

Learner Modeling to Predict Real-Time Affect in Serious Games

James Lester, Bradford Mott, Jonathan Rowe and Jennifer Sabourin

Department of Computer Science
North Carolina State University

INTRODUCTION

In recent decades research on human emotion has revealed that affect plays a central role in human cognition. Historically, emotion was viewed as separate from cognition, and only recently has there been consensus that affect is a central component of rational behavior and social interaction (Forgas, Wyland, & Lahan, 2006). These findings have motivated the goal of designing affect-sensitive computer systems capable of recognizing and responding to users' emotional states. This has led to the development of affective systems with a wide variety of capabilities and purposes, such as therapeutic virtual agents, mood detection of large social network groups, and empathetic tutorial systems (Picard, 1997). Applications such as these require that three critical decisions must be made to inform how affect will be modeled and detected: 1) How will affect be represented? 2) What signals will be used for affect detection? 3) How will affect predictions be utilized?

Devising frameworks for representing affect has been the subject of considerable study in psychology and cognitive science (Ortony, Clore, & Collins, 1990; Russell, 2003). Psychologists have proposed a broad range of solutions for representing affect. These representations have left computational researchers with a variety of frameworks from which to choose. For example, affect may be viewed as a concrete categorical state, such as *anger* or *happiness*. Alternatively, affect can be represented among a variety of dimensions such as valence or arousal, where a state is only differentiated by its position on the scale. Further, affect can be seen not as a state itself but as an emergent phenomena consisting of action tendencies, physiological responses, and conditional precedents. Each of these representations introduces benefits and challenges when utilized in computational systems, and these tradeoffs must be considered within the contexts of specific applications. This chapter explores the role of computational models of affect in serious games, with a focus on real-time affect prediction.

UNREALIZED CAPABILITIES AND GAPS

Computational Models of Affect

Affect is fundamentally a hidden state. Detecting affect, when performed by either humans or computers, requires inference from observable signals. Selecting signals to incorporate into an affect-detection system requires consideration of several factors such as cost, invasiveness, and predictive value (Calvo & D'Mello, 2010). Common signals include physiological information such as heart rate and skin conductance, which can be measured through physiological sensors. Expressive traits such as facial expressions, posture, gaze, and gesture are other informative channels that can be measured with web cams or specialized sensors. Behavioral evidence of affective states is a complementary modality, and it is often inferred from logs of user interactions with software. Further, typed statements or spoken sentiment can be analyzed in a variety of meaningful ways to infer emotion. Each of these channels can be considered independently, but many successful approaches to modeling combine input from multiple modalities to arrive at a more complete and accurate representation of the states they are endeavoring to model.

A primary concern that often guides investigations of affect is the application the inferred knowledge about affective state will be used for. In some cases it is useful to simply know how the user is feeling as an evaluation of the system itself. In other cases this information may be incorporated into run-time interactions to dynamically improve user experiences or other important outcomes. Differences in objectives may lead researchers to consider multi-dimensional approaches, which are particularly beneficial if a single trait is being monitored over time. However categorical definitions are better for communicating with the user about their state. Similarly, if the application is multi-platform and intended to be used by individuals in a broad range of settings, some physiological sensors may be impractical. However, if the objective is inferring highly accurate representations of states, incorporating as many signals as available may be the best route. In many ways the final application of the affective information drives all other decisions.

Modeling Affect in Serious Games

An important application of affective modeling capabilities is in serious games. Serious games combine two affectively charged activities, learning and gameplay, making them an interesting avenue for exploring issues of affective modeling. Affect is a critical component of learning and has been shown to influence how students process information, approach learning activities, and feel about themselves and their abilities (Baker, D’Mello, Rodrigo, & Graesser, 2010; Pekrun, Goetz, Titz, & Perry, 2002; Picard, et al., 2004). Meanwhile, though enjoyment and happiness are viewed as fundamental components of play, games can introduce degrees of frustration, confusion, sadness, and even anger. Together the rich affective experiences associated with learning and gameplay motivate the need for exploration of affective modeling in serious games.

Serious games have considerable capacity to evoke a broad range of emotions, because they can contextualize learning and training activities within meaningful contexts that approximate authentic settings. For example, games often incorporate narrative plots, which in many ways are fundamentally emotion-inducing: interesting stories are defined by conflict, uncertainty, suspense, and surprise. Furthermore, they are often host to interactions with believable virtual characters, fantasy or simulated settings, and dynamically unfolding situations. These types of learning environments lend themselves to a broad range of emotional responses that are likely to shape learning processes, but may be atypical of alternate educational environments that deliberately separate educational content and application context.

