Intelligent Narrative Technologies: Papers from the 2012 AIIDE Workshop
AAAI Technical Report WS-12-14

Toward Narrative Schema-Based Goal Recognition Models for Interactive Narrative Environments

Alok Baikadi, Jonathan P. Rowe, Bradford W. Mott, James C. Lester
Department of Computer Science
North Carolina State University
Raleigh, NC 27695
{abaikad, jprowe, bwmott, lester}@ncsu.edu

Abstract

Computational models for goal recognition hold great promise for enhancing the capabilities of drama managers and director agents for interactive narratives. The problem of goal recognition, and its more general form, plan recognition, have been the subjects of extensive investigation in the AI community. However, relatively little effort has been undertaken to examine goal recognition in interactive narrative. In this paper, we propose a research agenda to improve the accuracy of goal recognition models for interactive narratives using explicit representations of narrative structure inspired by the natural language processing community. We describe a particular category of narrative representations, *narrative schemas*, that we anticipate will effectively capture patterns of player behavior in interactive narratives and improve the accuracy of goal recognition models.

Introduction

Recent years have seen significant advances in real-time narrative generation. Narrative adaptation has 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 in the interactive narrative community is goal recognition. Goal recognition is a restricted form of the plan recognition problem. Both goal recognition and plan recognition are active areas of investigation (Armentano and Amandi 2009; Gold 2010; Sadilek and Kautz 2010; Kabanza, Bellefeuille, and Bisson 2010). Goal recognition involves identifying the specific objectives that a user is attempting to achieve, where the user's goals are hidden from the system and must be automatically inferred from user actions taken in the game environment. Within interactive narratives, player goals may comprise key plot points that must be satisfied in order to advance stories.

Goal recognition models offer several potential benefits to interactive narrative systems. First, they enable drama managers to preemptively augment story experiences in anticipation of player actions. Models that provide accurate predictions about players' actions are essential for interactive narratives to proactively mediate narrative conflicts, personalize story events, and manage player agency. Second, recognizing players' goals is important for narrative-centered learning environments. Interpretations of players' goals and plans contribute to assessments of learning progress, and goal recognition models can inform intelligent tutoring systems within narrative-centered learning environments. Third, goal recognizers can provide valuable information for post hoc data mining of interactive narrative log data. Providing detailed descriptions of players' objectives and problem-solving plans facilitates interpretation of raw game logs, and player goal data can be analyzed to inform subsequent interactive narrative designs.

Computationally representing narrative structures is an important consideration when designing interactive narrative systems. Several interactive narrative systems have incorporated special-purpose data structures inspired by narrative concepts, such as dilemmas or beats, in order to bundle story content for reactive delivery (Barber and Kudenko 2007; Mateas and Stern 2005). Classical planning techniques that align STRIPS-style plans with computational representations of plots and discourse (Riedl, Saretto, and Young 2003) have also been examined extensively. Alternatively, other work has modeled interactive narratives using decision-theoretic formalisms based on Markov decision processes (Roberts, et al. 2006) and dynamic decision networks (Mott and Lester 2006).

In this paper, we propose a narrative goal recognition framework that leverages *narrative schemas*—explicit representations of canonical story situations—to encode information about narrative structures that occur during interactive narrative experiences. The concept of narrative schemas is drawn from related work in the natural language processing (NLP) community that focuses on automatically capturing salient information about narrative sequences that recur across many related stories (Chambers and Jurafsky 2009). We anticipate that explicit representations of narrative schemas can be leveraged by goal recognition models to encode higher-level patterns of player behavior than those provided by raw game logs, and subsequently improve the performance of goal recognition models. We present an illustrative example using the CRYSTAL ISLAND narrative-centered learning environment in order to discuss how narrative schemas could be formulated for actual interactive narrative systems and subsequently incorporated into narrative goal recognition models.

Narrative Representation

Computationally modeling narrative structures poses significant representational challenges. One of the inherent challenges of narrative modeling is reasoning over time. Much of the work on temporal reasoning is based on interval logic, which was originally formulated in (Allen 1981). Interval logic provides a general, context-independent framework for reasoning about time that can be applied to narrative representation. The TimeML markup language (Pustejovsky et al. 2003) extends interval logic for natural language reasoning.

