Design and Development of an Adaptive Hypermedia-Based Course for Counterinsurgency Training in GIFT: Opportunities and Lessons Learned

Jonathan Rowe[1], Randall Spain[1], Robert Pokorny[2], Bradford Mott[1], Benjamin Goldberg[3], and James Lester[1]
North Carolina State University[1], Intelligent Automation, Inc.[2], U.S. Army Research Laboratory[3]

INTRODUCTION

There is broad recognition that intelligent tutoring systems (ITSs) are effective for enhancing student learning across a range of domains (Anderson, Corbett, Koedinger, & Pelletier, 1995; VanLehn, 2011; Ma, Adescope, Nesbit, & Liu, 2014). By leveraging computational models of adaptive pedagogical decision-making, ITSs create personalized learning experiences that are dynamically tailored to individual students. However, ITSs are resource-intensive to create. The amount of engineering effort required to develop one hour of ITS instruction is often estimated to be approximately 200 hours (Aleven, McLaren, Sewall, & Koedinger, 2009). To address this bottleneck, there have been several initiatives to devise tools for supporting ITS authors in creating adaptive training at reduced time and cost (Aleven et al., 2009; Sottilare, Baker, Graesser, & Lester, in press). These efforts hold the promise of making ITSs available across a broader range of subjects and contexts, enhancing the depth of current adaptive learning experiences, and enabling instructional designers and subject matter experts to create novel adaptive training solutions without requiring programming expertise.

Over the past several years, the Generalized Intelligent Framework for Tutoring (GIFT) has emerged as an important initiative to address the authoring challenges raised by ITSs. GIFT is an open source service-oriented framework of software tools, methods, and de-facto best practices for designing, developing, 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 (Rowe et al., 2016). 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. The Course Creator is also continuously improving with new capabilities and usability enhancements released several times a year. As GIFT transitions from the research lab to real-world use, these tools will be subject to new authorial demands and scalability challenges, which makes it a valuable test case for understanding how ITS technologies mature and scale.

In this paper, we describe our experiences and lessons learned from using GIFT to create an approximately 2-hour adaptive hypermedia-based training course for counterinsurgency (COIN) and stability operations. The course builds upon the UrbanSim Primer, which presents a range of multimedia training materials providing direct instruction on doctrinal concepts of COIN that accompanies the UrbanSim simulation-based training environment. The course serves as a showcase of recent enhancements to GIFT's Engine for Management of Adaptive Pedagogy (EMAP) that support adaptive assessment and remediation. Specifically, remediation features in GIFT are based on Chi's ICAP framework (2009). ICAP describes several modes of student engagement with learning materials, including passive, active, constructive, and interactive modes. The ICAP framework predicts that the interactive mode (e.g., peer dialogue) is more effective for learning than the constructive mode (e.g., writing an explanation), the constructive mode is more effective than the active mode (e.g., reading and highlighting a passage), and the active mode is more effective than the passive mode (e.g., reading a passage without doing anything else). But, there are tradeoffs between these pedagogical strategies, such as instructional time required and cognitive load imposed. We are utilizing the COIN hypermedia-based training course to gather data on student responses to passive, active, and constructive remediation activities, which is part of a broader research program on utilizing reinforcement learning to automatically induce intelligent tutoring policies for instructional remediation in GIFT.

Our course is notable in its scope: it utilizes nearly 40 adaptive course flow objects, more than 150 media objects, online videos, pre-post surveys, embedded assessments, adaptive feedback messages, glossaries, and other features of GIFT. Further, we are preparing the course for deployment to hundreds of users through a crowdsourcing study with Amazon Mechanical Turk, which requires preparation for remote deployment to dozens of concurrent users in a fashion that integrates seamlessly with tools and workflows from commercial crowdsourcing providers. We describe how this course was created with the GIFT Course Creator and associated ICAP-inspired functionalities; we describe methods for implementing key ITS features such as immediate feedback and scaffolding in hypermedia-based training with GIFT; and we describe challenges, solutions, and opportunities we have encountered from our experiences creating the course. Our findings point toward future directions for enhancing GIFT's capacity to reduce the authorial cost of creating ITSs and transitioning toward wider scale use.

