Adaptive Practice™

High-Efficiency Game-Based Learning

Thomas Butt

January 17, 2012

As technology plays a growing role in our society, education becomes increasingly critical to success, both individually and globally. More education is required to meet the minimum standards of knowledge-based jobs, and more frequent education is required to keep up with constant change. Now more than ever before, we need to find more effective ways to educate ourselves. Using the latest advances in fields like neuroscience, cognitive science, psychology, and gaming, Adaptive Practice offers a new approach to learning that significantly improves learning efficiency. In this paper, some of the basic principles of Adaptive Practice technology are described along with the underlying theory.

The Future of Learning

Imagine that you could learn anything you want by playing a computer game. And imagine that the game adjusts as you learn so that it keeps you interested, curious and challenged. And although it's challenging, it's never frustrating or difficult; in fact, you really enjoy playing. And when you finish playing you're surprised at how much you've learned. Topics you used to find very difficult now seem remarkably simple.

This vision describes the mission of Adaptive Practice technology - to make learning easy, fast and fun. Adaptive Practice uses game dynamics and the most up-to-date understanding of human learning to make learning remarkably efficient. Why is Adaptive Practice so effective? The rest of this paper will explain.

Learning, Thinking, and Remembering

To start, let's define what learning actually is. Learning is the process of creating memories that can be applied to new situations. Or said another way, learning is remembered experience. As humans we have the ability to learn from past experiences so we can be better equipped to handle future experiences. Learning involves storing information in our long-term memory so that it is readily available when we need it. To raise the efficiency of learning, we need to find the most efficient way to do this.

In his book, Why Don't Students Like School?, Daniel Willingham, offers the following idea - "memory is the residue of thought." With this quote he's making a really important point, namely, that we remember what we think about. If we don't think about something, we won't remember it, and what we remember will depend on what we thought. If you and I watch the same movie, our remembered experience, and therefore our learning, will be different, because even though we saw and heard the same thing, what we thought about will have been different. Good teachers understand that it's not what they say in the classroom that matters, it's what the students are thinking. So if learning comes from thinking, it's critically important that students are thinking about what you want them to learn when they are engaged in any learning activity. This is more easily said than done.

The good news is, we all like to think and learn. We're naturally curious. It's part of what makes us human. It gives us the ability to function in the world, and it's really quite remarkable. The bad news is, that thinking is hard, and we only do it successfully under certain limited conditions. To better understand those conditions, it's important that we distinguish thinking from remembering. Thinking involves placing your attention on something and reasoning about it, for example, by comparing it to something else. That something may be a concrete item like a ball, an idea like democracy, or a process like how to tie your shoes. Remembering, on the other hand, is simply recalling something, and while thinking is difficult and slow, remembering is quick and easy, at least it is when you've got a good strong long-term memory of something.

One of the really interesting aspects of learning is that although we need to think to learn, when we know something really well, we can use this knowledge without

thinking. When you first learn something, like tying your shoes, you think carefully about the details, but after a little bit of practice, something really interesting happens. You can tie your shoes without having to think at all. It becomes more or less automatic. What's more, you can do this while you're thinking about something else entirely. When you are really proficient at something, it's not because you have honed your thinking skills, it's because you have developed a lot of automatic memories so you don't have to think so much.

Another interesting thing occurs as we learn. The more things we know, the more easily we can learn new things. The rich get richer, so to speak. This is because our existing knowledge provides a framework for understanding new knowledge. In recent years, cognitive scientists are becoming more convinced that learning actually improves intelligence.

Practice - The Magic Pill

We've described learning as the process of building strong long-term memories. But what's the best way to build strong long-term memories? The answer is simple, practice. Practice when properly done reinforces memories, increases the speed of recall, and makes the memory last longer. Really strong memories can last a lifetime and be instantly recalled without apparent thought or effort. What's more, practice is the only reliable method for building strong memories.

