

Introduction to the Special Articles in the Fall and Winter Issues

Intelligent Learning Technologies Part 2: Applications of Artificial Intelligence to Contemporary and Emerging Educational Challenges

Vinay K. Chaudhri, Dave Gunning, H. Chad Lane, Jeremy Roschelle

■ *The fall and winter issues of AI Magazine present articles on some of the most interesting projects at the intersection of AI and education. Included are articles on integrated systems such as virtual humans, an intelligent textbook, and a game-based learning environment as well as technology-focused components such as student models and data mining. This issue concludes with an article summarizing the contemporary and emerging challenges at the intersection of AI and education.*

It is a time of remarkable transformation for education. Everyone recognizes the need to improve teacher effectiveness, to improve student engagement, and to create a twenty-first century education system that maximizes potential of every student. The challenges that must be addressed to make these improvements greatly exceed the scope of any single approach, whether it is educational technology, improved teacher training and better after school programs, and so on. In past research, AI — with its inextricable links to cognitive science, psychology, and mathematics — has proven a close fit for many of these challenging educational problems.

Educators have long known that the most effective teaching method is one-on-one tutoring. Ever since Benjamin Bloom's famous study,¹ education researchers have aspired to mimic the holy grail of one-on-one tutoring — to achieve a two-sigma improvement in student learning. Early AI researchers saw this as an opportunity to build intelligent tutoring systems (ITSs) that could adapt and tailor instruction to the individual needs of the student. Although today's systems fall short of the full two-sigma effect of a human tutor (roughly equivalent to two grade levels), intelligent tutors have demonstrated remarkable progress in that direction. In fact, researchers have struggled to replicate the two-sigma effect suggesting that ITSs may already be as effective as human tutors. This suggests that the contributions of AI to education are perhaps more profound than previously believed and leads us to wonder why AI-based learning tech-

nologies are not in every classroom, every home, every library, and on every mobile device.

Today, ITSs are in widespread use in K–12 schools and colleges and are enhancing the student learning experience. As a specific example, 600,000 students in more than 2600 middle or high schools use Carnegie Learning's Cognitive Tutor mathematics courses regularly. Some full-year evaluation studies of Cognitive Tutor algebra have demonstrated better student learning compared to traditional algebra courses. Debate has ensued, however, involving both the efficacy of intelligent tutoring as well as whether the ITS approach can scale up to meet the broad needs of educational systems throughout the world. More broadly, studies have consistently shown that when a system models the knowledge it seeks to teach and uses that knowledge to assess, track, and scaffold learning, that it is more effective than non-intelligent counterparts.

We face unprecedented challenges in science, technology, engineering, and mathematics (STEM) education. Projections on the availability and needed competencies of STEM workers are disconcerting: current workers with STEM training are retiring and not enough of the younger generation will either be available or adequately trained to meet emerging needs. Recent reports suggest important links between education and health, civic involvement, criminality, and even eligibility for military service. Although these challenges transcend STEM education, increasing student interest and preparing students for STEM remain key challenges for the future of education systems throughout the world.

The goal of the special articles in the fall and winter issues is twofold; first, to present some of the best work at the intersection of AI and education in a way that highlights the power of AI to promote human learning; and second, to define the needs and challenges facing STEM education today and align those definitions with modern and emerging AI research.

The idea of putting these issues together and many of the articles included here originated from the first cyberlearning summit sponsored by

the National Science Foundation.² The topics cover a significant cross section of AI and, in our opinion, represent some of the best interdisciplinary research in education today. Articles cover topics including uncertain reasoning, natural language processing, data mining, knowledge representation, explanation systems, automated reasoning, and more. Further, because it is critical to not lose sight that real, meaningful learning is the goal, authors also report empirical findings related to teaching efficacy and learning. We briefly introduce each of the articles in these two special issues.

In the fall issue, Cristina Conati and Samad Karden presented the latest on student modeling for novel (less structured, nonproblem-solving) domains. The essence of this problem is to model the complex ensemble of processes and states that constitutes human learning. The user-modeling framework presented in this article was an attempt to address this challenge. It uses clustering and class association rule mining to discover and recognize relevant interaction patterns during student interaction with educational software.

Kenneth R. Koedinger and his colleagues discussed the prospects for data-driven development and optimization of educational technologies, focusing on intelligent tutoring systems and illustrating techniques especially in the context of Cognitive Tutor courses. This article described the use of data-driven techniques to develop or optimize the key functions of an intelligent tutor such as select, evaluate, suggest, and evaluate.

Vasile Rus and his colleagues described ITSs with conversational dialogue, especially advances in macro- and microadaptivity made possible by the use of learning progressions and deeper dialogue and natural language processing techniques. Learning progressions are the central theme in their DeepTutor system around which everything else (domain modeling, assessment, and instructional tasks) is organized and aligned.

Finally, Vinay Chaudhri and coauthors described a new kind of intelligent textbook that answers students' questions, engages their interest, and

improves their understanding. The textbook leverages the knowledge representation of the textbook to offer functions such as concept summaries, suggested questions, and answers to a student's questions.

