Exploring the Benefits of Adaptive Learning Methods

Krista L. Ratwani, Alan S. Carlin[1], Evan Oster
Aptima, Inc.
Arlington, VA
kratwani@aptima.com, alan.carlin@stresearch.com, eoster@aptima.com
[1] Alan Carlin is now at STR.

ABSTRACT

One of the largest constraints that must be met in learning environments is ensuring that high numbers of learners can receive the necessary instruction in a fixed time period. Out of necessity, instructors must use the same material for all learners, with the goal of presenting the material in such a manner that it is not too challenging and not too easy for the average student. The risk, of course, is that students that do not quite grasp the material are left without the necessary remediation as the content continues to progress linearly, on a set timeline; conversely, more knowledgeable students may be bored and unable to achieve their full potential. Adaptive learning approaches were introduced to counter this "one-size-fits-all" approach and ensure that learners receive the content that is most appropriate for their current knowledge and skill level. However, the literature on aptitude by treatment interactions (ATIs; Cronbach & Snow, 1981) dictates that not all learning interventions may be equally effective for all groups of learners. Therefore, this research sought to explore the conditions under which adaptive learning approaches may be most effective.

A mathematical task focused on order of operations was developed and provided to 76 participants recruited through Amazon Mechanical Turk. Participants in a control condition received the same set of problems no matter their performance, whereas participants in an experimental condition received problems that aligned with their performance on previous problem sets. The results lend evidence to the idea that there may be boundary conditions when using adaptive learning approaches. Specifically, individuals that started the task at beginner and advanced skill levels were more likely to benefit from the adaptive learning approach. Individuals who possessed a more average skill level at the start performed better in the control condition. Practical implications and guidelines are discussed.

ABOUT THE AUTHORS

Dr. Krista Ratwani is a Principal Scientist and the Senior Director of the Learning and Training Systems Division at Aptima. Her work focuses on increasing the effectiveness of learning experiences for leaders through the development of assessment tools, as well as support tools to facilitate improved feedback. She also focuses on understanding how learning and development occurs throughout one's career. Dr. Ratwani holds a Ph.D. and M.A. in Industrial-Organizational Psychology from George Mason University and a B.A. in Psychology from Monmouth University.

Dr. Alan Carlin is a Principal Research Engineer and Team Lead at STR. His current research interests include cyber planning as well as analyzing group dynamics in social networks. Prior to STR, he worked at Aptima where his research interests included artificial intelligence in education and educational data mining. Dr. Carlin received a dual-B.A. in computer science and psychology from Cornell University, a Master's degree in Computer Science from Tufts University, and a Ph.D. in Computer Science from the University of Massachusetts. His dissertation studied AI decision-making, coordination, communication, and allocation of computing resources in decentralized, partially observable environments.

Mr. Evan Oster is a Scientist at Aptima with a background in courseware development, game-based training, and mixed reality technology. His current work centers around using cognitive science, instructional design, and analytics to increase training effectiveness in live, virtual, and constructive environments. Previous work includes designing game and simulation-based training using immersive virtual environments for the Navy. Mr. Oster has also created several augmented reality games with a university games lab while providing software development evaluation. Evan Oster holds an M.A. in Instructional Design and Technology and a B.S. in Human Development from Virginia Polytechnic Institute and State University as well as an M.S. in Curriculum and Instruction from Radford University.

Exploring the Benefits of Adaptive Learning Methods

Krista L. Ratwani, Alan S. Carlin[1], Evan Oster
Aptima, Inc.
Arlington, VA
kratwani@aptima.com, alan.carlin@stresearch.com, eoster@aptima.com
[1] Alan Carlin is now at STR.

INTRODUCTION

One of the largest constraints that must be met in learning environments is ensuring that high numbers of learners can receive the necessary instruction in a fixed time period. This emphasis on throughput can be experienced in large lecture-style college classrooms where there may be 100 students present in one class, as well as the military, where over 500 Soldiers go through Basic Combat Training at one time. Out of necessity, instructors must use the same material for all learners, with the goal of presenting the material in such a manner that it is not too challenging and not too easy for the average student. The risk is that students that do not fully grasp the material are left without the necessary remediation as the content continues to progress linearly, on a set timeline; conversely, more knowledgeable students may be bored and unable to achieve their full potential by moving to more advanced content levels.