Is practice the same thing as drill? The answer is no. Drill implies mindless repetition, and although repetition is an important element of practice, it's not the only element. To be effective, practice can't be mindless. When you practice, you must be thinking about what you want to learn and remember. Here's an example. I'm sure you've seen thousands of pennies over the course of your life and have seen them repeatedly over many years. Now here's a quick question? Is the head of Abraham Lincoln facing to the left or to the right? If you're like most people you're not really sure. How can that be with all the repetition? The answer is that you've never really had to think about this aspect of a penny, so you didn't develop any specific memories for this aspect.

Another important characteristic of effective practice is action. Although you might strengthen a memory by watching the same training video over and over again, it would not be a very efficient way to learn. One reason is that even with the best of intentions, you'll find it nearly impossible to stay mentally engaged for any length of time. We'll talk about why in the next section, but for now, it's important to know that you learn more quickly and effectively when you're doing - involved in an activity that requires action and thought.

The effectiveness of practice is also strongly influenced by timing. Studies have proven that practice is much more effective when it is spaced out over time. In other words, cramming is a terrible way to learn. This process of carefully spaced out practice is often referred to as spaced repetition. By spacing out practice, we minimize the amount of practice needed to develop strong memories.

I like to think of exercise as the magic pill for health. If someone invented a pill that had all of the health benefits of exercise,

that person would be immediately wealthy and the pill would be an overnight sensation. The problem is that exercise takes effort and for many people the inertia is just too difficult to overcome. In much the same way, practice is the magic pill for learning. To get really good at anything requires a lot of practice. Great athletes, who are the best in the world at what they do, consistently stress the importance of practicing the fundamentals. Those who excel at anything have found ways to make consistent practice enjoyable. This is one of the key characteristics of Adaptive Practice which we'll talk more about later.

The Goldilocks Point - The Sweet Spot of Learning

So, to learn anything we need to think about it, but thinking is hard, so we avoid it unless the conditions are right. If that's true, then how can we ever learn anything? The answer is that under the right conditions, we actually enjoy thinking and learning. What are those right conditions? We have to find the situation novel enough to be interesting and at the same time simple enough so that we are confident we can understand it. I like to call this the "goldilocks point". Not too easy. Not too hard. Just right. It represents the sweet spot for maximum learning. It's the balance between boredom and frustration and when new information strikes this perfect balance, we engage mentally and enjoy learning. When it doesn't, we rapidly disengage.

Try to teach me something I already know and I'll quickly disengage. I'll label your presentation as boring and I'll start thinking about something else. Give me too much new information at one time and a similar thing will happen. I'll label your presentation as too difficult or frustrating and my mind will wander. This process of assessing the novelty and difficulty of new information happens automatically, rapidly, and continuously, and is largely involuntary. You might have my attention and curiosity right now, but one quick drift from the goldilocks point and I'll just as quickly zone you out. And because we are always learning, the goldilocks point is always moving. Something that was novel 10 seconds ago is suddenly old news.

What's worse, our ability to pay attention to information that is not at the goldilocks point is severely limited, even when we really want to. Have you ever tried to pay attention to something that you found very complicated? It's virtually impossible. The reason why comes down to a serious limitation that we all have, namely the amount of space in our working or short-term memory. Even the brightest among us can only hold a few things in our awareness at any one time. When we are learning something new, this space is quickly filled with details about what we're learning and how it fits with what we already know. Like it or not, when we learn, we must do so one tiny bit at a time.

In addition, the goldilocks point is a very personal thing. Your point is not at the same place as my point. The strength of each of the thousands of tiny memories that represent our understanding of any topic are certainly different for every individual, and these differences greatly impact our ability to learn something new.

Flow

The goldilocks point is the sweet spot between boredom and frustration where our interest and attention is captured. When we are engaged in an activity that manages to keep our interest for a sustained period of time, we sometimes enter a very pleasurable mental state called flow. Athletes will sometimes describe this state as being in the zone. When we are in a state of flow, we become so engrossed in what we're doing that we lose track of time. We look up and are startled to see that several hours have passed even though it seemed like just a few minutes. Most of us have had this experience.

