Embodied Teachable Agents: Learning by Teaching Robots

Justin Werfel

Wyss Institute for Biologically Inspired Engineering,
Harvard University, Cambridge MA 02138, USA
justin.werfel@wyss.harvard.edu

Abstract. Robots have an untapped potential to contribute to education by acting as subordinate learners for students to teach. The benefits that the act of teaching provides for one's own learning have long been recognized; tutoring-associated improvements in measures like achievement scores, depth of understanding, and motivation are often far greater for the tutor than the tutee. Artificial agents can help students reap these benefits by providing surrogate pupils for them to teach, with potential advantages over human tutees. Robots have been observed to be more effective and compelling than virtual agents in a variety of contexts. However, research on teachable agents for education has been limited to virtual agents, while research on humans teaching robots has been concerned with learning for the benefit of the robot rather than that of the human. A new research direction exploring robots as teachable agents will lead to widespread benefits in education, and open new possibilities enabled by physical embodiment.

Keywords: educational technology, teachable agents, robotics in education

1 Introduction

Intelligent autonomous robots are making their way into an increasing range of application areas: manufacturing, transportation, surveillance and monitoring, rehabilitation, agriculture, service, and so on. One area which has so far been relatively neglected, and in which robots are uniquely qualified to contribute, is that of education. Robotic systems have the potential to increase student involvement and motivation and to improve the effectiveness of learning. Possible ways in which they could do so range from small personal robots used as teaching tools for areas like programming and computational thinking [1–3], to more science-fictional future scenarios with humanoid robots supplementing or even replacing teachers [4–6]; in this position paper I want to focus on another kind of possibility, one where robots act as subordinate learners for students to teach.

The observation that the act of teaching contributes to one's own learning goes back thousands of years [7]. Student teaching as a learning tool has been shown to provide benefits like improved retention and grasp of material, increased time and effort given to learning, and increased motivation and enjoyment [8–10]. The potential is vast and pervasive: some evidence suggests that teaching siblings contributes to the development of higher intelligence [11], and the benefits of learning by teaching have been observed in adults as well as children [12]. Recently virtual "teachable agents" have been used as a tool to help students learn [10, 13–15]: a student is tasked with training a simulated computer agent on course material, which can lead to a deeper and more committed understanding on the student's own part.

Robots have the potential to evoke particularly strong responses; robot helpers are perceived as more helpful, sociable, and enjoyable than virtual agents [16–18], and in some cases with children have been found to be preferable even to human interlocutors [19, 20]. Moreover, physical robots provide the opportunity for instruction on topics containing a physical component (e.g., motor tasks or demonstrations), where purely virtual agents are necessarily more limited. However, educational efforts with teachable agents have been focused exclusively on virtual agents. Research on the teaching of robots by humans [21,22] has instead been concerned with learning for the benefit of the *robot* rather than that of the *human*, often with interest in how human teaching can be characterized and leveraged.

I argue that robots as teachable agents for human education constitute an important and neglected research direction for the intelligent autonomous systems community. The most direct and obvious potential benefit of advances in this area is improved outcomes to education, at a time when education in general and STEM education in particular is identified as being in need of improvement [23, 24]. A robotic presence in classrooms may one day help to ameliorate problems of teacher shortages, as learning-by-teaching systems have historically been used to do (e.g., as in the case of the French *écoles mutuelles* in the eighteenth and nineteenth centuries [25]). And the data about human teaching styles that will be generated by widespread use of teachable robots in classroom settings will be invaluable in developing approaches to help robots learn more effectively from humans, where the goal is rapid and natural instruction of a robot in a new task [21, 22].

2 Learning by Teaching

Active learning, in which students participate in their own education, has long been recognized as being far more effective than passive learning [9]. When students are involved not just in listening to a teacher or reading a text but in discussion, synthesis, and problem-solving, the result is a deeper and fuller understanding and a stronger connection and commitment to the material.

