than three thousand subjects. The authors of the meta-analysis observe that analyzing the impact of pedagogical agents is challenging because of the myriad factors that bear on learning with agents. Agents can assume many forms. Their presentations can range from simple stick figures and talking heads to off-screen persona and full-bodied virtual humans. They can also be humanoid or non-humanoid, and they can be interactive videos of actual humans. In addition, agents have been designed for a multitude of learner populations and contexts, and they have played roles in both education and training. Learner populations include K-12 students, university students, and an exceptionally wide range of users for a multitude of training systems spanning business, defense, intelligence, and healthcare applications.

Despite the diversity exhibited by the studies along each of the dimensions noted above, the meta-analysis discovered statistically significant results among the studies. In particular, *the meta-analysis found that agents do enhance learning in comparison with learning environments that do not feature agents*. It also found that agents seemed to be more effective for science and math and less effective for the humanities.² Perhaps most interesting was the finding that, in formal education, *pedagogical agents seem to be more effective for younger learners than for older learners*. Specifically, it appears that agents seem to promote learning better for K-12 students than for post-secondary students.

Comparison and replication studies are helping to clarify the effect of various agent features, capabilities, and usage contexts. Mayer and his colleagues (Mayer and DePra 2012; Moreno et al. 2001) have conducted a series of careful studies comparing learning environments with different versions of agents. Graesser and his colleagues (Nye et al. 2014) have conducted a number of studies with different versions of AutoTutor, and have compared their effectiveness. Johnson and Wang (2010b) and McLaren et al. (2014) have tested the Politeness Effect with different learner populations and learning domains. Such studies have found, for example, that students interacting with pedagogical agents exhibit stronger learning outcomes when 1) pedagogical agents speak rather than communicate with text, 2) pedagogical agents use

² One of us (Johnson) takes some exception with this finding. Agents have proven to be highly effective in teaching foreign language and intercultural skills (Johnson 2010a; Johnson et al. 2011), which fall under the category of the humanities. So even though there is a growing body of research results to rely on, one must still be careful about overgeneralizing claims about pedagogical agents. Findings that hold in general may fail to hold in specific cases.

human-like gestures, 3) pedagogical agents communicate conversationally rather than formally, and 4) pedagogical agents use polite rather than direct phrasing. These studies are important because they help us to assess whether agent effects are robust or are somehow dependent on the experimental conditions of individual studies.

Since 2000 agents are increasingly being incorporated into game-based learning environments (Kim et al. 2009; Johnson 2010a; Rowe et al. 2011), where the agents play specific roles within the game. This is a natural use for animated agents, and in games there is usually a strong expectation on the learner's part that the characters will be animated.

Finally, experience has shown that pedagogical agents as coaches and tutors are most useful and effective with novice learners (e.g., Wang et al. 2008). This should come as no surprise since advanced learners have on average a higher level of motivation, self-efficacy, and self-regulation skills, and so have less need for the motivational support that virtual coaches and tutors offer. Virtual role-players in contrast provide useful practice for learners with a wide range of expertise.

What's Next for Pedagogical Agents?

In the time since the publication of 2000 paper pedagogical agents have proven themselves for a range of learner needs and learning domains. The technology is here to stay, and is sure to have a future. What will that future look like? Here are some speculations, based upon our experience to date.

In order for pedagogical agents to be adopted widely there need to be authoring tools that make it easy to create them. Earlier pedagogical agent work benefited from the availability of off-the-shelf agent authoring tools such as Microsoft Agent. Microsoft Agent has been discontinued, and no comparable tool has taken its place. But meanwhile human figures are becoming standard features of e-learning authoring tools such as Adobe Captivate and Articulate Studio. Alelo's VRP[®] MIL product makes it easy to populate training simulations with virtual role-players, drawn from reusable libraries. When these technologies are integrated into an easy-to-use package, pedagogical agents will become part of the standard repertoire that instructional designers use to create e-learning applications.

Although much progress has been made in our understanding of how pedagogical agents support learning, much remains unknown. The field needs to develop a broader empirically grounded research base on what types of pedagogical agents are most appropriate for what learner populations and subject matters and in what contexts. For example, are there learner populations or contexts for which pedagogical agents are not only unsupportive but in fact harmful? Are there learner populations or contexts for which pedagogical agents are very likely to be especially beneficial? Our own experience over the past decade over the course of many projects is that pedagogical agents seem to have a strong motivating effect for students ages 10–14 years, and they are especially helpful as non-player characters in cultural training applications. However, much remains to be learned about when pedagogical agents should be used and how their interactions with learners should be orchestrated. For every rule about pedagogical agents there are exceptions—perhaps even for this very rule!

