Adaptive Online Learning
The Present and Future of Education

November 2013

Dr. Bruce McLaren, Carnegie Mellon University

WILEY

The notion of adaptive learning is not new. In a classroom setting, an instructor often assesses how well the class is following as she moves through a lesson. She might ask a few questions to ensure that the class is following, initiate a discussion, or make subtle adjustments to her presentation based upon the body language and facial expressions of the students. She can create assignments to check understanding and then address areas of weakness the next day. Every instructor is a practitioner of adaptive learning to a greater or lesser extent based upon her training, experience, ability to diagnose student misunderstandings, and knowledge of the subject matter.

However, as class sizes increase, and more and more students arrive with different starting points and learn at different paces, the challenge to create an optimal learning experience increases. Instructors may wish that they had a tool that could do some of the work for them, assessing students and offering self-paced instruction based upon what each student needs to learn most.

Class Scenario

Consider again our instructor from above. Let's extend the experience in a fictional, yet realistic account of the myriad challenges she faces:

Our instructor is talented, particularly at understanding how well her students are grasping the material through reading their facial expressions and body language and through their general participation and engagement level in the class. That is, she was until her Algebra 1 class sizes started to double, then triple, in size. She had always had some students who needed to move more methodically through the material, or who had more difficulty grasping key concepts, while others were further ahead and needed to be challenged. But now the challenges of meeting the diverse needs of her students are magnified. She now has many more students needing additional support while a different, also large subset of students can move more quickly and focus on advanced materials only. Many students seem to hover somewhere between the two groups, and our instructor knows that she can help transform some of those students into high achievers, but with too many students and too little ability to personalize the course, the likelihood is that they will not excel. Let's look at some of the students in her class:

Giselle, Biology major
Motivation: Loves math and considers it one of her best subjects.
Challenges: She is from another country and her English is not as fluent as her classmates. She often struggles to keep up with the lessons in class although she does not find the homework very difficult. She is capable of moving more quickly through the material but cannot because the lectures are so difficult for her to follow.

Michael, Business major
Motivation: Can achieve well when he is shown how to do problems.
Challenges: He cannot grasp the material fully until he tries problems for himself. In class it is not always possible to work through the examples individually so although he is doing fine, he is frustrated that he is not doing better. He wishes he could go back and rewind the examples and then stop and try them out himself.

Mary, Physics major
Motivation: She is so far ahead she does not really need to attend classes but does so in order to not lose any grade points.
Challenges: She has already excelled at most of the class topics in previous classes. She spends most of the class texting and working on other homework.

Our instructor is sometimes aware of these disparities, but she can do little about them except to try to adjust her lessons accordingly, as much as possible. She wants to do the right things for her students but it is not always possible to know precisely what is happening, particularly when students are not engaged.

The institution is very focused on student outcomes and is looking at all of these classes in aggregate to ensure that all students are achieving to their highest potential. They cannot see where the points of failure lie; they only know that they exist so they can do little other than ask instructors to work harder to engage students and help them to be successful.

Our New Environment

What is the solution to scenarios such as these? How can instructors help students achieve their highest potential, not only in traditional classrooms but in blended and online classrooms as well? Welcome to adaptive learning, a software environment where technology, educational psychology, and cognitive science collaborate with big data to carve out customized pathways through curriculums for individual learners and free up teachers to devote their energies to more productive and scalable tasks (Stokes, 2013).

Technology, educational psychology, and cognitive science collaborate with big data to carve out customized pathways and free up teachers.

While the concept of adaptive learning is seemingly straightforward and simple, its implementation is not. As mentioned above, the traditional classroom does not support this kind of learning, as teachers struggle with large class sizes and the varying levels of baseline proficiency of each learner. Teachers simply do not have the time or bandwidth to customize instruction for all of their students. Into this challenging pedagogical scenario adaptive learning software is emerging to complement teachers by providing what they struggle to provide: instructional approaches and content that cater to the individual student's knowledge, motivation, and learning style.

