Exploiting Artificial Intelligence (AI) for Personalized Learning at Scale

The last generation has seen huge progress in both the science and the technology of education. On one hand, scientists now have a deeper understanding of the cognitive, social, and cultural processes of learning. This includes a greater understanding of the role of personalization, formative assessment, and metacognition in learning. On the other hand, the computing technologies of the Internet and social media are transforming education.

The Internet and social media have led to a variety of new online programs, including Georgia Tech's online Master of Science in Computer Science (OMSCS) and online Master of Science in Analytics (OMS Analytics) programs. Further, the combination of the new understanding of learning and the availability of online educational materials has led to broad adoption of pedagogies such as the flipped classroom and blended learning. Georgia Tech has been a leader in all of these spaces (e.g., see Day and Foley 2006; Madden et al., forthcoming).

In the science of education, we can expect continued progress in the cognitive and learning sciences, especially cognitive neuroscience. A generation from now, scientists should have a much deeper understanding of how the human brain processes information and how the human mind learns. In technology, too, we can expect rapid progress, especially in artificial intelligence (AI), which is already beginning to impact education in a myriad of ways, including through the use of intelligent tutoring systems and question-answering agents. The Georgia Tech Commission on Creating the Next in Education (CNE) posits that the current movement toward scale (exemplified by OMSCS) and personalization will not only continue but also accelerate, with AI acting as the key accelerator.

Further, Georgia Tech is well positioned to be leading in the growing movement toward exploiting AI for personalized learning at scale. In particular, the Commission recommends the following kinds of projects, classified here by when the Commission may expect the projects to yield rich outcomes (near, middle, and long terms).

Outcomes in the Near Term (1–2 years): Mastery Learning and Adaptive Learning Platforms

Mastery learning, which has been tested across subject matters and student populations for over thirty-five years, is a process in which learners can adjust their approach based on performance feedback and keep trying a new skill until they reach a desired level of performance. Coupling the pedagogy of mastery learning with online, adaptive learning technology can provide effective personalized learning solutions.

These adaptive platforms, due to their ability to provide flexible learning, remove time as a critical variable to a large extent and can be an effective response to Bloom's 2 sigma problem (Bloom 1984). Although mastery learning and personalized/adaptive learning are not new, learning environments are still challenged in their ability to bring a larger segment (70–80 percent) of students to the achievement levels that are conventionally expected of the top 10–20 percent of students.

Georgia Tech has an opportunity—and, as a state institution, a responsibility—to increase students' STEM success and engagement in courses by exploring these effective learning methodologies. Such explorations can also provide solutions for bottlenecks and large classes and create opportunities for blended learning that can generate alternative pathways to affordable college.

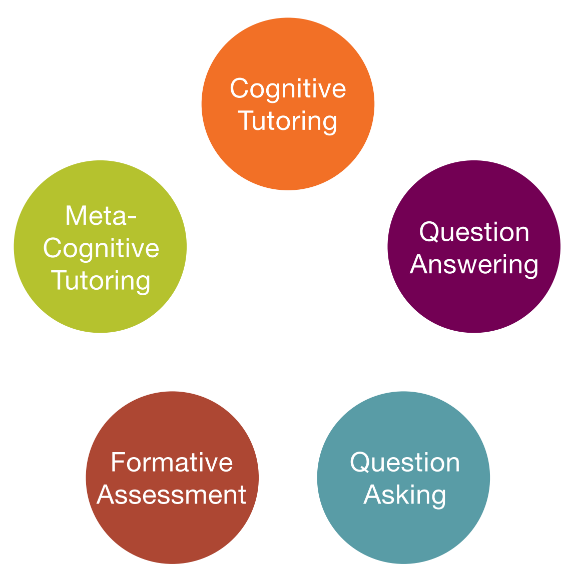

Figure 1: Some tasks in AI-based personalized learning.

Recently Georgia Tech conducted preliminary experiments in mastery and adaptive learning. From fall 2014 through spring 2017, Georgia Tech offered an OMSCS 7637 course on Knowledge-Based Artificial Intelligence that uses about 150 interactive exercises and 100 cognitive tutors embedded in about 18 hours of video lessons. That is approximately one interactive exercise every eight minutes of video. The format of this experiment supported mastery and adaptive learning (Goel and Joyner 2017). This same course also debuted "Jill Watson," an AI agent designed to answer routine questions within the class discussion forum (Goel and Polepeddi 2017). Initial evaluation shows that the online students (perhaps surprisingly) performed approximately the same in all the learning assessments as the residential students.

Similarly, in spring and summer 2017, Georgia Tech conducted a preliminary experiment for teaching a freshman-level course, CS 1301 Introduction to Computing Using Python, online using a personalized learning format that included mastery learning and an adaptive textbook on the McGraw Hill SmartBook platform (Joyner 2017). Initial assessments show that students' performance in the online version of the course was statistically the same as an on-campus traditional version of the course, yet the students in the online/personalized learning section reported that they spent less time per week on the course and that they liked it better than did the students in the traditional section of the course (Guzdial and Ericson 2015; Madden et al., forthcoming).

Based on the successes of these preliminary experiments, the Commission recommends pilot projects on appropriate adaptive learning platforms that could be customized by faculty who would insert the topical content. Some of these experiments may include interactive books and interactive videos, as well as AI agents like "Jill Watson" for many Georgia Tech classes, especially large, remedial, and/or online classes. Some of these adaptive learning platforms can also be transferred to K-12 education as well as many graduate classes (both online and in-person).

