Leveraging a Generalized Tutoring Framework in Exploratory Simulations of Ill-Defined Domains

James M. Thomas[1], Ajay Divakaran[2], Saad Khan[2]

[1]*Soar Technology, Inc., Ann Arbor, MI*
jim.thomas@soartech.com

[2]*SRI International Sarnoff, Princeton NJ*
{ajay.divakaran, saad.khan}@sri.com

Abstract. Generalized frameworks for constructing intelligent tutors, such as GIFT promise many benefits. These benefits should be extended to systems that work in ill-defined domains, especially in simulation environments. This paper presents ideas for understanding how ill-defined domains change the tutoring dynamic in simulated environments and proposes some initial extensions to GIFT that accommodate these challenges.

Keywords. Intelligent Tutoring Systems, Computer-Based Tutoring Systems, framework, GIFT, ITS, CBTS, student models, learner models, task models, ill-defined domains, simulated environments, exploratory learning.

1 Introduction

Intelligent Tutoring System (ITS) have been shown to enhance learning effectiveness in a wide variety of academic domains [1]. The ITS field has long drawn inspiration from studying strategies employed by human tutors in one-on-one engagement with students [2]. Success has spurred the research community to extend its aspirations into more complex, ill-defined domains. Ill-defined domains are those that lack clearly defined procedures to determine the correctness of proposed solutions to specific problems [3]. Our interest lies in exploratory training simulations of those domains.

To address the difficulty of guiding effective learning in these complex environments, it seems useful to develop and leverage generalized techniques. The GIFT architecture represents a comprehensive approach to facilitate reuse of common tools and methods for building ITS. Although much of the initial focus of GIFT has been directed toward well-defined domains, it we would like to consider how it could be extended to ill-defined domains as well [4], especially those rendered through exploratory simulations.

The authors' motivating example is a system we are building called "Master Trainer – Individualized" (MT-I). The goal for this system is to intelligently guide new military squad leaders in simulations that combine intercultural communication and negotiation skills with tactical challenges. This system integrates stand-off assessments of student affect to modulate the intensity of the simulation to optimize student challenge. One of the questions we are investigating is how to drive the relationship between the student and a simulated human to achieve pedagogically useful levels of anger, conflict or cooperation. We are interested in applying what we are learning toward the generation of useful domain-independent strategies that could be incorporated into GIFT.

2 Motivations for GIFT

Although the field of ITS research is imbued with a strong collaborative spirit, the field lacks common computational infrastructure. GIFT is a particularly promising approach toward a general reusable framework for intelligent tutoring could benefit the entire field.

Scientific research largely presumes the capability to make apples-to-apples comparisons of competing theories. Although they share some common concepts and goals, the majority of ITS research systems share little common architecture or code [1]. To make broad contributions to this field often requires a fairly full-featured ITS on which to perform analyses, yet bespoke software development is both time consuming and expensive. Shared platforms and plug-ins amortize development costs and grow communities of professionals who can more effectively collaborate and relocate between projects and organizations, accelerating the productivity of the field as a whole [5].

GIFT proposes common frameworks for alternative implementations of a broad set of ITS capabilities. Built on solid design principles and a comprehensive understanding of the work of the ITS community, GIFT promises to serve an increasingly useful role in accelerating the scientific and commercial success of the field. Three common challenges faced by the field: authoring, instructional management, and analysis form the core constructs of GIFT. Successful evolution of these constructs promises to accelerate scientific progress by sharing common evaluation methodologies, reducing the time and expense for reused software components, and promoting a more tightly integrated and collaborative community.

GIFT may help accelerate commercialization of scientific progress by facilitating the production of a common currency of evidence of learning effectiveness that can be used to sell the benefits of implemented systems. It can help provide a platform for rapid prototyping to more quickly cycle through alternative approaches to find those that work best. Much as Eclipse™ has accelerated software development [5], and Unity3D™ has democratized high fidelity game development [6], GIFT has the potential to grow into a common workbench that builds-in the ability to package and deploy new work to a full breadth of possible platforms.

3 Characteristics of Ill-Defined Domains

The current GIFT vision accommodates a wide range of ITS capabilities. However, ill-defined domains have not been a primary component of that vision [4]. This section begins with an irony-free definition of ill-defined domains, describes some of the challenges encountered by human tutors in a subset of these domains, and then considers issues and opportunities they present for ITS designers working with immersive simulation environments.

