

Enhancing Learning Outcomes Through Adaptive Remediation with GIFT

Randall Spain¹, Jonathan Rowe¹, Benjamin Goldberg², Robert Pokorny³, and James Lester¹

¹North Carolina State University

²U.S. Army CCDC – STTC

³Intelligent Automation Inc.

Raleigh, NC

Orlando, FL

Rockville, MD

{rdspain, jprowe, lester}@ncsu.edu

benjamin.s.goldberg.civ@mail.mil

bpokorny@i-a-i.com

ABSTRACT

Adaptive instructional systems (AISs) are envisioned to play a vital role in the Army's future training environment. A key feature of AISs is the capacity to automatically tailor instruction to fit the needs and skills of individual learners. Leveraging recent advances in artificial intelligence and machine learning, training and educational experiences can be tailored to the goals, learning needs, and preferences of individual learners and teams of learners. Tutorial planning, a critical component of adaptive training, controls how scaffolding and instructional interventions are structured and delivered to learners to create personalized learning experiences. Devising computational models that effectively scaffold learning experiences is a critical challenge for the field. For example, AISs need to determine when to scaffold, what type of scaffolding to deliver, and how scaffolding should be realized, all in real time. In this paper, we describe our work using the Generalized Intelligent Framework for Tutoring (GIFT), an open source framework for creating, deploying, and evaluating adaptive training systems, to create a web-based adaptive short course for teaching fundamental principles associated with counterinsurgency (COIN). The course presents students with a series of videos, integrated assessments, and remediation materials about doctrinal COIN concepts. The course's adaptive remediation features are based on the ICAP active learning framework to deliver constructive, active, and passive forms of remedial feedback to students. We report the results of a recent study in which 500 participants completed the adaptive training course along with pre- and post-training knowledge tests. The paper provides an analysis of learning gains and factors that moderated these gains and concludes with a discussion of future research as it pertains to the goals of the broader research program investigating applications of machine learning, and reinforcement learning in particular, to automatically generate policies for instructional remediation in AISs.

ABOUT THE AUTHORS

Dr. Randall Spain is a Research Psychologist in the Center for Educational Informatics at North Carolina State University where he uses principles, theories, and methods of applied psychology to design and evaluate the impact of advanced training technologies on learning and performance. He has conducted human factors research for the Department of Defense and the Department of Homeland Security for the past 10 years with a focus on adaptive training, performance assessment and measurement, user modeling, and human-automation interaction. Dr. Spain received his PhD in Human Factors Psychology from Old Dominion University. He currently serves on the editorial board for *Military Psychology*.

Dr. Jonathan Rowe is a Research Scientist in the Center for Educational Informatics at North Carolina State University, as well as an Adjunct Assistant Professor in the Department of Computer Science. He received his PhD and MS degrees in Computer Science from North Carolina State University, and a BS degree in Computer Science from Lafayette College. His research focuses on the intersection of artificial intelligence and human-computer interaction for advanced learning technologies, with an emphasis on game-based learning environments, intelligent tutoring systems, multimodal learning analytics, user modeling, and computational models of interactive narrative generation.

Dr. Benjamin Goldberg is a senior researcher in the Learning in Intelligent Tutoring Environments (LITE) Lab at the Combat Capabilities Development Command (CCDC) Soldier Center, Simulation and Training Technology Center (STTC) in Orlando, FL. He has been conducting research in modeling and simulation for the past eight years with a focus on adaptive learning and how to leverage artificial intelligence tools and methods for adaptive computer-based instruction. Currently, he is the LITE Lab's lead scientist on instructional strategy research within adaptive training environments. Dr. Goldberg holds a PhD from the University of Central Florida in Modeling & Simulation.

Dr. Robert Pokorny is Principal of the Education and Training Technologies Division at Intelligent Automation, Inc. He earned his PhD in Experimental Psychology at the University of Oregon in 1985, and completed a postdoctoral appointment at the University of Texas at Austin in Artificial Intelligence. His first position after completing graduate school was at the Air Force Research Laboratory, where he developed methodologies to efficiently create intelligent tutoring systems for a wide variety of Air Force jobs. At Intelligent Automation, he has led many cognitive science projects, including adaptive visualization training for equipment maintainers, and an expert system approach for scoring trainee performance in complex simulations.

Dr. James Lester is Distinguished University Professor of Computer Science at North Carolina State University, where he is the Director of the Center for Educational Informatics. His research centers on transforming education with technology-rich learning environments. With a focus on adaptive learning technologies, his research spans intelligent tutoring systems, game-based learning environments, affective computing, and tutorial dialogue. The adaptive learning environments he and his colleagues develop have been used by thousands of students in K-12 classrooms. He received his PhD in Computer Science from the University of Texas at Austin in 1994. He is a Fellow of the Association for the Advancement of Artificial Intelligence (AAAI).