Pedagogical experiments might be considered that examine where and how this personalized learning is effective. Besides being an integral part of a course, personalized learning modules can be used to support students of varying backgrounds and abilities, or to streamline a curriculum. For example, modules of adaptive learning coupled with competency tests could be developed and used at the beginning of a course so that students without the necessary experience might be able to obtain the essential skills required for the course. Such a practice would likely be advantageous for transfer students or for graduate students coming from other schools and having varying academic preparation. In a similar manner, personalized learning modules coupled with mastery-learning competency tests could be used to replace some course prerequisites in order to streamline a path through the curricula.

Outcomes in the Medium Term (2–5 years): Personalized Learning and Multifunctional Tutors

To keep ahead of the fast-moving envelope of innovation, The Commission proposes that Georgia Tech boldly develop a multifunctional virtual tutor that can perform many (not just one) of the typical functions of human tutors. While there has been considerable previous work on tutoring systems, there has been little work on a multifunctional virtual tutor that can perform cognitive tutoring on course concepts, question answering, question asking, evaluation of student progress through grading of assignments and examinations, or metacognitive tutoring on open-ended projects, etc.

Given recent advances in AI, learning and cognitive sciences, and data science and engineering, the time is ripe to build a multifunctional virtual tutor – a development that will put Georgia Tech at the forefront of AI use in education.

A multifunctional tutor will push the envelope on personalized learning in several ways. First, different students may need learning assistance of different kinds, and a given student may need different kinds of assistance at different points in a course. A multifunctional tutor will be able to personalize the different kinds of learning assistance to each student as needed. Second, a multifunctional tutor will be able to personalize learning assistance of the kinds that matter the most, such as formative assessment and metacognitive tutoring. Third, a given function will provide the context for other functions. For example, a "Jill Watson"–like AI agent operating as part of a virtual multifunctional tutor would be able to answer deeper questions about concepts taught by the cognitive tutors.

Outcomes in the Long Term (3–15 years): Personalized Learning and Human-Centered AI

In the long term, the development of a multifunctional virtual tutor fully capable of supporting personalized learning at scale would require advances in human-centered AI.

Human-centered AI refers to the development of interactive AI agents whose interactions with humans are informed by cognitive models of humans and the contexts of interactions with them. In this form of AI (also sometimes known as human-aware AI or context-aware AI), the AI agent builds, maintains, and uses models of humans with whom it is interacting and thereby enhances the quality of the human-AI interactions.

The focus of traditional research on AI is typically on building autonomous agents that can act independently of humans. This is also true of most research on machine learning. In contrast, research on human-centered AI focuses on developing interactive agents that can live, work, play, and learn with humans. Thus, human-centered AI is both adaptive—it adapts its interactions with humans depending on the context—and personalized—its interactions with humans are tailored to suit the backgrounds, profiles, goals, and needs of individual humans.

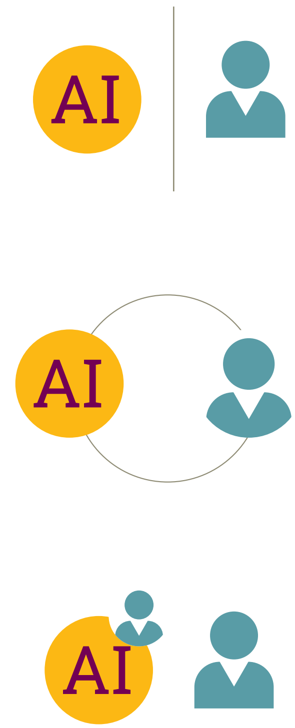

Figure 2: Research in AI has progressed from autonomous agents that act independently of humans (top of figure) to intelligent agents that interact with humans (middle of figure), to intelligent agents that have a cognitively informed model of the human with whom it is interacting (bottom side of figure).

Application to Specific Domains

While the Commission seeks general theories, architectures, representations, and methods for supporting adaptive and personalized learning, to make tangible progress we must focus on specific tasks in specific domains. Ideally, the domain of choice would be relevant to education at Georgia Tech as well as be relevant to the society and world at large. Equally important, the domain would be open ended. Potential domains may include environmental science, climate change, ecological modeling, and sustainable systems.

References

Bloom, Benjamin. 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.

Day, Jason A.., and James D. Foley. 2006. "Evaluating a Web Lecture Intervention in a Human–Computer Interaction Course." IEEE Transactions on Education 49 (4): 420–31.

Goel, Ashok, and David A. Joyner. 2017. "Using AI to Teach AI: Lessons from an Online AI Class." AI Magazine Summer: 48–58.

Goel, Ashok, and Lalith Polepeddi. 2017. Jill Watson: A virtual teaching assistant for online education. Presented to the Learning Engineering for Online Learning Workshop, Harvard University, June 2017; Georgia Tech Technical Report (http://hdl.handle.net/1853/59104); to appear as a chapter in Education at scale: Engineering online teaching and learning, edited by C. Dede, J. Richards, and B. Saxberg, New York: Routledge.

Guzdial, Mark, and Barbara Ericson. 2015. Introduction to Computing and Programming in Python. Fourth Edition. New Jersey: Prentice-Hall.

Joyner, David A. 2017. Introduction to Computing Using Python. New York: McGraw-Hill Education.

Madden, Amanda, Lauren Marguilieux, Rob Kadel, and Ashok Goel, (Editors.) Forthcoming. Blended Learning: A Guide to Researchers and Practitioners. Cambridge, MA: MIT Press.