

Probabilistic Goal Recognition in Interactive Narrative Environments

Bradford Mott, Sunyoung Lee, and James Lester

Department of Computer Science
North Carolina State University
Raleigh, NC 27695, USA
{bwmott, slee7, lester}@ncsu.edu

Abstract

Recent years have witnessed a growing interest in interactive narrative-centered virtual environments for education, training, and entertainment. Narrative environments dynamically craft engaging story-based experiences for users, who are themselves active participants in unfolding stories. A key challenge posed by interactive narrative is recognizing users' goals so that narrative planners can dynamically orchestrate plot elements and character actions to create rich, customized stories. In this paper we present an inductive approach to predicting users' goals by learning probabilistic goal recognition models. This approach has been evaluated in a narrative environment for the domain of microbiology in which the user plays the role of a medical detective solving a science mystery. An empirical evaluation of goal recognition based on n -gram models and Bayesian networks suggests that the models offer significant predictive power.

Introduction

Interactive narrative environments have been the subject of growing attention in recent years (Swartout *et al.* 2001; Riedl, Saretto, and Young 2003; Magerko *et al.* 2004; Mateas and Stern 2005; Si, Marsella, and Pynadath 2005). With the promise of dynamically crafting engaging stories with plots featuring a cast of believable characters whose behaviors are tailored for individual users, interactive narrative environments offer significant potential for education, training, and entertainment. We are beginning to see the appearance of increasingly sophisticated computational models of narrative generation. This work has explored techniques for mediating user interactions in narrative environments (Riedl, Saretto, and Young 2003), search paradigms for encoding author aesthetics with evaluation functions (Weyhrauch 1997), and creating believable characters for narrative environments (Loyall *et al.* 2004).

Despite recent advances, narrative environments currently lack a key functionality that could significantly increase their ability to create compelling stories: they

cannot recognize users' goals as they interact with virtual storyworlds. Providing narrative planners with the ability to accurately recognize users' goals would enable planners to ascertain whether a user's actions were threatening a plot (Riedl, Saretto, and Young 2003), take user goals into account in plot construction (Harris and Young 2005), and ensure that stories steadily progress (Magerko *et al.* 2004). However, narrative goal recognition is challenging because interactive narratives play out in highly dynamic virtual environments in which users have the freedom to explore complex storyworlds as they perform a broad range of actions in pursuit of their quests.

Goal recognition for interactive narrative should satisfy three requirements. First, because incorrectly predicting goals could significantly diminish the effectiveness of narrative planners, narrative goal recognizers should accurately infer users' goals. Moreover, as observations of users' activities become available, recognizers should make accurate "early" predictions (Blaylock and Allen 2003) – ideally these would be k -best predictions rather than a single predicted goal – and they should converge as quickly as possible on the most likely interpretation. Second, the real-time requirements of interactive narrative call for extraordinarily efficient recognizers. Any approach that depends on computations spanning more than a few milliseconds could be infeasible. Third, users in most narrative environments should not be interrogated about their current goal; their actions must speak for themselves. Because interrupting users to pose questions about their goals could interfere with the flow of the narrative and cause users to forgo their suspension of disbelief, a narrative goal recognizer should perform "keyhole goal recognition," i.e., it should unobtrusively observe a user as she interacts with the environment.

In this paper, we present an inductive approach to predicting users' goals in dynamic narratives by learning probabilistic goal recognition models. We report on two families of goal recognition models: n -gram models (unigrams and bigrams) and Bayesian network models. The models, which exploit knowledge of narrative structure as well as locational information about users' activities in the world, are induced from training data acquired from traces of users' performing narrative quests in a storyworld. We report on the empirical evaluation of probabilistic goal recognition models in a narrative-

centered environment for the domain of microbiology in which users play the role of a medical detective solving a science mystery. Experimental results suggest that probabilistic models can accurately predict users' goals, and that they converge on correct interpretations as observations of a user's activities become available over time. Probabilistic goal recognizers are also sufficiently efficient to meet the real-time performance demands of interactive narrative.