Enhancing Learning Outcomes Through Adaptive Remediation with GIFT

Randall Spain¹, Jonathan Rowe¹, Benjamin Goldberg², Robert Pokorny³, and James Lester¹

¹North Carolina State University

²U.S. Army CCDC – STTC

³Intelligent Automation Inc.

Raleigh, NC

Orlando, FL

Rockville, MD

{rdspain, jprowe, lester}@ncsu.edu

benjamin.s.goldberg.civ@mail.mil

bpokorny@i-a-i.com

INTRODUCTION

Adaptive instructional systems (AIS) will play a central role in the next generation of training systems for the military (TRADOC, 2017). AISs are computer-based systems that can readily tailor the content and focus of training to suit the needs of individual learners with the potential to alter aspects like time on task, content, practice examples and instructional strategy (Durlach, 2012; Sottolare, Barr, Robson, Hu, & Graesser, 2018). Numerous studies have shown that AISs, such as intelligent tutoring systems (ITS), have the potential to be as effective as a skilled human tutor at supporting student learning (VanLehn, 2011). Yet, determining what to adapt, when to adapt, and how to adapt instruction is a critical challenge facing AIS developers. This challenge stems in part from the wide range of pedagogical strategies and tactics that can be implemented in adaptive systems, as well as a lack of empirically grounded guidance about the relative contribution of different adaptive interventions on learning outcomes (Durlach & Ray, 2011).

Tutorial planning is a critical component of AISs, controlling how scaffolding is structured and delivered to learners. Tutorial planners operate at multiple levels, including the macro-level (e.g., selecting problems for learners to solve) and micro-level (e.g., delivering tailored hints and feedback about specific problems), and they perform tutorial actions that are both domain and learner-dependent, as well as independent. In the classical ITS architecture, a tutorial planner is a component of the pedagogical module. It contains a set of formalized pedagogical rules that an ITS uses to support learning. These rules are typically based on learning theories and provide general guidance on when to scaffold learners, what type of scaffolding to provide, and how scaffolding should be realized.

Despite their important role in devising adaptive training experiences, tutorial planners often suffer from three significant limitations. First, creating tutorial planners is expensive, requiring labor-intensive knowledge engineering processes that involves close collaboration between subject matter experts, education experts, and software developers (Murray, 2003). Second, tutorial planners often remain in a fixed state after they have been developed. The rules within the planner do not improve or change over time unless they are manually updated by a course author. Third, tutorial planners often model rules for scaffolding by using symbolic representational techniques, which unlike probabilistic representations, do not capture the uncertainty inherent in modeling student learning with AISs.

Recent developments in artificial intelligence and machine learning have introduced opportunities to reduce the authoring burden of AISs by devising data-driven tutorial planning policies that can automatically control how pedagogical support is structured and delivered to learners (Rowe & Lester, 2015; Williams et al., 2016; Zhou, Wang, Lynch, & Chi, 2017). In data-driven tutorial planning, the tutorial planner is an agent that makes a series of decisions about what pedagogical strategies to enact in order to optimize the trainees' learning outcomes. Unlike manually engineered tutorial planning policies, data-driven tutorial planners use student interaction data to induce models for pedagogical decision making that are explicitly designed to maximize learning. By analyzing students' interactions with an AIS and modeling accumulated rewards, the planner can make informed decisions about delivering feedback, remediation, and assessments that are tailored to the needs of individual learners.

In particular, reinforcement learning techniques have shown promise for automatically generating tutorial planning rules that optimize student learning outcomes without the need for experts or course authors to manually program or demonstrate these rules. Reinforcement learning is a type of machine learning that centers on devising software agents that perform actions in a stochastic environment using an optimal policy (Sutton & Barto, 1998). The agent generates the control policy by iteratively performing actions and observing their effects on the environment and its accumulated

rewards. However, reinforcement learning techniques are data intensive, which presents challenges for devising and evaluating reinforcement learning-based tutorial planners.

Over the past five years, the Generalized Intelligent Framework for Tutoring (GIFT) has become a key exemplar of how the challenges in developing AISs at scale can be addressed (Sottolare, Brawner, Goldberg, & Holden, 2012; Sottolare, Brawner, Sinatra, & Johnston, 2017). GIFT is an open-source domain-independent framework for designing, deploying, and evaluating adaptive training systems. GIFT provides instructors with a suite of web-based tools for rapidly creating intelligent tutors, and it is linked to several ongoing research efforts to devise methods for automating key elements of the adaptive training authoring process. Many of these tools are available through GIFT's Course Creator, which provides a drag-and-drop interface for devising adaptive training experiences across a range of domains.

