Editor: **Judy Kay,** University of Sydney, judy.kay@sydney.edu.au

AI and Education: Grand Challenges

Judy Kay, *University of Sydney*

The educational applications of AI—typical of much AI research—are a combination of what Pasteur's Quadrant describes as *use-inspired basic* and *pure applied* research.[1] Figure 1 illustrates the way educational problems drive AI in education (AIED) research to create new learning systems as well as the foundational knowledge that will enable them to be built. The figure also introduces the term *intelligent tutoring systems* (ITS), which is generally considered a near synonym of AIED. The figure shows how we can create AIED/ITS systems by drawing upon existing AI theories, tools, and techniques (represented by the downward arc at the lower left, showing pure applied research). Equally, the demands of the systems may drive researchers to make advances in fundamental AI, as indicated by the upward arc at the lower right, representing use-inspired basic research.

AIED research has had a similar relationship to models of human cognition, affect, motivation, and so on. The figure shows that these, too, can drive fundamental AI research and that there is a reciprocal relationship, with fundamental AI serving to create models that inform understanding of cognition. There are additional links between models of understanding of human cognition and AIED systems. Important AIED research includes studies of people, to gain understanding of what makes for effective learning and teaching. Similarly, the gold standard for evaluation of AIED systems measures whether they enhance learning outcomes; this can inform models of cognition.

This inaugural article in the new AI in Education department provides an overview of both the classic and emerging architectures for AIED systems, pointing to the many aspects of AI that play an important role in creating these systems. What is AIED? One way to answer this question is to look at the key publication venues for the AIED research community (http://ijaied.org/about). The 15th biennial AI in Education conference (AIED) will be held in 2013, and the 11th Intelligent Tutoring Systems conference, which runs in alternate years with AIED, was held earlier this year. The *International Journal of Artificial Intelligence in Education* began publishing in 1989. Much has stayed the same since then in terms of the goals of AIED. Of course, with the huge advances in AIED and technology, there have also been important changes. Figure 2 shows the typical architecture of AIED systems, with the key elements in boxes and the relevant AI research in unboxed italics.

The boxes with italicized labels were part of the vision of the earliest AIED researchers, who aimed to create personalized teaching systems. They were attempting to enhance learning outcomes by the huge margins that a highly skilled human teacher can achieve in one-to-one teaching, which can raise the performance of the average 50 percent–achieving student by two sigmas, to the levels achieved by the top 5 percent in typical classrooms.[2]

This vision had a focus on the solitary learner, whereas more recent work takes account of the broader learning context, in terms of the other people involved and the learning context. Among the other people involved are the learners' parents.

Not all AIED systems have all the elements in the figure, not even all the italicized ones. However, one of the most common is the *learner model*, which holds a representation of the learner's current knowledge, misconceptions, preferences, goals, and so on. The model needs to represent those aspects of the learner that will drive the whole system, including interpretation of the learner 's actions and the system's generation of its own actions. This is, arguably, the core of AIED/ITS systems because it enables the system to "care precisely" about the learner;[3] the model enables the system to better understand the learner, and it drives teaching personalization. The figure shows it in bold because its nature has changed.

1541-1672/12/$31.00 © 2012 IEEE
Published by the IEEE Computer Society

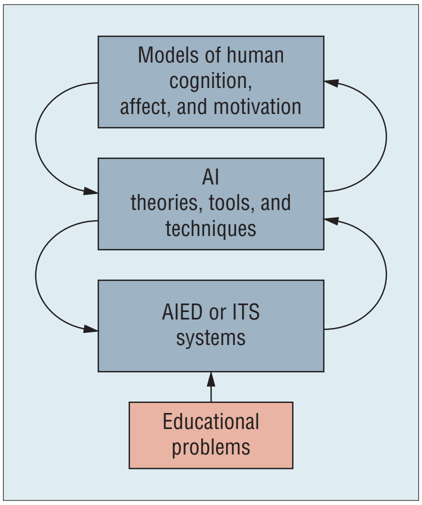

Figure 1. Problem-driven AI research. Educational needs drive AIED research, which both draws on and contributes to AI theories, tools, and techniques. A similar synergistic two-way relationship holds between AIED and research in human cognition with AI.

