UC Berkeley

UC Berkeley Previously Published Works

Title

Pedagogical agents to support embodied, discovery-based learning

Permalink

<https://escholarship.org/uc/item/9vr3990s>

ISBN

9783319674001

Authors

Abdullah, A
Adil, M
Rosenbaum, L
et al.

Publication Date

2017

DOI

10.1007/978-3-319-67401-8_1

Peer reviewed

Pedagogical Agents to Support Embodied, Discovery-based Learning

Ahsan Abdullah¹ ✉, Mohammad Adil¹, Leah Rosenbaum², Miranda Clemmons², Mansi Shah², Dor Abrahamson², and Michael Neff

¹ University of California, Davis**

aabdullah@ucdavis.edu, madil@ucdavis.edu, mpneff@ucdavis.edu

² University of California, Berkeley

leahr@berkeley.edu, mclemmons@berkeley.edu, emansishah@berkeley.edu,
dor@berkeley.edu

Abstract. This paper presents a pedagogical agent designed to support students in an embodied, discovery-based learning environment. Discovery-based learning guides students through a set of activities designed to foster particular insights. In this case, the animated agent explains how to use the Mathematical Imagery Trainer for Proportionality, provides performance feedback, leads students to have different experiences and provides remedial instruction when required. It is a challenging task for agent technology as the amount of concrete feedback from the learner is very limited, here restricted to the location of two markers on the screen. A Dynamic Decision Network is used to automatically determine agent behavior, based on a deep understanding of the tutorial protocol. A pilot evaluation showed that all participants developed movement schemes supporting proto-proportional reasoning. They were able to provide verbal proto-proportional expressions for one of the taught strategies, but not the other.

Keywords: pedagogical agents, discovery-based learning, dynamic decision networks

1 Introduction

Discovery-based learning is an educational activity paradigm whereby students are led through well-specified experiences that are designed to foster particular insights relevant to curricular objectives. It differs from many of the applications in which pedagogical agents have traditionally been used, in that the knowledge desired for the student is never explicitly stated in the experience. Rather, the child discovers it on her or his own. Testing also differs, as the goal is a deeper conceptual understanding, not easily measured by right or wrong answers. Although discovery-based learning has been a major approach in reform-oriented pedagogy for over a century in classrooms, only recently has begun to be incorporated into interactive technology. The broadest objective of the current paper

** The first two authors made equal contributions to technical aspects of the research.

is to highlight an approach, along with challenges and responses, for building autonomous pedagogical agents for discovery-based interactive learning.

This work explores the application of pedagogical agents to the experiential goal of discovery-based learning. In particular, we add a pedagogical agent to an embodied math learning environment, MITp, designed to teach children proportion. MITp is described in detail in Sec. 2. The basic idea is that a child is encouraged to move two markers on a screen with her fingers. As she does this, the screen changes color. If the height of the two markers is in a particular ratio, say 1:2, the screen will go green, and as the ratio varies away from this it will go to yellow, and then red. The child is never told anything about ratios or proportion or how to make the screen green. Rather she is guided to discover different ways to create a green screen, and by doing so, begins to build an understanding of proportion. This system has been used extensively in learning research with a human tutor guiding students. In this work, we seek to understand what is required to make an animated pedagogical agent effective in this tutoring role.

MITp is an embodied learning experience where the child learns through performing physical movements. Embodied pedagogical agents are particularly useful in this setting because of the engagement they engender and, most importantly, because they can enact virtual actions and gestures they wish the learner to perform.

This type of learning application creates unique computational challenges. Chief among them, it is very difficult to measure the student's progress when it is not possible to ask questions with right or wrong answers that are easy for a computer to grade. Our design process employed a deep analysis of the process used by human tutors, including reviewing many hours of video recorded interactions. We identified the key stages in the tutorial process, the types of actions tutors took and when they took them. From this analysis, we identified the following activity types the agent must engage in: *instructing* the child what to do; *valorizing* success; *waiting* so the child can explore on her own; *providing remedial training* when the child is blocked and *advancing* the child through the tutorial process.

