




Pedagogical Agents

H. Chad Lane & Noah L. Schroeder



Author note:

This is a preprint. The final article is published in “The Handbook on Socially Interactive Agents” by ACM.

Citation information:


H. C. Lane and N. L. Schroeder (2022). Pedagogical Agents. In B. Lugin, C. Pelachaud, D. Traum (Eds.), *The Handbook on Socially Interactive Agents – 20 Years of Research on Embodied Conversational Agents, Intelligent Virtual Agents, and Social Robotics*, Volume 2: Interactivity, Platforms, Application (pp. 307-329). ACM.

DOI of the final chapter: [10.1145/3563659.3563669](https://doi.org/10.1145/3563659.3563669)

DOI of volume 2 of the handbook: [10.1145/3563659](https://doi.org/10.1145/3563659)

Correspondence concerning this article should be addressed to

H. Chad Lane (hclane@illinois.edu) & Noah L. Schroeder (noah.schroeder@wright.edu)



Chapter 21: Pedagogical Agents

H. Chad Lane & Noah L. Schroeder

21.1 Background and motivation

A pedagogical agent (PA) can be defined as a virtual character or physical robot that seeks to promote learning, enhance motivation, and provide support to engage in an educational activity. PAs can be non-interactive or interactive, and typically seek to emulate naturalistic communication with learners through speech, gesture, emotions, and action. In this chapter, we outline the design considerations, summarize the state of empirical research, and suggest several important directions for future research for the PA field. While our emphasis is on research conducted with virtual agents (i.e., characters on a display), our position is that many of the fundamental design aspects of PAs are applicable whether the PA is physical or virtual. Similarly, from a learning sciences point of view, a known effective pedagogical strategy should be effective regardless of the medium. That said, we certainly recognize the fundamental practical differences between intelligent virtual agents (IVA) and social robots (SR), and the potential for one to have advantages over another in given contexts. Later in the chapter, we identify some recent research to explore these important differences. Our hope is that with growing interest in educational robots, the now extensive body of research on virtual PAs can be leveraged and applied so that mistakes are not replicated and that together, the two fields can advance with shared goals and common methodologies.

There are many reasons to think that PAs can be beneficial for learning. Given that much learning occurs with other people, such as in classrooms [Kumpulainen & Wray, 2003] and museums [Leinhardt et al., 2003], it is intuitive to consider learning as a social experience involving questions, dialogue, smiling, pointing, and much more. Introducing PAs may therefore activate a more natural, social frame for learners as they gain new skills or knowledge. From a more technological perspective, in an early review of the field, Johnson, et al. (2000) explained that “agents can demonstrate complex tasks, employ locomotion and gesture to focus students’ attention on the most salient aspect of the task at hand, and convey emotional responses to the tutorial situation” (p. 47). In other words, PAs greatly expand available communicative bandwidth beyond a text-only interaction *and* they introduce the potential to virtually inhabit a learning space with a learner. This suggests a significantly wider range of potential collaborative and shared learning interactions as well as the compelling possibility of modeling learning that occurs in the real world to a higher degree of fidelity.

However, the decision to include a PA in a computer-supported learning environment is a non-trivial one that involves both theoretical and technical considerations. Including a PA in a virtual environment could potentially increase the learner’s cognitive load and possibly hinder learning [Clark & Choi, 2007; Sweller et al., 2011]. Also, from a software engineering point of view, implementing a sufficiently robust, interactive, and appealing agent can be cost and time prohibitive. To achieve a meaningful level of interactivity with a PA, significant effort is often required to achieve speech (or text) input/output, animations of nonverbal behaviors, appropriate and recognizable gestures and movement, and more. Each of these aspects of PA design can require associated expertise on behalf of the designer, who must understand both the theoretical and technical implications of different design choices. Given the

complexity of PAs, it is important to understand the history of the technology and potential benefits they can provide learners.

21.2 History of pedagogical agents

From the earliest examples of mechanical teaching machines [Pressey, 1926] to the highly interactive, immersive, social, and digital environments of today, the development of educational technologies has brought profound changes in the ways people learn with technology. While some of these changes represent the application of more general technologies to educational problems (e.g., data mining, virtual reality), others have emerged from novel properties of learning itself. For example, one of the best forms of education known to humankind is to learn from an expert human tutor in a 1-1, face-to-face interaction [Bloom, 1984]. Using what we know about effective human learning to influence the design of educational technologies captures a large portion of modern educational technology research.

Thus, it is important to understand the strategies and techniques of effective teachers and tutors so that those same productive interactions can occur in interactions with educational technology. Because an effective human tutor can focus completely on a single learner, they are able to more precisely fine tune help that is given to that learner's specific emerging needs. They can constantly assess that learner's progress, offer personalized guidance and encouragement when that learner needs it, and optimize the time available for practice. When freed from the demand of simultaneously meeting the needs of all members in a group of learners, such as in a classroom, these kinds of specialized interactions become possible. Thus, over time, computer-based learning environments have sought to emulate many of the same tactics and provide similar interactive learning experiences that human tutors are able to provide [Merrill et al., 1992]. In particular, tutoring is fundamentally a *social* experience and so the trajectory of research on computer-based tutors has been to increase their social capabilities with the hope that their effectiveness would simultaneously increase. In other words, a clear trend in the evolution of research on computer-based tutoring environments has been one of making them more like humans – that is, more communicative, social, and emotionally aware.

The earliest intelligent tutoring systems (ITSs) tended to play the role of “homework helper” by focusing exclusively on problem solving support [Anderson et al., 1995; Shute & Psotka, 1996]. Limited by the technology at the time, these early ITSs communicated by displaying text on a screen and tended to focus exclusively on cognitive aspects of learning. For example, a typical ITS would flag an action as wrong by coloring it red, or after a student requested a hint, say “You should try to combine like terms.” By providing just-in-time help and attending immediately to the needs of the student, these and similar later systems achieved impressive learning gains of about 1 sigma, roughly a letter grade [Kulik & Fletcher, 2016].

Researchers were quick to realize that much more interactivity was possible, allowing these systems to pursue more human-like interactions. Inspired by the rich conversational tactics used by human tutors [Chi et al., 2001], a second generation of systems sought to leverage Natural Language Processing (NLP) techniques to allow learners to interact more naturally using their own words to explain and ask questions. Some systems emphasized analysis of student-generated explanations to tutor questions and rely exclusively on natural language interaction, whereas others use dialogue as a tactic during problem solving to support remediation of misconceptions and to better assess learner beliefs [Graesser et al., 2001].

While PA research has existed for more than 20 years, implementing a PA within an ITS can be technically challenging. However, recent advances in animation and sound has enabled modern ITSs to be embodied – that is, to have a virtual body [Johnson & Lester, 2016]. PAs typically help students learn by providing feedback that is explanatory as well as corrective [Moreno, 2005] and personalized [Kim & Baylor, 2016]. By using nonverbal communication in the set of pedagogical strategies (e.g., nodding, eye-brow raising) the PA can utilize, the bandwidth of communication between the PA and the learner is greatly expanded. For example, when explaining concepts to the student, agents can leverage relevant supporting gestures known to facilitate thinking and learning [Davis, 2018; Goldin-Meadow, 2003]. In digital environments such as virtual worlds, an embodied agent can even inhabit the same learning space, provide orientation actions (such as pointing), and demonstrate activities in the environment for the learner, enabling learning through observation.