To counter this traditional, linear path to learning, adaptive learning was introduced as a way to personalize instruction. Adaptive learning has become ubiquitous both in research and operational environments. For the former, research spans the gamut from generally examining and describing various intelligent tutoring systems (ITS; Vanlehn, 2006) to specifically analyzing the impact of targeted algorithms on personalizing learning (e.g., Mao, Lin, & Chi, 2018). For the latter, every branch of the military has called for the need to embrace personalized, adaptive learning methods tailored to each learner's unique profile of strengths and weaknesses (Department of the Army, 2011; Roberson & Stafford, 2017; U.S. Fleet Forces Command, 2017). In all instances, there are common themes that emerge in terms of characterizing adaptive learning. Specifically, it is characterized by providing learners a more personalized learning experience that is based on their current level of knowledge and skills, learning goals, and the gap between the two. In other words, adaptive learning techniques focus on moving beyond a one-size-fits-all approach to learning to one where learners can go at their own pace and access content that is more appropriate for their current level of knowledge and skill.

There are several constructs in the literature that describe why adaptive learning is an important concept to explore. First, the Zone of Proximal Development (ZPD; Vygotsky, 1978) articulates that there is a "sweet spot" for learning – if an activity is too challenging, learners may become frustrated and give up; similarly, if it is too easy, little learning will occur. Therefore, learning experiences must be just challenging enough to push them beyond their comfort zone. Adaptive learning techniques aim to keep learners within their ZPD by providing content that is at just the right level. Another important construct in the literature that has had significant bearing on the idea of adaptive learning is aptitude by treatment interactions (ATIs; Cronbach & Snow, 1981; Snow, 1991). ATIs describe how the outcome of a treatment (in this case, learning material) may differ based on the initial aptitude of the learner. The concept of ATIs is instrumental in understanding why adaptive learning is necessary, as it points to the need for more individualized instruction given that each learner begins the learning experience with a slightly different set of initial skills and knowledge. However, the concept of ATIs also highlights that individual differences may impact the effectiveness of any given adaptive learning approach.

Approach to Adaptive Learning

There are many different approaches to adaptive learning. Our approach is data-driven and combines Bayesian Knowledge Tracing (BKT; Anderson 1995) with Item Response Theory using Partial Credit Models (PCM; Masters 1982). PCM links performance measures together to create a complete assessment of the learner at a certain time point; BKT addresses how skills advance over time and uses a Hidden Markov Model (HMM) to describe progress according to the elements described in Table 1. By mathematical computation/reasoning about the skills and the measures taken to assess the current state of an individual, BKT can track which skills are trained and which skills are not. However, by itself, BKT only implements measures as a single probability.

Table 1. Components of the Hidden Markov Model

Model Component	Description
Init	The probability that the student has a skill prior to training
Transit	The probability the student learns the skill during training
Guess	The probability the student does not learn the skill, but guesses correctly anyway
Slip	The probability the student has the skill but makes a mistake anyway.

Our approach augments BKT in two ways. First, measures are adjusted by item difficulty using PCM. Whereas BKT focuses on the evolution of a skill over time, PCM is a model that can generate a better model of item performance, in that skill level can be related to task difficulty. The "two parameter" version of PCM models item performance as:

$$\Pr(X = x) = \frac{e^{\sum_{k=0}^{x}(S-D_k)}}{\sum e^{\sum_{k=0}^{j}(S-D_k)}}$$

where S is a parameter representing trainee ability and D_k is a number representing item difficulty at the credit threshold.

Second, developed models account for multiple skills and the relationships between them. To do this, our approach combines BKT and PCM into a POMDP (Partially Observable Markov Decision Process), which extends the above into a decision process. Previous work (Carlin, Dumond, Dean, & Freeman, 2013; Carlin, Oster, Nucci, Kramer, & Brawner, 2016; Levchuk, Gildea, Freeman, & Shebliske, 2007; Levchuk, Shebilske, & Freeman, 2012) introduced multiple training actions that were each able to train skills. Each training action is associated with a different Transit probability (see Table 1), and the learner is modeled as being in a high-skill, medium-skill, or low-skill state, with high-skill state as the training goal. The problem addressed by the POMDP is action selection – that is, to select the action that most quickly achieves this training goal.

Purpose of the Paper

The purpose of this paper is to explore the effectiveness of adaptive learning, specifically in regard to identifying conditions under which adaptive learning is most beneficial. As discussed, the concept of ATIs provides evidence that not all training and methods will be received equally across individuals. Therefore, there are likely conditions under which adaptive learning is more appropriate and effective than others. This research begins to investigate such conditions.