While the human tutor sits beside the learner, we placed our animated agent on the other side of the screen from the child, with access to the same touch surface as the learner (Figure 1). Dubbed *Maria*, our agent can execute a sequence of action blocks. Each block consists of any subset of spoken audio, facial animation, lip syncing and body animation, including arm and finger gestures. There are well over 100 actions that the agent can perform. We use Dynamic Decision Networks to decide when the agent should perform an action block and which action to perform.

Fig. 1: A child listens to the pedagogical agent explain concepts within the MITp learning environment.

We have conducted a pilot evaluation of the system. It showed that students were able to effectively explore the screen to find greens and enact a particular search strategy taught by the system. Students provided proto-mathematical descriptions of one solution strategy, although not the other. The results demonstrate good progress and also illuminate potential directions for future work.

2 MITp Learning Environment and Tutorial Protocol

Fig. 2: The MITp environment. The screen is green when the hands’ heights match a pre-programmed ratio.

The current study builds upon an earlier educational-research effort to design and evaluate embodied-interaction technology for mathematics instruction. Specifically, the pedagogical agent described herein was integrated into an existing activity design architecture called the Mathematical Imagery Trainer for Proportionality (MITp) [1, 2, 17], which we now present.

Proportional reasoning is important yet difficult for many students. It involves understanding multiplicative part-whole relations between rational quantities; a change in one quantity is always accompanied by a change in the other, and these changes are related by a constant multiplier [6, 26, 31].

Our MITp approach to support students in developing multiplicative understanding of proportions draws on embodiment theory, which views the mind as extending dynamically through the body into the natural-cultural ecology. Thus human reasoning emerges through, and is expressed as, situated sensorimotor interactions [3, 25]. Educational researchers informed by these theories have created technologies to foster content learning through embodied interaction (e.g., [5, 11]).

The MITp system (Figure 2) poses the physical challenge of moving two hands on a touch screen to make it green, a result which occurs when the ratio of hand heights matches the pre-programmed ratio of 1:2. Through this process, students can develop pre-symbolic mathematical understanding by engaging in this embodied activity and building particular movement schemes related to proportions. By introducing specific tools into the environment, here a grid and numbers, students are given progressively more mathematical tools with which to express those strategies.

The MITp system has been extensively tested for its educational effectiveness. Using qualitative analyses, the researchers demonstrated the variety of manipulation strategies students developed as their means of accomplishing the task objective of moving their hands while keeping the screen green [17]. Moreover, it was shown that students engaged in deep mathematical reflection as they were guided to compare across the strategies [2]. The studies have presented empirical

data of students shifting from naive manipulation to mathematical reasoning as they engage the frames of reference introduced into the problem space and the tutor’s critical role in facilitating this shift [1].

2.1 Interview Protocol

Maria is programmed to lead students through a series of activities on the MITp touchscreen, each supporting the development of particular movement strategies deemed relevant to proportional reasoning. Broadly, the two main phases are exploration, targeting the strategy “Higher-Bigger,” and “ a -per- b .” In “Higher-Bigger,” participants meet Maria on a screen with a red background, which is later overlaid with a grid. Participants are instructed to move the cursors up and down to make the screen green. At each green, Maria valorizes their work and asks them to make another green, either higher, lower, or elsewhere. Other than moving the cursors up and down, participants receive little guidance on particular movement strategies. With time, the grid is overlaid on the screen. The goal of this stage is for students to notice that when they make a green higher on the screen, the gap between their hands is bigger (“Higher-Bigger”). Next, in the “ a -per- b ” phase, participants are given instructions to start at the bottom of the screen, move their left hand up one grid unit, and then place their right hand to make green. Finally, the grid is supplemented with numerals. Participants are periodically asked to reflect on their rule for making green. Though Maria does not (yet) recognize speech, these reflections promote verbal description through which the developers can assess the participants’ proto-proportional understanding. Participants took about 20 minutes on average to complete the task.