For a range of educational settings, evoking these types of emotional processes during learning may be desirable, or even essential, to produce learners who will be capable of performing skills in realistic settings that will similarly evoke these emotions. However, creating educational environments that can evoke these types of emotions raises a number of issues: 1) how can we recognize students’ emotional states, including emotions that may not be commonly observed in non-game settings, such as fear, surprise, disorientation, anxiety, excitement, curiosity, sadness, so we that we can better understand which, and how, emotions impact student learning?, 2) how, and to what extent, can we effectively model students’ emotions of these types in real-time, including predictive models that converge on accurate predictions of student emotions in advance of their occurrence based on expected narrative states and affect-sensitive learner models?, and 3) how can we devise models that dynamically tailor events in serious games, including tutorial events, narrative events, and game parameters, to aid students in self-regulating their learning and affective processes?

Recent work in this area has explored many varied components of affect in serious games. For example, Conati *et al.* have explored probabilistic models for affect recognition in *Prime Climb*, a serious game for young children learning number factorization (Conati & Maclaren, 2009). These models incorporate in-

game context data along with physiological sensors to arrive at accurate predictions of student emotion. *FearNot!* is a serious game for teaching anti-bullying behavior to school children (Paiva, Dias, Sobral, & Aylett, 2004). This environment seeks to model the affective states of virtual characters to drive empathetic relationships with the learner. Robust emotion models have also been incorporated into the *BiLAT* serious game (Kim, et al., 2009). This environment seeks to teach military personnel skills for multicultural negotiation and uses affective information to drive character interactions and responses.

This work has highlighted the role that affect can play in the development and understanding of serious games and points to areas where future work is needed. First, much work in modeling user affect takes place in highly structured environments. In these systems users typically have a small set of available actions and have a clear indicator of correctness and progress. This simplifies the identification of features relevant to affect modeling. However, it will be important to begin exploring how affect can be modeled in more open-ended exploratory games. There may be different sets of affective phenomena at play in these environments and it is likely that a different set of tools and techniques will be needed to build successful predictive models in these systems. Another important area of future work is more exploration of how affect models can be incorporated back into serious games. While *FearNot!* and *BiLAT* utilize affective information to drive narrative and character interactions, there is less exploration on how this information can be used to reason about the learner or guide tutorial strategies. A deeper understanding of both sides of this affective picture is needed to fully describe the important role affect plays during interactions with serious games.

FUTURE CONCEPT

Given these gaps in the research on computational models of affect, we envision three major thrusts of research for real-time affect detection in serious games. First, we anticipate the potential for sizable advances in early prediction of student emotions during game-based learning. Second, we anticipate that data-driven models of student affect will be combined with theoretical frameworks of student learning and emotion to improve models for assessing student knowledge acquisition and transfer. Third, we anticipate that real-time models of student affect will be incorporated into tutorial systems capable of personalizing events in serious games to scaffold student learning and promote sustained engagement.

Predicting Student Affect in Guided Inquiry-Based Serious Games

Modeling student affect in serious games poses distinct challenges relative to other educational systems. These challenges primarily stem from two components of these systems: (1) user actions and goals often unfold in a multitude of different orders within rich simulated environments, and (2) models must meet the run-time performance requirements of games. The open-ended nature of many serious games means that learners are free, and encouraged, to approach the task in any way they choose. This makes selecting contextual features for modeling tasks significantly challenging. Furthermore, most contemporary theories of emotion suggest that affect is often generated in response to the success or failure of one's goals, and how this success or failure came about (Elliot & Pekrun, 2007; Ortony et al., 1990). In serious games, students' goals may be unclear and without this knowledge, it will be difficult for the system to reason about success or failure and resulting affective states.

Recent work has investigated predictive models of students' educational goals and plans in serious games. These models consider sequences of student actions in a game environment, and provide early predictions about the goal a student seeks to next achieve, or the problem-solving plan they are currently executing. Models of student intentions hold considerable promise for informing predictions of students' affective states. Affective experiences are often a direct result of an individual's intentions and how these are or are not being realized in the world. An accurate prediction of a student's goals or intentions enables a system

to make inferences on whether these goals are being achieved and offer the opportunity to reason about attribution of success or failure. Together, intention and attribution offer significant insight into an individual's affective state (Ortony, Clore, & Collins, 1990) and are an important component of a context-based affect recognition model.

