CONSIDERATIONS FOR LEARNER MODELING IN A DOMAIN-INDEPENDENT INTELLIGENT TUTORING SYSTEM FRAMEWORK

Anne M. Sinatra[(a)], Keith W. Brawner[(b)] , Gregory A. Goodwin[(c)]

[(a-c)] U.S. Army Research Laboratory

[(a)]anne.m.sinatra.civ@mail.mil, [(b)]keith.w.brawner.civ@mail.mil, [(c)]gregory.a.goodwin6.civ@mail.mil

ABSTRACT

This paper focuses on the learner modeling component of intelligent tutoring systems (ITSs). The majority of ITSs are domain-dependent with the domain content being closely tied to both knowledge about the learner, and the pedagogical strategies. However, the Generalized Intelligent Framework for Tutoring (GIFT) is a domain-independent ITS framework. This domain-independent advantage allows for significant reuse (instructional and learner models), reducing the amount of time it takes to generate ITSs. It also creates interesting challenges and considerations that need to be taken into account when determining what elements need to be included in the various ITS modules. While GIFT currently includes a learner module component, additional research is expected to be conducted to determine the ideal components to include in GIFT's learner module. The current paper discusses the unique challenges of developing domain-independent learner models, as well as concerns related to implementing and authoring.

Keywords: adaptive and predictive computer-based tutoring, adaptive training, learner modeling, intelligent tutoring systems

1. INTRODUCTION

Intelligent tutoring systems (ITSs) provide personalized and adaptive tutoring to individual learners. While classroom based learning has been the norm for many years, it is often difficult for students to get personalized attention, due to large class sizes. ITSs can be used either as a supplement to classroom learning, or as a primary means of learning. Allowing a student to engage with ITSs independently leads to a higher reliance on self-regulated learning, where individuals manage the pacing of their own learning. There are a number of strategies that lend themselves to successful self-regulation of learning, however, not all students spontaneously engage in them (Zimmerman, 1990). Therefore, one goal of ITSs are to provide information to students in such a way that they will be engaged with the system and exhibit patterns of interaction that will lead to long term learning gains. In order to lead to engagement, it may be useful to customize materials to the specific learner's characteristics, experiences, or current mood. In order to customize instruction it is important for the ITS to have a representation of the learner's state, which includes affective state (e.g., mood), as well as cognitive, and procedural assessments of the learner that are relevant to the domain area of interest (Pavlik, Brawner, Olney, and Mitrovic, 2013; Woolf, 2010).

1.1 Traditional ITS Components

ITSs traditionally have four software modules: the learner module, pedagogical module, domain module and tutor-user interface (Sottilare, Graesser, Hu, and Holden, 2013). The pedagogical module is responsible for the instructional strategies that are provided to the individual learner. The domain module is specific to the domain information (content, lessons, subject matter, etc.) being tutored. Naturally, the tutor-user interface is the way that the learner interacts with the system. The learner module is the software process where all the information about the individual learner (learner model) is stored and processed. It represents the previous knowledge about the learner, the current knowledge of the learner's state, and is traditionally updated throughout the learner's time engaging with the ITS.

In current terminology, a software "module" refers to an executable piece of software, running as a part of a total system. A "model" refers to the data and processes which run inside the module. As an example, a simplistic learner model may blindly communicate underperformance whenever it is made aware of it; this underperformance information would, of course, be communicated by a software module to other modules containing their own models of instructions.

1.2 The Generalized Intelligent Framework for Tutoring

The Generalized Intelligent Framework for Tutoring (GIFT) is a domain-independent framework that provides individuals with the ability to create ITSs, deliver training, and analyze data. The modules present in GIFT are similar to those of traditional ITSs; GIFT has a learner module, domain module, pedagogical module, sensor module, and a tutor-user interface (Sottilare, Brawner, Goldberg, and Holden, 2012). The addition of the sensor module provides assistance in measuring learner variables, and assists in updating the current state of the individual learner model, which is in the learner module (Sottilare, 2012). The development of GIFT's architecture is consistent with the Individual Learning Effect Model (Sottilare, 2012). Figure 1 shows the most recent version of this model. The learner model contains individual learner measures and states, which influence the instructional strategy that an individual receives. This strategy becomes an "instructional tactic" when it is implemented within a domain of instruction, the distinction being the addition of actual learning content. After the instructional intervention is received the performance of the individual learner updates the learner data and the state of the individual. This information will then once again influence the instruction that the individual receives. Therefore, the information about the learner's state drives pedagogy, which then drives the learner's state.