In this paper we describe results from a study that is part of a larger research program that aims to devise data-driven tutorial planning policies that can be used in GIFT to present learners with adaptive remediation. In particular, the paper presents preliminary results from a study involving over 500 learners who completed an approximately 90-minute hypermedia training course in GIFT that taught doctrinal concepts associated with counterinsurgency (COIN) and stability operations. The course contained integrated assessments and provided learners with adaptive remediation materials about doctrinal COIN concepts when needed. The paper addresses the following research questions:

- 1) Did participants demonstrate positive learning gains as a result of completing the adaptive training course?
- 2) Did any factors influence or moderate participants' learning gains?
- 3) Which forms of remediation were most effective at helping learners advance through the course?

Adaptive Training Course

Reinforcement learning techniques are data intensive, so a critical goal of this study was to design a training course that contained numerous opportunities for learners to receive instructional remediation. A second goal was to develop a system that was deployable through online crowdsourcing platforms to facilitate broad distribution to many learners. Key to this goal was capturing learning outcome data as well as learning trace data to be used for evaluating student learning behaviors and developing reinforcement learning-based tutorial policies. To meet these objectives, we developed an adaptive hypermedia-based training course in GIFT that built upon materials from the UrbanSim Primer. Originally developed by the USC Institute for Creative Technologies, the UrbanSim Primer is a self-paced, interactive tutorial that provides direct instruction on the core doctrinal principles of COIN and stability operations (McAlinden, Gordon, Lane & Pynadath, 2009). The Primer contains eight units that review the tenets of FM 3-07 (Stability Operations), FM 3-24 (Counterinsurgency) and the updated FM 3-0 (Operations). Each unit contains a series of videos in which subject matter experts review key concepts associated with leading COIN operations. The purpose of the UrbanSim Primer is to establish or refresh foundational knowledge regarding COIN doctrine prior to trainees engaging with the UrbanSim learning environment.

The training course we devised in GIFT contained content from the UrbanSim Primer and was organized into 4 chapters. Each chapter contained a series of short instructional videos, recall questions, and remedial training interventions designed to teach common themes, terminology, and principles of COIN operations (Rowe, Spain, Pokorny, Mott, Goldberg, & Lester, 2018). The course's instructional videos were approximately 90 seconds in length and covered topics such as "Identifying the center of gravity in COIN operations," "Defining intelligence preparation for the battlefield," and "Understanding lines of effort in COIN operations." The recall questions were presented in multiple choice format and were designed to assess students' knowledge and understanding of the content covered in the videos. The remediation interventions were implemented by leveraging recent enhancements to GIFT's Engine for Management of Adaptive Pedagogy (EMAP) that support adaptive assessment and remediation.

Specifically, new features in GIFT allow course authors to present learners with remediation interventions that are based on the ICAP active learning framework (Chi, 2009). This framework identifies four modes of student engagement: Interactive, Constructive, Active, and Passive; each of which is associated with different learning behaviors and cognitive processes. (1) *Passive* engagement, such as listening to a lecture or watching a video without taking notes, is the least active form of learning. A student engaged in passive learning is less likely to remember what has been discussed in the lecture than someone who engages in a more active learning strategy. (2) *Active* engagement is deeper than passive engagement. It calls for the student to engage in an overt action while learning, such as

highlighting or underlining text, copying lecture notes verbatim, or answering a multiple-choice question. In this mode of engagement, the student is attending to and assimilating knowledge during the learning experience, but the level of cognitive engagement is still somewhat limited because students are not generating new content. (3) *Constructive* engagement is a deeper form of learning. In constructive engagement, students produce new content beyond what is provided in the learning materials. A prime example of constructive engagement is self-explaining concepts, which require students to probe existing knowledge and explain concepts in their own words. (4) *Interactive* engagement activities are the most demanding form of cognitive engagement. They require students to engage in meaningful interaction with another peer, teacher, or system about a topic. During interactive engagement, the learner asks questions, explains concepts, and challenges or critiques her own or her partner's responses or rationale.

The remediation activities authored in the GIFT course required students to either passively, actively, or constructively engage with remedial content upon missing a recall question. Constructive remediation interventions required learners to read a transcription of the video they had just viewed and to provide a short answer response to the question they missed. Active remediation required learners to read the transcription of the video and highlight the passage of text that corresponded to the correct answer for the question they missed. Passive remediation required learners to only read the video transcription; they were not required to engage in an overt learning behavior (Interactive remediations required a system that could interact with the student and was beyond the scope of this project.) The ICAP model predicts that constructive remediation (e.g., writing an explanation), should be more effective than active remediation (e.g., reading and highlighting a passage), and that active remediation should be more effective than passive remediation (e.g., reading a passage without doing anything else) with respect to increasing student learning outcomes. However, there are tradeoffs between these pedagogical strategies, including impacts on instructional time, the amount of effort and cognitive load imposed on the learner, and course authoring burden.