Related Work

A key challenge posed by interactive narrative is user goal recognition. Providing narrative planners with the ability to recognize users' goals would enable planners to monitor users to determine if their goals were consistent with the plot and if sufficient plot progress were being made (Magerko *et al.* 2004). It would also enable them to detect whether the user was, either intentionally or inadvertently, interfering with critical plot objectives, which in turn could damage the story arc and thereby detract from the user's overall experience. User goal recognition could contribute to narrative planners' performing *reactive mediation* (Riedl, Saretto, and Young 2003), in which exceptional user actions are accommodated by changes to the plot, or in which a narrative planner intervenes to thwart undesirable plot threats. User goal recognition could also contribute to *proactive mediation* (Harris and Young 2005), in which a narrative planner seamlessly fuses a user's plan with the yet-to-be-executed narrative plan to steer clear of exceptional events. Moreover, it could support advice delivered by a narrator or by characters. In short, having a clear picture of the current narrative landscape by accurately inferring users' intentions could support a broad range of interactive story creation techniques.

The problem of goal recognition (Lesh 1998; Blaylock and Allen 2003) is a restricted form of the *plan recognition* problem. For decades, the natural language, user modeling, and agents communities have grappled with the problem of plan recognition (Carberry 1990; Charniak and Goldman 1993; Huber and Hadley 1997). Given a sequence of actions, plan recognition seeks to infer the plans that an agent is attempting to execute. In narrative understanding, plan recognition is used to infer characters' goals from their actions in stories (Charniak and Goldman 1993); in natural language processing, it is used to infer users' task-related plans for dialogue systems (Carberry 1990). Recent work has explored plan recognition for a variety of tasks ranging from security and collaborative filtering to robotics and software assistants. *Goal recognition* is the special case of plan recognition that focuses on inferring an agent's goals (Blaylock and Allen 2003), i.e., the specific objectives that the agent is attempting to achieve.

Because goal recognition and plan recognition are characterized by uncertainty, probabilistic solutions such as Bayesian models (Charniak and Goldman 1993) and

probabilistic grammars (Pynadath and Wellman 2000) have been investigated. Probabilistic solutions not only deal effectively with the inherent uncertainty, they may also circumvent the issues associated with manually constructing plan libraries, a labor intensive task that is prone to errors and whose scalability is questionable (Albrecht, Zukerman, and Nicholson 1998). Hidden Markov models have been used to recognize stereotypical team behaviors in an Unreal Tournament environment (Sukthankar and Sycara 2005), and Bayesian networks for off-line plan recognition have been studied for predicting a player's quest in a text-based Multi-User Dungeon from her current goal, action, and location (Albrecht, Zukerman, and Nicholson 1998). A promising utility-based approach to plan recognition has been proposed (Mao and Gratch 2004), though evaluation results have not yet been reported. Another approach focused on compiling agents' plans into Bayesian networks, which were subsequently used in plan recognition to support multi-agent coordination within a real-time arcade-style game (Huber and Hadley 1997). However, narrative goal recognition for the 3D virtual environments of interactive storyworlds has not previously been empirically investigated.

Probabilistic Narrative Goal Recognition

Interactive narratives play out in highly dynamic storyworlds in which users perform actions to achieve goals in the unfolding stories. Narrative goal recognizers can exploit three sources of information to infer users' goals:

- *Narrative State*: Narrative goal recognizers have intimate knowledge of a rich representation of the narrative, including the plot (typically represented in a *plot graph* (a partially ordered graph of plot elements) (Weyhrauch 1997)) or *narrative plan* (Riedl, Saretto, and Young 2003), the current focus of the story arc and its episodic structure, and the plans and goals of the synthetic agents who serve as (the other) characters in the story.
- *User Actions*: Narrative goal recognizers can observe users' actions in the world; recognizers also have access to auxiliary information about the interactions, e.g., any artifacts manipulated such as which objects have been picked up or which doors have been opened, as well as the characters with which users have interacted.
- *User Location*: Narrative goal recognizers can bring to bear a broad range of knowledge about the location in which users' actions are performed in virtual environments. In contrast to activity recognition in physical environments where recognizers must cope with noise and errors in sensors and perception (e.g., vision and speech), narrative goal recognition has access to precise locational information.

