

Using Learning Analytics in Personalized Learning

Ryan Baker

Traditional statistical methods for data analysis involved top-down and hypothesis-driven analysis of relatively small data sets. Although more exploratory, bottom-up approaches to working with data have been around for several decades (Tukey, 1977), the past few years have seen an explosion in the use of analytics and data mining, methods for making discoveries and extracting information from larger data sets, in a more bottom-up fashion (Han, Kamber, & Pei, 2011). Analytics and data mining methods specialized for use with educational data sets—and to answer educational questions—are referred to as *learning analytics* (Siemens & Long, 2011) and *educational data mining* (Baker & Yacef, 2009).

Learning analytics (LA) and educational data mining (EDM) have been used for a range of applications. For example, these methods have been used to determine when learners are disengaged within online learning (Baker, Corbett, & Koedinger, 2004; Pardos, Baker, San Pedro, Gowda, & Gowda, 2013), to make early predictions about long-term outcomes (Bowers, 2010; Jayaprakash, Moody, Lauría, Regan, & Baron, 2014; San Pedro, Baker, Bowers, & Heffernan, 2013), to understand how different students choose to use learning resources (Amershi & Conati, 2009; Beck & Mostow, 2008; Kizilcec, Piech, & Schneider, 2013), and for many other applications. Models (automated measurements produced using EDM/LA) of student cognition, engagement, and learning can predict not just student achievement within a specific school year (Pardos et al., 2013) but also can predict outcomes several years later, including college attendance (San Pedro et al., 2013) and college major (San Pedro, Baker, Heffernan, & Ocumpaugh, 2015).

This chapter discusses LA in the context of personalized learning, discussing both past successful examples and potential future opportunities, as well as action principles for how state education agencies (SEAs), local education agencies (LEAs), and schools can best put LA into practice.

Learning Analytics and Personalized Learning: State of the Art

The goal of individualizing learning to each student's needs is not a wholly new goal (e.g., Parkhurst, 1922), yet education is still a long way from achieving this goal. Indeed, despite attempts to introduce demonstrably effective practices such as mastery learning as early as the 1960s (see review in Airasian, Bloom, & Carroll, 1971), much learning remains focused on whole-group activities, such as lectures, that do not offer much scope for personalization. Even as we move to an era of greater usage of online learning resources, many contemporary resources such as xMOOCs (Breslow et al., 2013) and Khan Academy (Dijksman & Khan, 2011) still emphasize "one-size-fits-all" lectures and activities with limited scope for tailoring content or presentation to individual needs.

However, an increasing number of online and blended interactive learning systems are moving toward being more personalized. A blended or online learning environment can be, in general, personalized in two ways. First, students can be given options to personalize the environment

The goal of individualizing learning to each student's needs is not a wholly new goal (e.g., Parkhurst, 1922), yet education is still a long way from achieving this goal.

themselves. Some environments do offer students a considerable amount of choice about their learning experiences. For example, intelligent tutoring systems such as SQL-Tutor (used by hundreds of thousands of undergraduates to learn database programming) offer students choice about what topic to work on next (Mathews & Mitrovic, 2007). This type of personalization extends even to elementary school students, with systems such as the Project LISTEN Reading Tutor (for elementary school students). In Project LISTEN, students are allowed to choose what story they read next (Mostow et al., 2002). Other systems, such as gStudy, engage students in planning their learning experiences (Perry & Winne, 2006).

However, it is more common to see systems in which learning is made adaptive and personalized to the learner's needs by the teacher or by the learning system itself. This type of practice developed before the widespread use of computers in classrooms, with teachers using formative assessments to drive mastery learning practices in which students work through material on a given topic until they can demonstrate the skills relevant to that topic (Airasian et al., 1971). Indeed, some of the first individualizing of blended and online learning at scale involved replicating mastery learning practices through a computer. For example, Cognitive Tutors for Algebra, now used by hundreds of thousands of students a year, assessed student knowledge as students worked through mathematics problems (Corbett & Anderson, 1995) and used that information to implement mastery learning. Cognitive Tutors for Algebra has been effective at promoting positive learning outcomes (Pane, Griffin, McCaffrey, & Karam, 2014); its algorithms to assess student knowledge arguably represent the first widespread use of EDM/LA.

