Improving Goal Recognition in Interactive Narratives with Models of Narrative Discovery Events

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 of goal recognition hold considerable 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, has been the subject of extensive investigation in the AI community. However, there have been relatively few empirical investigations of goal recognition models in the intelligent narrative technologies community to date, and little is known about how computational models of interactive narrative can inform goal recognition. In this paper, we investigate a novel goal recognition model based on Markov Logic Networks (MLNs) that leverages *narrative discovery events* to enrich its representation of narrative state. An empirical evaluation shows that the enriched model outperforms a prior state-of-the-art MLN model in terms of accuracy, convergence rate, and the point of convergence.

Introduction

Recent years have seen significant advances in intelligent narrative technologies. Narrative adaptation has shown promise for dynamically tailoring interactive story experiences (Mateas and Stern 2005). Work on computational models of narrative has progressed along several dimensions, including believable characters (Porteous, Cavazza, and Charles 2010; McCoy, et al. 2011), plot conflict (Ware et al. 2012) and crowdsourced story generation (Li, et al. 2013). Computational representations of narrative are a key part of interactive narrative systems that adapt to individual players.

While there has been growing interest in user modeling in interactive narrative (Thue and Bulitko 2012; Yu and Riedl 2013), one area of user modeling that has received limited attention is goal recognition. Both goal recognition and the more general problem of plan recognition are active areas of investigation in the AI community (Armentano and Amandi 2009; Singla and Mooney 2011), and they have been examined in non-narrative games (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 interactive environment. Within interactive narratives, player goals may comprise key plot points that must be satisfied in order to advance the story.

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 of players' actions are essential for interactive narrative systems to proactively mediate narrative conflicts and personalize story events (Harris and Young 2009). For example, consider an interactive mystery in which the narrative's pacing depends on how clues about the mystery's solution are revealed to the player. If a director agent detects that the player is likely to skip an important clue, thereby jumping to a later point in the story, it can preemptively adjust the narrative's pacing by altering the clue's placement in the virtual environment. Alternatively, if the director agent recognizes that the player is pursuing a red herring that would disrupt the narrative's pacing, it could deliver pointed clues to lead the player back to a more promising path.

A second benefit of goal recognition models arises in educational applications of interactive narrative, such as 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. A third benefit of goal recognizers is their capacity to 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 logs, and player goal data can be analyzed to inform subsequent interactive narrative designs.

Within interactive narratives, computationally representing narrative structures holds significant promise for informing goal recognition models. 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).

In this paper, we investigate how computational representations of narrative plot impact goal recognition models for interactive narrative environments. We present *narrative discovery events*, which convey key pieces of information about a story to the player. Narrative discovery events are instrumental in driving a plot's action, particularly in narrative genres that deliberately hide information from the audience such as detective or crime stories. We leverage an explicit model of narrative discovery events to inform goal recognition models for predicting player actions in an interactive narrative environment. The goal recognition models are implemented as Markov Logic Networks (MLNs), a machine learning framework that blends first order logical reasoning and probabilistic inference. The models are evaluated on a corpus of user interactions with CRYSTAL ISLAND, an educational interactive narrative featuring a medical mystery scenario. An empirical evaluation indicates that incorporating narrative discovery events into the narrative representation leads to goal recognition models that perform better than prior state-of-the-art models in terms of accuracy, convergence rate, and the point of convergence. We conclude the paper by discussing promising directions for future work.

Computational Representations of Narrative

Computationally modeling narrative structures poses significant representational challenges. The *story intention graph* (Elson 2012) represents a narrative through propositional representations of the events and world state, along with annotations of character beliefs, goals, and affective states, and connections between the representation and the text. *Story Threads* (Gervás 2012) focus on one character at a time, weaving them together to form an overarching narrative structure. *Narrative Content Units* provide an annotated framework for capturing childrens' retellings (Passonneau, Goodkind and Levy 2007). The Scheherezade system used a representation of events to crowdsource narrative structure generation (Li, et al. 2013). *Plot Units* capture transitions in affective (Lenhert 1981) states, and they are used for narrative understanding and summarization.

