

Goal Recognition with Markov Logic Networks for Player-Adaptive Games

Eun Y. Ha, Jonathan P. Rowe, Bradford W. Mott, and James C. Lester

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

Abstract

Goal recognition in digital games involves inferring players' goals from observed sequences of low level player actions. Goal recognition models support player adaptive digital games, which dynamically augment game events in response to player choices for a range of applications, including entertainment, training, and education. However, digital games pose significant challenges for goal recognition, such as exploratory actions and ill defined goals. This paper presents a goal recognition framework based on Markov logic networks (MLNs). The model's parameters are directly learned from a corpus that was collected from player interactions with a non linear educational game. An empirical evaluation demonstrates that the MLN goal recognition framework accurately predicts players' goals in a game environment with exploratory actions and ill defined goals.

Introduction

Digital games have grown increasingly sophisticated in their graphical presentations, interaction designs, and simulation capabilities. One area in which games have experienced slower progress is in their ability to interpret and respond to players' low-level actions in virtual environments. Digital games without player models have limited capacity to adapt their behavior to individual player intentions, abilities, and preferences. For many years, digital games have embraced AI techniques for pathfinding (Sturtevant, 2011) and non-player character control (Orkin, 2005). However, player modeling is only now beginning to

gain recognition as a promising frontier for game AI (Yannakakis, 2012). One particularly important player modeling task is *goal recognition*, which involves automatically inferring a player's current gameplay objective based on observations of low-level actions in a virtual environment.

The prospective roles of goal recognition models in digital games are particularly apparent in open world (or "sandbox") games. Open world games, such as the popular Elder Scrolls series (Bethesda, 2011) and Minecraft (Mojang, 2009), feature vast environments, emergent gameplay, and multiple paths for accomplishing game objectives. In these games, players choose the goals they pursue, and develop their own plans to achieve goals. Goal recognition models introduce opportunities for adapting gameplay events based on the choices of individual players. For example, consider a scenario in which a player struggles to scale a steep cliff face. Utilizing a goal recognition model, the game could interpret the player's repeated jumps and backsliding as an attempt to reach the cliff summit. Armed with inferred knowledge about the player's goal, the game might inform the player about a stairway a short distance away, or redirect the player toward a different, more useful goal. Without a goal recognition model, the game is unable to interpret the player's efforts as anything other than a sequence of co-located repeating actions.

Goal recognition, as well as its sibling tasks, plan recognition and activity recognition, are long-standing AI problems (Kautz and Allen, 1986; Charniak and Goldman, 1993; Carberry, 2001; Singla and Mooney, 2011). The problems are cases of abduction: given domain knowledge and a sequence of actions performed by an agent, the task is to infer which plan or goal the agent is pursuing. Recent work has yielded notable advances in recognizing agents' goals and plans, including methods for recognizing multiple concurrent and interleaved goals (Hu and Yang, 2008), methods for recognizing activities in multi-agent settings (Sadilek and Kautz, 2010), and methods for

Copyright © 2012, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

This paper is an invited contribution to the AAAI-2012 Sub-Area Spotlights track. The paper is based on a paper by the same title, which received the Best Paper Award at the 7th AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment, 2011 (AIIDE-2011). Parts of the paper were re-written to make it more accessible to a general AI audience.

augmenting statistical relational learning techniques to better support abduction (Singla and Mooney, 2011).

Digital games pose significant computational challenges for goal recognition models. For example, in many games players' abilities change over time, both by unlocking new powers and improving motor skills over the course of gameplay. In effect, a player's action model may change, in turn modifying the relationships between actions and goals. Action failure is also a critical and deliberate design choice in games; in platform games, a poorly timed jump often leads to a player's demise. In multiplayer games, multi-agent goals arise that may involve players competing or collaborating to accomplish game objectives. Individual players may also pursue ill-defined goals, such as "explore" or "try to break the game." In these cases goals and actions may be cyclically related; players take actions in pursuit of goals, but they may also choose goals after they are revealed by particular actions. For these reasons, digital games offer promising testbeds for investigating different formulations of goal recognition tasks.

In conjunction with the challenges noted above, digital games offer notable benefits for investigating goal and plan recognition models. Digital games are highly configurable, particularly with the availability of low-cost game development tools and level editors. Digital games are popular and deployable, which aid in collection of real-world training data from human users. Games also simplify sensor issues in goal recognition. Virtual environment states are effectively fully observable, providing highly configurable sensor capabilities for monitoring player actions.