Below, we report the results from a recent study in which participants completed the adaptive training course described above, along with pre- and post-training knowledge tests. The paper provides an analysis of learning gains and factors that moderated these gains and concludes with a discussion of future research as it pertains to the goals of the broader research program investigating applications of machine learning, and reinforcement learning in particular, to automatically generate policies for instructional remediation in AISs.

METHOD

Participants

To administer the course to a broad distribution of learners, participants were recruited through Amazon's Mechanical Turk (MTurk) platform. A total of 533 participants (42% female) completed the study. Participants received \$8 dollars for completing the course which lasted approximately 90 minutes. To be eligible, participants had to be at least 18 years of age, reside in the United States, and consistently have demonstrated a high degree of success in competing tasks on MTurk, which was defined as having a 95% HIT completion rate (e.g., 5% dropout rate). Analysis of pre-test scores revealed that participants answered approximately a third of the pre-test questions correctly ($M = .35$, $SD = .18$) suggesting they had low prior knowledge of the concepts covered in the course. Course demographics showed 45% of participants were between 25-34 years of age; 39% were between 35 and 44 years of age. Eleven percent of participants had a master's degree, 35% had a bachelor's degree, 25% completed some college education, and 11% had a high school diploma.

Procedure

A brief description of the study was posted on the MTurk website. Individuals who were interested in completing the study accepted the task and advanced to an online informed consent that described the study's purpose, risks, benefits, and compensation requirements. After providing electronic consent, participants advanced to the training course which was hosted on a cloud-based instance of GIFT. Participants were greeted with a general welcome message and were then asked to read a brief set of instructions for completing the course. Then, participants completed a short demographic questionnaire that gathered information about their age, years of education, and familiarity with COIN topics and concepts. Next, participants completed a short goal orientation questionnaire that measured their task-based and intrinsic motivation to learn (Elliot & Murayama, 2008) followed by a 12-item pretest that assessed their prior knowledge of COIN principles and doctrine that aligned with content from the training course.

After completing the pre-training surveys, participants began the adaptive hypermedia portion of the course, which was organized into four chapters: (1) Introduction to COIN Operations; (2) Planning COIN Operations; (3) COIN Analysis Tools; and (4) COIN and Human Intelligence. Each chapter contained a series of narrated videos that covered lesson topics such as the importance of population support, processes for intelligence gathering, and issues in successful COIN operations. After watching each video, participants completed a series of recall questions that consisted of single or multi-concept review items that aligned with the course's learning objectives. Single concept review questions required learners to recall and apply concepts presented within the video lesson they had just viewed. Multi-concept review questions required learners to demonstrate a deeper understanding of the course's material by integrating concepts from multiple lessons. The course used a micro-sequencing adaptive training approach (Durlach & Spain, 2014) to "gate" learners' advancement through the course based on their demonstrated level of mastery. If participants missed a recall question, they received ICAP-inspired remediation that required them to: (1) passively re-read the narrated content that was just presented in the lesson video; (2) re-read the video content and actively highlight the portion of text that was most relevant to the quiz question that was missed; or (3) re-read the text and constructively summarize content related to the quiz question. The active and constructive remediation interventions prompted participants to self-evaluate the accuracy of their response compared to an expert response that was provided. The course also included a "no remediation" prompt wherein participants received a simple feedback message stating they incorrectly answered the question.

The course used a random policy to determine the type of remediation participants received after each missed recall question. Thus, participants were not assigned to a specific remediation condition (i.e., passive, active, constructive, or no remediation only) but instead, randomly received one of the four types of remediation each time they missed a recall question. After completing their encounter with a remediation intervention, participants were required to answer the recall question again. The ordering of the answer choices was different for each successive question attempt. Students continued to receive remediation until they demonstrated concept mastery (i.e., correctly answered the recall question).

In addition to the ICAP-inspired remediation prompts, the training course also provided feedback to learners who advanced through the videos too quickly or too slowly. For example, participants who tried to advance past a video before it ended received a feedback message stating that they should not rush through the training materials. Conversely, participants who spent too much time dwelling on the video (defined as more than 5 minutes on a video page) received feedback that they should complete the training materials at an efficient pace.

Upon finishing the final lesson, participants completed a series of post-training surveys in GIFT. These included a 12-item multiple-choice posttest to measure participants' recall of the concepts covered in the course and a short questionnaire to collect opinions about the training experience. After completing these activities, participants received a debriefing message and were thanked for their participation. Finally, participants followed MTurk procedures to receive credit for their time.

