

- New learning media and technology offer tremendous possibilities for improving learning outcomes and increasing productivity and scale. Myriad approaches across disciplines are being implemented and tested across the entire spectrum of two- and four-year colleges and universities and K-12 schools in the U.S. and around the world.

- Proven applications such as Carnegie Mellon's OLI courses incorporate feedback loops based on ongoing collection of all student learning data, clear analytical reports to guide teachers, and the integration of the cognitive sciences to create powerful, interactive learning environments.

- The George Lucas Educational Foundation focuses on telling stories through film to improve the educational process. One of the foundation's primary initiatives is to produce documentaries on successful innovative uses of technology that enhance deeper learning, by combining the development of academic skills and social and emotional learning skills.

- Leading universities in the United States are not just making their course content available through initiatives such as MIT's OpenCourseWare and Yale Open Courses, but, importantly, they are also experimenting with potentially game changing new learning media and technology for their own students. Adoption of new approaches at these universities harbors the potential to influence the direction and pace of adoption throughout higher education.

Effective Use of New Learning Media and Technology

STEVE ARNOLD, THE GEORGE LUCAS EDUCATIONAL FOUNDATION

KEVIN GUTHRIE, ITHAKA

CANDACE THILLE, CARNEGIE MELLON UNIVERSITY

New learning media and technology are being used and tested across the entire spectrum of two- and four-year colleges and universities and K-12 schools. Common goals of these widely varied efforts include improving learning outcomes, increasing efficiency and productivity, and making knowledge more accessible. Steve Arnold, vice-chair of the board of the George Lucas Educational Foundation, Kevin Guthrie, president of ITHAKA, and Candace Thille, director of the Open Learning Initiative (OLI) at Carnegie Mellon University, discuss the work of their organizations. Arnold describes the Lucas foundation's efforts at the K-12 level to identify, capture and share "what works in education" via its website edutopia.org. Guthrie reviews the case studies of online learning initiatives undertaken by leading universities, as described in *Unlocking the Gates*, written by Taylor Walsh for ITHAKA. Candace Thille explains the power of the data gathered through tracking all student interactions with OLI's courses. The data are part of a feedback loop of continuous improvement of the courses and, likewise, teaching and learning outcomes. All three organizations are engaged in research to assess the effectiveness of their work, so as to further a revolution in education enabled by new learning media and technology. Excerpts of Arnold's, Guthrie's and Thille's remarks at the Forum's June 2011 Aspen Symposium are reprinted here.

Steve Arnold

In the early 1980s I became convinced that personal computers might change the way kids learn how to think. So I started trying to understand what was going on with personal computers, and that led me to a career over the last 30 years or so in the technology business. I first worked at Atari and then ran a few divisions of George Lucas' company—and that's how he and I came to work together on his educational foundation. Then I ran a private company for Bill Gates for a little while, then started a venture capital firm with a couple of other partners. So I've been investing in early-stage technology businesses for the last 20 years or so.

But I've maintained a very strong interest in education, partly because of my early

background in developmental psychology, and because when George Lucas and I started thinking about the applications of technology to education in the early- and mid- 1980s, we discovered that we shared a strong passion for not only figuring out how to make education better, but also for trying to understand the ways that technology might be helpful.

Today, the George Lucas Foundation identifies great programs in the real world that embody important principles in education, and then we capture them on film and promote them on our edutopia.org website. George is unusual in the sense that he truly is deeply involved and passionate about telling stories through film to improve the educational process. He's definitely an engaged participant. For the last 10 years or

so, we've been focusing on documentary filmmaking around what works, what we call the "new world of learning." That new world combines academic skills and a higher order of interpersonal skills—social and emotional learning skills—that help kids become well-rounded and able to work with others. We look for innovative uses of technology, not for its own sake, but in the context of learning. We're also involved in research into what we're calling "deeper learning" and what happens between kids, teachers, and the learning environment that makes a difference and fosters that.

Edutopia is focused on six core strategies, including:

- integrated studies;
- project learning;
- technology integration;
- teacher development;
- social and emotional learning; and
- comprehensive assessment.

We make information related to these core strategies available free online. We look for accessible programs that are actually using these practices in the real world. We deconstruct the projects and identify what's interesting about what's happening, why it works, how it's generalizable, and how can somebody else do it.

George considers what the Lucas Foundation is doing to be "arming the foot soldiers of the revolution." What that really means is giving the people in the classroom and in the schools today an understanding of what models they might use, and the ways that they might innovate in their settings. Forty-four percent of the roughly 600,000 people a month that go to the Edutopia web site are classroom teachers. We have roughly 900,000 to a million page-views a month at Edutopia, and hopefully people are using what we're creating in ways that are constructive.