One of the most active possible ways to participate in learning is to teach someone else. It is a commonplace that teachers learn from their students, and more generally from the act of teaching. Teaching requires and develops a deep grasp of material, as the teacher reformulates concepts, understands connections, explains a subject to a new audience and answers novel questions raised in response. Not only does teaching a subject confer empirical benefits over just studying it, but the act of teaching results in more effective learning than does preparing for teaching without following through on that expectation [26].

Many studies and systems have observed and taken advantage of these benefits to the teacher as a pedagogical tool. Learning by teaching has been used from one-room schoolhouses [9] to more modern settings like Sudbury schools [27], Montessori schools [9], and the "Lernen durch Lehren" method [28]. Students teaching others has for decades been demonstrated to lead to significant improvements in achievement scores [8, 9]—often far greater on the part of the tutor than that of the tutee [9]—and has been reported to provide emotional as well as cognitive benefits for the tutors, improving interest, attentiveness, motivation, attitude, and confidence [9]. Importantly, the act of teaching helps the tutor learn more effective ways to learn.

With the goal of harnessing such benefits for students, recent work has explored the use of "teachable agents", providing computer-simulated virtual agents that students can teach [10,13–15]. Such systems have been demonstrated to help students achieve quantitative learning gains, improve their reasoning, and organize their understanding of the relative importance of different concepts. Of particular note is the observation that students training teachable agents spend more effort on behalf of their tutee than they do if working on their own behalf, and likewise learn more [10]. Teachable agents can potentially also provide advantages for the tutor over human tutees, with the opportunity to make their "thinking" explicit and visible, and the possibility of demonstrating productive learning behaviors as a model to emulate.

3 The Role for Robots

Robots can potentially confer the benefits of the best aspects of both human tutees and teachable agents, with unique advantages of their own. Like virtual agents, they can make details of their learning process clear and explicit, present any learning style that may be found to be optimally helpful for the tutor, and give every student the opportunity to benefit from being a teacher. Like human partners, they are physically embodied, which allows for a much greater range of learning activities and social interaction channels, and increases the appeal and effectiveness for the student.

The unreasonable effectiveness of robots has been observed in a variety of contexts. Robotic assistants have been shown to be significantly more effective than virtual ones, as well as than non-agent alternatives (e.g., a paper log or smartphone), in scenarios including weight-loss coaching [16], driving assistance [17], and exercise coaching [18]. Not only are benefits in outcome observed (e.g., users stick with a weight-loss program for longer), but robots are consistently reported to have social advantages, perceived as being more supportive, more helpful, closer allies, more enjoyable to work with, and providing a deeper and more attractive social presence.

Such advantages hold true for children as well as for adults. Robots are broadly compelling across age ranges—in some cases young children even prefer interacting with robots rather than with humans [19]—as well as across boundaries of sex and background; robots are frequently cited as a particularly promising teaching tool to help traditionally underrepresented demographics become more interested and active in STEM fields [5, 29, 30], and have shown successful teaching outcomes in such settings [31]. The breadth of appeal extends to other populations in need of intervention, such as children with autism spectrum disorders, for whom robot partners can be preferable to human ones [20, 32].

The physically embodied aspect of robots opens the possibility of their use for teaching subjects with a physical or mechanical component, such as handwriting [4] or table tennis [33], in addition to purely academic subjects. Moreover, physical robots can be better suited than virtual agents to not only recognizing [34] but also making use of nonverbal channels critical in human communication [35], such as pointing and other gestures, posture, personal space, and touch.

Widespread adoption of robots as teachable agents would over time make available large data sets characterizing how people naturally teach. Such data would be of great value both for research into how robots can most effectively be taught new tasks by human instructors [21, 22], and for the insights into human learning and cognitive science it could provide. One possible long-term outcome of such research that could contribute back to pedagogy could be a quantitative characterization of different human teaching styles, and the development of the ability by a robot to quickly recognize its tutor's style, and to adapt its own learning style accordingly to best help that student learn. Such adaptation could be a particularly effective way of providing the individualization identified as being key to learning [9].