RESULTS

To investigate learners' behaviors with the instructional content as well as to investigate the effectiveness of the course for promoting learning gains, participants' responses to the surveys, recall questions, and interaction trace logs (e.g., lesson start time, lesson complete time, responses to questions, over-dwell and under-dwell violations, feedback type provided, remediation type provided) were analyzed. Data were extracted using GIFT's Event Reporting Tool, and variables of interest were compiled using Python and Excel. All statistical tests reported in this paper were performed with SPSS version 25.

Prior to performing inferential statistics, data were screened for outliers and violations of normality. A total of 50 cases were removed from the sample because of a system error in which pre and posttest scores were not recorded, resulting in a final sample of 483 participants. Course level statistics were also examined to determine if there were any differences in the course implementation from the course design parameters. The course was designed to implement an equal proportion of remediation types, but due to a software error results showed that 40% of remediation interventions presented during the course were constructive, 40% were active, 10% were passive, and 10% were no-remediation. Thus, there was a bias towards more cognitively engaging remediation interventions.

Learning Gains

To address the first research question of whether participants' knowledge of COIN concepts and doctrine improved after completing the training course, we computed a repeated measures ANOVA using pre and posttest scores as the dependent variable. Results showed that posttest scores ($M = 8.68$, $SD = 2.50$) were significantly higher than pretest scores ($M = 4.35$, $SD = 2.25$), $F(1, 482) = 1590.88$, $p < .001$, suggesting that the course was successful in meeting its instructional objectives. In addition to examining differences in pre and posttest scores, we also examined participants' normalized learning gains, which were calculated to account for participants' pretest performance. Normalized learning gains reflect an individual's relative improvement in a course by calculating the ratio of actual improvement from pre to posttest over the maximum possible improvement. Normalized learning gain values range from -1 to +1 with values below 0 indicating negative learning (i.e., students performed worse on the posttest than pretest), zero indicating no gains, and positive values indicating higher learning gains. Results for our study suggest that participants made significant learning gains from completing the course, improving their post test scores by more than approximately 57% of the total possible gains available ($M = .57$; $SD = .28$).

Factors Influencing Learning Gains

To address the second research question, we conducted a set of exploratory analyses to identify factors that impacted participants' learning gains. To facilitate these analyses, we created a new variable that grouped participants into quartiles according to normalized learning gains. Grouping the data in this manner allowed for the examination of differences in learning behaviors between individuals at different ends of the learning gain spectrum. The average normalized learning gains for the 1st, 2nd, 3rd, and 4th quartile groups were 18%, 53%, 71% and 89%, respectively.

Table 1 presents the results of a series of one-way ANOVAs that examined differences between the quartile groups across several dependent variables. Specifically, it shows how each group performed on the pre and posttest, the average number of remediation instances received by participants in each quartile, and the average duration that participants in each quartile spent interacting with each form of remediation. The first analysis examined whether there were significant differences between the groups in their level of prior knowledge of COIN concepts and doctrine. Results showed a small, but significant difference in pretest scores between the quartile groups. Specifically, individuals in the 4th quartile group scored higher on the pretest than individuals in the 1st quartile group. However, results also confirmed that there were significant differences between each quartile group on the posttest, with participants in the 4th quartile scoring significantly higher than the other groups on the posttest, answering an average of 11 out of 12 post test questions correctly.

Table 1. Means, Standard Deviations, and Grouping Effect on Learning Outcomes and Behaviors

Dependent Variables	Normalized Learning Gains Quartiles				Overall mean	ANOVA F
	1 st Quartile (n = 125)	2 nd Quartile (n = 126)	3 rd Quartile (n = 113)	4 th Quartile (n = 119)		
Pre-test Scores (maximum score = 12)	3.95 (2.33)	4.37 (2.09)	4.28 (2.30)	4.82 (2.08)	4.35 (2.22)	3.14*
Post-Test Scores (maximum score = 12)	5.56 (2.21)	8.40 (1.18)	9.77 (0.78)	11.22 (0.57)	8.68 (2.50)	384.45**
Number of Remediations Instances	19.07 (20.69)	9.31 (6.83)	6.57 (3.87)	5.01 (3.19)	10.23 (12.73)	36.46**
Constructive Remediation Interaction Time (secs)	75.42 (42.55)	92.31 (52.56)	107.40 (74.49)	114.37 (82.72)	96.29 (65.39)	7.785**
Active Remediation Interaction Time (secs)	49.95 (27.07)	54.81 (32.09)	59.39 (31.75)	57.30 (34.21)	55.08 (31.30)	1.87
Passive Remediation Interaction Time (secs)	28.07 (37.71)	31.37 (25.25)	31.45 (16.62)	32.63 (33.02)	30.44 (30.07)	0.26
Level of Effort (maximum score = 7)	5.86 (1.23)	6.47 (0.72)	6.35 (0.88)	6.39 (0.86)	6.26 (0.98)	10.63**