More formally, we define narrative goal recognition as follows: Given a sequence of n observed user actions a_1 ,

a_2, \dots, a_n in a narrative environment, their associated narrative states n_1, n_2, \dots, n_n and user locations l_1, l_2, \dots, l_n , identify the most likely goal G^* from a set of candidate goals g_1, g_2, \dots, g_m that accounts for the action sequence in the given context.

To address the requirements for narrative goal recognition set forth above (accuracy, incremental recognition, and efficiency), and to cope with the uncertainty inherent in recognizing users' goals in interactive narrative environments, we investigate two families of probabilistic approaches to user goal recognition: n -gram models and Bayesian networks.

n -gram Models for Narrative Goal Recognition

Given an observation sequence O_1, O_2, \dots, O_n , the objective of narrative goal recognition is to identify the most likely goal G^* such that:

$$\begin{aligned} G^* &= \arg \max P(G | O_1, O_2, O_3, \dots, O_n) \\ &= \arg \max P(G | O_{1:n}) \end{aligned}$$

where each O_i is an observation encoding the current narrative state, the user's action, and the location at which her action was performed. The observation sequence O_1, O_2, \dots, O_n is denoted by $O_{1:n}$. Applying Bayes' rule yields:

$$G^* = \arg \max \frac{P(O_{1:n} | G)P(G)}{P(O_{1:n})}$$

which can be simplified by eliminating the constant term $P(O_{1:n})$ to obtain:

$$G^* = \arg \max P(O_{1:n} | G)P(G)$$

Applying the Chain Rule, the equation becomes:

$$\begin{aligned} G^* &= \arg \max P(O_n | O_{1:n-1}, G)P(O_{n-1} | O_{1:n-2}, G) \\ &\quad \cdot P(O_{n-2} | O_{1:n-3}, G) \dots P(O_1 | G)P(G) \end{aligned}$$

However, estimating these conditional probabilities is impractical – it would require exponentially large training data sets – so we make a Markov assumption that an observation O_i depends only on the goal G and a limited window of preceding observations.

Following an approach initially proposed for goal recognition in natural language dialogue (Blaylock and Allen 2003), we explore two n -gram narrative goal recognition models, a unigram model and a bigram model. The unigram model is based on the assumption that, given the goal G , O_i is conditionally independent of all other observations. Thus, the goal recognition formula for the unigram model can be simplified to:

$$G^* = \arg \max P(G) \prod_{i=1}^n P(O_i | G)$$

The bigram model is based on the assumption that, given the goal G and the preceding observation O_{i-1} , O_i is conditionally independent of all other observations. Thus, the goal recognition formula for the bigram model can be simplified to:

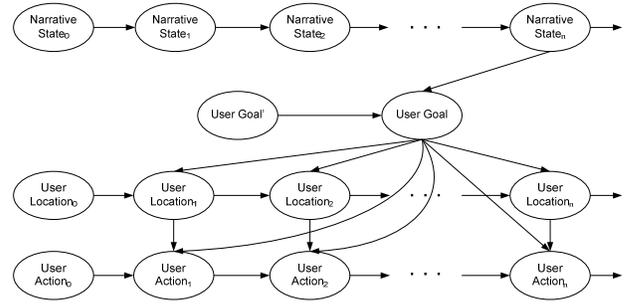


Figure 1. Bayesian Network Goal Recognition Model

$$G^* = \arg \max P(G) \prod_{i=1}^n P(O_i | O_{i-1}, G)$$

For the bigram model, O_0 is taken to be the *null* observation when a narrative begins and the previous narrative state, user action, and user location are all *null*. The resulting formulae for the unigram and bigram models are very efficient because updating the goal prediction for each new observation only requires computing the product of the probability returned by the previous prediction and the current conditional probability.