While participants interact with Maria, the presiding human interviewers try to minimize human-to-human interaction, albeit occasionally they respond to participant queries, confusion, or frustration.

3 Related Work on Pedagogical Agents

Our work finds its roots in previous work on Pedagogical agents and Intelligent Tutoring Systems. Intelligent tutoring systems are computer softwares designed to simulate a human tutor. Pedagogical agents aid the process by adding a human-like character to the learning process. Research over the past few decades [21] has validated the positive impact of having an embodied presence in virtual learning environment. They have been a success primarily because they add emotional and non-verbal feedback to the learning environment [22]. More expressive pedagogical agents tend to improve the learning experience [15].

Intelligent tutoring systems (ITS) have been developed for a wide range of topics. Cognitive Tutors [24] have been adapted to teach students mathematical and other scientific concepts like genetics. The Andes Physics tutor [36] focuses on helping students in introductory Physics courses at college level. Writing Pal [33] and iStart [19] help students in developing writing and reading strategies

respectively. Decision theoretic tutoring systems have also been very successful and range from generic frameworks like DT Tutor [29] to domain specific systems such as Leibniz [13]. The feedback and learning mechanism behind all these activities revolves around the tutor provided instructions or some sort of rule specification, followed by student responses, given as either text or multiple choice selections, to posed challenge questions. Our discovery based learning methodology differs fundamentally as the system never describes how to achieve the desired goal, and the student response has to be gauged in real-time based solely on the touch screen coordinates. There are not questions that can be used to directly gauge progress.

Pedagogical agents can interact with the student in various roles, such as interactive demonstrators, virtual teammates and teachers. Steve [20] is an early example of a demonstration based pedagogical agent to train people to operate ship engines. INOTS [7] teaches communication and leadership skills to naval officers. The agent is questioned about a case by officers during training, and their performance is evaluated by rest of the class watching this interaction. AutoTutor [14] is a modern system used to teach concepts in Science and Mathematics. The student agent works with the human student to solve the problems in different ways. Adele [34] and Herman the Bug [27] are two classic pedagogical agents designed to teach medicine and botanical anatomy respectively. Decision Networks have also previously been used in the development of adaptive pedagogical agents such as [8] and [32] and for narrative planning [28]. These decision theoretic agents work on concrete feedback from the user in form of biological signals or responses to questions. In fact, all the above mentioned agents use the standard teach and test framework in order to gauge student’s performance, allowing them to focus on specific concepts in the learning that the student is struggling with. Our pedagogical agent operates in a quite different, discovery-based learning paradigm. Some previous pedagogical agents have also targeted more open-ended learning environments. For example, a system designed for children with ASD allows children to interact with the system by telling stories, control the agent by selecting pre-defined responses or author new responses to create new stories [35]. Related work has sought to use agents not as instructors, but as virtual peers [23, 12].

4 System

4.1 System Overview and Architecture

An overview of our MITp autonomous agent system architecture is shown in Figure 3. It consists of a control system and Unity3D front end. Students interact with our agent using a touch screen. The screen is virtually divided into left and right halves around the agent, designating two large tracks where the learner can move the markers. Maria is standing in the middle as shown in Figure 1 and can reach most of the screen. The Unity client sends the system state consisting of the two touch locations to the control system, which then instructs the agent to perform particular actions by specifying Action IDs.

Fig. 3: System Architecture

The control system employs dynamic decision networks [10] to model the behavior of our pedagogical agent. The decision networks are updated based on the evidence received from Unity and history maintained throughout the interaction. They may decide to do nothing or have the agent perform one of many pre-designed actions, depending on what appears most efficacious in the current context. Triggered actions are sent to Unity as action block IDs. Each action block consists of an audio file, a facial animation and a body animation. Our database contains 115 different action blocks.