Adaptive learning is all about reacting to what a student knows, analyzing how they learn, and catering instruction to that knowledge. It can recognize this need and provide the targeted practice that a student needs. Looking at the example of Mary in the case above, she may be well ahead of the fractions level of mathematical understanding, perhaps already tackling algebraic concepts, and thus needs to be challenged with material that is commensurate with her level. Again, an adaptive system should – and can – recognize this need and present more advanced materials to this student. The system can know the difference between these students by presenting a variety of problems to students, of different types and at different levels of difficulty, and evaluating their performance.

Emerging K-12 standards, such as the Common Core Standards (National Governors Association Center for Best Practices and Council of Chief State School Officers, 2010), push educational institutions to make sure all of their students attain a basic level of achievement and outcomes in preparation for college. This is only possible by paying close attention to what individual students are learning and addressing the specific needs of those students. At the same time, nearly 1 in 5 US College Presidents indicated that they are strongly interested in pursuing adaptive online learning and prior learning assessment over other education technology innovations such as MOOCs (Massive Open Online Course). They know that as funding becomes increasingly tight, any technology that can help to improve the outcomes and engagement of their students is of the highest priority (Lederman, 2013).

Adaptive Learning and the Research Base

Adaptive learning technology is the product of many years of research, conceptualization, and practice. For instance, intelligent tutoring systems (Koedinger et al, 1997; VanLehn, 2011; Woolf, 2009), educational games (Conati & Manske, 2009; Gee, 2008; Hapgood & Ainsworth, 2011), and item response theory (Fox, 2010; Hambleton, Swaminathan, & Rogers, 1991) have all influenced the current state of adaptive learning technology. Adaptive learning systems have been shown to lead to impressive improvement in student learning across a wide range of domains, including math (Beal et al, 2007 Koedinger et al, 1997), reading (Mostow & Beck, 2007), science (VanLehn et al, 2005), physical procedural tasks (Rickel & Johnson, 1999), medicine (Eliot & Woolf, 1994), computer programming (Corbett & Anderson, 1992) and even across domains (Graesser et al, 2005). These adaptive learning systems are typically grounded in established cognitive theory, cognitive task analysis, and cognitive modeling[1].

Adaptivity is at the core of how people learn. Expecting students to learn by presenting them all with the same material is not only unrealistic, it is unproductive. There is clear evidence that many students get turned off by the process of learning and lose their motivation if they are not given instructional material and guidance that is commensurate with their goals and level of understanding, and that challenges them (Dweck, 1986; Pintrich, 2000). On the other hand, students who need help should be given that help, also called scaffolding in the Learning Sciences, until they have reached a more independent state of understanding.

Vygotsky (1978), a Russian psychologist who lived in the early part of the 20th century and is still quite influential in today's world of educational research, proposed that young learners be given help by adults until they are ready to work more independently. Vygotsky called the learning level at which students could solve problems, if assisted by adults, the "Zone of Proximal Development" (ZPD). This theory, which has been used and extended by a variety of educational theorists, in both independent and social learning, since it was first proposed in the 1930s, is a key concept underlying adaptive learning. The way the theory is applied in educational practice today -- and adaptive learning, in particular -- is that a student is given material just at the edge, or slightly beyond, his or her current level of understanding. In this way, the learner is challenged, but at the same time can use what he or she currently knows to create new knowledge and understanding.

An adaptive learning system tries to evaluate what a student knows and present that student with new material and/or problems to solve that are just at the edge, or just beyond the edge, of the student's demonstrated knowledge and understanding. The adaptive learning system then supports (or scaffolds) learning at this point. Such an approach is depicted in Figure 1.

---

[1] Most of the academic research has been done on a relatively small scale (e.g., lab studies, one or a few schools) in focused research studies and/or hasn't made its way into mainstream use in education, with some notable exceptions. For instance, Carnegie Learning, a spin-off company from Carnegie Mellon University, and now wholly owned by the Apollo Group (also the owners of the University of Phoenix), has had its Cognitive Tutors (TM) on the market for more than 10 years, helping middle and high school children learn mathematics. The software identifies weaknesses in student understanding and tailors the curricula to those weaknesses. The approach that Carnegie Learning takes is to leverage the best academic research on how people learn and, over time, incorporate the research recommendations into their software (Aleven & Koedinger, 2002; Corbett, Koedinger, & Hadley, 2001; Koedinger et al, 1997). There have been a number of larger-scale studies of Cognitive Tutors to validate their effectiveness (Ritter et al, 2007).

