The Curiosity Notebook: A Platform for Supporting Learning with Teachable Robots

Edith Law
edith.law@uwaterloo.ca
University of Waterloo
Waterloo, Ontario, Canada

Ken Jen Lee
kj24lee@edu.uwaterloo.ca
University of Waterloo
Waterloo, Ontario, Canada

Parastoo Baghaei Ravari
parastoo.baghaei.ravari@uwaterloo.ca
University of Waterloo
Waterloo, Ontario, Canada

1 INTRODUCTION

Learning by teaching is a popular and well studied pedagogical technique that has been shown to produce the *protégé effect* [7]—students who are asked to teach others are better able to synthesize and structure materials, become aware of their own learning process, and expend more effort to learn. Despite extensive research, however, our understanding of the exact conditions that make learning by teaching effective is limited, mostly because there can be high variability in the tutor and tutee behaviour. As Roscoe [18] recommends, one way to more systematically test hypotheses about learning by teaching is to develop *teachable agents* with manipulable characteristics and behaviour. In this work, we introduce a learning-by-teaching web application called the Curiosity Notebook that supports learning by teaching through dialogue with virtual agents and physical robots, describe the design evolution of the system based on insights drawn from a 4-week exploratory study involving elementary school students teaching a conversational robot how to classify objects, and conclude with a discussion about limitations and future work.

2 BACKGROUND

In computer-mediated learning applications, agents have mostly served as peers [13, 19] or tutors [10, 16], with only a handful of systems positioning the agent as a less intelligent or knowledgeable peer that students teach [3, 5]. SimStudent [15], for example, is a simulated learner used to study student-tutor learning in mathematics problem solving. Most extensively studied is the Betty's Brain [3, 4] platform, a learning environment in which students read articles, then teach and quiz a virtual agent (i.e., an avatar called Betty) about causal relationships (e.g., burning fossil fuels increases CO_2) in science by manipulating concept maps. Betty's Brain has been used to study self-regulated learning [17], collaborative teaching [8], the role of feedback [20, 21], scaffolded learning of multiple representations [1], to name a few.

Other teachable agent research involves physical robots. In Tanaka and Matsuzoe [23], young children (i.e., 3-6 years old) taught a humanoid robot (i.e., NAO) English words, while simultaneously interacting with a human teacher. In a later study, they [22] also investigated how preschool children learn English by teaching Pepper, an adult-size humanoid robot while receiving guidance from a human teacher demonstrating vocabulary-related gestures on a small screen attached to Pepper's chest. Yadollahi et al. [24] developed a collaborative story reading environment, where children (aged 6-7) can correct the robot's mistakes as it reads aloud. Several works have designed teachable agents to help primary school aged children—working individually, in pairs and groups—to improve their handwriting [6, 11]. Here, students use a card to show the robot a 3-letter word, which the robot then writes on a tablet and asks for feedback. Students can follow up by rewriting on the tablet any letters they felt were incorrect. Other studies have focused on older student populations. One study had students (mean age=20) teach a NAO robot to solve math problems, and investigated the effects of dyadic stance formations (e.g., face-to-face stance vs. side-by-side stance) on student attitudes. In general, learning by teaching *human-like* physical robots invites a whole new set of research questions, e.g., related to the physical interaction between the students and the robot, that are distinct from those related to virtual agents [12].

Our work diverges from existing learning-by-teaching platforms in several ways. Similar to Betty's Brain, our learning by teaching platform engages students in open domain knowledge learning; distinct to our platform, however, is the specific focus on *classification* as the learning/teaching task. The Curiosity Notebook includes a multimodal conversational agent that can take the form of a text-based chatbot, voice-only agent, or physical robot (i.e. NAO). The agent is designed to facilitate both individual- and group-based teaching by controlling turn-taking and encouraging discussions amongst the student teachers. These platform features together enable a wide range of learning by teaching scenarios—individual versus group-based teaching, in-person versus online teaching, and the ability to teach agents with different embodiments.

