Modeling User Knowledge with Dynamic Bayesian Networks in Interactive Narrative Environments

Jonathan P. Rowe and James C. Lester

Department of Computer Science
North Carolina State University
Raleigh, NC 27695
{jprowe, lester}@ncsu.edu

Abstract

Recent years have seen a growing interest in interactive narrative systems that dynamically adapt story experiences in response to users' actions, preferences, and goals. However, relatively little empirical work has investigated runtime models of user knowledge for informing interactive narrative adaptations. User knowledge about plot scenarios, story environments, and interaction strategies is critical in a range of interactive narrative contexts, such as mystery and detective genre stories, as well as narrative scenarios for education and training. This paper proposes a dynamic Bayesian network approach for modeling user knowledge in interactive narrative environments. A preliminary version of the model has been implemented for the CRYSTAL ISLAND interactive narrative-centered learning environment. Results from an initial empirical evaluation suggest several future directions for the design and evaluation of user knowledge models for guiding interactive narrative generation and adaptation.

Introduction

Artificial intelligence technologies for real-time narrative generation and adaptation have shown promise for dynamically tailoring interactive story experiences (Mateas and Stern, 2005). Interactive narrative systems can benefit from user models that inform real-time narrative adaptation decisions. Over the past several years, the interactive narrative community has investigated empirical models of canonical player types for dynamic quest selection (Thue et al. 2007), models of user behavior for informing proactive story mediation decisions (Magerko 2006), and predictive assessments of user responses to dilemma scenarios (Barber and Kudenko 2007).

One area of user modeling that has received comparatively little attention by the interactive narrative community is modeling user knowledge. User knowledge models assess understanding of important narrative events, concepts, facts, and strategies. User knowledge is important in several genres of narratives, such as mystery and detective scenarios, where interpreting clues and prior events is critical. Interactive narratives can employ knowledge to drive run-time decisions about story adaptations, such as varying non-player character dialogues, introducing side quests, or dynamically adjusting degree of difficulty. For example, consider an interactive mystery where the user has overlooked an important clue. A narrative director agent could query a user knowledge model to identify gaps in the user's understanding, and then assign a quest that will re-reveal the necessary information, or direct a non-player character to provide a hint. Similarly, knowledge models are valuable for providing run-time assistance in education and training applications of interactive narratives.

Devising effective models of user knowledge poses significant computational challenges. First, models of user knowledge must cope with multiple sources of uncertainty inherent in the modeling task. Second, knowledge models must dynamically model knowledge states that change over the course of a narrative interaction. Third, the models must concisely represent complex interdependencies among different types of knowledge, and naturally incorporate multiple sources of evidence about user knowledge. Fourth, the models must operate under the real-time performance constraints of interactive narratives.

To address these challenges, this paper proposes a dynamic Bayesian network approach to modeling user knowledge during interactive narrative experiences. Dynamic Bayesian networks offer a unified formalism for representing temporal stochastic processes such as those associated with knowledge modeling in interactive narrative environments. The framework provides a mechanism for dynamically updating a set of probabilistic beliefs about a student's understanding of narrative, scenario solution, strategic, and curricular knowledge components that are accumulated and demonstrated during interactions with a narrative environment. An initial version of the model has been implemented in CRYSTAL ISLAND, a testbed narrative-centered learning environment featuring an educational science mystery. An empirical evaluation offers preliminary support for the framework's potential for interactive narrative environments and run-time knowledge models.

Related Work

A majority of interactive narrative work has focused on computational models for dynamically adjusting plot and discourse during user interactions with narrative environments (Mateas and Stern 2005; Riedl and Stern 2006), as well as driving believable non-player character behaviors (Si, Marsella, and Pynadath 2006). Comparatively little work has investigated devising and validating user models to inform the narrative generation process. The Interactive Drama Architecture (IDA) implements a rule-based player model to predict future user actions in the *Haunt 2* virtual environment (Magerko 2006). Other work has classified players into predetermined behavioral categories (Thue et al. 2007), structured user models around small sets of possible behaviors (Roberts et al. 2006), or has not been empirically validated. Prior work on the CRYSTAL ISLAND narrative environment used n-grams and Bayesian networks to induce models for goal recognition (Mott, Lee, & Lester, 2006).