State-of-the-art serious games often employ computationally intensive graphics rendering and simulation capabilities. These technologies require substantial computational resources, and therefore limit the resources available for other system components such as real-time affect prediction models. While many successful models of affect prediction have been achieved by combining multiple data streams, it will be important to identify the data sources that provide the most benefit while imposing the least amount of computational overhead. For example, rich data from physiological sensors may be highly informative, but unless it can be effectively processed along with games' other required computational requirements, it may not be useful in practice. This highlights the importance of developing algorithms and data streams that are able to utilize computational resources as they become available and make predictions efficiently in real-time.

Affect-Informed Models of Student Learning

While affect predictions alone can provide meaningful information about student engagement while interacting with a serious game, another promising direction is identifying how affect predictions can be incorporated back into the understanding of the user's learning. Learner models are used to guide interactions and understanding of the student in a variety of systems. These models first seek to assess student knowledge, and then adapt tutorial interactions accordingly to maximize learning gains. A student model that incorporates affect may provide a richer picture of student learning. For example, affect has been shown to influence whether students are more likely process information in a bottom-up or top-down fashion (Pekrun et al., 2002). By considering a student's affective state as part of the learner model, considerations can be made about how a student is processing information, the connections that are likely being made, and how further information should be presented.

Incorporating affective channels into models of student knowledge acquisition and transfer for serious games hold considerable promise, particularly in domains where real-world applications of the learned knowledge are affect-intensive. Affect has been shown to influence cognition in many ways, and exploration of how these relationships can be utilized along with students' affective states is an important direction for developing comprehensive learner models.

Affect-Informed Tutorial Planning in Serious Games

Real-time models of students' affective processes could inform models for dynamically enhancing students' learning experiences in serious games. By combining serious games and intelligent tutoring systems, it is possible to devise models that tailor game mechanics, missions, rewards, and difficulty levels to scaffold student learning. In the case of story-based games, narrative-centered tutorial planners are a form of integrated pedagogical planner and interactive narrative director agent that discreetly support students' learning processes by tailoring story events. Narrative-centered tutorial planners consider the state of the student, interactive narrative, and learning progression to make decisions about how to adapt event sequences in the storyworld to support students' learning, problem solving, and engagement. For example, a narrative-centered tutorial planner may modify a virtual characters' behavior to provide additional hints and explanations to a student requiring special assistance. Alternatively, the planner could introduce novel goals and sub-plots to provide remediation for students, or opportunities for assessing advanced skills.

Real-time models of students' emotional states are well positioned to enhance the capabilities of these narrative-centered tutorial planners. However, there has been little work to-date investigating how real-time affect models can enhance narrative-centered tutorial planners' effectiveness. By making accurate and early predictions of students' emotions, narrative-centered tutorial planners could direct characters to provide personalized affective assistance, such as encouragement when a student is feeling confused, empathy when a student is feeling anxious, or advice to take a break when a student has experienced prolonged frustration while investigating a complex problem-solving task. These types of capabilities are poised to significantly enhance students' cognitive-affective processes in story-based serious games.

Illustrative Scenario

In order to illustrate these future concepts about real-time affect modeling in serious games, we describe a vision for affect-driven models in a game-based learning environment for middle school microbiology, CRYSTAL ISLAND (Rowe, Shores, Mott, & Lester, 2011). CRYSTAL ISLAND (Figure 1) features a science mystery, in which the student has arrived on a remote island to discover that the research team that has been established there has fallen ill. The camp nurse explains that they have not been able to identify the cause or type of illness and asks for the student's help. The student then works to collect clues by talking with virtual characters, running tests on objects in the world, and reading related books and posters. Once the student arrives at the correct source and type of illness and proposes a diagnosis, they have solved the mystery and completed the game.

As an open-ended game-based learning environment, CRYSTAL ISLAND features many goals and objectives that students may be working towards at any given time. For example, one student may be actively trying to identify the common symptoms among patients, while another may be trying to identify common food items that may be the source of the illness. Distinguishing between these two goals is critical for identifying whether a student is being successful at their goals and how they may feel as a consequence. If the student is trying to identify common food sources, but only hears about symptoms, they will likely be frustrated, while the other student would be feeling confident and hopeful that they will progress successfully.