METHOD

Participants

Participants were recruited through Amazon's Mechanical Turk (AMT). AMT is a crowdsourcing marketplace that makes it easy for individuals and businesses to outsource their processes and jobs to a distributed workforce who can perform these tasks virtually. Data were collected from 76 participants. Additional participants attempted the task; however, after engaging in data cleaning procedures, their data were not used in the final analyses. Specifically, data from participants who correctly answered all of the pretest questions were not utilized (n = 26); also participants who demonstrated a lack of effort (as defined by spending less than 10 seconds on any given problem set) were also dropped from the final data set (n = 8). Participants were distributed using a randomizer across a control condition (n = 49) and an adaptive training condition (n = 27)[1].

The Task

To conduct this research, we used problems from http://www.math-aids.com/ to create a basic math task that tests the participant's knowledge of basic order of operations. We chose various problems that could be grouped into different levels of difficulty based on the number of skills exercised, difficulty with respect to those skills, and the number of steps required to complete the problem. For example, an easy problem [(30 - 6) ÷ 4 – 42] contains two skills (subtraction and division), basic levels of skill difficulty, and requires three steps to complete. A more difficult problem [(-7) - { (-12)x2 ÷ (-6)x2 }2] • (-9)x + (-11)x] contains eight skills (addition, subtraction, multiplication, division, exponents, negative numbers, parenthetical operations, and variables), advanced levels of skill difficulty, and requires five steps to complete. Altogether, problem sets at 20 levels of difficulty were chosen. Each participant was given a diagnostic pretest as well as a posttest, each with three problems representing three difficulty levels (levels 5, 11, and 17).

Control Condition

Participants randomly assigned to the control condition started with problems at a difficulty level of 1 and were advanced to the next level once they successfully completed all three problems in the set. Lesson groups were numbered to reflect their difficulty level – for example, Lesson Group 1.1 and 1.2 have an overall difficulty level of 1, Lesson Group 2.1 and 2.2 have an overall difficulty level of 2, and so on. If all three problems were not answered correctly, a second set of problems at the same level of difficulty was presented, continuing in this manner for all problem sets. Table 2 provides example problems, as well as the skills addressed in those problems. Participants in the control condition were given a total of five sets of three problems.

Table 2. Example Problems, Skills Addressed, and Difficulty Levels

		Targeted Skills and Associated Difficulty			
Lesson Group	*Example Problem*	*Addition*	*Subtraction*	*Multiplication*	*Division*
1.1	2 + 6 + 2 * 9	X		X	
1.1	5 + 15 * 16 +12	X		X	
1.1	14 - 12 * 3 + 2	X	X	X	
1.2	3 + 2 +19 * 2	X			
1.2	8 - 6 * 12 + 9	X	X	X	
1.2	3 + 9 + 15 * 5	X		X	
2.1	78 ÷ 13 + 4 • 11	X		X	X

Experimental Condition

We first estimated the initial skill level of participants randomly assigned to the experimental condition based on scores from the pretest. Specifically, the results of the pretest were used to classify participants to direct them to an initial problem set that best aligned with their initial skill level. As shown in Table 3, participants in the experimental condition were classified as either Proficient (getting all three pretest problems correct, and thus testing out of the experiment); Advanced (getting two problems correct and the third incorrect); Intermediate (only getting one problem correct); or Beginner (not getting any of the three problems correct). There were also several profiles of users that did not correlate to one of those groups (the "Mixed" profile in Table 3); those individuals were given an initial problem set that best represented (to the extent we could) their initial ability. Regardless of performance, participants in the experimental condition were given a total of four sets of three problems[2].

---

[1] Following an initial set of participants, changes were made to the adaptive model to better align the number of exercises between the two conditions, as well as to ensure that the problems in the experimental condition were sufficiently challenging. Due to the change, less usable data were available in the experimental condition.

[2] A slightly smaller number of problem sets was chosen in the experimental condition versus the control condition to minimize participant fatigue. Pilot data demonstrated that participants in the experimental condition were taking longer on the problems, likely due to the increased difficulty.