3.1 Defining Ill-Defined Domains

Much of the historical grounding of ITS research is focused on guiding students through well-structured discrete learning tasks, to impart deeply decomposable knowledge [5] from well-defined domains. Fournier-Viger et al. [8] declare ill-defined domains to be those "where traditional approaches for building tutoring systems are not applicable or do not work well". Lynch and Pinkus [9] characterize problems in ill-defined domains as lacking definitive answers, having answers heavily dependent upon the problem's conception, and requiring students to both retrieve relevant concepts and map them to the task at hand. Mitrovic [10] underscores the important distinction between ill-defined domains and ill-defined tasks, anticipating Sottilare's [4] observation that ITS authoring in ill-defined domains is complicated by the multiplicity of "paths to success" compared to the more well-defined domains in which of ITS research has been situated.

3.2 Tutoring Challenges Posed by Ill-Defined Domains

Human tutors have served as both an inspiration for ITS behavior and benchmark and a benchmark for ITS performance [1]. Because no one has yet made a comprehensive study comparing human tutor behaviors in traditional domains with those in ill-defined domains to identify the most necessary extensions to tutorial reasoning, our work on the MT-I system is inspired by specific analogues of human tutors in the domains of live military training for tactics and intercultural effectiveness.

Live training in environments that combine military tactics and interpersonal challenges often spans many hours or days, ranges through confined indoor and expansive outdoor spaces, and requires dozens or hundreds of live role players. Interactions with these role players are often guided by scripted prompts, but involve a lot of improvisation as well. Examples include resolving disputes between armed civilians, negotiating with civic or spiritual leaders, as well neutralizing threats posed by snipers or potential ambushes. Trainers are usually embedded within the environment and have the ability to provide guidance to students during the simulation.

When comparing the behavior of the trainers/tutors in these exercises to that of academic tutors, a striking contrast is immediately evident. Feedback is often deferred over much longer intervals than what one would typically see in one-on-one tutoring in well-defined academic domains. Because is often unsafe or impossible to suspend exercises involving moving/flying vehicles and timed explosions, most incorporate extensive after-action review (AAR) as the primary conduit for feedback and guidance. To some extent, the tutors may elect to integrate feedback within the broader context of a scheduled AAR. In other cases, immediate feedback cannot be given on an individual student action choice because multiple student actions are required before a judgment can be made. Some of this deferral is linked to the interplay between student and role players, as it is difficult to guide the student without impacting the on-going social exchange. Finally, unlike many academic tasks, it is difficult to reset the problem state after an incorrect student action, as the training is situated in a narrative context with a fixed rate of flow to coordinate the many moving parts.

The immediate feedback tutors do provide in these simulations is often constrained to ensuring that the student is engaged and taking actions that move the implicit narrative forward. The deferred feedback is often a holistic reflection involving multiple learning objectives, student affect and metacognitive guidance on productive application of the feedback to future performance.

3.3 Tutoring in Computer Simulations of Ill-Defined Domains

Many of the challenges encountered by humans in ill-defined domains carry over into computer-based tutoring. The assessment granularity sometimes spans multiple actions, can sometimes be entwined in social interactions, and can sometimes be entwined in narrative. Each of these specific constraints can be viewed more generally.

What we commonly describe as narrative in simulation environments is more generally described as a meaningful continuity of state over time. Narrative-centered learning environments [11, 12] can vary in the extent to which they support alternative branches toward "success" or even emergent run-time generation. Yet they share the constraint that the continuity associated with the progression of states cannot be broken or the reversed without consequence, which in turn places limitations on the action choices available to both student and tutor.

Similarly, what we commonly perceive as social interaction between students and non-player characters (NPCs) in simulations is one particular case of an addition of simulation-based elements to tutorial state. In this case, it is the game-state data associated with the NPCs attitudes toward the student that persistent over sequences of tutorial actions. Other examples of game/simulation state variables that influence tutorial state include consumable or non-replenishable resources in the simulation which may affect the span of future tutorial choices.

Finally, the dependency on multiple student actions for student assessment is a specific manifestation of the well-recognized and more general problem of assessing student correctness at all in ill-defined domain. Yet while these challenges complicate the job of intelligent tutoring, they also introduce new tools. Narrative continuity can be exploited both to scaffold instruction and provide context for interpretation of actions. NPCs and other simulation based entities can be manipulated for pedagogical purposes, providing implicit guidance or challenge to the student. The complexity of interpretation of student action affords the intelligent tutor the opportunity for more nuanced and complex forms of guidance that may have more profound and lasting effects on learning.

