The Asymmetric Relationship of Artificial Intelligence and Education

Jim Greer

ARIES Laboratory
Department of Computer Science
University of Saskatchewan
Jim.greer@usask.ca

Abstract: In this short paper I argue that the relationship of AI and Education could not have been construed to resemble a marriage. The intersection of the two disciplines is in fact fairly minimal and even within that intersection there is and always has been an asymmetric relationship between the disciplines. While many examples from the literature support such a conjecture, I offer a brief new case study of how the relationship continues bump along as we work to support classroom-based learning at scale.

1 Some areas of Education benefit from AI techniques and methodologies.

Traditionally education has been derived from a relationship between teacher and student. The learner is coaxed to construct knowledge resulting from planned curricula, planned communication, practice, feedback, and assessment – all provided or guided by the teacher. This works well when a teacher interacts with a handful of learners. But as we move toward increasing the ratio of learners to teachers and attempt to scale up or automate the teaching and learning environment, supports derived from technology become more essential. In fact we can claim that AI technologies become beneficial when a more subtle type of technology-based coaching or tutoring is needed. When human teacher supervision is unavailable or in short supply, the AI technologies may even become essential.

There is ample evidence in the AIEd literature that a dose of AI can help with automated instructional support. This technology-based support can be in the form of ITS style one-on-one tutoring or adaptive coaching. Any time it is less feasible to have intensive interaction with a human teacher, more inferential capability is required in the technology support system.

It is also becoming evident that "learning at scale" requires some AI technology. Managing the large amounts of data derived from fine-grained observation of myriad learners is beyond the capability of human teachers. Mining the data, finding relevant actionable patterns, and even delivering guiding advice at scale can be accomplished only with the support of AI tools and techniques. Below I relate a case study to support this claim.

Yet much of the educational activity we continue to see in school, college and university classrooms continues to operate as "AI-optional". Some non-AI technologies, such as new media, online resources, or learning management systems, offer conveniences to classroom teachers and learners. But we are still far from a scenario in which AI is perceived as beneficial, let alone necessary to all of education. So even after more than 25 years, AIEd remains at best near the periphery of educational practice, and at worst a niche sideline area of educational practice of interest to only a few educational practitioners.

2 Education is an application area potentially fruitful to AI.

It is quite clear that problems in the domain of education tend to challenge and stretch AI techniques. The nature of human learning itself, whether building up inductive knowledge, coping with incompleteness and inconsistency, shedding misconceptions, reaching impasses, or remembering and forgetting, offer wonderful challenges to AI reasoning systems. AI work in areas such as knowledge representation, ontologies and semantics, user modeling, natural language processing, and planning has been inspired in part by educational problems. But it also seems that education problems are some of the harder problems for AI to solve. The social nature of education offers a degree of fuzziness and imprecision that encourages many AI researchers to turn to other application domains. Furthermore, despite the fact that huge amounts of money are devoted to education each year, it mainly goes to teachers and schools and no significant proportion is likely to find its way into the pockets of AI innovators and researchers.

Many have postulated that education is the perfect application area for AI. Education offers a relatively low-stakes testbed for experimentation with AI techniques, whether they be modeling and recognition, adaptive decision making and diagnosis, planning and interacting, or natural language processing. It is low stakes because errors in recognition or decision-making are unlikely to have serious adverse effects on the learner or his/her learning. So AI techniques can be safely tested in educational environments without much risk.

At the same time, there continues to be the 2-sigma promise that AI techniques can dramatically improve learning gains and so AI could one day prove to be a necessary and sufficient ingredient in learner success. Unfortunately we have so far failed to show that it is the AI in our technological systems that drives the learning gains. About 20 years ago Elliot Soloway proclaimed that AIEd sought to address the "level 16 problems in education" and that there was "so much simpler technology that could be used to address the level 1, 2 and 3 problems" and we should focus there first. It was with this proclamation that Elliot left AIEd to work on what he considered to be more urgent and relevant problems in technology-based learning.

3 AI and Education do not fundamentally need one another

The story so far claims that most of the world of Education can operate without AI and that most of the world of AI can move forward without Education. This would indicate that there never was a true marriage of the two disciplines. They are not and never have been inextricably intertwined. Education, though an interesting area of application for AI, has not been well funded or afforded much credibility by mainstream AI researchers. Similarly AI, though a potentially valuable resource to Education, has not been well funded or afforded much credibility by mainstream Education researchers.

Beyond this, another interesting observation about the disciplines of AI and Education can be made. Both share a sort of inferiority complex. Education is often perceived to be the poor cousin of the social sciences. Education researchers seem to feel that they are near the bottom of the pecking order of the social sciences. Interestingly, AI researchers share that view regarding a lowly place in the pecking order of the sciences. Between the 1970s and 2000, AI moved from an over-hyped savior for humanity to colossal failure. Now as it slowly seeks to gain favour with much more modest aims, its credibility may be improving a bit. But if considering a serious partnership, this would be like the marriage of one poor cousin with another.

Of course there are a few people, like us, who over the years have seen the benefits of a closer connection between AI and Education and who have sought to support a permanent relationship through the field of AI in Education. But we do realize that the intersection of AI and Ed is relatively small. We represent a small sliver of AI and a small sliver of Education who care enough about the other to try to maintain an equitable and mutually beneficial relationship.