Table 3. Guidelines for Placing Experimental Participants in their First Problem Set

	Level of Pretest Problem			
Participant Profile	Level 5	Level 11	Level 17	*Recommendation for First Problem Set*
Proficient	C	C	C	Test Out – no further problems
Advanced	C	C	I	Level 11
Intermediate	C	I	I	Level 4
Beginner	I	I	I	Level 1
Mixed 1	C	I	C	Level 7
Mixed 2	I	I	C	Level 4
Mixed 3	I	C	C	Level 11

C = Correct; I = Incorrect

Within the experimental condition, the next problem set was chosen by an adaptive training policy based on how well a participant did on the previous problem set. The adaptive training policy was constructed by modeling student progression as a POMDP and selecting problems by optimizing a policy using that model. The model itself consisted of eight skills (Addition, Subtraction, Multiplication, Division, Positive/Negative Numbers, Variables, Parentheses, Exponents), and each problem was tagged for which knowledge applied to it, as well as a difficulty level of that knowledge. The adaptive training policy software performed two operations after each problem set. In the first operation, the software would reassess learner progress on each of the relevant skills using PCM. Thus, learner assessments were upgraded when problems were answered correctly and downgraded when problems were answered incorrectly; the magnitude of the reassessment depended on the difficulty level of the problem (i.e., learners who answered easy questions incorrectly were more likely to be classified as novices than learners who answered difficult questions incorrectly). In the second operation, the next problem for each learner was selected based on the problem profile that best matched the learner's updated assessed state on the tested skills.

Because adaptive policy recommendations are deterministic (e.g., all students who received 2 out of 3 correct on the pretest, followed by getting 3 out of 3 on the first problem set correct, received the same recommendations), it is possible to analyze adaptive recommendations as shown in Table 2. Thus, the adaptive policy specified that if a participant got all three problems in one set correct, the difficulty level of the next set increased three groups (e.g., from 1.1 to 4.1); if the participant got two of the three problems correct, the difficulty level of the next set went up one group; and if the participant only got one of three problems correct, the next set of problems remained at a similar difficulty group to the problem set just completed. Finally, not getting any problems correct resulted in a reduced difficulty level for the next problem set.

RESULTS

Participants across both conditions spent an average of 15 minutes and 19 seconds (*SD* = 7:41) on the total task, including the time spent on the pre and posttests. Participants in the control condition spent an average of 14:27 minutes (*SD* = 7:20) on the full task, and Participants in the experimental condition spent an average of 16:11 minutes (*SD* = 8:32). This difference was not significant.

Participants' scores were computed in terms of percent correct (e.g., a score of 50 is 50% correct). In the control condition, the average score on the pretest was 36.08 (*SD* = 28.07), with a similar average pretest score obtained by those participants in the experimental condition ($M = 39.52$; $SD = 26.39$). This difference was not significant, $t(74) = -.52$, $p = .60$, indicating that participants in the control and adaptive conditions began the task with a similar level of knowledge. For the posttest scores, those in the control condition scored an average of 48.98 (*SD* = 39.21), and those in the experimental condition scored an average of 59.22 (*SD* = 32.59). Although these results show no main effect for the learning condition ($t(74) = -1.15$, $p = .25$), the trend shows greater learning gains for those participants in the experimental condition. See Figure 1 for a depiction of the results. Specifically, individuals in the control condition experienced an average gain of around 12% from pre to posttest, whereas the gain for participants in the experimental condition was approximately 20%. Given the small sample size, this trend is promising in terms of illustrating the general benefits of the adaptive learning approach.