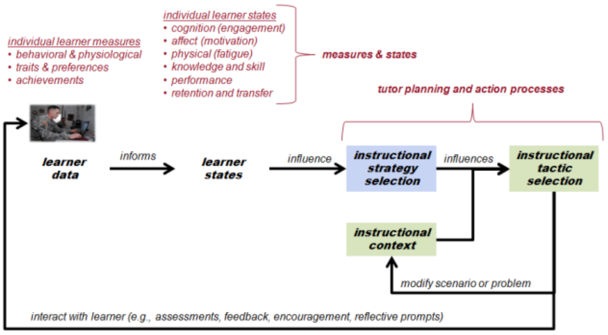

Figure 1: Individual Learning Effect Model

While the GIFT architecture contains these components, the framework and software itself are still under active development. The functionality that is currently in GIFT will expand as GIFT continues to be developed and improved. As GIFT is an open-source project, the needs of the user base can influence the directions taken for future design decisions.

The current paper discusses the unique considerations in the development of a domain-independent learner model, highlights the supporting technology that can interface with the learner module, and discuss approaches to authoring. Additionally, the paper discusses future directions that can be taken in adding additional functionality to GIFT's learner module.

2. CONSIDERATIONS FOR DOMAIN-INDEPENDENT LEARNER MODELING IN GIFT

Traditionally in ITSs the learner module contains information about the individual learner who interacts with the system. While learner models that are used in different ITSs have much in common, they are generally unique to the development of the individual systems. Most ITSs are domain-specific, and are further designed to teach one specific well-defined domain, resulting in the majority of ITSs being for instruction in mathematics and physics. This is partially a function of demand, and partially a function of the relative difficulty of developing an ITS for well-defined as opposed to ill-defined problems (Sottilare and Holden, 2013). For these domains, the author of the ITS decides what learner assessments are useful and relevant to the instructional domain when creating the learner model. In the case of mathematics this may include prior math courses taken, grade point average, and scores on relevant standardized tests. It might also include factors like the learner's intelligence, metacognitive skills, conscientiousness, and grit. As can be seen, some of these attributes are domain dependent (e.g., prior relevant coursework, previous performance) and others are domain independent (e.g., intelligence)

Using GIFT's domain-independent framework has both challenges and advantages. The advantages include the reusability of the authoring tools, content, and portions of the created ITS. Additionally, questions and assessments that are authored in the domain-independent ITS can be edited and used as a foundation of other courses. These advantages also lead to a savings of money, as new systems do not need to be developed for every type of ITS course that is generated. When developing a domain-independent architecture such as GIFT, a challenge is to create a learner model flexible enough to work with any domain. One challenge is in determining ways in which domain independent learner attributes can be used to adapt training across domains. For example, learner intelligence may be used to adjust the difficulty or pace of the training. Another challenge is determining ways in which assessments of the learner's past experiences, training, assignments, goals, and interests can be collected from existing data sources such as personnel, training, and learning management systems. For example, college transcript grades could potentially be utilized to help determine the difficulty level of current training that should be assigned. Currently, GIFT is limited to collecting this kind of information about the learner at the beginning of a course. However, it is not practical to collect this amount of information from learners each time they start a new course. Automating the process of collecting that information when a learner begins training will both improve the learner's experience by reducing or eliminating lengthy pre-training surveys and will facilitate the development of predictive models of learner performance and training effectiveness.

2.1. Domain-Independency and Time of Assessment

There are two main types of assessments that are commonly supported by learner models: pre-training assessments, and in-training assessments. Further, these types of assessments can be further sub-divided into domain-dependent and domain-independent categories. The following sections describe current features of GIFT, and ways that these assessments can be accomplished using GIFT.

2.1.1. States vs. Traits

Psychologists generally distinguish between states and traits. While not strictly the same thing, ITSs distinguish between long-term and short-term learner models, and these terms have previously been used in the context of ITSs (Pavlik, Brawner, Olney, and Mitrovic, 2013). States are short-term and specific to how the individual is currently feeling or performing (i.e., the short-term learner model).Traits are associated with longer-term characteristics, such as an individual's personality scores (i.e., the long-term learner model). While there may be some variation in the mood of an individual that will fluctuate from hour to hour or day to day, an individual's overall level of neuroticism is not expected to shift dramatically in a short period of time. In the context of adaptive tutoring, competencies and aptitudes are relatively stable "traits" that are generally assessed pre-training, and are used to make decisions about the type of material the learner will receive. Further, "state" measures, such as the learner's current mood, or current performance have an influence over in-training materials that are provided and update the learner model while the learner is actively engaged with the tutor. Figure 2 is a proposed assessment framework for a learner model, which accounts for both domain-independent, and domain-dependent components, and serves as a basis for our discussion of learner modeling.