The systems that have followed the Cognitive Tutor use models developed based on learning analytics to adapt to students in many more ways than simply implementing mastery learning. For example, the ALEKS system for algebra and chemistry, used by more than 100,000 students a year and shown to be effective (Craig et al., 2013), also uses EDM/LA to determine what prerequisite skills the student is lacking in order to shift the student's work to prerequisite skills when necessary (Doignon & Falmagne, 1999). This type of practice helps to avoid situations in which a student "wheel spins" (Beck &

Gong, 2013), working continuously on material with no success and little potential of success due to not knowing the prerequisites for the current material—or worse, when the student continually advances to harder material, failing topic after topic.

Learning analytics about student knowledge is used for more than just automated adaptation. Many online learning providers use it to support instructor practice as well. For example, automated data on student success in mathematics is presented to elementary school classroom teachers by the Reasoning Mind system, used by more than 100,000 students a year and shown to be effective (Waxman & Houston, 2008, 2012). This system also provides teachers with professional development that shows them how to use the system's analytics to inform proactive remediation, in which the instructor selects students or groups of students for one-on-one or small-group tutoring during class (Miller et al., 2015). A similar approach is taken at the undergraduate level by the Course Signals system, which tracks student course participation and performance on early assignments and integrates these data sources into systems that predict eventual student course failure and dropout in order to provide instructors with reports on which students are at risk and why (Arnold & Pistilli, 2012). The reports in Course Signals are combined with recommendations for instructors on how to use them, including templates for emails that automatically fill in the student's name and performance factors that indicate risk. The use of Course Signals was shown in a study at Purdue University to lead to significantly lower dropout rates (Arnold & Pistilli, 2012).

Learning Analytics and Personalized Learning: Future Potentials

Modern LA for personalization extends further than simply assessing and supporting learning and performance to attempting to enhance engagement and affect. Although the evidence for effectiveness is still preliminary, involving small studies rather than national-level randomized controlled trials, some pilot projects have shown evidence that these approaches can be beneficial.

For example, some of the first work using EDM involved systems that could automatically infer when a student was “gaming the system,” misusing a learning system in order to succeed without learning, for instance, by clicking through hints at high speed or systematically guessing (Baker et al., 2004). Automated measurements (often termed “models”) of gaming the system have been used to trigger automated intervention, reducing gaming behavior and improving learning (Baker et al., 2006). They have also been used in interventions that teach students why gaming is ineffective and reduces their learning, also reducing gaming behavior and improving learning (Arroyo et al., 2007). Similarly, models that can automatically infer student emotion have been used in systems at the undergraduate level, responding in supportive ways to struggling students and in sarcastic ways to students who are generally successful but are not putting in enough effort (D’Mello et al., 2010). Further work has developed approaches that not only attempt to support students but also actually attempt to *increase* student confusion in some situations, increasing challenge and improving learning outcomes (D’Mello, Lehman, Pekrun, & Graesser, 2014).

The potential for enhancing self-regulated learning (SRL) is somewhat less certain. For example, LA was used to study which hint-use strategies led to better learning in Cognitive Tutors (Aleven, McLaren, Roll, & Koedinger, 2004). Teaching students more effective SRL strategies and providing immediate feedback on ineffective or inappropriate

hint use led to lasting changes in student behavior but no difference in learning outcomes for the mathematics material students were expected to learn (Roll, McLaren, Alevan, & Koedinger, 2011). This result has been replicated in another study conducted by Albert Corbett at Carnegie Mellon University. Overall, there has been insufficient research to know whether the relative lack of success of this approach is indicative of the general difficulty of improving learning through using analytics-based SRL interventions or whether some aspect of the design of this intervention led to a lower impact on learning outcomes.

Determining the eventual impact of EDM/LA is difficult. In general, LA and EDM are still advancing relatively rapidly. The past decade has seen construct after construct that seemed difficult to measure turn out to be feasible to measure effectively using EDM/LA. However, work to use EDM/LA for personalization is still ongoing and lags a few years behind the work on measuring constructs, such as gaming the system and emotion, for the simple reason that it is not possible to use an automated measure of a learning-related construct to enhance learning before that measure exists. In addition, the individuals who are skilled in interaction and educational design—in developing interventions that use EDM/LA to improve outcomes—are not the same individuals who are good at using LA and EDM to build the measurements on which those interventions depend. As a result, readers of this chapter may find, scant years from now, that the personalization technologies that are available at the time they are reading this chapter far surpass the technologies reported today (or, alternatively, the technologies may be very similar; see some of the action principles discussed below).

