The Science Behind the ScootPad Mastery Learning Platform

WHITE PAPER

Table of Contents

1. Introduction

Recent advancements in technology have made it possible for individualized learning to be more affordable and accessible to all students like never before. This white paper outlines how the ScootPad adaptive learning platform incorporates proven research-based strategies, advanced learning science, modern design principles, and cutting-edge technologies to deliver a deeply personalized learning experience for every student—anywhere, anytime, on any device, and at any scale.

a. Background

Interest in what is now called "adaptive learning" can be traced back to educational research in the 1960s. This interest was further stimulated when the large positive impact that one-on-one tutoring can have on student performance was experimentally quantified. For example, one influential study by Benjamin S. Bloom (1984) reported that individual tutors can generate an improvement of two standard deviations. This dramatic effect provided the motivation for a large body of work in adaptive tutoring research (Evens and Michael 2006; reviewed by VanLehn 2011): If a human tutor can improve learning outcomes so radically, then many of the benefits (though likely not all) might be captured by an automated system.

Influenced and inspired by decades of research, the ScootPad adaptive learning platform builds on key findings and lessons learned from several interweaving fields, including adaptive tutoring systems, psychometrics, and cognitive learning theory, but it also introduces several key innovations that make it possible to scale and extend these benefits to far more students than was possible before, and in a manner that covers multiple domains.

b. Stanford University Research Advisors

We are grateful to our research advisors for their thought leadership and continued contributions towards conceiving and realizing the vision and roadmap for the ScootPad platform.

Karin Forsell, Ph.D., Department of Education, Stanford University

Karin Forsell is the director of the Learning Design and Technology (LDT) master's program and senior lecturer at the Stanford Graduate School of Education. Her research focuses on the choices people make in learning about and using new digital tools, with a special interest in teachers as learners and the design process. Her recent publications include "Making Meaningful Advances: TPACK for Designers of Learning Tools." Find out more: Karin Forsell

Robert Tibshirani, Ph.D., Department of Statistics, Stanford University

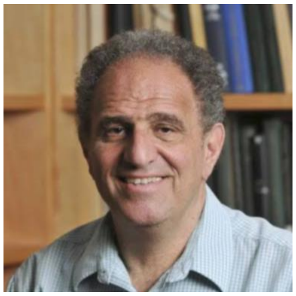

Rob Tibshirani is a professor of Statistics and Biomedical Data Science at Stanford University where he leads research in data science and statistics and how they can bring scalable social change. His research focuses on the challenges of sorting through large amounts of data and separating out the consistent patterns from the noise. His recent publications include: "The Elements of Statistical Learning: Data Mining, Inference and Prediction." Find out more: Rob Tibshirani

Emma Brunskill, Ph.D., Department of Computer Science, Stanford University

Emma Brunskill is a computer science professor at Stanford University where she is part of the Stanford Artificial Intelligence Lab and the Statistical Machine Learning Group. Her research focuses on reinforcement learning in high-stakes scenarios: how can an agent learn from experience to make good decisions when experience is costly or risky, such as in educational software. Publications include: "The Misidentified Identifiability Problem in Bayesian Knowledge Tracing." Find out more: Emma Brunskill

c. Overview of the ScootPad Approach

Powered by a continuous cycle of mastery, ScootPad enables personalized learning based on real-time adaptive orchestration of diagnostic-driven practice, just-in-time instruction, and automatic assessment.

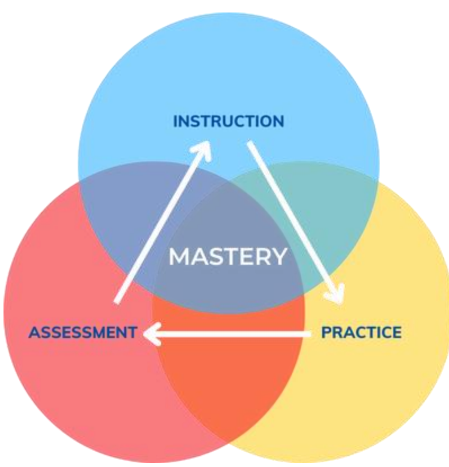

Adaptive Diagnostic & Practice
Students start with an adaptive diagnostic and continue with ongoing adaptive practice in concepts.

Just-In-Time Instruction
Bite-sized videos offer instruction in each concept - available on-demand and pushed when required.

Automatic Mastery Assessment
Strategically spaced automatic mastery checks assess student knowledge, including retention and recollection.

2. Underlying Research and Effectiveness Studies

The breadth and depth of research into intelligent tutoring systems (ITS) is evidence not just of their utility and promise, but of the difficulty of their construction. Many well-known tutoring systems, e.g., ALEKS (Doignon and Falmagne 1999), AutoTutor (Graesser et al. 1999), Andes (Gertner and VanLehn 2000), and Guru (Olney et al. 2012), represent years of laboratory effort and specialize in one or a few specific domains of study, such as Mathematics, physics, and biology.

This research has yielded learning systems with measurable positive effects on student outcomes, as well as useful insights into learning itself, such as the observation that tutoring systems that present content in a scaffolded way, where content is structured to build up more complex ideas in stages, have an advantage over those that do not (VanLehn 2011). While these approaches were necessary and important steps in the development of modern intelligent tutoring systems, none were developed with the intention of bringing the benefits of adaptive learning to as many students and across as many content domains as possible.

ScootPad's novel approach to creating a scalable intelligent tutoring system was built on three interweaving threads of research (described below) and their measurable positive effects on student outcomes. This resulted in a platform capable of delivering adaptivity and individually targeted learning experiences at scale across any content domain.

a. Mastery Learning

By definition, mastery learning is a method of instruction where the focus is on the role of feedback in learning. Furthermore, mastery learning refers to a category of instructional methods which establishes a level of performance that all students must "master" before moving on to the next unit.

Benjamin S. Bloom was an American educational psychologist who made significant contributions to the classification of educational objectives and to the theory of mastery learning. In 1956, Bloom edited the first volume of *Taxonomy of Educational Objectives: The Classification of Educational Goals*, which outlined a classification of learning objectives that has come to be known as Bloom's taxonomy and remains a foundational and essential framework within the educational community.

Dr. Bloom's research highlighted an educational phenomenon which was initially reported in 1984 in the journal *Educational Researcher*. Bloom found that the average student who was tutored one-to-one using mastery learning techniques performed two standard deviations better than students who learned via conventional instructional methods—that is, "the average tutored student was above 98% of the students in the control class." Additionally, the variation of the students' achievement changed: "about 90% of the tutored students . . . attained the level of summative achievement reached by only the highest 20%" of the control class. Hence this phenomenon, Bloom's 2 Sigma Problem, is named after him.