In this issue, William Swartout and his colleagues discuss how virtual humans represent a new metaphor for interacting with computers, one in which working with a computer becomes much like interacting with a person, and this can bring social elements to the interaction that are not easily supported with conventional interfaces. These emerging technologies focus on informal science education and represent an intersection of AI with modern video game and computer graphics techniques.

James Lester and coauthors describe work in intelligent game-based learning environments that integrate commercial game technologies with AI methods from intelligent tutoring systems and intelligent narrative technologies. Crystal Island features a science mystery where students attempt to discover the identity and source of an infectious disease that is plaguing a research team stationed on a remote island.

Bert Bredeweg and his colleagues describe DynaLearn, an intelligent learning environment that allows learners to acquire conceptual knowledge by constructing and simulating qualitative models of how systems behave, where students learn by constructing their own models of the domain.

Finally, Beverly Park Woolf and her coauthors wrap up this issue by unpacking contemporary educational challenges and arguing that they are ideal targets for AI research. They describe five grand challenges for AI and education that encourage an interdisciplinary perspective on needed future research. AI-based learning technologies are discussed as a critical piece of the future of education. Leveraging what is known about human learning and building on the substantial progress that has been made in the last four decades, the vision for the future includes a heavy role for group and collaborative learning, use of large data sets, models of continuous learning

Introducing a New AI Magazine Column in This Issue!

Educational Trends in Artificial Intelligence

The emergence of massive open online courses has initiated a broad national-wide discussion on higher education practices, models, and pedagogy. Artificial intelligence and machine learning courses were at the forefront of this trend and are also being used to serve personalized, managed content in the back-end systems.

Massive open online courses are just one example of the sorts of pedagogical innovations being developed to better teach AI. This column will discuss and share innovative educational approaches that teach or leverage AI and its many subfields, including robotics, machine learning, natural language processing, computer vision, and others at all levels of education (K-12, undergraduate, and graduate levels). In particular, this column will serve the community as a venue to learn about the Symposium on Educational Advances in Artificial Intelligence (EAAI) (colocated with AAAI for the past four years); introductions to innovative pedagogy and best practices for AI and across the computer science curricula; resources for teaching AI, including model AI assignments, software packages, online videos and lectures that can be used in your classroom; topic tutorials introducing a subject to students and researchers with links to articles, presentations, and online materials; and discussion of the use of AI methods in education shaping personalized tutorials, learning analytics, and data mining.

– Laura Brown

throughout life, and a focus on skills needed for success in the modern world.

Acknowledgments

Several of our colleagues wrote thoughtful reviews for the papers in this issue, and we sincerely thank them for their efforts. The reviewers included Tiffany Barnes, Lee Becker, Daniel Bobrow, Ryan Carlson, Toby Dragon, Peter Khooshabeh, Jack Mostow, Ray Perrault, Jonathan Rowe, Vasile Rus, Rory Sie, Bryan Wiltgen, Beverly Woolf, Ruth Wylie, and Louis Yarnall. We are grateful for their help and for the editors at *AI Magazine* in making this special issue possible.

This material is based upon work supported by the National Science Foundation under Grant No. 1132393. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

Note

1. See Benjamin S. Bloom, 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.

2. See cyberlearning.sri.com.

Vinay K. Chaudhri received his Ph.D. in computer science from University of Toronto, Canada, and is currently a program director in the Artificial Intelligence Center at SRI International. His research focuses on the science, engineering, and application of large knowledge base systems.

David Gunning is the program director for Enterprise Intelligence at the Palo Alto Research Center (PARC). He directs PARC's efforts in artificial intelligence and predictive analytics focused on the enterprise. These include projects in anomaly and fraud detection, contextual intelligence, recommendation systems, and tools for smart organizations. Prior to PARC, Gunning was a senior research program manager at Vulcan Inc., a program manager at DARPA (twice), SVP of SET Corp., vice president of Cycorp, and a senior scientist in the Air Force Research Labs. At DARPA, he managed the Personalized Assistant that Learns (PAL) project that produced Siri, which was acquired by Apple, and the Command Post of the Future (CPoF) project, which was adopted by the U.S. Army for use in Iraq and Afghanistan. Gunning holds an M.S. in computer science from

Stanford University, an M.S. in cognitive psychology from the University of Dayton, and a B.S. in psychology from Otterbein College.

H. Chad Lane is a research scientist at the University of Southern California's Institute for Creative Technologies. His research is highly interdisciplinary and involves the application of entertainment and intelligent technologies to a variety of challenges, including education and health behavior. A significant portion of this work investigates the many roles virtual humans can play in virtual learning environments, such as coach and colearner, as well as in the context of mobile educational games. Lane received his Ph.D. in computer science from the University of Pittsburgh in 2004. Currently, he serves on the Artificial Intelligence in Education (AIED) Society executive committee and as program cochair for the AIED 2013 Conference.

Jeremy Roschelle specializes in the design and development of integrated interventions to enhance learning of complex and conceptually difficult mathematics and science; learning sciences-based research in mathematics education, on collaborative learning, and with interactive technology; and the management of large-scale multi-year, multiinstitutional research and evaluation projects.