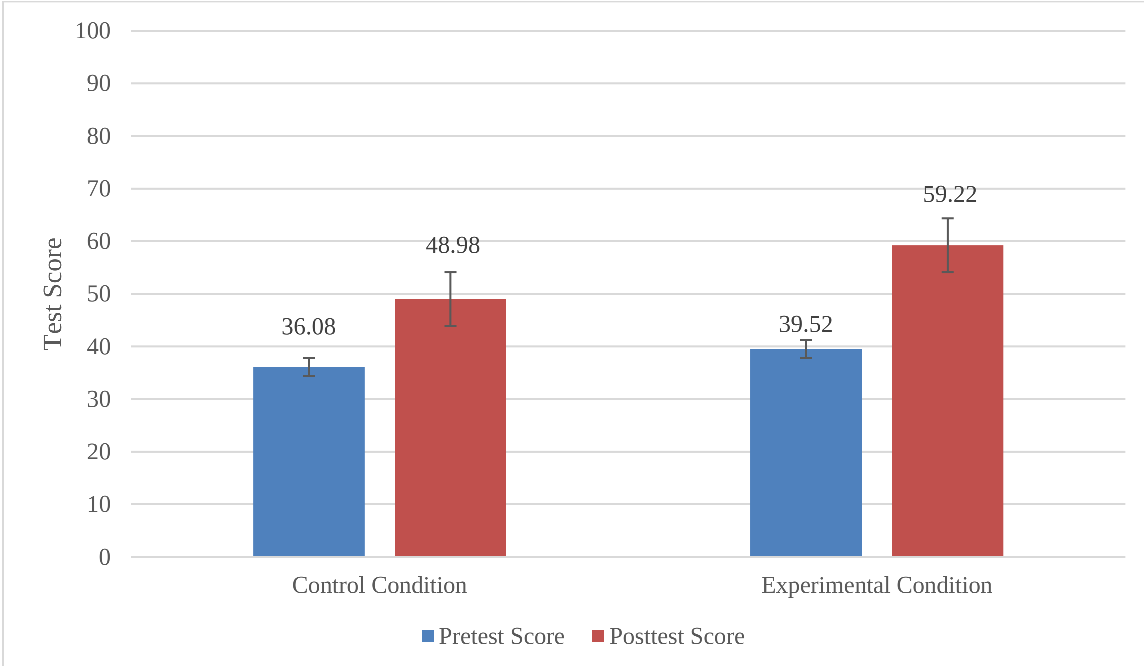

Figure 1. Pre vs. posttest scores for participants in the control and experimental conditions. Participants in the experimental condition demonstrated greater gains (although not statistically significantly so) in scores compared to those in the control condition.

To continue exploring the impact that the adaptive training approach may have had on participant learning, participants were grouped according to their performance on the pretest. Participants were segmented into three groups: *Beginner* (participants that did not correctly answer any questions on the pretest); *Intermediate* (participants that correctly answered one-third of the pretest questions); and *Advanced* (participants that correctly answered two-thirds of the pretest questions). This grouping variable was used as a moderator variable to examine the impact of participant's initial skill level on pre to posttest performance within each condition.

The interaction plot is in Figure 2. Although this interaction is not statistically significant ($F(2,75) = 1.34$, $p = .29$; $\eta^2 = .031$), that is likely due to the small sample size. First, as can be seen, the greatest gains in scores were for those participants in the experimental condition who started as Beginners; for that group, the difference in pre to posttest scores was 44.33 (or 44.33%). This gain is in comparison to a change in score for the Intermediate group by 13.50% and 11.91% for the Advanced group for the experimental condition. In both conditions, the smallest differences in learning gains were demonstrated by those that were classified as Advanced. In fact, in the control condition, there was very little difference between pre and posttest scores for those Advanced individuals. Second, not only did the magnitude of learning gains differ across the three groups of participants, but the effectiveness of each condition differed depending on the starting state of the participant. For individuals in the Intermediate group, the largest gains in knowledge were seen for participants in the control condition, whereas the opposite is true for those individuals who were classified as either Beginners or Advanced. Consistent with the literature on ATIs, this finding demonstrates that personalized training may have differential benefits depending upon one's initial level of expertise or knowledge in the area to be trained.

