Adaptive Learning: The Premise, Promise, and Pitfalls

Mr. Petr Johanes, Stanford University

Petr Johanes is currently a PhD student in Learning Sciences and Technology Design (LSTD) at the Stanford University Graduate School of Education. He holds a B.S. and M.S. from the Department of Materials Science at Stanford University and has experience teaching engineering courses as well as researching engineering education, especially in the context of online learning. Right now, Petr is investigating the role of epistemic cognition in learning.

Dr. Larry Lagerstrom, Stanford University

Larry Lagerstrom is an Associate Dean and Director of Summer Session at Stanford University (and previously was the Director of Academic Programs at the Stanford Center for Professional Development). Before coming to Stanford he taught computer programming and electrical engineering for sixteen years at U.C. Berkeley and U.C. Davis. He has degrees in physics, math, history, and interdisciplinary studies, including a PhD in the history of science and technology. He also has developed a MOOC on "Understanding Einstein: The Special Theory of Relativity."

Adaptive Learning: The Premise, Promise, and Pitfalls

Introduction

In a 2015 speech before the American Council of Education, John Hennessy, Professor of Engineering and President at Stanford University, laid out a vision for how new technological tools and pedagogical methods can improve higher education. He especially highlighted the opportunity to revitalize courses by crafting online and hybrid learning materials that adapt their speed, depth, and approach to the individual student. [21] Others have made the same point. The National Academy of Engineering, for example, has listed "personalized and adaptive learning" as one of its Grand Challenges, and a Learning Analytics Workgroup, composed of thirty-seven representatives from universities, foundations, government entities, non-profit organizations, and for-profit companies, has put forth an "endgame" vision of "personalized cyber learning at scale for everyone on the planet for any knowledge domain." [34] Given that companies such as Knewton, Acrobatiq, Coursera, and Udacity are either commercializing or implementing adaptive learning technology, and online higher education institutions such as Western Governors University are building it into their courses, it is likely in the near future that engineering schools and faculty will face questions about their use of this and similar technologies that enhance learning. These questions may come from students and parents, of course, but also from the media and perhaps even accreditors.

In this review paper, we aim to provide guidance to engineering education leaders and engineering faculty via three main goals. First, to explain what adaptive systems are and what kinds of data they require. Second, to categorize the main use cases and possibilities of adaptive systems. Third, to outline the current limitations and concerns surrounding adaptive systems. Engineering leaders and instructors can then determine if their pedagogical context is amenable to deploying these systems, and education researchers can navigate the current systems' characteristics to find areas where to make impactful contributions.

The Premise of Adaptive Learning

Changing what a single learner – as opposed to a whole class – experiences while learning is neither a new goal for educators, nor an unrealized method. After all, a teacher providing student-specific feedback or a teaching assistant helping each student differently during a laboratory session is not science fiction. What has been science fiction up until recently is having the ubiquitous computational power, large user communities, and scalable analytical algorithms to change what a single learner experiences based on learner input rather than teacher intuition, and to do so at scale. In this section, we present a short history of where we can situate adaptation in the breadth of education research and also what modern adaptation in education broadly requires.

A History of Adaptation in Education

If personalization of education is the endgame, then adaptation of education is currently considered the winning strategy. One historical root arises from the advent of the personal computer in the 1980s and the obvious possibility of using the computer as an automated form of tutor, or as an “intelligent tutoring system” (ITS). [42] An ITS is “any computer system that performs teaching or tutoring functions (e.g., selecting assignments, asking questions, giving hints, evaluating responses, providing feedback, prompting reflection, providing comments that boost student interest) and adapts or personalizes those functions by modeling students’ cognitive, motivational or emotional states.” [31] As might be expected, STEM topics – and computer science in particular – proved well-suited to these modeling efforts. Not only were computer scientists the ones designing the computers in the first place, but they were also operating in a knowledge domain that lends itself nicely to a clear and computer-understandable separation of declarative knowledge (e.g., knowledge of concepts and facts) and procedural knowledge (e.g., knowledge of methods and approaches). [13,15] Unlike in some other fields of learning, teaching computer programming without a computer is difficult and the point of computer programming is that the computer needs to be able to understand the code. Using the computer as a medium for learning as well as a medium for adaptation was a logical and natural union.

Another historical root also happens to be the most recent, thanks to the parallel rise of massively open online courses (MOOCs) in education and of machine learning methods in data science. Online environments that attracted thousands of students to read material, attempt assessments, and watch videos venture beyond traditional ITS in terms of their activity variety and learning scope. Many of the early MOOCs were also facilitated by startups, and so the need for data collection and adaptation took on a more commercial rather than academic level of urgency. Part of the interest in adaptation here stems from a well-documented dropout problem [20,26] due to different users entering MOOCs with, for instance, different motivations while also displaying differences in video watching [18,27] and course navigation [19,25], among many other patterns. Analyzing all of these clickstream data quickly for the purpose of increasing learner engagement, performance, and persistence is a fundamental challenge for online learning platforms and providers. To help with this task, researchers are turning to machine learning methods, including variations on methods from ITS [8], such as Bayesian Networks (BNs). Experimenting with newer analytical methods, such as Recurrent Neural Networks (RNNs), and benchmarking their performance in describing and predicting learning in these environments is set to become a logical evolution of research into online learning. (See [9] for an introduction to BNs and [35] for an introduction to RNNs.)