Body movement, gestures, and facial expressions also expand the communicative bandwidth for PAs to convey emotional messages, such as empathy and excitement. This enables implementation of more socially adept agents, which has been shown to benefit student learning in many cases (Atkinson et al., 2005), increase self-efficacy [Lane et al., 2013], and more broadly, enable systems to better motivate and engage learners [Krämer & Bente, 2010; Lane, 2016; Schroeder & Adesope, 2014]. A growing body of evidence supports claims that when designed properly, PAs enhance students' motivational and learning outcomes [Kim & Baylor, 2016; Schroeder et al., 2013]. Many of these effects are believed to be due to overt attempts to convey empathy and build rapport with human partners and engage in social conversation, which distinguish PA-based learning environments from many (non-PA enhanced) computer-based learning environments.

21.3 Designing and implementing pedagogical agents

Incorporating a PA into a computer-based learning environment entails a substantial amount of additional design and implementation work. Designing effective PAs requires an in-depth understanding of the context they are being implemented in, such as the learners who will be interacting with the PA, the information being taught, and the interaction between the learners' prior knowledge and the learning content. As such, not only is a strong understanding of learning theory necessary to know when and how to implement a PA, but a designer or design team must also have the technical skills to implement an effective PA. More specifically, in addition to deciding on the *external* properties such as the appearance of the agent as well as its role in the learning environment, defining *internal* models and properties (such as expert knowledge, learner model, and pedagogical policies) are also essential [Dehn & van Mulken, 2000]. We provide an overview of these topics below, but refer readers to Heidig and Clarebout (2011) for more detailed frameworks.

21.3.1 External properties - roles and appearance

A consistent argument given for the use of PAs is that they are believed to create a more social learning experience for a learner. While much of the success behind creating a social experience depends on the underlying technologies (e.g., natural language processing, animation quality), the basic design of the agent can have a profound impact on the experience. Two of the most fundamental design choices that must be made are (1) the role that PA will play in the learning experience and (2) what the agent will look like, how it will speak, and the agent's "personality". Table 1 provides examples of some of these important design choices.

Table 1. Examples of important elements of the design space of pedagogical agents.

Role	expert (teacher, tutor, coach) learning companion (peer) simulated role player (e.g., virtual patient for doctor training)
Appearance	fidelity (cartoon, photoreal) shape, size, species, race/ethnicity (humanoid, animal, “living” object) gender, age, size, & shape hair, eye, skin color; clothing & accessories non-verbal behaviors (gestures, body language)
Language/speech	voice type (machine-generated or recorded human speech) voice qualities (pitch, prosody, rate, disfluencies) word choice (formal, informal, colloquial/culturally aligned)
Backstory/personality	place of origin likes/dislikes sense of humor

A systematic review revealed that a large percentage of PAs reported in the literature play the role of teacher or coach [Schroeder & Gotch, 2015]. In this case, the PA takes on the traditional role of guiding the learner through learning tasks, such as solving problems or gaining conceptual understanding of some topic. PAs in this category are often responsible for structuring a learning experience by selecting problems to solve or topics to discuss, providing hints and feedback, and generally managing the learner’s experience. The next most popular role for PAs falls under the pedagogical approach of *reciprocal teaching* [Palincsar & Brown, 1984] by positioning the student as a simulated learner with the human learner acting as teacher. Such *teachable agents* implement the *learning-by-teaching* paradigm and have been explored in a variety of contexts [Biswas et al., 2016; Matsuda et al., 2013]. Sometimes a teachable agent can take the role of a peer collaborator as well, meaning that the human learner may both learn from and teach during simulated collaborative interactions, see also Chapter 22 on “Socially Interactive Agents as Peers” [Cassell, 2022] of this volume of this handbook.

The appearance of an agent is a similarly important design choice that can have consequences on learning and motivational outcomes [Veletsianos, 2010]. Basic design choices like age, gender, or ethnicity may impact how a PA is perceived and in what implicit assumptions are made by a learner. Further, the level of realism in the appearance, or whether the agent is humanoid or not, also likely influences the nature of the interactions a learner may or choose not to have with an agent. For example, research has shown that age, gender, or ethnicity of an agent may challenge or reinforce stereotypes that a learner may have about the role the agent is playing [Baylor & Kim, 2004], and research has also shown that learners can have biases against PAs based on the PAs skin tone [Zipp & Craig, 2019]. Yet, research does not always provide clear answers as to how agents should be designed in regards to their appearance. There may be benefits to offering learners a choice of agent, specifically for learners from minority and underrepresented groups [Baylor, 2005]. Relatedly, evidence has been found for the *similarity attraction hypothesis* which says that people are naturally drawn to others that look and behave in ways that are similar to them. In a study of over 500 middle school aged learners, there were motivational benefits for gender matching between a learner and a PA [Ozogul et al., 2013]. However, another study found that learners who chose to learn with PAs that

were the same ethnicity as themselves performed worse on learning tests than those who chose to learn with PAs of a different ethnicity than themselves [Moreno & Flowerday, 2006]. As such, it is important to note that such findings are challenging to generalize due to the vast individual differences between learners, including important factors like age, interests, and background.

As noted, we have only shared a few examples of a vast body of complex empirical research on the design of PAs for learning. Heidig and Clarebout (2011) provide a far more in-depth framework capturing the design aspects of implementing PAs.

23.3.2 Internal models and properties - policies, actions, and behaviors

While the external properties of a PA can be important, there is substantial evidence that the pedagogical activities of intelligent systems can and do impact learning. The best ITSs are known to improve learning by effect sizes of 1.0-1.2 standard deviations [Kulik & Fletcher, 2016; VanLehn, 2011]. Given that most of the ITSs in these studies do not include an embodied agent, it is likely that the learning gains are largely due to the pedagogical interventions provided by the system.

Dehn and van Mulken (2000) refer to these as *internal* properties of a system given that they involve pedagogical activities (e.g., assessing learner actions, deciding what feedback to give, delivering help) and do not require use of an embodied agent. The pedagogical decision space of intelligent tutors includes critical decisions points such as how to select problems, how often to give feedback and hints, what level of detail to include in hints, how to respond to emotional changes, and more [Anderson et al., 1995; VanLehn, 2006]. In this section we focus primarily on how PAs enhance and extend this design space and explore some possible features of PAs that may improve upon a disembodied tutoring system (Figure 1).