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 that focuses on the interactive nature of the experience.

In this work, we propose a representation based on *narrative discovery events*. Narrative discovery events advance the audience's knowledge, or player's knowledge, about key information in the story, and they drive plot progress in narratives where story information is deliberately hidden from the audience for narrative effect. Narrative discovery events are used in a similar way to Plot Units to contextualize different states in the narrative. In contrast to previous STRIPS-style plan-based representations, narrative discovery events capture the state of the player's knowledge, rather than the state of the world in which the narrative is occurring. By providing narrative context to actions, narrative discovery events are able to support goal recognition by providing a measure of what the player has experienced in an interactive narrative.

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. Goal recognition, which focuses on inferring users' goals without concern for identifying their plans or sub-plans, has been studied using statistical corpus-based approaches without the need for hand-authored plan libraries (Blaylock and Allen 2003). In this work, we cast goal recognition as a classification problem. A classifier is machine learned from a corpus of player interaction data off line. The computational demands of runtime goal recognition are dependent on the time complexity of assigning a goal to each observed player action. In our framework, classifying an action with a goal involves probabilistic inference with a Markov network.

Recent work on goal recognition has focused on efficient and early online prediction. Armentano and Amandi (2009) predicted software usage behaviors by learning a variable order Markov model classifier for each user goal. Input-Output HMMs were used to determine players' objectives in an action-adventure game (Gold, 2010). The PHATT algorithm (Geib and Goldman 2009) was used to create adaptable computer-controlled opponents (Kabanza, Bellefeuille, and Bisson 2010). Markov Logic Networks (MLNs) have been used to perform plan recognition through abductive reasoning (Singla and Mooney 2011), as well as recognize multi-agent behaviors in Capture the Flag games (Sadilek and Kautz 2012). Confidence Bounds applied to Trees was used to predict rush behavior in digital football games (Laviers and Sukthankar 2011). MLNs and Dynamic Bayesian Networks have been investigated in narrative centered learning environments (Mott, Lee and Lester 2006; Ha et al. 2011), and an unsupervised learning method was used to analyze player interactions in The Restaurant Game (Orkin et al. 2010).

Recognizing player goals in interactive narratives raises a number of computational challenges. If goals are not explicitly presented to the player, she is likely to engage in exploratory behavior within the interactive narrative environment. Typically, goals prompt the player's selected course of action. But in the case of exploratory behavior, a player's actions may lead to the discovery of new goals, thus creating an ambiguous relationship between observed actions and goals. In our work, undirected probabilistic models are used to capture those ambiguous relationships.

In many cases, goals are not independent of one another. In interactive narratives, goals are often partially-ordered narrative milestones that the user is expected to accomplish during the course of the story. Even among goals that have no order restrictions, proximity within the virtual environment often influences the order in which goals are completed. To address this challenge, we focus on models that are capable of explicitly encoding the relationships between multiple goals. However, we make a simplifying assumption that players pursue only one goal at a time.

Encoding Narrative Plot with Narrative Discovery Events

A defining characteristic of narrative is its capacity to communicate a sense of progression. One of the fundamental purposes of narrative structure is to measure this story progression. Measures of narrative progress can take many forms, depending on the task or genre of the narrative. In the classical three-act and five-act structures, shifts in dramatic tension mark key stages of plot progress. In computational models of narrative, several measures of progress have been investigated. Some narrative representations denote transitions between affective states (Lenhart 1981). Other narrative representations focus on characters, such as their beliefs, goals, or conflicts between them (Ware et al. 2012). Logical representations of world state or fabula have been investigated extensively (Elson 2012; Riedl and Young 2005), and narrative is conceptualized as the plan that transitions between states.