The Science behind Adaptation in Education

Brusilovsky and Millán review a variety of research arising from this history, especially that related to the use of knowledge and user modeling in the fields of adaptive hypermedia systems and adaptive educational systems. [9] The adaptiveness of these systems stems from their ability to change what each individual user experiences in the system based on information the system has collected and processed. The information that the system uses to make the decision on what and how to adapt – the adaptation model – comes from two sources: the domain model (e.g., different knowledge domains such as mathematics, grammar; different knowledge types such as declarative, procedural) and the user model (e.g., user's knowledge in relation to the domain knowledge model; user's characteristics and traits in relation to user behavior). In adaptive educational systems, since the domain being modeled is a knowledge domain, the model is often called the knowledge model, and, since the users are students, the user model is often called the student model. The critical considerations for any creator or consumer of adaptive systems, then, are: what is being modeled; how it is modeled; and how the models are maintained. While various methods for modeling exist, the most used method today is overlay modeling. The core principle behind overlay modeling is simply that there is some underlying model of a domain and that the model of some user is a subset of that domain model. The objective of adaptive systems operating in this paradigm is two-fold: (1) to adapt what the user encounters such that eventually the user's overlaying subset matches the system's underlying whole and (2) to adapt the system's underlying whole to become a more accurate representation of the domain. [9]

In this conception of modeling, we engage in what is called uncertainty-modeling, which effectively states that we cannot know with 100% certainty what is being modeled. A very simple example of this is that if a learner answers correctly on a multiple-choice assessment item, there is still a possibility that the learner guessed the answer and the model therefore cannot state with 100% certainty that the learner knows the knowledge evidenced by that assessment item. The corollary to this example is that even if the model has accumulated enough evidence to be satisfied that the learner knows some knowledge, the probability of the learner knowing that knowledge might be 0.99 instead of 1.00 to account for forgetting or slipping. [4,15] Because of the admission of uncertainty, most models are probabilistic in nature and the most popular probabilistic model is a Bayesian Network. A Bayesian Network uses a knowledge model that consists of individual knowledge elements (e.g., concepts, procedures, rules) connected to evidential elements (e.g., responses, behaviors, correctness). Each connection has a certain weight, indicating how strongly the evidence correlates to knowing the knowledge and how strongly knowing the knowledge is needed to elicit the evidence. As the user interacts with the system, the system can update the user's knowledge state (the probability, based on user input until that time, that the user knows certain knowledge elements) after each interaction to suggest what activities would serve as further evidence of the user being in a particular knowledge state. This method of updating the probabilities of a learner being in a certain learned state after a set of exercises, activities, or assessments (rather than after the whole curriculum) is referred to as knowledge tracing. [15]

Beyond modeling user knowledge states, adaptive systems designers are also interested in, for instance, modeling user goals, interests, problem-solving strategies, affective states, and social psychological states. [9] Of course, there is a trade-off between a model's conceptual completeness and a model's computational efficiency, but the ultimate hope is that one day these systems will be able to process a multitude of complex interactive behaviors about any given user and adapt the user's experience in a precisely personalized manner.

The (Potential) Promise of Adaptive Learning

More personalized instruction would be especially beneficial for large lecture classes, where the amount of student-instructor interaction is often limited. In an oft-cited study, “The 2 sigma problem: The search for methods of group instruction as effective as one-to-one tutoring,” Benjamin Bloom (of Bloom’s taxonomy fame) compared learning outcomes between tutoring and regular classroom instruction and found an effect size of 2.0 in favor of tutoring. That is, the learning attained by the average tutored student was two standard deviations above that of the average student in the regular classroom (or, to put it another way, the average tutored student would have been in the 98th percentile among the classroom students). [5] More recently, a meta-analysis by VanLehn has questioned whether the difference is closer to one standard deviation than two. [40] But the point remains: is it possible to replicate – or actually surpass – the effectiveness of high-quality tutoring and provide its benefits to greater numbers of students?

The Landscape of Adaptive Learning

With both the academic and popular media providing lists of disruptive adaptive learning companies and in-depth profiles of adaptive learning in public school districts, there is now a mainstream effort to find out. [12,38,41] And there is a corresponding effort to verify if adaptive learning can be lucrative, as a number of recently founded companies hope: Acrobatiq (2013), Knewton (2008), CogBooks (2006), Cerego (2000), Realizeit (2007), LoudCloud (2010), and Smart Sparrow (2010), to name a few. Services offered by these companies include full off-the-shelf courses, supplemental materials for popular subjects (especially introductory and developmental topics in languages and STEM fields), authoring platforms for instructors and course designers, and assistance to universities or departments in developing a curriculum based on adaptive learning. At least three reviews provide a useful survey of dozens of such service providers. [7,10,43] Educational publishers have entered the market by either partnering with an adaptive learning company (e.g., Pearson with Knewton) or buying one (e.g., McGraw-Hill and ALEKS, Barnes & Noble and LoudCloud). And most providers of learning management systems—such as Blackboard, Canvas, D2L, and Moodle—are either working on their own adaptive learning components or offering connectivity to third-party ones, or both. In sum, entrepreneurs, publishers, and back-end systems providers are all betting their revenue and their clients' satisfaction on adaptive learning.