RESEARCH CONTEXT

Tutorial planning, a critical component of adaptive training, controls how scaffolding and instructional interventions are structured and delivered to learners. Devising computational models that scaffold effectively, i.e., determining when to scaffold, what type of scaffolding to deliver, and how scaffolding should be realized, is a critical challenge for the field of ITSs. Recent years have seen growing interest in data-driven approaches to tutorial planning (Rowe & Lester, 2015; Williams et al., 2016; Zhou, Wang, Lynch, & Chi, 2017). In particular, reinforcement learning techniques have shown promise for automatically inducing tutorial policies that optimize student learning outcomes and do not require pedagogical policies to be manually programmed or demonstrated by expert tutors. These techniques are complementary to advances in ITS authoring, including authoring tools implemented in GIFT, to address challenges inherent in constructing adaptive training materials.

Reinforcement learning is a category of machine learning that centers on devising software agents that perform actions in a stochastic environment to optimize some concept of numerical reward (Sutton & Barto, 1998). In reinforcement learning, the agent induces a control policy by iteratively performing actions and observing their effects on the environment and accumulated rewards. Tutorial planning can be formalized as a reinforcement learning task by conceptualizing the tutor as the agent: the tutor seeks to enact pedagogical decisions (i.e., actions) that will affect its environment (i.e., the trainee and his/her learning environment) in order to optimize student learning outcomes (i.e., rewards). In our case, the pedagogical decisions are choosing between ICAP-inspired remediation activities, and the tutorial planner’s objective is to optimize student learning in an adaptive hypermedia-based training course for COIN.

To investigate a reinforcement learning framework for ICAP-inspired remediation in GIFT, we plan to obtain a large dataset consisting of trainee responses to different types of instructional remediation activities as well as pre-post learning outcomes. The purpose of the dataset is to serve as a training corpus for inducing and evaluating reinforcement learning policies for tutorial planning (Rowe & Lester; Wang et al., 2017). Reinforcement learning techniques are data-intensive, so in order to collect sufficient data, we have devised a training course that is designed to meet three objectives: (1) the course contains numerous opportunities for learners to receive instructional remediation, which will serve as the training data for reinforcement learning; (2) the course is deployable through online crowdsourcing platforms, which will facilitate broad distribution to many learners efficiently; and (3) the course enacts an exploratory (i.e., random) remediation policy in order to broadly sample the space of possible pedagogical decisions. To meet these objectives, we developed an adaptive hypermedia-based training course in GIFT that builds upon materials from the UrbanSim Primer.

UrbanSim Primer

The UrbanSim Primer is a hypermedia-based learning environment that was developed by the USC Institute for Creative Technologies to provide direct instruction on COIN doctrine and principles. Major topics include the importance of population support, processes for intelligence gathering, and issues in successful COIN operations. The UrbanSim Primer's training materials are divided across seven lessons that interleave hyperlinked video, audio, text, and diagrams delivered using Adobe Flash.

In our project, we focus on a subset of training materials from UrbanSim Primer Lessons 1-4. We have extracted video, audio, and text content from the UrbanSim Primer, and we have reconfigured these materials for web-based presentation using GIFT. Specifically, GIFT enables the delivery of UrbanSim Primer materials via web browsers, it enables interleaved training materials that include embedded assessments and instructional remediation, and it supports automatic logging of learner actions within the training course.

