Adaptive Learning in Simulation

Robert Pokorny, Wilbur Peng, Jacqueline Haynes,
Intelligent Automation, Inc.
Rockville, Maryland
bpokorny, wpeng, jhaynes @email.com,

Jonathan Rowe, Bradford Mott, James Lester,
North Carolina State University
Raleigh, North Carolina
jprowe, bwmott, lester@ncsu.edu

Benjamin Goldberg
Army Research Laboratory
Orlando, Florida
benjamin.s.goldberg.civ@mail.mil

ABSTRACT

To enable adaptive training for Soldiers, Sailors, and Airmen in simulations and immersive environments, adaptive training systems must become affordable. Currently, they require handcrafting by experts with specialized knowledge of Artificial Intelligence working closely with domain experts. We focus on three complementary tasks that address affordable adaptive tutoring: (1) identifying performance weaknesses on which trainees must improve; (2) making better tutoring decisions; and (3) integrating adaptive instruction into a common architecture enabling capabilities developed once to be applied elsewhere. We report experimental results from studies investigating these tasks.

We have applied *expert policy capture* to identify student performance weaknesses. Experts (1) reviewed many students' performance in a simulation, (2) critiqued that performance, and (3) rated performance for overall quality. Analysts constructed scoring rules that mimicked experts' judgments. For each identified area of student weakness, an instructional intervention was constructed.

To determine which instructional interventions students should interact with, we are using a modular reinforcement learning framework that uses the history of many learners to determine which instructional approaches work best for which types of learners and contexts. This approach decomposes adaptive instruction in terms of *adaptable event sequences*, which are each separately modeled as a Markov decision process (MDP). The instructional policies for the MDPs are induced using off-line reinforcement learning techniques with a corpus of prior learner interaction data.

The instruction is presented within a unifying framework for adaptive instruction called the Generalized Intelligent Framework for Tutoring (GIFT). GIFT enables instructional capabilities that are developed within GIFT to be applied in the future to other instructional systems and training environments. GIFT authoring tools support domain independent principles of instruction, as well as domain specific instruction.

ABOUT THE AUTHORS

Dr. Bob Pokorny is Director of the Education and Training Technology 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. Wilbur Peng received his B.S. in Electrical Engineering from Cornell University, Ithaca N.Y. in 1992, and received his PhD in Electrical Engineering from the University of Maryland at College Park in 2002. His dissertation research addressed the problem of representing and using expert knowledge in the form of graph-based similarity structures. At IAI, Dr. Peng leads research and development projects in several areas, including cluster and distributed computing; autonomous aerial vehicle software control architecture; parallel discrete event simulation on multicore distributed computing clusters; network simulation, emulation and modeling; and semantic and ontological modeling methodologies, frameworks and applications.

Dr. Jacqueline Haynes is co-founder and Executive Vice President at Intelligent Automation, Inc. Her background combines education and psychology with artificial intelligence applications. She received her Ph.D. from the University of Maryland in Curriculum and Instruction, and did post-doctoral work there in artificial intelligence and intelligent tutoring systems. Previously she was a faculty member at the University of Maryland, College of Education. Her research interests include research-based instructional design, tools for web-based instruction, and reading comprehension.

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.

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. 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).

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) in Orlando, FL. He has been conducting research in the Modeling & Simulation community for the past eight years with a focus on adaptive learning in simulation-based environments and how to leverage Artificial Intelligence tools and methods to create personalized learning experiences. Currently, he is the LITE Lab's lead scientist on instructional management research within adaptive training environments and is a co-creator of the Generalized Intelligent Framework for Tutoring (GIFT). Dr. Goldberg is a Ph.D. graduate from the University of Central Florida in the program of Modeling & Simulation. His work has been published across several well-known conferences, with recent contributions to the Human Factors and Ergonomics Society (HFES), Artificial Intelligence in Education and Intelligent Tutoring Systems (ITS) proceedings. Dr. Goldberg has also recently contributed to the journal Computers in Human Behavior and to the Journal of Cognitive Technology.

Adaptive Learning in Simulation

Robert Pokorny, Wilbur Peng, Jacqueline Haynes,
Intelligent Automation, Inc.
Rockville, Maryland
bpokorny, wpeng, jhaynes @email.com,

Jonathan Rowe, Bradford Mott, James Lester
North Carolina State University
Raleigh, North Carolina
jprowe, bwmott,lester@ncsu.edu

Benjamin Goldberg
Army Research Laboratory
Orlando, Florida
benjamin.s.goldberg.civ@mail.mil

INTRODUCTION

Students can learn effectively from simulations. To optimize learning efficiency in simulations, students' performance should be assessed, and instruction should adapt to what students need to know.