Universities are not far behind. In fact, several of the companies on the list above spun out of university implementation programs and research groups. For instance, Acrobatiq emerged in 2013 thanks to the Open Learning Initiative (OLI), started at Carnegie Mellon University in 2002, and Smart Sparrow emerged in 2010 thanks to work started in 2007 in the engineering school at the University of New South Wales. Furthermore, many of the online learning platforms currently developing MOOCs and increasingly deploying adaptive learning, such as Coursera, Udacity, and Open EdEx, also emerged from universities, as have usually the algorithms these companies are now using. For example, Montana State University is in the middle of a multi-year project to introduce adaptive learning into its digital logic courses. [28] Universities have also invested in developing predictive student models for use in early warning systems for at-risk students. These models assist in the design and implementation of adaptive interventions that can reduce failure and drop-out rates. The Course Signals system at Purdue University is perhaps the best known. But similar systems are in use or in development at the University of Michigan, the University of Alabama, Northern Arizona University, Georgia State, Delaware State University, and the University of Phoenix, among others. [3,17,29,32]

The Possibilities of Adaptive Learning

However, one of the byproducts of such a heavy commercial presence in the adaptive learning space is that many of the algorithms, results, and data remain proprietary until a given company makes the choice to release them via a research paper, or more often a white paper or press release. Unless the industry and academy find a way to rectify this issue, it will likely impede adaptive learning development. At the same time, thanks to several collaborations between industry and academia, we can identify various learning results and successful implementations. To get a sense of what these collaborations are pursuing, consider the meta-analysis by VanLehn, which concluded that intelligent tutoring systems are "just as effective as adult, one-on-one human tutoring for increasing learning gains in STEM topics" (given certain caveats concerning the ITS design and use). In particular, VanLehn found an effect size of 0.79 for human 1-1 tutoring as compared to no tutoring, and an almost identical effect size of 0.76 for ITS-based tutoring. [40] The list below shows some corporate-academic partnerships in adaptive learning and their major findings:

- **Time-to-learn reduction**
    - OLI and Carnegie Mellon University. Studies showed, for example, that OLI students were able to complete the OLI statistics course "in *half the time with half the number of in-person course meetings*," while showing "*significantly greater learning gains* on the national standard 'CAOS' test for statistics knowledge and similar exam scores.... There was no significant difference between OLI and traditional students in follow-up measures given 1+ semesters later." [30] A follow-up "trial compared a hybrid version of the OLI statistics course with a traditional face-to-face statistics course with randomly assigned students at six [public] institutions. Students in the hybrid format had *comparable or better learning gains and took 25% less time to learn the same outcomes.*" [6]
- **Closing achievement and engagement gaps**
    - OLI and Carnegie Mellon University. A study of the OLI psychology, anatomy & physiology, biology, and statistics courses in community college settings found that "faculty use of and experience with the OLI course was associated with *higher student achievement gains and may help smooth out expected negative outcomes associated with race.*" [24]
    - Realizeit and University of Central Florida. Realizeit developed an adaptive learning system for psychology, nursing, and algebra courses at the University of Central Florida. The results showed a moderate increase in performance and high student satisfaction—83% reported that the system helped them learn better. [22]
- **Increasing passing rates**
    - Smart Sparrow and Australian Universities. Adaptive tutorials were developed to assist in the teaching of introductory mechanics. Failure rates dropped from 31% to 19% in the first year of use and, as the curriculum was tuned, continued to decrease over the next two years to under 10%. They also observed an improvement in performance by students who subsequently took third-year mechanical engineering courses. The tutorials are now used in the teaching of introductory mechanics at several universities in Australia. [37]
    - Realizeit and Colorado Technical University. Colorado Technical University, using a system developed by Realizeit, was reportedly able to improve pass and retention rates for introductory-level online courses by around 5-10%, or better. In one algebra course the failure rate decreased from 30% to 18%. They also used the system in blended courses in trigonometry and pre-calculus, dividing the student work 50/50 between online and in-class activities. The trigonometry pass

rate increased from 76% to 98% and the pre-calculus pass rate from 66% to 98%. One of the keys to their success was to take a cautious development approach, first implementing the system in three courses the first year, then sixteen the next, twenty-five in the third year, and sixty-three in the fourth year. [16]

- Knewton and Arizona State University. Arizona State University has used Knewton's adaptive learning technology in introductory math classes. Pass rates increased by 18% and withdrawals decreased by 47%. [11]

Again, the list above is purely illustrative in order to show the kinds of learning benefits these systems can provide when universities and companies collaborate. Thus far all of them are in line with the results and potential outlined in Corbett and Anderson's seminal knowledge tracing papers. [2,15] Without going into detail on these and many other studies, one can begin to understand the impact that these adaptive learning systems can have on the educational landscape, some major points of which we summarize in **Figure 1**.