Figure 1 The Zone of Proximal Development and Adaptive Learning Systems

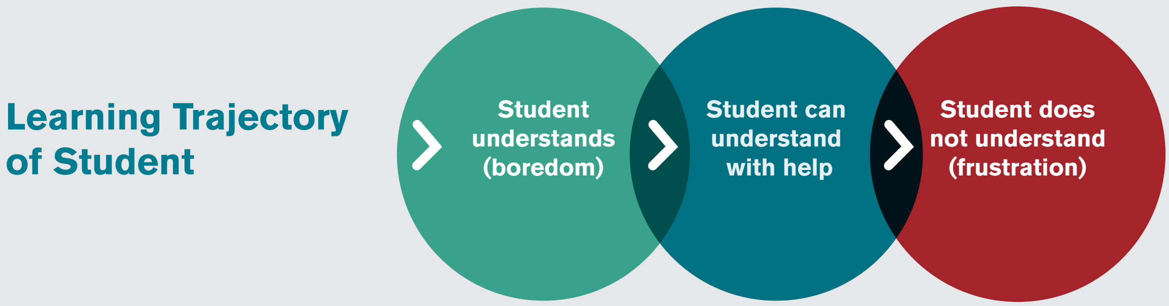

In Figure 1 three key regions of "understanding" are highlighted: a region where the student knows the content being presented, a region where the student could understand, with help, and a region that is well beyond the student's current understanding. It is not productive for the student to work with material he already understands, the region shown on the left. This is likely to lead to boredom and disinterest. It is also not productive for the student to be presented with material well beyond his understanding, as depicted by the region on the right. This will lead to frustration. What an adaptive learning system tries to do is keep the student working with material just beyond what he knows, in a region where some support and feedback could lead the student to learning and new insights. This is adaptivity as Vygotsky imagined it.

It is also important for students to use what they know to help them learn more. Essentially, this is like "bootstrapping" new knowledge from existing knowledge. From a cognitive science perspective, this involves using long-term memory, what a learner already knows, to incorporate and associate new concepts and knowledge from short-term, working memory (Sweller, Van Merriënboer, & Paas, 1998). We have long known that working memory is relatively small (Miller, 1956) and thus can be easily taxed to the point of blocking new knowledge from transferring to long-term memory. A key to adaptive learning is to maintain low stress on working memory, so the student can easily and actively extend her long-term network of associations and strengthen her understanding of partially understood concepts. Again, by keeping the student's learning trajectory towards the middle of Figure 1, working memory is less taxed and the important associations and meanings that extend long-term memory are given the opportunity to develop. Furthermore, by extending long-term memory and associations, the student may also be able to apply learned concepts to new contexts of the student's demonstrated knowledge and understanding.

Benefits of Adaptive Learning for Instructors and Institutions

Adaptivity in learning systems holds benefits for the individual, the instructor, and the institution. In each case the benefit is different but equally important.

Adaptive systems provide both formative and summative assessments... allowing the teacher to customize classroom teaching.

Individual students benefit from adaptive learning systems by being presented with challenges and problems that match up with their level of understanding and advancement in a particular topic, as illustrated in Figure 1. This type of individualization is virtually impossible to achieve in the traditional classroom, as previously mentioned. Not only is it impossible for a teacher to constantly and carefully evaluate where each individual student stands in his or her understanding of the content, but it is equally impossible for the teacher to cater classroom instruction to every student, separately and individually. Yet, at the same time, it is well known that students benefit the most when they learn from human tutors (Bloom, 1984) – who ultimately must be adaptive in their approach to teaching – as opposed to traditional classroom learning. Adaptive learning systems attempt to emulate human tutors, recognizing the needs of their tutees and pacing the presentation of new materials accordingly, and have been shown to come very close to their level of success (VanLehn, 2011).