Student knowledge modeling has been the subject of extensive investigation in the intelligent tutoring systems community. Student models are critical in intelligent tutoring systems' for dynamically selecting practice problems and delivering tailored hints (VanLehn, 2006). A number of intelligent tutoring systems utilize probabilistic representations in order to cope with the uncertainty inherent in student modeling. The Cognitive Tutor family of intelligent tutors uses knowledge tracing, a probabilistic overlay technique for assessing students' domain knowledge (Corbett and Anderson 1994). Andes is an intelligent tutor for introductory college physics, which utilizes Bayesian networks to probabilistically model students' domain knowledge and problem-solving plans (Conati et al., 2002). Prime Climb is an educational game for young students learning number factorization, and has produced several iterations of dynamic Bayesian networks for modeling student knowledge and affect (Manske & Conati, 2005). While intelligent tutoring systems provide a useful reference point for effective knowledge modeling, interactive narratives pose additional computational challenges that merit investigation.

Representational Requirements

Knowledge modeling is a problem of partial observability and is characterized by multiple sources of uncertainty. One source of uncertainty has been described as the *assignment of credit problem* in the intelligent tutoring systems literature (VanLehn, 2006). The problem is associated with situations in which a single sequence of observations corresponds equally well to several qualitatively distinct models of the user's knowledge. For example, consider an interactive detective scenario where a user decides to accuse a particular non-player character of being the mystery's main culprit. This action could be consistent with multiple possible models of the user's knowledge: the user may confidently know of the character's guilt, or the user may have luckily guessed the culprit's identity without the requisite knowledge, or the user may suffer from a misconception that fingered the correct character but for the wrong reasons. Without additional information, it can be difficult to identify the type of knowledge evidenced by the user's action.

Another representational requirement is coping with student knowledge states that change over the course of a narrative interaction. As users successfully explore the environment and uncover plot details, their knowledge of the scenario and narrative environment is likely to increase. Further, these changes often involve complex combinations of declarative, procedural, and narrative knowledge evidenced across diverse sets of information gathering and problem solving actions.

Many interactive narratives provide multiple resources for learning about important narrative events and concepts (e.g., conversations with non-player characters, environmental props such as posters or televisions), as well as methods for demonstrating knowledge (e.g., narrative actions, solutions to diegetic puzzles). A knowledge model should be able to integrate these different sources of evidence of student knowledge, differentiate the degree of evidence provided by each source, and support parameters to represent individual users' knowledge abilities.

Further, the real-time performance constraints of interactive narratives require that any computational formalism for modeling user knowledge support efficient algorithms for updating and querying the model. A knowledge model need not necessarily update during each clock cycle of an interactive narrative's underlying game engine, but it must update quickly enough to be useful for real-time narrative decision making.

A Dynamic Bayesian Network Approach to Modeling User Knowledge

Dynamic Bayesian networks (DBNs) provide a natural representation to meet the requirements of user knowledge modeling in interactive narratives. Dynamic Bayesian networks are an extension of Bayesian networks, which provide a concise, graphical representation for quantifying and reasoning about uncertainty. A Bayesian network is a directed acyclic graph with nodes representing individual random variables, and directed links representing dependencies between variables (Pearl, 1988). Each node is annotated with a conditional probability distribution over possible states given all possible combinations of its parent nodes' states. Parent nodes influence child node values in the direction of their connecting links. The absence of a directed link between any two nodes indicates an independence relationship between the respective random variables. Whenever a direct observation is made of the system being modeled, the associated node in the Bayesian network is clamped to the appropriate value, and the evidence is propagated through the network via Bayesian inference. Several exact and approximate Bayesian inference techniques have been devised, and are adequately efficient for practical use.

Dynamic Bayesian networks extend static Bayesian networks by dynamically inserting and removing nodes over the course of a modeling task (Dean & Kanazawa, 1989). Temporal DBNs typically follow the form of a first-order Markov process, which assumes that future states are independent of the model's prior history, given its current state. A significant advantage of DBNs over static networks is their ability to explicitly represent changes in a model's belief state over time. By structuring contemporaneous state variables into discrete time slices, and progressively appending new slices to the network, DBNs maintain a growing history of the process being modeled. Figure 1 provides a conceptual illustration of such a DBN structure.

A DBN knowledge-tracing model functions as a probabilistic overlay model, monitoring the user's actions during the narrative scenario and updating values representing her mastery over a set of target knowledge components. As the user's knowledge changes over the course of an interaction, those changes are reflected through dynamic updates of the knowledge model. Knowledge values are represented as probabilities, which indicate the model's beliefs about the user's understanding of the associated knowledge components.