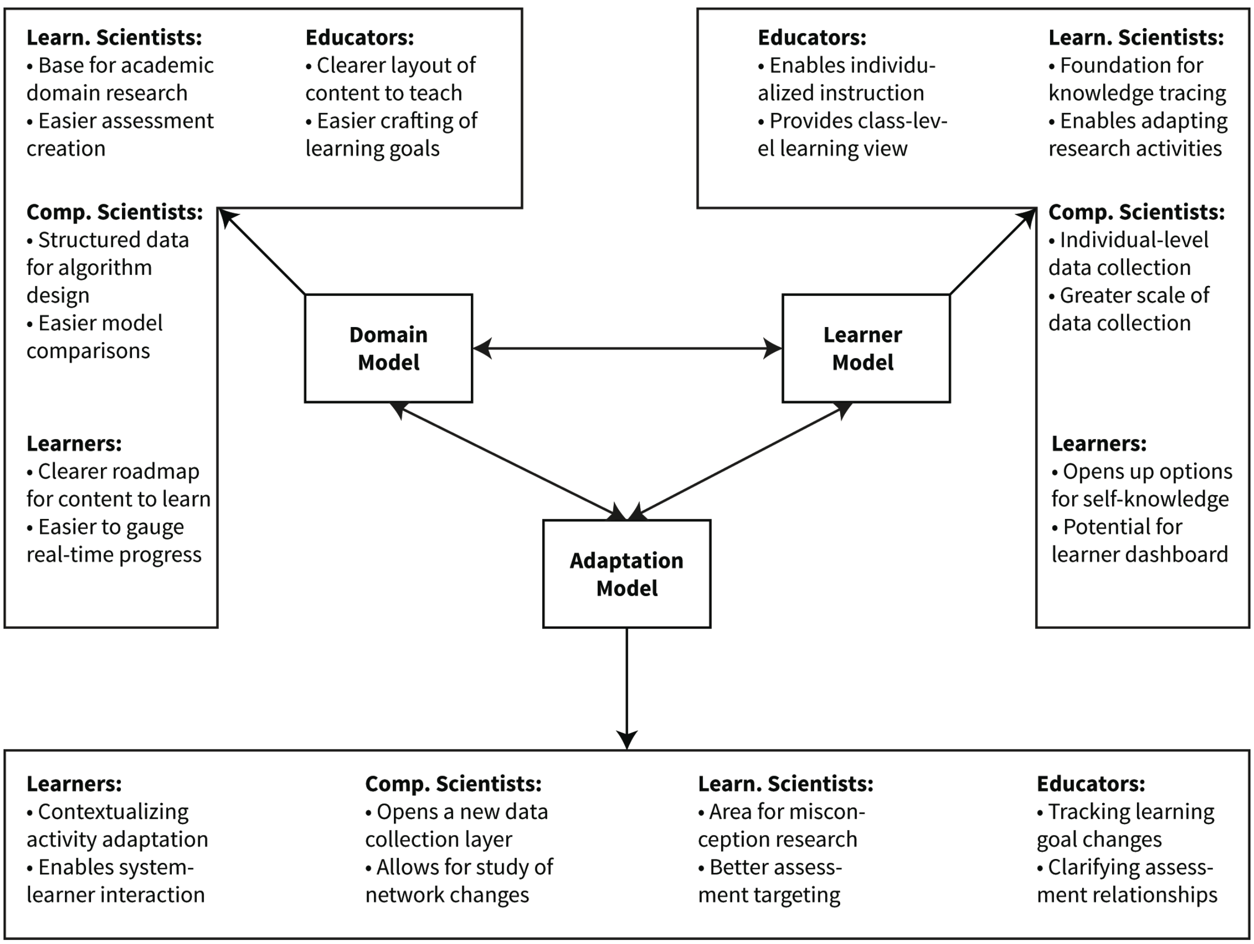

Figure 1. Organizing the potential benefits of adaptive educational systems to four stakeholder groups: learners; educators; learning scientists; computer scientists.

The (Potential) Pitfalls of Adaptive Learning

Given these empirical results and general intuitions about the promise of adaptive learning for education, both in the classroom and in the system, it could be easy to fall prey to the hype. It is therefore important not to overlook the potential and actual pitfalls, limitations, and/or concerns surrounding adaptive learning.