Assessing student performance in open-ended simulations is often seen as difficult. The assessment system should ideally evaluate overall performance, identify the content that students need to learn, and link to instructional content that remedies poor performance. Further, each trainee will likely have many areas in which his/her performance can be improved. The assessment should provide information that reflects the severity of the student's weaknesses.

A second challenge for training with simulations is to select and sequence topics that should be remediated. Major instructional questions are (1) when to intervene, (2) what topic to address, and (3) how to intervene (with what instructional strategy or tactic). In regards to instructional strategies and tactics, there are many ways to support student learning in simulations. The instruction could provide notices to human instructors, and allow instructors to mediate the instruction. The instruction could explain the reasons underlying preferred performance within the simulation directly. The instruction could use a scaffolding approach by leading the student to construct his/her own understanding of preferred performance within the simulation. The instructional system could monitor student attention, and alter the remediation based on whether the student is paying attention and performing poorly, or is distracted and performing poorly.

A third challenge is to integrate the capabilities developed for one computer-based training system (CBTS) into other CBTSs. Ideally, once an instructional capability has been constructed for one domain, it would be applicable to other CBTS. Domain-specific information must be incorporated into the CBTS, but the structures, methods, and general-purpose capabilities should be shared. For example, the ability to track students' attention should apply across domains. Instructional interventions should differ if poor performance is due to (1) not paying attention versus (2) paying attention but not understanding. A framework that supports effective assessment and tailored instructional interventions can enable CBTSs to be more powerful and efficient.

In this paper, we describe how training simulations can reasonably address the issues stated above. We will report preparation for a study that integrates current approaches to assessing performance, selecting and sequencing instruction, and providing a framework through which intelligent tutoring capabilities can be applied.

Approaches to these three challenges

We are investigating methods for (1) developing assessment tied to instructional content; (2) selecting, sequencing, and supporting instructional interactions for students; and (3) integrating key functionalities of adaptive training systems with a generalized software framework. To develop assessment and linked instruction, we apply a method in which experts review and critique work samples of many performers. Task analysts integrate the critiques into scoring rules and validate them against expert judgments of new work samples. Associated with each assessment-identified weakness is an instructional module. To adaptively control when instructional content is presented, as well as how it is presented, we have developed an intelligent tutoring framework that is based on concurrent Markov decision processes, which are solved using modular reinforcement learning techniques. These methods can take existing rules for selecting and sequencing content, and then use machine learning methods to improve the quality with which instruction is sequenced. Finally, to integrate various tutoring capabilities into a common framework, we are using the Generalized Intelligent Framework for Tutoring (GIFT), developed by the Army Research Laboratory. GIFT provides a general software framework for intelligent tutoring systems in which modular components of tutoring systems can be integrated, coordinated, and reused across training environments.

The use case in which we developed this adaptive training was UrbanSim, a practice environment in which students play the role of a Battalion Commander in a Full Spectrum operation. While UrbanSim was originally created for Battalion and Brigade Commanders before they took those positions in Iraq in the mid-2000s, UrbanSim is now used for educating lower ranking Soldier about how to conceptualize Counter-Insurgency (COIN) operations. We developed an assessment system for trainees within UrbanSim which identifies strengths and weaknesses in performance. We also created instructional modules linked to each instructional weakness.

For purposes of this investigation, we have focused on two of six dimensions identified by the instructional assessment. These are (1) the degree and severity of security actions that students take, and (2) the communications and meetings that students schedule with host nation leaders. The instructional modules focus on specific weaknesses that were detectable from the trainee's performance (e.g., allowing all three of the insurgent groups to exist through turn 5, which indicates little progress toward the goal of eradicating insurgents, or meeting the Deputy Mayor before the Mayor, which is disrespectful to the Mayor). In addition, the instructional modules address the general principles of security balance with other factors necessary to conduct successful COIN operations. Finally, they provide feedback on learning behaviors that are off task, such as shooting a host nation leader, as well as other behaviors that indicate a student is not taking the instruction seriously.

For sequencing instruction, we use Markov decision processes to model (1) which pedagogical strategy to use, such as whether to focus on one performance weakness or multiple categories of weaknesses; and (2) which pedagogical tactics to apply. We developed instructional tactics inspired by three of four instructional elements of Micheline Chi's ICAP framework (Chi, 2009), which distinguishes between four levels of cognitive engagement with learning activities: passive, active, constructive, and interactive. In our instruction, depending on the pedagogical tactic chosen, the student either (1) passively reads an instructional message about one (or multiple) performance weaknesses that he/she has shown in UrbanSim; (2) reads and actively highlights key phrases in the instructional message; or (3) reads and constructively summarizes key elements of the instructional message. We chose not to operationalize the fourth *interactive* category of the framework, as it generally involves engaging in tutorial dialogues with humans, which is beyond the scope of this project. But it should be noted that students *do* engage in highly interactive learning while using UrbanSim. The use of the ICAP framework here simply describes how students interact with remedial instructional content that is delivered during training to address performance weaknesses.