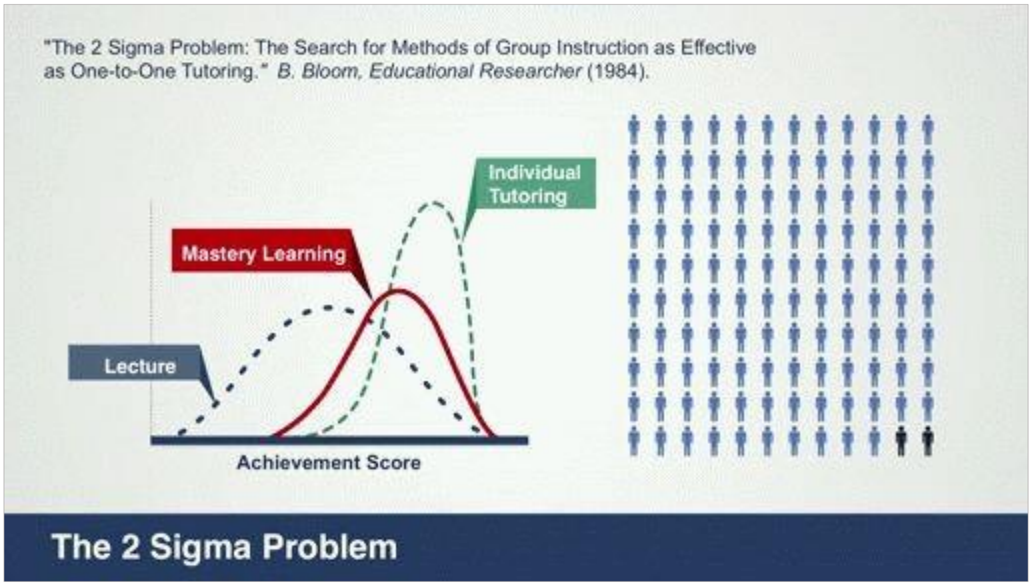

Though Bloom concluded that one-to-one tutoring is "too costly for most societies to bear on a large scale," Bloom conjectured that a combination of two or three altered variables may result in a similar performance improvement. Bloom thus challenged researchers and teachers to "find methods of group instruction as effective as one-to-one tutoring."

Bloom classified alterable variables that may have, in combination, a 2 sigma effect as the following "objects of change process": learner, instructional material, home environment or peer group, and teacher. Bloom considered and tested various combinations of these variables, focusing only on those variables that individually had a 0.5 or higher effect size. These included:

Effect of selected alterable variables on student achievement. Adapted from Walberg (1984).

Object of change process	Alterable variable	Effect size	Percentile equivalent
Teacher	Tutorial instruction	2.00	98
Teacher	Reinforcement	1.2	
Learner	Feedback-corrective (Mastery Learning)	1.00	84
Teacher	Cues and explanations	1.00	
Teacher, Learner	Student classroom participation	1.00	

Learner	Student time on task	1.00	
Learner	Improved reading/study skills	1.00	
Home environment / peer group	Cooperative learning	0.80	79
Teacher	Homework (graded)	0.80	
Teacher	Classroom morale	0.60	73
Learner	Initial cognitive prerequisites	0.60	
Home environment / peer group	Home environment intervention	0.50	69

ScootPad incorporates the most impactful mastery learning techniques from Dr. Bloom's recommendations into the platform, including tutorial instruction (1:1), reinforcement, feedback-corrective (mastery learning), cues and explanations, initial cognitive prerequisites, and home environment intervention. Hence, the ScootPad platform is a practical realization of Dr. Bloom's mastery learning techniques, one that is as effective as one-to-one tutoring and capable of delivering a 2 sigma effect on student achievement.

b. Personalized System of Instruction

Personalized System of Instruction (PSI), also known as the Keller Plan, was developed in the mid-1960s by Fred Keller and colleagues. It was developed using mastery learning theory and is based on the idea of reinforcement in teaching processes. While traditional teaching is "same pace, different learning," a key distinguishing factor of PSI is that it instead advocates "different pace, same learning."

A traditional course might have all students follow the same weekly lectures, exercises, etc., followed by an end-of-course exam that all take at the same time – but possibly with huge variations in learning outcomes (e.g., 95% achievement for a strong student, but just 55% for a weak one). In a course run according to PSI, all students must pass a high threshold of achievement on each module within the course (for instance 90%). The difference between weak and strong students would then be that the stronger ones are able to finish the course faster, while the weaker ones would need more time.

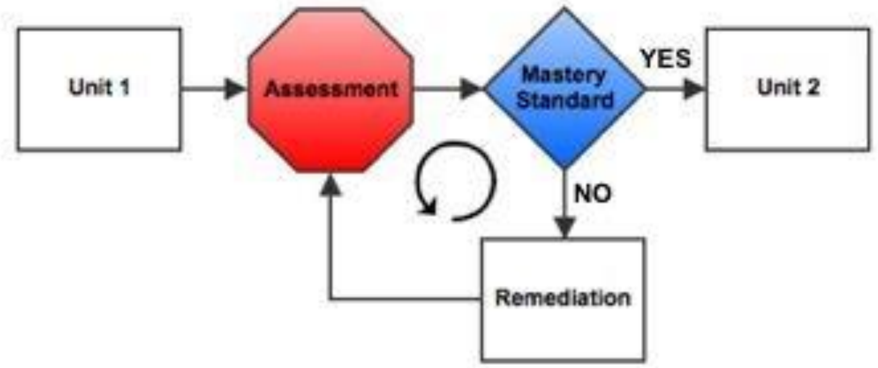

Keller argued that effective instruction should incorporate five principles, which are the essential elements of the Keller Plan:

1. Written materials: The primary presentation of new content should be through written texts. Given the forms of media available at the time when the Keller Plan was developed (e.g., lectures, movies, audio records, television, radio, paper-based text, etc.), paper-based texts gave students the greatest freedom: books and texts are portable, can be read at one's own pace, can be started and stopped at any time, can be easily reviewed, and can be marked by the reader. As an application of behaviorism, the Keller Plan was meant to maximize the number of operant behaviors that could be reinforced; this could best be done with written materials rather than have the learner be a passive observer of other media. Digital media available today could provide the same kinds of learner control and presumably could be incorporated into a contemporary implementation of PSI.
2. Units of content: Subject matter material should be broken down into separable, meaningful units. These units could have various kinds of relationships; for example, one could provide prerequisites for understanding a second, or the second could provide deeper elaboration of a preceding unit. In any case, specific learning objectives should be definable for each discrete unit of content.
3. Self-paced instruction: Students should be allowed to advance through the course material at their own pace. While an instructor might specify the order in which learning units are completed, the learners should decide when and at what rate they learn. This means that learners can move through a course as quickly or as slowly as they choose.
4. Unit mastery: Students must satisfy a mastery requirement in one unit before proceeding to the next. Students must demonstrate mastery of a unit's objectives to a certain level of quality. If the student does not reach the threshold, they are redirected to unit materials (or supplements if provided) and then take an equivalent form of the unit assessment. From the point of view of behaviorism, demonstrating mastery and being allowed to continue to a subsequent unit was presumed to be reinforcing.