While the methods of learning analytics have considerable potential to enable high-quality adaptive personalization to learning, there are several challenges ...

Some Considerations on Using Learning Analytics in Personalized Learning

While the methods of learning analytics have considerable potential to enable high-quality adaptive personalization to learning, there are several challenges that must be taken into consideration for these methods to reach their full potential to enhance student outcomes. The following sections of this chapter discuss the role played by issues such as privacy and model validity and how these challenges can be appropriately addressed. The chapter also discusses the essential role played by stakeholders such as teachers, school leaders, and parents and how LA-based personalization can effectively support these stakeholders.

Privacy

In recent years, there has been considerable concern about student privacy in this emerging era of analytics (Slade & Prinsloo, 2013). There are reasons for concern when student data may be used for marketing or may be disclosed unnecessarily. Regrettably, some of this concern has led to suggesting policies that are very likely to hinder the use of educational data for educational improvement. For example, as of this writing, the U.S. Department of Education has recommended terms of use for online learning that forbid “data mining” (Privacy Technical Assistance Center, 2015) based on the apparent misconception that “data mining” is equivalent to advertising. Recent legislation has also proposed policies for handling educational data that require that no personally identifiable information be available or indeed that require that all data be discarded at the end of each school year. Discarding all data essentially destroys the potential for using analytics

and data mining to enhance education, for little reduction in risk. Even the seemingly reasonable compromise of removing all personally identifiable information from data has the potential to reduce the degree to which we can improve education through personalization driven by LA. Data that do not include personally identifiable information cannot be used to conduct longitudinal research in which performance and behavior are linked to eventual learner outcomes. If it is impossible to verify long-term outcomes, technologies may be selected that enhance learning in the short term but do not produce positive outcomes in the long term.

Several possible solutions remove the drawbacks of full anonymization while protecting student privacy and maintaining compliance with the Family Educational Rights and Privacy Act (FERPA), the federal law that protects the privacy of student education records. For example, SEAs and LEAs can store personally identifying information in trust, with an individual within the LEA or SEA holding a strictly guarded key to the data sets and the links between them, allowing access only for legitimate educational research and enhancement purposes. Alternatively, a trusted broker can be selected to protect this information, as the National Student Clearinghouse does for undergraduate enrollment data. Modern technologies for data mining and analytics can support analysis by remote researchers in which analyses can be conducted using sensitive data but in which the sensitive data itself are never exposed to the remote researcher. All data would be retained by organizations entrusted to protect students, and thus it would be possible to use LA to its full potential. Modern systems for educational data, such as the MARi platform and OpenLAP, limit access to data, keep control over data in the hands of students and their parents, and do not inhibit educational improvement. SEAs, LEAs, and schools have a role to play in realizing the potential of LA by partnering with reliable commercial and nonprofit entities to insist on systems that protect privacy but do not prevent students from having access to high-quality personalized education.

Model Validity

When using LA to impact educational practice, it is important to ensure that the LA models are valid. Although a great deal of high-quality software is available, there is also considerable software that is low quality. Schools should be prepared to ask good questions of developers. Traditionally, school purchasing decisions have been based on relatively light evidence, such as testimonial evidence provided by developers. The What Works Clearinghouse (<http://ies.ed.gov/ncee/wwc/>) encourages schools to ask, “Does it work?” and to insist on evidence from randomized controlled trials. In a randomized controlled trial, a system is compared with some existing pedagogical practice in a study with random assignment. As schools increasingly work with vendors that provide personalized learning systems and analytics, the schools should ask to see evidence on how the personalization and analytics were developed. Scientific papers in reputable, peer-reviewed journals and conferences can provide evidence that the system under consideration was developed according to valid principles. For example, schools and school districts should examine these publications for evidence on whether models were tested on the same students they were developed for or whether the models are shown to function appropriately for students other than those for which the models were developed.