Much of the national debate on educational reform today is often focused on political issues and organizational issues such as charter schools and teachers unions. Those are important issues, but at the end of the day, after we figure out funding and the employment structure, we still need to know what needs to happen between kids and teachers, and what learning environment is best for kids. Edutopia focuses on the relationship between the learner and the teacher—what it is that teachers need to know to be able to do their jobs differently and most effectively.

Another aspect of the Lucas Foundation's work is research. In addition to telling the stories, we identify and collaborate with researchers to investigate some of the fundamental principles that we're advocating for. The first project that we've undertaken is co-funded by the George Lucas Foundation and the Bill and Melinda Gates Foundation. We're collaborating with researchers at the University of Washington's LIFE Center

to rigorously assess the effect of project-based learning in upper-level high school courses. The focus is on AP courses, not because we're endorsing AP courses in particular, but because they are widely accepted and standardized courses. For the last three years, we've refined the curriculum for AP American Government, and taught kids in a variety of settings using the project-based learning approach and compared it with the traditionally taught course on two dimensions. First, can they do as well or better on the AP test? Turns out they can. But we also want to know if they can demonstrate a greater sense of deeper learning through a different kind of assessment, one that allows the kids to demonstrate their ability to use the information they've learned in the course creatively.

So far this is a very promising project. We've done it with American Government, and are now in the process of re-designing AP Environmental Science. We're moving with the redesign of some other courses as well, to assess whether this project-based model of instruction can be both rigorous and more deeply informative and useful to kids.

Interesting story about one of the byproducts of this course: We did some fishbowl interviews with these kids at the end of one of the courses, and asked them to talk to us frankly about their experiences. It was literally a jaw-dropping experience for some of the people who had been teaching high school kids for 30 years, because kids said things like, "You know, this class really annoyed us because in the beginning you didn't tell us what we were supposed to know. And usually in this school you tell us what we're supposed to know, and then you tell it to us, and then you ask us to tell it back. But you didn't do that here, and it made us angry. And then we figured out we have to teach each other, and then we knew how to do it." Several of the students were able to describe a different understanding of their own learning. We also got feedback from kids who went off to college and said that the project-based AP course was the one that prepared them best for college because they knew how to collaborate with other people, how to ask hard questions, and when nobody's giving them answers, how to figure out what they're supposed to do.

It's very promising early work. We haven't drawn any final conclusions yet, but we have indications that are very promising.

Kevin Guthrie

For those of you who may not be familiar with ITHAKA, it's an independent, not-for-profit organization established originally with grant funding from the Mellon, Hewlett and Niarchos Foundations. Our mission is to help the academic community use digital technologies to preserve the scholarly record and to advance research and teaching in sustainable ways. Today I am going to talk about the impact of these technologies on

teaching and learning, and some of the early projects that are using these technologies in new ways.

In his Romanes Lecture at Oxford in 2000, Bill Bowen said:

...for the purposes of this talk, I use digitization to mean the electronic assembling, disassembling and transmitting of the basic elements of intellectual capital—these include words, sounds, pictures and data. The ability to take these sources apart, send them easily over distances and reconstruct them renders the walls around universities far more porous. Once those walls are pierced in this way—that is to say, once both the basic materials and the fruits of the work of academic institutions are easily gathered and sent, the very currency of the university becomes dramatically more accessible, and these institutions find themselves drawn increasingly into the realm of commerce.

Without question, face-to-face interaction remains important, but we also have to remember that the core mission of these institutions is built around the dissemination and the creation of knowledge. The ability to disseminate information and knowledge has moved to the network, and that has huge implications because not only does it draw universities towards commerce, it draws commerce into universities. That is revolutionary.

Unlocking the Gates was written for ITHAKA by Taylor Walsh and published in early 2011. It's about how leading universities in the United States, and in India, too, are opening up access to their course content by making it available online. It includes seven case studies. Taylor Walsh did a brilliant job writing the book. I've talked to people at the various places that were studied, and they all say that what is written is accurate and consistent with their recollections, that it's excellent at representing what actually happened. I should say that the case studies don't necessarily fit together thematically, largely because they represent activity during a time of experimentation, so institutions were trying different things. So, the purpose of the book is not to make an argument for a particular approach to online courseware or online learning; rather, it attempts to document how these projects came to be and how they were executed. The book chronicles the development of each of the seven initiatives. The goal was to provide something actionable, to put senior administrators in the space of those cases and see why decisions were made as they were, so they could learn from those experiences and apply those ideas in their own circumstances at their own institutions.