DBN model parameters can be hand engineered, or they can be machine learned from data. The size and complexity of a user knowledge model can significantly impact the efficacy of techniques for automatic parameter learning; for this reason, an investigation of automated techniques for parameter learning is left for future work. This work presents a four-stage procedure for hand authoring a DBN knowledge-tracing model for an interactive narrative. First, a set of relevant knowledge components is identified as the subject of modeling. Second, the relationships between different knowledge components are specified, characterizing a template network structure. Third, a set of observable user actions is identified, as are relationships between the observable actions and the model's knowledge components. Fourth, the model's parameters are defined in terms of conditional probability tables for each of the hidden and observable nodes in the network.

Identification of knowledge components. The first phase of the DBN model's construction requires identifying a relevant set of knowledge components. Knowledge components are represented as binary random variables, each designating the probability that the student understands a particular concept or fact. Several categories of knowledge may need to be included in the model, such as narrative knowledge, interaction strategy knowledge, and scenario solution knowledge. *Narrative knowledge* consists of scenario-specific facts that are gathered during the interaction. *Strategic knowledge* represents a student's understanding of the activities necessary for completing the interactive narrative scenario. *Scenario solution knowledge* consists of scenario-specific propositions that are deduced by the student as she works toward completion of the narrative, and are relevant to its ultimate resolution. In the case of educational narratives, it may be necessary to also define *Content knowledge* components, which consist of facts and concepts from a learning domain.

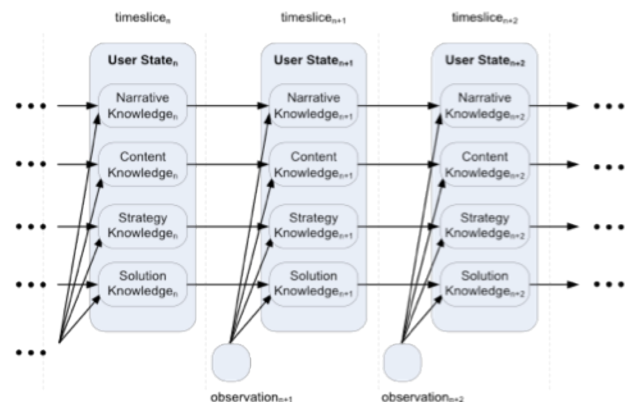

Figure 1. Conceptual illustration of three slices from a dynamic Bayesian network.

Specification of network structure. After identifying a set of relevant knowledge components, the network's directed link structure must be defined. Connections between knowledge components within a single slice (contemporaneous) follow precondition relations. Child nodes represent knowledge components that can be deduced from mastery of the parent nodes' knowledge components. For example, if a student knows the symptoms of one sick character, as well as the symptoms of another sick character, she is equipped to deduce the shared symptoms of the two ill characters. Network connections across adjacent slices (non-contemporaneous) connect subsequent instances of the same knowledge component. For example, the node representing knowledge of a sick patient's symptoms in slice$_t$ is the source of a directed link whose destination is the corresponding symptom node in slice$_{t+1}$. This structure permits the model's values for a given component to be conditioned on its prior values, in addition to newly available evidence.

Incorporation of observable features. The third stage of model construction specifies a set of in-game actions that provide evidence of user knowledge. The relationships between these observable features and the model's knowledge components are also defined during this stage. It should be noted that each type of observation offers a distinct amount of evidence of the student's underlying knowledge. Furthermore, some types of in-game actions may be diagnostic of student knowledge, while other actions may actually modify student knowledge. Each of these distinctions requires adjustments to the choice of directed links and conditional probabilities associated with observation nodes and their neighbors.

Specification of network parameters. During the fourth stage of model construction, conditional probability tables (CPTs) are manually defined for each of the network's nodes. In conjunction with directed links, the CPTs specify the relationships among adjacent hidden knowledge components, as well as between hidden and observable variables, encoding policies for propagating evidence through the network. When conditional probability tables describing knowledge precondition relationships are manually specified, knowledge components may be individually weighted based on their perceived importance in the associated deduction, or more systematically weighted using a consolidated set of model-wide parameters.