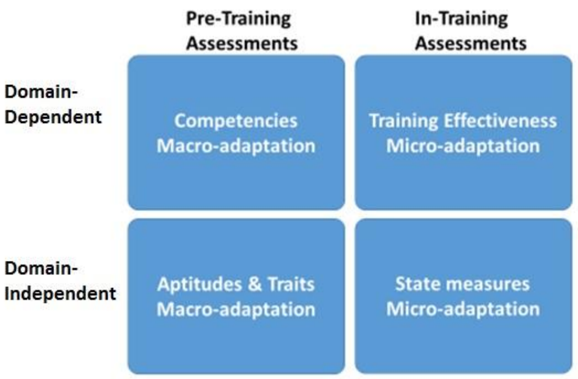

Figure 2: Assessment framework for a learner model

2.1.2. Pre-Training Assessments: Questionnaires

In an ideal training environment, there are stable characteristics of a learner that might influence the type of training that the learner is given. In the ideal learner model all of the learner's relevant scores, questionnaire data, and survey data would be stored. This information could be gathered from existing records, or if a relevant score is missing the learner could be prompted to take a survey or provide information that would update their learner model. Further, rather than having the learner repeatedly take a survey like a personality test or a working memory assessment, which is time-consuming, the output scores should be stored for later retrieval for the learner model. This information could then be referenced in the future for making of pedagogical decisions between different instructional strategies. In current implementations of GIFT, survey outputs and scores are associated specifically with courses that have been developed, as opposed to being associated with the individual learner, but it would take little overall work to retrieve them for future use in new courses.

Questionnaires are very useful in measuring information that could be of interest in instruction (e.g., reading level, working memory, special ability). These scores are unlikely to change and by storing them it could improve the ITSs ability to adapt to the individual. Further, by storing interest preference based information that could be gathered in surveys it would allow for additional opportunities to customize instruction to the individual. If more detail is stored about the individual learner it can lead to more accurate selection of relevant instructional strategies, and better learning outcomes.

2.1.3. In-Training Assessments: Sensors and Questionnaires

There various ways that state of the learner can be assessed during training in order to make adjustments to the materials the learner is receiving. In GIFT, the current state or performance of the individual is received from the learner, and then used to update the learner model. Sensors are primarily a state based measure, which measure real-time information such as attention that can change relatively quickly. Input from the sensors in GIFT is provided to the learner module, which updates the current state of the individual in the model it contains. There are a number of different sensors currently integrated with GIFT, including the Microsoft Kinect (which can examine movement and provides camera tracking) and the Q-sensor (which measures skin conductance, a proxy for anxiety). The information provided from these sensors can help to determine the individual learner's current affective state (Paquette, et al., 2015). While sensors are a passive way of determining state, which in general are not disruptive to the flow of a tutoring session, they can sometimes be difficult to work with in real-time. Many considerations have to be put into place in order to provide state adjustment based on sensor data, and calibration of sensors may be difficult or unrealistic in distance-learning environments. An additional way of determining the current affective state of the learner is through direct user query, such as the Self-Assessment Manikin (SAM) survey (Bradley and Lang, 1994), and through asking individuals to rate their mood. This method is not without its drawbacks, as it can disrupt the flow of the tutoring and is reliant on the learner's own self-assessment. However, surveys and questionnaires are a relatively stable way to get information about the learner's state that can be used for adaptations in learning material.

2.1.4. Domain-Dependent vs. Domain-Independent Assessments

In the case of both pre- and in-training assessments, there is domain-dependent and domain-independent information that is useful for adaptation. Traditional ITSs are tied to one specific domain, therefore, all of the data that they store in their learner model has been selected to cover relevant information for the specific context. However, in a domain-independent framework, it is necessary to include domain-independent measures that are relevant for a number of different domains, and provide the ability for authors to incorporate the measures that they wish to use for their assessments. For instance, reading level is domain-independent, but is relevant for numerous domains including: reading comprehension, math, physics, and computer programming. Another example is spatial ability, which is highly relevant for a number of domains including navigation, mental rotation, and drawing. However, in the case of domains such as math, spatial ability may not be as relevant and the learner model should provide the author the flexibility to select the elements that are relevant to the domain of interest.