I'm going to give you just some quick highlights of the different cases so that you understand what's in them. First, for context, there are two cases that weren't part of the open

education movement, but preceded it. Fathom, which was at Columbia, and AllLearn, which was led by Yale but also included Princeton, Stanford and Oxford. These projects were essentially reactions to the commercial hysteria of the dot.com bubble of the late 1990s. The concern was that universities were going to miss out commercially if they didn't get online.

Fathom was started up as a for-profit entity. AllLearn was a not-for-profit. Both tried to create charging models to provide access to university-created content to the general public. Neither was successful or survived, and some people think that's because they were charging for access. I think that that's a part of it, but a lot of it was just not getting to a constituency that valued enough what they were doing. In the case of AllLearn, it was about trying to reach alumni with its courses. With Fathom, it was a bit unclear who they were trying to reach. In both cases, the content was provided for self-directed education; no credential or other certification was to be offered for using or reading material on the sites.

MIT OpenCourseWare is another case that's profiled in the book. As MIT OpenCourseWare was getting started, it was clear that AllLearn and Fathom weren't doing well. MIT also was evaluating whether to get engaged in online learning and whether to attempt to make money from doing so. After considerable debate among faculty, it was decided that there was much more to be gained from sharing MIT course content without charge, if it could be done, than from trying to sell access to it. Chuck Vest, MIT's president at the time, considered it a much better fit with MIT's mission and history to find a way to make the MIT course materials available for all the world.

There's been a lot of confusion about MIT's motivation for OpenCourseWare. It's actually not online learning. There is no instruction provided. It simply makes course materials available online. Chuck Vest's vision for it was as a kind of a B2B application—teacher to teacher. He thought MIT could help improve teaching by providing instructors with high quality raw materials for instruction. These teachers would then have a head start when building their own courses. The idea was to provide support to that process. One of the unintended consequences of putting these materials online, though, was that they went right to end-users and allowed people to learn by going through the content in a self-directed way, which was a great outcome, but that led some people to conclude that MIT intended for the materials to be courses, which they were not. Fundamentally, it was about putting materials that were already being created on the web.

An important point to note about all of these entities is that their motivations were, for the most part, supply-side-originated. They didn't define a need and then try to figure out

how to address it. In that sense OpenCourseWare was almost purposefully not self-disruptive, which I'll touch on in a bit.

I'm not going to talk about Carnegie Mellon's Open Learning Initiative because Candace Thille is here to present about that, except to say it's very different from the other cases because it is not a project to provide access to course materials, it is a new kind of online infrastructure to support a new kind of teaching and learning environment. The system is intended to provide a full suite of what it takes to learn: content, assessments, and evaluation, along with feedback loops that help the student know how he or she is doing, feedback loops that help the teacher evaluate the student, and feedback from the overall system to cognitive scientists who then can improve the system.

The Yale Open Courses is our fifth case in the book. The primary motivation behind it was to represent Yale globally, to reach out beyond the campus borders and portray Yale around the world in a particular way. Courses to be included in the program were selected strategically; they were the most popular courses taught by Yale's star faculty, and they were produced at very high quality. The approach was designed to introduce people around the world to the very best that Yale had to offer on campus. In contrast with Yale's approach to select courses to represent itself, MIT's OpenCourseWare goal was to provide access to every single course.

The sixth case in the book is from University of California at Berkeley, a project called `webcast.berkeley`. It was the brainchild of a single faculty member who wanted to, back in the early '90s, make it easier for his students to get access to what he was doing in class. It was not a project with top-down strategic oversight and guidance; it was a grassroots effort. So at very low cost he put up a video camera up, taped his class, and made it available to the students. He didn't have the IT resources to attempt to restrict access to only his students, so he put it on the basic network without passwords or authentication. He wasn't ideologically trying to make his course open to the world—it was just too expensive to develop the systems to restrict it. So it was available to the world. Students liked the convenience of being able to see the lectures when they wanted, and the idea caught on. UC-Berkeley started to fix cameras in different locations, and more courses became available. In that sense, selection of what courses became available was not strategic, it was really a function of whether there were cameras in the room where the class was being taught. If you are a professor and there is a camera in the room in which you are teaching, the default is that the course will be recorded and available unless you decline. So the motivation and goals of Berkeley's effort were completely different from the others; it was just a very basic, low cost effort to provide a convenience to students.