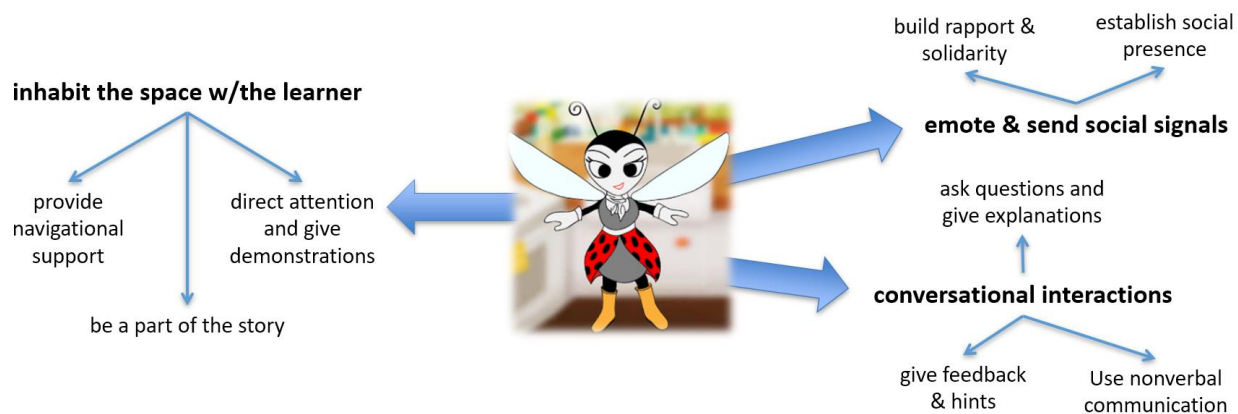


Figure 1. Examples of how pedagogical agents can interact with learners that enhance or extend the capabilities of a traditional (non-embodied) intelligent tutor. The PA shown is “Dotty” from the Virtual Sprouts project [Bell et al., 2018].

Inhabiting the same space as the learner

In the physical world, a teacher or coach exists in the space and can interact with the same objects as the learner. For example, if a learner is learning how to paint, an instructor can step in to demonstrate

how to hold a brush in their hand or perform a stroke for that learner on the canvas. PAs that exist inside virtual environments alongside a learner's avatar also have this ability. A PA might lead the learner to a specific location in a virtual world by having the learner follow them (navigational support). Naturalistic supports are possible as well by gaze changes (where the agent looks) or pointing in the virtual environment. A PA can also demonstrate a skill enabling the learner to play more the role of apprentice, or provide additional scaffolding during problem solving for a learner by doing a step for them or undoing a mistake. One of the first embodied PAs, STEVE, exhibited these actions for repair of HVAC (Heating, Ventilation, and Air Condition) systems on ships [Rickel & Johnson, 1997].

PAs can also play meaningful roles in narrative-based learning environments by portraying part of a larger story and collaborating with the learner to solve problems. A good example of this approach can be found in *Crystal Island* [Rowe et al., 2011] a multi-agent educational game for middle school biology. In the game, the learner is trapped on an island with everyone getting sick from a mysterious illness. The pedagogical agents play various roles, including scientists and medical personnel. As one the last remaining conscious inhabitants of the island, the learner must conduct experiments and work with the remaining agents to learn about their progress and unravel the mystery.

Emoting and social signals

While a certain level of emotion is certainly possible in text, such messages can be more impactful when combined with corresponding physical gestures and facial expressions. One such example can be found in *Coach Mike*, a virtual agent who supports learners at a museum exhibit for computer programming [Lane et al., 2013; Lane et al., 2011]. Coach Mike used synthesized speech along with a number of gestures and animations to build rapport with museum visitors and convey emotional messages such as empathy and excitement (figure 2). Gestures are also known to play an important role in learning, especially when they align conceptually with the content. For example, pedagogical agents that used gestures during a lesson on mathematical equivalence were more effective than agents that did not [Cook et al., 2017]. Importantly, this preliminary evidence suggests the fundamental role of gesture in promoting learning, and is impossible without the use of a PA (for computer-based learning).



Figure 2. Coach Mike, a pedagogical agent for informal computer science education that seeks to be welcoming and excited for the learner when successfully completing tasks (Lane et al., 2011).

Conversational interactions

One of the most obvious benefits of using a PA is the potential for more naturalistic communication and to achieve learning gains that approach human tutoring [Bloom, 1984]. Basic research on human communication tells us that humans instinctively naturally bring a social frame to interactions with machines [Reeves & Nass, 1996], which is promising for work on agents given that the aim is also to emulate human communication, both verbal and physical. As mentioned in section 21.2, intelligent tutors based on conversational interaction were around long before agents were possible [Graesser et al., 2001]. What is possible with a PA is to go beyond the textual components of a conversation to incorporate additional conversational behaviors such as nodding, head-shaking, eye-brow raising, and cognitive gestures (e.g., hand to chin, head scratching). These behaviors all contribute to establishing a sense of social presence, which has been positively linked to learning [Frechette & Moreno, 2010]. While not immediately applicable to achieving any particular learning goal in a conversation, such interactions do contribute to at least two related aims. First, improved conversational skills can help an agent build *rapport*, an important element of any conversation, but especially teaching and learning [Estep & Roberts, 2015; Gratch et al., 2007]. Second, building *trust* between a human and agent is widely regarded as important [Rheu et al., 2021] and related to conversational skills and relationship building. Early studies have shown, however, that trust primarily influences perceptions of an agent and, so far, has minimal impact on learning outcomes [Schroeder et al., 2021].

21.4 Empirical status of pedagogical agents

Researchers have explored the use of PAs in a wide variety of learning environments, for a diverse population of learners, and in varied contexts for more than twenty years. Meta-analyses of PA research have consistently showed that PAs can have small, positive effects on learning when compared to conditions that do not have PAs present [Castro-Alonso et al., 2021; Schroeder et al., 2013]. In contrast, systematic reviews of the literature have not found consistent, statistically significant benefits of pedagogical agents for facilitating learners' motivation [Heidig & Clarebout, 2011; Schroeder & Adesope, 2014], although Schroeder and Adesope (2014) did find more positive effects than negative. Finally, in regards to cognitive load, Schroeder and Adesope's systematic review found few significant effects, and non-significant effects were also mixed with some studies finding PAs increased cognitive load while others found PAs decreased cognitive load.

It is important to note that researchers have argued that examining the overall effectiveness of PAs across these varied contexts is too broad a question that overlooks the nuances of context and PA design [Heidig & Clarebout, 2011; Schroeder & Craig, 2021]. This is exemplified by the findings of Schroeder et al.'s (2013) meta-analysis, which found strong, positive effects when K-12 students learned from PAs. While the number of studies included in the analysis was relatively small, the effect stands in contrast to the small positive effect found for postsecondary learners.

In order to help understand the design and contextual factors that may influence the impacts of PAs, many researchers have moved beyond comparing PAs to non-PA conditions, and instead have begun investigating specific design aspects of PAs, the impacts of PAs on various types of outcomes other than learning, or exploring the use of PAs in novel contexts.

21.4.1 Research-based PA Design – Exploring Agent Voice

Designing an effective PA requires attention to a wide range of details, as illustrated in this chapter and also by Heidig and Clarebout's (2011) frameworks. Discussing all of these components in detail is outside the scope of this chapter, but we feel it is important to understand the depth of the research that takes place within each major design component. As such, in this section we explore one aspect of PA design that has been investigated in current research and is also broadly applicable: the way the agent communicates with the learner. This aspect of PA design is multi-faceted, spanning various design components such as the role of the PA, its voice and language, and its backstory. For the purposes of our example in this section, we explore only the voice and language used by a PA in detail.

A critical question when designing PAs is how they will communicate with the learner, and therefore a designer has a few critical choices to make. First, one must choose a type of voice for the PA (i.e., machine-generated or recorded human speech), which will have ramifications with regards to the qualities of the voice itself, and they also must consider different types of speech patterns.