In the current default state, GIFT's learner model tracks state data such as anxiety, boredom, confusion, and surprise. Additionally, it can adjust based on trait data such as locus of control, learning style, self-efficacy, grit, and goal orientation. However, the learner model is flexible and can allow for authors and researchers to follow and adjust based on specific states of interest in their specific domain. This is advantageous as it allows for customizable tutoring, as well as the ability to conduct experiments that examine what learner model elements are relevant in the domain of interest.

2.1.5 Competency Measurement

A major challenge for any ITS is the development of learner competencies. Competency is the set of knowledge, skills, and abilities that comprise competence in a specific job or role. Competencies develop over weeks months or years. To date, learner models have not typically incorporated competencies into their frameworks. Competencies are generally domain-dependent, and can be used during pre-assessment to adapt the materials that the learner receives.

2.2. Performance

The performance of the individual as they engage with the ITS is also an important component of the learner model. This performance is generally domain-dependent, and in the case of GIFT, while the learner model remains domain-independent, the domain model provides the information to the system to make sense of the learner performance in context. There are a number of different strategies that can be implemented to establish the learner's current domain-dependent performance state, as compared to the ideal learner state. Among these strategies are using overlay student models (e.g., rule space student models, model tracing student models), knowledge space models, and dialogue student models (Pavlik, Brawner, Olney, and Mitrovic, 2013). The learner's established state influences the pedagogical and instructional strategies that are selected for them. Once they have engaged in their interactions with the system their performance can inform the learner model and update their knowledge state. Therefore, the performance of the individual within the tutoring environment is clearly an important component of learning modeling.

2.3. Supporting Technology

There are a number of supporting technologies in order to capture, store, and utilize learner information. The basic types of information which are interchanged by GIFT are information of performance, captured with system assessments, as well as state and trait information, captured via sensors and surveys. This information is stored in log files for processing, and, depending on its nature, reported out to various external systems. A brief synopsis is included below.

2.3.1 Assessment

GIFT includes features which allow for the assessment of learner mastery of individual concepts for instruction. All of GIFT's modules are domain-independent except for the domain module. For instance, the learner module includes general statements about the learner, and the pedagogical module includes general instructional strategies. In GIFT the link is made between these general statements and the domain-specific content by the course author. The first of these links is the Domain Knowledge File (DKF). In the DKF, the course author provides assessment logic which defines messages in terms of the domain, and allows the Domain Module to receive messages which are passed to it using pre-coded interfaces. For instance, the pedagogical module may indicate that the strategy of "provide feedback 1" should be used. The DKF is where the author will make the connection that "provide feedback 1" should say "Make sure that you are using the correct order of operations: Add, Subtract, Multiply, Divide." The DKF provides the link between the general and the specific in GIFT's domain-independent framework. Further, the DKF defines concepts and the level of achievement the learner currently has based on performance (above expectation, at expectation, below expectation). Specific feedback and actions can be taken as a result of changing to different performance states. Examples of this assessment are demonstrated in courses included with GIFT software releases that interface with PowerPoint (assessing dwell time on slides), and in VBS2 (assessing individual markers).

The second manner in which performance is assessed is through the use of an external assessment engine. In this instance, the DKF simply indicates that messages of a certain type should be forwarded to another engine for assessment. An example of the use of this type of external engine can be seen in the SIMILE engine, used for vMedic training, and its authoring tool which is included with the GIFT software.

2.3.2 Log files and Databases

For the purposes of experimentation and evaluation, GIFT records nearly every transaction as part of its log files. Given that each of these messages is effectively stored for long term analysis, there is a tool for the extraction of specific elements within this sea of data. The Event Reporting Tool (ERT) allows for the extraction of key items of interest within a training scenario.

In addition to this storage, there is an amount of information which is stored within a database (mySQL, or Derby, depending on configuration) for later use. Examples of information which is stored include learner trait information, survey entry values, and "scored" information from interactions within an environment. The latter portion of this data is used in a simple course recommendation engine, which recommends courses based upon unsatisfactory completion.