A Brake on the Hype: The ALMAP Study

The positive results and press for adaptive learning prompted the Bill and Melinda Gates Foundation in 2013 to invite proposals for a new Adaptive Learning Market Acceleration Program (ALMAP), with the goal of expanding its use. [45] Fourteen higher education institutions were part of the successful applications, representing over twenty courses and a range of adaptive learning products. Subjects included basic math, algebra, English language arts, business, marketing, economics, psychology, and biology. Vendors or courseware suppliers included Pearson/Knewton, Cerego, Smart Sparrow, CogBooks, Adapt Courseware, Learn Smart/Connect, Assessment and Learning in Knowledge Spaces (ALEKS), and the Open Learning Initiative. The courses implemented a variety of pedagogical changes. Over 23,000 students participated in the courses and corresponding studies from 2013 to 2015, of which approximately 10,000 used the adaptive courseware (the others were in control or comparison courses). Pell Grant (low-income) students represented 40% of the total. SRI International was hired to evaluate the various implementations and report on the "learning impact, cost and satisfaction findings." [45] The results were decidedly mixed [44]:

- **Course Performance.** Four of fifteen courses showed a slight improvement in course grades, but most had "no discernible effect." (For any given measure, not all the courses had enough data or an appropriate study design to come to statistically significant conclusions.) On the other hand, there were seven instances where "side-by-side comparisons of scores on common learning assessments" were possible, and the results were "modest but significantly positive." The most positive impact was seen in those cases that transitioned from a traditional lecture to an adaptive learning format.
- **Course Completion.** In terms of course completion rates, the adaptive courseware had no measureable effect. Only two of sixteen cases showed an increase. Some previous studies had indicated that low-income students tended to perform less well in online learning and/or blended learning environments. [45] In these ALMAP studies, however, their performance was equal to other students.
- **Instructor Satisfaction.** Instructor perceptions and satisfaction varied. Overall, 74% of instructors reported they were satisfied with the project, and they especially appreciated the real-time dashboards that tracked student progress. There was a divide, however, between those teaching developmental (remedial) courses and general education gateway courses. While 67% of the developmental course instructors planned to use the adaptive courseware in the future, only 49% of the gateway course instructors did. A major concern in all cases was simply getting students to use the adaptive courseware enough to be of benefit. If it is just seen as a fancy online textbook or additional materials and exercises rather than as a core part of the course, students are likely to treat it as they do many textbooks—with indifference.

- **Student Satisfaction.** Student perceptions and satisfaction varied by the type of student and course. Students in two-year college programs favored the adaptive courseware more than students in four-year programs: 77% of the two-year students indicated it helped their learning while only 51% for the four-year college students did so. Furthermore, 56% of the two-year students reported an overall satisfaction with it versus 33% of the four-year students. Over 90% of students in developmental course recognized an improvement in their learning, and 60% of them felt that they were more engaged in the adaptive course material. But only 25% of students in gateway courses felt more engaged and only 33% reported a positive learning effect.

As might be expected, implementation costs were higher at the start, mainly due to increased work required by the instructor. But once past the initial course run, seven of ten courses that had the appropriate data showed lower ongoing costs.

The most important conclusion from the ALMAP studies is that the nature and quality of the implementation is crucial. The same adaptive learning product may yield more or less positive results in different implementations. It is highly recommended that "institutions planning large-scale adoptions of adaptive courseware should conduct their own internal analyses of student outcomes with that courseware compared to other alternatives." [44] Attempts should be made to measure results on a more precise level than just course grades and completion rates, and attention must be paid to baseline equivalence when comparing results across students and courses: "multiple factors affect learning outcomes and to make sense of student outcomes, analyses need to incorporate student characteristics [including prior knowledge and skills], specifics of how the adaptive courseware is used, aspects of the course beyond the courseware product, and the way learning is measured to make sense of student outcomes." [44] The rationale and expectations for adaptive courseware should also be explicitly addressed with students in order to create an effective learning environment that motivates students to engage with the course material.

The Pitfalls of Adaptive Learning

As the ALMAP study indicates, not all might be well in the state of adaptation. Other pitfalls and concerns also arise, ranging from the technological to the philosophical, and below we choose to highlight three:

- **Discrimination and labeling of students, creating consequential feedback loops.** Let us assume that an algorithm classifies students based on their performance as low, medium, or high achievers in a particular subject. What happens based on that classification? Are students locked into certain learning trajectories based on it? If so, then who deals with edge or outlier cases? An algorithm is not racially, politically, or otherwise inherently neutral because it is designed by humans who are both most likely biased and most likely unaware of their bias. While this might seem to be a hypothetical or abstract point, it is not, and Cathy O'Neil documents in her book *Weapons of Math Destruction* the various ways in which algorithmic design is eating away at financial and political equality. [33] O'Neil sums up the core question as "whether we've eliminated human bias or simply camouflaged it with technology" [33:23]. The camouflage is not just a turn of phrase because the kinds of harmful algorithms she discusses are “by design, inscrutable black boxes” [33:28]. She devotes a whole chapter to the impact of the *U.S. News* college ranking system and how its use of known proxies creates a system prone to gaming. This possibility for gaming then lay the foundation for a consulting industry that further privileges students with a high socioeconomic status (usually white, urban, upper-class families). Therefore, even if the model itself might not be the main discriminatory tool, it can set clear conditions for discriminatory practices. This is a salient point because, as O’Neil illustrates, one of the reasons we wish to create models is so that we can create real-world feedback loops based on the models’ measurements. Some forms of discrimination might be overt, but not all are, and unless we are vigilant, we might be engaged in feedback loops in which learning might become more efficient in terms of time and money, but not in terms of civic principles and integrity. [33]