Mihaly Csikszentmihalyi, world renowned professor of psychology, is best known for his work on the concept of Flow, which he describes in detail in his book, *Flow, the Psychology of Optimal Experience*. Here is how he described flow in an article in Wired magazine (September 1996) - "being completely involved in an activity for its own sake. The ego falls away. Time flies. Every action, movement, and thought follows inevitably from the previous one, like playing jazz. Your whole being is involved, and you're using your skills to the utmost."

How can we help someone achieve a state of flow? Here is how it is described in the Mihaly Csikszentmihalyi page on Wikipedia - "To achieve a flow state, a balance must be struck between the challenge of the task and the skill of the performer. If the task is too easy or too difficult, flow cannot occur. Both skill level and challenge level must be matched and high; if skill and challenge are low and matched, then apathy results." Sound familiar? The conditions for entering the state of flow sound a lot like the goldilocks point and the implications for learning are remarkable.

We mentioned that learning requires thinking, thinking requires attention, and attention is drawn by the goldilocks point. It's not surprising then, that these same conditions are required for entering this zen-like state of flow which might be characterized as sustained attention. This relaxed, focused state is ideal for learning and remembering, but it gets even better. Mihaly Csikszentmihalyi believes through his research that people are happiest when in a state of flow. As a result they are intrinsically motivated to create and maintain these states for as long as possible. If we can get students in a state of flow while learning, both the motivation to learn and the learning outcomes improve.

Limitations of Traditional Learning Methods

Now that we understand some of most important conditions for optimal learning, let's explore some of the challenges that we face when using traditional learning methods. These challenges, as we'll see, make it virtually impossible to maintain optimal learning conditions for any individual student for any length of time. The amount of inefficiency in this process is both troubling and exciting. It's troubling that we continue to use a system with so many obvious flaws, but the degree of improvement that we can expect in the not too distant future is really exciting.

First, it is incredibly difficult to know in advance what a student might find easy or hard. Memories are complex and as we discussed, are continuously evolving. Each learning objective that we identify in a les-

son plan involves building hundreds and perhaps thousands of separate tiny memories. Teachers use their best judgement based on how well they know their students, but even the best teacher simply cannot know in great detail what a student actually knows.

Second, even if you get the "Goldilocks" point exactly right for an individual, classes by design have more than one student and those students won't all be at the same starting place. Every student is different and what they know is equally different. To make matters worse, classes are typically organized by age, not by knowledge, skill, or ability. The bottom line is that a lesson created for any large class is sub-optimal for every student in the class. Teachers try to "differentiate" their teaching to accommodate these differences, but for practical reasons, can only do so in a very limited manner.

Third, even if every student in the class started at the same place, every student learns new material at a different pace. The speed and the manner in which you learn new material is again a highly individualized trait. The process of thinking, fitting new ideas into what we already know, and reaching those "aha" moments of learning is a process that must occur in each individual. The timing of this process and the path it takes is likewise unique. Even if every student in a class learned at the same pace "on average", they would differ moment to moment. So a class that started together at the beginning of a lesson would very quickly fall out of sync.

Fourth, even when the conditions are perfect, some students won't be paying attention some of the time. None of us are robots. Sometimes, probably more often than we care to admit, our minds wander. If a teacher is presenting a lesson, then information that is presented when someone is not paying attention would need to be repeated for that person. Since in any typical class there is always someone not paying attention, material must be continuously repeated. While it's true that repetition is good for learning, this type of repetition is highly inefficient. And even with the repetition, some students will miss some material some of the time just because of the randomness of their attention.

Fifth, even when they're paying attention, the process of thinking which you are striving to achieve will by necessity divert their attention. Ok, let's suppose you've gotten everything else right and you've finally gotten the students to think about the right things so they learn the right things. Well, guess what, when they're thinking about those things, they are no longer paying attention. Why? Because they can't keep that many things going on in their mind simultaneously. This loss of attention might be more noble than that caused by thinking about what they're going to do after school, but the effect is the same.