During training, we estimate $P(G)$, $P(O_i | G)$, and $P(O_i | O_{i-1}, G)$ using training data acquired with a narrative environment as described below. Because training data is necessarily sparse, i.e., we are unlikely to observe all narrative states, actions, locations, and goals, the unigram and bigram models employ a standard smoothing technique (a flattening constant) to reevaluate zero-probability and low-probability n -grams.

Bayesian Network Narrative Goal Recognition

Narrative goal recognition can alternatively be modeled with Bayesian networks. Following an approach initially proposed for keyhole plan recognition in a text-based adventure game (Albrecht, Zukerman, and Nicholson 1998), we explore a Bayesian network model for narrative goal recognition (Figure 1). In contrast to the aggregate observation variables O_i of the n -gram models, the Bayesian network goal recognizer explicitly models dependencies between the constituent variables, i.e., between narrative state, user action, and user location. Thus, it represents the influences of the following variables on the user's goal G : the user's previous goal G' , the sequence of narrative states n_1, n_2, \dots, n_n , the sequence of user actions a_1, a_2, \dots, a_n , and the sequence of user locations l_1, l_2, \dots, l_n . While the narrative states influence G because the plot elements and story arc affect the goals of the user in the story, G itself influences the locations where the user performs her actions, as well as the actions that she performs there. The locations also influence the user actions directly.¹ As with the bigram model, N_0, A_0 , and L_0 are taken to be *null* when a narrative begins.

¹ Locations are modeled as influencing actions because particular locations *afford* particular types of actions. However, the converse model

From the Bayesian network, we have the following:

$$G^* = \arg \max P(G | G', A_{0:n}, L_{0:n}, N_{0:n})$$

Given a sequence of actions $A_0, A_1, A_2, \dots, A_n$, a sequence of locations $L_0, L_1, L_2, \dots, L_n$, and a sequence of narrative states $N_0, N_1, N_2, \dots, N_n$, the current action A_n depends only on the previous action A_{n-1} , the current location L_n , and the current goal G ; the current location L_n depends only on the previous location L_{n-1} and the current goal G ; the current narrative state N_n depends only on the previous narrative state N_{n-1} ; and the current goal G only depends on the previous goal G' and the current narrative state N_n . Therefore:

$$P(A_n | G, G', A_{0:n-1}, L_{0:n}, N_{0:n}) = P(A_n | G, A_{n-1}, L_n)$$

$$P(L_n | G, G', A_{0:n-1}, L_{0:n-1}, N_{0:n}) = P(L_n | G, L_{n-1})$$

$$P(N_n | G', A_{0:n-1}, L_{0:n-1}, N_{0:n-1}) = P(N_n | N_{n-1})$$

$$P(G | G', A_{0:n-1}, L_{0:n-1}, N_{0:n}) = P(G | G', N_n)$$

where $n \geq 1$.

By applying the Chain Rule and the above equations, the goal recognition formula for the Bayesian network becomes:

$$\begin{aligned} G^* &= \arg \max P(G | G', A_{0:n}, L_{0:n}, N_{0:n}) \\ &= \arg \max P(G') P(G | G', N_n) P(A_0) P(L_0) P(N_0) \\ &\quad \cdot \prod_{i=1}^n P(A_i | G, A_{i-1}, L_i) P(L_i | G, L_{i-1}) P(N_i | N_{i-1}) \end{aligned}$$

As with the n -gram models, the resulting formula is very efficient because updating the goal prediction for each new observation requires only noting the previous prediction, looking up a few CPT entries, and computing their product. During training, we estimate $P(A_i | G, A_{i-1}, L_i)$, $P(L_i | G, L_{i-1})$, $P(N_i | N_{i-1})$, $P(G | G', N_i)$, and $P(G)$ using training data acquired with a narrative environment as described below.

Evaluation

While the accuracy, incremental recognition capabilities, and efficiency of a proposed approach to narrative goal recognition can be analytically evaluated, because narrative is an HCI phenomenon that is the by-product of users' interactions with narrative planners in storyworlds, it is instructive to empirically investigate these issues. We adopt the following 3-phase evaluation methodology to study the probabilistic narrative goal recognizers introduced above.