However, even this type of validation is sometimes not enough. Ideally, models should also be validated for accuracy in contexts similar to the schools where they will be used.

A rural school should be wary of using software developed for suburban students; there is evidence that the same behaviors do not always predict emotion or engagement in different populations (Ocumpaugh et al., 2014). It is increasingly considered best practice at the higher education level to validate models for individual universities, a practice adopted, for instance, by the company ZogoTech. Although it may not be feasible at the current time to validate models for each and every school in the United States, it is feasible to ask whether a model being used was validated on students similar to those in the school considering adoption. There are even metrics for the similarity between schools that can be used to inform consideration of the relevance of study evidence for a given school (Tipton, 2014).

Leveraging All the Relevant Stakeholders

Often, schools rely solely on teachers to personalize students' education beyond what online and blended learning can provide. Teachers have a key role to play in making learning effective for students, and most LA reports are targeted toward them. However, many other stakeholders also have roles to play. Guidance counselors can access LA reports and use automated predictions to identify students who are engaged by the subject they are studying but who might not be considering careers in this area, due, for instance, to demographic factors. These students can be encouraged to participate in summer or afterschool enrichment programs that give them experience in the area of study. So, too, students who are engaged by a subject but struggling with it and are not on track to be able to go into the careers they are interested in are ideal candidates for afterschool tutoring or other support. By contrast, a student who is performing well at a subject but who does not seem to be particularly engaged with it should probably be encouraged to place his or her efforts into other subjects. As such, guidance counseling can be made more personal and potentially more effective.

Similarly, school leaders—particularly those whose task it is to deal with disciplinary problems—also may benefit from LA from the systems students are using. Although a considerable proportion of disciplinary incidents involves factors outside the direct control of schools (e.g., Bachman & Schulenberg, 1993; Murray, Farrington, & Sekol, 2012), it may be beneficial for a school leader to see evidence that a student who is getting into trouble is nonetheless remaining engaged in learning one or more subjects. This may suggest positive behavior supports (Bambara, Nonnemacher, & Kern, 2009) that the school leader can consider applying, including activities to reengage the student with schooling through his or her preferred subjects.

Finally, parents can be empowered to help support their children's learning. Currently, efforts to incorporate parents in their children's learning are often very limited, with report cards only provided at occasional intervals and reports containing relatively limited information about how to help their specific child. If anything, the move to online learning has disempowered parents further because many parents cannot help students with their homework as easily as before (because it occurs within an unfamiliar online system rather than on paper). When resources are given to parents, they are often provided to every parent in a class, ignoring whether that student needs the resource or how to individualize it for that child. By contrast, reports from personalized learning systems that collect considerable data about each child can be provided to parents. For example, the ASSISTments system sends text messages and emails to parents, telling them what

their children are currently struggling with (Broderick, 2011). Even simple systems that notify parents about missed assignments can lead to positive impacts on student academic outcomes (Bergman, under review).

Action Principles for States, Districts, and Schools

- a. Develop data policies that make learning analytics possible. Schools, LEAs, and SEAs have an important role to play in making it possible for LA to be used to benefit students. By partnering with organizations that handle student data responsibly and by adopting policies that protect privacy but preserve data and ways to link student learning data to future data on their success, schools, LEAs, and SEAs can increase the potential for personalized learning to benefit their students.
- b. Mitigate the data loss stemming from student mobility. School mobility is a fact of 21st century education; because American society is highly mobile, students are likely to change schools repeatedly during their education. Although school mobility may not be problematic for students of high socioeconomic status (SES), it is associated with poorer outcomes among lower SES and minority students, especially if a student changes schools several times (Xu, Hannaway, & D'Souza, 2009). Districts and SEAs need to work with technology providers so that a student's data in one school can follow him or her to another school. A student who is halfway through the school year and has used a high-quality system such as ALEKS all year should not start from square one if his or her new school also uses ALEKS. By coordinating between schools and technology providers, a student's account can be transferred between schools, and the student can pick up in Paterson (NJ) where he or she left off in Orange (NJ).