3 CURIOSITY NOTEBOOK

3.1 Teaching Interface

The Curiosity Notebook provides a web interface that students use to read articles and teach a conversational agent how to classify objects, e.g., classifying paintings as impressionist, cubist, or realist art; classifying animals into mammals, insects, and reptiles; or classifying rocks as metamorphic, igneous or sedimentary. We chose to focus on classification tasks because they are well structured—learning how to classify objects involves mainly identifying features and mapping them to categories—which means that the teaching conversations can be highly structured as well. Classification tasks are also amenable to machine learning, allowing computational models of learning to be eventually implemented in the agent. Finally, classification tasks can be made arbitrarily simple or complex to adapt to the age and ability of the students; for example, rock classification is a topic in both Grade 4 and college-level curriculum.

Upon logging onto the Curiosity Notebook, students are presented with a set of topics (e.g., rock, animal or painting classification) to choose from, which subsequently brings them to the main teaching interface (Figure 1). The teaching interface consists of a

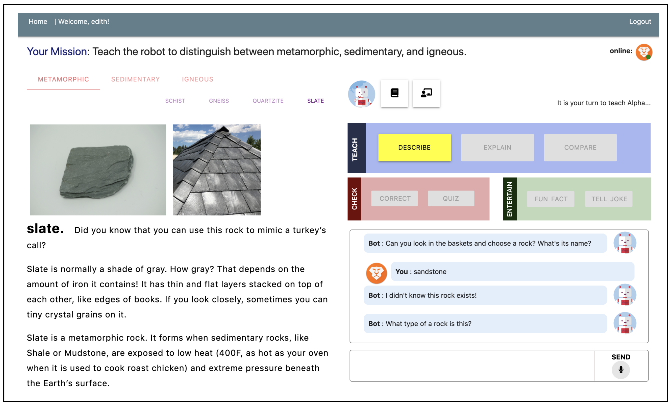

Figure 1: Teaching Interface

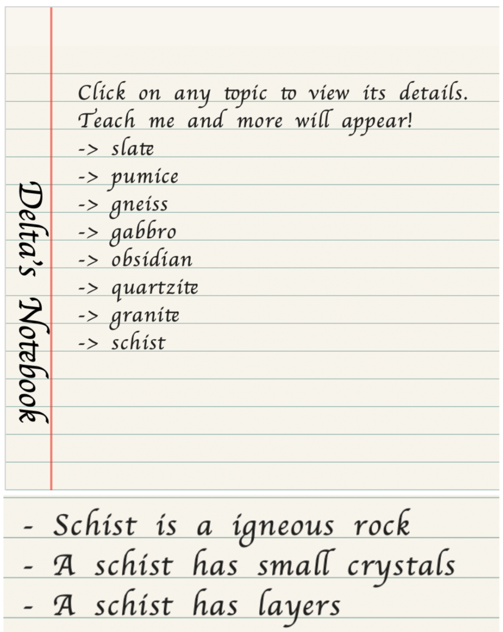

Figure 2: Robot's Notebook, including an index page outlining the list of learned entities (top) and specific knowledge learned for a particular entity (bottom).

reading panel (left), containing articles about objects (e.g., Limestone) belonging to different categories (e.g., Sedimentary Rocks). Interactive functions allow students to highlight sentences to teach the agent. The teaching panel (right) contains buttons that students can click on to teach (blue), entertain (green) and check the progress (red) of the agent. When a button is clicked, the system locks the agent and the student into a mini-conversation that involves 4-6 rounds of question and answering. A chat window (below the buttons) displays a transcript of the conversation and can be made visible or hidden depending on whether the students are interacting with a text-based chatbot, voice-only agent or physical robot.

To teach the agent, students can click on the *describe* button to teach the agent about an object's features (e.g., "Obsidian does not