- **Narrow constraints of knowledge, knowing, and learning.** ITS and adaptive learning systems can be considered knowledge-based systems (KBS) in that they model a system of knowledge along with a student’s state of knowledge and then recommend a knowledge-creation trajectory. [13,15] These kinds of tasks fall into the jurisdiction of the study of knowledge and knowing, or epistemology. Most of these systems to date, however, follow an epistemological view based on the Adaptive Character of Thought (ACT-R) theory. [1] The theory conceives of human cognition as consisting of encoding objects in the environment into knowledge units (via chunks) on one hand and encoding transformations on the environment (via production rules) on the other. This effectively means that most of the systems we use today hold a view of knowledge as existing in two forms: declarative and procedural. And the systems see all relevant knowledge as not only being model-able, but also being explicit. [1] These are not the only epistemological models available, and a long-standing argument for tacit knowledge being necessary in learning demonstrates that we cannot take this aspect of these systems lightly. [14,36] What if, by over-applying these systems beyond their use cases, we actually begin constraining various domains of knowledge to the epistemologies of these systems? What if, by not investigating alternative epistemological constructions, we marry the systems to a conception of knowledge and knowing that might be counterproductive in the long-term in a discipline or across disciplines? It is possible that by narrowing the epistemological underpinnings, the systems are missing out on certain education and research upsides. [23]
- **Transparency, availability, and security of data.** A fundamental question arises with the use of adaptive learning systems: who owns the data? Data generated by student input can not only be used to grade, assess, and certify, but can now also be breached, searched, and sold. How secure are the data? How available are they (or should they be) to students, instructors, and universities? More importantly, online learning providers might be collecting data that the users are not even aware of, such as keystroke cadence (which could be used for user identification even outside of the learning context). While this might seem harmless, as online learning platforms begin deploying more varied and more complex data collection schemes around affect and social psychology, for instance, the platforms might be able to infer more about us than we might be comfortable with.

We do not mean to imply that these pitfalls are inherent, irreversible, or even fatal to the success of adaptive learning systems. They are three prominent areas of concern to be vigilant of because even the systems with the best of intentions for our students may have unintended consequences.

Through them, we wish to encourage an informed and inclusive engineering philosophy rather than an unaware and uninviting one.

Conclusions

In this paper, our aim is to provide a brief introduction to the what, why, why not, and how of adaptive learning, so that individual educators and administrators can better navigate any potential choice to consider, implement, and evaluate an adaptive learning system. In **Table 1**, we summarize some of the major promises and pitfalls currently apparent in adaptive learning based on empirical evidence as well as theoretical intuition. Based on the proliferation of adaptive learning in the corporate and academic worlds, it is certain that they are here to stay. The question for most educators, therefore, is not whether we are willing to adopt them, but how willing we are to adapt them. As we point out in the section on pitfalls, any limitations and weaknesses of adaptive learning systems do not exist because they are essential, but because they are engineered, intentionally or not. The more we as educators and researchers begin to work with – rather than against – the design of these systems, the more we can do to influence and inform their evolution. They are neither a panacea, nor a plague, but they can be powerful. We currently have a window of opportunity to make them powerful for pedagogical purposes first. And higher education institutions are uniquely positioned to do so by engaging in collaborations between educators, learning scientists, and computer scientists, all of whom (at major institutions at least) tend to reside on the same campus. If we take charge of consciously creating learning environments that then employ learning systems to facilitate learning for students, we can redefine the practice of teaching, researching, and learning. Adaptive learning systems can serve as a powerful catalyst for exactly these kinds of impactful collaborations.

Promises	**Pitfalls**
Clarification Promise. Clarify the underlying content, skills, and dispositions needed to master a certain domain	*Epistemological Pitfall.* Limiting instructor, learner, and researcher conceptions of knowledge and knowing
Personalization Promise. Find personalized paths through the learning process for each and every student	*Ownership/Security Pitfall.* Mishandling learner data legally, ethically, and economically (intentionally or not)
Optimization Promise. Increase learning gains while reducing the time durations needed to achieve them	*Development Pitfall.* Creating an adaptive learning system can bankrupt an institution due to high cost in expertise, time, and capital
Equalization Promise. Learners from all backgrounds can receive the education they want in a manner they need	*Discrimination Pitfall.* Biased, opaque, and inscrutable models can discriminate against certain learners via labeling
Instructional Promise. Teachers can be empowered and better supported in facilitating high-quality learning	*Learning Pitfall*. The models focus on a particular interpretation of learning that can be neglectful of social and physical learning
Research Promise. The scale and nature of the collected data open up new research avenues in data as well as learning science	*Deluge Pitfall.* While the systems are collecting a lot of data, much of that data might not be mission-critical/meaningful for analysis

Table 1. Summary of the promises and pitfalls currently apparent in adaptive learning.