Choosing what type of voice to use for a PA used to be a relatively simple process. Research around the voice principle [Mayer, 2014] had shown that in multimedia learning situations, recorded human speech consistently lead to higher learning outcomes than machine-generated, or text-to-speech (TTS), voices [Atkinson et al., 2005; Mayer et al., 2003]. However, there is a growing body of research investigating the type of voice one should use to aid learning in multimedia learning situations [Craig & Schroeder, 2019] and, more specifically, the type of voice that a PA should use to communicate with the learner [Chiou et al., 2020; Craig & Schroeder, 2017], which is showing that modern TTS voices can facilitate learning as well as recorded human voices. As such, PA designers should consider their software and its technical requirements as part of their decision-making process when picking a type of voice for the PA to use. Specifically, if a system needs to be able to reply with a wide variety of responses, or can create its own responses rather than replying with pre-made phrases, this may be more easily implemented with a TTS engine rather than a recorded human voice. In contrast, if one is creating a video that contains a PA but the learner does not interact or converse with the PA, they simply listen to a recording, a recorded human voice may be preferable in some circumstances. To conclude our discussion of voice type, it is important to note that research cited in this section was often dealing with PAs that did not naturally and adaptively converse with the learner, but rather conveyed either pre-scripted speech or were simple recordings of PA delivering a narrative. It is not clear if TTS or recorded human voices would be superior for facilitating learning in systems that generate unique responses in the course of a conversation with a learner.

While the choice of TTS compared to recorded human speech may seem relatively straightforward due to software requirements, one must also consider the qualities of the voice itself. For example, the prosody of the voice may be important. Prosody refers to the stress, intonation, and rhythm of speech that can provide important information to a listener [American Psychological Association, 2021; Davis et al., 2019; Shintel et al., 2014]. This is potentially an important factor in PA design, especially if the learner is communicating in their non-native language [Davis et al., 2019]. Few studies have examined voice prosody in relation to PAs specifically. However, Davis et al. (2019) investigated the use of PAs with strong or weak voice prosody for teaching non-native speakers. Their results revealed no statistically significant differences on learning outcomes, but they concluded that more research is

needed to better understand what features of human voices are important for facilitating learning. A PA designer must then decide how important prosody is in their context. This potentially has implications for if they can use a TTS voice or would be better served by a recorded human voice, depending on the software they have available. Please see Chapter 6 on “Building and Designing Expressive Speech Synthesis” [Aylett et al., 2021] of volume 1 of this handbook [Lugrin et al., 2021] on expressive speech synthesis for further details.

21.4.2 Communication Styles

A PA designer must also decide how the PA will use language when communicating with the learner. In the research literature, this is often characterized by using a conversational style compared to a more formal style. In the broader literature outside of PAs specifically, researchers have investigated this question in some depth. A meta-analysis found that using a conversational style aided both retention and transfer outcomes, with a noteworthy effect on learning transfer outcomes, $d = .54$ [Ginns et al., 2013]. Importantly, the transfer outcomes were not significantly moderated by either the specific language (e.g., English, German, etc.) nor the learners’ grade level, and Ginns et al. found that the conversational style provided strong benefits for effective cognitive processing ($d = .62$). It is important to note however, that Ginns et al. found that the length of the intervention significantly moderated the transfer outcomes, and few studies were longer than 35 minutes in length. As such, it seems that, especially for relatively short interventions, conversational style language may benefit learning.

21.4.3 Summary of Empirical Status

Veletsianos and Russell (2014) observed that while PAs have often been associated with providing learning benefits to learners, it is not necessarily their presence alone that typically facilitates learning, but rather, it is often the pedagogical strategies that the PA facilitates that may benefit learning compared to non-PA conditions. In addition, Craig and Schroeder (2018) suggested that PAs may benefit learning most when the social aspects of their use can be leveraged, such as when they can act as conversational partners that convey emotions and model and facilitate effective learning strategies. As such, it is prudent that designers consider what educational benefits a PA could provide in their specific system within their context, and design each aspect of the PA as appropriate.

There is a vast amount of research around many of the design choices one needs to make when designing a PA. It is also important to remember that there can be interactions between different PA design components. For example, the role the agent plays in the environment may also influence how an agent should communicate, either conversationally or more formally, and thus a designer should consider all aspects of their PA design together rather than in isolation.

21.5 Relationship with IVA and SR research

Research on IVA and SR has many shared goals and similar challenges, with the key (obvious) difference being the use of a physical or virtual agent/robot. As mentioned in chapter 1 on “Introduction to Socially Interactive Agents” [Lugrin, 2021] of volume 1 of this handbook [Lugrin et al., 2021], comparisons between these two modalities are revealing strengths and weaknesses of both. The challenges, costs,

and limitations of a robot may present challenges to implementation, however the benefits of interacting with the physical world of a learner has great potential for supporting learning in deeper ways. For example, learning very often occurs in a social context (classrooms, museums, outdoors) and so bringing a PA from a computer screen into these physical contexts has obvious benefits, such as in-context guidance and coordinating group work on a collaborative, physical task (e.g., planting a garden). Emerging research on using SR for education is beginning to explore these benefits [Belpaeme et al., 2018; Mubin et al., 2013].

Research in a wide range of areas has clear applications to both IVAs and SRs. For example, the challenge of increasing naturalistic interactions with speech and language understanding, natural language generation, dialogue modeling, animation of nonverbal behaviors and more have value broadly for IVA and SR, and specifically for PAs. In one study of college students, robots were viewed as credible instructors by students (Edwards et al., 2016), due in part to believable and natural interactions that resulted from remote human control (e.g., “Wizard of Oz”). Researchers have also shown that robots that emote are more effective in terms of supporting learners, due most likely to feelings of empathy and concern that learners experience [Obaid et al., 2018; Saerbeck et al., 2010]. Clearly, underlying algorithms that detect, express, and enable naturalistic interactions can contribute to both virtual and physical PAs.

21.6 Current challenges

Pedagogical agent researchers face a variety of challenges in their work. Notably, PAs can be time and resource-intensive to develop. One positive outcome of this challenge is that researchers have developed systems that greatly vary, from individual agents delivering a narrative without any back and forth dialogue with the user [Schroeder, 2017], through intelligent tutoring systems with one or more PAs [Graesser et al., 2017], teachable agents [Silvervarg et al., 2021], and even virtual worlds [Rowe et al., 2011]. However, the challenge of creating PAs persists. To date, there are few user-friendly software programs with a graphic-user interface that can help a researcher create a PA, and likely fewer yet that allow for the development of interactive PAs. While some platforms do exist to facilitate the interactive PA creation process [Hartholt et al., 2013; Lane et al., 2015], the general process of authoring a new intelligent agent, embodied or not, still represents one of the most persistent and complex challenges for the field [Dermeval et al., 2018; Murray et al., 2003; Sottolare et al., 2017], in part due to its interdisciplinary nature.