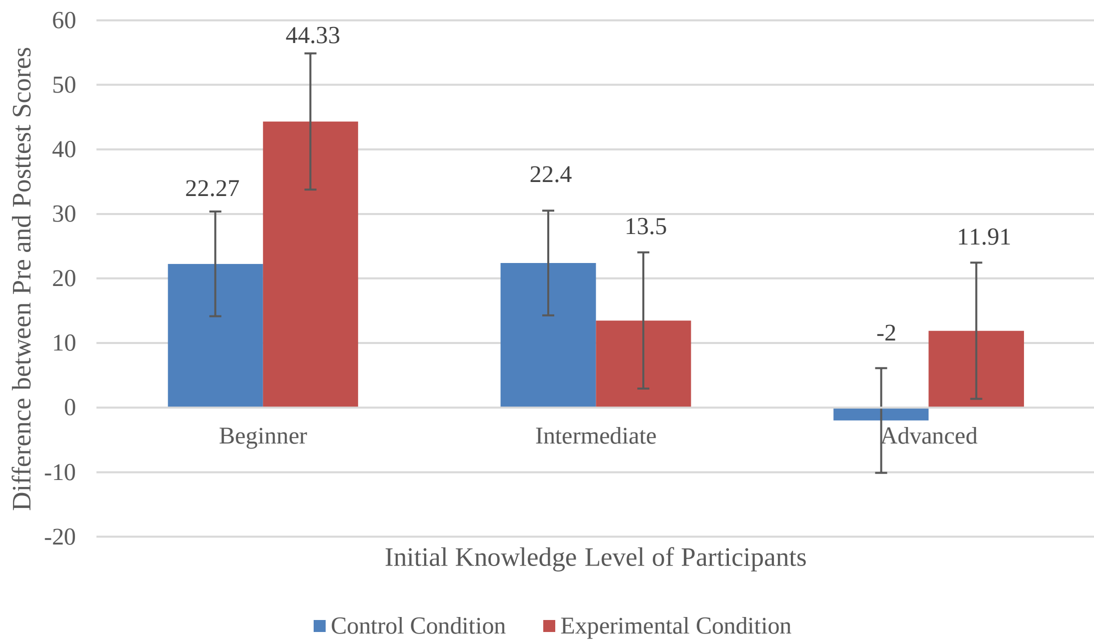

Figure 2. Differences in gains from pre to posttest among three groups of participants. Participants that began the experiment with the least level of knowledge and were in the experimental condition demonstrated the greatest gains in knowledge.

To further explore these findings, total time on task was calculated for the above three groupings of participants according to their experimental condition. As can be seen in Figure 3, Beginner participants spent about the same time on task regardless of condition; Intermediate participants in the Control condition spent more time on the task than Intermediate participants in the experimental condition, whereas the opposite pattern emerged for the Advanced participants. Taken together with Figure 2, these results demonstrate that those participants with the largest skill gains (Beginner participants in the experimental condition) took about the same time to complete the task as Beginner participants in the control condition *and* experienced a larger gain in knowledge. This result is important, as it has implications for both the efficiency and effectiveness of the adaptive learning approach.

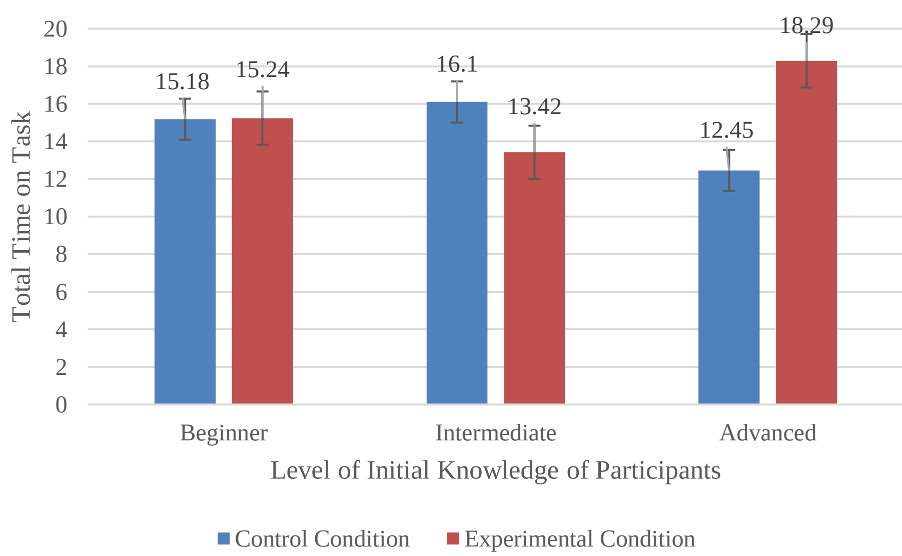

Figure 3. Total time on task.

DISCUSSION

This research sought to demonstrate the impact of adaptive learning approaches. The resulting trends show that, overall, participants in the experimental condition demonstrated greater learning gains than participants in the control condition. Although the differences were not statistically significant (likely due to small sample sizes), the results begin to demonstrate support for our adaptive learning approach advancing individuals to increased levels of expertise compared to a non-adaptive approach. As exhibited by the graph in Figure 2, for all participants in the experimental condition (no matter their initial level of knowledge), the minimum learning gain from pre to posttest was around 12%. The same cannot be said for the control condition, where one group of participants (the Advanced participants) did not demonstrate any learning gains at all.