An Implemented DBN User Knowledge Model

A preliminary dynamic Bayesian network for modeling user knowledge has been implemented in the CRYSTAL ISLAND interactive narrative environment. After describing CRYSTAL ISLAND, as well as an initial implementation of the DBN knowledge model, results from an empirical study using data collected from user interactions with the CRYSTAL ISLAND environment are presented.

CRYSTAL ISLAND narrative environment

Now in its third major iteration, CRYSTAL ISLAND (Figure 2) is an interactive narrative-centered learning environment built on Valve Software's Source™ engine, the 3D game platform for Half-Life 2. Students play the role of the protagonist, Alyx, who is attempting to discover the identity and source of an infectious disease plaguing a newly established research station. Several of the team's members have fallen gravely ill, and it is the student's task to discover the nature and cause of the outbreak.

CRYSTAL ISLAND's narrative takes place in a small research camp situated on a recently discovered tropical island. Students investigate the island's spreading illness by forming questions, generating hypotheses, collecting data, and testing hypotheses. Throughout their investigations, students interact with virtual characters offering clues and microbiology facts via multimodal "dialogues" delivered through student menu choices and spoken language. The dialogues' content is supplemented by virtual books, posters, and other resources encountered in several of the camp's locations. To solve the mystery, students complete a *diagnosis worksheet* to manage their working hypotheses and record findings about patients' symptoms and medical history, as well as any findings from tests conducted in the camp's laboratory.

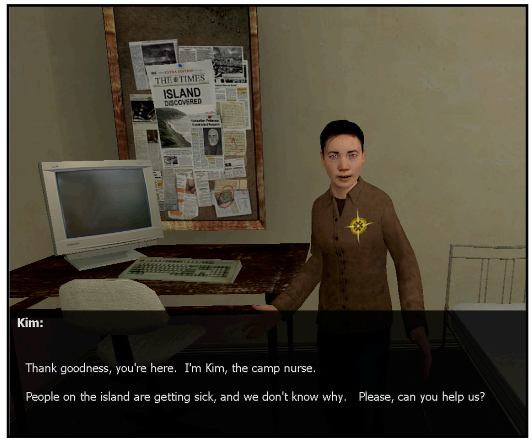

Figure 2. CRYSTAL ISLAND narrative environment.

Implementation

The dynamic Bayesian network for knowledge tracing has been implemented in C++ using the SMILE Bayesian modeling and inference library developed by the University of Pittsburgh's Decision Systems Laboratory (Druzdzel, 1999). The model maintains approximately 135 binary nodes, 100 directed links, and more than 750 conditional probabilities. As the knowledge-tracing model observes student actions in the environment, the associated evidence is incorporated into the network, and a Bayesian update procedure is performed. The update procedure, in combination with the network's singly-connected structure, yields updates that complete in less than one second. Initial probability values were fixed across all students; probabilities were chosen to represent the assumption that students at the onset had no understanding of scenario-specific knowledge components and were unlikely to have mastery of curriculum knowledge components.

Preliminary Empirical Evaluation

A human participant study was conducted in Fall 2007 with 116 eighth-grade students from a North Carolina middle school interacting with the CRYSTAL ISLAND environment. During the study, students interacted with CRYSTAL ISLAND for approximately fifty minutes. Logs of students' in-game actions were recorded, and have been subsequently used to conduct a preliminary evaluation of the DBN knowledge-tracing model. The evaluation aims to assess the model's ability to accurately predict students' performance on a content knowledge post-experiment test. While the evaluation does not inspect the user knowledge model's intermediate states during the narrative interaction, nor students' narrative-specific knowledge, an evaluation of the model's final knowledge assessment accuracy provides a useful starting point for refining the model and evaluating its accuracy.

The DBN knowledge-tracing model used the students' recorded trace data as evidence to approximate students' knowledge at the end of the learning interaction. This yielded a set of probability values for each student, corresponding to each of the knowledge-tracing model's knowledge components. The resultant data was used in the analysis of the model's ability to accurately predict student responses on post-experiment content test questions. The mapping between the model's knowledge components and individual posttest questions was generated by a researcher, and used the following heuristic: if a posttest question or correct response shared important content terms with the description of a particular knowledge component, that knowledge component was designated as necessary for providing an informed, correct response to the question. According to this heuristic, several questions required the simultaneous application of multiple knowledge components, and a number of knowledge components bore on multiple questions. This yielded a many-to-many mapping between knowledge components and posttest questions.