Extensions of the TimeML system for temporal reasoning have been examined for narrative event reasoning (Chambers and Jurafsky 2008; Chambers and Jurafsky 2009). The *narrative event chains* model extracts event information from texts and analyzes co-occurrence information, along with temporal reasoning supported by TimeML, to discover likely sequences of narrative events (Chambers and Jurafsky 2008). However, narrative event chains impose representational restrictions by assuming that chains involve only a single protagonist. The narrative schemas model extends narrative event chains by allowing for multiple protagonists and general protagonist roles (Chambers and Jurafsky 2009).

Work on the Scheherazade project has examined story graph techniques for encoding narrative events and structures (Elson and McKeown 2010; Elson 2012). A timeline layer encodes the events and state information in a propositional representation. A textual layer represents text spans that are connected to the timeline layer. The *story intention graph* extends the model to include an interpretive layer, which is used to annotate the beliefs, goals and affective state of the characters (Elson 2012).

Within the computational linguistics community, it has become popular to represent events in terms of semantic frames (Harabagiu and Bejan, 2010 ; Manshadi, Swanson, and Gordon, 2008). Semantic frames are a description of a type of event, relation, or entity and the participants in it. When understanding a more goal-oriented text, explicitly representing the goals, actions, and context has been shown to be effective (Jung et al. 2010).

Interactive narratives impose different representational requirements than linear narratives. The intelligent narrative technologies community has examined several alternate representations for dynamically managing interactive story experiences. Architectures for performing interactive narrative planning are often characterized in terms of story-centric approaches and character-centric approaches (Porteous, Cavazza, and Charles 2010; Si, Marsella, and Riedl 2008). Story-centric approaches use explicit models of stories in order to guide interactive narrative generation. Common approaches include classical planning techniques (Riedl, Saretto, and Young 2003) or reactive planning with special-purpose data structures inspired by narrative concepts, such as dilemmas or beats, in order to bundle story content for reactive delivery (Mateas and Stern 2005; Barber and Kudenko 2007; Roberts, et al. 2006). Search-based approaches attempt to find plot sequences that optimize designer-specified evaluation functions (Weyhrauch 1997). The EAT annotation framework (Orkin et al. 2010) instead chose a Task oriented representation, which focus on the interactive nature of the experience.

Goal Recognition

Recognizing the goals and plans of players offers significant promise for increasing the effectiveness of interactive narrative environments for entertainment, training, and education. Plan recognition, which seeks to infer users' goals along with their plans for achieving them from sequences of observable actions, has been studied for tasks ranging from natural language understanding to collaborative problem solving and machine translation (Carberry 2001; Kautz and Allen 1986). In story understanding, plan recognition is used to infer characters' goals from their actions (Charniak and Goldman 1993); in dialogue systems, it supports natural language understanding and intention recognition (Blaylock and Allen 2003). Because plan recognition is inherently uncertain, solutions supporting reasoning under uncertainty such as Bayesian models (Charniak and Goldman 1993), probabilistic grammars (Pynadath and Wellman 2000), and variations on hidden Markov models (Bui 2003) have been investigated. In the restricted form of plan recognition focusing on inferring users' goals without concern for identifying their plans or sub-plans, goal recognition models have been automatically acquired using statistical corpus-based approaches without the need for hand-authored plan libraries (Blaylock and Allen 2003).

Recent work focused on real-world applications of goal recognition has emphasized efficient and early online prediction. Armentano and Amandi (2009) describe an approach for predicting software usage behaviors by learning a variable order Markov model classifier for each user goal. Sadilek and Kautz (2010) use Markov logic to investigate multi-agent applications in the related area of activity recognition. Within digital games, recent work has explored goal recognition to determine players' objectives in a simple action-adventure game (Gold 2010) and create adaptable computer-controlled opponents (Kabanza, Bellefeuille, and Bisson 2010).

Leveraging Narrative Schemas for Goal Recognition

Many computational representations of narrative are designed to encode canonical story situations. If effectively encoded, these situations can then be sequenced to form overarching narrative structures. In the context of goal recognition, these situations describe events that lead to the completion of goals. This is the knowledge that is engineered into *plan libraries* in classical plan or goal recognition problems. When devising plan recognition models for interactive narrative systems, we aim to leverage explicit representations of common narrative sequences in order to identify relationships between higher-level story structures and player goals. Our intention is to automatically mine interactive narrative schemas from logs of players' interactive narrative interactions rather than undertaking labor-intensive knowledge engineering practices to devise them.