Generating meaningfully labelled data for learning agent behavior is a challenge in a discovery-based setting. At this point, we do not make assumptions about student’s learning during interaction and instead rely on a robust tutorial process which can adapt and provide remedial instruction if the student is struggling. Modeling student’s learning on the fly requires mapping from patterns of finger movement to their mental state, which remains future work. Due to these issues, learning based techniques such as [18] are not a good fit for the current problem. DDNs allow us to leverage our strong understanding of the tutorial process by pre-encoding it into system parameters.

4.2 Decision Networks

Decision networks find their roots in *Bayesian Networks* [30], which are graphical models consisting of chance nodes representing a set of random variables. Random variables are events that could possibly occur in the given world. Chance nodes, drawn as ovals in the graph, can take any type of discrete value, such as a boolean or an integer. These values are generally finite and mutually exclusive. Arcs between these nodes capture conditional dependencies. Cycles are not allowed. Given the conditional probabilities, prior and evidence, inferences can be made about any random variable in the network.

Decision networks [16] extend Bayesian Networks by using Bayesian inference to determine a decision that maximizes expected utility. Decision networks add decision and utility nodes, represented by a rectangle and diamond respectively. Decision nodes represent all choices the system can make while the utility node captures our preferences under different factors impacting the decision. This concept can be further extended to *Dynamic Decision Networks* (DDNs)

[10]. DDNs consist of time varying attributes where there are conditional dependencies between nodes at different time steps. Decisions made in previous time steps can impact the probability distribution or state of the network in future steps. DDNs provide a useful way of modeling evolving beliefs about the world and changing user preferences.

4.3 Building Decision Networks for Experiential Learning

Discovery-based learning is challenging compared to other learning settings because the pedagogical agent must make decisions with very impoverished information as there is no continuous stream of concrete verbal or q&a based feedback that the agent can use to assess the student’s understanding. For example, in our case, we only have student’s touch coordinates on the screen as input. Our agent guides the learner through the discovery process using this input and a deep understanding of the tutorial process, encoded in the DDNs. Per the tutorial analysis, the agent must *instruct*, *valorize*, *wait*, *provide remedial training* and *advance* the child to the next stage. To decide if the agent should *wait*, the child must have a notion of the passage of time, which can be combined with the child’s activity pattern to determine if allowing time for exploration is appropriate. *Remedial* actions are triggered when the child is not executing desired movement patterns after repeated instructions and ample trial time. An example remedial strategy involves the agent displaying a marked location for the child’s one hand and then encouraging them to find green by only moving the other hand. Based on their performance, the agent may ask students to repeat activities to deepen their learning.

The learning experience progresses through multiple activities tied to particular interaction strategies: “Higher-Bigger” exploration without and with the grid, “a-per-b” without and with numbers, and an optional “speed” activity. All the activities are modeled using dynamic decision networks. Space limitations preclude discussion of each, but we will explain our approach using the “Higher-Bigger” task as an exemplar.

Exploring “Higher-Bigger” is the introductory activity students go through, as outlined in Sec. 2. The activity guides the student to find greens at several locations on the screen, with the goal of having the child realize that the separation between her hands is larger for greens higher on the screen than greens that are lower. Figure 4 and 5 show the decision network which governs the behavior of our pedagogical agent for this activity at two different points during an interaction.

The decision network is updated multiple times each second to ensure real time responses. The network encodes our agent’s belief about the state of tutorial process and student’s interaction at each time step. Factors impacting the agent’s decision for the network shown in Fig. 4 and 5 are:

Goal : Models agent’s temporally evolving expectations for the student. During this activity the student is first expected to find a couple of greens anywhere on

Fig. 5: At this point, the child is either actively exploring, but failing to satisfy the goal, or has stopped interacting with the system. The agent queries the network for the optimal decision to make after setting evidence for Timeout, screen Color and current exploration Location on the screen (bold words). The decision network suggests Valorization and Educating the child about the task as actions of equal and optimal utility (shown in blue). We choose either action randomly. Remedial activities override the decision network and are triggered if there have been multiple continuous timeouts and student hasn't been able to achieve a goal for an extended time.