1. **Narrative Quest Trace Acquisition:** In an interactive narrative environment, collect traces of

users' performing narrative quests. Quest traces encode extensive sequences of narrative states and user goals, locations, and actions.

2. **Goal Recognizer Induction:** Learn narrative goal recognizers from the quest traces by using the narrative state sequences, user location sequences, and user action sequences to induce goal classifiers.
3. **Predictive Recognition Evaluation:** With a cross validation approach, determine the accuracy and the incremental recognition abilities of each goal recognizer.

This methodology has been used to study the unigram, bigram, and Bayesian network narrative goal recognition models. After briefly describing the narrative environment testbed in which the experiments were carried out, we describe the collection of narrative quest traces, and report the experimental results.

Crystal Island Storyworld

To serve as an effective "laboratory" for studying user goal recognition in interactive narrative, a testbed should pose the same kinds of challenges that goal recognizers are likely to encounter in future interactive narrative environments. It should offer users a broad range of actions to perform and provide a rich storyworld in a non-trivial 3D environment. The narrative should exhibit some complexity, and the storyworld should be populated by manipulable artifacts and be inhabited by multiple characters. To this end, we have devised CRYSTAL ISLAND, a narrative environment testbed in the science mystery genre (Figure 2).

The CRYSTAL ISLAND testbed environment features a science mystery set on a recently discovered volcanic island where a research station has been established to study the unique flora and fauna. The user plays the protagonist who is attempting to discover the origins of an unidentified illness at the research station. The story opens by introducing her to the island and the members of the research team for which her father serves as the lead scientist. As members of the research team fall ill, it is her task to discover the cause of the outbreak. She is free to explore the world to collect physical evidence and interact with other characters. Through the course of her adventure she must gather enough evidence to correctly choose among candidate diagnoses including botulism, cholera, giardiasis, paralytic shellfish poisoning, salmonellosis, and tick paralysis as well as identify the source of the disease.

The virtual world of CRYSTAL ISLAND, the semi-autonomous characters that inhabit it, and the user interface were implemented with Valve Software's Source™ engine, the 3D game platform for Half-Life 2. The Source engine also provides much of the low-level (reactive) character behavior control. The narrative planner of CRYSTAL ISLAND has been implemented with an HTN planner. In CRYSTAL ISLAND, the user can perform a broad range of actions including performing experiments in the laboratory, interacting with other characters, reading

in which actions influence locations is also plausible because a user might journey to a location in the storyworld to perform an intended action there. Since Bayes nets are by definition acyclic, one of these directions of causality must be selected. While the former is chosen here, the latter offers an interesting direction for future work.



Figure 2. The CRYSTAL ISLAND Testbed

“virtual books” to obtain background information on diseases, and collecting data about the food recently eaten by the members of the research team. Throughout the mystery, users can walk around the island and visit the infirmary, the lab, the dining hall, and the living quarters of each member of the team. In the current testbed, there are twenty goals the users can achieve, three hundred unique actions the user can perform, and over fifty unique locations in which the actions can be performed.

To obtain baseline readings on the overall entertainment value of the CRYSTAL ISLAND environment, a formal user evaluation was conducted with fifty-three subjects. Given that the bar for interactive entertainment is relatively high for the college-age subjects in the study, subjects’ rating their experiences as “enjoyable” (Mean 3.89/5.00, SD 0.80) on a five-point Likert scale suggests that the system is reasonably entertaining. Similar results were obtained for “engagement,” “taking time to explore the virtual environment,” and “interacting with characters to learn more about them.” Exit interviews revealed that the users clearly enjoyed their interactive experience.