Narrative schemas (Chambers & Jurafsky 2009) offer a promising example of a narrative structure representation that originated in the natural language processing community, and they are attractive because they can be learned (unsupervised) from narrative corpora. Narrative schemas provide a structured representation of sequences of events and their participants as an extension of *scripts* (Shank & Ableson 1977). Narrative schemas are a partially ordered probabilistic model which captures information about the situation by analyzing how often events appear in the same narrative. A schema is comprised of several narrative event chains, which each center around a common protagonist. Each chain is associated with a set of protagonist roles, which constrain the types of actors that can participate in the chain. A narrative event chain is composed of several event slots, which capture the types of events by means of the verbs used in the text, as well as whether the protagonist is the subject or object of the verb.

In an analogous manner, in interactive narrative systems, a director agent observes the player's actions and the arguments of those actions. The events themselves are conveyed by actions performed by the player and non-player characters in the virtual environment. Instead of determining textual subject and object role information, an

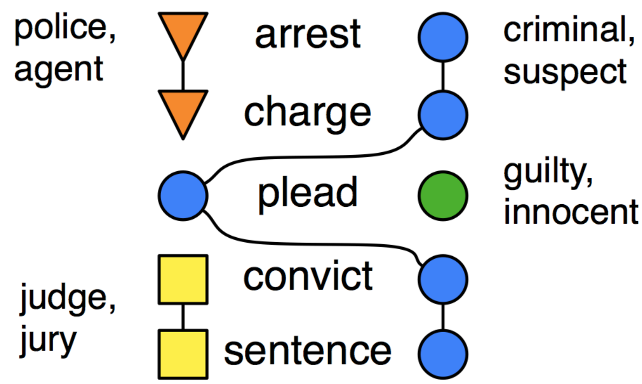

Figure 1. Example of narrative schema and constituent narrative event chains. Figure obtained from (Chambers and Jurafsky, 2009).

interactive narrative system must identify the action's *agent*, as well as any characters that appear in the action's arguments. These actors populate the narrative schema's event slots, allowing the representation to be used within an interactive narrative context.

Figure 1 shows an example of a narrative schema, as well as it's constituent narrative chains. The schema contains 4 chains, one centered around the police officer (the red triangles), one centered around the criminal (the blue circles), one centered around the judge (the yellow squares), and one single-event chain involving the verdict (the green circle). The vertical sequence of words in the middle of the diagram represents the schema's verbs, with each verb's subject on the left and object listed on the right. Event chain slots that co-refer to the same event are merged. This merge process occurs by counting how frequently the slots co-occur with co-referring arguments. For example, the Arrests verb is shown with a subject slot filled by the *police* protagonist. Its object slot is filled by the *criminal* protagonist. Since these slots often occur together in a corpus of criminal sentencing narratives, they were merged together as shown in the figure.

In the analogous context of a typical single-user interactive narrative system, there are several simplifications that can be made to the model. Since there is only a single player agent, we can restrict our set of chains to only those involving the user, as either the agent or argument of an action, rather than having several chains within a single schema. In addition, the set of protagonist role descriptors is used to help restrict which events match a chain in the textual domain. In a single-user interactive narrative experience, the only protagonist candidate is the user-controlled agent.

In the context of natural language processing, one assumption of the narrative schemas model is that a single document corresponds to a single narrative. However, a variation on this assumption is necessary when using narrative schemas with interactive narratives. In goal recognition, the system must observe how the narrative progresses as the player moves through the environment and performs actions. In some interactive narrative genres, such as mysteries, the shape of a narrative experience is defined by the information gained by the player throughout the course of the investigation. As each clue is uncovered, the player gets closer to solving the final mystery, which is the end of the narrative experience. Each clue is uncovered by a sequence of events. These sequences share the same structure within the game as a narrative schema. Each sequence is a repeatable situation, which can be recognized within the goal recognition context.

Within a mystery narrative, the player's goals are also modeled in terms of the information gained. In a mystery, there are two types of clues that can be revealed. Some offer evidence that supports a possible solution, and some offer evidence that refutes a possible solution. While all of these clues are candidates for schema, the goals are defined by the set of clues that support the correct solution.

Goal recognition models are poised to take advantage of narrative schema representations to enhance their predictive accuracies. The process of learning and recognizing narrative schemas based on low level action observations transforms interactive narrative log data. The goal recognition model can then operate at a higher level of abstraction, reasoning about players' role and progress within interactive narrative plot structures rather than low-level actions.