This paper focuses on an investigation of goal recognition models in an open-ended game environment with ill-defined goals and exploratory player behaviors. Given its relationship to abduction, goal recognition appears well suited for logical representation and inference. However, goal recognition in digital games also involves inherent uncertainty. For example, a single sequence of actions is often explainable by multiple possible goals. Markov logic networks (MLNs) provide a formalism that unifies logical and probabilistic representations into a single framework. To address the problem of goal recognition with exploratory goals in game environments, a Markov logic goal recognition framework is investigated.

The MLN goal recognition model is trained on a corpus collected from player interactions with an open-ended adventure game. The game environment used to train and evaluate the goal recognition model is CRYSTAL ISLAND, a story-centric educational game for middle grade science. In CRYSTAL ISLAND, players are assigned a single high-level objective: solve a science mystery. Players interleave periods of exploration and deliberate problem solving in order to identify a spreading illness that is afflicting residents on the island. In this setting, goal recognition involves predicting the next narrative sub-goal that the

player will complete as part of investigating the mystery. We present findings that suggest the MLN goal recognition framework yields significant accuracy gains beyond alternative probabilistic approaches for predicting player goals in a nonlinear game environment.

Related Work

Recognizing players' goals and plans offers significant promise for increasing the effectiveness of digital game environments for education, training, and entertainment. 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 that focuses 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).

The classic goal recognition problem assumes that a single agent is pursuing a single goal using deterministic actions, and it assumes that a user's plan can be identified using a given plan library. A major focus of recent work on goal and plan recognition has been probabilistic approaches that relax several of these assumptions. For example, Ramirez and Geffner (2010) describe a plan recognition approach that does not require the provision of an explicit plan library. Hu and Yang (2008) describe a two-level goal recognition framework that uses conditional random fields and correlation graphs to support recognition of multiple concurrent and interleaving goals. Geib and Goldman (2009) have devised the PHATT algorithm, which is a Bayesian approach to plan recognition that focuses on plan execution. PHATT provides a unified framework that supports multiple concurrent goals, multiple instantiations of a single goal, partial ordering among plan steps, and principled handling of unobserved actions.

Recent work has examined using statistical relational learning frameworks for plan and activity recognition. Sadilek and Kautz (2010) use Markov logic to investigate activity recognition in multi-agent applications. Sadilek

and Mooney (2011) propose a Hidden Cause model and novel model construction method to improve the efficiency and effectiveness of MLNs for abductive inference. The work presented here also uses MLNs, but it focuses on goal recognition in complex, nonlinear game environments, which often include ill-defined sub-goals and cyclical relationships between goals and actions.

There have been several investigations of goal and plan recognition in digital games. Recent work has explored goal recognition to determine players' objectives in an action-adventure game, support dynamic narrative planning, and create adaptable computer-controlled opponents. Gold (2010) describes training an Input-Output Hidden Markov Model to recognize three high-level player goals in a simple action-adventure game. Mott, Lee, and Lester (2006) explore several probabilistic goal recognition models to support dynamic narrative planning. Kabanza, Bellefeuille, and Bisson (2010) explore challenges with behavior recognition in real-time strategy games and present preliminary results for creating adaptable computer-controlled opponents. The current work investigates a Markov logic network goal recognition framework for an educational game environment, with the eventual aim of dynamically tailoring game experiences to players.

Observation Corpus

In order to investigate goal recognition in a nonlinear game environment involving many possible goals and player actions, data collected from player interactions with the CRYSTAL ISLAND educational game were used.

CRYSTAL ISLAND (Figure 1) is an educational game about middle grade microbiology. It is built on Valve Software's Source™ engine, the 3D game platform for Half-Life 2. The environment features a science mystery where players attempt to discover the identity and source of an infectious disease that is plaguing a research team stationed on the island. Players adopt the role of a visitor who has recently arrived, but was promptly drawn into a mission to save the research team from the outbreak. Players explore the camp (Figure 2) from a first-person viewpoint and manipulate virtual objects, converse with characters, and use lab equipment and other resources to solve the mystery. Now in its fourth major iteration, CRYSTAL ISLAND has been the subject of extensive empirical investigation, and has been found to provide substantial learning and motivational benefits (Rowe et al. 2011). Middle school students consistently demonstrate significant learning gains after using CRYSTAL ISLAND, and they report experiencing boredom less frequently than in alternative instructional software. CRYSTAL ISLAND is also challenging, with fewer than 50% of students solving the mystery in less than an hour. The current investigation of



Figure 1. CRYSTAL ISLAND game environment

goal recognition models contributes to an overarching research agenda to devise user-adaptive computational models to dynamically shape players' interactions with game-based learning environments. Prior work has focused on a range of computational modeling tasks, including probabilistic representations for user knowledge modeling (Rowe and Lester 2010) and machine learning frameworks for driving characters' affective behaviors (Robison, McQuiggan, and Lester 2009).