* $p < .05$; ** $p < .001$

Next, we examined the total number of remediation interventions received across quartile groups. The purpose of this analysis was to determine if participants who demonstrated higher learning gains also experienced more remediation interventions. Results showed that individuals in the 1st quartile group received, on average, 19 instances of remediation, which was significantly higher than all other groups (Table 1). Post-hoc analysis using Bonferroni corrections indicated that individuals with the highest learning gains (those in the 4th quartile) received significantly fewer instances of remediation than any other group. Thus, learning gains appeared to be inversely related to the number of remediation attempts received. Correlational analyses supported this assertion, showing a significant negative relationship between normalized learning gain scores and the total number of remediation interventions received ($r(469) = -.45, p < .001$).

The next set of analyses examined how long each group spent interacting with each form of instructional remediation. The goal was to determine if individuals in the lower learning gain group displayed different learning behaviors while completing the remediation activities compared to individuals in the higher learning gain quartile groups. Results revealed that individuals in the 4th quartile group spent significantly more time interacting with the constructive remediation content compared to individuals in the 1st quartile group. On average, participants in this higher performing group spent approximately two minutes completing each constructive remediation activity, whereas participants in the 1st quartile group spent an average of 1 minute and 15 seconds completing constructive remediation activities. There were no differences between the groups for the active or passive remediation interventions. Collectively, these results suggest that individuals who achieved higher learning gains may have done so because they spent more time and effort interacting with the course's remediation content. Indeed, an analysis of participants' self-reported levels of effort show that participants in the lowest quartile group reported the lowest level of effort. However, it is worth noting that participants in this group still reported putting modest effort into the activity, with average scores being above the measure's midpoint. An alternative explanation for the differences observed between the groups is that participants in the lower quartile spent less time, on average, interacting with the constructive content because they became fatigued with the intervention activity. A review of course frequency statistics shows participants in the 4th quartile group completed an average of two constructive remediation activities (range = 0 to 9 instances of constructive remediation), compared to participants in the 1st quartile group who completed, on average, seven constructive remediation activities on average (range = 0 to 39 instances of constructive remediation), which may introduce a confounding factor. Nevertheless, these results provide valuable insights about the tradeoffs between pedagogical interventions and learner behavior, particularly for those individuals that need the most support (i.e., individuals with lower learning gain).

Remediation Effectiveness

To address our third research question, we conducted a binary logistic regression to identify which form of remediation was most likely to result in a learner correctly answering a recall question following the first remediation attempt. For this analysis, the binary outcome of remediation effectiveness (Yes | No) served as the dependent variable, and remediation type served as the predictor. Three dummy variables were created to define membership for the passive, active, and constructive remediation categories; the no-remediation condition served as the reference category. As previously mentioned, learners continued to receive remediation interventions until they correctly answered the recall question, so this analysis aimed to identify which form of remediation was the most successful, on average, at helping students overcome an impasse on the first attempt. According to the ICAP model, constructive should be the most effective form of remediation.

Results largely supported this prediction; the overall regression model using remediation type as the predictor was significant ($\chi^2 = 109.73, p < .000$ with $df = 3$). An examination of odds ratios showed constructive remediation was the strongest predictor of remediation effectiveness. Specifically, learners were 3.58 times more likely to correctly answer a recall question after receiving constructive remediation than if they received no remediation (Table 2). The second strongest predictor of remediation effectiveness was active remediation with results showing that learners were 2.99 times more likely to answer a recall question correctly after receiving this form of remediation than if they received no remediation. Learners who received passive remediation were 1.72 times more likely to advance to the next question compared to if they received no remediation. These results align with guidance provided by the ICAP model and provide additional evidence of the benefits of engaging learners in active learning interventions such as self-explanation and highlighting, which classify as constructive and active remediation activities, respectively.

Table 2. Logistic Regression Predicting Remediation Effectiveness

Predictor	Remediation Effectiveness (<i>n</i> = 483)		
	B	WALD	EXP(B)
Constructive Remediation	1.28	99.46**	3.58
Active Remediation	1.10	76.39**	2.99
Passive Remediation	0.54	10.81**	1.72
No Remediation	0.41	14.40**	1.50
$\chi^2 = 109.73, p < .000$ with <i>df</i> = 3			

** *p* < .001

DISCUSSION AND DIRECTIONS FOR FUTURE RESEARCH

Results of the current effort highlight several important findings as they relate to the design and development of AISs. First, the results demonstrate that the adaptive course developed and administered through GIFT was effective at improving participants' knowledge of COIN concepts. Learning gains were particularly strong for learners who spent more time interacting with constructive remediation activities. These results support the ICAP model of cognitive engagement, and they suggest that learners can benefit from investing effort in completing constructive remediation interventions.