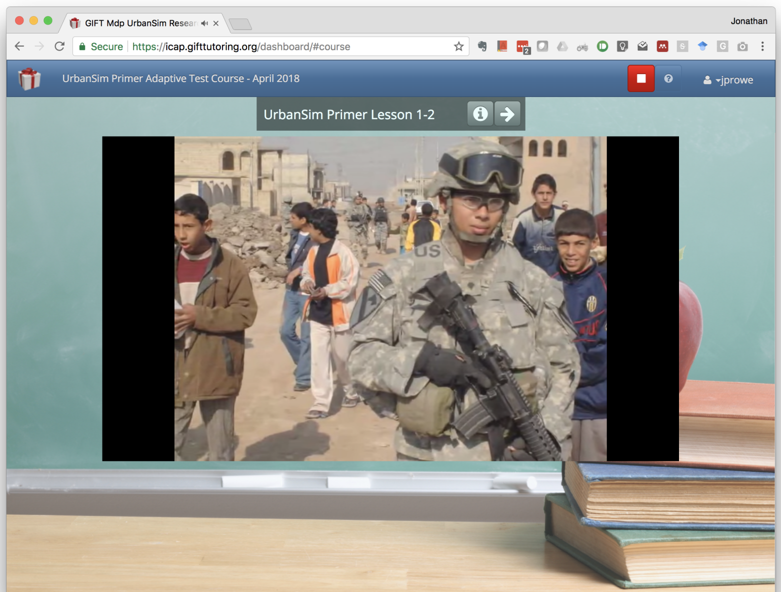

Figure 1: Screenshot of UrbanSim Primer training video presented in GIFT.

An Implemented Adaptive Hypermedia-Based Training Course for COIN

We have designed an adaptive hypermedia-based training course based on the UrbanSim Primer using a branch of GIFT Cloud that supports recent enhancements to the GIFT EMAP to support ICAP-based instructional remediation functionalities. The course builds upon the doctrinal lessons presented in the UrbanSim Primer and includes a series of short videos, instructional texts, quiz questions, remedial content, and glossaries related to the fundamental principles of COIN and stability operations. Trainee experiences with the COIN training course proceed as follows.

The course begins with a general message that welcomes students to the training course. Following this introduction, participants complete a demographic questionnaire that asks them about their age, years of education, and familiarity with COIN topics and concepts, followed by a goal orientation questionnaire that measures students' task-based and intrinsic motivation to learn (Elliot & Murayama, 2008). Next, students complete a 12-item pretest that measures prior knowledge of COIN principles and doctrine.

After completing the pre-training surveys, participants begin the adaptive hypermedia portion of the course, which is 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 contains a series of narrated videos and text-based content that cover lesson 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.” Each lesson is implemented as a series of adaptive course flow objects within the GIFT course.

After each video from the UrbanSim Primer, participants complete a brief multiple-choice quiz. Quiz questions consist of single concept and multi-concept review items that align with the course’s learning objectives. Single concept review questions require learners to recall and apply concepts presented within the lesson. Multi-concept review questions require learners to demonstrate a deeper understanding of course material by integrating concepts from multiple lessons. The course uses a micro-sequencing adaptive training approach (Durlach & Spain, 2014) to “gate” progress according to learners’ demonstrated level of mastery. Learners who correctly answer a quiz question are allowed to advance to the next question or lesson, whereas learners who incorrectly answer a question receive ICAP-inspired supplemental remediation.

When a learner receives supplemental remediation following a missed question, GIFT prompts the learner to either: (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 is 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 remediation prompts also include expert highlighting/summaries that students can use to self-evaluate the accuracy of their responses. The course also includes a “no remediation” prompt that only provides students with minimal feedback before being asked to re-answer the quiz question. The course uses a random assignment policy at the item level to determine whether students receive passive, active, constructive, or no remediation after each incorrect item response. Students continue to receive supplemental remediation until they demonstrate concept mastery (i.e., correctly answer the quiz question).

In addition to the ICAP-inspired remediation prompts, the training course also monitors how long students engage with the video-based lessons and provides prompts to those participants who advance through the videos too quickly or too slowly. For example, participants who click past a video before it ends receive the following message, “It appears that you clicked past the video before enough time elapsed for it to play in entirety. Please do not rush through the training materials, or else you may not achieve the course learning objectives.” Conversely, participants who spend too much time dwelling on the video (defined as more than 5 minutes on a video page) will receive the following message, “It appears that you spent an unusually long amount of time on this video. Please attempt to complete the training materials at an efficient pace.” The maximum video length is approximately 1.5 minutes.