Finally, GIFT has been used in multiple ways to integrate various instructional approaches with UrbanSim. We have previously applied this instruction to work with GIFT so that the focused coaching that reflects the student's weakest performance is presented to students. GIFT provides the framework in which different assessment systems can be used and tied to different instruction, while using the available infrastructure to present the instruction and record and analyze the data.

With this introduction, we next present the issues surrounding each of these three major concerns when building adaptive simulation-based training, and we discuss how we address them within the context of UrbanSim training.

ASSESSMENT AND INSTRUCTION

The Challenge of Assessment and Instruction

To optimize learning efficiency in games and simulations, students' performance should be assessed, and instruction should be tailored to students' needs. However, designers of Intelligent Tutoring Systems (ITS) report that developing assessment systems for constructing student models and using them to guide instruction is complicated and time-consuming. Woolf (2010) states, "Representing and reasoning about students, domain, and teaching knowledge can involve complex and difficult processes . . . Yet these difficulties are inherent in the process of creating student and teaching models." Mislevy has developed an assessment approach for simulations he calls "Evidence Centered

Design." Mislevy and Riconscente (2006) describe their work in Evidence Centered Design as being labor intensive and time consuming: "Initial applications of the ideas encompassed in the Evidence Centered Design framework may be labor intensive and time consuming. Nevertheless, the import of the ideas for improving assessment will become clear."

Existing experimental procedures and statistical methods do not easily apply to assessing performance in games and simulations. As pointed out by Honey and Hilton (2011), most assessment methods do not address how to generalize student capabilities when measurement occurs across the variety of contexts that students find themselves in. To quote from Chapter 5, "Fundamentally, assessment based on standard psychometric approaches requires drawing inferences in real time about student learning from these diverse behaviors and performances. However, most conventional psychometric theory and methods are not well suited for such modeling and interpretation."

One approach to assess student performance in complex environments is the previously mentioned Evidence Centered Design, (ECD), championed by Robert Mislevy. In ECD, assessment begins by designing the competencies and levels of mastery in performance that a student currently has. An evidence model is developed to relate how performance in an environment could be related to estimates of the student's competency or current mastery level. Then a game or simulation environment is created so that the situations defined by the evidence model can be observed. When the student is unaware that the performance he/she executes is being used for assessment, the assessment is called "stealth assessment." (Shute, 2011)

The goals of ECD and Stealth Assessment are admirable goals, and adaptive training does need a way to assess performance in order to drive adaptive training. While the goals are exactly what we seek, there are a few impediments. One challenge of using ECD is the cost. Wang, Shute, and Moore (2015) state, "The major limitation of implementing stealth assessment using ECD is the cost in terms of time and effort, whether it is a commercial or a homemade game." They continue "the process usually spans one or more years involving learning scientists, psychometricians, game designers, programmers, and possibly others (e.g., content experts)." In a 2013 review of the state of ECD, Kristen DiCarbo (2013) wrote, "I would want to see a game in production that uses assessment models to provide valid, actionable feedback to students and teachers about specific knowledge, skills, or attributes of interest in real time. We're working on it, and it isn't too far in the future, but we're not there yet."

The combination of ECD and stealth assessment is currently best used for projects that have considerable time and money to apply the approach. In order to create adaptive training systems with simulations that can be widely applied, assessment systems would ideally be affordable and lead reliably to instructional interventions that target identified student weaknesses. Evidence Centered Design for Learning (ECDL) enhances ECD with (1) a pedagogical model that aims to describe how to improve learning, and (2) suggests ways to improve associated learning by incorporating engagement, accessibility, and learning efficiency (Feng, Hansen, and Zapata-Rivera (2009).

An Approach to Creating Simulation-based Assessment and Instruction

To develop an assessment system where each assessment category is linked with an instructional topic, we developed Performance Evaluation by Expert Evaluation (PEER). The inspiration for PEER was to evaluate an ITS that taught troubleshooting of complex equipment (Pokorny, et. al., 2010). We needed to assess performance of technicians as they completed troubleshooting tasks. We originally had experts provide holistic judgments of troubleshooting quality. Even though the troubleshooting actions that a trainee could take were unbounded, expert reviewers' scores of overall quality were always correlated statistically significantly, and were typically in the .80s. We had conducted task analyses of the domain to develop the ITS, and we had learned how experts grouped the equipment and classified types of actions and results that might occur on any action on any equipment section. When having to conduct troubleshooting studies when we did not have access to known experts, we investigated if we could develop a scoring system that mimicked experts' scores based on experts' critiques of work samples.

We used PEER to elicit expert knowledge and generated scoring rules for assessing performance. The scoring rules identify student performance strengths and weaknesses. Generally, experts reviewed work samples from a broad range of performers, assigned scores of overall quality, and then provided critiques that justified their scores. These critiques were turned into scoring rules and later validated. From experts' critiques, we constructed a rough draft of instructional content tied to each performance weakness.