Action Principles for States and Districts

- a. Ask for raw data and student models from providers. The data being collected by personalized learning systems is useful, not just within that specific learning system but more broadly as well. Models of constructs such as engagement can be processed by states or school districts into reports for guidance counselors that predict student long-term outcomes and help the guidance counselors advise students how to stay on track.
- a. Incorporate these models into state or city early-warning systems, complementing traditional data sources, such as grade data, disciplinary incidents, standardized examination scores, and demographic data. Data from personalized learning systems are a treasure trove for SEAs and LEAs wanting to improve student outcomes.

Action Principles for Districts and Schools

- a. Seek appropriate professional development for teachers working with analytics. Teaching with blended learning and online homework differs from traditional pedagogical approaches, and different teacher practices are relevant (Ronau, Rakes, & Niess, 2012). There is considerable evidence that these new approaches to teaching are more effective in the hands of teachers who have received appropriate professional development (see review in Lawless & Pellegrino, 2007). Also, instructors who have received sufficient professional development are more likely to adopt effective practices, such as viewing reports on student knowledge and success and

using proactive remediation strategies to help struggling students (Miller et al., 2015). Professional development for working with modern personalized learning technologies is available from technology and curriculum providers and from universities ranging from Teachers College Columbia University to Framingham State University. Students will benefit considerably if schools make resources available for teachers to partake in these programs.

- b. Leverage multiple stakeholders to participate in personalization. Personalization is not something that an online learning or blended learning system does alone. It works most effectively when it leverages—and empowers—what teachers, guidance counselors, school leaders, and parents have to offer. Extending analytics reports to all these individuals—and when appropriate, providing them with training on how to use reports—has the potential to considerably improve student outcomes.

Action Principles for Schools

- a. Be an educated consumer of personalized learning software. School officials should insist on seeing evidence in appropriate peer-reviewed conferences and journals that the systems under consideration have been validated to work for students similar to the ones in their school. Failing to check this risks that students will receive ineffective learning support.

References

- Airasian, P. W., Bloom, B. S., & Carroll, J. B. (1971). *Mastery learning: Theory and practice*. J. H. Block (Ed.). New York, NY: Holt, Rinehart, & Winston.
- Aleven, V., McLaren, B., Roll, I., & Koedinger, K. (2004). Toward tutoring help seeking. *Proceedings of the International Conference on Intelligent Tutoring Systems* (pp. 227–239). Maceió, Alagoas, Brazil.
- Amershi, S., & Conati, C. (2009). Combining unsupervised and supervised classification to build user models for exploratory learning environments. *Journal of Educational Data Mining*, 1(1), 18–71.
- Arnold, K. E., & Pistilli, M. D. (2012). Course signals at Purdue: Using learning analytics to increase student success. *Proceedings of the 2nd International Conference on Learning Analytics and Knowledge* (pp. 267–270). Washington, DC: Association for Computing Machinery.
- Arroyo, I., Ferguson, K., Johns, J., Dragon, T., Meheranian, H., Fisher, D., & Woolf, B. P. (2007). Repairing disengagement with non-invasive interventions. *Proceedings of the International Conference on Artificial Intelligence in Education* (pp. 195–202). Los Angeles, CA.
- Bachman, J. G., & Schulenberg, J. (1993). How part-time work intensity relates to drug use, problem behavior, time use, and satisfaction among high school seniors: Are these consequences or merely correlates? *Developmental Psychology*, 29(2), 220.
- Baker, R. S., Corbett, A. T., & Koedinger, K. R. (2004). Detecting student misuse of intelligent tutoring systems. *Proceedings of the 7th International Conference on Intelligent Tutoring Systems* (pp. 531–540). Maceió, Alagoas, Brazil.
- Baker, R. S. J. D., Corbett, A. T., Koedinger, K. R., Evenson, S. E., Roll, I., Wagner, A. Z.,...Beck, J. (2006). Adapting to when students game an intelligent tutoring system. *Proceedings of the 8th International Conference on Intelligent Tutoring Systems* (pp. 392–401). Jhongli, Taiwan.
- Baker, R. S. J. D., & Yacef, K. (2009). The state of educational data mining in 2009: A review and future visions. *Journal of Educational Data Mining*, 1(1), 3–17.
- Bambara, L. M., Nonnemacher, S., & Kern, L. (2009). Sustaining school-based individualized positive behavior support: Perceived barriers and enablers. *Journal of Positive Behavior Interventions*, 11(3), 161–176.