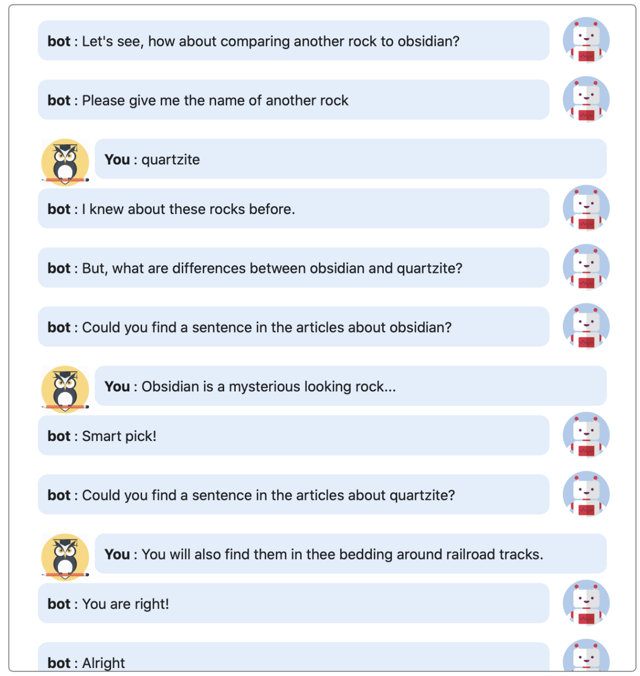

Figure 3: The "compare" conversation, where the student teacher is describing the similarities and differences between different entities to the agent.

crystals") and the *explain* button to explain why an entity has that feature (e.g., "Obsidian forms when lava cools; however, because it forms above the surface, the lava hardens too quickly that crystals did not have time to form"). The "compare" conversation (Figure 3) asks the students to find rocks that belong to either the same category or different categories, and explain how they are similar or different. Upon teaching a fact, students can toggle to the robot's notebook to inspect the list of facts that the agent has learned for each object (Figure 2).

Students can click on the *correct* button to update a fact that was previously taught to the agent (Figure 5). The agent will first ask what object the student wants to focus on, then present the student with a list of its notebook entries associated with that object to choose from, and finally, use questions to elicit a specific kind of correction. To probe the current performance of the agent, students can click on the *quiz* button, and select an object to test the agent on. Based on its current knowledge model, the agent will attempt to classify the object (Figure 4). Lastly, students entertain the agent by clicking on the "fun fact" and "tell joke" buttons.

Together, these interfaces allow students to partake in different teaching activities, such as describing and explaining an object's features, comparing and contrasting objects, monitoring the agent's learning progress (using the robot's notebook), building rapport with the agent by telling fun facts and jokes. Having a choice of when and how to teach allows us to observe the different teaching strategies students might take, and how they allocate time between teaching, checking and entertaining.

Students can teach the agent individually or in groups of arbitrary size, and their group placement can be configured by teachers

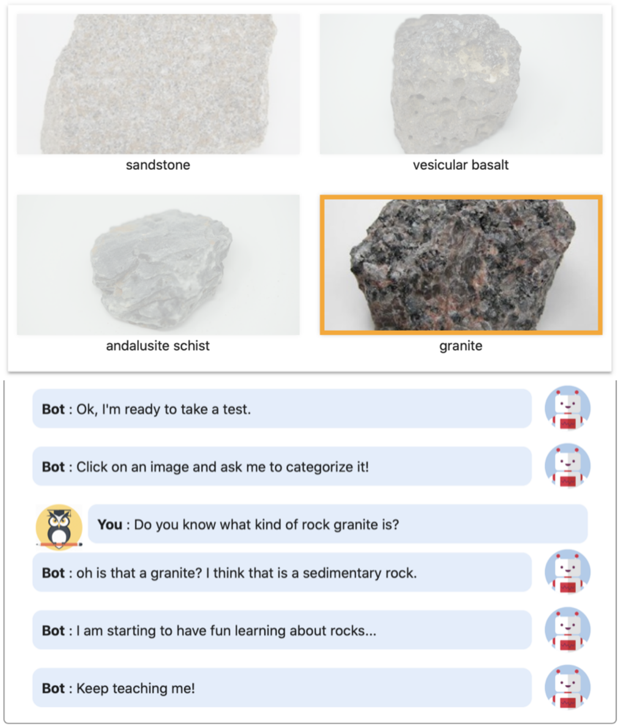

Figure 4: The student teacher selects an entity (top) to quiz the agent on through conversation (bottom).