Upon finishing the final lesson, participants complete a series of post-training surveys. These include a multiple-choice posttest to measure retention of foundational COIN concepts and a short questionnaire to collect opinions about the training experience. After completing these activities, participants receive a debriefing message and are thanked for their participation. In addition, participants who access the course through an online crowdsourcing platform (e.g., Mechanical Turk) receive a unique code that they can provide to the crowdsourcing vendor to receive payment for participation. The code is randomly generated by a customized survey implemented as the final course object in the training course.

In order to collect data on learner interactions with ICAP-inspired remediation activities, we plan to conduct a human subjects study with a sample of 300-500 participants recruited through Amazon Mechanical Turk. A short description of the study will be posted on the Mechanical Turk website. Individuals interested in completing the training course will first complete an electronic informed consent before being hyperlinked to the published training course hosted on the cloud-based version of GIFT. Once in the course, participants will proceed through the course activities described above. At the end of the training course, participants will receive a unique 7-digit code that they must enter into the Mechanical Turk site to receive payment for their participation. Using the data gathered from the Mechanical Turk study, we will begin to investigate data-driven models of tutorial planning using reinforcement learning techniques.

DESIGNING AN ADAPTIVE HYPERMEDIA-BASED TRAINING COURSE IN GIFT

To develop an adaptive hypermedia-based training course that can be delivered through the web, we made extensive use of the GIFT Course Creator. The GIFT Course Creator is a web-based GUI authoring tool that enables instructional developers to construct training workflows that encode sequences of online learning activities using a drag-and-drop interface. The Course Creator enables instructional developers to specify *fixed course flows*, which are course-object sequences that unfold the same way for every learner, as well as *adaptive course flows*, which utilize the GIFT EMAP to drive macro-adaptive pedagogical decisions about content sequencing based on student performance. A key component of our work on the adaptive hypermedia-based COIN training course is utilization of an enhanced version of EMAP that supports ICAP-inspired remediation functionalities. Specifically, the course includes 39 adaptive course flow objects, each linked to a range of supporting media files including videos, text passages, feedback statements, ICAP-inspired remediation prompts, and quiz questions that align with course concepts. In this section, we briefly describe how these adaptive courseflow objects are configured to provide direct instruction, embedded assessment, immediate feedback, and adaptive remediation on COIN concepts.

In GIFT, adaptive courseflow objects are deeply grounded in Component Display Theory (Merrill, Reiser, Ranney & Trafton, 1992). Component Display Theory describes a process for learning the rules of a domain, examining relevant examples, testing recall of knowledge, and engaging in guided practice. These four types of learning activities delineate quadrants in an adaptive courseflow object: Rules, Examples, Recall, and Practice. During a typical interaction with an adaptive courseflow object, the learner begins by viewing multimedia training materials associated with a set of target domain concepts; this is the learner's experience of the Rules Quadrant. After viewing these materials, the learner transitions to the Examples Quadrant in which she views additional training materials that illustrate examples of the target concepts. Afterward, the learner transitions to the Recall Quadrant, where her understanding of the target concepts is assessed through a series of quiz questions. After successfully completing the quiz, the learner optionally transitions to the Practice Quadrant, where she interacts with an external training simulation to apply her relevant knowledge in a hands-on manner. In our course, we do not currently make use of the Practice Quadrant.

In the ICAP-enhanced version of the GIFT EMAP, the four quadrants are augmented with an additional fifth quadrant: Remediation. The Remediation Quadrant houses logic and training materials for presenting instructional feedback and ICAP-inspired remediation to learners with below-threshold performance in the Recall Quadrant. In other words, when learners miss too many embedded quiz questions, they receive immediate feedback and remediation. The Remediation Quadrant is populated with multimedia training materials that are distinct from those presented in the Rules and Example Quadrants. Remediation materials can be conceptualized in terms of three categories: (1) Constructive-response remediation, (2) Active-response remediation, and (3) Passive-response remediation.[1] Constructive- and active-response remediation materials are created using built-in GIFT authoring templates, whereas passive remediation materials can be constructed with a range on supported media types, including videos, text passages, web pages, and slide decks. In the case of our course, all remediation materials are text based, and we specifically utilize text-based local web pages to implement Passive-response remediation. At present, the presentation of these three different types of remediation is performed according to a uniform random policy. This

[1] The fourth category of ICAP, interactive-response remediation, is not currently supported by the GIFT EMAP.

control policy will be replaced with an adaptive policy induced using reinforcement learning following the completion of the Mechanical Turk study, subsequent data analysis and model creation.