- Beck, J. E., & Gong, Y. (2013). Wheel-spinning: Students who fail to master a skill. *Artificial Intelligence in Education, 7926*, 431–440.
- Beck, J. E., & Mostow, J. (2008). How who should practice: Using learning decomposition to evaluate the efficacy of different types of practice for different types of students. *Proceedings of the International Conference on Intelligent Tutoring Systems* (pp. 353–362). Montreal, Canada.
- Bergman, P. (under review). *Parent-child information frictions and human capital investment: Evidence from a field experiment*. Prepublication draft retrieved from <http://www.columbia.edu/~psb2101/BergmanSubmission.pdf>
- Bowers, A. J. (2010). Grades and graduation: A longitudinal risk perspective to identify student dropouts. *Journal of Educational Research, 103*(3), 191–207.
- Breslow, L., Pritchard, D. E., DeBoer, J., Stump, G. S., Ho, A. D., & Seaton, D. T. (2013). Studying learning in the worldwide classroom: Research into edX's first MOOC. *Research & Practice in Assessment, 8*(1), 13–25.
- Broderick, Z. (2011). *Increasing parent engagement in student learning using an Intelligent Tutoring System with automated messages*. Unpublished master's thesis, Worcester Polytechnic Institute, Worcester, MA, USA.
- Corbett, A. T., & Anderson, J. R. (1995). Knowledge tracing: Modeling the acquisition of procedural knowledge. *User Modeling and User-Adapted Interaction, 4*(4), 253–278.
- Craig, S. D., Hu, X., Graesser, A. C., Bargagliotti, A. E., Sterbinsky, A., Cheney, K. R., & Okwumabua, T. (2013). The impact of a technology-based mathematics after-school program using ALEKS on student's knowledge and behaviors. *Computers & Education, 68*, 495–504.
- Dijksman, J. A., & Khan, S. (2011, March). Khan Academy: The world's free virtual school. In *APS Meeting Abstracts, 1*, 14006.
- D'Mello, S., Lehman, B., Pekrun, R., & Graesser, A. (2014). Confusion can be beneficial for learning. *Learning and Instruction, 29*, 153–170.
- D'Mello, S., Lehman, B., Sullins, J., Daigle, R., Combs, R., Vogt, K., & Graesser, A. (2010). A time for emoting: When affect-sensitivity is and isn't effective at promoting deep learning. *Proceedings of the International Conference on Intelligent Tutoring Systems* (pp. 245–254). Pittsburg, PA.
- Doignon, J. P., & Falmagne, J. C. (1999). *Knowledge spaces*. Berlin, Germany: Springer.
- Han, J., Kamber, M., & Pei, J. (2011). *Data mining: Concepts and techniques*. Waltham, MA: Morgan Kaufmann.
- Jayaprakash, S. M., Moody, E. W., Lauria, E. J., Regan, J. R., & Baron, J. D. (2014). Early alert of academically at-risk students: An open source analytics initiative. *Journal of Learning Analytics, 1*(1), 6–47.
- Kizilcec, R. F., Piech, C., & Schneider, E. (2013, April). Deconstructing disengagement: Analyzing learner subpopulations in massive open online courses. *Proceedings of the Third International Conference on Learning Analytics and Knowledge* (pp. 170–179). Washington, DC: Association for Computing Machinery.
- Lawless, K. A., & Pellegrino, J. W. (2007). Professional development in integrating technology into teaching and learning: Knowns, unknowns, and ways to pursue better questions and answers. *Review of Educational Research, 77*(4), 575–614.
- Mathews, M., & Mitrovic, A. (2007). Investigating the effectiveness of problem templates on learning in ITSs. In R. Luckin, K. Koedinger, & J. Greer (Eds.), *Proceedings of Artificial Intelligence in Education* (pp. 611–613). Los Angeles, CA.
- Miller, W. L., Baker, R. S., Labrum, M. J., Petsche, K., Liu, Y. H., & Wagner, A. Z. (2015, March). Automated detection of proactive remediation by teachers in reasoning mind classrooms. *Proceedings of the Fifth International Conference on Learning Analytics and Knowledge* (pp. 290–294). Washington, DC: Association for Computing Machinery.