While benefits to the individual learner of adaptive learning are the most apparent, there are also clear benefits to instructors. First, adaptive learning systems relieve some of the burden on the instructor, offloading that burden to an intelligent, automated helper. Although a teacher would ideally spend lots of time with every student in a class, this simply isn't practical in today's classrooms. Adaptive learning systems are tireless companions to students, confederates who act in place of teachers and who pay attention to everything the student does, down to the last keystroke. Another key benefit to instructors is the feedback that adaptive systems provide, in the form of student reports. Since all of the activity that students undertake while using an adaptive system – the problems they solve, the mistakes they make, the help they request, and the questions they answer – and the responses the learning system provides (error messages, hints, and other feedback) are recorded and summarized, teachers who use adaptive learning software are able to get a comprehensive picture of where all of their students stand and what their individual levels of understanding are, without having to be at the student's shoulder all the time. These adaptive system evaluations often provide both formative assessments – indications of interim progress that can be used to make adjustments during learning – and summative assessments – indications of achievement at the completion of a significant piece of the curriculum. This, of course, allows the teacher to better customize classroom teaching and identify students who may need special attention.

Finally, adaptive learning systems benefit the institution in several ways. Students who use adaptive learning systems improve their grades and the passing rates of these students will often be higher than those who do not use such systems. The adaptivity of learning systems, which allows students to progress at their own rates, also supports faster completion rates, for those students who can handle the material more easily. Lastly, an institution that employs adaptive learning software can be said to be keeping with the times and technology – something that is absolutely critical in attracting and retaining students.

ORION - A Real-World Solution to the Need for Adaptive Online Learning

ORION, Wiley's answer to adaptive learning, is real-world educational software that addresses the need to adapt instruction to individual students. It relies on much of what has been learned in academic research, putting that research into practice. ORION is able to adapt its curricula through the student's performance on given questions; the student's responses are used to update a proficiency analysis using a sophisticated statistical model. After the student takes a brief diagnostic, the student is presented with material and problems that match his or her (current) level of understanding. In typical mastery learning fashion (Bloom, 1985), the system starts the student with materials that the model indicates they are ready to tackle and then, over time and as they demonstrate readiness, presents them with more challenging material.

ORION's initial diagnostic presents the student with a variety of questions at different levels of difficulty to assess which of the many different possible pathways of progression they are most suited to follow. Looked at another way, this helps ORION assess where the student is in her Zone of Proximal Development, as per Figure 1, and which questions or problems would be most appropriate to present to the student to keep her engaged and continually learning.

ORION also provides students with information about how they are progressing and what they need to study or practice more.

ORION in its adaptive algorithm, also takes into consideration the degree of confidence a student reports she has with her mastery of a certain topic. A student who is very confident, yet does not display mastery of the material, is presented with different problems and challenges than a student who is not confident but displays more mastery. Students are able to see how their performance matches their confidence in a unique "meta-cognitive" graph displayed by the system.

In addition to adaptation to the student, ORION provides the instructor with teacher reports that allow her to review quickly how all of her students are doing, to be sure they are progressing as expected. These reports in ORION can also help the teacher revise classroom lectures accordingly. For instance, if most of the students in a class are struggling with a particular topic, that topic can be the emphasis and focus of the teacher's next lecture. To tie this back to our instructor, with ORION's reports she would be able to see that Giselle is actually quite proficient at math and her performance is very high when she has the opportunity to move through the material at her own pace. Michael, too, excels when he has access to more of the material and can work through practice items as he learns. Mary, the most advanced student, is presented with challenge material. This additional insight makes all of the difference to the instructor who can then focus her time more effectively.
ORION is domain-independent adaptive learning software. This means that the software is not designed and developed to work for a particular academic subject or area of content – an important feature of software that one hopes can be generally used across topics. Currently, Wiley has customized ORION to work for students in the domains of anatomy and physiology, psychology, introduction to business, and financial accounting. Wiley plans to extend ORION to many more domains, including many in the life sciences, physical sciences, business, accounting, mathematics, and social sciences.

ORION's Adaptive Algorithms

An adaptive learning system like ORION must first assess what a student knows before presenting instructional material to help the student learn. This is typically done by giving the student a test (or tests) that assesses the skills and/or knowledge the student has. The results of the tests are accumulated and represented in what is known as a "student model" (VanLehn, 2006; Woolf, 2009), a representation of the state of that student's knowledge, broken down into components of knowledge within a particular domain or subject.