As we discussed above, computer animation is not an essential feature of pedagogical agents, and in the future we see the technology migrating to other types of mixed-reality interfaces. One example is the RALL-E (Robot-Assisted Language Learning in Education) project at Alelo. Alelo has migrated its agent technology to the Robokind R25 robot, and the resulting prototype robot can engage in conversation in Chinese (Fig. 2). The concept is being tested in collaboration with the Virginia Department of Education and the Thomas Jefferson High School for Science and Technology in Alexandria, Virginia. Embedding agent technology into the robot makes it possible for learners to engage naturally in mixed-initiative conversation with the robot, much more so than is the case when the agent is displayed on a computer screen. Technically the approach is a direct descendent of the Steve pedagogical agent work described in Johnson et al. (2000), in that the agent and the learners cohabit a shared world. But while learners had to enter STEVE's virtual world to interact with him, the robot hardware makes it possible for the pedagogical agent to enter the real world of the learners and interact with them (Fig. 2).

As pedagogical agents become more widespread we will develop a better understanding of how best to combine interactive agents with human teachers and peers, to make learning most effective. Pedagogical agents complement the roles of humans in the learning process, and should not be viewed as taking the place of them. An important research challenge is to understand how best to combine the strengths of real people and artificial agents in blended learning environments. As tools for creating pedagogical agents become more prevalent we anticipate that teachers and students will play a significant role in designing, creating, and participating in those blended learning solutions. This will place a premium on agents that are not just technically sophisticated but also easy to author and use.

Pedagogical agents for personalized education are particularly intriguing. One can imagine a future in which every learner has her own pedagogical agent—or perhaps a cast of pedagogical agents—that accompanies her from the time she is young through

Fig. 2 The RALL-E Robot

adulthood and on into senescence. Her agents could provide highly customized support, delivered ubiquitously in all of her activities. And rather than being limited to a particular subject matter, they could support all subject matters, and expand to metacognitive and self-regulatory skills and beyond. They could also support complex collaborative problem solving in which teams of learners and their agents coordinate their learning activities. Similarly sophisticated capabilities can be envisioned for training, with the lines between education and training blurring and eventually disappearing altogether. These developments will require significant advances in agents' communicative abilities, as well as substantial increases in computing power than is now available, but both of these types of developments are clearly already underway.

As agents interact with learners over longer periods of time, they will not simply be interacting with learners but will be forging on-going relationships with them, just as people do. This will raise new questions about how these agents will establish rapport with learners and earn their trust and respect. Will learners come to trust and find common ground with agents as they do with their real teachers and mentors? As the boundary between interactive agents and portrayals of real people dissolves, will learners treat agents as fictional or real? Will this still be a meaningful distinction? These are questions that the next generation of pedagogical agents may have to address.

Conclusions

In the 2000 article we claimed that animated pedagogical agents offer enormous promise for interactive learning environments. Fifteen years later, our predictions have been largely borne out. Pedagogical agents have proven to be useful and sometimes highly effective in promoting learning, in a broad range of applications. And we continue to be optimistic for the future of pedagogical agents, and remain convinced that their full potential has yet to be realized, as learning technologies and as components of blended learning solutions.

Yet although there is now a significant body of experience and research findings relating to pedagogical agents, there is much that we still do not know. Many questions remain about when pedagogical agents are most effective, and how to design them and use them to maximize effectiveness. As agents become capable and more widespread, new questions will arise as to how to engender trust in this technology and promote its broad adoption. The findings of comparative studies such as Mayer and DePra (2012) and Nye et al. (2014) remind us that agent technology is only part of what makes a learning environment effective, and often not the most important part. Agent technology will not turn a badly designed learning environment into a good one. But it can help make a good learning environment better.

Acknowledgments The authors wish to acknowledge the contributions of our third co-author, Jeff Rickel, who passed away a few years after the publication of the 2000 article. Jeff's ideas and research contributions laid much of the groundwork for subsequent work on pedagogical agents. This research was supported in part by the National Science Foundation under Grants DRL-0822200, DRL-1020229, IIS-1138497, IIS-1321056, IIS-1344803, and IIS-1409639. Any opinions, findings, and conclusions expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