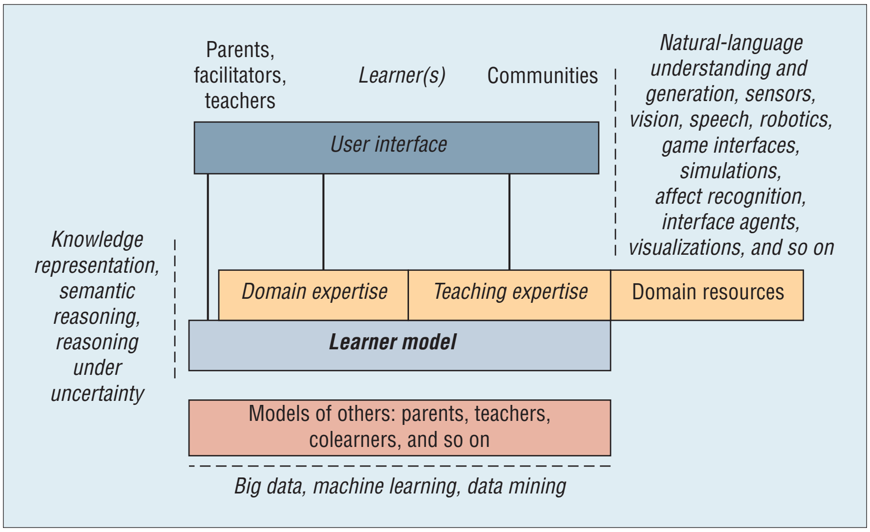

Figure 2. Key elements of AI in Education systems. From its foundation, AIED research has had four core elements: a model of the learner, domain expertise, teaching expertise, and interfaces. Recent research takes a broader view, recognizing the importance of other people and rich collections of digital learning resources.

Figure 2 shows the learner model as closely linked to two other key elements: *domain expertise* and *teaching expertise*. In practice, AIED systems need to have a tight coupling across these elements: all three need to operate together for effective teaching. The left side of the figure characterizes aspects of AI that are important for all three elements, reflected in the core topics of recent conferences.[4,5] For example, one strand aims to improve AI theories, tools, and techniques for knowledge representation, semantic reasoning, and reasoning under uncertainty. AIED systems need to operate with uncertain, inconsistent, and noisy sources of information about the learner. For example, learners make slips as well as guesses that happen to be the correct answer. Above all, AIED systems are designed to be nonmonotonic; that is, the model is expected to change as the student learns!

The fourth element of the classic AIED systems view is the user interface. This encompasses a very wide range of AIED research across perception, interface generation, and intelligent immersive environments, as Figure 2 indicates. Many AIED researchers strive to create interfaces that enhance learning effectiveness by making the interaction natural and compelling. For example, a long-term, important strand of AIED research has aimed to provide natural-language interfaces, calling for natural-language understanding and generation and, in some cases, speech understanding and generation. Recent AIED research has made strong contributions to creating interface agents in the form of an avatar that provides both language, facial expression, and an identity associated with aspects such as race and gender.[6] There has also been considerable work on creating interfaces that interpret the learner's affective and motivational state. The figure mentions some of the other, diverse interface approaches, such as simulations and game environments.[7]

Figure 2's depiction of an AIED architecture shows teaching expertise driving interface actions and interpreting the information from the learner, drawing upon the learner model and domain expertise. There has been a large body of diverse AIED research up to now. Importantly, some strands of work have matured to the point where they can "move into the wild," into classrooms and widely used online learning systems. One important development has been the establishment of systems founded on cognitive theories, and there has been real progress in creating tools that make it easier to build upon components of previous systems. This is important in giving a foundation for exploring new challenges and for moving towards real-world deployment. Examples include the widely deployed cognitive tutors and constraint-based tutors in daily use by thousands of students.[8,9]

The remaining elements of the figure are not part of the orthodox view described so far. Rather, they reflect more recent and emerging trends, notably the way it treats the learner model as more than simply part of an AIED system. It can have a direct link to an interface as an open learner model.[10] More than this, it can be a first-class citizen and exist independently of any single system.[11] The figure also shows that there is a role for models of other actors. This is part of the task of supporting the individual learner as part of a broader social and learning group, but it involves challenges of sharing data. Increasingly large collections of data about learners and their learning processes create an important and growing role for machine learning and data mining, reflected in the emergence of the educational data mining community (www.educationaldatamining.org) with its own annual conference and journal.[12]

The remaining part of the figure is the *domain resources*, which represent the wealth of learning objects and tools available, especially on the Web. So, for example, recommender technologies can play an important role for AIED.[13] Broadly, most areas of AI research are relevant to and can be informed by AIED research.

The importance of AIED research has been increasingly recognized, as reflected in identified "grand challenges." For example, the Computing Research Association identified just five Grand Research Challenges in Information Systems (http://archive.cra.org/Activities/grand.challenges). One of these, A Teacher for Every Learner, described the potential for computing to transform education, enhancing learning outcomes for all and transforming the way that people learn throughout their lives—the Grand Research Challenge aimed to provide all learners with learning environments that "approach the effectiveness of a one-on-one human tutor." The challenge proposed specific research toward that goal in terms of highly effective personalized teaching systems, simulation-based educational software, massive multiplayer online games, collaborative authoring, learning in context, and just-in-time learning.

The National Academy of Engineering has identified 14 Grand Challenges for Engineering, one of which is to Advance Personalized Learning (www.engineeringchallenges.org/cms/8996/9127.aspx). The challenge description points to successful Web-based personalized learning materials and systems, as well as recommender systems that help the learner find the right materials from the vast array available. It points to the potential of educational data mining, which can exploit the digital traces from digital learning activities. Like the figure above, the challenge recognizes the two-way link between the human brain and learning-systems engineering, based on neuroscience and medical measurement technology.