Like many fields, the PA literature is also limited by the lack of psychometrically strong measurement tools within the field. These limitations were noted by researchers more than 15 years ago [Clark & Choi, 2005], however relatively few validated measures developed specifically for use with PAs exist in the literature. In fact, much of the literature still relies on the Agent Persona Instrument [Ryu & Baylor, 2005] or Schroeder et al.’s (2018) revision of it. While this instrument has been useful for examining different aspects of PA persona, researchers have found that it does not strongly predict or influence learning outcomes [Schroeder, 2017; Schroeder et al., 2018]. The field would benefit from new instruments being developed which can help push the field forward. For example, while a PA may only facilitate a small learning benefit in some circumstances, which may not rationalize the cost and effort required to create it, there may be other benefits (e.g., social or emotional outcomes) from the agent that are viewed as more important. However, without adequate measurement techniques it is challenging to approach those areas of work.

Finally, the field has also found challenges in creating PAs that express emotions well, and understanding if and how these emotions impact learners. However, recent work is making advancements in this area. For example, Lawson et al. (2021b) compared a human instructor and a PA, and found that learners could differentiate positive and negative emotions in both cases. In a different study, Lawson et al. (2021a) found that learners were able to differentiate a PA's emotions, and generally perceived positive PAs better than negative PAs. Meanwhile, other lines of work have examined how to best create believable PA emotions. Anasingaraju et al. (2020) found that the PAs body was the most important aspect of creating the believable expression of emotions. Together with basic research in affective computing to model and recognize emotions [Calvo et al., 2015], these represent significant ongoing challenges for the field.

21.7 Future directions

While many ongoing research areas related to PAs will continue to see much attention in coming years, some areas remain elusive. In this section we provide our opinions on what the field should emphasize in the coming years.

21.7.1 PA Support for lifelong learning

Much of what we know about learning from PAs and specific aspects of PA design comes from studies that are of relatively short duration. For example, many studies investigating specific design aspects of PAs contain interventions less than 10 minutes in duration [Chiou et al., 2020; Davis, 2018], whereas a longer duration study may be a few hours [Li & Graesser, 2021]. This being the case, a long-standing question in PA research is how learners will learn from [Choi & Clark, 2006; Gulz, 2004] and develop relationships with PAs over longer durations [Veletsianos & Russell, 2014]. For example, one question posed by researchers is if the benefits of PAs are simply due to novelty effects that may wear off over time [Choi & Clark, 2006]. Importantly however, Veletsianos and Russell (2014) raised a number of questions about the ethics related to designing agents that make emotional connections between the learner and the agent. In terms of technical demands of supporting lifelong learning, a PA would require a long-term model of a learner that spans different domains and contexts. This challenge is well-known [Kay, 2008] and different approaches have been pursued that seek to inform pedagogical decision making, such as meeting learner preferences and modeling memory decay [Kay & Kummerfeld, 2012; Pavlik et al., 2020]. This area surely needs more research to address it, especially if longer-term PA interventions are pursued by researchers.

21.7.2 PA Support for medical and health literacy

Research on building conversational agents for health literacy and supporting medical information needs has grown substantially in recent years [Montenegro et al., 2019; Morrow et al., 2021]. Given the cost of human-human support, and the challenges associated with printed materials, this line of research seeks to provide intuitive and natural ways for acquiring information about diagnoses, treatments, and self-care. In overlapping with support for lifelong learning, Bickmore et al. have implemented numerous agents that provide such support [Bickmore & Picard, 2005]. This work has provided a framework for agent-human relationship building as well as methods for behavior change [Bickmore et al., 2013]. Agents have particularly been successful with aging adults and supporting complex medication schedules and self-care techniques [Azevedo et al., 2018]. While such systems have implicitly recognized patients as learners, such as by answering questions and providing information,

there remains a long list of techniques from PA research that focus on supporting learning that have potential to transform agents in the area of health literacy. Please consult chapter 24 on “Health-Related Applications of Socially Interactive Agents” [Bickmore, 2022] of this handbook for further details.

21.7.3 New and Emerging Areas for PA Research

While researchers have long-examined the effects of PAs on learning, researchers have also began looking at factors other than learning that PAs may influence, such as learners’ trust in the agent. For example, Chiou et al. (2020) examined how different types of PA voices influenced learners trust in the PA. Outcomes such as trust and credibility, as well as other outcomes such as social presence, may be particularly important as researchers are moving PAs beyond science and mathematics teaching contexts to more sensitive topics such as health (section 21.7.2) and counseling [King et al., 2020].

Another emerging line of research is how various learner perceptions of the PA can influence learning outcomes [Schroeder et al., 2021; Schroeder et al., 2017; Schroeder et al., 2018]. Presumably, the goal of these lines of work is work towards identifying the specific outcomes that most influence learning, then researchers can isolate specific design features to help facilitate these outcomes and improve learning from PAs.

21.8 Summary

In this chapter we have summarized research on PAs by describing the “external” design considerations (age, gender, ethnicity, form, etc.), the “internal” properties that define how such agents can promote learning, and documented specific interaction techniques for which PAs are unique or can improve upon previous non-embodied intelligent tutoring systems. We summarized current research on PAs and provided an overview of findings as they relate to the design of agents, pedagogical strategies, and roles. While empirical evidence is modest to date, it is nonetheless positive for many contexts and learners (young and older). In the future, we suggest that PA research should focus on providing longer term support, across domains, integrate learning technologies into health literacy agents, and continue to focus on specific ways that PAs connect with learners either via enhanced trust, social presence, and quality of learner experiences.