5. Proctors: Human proctors are an important part of the Keller Plan. The proctors could be "external" to the course (adults or peers brought to the course from external sources) or "internal" (advanced students in the course who are doing well, have completed all units to date, and have good interpersonal skills). Proctors are the arbiters of unit mastery; they "certify" mastery, discuss areas of weakness, and direct students to the next units. Behaviorists were always concerned about bringing conditioned behaviors under the control of "natural" reinforcers; interactions with the proctors were presumed to provide natural social reinforcers that encouraged learning behaviors and perseverance in the course.

ScootPad's learning science and pedagogy (see section 3 for more details) incorporates these five principles of the Keller Plan and ensures reinforcement throughout the learning process.

c. Item Response Theory

Imagine that you're teaching a math remediation class full of fourth graders. You've just administered a test with 10 questions. Of those 10 questions, two questions are very simple, two are incredibly hard, and the rest are of medium difficulty. Now imagine that two of your students take this test. Both answer nine of the 10 questions correctly. The first student answers an easy question incorrectly, while the second answers a hard question incorrectly. Which student has demonstrated greater mastery of the material?

Under a traditional grading approach, you would assign both students a score of 90 out of 100, grant both an A, and move on to the next test. But this approach illustrates a key problem with measuring student ability via testing instruments: test questions do not have uniform characteristics. So how can we measure student ability while accounting for differences in questions?

Item response theory (IRT) attempts to model student ability using question-level performance instead of aggregate test level performance. Instead of assuming all questions contribute equally to our understanding of a student's abilities, IRT provides a more nuanced view of the information each question provides about a student. It is founded on the premise that the probability of a correct response to a test question is a Mathematical function of parameters such as a person's latent traits or abilities, as well as item characteristics (such as difficulty, "guessability," and specificity to topic).

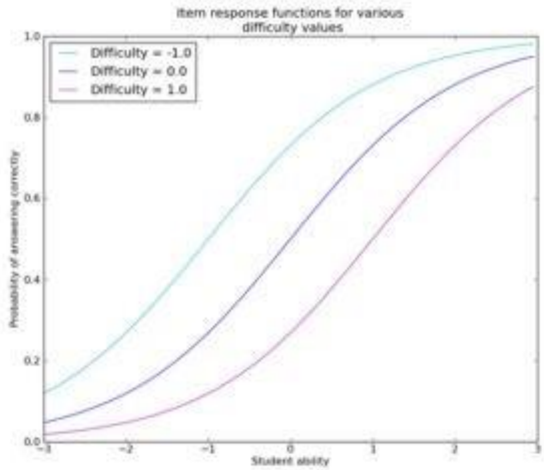

While IRT is very popular and has a proven track record in psychometric testing, many common IRT models make assumptions that can be limiting in the context of adaptive learning. For example, a common assumption is that a student's proficiency is constant, or at least that it is being measured at a single instant. This does not match the expectations of an adaptive learning system, where assessments are continuous over a long period of time in which the student is expected to learn.

Another common limiting assumption in IRT models is that the overall ability of a student can be condensed into a single parameter. While this assumption might be reasonable when estimating student abilities with a single assessment, it is much less likely to be true in an adaptive learning environment where students interact with content across multiple fine-grained concepts.

Modeling these finely grained skills requires extending IRT to support multiple proficiencies and providing a way for these proficiencies to interact. This interaction may be compensatory (where a high ability in any relevant concept can be used) or non-compensatory (where sufficiently high proficiency in every relevant concept is necessary).

The knowledge map (see section 3.a for more details) allows ScootPad to work with concepts which are very fine-grained. The edges in the map are used to describe relationships between individual concepts such that it becomes possible to draw inferences about, for instance, how a student's proficiency in a prerequisite of a concept changes given changes in her proficiency with respect to the original concept.

ScootPad makes learning dynamic by allowing proficiencies to change over time. Similarly, because skills and knowledge increase over time, giving more weight to more recent assessments can facilitate the learning process and encourage teaching practices that are focused on learning growth.

These modifications to the IRT framework, coupled with the knowledge map and real-time student proficiency computations, enable ScootPad to measure student knowledge more accurately to deliver deeply personalized learning experiences.

3. Learning Science-Fueled Innovative Pedagogies

By harnessing technology and applying what we know about learning science, we have implemented innovative pedagogies to deliver powerful personalized learning experiences. Guided by distinguished advisors from Stanford University, we bring the best in educational research, instructional methods, and data analytics to continually develop, refine, and deliver the most effective and impactful solutions. Fueled by learning science research, ScootPad is built on six pillars of pedagogical innovations, outlined below:

a. The Knowledge Map

The notion of representing content relationships in a structured format is not new, and the ScootPad knowledge map builds on the graph theory and the Mathematical structures used to model pairwise relations between objects. However, where the ScootPad knowledge map innovates is in its flexible and expressive ontology, which allows diverse content to be easily represented and connected with each other. This makes the ScootPad knowledge map a core instrument in enabling adaptivity at scale across content domains.

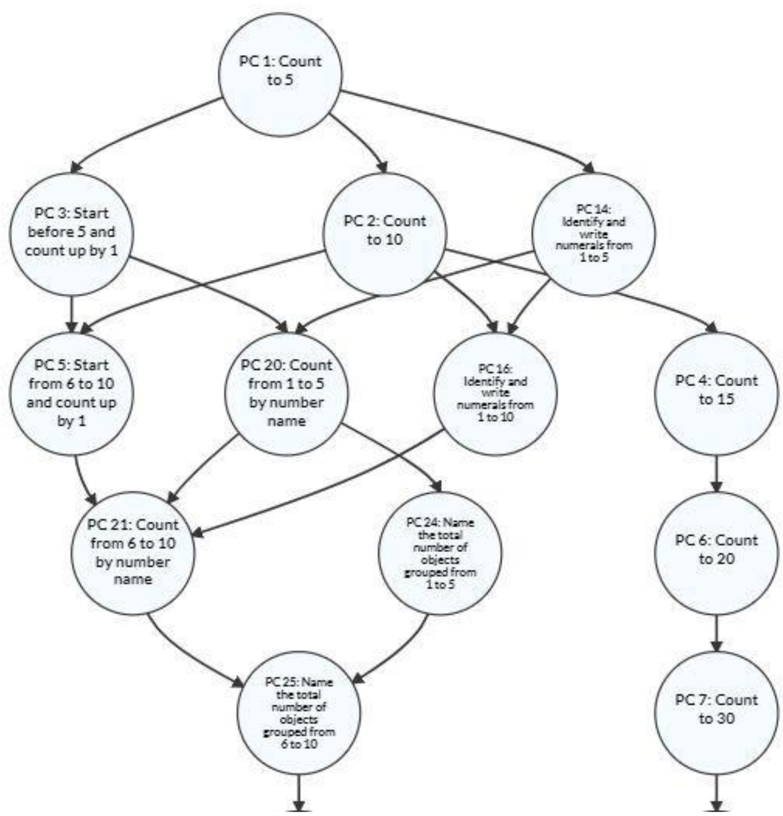

The ontology of ScootPad's knowledge map is made up of concepts (i.e. nodes representing fine grained skills) and the relationships between them (edges representing the dependencies).