The cumulative effect of these problems is that with traditional learning, optimal conditions do not exist for most students most of the time. This is depressing, but before we talk about solutions, let's consider one other standard practice that makes things even worse.

If it is so hard to get students to learn, what do we do when they don't? What happens when a student "completes" a

year, or a semester, or a day of school without learning what they were supposed to learn? The answer in nearly every school is this - they move right along as if they did! If we thought teachers had a hard job trying to get students to learn before, let's now use a curriculum that assumes that every student in the class has learned all the prior material, even if many of them haven't.

Optimal Learning

Now that we understand a bit more about learning and some of the challenges that we face using traditional methods, let's talk about how we might create an optimal learning tool.

Our design goals are simple. We want to make learning fast and fun. To be fast, students must learn in the least amount of time, with the least effort. We want to respect the value of every second of time that a student devotes to learning. If it is possible that a student can effectively learn something in 10 seconds, it shouldn't take 11. Of course it's also important to make it as fun as possible so that students want to spend more time learning.

As we've seen, learning does not depend on what we teach, it depends on what students think. To maximize the speed of learning we must capture and maintain students' attention and get them to think about what we want them to learn. If we can sustain this for a while under just the right conditions, they will enter the magical state of flow where learning is optimized.

How can we do this? As it turns out, video games provide the answer. Think about how video games work. They present the player with a challenge and provide tools to be used in overcoming the challenge. The challenge automatically becomes more difficult as the player's skill improves. To maintain the player's interest, video game designers understand that they must balance boredom and frustration; in other words, they must find and maintain the goldilocks point.

The best kind of video games for learning are those that keep a student continuously active, where they must respond rapidly, every few seconds during play. There are several reasons for this. First, the more frequently we make the student respond, the more quickly and accurately we can adjust to their needs. This helps to keep the game at the goldilocks point and maintain a student's interest. Second, rapid-response games demand attention. It is simply impossible to play if your mind wanders. So rapid response video games are ideal for solving some of the most difficult challenges that we face when designing learning tools.

Whenever someone plays a video game, they are in fact learning, and the learning is efficient and fun. The problem is that *what* they are learning is not very useful or practical. But the question is, can we design video games so that when students play they learn what we want them to learn? The answer is a resounding yes.

Adaptive Practice represents both a set of principles for designing high-efficiency learning games and the technology behind those games. This technology includes a software engine designed to rapidly and accurately adapt during play to maintain the ideal conditions for learning.

Adaptive Practice games, like all video games, present the player with challenges. In this case though, the challenges are designed specifically to make the player think about something we want them to learn. In addition, the player's response is designed to reveal his or her level of understanding. We can use this response to adjust the challenge so that it perfectly matches what this player needs in this moment. The result is learning that is fun, surprisingly fast, and virtually effortless.

How long does it take someone to learn using Adaptive Practice? The answer it turns out is - it depends. Which brings us to another important concept - mastery based learning.

Mastery Based Learning

Adaptive Practice games use mastery based learning as opposed to grade based learning. To better understand what this means, let's look at an example.

Let's say we're in a typical fifth grade math class and the current topic of study is fractions. There are thirty students in the class and the teacher has allocated three weeks to complete this topic. There will be lectures, in-class practice, homework assignments, and quizzes. Then, at the end of the three weeks, there will be a test. The test will be used to determine how well each student has learned the material. Some will get As, but others will get Bs, Cs, or worse. The students who got As are ready to move on to the next topic. The ones who didn't are not ready, but in most cases, they all move on anyway. The cumulative effect of this is devastating to student performance and morale.

Consider what happens to the student who got a C? He is now behind and needs to catch up. The problem is that he's already shown that he is having trouble keeping up with the class at the normal pace. Now, we've given him the extra burden of having to complete additional learning to catch up. What are the chances that he will be able to do this successfully? Very small. How likely is it that any new material being introduced hits his goldilocks point? Almost nil.

