Pedagogical Management in Support of a Generalized Framework for Intelligent Tutoring

Benjamin Goldberg[1]
U.S. Army Research Laboratory[1]

INTRODUCTION

A current goal associated with the development of the Generalized Intelligent Framework for Tutoring (GIFT) is providing a set of tools for training practitioners to rapidly build adaptive instructional materials based on an interplay of knowledge acquisition and skill development. To accommodate this guiding requirement, an instructional management research vector was devised as a means to coordinate resources and efforts to meet the needs of end users. While the science surrounding intelligent tutoring system (ITS) development is multidisciplinary, a guiding assumption is that targeted GIFT developers in the military community will be subject matter experts (SMEs) within their respected fields, but lack many of the technical disciplines that go into ITS development. Accordingly, it is also safe to assume these SME developers will also require assistance in authoring training content that adheres to learning science principles. As such, authoring workflows and ITS methods must be developed to compensate for the skills a GIFT user lacks when creating a lesson or course.

Due to this challenge, GIFT development in the instructional management vector aims at providing a means for embedding pedagogical theory into GIFT schemas and authoring workflows. The goal is to develop enabling technologies that allow SMEs to author GIFT-based lesson materials that are empirically informed and grounded in instructional design theory. In this paper, we present the current state of GIFT as it relates to instructional management capabilities and associated research to extend how GIFT can manage and personalize training interactions. As an organizing function, we arrange GIFT pedagogy into three temporal categories: (1) instructional management at the lesson level, which personalizes content and adapts course sequencing based on performance and persistent learner attributes (i.e., outer-loop adaptation), (2) instructional management at the interaction level, which deals with real-time coaching and scenario manipulation across an array of practice events (i.e., inner-loop adaptation), and (3) instructional management at the after-action review level, which focuses on reflection and remediation practices following completion of a learning event. Each category has a set of unique features dictated by instructional theory, with on-going projects informing their design and evaluating their utility. We will present the foundations informing the methods applied, and the current state of their practice. We will conclude with future directions of GIFT development, and how the instructional management research vector is aligned to meet training requirements road mapped for future training applications and methods.

Dimensions of Instructional Management Research

In November 2015 members of the GIFT team published a research outline that examined specific goals and interests associated with instructional management in ITS type environments (Goldberg, Sinatra, Sottilare, Moss & Graesser, 2015; Goldberg, Sottilare, Moss & Sinatra, 2015). The authors identified the following dimensions as critical benchmarks in driving capability enhancements to GIFT pedagogical practices:

- Guidance and Scaffolding: focuses on identifying a set of pedagogical best practices that adhere to the tenets of learning and skill development. The challenge is identifying methods that generalize across domains and task environments, and providing tools flexible enough to create scaffolding that can be represented in domain-agnostic terms. Current research aims at creating logic to manage timing, specificity, and modality determinations of intervention content at the individual level.

- Social Dynamics and Virtual Humans: focuses on the social component of learning, and building tools and methods that adhere to the social cognitive tenets of how individuals interact to instill knowledge and solve problems (Bandura, 1986; Vygotsky, 1978). From an adaptive instructional management standpoint, social dynamics is concerned with: (1) using technology to replicate interactive discourses common in learning and operational settings, (2) using technology to create realistic and reactive virtual humans as training elements in a simulation or scenario, and (3) using technology to create social networks for the purpose of supporting peer-to-peer and collaborative learning opportunities.

- Metacognition and Self-Regulated Learning (SRL): focuses on instructional management practices that aim at building habits linked to successful regulation of learning practices and that promote metacognitive applications. This approach to instructional management varies from traditional guidance and scaffolding techniques as it focuses on behavior and application of strategy, rather than on task dependent performance. This research area is of interest as it is based around GIFT supporting SRL, and the efficacy of defining and modeling persistent metacognitive strategies that can be applied across domain applications. The goal is to embed instructional supports that promote situational awareness, and guide learners in planning, monitoring, and reflection based activities.

- Personalization (Occupational and Non-Cognitive Factors): focuses on the use of learner dependent information to personalize a training experience. This can involve personalizing content based on interests, with the goal of inducing a higher level of motivation when the context of a learning event is framed within a use case the learner cares about. In addition, the personalization dimension is also interested in identifying ways to automatically personalize training interactions based on occupational factors that are unique to their upcoming assignment or current job description. All of these instructional management practices require research to identify mechanisms for easily implementing personalization techniques, along with empirical evidence supporting their application for wide GIFT application.