When broken down further, we see that the positive impact of the adaptive training approach was more pronounced for individuals who began the task with a more basic knowledge of the domain (i.e., order of operations). Overall, those participants classified as a Beginner demonstrated greater gains in performance compared to those participants who were more advanced in their knowledge, demonstrating that the tailored learning path provided problems that kept them in their learning “sweet spot.” Learners classified as Advanced also benefited more from the adaptive learning approach compared to the linear approach of the control condition, albeit not as much as the more novice learners. These findings are not surprising given that (a) the problems presented in the control condition were too easy for Advanced learners, thereby producing a null effect on learning; and (b) smaller gains in learning are understandable given that a ceiling effect likely exists (meaning, Advanced learners did not have as much room to improve given they began the task with a higher skill level). Interestingly, those individuals that started the task with an intermediate level of knowledge benefited more from the linear control condition than the adaptive learning path. This finding aligns directly with the idea that most curriculum is designed to benefit the “average” learner. In this experiment, individuals with an intermediate level of knowledge represent those “average” learners, whose level of knowledge is most well-aligned with the problems presented in the more traditional, linear approach to learning.

We also delved briefly into the idea of efficiency of using an adaptive learning approach. Such an approach is more costly and may perhaps also take more time for learners to progress through the learning material. Our data illustrate, however, that our adaptive learning approach produced greater gains in learning *in about the same amount of time* as the control condition for participants classified as beginners. This finding provides initial support for the idea that adaptive learning approaches are useful for bringing individuals to a higher level of expertise within the same amount of time. This pattern was not replicated for the other group of participants who also benefited from the adaptive learning approach (the Advanced participants). For this group, while they ultimately experienced greater gains in learning with the adaptive approach, it took them longer to do so compared to the control condition. However, given that the Advanced participants in the control condition essentially demonstrated no gains in learning, the data demonstrate that adaptive learning produces a greater return on investment for those more advanced learners even with the additional time needed for learning.

Limitations and Challenges

The main limitations of this research were related to the sample, as well as the time available with the sample to conduct the task. First, while AMT provided a viable way to source participants, it also has drawbacks. On several instances there were network glitches, locking participants out from the task they started, and on one occasion there was a cyber spam that registered over 500,000 requests for the task when only 30 should have been attempted. At that point, we updated some of our software backend and resumed the experiment after a week's delay. There were also some abnormal patterns in the data that produce some uncertainty in the results. For example, there were participants that spent a long time on problem sets but it is not clear if that time was spent in an effortful way on the task, or possibly the participant simply stepped away from the task for some time. When using AMT, it is not possible to control for the variety of extraneous factors that may impact participant behavior.

One of the other limitations to this work is associated with the time period – can we really expect learning will occur over such a short period of time? While the results demonstrate trends that show that the adaptive approach produces greater knowledge gains, how this translates to “real world” tasks that occur over a longer period of time is unknown. Ideally, this same type of approach would be tested with learners over a course of many weeks.

Practical Implications and Guidelines

This research lends credence to the use of adaptive learning approaches for increasing the learning gains that can be made in learners. As many organizations, both military and non-military, are interested in using adaptive learning approaches, it is important to discuss the practical guidance that can be derived from this work, especially given that adaptive learning approaches are more difficult to implement and may not produce the same level of throughput (an important consideration when there is a large number of learners at any given time). The following represent some of the more practical guidelines to emerge from this work, both from operational and research perspectives, in relation to adaptive learning:

1. *Understand your learner*. The data from this research align well with the classic research on ATIs. Specifically, the data demonstrate that not all learners will respond in the same way to all learning interventions. Therefore, it is important to understand the learner population with whom you are working. If all learners fall into the "average" category, or possess the same skill level as one another, investing in adaptive learning processes and approaches may not be beneficial. As learning activities are crafted, it is important to develop an approach that aligns with the needs of the targeted learner population.

2. *Understand your training content.* Like needing to understand the learner population, it is also necessary to understand the goals of your training content. Adaptive learning methods may not be necessary for learning objectives that are heavily focused on lower levels of learning. For example, if the primary objective of a learning experience aligns with a lower level of learning (e.g., as put forth in Bloom's revised taxonomy; Anderson, Krathwohl, & Bloom, 2001) such as basic declarative knowledge, a linear approach to teaching (e.g., a lecture) may be sufficient. If, however, a deeper level of learning is required, and the goal is to teach learners how to evaluate or create information, adaptive learning approaches may be more beneficial. To our knowledge, no work has yet examined the levels of learning where adaptive learning approaches may have more or less of an effect. Future research should examine this, and those managing and creating learning experiences should consider the desired learning level when making the decision as to the benefit of adaptive learning approaches.