In our course, we utilize adaptive courseflow objects to provide immediate feedback and remediation after each quiz question. We devise a unique adaptive courseflow object for each embedded quiz item in the course. Each adaptive courseflow object contains a Recall Quadrant with a single question, as well as a Remediation Quadrant with four associated media files: a passive-response remediation intervention, an active-response remediation intervention, a constructive-response remediation intervention, and a non-remediation intervention. Each of these remediation media files contains a feedback statement about the quiz question that the learner must have missed prior to receiving the remediation. Because our course presents multiple embedded quiz questions after each video from the UrbanSim Primer, a subset of adaptive courseflow objects contain links to YouTube videos in their Example Quadrants. However, not all adaptive courseflow objects contain these videos, or contain Example Quadrant media files at all. In these cases, we repurpose the Example Quadrant to present a local webpage containing positive feedback about the quiz question that the student just answered; when a learner transitions to a new adaptive courseflow object, she must have just answered a quiz question correctly, so we provide positive feedback. The Rule Quadrant of each adaptive courseflow object is generally not used in our course, except to present transition text in a handful of locations. For an illustration of all of the materials and media files associated with each adaptive courseflow object in our course, please see Figure 2.

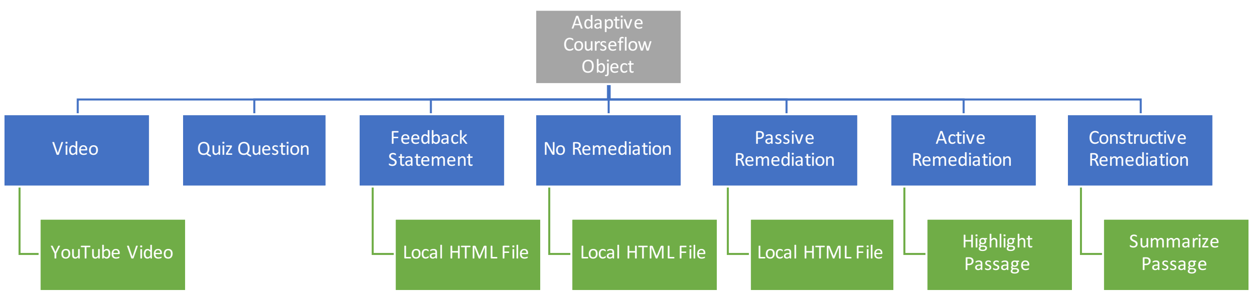

Figure 2. Overview of training materials associated with a single adaptive courseflow object.

BEST PRACTICES AND LESSONS LEARNED

As noted above, the course is relatively large in comparison to many GIFT courses that have been created to date. The course includes approximately 40 adaptive course flow objects, more than 150 media objects, several pre-post surveys, numerous embedded assessments, adaptive feedback messages, glossaries, and other features of GIFT.

As we developed the course, we found that designing a large training course in GIFT required significant preparation and planning. A key practice that we utilized was to develop a course prototype outside of GIFT prior to constructing the training course inside of GIFT. In our case, we developed a rough course prototype in PowerPoint, which served two functions: (1) It provided the team with an easily editable instructional design map of the training course, including an overview of the course flow for each chapter and subsequent lessons, and (2) It allowed the team to quickly edit and refine the training content, embedded assessments, and remediation content before implementing the full user-ready version in GIFT. We also found that the prototype served as a useful reference for authoring remediation prompts. For each lesson, we created a series of slides that showed the quiz question that aligned with the lesson, transcripts of the narrated text from the video, text for the passive remediation prompts, text and suggested highlighting for the active remediation prompts, and content for the constructive remediation prompts. Organizing all of this information in a format that could be rapidly generated, easily edited, shared between collaborators, and which did not require perfect precision in specifying courseflows played a vital role in the early stages of authoring the adaptive training course.

