

The Holy Grail of Education: Personalized Learning

by Maria H. Andersen, Learning Futurist

The LIFT Institute at Muskegon Community College

A Case for Disruptive Innovation

Humans have always been learning, but the learning methodologies have changed over time. The earliest means of education were highly personal: oral histories passed from adults to children, informal or formal apprenticeships, and one-on-one tutoring all took form in the history of most cultures. Only in relatively recent history have formalized systems of mass education (aka industrialized education) been designed.

Industrialized education worked so long as the students were highly motivated to learn and get ahead of their peers. As Clayton Christiansen pointed out in his book *Disrupting Class* (2008), “prosperity can be an enemy to motivation.” In a developing country, success in education translates to larger steps on the ladder of personal economic prosperity. When the relative size of extrinsic gains from education are smaller (developed countries), Christiansen argues that the key to learning is that it must be intrinsically motivating. With education still designed for delivery en masse and learners in need of intrinsic motivations, there is quite a mismatch between the system and the participants. To reengage students in developed countries, we will have to customize learning to the personal interests and motivations of each individual. Indeed, research supports the effectiveness of personalized learning. In 1984, educational researcher Benjamin Bloom found that if the average student were tutored one-

on-one, they outperformed 98% of the students who were learning via conventional methods (this is referred to as Bloom's 2 sigma problem).

The vision of an education built around personalized learning is not new, but it is definitely tantalizing. Neal Stephenson's novel, *The Diamond Age*, shares a vision of personalized learning in the future via an interactive book that possesses a conversational interface and "pseudo-intelligence," a kind of artificial intelligence inferior to human intelligence. It's likely that we'll see decent conversational interfaces within the next decade, and certainly applications like Google Voice are moving us much closer to this reality. However, an artificial intelligence (AI) capable of directing the learning needs of a human will take much longer, developing in the next 20-50 years. The needs for personalized learning exist in the here and now. So how does one bridge this vision of the future with the realities of the present?

Let's start by taking stock of what we currently have. We already have software that stores the content we like (e.g. Evernote, Posterous), and software that merely stores the *location* of that content (e.g. Diigo or Delicious). Even traditional media, like books, have parallel digital systems which allow for notetaking, highlighting, and bookmarking (e.g. Kindle, Nook, or iPad). While it's useful to store and search information, I would venture that we rarely go back to look at the information we mark for storage. This is a problem, because for deep learning to occur we need to have repeated structured exposure combined with time for reflection - a chance for information to become knowledge, understanding, and wisdom.

Current educational systems provide students with repeated structured exposure required for deeper learning, but fall short at providing it in the personalized manner necessary

for motivating students. Educational institutions form a vast interconnected network, and while small changes can occur within the system, individual parts only have the ability to flex within their existing boundaries. For personalized learning to take hold in education, disruptive change will be required.

Consider some of the disruptive changes we've seen in the last six years that affect what we do with our free time. Facebook, now with 500 million users, has disrupted normal social interactions in a little over six years by simply asking us to answer the question: *What's on your mind?* Twitter users now send more than 50 million tweets per day, answering the question: *What's happening?* Big news stories now break first on Twitter, in real time, with eyewitness accounts. As big as Twitter is, there are now more people playing Farmville (a social media game on Facebook) than there are active Twitter users – a fact that has not gone unnoticed by game designers and educators. These Farmville players are choosing to spend their free time for collaborative activities (their “cognitive surplus”) plowing virtual soil and planting virtual crops.

These innovative social disruptions have happened quickly, but not from within the existing organizational structures. For example, Facebook did not disrupt phone communication by changing the nature of phone calls or phones. Facebook built an entirely *new* system that eventually circled back around to phones by the way of phone apps. In the same way, the trick to developing a personal learning system is to abandon thinking about how to build it from *within* the existing educational system and to begin pondering how such a system could be developed outside of education.

A Simple Idea: Learn This

Let me propose a realistic scenario of what a true personalized learning system might look like and how it would function. We first have to create (1) a new layer of learning media in the background of the existing Internet and (2) an ecosystem of software to easily manage the learning media we engage with. In the same way we've integrated buttons like Twitter's "Tweet this" and Facebook's "Like" at the end of videos, articles, and other media, imagine we now add a button that says "Learn." Clicking this button (anywhere you find it) would bring you into an interface to help you learn the content.