3. *Keep the human in the loop.* During the initial phases of this research, we made several adjustments to the parameters in our adaptive learning model after doing some initial testing. Prior to the experiment, we played the roles of different possible representative students, walked them through the adaptive curriculum, and validated that the recommendations were reasonable. In other words, we played the role of a student who got every problem right, and looked at the recommendations, followed by playing the role of a student who got every problem wrong, and reviewing those recommendations, etc. Based on the recommendations made, changes were instantiated to ensure that the adaptive recommendations made by the algorithms were reasonable. Thus, while one benefit of adaptive learning approaches is that models and algorithms can be built to automatically determine the next best learning experience (as opposed to having a human handpick the next experience), it is always beneficial to ensure that the adaptive learning policy has some review and oversight by a human in the loop.

4. *Choose your adaptive learning model wisely.* There are various modeling approaches that can be used to make learning recommendations. To more easily and efficiently utilize an adaptive learning approach, it is necessary to use a modeling approach that is simple, extensible, and repeatable. One advantage of the approach used here is that it has fewer "moving pieces" (i.e., models) that need to be implemented to make the recommendation. Such an approach increases the flexibility of the system where it is implemented and makes it more repeatable over time.

Overall, adaptive learning methods are powerful and can create more effective and efficient learning experiences for individuals if used in ways that are appropriate for the situation. The use of adaptive learning methods must be thoughtfully considered and implemented, just as instructional designers determine the overall appropriate instructional strategy. Future research should continue to explore the situations where adaptive learning may be most useful.

ACKNOWLEDGEMENTS

We would like to acknowledge Aptima, Inc. for supporting this work through an Internal Research and Development (IR&D) investment.

REFERENCES

Anderson, J. R., Corbett, A. T., Koedinger, K. R., & Pelletier, R. (1995). Cognitive tutors: Lessons learned. *The Journal of the Learning Sciences, 4*(2), 167-207.

Anderson, L. W., Krathwohl, D. R., & Bloom, B. S. (2001). *A taxonomy for learning, teaching, and assessing: A revision of Bloom's taxonomy of educational objectives*. New York, NY, USA: Allyn & Bacon.

Carlin, A., Dumond, D., Dean, C., & Freeman, J. (2013). *Higher Automated Learning*. Short paper at Conference Artificial Intelligence in Education (AIED).

Carlin, A., Oster, E., Nucci, C., Kramer, D., & Brawner, K. (2016). *Educational Data Mining Using GIFT Cloud*. GIFTSym 5.

Cronbach, L. J., & Snow, R. (1981). *Aptitudes and Instructional Methods: A Handbook for Research on Interactions.* Irvington Publishers.

Department of the Army. (2011). *Army Learning Concept for 2015*. TRADOC Pamphlet 525-8-2. Fort Monroe, VA: Training and Doctrine Command.

Levchuk, G.M., Gildea, K., Freeman, J., and Shebilske, W. A. (2007). *Benchmarked Experiential System for Training (BEST)*. Proceedings of the 12th International Command and Control Research and Technology Symposium, Newport, RI.

Levchuk, G., Shebilske, W., & Freeman, J. (2012). *A Model-Driven Instructional Strategy*. Adaptive Technologies for Training and Education, 303.

Mao, Y., Lin, C., & Chi, M. (2018). Deep learning vs. bayesian knowledge tracing: Student models for interventions. *Journal of Educational Data Mining*, *10*(2):28–54.

Masters, G. N. (1982). A Rasch model for partial credit scoring. *Psychometrika*, *47*(2), 149-174.

Roberson, D. L., & Stafford, M. C. (2017). *The Redesigned Air Force Continuum of Learning*. Air University Press, Air Force Research Institute.

Snow, R. E. (1991). Aptitude-treatment interactions as a framework for research on individual differences in psychotherapy. *Journal of Consulting and Clinical Psychology, 59*(2), 205-216.

U.S. Fleet Forces Command (2017). *Vision and Guidance for Ready Relevant Learning*.

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

Vygotsky, L. (1978). Interaction between learning and development. *Readings on the Development of Children, 23*(3), 34-41.