The following scenario illustrates a typical interaction with CRYSTAL ISLAND. The scenario begins with the player's arrival at the research camp. The player approaches the first building, an infirmary, where several sick patients and a camp nurse are located. A conversation with the nurse is initiated when the player approaches the character and clicks the mouse. The nurse explains that an unidentified illness is spreading through the camp and asks for the player's help in determining a diagnosis. The conversation with the nurse takes place through a combination of dialogue menu selections and multimodal character dialogue integrating spoken language, gesture, facial expression, and text. Character dialogue is provided by voice actors and utilizes a deterministic branching structure.

After speaking with the nurse, the player has several options for investigating the illness. Inside the infirmary, the player can talk to sick patients lying on medical cots. Clues about the team members' symptoms and recent eating habits can be discussed and recorded using in-game note-taking features. Alternatively, the player can move to the camp's dining hall to speak with the camp cook. The cook describes the types of food that the team has recently eaten and provides clues about which items warrant closer investigation. In addition to learning about the sick team members, the player has several options for gathering information about disease-causing agents. For example, the player can walk to the camp's living quarters where she will encounter a pair of virtual scientists who answer questions about viruses and bacteria. The player can also learn more about pathogens by viewing posters hanging inside of the camp's buildings or reading books located in

Goal Recognition with MLN

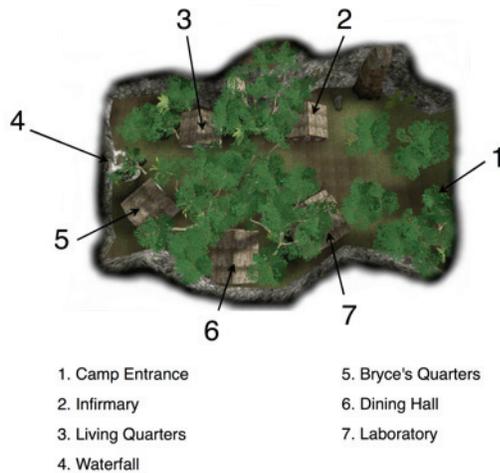


Figure 2. Map of the CRYSTAL ISLAND research camp

a small library. In this way, the player can gather information about relevant microbiology concepts using resources that are presented in multiple formats.

Beyond gathering information from virtual scientists and other instructional resources, the player can conduct tests on food objects using the laboratory's testing equipment. The player encounters food items in the dining hall and laboratory, and she can test the items for pathogenic contaminants at any point during gameplay. A limited number of tests are allocated to the player at the start of the scenario, but additional tests can be earned by answering microbiology-themed questions posed by the camp nurse.

After running several tests, the player discovers that the sick team members have been consuming milk that is contaminated with bacteria. The player can use the camp's books and posters in order to investigate bacterial diseases that are associated with symptoms matching those reported by the sick team members. Once she has narrowed down a diagnosis and recommended treatment in a diagnosis worksheet, the player returns to the infirmary in order to inform the camp nurse. If the player's diagnosis is incorrect, the nurse identifies the error and recommends that the player keep working. If the player correctly diagnoses the illness and specifies an appropriate treatment, the mystery is solved.

All player actions are logged by the CRYSTAL ISLAND software and stored for later analysis. The data used for creating the MLN goal recognition system was collected from a study involving the eighth grade population of a public middle school. There were 153 participants in the study. Data for sixteen of the participants was removed from the analysis due to incomplete data or prior experience with CRYSTAL ISLAND. Participants whose data was included had no prior experience with the software.

Following previous work on goal recognition (Blaylock and Allen 2003; Mott, Lee, and Lester 2006), this work defines goal recognition as the task of predicting the most likely goal for a given sequence of observed player actions in the game environment. Current work assumes that a given sequence of actions maps to a single goal, and no interleaving occurs between actions associated with different goals. Under these conditions, goal recognition is cast as a classification problem, in which a learned classifier predicts the most likely goal associated with the currently observed player action.