We don't need artificial intelligence to begin this journey. In the first version, learning should simply be by way of Socratic questioning, where questions are rephrased as answers. Such questions would be used to analyze concepts, to prod at the depth of knowledge, and focus on concepts, principles, issues, or problems. Socratic questions are elegant because, unlike other formats (e.g. multiple-choice), the learner must self-generate the answers rather than relying heavily on the ability to recognize a correct answer when they see it. The personal learning system would use a spaced repetition algorithm (SRA) to reintroduce the Socratic questions over time so that biological memory is more likely to grasp onto the ideas and information. For now, let's call this system SOCRAIT (a play on "Socratic" that includes "AI" within its text).

For example, suppose I read an article outlining the recent US Justice Department's decision about digital rights management (DRM, i.e. jailbreaking iPhones), and I decide that it's important for me to remember the generalities of this decision. At the end of this article, I choose "Learn" to add a question to my SOCRAIT question bank. Two options would appear:

(1) write your own question and (2) choose from a list of questions written by others. If I choose the first option, I might write a simple question and answer for myself: “On July 27, 2010, the US Justice Department made a decision about DRM. What was the decision and what are the implications?” Following this, I’d write a short summary or clip a few sentences of content from the article to summarize the findings. Along with the question and answer, SOCRAIT would save the source URL (link to the content) and I could tag the question with metadata tags I indicate (e.g. Copyright and DRM).

Later in the day or the week, when I have some down time, I could reengage with SOCRAIT. Here’s how it would work: I read or listen to a question, answer it in my head or out loud, view or listen to the answer, rate my understanding, and go to the next question. There is no need to develop software to verify the answers. Since the learning is tailored to intrinsic motivations, the learner could rate their own ability to answer a question (e.g. 1 = I have no clue, 2 = I knew some of it, and 3 = I nailed it!), and SOCRAIT could make decisions based on these ratings. If your rating of understanding is low or spotty, the system would offer to send you back to the original media for another look.

With a rudimentary CI (like the one implemented in Google Voice), there’s no reason why the SOCRAIT interface couldn’t be vocal and available anywhere we interact with computers (e.g. cell phones, tablets, auto navigation systems). This would allow us to improve our learning while performing other tasks: driving, making dinner, or walking the dogs. Initially, the so-called “Pareto’s Vital Few” (the 20% of people who get 80% of the work done) would be the ones who would be most interested in creating and engaging with questions. But as the connectedness of the system matures, the need to write your own Socratic questions would

lessen. Authors and media creators would write their own questions, targeting comprehension of important ideas and facts. Media consumers would be able to choose from a list of questions, perhaps seeing a sorted list based on their indicated learning priorities. Two readers of the same article would see different questions at the top of their “suggested questions” based on tags on the content. In some cases, the user might choose to pay for curated or reputable content so that their learning can later be certified by an employer, educational body, or organization.

Implications

Now let’s take a step back and look at the big picture. Any content that exists on or is connected to the Internet would be tagged with Socratic learning questions and metadata for subjects. Each learner would have their own bank of questions, personalized to their own learning interests. This system I describe is does not have artificial intelligence (at least, not initially). It can’t rate how well you’ve learned the content. But this is one of the first steps on the path towards the intelligent conversational interfaces of the future.

Instead of learning that is designed around a physical place (i.e. schools), an educational space (i.e. learning management systems), or a person of authority (i.e. instructor), this system is designed around the learner. It goes without saying that the implications for education are huge. In the space of a few years (think: Facebook) we could develop a completely separate content learning system that’s incredibly flexible and personalized to the interests of the learner. The architecture needs to develop organically around web-based content and grow tendrils into everything we produce in the future. It will take some time to go back and create a

learning layer to overlay all the content that we already have. But we have seen from projects like Wikipedia, which became the largest encyclopedia ever assembled a mere six years after its creation, that there are people willing to contribute their time and energy to these kinds of tasks. For example, Clay Shirky estimated that Wikipedia was built with the time of less than 1% of the time that Americans spend watching TV every year.

Modern journalism has been struggling with a problem of income stream, and while revenue has shifted to online advertising, it is not enough to shore up the industry. At present, the vast majority of Internet content is free and, as Chris Anderson argues in his book, *Free*, it's not likely to change. How do you get readers (or viewers) to pay for something that they already get for free? Add something to the content that's not already there. If readers or viewers had the ability to quickly add reputable questions to their learning bank, this would be a value-added service. Cleverly, the media content would remain free, but access to the question bank would require a one-time payment or ongoing subscription by the consumer.