In many narratives, a central question often drives the narrative events that move the story forward. Central questions can be about the nature of the setting, or some fact about the world, what characters will do or have done, and what the consequences will be of a major event (Card, 1988). For example, *the Lord of the Rings* trilogy asks, among other things, "What lies beyond the Shire?" Murder mysteries often focus on the question of "Who did it?" Romance stories ask "Will they, or won't they?" Many disaster stories ask how the world will change after some great disaster.

In this work, we focus on a class of interactive narrative events where the player discovers key information necessary to resolve the central question that drives the interactive narrative's plot. These *narrative discovery events* may reveal information about the backstory, about the connections between characters and objects in the world, or provide support for competing explanations.

In our work, narrative is represented by the set of narrative discovery events that have occurred during a player interaction with the virtual environment. Narrative discovery events are defined as hand-authored predicates over player actions and world events. By explicitly encoding the set of narrative discovery events encountered by the player, goal recognition models can take advantage of contextual information about the player's progress towards resolving the narrative.

To illustrate the concept of a narrative discovery event, consider the following medical mystery scenario. The protagonist, a medical detective, is investigating a disease afflicting a group of sick patients. In order to learn more about the illness, the protagonist asks a sick patient about his symptoms and recent medical history. This conversation is a narrative discovery event; the protagonist has gained valuable information to diagnose the illness and resolve the story's central question: "What is making the patients sick?" Later in the investigation, the protagonist runs a laboratory test on some eggs that the sick patients recently consumed. The eggs test positive for Salmonella. This is another narrative discovery event that reveals the likely transmission source for the disease. In addition to plot revelations, narrative discovery events can include first-time demonstrations of game mechanics that are required to advance the interactive narrative's plot. For example, when the protagonist demonstrates that she can successfully use the laboratory's testing equipment for the first time, it could be considered a narrative discovery event. In this manner, narrative discovery events encompass a broad range of events that transpire in interactive narratives.

Interactive Narrative Observation Corpus

In order to investigate goal recognition models in an interactive narrative environment involving many possible goals and player actions, we are utilizing data collected from player interactions with CRYSTAL ISLAND. 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

$$\forall t,g: |\exists\, t_2 < t : action(t_2, Worksheet)| \geq 1 \Rightarrow goal(t,g) \quad (14)$$

$$\forall t,g: |\exists\, t_2 < t : action(t_2, Test)| \geq 1 \Rightarrow goal(t,g) \quad (15)$$

$$\forall t,g: \left|\begin{matrix}\exists\, t_2 < t : action(t_2, Read) \\ \wedge\, arg(t_2, disease)\end{matrix}\right| \geq 1 \Rightarrow goal(t,g) \quad (16)$$

$$\forall t,g: \left|\begin{matrix}\exists\, t_2 < t : action(t_2, Dialog) \\ \wedge\, arg(t_2, symptoms)\end{matrix}\right| \geq 1 \Rightarrow goal(t,g) \quad (17)$$

$$\forall t,g: \left|\begin{matrix}\exists\, t_2 < t : action(t_2, Dialog) \\ \wedge\, arg(t_2, eating)\end{matrix}\right| \geq 1 \Rightarrow goal(t,g) \quad (18)$$

$$\forall t,g: |\exists t_2 < t : \mathrm{arg}(t_2, bacteria)| \geq 1 \Rightarrow goal(t,g) \quad (19)$$

Figure 3: Narrative discovery event formulae

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.

Gameplay data from CRYSTAL ISLAND was gathered from a study involving 153 eighth-grade students. Sixteen students were removed due to incomplete data or prior experience with the game. The students interacted with the system for a maximum of 60 minutes. A full description of the study methods can be found in (Rowe et al. 2011).