Image what it's like for students in first grade when this happens. They begin to think they are stupid and slow. What happens as the year progresses? They fall behind even further. Either they quickly become hopelessly left behind, or the teacher begins to slow down the rest of the class to accommodate the slower students. How about the student who is a really quick learner? They understand a new topic in the first few minutes of class and are immediately bored. This student continuously gets A's but will also lose interest and motivation. As these students progress through school, learning is no longer fun and engaging. It's more like drudgery.

Adaptive Practice games incorporate mastery based learning. Every student learns at his or her own pace. Students play each lesson (game) until it has been fully mastered. Some students do this quickly, and some take longer, but everyone masters every topic before they move on to the next topic. There is no need for grades; we just need to know what has been mastered and what hasn't. In other words, everyone always gets an A. For students who are not used to getting As, the pride that comes from mastering something new

can have a profound affect on their attitudes and self-confidence.

But we use the term grades in two different ways. We've been discussing the meaning that corresponds to measuring a student's understanding and assigning a letter grade A, B, and so one. We also use grades to determine where a student is in school, first grade, second grade, etc. It turns out that when each student can move at an entirely independent pace; this concept of grades also loses much of its meaning and value.

Learning About Learning

Although Adaptive Practice is based on our latest understanding of human learning, I firmly believe that what we know about learning has only just begun to scratch the surface of what there is to learn. Adaptive Practice includes a mechanism for continuously improving our understanding of human learning.

One of the hottest trends in industry these days is big data analytics. Big data analytics refers to the ability of companies to collect massive amounts of data about business operations and then, using modern statistical analysis and data mining techniques, to find a wealth of valuable intelligence from that data. This intelligence can then be used to better manage the business. It can be used to improve things like marketing, inventory control, cash flow, sales, etc...

One of the benefits of cloud-based online learning is the ability to collect data. Not just a little bit of data, a LOT of data. During play, every student response is recorded to the cloud. This data tells us each student's current proficiency level on each learning objective, as well as how much time they've actually been practicing. We can then, for example, generate reports for teachers or parents. An even more interesting use of the data though, is to help us learn more about learning. There are two primary ways to do this.

The first involves experimenting with different techniques to find the ones that work best. One great way to do this is through split-testing. Split testing involves creating two different versions of a lesson or portion of a lesson and then letting some students use one while the rest use the other. We can then look at the response data to determine if one technique clearly works better than the other. Whatever we discover will of course lead to follow-up testing. As this process continues, we learn more and more about learning. Now this might seem like a slow and cumbersome process, but it can be amazingly rapid. With a large user base, we will be able to collect a statistically valid data set very rapidly, perhaps in as little as a few minutes. Also, we could automate a lot of the split testing so that we are running many simultaneous tests while minimizing the need for human intervention and control.

The second way to use the data to learn more about learning is through data mining. Data mining involves using computers to sift through large amounts of data looking for patterns. These patterns can tell us a lot about learning. Sometimes, we use data mining to answer questions that we ask, but data mining can also find the answers to questions we didn't even know we should be asking. Many different industries are now taking advantage of the great

value that data mining provides. Education is just now beginning this process.

Conclusion

The more we understand about how we learn, the clearer it becomes that traditional methods of learning cannot meet our society's ever-expanding need for knowledge and skill. As we evolve from teacher-centric, classroom-based learning to more student-centric, online learning, the technologies must evolve to maximize the possibilities of this new paradigm. Adaptive Practice offers a new approach that promises to fulfill the vision of high-efficiency learning through a process we actually enjoy.

About the Author

Thomas Butt is the President of Brainrush and the creator of Adaptive Practice technology. Tom has a diverse technology background which includes engineering and management positions in nuclear power, spacecraft design, and computer technology. Tom is passionate about education and accelerated learning.

www.brainrush.com