5.1 Experiment Description

The agent followed the protocol outlined in Sec. 2.1.

Participants included 10 children (5 male, 5 female) aged 9 - 12 years old.

Data Gathering: As participants worked with Maria, three sets of data were collected simultaneously. One researcher used an objective-observation instrument, detailed below, to note the occurrence and frequency of key participant movements and expressions. A second researcher took qualitative field notes of the participant's interaction with Maria. Finally, the interview session was video and audio recorded. The qualitative field notes and video records were used to verify key moments indicated in the observation instruments.

Note that these data sources are consistent with qualitative methods. As this system represents a new genre of agent-facilitated, discovery-based learning, we first seek to understand the types of interactions and insights that emerge as children work with the system. Such qualitative work is not consistent with statistical analyses or pre/post comparison, which pre-suppose which ideas will emerge as salient for participants and considers only those insights that can be quantitatively evaluated. With a qualitative understanding in hand, later work can focus the system on insights deemed most mathematically productive, at which stage a quantitative evaluation would be appropriate. This progression is well established within grounded theory [9] and design based research [4].

Observation Instrument: A one-sheet observation instrument was used during each interview to code, in real time, when each participant exhibited particular benchmark movements and expressions. These movements and expressions were determined in consultation with researchers familiar with analogous, human-run interviews. For example, a participant expression of "the higher I go on the screen, the bigger the distance must be between my hands to make green"

is a proto-proportional expression that reflects the changing additive difference between numerator and denominator. Additionally, a movement scheme of raising the right hand 2 units for every 1 unit increase in the left hand reflects the constant multiplicative relationship between equivalent ratio (e.g., 4:8 and 5:10). These two strategies are termed “Higher-Bigger” and “*a-per-b*,” respectively.

5.2 Results

Subject	Greens by Screen Region			“ <i>a-per-b</i> ”	“Higher-Bigger”	“ <i>a-per-b</i> ”	Other Insights
	Low	Middle	High				
1	14	4	5	1	0	1	1
2	12	7	7	8	0	1	3
3	6	9	3	4	0	2	2
4	10	16	7	1	0	0	5
5	11	4	6	7	0	2	1
6	11	15	7	9	0	0	4
7	9	11	5	4	0	2	2
8	7	5	5	5	0	1	3
9	8	15	8	5	0	2	2
10	2	7	10	4	0	0	3

Table 1: Frequency of greens, by region, and “*a-per-b*” performance. Table 2: Participant expressions by category.

The observation instrument was designed to measure study participants’ physical movements and/or verbal and gestural utterances that would imply they are engaging in “the higher, the bigger” or “*a-per-b*” strategies as their means of solving the bimanual manipulation problem. Performance of a particular movement pattern suggests that the participant is enacting a perceptuo-motor strategy that could result in conceptual learning. Verbal description of those movements in proto-mathematical terms indicates further progress along that learning pathway per our instructional design.

As indicated in Table 1, all participants produced green low down, in the middle, and high up on the screen. Interestingly, no participant described the changing distance between their hands at these various regions (Table 2). Additionally, all participants performed the “*a-per-b*” strategy (Table 1 indicates the number of times this was observed), and 7 participants verbally described it, such as “for every 1 my left goes, my right goes 2.” Notably, all participants developed other insights into the system’s functioning, which fell into 1 of 4 categories: observational (“Generally, my right hand has to be on top”), feedback-based (“If you see the red screen flash [to green], move back to where you were”), memorization (“4 and 7 or 8 will make green. 6 and 11 or 12 will make green. 2 and 3 makes green.”), and procedural (“Keep one hand in one spot and move your other hand around”).

Unanticipated was that researchers were obliged to interact with participants on average twice per interview. In 7 of the 10 interviews, this interaction involved researchers restating the “*a-per-b*” instructions with similar phrasing to Maria’s.

5.3 Discussion

We are encouraged by the widely adopted movement strategies, both for making green in all regions of the screen and in adopting the “*a-per-b*” strategy. Instructed by Maria, and largely independent of human intervention, all participants developed movement schemes supporting proto-proportional reasoning.