A second lesson we learned was that during course development and revision, there were several occasions in which we needed to revise course content (e.g., quiz questions and prompts) to improve the clarity of the training materials. These changes were based on upon user feedback from pilot testing of the course. We found it helpful to keep track of these revisions in the PowerPoint prototype of the course, which allowed the project team to easily track changes made to the course over the development cycles.

A third lesson we learned was that it is important to implement a naming and organization convention for the media files used in the course. As a best practice, we used an object + lesson naming scheme (e.g., Remediation 2-3 Constructive; Remediation 2-3 Active, etc.) to provide structure and consistency among all of our training assets. This organizational scheme was particularly useful in managing the feedback statements, passive remediation files, and no-remediation files associated with each adaptive courseflow object. This allowed us to quickly review which objects were included in each course object's Remediation Quadrant. It also helped us manage the large number of training assets saved in the course's media folder. As previously noted, our course includes over 150 content files. During the authoring process, there were many occasions in which we needed to either preview or edit passive and/or non-remediation files associated with the course. In the current implementation of GIFT, the only way to preview and update these files is by accessing the file through the media content organizer, which lists all of a courses' media objects (see Figure 3). Using a pre-established naming convention allowed us to quickly locate and replace old course objects when we needed to make changes to the training course, which occurred several times during the iterative course development and refinement process.

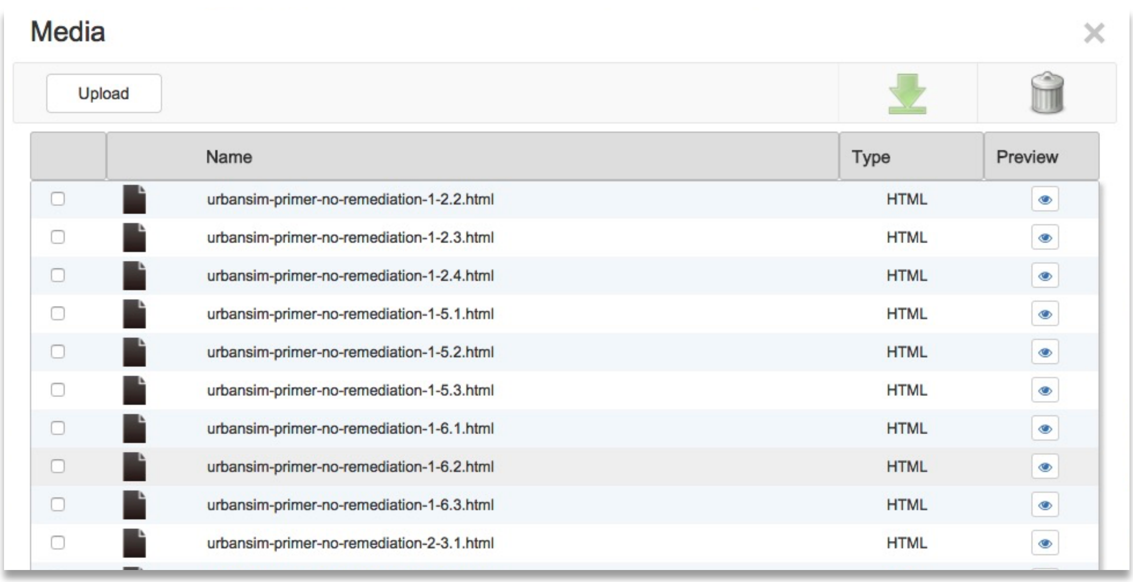

Figure 3. Training Assets in the media folder of GIFT's Course Creator.

A fourth lesson is the importance of developing a large hypermedia course such as this one in an iterative fashion. As a best practice, we developed the course one chapter at a time and conducted internal pilot testing between development cycles to ensure the course workflow and remediation materials were being implemented properly. During our pilot testing sessions, we examined extreme, correct, and incorrect responses to the quiz questions to ensure the course logic was correct, and we examined whether the remediation prompts were being executed correctly in order to tune course parameters and functionality. In addition to reviewing the behavior of these system level features, we also used testing as an opportunity to make any changes to the visual design of the course, such as making changes to font sizes and line spacing in our remediation prompts and messages prior to developing the rest of the course's media objects; a change in the presentation style of one feedback message could potentially propagate to more than a hundred additional files if an author is not careful about phased development.