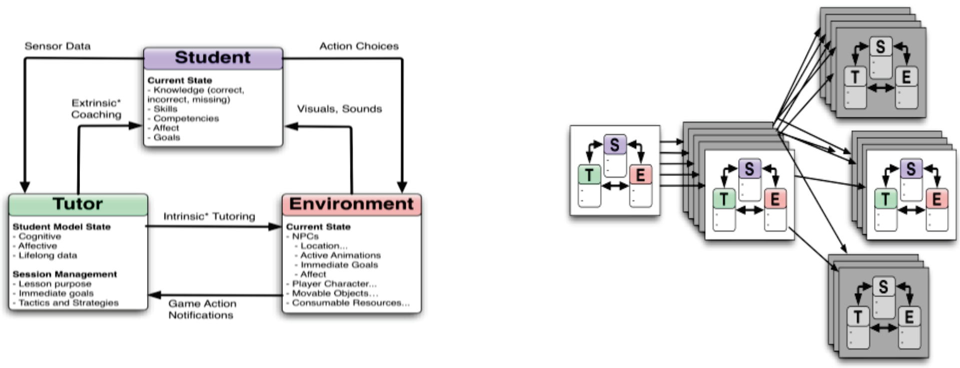

Fig. 8. Model of Tutoring Dynamic in Simulated Ill-Defined Domains

To best confront the challenges and make use of the opportunities of learning based in simulation environments over ill-defined domains, we need models that understand the tutoring dynamic as more than a one-on-one exchange. Rather, the model must recognize that the persistent state, continuity constraints, and assessment ambiguity of the simulation environment continually shape the interactions between tutor and student. Figure 1 is a depiction of such a model. At any point in time, the state of an ITS can be described as a combination of the state data associated with the student, the tutor and the simulation environment. Arrows depict the flow of state-changing actions between these three components. Note that some of these actions may proceed in parallel and may last for human-perceptible durations; perhaps with sufficient frequency that the overall state of the system may be more often in flux than it is quiescent.

This expanded interaction model complicates the prescription of the “learning effect chain” [4]. Because any change to student, tutor, or environment is represented as a new state, the progression toward learning gains involves navigating through a broad space of potential alternative paths. As shown in the rightmost half of Figure 1, one particular progression (the sequence of colored tutorial state snapshots depicted against white backgrounds) is merely one path through a rapidly expanding profusion of alternatives.

This tableau of interwoven learning progressions and alternatives gives an ambitious tutorial agent a lot to think about. Sufficiently inspired agents may perform plan-based reasoning to map the possibilities and nudge the learning experience in the most fruitful directions. In fact, tutorial agents have been constructed that mine the space of alternative actions sequences [12] to devise remediation strategies. Advanced agents might even consider choreographing multiple sessions, altering emphasis and tactics as it varies the pedagogical purpose of each session.

Alternatively, the profusion of possibilities can influence time-sensitive developers to move in the other direction, building “knowledge-lean” [13tutorial agents. As a consequence, ITS developers in these environments often eschew deep knowledge-tracing expert models in favor of less precise, but more easily authored constraint-based approaches [8]. This suggests that a generalized intelligent framework, such as

GIFT, should consider supporting a variety of modeling approaches. In fact, our current implementation of MT-I, which features ill-defined tasks within overlapping ill-defined domains, we have found it useful to author constraint-based models to characterize the correctness of individual student tasks in a wide range of potential contexts, where that model feeds a higher-level knowledge tracing model of higher-level, more abstract learning objectives that operates over longer time spans.

4 Enhanced Knowledge Representations and Reasoning

Not surprisingly, some of the challenges posed by ill-defined domains in simulated environments can be addressed by providing tools to create better definitions. Flexible and knowledge representations (KR) can serve as the definitional "handles" that tutorial agents can use to enhance reasoning about the state of the student and simulation. That reasoning can be converted to action if the simulation is instrumented with "levers", software hooks that cause pedagogically useful changes expressed through those handles. This section proposes three levers that use non-traditional extensions to tutorial knowledge representations to provide enhance tutorial reasoning and more effective student guidance.

Lever #1: Enriched Definitions of Learning Objectives. Trainers in the sophisticated simulations involving role players discussed earlier are often trying to steer their students toward states of mind that go beyond a prescribed set of factual knowledge to include social, narrative and affective dimensions, as shown in Figures 1 and 2. To achieve similar results in simulated environments, tutorial agents must reason about those dimensions of learning objectives as well. The KR should be able to qualify, for example, not only that the trainee know how to greet respectfully a village leader, but that the student can perform that greeting is accomplished while in a highly agitated state.