ScootPad's knowledge map is a true map of skills: new learning is built on previous, foundational understandings, thus creating an interconnected web of prerequisite skills. It's important to recognize that a learning progression as comprehensive and interrelated as ScootPad's knowledge map takes years to develop and could only come to fruition through a continuous process of research, expert review, and iterative revision.

The knowledge map allows our adaptive algorithms to generate personalized pathways based on pedagogical criteria, prerequisite scaffolding, and remediation assistance. The map is also a key input to diagnose student understanding (and misunderstanding), recommend intelligent intervention strategies, and bridge between prior knowledge and new knowledge.

b. Learning Paths and Mastery-Based Progression

ScootPad's Learning Path is an intuitive and flexible tool to define the scope of any desired curriculum (example: 2nd grade math, 3rd grade English etc.). It provides a logical structure to organize the scope of the content that students must master by breaking it down into separable, meaningful units. These units enable the sequence in which students are introduced to the content. Each learning path can have an unlimited number of units and each unit can have one or more concepts across any content domains.

ScootPad's content bank offers curriculum and content to support Common Core, state-specific learning standards, and selected international standards. Pre-made learning paths are available for each grade level and are aligned to each curriculum (example: the Common Core Grade 5 math Learning Path covers 100% of the Common Core standards a 5th grader must master). The ScootPad platform offers districts, schools, and teachers the opportunity to design their own custom learning paths using an intuitive drag-and-drop user interface. During the creation of a custom learning path, ScootPad applies the knowledge map to group and sequence the selected concepts into recommended units and concept groupings; users have the option to overwrite such recommendations and define their own sequence and grouping. Custom learning paths can then be shared with peers across the school, district, and the larger ScootPad.

With ScootPad's mastery-based model, students are assigned a learning path that is appropriate to their learning level and are supported to work at their own pace, so they can take the time they need to fully understand and master the material. Simply put, mastery-based learning progressions are descriptions of how learning typically advances in a subject area. "Empirically based learning progressions can visually and verbally articulate a hypothesis, or an anticipated path, of how student learning will typically move toward increased understanding over time with good instruction" (Hess, Kurizaki, and Holt, 2009). Finally, Heritage (2011, p. 3) suggests that learning progressions provide descriptions of "how students' learning of important concepts and skills in a domain develops from its most rudimentary state through increasingly sophisticated states over a period of schooling."

Inherent in these views of progressions is the idea of a coherent and continuous pathway along which students move incrementally through states of increasing competence in a domain. Every incremental state builds on and integrates the previous one as students accrue new levels of expertise with each successive step in the progression. As Herman (2006, p. 122) observes, "whether and how children are able to engage in particular learning performances and the sequence in which they are able to do so are very much dependent on previous opportunities to learn." The benefit of progressions is that they lay out a continuum to guide teaching and learning over time so that student competence in the domain can be advanced coherently and continuously.

Traditionally, educational products claiming personalized learning capabilities identify a student's level of proficiency in a subject matter only at a single point in time and only by using a fixed number of questions (with a diagnostic test, knowledge check, quiz, assessment etc.). Then, it will seek to develop a curriculum to fill the gaps of knowledge. In such cases, the rigor and the curriculum remain static. ScootPad, however, provides *continuous personalization*. ScootPad analyzes each student response in real-time taking into account their latest performance and will adjust the rigor (number of questions to practice and their respective depth of knowledge) and adjust the curriculum (the next concept to learn). In other words, ScootPad's adaptive learning process is dynamic and thereby continuously adapting to each student's needs and optimizing the path in real-time to achieve mastery in the desired concepts.

With ScootPad's mastery-based progression model, students work at their own pace and advance to the next unit only when they have learned all concepts in the preceding unit. Once a student masters all concepts from the learning path across all units, they're deemed to have mastered the assigned curriculum and are ready for new and advanced curriculum.

c. Spaced Learning and Distributed Practice

In a spiral curriculum, material is revisited repeatedly over time and across many sessions. Different terms are used to describe such an approach, including “distributed” and “spaced.” Humans and animals learn items in a list more effectively when they are studied in several sessions spread out over a long period of time, rather than studied repeatedly in a short period of time, a phenomenon called the spacing effect. The opposite is called massed practice and consists of fewer, longer practice sessions where learning is concentrated in continuous blocks. It is generally a less effective method of learning. For example, when studying for an exam, studying more frequently over a larger period of time will result in more effective learning than intense study the night before.

Findings about distributed learning are among the most robust in the learning sciences, applying across a wide range of content and age groups. “Space learning over time” is the first research-based recommendation in a recent practice guide from the U. S. Department of Education’s Institute of Educational Sciences (Pashler et al., 2007). In a recent review of the literature, Lisa Son and Dominic Simon write, “On the whole, both in the laboratory and the classroom, both in adults and in children, and in the cognitive and motor learning domains, spacing leads to better performance than massing” (2012).

Massing encourages quick and effortless processing of information leading to reduced attention and short-term learning (Dempster, 1988; Rohrer, 2009). Easy learning often doesn’t lead to the best retention; more difficult learning can lead to more robust encoding of information and better long-term learning (Schmidt & Bjork, 1992). This explanation identifies the spacing effect as an example of a “desirable difficulty” that enhances learning.

Furthermore, spiraling helps learners make connections over time, which creates more robust pathways for recalling information. Multiple, strategically spaced and strategically progressing learning experiences may produce deeper, more conceptual learning.