The dimensions reviewed above provide a means for organizing and prioritizing efforts to enhance GIFT’s current instructional management support. While the research outline mapped out desired end-state functions of GIFT, it is also important to capture the current state of practice, as those piece parts are the ultimate methods rolled out to the community at large. In the remainder of this paper, we identify GIFT’s instructional management functions, and how they apply to future enhancements that aim to meet the goals of the overarching instructional management capability dimensions.

CURRENT PRACTICE OF INSTRUCTIONAL MANAGEMENT IN GIFT

In the remainder of the paper, we present the current state of practice as it pertains to instructional management capabilities built within GIFT. Much of the projects over the past year have been influenced by the road mapping exercise captured in the research outline described above. GIFT pedagogy and instructional management will be described at three levels of interaction, each with a distinct set of adaptive options supported. These include: (1) the lesson level, (2) the interaction level, and (3) the after-action level.

Instructional Management at the Lesson Level

At the lesson level GIFT provides the tools to build a sequence of course objects that dictate the learner experience (see Figure 1). Course objects are designed to either inform the learner, collect information from the learner, or manage execution of content delivery and assessment across problem-sets and scenarios configured during the authoring process. From an instructional management standpoint, pedagogy at the lesson level focuses primarily on lesson sequencing and personalization. It manages adaptations at the macro outer-loop level, where models have been established that dictate what a learner will experience next based on an established learner model and performance data captured during prior interactions. The lesson level pedagogy logic is currently captured within GIFT's first generalized pedagogical model called the Engine for Management of Adaptive Pedagogy (EMAP).

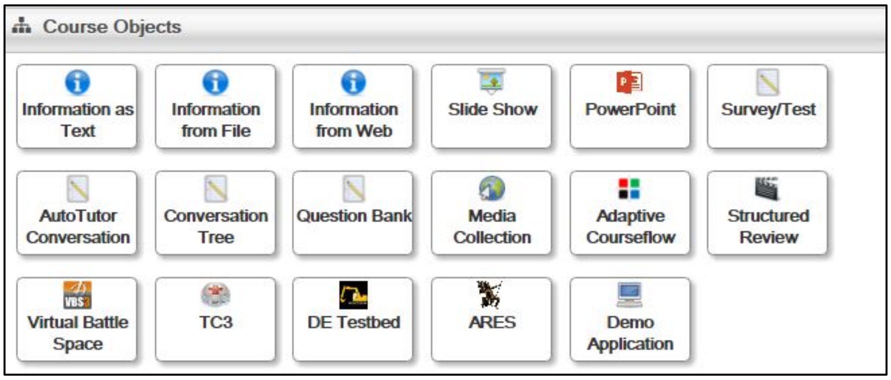

Figure 1. Current list of supported GIFT course objects

EMAP

The EMAP is GIFT's first domain-independent pedagogical model informed by instructional theory, with much of the research and development documented over the years (Goldberg et al., 2012; Wang et al., 2013; Goldberg, Hoffman & Tarr, 2015). What the EMAP provides is a theoretical instructional design framework embedded in GIFT that guides the authoring and configuration of adaptive learning experiences. The EMAP is based on David Merrill's component display theory (CDT), with learning broken up into four primary categories: (1) presenting rules of a domain, (2) presenting examples of those rules applied, (3) asking the learner to recall information as it relates to the domain, and (4) allowing the learner to practice the application of those rules in a novel context for the purpose of skill development (see Figure 2).

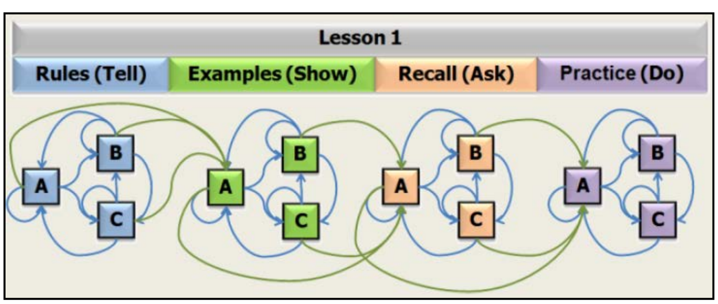

Figure 2. EMAP representation for a lesson teaching three overall concepts and all possible permutations based on variations in assessment outcomes

These four categories are the building blocks of the Adaptive Courseflow GIFT course object. Each category is embedded within the object's schema, where an author configures the content, assessments, and practice events in each categorical bin. This includes loading in all relevant lesson content into the rules and examples bins, and establishing metadata tags for each file. The metadata tags are currently informed from IEEE's Learning Object Metadata (LOM) standard (Mitchell & Farha, 2007), and are used to describe what that associated piece of content covers and the type of materials established within (e.g., videos, figures, worked examples, etc.). GIFT has back end logic informed from a populated pedagogical configuration file that matches learner model attributes (e.g., prior knowledge, motivation, etc.) with metadata descriptors that associate with content, difficulty, and scaffolding type recommendations.