A New Learning Ecosystem

Books like *The Shallows*, by Nicholas Carr, cause us to question whether we might be trapped on the information superhighway – stuck on the line between data lanes and unable to scoot forwards or backwards. Twitter users regularly use the phrase “drink from the firehose” when referring to their experience of dipping into the live data stream. Information, whether it be from radio, television, print, web-media, or social networks, is coming at us too quickly and all most of us can do is surface-skim this information, rarely pausing to reflect or think deeply.

To learn, to analyze, to innovate, and to think creatively we must internalize some of the information we process.

An entirely new ecosystem could grow up around this Socratic learning system.

Certainly a ratings system for question authors could be built using the technology developed by companies like Netflix. “Your friends John and Iveta chose this question. Would you like to see other questions/media they chose for this topic?” If you chose to do so, the questions you see when you add content to your question bank could be filtered by your existing social networks. Rather than showing all the possible questions in existence for that media (which could become a fairly lengthy list), you could choose to see only the ones people in your social network have also used.

So far I’ve discussed how the system would work if you engaged in reading and watching media as you do today. However, such a system could also shift how and when we seek out content. After all, a lot of time is wasted in modern education by reteaching content that *some* of the learners already know. There is no incentive for a student to get ahead when the reward is sitting through a lecture on something they’ve already learned. Imagine: when you need to learn something new, you could subscribe to a curated collection of questions on that topic. For example, “Digital Rights Media 101” might be a collection of questions developed by somebody who teaches digital copyright to beginners. The truly fascinating shift is that you wouldn’t necessarily start by consuming the media that goes with the questions. Instead, you would simply start answering the questions in your bank. As you encounter learning questions that you can’t answer, you could dive into the content at those points in time – this is the exact point between boredom and frustration – the balance point to engage in learning.

Almost immediately after the personalized learning architecture is in place, we will need a new educational industry tasked with certifying knowledge and understanding. For lack of a better name, let's call these folks "Socratic scholars." Their job will be to rate how well you know what you claim to have learned. For example, let's say I've engaged with and theoretically learned 500 tagged questions on biochemistry to prepare for teaching a new class. In order for this to count towards professional development hours, my college asks me to certify the learning. I pay for a Socratic scholar who specializes in chemistry to rate my knowledge. We meet either in person or via the web (more likely) and have a discussion about the questions in my learning bank on biochemistry.

The scholar has access to the 500 questions I say I've mastered and randomly asks me to answer a random selection. Of course this is where it would be valuable to have reputable questions in my learning bank (from authors, researchers, scientists, leaders in their field). Since the scholar can see both my questions and the answers (linked back to original content) it should not be difficult to ascertain whether I have, in fact, mastered the knowledge and concepts as I have claimed. Because the certification is human-to-human, and not human-to-machine, the nuances of human language would be understood. So if the language of the verbal answer and the language of the written answer don't match up exactly, that's not going to be a problem. At the end of the session, the scholar would "grade" my understanding of the 500 questions on biochemistry and I could provide this certification to the human resources department.

In many respects this is a much better system that we have today. For most certification of learning, we simply look at a transcript. If the class is listed, we assume the

learner has that knowledge. Of course, we all know that knowledge ages. The fact that I got a chemistry degree in 1996 does not mean you would want to hire me as a chemist today.

Ideally, you'd want me to recertify before I entered the "chemist" job pool. Biological memory is not reflected in the metrics of transcripts or grade point averages. I should pause for a moment to clarify that I'm not saying that content knowledge equals ability to function as a practitioner in the discipline. Even a diploma only indicates that the educational system has walked you through some series of appropriate paces for the discipline. Skills like critical thinking and creativity are often lost in education (especially in the STEM fields) because there is such an incredible amount of content to cover. If the content knowledge moved outside the educational system, educators could focus on the skills that surround technical knowledge instead.

Let's imagine what would happen if a robust Socratic learning system was at the heart of the educational system? A learning coach (which has replaced the teacher or instructor) will designate some core of material they want you to learn. For example, in calculus I might use a set of 500 curated concept-oriented questions from a well-known Calculus textbook author, each question linking to supporting media. Every student would be working on this set of questions, and so, as a learning community, we'd all work on that all together. Hopefully this doesn't sound like too radical a departure from normal. This is where it changes: because every student has different interests and career ambitions, I would also require that each student find an additional 100 questions tagged with both "Calculus" and questions that are of interest to them. For a student studying to be a doctor, questions tagged with "medicine" or

“epidemiology” might be appropriate. For a student going in to business, questions tagged with “marketing” or “management” might be more appropriate.