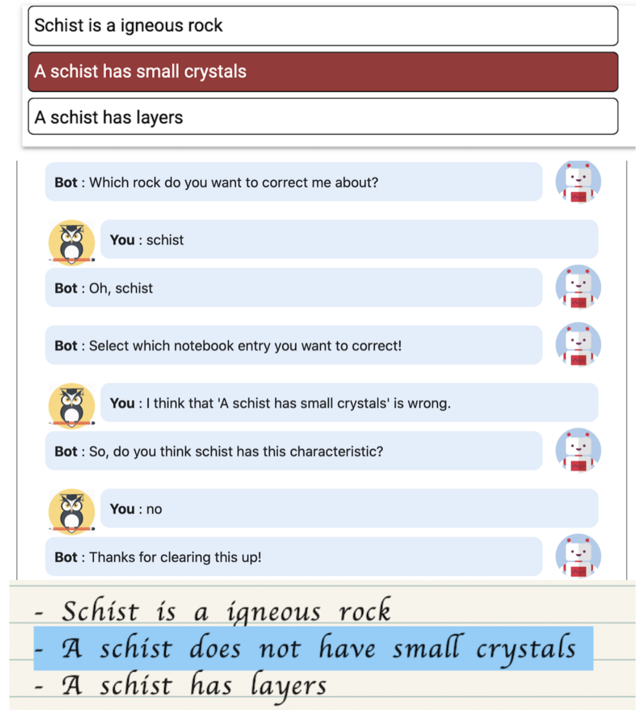

Figure 5: The student teacher selects a previously taught fact (top) and engages in a conversation with the agent (middle) to correct its knowledge. The corrected entry is highlighted in the Robot's Notebook (bottom) after the conversation.

or researchers through an administrative interface. If a student is placed in a group and their group members are present, their view of the system is synchronized—that is, if one student clicks a button to activate a particular mini-conversation, all students will be automatically locked into the same dialog.

3.2 Dialog System

Each mini-conversation, as invoked by one of the 7 buttons, is controlled by a state machine, which specifies the flow of the conversation. Each state specifies the set of possible `actions`, corresponding to what the agent would say, and an `action selector`, which specifies how the action is chosen (e.g., randomly, according to the current knowledge state of the agent, etc). Each action, in turn, consists of (1) a set of possible message `templates`, i.e., sentences with placeholders expressing the same meaning, and a `template selector` for specifying how to select the template (e.g., randomly, based on the classification task, etc), and (2) a set of possible `next states` the dialog can transition to after the action is taken, and a `state selector`, which specifies how that next state is selected (e.g., randomly, based on a probability distribution, etc). In addition to the main dialog for each mini-conversation, there are side dialogues that the agent can transition to; these side dialogues enable the agent to diverge from the main teaching conversation and return to it afterwards. As an example, when the agent learns a new concept, it can (with some probability) enter a side dialog to tell a joke, to provide comic relief during a long teaching session. These dialog state machines are entirely specified in a JSON format, enabling the rapid prototyping of conversations without having to modify the implementation. Furthermore, the action, template and state selectors together provide a flexible way to specify a dialog flow that generalizes across different classification tasks (i.e., the same state machine is used for rock, animal and painting classification).

The dialog system is tightly connected to the agent's knowledge model. For example, the dialog will proceed differently depending on whether the agent knows about a particular object. If the agent has never been taught Limestone before, it will say "Thanks! I never knew this rock exists!"; conversely, if the agent knew about Limestone but not its characteristics, it will say "I knew about Limestone, but don't know what it looks like." and proceed to ask about its features.

3.3 Administrative Interfaces

Configuring Agent Embodiment. The Curiosity Notebook supports a clean separation between agent logic and embodiment, thereby allowing the teachable agent to take on different types of embodiment. This is accomplished by keeping logic of the teachable agent—e.g., how it learns, how it feels and what it says—inside the Curiosity Notebook web application, and having an external program (e.g., a python script) ping the database for chat messages that the physical robot should say out loud. Each chat message is associated with an emotion tag (e.g., curious), which can be used to control the movements/gestures of the robot (e.g., rubbing its head or chin) to convey that emotion. Similarly, the external program can push sensing events to the Curiosity Notebook. The NAO robot, for example, has tactile sensors on the head, hands and feet, which can serve as alternative ways for students to provide feedback to the robot (e.g., patting its head when it answers a quiz question correctly). The external program transmits the sensing event to the Curiosity Notebook using https protocol, and the teachable agent logic inside the web application handles this event and decides how the agent should respond. A similar mechanism can be used to handle camera and microphone events.

Configuring Learning Activities. As we envision the Curiosity Notebook to be eventually used by teachers to organize learning-by-teaching activities for their class, the platform provides a set of web-based administrative tools for adding/removing users, updating user information, assigning users to groups, as well as configuring classification tasks and materials (e.g., articles, images).