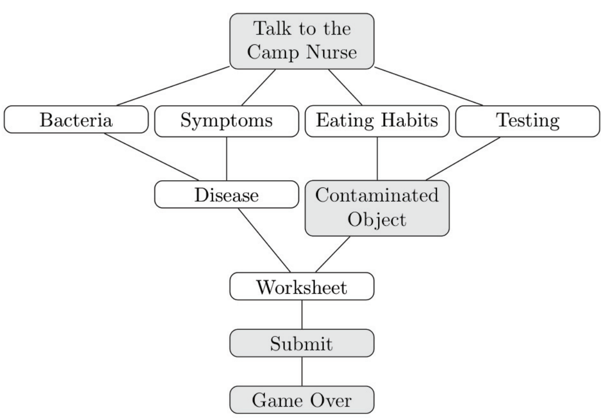

Figure 1: Narrative plot graph

MLN-based Goal Recognition with Narrative Discovery Events

We have identified seven goals that are central to the mystery in CRYSTAL ISLAND. These goals are never directly presented to the player, and are instead discovered during the course of gameplay. Five of the goals involve gathering information from virtual characters. One goal involves successfully running a laboratory test on the disease's transmission source. The final goal involves submitting a correct diagnosis to a virtual character, thereby solving the mystery.

There are three types of observations that are available to a goal recognition system for CRYSTAL ISLAND: actions, locations, and narrative states. There are nineteen different types of actions in CRYSTAL ISLAND. These actions allow students to interact with objects, talk with characters, gather and record information, and interact with the lab equipment. The arguments observed include the objects in the environment and topics that are covered by the virtual characters and recorded in a diagnosis worksheet. These topics include microbiology concepts, such as viruses and bacteria, as well as clues to the mystery, such as symptoms and eating habits of the patients.

There are seven major regions of the island, including the living quarters, an infirmary, dining hall, laboratory, lead scientist's quarters, waterfall, and outdoor camp. Each of these locations is further divided into several sub-areas, comprising a rough grid over the map. In total, there are 39 distinct, non-overlapping locations. The player's coordinates within the game space are mapped into these regions, and recorded as an observation.

The narrative state characterizes the player's progress in the plot. Ha et al. (2011) leverage a narrative state that captures four milestones in the narrative. The first milestone marks whether or not the student has been informed of their role in solving the mystery. The next milestone indicates whether or not the player has successfully identified the spreading disease's transmission source in the laboratory. The third milestone is achieved when the player first submits her diagnosis worksheet for review, which is an attempt to solve the mystery. The final milestone is achieved by solving the mystery scenario, which involves correctly diagnosing the illness and recommending a course of treatment. Figure 1 shows these milestones in colored boxes.

In this work, the narrative state has been expanded to include several additional milestones from *narrative. discovery events* Two of the milestones are delivered through interactions with the characters: a list of patient symptoms, and the foods they have eaten. Two of the milestones are delivered through interactions with the system: feedback from testing objects, and recording findings in the diagnosis worksheet. The last two involve the learning content: learning about bacteria from books or characters, and reading about the correct disease. These milestones are shown in Figure 1 in white boxes.

$$\forall t,a: action(t,a) \Rightarrow |\exists g: goal(t,g)| = 1 \quad (1)$$
$$\forall t,g: goal(t,g) \quad (2)$$
$$\forall t,a,g: action(t,a) \Rightarrow goal(t,g) \quad (3)$$
$$\forall t,l,g: loc\ (t,l) \Rightarrow goal(t,g) \quad (4)$$
$$\forall t,s,g: state(t,a) \Rightarrow goal(t,g) \quad (5)$$
$$\forall t,a,s,g: action(t,a)\ \&\ state(t,s) \Rightarrow goal(t,g) \quad (6)$$
$$\forall t,a,g: action(t-1,a) \Rightarrow goal(t,g) \quad (7)$$
$$\forall t,l,g: loc\ (t-1,l) \Rightarrow goal(t,g) \quad (8)$$
$$\forall t,s,g: state(t-1,a) \Rightarrow goal(t,g) \quad (9)$$
$$\forall t,a,s,g: action(t-1,a) \wedge state(t-1,s) \Rightarrow goal(t,g) \quad (10)$$
$$\forall t,a_1,a_2,g: action(t-1,a_1) \wedge action(t,a_2) \Rightarrow goal(t,g) \quad (11)$$
$$\forall t,g_1,g_2: goal(t-1,g_1) \Rightarrow goal(t,g_2) \quad (12)$$
$$\forall t,a_1,a_2,g_1,g_2: action(t-1,a_1) \wedge goal(t-1,g_1) \wedge action(t,a_2) \Rightarrow goal(t,g_2) \quad (13)$$