As the learning coach, I might help students search for good questions, help them to understand the questions they find, and foster discussion with other students on these topics. When it comes time to certify the learning for each student, it is done by an oral interview in which the learning coach has access to the common questions and the personalized questions for each student. Even if the learning coach isn’t an expert on all the personalized questions, the answers are provided and the content is related to a subject they do have expertise in. Again, the learning coach would only ask about a random selection of questions. At the end of the semester, each student has learned their own personal version of Calculus, while still learning a core of common material.

Such a system has implications for lifelong learning “on the job” too. Instead of holding mandatory training, a human resources department could push out a bank of Socratic questions to all their employees about safety, new initiatives, mission statements, etc. For example, to train employees on OSHA compliance, the employees would be invited to add a curated list of 40 questions about OSHA policies, each question leading back to source media that provides the necessary content to answer the question. After two weeks, someone in HR can act as the Socratic Scholar and spend 5 minutes with each employee to test their knowledge of the policies with a random selection of questions.

A Game Layer for Learning

The futurist John Smart writes about a coming “valuecosm” on the 5-10 year horizon, when we’ll be able to program our apps or avatars to make decisions for us based on what we say is our set of values. The real question is whether learning can become one of our new values, especially in the U.S. In 2009, The U.S. Bureau of Labor Statistics estimated that the average American adult spent more than 5 hours per day on leisure activities (close to 3 of that watching television) and about 30 minutes per day on educational activities. Given the 10:1 ratio of leisure to educational activities, is our culture likely to embrace learning as a choice? Initially my answer was no, but then I began to think about video game design.

Seth Priebatsch spoke at TEDxBoston (2010) about building a “game layer on top of the world” – what if one of the game layers we create surrounds learning? The same game dynamics that are used to build successful video games (e.g. appointment, influence and status, progression dynamics) could be deployed to make *learning* the game itself. While this might still be a hard sell for adults, there will be subpopulations, like the early technology adapters, that will see the immediate value in cultivating and learning from their own question banks. Children who grow up learning with a Socratic question system might gain learning values naturally, and would hopefully carry these to their adult lives.

A successful Spaced Repetition Socratic Learning System (SRSLs) would have to entice you to keep to specific goals, like answering 50 questions per week or answering 100 questions with a certain tag in the next month. Any of these goals could be incentivized with points (1 question answered correctly = 1 point), incentive rewards for meeting certain goals (you’ve earned your Silver Calculus badge for 100 questions learned), and social status levels (Maria has

just become a Calculus Master, can you do it too?). Those engaged in formal education would participate with a far greater intensity of daily questions than those who are post-formal education. However, the wise worker would continue to learn, albeit at a slower pace. Resumes would boast levels of knowledge on particular topics, and stats on the intensity at which you participate in learning.

Let's Build It

A diploma has become a social signal to stop learning, and in today's world, where technical knowledge doubles every two years [citation?], this is absolutely the wrong thing to do. Careers shift overnight and industries collapse rapidly. We have to learn, and learn faster than we ever have before in order to stay ahead of the problems we are now creating. The *content* for a spaced-repetition Socratic learning system already exists; it is the architecture and interface we are missing. This new learning media needs to be an interconnected network of user-generated or author-generated Socratic questions with an interface that manages them. The architecture needs to remain open so that anyone can create questions on any content, and any developer can build applications for computing devices.

A system for personalized learning will not grow from inside formal education. Education is like a field that's been overplanted, with only patches of fertile soil. Too many stakeholders (parents, Unions, administration, faculty) compete with each other with various ideas about how to change, acting like weeds or plagues that choke off plant growth. The fresh and fertile soil of the open web can foster the quick growth of a personalized learning system.

Then, like a good fertilizer, it can be used to replenish the soil of formal education, and helping us to reach that “Holy Grail” of education ... personalized learning for all.

About the Author

Maria H. Andersen is the Learning Futurist for The LIFT Institute at Muskegon Community College. She has degrees in Mathematics, Chemistry, Biology, Business and (ABD) Higher Education Leadership. Maria is considered an expert in educational technology, and has been studying, researching, speaking, and writing about the future of education and learning for several years.