The resulting decision tree was created following the completion of a literature review that aimed to capture instructional management best practices that could be organized at a domain-independent level (Goldberg et al., 2012). A recognized limitation from the literature review was a lack of generalized findings across numerous studies. This is due in part to the nature in which instructional strategies are defined. Each generalized strategy must be contextualized into the domain and application context it will be delivered within, thus producing confounding factors for defining the specific characteristics associated with an investigated intervention. However, a substantial finding from this project was the recognition of four learner attributes found to account for consistent variance in performance outcomes across studies, including (1) prior knowledge/skill, (2) motivation, (3) self-regulatory ability, and (4) grit/perseverance (Wang et al., 2013). These variables served as the moderators to base the first EMAP instantiation around. Access the GIFT documentation for a complete breakdown of the EMAP and its underlying logic (https://gifttutoring.org/projects/gift/wiki/Engine_For_Management_of_Adaptive_Pedagogy_(eMAP)_2020-1).

A recognized limitation of the current EMAP is the deterministic nature in which the model was developed. The EMAP functions as a decision tree that maps learner traits and attributes with content descriptors through GIFT's pedagogical configuration file. The configuration file is referenced at runtime, with a content selection algorithm in place that identifies the best piece of content based on concept coverage and the most metadata matches based on a learners profile and available rule and example content. In addition, the remediation logic was also recognized as being rather simplistic, where it would select a new piece of content, if one existed, covering the concept not meeting assessment criteria in either the recall or practice phase of the EMAP interaction. In these instances, the remediation logic passes a learner back into the rule or example category where content is selected for presentation. This remediation approach assumes there is material designated to support an intervention and that the learner has the required understanding to use the new information to correct the misconceptions or impasses identified during assessment.

EMAP Enhancements

A current effort led by North Carolina State University and Intelligent Automation, Inc. is applying methods grounded in tutorial planning and reinforcement learning to extend the current state of the EMAP to support a stochastic modeling approach that introduces probabilistic reasoning in the selection of tutorial actions carried out by GIFT at the lesson level. The approach incorporates embedding a learning activity framework put forth by Chi (2009) that differentiates activities undertaken by learners into constructive, active, and passive (CAP) interactions. This CAP activity framework enables the incorporation of multiple remediation types for a given concept, and is well-suited for markov decision processes (MDPs) to dictate what intervention is best for what type of learner. In addition, this approach supports a back-end reinforcement learning method that trains and optimizes the MDP policies over time as more data and evidence is made available following interactions from a large set of learners.

To accommodate this approach, the adaptive courseflow GIFT course object is being re-factored to support a category of content designated for remediation purposes. Now, an author has the ability to establish core rule and example lesson materials (i.e., content all learners will see regardless of learner profiles), with a remedial bin that enables a developer to build CAP-based tutorial interactions for remediation support following assessment events managed by the EMAP. When performance states are available, the EMAP will select remedial materials based on outputs from the MDP policies. To support policy optimization, a GIFT reinforcement learning tool is being developed for the purpose of updating policies as evidence is made available on the utility of their application. The described enhancements are currently under development during the writing of this chapter, with data collections planned for model training and evaluation over the next twelve months.

EMAP and Massive Open Online Courses (MOOCs)
With the EMAP providing a generalized pedagogical framework to structure personalized content selection and remediation, one area of interest is how GIFT fits within the context of MOOCs. A current effort in collaboration with the University of Pennsylvania and Carnegie Mellon University is looking at the utility of GIFT in providing a standardized framework for personalizing MOOC interactions. For the effort, GIFT is being integrated with the MOOC platform edX (https://www.edx.org/). By making GIFT compliant with the Learning Tools Interoperability (LTI; Severance, Hanss & Hardin, 2010) standard, developers in edX can reference GIFT configured lessons within the courseflow of their MOOC. In this instance, a MOOC can handover control to GIFT for a designated lesson where the EMAP can be applied to personalize the experience an individual receives based on the content and assessment made available by the lesson creator. Following completion of that lesson, GIFT can communicate results and behavior data back to edX for performance tracking and accreditation purposes. These described features are currently being developed, with experimentation planned to determine if personalization practices improve MOOC usage and overall learning outcomes.