Figure 2: Baseline formulae for goal recognition

In order to use player interaction data from CRYSTAL ISLAND to evaluate our goal recognition model, we first had to associate a goal with each sequence of actions. First, each of the actions that achieved a goal were identified. These were then used to label the actions between the completion of the last goal and the completion of the current goal as working towards the current goal. The actions that achieved the goal were then removed, since it would be trivial to predict the goal given the goal-achieving action. Lastly, all actions performed after the completion of the final goal, but before the end of the interaction time limit were removed, since the system was unable to attribute them to any goal.

The goal recognition models were implemented as *Markov Logic Networks* (MLNs) (Domingos et al. 2006). Markov logic is a template language for specifying a Markov network in terms of first-order logic and probabilistic weights. These weights offer intuition for how deterministic the relationship is. As the weights approach infinity, the reasoning approaches first order logical entailment. A *Markov Network* is a probabilistic graphical model where each node represents a random variable. Each undirected edge represents a (conditional) dependency between two variables, and are derived from the logical connectives in the first-order formulae. The joint probability function is represented as the product of *potential functions*. Potential functions assign probabilities over cliques in the Markov network's graph structure. The models were trained using Markov: The Beast, an open source implementation of Cutting Plane Inference for the Markov Logic Network learning framework (Riedel 2008).

Three sets of formulae were designed for the goal recognition models. The baseline set models the dependencies between the current time step and the previous time step, as well as joint probabilities between the two time steps. It achieved state-of-the-art performance on the observation corpus (Ha et al. 2011). The formulae are shown in Figure 2. The ablated set is a subset of the baseline that only considers the current time step in making a prediction. It is formed by equations 1–6 in Figure 2. The narrative discovery event set captures whether or not each narrative discovery event has been observed in any prior time step. These formulae recognize actions or arguments that indicate that the player has discovered the requisite information. The formulae are given in Figure 3. MLNs reason by drawing connections between the predicates involved in the formulae. By separating this knowledge into a separate set of formulae, the system is able to reason about how the narrative states connect to the rest of the observations. This is an advantage over the representation that was used in the baseline model.

Model	F1	Conv. Rate	Conv. Point
Ablated + ND	0.546[1,2]	50.056[1,2,3]	35.862[1,2]
Baseline + ND	0.537[3]	32.905[3]	51.298[1,3]
Baseline	0.488[2]	30.906[2,4]	50.865[2,4]
Ablated	0.477[1,3]	37.001[1,4]	37.628[3,4]

Table 1: Summary of results. Pairs with the same superscript are significant at $p < 0.05$

Evaluation

The evaluation of our goal recognition model focuses on three principal metrics. F1 measures the predictive accuracy of the models. Convergence rate is the percentage of sequences that eventually find the correct goal. Any sequence whose final action is predicted as belonging to the correct goal is said to have converged on the goal. Convergence point measures the percentage of a converged sequence that was observed before the correct goal was consistently predicted.

We compare four different models formed from the formulae sets. The baseline and ablated models are formed from their respective set of formulae. The other two models are formed by including the narrative discovery formulae to each of the baseline and ablated models.

The models were trained using 10-fold cross validation. The folds were calculated using the number of students, rather than the number of goal sequences, to ensure independence between folds. The three evaluation metrics were computed across each of the ten evaluations, and the four models were compared using One-Way ANOVA. Results are given in Table 1.

There are several interesting trends in the results. On the F1 metric, both models with the narrative discovery formulae significantly outperform the baseline and ablated models. On the convergence metrics, the ablated models converged faster and more often than the baseline.