A Method Used to Create Simulation-based Assessment and Instruction

PEER elicits experts' assessment knowledge by having experts review performance data collected from the simulation's users. Experts are asked to do the following:

1. review and rank order user performance by overall quality;
2. assign scores that reflect overall quality; and
3. explain their scores in the form of critiques of performance.

We verify that experts are using a common set of policies in assessing performance by confirming that experts' scores of overall quality are correlated with each other. We then collect the expert knowledge in the form of critiques of other performers' work samples, and capture it in the form of assessment policies and rules. We verify that the assessment knowledge captured is correct by executing the scoring rules on new performance data and comparing the results to experts' judgments. The assessment system can be refined by discussing those work samples with experts on which scores from the rules and experts were far apart. These discussions identify how the policies incorporated in scoring rules diverge from experts' understanding of good and poor performance, which leads to rule refinement. The refinement can continue until the rules yield scores as close to experts' scores as desired.

While the scoring rules are tuned using existing work samples scored by experts, the goal is to apply the scoring rules so they accurately score new work samples. To verify the rules, new work samples are scored by the rules and by experts. Scores from the rules and experts are compared using correlations. If the scoring rules address performance deficiencies as experts do, the correlations between rules and experts should be high. These correlations reflect how well the scoring rules and experts' judgments correlate with work samples that will be collected in the future.

One limitation of PEER relative to ECD is that ECD has been used to assess many dimensions of performance, such as "creativity exhibited" or "persistence," while PEER has only been used to collect scores for overall quality. Overall quality holds a special place within the application of PEER because quality is the most natural dimension on which performance is judged. Earlier studies (Pokorny, et. al., 2010) reported that experts' judgments of overall quality were highly correlated with the experience level of the performers whose work samples were being judged. Further studies could investigate experts' ability to judge other features of performance besides overall quality of performance, such as creativity or persistence.

One reason frequently cited for the difficulty of assessing performance is the difficulty of scoring how each action informs an underlying cognitive variable, such as the knowledge that a student possesses. PEER scoring policies overcome this complexity by looking at patterns of actions rather than individual actions. Consider the following example from UrbanSim, the COIN simulation that we use in this work: a trainee acting as a virtual commander should order police to be recruited from both Shiite and Sunni districts. If the trainee recruits only in Shiite areas during the first turn of the simulation (in which the trainee issues a set of fragmentary orders), expert reviewers will not assert a problem. However, by the trainee's fifth turn, if the trainee has only recruited in Sunni regions, the expert will identify this pattern as a violation of good policy. In PEER, experts take a conservative approach to identifying problems in performance—they only identify a performance weakness when they are confident that their judgment of poor performance is clearly indicated by the student's pattern of performance.

SELECTING AND SEQUENCING INSTRUCTION

The Challenge of Selecting and Sequencing Instruction

Even after a training simulation has an assessment system and instructional content tied to that assessment system, questions remain about what is the most effective way to present the instructional content. When should an instructional intervention be presented? How should training content be presented, assuming multiple formats have been prepared? Frequently these questions are answered by pedagogical theory.

A Data-Driven Approach to Selecting and Sequencing Instruction

Our approach answers selection and sequencing of instruction by encoding the task in terms of Markov decision processes (MDPs), which can be solved using modular reinforcement learning techniques. MDPs provide a mathematical formalism that enables the use of historical data from student learning to plan the best time and format for presenting instructional content (Rowe & Lester, 2015). Planning tutorial actions is a critical component of ITSs. Tutorial planners determine how pedagogical tasks and scaffolding are tailored to learners at run-time. Devising tutorial planners that generalize across students, learning environments, and domains is an important challenge for the field. These challenges are amplified by the increasing complexity of advanced learning technologies, such as simulations (Johnson 2010) and digital games (Shute, Ventura, & Kim, 2013; Rowe & Lester, 2015; Baker, Clarke-Midura, & Ocumpaugh, 2016), which feature ill-defined domains, open-ended problem scenarios, and vast decision spaces.

Data-driven methods for devising tutorial planners, such as reinforcement learning, show great promise for addressing these challenges due to their capacity to automatically induce pedagogical models from large datasets characterizing student behavior and learning outcomes. Further, they introduce the possibility of devising ITSs that automatically refine and improve their pedagogical strategies over time, leveraging data on student outcomes to update their tutorial policies.