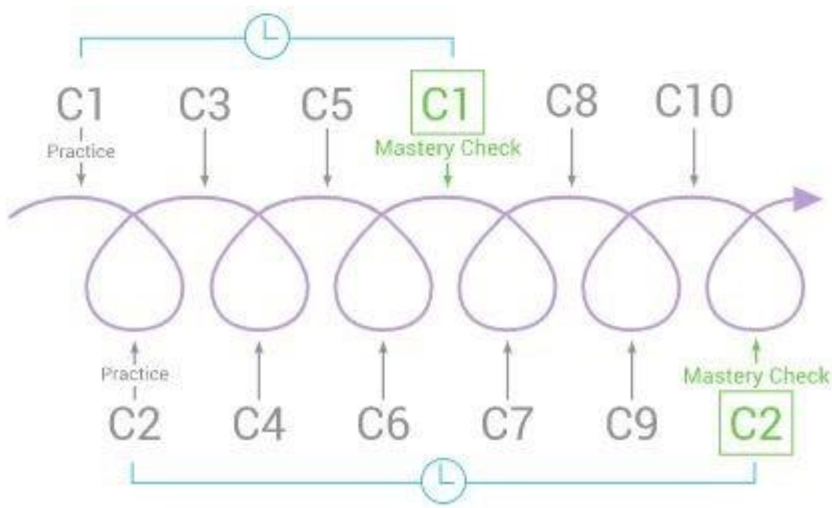

Most curricula are not designed to take advantage of the spacing effect, much to the frustration of the psychologists who have documented its power (Dempster, 1988; Rohrer, 2009). One reason is that the spacing effect is counterintuitive: People feel that massing leads to higher performance, which is true in the short term but is not if the goal is long-term learning. People confuse short-term performance with long-term learning and inaccurately predict that massed practice will lead to better long-term results than spaced practice. UCLA psychologist Robert Bjork uses the term “illusion of competence” to describe this feeling (1999). Another reason spiraling is not common in curriculum design is that many teachers are unaware of the benefits of spacing learning over time. Teachers may also be discouraged when confronted with how much their students forget, something that is more apparent with spacing (in which topics are revisited after students have had time to forget) than with massing (in which topics are not revisited so that forgetting is not as obvious). A third reason is that students find spaced learning harder than massed learning, so they tend to prefer a massed approach even though it’s less efficient. A final reason that spiral curricula are not common is that building such curricula is more difficult and complicated.

With ScootPad’s spiral curriculum and distributed learning, concepts are practiced simultaneously over time ("spiraling") and revisited across multiple sessions ("spacing"). Unlike the "sequential" or "streak" method where each concept is practiced and mastered one at a time, the ScootPad approach helps students make connections over time creating more robust pathways for recalling information leading to increased attention and long-term learning (Dempster, 1988; Rohrer, 2009).

d. Scaffolding to Fill Knowledge Gaps

Closely related to achievement gap and opportunity gap, a knowledge gap is the difference between what a student has learned, i.e., the academic progress he or she has made, and what the student was expected to learn at a certain point in his or her education, such as a particular age or grade level. A knowledge gap can be relatively minor, e.g., the failure to acquire a specific skill or meet a particular learning standard, or it can be significant and educationally consequential, as in the case of students who have missed large amounts of schooling.

One of the more consequential features of knowledge gaps is their tendency, if left unaddressed, to compound over time and become more severe and pronounced, which can increase the chances that a student will struggle academically or socially, possibly even dropping out of school. In addition, if foundational academic skills such as math and reading are not acquired by students early on in their education, it creates a domino effect and causes the knowledge gap to compound and widen across content domains in the subsequent years of student learning. As students progress through their education, remediating knowledge gaps tends to become more difficult because students may have fallen well behind their peers, or because middle school or high school teachers may not have specialized training or expertise in teaching foundational academic skills. For these and other reasons, many educators and researchers have called for greater investments in early childhood education and remediation.

ScootPad's continuously adaptive learning system constantly mines student performance data and responds in real time to a student's activity using the platform. When a student initially struggles with a particular concept, ScootPad immediately directs the student to a time-out lesson in order to instantly and independently fill that knowledge gap. This lesson is immediately followed by a quick practice to assess the newly gained knowledge. If the student continues to struggle, ScootPad will know where that particular student's weaknesses lie in relation to all prerequisite concepts from the knowledge map and can automatically scaffold the student to the necessary prerequisite concept(s) for instruction and practice. This process of automatic scaffolding continues until the student demonstrates proficiency in all prerequisite concepts and eventually in the original concept itself. This process leads to dynamically filling all knowledge gaps.

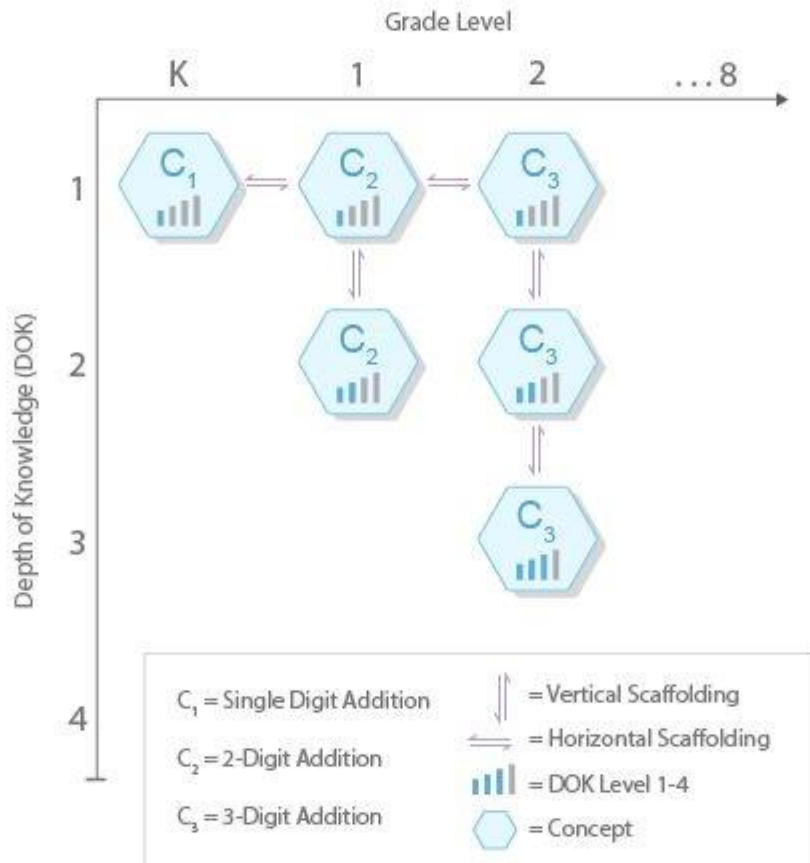

Scaffolding of content in the ScootPad platform spans across depth of knowledge (DOK), also known as “vertical scaffolding,” and across grade levels, also known as “horizontal scaffolding.” In this way, ScootPad’s continuously adaptive learning system provides each student with a truly personalized learning experience every step along the way, effectively remediating all knowledge gaps in the desired curriculum.

e. Reinforcement-Based Memory Retention

Inspired by Hermann Ebbinghaus’s work on memory retention and learning curves, ScootPad uses the learning and forgetting curves to model changes in student ability. The forgetting curve hypothesizes the decline of memory retention, demonstrating how information is lost over time when there is no attempt to retain it. A related concept is the strength of memory, which refers to the durability of certain memory traces in the brain. In essence, the stronger the memory, the longer a person is able to recall it. A typical graph of the forgetting curve purports to show that humans tend to halve their memory of newly learned knowledge in a matter of days unless they consciously review the learned material.