This work surfaced a gap, however, between participants’ performed movements and their descriptions thereof. Participants adeptly made green in all regions of the screen and performed the “*a-per-b*” strategy, yet they did not develop proto-mathematical descriptions of “Higher-Bigger,” only of “*a-per-b*.” Comparing the conditions of “*a-per-b*” work with the conditions for “Higher-Bigger” work suggests sources of this disparity. In the “*a-per-b*” phase, participants received verbal instructions for the target movement strategy as well as a visual grid and numerals. The verbal instructions highlighted discrete, alternative hand movements (laying the “—per—” foundation) while the grid and numerals drew attention to unit quantities (supporting specifically “1-per-2”). In contrast, instructions for the “Higher-Bigger” phase did not explicitly mention the distance between participants’ hands. Additionally, students were not instructed to follow a particular green progression, for example making green low down, in the middle, and up high, that would facilitate noticing a growing distance. Future efforts should focus on understanding how to embed in the pedagogical agent’s models and actions the nuances of human-tutor actions that have led students to attend to the interval between their hands.

Patterns in researcher intervention suggest another area for design iteration. Researchers consistently interacted with participants by repeating the “*a-per-b*” instructions after participants worked unsuccessfully for 3 or more minutes. During this time, participants developed a host of less effective movement patterns - alternating left and right but moving each 1 unit or raising the left hand to the top of the screen then raising the right hand. Though Maria repeated fragments of her original instruction, the timing and particular fragment selected often did not correct the participant’s movement. And while researchers tried to mimic Maria’s exact instructional language, their choice of *when* to give those instructions and *which words to emphasize* gave more information than Maria is programmed to do. Further work is required to analyze these researcher interventions and convert their decisions into procedures for the autonomous agent.

Overall, we find this work to tentatively support the added value of a virtual agent in discovery-based learning environments. The agent provides feedback on student work and suggests corrective or novel movement strategies that would likely not arise in agent-free work. In particular, the agent draws the student’s attention to multiple parametric regions of the problem space, such as particular spatial locations on the monitor, that had not occurred to the student in their free exploration, and the agent suggests new spatial-temporal interaction schemes, such as introducing a sequential bimanual manipulation regime where the student was trying only simultaneous actions. The agent also provides encouragement, validating the students’ efforts and encouraging them to explore in new ways. As none of the participants noticed the “Higher-Bigger” relationship,

this first prototype was somewhat less successful than human tutors. However, this is not surprising. Human tutors perceive a wider range of student behaviors (posture, oral expressions, facial expressions) contemporaneous with their on-screen actions, giving more information upon which to determine the pacing and content of guidance. Additionally, human tutors enjoy the full range of their gestural and verbal vocabularies in responding to and guiding participants. Consequently, we did not expect that the virtual agent would perform to the same benchmark as human tutors. Nevertheless, we see the results of this work as a success, then, in that all participants performed, and almost all expressed, the proportional “*a-per-b*” relationship under the virtual agent’s guidance.

6 Conclusion

The MITp system presents a very challenging application for pedagogical agents as they must determine appropriate actions based on very little feedback from the learner, in this case, only the location of two markers on the screen. Such constraints are typical of discovery-based learning, where the asking of concrete questions is limited, and the learner is given freedom to explore. The system performed quite well on the task, leading to appropriate movement patterns in all cases and desired verbal expressions for one of the two movement strategies taught. This suggests that the potential for pedagogical agents in discovery-based learning is high and that DDNs represent an effective control strategy.

The system was effective due to a very thorough understanding of the tutoring protocol that was then encoded in the DDNs. In our case, this was based on an analysis of human-led tutoring of the same task. Two significant shortcomings were noted in the study, students failed to verbally explain the “Higher-Bigger” pattern and some amount of human intervention was required, generally when students failed to progress later in the last task. Both of these suggest the need to further refine the protocol encoded in the DDNs. Future work should consider using verbal input from the learners.