ORION uses item response theory (IRT - Fox, 2010; Hambleton, Swaminathan, & Rogers, 1991) and its variant, continuous response to estimate a student's skill proficiency, across a variety of skills. IRT is based on the idea that the probability of a correct response to an item, which can be thought of as a specific "skill", is a function of how a particular student has so far performed (e.g., correct responses on similar and related items) and item parameters (e.g., difficulty of the item). The IRT item probabilities are effectively used to estimate a student's skill proficiencies. Using IRT, the ORION system has produced an overall prediction accuracy of an estimated 75%-80%.

In order to be as precise as possible about what a student really knows, or the skills they have, ORION takes multiple factors into account aside from correct or incorrect answers. For example, the system factors into student proficiency whether they request hints (indicating they still need help to demonstrate this skill), when they only get a question partially correct (in the case of a multi-part question), and when they express a lack of confidence in their answer. (Students can self-report a confidence level, and less than "Fully Confident" can indicate guessing.)

It is also important to break down a student's skills and knowledge at different levels of granularity, something that ORION is also designed to do. Such a representation allows an adaptive learning system such as ORION to know what a student knows at different levels of specificity. This is achieved in Wiley's ORION system by breaking knowledge down according to various curricula levels, from largest grain size to smallest grain size, at the level of the course, sections of the course (chapters, corresponding to books), and learning objectives, including the ways in which those individual learning objectives interact with each other.

1. Course
2. Chapter
3. Learning Objective (what a student should be able to do at the end of a section)
4. Enabling Learning Objectives (more granular steps that facilitate the overall learning objective)

In summary, ORION uses a sophisticated, statistical approach to estimate skill proficiency based on a variety of information, including how a student has done on related material, his or her confidence in provided responses, and whether or not he or she has requested help. This leads to the possibility that performance on related topics, tackled in completely different courses, can impact skill assessments in a new course. Estimates of skill proficiency, in turn, guide ORION in selecting the right practice questions or instructional material to present to a student, aiming to stay within the student's "zone of proximal development."

Using Adaptive Technology in the Classroom

A Look into the Future

If we fast forward to look at our instructor's class in the future it might resemble something like this: All students have a persistent digital profile that allows instructors to better understand who is coming into their classes, and at what level of proficiency. In fact, they may be grouped by need: high achieving students with language difficulties might be one sub-section; poorly motivated students with only one math requirement might be another. Students may be offered a choice of an on-site or online course, or a combination of both but no matter what their preference, they will be able to access material that moves at their own pace, with the instructor available to remediate, to challenge, and to motivate.

We also see a future in which more and more of a student's instruction will be available in informal settings in digital fashion, on devices the student keeps with him and accesses "on the move." Certainly the smart phones and tablet devices are the obvious first steps in this direction – and work in this direction is already well underway (Echeveria et al, 2011; Sharples et al, 2009; Vavoula, et al, 2009) – but there is likely to be other "wearable" technology that will also be part of that transformative future (see, for instance, Google Glass). Although the mode of delivery will be dramatically different from today's print and common digital books, adaptive learning technology can apply just as easily to a variety of emerging instructional technologies.

A key trend is the use of "big data" in supporting adaptive learning software. Educational Data Mining (Baker and Yacef, 2009; Romero and Ventura, 2010) and learning analytics (Ferguson, 2012; Siemens, 2010) are two emerging fields of research that, together, are setting the agenda for a future in which everything the student does – not only their academic achievement but their preferences, likes/dislikes, motivations, friendships, etc. – will play a role in adapting instructional content. Researchers are already able to tell when students are "gaming the system" – engaging in unproductive learning behavior, such as rapidly moving through a curriculum – through the use of machine-learned classifiers that use various data to identify behaviors (Baker et al, 2006). All of this wide variety of collected data – and copious amounts of it, collected from the increasing number of students engaging in online learning of many types – will eventually allow us to better calibrate and adapt our educational technology and to develop new and better adaptive technology.