References

- Anasingaraju, S., Adamo-Villani, N., & Dib, H. N. (2020). The Contribution of Different Body Channels to the Expression of Emotion in Animated Pedagogical Agents. *International Journal of Technology and Human Interaction (IJTHI)*, 16(4), 70-88.
- Anderson, J. R., Corbett, A., Koedinger, K., & Pelletier, R. (1995). Cognitive tutors: Lessons learned. *Journal of the Learning Sciences*, 4(2), 167-207.
- Atkinson, R. K., Mayer, R. E., & Merrill, M. M. (2005). Fostering social agency in multimedia learning: Examining the impact of an animated agent's voice. *Contemporary Educational Psychology*, 30(1), 117-139. <https://doi.org/http://dx.doi.org/10.1016/j.cedpsych.2004.07.001>
- Aylett, R. (2022). Interactive narrative and story-telling. In B. Lugrin, C. Pelachaud, and D. Traum (Eds.), *The Handbook on Socially Interactive Agents: 20 years of Research on Embodied Conversational Agents, Intelligent Virtual Agents, and Social Robotics Volume 2: Interactivity, Platforms, Application*. ACM Press, 463–491. DOI: <http://dx.doi.org/10.1145/3563659.3563674>.
- Azevedo, R. F. L., Morrow, D., Graumlich, J., Willemsen-Dunlap, A., Hasegawa-Johnson, M., Huang, T. S., . . . Halpin, D. J. (2018). Using conversational agents to explain medication instructions to older adults. *AMIA ... Annual Symposium proceedings. AMIA Symposium, 2018*, 185-194.
- Baylor, A. L. (2005). The impact of pedagogical agent image on affective outcomes. International Conference on Intelligent User Interfaces, San Diego, CA,
- Baylor, A. L., & Kim, Y. (2004). Pedagogical Agent Design: The Impact of Agent Realism, Gender, Ethnicity, and Instructional Role. In J. C. Lester, R. M. Vicari, & F. Paraguaçu, *Intelligent Tutoring Systems* Berlin, Heidelberg.
- Bell, B. M., Martinez, L., Gotsis, M., Lane, H. C., Davis, J. N., Antunez-Castillo, L., . . . Spruijt-Metz, D. (2018). Virtual Sprouts: A Virtual Gardening Pilot Intervention Increases Self-Efficacy to Cook and Eat Fruits and Vegetables in Minority Youth. *Games for Health Journal*, 7(2), 127-135. <https://doi.org/10.1089/g4h.2017.0102>
- Belpaeme, T., Kennedy, J., Ramachandran, A., Scassellati, B., & Tanaka, F. (2018). Social robots for education: A review. *Science robotics*, 3(21).
- Bickmore, T. W. (2022). Health-related applications of socially interactive agents. In B. Lugrin, C. Pelachaud, and D. Traum (Eds.), *The Handbook on Socially Interactive Agents: 20 years of Research on Embodied Conversational Agents, Intelligent Virtual Agents, and Social Robotics Volume 2: Interactivity, Platforms, Application*. ACM Press, 403–435. DOI: <http://dx.doi.org/10.1145/3563659.3563672>.
- Bickmore, T. W., & Picard, R. W. (2005). Establishing and maintaining long-term human-computer relationships. *ACM Transactions on Computer-Human Interaction*, 12(2), 293-327. <https://doi.org/10.1145/1067860.1067867>
- Bickmore, T. W., Schulman, D., & Sidner, C. (2013). Automated interventions for multiple health behaviors using conversational agents. *Patient Education and Counseling*, 92(2), 142-148. <https://doi.org/https://doi.org/10.1016/j.pec.2013.05.011>
- Biswas, G., Segedy, J. R., & Bunchongchit, K. (2016). From design to implementation to practice a learning by teaching system: Betty's Brain. *International Journal of Artificial Intelligence in Education*, 26(1), 350-364.
- Bloom, B. S. (1984). The 2 sigma problem: The search for methods of group instruction as effective as one-to-one tutoring. *Educational Researcher*, 13(6), 4-16.
- Calvo, R. A., D'Mello, S., Gratch, J. M., & Kappas, A. (2015). *The Oxford handbook of affective computing*. Oxford Library of Psychology.
- Cassell, J. (2022). Socially interactive agents as peers. In B. Lugrin, C. Pelachaud, and D. Traum (Eds.), *The Handbook on Socially Interactive Agents: 20 years of Research on Embodied Conversational*

- Agents, Intelligent Virtual Agents, and Social Robotics Volume 2: Interactivity, Platforms, Application*. ACM Press, 331–365. DOI: <http://dx.doi.org/10.1145/3563659.3563670>.
- Castro-Alonso, J. C., Wong, R. M., Adesope, O. O., & Paas, F. (2021). Effectiveness of Multimedia Pedagogical Agents Predicted by Diverse Theories: a Meta-Analysis. In: Springer.
- Chi, M. T. H., Siler, S. A., Jeong, H., Yamauchi, T., & Hausmann, R. G. (2001). Learning from human tutoring. *Cognitive Science*, 25(4), 471-533.
- Chiou, E. K., Schroeder, N. L., & Craig, S. D. (2020). How we trust, perceive, and learn from virtual humans: The influence of voice quality. *Computers & Education*, 146, 103756. <https://doi.org/https://doi.org/10.1016/j.compedu.2019.103756>
- Choi, S., & Clark, R. E. (2006). Cognitive and Affective Benefits of an Animated Pedagogical Agent for Learning English as a Second Language. *Journal of Educational Computing Research*, 34(4), 441-466. <https://doi.org/10.2190/a064-u776-4208-n145>
- Clark, R. E., & Choi, S. (2005). Five design principles for experiments on the effects of animated pedagogical agents. *Journal of Educational Computing Research*, 32(3), 209-225.
- Clark, R. E., & Choi, S. (2007). The questionable benefits of pedagogical agents: Response to Veletsianos. *Journal of Educational Computing Research*, 36(4), 379-381.
- Cook, S. W., Friedman, H. S., Duggan, K. A., Cui, J., & Popescu, V. (2017). Hand gesture and mathematics learning: lessons from an Avatar. *Cognitive Science*, 41(2), 518-535.
- Craig, S. D., & Schroeder, N. L. (2018). Design principles for virtual humans in educational technology environments. In *Deep Comprehension* (pp. 128-139). Routledge.
- Davis, R. O. (2018). The impact of pedagogical agent gesturing in multimedia learning environments: A meta-analysis. *Educational Research Review*, 24, 193-209.
- Dehn, D. M., & van Mulken, S. (2000). The impact of animated interface agents: a review of empirical research. *Int. J. Hum.-Comput. Stud.*, 52(1), 1-22. <https://doi.org/http://dx.doi.org/10.1006/ijhc.1999.0325>
- Dermeval, D., Paiva, R., Bittencourt, I. I., Vassileva, J., & Borges, D. (2018). Authoring Tools for Designing Intelligent Tutoring Systems: a Systematic Review of the Literature. *International Journal of Artificial Intelligence in Education*, 28(3), 336-384. <https://doi.org/10.1007/s40593-017-0157-9>
- Edwards, A., Edwards, C., Spence, P. R., Harris, C., & Gambino, A. (2016). Robots in the classroom: Differences in students' perceptions of credibility and learning between "teacher as robot" and "robot as teacher". *Computers in Human Behavior*, 65, 627-634. <https://doi.org/https://doi.org/10.1016/j.chb.2016.06.005>
- Estep, C. M., & Roberts, T. (2015). Teacher immediacy and professor/student rapport as predictors of motivation and engagement. *NACTA Journal*, 59(2), 155-163.
- Frechette, C., & Moreno, R. (2010). The Roles of Animated Pedagogical Agents' Presence and Nonverbal Communication in Multimedia Learning Environments. *Journal of Media Psychology*, 22(2), 61-72. <https://doi.org/10.1027/1864-1105/a000009>
- Ginns, P., Martin, A. J., & Marsh, H. W. (2013). Designing instructional text in a conversational style: A meta-analysis. *Educational Psychology Review*, 25(4), 445-472.
- Goldin-Meadow, S. (2003). *Hearing gesture: How our hands help us think*. Harvard University Press.
- Graesser, A. C., Forsyth, C. M., & Lehman, B. A. (2017). Two Heads May Be Better than One: Learning from Computer Agents in Conversational Dialogues. *Teachers College Record*, 119(3), 1-20.
- Graesser, A. C., VanLehn, K., Rose, C. P., Jordan, P. W., & Harter, D. (2001). Intelligent tutoring systems with conversational dialogue. *AI Mag.*, 22(4), 39-51.
- Gratch, J., Wang, N., Gerten, J., Fast, E., & Duffy, R. (2007). Creating Rapport with Virtual Agents. In C. Pelachaud, J. Martin, E. Andre, G. Chollet, K. Karpouzis, & D. Pele (Eds.), *Proceedings of the 7th international conference on Intelligent Virtual Agents* (pp. 125-138). Springer-Verlag. https://doi.org/http://dx.doi.org/10.1007/978-3-540-74997-4_12