Instructional Management at the Interaction Level

At next level of GIFT pedagogy, instructional management has to do with monitoring real-time interaction and managing specific events through coaching and scenario adaptation practices. A major function of GIFT is its ability to capture real-time interaction data from external training environments through an interop configuration. These environments are used to support experiential practice type opportunities for learners, and range from PowerPoint slide presentations to interactive first-person shooter type game environments. With an established gateway module, GIFT can capture data produced from any external system and route specified information into the domain module for assessment purposes (Sottilare, Goldberg, Brawner & Holden, 2012). In the domain module, a domain knowledge file (DKF) is configured for the purpose of contextualizing raw system data around a set of concepts represented in a task ontology. The ontology organizes a scenario into a set of tasks a learner will be asked to perform, a set of concepts for each task a learner will be measured against, and a set of conditions defined to inform the performance measurement of those concepts. In this instance, raw data is used for the purpose of informing measures to gauge performance and infer competency.

From an instructional management perspective, there are currently four supported pedagogical requests communicated from the pedagogical module to the domain module when interacting with an external training application. These include, (1) request-instructional-intervention (e.g., provide guidance through the form of a hint or prompt; see Figure 3 for an example of a coaching hint communicated by GIFT), (2) request performance assessment (e.g., ask a question of the learner to update learner state), (3) request scenario adaptation (e.g., modify the scenario or problem to adjust difficulty), and (4) do nothing. These message types are currently informed by observable performance state transitions across the concepts being tracked in a DKF (e.g., performance on concept 2a transitioned from at-expectation to below-expectation). These transitions are reported out to the learner module, where the learner state is defined, including performance and all other relevant attributes in the learner model. The learner state is then passed to the pedagogical module for determining intervention type at the individual level.

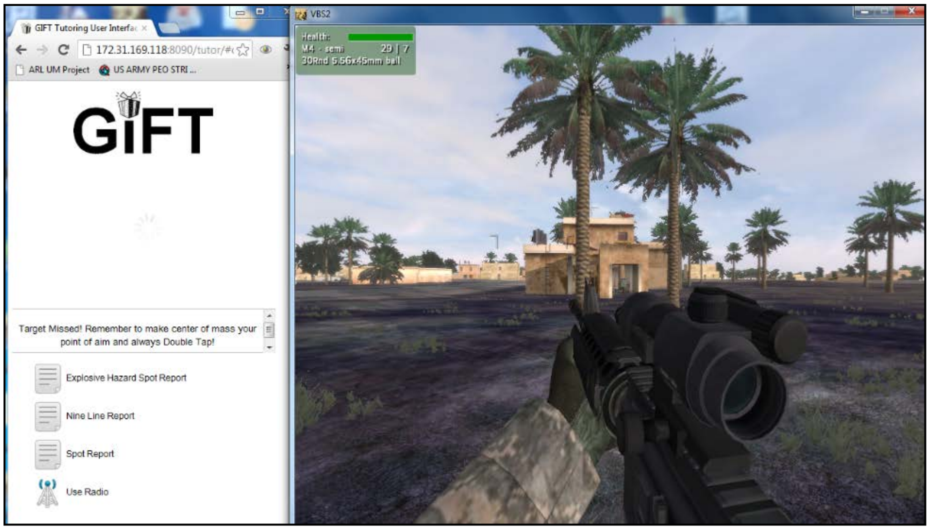

Figure 3. GIFT providing a real-time hint on 'Rules of Engagement' concept in a Virtual Battle Space scenario

In its current state, a GIFT developer configures the strategy types to be enacted when a specific concept transition is observed. The developer is then responsible for translating that strategy request into a specific tactic to be executed at run-time (i.e., define the exact hint to be presented when a request-instructional-intervention request is received for any given concept in the DKF). While the tools and methods in GIFT for interaction level pedagogy support real-time interventions when performance conditions are met, there is much work left to be done in determining how best to intervene and adapt based on individual differences and learner profile information.