The data described in the previous section poses significant challenges for goal recognition. First, individual goals are not independent of one another. Goals in our data represent milestone activities players take in the course of solving the science mystery. Some of these activities naturally occur in a sequential manner. The layout of the island can also impose co-occurrence patterns among goals. To model these associations among the milestone activities, goals should be inferred in relation with one another rather than in isolation. Second, the causality between actions and goals is ambiguous. In CRYSTAL ISLAND, players are not provided explicit goals to achieve. Instead, players discover goals while interacting with the virtual environment. Thus, causality between player actions and goals is bidirectional: a goal could influence a player's current action if she has a particular goal in mind, and it is also possible that the player's current action reveals which goal needs to be pursued. For instance, a player can enter a new location without a particular goal in mind, and afterward she can engage a character in conversation which reveals a new goal.

To address these challenges, the current work utilizes the statistical relational learning framework provided by *Markov logic networks* (MLNs) (Richardson and Domingos 2006). Statistical relational learning frameworks effectively handle machine learning tasks in domains with complex associations among modeled entities. MLNs, in particular, encode undirected probabilistic graphical models with structures that are determined by first-order logic formulae and associated weights. In contrast to directed graphical models (e.g., Bayesian networks, hidden Markov models), undirected graphical models are well suited for representing bidirectional relationships between entities, such as ambiguous causality between actions and goals in our data. In addition, MLNs offer the representational expressiveness of first-order logic. This capability allows for more compact representation of the domain, compared with traditional machine learning frameworks that rely on propositional expressions, such as Bayesian networks and hidden Markov models.

MLNs have recently been applied to tasks that are related to goal recognition, such as probabilistic abduction for plan recognition (Kate and Mooney 2009) and multi-agent activity recognition (Sadilek and Kautz 2010).

Features

Similar to previous work by Mott, Lee, and Lester (2006), the current work encodes player actions in the game environment using three properties: *action type*, *location*, and *narrative state*.

- **Action Type:** Type of action taken by the player, such as *moving to a particular location*, *opening a door*, and *testing an object using the laboratory’s testing equipment*. Our data includes 19 distinct types of player actions.
- **Location:** Place in the virtual environment where a player action is taken. This includes 39 fine-grained and non-overlapping sub-locations that decompose the seven major camp locations in CRYSTAL ISLAND.
- **Narrative State:** Representation of the player’s progress in solving the narrative scenario. Narrative state is encoded as a vector of four binary variables, each of which represents a milestone event within the narrative.

MLN for Goal Recognition

A *Markov Logic Network (MLN)* consists of a set of weighted first-order logic formulae. A weight reflects the importance of the constraint represented by its associated logic formula in a given model. Figure 3 shows 13 MLN formulae that are included in the current goal recognition model. Formula 1 represents a hard constraint that needs to be satisfied at all times. This formula requires that, for each action a at each time step t , there exists a single goal g . The formulae 2-13 are soft constraints that are allowed to be violated. Formula 2 reflects prior distribution of goals in the corpus. Formula 3-13 predict the player’s goal g at time t based on the values of the three action properties, action type a , location l , and narrative state s , as well as previous goal. The weights for the soft formulae were learned with *theBeast*, an off-the-shelf tool for MLNs that uses cutting plane inference technique (Riedel 2008).

Evaluation

To evaluate the MLN goal recognition model, the collected data was preprocessed in several steps. First, all player actions that achieve goals were identified. Second, all actions in the observation sequence that precede the current goal but follow the previous goal were labeled with the current goal. 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-

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

Figure 3. Formulae for MLN goal recognition model

achieving actions. Finally, all actions that were taken after achievement of the last goal were removed, since those actions have no direct mapping to any goal. Table 1 shows summary statistics of the resulting corpus, which includes 77,182 player actions and 893 achieved goals, with an average of 86.4 player actions per goal. Table 2 shows the set of goals considered in this work and their frequencies in the processed corpus data.