Second, results showed that despite receiving more instructional remediation interventions, learners in the lower learning gains quartile group did not perform as well on the posttest compared to the other quartile groups. Ideally, instructional remediation would be as effective, or even more effective, for low-scoring students because they have the greatest need for support. As noted above, the pedagogical policy implemented in the course randomly chose whether participants received constructive, active, passive, or no remediation. Our results suggest that although a random policy may promote learning gains when course effectiveness is examined at a global level, it may not be effective at supporting learners who are in the most need of scaffolding and instructional support (lower scoring or lower motivated learners). Further, it highlights the potential for using data-driven techniques, such as reinforcement learning, to develop pedagogical policies that provide tailored forms of remediation to learners. Because motivation also plays a significant role in student learning (D'Mello, 2013), future research should explore how motivational and cognitive interventions could be embedded in tutorial planners to provide learners with tailored scaffolding and affective support as they advance toward completing training.

Third, the hypermedia course required learners to demonstrate course mastery before they could advance to the next recall question or lesson. Results show that learners who immediately received constructive remediation were three times more likely to demonstrate concept mastery on subsequent recall questions compared to learners who received no remediation after a missed recall question. These results are encouraging, but only investigate remediation effectiveness at the course level. Future research should investigate the effectiveness of different remediation prompts and policies at a more granular level to examine trade-offs between cognitive engagement and instructional effectiveness, particularly as learners spend more time in a course. The research team is currently investigating these questions with the dataset by applying reinforcement learning to induce data-driven remediation policies, which will enable exploration of whether remediation policies are impacted by the richness of a model's state data about the course and the student (e.g., how many times has the learner received remediation, what type of remediation did the learner most recently receive, what was the learner's level of prior knowledge for the topic?).

One item of note regarding the course implementation is that participants were not assigned to receive a specific type of remediation, which meant that we did not have a traditional baseline or control group to make relative comparisons between the different remediation types. Rather, remediation interventions were randomly administered to learners on an item-by-item basis. The purpose of using this design was to facilitate collecting a large corpus of interaction data for training data-driven tutorial planning policies using reinforcement learning. To devise data-driven models we needed to collect data from learners who interacted with different types and sequences of remediation. While this approach is necessary for reinforcement learning-based tutorial policies that can deliver ICAP inspired remediation in GIFT, from an experimental design standpoint, it means that the learning gains that were observed in our analysis

could be attributed to extraneous factors. We are currently addressing this limitation by collecting additional data that includes a traditional control group wherein participants do not receive any form of remediation.

From a practical standpoint, the results of the current study indicate that there are trade-offs to consider when devising rules for delivering remediation in adaptive hypermedia-based training environments. Providing learners with active learning activities is an effective remediation strategy for promoting learning gains. However, this can be a double-edged sword. Our results showed constructive remediation was the most effective form of remediation for helping students overcome an impasse, but it also required the highest level of cognitive engagement, which requires additional time and could gradually contribute toward accelerated fatigue. One question that course designers might have is whether passive and active remediation can be substituted for constructive remediation without impacting course effectiveness. Using a data-driven tutorial planning approach to determine the right sequence and combination of remediation interventions offers a promising approach for addressing this question.

CONCLUSION

Tutorial planners are critical to shaping student learning outcomes during training with AISs because they control how and when pedagogical interventions are delivered to students. Recent advances in machine learning, and reinforcement learning in particular, show significant promise for enabling data-driven methodologies for the creation of effective tutorial planners. To collect data for investigating reinforcement learning-based tutorial planning, we conducted a study on Amazon MTurk in which learners engaged with an online adaptive hypermedia-based training course on COIN doctrine that was implemented with GIFT. During the course, students received adaptive remediation based on the ICAP model of cognitive engagement. Results showed that students achieved significant learning gains from completing the training course, and furthermore, the results supported the ICAP model: student interactions with constructive remediation activities were most effective at improving student performance on embedded recall questions, followed by active remediation activities, and then passive remediation activities. However, low-scoring students, despite receiving a greater quantity of remediation, spent less time engaging with constructive remediation, and their learning gains were substantially reduced compared to their high-scoring counterparts. Notably, these results stem from remediation that was delivered according to a random policy; no adaptive personalization was implemented. These findings underscore the important role of cognitive engagement in training with AISs and the promise of tutorial planners that can effectively tailor pedagogical support and remediation to students based upon their individual needs.