DKF Enhancements to Support CAP Tutorial Planning

A recognized limitation in the current run-time instructional support of GIFT for individuals interacting with an external training application is the lack of logic to drive personalized coaching based on individual differences. At the moment, regardless of information associated with an individual's learner model, GIFT is set up to provide the same real-time tutorial actions based on observed transitions defined in the DKF. In lieu of this technical short-fall, current work is investigating the application of the aforementioned CAP instructional activity model to fit within the context of DKF pedagogical practice. This will enable GIFT to provide passive feedback information during a training event when appropriate, while also enabling GIFT to intervene with targeted activity exercises that are aimed at coaching a specific concept or skill when assessment logic (i.e., tutorial planning MDPs) deems it appropriate. The policies are designed to optimize overtime through a reinforcement learning method applied to the back-end as data is made available across practice scenarios. This also supports a reuse philosophy of intervention materials, where CAP associated activities can be triggered at both the lesson and interaction level, reducing the requirement to author interventions across both instances.

Metacognitive Training

Another area of instructional management the GIFT system aims to support is using pedagogical practices to aid learners in developing metacognitive skills that associate with self-regulated learning, critical thinking, and on-the-spot problem solving (Biswas, Segedy & Kinnebrew, 2014). For this reason, a current effort in collaboration with Vanderbilt University is investigating how to use GIFT to model self-regulated learning behaviors and metacognitive skills for the purpose of driving interventions aimed at improving how individuals go about solving a problem, rather than focusing on the problem solution itself. While the main component of this research is focused on the learner modeling aspect to start, the overarching theme of the project is driven by eventual instructional support currently not provided in GIFT. The resulting effort is restructuring the hierarchical schematics of the DKF to support a layered inference procedure (see Rajendran et al.'s technical discussion on the learner modeling work further down in the proceedings). With a learner modeling approach in place, GIFT can infer an individual's understanding of metacognition and self-regulated learning for the purpose of driving focused interventions that guide users in the application a identified behaviors congruent with effective problem-solving applications.

Psychomotor Skill Development and Coaching

Another exciting area of instructional management research is seen in examining GIFT's utility to train psychomotor skill domains. This is a complicated application of ITS as it breaks away from common cognitive problem spaces these systems are traditionally developed within. As such, there is still much research to be done on the modeling and pedagogy components of training a psychomotor task in the absence of human instructors. From an instructional management perspective, work is being performed to account for psychomotor training at both the lesson and interaction level (see Brown, Bell & Goldberg paper in the proceedings for a full breakdown of conceptualized approach being implemented). At the lesson level, this involves creating course objects based on the abstraction of the EMAP that takes into account instructional theory models grounded in psychomotor application. This breaks away from the EMAP's dependency on the CDT, where these objects can now associate with any number of instructional models that a developer wants to base their interactions within. At the interaction level, research is required to determine how best to manage feedback and remediation practices. This involves creating data capture techniques that provide inputs used for modeling physical behaviors, and building representations of those behaviors that are used to guide assessment practices that ultimately inform pedagogical decisions.

Instructional Management and After Action Review

While much of the current instructional management functions in GIFT focus on real-time interactions at varying levels of granularity, new tools and methods are being developed to support personalized after-action review (AAR) materials. For complex skill domains, AARs serve as critical functions in the training process. In these instances, learners have the opportunity to reflect on problems and scenarios undertaken for the purpose of critiquing their own interaction and understanding the implications of their actions on reaching scenario and specific task objectives. A current effort (see Carlin, Brawner, Nucci, Kramer & Oster in the proceedings for details on AAR approaches being investigated in GIFT) is examining how to apply modeling methods to create individualized AAR interactions based on what is observed across a GIFT managed lesson. The goal is to develop technologies that automatically identify critical errors and misconceptions by the learner(s) and automatically select an optimal instructional path and associated instructional content to construct an AAR. The project is applying MDP inference procedures for identifying concepts to personalize an AAR around, along with the organization of content and activities that target the goals of the AAR interaction. These goals include reinforcing learning objectives, addressing impasses, and contextualizing the lesson and training with real world application through mental reflection exercises.

CONCLUSIONS

In this paper, we presented a snapshot of current instructional management capabilities within GIFT, along with ongoing efforts aimed at enhancing applied methods. GIFT is a moving target in terms of development, so it is important to document the methods applied within GIFT and the research that went into its implementation. The pedagogical infrastructure in GIFT is maintained at the lesson level, the interaction level, and the after action level where varying modeling techniques are applied to determine the instructional adaptation/intervention to enact. There is still much work to be done before GIFT's pedagogical practices are easily implemented in an operational context. In addition, GIFT must be able to adapt pedagogical practice as future training instances and applications are developed and transitioned to the Warfighter.