Use of Markov Decision Processes

Our tutorial planning system is based upon a modular reinforcement learning framework, which encodes decisions about selecting and sequencing instructional content in terms of MDPs. MDPs consist of a set of states, actions, a probabilistic state transition model, and reward model. For each MDP, state consists of data about the learner and the learner's history, as well as the learning environment; actions represent the pedagogical decisions the planner can perform; the probabilistic state transition model encodes how learners and the learning environment respond to the planner's tutorial decisions; and a reward model encapsulates measures of trainees' learning outcomes, which the tutorial planner seeks to optimize. The solution to a modular reinforcement-learning problem is a set of policies, or mappings between states and tutorial actions, that govern how the tutorial planner scaffolds trainees' learning. Different types of tutorial decisions are modeled separately as MDPs. If two policies conflict, externally defined arbitration procedures specify which policy prevails.

The tutorial planner will support a broad range of tutorial interventions which share a generalized encoding of instructional strategies and tactics across multiple learning environments. By decomposing tutorial planning into multiple sub-problems, the complexity of reinforcement learning is reduced, both in terms of training data required, as well as training time. To perform this decomposition, we employ the concept of an adaptable event sequence (AES), an abstraction for a recurring series of one or more instructionally related events that, once triggered, can unfold in several different ways within the learning environment (Rowe and Lester, 2015).

For our current UrbanSim example, we use two nested levels of AESs. A set of high-level AESs will control decisions about which instructional technique to deploy at particular points of training, such as single-topic coaching, multi-topic summaries, feedback on unproductive learning behaviors, and non-interventions. A set of lower-level AESs will control decisions about which instructional tactic to utilize for presenting instructional content. Our instructional tactics are based upon Chi's ICAP framework (2009), in which instructional activities are conceptualized in terms of interactive tasks, constructive tasks, active learning tasks, and passive tasks. The performance content selected for remediation is the concept that the student scored worst on amongst those concepts being tracked by the assessment framework.

A FRAMEWORK IN WHICH TO INTEGRATE TUTORING CAPABILITIES

The Challenge and Need for a Tutoring Framework

One inspiration for developing the Generalized Intelligent Framework for Tutoring (GIFT) is the observation that there are many adaptive instructional capabilities that address different elements of tutoring, including many efforts and approaches to (1) assess performance; (2) assess student engagement; (3) present content; and (4) support metacognition. Developing an ITS that uses any set of capabilities requires building a whole system that integrates many capabilities that support student learning.

Producing ITSs would be streamlined by usage of a common software framework and tools. By reusing existing ITS components, researchers and developers can focus their efforts on improving specific aspects of ITSs, rather than creating customized ITSs in their entirety for each domain and learning environment. In this model, studies can be more easily conducted to investigate the impacts of individual components of ITSs, such as specific instructional strategies, contents, or assessments. Further, as new capabilities are developed, existing ITSs could more easily take advantage of these new features. The use of a common software framework would enable developers to take advantage of existing ITS functionalities.

An Approach to Constructing a Tutoring Framework

To build such a framework, the U.S. Army Research Laboratory (ARL) has hosted a series of meetings with ITS researchers, developers, and instructional designers. The intent was to develop the requirements for a general framework for intelligent tutoring. This led to the design of the Generalized Intelligent Framework for Tutoring (GIFT), which is the subject of iterative design, development, and refinement. ARL funds various research groups to use and elaborate the framework.

Using the GIFT Framework.

The GIFT framework is shown in Figure 1. As we develop the capabilities for this project, the set of assessment rules for COIN as applied in UrbanSim are conceptualized as domain-specific knowledge. The use of MDPs to select the pedagogical methods are part of a computational model for instructional strategies. As more capabilities are used, the instructional platform that uses GIFT for UrbanSim becomes more powerful. Further, instruction treatments can be tested. Components and their communication are shown in Figure 1.