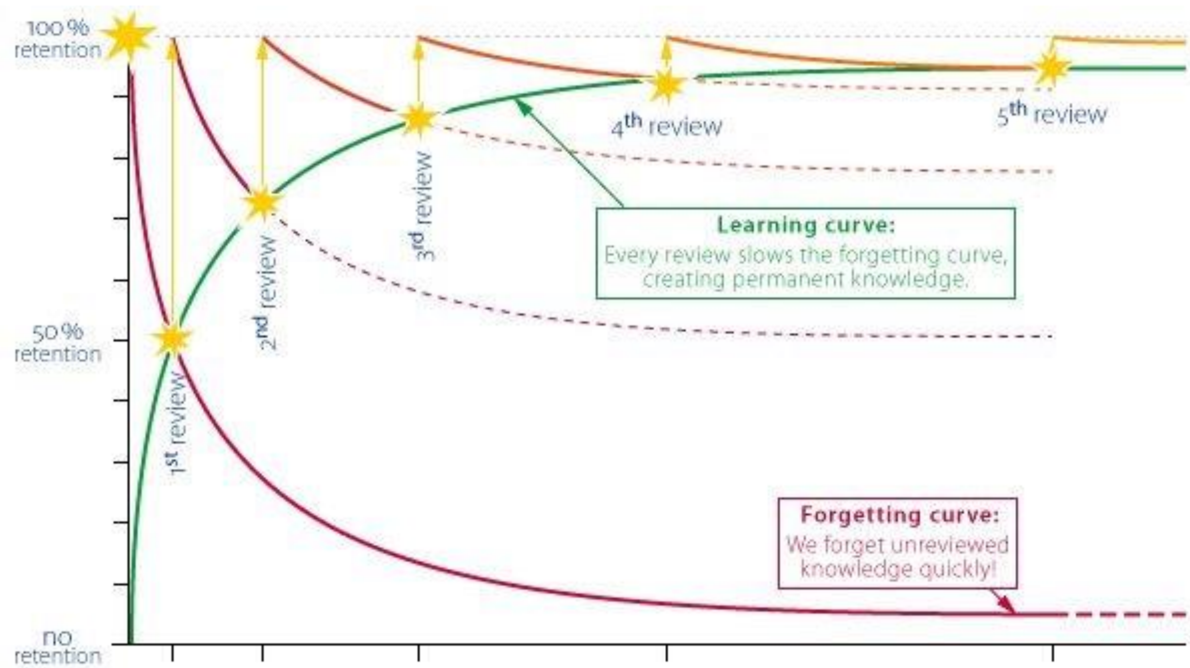

The forgetting curve itself that governs rate of retention is roughly described by the following formula:

$$R = e^{\frac{-t}{s}}$$

where R is retrievability (a measure of how easy it is to retrieve a piece of information from memory), s is stability of memory (determines how fast R falls over time in the absence of training, testing or other recall), and t is time.

By integrating this curve into the reinforcement algorithms, ScootPad can strategically space mastery checks depending on how and when they practiced each new concept—helping students make connections over time and creating more robust pathways for retaining and recalling information over a longer term.

f. Student Learning Profile

A key personalized learning strategy is using data—specifically, data from multiple sources, such as diagnostic tests, practices, lessons, assessments, and interventions, as well as non-achievement data and learning goals—to understand student progress, inform instructional decisions, and implement intervention strategies. With ScootPad, students can visualize a continuously updated learning profile that contains information on what the student knows and how they achieved progress towards mastery. The profile is progressive, which means it keeps getting smarter the longer the student interacts with the platform.

For instance, if a student has already mastered a grade level curriculum on ScootPad and moves to the next grade level curriculum, they start off "warm" with their learning path (as opposed to starting "cold" with no data). The new learning path takes into account the student's previously mastered concepts and their unique trajectory through the path and uses this knowledge to maximize student learning continuously from that point forward. Once enough data is collected, the platform will uncover patterns in the student's learning, including granular strengths and weaknesses. The more data in a student's knowledge profile, the more effective the platform becomes at serving up targeted learning material for that student.

ScootPad emphasizes visualizing student data and progress, with a focus on making it clear, accurate, and user-friendly. This kind of visualization has been proven to directly impact student achievement. In *Classroom Assessment and Grading that Work* (2006), Robert Marzano reminds us that graphic representation of student progress and achievement is so motivating and helpful that it results in an increase of 26 percentile points in student achievement when used regularly.

Nothing is more dissatisfying to students than feeling like the challenges they face are essentially arbitrary and culminate in nothing. The student profile answers the student need for continuity and meaning by affording students a sense of long-term investment in the learning process.

"Presumably, seeing a graphic representation of students' scores provides teachers with a more precise and specific frame of reference for making decisions about next instructional steps." Fuchs and Fuchs (1986). To this end, ScootPad mines students' data (within a classroom), analyzes patterns of predictable interventions, organizes actionable insights into the three response-to-intervention (RTI) tiers and helps educators implement intervention strategies from student data quickly and effectively.

4. Implementation at Scale

Storing and processing massive amounts of data produced by students requires a modern, scalable architecture that is sophisticated enough to run complex statistical operations, reliable enough to keep all services available at all times, and fast enough to perform necessary calculations in milliseconds. ScootPad has established state-of-the-art infrastructure which offers high-availability and high-performance powered by an infinitely scalable serverless architecture.

a. Platform Architecture

The ScootPad platform uses a proprietary cloud-based service-oriented architecture where each service is responsible for solving a modular and well-defined problem. The services themselves communicate both with each other and with the ScootPad user application, offering unparalleled flexibility and customization.

b. Tech-Enhanced Content Bank

ScootPad's extensive content bank enables non-technical users to build and maintain standards-aligned and DOK-leveled content across grade levels and subject areas using over 50 technology enhanced test question (item) formats. ScootPad's own content bank includes content covering grade levels K through 8 in math and English; this content includes instructional videos with multiple videos per concept; practice problems (over 40 for each concept); and assessment problems (over 10 for each concept).