The seventh and final case is very different from the others, for two reasons. First, it is not a study of a single institution providing instruction or courses online, it is an effort to approach the challenge from a systemwide level. Second, it is in India, which poses a whole other magnitude of a challenge. They have an enormous number of students they want to teach, and they have very few institutions in which to teach them. Moreover, there are a small number of institutions that are incredibly strong, but the gap between those institutions and all of the others is very, very large. They are therefore trying to figure out if online technologies can help them extend the high quality educational experience to more students. The case documents how they are working at the system level to put in place technology that can expand the reach of the great IITs and IIMs to that vast majority of schools that don't have access to the resources of the top schools. The case documents the challenges they are facing as they try to address a problem of great scale.

What would I say are the findings based on these case studies? First, the specific context and goals within the particular environment are deterministic of what the entities end up looking like. One of the difficult things about addressing these issues at your institutions is the complexity of both the activities that you have on your campuses and the complexity of your mission. It becomes very difficult to say what you are really trying to accomplish. MIT OpenCourseWare was very successful because it was able to identify clearly what it aimed to accomplish, and it went after that with a laser focus. They also created a common vocabulary on the campus for the faculty, and got everyone motivated. Yale's focus was to extend Yale's reputation. Berkeley's was to make access to course lecture content more convenient. All of these initiatives have succeeded in part because of the simplicity of their basic objectives.

The other point I would make is that it's been extremely difficult to measure the impact of these projects. We hear the same sort of statement in every single case: "It is hard to measure the impact, but we get a lot of emails from people who really love us." It's anecdotal. We believe that it's impactful, but it's very hard to measure the impact. And unlike a for-profit organization that typically will focus on identifying a demand and developing the services for that and keep iterating on it, these were in large part supply-side efforts in response to grant funding or just the opportunity to make available something that was already being done.

Let me return to a comment I made earlier; that MIT's approach is nondisruptive. It is an important point as leaders think about these kinds of projects. In developing the strategy for OCW, MIT identified two things that are special about MIT: 1) the credential; and 2) getting people together on their campus with other talented students. MIT decided not to

award credentials. In their minds, they were giving away the part of the institution that didn't create a threat or that would be competitive. I should add that all of these entities I have discussed do not offer credentials, which is important. Yet the people involved with AllLearn and Fathom say the number one reason they really couldn't get traction was because they couldn't offer the credential.

The last and final important point to be made about these efforts is that none of them has developed a self-generating financial model to sustain their initiatives on an ongoing basis. The primary sustainability model they've pursued is to try to get the project to become so important to the campus itself that it becomes a budget line item. MIT's aim was to make OpenCourseWare so important internally that there would be no question that the institution would support it. They've accomplished that to a considerable degree, and they've established an amount of money that they're willing to put into it, equal to about half of its costs. But the other half has to be raised. The worldwide impact and awareness of MIT's OCW makes a compelling argument for continuing support. Similar indirect benefits can be cited for the other projects as well, but as the budget situation at colleges and universities continue to grow more difficult, will it be possible to continue to invest in these efforts where the benefits in large part accrue to people outside that university's community?

Before concluding, I'd like to describe very briefly the work ITHAKA is doing in what we call "interactive learning online (ILO)." By ILO, we mean highly sophisticated, interactive technologies in which instruction is delivered online and is largely machine-guided. The best of these systems rely on increasingly sophisticated forms of artificial intelligence, drawing on usage data collected from many students to deliver customized instruction tailored to individual students.

These systems are increasingly immersive and merge the network and infrastructure, the pedagogy, and the content. Traditionally, those three things have been in three different silos, provided by three different providers. If these systems take hold, infrastructure, pedagogy and content will all be delivered as part of one system. This is inherently disruptive and existing providers see the challenge. Educational publishers like Pearson, for example, are aggressively entering the online learning market. Why? Because the first thing people think about online learning is they won't have to buy

a textbook. As this convergence increases, it is going to have a dramatic impact on the way your institutions operate, because the roles of the faculty and every system designed to support them will be affected.

ITHAKA is doing two studies now. One is a rigorous side-by-side comparison of learning outcomes—a randomized, highly academic study of an OLI statistics course provided at eight selective campuses. Students will be randomly assigned to a traditionally taught version of the course or to the online course, and we'll measure both the outcomes and the cost of

delivery. We've already completed two pilots working out the logistics, which are difficult. The actual study will begin this fall and it should be very interesting. Bill Bowen is leading that project for us.

The second project is a qualitative study of the barriers to adoption of these new systems at U.S. colleges and universities. Led by Larry Bacow, the recently retired president of Tufts University, we will interview a cross-section of institutions, some that are deeply engaged in online learning and some that are not, to ascertain what the barriers are to their using these new forms of more sophisticated OLI-type systems. Why, for example, are faculty going to support or resist new initiatives? What will it take to overcome these challenges? We expect these studies to be released in mid-2012.