Acknowledgements

Financial support for this work was provided by the National Science Foundation through grants IIS 1320029 and IIS 1321042. We gratefully acknowledge Huaguang (Chad) Song, Arman Kapbasov, Rhea Feng, Austin Berbereia, Quan Pham, Gregory Moore and Jessica Sheu for helping develop character content and Alyse Schneider, Virginia Flood and Seth Corrigan for valuable analysis and for providing voice and motion data. Finally, we thank the many children and their parents who participated in our IRB-approved studies.

References

1. Abrahamson, D., Gutiérrez, J., Charoenying, T., Negrete, A., Bumbacher, E.: Fostering hooks and shifts: Tutorial tactics for guided mathematical discovery. *Technology, Knowledge and Learning* 17(1-2), 61–86 (2012)

2. Abrahamson, D., Lee, R.G., Negrete, A.G., Gutiérrez, J.F.: Coordinating visualizations of polysemous action: Values added for grounding proportion. *ZDM* 46(1), 79–93 (2014)
3. Anderson, M.L., Richardson, M.J., Chemero, A.: Eroding the boundaries of cognition: Implications of embodiment. *Topics in Cognitive Science* 4(4), 717–730 (2012)
4. Anderson, T., Shattuck, J.: Design-based research: A decade of progress in education research? *Educational researcher* 41(1), 16–25 (2012)
5. Antle, A.N.: Research opportunities: Embodied child–computer interaction. *International Journal of Child-Computer Interaction* 1(1), 30–36 (2013)
6. Boyer, T.W., Levine, S.C.: Prompting children to reason proportionally: Processing discrete units as continuous amounts. *Developmental psychology* 51(5), 615 (2015)
7. Campbell, J., Core, M., Artstein, R., Armstrong, L., Hartholt, A., Wilson, C., Georgila, K., Morbini, F., Haynes, E., Gomboc, D., et al.: Developing inots to support interpersonal skills practice. In: *Aerospace Conference, 2011 IEEE*. pp. 1–14. IEEE (2011)
8. Conati, C.: Probabilistic assessment of user’s emotions in educational games. *Applied Artificial Intelligence* 16(7-8), 555–575 (2002)
9. Creswell, J.W.: *Qualitative inquiry and research design: Choosing among five approaches*. Sage Publications (2012)
10. Dean, T., Kanazawa, K.: A model for reasoning about persistence and causation. *Computational intelligence* 5(2), 142–150 (1989)
11. Duijzer, C.A., Shayan, S., Bakker, A., Van der Schaaf, M.F., Abrahamson, D.: Touchscreen tablets: Coordinating action and perception for mathematical cognition. *Frontiers in Psychology* 8 (2017)
12. Finkelstein, S., Yarzebinski, E., Vaughn, C., Ogan, A., Cassell, J.: The effects of culturally congruent educational technologies on student achievement. In: *International Conference on Artificial Intelligence in Education*. pp. 493–502. Springer (2013)
13. Gluz, J.C., Cabral, T., Baggio, P., Livi, P.: Helping students of introductory calculus classes: the leibniz pedagogical agent
14. Graesser, A.C., Lu, S., Jackson, G.T., Mitchell, H.H., Ventura, M., Olney, A., Louwerse, M.M.: Autotutor: A tutor with dialogue in natural language. *Behavior Research Methods* 36(2), 180–192 (2004)
15. Hasegawa, D., Shirakawa, S., Shioiri, N., Hanawa, T., Sakuta, H., Ohara, K.: The Effect of Metaphoric Gestures on Schematic Understanding of Instruction Performed by a Pedagogical Conversational Agent, pp. 361–371. Springer International Publishing, Cham (2015)
16. Howard, R.A.: *Readings on the principles and applications of decision analysis*, vol. 1. Strategic Decisions Group (1983)
17. Howison, M., Trninic, D., Reinholz, D., Abrahamson, D.: The mathematical imagery trainer: From embodied interaction to conceptual learning. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. pp. 1989–1998. ACM (2011)
18. Jaakkola, T., Singh, S.P., Jordan, M.I.: Reinforcement learning algorithm for partially observable markov decision problems. In: *Advances in neural information processing systems*. pp. 345–352 (1995)
19. Jackson, G.T., Boonthum, C., McNAMARA, D.S.: istart-me: Situating extended learning within a game-based environment. In: *Proceedings of the workshop on intelligent educational games at the 14th annual conference on artificial intelligence in education*. pp. 59–68. AIED Brighton. UK (2009)