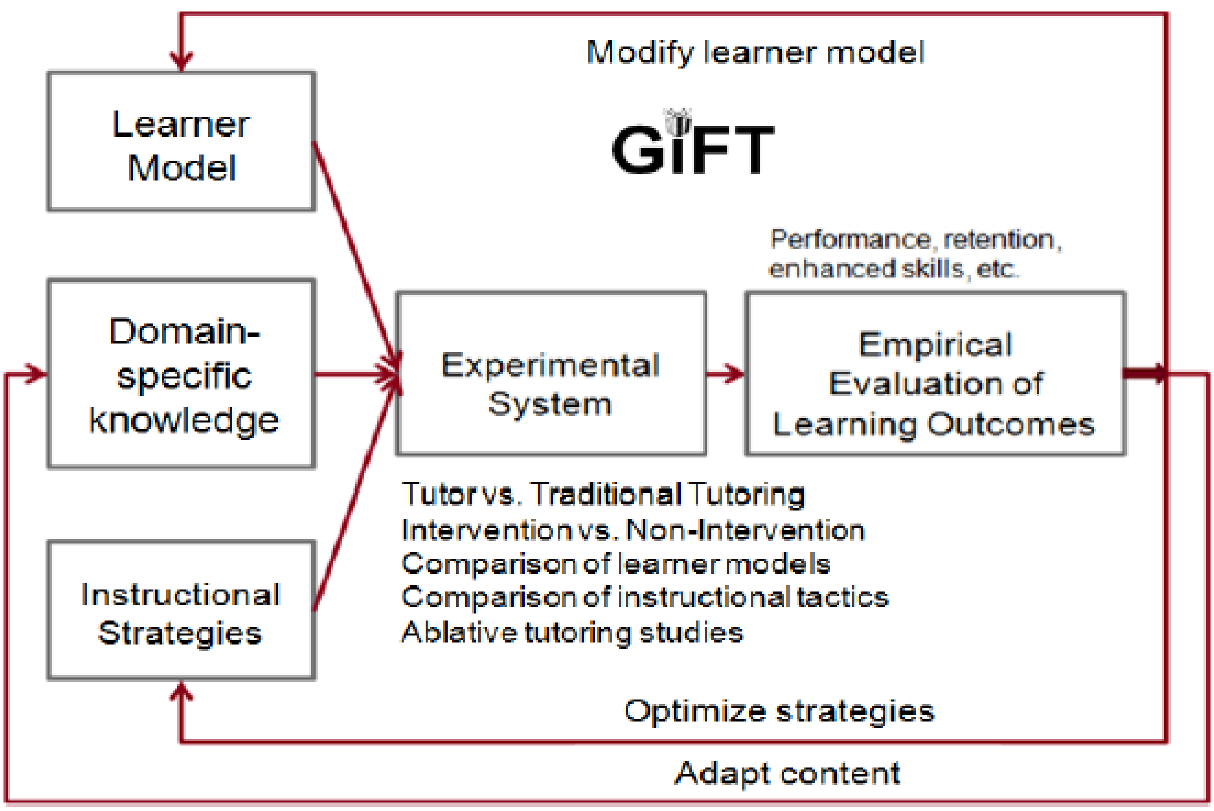

Figure 1. The GIFT Framework

RESULTS FROM CREATING ASSESSMENT AND INSTRUCTION

The integration of (1) PEER to create an assessment system that identifies strengths and weaknesses of learners and links these to instructional content; (2) MDPs that guide selection of instructional strategies and tactics, and (3) GIFT as an intelligent tutoring framework is a general solution to developing ITS across multiple domains and learning environments. The MDP approach and the GIFT framework can be the utilized across ITSs for different domains.

What will change between the tutors will be (1) the environment simulated, (2) the assessment system (scoring rules) and instructional content devised, and (3) the training data utilized to induce pedagogical policies for the MDPs. Next we will report results of applying PEER to inform assessment and instructional systems across many domains.

PEER Applied to Complex Equipment Troubleshooting

PEER was originally developed for assessing performance of technicians troubleshooting faults in complex equipment. This application of PEER has the most data and has been investigated in depth for statistical properties. We first developed scoring rules on one set of data. Then we applied it to another set of data. The correlations between the expert scores of overall quality and scores from the scoring rules was always significant, and ranged from .59 to .88 (Pokorny, Gott, Haynes, 2010). So the rules provided scores consistent with experts' policies.

The scoring system identified divisions of the equipment, the sequence that the equipment sections should be investigated in, and methods that could be applied to the active path of the circuit in each equipment section. This decomposition guided both the assessment system and the instructional content.

PEER Applied to Air-to-Air Dogfight

PEER was applied to fighter pilot performance. We collected performance of pilots attacking an enemy plane in a one on one dogfight. In a dogfight, time is critical. Experts' critiques reflected the time critical nature of the domain: when a pilot did not respond quickly enough to an observed situation that demanded action, experts criticized the performance.

Experts also criticized where pilots looked. While the replay from pilot performance included many forms of data, the expert reviewers chose to look at the scene from the pilot's point of view through the cockpit window. The replay of the dogfight that experts reviewed showed (1) the view out the cockpit window; (2) the cockpit console and instrument; and (3) the direction in which pilots gazed. An eye tracking system collected where the pilot was looking in the simulation; the replay showed this direction by superimposing a red dot where the pilot was looking.

The expert reviewers identified multiple policies. These contain trigger events which signify changes in combat state. Different rules are in effect for specified periods of time after a trigger event. The assessment system could easily guide instruction for pilots.

PEER Applied to Army Mission Command

PEER is currently being applied to Army Mission Command. Mission Command is a process in which higher level commanders do not specify a specific target and objective (e.g., Attack and Take Hill #3) but present the outcome of the mission that they want, and then allow lower level commanders to select the method by which to achieve the specified end state. The higher level commander tracks the performance for the lower level teams, and intervenes when necessary. We had experts review the commands issued, as well as the communications and decisions made by higher level commanders as they communicated with and observed the progress of lower level units (e.g., Companies). Experts again had little trouble generating critiques of Commanders as they executed Mission Command in a complex, simulated environment.