Configuring Experiments. The Curiosity Notebook serves both as an educational tool for students, and a research infrastructure for studying learning by teaching phenomena. As such, the platform provides researchers with the ability to add/remove experiments, add/remove conditions to/from experiments, and assign users to a specific experiment and condition. The platform also provides functionalities for researchers to configure the verbal behaviour of the agent, and associate different verbal behaviour with different experimental conditions.

3.4 Design Evolution: Past, Present and Future

The development of the Curiosity Notebook has undergone multiple design iterations. In the initial design (Figure 6), the agent begins each teaching conversation by asking the students to pick an artifact (e.g., an animal figurine, a rock or mineral, and a painting postcard from NYC Metropolitan Museum). The teaching conversation proceeds with the agent highlighting a knowledge bubble (which represents a feature used in classification, e.g., "has layers") and asking a series of 4 or 5 questions about the corresponding feature. As the conversation concludes, the knowledge bubble is filled and students are rewarded with confetti on the screen letting them know that the agent has "learned" that feature.

Different from the initial design, the current system allows students to choose *how* to teach, as opposed to the agent directing the entire sequence of teaching interactions by posing questions for the student teachers to answer. The current system also provides a more transparent representation of the agent's learned knowledge—as explicitly written facts about each object in the robot's notebook—which allows the students to both inspect and debug the agent's knowledge. This is in contrast to knowledge bubbles, which are abstract and had the unintended effects of extrinsically motivating students to "fill all the bubbles as quickly as possible".

These design changes were based on insights drawn from a 4-week exploratory study [14] that we conducted with 12 fourth and fifth grade students at a local school using the initial version of the curiosity notebook. Students (7M/5F) participated in the study over 4 weeks. The study was conducted in an after-school club, which ran once a week for 1.5 hours each. Four NAO robots were used in each session. Students formed groups of 3, and taught the

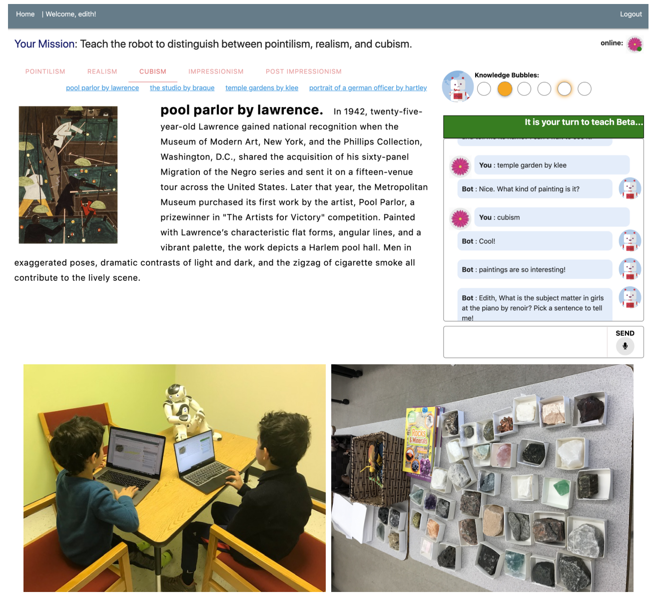

Figure 6: Students teaching a NAO robot (bottom left) using physical artifacts (bottom right) using the initial design of the Curiosity Notebook (top).

robot about a different topic (i.e., animals, rocks, paintings) each week, then all topics during the last week. Each student was given a chromebook, and sat together with their group members facing the robot, which was positioned in a sitting posture in front of the students on the table (Figure 6).

Overall, results [14] show that the Curiosity Notebook enabled multiple student groups of different sizes to simultaneously teach different robots in the same classroom, and provided insights into some surprising factors that can affect the group-based learning by teaching experience. For example, when asked which classification task students liked the most, the responses can be clustered into: (1) students who liked teaching a topic because it was easier, because they knew more about it, and because they perceived the robot to be learning more/better about that topic, and (2) students who liked teaching a topic because they knew less about it. This observation implies that personality traits (e.g., the desire for challenge, growth vs fixed mindset) can critically affect students' preferences of topics to teach and how much they enjoy the teaching experience. As another example, the amount of attention that the robot gives to each student teacher seems to affect students' perception of their teaching ability; one student said "Student X teaches way better because the robot chooses X more.". This suggests a more personalized approach to managing group-based teaching that takes into account each student's unique need for attention from the agent.