2.3.3 Learning Record Stores and xAPI

Information which is stored within this database is additionally communicated externally to a Learning Record Store, using the xAPI encoding. xAPI is a manner of encoding learner "experiences" for the provision to other systems in the "subject verb noun" fashion. An example of an English xAPI statement would be that "John Mastered Italian", each of these objects has supporting field information (email address for John, wordnet definition of mastered, competency ontological link to Italian). Further information on xAPI statements can be found at http://www.adlnet.gov/expapi/.

xAPI statements require a storage location, which is the key feature that allows various systems to output xAPI information into a central repository which can be read from and written to. Because of the standardization of the xAPI statements, a wide variety of Learning Record Store (LRS) systems have been created, many of which have built from the Advanced Distributed Learning (ADL) open source reference implementation. Further information on xAPI statements can be found at http://tincanapi.com/learning-record-store/, which also provides a freely available and hosted LRS implementation.

GIFT makes use of both of these technologies in simplistic fashion. It rephrases its traditional score reporting to be compatible with the xAPI standard. It redirects its learner information to an LRS instead of a simplistic database. Each of these technologies has the potential to read and write much more data than is currently being broadcast.

3. AUTHORING CONSIDERATIONS IN A DOMAIN-INDEPENDENT INTELLIGENT TUTORING SYSTEM FRAMEWORK

Creating authoring tools that are domain-independent is a unique challenge. The tool needs to be easy to use, but also general enough that it can be used by individuals of varying skill levels and experience. GIFT includes a number of different authoring tools that have been structured to allow for the authoring of courses, and models that are not specifically tied to any one domain. Surveys and questionnaires are authored in the Survey Authoring System (SAS), which provides many of the features that are standard in survey creation utilities. These surveys are domain-specific and can be associated with specific courses by the author using the GIFT Authoring Tool (GAT). GIFT provides editing capabilities for the different models including the learner model, sensor model, and pedagogical model. The course author can include the components of the learner model that he or she feels is relevant for their specific ITS. Additionally, if the author would like to conduct research into which elements are relevant in their domain they can quickly swap out the tracked/adjusted for elements with little work. The DKF authoring tool allows for the linkage of general material to the specific domain. The author defines the different concepts that will be monitored, and then associates different types of feedback with transitions that happen based on performance. Once all of the modules are configured and all materials have been gathered, the author creates their overall course flow which references domain-specific materials and training applications (e.g., VBS2, PowerPoint). See Figure 3 for a screenshot of the GAT in GIFT 2015-1.

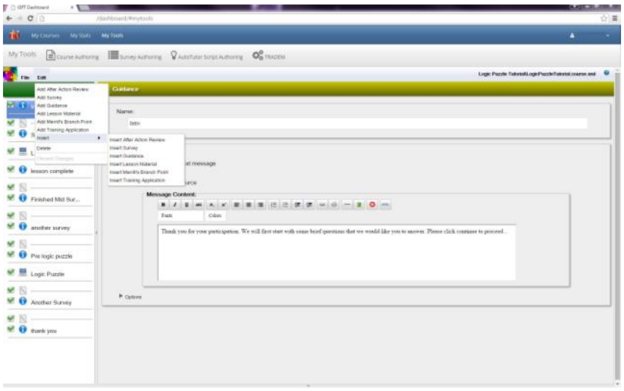

Figure 3: Interface for the GIFT Authoring Tool. After configuration has been done, this authoring tool combines both the domain-independent and domain-dependent components into one course flow.

4. CONCLUSIONS

The learner module is a traditional component of ITSs. However, it is also traditionally very tightly coupled with the domain being tutored. This provides challenges in ITS development in a domain-independent intelligent tutoring framework such as GIFT. However, it also provides the ability to separate the pedagogy, learner state, and the domain of interest. By doing so it allows more flexibility in the development of the ITS system, and promotes reusability. Further, it provides a way for researchers to compare the relative benefit of including specific elements in their learner models and conducting experiments.

Future directions of research intend to further expand GIFT's learner modeling into the area of teams and assess what elements should be examined on the team level, individual level, and at both levels. Additionally, as GIFT's domain-independent framework continues to mature it would be relevant to assess how much time it takes to author learner models and associated domain-dependent materials. This line of research could further refine and inform the process that is used in GIFT to create learner models. Learner models are an important part of any ITS. It is important to consider the benefits and challenges of developing domain-independent learner models which are reusable, interoperable, and allow for easy editing.