ACKNOWLEDGEMENTS

The authors wish to thank Mike Hoffman of Dignitas Technologies for providing extensive assistance in developing GIFT software features that enabled this research. We would also like to thank Yonatan Vaknin of Intelligent Automation, Inc. for his assistance with data extraction and preparation. The research described herein has been sponsored by the U.S. Army Research Laboratory under cooperative agreement W911NF-15-2-0030. The statements and opinions expressed in this article do not necessarily reflect the position or the policy of the United States Government, and no official endorsement should be inferred.

REFERENCES

- Chi, M.T.H. (2009). Active-Constructive-Interactive: A conceptual framework of differentiating learning activities. *Topics in Cognitive Science*, 1(1), 73-105.
- Chi, M. T., & Wylie, R. (2014). The ICAP framework: Linking cognitive engagement to active learning outcomes. *Educational psychologist*, 49(4), 219-243.
- D'Mello, S. (2013). A selective meta-analysis on the relative incidence of discrete affective states during learning with technology. *Journal of Educational Psychology*, 105(4), 1082.
- Durlach, P. J. (2012). A road ahead for adaptive training technology. In P.J. Durlach & A.M. Lesgold (Eds). *Adaptive Technologies for Training and Education*, (pp 331-341). New York: NY. Cambridge University Press.
- Durlach, P. J., & Ray, J. M. (2011). *Designing adaptive instructional environments: Insights from empirical evidence*. (Technical Report No. 1297) U.S. Army Research Institute for the Behavioral and Social Sciences; Arlington, VA.

- Durlach, P., & Spain, R. (2014). *Framework for Instructional Technology: Methods of Implementing Adaptive Training and Education* (Technical Report No. 1335). U.S. Army Research Institute for the Behavioral and Social Sciences. Fort Belvoir, VA.
- Elliot, A. J., & Murayama, K. (2008). On the measurement of achievement goals: Critique, illustration, and application. *Journal of Educational Psychology*, 100(3), 613.
- McAlinden, R., Gordon, A., Lane, H. C., & Pynadath, D. (2009). UrbanSim: A game-based simulation for counterinsurgency and stability-focused operations. In *Workshop on Intelligent Educational Games at the 14th International Conference on Artificial Intelligence in Education* (pp. 41-50).
- Rowe, J., & Lester, J. (2015). Improving Student Problem Solving in Narrative Centered Learning Environments: A Modular Reinforcement Learning Framework. In *Proceedings of the Seventeenth International Conference on Artificial Intelligence in Education*, pp. 419-428.
- Rowe, J., Spain, R., Pokorny, B., Mott, B., Goldberg, B., & Lester, J. (2018). Design and development of an adaptive hypermedia course for counterinsurgency training in GIFT: Opportunities and lessons learned. *Proceedings of the Sixth Annual GIFT User Symposium (GIFTSym6)*, pp. 229-239.
- Sottolare, R., Barr, A., Robson, R., Hu, X., & Graesser, A. (2018). Exploring the opportunities and benefits of standards for adaptive instructional systems (AISs). In *Proceedings of the Adaptive Instructional Systems Workshop in the Industry Track of the 14th International Intelligent Tutoring Systems (ITS) Conference*, pp. 49-53.
- Sottolare, R. A., Brawner, K. W., Goldberg, B. S., & Holden, H. K. (2012). The generalized intelligent framework for tutoring (GIFT). *U.S. Army Research Laboratory–Human Research & Engineering Directorate (ARL-HRED)*.
- Sottolare, R. A., Brawner, K. W., Sinatra, A. M., & Johnston, J. H. (2017). An updated concept for a generalized intelligent framework for tutoring (GIFT). Retrieved from www.gifttutoring.org
- Sutton, R. S., & Barto, A. G. (1998). Reinforcement learning: an introduction MIT Press. *Cambridge, MA*.
- Williams, J. J., Kim, J., Rafferty, A., Maldonado, S., Gajos, K., Lasecki, W., & Heffernan, N. (2016). AXIS: Generating Explanations at Scale with Learner sourcing and Machine Learning. *Proceedings of the 3rd Annual ACM Conference on Learning at Scale*, Edinburgh, SCT, pp. 379-388.
- TRADOC (2017). The U.S. Army Learning Concept for training and education 2020-2040. Retrieved from <https://adminpubs.tradoc.army.mil/pamphlets/TP525-8-2.pdf>
- VanLehn, K. (2011). The relative effectiveness of human tutoring, intelligent tutoring systems, and other tutoring systems. *Educational Psychologist*, 46(4), 197-221.
- Zhou, G., Wang, J., Lynch, C., & Chi, M. (2017). Towards closing the loop: Bridging machine-induced pedagogical policies to learning theories. In *Proceedings of the Tenth International Conference on Educational Data Mining* (pp. 112-119).