The UK Computing Research Committee identified nine Grand Challenges in Computing Research. One, Learning for Life, calls for the creation of technology that lets each learner learn as he or she can and wants to do, and that connects the right teachers and learners. This challenge notes the need for effectiveness at the levels of the individual, groups, and society. Unlike the other Grand Challenges, this one goes beyond the individual learner. Another UK Grand Challenge is Memories for Life, recognizing that we are defined by our memories. Augmented cognition is linked to a broad view of AIED, since our augmented memories are critical to lifelong and lifewide learning. A third challenge, called Bringing the Past to Life for the Citizen, lays out a bold vision of exciting learning in history and cultural heritage. In all, about half of these Grand Challenges relate strongly to AIED research.

One last example of such visionary statements comes from the report of the Microsoft Being Human: Human-Computer Interaction in the Year 2020 forum (http://research.microsoft.com/en-us/um/cambridge/projects/hci2020/). The report points to a range of aspects of learning, including formal learning contexts as well as learning with emerging mobile and ubiquitous-interaction devices. It recognizes the potential importance of rich and fine-grained data from the learner 's digital footprints. This can change the nature of assessment and provide parents with richer understanding of a child's progress. It also recognizes the importance of learning in relation to lifelong health needs, particularly as people age and need to take account of challenges to health and wellness. Emerging technology can capture valuable data about these and other aspects of life, but we need to enable people to make effective use of that information. While this vision is cast as dealing with human-computer interaction, it is replete with AI and AIED challenges—which is characteristic of the nature of AIED.

AIED has already achieved much. At the same time, its progress has highlighted the potential to achieve practical and profound improvements in education. Much work lies ahead in tackling the Grand Challenges. We need to both exploit and contribute to research in cognition and AI to establish validated AIED tools and techniques. When we have established collections of these, it will be feasible, perhaps even easy, to build the new elements in the AIED architecture and to modify existing elements to tune them to different learning contexts. The challenges include a combination of building intelligent systems with empirical studies, both in lab and in the wild, to accumulate evidence of the learning benefits of these tools, techniques, and systems. Upcoming articles will address these and other emerging issues.

References

1. D. Stokes, *Pasteur's Quadrant: Basic Science And Technological Innovation,* Brookings Inst. Press, 1997.

2. B. Bloom, "The 2 Sigma Problem: The Search for Methods of Group Instruction as Effective as One-To-One Tutoring," *Educational Researcher*, vol. 13, no. 6, 1984, pp. 4–16.
3. J. Self, "The Defining Characteristics of Intelligent Tutoring Systems Research: ITSs Care, Precisely," *Int'l J. AI in Education*, vol. 10, no. 3, 1999, pp. 350–364.
4. G. Biswas et al., eds., *Proc. 15th Int'l Conf. AI in Education* (AIED 11), LNCS 6738. Springer, 2011.
5. S. A. Cerri et al., eds., *Proc. 11th Int'l Conf. Intelligent Tutoring Systems* (ITS 12), LNCS 7315. Springer, 2012.
6. I. Arroyo et al., "Affective Gendered Learning Companions," *Proc. 14th Int'l Conf. AI in Education* (AIED 09), IOS Press, 2009, pp. 41–48.
7. W. Johnson, H. Vilhjalmsson, and S. Marsella, "Serious Games for Language Learning: How Much Game, How Much AI?" *Proc. 12th Int'l Conf. AI in Education* (AIED 05) IOS Press, 2005, pp. 306–313.
8. K. Koedinger et al., "Intelligent Tutoring Goes to School in the Big City," *Int'l J. AI in Education*, vol. 8, no. 1, 1997, pp. 30–43.
9. A. Mitrovic, "An Intelligent SQL Tutor on the Web," *Int'l J. AI in Education*, vol. 13, no. 2–4, 2003, pp. 173–197.
10. S. Bull and J. Kay, "Student Models That Invite The Learner In: The SMILI Open Learner Modelling Framework," *Int'l J. AI in Education*, vol. 17, no. 2, 2007, pp. 89–120.
11. J. Kay, "Lifelong Learner Modeling for Lifelong Personalized Pervasive Learning," *IEEE Trans. Learning Technologies*, vol. 1, no. 4, 2008, pp. 215–228.
12. J. Stamper and K. Koedinger, "Human-Machine Student Model Discovery and Improvement Using Datashop," *Artificial Intelligence in Education*, Springer, 2011, pp. 353–360.
13. P. Brusilovsky and E. Millan, "User Models For Adaptive Hypermedia And Adaptive Educational Systems," *The Adaptive Web*, LNCS 4321, Springer, 2007, pp. 3–53.

Judy Kay is a professor of computer science in the School of Information Technologies at the University of Sydney and a principal in the CHAI: Computer Human Adapted Interaction Research Group. Contact her at judy.kay@sydney.edu.au.