4 Scenarios and Challenges

Teachable robots for education will first be adopted as classroom tools, with small numbers initially available for students to take turns using; as robots become increasingly ubiquitous and affordable, usage may move toward every member of a class teaching their own robot at once; eventually, personal robots at home may be available for homework exercises. While the latter scenarios may be far off, sophisticated humanoid robots are already being purchased and used in public school systems [5].

Different interaction models will be possible. A long-term goal may be for students to teach robots through natural interactions, conversationally and through gesture, the way they would teach classmates or younger peers. Such modes will become increasingly feasible as work in human-computer interaction and human-robot interaction advances more generally. A specific direction of study for this application area may be a careful characterization of how human learners behave while being taught, to let robots emulate this behavior in order to provide appropriate and effective feedback to the students teaching them, in the interest of making the experience feel consistently rewarding and engaging to the tutor. In the shorter term, robot learners could be taught academic subjects with interfaces like those currently used for virtual teachable agents, e.g., graphical causal maps constructed by the students to explicitly represent the tutee's reasoning about key concepts and their relations [10, 14, 15].

The range of subjects that could potentially be taught encompasses any subject we may want students to learn, subject only to interface and hardware limitations. Work on virtual teachable agents has explored the teaching of science and math topics (e.g., river ecosystems, trip planning, climate change), using interfaces like concept maps and multiple-choice questions [13–15]. Virtual agents have been used as tutors to teach strategies for improving reading comprehension, based on analysis of text entered by a student by keyboard [36]; an inversion of this approach could make the student responsible for teaching such strategies to the artificial agent. Younger students could teach robots reading, to reinforce their own skills, as students currently sometimes do with younger peers—reading to the robot gives the student direct practice, and could let the robot evaluate their strengths and weaknesses for later reference; supervising the robot reading, and correcting its (deliberate and targeted) mistakes or working with it to sound out words together, would put the student in a more explicitly mentorial role. Physical robots could also potentially engage in more active hands-on pursuits like sports and crafts, as learning-through-teaching programs have historically encompassed [9].

The way a student teaches will carry information about their understanding of and approach to a subject; that information could potentially be used by a robot learner to highlight points of weakness, and to strengthen them. For instance, for topic learning, the robot can ask questions that focus on less well-covered areas; for an activity like learning handwriting, the robot could demonstrate behavior that exaggerates the student's poor practices, to make it clearer to them what they need to improve. By mirroring the student's own understanding, the robot can guide their efforts in a personalized and non-threatening way.

As progress in robotics in general makes robots more capable and self-reliant, the assistance of such robots may eventually help teachers more effectively handle larger groups of students at once. Such a capability may be especially useful given perennial concerns about teacher shortages, notably in STEM fields [24]. The potential value of the use of robots may be of particular interest in special education [37], considering particular teacher shortages identified there [38] and the observed effectiveness of robotic technology for at least some special-needs populations [20,32]. More broadly, proliferation of educational robots may better enable more extensive use of one-on-one training, identified as more effective for learning but not feasible with typical teacher-student ratios [36].

One set of key challenges to developing the potential of robots as teachable agents will involve research on how the robot can act as a learner in ways that best help the student. The work on virtual teachable agents will serve as a natural starting point for this area of study. Research on how students teaching robots learn, and exploration of different approaches they might use, will help advance the understanding of how teaching and learning interact, which will in turn improve our ability to design effective robot learners and optimize the experience for individual tutors. General progress in human-robot interaction research will advance both the ability of robots to act effectively for students' benefit and the ease with which students can interact with them.