Candace Thille

I'm going to start with a quote from Clay Christensen, from his February 2011 article, "Disrupting College." Christensen wrote,

"Changing circumstances mandate that we shift the focus of higher education policy away from how to enable more students to afford higher education to *how we can make a quality postsecondary education affordable.*"

Two years ago here in Aspen, Joel Smith and I presented about the Open Learning Initiative. We framed the conversation around the productivity question, and looked at Baumol-Bowen's cost disease. For those of you who don't remember, Baumol-Bowen's cost disease explains how difficult it is to increase productivity in labor-intensive endeavors. The classic example that Baumol and Bowen used was the horn quintet, saying essentially that the quintet's productivity cannot be increased without decreasing quality. The other example they used in that original article in 1967 was higher education.



The best of these [interactive learning online] systems rely on increasingly sophisticated forms of artificial intelligence, drawing on usage data collected from many students to deliver customized instruction tailored to individual students.

Joel and I posited that that was a false dichotomy, that higher education does not have to make a choice between either high quality and low productivity, or high productivity and low quality. We presented the work we're doing at Carnegie Mellon, where we're showing that we can create systems that increase productivity and increase or maintain quality. We do so by using enabling technologies. The enabling technology is not just the online technology, but also our increasing understanding of how people learn. We also recognize that context is important; that is, by creating environments in which students learn to synthesize all they're learning and apply it to real-world problems. The OLI environments are designed not to drill content into students, but rather to support students to achieve higher level learning outcomes.

We've done studies of our learning environments. One of the studies of our chemistry course showed that the number of actions the students engaged in in the virtual lab was the most predictive factor of learning gain. The virtual lab is a big, open exploratory space. Exploratory actions in the virtual lab blew out all other predictors, including SAT score and gender, in predicting learning gain in chemistry.

The real key, though, lies in mining the data and the feedback we can get from tracking students' interactions in the online environment. By giving good, timely feedback to the students and to the instructors, we can increase productivity. Traditionally, people have tried to solve Baumol-Bowen's cost disease in higher education by using the same technology that they used to solve the cost disease for the horn quintet, which was to widely disseminate high quality recordings of the quintet. The claim I'm making is, while it is fine to record and disseminate lectures, the performance of the lecture is not the

service that higher education provides. The service that higher education provides is to support the change in the knowledge state of the learner. It's what students do in practice that changes their knowledge state and makes a difference.

So what's happened since we presented here two years ago? First, global higher education continues to record and distribute lectures and continues to create open content. Many institutions are getting into making content openly and freely available, which is a wonderful thing. But there are the other pieces of education that people are now stepping in to unbundle and produce as well, outside of the higher education system. There are new open-learning networks freely available. A network called "open study," for example, was developed so that people using the free, open content could form peer-to-peer study groups where they can learn from each other in a networked community.

The other major development in the past two years is credentialing—the other service that higher education provides. There are new systems out there to provide credentials. For example, Peer 2 Peer University's new School of Webcraft offers badges for knowledge that will be accepted by firms to indicate that the holder knows web design.

That's what the world's been doing. What we've been doing at OLI is several new projects with many new partners. We're no longer just doing development with Carnegie Mellon faculty at Carnegie Mellon. We're working with faculty domain experts across multiple kinds of institutions, including a big mix of community colleges and land grants, to create shared OLI courses. We're also partnering with ITHAKA on evaluation studies, and we're continuing research on how to improve the collection and representation of student learning data to give feedback to students and instructors, which I consider crucial.

Figure 1: Screen Capture of a Section of an Instructor Learning Dashboard from an OLI Statistics Course