Successful affect recognition in CRYSTAL ISLAND introduces opportunities for affect-driven learner models. Affective information can provide insight into how a student is learning, and the activities they are likely to pursue. For example, a student who is feeling confused is likely missing some piece of information that is critical to their understanding. A learner model that is able to assess knowledge and affect may be able to identify these gaps in understanding and drive tutorial planning so that the student is able to overcome this cognitive dissonance.

Another possibility for incorporating affect in tutorial and narrative planning revolves around the system's ability to foster positive affective states and engagement. For example, if the system detects that a student is frustrated or bored with the learning task, it may be advantageous to allow the student some respite from the difficult material. Game-based learning environments like CRYSTAL ISLAND allow opportunity for students to be engaged in the environment but not on the learning task specifically. Students may interact with the game environment's physics simulation by jumping on and stacking objects, or by simply exploring the rich 3D world. These types of actions may have positive affective benefits that can be encouraged if the system recognizes the need for affective regulation.



Figure 1. Screenshot of the CRYSTAL ISLAND game-based learning environment.

DISCUSSION

The opportunities that we have outlined for real-time affect detection in serious games highlight a number of implications for research and design of intelligent tutoring systems and related learning technologies. In particular, the key capabilities and tools provided by the GIFT framework are synergistic with the requirements and opportunities that we have discussed for modeling student affect. GIFT provides three primary functions: authoring capabilities for constructing novel learning technologies, instruction that integrates tutorial principles and strategies, and support for evaluating novel educational tools and frameworks. These capabilities provide a foundation for investigating real-time affect modeling in serious games, and further extensions to GIFT will enable studies of generalizable affect modeling frameworks.

Affect in Serious Games

Creating serious games can pose significant costs, due to the need for aesthetic 3D assets, novel interactive narratives, robust simulation models, and believable virtual agents, all meeting the run-time performance requirements of games. While serious games often do not compete directly with commercial entertainment-focused games, the continually rising bar for games' production values also increases expectations of serious games' complexity and aesthetics. Creating serious games requires close collaboration between inter-disciplinary teams of digital artists, computer scientists, subject matter experts, game designers, and instructors. Consequently, devising tools that can reduce the authoring costs associated with integrating adaptive tutoring capabilities for real-time affect modeling and scaffolding will represent a substantial advance. Further support for reducing authoring costs and increasing component reusability through GIFT, as well as research demonstrating the authoring benefits of these technologies, would be valuable to serious game developers.

Many serious games take advantage of advanced 3D graphics and real-time agent behavior algorithms to create immersive, believable virtual environments. However, these characteristics are computationally

intensive, and typically must be performed at more than 30 frames per second to preserve visual fidelity. In many cases, this leaves limited local resources for updating computationally sophisticated models of student emotion or learning processes. Off-loading these modeling and reasoning capabilities to external modules, especially modules hosted on external servers, is a promising approach for supporting computationally intensive probabilistic models of students' learning and affective processes, while meeting the run-time performance requirements of serious games that run on students' local computing hardware. GIFT's modular framework and service-oriented architecture are conducive to de-coupling learner modeling and pedagogical planning decisions from students' client machines.

GIFT also provides broad support for external sensors, which can supply real-time data on students' physiological state such as skin conductance, heart rate, posture, and eye gaze. Self-reports provide a useful window into students' affective processes, but they may be disruptive if presented during gameplay, and have limited accuracy particularly in cases where they are provided after an experience has concluded, and thus temporally removed from the actual occurrence of the emotion. Physiological sensors in many cases may be able to provide complementary information about students' affective states without disrupting gameplay and learning experiences. Furthermore, sensor data is objective, which removes some of the limitations in depending on students' abilities to recognize their own emotions and effectively and precisely report them.

Affect and the GIFT Architecture

The GIFT framework encompasses a modular architecture, which includes the following tutor components: a sensor module, a learner module, a pedagogical module and a domain module (Sottilare, Goldberg, Brawner, & Holden, 2012). Given this architecture, there are a number of questions and opportunities concerning how to effectively incorporate affect sensitivity within the GIFT framework. Future extensions to GIFT in service of real-time affect modeling will likely touch upon each of these four modules.