The findings indicate that the richer narrative representation, which includes an encoding of narrative discovery events, has a beneficial effect on goal recognition accuracy. One possible explanation is that the richer narrative representation offers an improved context within which the goal recognition framework can interpret the actions and attribute them to the goal. By capturing which narrative discovery events have passed, the system knows which goals the player has left to accomplish and whether or not she has enough information to pursue the goal. Knowing which foods the patients have eaten, for example, is necessary for selecting which objects to test.

While additional features about the observations in the previous time step enhanced the goal recognition models' accuracy, the simpler ablated models performed better on the convergence metrics. With many statistical modeling approaches, more features often result in more noise and a greater risk of over-fitting. With MLNs, more formulae results in a greater number of cliques, which impacts the weight learning algorithms. While accuracy is not significantly altered through ablation, the predictions are more consistent, resulting in better convergence performance. With the addition of the extended narrative state representation, the risk of accuracy loss is mitigated.

Overall, the ablated model with the narrative discovery events offers the best performance. It achieved the highest F1 score, though its nearest neighbor was not significantly different, and significantly outperforms competing models on the convergence metrics. It leverages both the extended narrative state and the ablated model to capture the goal behavior of players.

While these findings are encouraging, the work does have several limitations that should be noted. One of the limitations of corpus-based goal recognition is that it requires a significant amount of data in order to produce high-fidelity models. In order to address this, we are adopting an iterative approach to software development: a prototype game is used to collect the required data, models are trained and evaluated with the data, the models are integrated back into the game, and new data about their performance is collected to repeat the cycle. This process can be repeated multiple times to refine the model and to capture interaction effects.

While we hypothesize that our goal recognition framework can be generalized to other games, we should note several assumptions in its application to CRYSTAL ISLAND. Our conceptualization of narrative discovery events is likely to be most useful for games with a branching narrative structure. We assume players behave in a goal-directed manor. If this were relaxed, there may be insufficient structure in the data for machine learning techniques to exploit. We also assume that the player is pursuing one goal at a time without interleaving goals. The first formula in the baseline model enforces this by stating that only one goal can be active at any given time. Relaxing these assumptions by investigating alternate games with minimally-structured narratives, interleaved goal sequences, and which encourage exploration without defined goals, are promising directions for future work.

Conclusions and Future Work

In this paper, we have proposed a narrative representation based on narrative discovery events. Narrative discovery events represent key information about an interactive narrative that a player needs to resolve the central question driving a plot. By incorporating a representation of narrative discovery events into goal recognition models, we have produced a model that outperforms a prior state-of-the-art approach in terms of accuracy, convergence rate, and point of convergence metrics.

There are several promising directions for future work. Exploring alternate measures of narrative structure beyond narrative discovery events holds promise for yielding further improvements to goal recognition models. The narrative discovery event model captures basic structural information about CRYSTAL ISLAND's plot and gameplay, but another interactive narrative system may lend itself to an alternate representation, such as an encoding of character emotions or dramatic tension. Another promising direction is automated feature selection for goal recognition. Since the ablated model with narrative discovery events achieved the greatest performance, investigating which additional features add information, and which introduce noise, could be informative. Finally, investigating the runtime impact of goal recognition models on player experiences in interactive narratives is an important next step. Different experience management tasks may have different requirements for goal recognition, and empirical investigation of goal recognizers' utility in run-time settings, as well as their capacity to transfer to alternate domains, would yield valuable insights.

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, EDUCAUSE, and the Social Sciences and Humanities Research Council of Canada.

References

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.

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.

Card, O.S. 1988. *Characters and Viewpoint*. Cincinnati, OH: Writer's Digest Books.

Domingos, P.; Kok, S.; Poon, H.; Richardson, M.; Singla, P. 2006. Unifying logical and statistical AI. In *Proceedings of the Twenty-First National Conference on Artificial Intelligence*, 2–7. Menlo Park, CA: AAAI Press.

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

Geib, C.W., and Goldman, R.P. 2009. A probabilistic plan recognition algorithm based on plan tree grammars. *Artificial Intelligence* 173: 1101–1132.