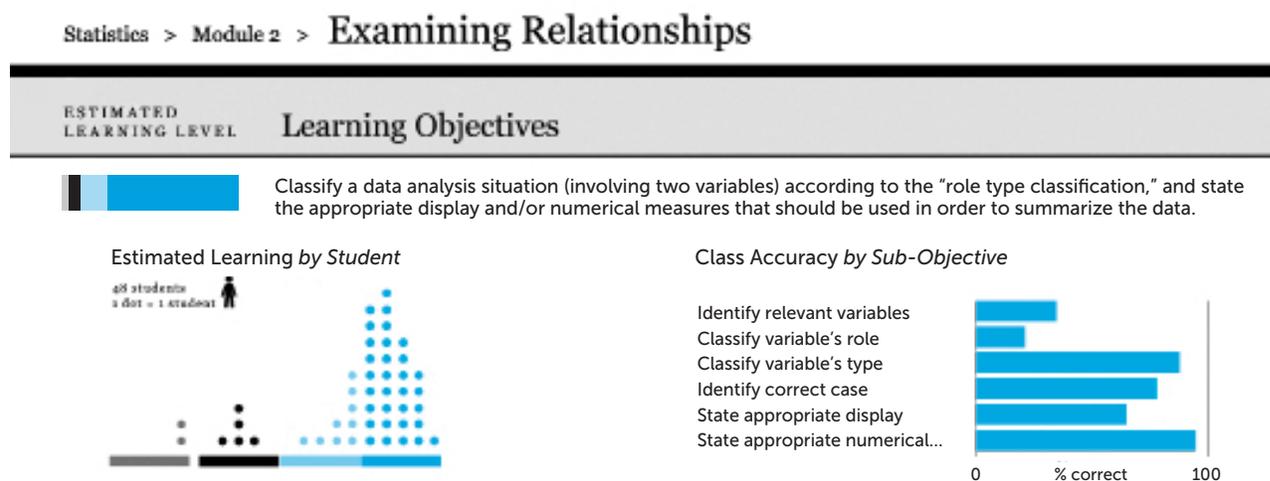


Figure 1 is an example of one representation of the student learning data that the OLI system presents to instructors: the OLI Instructor Learning Dashboard. The Instructor Learning Dashboard is a research project that is being led by Dr. Marsha Lovett, a cognitive scientist in our Eberly Center for Teaching Excellence.

When we design a course, we articulate student-centered, measurable learning outcomes that we want the students to achieve. Next we think about what kinds of activities students can engage in to help them to achieve those outcomes. We collect all the data about what the students are doing, and we make inferences from that data about what they're learning. We then have the learning scientists analyze the data from the students' interactions with the environment to create the underlying cognitive models that drive the feedback reports.

If I'm an instructor, I can have my students work through, say, module two of the OLI course in statistics. Then, before class, I can check the instructor feedback report to see how my students are performing on each learning outcome. Each horizontal bar indicates the current state of my class on each outcome. We can predict that students in the dark blue part of the bar would do very well on an assessment of that outcome today; students in light blue we predict would do OK; students in black we predict would really struggle; and the students in grey haven't done enough work interacting with the environment for us to make any kind of reliable prediction. Instructors can get more detailed information by clicking on each horizontal bar.

Each one of the dots represents a student in the class. The instructor could click further to find out who those students are, and what they are doing that's driving that prediction. The instructor can drill down even further to view information about specific subskills and activities to develop a detailed understanding of what the students understand and where they're having trouble.

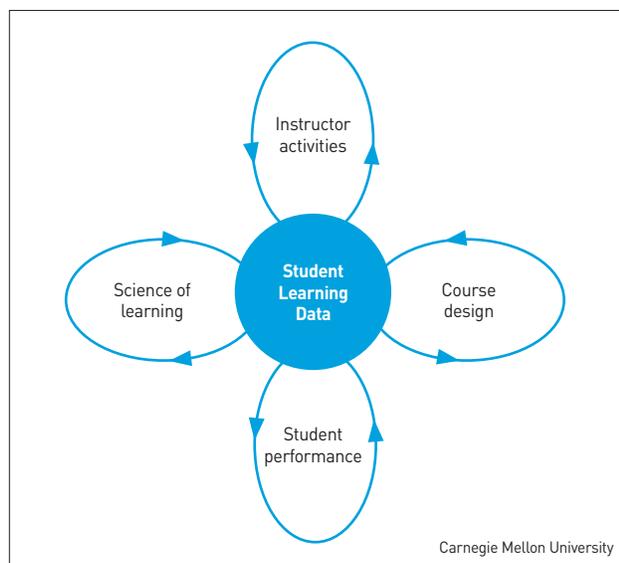
The killer app, as shown in Figure 2, is collecting and representing student interactive learning data to provide actionable information for the actors in the system and guide what they're doing.

We're fortunate that the sister project to OLI is the Pittsburgh Science of Learning Center whose whole purpose, as an NSF-funded learning research center, is to define and refine theories of human learning. They use the OLI courses as part of their research environment. Historically, we unpack expertise by applying learning science techniques, one of which is a cognitive task analysis. That's usually a very labor-intensive effort involving structured interviews, protocols with experts and novices, factor analysis, etc. It's very expensive to do—and part of the reason why building effective learning support systems is so expensive.

The data we mine from OLI courses, though, can be used to do this cognitive task analysis far more efficiently. We can discover new cognitive models and create visualizations that aid our human scientists. The Pittsburgh Science of Learning Center is creating analysis tools and software that can analyze this data and propose knowledge models that fit the data, as well as the models proposed by human cognitive scientist.