- Gulz, A. (2004). Benefits of virtual characters in computer based learning environments: Claims and evidence. *International Journal of Artificial Intelligence in Education*, 14(3), 313-334.
- Hartholt, A., Traum, D., Marsella, S. C., Shapiro, A., Stratou, G., Leuski, A., . . . Gratch, J. (2013). All Together Now. In R. Aylett, B. Krenn, C. Pelachaud, & H. Shimodaira, *Intelligent Virtual Agents* Berlin, Heidelberg.
- Heidig, S., & Clarebout, G. (2011). Do pedagogical agents make a difference to student motivation and learning? *Educational Research Review*, 6(1), 27-54.
- Johnson, W. L., & Lester, J. C. (2016). Face-to-Face Interaction with Pedagogical Agents, Twenty Years Later [journal article]. *International Journal of Artificial Intelligence in Education*, 26(1), 25-36. <https://doi.org/10.1007/s40593-015-0065-9>
- Johnson, W. L., Rickel, J., & Lester, J. C. (2000). Animated Pedagogical Agents: Face-to-Face Interaction in Interactive Learning Environments. *International Journal of Artificial Intelligence in Education*, 11, 47 - 48. http://ihelp.usask.ca/iaied/ijaied/abstracts/Vol_11/johnson.html
- Kay, J. (2008). Lifelong learner modeling for lifelong personalized pervasive learning. *IEEE Transactions on Learning Technologies*, 1(4), 215-228.
- Kay, J., & Kummerfeld, B. (2012). Lifelong learner modeling. *Adaptive Technologies for Training and Education*, 140-164.
- Kim, Y., & Baylor, A. L. (2016). Research based design of pedagogical agent roles: A review, progress, and recommendations. *International Journal of Artificial Intelligence in Education*, 26(1), 160-169.
- King, A. C., Campero, M. I., Sheats, J. L., Castro Sweet, C. M., Hauser, M. E., Garcia, D., . . . Bickmore, T. (2020). Effects of Counseling by Peer Human Advisors vs Computers to Increase Walking in Underserved Populations: The COMPASS Randomized Clinical Trial. *JAMA Internal Medicine*, 180(11), 1481-1490. <https://doi.org/10.1001/jamainternmed.2020.4143>
- Krämer, N., & Bente, G. (2010). Personalizing e-Learning. The Social Effects of Pedagogical Agents. *Educational Psychology Review*, 22(1), 71-87. <https://doi.org/10.1007/s10648-010-9123-x>
- Kulik, J. A., & Fletcher, J. (2016). Effectiveness of intelligent tutoring systems: a meta-analytic review. *Review of Educational Research*, 86(1), 42-78.
- Kumpulainen, K., & Wray, D. (Eds.). (2003). *Classroom Interaction and Social Learning: From Theory to Practice*. Routledge.
- Lane, H. C. (2016). Pedagogical Agents and Affect: Molding Positive Learning Interactions. In S. Y. Tettegah & M. Gartmeier (Eds.), *Emotions, Technology, Design, & Learning* (pp. 47-61). Academic Press.
- Lane, H. C., Cahill, C., Foutz, S., Auerbach, D., Noren, D., Lussenhop, C., & Swartout, W. (2013). The Effects of a Pedagogical Agent for Informal Science Education on Learner Behaviors and Self-efficacy. In K. Yacef, H. C. Lane, & J. P. Mostow, Phillip (Eds.), *Proceedings of the 16th International Conference on Artificial Intelligence in Education (AIED2013)* (Vol. 6738, pp. 309-318). Springer. https://doi.org/10.1007/978-3-642-39112-5_32
- Lane, H. C., Core, M. G., Hays, M. J., Auerbach, D., & Rosenberg, M. (2015, 2015//). Situated Pedagogical Authoring: Authoring Intelligent Tutors from a Student's Perspective. *Artificial Intelligence in Education*, Cham.
- Lane, H. C., Noren, D., Auerbach, D., Birch, M., & Swartout, W. (2011). Intelligent tutoring goes to the museum in the big city: A pedagogical agent for informal science education. In G. Biswas & S. Bull (Eds.), *Proceedings of the 15th International Conference on Artificial Intelligence in Education (AIED2011)* (Vol. 6738, pp. 155-162). Springer. https://doi.org/10.1007/978-3-642-21869-9_22
- Lawson, A. P., Mayer, R. E., Adamo-Villani, N., Benes, B., Lei, X., & Cheng, J. (2021a). Do Learners Recognize and Relate to the Emotions Displayed By Virtual Instructors? *International Journal of Artificial Intelligence in Education*, 31(1), 134-153. <https://doi.org/10.1007/s40593-021-00238-2>

- Lawson, A. P., Mayer, R. E., Adamo-Villani, N., Benes, B., Lei, X., & Cheng, J. (2021b). Recognizing the emotional state of human and virtual instructors. *Computers in Human Behavior*, *114*, 106554. <https://doi.org/https://doi.org/10.1016/j.chb.2020.106554>
- Leinhardt, G., Crowley, K., & Knutson, K. (2003). *Learning conversations in museums*. Taylor & Francis.
- Li, H., & Graesser, A. C. (2021). The impact of conversational agents' language on summary writing. *Journal of Research on Technology in Education*, *53*(1), 44-66. <https://doi.org/10.1080/15391523.2020.1826022>
- Lugrin, B. (2021). Introduction to socially interactive agents. In B. Lugrin, C. Pelachaud, and D. Traum (Eds.), *The Handbook on Socially Interactive Agents: 20 years of Research on Embodied Conversational Agents, Intelligent Virtual Agents, and Social Robotics Volume 1: Methods, Behavior, Cognition*. ACM Press, 1–18. DOI: <http://dx.doi.org/10.1145/3477322>. 3477324.
- Lugrin, B., Pelachaud, C., and Traum, D. (Eds.). (2021). *The Handbook on Socially Interactive Agents: 20 years of Research on Embodied Conversational Agents, Intelligent Virtual Agents, and Social Robotics Volume 1: Methods, Behavior, Cognition*. ACM Press, 538 pages. DOI: <https://doi.org/10.1145/3477322>.
- Matsuda, N., Yarzebinski, E., Keiser, V., Raizada, R., Cohen, W. W., Stylianides, G. J., & Koedinger, K. R. (2013). Cognitive anatomy of tutor learning: Lessons learned with SimStudent. *Journal of Educational Psychology*, *105*(4), 1152.
- Merrill, D. C., Reiser, B. J., Ranney, M., & Trafton, J. G. (1992). Effective tutoring techniques: A comparison of human tutors and intelligent tutoring systems. *The Journal of the Learning Sciences*, *2*(3), 277-305.
- Montenegro, J. L. Z., da Costa, C. A., & da Rosa Righi, R. (2019). Survey of conversational agents in health. *Expert Systems with Applications*, *129*, 56-67. <https://doi.org/https://doi.org/10.1016/j.eswa.2019.03.054>
- Moreno, R. (2005). Multimedia learning with animated pedagogical agents. In R. Mayer (Ed.), *Cambridge Handbook on Multimedia Learning* (pp. 507-524). Cambridge University Press.
- Moreno, R., & Flowerday, T. (2006). Students' choice of animated pedagogical agents in science learning: A test of the similarity-attraction hypothesis on gender and ethnicity. *Contemporary Educational Psychology*, *31*(2), 186-207.
- Morrow, D. G., Lane, H. C., & Rogers, W. A. (2021). A Framework for Design of Conversational Agents to Support Health Self-Care for Older Adults. *Human Factors*, *63*(3), 369-378. <https://doi.org/10.1177/0018720820964085>
- Mubin, O., Stevens, C. J., Shahid, S., Al Mahmud, A., & Dong, J.-J. (2013). A review of the applicability of robots in education. *Journal of Technology in Education and Learning*, *1*(209-0015), 13.
- Murray, T., Blessing, S., & Ainsworth, S. (2003). *Authoring Tools for Advanced Technology Learning Environments*. Kluwer Academic Publishers.
- Obaid, M., Aylett, R., Barendregt, W., Basedow, C., Corrigan, L. J., Hall, L., . . . Castellano, G. (2018). Endowing a Robotic Tutor with Empathic Qualities: Design and Pilot Evaluation. *International Journal of Humanoid Robotics*, *15*(06), 1850025. <https://doi.org/10.1142/s0219843618500251>
- Ozogul, G., Johnson, A. M., Atkinson, R. K., & Reisslein, M. (2013). Investigating the impact of pedagogical agent gender matching and learner choice on learning outcomes and perceptions. *Computers & Education*, *67*, 36-50.
- Palincsar, A. S., & Brown, A. L. (1984). Reciprocal Teaching of Comprehension-Fostering and Comprehension-Monitoring Activities. *Cognition and Instruction*, *1*(2), 117-175. <http://www.jstor.org/stable/3233567>
- Pavlik, P., Eglinton, L. G., & Harrell-Williams, L. M. (2020). Generalized Knowledge Tracing: A constrained framework for learner modeling.