The evaluation procedure required the definition of a threshold value to discriminate between mastered and un-mastered knowledge components: knowledge components whose model values exceeded the threshold were considered *mastered*, and knowledge components whose model values fell below the threshold were considered *un-mastered*. The model predicted a correct response on a posttest question if all of the question's associated knowledge components were considered mastered. The model predicted an incorrect response on a posttest question if one or more associated knowledge components were considered un-mastered. The use of a threshold to discriminate between mastered and un-mastered knowledge components mirrors how the knowledge model might be used in a runtime environment to inform interactive narrative decision-making.

Rather than choose a single threshold, a series of values ranging between 0.0 and 1.0 were selected. For each threshold, the DBN knowledge model was compared to a random model, which assigned uniformly distributed, random probabilities for each knowledge component. New random probabilities were generated for each knowledge component, student, and threshold. Both the DBN model and random model were used to predict student posttest responses, and accuracies for each threshold were determined across the entire test. Accuracy was measured as the sum of successfully predicted correct responses plus the number of successfully predicted incorrect responses, divided by the total number of questions. The results of this analysis are displayed in Figure 3.

The DBN model outperformed the random baseline model across a range of threshold values. The DBN model most accurately predicted students' posttest responses at a threshold level of 0.32 (M = .594, SD = .152). A Wilcoxon-Mann-Whitney U test verified that the DBN knowledge-tracing model was significantly more accurate than the random model at the 0.32 threshold level, $z = 4.79$, $p < .0001$. Additional Mann-Whitney tests revealed that the DBN model's predictive accuracy was significantly greater than that of the random model, at the $\alpha = .05$ level, for the entire range of thresholds between .08 and .56.

Discussion

An evaluation of the implemented user knowledge model's predictive accuracy on the post-experiment content test

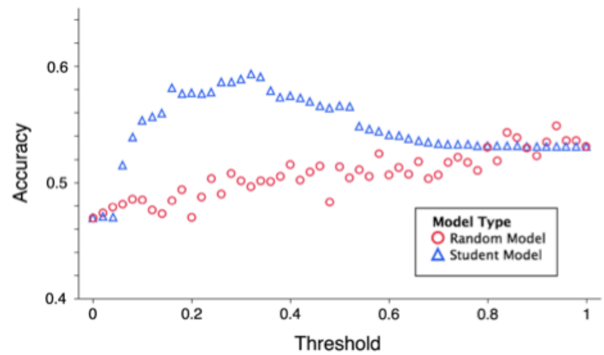

Figure 3. Comparison of DBN knowledge model accuracy vs. Random model accuracy for post-test answer prediction across mastery threshold values.

shows initial promise. The DBN model significantly outperformed a baseline model at predicting student responses on the content test across a range of mastery threshold levels. The predictive accuracy gains were modest, but the findings provide a meaningful floor against which future models of user knowledge in interactive narratives can be compared. Although the range of thresholds for which the knowledge-tracing model outperformed the random model may appear low, the majority of the model's knowledge component values in fact fell between .15 and .40. The model was most effective at predicting student test responses at thresholds between .08 and .56, which is the same range in which the majority of knowledge component values fell.

The range of observed values could be attributed to several factors. The values might stem from an underlying issue in the model's parameters. Hand-authoring Bayesian network structures and conditional probabilities can be challenging and resource intensive; machine learning model parameters, using techniques such as expectation maximization, is a promising direction for future work. An alternative explanation is inherent weakness in the trace data's bandwidth. Weak evidence would elicit correspondingly low probability values for the model's knowledge components, and would be consistent with the modest but significant boost in predictive accuracy encountered in the evaluation. Identifying gameplay actions that both support the narrative and provide increased evidence of user knowledge is an important direction for future investigation.

Conclusions and Future Work

This work represents an initial step toward devising accurate user knowledge models to inform runtime story adaptations in interactive narrative environments. A preliminary dynamic Bayesian network knowledge model has been implemented for the CRYSTAL ISLAND narrative-centered learning environment. The model maintains probabilistic assessments of students' knowledge as they navigate a story-centric problem-solving scenario. An empirical evaluation indicates that the model assesses users' content knowledge more accurately than a baseline approach, and introduces opportunities for further improvements in user knowledge modeling for interactive narratives.