How will ORION capitalize on this trend towards using big data to improve educational technology? Fundamentally, ORION's approach is already capitalizing on these developments. A statistical approach such as that taken in ORION relies, in essence, on past data and performance and how that helps to assess current status and future possibilities. Yet the types of data being evaluated now by educational data miners – preferences, motivations, connections to others has not been leveraged until now in ORION. We see a future in which ORION, and its successor technologies, uses all of this data to adapt instruction to the student, not only to their skills and knowledge, but also to "softer" aspects of learning, such as preferences. It may be used then not just to improve the experience for learners, instructors, and institutions, changing and adapting the ways in which content is delivered, but may fundamentally change the way learning materials are created. This will allow ORION to ensure that only those materials that have the most impact on learning are delivered to the learners, and that all learning assets can be tailored by individual, by class, by institution, by language, by learning style. The possibilities in the future are endless.

References

Aleven, V.A.W.M.M.,& Koedinger, K. R. (2002). An effective metacognitive strategy: Learning by doing and explaining with a computer-based Cognitive Tutor. *Cognitive Science*, 26(2).

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*.

Baker, R.S.J.D., Corbett, A.T., Koedinger, K.R., Evenson, E., Roll, I., Wagner, A.Z., Naim, M., Raspat, J., Baker, D.J., Beck, J. (2006) *Adapting to when students game an intelligent tutoring system. Proceedings of the 8th International Conference on Intelligent Tutoring Systems*, 392–401.

Beal, C. R., Walles, R., Arroyo, I., and Woolf, B. P. (2007). Online tutoring for math achievement: A controlled evaluation. *Journal of Interactive Online Learning*, 6, 43–55.

Bloom, B. (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).

Bloom, B. (1985). *Developing talent in young people* (1st ed). Ballantine Books Publishing.

Conati C. and Manske M. (2009). Evaluating adaptive feedback in an educational computer game. In: *Proceedings of IVA 2009, 9th International Conference on Intelligent Virtual Agents, Lecture Notes in Artificial Intelligence* 5773. Springer Verlag, p. 146–158.

Corbett, A. T. & Anderson, J. R. (1992). The LISP intelligent tutoring system: Research in skill acquisition. In J. Larkin, R. Chabay, and C. Scheftic (Eds.), *Computer Assisted Instruction and Intelligent Tutoring Systems: Establishing Communication and Collaboration*. Hillsdale, NJ: Erlbaum.

Corbett, A. T., Koedinger, K. R., & Hadley, W. H. (2001). Cognitive Tutors: From the research classroom to all classrooms. In Goodman, P. S. (Ed.) *Technology Enhanced Learning: Opportunities for Change*, (pp. 235–263). Mahwah, NJ: Lawrence Erlbaum Associates.

Dweck, C. (1986) Motivational processes affecting learning. *American Psychologist*. 41(10), 1040–1048.

Echeverría, A. Nussbaum, M., Calderón, J. F., Bravo, C., Infante, C., and Vásquez, A. (2011). *Face-to-face collaborative learning supported by mobile phones.*
Interactive Learning Environments Vol. 19(4), Pages 351–363, 2011

Eliot, C., & Woolf, B. (1994). Reasoning about the user within a simulation-based real-time training system. In *Proceedings of the fourth international conference on user modeling,* 121–126.

Fain, P. New Player in Adaptive Learning. *Inside Higher Ed*, July 29, 2013.

Ferguson, R. (2012). The state of learning analytics in 2012: A review and future challenges. Technical Report. Knowledge Media Institute: The Open University, UK, 2012.

Fox, Jean-Paul (2010). *Bayesian item response modeling: Theory and applications*. Springer. ISBN 978-1-4419-0741-7.

Gee, J. P. (2008). *What video games have to teach us about learning and literacy*, revised and updated, Basingstoke: Palgrave Macmillan.

Graesser, A. C., Chipman, P., Haynes, B.C., and Olney, A. (2005). AutoTutor: An intelligent tutoring system with mixed-initiative dialogue. *IEEE Transactions in Education*, 48, 612–618.

Habgood, M.P.J., and Ainsworth, S. (2011). Motivating children to learn effectively: Exploring the value of intrinsic integration in educational games. *Journal of the Learning Sciences*, 20/2, 169–206.