CONCLUSIONS AND RECOMMENDATIONS FOR FUTURE RESEARCH

Adaptive training systems show considerable promise for enhancing student learning across a range of domains. Recent advances in ITS authoring tools, as well as data-driven tutorial planning, are showing significant progress toward reducing the effort required to create personalized learning experiences. Reinforcement learning is a natural formalism for automatically inducing tutorial planning models to drive pedagogical decisions about instructional feedback and remediation. In order to utilize reinforcement learning techniques for data-driven tutorial planning, we have constructed an adaptive hypermedia-based training course in GIFT that is based on the UrbanSim Primer to teach foundational principles and doctrine on COIN operations. We utilize ICAP-inspired enhancements to GIFT's EMAP to provide immediate feedback and remediation during the adaptive training course. Based on our experience creating the course, we have identified several best practices and lessons learned for adaptive course creation in GIFT. These include the importance of external prototyping, carefully tracking course revisions, devising consistent file-naming schemes, and emphasizing iterative design and development throughout course creation.

As a next step, we will deploy the adaptive training course in a human subject study using the Amazon Mechanical Turk crowdsourcing platform in order to collect a training corpus for investigating reinforcement learning-based tutorial planning. Following the study, we will utilize the dataset to induce control policies for adaptively personalizing remediation and feedback decisions to individual learners. In the future, we plan for these models to be incorporated back into the run-time adaptive training course and evaluated with a new cohort of learners in order to evaluate the effectiveness of reinforcement learning techniques for data-driven tutorial planning in GIFT.

There are several promising avenues for future enhancements to GIFT. One recommendation is to include advanced previewing capabilities within the GIFT Course Creator. In particular, adding features that allow authors to preview adaptive course flow objects, and in particular, Remediation Quadrant materials, would be highly valuable. Currently, course authors can access and edit the content of the constructive and active remediation prompts, but they cannot preview how these prompts appear at run-time when they are presented by GIFT. The same previewing functionality would be useful for passive remediation content as well, such as local web pages, particularly if they could be previewed directly from adaptive courseflow objects in the Course Creator.

Enhancements related to viewing and managing large numbers of media files would also be helpful to course creators. Including a feature that allows course authors to quickly view all of the media file labels attached to an adaptive courseflow object would significantly facilitate authoring for large courses. Currently, authors have to open each adaptive courseflow object and individually click on each quadrant to see which media files are linked to each quadrant. Including a feature that could quickly export or summarize this information at a high level would eliminate this process and would be a valuable tool for evaluating and refining the training course.

ACKNOWLEDGMENTS

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

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

Aleven, V., Mclaren, B.M., Sewall, J., & Koedinger, K.R. (2009). A new paradigm for intelligent tutoring systems: Example-tracing tutors. *International Journal of Artificial Intelligence in Education*, 19(2), 105-154.

Chi, M.T.H. (2009). Active-Constructive-Interactive: A conceptual framework of differentiating learning activities. *Topics in Cognitive Science*, 1(1), 73-105.

Durlach, P. J., & Spain R. D. (2013). *Framework for instructional technology: Methods of implementing adaptive training and education*. (Technical Report 1335). Arlington VA: U.S. Army Research Institute for the Behavioral and Social Sciences.

Elliot, A. J., & Murayama, K. (2008). On the measurement of achievement goals: Critique, illustration, and application. *Journal of Educational Psychology*, *100*(3), 613.

Ma, W., Adesope, O.O., Nesbit, J.C., & Liu, Q. (2014). Intelligent tutoring systems and learning outcomes: A meta-analysis. *Journal of Educational Psychology*, 106(4), 901.