20. Johnson, W.L., Rickel, J.: Steve: An animated pedagogical agent for procedural training in virtual environments. *SIGART Bull.* 8(1-4), 16–21 (Dec 1997)
21. Johnson, W.L., Lester, J.C.: Face-to-face interaction with pedagogical agents, twenty years later. *International Journal of Artificial Intelligence in Education* 26(1), 25–36 (2016)
22. Johnson, W.L., Rickel, J.W., Lester, J.C.: Animated pedagogical agents: Face-to-face interaction in interactive learning environments. *International Journal of Artificial intelligence in education* 11(1), 47–78 (2000)
23. Kim, Y., Baylor, A.L.: A social-cognitive framework for pedagogical agents as learning companions. *Educational Technology Research and Development* 54(6), 569–596 (2006)
24. Koedinger, K.R., Corbett, A., et al.: Cognitive tutors: Technology bringing learning sciences to the classroom (2006)
25. Lakoff, G., Núñez, R.E.: Where mathematics comes from: How the embodied mind brings mathematics into being. Basic books (2000)
26. Lamon, S.J.: Rational numbers and proportional reasoning: Toward a theoretical framework for research. *Second handbook of research on mathematics teaching and learning* 1, 629–667 (2007)
27. Lester, J.C., Stone, B.A., Stelling, G.D.: Lifelike pedagogical agents for mixed-initiative problem solving in constructivist learning environments. *User modeling and user-adapted interaction* 9(1), 1–44 (1999)
28. Mott, B.W., Lester, J.C.: U-director: a decision-theoretic narrative planning architecture for storytelling environments. In: *Proceedings of the fifth international joint conference on Autonomous agents and multiagent systems*. pp. 977–984. ACM (2006)
29. Murray, R.C., VanLehn, K.: Dt tutor: A decision-theoreticdynamic approach for optimal selection of tutorial actions. In: *International Conference on Intelligent Tutoring Systems*. pp. 153–162. Springer (2000)
30. Pearl, J.: Probabilistic reasoning in intelligent systems (1988)
31. Piaget, J., Inhelder, B.: The psychology of the child, vol. 5001. Basic books (1969)
32. Prendinger, H., Ishizuka, M.: The empathic companion: A character-based interface that addresses users’ affective states. *Applied Artificial Intelligence* 19(3-4), 267–285 (2005)
33. Roscoe, R.D., McNamara, D.S.: Writing pal: Feasibility of an intelligent writing strategy tutor in the high school classroom. *Journal of Educational Psychology* 105(4), 1010 (2013)
34. Shaw, E., Ganeshan, R., Johnson, W.L., Millar, D.: Building a case for agent-assisted learning as a catalyst for curriculum reform in medical education. In: *Proceedings of the International Conference on Artificial Intelligence in Education*. pp. 509–516 (1999)
35. Tartaro, A., Cassell, J.: Authorable virtual peers for autism spectrum disorders. In: *Proceedings of the Combined workshop on Language-Enabled Educational Technology and Development and Evaluation for Robust Spoken Dialogue Systems at the 17th European Conference on Artificial Intelligence* (2006)
36. Vanlehn, K., Lynch, C., Schulze, K., Shapiro, J., Shelby, R., Taylor, Treacy, Weinstein, A., Wintersgill, M.: The andes physics tutoring system: Lessons learned. *International Journal of Artificial Intelligence in Education* 15(3), 147–204 (2005)