A second set of challenges will be associated with hardware design. Humanoid robots [5] will likely be most generally effective as sympathetic learners, but alternative designs may be more effective for certain purposes [39], potentially in a culture-dependent way [40]. A need for broad physical capabilities will be opposed by a need for low cost in order to enable widespread adoption; one possibility is accordingly a general-purpose robot with minimal physical facility for use in most cases, plus specialized or reconfigurable robots for specialized physical learning tasks as called for. A single consistent "personality" developed through ongoing association with a particular student can accompany them to activities with different physical bodies. An intermediate approach for a limited-cost, general-purpose educational robot might be to use a physical robot head for social interaction, plus a virtual body manipulating animated objects on a video screen. Such a system could flexibly interact with a student for teaching a wide variety of geometric and object-based topics, without needing to solve major open problems of real-world manipulation in the short term, while still providing the physical presence that has been found empirically to be more effective than purely virtual agents.

5 Summary

In this position paper I have argued that an important new research direction for the intelligent autonomous systems community is education, in the form of developing theory and practice for embodied learning agents which students can teach in order to gain personal benefit. This area leverages the established benefits of learning by teaching, together with the established benefits of robots as social and educational tools. The many challenges associated with making this vision a reality will provide many opportunities for advancing the state of the art of robotics and its relevance to society, with critical impact given the universal importance of education.

References

1. Papert, S.: Mindstorms: children, computers, and powerful ideas. Basic Books, New York (1980)
2. Balch, T. et al.: Designing personal robots for education: hardware, software, and curriculum. IEEE Pervasive Computing 7(2), 5–9 (2008)
3. Riedo, F., Rétornaz, P., Bergeron, L., Nyffeler, N., Mondada, F.: A two years informal learning experience using the Thymio robot. Advances in Autonomous Mini Robots, 37–48 (2012)
4. AFP: Calligraphy robot teaches Japan's schoolchildren the art of 'shodo' writing. August 1, 2013
5. Stansbury, M.: Classroom robot boosts STEM, engages children with autism. eSchool News. April 10, 2013
6. Asimov, I: The fun they had. Boys and Girls Page, NEA Service (1951)
7. Seneca: Epistulae morales ad Lucilium 7 (65)
8. McHugh, W.: Team learning in skills subjects in intermediate grades. Boston University Journal of Education 142, 22–51 (1959)
9. Gartner, A., Kohler, M., Riessman, F.: Children teach children: Learning by teaching. Harper & Row, New York (1971)
10. Chase, C., Chin, D., Oppezzo, M., Schwartz, D.: Teachable agents and the protégé effect: increasing the effort towards learning. J. Sci. Educ. Technol. 18, 334–352 (2009)
11. Kristensen, P., Bjerkedal, T.: Explaining the relation between birth order and intelligence. Science 316, 1717 (2007)
12. Cortese, C.: Learning through teaching. Management Learning 36, 87–115 (2005)
13. Biswas, G., Katzlberger, T., Bransford, J., Schwartz, D.: Extending intelligent learning environments with teachable agents to enhance learning. In J. Moore, C. Redfield, W. Johnson (eds.), Artificial Intelligence for Education 01. IOS Press, Amsterdam, 389–397 (2001)
14. Biswas, G., Leelawong, K., Schwartz, D., Vye, N.: Learning by teaching: a new agent paradigm for educational software. Applied Artificial Intelligence 19, 363–392 (2005)
15. Kinnebrew, J., Biswas, G., Sulcer, B., Taylor, R.: Investigating self-regulated learning in teachable agent environments. In R. Azevedo, V. Aleven (eds.), International Handbook of Metacognition and Learning Technologies. Springer, New York, 451–470 (2013)
16. Kidd, C., Breazeal, C.: Robots at home: understanding long-term human-robot interaction. IEEE Int. Conf. on Intelligent Robots and Systems, 3230–3235 (2008)
17. Williams, K., Breazeal, C.: Reducing driver task load and promoting sociability through an affective intelligent driving agent (AIDA). Human-Computer Interaction (INTERACT), 619–626 (2013)
18. Fasola, J., Matarić, M.: A socially assistive robot exercise coach for the elderly. J. Human-Robot Interaction 2, 3–32 (2013)
19. Breazeal, C.: Personal robots. Lecture, Robots: From Imagination to Market, Cambridge, MA, April 17, 2014
20. Diehl, J. et al.: Clinical applications of robots in autism spectrum disorder diagnosis and treatment. In V. Patel et al. (eds.), Comprehensive Guide to Autism. Springer, New York, 411–422 (2014)
21. Knox, W., Glass, B., Love, B., Maddox, W., Stone, P.: How humans teach agents: a new experimental perspective. Int. J. Soc. Robot. 4, 409–421 (2012)