As this project continues, a major area of interest will be processing chat communication during the assessment of trainee performance. If communication can be understood in real time to classify messages into categories that experts identified as critical for scoring performance and identifying strengths and weaknesses of trainees, then building an ITS for Mission Command will be straightforward. At this point, we cannot say how easily language interpretation will handle the communication used to convey mission command.

PEER Applied to Team Assessment

PEER was applied to assess a squad working on a scenario to clear a road of IEDs. The scenario was conducted in VBS2, a common small unit training platform. For this training, teams of 7-8 Soldiers were briefed on a Road Clearing task that they would execute in VBS2. Each member of the team was assigned to drive a vehicle on the team. Additionally, one member of the team was the Commander, who rode in one vehicle. The training took about 3 hours; the team drove along a road for about 2.5 hours, and encountered four IEDs. After each IED, the instructor reviewed what the team did well and what they did poorly. Interestingly, the instructor easily identified the member(s) of the team responsible for the poor performance. Sometimes the responsibility was one individual. For example, the player whose job it was to detonate the IED might have made a mistake. Other times, the responsibility was shared between the Commander and an individual, such as when one vehicle drove too close or too far from the vehicle in front of it. Other times, multiple Soldiers and the Commander were responsible for poor performance, such as when a group of Soldiers formed a cordon incorrectly. While the instructor did not use the term, the instructor sometimes identified taskwork as a weakness, and sometimes identified teamwork as a weakness (Kluge, 2014).

Based on the discussion of the instructor with the trainees, we created a set of scoring rules that identified performance strengths and weaknesses. We discussed the scoring rules we developed with the instructor later, and defined a set of rules that could be used to automatically score trainee performance. While this study did not complete the PEER process, the ease with which the rules were created and their simplicity suggests that this approach is worth investigating as a means to assess team performance.

Creating Instruction from Critiques

The critiques of experts that identified performance weaknesses lead directly to content areas in which performance should improve. For example, if trainees' security actions were too lax, the target performance should lead trainees to use more security actions, which would approximately equal the security actions of an expert. The instructional target would aim to have the trainee think similarly as the expert in regards to how many resources should be assigned to security. Given this goal, there are many different instructional strategies and tactics that might be applied. One strategy would explain how the expert thinks about security actions. A second instructional strategy would use cognitive scaffolding: the trainee's attention would be brought to a particular target, and the training system would provide hints by which the trainee constructs an improved model of security actions that more closely resembles an expert's.

A structure associating a performance weakness with instruction to support adaptive training has been suggested by Lesgold and Nahemow (2001). They present their view of a Process Observer. It describes conditions that imply specified instruction, regardless of when the condition occurs. For each identified performance weakness, content is available for presentation and interaction with the student.

INSTRUCTION SYSTEM FOR URBANSIM

When applying PEER to trainees' performance from UrbanSim, data from trainees was formatted so experts could easily interpret their actions and the current state of the environment. Then experts judged the overall quality of trainees' execution of the counter-insurgency mission, assigning each trainee a score reflecting the quality of his/her performance on a scenario. The experts then critiqued the performance; in a way, the critiques justified their scores. The task analyst used the critiques and scores to created scoring rules that awarded points for complying with good performance quality and deducted points for violating good performance quality.

With data collected from trainees at Ft. Leonard Wood, US Army experts agreed on judgments of overall quality of COIN mission execution. The correlations of three Army experts were generally in the .70s; all correlations between Army experts were statistically significant.

Task analysts divided the critiques into six classes. These are (1) security; (2) meeting with host nation leaders; (3) supporting host nation governance, (4) information operations; (5) infrastructure; and (6) consistency of mission. Most of these classes had subclasses as well. For example, one sub-rule for security was "Battalion Commanders' actions were too lax." The score associated with "too lax" or any other action depended on the degree to which the violation of good policy was exhibited. For example, if there are three active insurgent groups, and the Commander ordered only one security action, the violation score for that instance would be much higher than the violation score if there were two active insurgent groups and the student ordered three security actions.

The Challenge of Creating Instruction for UrbanSim

We selected UrbanSim as a target for adaptive instruction because (1) it addressed a topic that was seen as very difficult: COIN is an ill-structured domain, and most ITS research has focused on well-structured domains, such as math or physics. Second, UrbanSim was an important topic for Army training. It was initially developed for Battalion commanders to train for leadership in Iraq. Since then, UrbanSim has been used for lower ranking Soldiers to learn about COIN. UrbanSim is currently one of the training games used in the Army Games for Training suite.

An Approach to Developing Instruction for UrbanSim

We used PEER to develop an assessment system. The set of performance weaknesses led to content that could be used for instruction. We are developing MDPs to select training interventions. ARL is funding other researchers to create other elements of training that can be integrated with UrbanSim.