References

- Arroyo, I., Cooper, D.G., Burleson, W., Woolf, B.P., Muldner, K. & Christopherson, R.M. (2009). Emotion Sensors Go To School. In *14th International Conference on Artificial Intelligence in Education*, 17–24.
- Artstein, R., Traum, D., Alexander, O., Leuski, A., Jones, A., Georgila, K., Debevec, P., Swartout, W., Maio, H., & Smith, S. (2014). Time-offset interaction with a Holocaust survivor. In *IUI '14: Proceedings of the 19th International Conference on Intelligent User Interfaces* (163–168). New York: ACM Press. doi:10.1145/2557500.2557540#_blank.
- Baylor, A. L., Ryu, J., & Shen, E. (2003). The effects of pedagogical agent voice and animation on learning, motivation, and perceived persona. Paper presented at the Annual World Conference of Educational Multimedia, Hypermedia, & Telecommunication, Honolulu, Hawaii, 2003.
- Biswas, G., Leelawong, K., Schwartz, D., & Vye, N. (2005). Learning by teaching: a new agent paradigm for educational software. *Applied Artificial Intelligence*, 19(3–4), 363–392. doi:10.1080/08839510590910200.
- Chi, M., Bassok, M., Lewis, M., Reimann, P., & Glaser, R. (1989). Self-explanations: how students study and use examples in learning to solve problems. *Cognitive Science*, 13(2), 145–182. doi:10.1207/s15516709cog1302_1.
- D'Mello, S. K., & Graesser, A. C. (2010). Multimodal semi-automated affect detection from conversational cues, gross body language, and facial features. *User Modeling and User-Adapted Interaction*, 20(2), 147–187. doi:10.1007/s11257-010-9074-4.
- Grafsgaard, J., Wiggins, J., Boyer, K., Wiebe, E., & James, L. (2014). Predicting Learning and Affect from Multimodal Data Streams in Task-Oriented Tutorial Dialogue. *Proceedings of the Seventh International Conference on Educational Data Mining*, London, England, 122–129.
- Hayes-Roth, B., & Doyle, P. (1998). Animate characters. *Autonomous Agents and Multi-Agent Systems*, 1(2), 195–230.
- Johnson, W. L. (2010a). Serious use of a serious game for language learning. *International Journal of Artificial Intelligence in Education*, 20(2), 175–195.
- Johnson, W. L., & Wang, N. (2010b). The Role of Politeness in Interactive Educational Software. In C.C. Hayes & C.A. Miller (Eds.), *Human-Computer Etiquette*, 91–113. New York: Taylor and Francis. doi:10.1201/b10420-8.
- Johnson, W.L. (2015a). Cultural training as behavior change. *Proceedings of the 4th International Conference on Cross-Cultural Decision Making*. London: CRC Press.
- Johnson, W.L. (2015b). Constructing Virtual Role-Play Simulations. in R Sottolare, A. Graesser, X. Hu, K. Brawner (Eds.), *Design Recommendations for Adaptive Intelligent Tutoring Systems: Authoring Tools*, 3, 211–226. Orlando: US Army Research Laboratory.
- Johnson, W. L., Rickel, J. W., Stiles, R., & Munro, A. (1998). Integrating pedagogical agents into virtual environments. *Presence: Teleoperators and Virtual Environments*, 7(6), 523–546. doi:10.1162/105474698565929.
- Johnson, W. L., Rickel, J. W., & Lester, J. C. (2000). Animated pedagogical agents: face-to-face interaction in interactive learning environments. *IJAIED*, 11, 47–78.
- Johnson, W.L., Friedland, L., Schridder, P., Valente, A., & Sheridan, S. (2011). The Virtual Cultural Awareness Trainer (VCAT): Joint Knowledge Online's (JKO's) Solution to the Individual Operational Culture and Language Training Gap. In *Proceedings of ITEC 2011*. London: Clarion Events.
- Kapoor, A., & Picard, R.W. (2005). Multimodal Affect Recognition in Learning Environments. In *Proceedings of the 13th Annual ACM International Conference on Multimedia*, 677–682. doi:10.1145/1101149.1101300.
- Kim, Y., & Baylor, A. (2015). Research-based design of pedagogical agent roles: a review, progress, and recommendations. *International Journal of Artificial Intelligence in Education*, 25.
- Kim, J., Hill, R. W., Durlach, P. J., Lane, H. C., Forbell, E., Core, M., Marsella, S. C., et al. (2009). Bilat: a game-based environment for practicing negotiation in a cultural context. *International Journal of Artificial Intelligence in Education*, 19(3), 289–308.
- Laurel, B. (1990). Interface agents: Metaphors with character. In B. Laurel (Ed.), *The art of human-computer interface design*, 355–365. New York: Addison-Wesley.
- Lester, J.C., Converse, S.A., Kahler, S.E., Barlow, S.T. Stone, B.A., & Bhogal, R.S. (1997a). The persona effect: Affective impact of animated pedagogical agents. In *Proceedings of CHI'97*, 359–366. doi:10.1145/258549.258797.
- Lester, J.C., Converse, S.A., Stone, B.A., Kahler, S.E., & Barlow, S.T. (1997b). 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). Amsterdam: IOS Press.
- Lester, J. C., Stone, B. A., & Stelling, G. D. (1999). Lifelike pedagogical agents for mixed-initiative problem solving in constructivist learning environments. *User Modeling and User-Adapted Interaction*, 9, 1–44. doi:10.1007/978-94-017-1118-0_5.
- Mayer, R. E., & DaPra, C. S. (2012). An embodiment effect in computer-based learning with an animated pedagogical agent. *Journal of Experimental Psychology: Applied*, 18, 239–252. doi:10.1037/a0028616.