We can also use these data to look at other factors besides cognitive tasks. Researchers at the Pittsburgh Science of Learning Center have developed computer tutors that look at help-seeking behavior. How do you teach students to seek help ap-

Figure 2: The "Killer App" Feedback Loops for Continuous Improvement



propriately? And how do we give students feedback on that? We can also use machine-learning algorithms as detectors of motivation, reflection and affect. Clearly, the data are a tremendous resource. We're talking about business intelligence systems for education.

The real challenge is that in order to get to better teaching practice or better tools to support teachers and students, we need much better learning theory. In order to develop better learning theory, we need a lot more data about many more students learning in many different contexts. We can use our current learning theory to build the best educational technology, and then we can use the power of the educational technology to collect data, and use that data to refine our theory. Setting up this virtuous cycle is what is going to allow us to increase both quality and productivity in higher education.

Back in 1967, Baumol said, "Without a complete revolution to our approach to teaching, we cannot go beyond [current levels] of productivity." My message is, such a revolution is clearly

possible. But my question to you is, who's going to lead it? My more pointed question is, can the not-for-profit higher education lead it, or will the for-profit sector be out front on this?

Discussion

Q: People in the K-12 space have been thinking about how to get at the kind of data-driven instruction that you presented, Candace. One of the issues is that there's a lot of training—regular weekly training with master teachers—that needs to happen for the new approach to be really affective in the classroom. My question for you is, what are you learning in the college context? And what lessons can we take from K-12, if any, on this?

Ms. Thille: I can speak to what we're learning in the college context. The statistics course I mentioned was used in about 50 or 60 different classrooms at different institutions last academic year, ranging from Santa Ana Community College to Bryn Mawr—a nice swath of institutions. We do virtually no training. We'll do a webinar or sometimes we'll do a workshop if they're participating in a study.

But for the most part a faculty member can just come to our site, look through the course, and decide to use it. They send us an email asking for an instructor account. We check to make sure they're an instructor and give them an account. Then they use our web-based interface to select and sequence the content so that they can customize it. Then they create a course key for their students to use when they log on so the system knows to put those students in that person's instance of the course so that they are reflected in that instructor's dashboard.

How faculty use the course and reporting system is completely up to them. We have found that it takes about a semester to really start to shift faculty behavior. At first, faculty think of it—and I think it's probably a good way to think about it—as a replacement for the textbook. They can give their students the OLI course, and they don't have to spend \$150 on a statistics textbook—that's a win. So that's what they do. Maybe they look at the instructor reports, maybe they don't. Maybe when they look at the reports and they don't have a clue what to do with the information. So they might just still give the same lecture on week three that they always gave on week three. But then they start to see a difference in their students' performance on their assessments, and they start to use OLI differently, and get the full benefit

of the learning environment and the data available on the instructor reports. We are now looking at what sorts of additional support we should be providing instructors.

Q: I would like to pursue the issue a little further, Candace, about the cumulative data on learning effectiveness, if I can call it that, and how having cognitive science faculty examine those patterns can lead to changes that will enhance and improve the learning. The notion of continuous improvement by integrating the cognitive science assessment of a growing amount of data holds enormous potential.

Ms. Thille: Thank you. You said that very succinctly. That's exactly what I mean. The way most people are thinking about online education right now is as a distribution mechanism. That's part of it, but the real power, I believe, is in this data collection piece, in the information we can gather to refine our understanding of how people learn, and use that to progressively improve how we're delivering our service. I believe that is the only way we're going to be able to address issues of scale. Kevin talked about the issues of scale in India; we have them here in the United States too. How are we going to increase productivity without doing that? I can't think of another way.

Q: I want to mention an unintended consequence of OCW at MIT that's becoming clearer and clearer. That is, professors started looking at the course materials of the professors for the courses that they were teaching prerequisites for. It can be very hard to align those courses, but now they're looking at each other's materials, and alignments of the prerequisite courses has improved. OCW is having a very interesting effect. The faculty are learning from each other, for example, realizing how another faculty member is teaching differential equations and supporting that with adjustments in their own teaching.

Q: There are other unintended consequences. One is that it's become a very important course guide for the students to understand what the courses are teaching and how to align their own programs. And then there's the competitive piece, which is that when all the faculty material is posted, everyone wants theirs to look as good as everyone else's. Part of it is the improved alignment, and part of it is just that good old-fashioned competitive spirit.