Each tutor module has its own distinct opportunities for processing affective information and communicating this to other modules. The sensor module is responsible for collecting and synthesizing information from various sources, such as webcams, electrodermal activity sensors and pressure-sensitive mice. This module must handle raw multimodal input from multiple concurrent sensors and provide output metrics that are useful for modeling the affective states of users. For example, a webcam capturing a large stream of video data could be analyzed to identify facial expressions, posture or non-verbal gestures indicative of affective states. In many cases, these sensors can produce large quantities of raw data. Consequently, as additional sensors are incorporated into GIFT, the sensor module must be capable of handling the memory and processing demands imposed by these new types of data, which are distributed across the collection, storage, cleaning and transformation stages of data management.

The learner module is responsible for representing the cognitive and affective states of learners using processed sensor data as well as learner performance data from the environment. Affective representations can include a variety of features, such as emotional state, self-efficacy, motivation, interest and intention. Accurately modeling each of these components may require detailed knowledge of the task and learning environment in addition to inputs provided by the sensor module. For example, modern appraisal-based theories of emotion depend heavily on how learners' actions and intentions play out in particular environments. Consequently, learner modules that implement these theories need to have access to information about the learning environment. Furthermore, affective and cognitive states are highly intertwined and must be considered in concert. For example, a learner's performance and knowledge may influence her emotional state and self-efficacy, which will then impact how she approaches the task moving forward.

The pedagogical module receives information about learners' current and predicted states, and it utilizes this to guide instructional strategies. One challenge that must be addressed by this module is balancing the tutoring system's affective and cognitive goals. For example, a more difficult learning task may be beneficial for increasing knowledge, yet an easier task may increase confidence and enjoyment. Furthermore, this module must consider a variety of strategies to bring about cognitive and affective improvements. Hints and feedback are commonly delivered to guide students' knowledge acquisition and problem solving, however these strategies may also include affective content. For example, empathetic feedback may encourage students to continue feeling positively in spite of poor performance. Alternatively, hints on effective emotion regulation strategies could be provided to students who appear to be struggling. Tailored support of learners' cognitive-affective states are likely to have substantial positive impacts on increasing time-on-task and sustaining learner engagement.

The domain module contains information about both the content area and the task environment in which learning interactions take place. It includes explicit representations of the types of feedback and adaptation capabilities supported by the learning environment, which can be used to enact the proposed strategies suggested by the pedagogical module. Such adaptations may involve virtual agents capable of verbal and non-verbal affective expression, tailored events in the learning environment's narrative, or dynamic adjustments to task difficulty. This module is also responsible for assessing student performance and identifying learner behaviors that are indicative of cognitive-affective states. For example, student off-task or gaming behaviors may serve as powerful indicators of student engagement and motivation.

Devising computational models for real-time affect prediction in serious games offers significant promise. Initial versions of GIFT have implemented modules primarily devised to effectively scaffold student knowledge acquisition and skill mastery. Future efforts to endow GIFT with affect modeling capabilities, such as the features described above, will require additional research and development efforts to expand the capabilities of each tutor module.

RECOMMENDATIONS AND FUTURE RESEARCH

GIFT provides a technological foundation for investigating predictive models of affect that can be generalized across different serious games. There are several promising directions for extending the capabilities of GIFT's modules to integrate comprehensive support for recognizing, understanding, and expressing affect in support of learning. Specifically, achieving the vision we have outlined for real-time affect detection in serious games through GIFT will require solutions to several research questions:

1. How can generalizable affect recognition, understanding, and expression capabilities be implemented within each of the modules of tutoring architectures such as GIFT?
2. How should interfaces between tutoring modules be developed to ensure robustness in handling different configurations of affect sensors and training environment capabilities?
3. How can generalizable, modular implementations of predictive affect models operate within the run-time performance requirements of serious games?
4. How can game-specific models for affect support be incorporated into generalizable tutoring architectures and transfer to alternate tasks and domains?