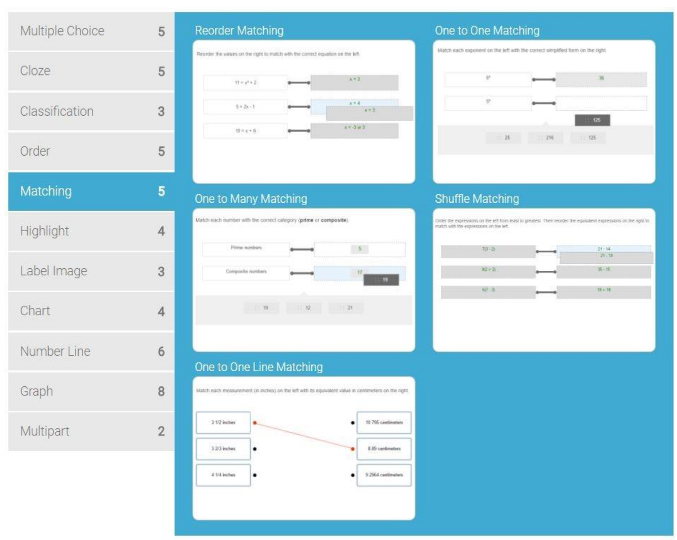

With ScootPad, educators can build their own content with tech-enhanced items (TEIs) across any curricula, any standards, any grade levels, and any subject areas. The platform supports an extensive and ever-expanding list of 50+ TEI types—all the technology enhanced item types needed to build interactive student experiences.

c. Content Agnostic Adaptive Engine

ScootPad's adaptive learning engine is "agnostic" of the content and can support any curricula, as educators can use preloaded curricula or build their own content. Educators can also utilize a mixture of both. ScootPad is designed to support multiple content banks and provides the ability for content teams to seamlessly collaborate. Through collaboration between our content designers and the respective trained subject matter experts, ScootPad knowledge maps for any new content domains can be developed in just a few weeks using our proven approach. Automated methods also assist in generating a knowledge map by combining historical content data and 3rd-party content maps. When faced with a large content collection, ScootPad works with partners to define and staff a scalable content mapping and validation process.

ScootPad currently provides content across two subject areas: math and English. This content is aligned to offer 100% coverage of over 15 international curricula including Common Core State Standards, 16 U.S. state standards, and England's National Curriculum.

5. Common Misconceptions

There are many different degrees and types of adaptive learning (single point vs. continuous, adaptive testing vs. adaptive learning), but often these distinctions aren't made clear. As the quest for personalized learning gains traction among educators, and more educational products claim "adaptive learning" capabilities, a certain fuzziness has emerged around the term. The ScootPad team is entirely focused on adaptivity, as we iterate on and improve our adaptive learning platform. In this section, we'll attempt to clearly articulate our interpretation of the value of adaptive learning and help to dispel some of the common misconceptions and misunderstandings in this area.

a. Personalized Learning vs. Adaptive Learning

This may seem like a matter of semantics, but "personalized" is not the same as "adaptive." For example, you may buy your child a new pair of shoes that fits their size, style, and interests perfectly. However, a year later you will likely find that the child has changed in a number of ways and now needs new shoes. The shoes were certainly personalized when you purchased them, but they can't adapt as the child grows.

This is exactly how learning works! Personalized refers to "targeted" content offering based on a student's specific needs at a particular point in time. Every student who meets this criterion will receive the same experience. The content, however, cannot adapt to meet the less well-defined nuances of their individual needs. For example, it cannot determine how quickly the student is picking up new information or whether they have retained it over the long term. And that means it can't adjust content in real time to focus on areas where the student needs more help, nor can it advance the content to challenge the student.

When exploring educational products that promise to provide "personalized learning" or "adaptive learning," dig in to understand how the learning experience evolves with the student over the long term. If the product is only offering opportunities to personalize content (either manually or systematically) for students based on certain predefined attributes, it is providing personalized—not adaptive—learning.

b. Single Point Adaptivity vs. Continuous Adaptivity

When most people use the "adaptive" buzzword, what they're really discussing is either a) single point adaptivity, which evaluates a student's performance at one point in time in order to determine the level of instruction or material he receives from that point on, or b) adaptive testing, which determines a student's exact proficiency level using a fixed number of questions.

When we refer to adaptive learning in the context of ScootPad, it means a system that is continuously adaptive—one that responds in real time to each student's performance and activity in the platform and that maximizes the likelihood a student will obtain her learning objectives by providing the right instruction, at the right time, and about the right thing. In other words, while adaptive testing answers the question, "How do I get the most accurate picture of a student's state of knowledge with a fixed number of questions?," adaptive learning answers the question, "Given what we understand about a student's current knowledge, what should that student be working on right now?"

c. Machine Learning vs. Adaptive Learning

When most people hear the term "adaptive learning" they automatically assume that it inherently uses "machine learning." Most traditional adaptive learning systems are based on approaches and models that are derived from data, critical parameters of student modeling, adaptation algorithms, etc. However:

Machine learning is not necessary. Many successful adaptive learning systems are based on knowledge engineering where necessary data, models, and parameters are provided by domain experts rather than "learned" from data. Moreover, there are many relatively simple approaches in the area of adaptive learning that do not need anything close to machine learning and are still very efficient. For example, classic overlay student modeling could be implemented with simple quantitative or qualitative rule-based approaches and offer a high quality of adaptation.

Machine learning is not always possible. Machine learning needs a considerable volume of data to work and in many cases the right volume of data simply might not be available. And yet, what is most exciting about modern adaptive learning is that it is possible to use stronger and more reliable quantitative approaches since we now have access to large volumes of learners' data.

6. Ongoing Research and Development

The advent of "big data" in areas such as internet search engines and online shopping has disrupted existing industries, created new industries, and led to the extraordinary success of companies such as Google and Amazon. Big data unleashes a range of productive possibilities in the education domain in particular, since data that reflects cognition is structurally unique and offers a high degree of correlation (mastery of fractions and mastery of exponentiation, for example). Educational data offers a tremendous potential to optimize user experiences over time and provide tangible value for students.