Q: Does it have to be a competition between the not-for-profit and the for-profit sectors in providing robust online learning environments, or is there crosstalk? With all of the



The real power [of online education], I believe, is in this data collection piece, in the information we can gather to refine our understanding of how people learn, and use that to progressively improve how we're delivering our service. I believe that is the only way we're going to be able to address issues of scale.

proprietary software that's being developed by the for-profits, there must be some pretty powerful learning tools out there. Are you engaged in any conversations with those organizations and developers?

Mr. Guthrie: That's a really important question. This shouldn't pivot around ideological issues between for-profits and non-profits. One of the key issues is when you start thinking about Google or Netflix or Amazon, these enterprises build scale on their data and they improve their services based on their data. When they become powerful—think of Facebook—and they have huge constituencies, that data drives improvement in a way that that drives off competition.

If some organization owns all the data accumulated from all the students behaving inside these environments, and can pull those data into their system, that is a very powerful place to be. I think the real question is, what if that happens? That's a challenge. In the short run, I think there are a lot of presidents who would say they would rather see the commercial marketplace provide these products because they think that they'll get a better product faster. But there are some issues around the long-run motivations of commercial enterprises and their need for growth and scale and potentially domination of a relatively small market that I think people should think about.

Q: So I want to see if I understand the way that feedback works: If you had a student and the student has a unique identifying number in the system, and say the student takes three or four OLI courses, it seems to me that at the end of the third or fourth course, you're going to have a profile of the way in which that student learns best. And if you can accumulate that, there are important things you can tell that student, or that maybe the student will be able to see for herself, that would make for better course selection—whether, for example, to take a project-based course or a more of a theory course. Isn't that ultimately what this feedback would relate to?

Ms. Thille: Yes, and I would put a refinement on that, that is, this notion that students have a particular modality in which they're going to learn best all the time, it's not really an attribute of the student. It is a complex interaction of that student's prior knowledge, relevant skills, future goals, and the challenges of the domain. There are many factors that go into the decision about the best move for a student at any given point in their learning process, to help them move along some trajectory. And that's why it's a science—because knowing what is best to do to support a learner is a very complex thing.

One of the challenges that we've had—and part of the reason we do higher education the same way we've done it for 200 years—is that we're told that we have to do 14 different things differently, to radically change what we're doing, but we don't know what the mechanisms are, so that's all we can say.

It's almost hopeless to say, take this practice that this person is doing this way and implement it over here with absolute fidelity—that's not going to happen, and maybe it's not even necessary that it does happen.

The categories that people came up with like visual learner, auditory learner, were a way of trying to manage complexity. Information technology allows us to manage complexity in a very different way so we don't have to artificially categorize students and say, OK, you learn this way so this is what you get. That's not personalization. Personalization is really having a bead on where that student is at this point in time, and what in this moment is the next best action for that student to take.

● **STEVE ARNOLD** is CFO and vice-chair of the board of the George Lucas Educational Foundation. In the early 1990s, he served as vice president of Broadband Media Applications at Microsoft Corporation, and as president and CEO of Continuum Productions (now Corbis), a private company founded by Bill Gates to pioneer the creation of large digital libraries for online distribution. Prior to Continuum, Arnold served as vice president and general manager of LucasArts Games and Learning divisions, and vice president of the New Media Group at Lucasfilm Ltd. He co-founded Polaris Venture Partners in 1996, and focuses on investments in information technology and digital media. Arnold can be reached at stephen.arnold@edutopia.org.

● **KEVIN GUTHRIE** is president of ITHAKA. He was the founding president of both JSTOR (1995) and Ithaka (2004). JSTOR and Ithaka merged in 2010 to form a new organization, ITHAKA. Previously, Guthrie started a software development company serving the needs of college and professional football teams, and later served as a research associate at The Andrew W. Mellon Foundation, where he authored *The New-York Historical Society: Lessons from One Nonprofit's Long Struggle for Survival* (1996). Guthrie has also been a professional football player, sports broadcaster, and producer. Guthrie can be reached at kg@ithaka.org.

● **CANDACE THILLE** is director of the Open Learning Initiative (OLI) at Carnegie Mellon University, a position she has held since the program's inception in 2002. She is also co-director of OLnet, an open educational research network, a collaboration between Carnegie Mellon and the Open University, UK. Thille serves as a redesign scholar for the National Center for Academic Transformation; as a fellow of the International Society for Design and Development in Education; and on the Global Executive Advisory board for Hewlett Packard's Catalyst Initiative. She also serves on a working group of the President's Council of Advisors on Science and Technology (PCAST) to write a report for the Obama Administration on improving STEM Higher education. Thille can be reached at cthille@cmu.edu.