These questions highlight the challenges associated with investigating generalizable intelligent tutoring architectures capable of real-time affect modeling in serious games. With regard to GIFT specifically, it

will be important to identify the key issues, and solutions, for incorporating affect sensitivity within each of the tutor modules in the GIFT architecture. When designing these components, one must consider how these components communicate with one another, and how the system should be configured to support cases where components are missing. For example, physiological sensors are highly beneficial for affect recognition, but may not be available in all cases. Consequently, a learner model relying on output from such a sensor would need to be adapted, or gracefully deactivated, in a manner that minimizes negative impacts on other modules. Similarly, different genres of serious games have distinct capabilities and affordances. For example, serious games with believable virtual agents or rich narrative contexts may present different opportunities for affective feedback than serious games without these features. In cases such as these, pedagogical modules that recommend empathetic character behaviors or story event adaptations to provide affective support require mechanisms for handling cases where learning environments do not support these types of intervention naturally.

Additional challenges are raised by the computational demands of serious games. Games impose significant run-time performance requirements because they are computationally intensive, and they must balance graphics rendering and simulation capabilities alongside affective and tutorial modeling. Predictive affect models must balance efficiency, accuracy, and relevance to shape effective pedagogical interventions embedded within game environments. Finally, it will be important to investigate the ways in which findings related to the benefits and constraints of supporting affect in serious games extends to other tutoring environments. Advancements such as these would likely reduce the costs of devising new intelligent tutoring systems, as well as enhance their effectiveness for different learning settings.

REFERENCES

- Baker, R. S., D’Mello, S., Rodrigo, S. K., & Graesser, A. C. (2010). Better to Be Frustrated than Bored: The Incidence, Persistence, and Impact of Learners’ Cognitive-Affective States during Interactions with Three Different Computer-Based Learning Environments. *International Journal of Human-Computer Studies*, 68(4), 223–241.
- Calvo, R. A., & D’Mello, S. (2010). Affect Detection: An Interdisciplinary Review of Models, Methods, and Their Applications. *IEEE Transactions on Affective Computing*, 1(1), 18–37.
- Conati, C., & Maclaren, H. (2009). Empirically building and evaluating a probabilistic model of user affect. *User Modeling and User-Adapted Interaction*, 19(3), 267–303.
- Elliot, A. J., & Pekrun, R. (2007). Emotion in the Hierarchical Model of Approach-Avoidance Achievement Motivation. In P. A. Schutz & R. Pekrun (Eds.), *Emotion in Education* (pp. 57–74). London: Elsevier.
- Forgas, J., Wyland, C., & Lahan, S. (2006). Hearts and minds: An introduction to the role of affect in social cognition and behavior. In J. P. Forgas (Ed.), *Affect in Social Thinking and Behavior* (pp. 3–18). New York: Psychology Press.
- Kim, J., Hill, R. W., Durlach, P. J., Lane, H. C., Forbell, E., Core, M., Marsella, S. C., & et al. (2009). BiLAT: A game-based environment for practicing negotiation in a cultural context. *International Journal of Artificial Intelligence in Education*, 19(3), 289–308.
- Ortony, A., Clore, G., & Collins, A. (1990). *The cognitive structure of emotions*. Cambridge University Press.
- Paiva, A., Dias, J., Sobral, D., & Aylett, R. (2004). Caring for agents and agents that care: Building empathic relations with synthetic agents. *Proceedings of the 3rd International Joint Conference on Autonomous Agents and Multiagent Systems* (pp. 194–201).

Design Recommendations for Adaptive Intelligent Tutoring Systems Learner Modeling (Volume I)

- Pekrun, R., Goetz, T., Titz, W., & Perry, R. (2002). Academic Emotions in Students' Self-Regulated Learning and Achievement: A Program of Qualitative and Quantitative Research. *Educational Psychologist, 37*(2), 91–105.
- Picard, R. (1997). *Affective Computing*. Boston: MIT Press.
- Picard, R. W., Papert, S., Bender, W., Blumberg, B., Breazeal, C., Cavallo, D., Machover, T., et al. (2004). Affective Learning — A Manifesto. *BT Technology Journal, 22*(4), 253–269.
- Rowe, J. P., Shores, L. R., Mott, B. W., & Lester, J. C. (2011). Integrating learning, problem solving, and engagement in narrative-centered learning environments. *International Journal of Artificial Intelligence in Education, 21*(2), 115–133.
- Russell, J. (2003). Core affect and the psychological construction of emotion. *Psychological Review, 110*, 145–172.
- Sottolare, R. A., Goldberg, B. S., Brawner, K. W., & Holden, H. K. (2012). A modular framework to support the authoring and assessment of adaptive computer-based tutoring systems (CBTS). *Proceedings of the Interservice/Industry Training, Simulation, and Education Conference*.