In the context of ScootPad, we're most excited about leveraging the large volume of learner data that is now available to build a modern adaptive learning platform that utilizes stronger and more reliable quantitative approaches. We can build models, make predictions, and run reliable experiments using the massive amounts of data produced by students (anonymously). Our current research is focused on three key areas of the platform and their tangible impact on student outcomes, as outlined below:

a. Predictable Knowledge Heatmaps

Knowledge Heatmap is a data visualization model based on continuous analysis of thousands of real-world student experiences to predict which nodes (i.e., concepts or prerequisite skills) within a specific content domain (ex: 3rd grade math) are the most difficult for students to master. These heatmaps are designed to help teachers and administrators identify critical areas for instruction, supplemental resources, teacher support and professional development.

b. Data-Driven Content Improvement

A continuous analysis of learners' performance data powered by machine learning can be incredibly valuable in identifying opportunities for content improvement. Such analysis can assist in modifying the level of difficulty, alignment to concept, accuracy, quality, and relationship to pedagogy.

c. Enhanced Learning Progression

As Herman (2006, p. 122) observes, "whether and how children are able to engage in particular learning performances and the sequence in which they are able to do so are very much dependent on previous opportunities to learn." The benefit of progressions is that they lay out a continuum to guide teaching and learning over time so that student competence in the domain can be advanced coherently and continuously.

Using machine learning-powered models, we plan to empirically analyze thousands of real-world student progressions to enhance each concept's learning progression criterion and scaffolding sequence, thus enabling students to efficiently and effectively master the curriculum.

7. Conclusion

ScootPad's approach to modern adaptive learning and scalability across content domains, as well as its emphasis on leveraging large amounts of student data to improve outcomes, represent innovations in the field of intelligent adaptive learning. The science and technology built to support these innovations enable personalized recommendations and up-to-date predictions on student outcomes. As the data available to ScootPad increases, and as educational approaches continue to evolve, this approach allows ScootPad to continue to grow both in its ability to scale across content, and its ability to meet student needs in ever more innovative ways. This growth will offer more students the opportunity to experience the transformational power of adaptive learning that is both highly effective and engaging.

8. References and Works Cited

Bloom, B. S. "The 2 Sigma Problem: The Search for Methods of Group Instruction as Effective as One-to-One Tutoring." Educational Researcher 13.6 (1984): 4-16. Web.

Keller, F. S. (1967). Engineering personalized instruction in the classroom. Revista Interamericana de Psicologia/Interamerican Journal of Psychology, 1(3).

Reckase, M. D. (1997). The Past and Future of Multidimensional Item Response Theory. Applied Psychological Measurement, 21(1):25–36.

Anderson, J. R., Boyle, C. F., Corbett, A. T., and Lewis, M. W. (1990). Cognitive Modeling and Intelligent Tutoring. Artificial Intelligence, 42(1):7–49.

Smith, S. G. and Sherwood, B. A. (1976). Educational Uses of the PLATO Computer System. Science, 192(4237):344–352.

Evens, M. and Michael, J. (2006). One-on-One Tutoring by Humans and Machines. Computer Science Department, Illinois Institute of Technology.

VanLehn, K. (2011). The Relative Effectiveness of Human Tutoring, Intelligent Tutoring Systems, and Other Tutoring Systems. Educational Psychologist, 46(4):197–221.

Doignon, J. P. and Falmagne, J. C. (1999). Knowledge Spaces. Springer.

Graesser, A. C., Wiemer-Hastings, K., Wiemer-Hastings, P., and Kreuz, R. (1999). AutoTutor: A Simulation of a Human Tutor. Cognitive Systems Research, 1(1):35–51.

Gertner, A. S. and VanLehn, K. (2000). Andes: A Coached Problem Solving Environment for Physics. In Gauthier, G., Frasson, C., and VanLehn, K., editors, Proceedings of the 5th International Conference on Intelligent Tutoring Systems, volume 1839 of Lecture Notes in Computer Science, pages 133–142, London. Springer-Verlag.

Graesser, A. C., Conley, M. W., and Olney, A. (2012). Intelligent Tutoring Systems. APA Handbook of Educational Psychology. Washington, DC: American Psychological Association.

Fok, Apple W. P., and Horace H. S. Ip. "Personalized Education: An Exploratory Study of Learning Pedagogies in Relation to Personalization Technologies." Advances in Web-Based Learning – ICWL 2004 Lecture Notes in Computer Science (2004): 407-15. Web.

Mary Brabeck, Jill Jeffrey, and Sara Fry. Practice for Knowledge Acquisition (Not Drill and Kill). American Psychological Association. Web.

Denny, Paul. "The Effect of Virtual Achievements on Student Engagement. "Proceedings of the SIGCHI Conference on Human Factors in Computing Systems - CHI '13 (2013): n. pag. Web.

Boger-Mehall, S.R. (1996). Cognitive Flexibility Theory: Implications for Teaching and Teacher Education. In B. Robin, J. Price, J. Willis & D. Willis (Eds.), Proceedings of Society for Information Technology & Teacher Education International Conference 1996 (pp. 991-993). Chesapeake, VA: Association for the Advancement of Computing in Education (AACE).

John F. Pane, Elizabeth D. Steiner, Matthew D. Baird, Laura S. Hamilton. Continued Progress: Promising Evidence on Personalized Learning. DOI: 10.7249/RR1365 (2015). Web.

Herman, J. L. (2006). Challenges in integrating standards and assessment with student learning. Measurement, 14(1 & 2), 119–124.

Heritage, M. (2009). The case for learning progressions. San Francisco, CA: The Stupski Foundation.

Heritage, M. (2011, April). Developing learning progressions. Paper presented at the Annual Meeting of the American Educational Research Association, New Orleans, LA.

Hess, K., Kurizaki, V., & Holt, L. (2009). Reflections on tools and strategies used in the Hawaii Progress Maps Project: Lessons learned from learning progressions.

"Memory: A Contribution to Experimental Psychology -- Ebbinghaus (1885/1913)." Retrieved 2007-08-23.

Classroom Instruction That Works (2001): Research-Based Strategies for Increasing Student Achievement [Robert J. Marzano, Debra Pickering, Jane E. Pollock]

Tutorial Instruction: Anania, J. (1982). The effects of quality of instruction on the cognitive and affective learning of students (Doctoral dissertation, University of Chicago, 1981). Dissertation Abstracts International, 42, 4269A.

Reinforcement: Lysakowski, R.S., & Walberg, H.J. (1981). Classroom reinforcement: A quantitative synthesis. Journal of Educational Research, 75, 69-77.

Feedback-Corrective, Cues & Explanations, and Student Classroom Participation: Lysakowski, R.S., & Walberg, H.J. (1982). Instructional effects of cues, participation, and corrective feedback: A quantitative synthesis. American Educational Research Journal, 19, 559-578.

Initial Cognitive Prerequisites: Leyton, F.S. (1983). The extent to which group instruction supplemented by mastery of the initial cognitive prerequisites approximates the learning effectiveness of one-to-one tutorial instruction (Doctoral dissertation, University of Chicago, 1983). Dissertation Abstracts International, 44, 974A.

Home Environment Intervention: Iverson, B.K., & Walberg, H.J. (1982). Home environment and learning: A quantitative synthesis. Journal of Experimental Education, 50, 144-151.