References

1. John R. Anderson. 1996. ACT: A simple theory of complex cognition. *American Psychologist* 51, 4: 355.
2. John R. Anderson, Albert T. Corbett, Kenneth R. Koedinger, and Ray Pelletier. 1995. Cognitive tutors: Lessons learned. *The journal of the learning sciences* 4, 2: 167–207.
3. Kimberly E. Arnold and Matthew D. Pistilli. 2012. Course signals at Purdue: Using learning analytics to increase student success. In *Proceedings of the 2nd international conference on learning analytics and knowledge*, 267–270. Retrieved February 13, 2017 from http://dl.acm.org/citation.cfm?id=2330666
4. Ryan SJ Baker, Albert T. Corbett, and Vincent Aleven. 2008. More accurate student modeling through contextual estimation of slip and guess probabilities in bayesian knowledge tracing. In *International Conference on Intelligent Tutoring Systems*, 406–415. Retrieved February 12, 2017 from http://link.springer.com/chapter/10.1007/978-3-540-69132-7_44
5. 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.
6. William G. Bowen, Matthew M. Chingos, Kelly A. Lack, and Thomas I. Nygren. 2014. Interactive Learning Online at Public Universities: Evidence from a Six-Campus Randomized Trial: Interactive Learning Online at Public Universities. *Journal of Policy Analysis and Management* 33, 1: 94–111. https://doi.org/10.1002/pam.21728
7. Jessie Brown. 2015. *Personalizing Post-Secondary Education: An Overview of Adaptive Learning Solutions for Higher Education*. Ithaka S+R. Retrieved from http://www.sr.ithaka.org/publications/personalizing-post-secondary-education/
8. Peter Brusilovsky. 2000. Adaptive hypermedia: From intelligent tutoring systems to web-based education. In *International Conference on Intelligent Tutoring Systems*, 1–7. Retrieved February 12, 2017 from http://link.springer.com/chapter/10.1007/3-540-45108-0_1
9. Peter Brusilovsky and Eva Millán. 2007. User models for adaptive hypermedia and adaptive educational systems. In *The adaptive web*, 3–53. Retrieved February 12, 2017 from http://dl.acm.org/citation.cfm?id=1768199
10. Gates Bryant, Adam Newman, Brian Fleming, and Laura Srkisian. 2016. *Learning to Adapt: The Evolution of Adaptive Learning in Higher Education*. Tyton Partners. Retrieved from http://tytonpartners.com/library/learning-to-adapt-2-0-the-evolution-of-adaptive-learning-in-higher-education/
11. Gates Bryant, Adam Newman, and Peter Stokes. 2013. *Learning to Adapt: A Case for Accelerating Adaptive Learning in Higher Education*. Tyton Partners. Retrieved from http://tytonpartners.com/library/accelerating-adaptive-learning-in-higher-education/
12. Peter Burrows. 2016. Adaptive learning: Can industry's most provocative software make a difference in the classroom? *EdSurge*. Retrieved December 5, 2016 from https://www.edsurge.com/research/special-reports/adaptive-learning/
13. Balakrishnan Chandrasekaran, Todd R. Johnson, and Jack W. Smith. 1992. Task-structure analysis for knowledge modeling. *Communications of the ACM* 35, 9: 124–137.
14. Harry Collins. 2010. *Tacit and explicit knowledge*. University of Chicago Press.
15. Albert T. Corbett and John R. Anderson. 1994. Knowledge tracing: Modeling the acquisition of procedural knowledge. *User modeling and user-adapted interaction* 4, 4: 253–278.