22. Thomaz, A., Breazeal, C.: Teachable robots: understanding human teaching behavior to build more effective robot learners. Artificial Intelligence 172, 716–737 (2008)
23. Strauss, R.: Remedial education: federal education policy. Council on Foreign Relations Press (2013)
24. U.S. Department of Education: Science, technology, engineering and math: education for global leadership. `http://www.ed.gov/stem` (2014)
25. Lesage, P.: La pédagogie dans les écoles mutuelles au XIXe siècle. Revue française de pédagogie, 62–70 (1975)
26. Fiorella, L., Mayer, R.: The relative benefits of learning by teaching and teaching expectancy. Contemporary Educational Psychology 38, 281–288 (2013)
27. Greenberg, D.: Sudbury Valley's secret weapon: allowing people of different ages to mix freely at school. In The Sudbury Valley School Experience (3rd ed.). Sudbury Valley School Press, Framingham (1992)
28. Martin, J.-P., Kelchner, R.: Lernen durch Lehren. In J.-P. Timm (ed.), Englisch lernen und lehren—Didaktik des Englischunterrichts. Cornelsen, Berlin, 211–219 (1998)
29. Arroyo, I., Arroyo, A., Schwartz, E.: The integration of robotics in grade school to promote interest in math and science. Florida Conference on Recent Advances in Robotics (2003)
30. Yuen, T., Ek, L., Scheutze, A.: Increasing participation from underrepresented minorities in STEM through robotics clubs. IEEE Int. Conf. on Teaching, Assessment and Learning for Engineering, 24–28 (2013)
31. Kazakoff, E., Sullivan, A., Bers, M.: The effect of a classroom-based intensive robotics and programming workshop on sequencing ability in early childhood. Early Childhood Educ. J. 41, 245–255 (2013)
32. Bekele, E. et al.: A step toward developing adaptive robot-mediated intervention architecture (ARIA) for children with autism. IEEE Trans. Neural Syst. and Rehabilitation Eng. 21, 289–299 (2013)
33. Matsushima, M., Hashimoto, T., Takeuchi, M., Miyazaki, F.: A learning approach to robotic table tennis. IEEE Trans. Robotics 21, 767–771 (2005)
34. Yang, H.-D., Park, A.-Y., Lee, S.-W.: Gesture spotting and recognition for human-robot interaction. IEEE Trans. Robotics 23, 256–270 (2007)
35. Knapp, M., Hall, J., Horgan, T.: Nonverbal communication in human interaction (8th ed.). Wadsworth, Boston (2014)
36. Jackson, G., Dempsey, K., McNamara, D.: The evolution of an automated reading strategy tutor: from the classroom to a game-enhanced automated system. In M. Khine, I. Saleh (eds.), New Science of Learning. Springer, New York, 283–306 (2010)
37. Kimbler, D.: Robots and special education: the robot as extension of self. Peabody J. Education 62, 67–76 (1984)
38. Boe, E.: Teacher demand, supply, and shortage in special education: a national perspective. In P. Sindelar et al. (eds.), Handbook of Research on Special Education Teacher Preparation. Routledge, New York, 67–93 (2014)
39. Woods, S., Dautenhahn, K., Schulz, J.: The design space of robots: investigating children's views. Proc. IEEE Int. Workshop on Robot and Human Interactive Communication, 47–52 (2004)
40. Li, D., Rau, P., Li, Y.: A cross-cultural study: effect of robot appearance and task. Int. J. Soc. Robot. 2, 175–186 (2010)