Several directions for future work appear promising. First, enhancements to the narrative environment, logging facilities, and user knowledge model can address shortcomings in the model's knowledge components and network structures, as well as increase the bandwidth of trace data elicited from student interactions. Second, machine learning techniques for inducing DBN conditional probabilities from data can be investigated to address the considerable hand authoring requirements of the current methodology. Finally, the knowledge model can be directly incorporated into an interactive narrative planner to drive story decisions in a runtime environment. The model could be incorporated into a runtime narrative planner by extending rule-based and plan-based drama managers to query user knowledge models when considering candidate narrative adaptations. Alternatively, DBN knowledge-tracing models lend themselves to integration in decision-theoretic director agent architectures, particularly those that use dynamic decision network approaches.

Acknowledgments

The authors wish to thank members of the IntelliMedia Group of North Carolina State University for their assistance, Omer Sturlovich and Pavel Turzo for use of their 3D model libraries, and Valve Software for access to the Source™ engine and SDK. This research was supported by the National Science Foundation under Grants REC-0632450 and DRL-0822200. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

References

Barber, H., and Kudenko, D. 2007. Dynamic Generation of Dilemma-based Interactive Narratives. In *Proceedings of the Third Artificial Intelligence for Interactive Digital Entertainment Conference*, 2–7. Menlo Park, CA: AAAI Press.

Conati, C., Gertner, A., and VanLehn, K. 2002. Using Bayesian Networks to Manage Uncertainty in Student Modeling. *User Modeling and User-Adapted Interaction 12*(4): 371–417.

Corbett, A., and Anderson, J. 1994. Knowledge Tracing: Modeling the Acquisition of Procedural Knowledge. *User Modeling and User-Adapted Interaction 4*(4): 253–278.

Dean, T. and Kanazawa, K. 1989. A Model for Reasoning About Persistence and Causation. *Computational Intelligence 5*(3): 142–150.

Druzdzel, M. 1999. SMILE: Structural Modeling, Inference, and Learning Engine and Genie: A Development Environment for Graphical Decision-Theoretic Models. In *Proceedings of the Sixteenth National Conference on Artificial Intelligence,* 342–343. Menlo Park, CA: AAAI Press.

Magerko, B. 2006. Player Modeling in the Interactive Drama Architecture. Ph.D. diss., Dept. of Computer Science, University of Michigan, Ann Arbor, MI.

Manske, M. and Conati, C. 2005. Modelling Learning in an Educational Game. In *Proceedings of the Twelfth International Conference on Artificial Intelligence in Education,* 411–418. Amsterdam, The Netherlands: IOS Press.

Mateas, M. and Stern, A. 2005. Structuring Content in the Façade Interactive Drama Architecture. In *Proceedings of the First Artificial Intelligence in Interactive Digital Entertainment Conference,* 93–98. Menlo Park, CA: AAAI Press.

Mott, B., Lee, S., and Lester, J. 2006. Probabilistic Goal Recognition in Interactive Narrative Environments. In *Proceedings of the Twenty-First National Conference on Artificial Intelligence,* 187–192. Menlo Park, CA: AAAI Press.

Pearl, J. 1988. *Probabilistic Reasoning in Intelligent Systems*. San Mateo, California: Morgan Kaufmann.

Riedl, M., and Stern, A. 2006. Believable Agents and Intelligent Scenario Direction for Social and Cultural Leadership Training. In *Proceedings of the Fifteenth Conference on Behavior Representation in Modeling and Simulation.*

Roberts, D., Nelson, M., Isbell, C., Mateas, M., and Littman, M. 2006. Targeting Specific Distributions of Trajectories in MDPs. *Twenty-First National Conference on Artificial Intelligence,* 1213–1218. Menlo Park, CA: AAAI Press.

Si, M., Marsella, and S., Pynadath, D. 2006. Thespian: Modeling Socially Normative Behavior in a Decision-Theoretic Framework. In *Proceedings of the Sixth International Conference on Intelligent Virtual Agents*, 369–382. Berlin/Heidelberg: Springer.

Thue, D., Bulitko, V., Spetch, M., and Wasylishen, E. 2007. Interactive Storytelling: A Player Modelling Approach. In *Proceedings of the Third Artificial Intelligence for Interactive Digital Entertainment Conference,* 43–48. Menlo Park, CA: AAAI Press.

VanLehn, K. 2006. The Behavior of Tutoring Systems. *International Journal of Artificial Intelligence in Education 16*(3): 227–265.