Narrative Quest Traces

In a formal evaluation, more than fifteen hours of narrative trace data was collected from forty subjects interacting with the testbed environment. Subjects were first situated in the narrative world and given an overview of the kinds of activities they could perform. Next, they were told how their character could be controlled, and once they entered the virtual environment, they were successively given goals they were expected to achieve via an onscreen message. Based on the narrative structure, the order of the goals differed from session to session. Users completed each goal, and upon completing the final goal, they were complimented on their successful performance. Detailed quest traces were recorded of all sequences of actions, goals, locations, and narrative states. Narrative states were represented with the episodic structure of the unfolding story and the narrative arc in which it was situated. There were eighty training sessions collected (two sessions per subject), which generated just over twenty thousand

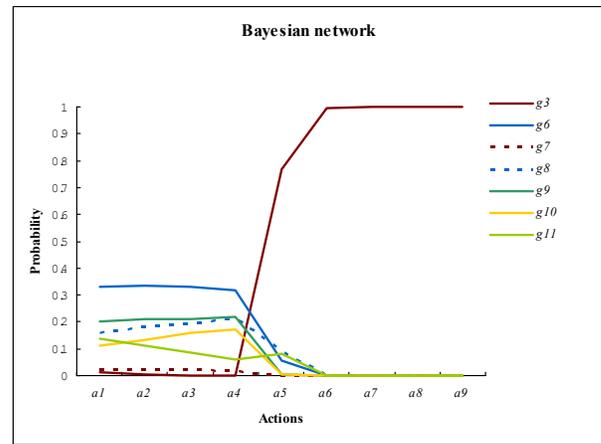


Figure 3. Bayesian Network Convergence Graph

training records. The number of training records was high because of the frequency of sampling and the length of action sequences per goal.

Results

Unigram, bigram, and Bayesian network goal recognition models were learned from the resulting traces. For the Bayesian network model, the structure was fixed (as shown in Figure 1); the conditional probabilities were learned. Smoothing was used as described above. The goal recognition models were evaluated using the following criteria (Blaylock and Allen 2003):

- *Accuracy*: The ratio of correct predictions to the total number of observations.
- *Converged*: The percentage of observation sequences in which the goal recognizer’s final prediction is correct.
- *Convergence Point*: For observation sequences which converged, the point within the sequence when the goal recognizer started making the correct prediction and continued to make the correct prediction for the remainder of the sequence.
- *Average Actions of Converged*: The average number of actions within observation sequences which converged.

The induced models were tested using a 10-fold cross validation. (In each fold, nine segments are used for training and one, which was not used for training, is used for testing.) The results are shown in Table 1.

	<i>Unigram</i>	<i>Bigram</i>	<i>Bayesian</i>
<i>Accuracy</i>	54.8%	51.5%	53.7%
<i>Converged</i>	83.7%	79.3%	67.2%
<i>Convergence Point</i>	50.5%	48.5%	43.9%
<i>Avg Actions of Converged</i>	16.3	16.9	17.1

Table 1. Evaluation Results

Each of the models performed at a reasonable level. Although the 51% to 54% accuracy may at first appear low, the recognizers performed significantly better than chance, which would be 5%. (There were 20 candidate goals.) The resulting accuracies are promising and are consistent with the results seen in other statistical

approaches, e.g., (Blaylock and Allen 2003). In addition to accuracy, the models also exhibited the ability to make early predictions, i.e., predictions based on a few observations, and to converge reasonably quickly to the correct goal. Figure 3 shows the top seven predictions of the Bayesian goal recognition model as it converges to g_3 after five actions. Because all of the models are probabilistic, they can provide the k -best predictions.

Conclusion

User goal recognition is a central problem in interactive narrative. Equipping narrative planners with the ability to infer users' intent on a moment-by-moment basis could contribute to their ability to craft story experiences that are all the more engaging. With effective goal recognition, narrative planners might be able to more accurately assess plot threats, opportunistically interleave narrative objectives with user goals in dynamic plot construction, and ensure that steady plot progress is made. Probabilistic approaches to recognizing users' goals can cope with the uncertainty inherent in the task. They also offer the advantage of being automatically acquired rather than being manually constructed. Empirical studies of two families of probabilistic narrative recognizers, n -gram models and Bayesian networks, suggest that probabilistic approaches can perform keyhole user goal recognition that is accurate and incrementally converging.