Total Number of Observed Player Actions	77,182
Total Number of Goals Achieved	893
Average Number of Player Actions per Goal	86.4

Table 1. Number of actions and goals in corpus

<i>Running laboratory test on contaminated food</i>	26.6%
<i>Submitting a diagnosis</i>	17.1%
<i>Speaking with the camp’s cook</i>	15.2%
<i>Speaking with the camp’s bacteria expert</i>	12.5%
<i>Speaking with the camp’s virus expert</i>	11.2%
<i>Speaking with a sick patient</i>	11.0%
<i>Speaking with the camp nurse</i>	6.4%

Table 2. Goals and their frequencies in corpus

The performance of the proposed MLN goal recognition model was compared to one trivial and two non-trivial baseline models. The trivial baseline was the majority baseline, which always predicted the goal that appears most frequently in the training data. The non-trivial baselines were two n -gram models, *unigram* and *bigram*. The unigram model predicted goals based on the current player action only, while the bigram model considered the previous action as well. In the n -gram models, player actions were represented by a single aggregate feature that combined all three action properties: action type, location,

and narrative state. N -gram models have been used in previous goal recognition work for spoken dialogue systems (Blaylock and Allen 2003) and interactive narrative game environments (Mott, Lee, and Lester 2006). Although simplistic, the n -gram models were shown to be effective. Mott, Lee, and Lester (2006) found that unigram and bigram models achieved higher prediction accuracies than a more sophisticated Bayesian Network model. The n -gram comparison models were also created as simple MLNs. The unigram model consisted of the single weighted formula (14). The weighted formula defined for the bigram model was similar but considered two consecutive player actions at the same time.

$$\forall t,a,l,s,g : action(t,a) \wedge location(t,l) \wedge state(t,s) \Rightarrow goal(t,g) \quad (14)$$

The two n -gram models and the proposed MLN model were evaluated with ten-fold cross validation. The entire data set was partitioned into ten non-overlapping subsets, ensuring data from the same player did not appear in both the training and the testing data. Each subset of the data was used for testing exactly once. The models' performance was measured using $F1$, which is the harmonic mean of *precision* and *recall*¹. Table 3 shows the average performance of each model over ten-fold cross validation. The MLN model scored 0.484 in $F1$, achieving an 82% improvement over the baseline. The unigram model performed better than the bigram model. A one-way repeated-measures ANOVA confirmed that the differences among the three compared models were statistically significant ($F(2,18) = 71.87, p < 0.0001$). A post hoc Tukey test revealed the differences between all pairs of the three models were statistically significant ($p < .01$).

	Baseline	Unigram	Bigram	MLN
F1	0.266	0.396	0.330	0.484
Improvement over Baseline	N/A	49%	24%	82%

Table 3. $F1$ scores for MLN and baseline goal recognition models

Discussion

All three models performed better than the baseline. The best performance was achieved by the MLN model, which suggests that the proposed MLN goal recognition framework is effective at predicting player goals from actions in a complex game environment. The $F1$ score of 0.484 achieved by the MLN model may appear somewhat low. However, this is an encouraging result given the challenges posed by the data. The superiority of the MLN

¹ It should be noted that in the current work the values of *precision*, *recall*, and $F1$ are the same, because each observed player action is associated with a single goal in our data and the goal recognition model predicts a single most likely goal for each player action.

model over the n -gram models can be partially explained by the MLN's relational learning framework, which facilitates explicit modeling of associations between goals. Furthermore, the structural flexibility of undirected graphical models, which permit bidirectional causality, enables MLNs to model richer relations between actions and goals than n -gram models. The unigram model achieved higher performance than the bigram model, which is consistent with the result reported by Mott, Lee, and Lester (2006). Among the possible reasons for this is data sparsity; the bigram model considers two consecutive previous goals, which would lead to the training instances for each bigram become sparser than in the unigram model.

Inducing accurate goal recognition models has several prospective benefits for intelligent game-based learning environments. First, goal recognizers can be used to inform player-adaptive decisions by *narrative centered tutorial planners*, which comprise a particular class of player-adaptive systems that tailor events during students' game-based learning experiences in order to individualize pedagogical scaffolding and promote student engagement. Data-driven approaches to narrative-centered tutorial planning are the subject of active investigation by the CRYSTAL ISLAND research team. Second, goal recognizers can be used during data mining to inform the iterative refinement of intelligent game-based learning environments. By automatically recognizing players' goals, and identifying which actions are likely to be associated with those goals, goal recognition models can enable educational game designers to better understand common families of problem-solving paths and to identify key challenges encountered by students. Finally, recognizing players' goals will enrich in-game assessments of student learning, problem solving, and engagement, which are critical challenges for the educational games community.