Hambleton, R. K., Swaminathan, H., & Rogers, H. J. (1991). *Fundamentals of item response theory*. Newbury Park, CA: Sage Press.

Koedinger, K. R., Anderson, J. R., Hadley, W. H., & Mark, M. A. (1997). Intelligent tutoring goes to school in the big city. *International Journal of Artificial Intelligence in Education*, 8, 30–43.

Lederman, D. Affirmative action, innovation and the financial future: A survey of presidents. *Inside Higher Ed* . March 1, 2013.

Miller, G.A. (1956), The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological Review*, 63, 81–97.

Mostow, J., and Beck, J. (2007). When the rubber meets the road: Lessons from the in-school adventures of an automated Reading Tutor that listens. In B. Schneider and S.-K. McDonald (Eds.) *Conceptualizing scale-Up: Multidisciplinary perspectives* (pp. 183–200). Lanham, MD: Rowman and Littlefield.

National Governors Association Center for Best Practices and Council of Chief State School Officers. *Common core state standards initiative: Preparing America's students for college & career*, 2010.

Pintrich, P. (2000). An achievement goal theory perspective on issues in motivation terminology, theory, and research. *Contemporary Educational Psychology*, 25, 92–104.

Rickel, J., and Johnson, W.L. (1999). Animated agents for procedural training in virtual reality: Perception, cognition, and motor control. *Applied Artificial Intelligence,* 13, 343–382.

Romero, C. & Ventura, S. (2010). Educational data mining: A review of the state of the art. *IEEE Transactions on Systems, Man, and Cybernetics - Part C: Applications and Reviews, Volume 40*, Issue 6: 601–618.

Sharples, M., Amedillo Sanchez, I.; Milrad, M. and Vavoula , G. (2009). Mobile learning: small devices, big issues. In: Balacheff, N.; Ludvigsen, S.; Jong, T. de and Barnes, S. eds. *Technology Enhanced Learning: Principles and Products*. Heidelberg, Germany: Springer, pp. 233–249.

Siemens, G. (2010). *What are learning analytics?* Elearnspace, August 25, 2010.

Stokes, P. Learning to Adapt. *Insider Higher Ed*. April 4, 2013.

Sweller, J., Van Merriënboer, J., & Paas, F. (1998). Cognitive architecture and instructional design. *Educational Psychology Review* 10 (3): 251–296.

VanLehn, K. (2006) The behavior of tutoring systems*. International Journal of Artificial Intelligence in Education*. 16, 3, 227–265.

VanLehn, K. (2011). The relative effectiveness of human tutoring, intelligent tutoring systems and other tutoring systems. *Educational Psychologist*, 46, 4, 197–221.

VanLehn, K., Lynch, C., Schulze, K. Shapiro, J. A., Shelby, R., Taylor, L., Treacy, D., Weinstein, A., & Wintersgill, M. (2005). The Andes physics tutoring system: Lessons Learned. *International Journal of Artificial Intelligence and Education*, 15 (3), 1–47.

Vavoula, G., Sharples, M., Rudman, P., Meek, J., and Lonsdale, P. Myartspace: Design and evaluation of support for learning with multimedia phones between classrooms and museums*. Computer & Education*, vol. 53, Pages 286–299, 2009.

Vygotsky, L.S. (1978). *Mind and society: The development of higher psychological processes*. Cambridge, MA: Harvard University Press.

Woolf, B. (2009). *Building intelligent interactive tutors: Student-centered strategies for revolutionizing e-learning*. Morgan Kaufmann, Burlington, MA.

Dr. Bruce M. McLaren is a Senior Systems Scientist at Carnegie Mellon University who is passionate about how technology can support education and has dedicated his work and research to projects that explore how students can learn with educational software, in particular, software that runs on the world-wide web. He is particularly interested in intelligent tutoring systems, e-learning principles, and collaborative learning. Dr. McLaren has over 100 publications spanning peer-reviewed journals, conferences, workshops, symposiums and book chapters.

Other contributors to this whitepaper include Tara Hein-Phillips of Vestar Consulting and Barbara Heaney.