This paper represents a first step towards goal recognition for interactive narrative environments. Several directions for future research appear promising. First, the models investigated here are based on the simplifying assumption that the user is pursuing a single goal. Users often pursue more than one goal at the same time, so models accommodating multiple simultaneous goals need to be studied. Second, it will be interesting to develop techniques for dynamically relaxing the keyhole requirement. In some narrative situations, it is appropriate for a character to approach the user and ask her what she is doing. It will be important to create techniques for identifying such situations and integrating the resulting information into the goal monitoring system. Third, while narrative goal recognition provides an important source of information about the user to narrative planners, narrative plan recognition would also be beneficial, particularly in domains with large plan spaces.

References

- Albrecht, D., Zukerman, I., Nicholson, A. 1998. Bayesian Models for Keyhole Plan Recognition in an Adventure Game. *UMUAI* 8(1-2): 5-47.
- Blaylock, N. and Allen, J. 2003. Corpus-based, Statistical Goal Recognition. In *Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence*, 1303-1308, Acapulco, Mexico.
- Carberry, S. 1990. *Plan Recognition in Natural Language Dialogue*. Cambridge, MA: The MIT Press.
- Charniak, E. and Goldman, R. 1993. A Bayesian Model of Plan Recognition. *Artificial Intelligence* 64(1):53-79.
- Harris, J. and Young, M. 2005. Proactive Mediation in Plan-Based Narrative Environments. In *Proceedings of the Fifth International Conference on Intelligent Virtual Agents*, 292-304. Kos, Greece.
- Huber, M. and Hadley, T. 1997. Multiple Roles, Multiple Teams, Dynamic Environment. In *Proceedings of the First International Conference on Autonomous Agents*, 332-339, Marina del Rey, CA.
- Lesh, N. 1998. Scalable and Adaptive Goal Recognition. Ph.D. diss., Computer Science & Engineering, University of Washington.
- Loyall, B., Reilly, S., Bates, J., and Weyhrauch, P. 2004. System for Authoring Highly Interactive, Personality-Rich Interactive Characters. In *Proceedings of the 2004 ACM SIGGRAPH / Eurographics Symposium on Computer Animation*, 59-68. Grenoble, France.
- Magerko, B., Laird, J., Assanie, M., Kerfoot, A., and Stokes, D. 2004. AI Characters and Directors for Interactive Computer Games. In *Innovative Applications of Artificial Intelligence Conference*, 877-883. San Jose, CA.
- Mao, W. and Gratch, J. 2004. A Utility-Based Approach to Intention Recognition. In *Proceedings of the AAMAS 2004 Workshop on Agent Tracking: Modeling Other Agents from Observations*. New York, NY.
- Mateas, M. and Stern, A. 2005. Structuring Content in the Façade Interactive Drama Architecture. In *Proceedings of AIIDE 2005*, 93-98. Marina del Rey, CA.
- Pynadath, D. and Wellman, M. 2000. Probabilistic State-Dependent Grammars for Plan Recognition. In *Proceedings of the 16th Annual Conference on Uncertainty in Artificial Intelligence*, 507-514. San Francisco, CA.
- Riedl, M., Saretto, C., and Young, M. 2003. Managing Interaction Between Users and Agents in a Multi-agent Storytelling Environment. In *Proceedings of the Second International Conference on Autonomous Agents and Multi-Agent Systems*, 741-748. Melbourne, Australia.
- Si, M., Marsella, S., and Pynadath, D. 2005. Thespian: Using Multi-Agent Fitting to Craft Interactive Drama. In *Proceedings of AAMAS 2005*, 21-28. Utrecht, Netherlands.
- Sukthankar, G. and Sycara, K. 2005. Automatic Recognition of Human Team Behaviors. In *Proceedings of the International Joint Conference on Artificial Intelligence Workshop on Modeling Others from Observations*. Edinburgh, Scotland.
- Swartout, W., Hill, R., Gratch, J., Johnson, L., et al. 2001. Towards the Holodeck: Integrating Graphics, Sound, Character and Story. In *Proceedings of the Fifth International Conference on Autonomous Agents*, 409-416. Montreal, Canada.
- Weyhrauch, P. 1997. Guiding Interactive Drama. Ph.D. diss., Dept. of Computer Science, Carnegie Mellon Univ.