Future Directions

Current trends in ITS research as it relates to the GIFT project is focusing on two fronts: (1) using adaptive training and education practices to support team development and cohesion and (2) using adaptive training and education practices within mobile applications to support ease of access and on-the-spot training support. These themes will be addressed in the coming years as team-based and mobile-based ITS applications mature. Both themes are being addressed in current projects, but the instructional management components associated with their instantiations have yet to be examined. As the future of training and education evolves, GIFT is set up to instructionally support all facets of learning and skill development.

REFERENCES

Bandura, A. (1986). *Social Foundations of Thought and Action: A Social Cognitive Theory*. Englewood Cliffs, N.J.: Prentice Hall.

Biswas, G., Segedy, J. R., & Kinnebrew, J. S. (2014). A Combined Theory-and Data-Driven Approach for Interpreting Learners' Metacognitive Behaviors in Open-Ended Tutoring Environments. In R. Sottilare, A. Graesser, X. Hu & B. Goldberg (Eds.), *Design Recommendations for Intelligent Tutoring Systems (Volume 2: Instructional Management)*.

Chi, M. T. (2009). Active-constructive-interactive: A conceptual framework for differentiating learning activities. *Topics in Cognitive Science*, 1(1), 73-105.

Goldberg, B., Brawner, K. W., Sottilare, R., Tarr, R., Billings, D. R., & Malone, N. (2012). Use of Evidence-based Strategies to Enhance the Extensibility of Adaptive Tutoring Technologies. Paper presented at the Interservice/Industry Training, Simulation, and Education Conference (I/ITSEC) 2012, Orlando, FL.

Goldberg, B., Hoffman, M., & Tarr, R. (2015). Authoring Instructional Management Logic in GIFT Using the Engine for Management of Adaptive Pedagogy (EMAP). In R. Sottilare, A. Graesser, X. Hu & K. Brawner (Eds.), *Design Recommendations for Intelligent Tutoring Systems: Authoring Tools (Volume 3)*: U.S. Army Research Laboratory.

Goldberg, B., Sinatra, A., Sottilare, R., Moss, J., & Graesser, A. (2015). Instructional Management for Adaptive Training and Education in Support of the US Army Learning Model-Research Outline: DTIC Document.

Goldberg, B., Sottilare, R., Moss, J., & Sinatra, A. (2015). *Dimensions of Instructional Management Research in Support of the Generalized Intelligent Framework for Tutoring (GIFT).* Paper presented at the Defense and Homeland Security Simulation (DHSS) Workshop, Bergeggi, Italy.

Mitchell, J. L. & Farha, N. (2007). Learning Object Metadata: Use and Discovery. In K. Harman & A. Koohang (Eds.), Learning Objects: Standards, Metadata, Repositories, and LCMS (pp. 1-40). Santa Rosa, CA: Informing Science Press.

Severance, C., Hanss, T., & Hardin, J. (2010). Ims learning tools interoperability: Enabling a mash-up approach to teaching and learning tools. *Technology, Instruction, Cognition and Learning*, *7*(3-4), 245-262.

Sottilare, R., Goldberg, B., Brawner, K. W., & Holden, H. (2012). *Modular Framework to Support the Authoring and Assessment of Adaptive Computer-Based Tutoring Systems*. Paper presented at the Interservice/Industry Training, Simulation, and Education Conference (I/ITSEC), Orlando, FL.

Vygotsky, L. S. (1978). *Mind in Society: The development of higher psychological processes*. Cambridge, MA: Harvard University Press.

Wang-Costello, J., Tarr, R. W., Cintron, L. M., Jiang, H. & Goldberg, B. (2013). Creating an Advanced Pedagogical Model to Improve Intelligent Tutoring Technologies. Paper presented at the The Interservice/Industry Training, Simulation & Education Conference (I/ITSEC).

ABOUT THE AUTHORS

Dr. Benjamin Goldberg *is an adaptive training scientist at the Army Research Laboratory's Human Research and Engineering Directorate. He leads research focused on instructional management within ARL's Learning in Intelligent Tutoring Environments (LITE) Lab and is a co-creator of the Generalized Intelligent Framework for Tutoring (GIFT).*