Lever #2: Enriched Definitions of Tutorial Purpose. Sophisticated simulations can be used in a broader set of pedagogical contexts that traditional systems, ranging from direct instruction of material to which the student has not previously been exposed, to consolidation of previously taught material, to transfer of knowledge to new domains, to assessment of knowledge and performance, to building confidence, teamwork or skills that apply acquired knowledge. Thus, the purpose of a given tutorial session can vary more widely than in traditional instruction, which demands that tutorial strategies and tactics be labeled according to their relevance for these various purposes. For example, a particular tutorial action may have a stronger positive effect on student self-efficacy that an alternative which may have a stronger positive effect on didactic specificity. An enhanced KR enables the tutorial agent to choose between these alternatives based on whether the purpose of the current session is to build confidence or impart knowledge.

Lever #3: Persistently-labeled Learner Data. To maximally leverage the opportunities of sophisticated learning environments, in which multiple learning sessions for varying learning purposes may span arbitrary time periods, individual student data must be persistent and pervasive: accessible and publishable at any level by any component of the tutorial framework. This allows tutorial agents running at various levels with various time horizons to tie together data collected on individual students across multiple sessions. For example, it could prove useful to know how quick a student is to anger, or which immediately reachable emotional states are most conducive to learning for a particular student, where that data may have been collected and stored in an earlier tutorial session by an agent using the same generalize frame work. All student model data should be tagged with its expected lifespan: step, task, session, application, or beyond. This enhances the ability of any particular tutorial agent to perform macro-level adaptation [14] to evolve learning across multiple domains that enhance domain-independent competencies.

5 Conclusions

A generalized framework like GIFT holds significant promise for accelerating scientific and commercial success of ITS. Yet one of the areas in which that acceleration is most desperately needed, ill-defined domains in simulated environments, are not addressed in depth by the current approach to GIFT. We suggest that a first step in this direction would to explore several extensions to the knowledge representations in GIFT to meet the demands of those environments.

6 Acknowledgements

This research effort was supported in part by contract W911QX-12-C-0149 from the US Army Research Laboratory. We gratefully acknowledge the interest and encouragement from Dr. Bob Sottilare at ARL. In addition, Scott Flanagan of Sophia Speira, LLC, has provided invaluable understanding of military training practices and domains.

7 References

1. VanLehn, K. (2011). The relative effectiveness of human tutoring, intelligent tutoring systems, and other tutoring systems. *Educational Psychologist*, *46*(4), 197-221.
2. Wood, D., Bruner, J. S., & Ross, G. (1976). The role of tutoring in problem solving*. *Journal of child psychology and psychiatry*, *17*(2), 89-100.
3. Minsky, M. (1995). Steps to artificial intelligence. In Luger, G. F., editor, *Computation and Intelligence, Collected Readings*, chapter 3, pages 47–90. AAAI, Menlo Park CA, and MIT Press.
4. Sottilare, B. (2012). Considerations in the development of ontology for a generalized intelligent framework for tutoring. *Proceedings of the International Defense and Homeland Security Simulation Workshop.*
5. Kidane, Yared H., and Peter A. Gloor. "Correlating temporal communication patterns of the Eclipse open source community with performance and creativity." *Computational and mathematical organization theory* 13.1 (2007): 17-27.
6. Torrente, Javier, et al. "Game-like simulations for online adaptive learning: A case study." *Learning by Playing. Game-based Education System Design and Development*. Springer Berlin Heidelberg, 2009. 162-173.VanLehn, K. (1990). Mind bugs: the origins of procedural misconception. MIT press.