- Pressey, S. L. (1926). A simple apparatus which gives tests and scores - and teaches. *School and Society*, 23(586), 373-376.
- Reeves, B., & Nass, C. (1996). *The media equation: how people treat computers, television, and new media like real people and places*. Cambridge University Press.
- Rheu, M., Shin, J. Y., Peng, W., & Huh-Yoo, J. (2021). Systematic Review: Trust-Building Factors and Implications for Conversational Agent Design. *International Journal of Human-Computer Interaction*, 37(1), 81-96. <https://doi.org/10.1080/10447318.2020.1807710>
- Rickel, J., & Johnson, W. L. (1997). *Integrating pedagogical capabilities in a virtual environment agent* Proceedings of the first international conference on Autonomous agents, Marina del Rey, California, USA.
- Rowe, J. P., Shores, L. R., Mott, B. W., & Lester, J. C. (2011). Integrating Learning, Problem Solving, and Engagement in Narrative-Centered Learning Environments. *International Journal of Artificial Intelligence in Education*, 21(1), 115-133. <https://doi.org/10.3233/jai-2011-019>
- Ryu, J., & Baylor, A. L. (2005). The psychometric structure of pedagogical agent persona. *Technology Instruction, Cognition, and Learning*, 2(4), 291-315.
- Saerbeck, M., Schut, T., Bartneck, C., & Janse, M. D. (2010). Expressive robots in education: varying the degree of social supportive behavior of a robotic tutor. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 1613–1622). Association for Computing Machinery. <https://doi.org/10.1145/1753326.1753567>
- Schroeder, N. L. (2017). The Influence of a Pedagogical Agent on Learners' Cognitive Load. *Journal of Educational Technology & Society*, 20(4), 138-147. <http://www.jstor.org/stable/26229212>
- Schroeder, N. L., & Adesope, O. O. (2014). A Systematic Review of Pedagogical Agents' Persona, Motivation, and Cognitive Load Implications for Learners. *Journal of Research on Technology in Education*, 46(3), 229-251.
- Schroeder, N. L., Adesope, O. O., & Gilbert, R. B. (2013). How effective are pedagogical agents for learning? A meta-analytic review. *Journal of Educational Computing Research*, 49(1), 1-39.
- Schroeder, N. L., Chiou, E. K., & Craig, S. D. (2021). Trust influences perceptions of virtual humans, but not necessarily learning. *Computers & Education*, 160, 104039. <https://doi.org/https://doi.org/10.1016/j.compedu.2020.104039>
- Schroeder, N. L., & Craig, S. D. (2021). Learning with virtual humans: Introduction to the special issue. *Journal of Research on Technology in Education*, 53(1), 1-7. <https://doi.org/10.1080/15391523.2020.1863114>
- Schroeder, N. L., & Gotch, C. M. (2015). Persisting issues in pedagogical agent research. *Journal of Educational Computing Research*, 53(2), 183-204.
- Schroeder, N. L., Romine, W. L., & Craig, S. D. (2017). Measuring pedagogical agent persona and the influence of agent persona on learning. *Computers & Education*, 109, 176-186. <https://doi.org/https://doi.org/10.1016/j.compedu.2017.02.015>
- Schroeder, N. L., Yang, F., Banerjee, T., Romine, W. L., & Craig, S. D. (2018). The influence of learners' perceptions of virtual humans on learning transfer. *Computers & Education*, 126, 170-182. <https://doi.org/https://doi.org/10.1016/j.compedu.2018.07.005>
- Shute, V. J., & Psotka, J. (1996). Intelligent tutoring systems: Past, present, and future. In D. H. Jonassen (Ed.), *Handbook for research for educational communications and technology* (pp. 570-599). Macmillan.
- Silvarg, A., Wolf, R., Blair, K. P., Haake, M., & Gulz, A. (2021). How teachable agents influence students' responses to critical constructive feedback. *Journal of Research on Technology in Education*, 53(1), 67-88. <https://doi.org/10.1080/15391523.2020.1784812>
- Sottolare, R. A., Brawner, K. W., Sinatra, A. M., & Johnston, J. H. (2017). An updated concept for a Generalized Intelligent Framework for Tutoring (GIFT). *GIFTtutoring.org*, 1-19.

- Sweller, J., Ayres, P. L., & Kalyuga, S. (2011). Cognitive load theory. <http://site.ebrary.com/id/10461613>
- VanLehn, K. (2006). The Behavior of Tutoring Systems. *International Journal of Artificial Intelligence in Education*, 16(3), 227-265.
- VanLehn, K. (2011). The Relative Effectiveness of Human Tutoring, Intelligent Tutoring Systems, and Other Tutoring Systems. *Educational Psychologist*, 46(4), 197-221. <https://doi.org/10.1080/00461520.2011.611369>
- Veletsianos, G. (2010). Contextually relevant pedagogical agents: Visual appearance, stereotypes, and first impressions and their impact on learning. *Computers & Education*, 55(2), 576-585.
- Veletsianos, G., & Russell, G. (2014). Pedagogical Agents. In J. M. Spector, M. D. Merrill, J. Elen, & M. J. Bishop (Eds.), *Handbook of Research on Educational Communications and Technology* (pp. 759-769). Springer New York. https://doi.org/10.1007/978-1-4614-3185-5_61
- Zipp, S. A., & Craig, S. D. (2019). The impact of a user's biases on interactions with virtual humans and learning during virtual emergency management training. *Educational Technology Research and Development*, 67(6), 1385-1404.