Current State of the UrbanSim Tutoring

To test the viability of using PEER-based assessment, we use MDPs to model decisions about instructional sequencing within GIFT, and we will conduct tests of the instructional effectiveness of the integrated system. Our initial studies will include two of the six dimensions of performance for assessing performance in UrbanSim. Modular reinforcement learning techniques will be used to induce policies for selecting instructional strategies and instructional tactics. Studies with students from many colleges ROTC units are currently being planned for Spring 2017. If these studies yield evidence of effective instruction, these studies will demonstrate that complex performance can be assessed by PEER-based domain analysis, and scaffolded using MDP processes to sequence instruction that is presented within GIFT.

Recommendations for Future Research and Application

After the currently planned training studies, a next research step should include integrating other features of GIFT functionality into the instruction with UrbanSim. In particular, other capabilities being built into GIFT include assessment of student affect, assessment and instruction of metacognitive abilities, and use of Natural Language Processing to interact with students. Each of the capabilities could be integrated to show how GIFT can support a variety of methods that support student learning.

Further field applications of the integration we have completed should provide training to soldiers using the combination of PEER-based assessment and instructional content, as well as MDP instructional sequencing and the GIFT framework. These applications should be fielded to students and evaluated for their effectiveness.

ACKNOWLEDGEMENTS

We thank organizations and individuals who have contributed to this research. The Advanced Distributed Learning Laboratory funded research to apply the PEER method to UrbanSim. The Army School of Engineering provided trainees and instructors who supported the assessment and training content for UrbanSim and for the team task of Clearing a Road of IEDs. North Carolina State's University ROTC post provided trainees who participated in pilot studies in which they learned COIN from UrbanSim. This research is sponsored by the U.S. Army Research Laboratory (ARL) under contract number W911NF-15-2-0030; the content of this paper does not necessarily reflect the position or the policy of the Government, and no official endorsement should be inferred.

REFERENCES

Baker, R. S., Clarke‐Midura, J., & Ocumpaugh, J. (2016). Towards general models of effective science inquiry in virtual performance assessments. *Journal of Computer Assisted Learning*, *32*(3), 267-280

Chi, Michelene T. H. (2009) Active-Constructive-Interactive: A Conceptual Framework for Differentiating Learning Activities. *Topics in Cognitive Science*. 1:73-105

DiCerbo, Kristen (2013). Stealth Assessment: Let's be Clear Where We Are. Downloaded from http://researchnetwork.pearson.com/digital-data-analytics-and-adaptive-learning/stealth-assessment, May 31, 2016.

Feng. M, Hansen, E., and Zapata-Rivera, D, (2009) Using Evidence-Centered Design for Learning (ECDL) to Examine the ASSISTments Systems. Proceedings of American Educational Research Association.

Honey, M & Hilton, M. (2011) *Learning Science Through Computer Games and Simulations*, National Academies Press. Washington, District of Columbia.

Johnson, W. L. (2010). Using immersive simulations to develop intercultural competence. In *Culture and computing* (pp. 1-15). Springer Berlin Heidelberg.

Kluge, A. (2014) *The Acquisition of Knowledge and Skills for Taskwork and Teamwork to Control Complex Technical Systems: A Cognitive and Macroergonomics Perspective*, Dordrecht, Germany: Springer.

Lesgold, A. & Nahemow, M. (2001). Tools to assist learning by doing: Achieving and assessing efficient technology for learning. In D. Klahr & S. Carver (Eds.), *Cognition and instruction: Twenty-five years of progress*. Mahwah, NJ: Erlbaum

Mislevy, R. J., & Riconscente, M. M. (2006). Evidence-centered assessment design. *Handbook of test development, 61-90*

Pokorny, R., Haynes, J., Gott, S. (2010). Performance Assessment in Complex Simulations. *In Proceedings from Interservice/Industry Training, Simulation and Education Conference*, Orlando, Fl.

Rowe, J. P., & Lester, J. C. (2015, June). Improving student problem solving in narrative-centered learning environments: A modular reinforcement learning framework. In *International Conference on Artificial Intelligence in Education* (pp. 419-428). Springer International Publishing.

Shute, V, (2011) Game, Learning, and Stealth Assessment, Downloaded from http://www.slideshare.net/g4li/what-is-stealth-assessment, May 28.2016.

Shute, V. J., Ventura, M., & Kim, Y. J. (2013). Assessment and learning of qualitative physics in newton's playground. *The Journal of Educational Research*, *106*(6), 423-430

Wang, L., Shute, V., & Moore, G. (2015). Lessons learned and best practices of stealth assessments. *International Journal of Gaming and Computer Mediated Simulations, 74*(4), 66-87

Woolf, B. P. (2010). *Building intelligent interactive tutors: Student-centered strategies for revolutionizing e-learning*. Morgan Kaufmann