Gervás, P. 2012. From the Fleece of Fact to Narrative Yarns : a Computational Model of Composition. In *Proceedings of the Third Workshop on Computational Models of Narrative*, 123–131. Paris, France: Eruopean Language Resource Association.

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.P.; Mott, B.W.; and Lester, J.C. 2011. Goal Recognition with Markov Logic Networks for Player-Adaptive Games. In *Proceedings of the Seventh Conference on Artificial Intelligence in Interactive Digital Entertainment*. Menlo Park, CA: AAAI Press.

Harris, H., and Young, R. M. 2009. Proactive mediation in plan-based narrative environments. *IEEE Transactions on Computational Intelligence and AI in Games*, 1(3): 233–244.

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)*. Melo Park, CA: AAAI Press.

Laviers, K., and Sukthankar, G. 2011. A real-time opponent modeling system for rush football. In *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence*, 2476–2481. Menlo Park, CA: AAAI Press.

Lenhert, W.G. 1981. Plot Units and Narrative Summarization. *Cognitive Science* 5: 293–331.

Li, B.; Lee-Urban, S.; Johnston, G.; and Riedl, M. 2013. Story generation with crowdsourced plot graphs, In *Proceedings of the 27th AAAI Conference on Artificial Intelligence*. Menlo Park, CA: AAAI Press.

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. Menlo Park, CA: AAAI Press.

McCoy, J.; Treanor, M.; Samuel, B.; Wardrip-Fruin, N.; and Mateas, M. 2011. Comme il Faut: A system for authoring playable social models. In *Proceedings of the 7th AAAI Conference on Artificial Intelligence in Interactive Digital Entertainment*, 158–163, 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 Intelligenc*e, 187-192. Menlo Park, CA: AAAI Press.

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*. New York, NY: ACM.

Passonneau, R.J.; Goodkind, A.; and Levy, E.T. 2007. Annotation of Children's Oral Narrations: Modeling Emergent Narrative Skills for Computational Applications. In *Proceedings of the Twentieth Annual Meeting of the Florida Artificial Intelligence Research Society*. Menlo Park, CA: AAAI Press.

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*, 1297-1304. Toronto, Canada.

Riedel, S. 2008. Improving the Accuracy and Efficiency of MAP Inference for Markov Logic. In *Proceedings of the Twenty-Fourth Annual Conference on Uncertainty in Artificial Intelligence*, 468–475. Corvallis, OR: IUAI Press

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. New York, NY: ACM.

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.

Rowe, J.P.; Shores, L.R.; Mott, B.W.; and Lester, J.C. 2011. Integrating Learning, Problem Solving, and Engagement in Narrative-Centered Learning Environments. *International Journal of Artificial Intelligence in Education* 21(1-2): 115-133.

Sadilek, A., and Kautz, H. 2012. Location-based reasoning about complex multi-agent behavior. *Journal of Artificial Intelligence Research*. 43:87–133.

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*. Vancouver, Canada.

Singla, P., and Mooney, R.J. 2011. Abductive markov logic for plan recognition. In *Twenty-fifth National Conference on Artificial Intelligence*, 1069–1075. Menlo Park, CA: AAAI Press .

Thue, D., and Bulitko, V. 2012. Procedural Game Adaptation: Framing Experience Management as Changing an MDP. In *Proceedings of the AIIDE'12 Workshop on Intelligent Narrative Technologies V*, 44–50. Menlo Park, CA: AAAI Press.

Ware, S. G.; Young, R. M.; Harrison, B.; and Roberts, D. L. 2012. Four Quantitative Metrics Describing Narrative Conflict. In *Proceedings of the Fifth International Conference on Interactive Digital Storytelling*, 18–29. San Sebastián, Spain.

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

Yu, H., and Riedl, M. 2013. Data-driven personalized drama management. In *Proceedings of the 9th AAAI Conference on Artificial Intelligence in Interactive Digital Entertainment.* Menlo Park, CA: AAAI Press.