4 The asymmetry of AI and Education in AIEd

I think it was John Self who promoted the term "AI in Education" as the moniker for AIEd. The word "in" represents an interesting juxtaposition of terms where Education is the goal or purpose and AI is brought to the world of Education as an enabler to presumably create something new and better. While I will resist making sexual innuendos and avoid talk of consensual relations, it is fairly easy to see the asymmetry of the relationship. Existing AI techniques and tools are brought to bear on Education in order to enrich Education in some way. For the most part, this type of relationship has less emphasis on trying to change, influence or evolve AI, but rather to make use of its available tools and techniques to advance Educational goals and even potentially revolutionize Educational practice.

I suppose that insertion of AI into Education was seen to some defenders of the status quo as an attempt to supplant the role of the teacher in education. And this mission was destined to fail from the outset. Yet AI has been insinuating itself into educational systems steadily and pervasively and sometimes quietly over the years. Adaptive exercises, workbooks and e-textbooks are at the center of a new business model for the textbook publishers of old. Data-driven approaches to early-alert systems are raising flags for teachers, academic advisors, and learners in time to avert failure.

5 The case of Sara

For the past year I have been leading a project to enable classroom instructors to provide personalized, weekly advice to each individual student in their large courses. We have been working with freshman STEM courses where class section sizes are between 300 and 600 students. A system called the Student Advice Recommender Agent (SARA) was developed and first piloted in 2014 in the three sections of a 1200-student freshman Biology course. (Greer et al, 2015)

Demographic data from the student information system was linked to data from a 75-item entry census completed by most students prior to their arrival on campus. The entry census includes questions about goals and expectations, hours working off campus, family and living arrangements, financial sources, as well as some standardized subscales to measure motivation, approaches to studying, and grit. Using academic records and this rich demographic data a predictive model was developed for expected student success overall and in individual courses. Based on success predictions and other demographic variables, advice templates were constructed using stereotypical student personas and advice was personalized and delivered by our agent, SARA. As the course progressed, student activity (in the LMS) and achievement in labs and assignments was factored into the predictions for each learner and SARA's advice was adjusted accordingly. Advice focused on directing students toward extra help, academic advising, online resources, reflective activities, peer-led study groups, clubs, and sometimes meetings with the instructor. Some effort was made to direct students toward the resources that they would most likely find useful and to ensure that those directed toward the more expensive one-on-one sessions with a human advisor would be those who could benefit most from such encounters.

At the end of the one-semester, course raw grades were compared against the corresponding cohorts from the prior year and we found 23.3% fewer D and F grades and 28.3% more A grades. The mean raw grade in 2014 was significantly higher (2.6 percentage points) than in 2013. Six items from the 2013 final exam, having good point-biserial correlations with the overall score, were repeated on the 2014 final exam. Instructors were blinded as to which items would be repeated. Significant improvement year over year was also detected on these six items (about 3 percentage points). The evaluation rubric and the instructional team had not changed significantly in 3 years, but there was one instructional change. In 2014 students were given access to the OLI Biology system and required to complete a weekly online quiz in OLI. Instructors felt the weekly quizzes contributed to the overall improvement in grades. Analysis continues to try to determine whether the advice from SARA made a difference or not.

Since the predicted model of student success in 2014 was developed from the 2013 student demographics and achievement (where weekly quizzes and SARA's advice did not exist), we felt that the difference between predicted grade and actual grade might be a better measure of the impact of our interventions in 2014. Mean predicted grade in 2014 (based on the 2013 model) was significantly lower than the actual 2014 grade (by about the same amount as the grade increase from 2013 to 2014). Since all students were required to do the weekly quizzes, a constant increase across the board could be attributed to the quizzes. To dig a little deeper, students were asked on an end of term survey whether or not they paid attention to the advice from SARA. Those who claimed to pay attention to SARA received higher grades than those who did not. Of course this could be simply selection bias – better students may be more likely to pay heed to advice than poorer students. When considering actual minus predicted grade, no difference was observed between those who said they ignored SARA versus those who considered SARA's advice. Taken together, we cannot prove that SARA helped increase student achievement levels, but it might have and surely something happened to improve achievement. Surely, additional study is surely warranted.

This little case study shows how fairly traditional classroom learning can potentially benefit from a small dose of AI. The machine learning and predictive modeling, adaptive personalization of advice, and the personalized delivery of advice through an advice agent made use of modest and fairly simple AI techniques. The results led to a scenario where instructors and students were delighted with the outcomes whether or not the AI made a difference. Further, it seems they did not feel overly threatened by the intrusion of technology.

6 Conclusions

This paper has attempted to show that the relationship between AI and Education could hardly be construed as ever having been a healthy marriage. Like AI in Manufacturing or AI in Medicine, or AI in Law, the story here is bringing tools and techniques to further the capabilities of practitioners in a discipline. The success of AI in discipline "X" is determined largely by those who regulate, practice, or finance the activity of discipline "X". The recipients of discipline X (be they learners, consumers, patients, or clients) need not care or even know that AI has been the reason for a better performing system.

Going forward, I believe that bringing AI into Education will never be fully embraced with open arms by the patriarchs and protectors of general education. But those who understand the benefits that AI can bring to Education should not be discouraged or deterred. For the good of future learners we must continue to work to bring old and new AI tools and techniques to bear on problem areas within the educational landscape.

References

Greer, J., Frost, S., Thompson, C., Banow, R., Kuleza, S. and Koehn, G., (2015) The Student Advice Recommender Agent (Sara). Workshop on Personalization Approaches in Learning Environments, UMAP, 2015, 5pp., submitted.