4 LIMITATIONS AND FUTURE WORK

Curiosity Notebook is designed with the goal of enabling researchers to more systematically test learning-by-teaching hypotheses; however, it remains unclear whether teaching a virtual agent or robot is exactly equivalent to teaching a human tutee. Caution must be exercised when generalizing research results to the human-to-human learning by teaching context. Second, teachable agents are, at its core, a persuasive technology, "designed to change people's attitudes or behavior" [9]. Berdichevsky and Neuenschwander suggested eight ethical principles for designing persuasive technology [2]; the last principle, called the "Golden Rule of Persuasion", states that system creators should never persuade users of something the creators themselves would not like to be persuaded of. Our Curiosity Notebook agent was designed to encourage certain learning behaviour and attitude in students. Though a benevolent objective, students were not explicitly told that the agent, which is acting as a less knowledgeable peer, is in fact pedagogical in nature. One can potentially address this ethical dilemma by divulging to the student teachers the intent of the learning-by-teaching exercise—that students are expected to learn through the process of teaching the agent.

In collaboration with multiple elementary schools, we are currently conducting studies using the Curiosity Notebook to understand how *tutee characteristics* (e.g., sense of humour, ability to adapt the conversation to changing group dynamics) affect the way students teach and learn through teaching. Future work includes developing better end-user functionalities to facilitate researchers and teachers alike in rapid deployment of learning-by-teaching activities, and building more intelligence into the teachable agent, enabling it to automatically retrieve materials from the Web, create novel classification tasks, train ML-based knowledge models, and strategically learn from student teachers.

REFERENCES

[1] Satabdi Basu, Gautam Biswas, and John S. Kinnebrew. 2016. Using multiple representations to simultaneously learn computational thinking and middle school science. In *30th AAAI Conference on Artificial Intelligence, AAAI 2016*. 3705–3711.

[2] Daniel Berdichevsky and Erik Neuenschwander. 1999. Toward an Ethics of Persuasive Technology. *Commun. ACM* 42, 5 (May 1999), 51–58. https://doi.org/10.1145/301353.301410

[3] Gautam Biswas, Krittaya Leelawong, Daniel Schwartz, Nancy Vye, and The Teachable Agents Group at Vanderbilt. 2005. Learning by teaching: A new agent paradigm for educational software. *Applied Artificial Intelligence* 19, 3-4 (2005), 363–392.

[4] Gautam Biswas, James R. Segedy, and Kritya Bunchongchit. 2016. From Design to Implementation to Practice a Learning by Teaching System: Betty's Brain. *International Journal of Artificial Intelligence in Education* 26, 1 (2016), 350–364.

[5] Sean Brophy, Gautam Biswas, Thomas Katzlberger, John Bransford, and Daniel Schwartz. 1999. Teachable agents: Combining insights from learning theory and computer science. In *Artificial intelligence in education*, Vol. 50. Citeseer, 21–28.

[6] Shruti Chandra, Raul Paradeda, Hang Yin, Pierre Dillenbourg, Rui Prada, and Ana Paiva. 2017. Affect of Robot's Competencies on Children's Perception. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, 1490–1492.

[7] Catherine C. Chase, Doris B. Chin, Marily A. Oppezzo, and Daniel L. Schwartz. 2009. Teachable Agents and the Protégé Effect: Increasing the Effort Towards Learning. *Journal of Science Education and Technology* 18, 4 (01 Aug 2009), 334–352.

[8] Mona Emara, Ramkumar Rajendran, Gautam Biswas, Mahmod Okasha, and Adel Alsaeid Elbanna. 2018. Do Students' Learning Behaviors Differ when They Collaborate in Open-Ended Learning Environments? *Proc. ACM Hum.-Comput. Interact.* 2, CSCW, Article 49 (Nov. 2018), 19 pages.

[9] B.J. Fogg. 2003. *Persuasive Technology: Using Computers to Change What We Think and Do.* Morgan Kaufmann Publishers Inc., San Francisco, CA, USA.

[10] Arthur C Graesser, Katja Wiemer-Hastings, Peter Wiemer-Hastings, Roger Kreuz, Tutoring Research Group, et al. 1999. AutoTutor: A simulation of a human tutor. *Cognitive Systems Research* 1, 1 (1999), 35–51.