- McLaren, B. M., Sosnovsky, S., & Aleven, V. (2014). Preface – emerging technologies and landmark systems for learning mathematics and science: dedicated to the memory of Erica Melis – Part 1. *International Journal of Artificial Intelligence in Education*, 24(3), 211–215. doi:10.1007/s40593-014-0021-0.
- McQuiggan, S., Lee, S., & Lester, J. (2007). Early prediction of student frustration. In *Proc. of the 2nd Intl. Conf. on Affective Computing and Intelligent Interaction*, 698–709. doi:10.1007/978-3-540-74889-2_61.
- Moreno, R., Mayer, R. E., Spires, H., & Lester, J. (2001). The case for social agency in computer-based teaching: do students learn more deeply when they interact with animated pedagogical agents? *Cognition and Instruction*, 19, 177–214. doi:10.1207/s1532690xci1902_02.
- Nagao, K., & Takeuchi, A. (1994). Social interaction: Multimodal conversation with social agents. In *Proceedings of the Twelfth National Conference on Artificial Intelligence (AAAI-94)* (pp. 22–28). Menlo Park, CA: AAAI Press.
- Nye, B. D., Graesser, A. C., & Hu, X. (2014). AutoTutor and family: a review of 17 years of natural language tutoring. *International Journal of Artificial Intelligence in Education*, 24(4), 427–469. doi:10.1007/s40593-014-0029-5.
- Rickel, J., & Johnson, W.L. (1999). Virtual humans for team training in virtual reality. In *Proceedings of the Ninth International Conference on Artificial Intelligence in Education*. Amsterdam: IOS Press.
- Rowe, J., Shores, L., Mott, B., & Lester, J. (2011). Integrating learning, problem solving, and engagement in narrative-centered learning environments. *International Journal of Artificial Intelligence in Education*, 21(1), 115–133.
- Sabourin, J., Mott, B., and Lester, J. (2011). Modeling learner affect with theoretically grounded dynamic bayesian networks. *Proceedings of the Fourth International Conference on Affective Computing and Intelligent Interaction*, Memphis, Tennessee, 286–295. doi:10.1007/978-3-642-24600-5_32.
- Schank, R. (2010). Learning through storytelling, not documents: knowledge management meets AI. *eLearn Magazine*, October 2010. Retrieved June 9, 2015 from <http://elearnmag.acm.org/featured.cfm?aid=1872819>.
- Schroeder, N., Adesope, O., & Gilbert, R. (2013). How effective are pedagogical agents for learning? A meta-analytic review. *Journal of Educational Computing Research*, 49(1): 1–39. doi:10.2190/ec.49.1.a.
- Shaw, E., Johnson, W.L., & Ganesan, R. (1999). Pedagogical agents on the Web. In *Proceedings of the Third International Conference on Autonomous Agents*, 578–585. New York: ACM Press. doi:10.1145/301136.301210.
- Swartout, W., Artstein, R., Forbell, E., Foutz, S., Lane, H. C., Lange, B., Morie, J., Noren, D., Rizzo, S., & Traum, D. (2013). Virtual humans for learning. *AI Magazine*, 34(4), 13–30.
- Veletsianos, G. (2010). Contextually relevant pedagogical agents: visual appearance, stereotypes, and first impressions and their impact on learning. *Computers & Education*, 55(2), 576–585. doi:10.1016/j.compedu.2010.02.019.
- Wang, N., Johnson, W. L., Mayer, R. E., Rizzo, P., Shaw, E., & Collins, H. (2008). The politeness effect: pedagogical agents and learning outcomes. *International Journal of Human Computer Studies*, 66, 96–112. doi:10.1016/j.ijhcs.2007.09.003.
- Woolf, B. P., Arroyo, I., Muldner, K., Burleson, W., Cooper, D. G., Dolan, R., & Christopherson, R. M. (2010). The effect of motivational learning companions on low achieving students and students with disabilities. In *Proceedings of the 10th International Conference on Intelligent Tutoring Systems* (pp. 327–337). Berlin: Springer. doi:10.1007/978-3-642-13388-6_37.