CRYSTAL ISLAND Testbed Environment

In order to investigate narrative schemas and goal recognition models in an interactive narrative environment involving many possible goals and user actions, we are utilizing data collected from player interactions with the CRYSTAL ISLAND interactive narrative. CRYSTAL ISLAND is an educational interactive narrative for middle school science. It is built on Valve Software's Source™ engine, the 3D game platform for Half-Life 2. The environment features a science mystery where students attempt to discover the identity and source of an infectious disease that is plaguing a research team stationed on a remote island. Students play the role of a visitor who recently arrived in order to see her sick father, but they are promptly drawn into a mission to save the entire research team from the outbreak. Students explore the research camp from a first-person viewpoint and manipulate virtual objects, converse with characters, and use lab equipment and other resources to solve the mystery.

Within the CRYSTAL ISLAND interactive narrative environment, goal recognition models have several potential applications. The primary set of narrative goals within CRYSTAL ISLAND's mystery plot distinguish among key stages of the problem solving process that the learner must complete in order to solve the mystery. Recognizing these goals could allow the system to provide scaffolding to the scientific method by providing dynamic, user-tailored support. If the goal recognition system is augmented with a component that models the user's knowledge, the system could also provide an adaptive narrative experience by recognizing where the user is searching for required information, and adjust the clues given to the student in accordance with the actions taken. If the goal set is augmented, the framework could also be used to investigate off-task behavior, and can adapt the narrative experience to guide the player back on track. Each of these dynamic adaptations, and other similar ones, can be used to dynamically shape the narrative to enhance the user's experience.

Prior Findings

CRYSTAL ISLAND has been used in several prior investigations of goal recognition models. In one study with an earlier version of CRYSTAL ISLAND, forty participants interacted with the interactive narrative system and completed successively assigned goals that were presented as onscreen messages (Mott, Lee and Lester, 2006). Based on CRYSTAL ISLAND'S interactive narrative structure, the order of the goals differed from session to session. Players completed each goal in succession, 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. Eighty training sessions were collected (two sessions per subject), which generated just over twenty thousand training records.

Goal recognition was performed using a Bayesian network model. The model considered features regarding users' actions, locations, and a feature called *narrative state*. The narrative state encoded the sequence of story milestones the user had already experienced up to the point in time the action was being observed. The Bayesian model was compared to unigram and bigram models. All three models outperformed a baseline of random chance, but unigram models proved to achieve the highest accuracy.

A follow-up study examined alternate goal recognition techniques in a more recent version of the CRYSTAL ISLAND environment, (Ha et al. 2012). Rather than explicitly presenting goals to the human participant, goals were identified based on player actions in the environment. First, player actions that achieved goals were identified. Second, all actions in the observation sequence that preceded the achieved goal but followed the previous goal were labeled with the achieved goal's name. Third, actions that achieve goals were removed from the data. Removing goal-achieving actions was necessary to ensure that model training was fair, because it would be trivial to predict goals from the goal-achieving actions. Finally, all actions that were taken after achievement of the last goal were ignored, since those actions have no direct mapping to any goal. The corpus was gathered from 173 play traces, generating 77,182 actions and 893 goals achieved.

The goal recognition model was represented as a Markov Logic Network. Markov Logic Networks (MLNs) define a probabilistic graphical model through the use of first order predicate calculus formulae. These formulae are assigned weights during the learning process, which allows the model to perform statistical relational inference. The same categories of features were used as in the prior goal recognition work. The MLN model was able to outperform unigram and bigram approaches in an empirical evaluation, and it achieved an 80% improvement in recognition accuracy compared to baseline approaches. We aim to further improve the performance of these goal recognition models by incorporating formulae that model narrative schemas.

Narrative Schema-Based Goal Recognition in CRYSTAL ISLAND

The CRYSTAL ISLAND interactive narrative environment is shaped by discoveries made by the player during the course of solving the mystery. There are various sources for this information that can be leveraged as the narrative contexts for the goal recognition problem. One of the primary sources of information is through dialogues with virtual characters. Several characters play the roles of virtual scientists who have expertise on disease-causing agents, such as viruses, bacteria and carcinogens. In addition, some characters provide clues about the events that led up to the illness that the player is attempting to diagnose. Books and posters placed in the environment also provide information that the player needs in order to solve the mystery. Testing the various objects for harmful microbes is also a crucial step in gathering the required information.