REFERENCES

Bradley, M.M, and Lang, P.J., 1994. Measuring emotion: The self-assessment manikin and the semantic differential. Journal of Behavior Therapy and Experimental Psychiatry, 25 (1), 49 – 59.

Paquette, L., Rowe, J., Baker, R., Mott, B., DeFalco, J., Brawner, K.W., Sottilare, R., & Georgoulas, V., 2015. Sensor-free or sensor-full: A comparison of data modalities in multi-channel affect detection, Proceedings of the Educational Data Mining Conference. Madrid, Spain, June 2015.

Pavlik, P.I., Jr., Brawner, K.W., Olney, A., and Mitrovic, A, 2013. A review of student models used in intelligent tutoring systems. In: Sottilare, R., Graesser, A.C., Hu, X., and Holden, H., eds. Design Recommendations for Intelligent Tutoring Systems, Volume 1, Learner Modeling. Orlando, FL: Army Research Laboratory, pp. 39 - 67.

Sottilare, R., 2012. Considerations in the development of an ontology for a Generalized Intelligent Framework for Tutoring. International Defense & Security Simulation Workshop in Proceedings of the I3M Conference, pp. 19 – 25. September 19 – 21, 2012, Vienna, Austria.

Sottilare, R., Brawner, K.W., Goldberg, B.S., and Holden, H.K, 2012. The Generalized Intelligent Framework for Tutoring (GIFT). Orlando, FL: US Army Research Laboratory – Human Research & Engineering Directorate.

Sottilare, R., Graesser, A.C., Hu, X., and Holden, H. Preface. In: Sottilare, R., Graesser, A.C., Hu, X., and Holden, H., eds. Design Recommendations for Intelligent Tutoring Systems, Volume 1, Learner Modeling. Orlando, FL: Army Research Laboratory, pp. i – xiii.

Sottilare, R., and Holden, H.K., 2013. Motivations for a generalized intelligent framework for tutoring (GIFT) for authoring, instruction and analysis. In: Artificial Intelligent in Education (AIED) Workshop Proceedings, Memphis, TN, July 9 – 13, 2013, pp. 1 – 9.

Woolf, B.P., 2010. Student Modeling. In: Nkambou, R., Bourdeau, J., and Mizoguchi, R. eds. Advances in Intelligent Tutoring Systems. Springer-Verlag Berlin Heidelberg, pp. 267 – 279.

Zimmerman, B.J., 1990. Self-regulated learning and academic achievement: An overview. Educational Psychologist, 25 (1), 3 – 17.

ACKNOWLEDGMENTS

The authors thank the Generalized Intelligent Framework for Tutoring and individual learner and team modeling vector teams including Dr. Robert Sottiliare, Dr. Benjamin Goldberg, and Dr. Joan Johnston for their input to this paper through numerous meetings and discussions.

AUTHOR BIOGRAPHIES

Anne M. Sinatra, Ph.D. is a researcher within the Learning in Intelligent Tutoring Environments (LITE) Lab at the U.S. Army Research Laboratory's Human Research and Engineering Directorate (ARL-HRED) in Orlando, FL. The focus of her research is in cognitive psychology, human factors psychology, and adaptive tutoring. She has interest in how information relating to the self and interest personalization can aid in memory, recall, and tutoring. She is currently a Research Psychologist and Adaptive Tutoring Scientist at ARL, and previously completed two years as an Oak Ridge Associated Universities/ARL Post Doctoral Fellow.

Keith W. Brawner, Ph.D. is a researcher within the Learning in Intelligent Tutoring Environments (LITE) Lab within the U. S. Army Research Laboratory's Human Research & Engineering Directorate (ARL-HRED) in Orlando, Florida. He has 9 years of experience within U.S. Army and Navy acquisition, development, and research agencies. He holds a Masters and PhD degree in Computer Engineering with a focus on Intelligent Systems and Machine Learning from the University of Central Florida. His current research is in machine learning, active and semi-supervised learning, ITS architectures and cognitive architectures. He manages research in adaptive training, semi/fully automated user tools for adaptive training content, and architectural programs towards next-generation training.

Gregory A. Goodwin, Ph.D. is a Senior Research Scientist at the Army Research Laboratory – Human Research and Engineering Directorate at the Paul Ray Smith Simulation and Training Technology Center in Orlando, FL. For the last decade, he has worked for the Army researching ways to improve training methods and technologies. He holds a Ph.D. in Psychology from Binghamton University and an M.A. in Psychology from Wake Forest University.