[11] Deanna Hood, Séverin Lemaignan, and Pierre Dillenbourg. 2015. When Children Teach a Robot to Write: An Autonomous Teachable Humanoid Which Uses Simulated Handwriting. In *ACM/IEEE International Conference on Human-Robot Interaction (HRI)*.

[12] Frank Jamet, Olivier Masson, Baptiste Jacquet, Jean-Louis Stilgenbauer, and Jean Baratgin. 2018. Learning by Teaching with Humanoid Robot: A New Powerful Experimental Tool to Improve Children's Learning Ability. *Journal of Robotics* (2018), 1–11.

[13] Takayuki Kanda, Takayuki Hirano, Daniel Eaton, and Hiroshi Ishiguro. 2004. Interactive robots as social partners and peer tutors for children: A field trial. *Human–Computer Interaction* 19, 1-2 (2004), 61–84.

[14] Edith Law, Parastoo Baghaei Ravari, Nalin Chhibber, Dana Kulic, Stephanie Lin, Kevin D. Pantasdo, Jessy Ceha, Sangho Suh, and Nicole Dillen. 2020. Curiosity Notebook: A Platform for Learning by Teaching Conversational Agents. In *Extended Abstracts of the SIGCHI Conference on Human Factors in Computing Systems (Late Breaking Work) (CHI '20)*. 1–8.

[15] Noboru Matsuda, Evelyn Yarzebinski, Victoria Keiser, Rohan Raizada, William W. Cohen, Gabriel J. Stylianides, and Kenneth R. Koedinger. 2013. Cognitive anatomy of tutor learning: Lessons learned with simstudent. *Journal of Educational Psychology* 105, 4 (2013), 1152–1163.

[16] Roxana Moreno, Richard E Mayer, Hiller A Spires, and James C Lester. 2001. The case for social agency in computer-based teaching: Do students learn more deeply when they interact with animated pedagogical agents? *Cognition and instruction* 19, 2 (2001), 177–213.

[17] Anabil Munshi, Ramkumar Rajendran, Allison Moore, Gautam Biswas, and Jaclyn Ocumpaugh. 2018. Studying the interactions between components of self regulated learning in open ended learning environments. In *Proceedings of International Conference of the Learning Sciences, ICLS*. 1691–1692.

[18] Rod D. Roscoe and Michelene T.H. Chi. 2007. Understanding tutor learning: Knowledge-building and knowledge-telling in peer tutors' explanations and questions. *Review of Educational Research* (2007).

[19] Jong-Eun Roselyn Lee, Clifford Nass, Scott Brenner Brave, Yasunori Morishima, Hiroshi Nakajima, and Ryota Yamada. 2006. The case for caring colearners: The effects of a computer-mediated colearner agent on trust and learning. *Journal of Communication* 57, 2 (2006), 183–204.

[20] James R. Segedy, John S. Kinnebrew, and Gautam Biswas. 2012. Supporting student learning using conversational agents in a teachable agent environment. In *10th International Conference of the Learning Sciences: The Future of Learning, ICLS 2012 - Proceedings*.

[21] Jason Tan, Gautam Biswas, and Daniel L Schwartz. 2006. Feedback for Metacognitive Support in Learning by Teaching Environments. *Proceedings of the Annual Meeting of the Cognitive Science Society, 28* (2006), 828–833.

[22] Fumihide Tanaka, Kyosuke Isshiki, Fumiki Takahashi, Manabu Uekusa, Rumiko Sei, and Kaname Hayashi. 2015. Pepper Learns Together with Children : Development of an Educational Application. In *2015 IEEE-RAS 15th International Conference on Humanoid Robots (Humanoids)*. IEEE, 270–275.

[23] Fumihide Tanaka and Shizuko Matsuzoe. 2012. Children Teach a Care-Receiving Robot to Promote Their Learning: Field Experiments in a Classroom for Vocabulary Learning. *Journal of Human-Robot Interaction* 1, 1 (2012), 78–95.

[24] Elmira Yadollahi, Wafa Johal, Ana Paiva, and Pierre Dillenbourg. 2018. When deictic gestures in a robot can harm child-robot collaboration. (2018), 195–206.