Each of these activities involves its own canonical situations that can be captured by the narrative schema representation. The set of goals used in our prior goal recognition work capture a subset of these schemas: the goals that are along the critical path to solving the mystery. The other schemas provide supportive information to allow the user to converge on the correct answer, but are not strictly required to solve the mystery. They can still be used in the goal recognition context, however, because a pattern of reading books and testing objects leading up to the provided goals would provide some structure to the narrative that can help the system reason about those goals. In the CRYSTAL ISLAND narrative environment, it is common for a student to test several objects before discovering the disease's transmission source. This sequence can be captured through several instances of the *testing* narrative schema, which contains actions such as picking up the object, moving to the testing site, and performing lab tests. The key intuition is that the user is establishing a pattern of testing objects, and so it becomes increasingly likely that the next goal to be achieved will be the testing of the correct object.

In order to incorporate this intuition into the most recent goal recognition framework based on Markov logic networks, several changes to the model must be investigated. First, the narrative schema must be obtained from the low-level action data. Once the narrative schema have been identified, the MLN goal recognition model must be updated so that the schema, rather than the observation features, become features for the classifier. Additional allowances for the structure between features may also be required to improve the goal recognition model. To address these issues, we will investigate separated and joint modeling approaches for inserting narrative schema as additional features into the MLN goal recognition framework outlined above.

Conclusions and Future Work

In this paper, we have proposed incorporating explicit representations of narrative schemas into goal recognition models for interactive narrative systems. While goal recognition has many applications in the design and execution of interactive narratives, many current approaches only consider surface observations for clues to underlying goal structure. By utilizing the insight that the intended experience is fundamentally narrative in nature, we proposed an approach that leverages the narrative information captured in narrative schemas to enhance goal recognition models.

The approach offers several avenues for investigation. To empirically explore these, narrative schemas will be learned from a corpus of player log data from the CRYSTAL ISLAND environment. Existing goal recognition models will also be extended in order to incorporate the new narrative schema information. An investigation of a joint modeling task, where the narrative schema and their relationships to the goals are learned simultaneously, is complementary direction for research that is also promising. The framework may also allow goals to be automatically mined from the text, by considering sets of schemas and how they may relate to each other. This may allow for the discovery of goals, such as off-task behavior or other behavior that is counter to the principal objectives in the narrative. One of the limitations of the current approaches to goal recognition in CRYSTAL ISLAND is the assumption that a user is focused on a single goal at a time. Interleaving goals, or interleaving actions which support two different schemas, offer their own challenges. One promising approach is to leverage temporal reasoning methods used by the linguistics community to reorder textual events in order to group together interleaved events in their correct clustering.

Acknowledgements

This research was supported by the National Science Foundation under Grant DRL-0822200 Any opinions, findings, and conclusions expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation. Additional support was provided by the Bill and Melinda Gates Foundation, the William and Flora Hewlett Foundation, and EDUCAUSE.

References

Allen, J. F. 1981. An interval-based representation of temporal knowledge. In *Proceedings of International Joint Conference on Artificial Intelligence*, 221–226.

Armentano, M. G. and Amandi, A. 2009. Goal Recognition with Variable-Order Markov Models. In *Proceedings of the Twenty-First International Joint Conference on Artificial Intelligence*, 1635-1640. Menlo Park: AAAI Press.

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.

Bejan, C.A., Harabagiu, S. 2010 *Unsupervised Event Coreference Resolution with Rich Linguistic Features*. In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*. 1412-1422

Blaylock, N. and Allen, J. 2003. Corpus-based, Statistical Goal Recognition. In *Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence*, 1303-1308. San Francisco, CA: Morgan Kaufmann Publishers Inc.

Bui, H. H. 2003. A General Model for Online Probabilistic Plan Recognition. In *Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence*, 1309-1315. San Francisco, CA: Morgan Kaufmann Publishers Inc.

Carberry, S. 2001. Techniques for plan recognition. User Modeling and User-Adapted Interaction 11(1-2):31–48.

Chambers, N., and Jurafsky, D. 2008. Unsupervised Learning of Narrative Event Chains. In *Proceedings of the 46th Annual Meeting of the Association for Computational Linguistics*, 789–797.

Chambers, N., and Jurafsky, D. 2009. Unsupervised Learning of Narrative Schemas and their Participants. In *Proceedings of the Joint conference of the 47th Annual Meeting of the Association for Computational Linguistics and the 4th International Joint Conference on Natural Language Processing of the Asian Federation of Natural Language Processing*, 602–610.