Merrill, D., Reiser, B., Ranney, M., & Trafton, J. (1992). Effective Tutoring Techniques: A Comparison of Human Tutors and Intelligent Tutoring Systems. *Journal of the Learning Sciences*, 2(3), 277-305.

Rowe, J., Frankosky, M., Mott, B., Lester, J., Pokorny, B., Peng, W., & Goldberg, B. (2016). Extending GIFT with a Reinforcement Learning-Based Framework for Generalized Tutorial Planning. *Proceedings of the Fourth Annual GIFT User Symposium* (GIFTSym4), pp. 87-97.

Rowe, J. P., & Lester, J. C. (2015). Improving Student Problem Solving in Narrative-Centered Learning Environments: A Modular Reinforcement Learning Framework. *Proceedings of the 17th International Conference on Artificial Intelligence in Education* (pp. 419-428), Madrid, Spain.

Sottilare, R.A., Baker, R., Graesser, A.C. & Lester, J. (In press). Special Issue on the Generalized Intelligent Framework for Tutoring (GIFT): Creating a Stable and Flexible Platform for Innovations in AIED Research. *International Journal of Artificial Intelligence in Education.*

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

Williams, J. J., Kim, J., Rafferty, A., Maldonado, S., Gajos, K. Z., Lasecki, W. S., & Heffernan, N. (2016). Axis: Generating Explanations at Scale with Learnersourcing and Machine Learning. *Proceedings of the 3rd ACM Conference on Learning@ Scale* (pp. 379-388), Edinburgh, UK.

Zhou, G., Wang, J., Lynch, C. & Chi, M. (2017) Towards Closing the Loop: Bridging Machine-induced Pedagogical Policies to Learning Theories. *Proceedings of the 10th International Conference on Educational Data Mining*, (pp.112-119), Wuhan, China.

ABOUT THE AUTHORS

Dr. Jonathan Rowe *is a Research Scientist in the Center for Educational Informatics at North Carolina State University. He received the Ph.D. and M.S. degrees in Computer Science from North Carolina State University, and* *the B.S. degree in Computer Science from Lafayette College. His research is in the areas of artificial intelligence and human-computer interaction for advanced learning technologies, with an emphasis on game-based learning environments, intelligent tutoring systems, user modeling, educational data mining, and computational models of interactive narrative generation.*

Dr. Randall Spain *is a Research Psychologist in the Center for Educational Informatics at North Carolina State University. He earned his Ph.D. in Human Factors Psychology from Old Dominion University in 2009 and has been actively conducting training and human factors research for the Department of Defense and the Department of Homeland Security for the past 10 years. His research uses principles, theories, and methods of applied psychology (human factors, educational psychology, personnel psychology, experimental psychology, and psychometrics) to evaluate the impact of advanced training technologies on learning and performance.*

Dr. Robert Pokorny *is Principal of the Education and Training Technologies Division at Intelligent Automation, Inc. He earned his Ph.D. in Experimental Psychology at University of Oregon in 1985, and completed a postdoctoral appointment at University of Texas at Austin in Artificial Intelligence. Bob's 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, Bob 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. Bradford Mott *is a Senior Research Scientist in the Center for Educational Informatics at North Carolina State University. Prior to joining North Carolina State University, he served as Technical Director at Emergent Game Technologies where he created cross-platform middleware solutions for video game development, including solutions for the PlayStation 3, Wii, and Xbox 360. Dr. Mott received his Ph.D. in Computer Science from North Carolina State University in 2006, where his research focused on intelligent game-based learning environments. His current research interests include computer games, computational models of interactive narrative, and intelligent game-based learning environments.*

Dr. Benjamin Goldberg *is a member of the Learning in Intelligent Tutoring Environments (LITE) Lab at the U.S. Army Research Laboratory's (ARL) Human Research and Engineering Directorate (HRED), Simulation and Training Technology Center (STTC) in Orlando, FL. He has been conducting research in modeling and simulation for the past five 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 is a Ph.D. graduate from the University of Central Florida in the program of Modeling & Simulation.*

Dr. James Lester *is Distinguished Professor of Computer Science at North Carolina State University, where he is 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 Ph.D. 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).*