16. Jennifer B. Daines, Tonya Troka, and John M. Santiago Jr. 2016. Improving Performance in Trigonometry and Pre-Calculus by Incorporating Adaptive Learning Technology into Blended Models on Campus. https://doi.org/10.18260/p.25624
17. Manuela Ekowo and Iris Palmer. 2016. *The Promise and Peril of Predictive Analytics in Higher Education: A Landscape Analysis*. New America.
18. Philip J. Guo, Juho Kim, and Rob Rubin. 2014. How video production affects student engagement: an empirical study of MOOC videos. 41–50. https://doi.org/10.1145/2556325.2566239
19. Philip J. Guo and Katharina Reinecke. 2014. Demographic differences in how students navigate through MOOCs. 21–30. https://doi.org/10.1145/2556325.2566247
20. Sherif Halawa, Daniel Greene, and John Mitchell. 2014. Dropout prediction in MOOCs using learner activity features. *Experiences and best practices in and around MOOCs* 7.
21. John L. Hennessy. 2015. Information Technology and the Future of Teaching and Learning. Retrieved from https://www.youtube.com/watch?v=QLLMLG-jqxc
22. Colm Howlin. 2015. *Realizeit: System Overview*. Realizeit.
23. Petr Johanes. 2017. Epistemic cognition: A promising and necessary construct for enriching large-scale online learning analysis. In *Proceedings of the Fourth ACM Conference on Learning at Scale, L@S 2017*. https://doi.org/http://dx.doi.org/10.1145/3051457.3051462
24. Julia Kaufman, Sarah Ryan, Candace Thille, and Norman Bier. 2013. *Open Learning Initiative Courses in Community Colleges: Evidence on Use and Effectiveness*. Open Learning Initiative. Carnegie Mellon University. Retrieved from http://www.hewlett.org/library/open-learning-initiative-courses-in-community-colleges-evidence-on-use-and-effectiveness/
25. Juho Kim, Philip J. Guo, Carrie J. Cai, Shang-Wen (Daniel) Li, Krzysztof Z. Gajos, and Robert C. Miller. 2014. Data-driven interaction techniques for improving navigation of educational videos. 563–572. https://doi.org/10.1145/2642918.2647389
26. René F. Kizilcec, Chris Piech, and Emily Schneider. 2013. Deconstructing disengagement: analyzing learner subpopulations in massive open online courses. In *Proceedings of the third international conference on learning analytics and knowledge*, 170–179. Retrieved February 1, 2016 from http://dl.acm.org/citation.cfm?id=2460330
27. Larry Lagerstrom, Petr Johanes, and Mr Umnouy Ponsukcharoen. The Myth of the Six Minute Rule: Student Engagement with Online Videos. *age* 26: 1.
28. Brock J. LaMeres. 2015. Deploying Adaptive Learning Environments to Overcome Background Deficiencies and Facilitate Mastery of Computer Engineering Content. https://doi.org/10.18260/p.23783
29. Steven Lonn, Andrew E. Krumm, R. Joseph Waddington, and Stephanie D. Teasley. 2012. Bridging the gap from knowledge to action: Putting analytics in the hands of academic advisors. In *Proceedings of the 2nd International Conference on Learning Analytics and Knowledge*, 184–187. Retrieved February 13, 2017 from http://dl.acm.org/citation.cfm?id=2330647
30. Marsha Lovett, Oded Meyer, and Candace Thille. 2008. JIME-The open learning initiative: Measuring the effectiveness of the OLI statistics course in accelerating student learning. *Journal of Interactive Media in Education* 2008, 1. Retrieved February 12, 2017 from http://jime.open.ac.uk/article/view/142/jime-2008-14.pdf/
31. John C. Nesbit, Qing Liu Arita Liu, and Olusola O. Adesope. 2015. Work in Progress: Intelligent Tutoring Systems in Computer Science and Software Engineering Education. In

Proc. 122nd Am. Soc. Eng. Education Ann. Conf. Retrieved February 12, 2017 from https://www.asee.org/file_server/papers/attachment/file/0005/4286/ITS_in_CS_and_SE_Education.pdf
32. Sandra Nunn, John T. Avella, Therese Kanai, and Mansureh Kebritchi. 2016. Learning analytics methods, benefits, and challenges in higher education: A systematic literature review. *Online Learning* 20, 2. Retrieved February 12, 2017 from http://olj.onlinelearningconsortium.org/index.php/olj/article/view/790
33. Cathy O'Neil. 2016. *Weapons of Math Destruction: How big data increases inequality and threatens democracy*. Crown Publishing Group, New York, NY.
34. Roy Pea. 2014. *The learning analytics workgroup: A report on building the field of learning analytics for personalized learning at scale*.
35. Chris Piech, Jonathan Bassen, Jonathan Huang, Surya Ganguli, Mehran Sahami, Leonidas J. Guibas, and Jascha Sohl-Dickstein. 2015. Deep knowledge tracing. In *Advances in Neural Information Processing Systems*, 505–513. Retrieved October 24, 2016 from http://papers.nips.cc/paper/5654-deep-knowledge-tracing
36. Michael Polanyi. 2009. *The tacit dimension*. University of Chicago Press.
37. B. Gangadhara Prusty and Carol Russell. 2011. Engaging students in learning threshold concepts in engineering mechanics: adaptive eLearning tutorials. In *17 th International Conference on Engineering Education (ICEE)*. Retrieved February 12, 2017 from http://www.ineer.org/Events/ICEE2011/papers/icee2011_submission_250.pdf
38. Roger Riddell. 2013. Adaptive learning: The best approaches we've seen so far. *Education Dive*. Retrieved December 14, 2016 from http://www.educationdive.com/news/adaptive-learning-the-best-approaches-weve-seen-so-far/187875/
39. Ying Tang, Kauser Jahan, and Talbot Bielefeldt. 2015. The Effectiveness of an Adaptive Serious Game for Digital Logic Design. https://doi.org/10.18260/p.24861
40. Kurt VanLehn. 2011. The Relative Effectiveness of Human Tutoring, Intelligent Tutoring Systems, and Other Tutoring Systems. *Educational Psychologist* 46, 4: 197–221. https://doi.org/10.1080/00461520.2011.611369
41. Tony Wan. 2016. Adaptive Learning's Potential and Pitfalls. *EdSurge*. Retrieved December 5, 2016 from https://www.edsurge.com/news/2016-02-08-adaptive-learning-s-potential-and-pitfalls
42. Beverly P. Woolf. 2010. *Building intelligent interactive tutors: Student-centered strategies for revolutionizing e-learning*. Morgan Kaufmann.
43. 2016. *Decoding Adaptive*. EdSurge, London.
44. 2016. *Lessons learned from early implementations of adaptive courseware*. SRI International.
45. Adaptive Learning Market Acceleration Program (ALMAP) Evaluation. *SRI International*. Retrieved from https://www.sri.com/work/projects/adaptive-learning-market-acceleration-program