Charniak, E., and Goldman, R. P. 1993. A Bayesian model of plan recognition. *Artificial Intelligence*, 64(1):53–79.

Elson, D. K. 2012. DramaBank : Annotating Agency in Narrative Discourse. In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC 2012)*.

Elson, D. K., and McKeown, K. R. 2010. Building a Bank of Semantically Encoded Narratives. In *Proceedings of the Seventh International Conference on Language Resources and Evaluation* (LREC 2010).

Gold, K. 2010. Training Goal Recognition Online from Low-Level Inputs in an Action-Adventure Game. In *Proceedings of Artificial Intelligence and Interactive Digital Entertainment Conference*, 21– 26. Menlo Park, CA: AAAI Press.

Ha, E. Y; Rowe, J; Mott, B.; and Lester, J. 2012. *Goal Recognition with Markov Logic Networks for Player-Adaptive Games*. In *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence*, Toronto, Canada, pp. 2113-2119.

Kabanza, F.; Bellefeuille, P.; and Bisson, F. 2010. Opponent Behaviour Recognition for Real-Time Strategy Games. In *Proceedings of AAAI Workshop on Plan, Activity and Intent Recognition (PAIR)*.

Kautz, H. and Allen, J. F. 1986. Generalized Plan Recognition. In *Proceedings of the Fifth National Conference on Artificial Intelligence*, 32-37. Menlo Park, CA: Morgan Kaufmann.

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

Manshadi, M., Swanson, R. and Gordan, A. S. 2008. *Learning a Probabilistic Model of Event Sequences from Internet Weblog Stories*. In *Proceedings of the Twenty-first International Conference of the Florida AI Society, Applied Natural Language Processing track*. 159-165.

Mateas, M., and Stern, A. 2005. Structuring Content in the Facade Interactive Drama Architecture. In *Proceedings of the First Conference on Artificial Intelligence and Interactive Digital Entertainment*, 93–98.

Mott, B. W. and Lester, J. C. 2006. U-Director: A decision- theoretic narrative planning architecture for storytelling environments. In *Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multiagent Systems*, 977-984. New York, NY: ACM.

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 Intelligenc*e, Boston, Massachusetts, pp. 187-192.

Orkin, J., Smith, T., Reckman, H., and Roy, D. 2010. *Semi-Automatic Task Recognition for Interactive Narratives with EAT & RUN*. In *Proceedings of the FDG'10 Workshop on Intelligent Narrative Technologies III*.

Porteous, J.; Cavazza, M.; and Charles, F. 2010. Narrative Generation through Characters Point of View Categories and Subject Descriptors. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems* (AAMAS).

Pustejovsky, J.; Castano, J.; Ingria, R.; Sauri, R.; Gaizauskas, R.; Setzer, A.; and Katz, G. 2003. TimeML : Robust Specification of Event and Temporal Expressions in Text. In *Proceedings of the Fifth International Workshop on Computational Semantics*, 1–12.

Pynadath, D. V., and Wellman, M. P. 2000. Probabilistic state-dependent grammars for plan recognition. In *Sixteenth conference on Uncertainty in Artificial Intelligence*, 507–514. San Francisco, CA: Morgan Kaufmann.

Riedl, M. O.; Saretto, C. J.; and Young, R. 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*, 186–193.

Roberts, D. L.; Nelson, M. J.; Isbell, C. L. J.; Mateas, M.; and Littman, M. L. 2006. Targeting specific distributions of trajectories in MDPs. In *Proceedings of the twenty-first National Conference on Artificial Intelligence*, 1213–1218. Menlo Park, CA: AAAI Press.

Sadilek, A., and Kautz, H. 2010. Recognizing multiagent activities from gps data. In Proceedings of the *Twenty-Fourth AAAI Conference on Artificial Intelligence*, X – x. Menlo Park, CA: AAAI Press.

Si, M.; Marsella, S. C.; and Riedl, M. O. 2008. Integrating Story-Centric and Character-Centric Processes for Authoring Interactive Drama. In *Proceedings of the 4th Artificial Intelligence for Interactive Digital Entertainment Conference* (AIIDE-08).

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.

Weyhrauch, P. (1997). *Guiding Interactive Drama*. Ph. D. Thesis, Carnegie Mellon University Department of Computer Science.