7. VanLehn, K. (1990). Mind bugs: the origins of procedural misconception. MIT press.
8. Fournier-Viger, P., Nkambou, R., & Nguifo, E. M. (2010). Building Intelligent Tutoring Systems for Ill-Defined Domains. In *Advances in Intelligent Tutoring Systems* (pp. 81-101). Springer Berlin Heidelberg.
9. Lynch, C., Ashley, K., Aleven, V., & Pinkwart, N. (2006). Defining ill-defined domains; a literature survey. In Proceedings of the Workshop on Intelligent Tutoring Systems for Ill-Defined Domains at the 8th International Conference on Intelligent Tutoring Systems (pp. 1-10).
10. Mitrovic, A., & Weerasinghe, A. (2009). Revisiting Ill-Definedness and the Consequences for ITSs. Artificial Intelligence in Education: Building Learning Systems That Care: from Knowledge Representation to Affective Modelling, 200, 375.
11. Mott, B., McQuiggan, S., Lee, S., Lee, S. Y., & Lester, J. C. (2006). Narrative-centered environments for guided exploratory learning. In *Proceedings of the AAMAS 2006 Workshop on Agent-Based Systems for Human Learning* (pp. 22-28).
12. Thomas, J. M., & Young, R. M. (2009, July). Using Task-Based Modeling to Generate Scaffolding in Narrative-Guided Exploratory Learning Environments. In *Proceeding of the 2009 conference on Artificial Intelligence in Education: Building Learning Systems that Care: From Knowledge Representation to Affective Modeling* (pp. 107-114).
13. Lane, H. C., Core, M. G., Gomboc, D., Karnavat, A., & Rosenberg, M. (2007, January). Intelligent tutoring for interpersonal and intercultural skills. In *The Interservice/Industry Training, Simulation & Education Conference (I/ITSEC)*(Vol. 2007, No. 1). National Training Systems Association.
14. Sottilare, R. A., Goldberg, B. S., Brawner, K. W., & Holden, H. K. (2012, January). A Modular Framework to Support the Authoring and Assessment of Adaptive Computer-Based Tutoring Systems (CBTS). In *The Interservice/Industry Training, Simulation & Education Conference (I/ITSEC)*(Vol. 2012, No. 1). National Training Systems Association.

Authors

James M. Thomas: James M. (Jim) Thomas is a Research Scientist at Soar Technology. His current work includes intelligent training and assessment systems that aid children on the autism spectrum and guide soldiers to integrate socio-cultural and tactical skills. He received his Ph.D.degree in Computer Science from North Carolina State University in 2011. His dissertation entitled "Automated Scaffolding of Task-Based Learning in Non-Linear Game Environments", explored general mechanisms to generate intelligent tutorial planning within exploratory learning environments. Jim has authored more than a dozen papers in the area of intelligent tutoring systems, automated planning, and computer games. His graduate studies were supported by a National Science Foundation Graduate Research Fellowship in Artificial Intelligence. Concurrent with his doctoral studies, Jim developed game-based learning systems designed to improve children's social skills at the 3-C Institute for Social Development, including work which was published in the journal *Child Development*. Jim also benefits from 15 years of experience in the computer and telecommunications industries, including software development, management, and senior marketing management with IBM, BNR and Nortel Networks.

Ajay Divakaran: Ajay Divakaran, PhD is a Technical Manager at SRI International Sarnoff. He has developed several innovative technologies for multimodal systems for both commercial and government programs over the past 16 years. He currently leads SRI Sarnoff's projects on Modeling and Analysis of Human Behavior for the DARPA SSIM project, ONR Stress Resiliency project, Army "Master Trainer" Intelligent Tutoring project among others. He worked at Mitsubishi Electric Research Labs for ten years where he was the lead inventor of the world's first sports highlights playback enabled DVR, as well as a manager overseeing a wide variety of product applications of machine learning. He was elevated to Fellow of the IEEE in 2011 for his contributions to multimedia content analysis. He established a sound experimental and theoretical framework for human perception of action in video sequences, as lead-inventor of the MPEG-7 video standard motion activity descriptor. He serves on TPC's of key multimedia conferences and served as an associate editor of the IEEE transactions on Multimedia from 2007 to 2011 and has two books and over 100 publications to his credit as well as over 40 issued patents. He received his Ph.D. degree in Electrical Engineering from Rensselaer Polytechnic Institute in 1993.

Saad Khan: Saad Khan is a Senior Computer Scientist at SRI International with 10 years of experience developing computer vision algorithms. He has led the design and development of advanced military training systems that can adapt to both training scenarios and learners' behavior. He serves as PI/ Co-PI and technical lead on programs in multimodal sensing algorithms for immersive training for DARPA, ONR and PMTRASYS. He led the development and transition of APELL (Automated Performance Evaluation and Lessons Learned) training system. APELL is an immersive, interactive, Mixed Reality training system that has been successfully deployed at the Marines Camp Pendleton training facility. Prior to joining SRI Sarnoff, Dr. Khan conducted research on 3D model based object tracking and human activity analysis in the IARPA VACE program. His work in automated image based localization earned an Honorable Mention award at the International Conference of Computer Vision 2005. He has authored over 20 papers and has 2 issued patents. His work on multiple view human tracking is one of the most highly cited works in recent tracking literature. He received his PhD in Computer Science from University of Central Florida in 2008.