Conclusions

Effective goal recognition holds considerable promise for player-adaptive games. Accurately recognizing players' goals enables digital games to proactively support gameplay experiences that feature nonlinear scenarios while preserving cohesion, coherence and believability. This paper has introduced a goal recognition framework based on Markov logic networks that accurately recognizes players' goals. Using model parameters learned from a corpus of player interactions in a nonlinear game environment, the framework supports the automated acquisition of a goal recognition system that outperforms three baseline models.

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. The authors are grateful to Ray Mooney and Pat Langley for insightful discussions on goal recognition.

References

- Bethesda Softworks. 2011. *The Elder Scrolls V: Skyrim*. Retrieved from <http://www.elderscrolls.com/skyrim>.
- 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.
- Charniak, E. and Goldman, R. P. 1993. A Bayesian Model of Plan Recognition. *Artificial Intelligence*, 64(1): 53-79.
- Cramer, K. and Y. Singer. 2003. Ultraconservative online algorithms for multiclass problems. *The Journal of Machine Learning Research*, 3: 951-991.
- Geib, C. W. and Goldman, R. P. 2009. A Probabilistic Plan Recognition Algorithm Based on Plan Tree Grammars. *Artificial Intelligence*, 173(11): 1101-1132.
- Gold, K. 2010. Training Goal Recognition Online from Low-Level Inputs in an Action-Adventure Game. In *Proceedings of the Sixth International Conference on Artificial Intelligence and Interactive Digital Entertainment*, 21-26. Menlo Park, CA: AAAI Press.
- Hu, D. H. and Yang, Q. 2008. CIGAR: Concurrent and Interleaving Goal and Activity Recognition. In *Proceedings of the Twenty-Third National Conference on Artificial Intelligence*, 1363-1368. Menlo Park, CA: AAAI Press.
- Jennings-Teats, M., Smith, G. and Wardrip-Fruin, N. (2010). Polymorph: A Model for Dynamic Level Generation. In *Proceedings of the 6th Artificial Intelligence and Interactive Digital Entertainment Conference*, 138-143. Menlo Park, CA: AAAI Press.
- Kabanza, F., Bellefeuille, P. and Bisson, F. 2010. Opponent Behavior Recognition for Real-Time Strategy Games. In *Proceedings of the AAAI-10 Workshop on Plan, Activity, and Intent Recognition*. Atlanta, GA.
- Kate, R. J. and Mooney, R. J. 2009. Probabilistic Abduction using Markov Logic Networks. In *Proceedings of the IJCAI-09 Workshop on Plan, Activity, and Intent Recognition*. Pasadena, CA.
- 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.
- Li, B., and Riedl, M. O. (2010). An Offline Planning Approach to Game Plotline Adaptation. In *Proceedings of the 6th Artificial Intelligence and Interactive Digital Entertainment Conference*, 45-50. Menlo Park, CA: AAAI Press.
- Mateas, M. and Stern, A. 2005. Structuring Content in the Façade Interactive Drama Architecture. In *Proceedings of the First Artificial Intelligence and Interactive Digital Entertainment Conference*, 93-98. Menlo Park, CA: AAAI Press.
- Mojang. 2009. *Minecraft*. Retrieved from <http://www.minecraft.net/>.
- Mott, B., Lee, S., and Lester, J. 2006. Probabilistic Goal Recognition in Interactive Narrative Environments. In *Proceedings of the Twenty-First National Conference on Artificial Intelligence*, 187-192. Menlo Park, CA: AAAI Press.
- Orkin, J. 2005. Agent Architecture Considerations for Real-Time Planning in Games. In *Proceedings of the First AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, 105-110. Menlo Park, CA: AAAI Press.
- Pynadath, D. V. and Wellman, M. P. 2000. Probabilistic State-Dependent Grammars for Plan Recognition. In *Proceedings of the Sixteenth Conference on Uncertainty in Artificial Intelligence*, 507-514. Stanford, CA: Morgan Kaufmann Publishers Inc.
- Ramirez, M. and Geffner, H. 2010. Probabilistic Plan Recognition using off-the-shelf Classical Planners. In *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence*, 1121-1126. Menlo Park, CA: AAAI Press.
- Richardson, M. and Domingos, P. 2006. Markov Logic Networks. *Journal of Machine Learning*, 62(1-2): 107-136.
- Riedl, M. O., Stern, A., Dini, D. M., and Alderman, J. M. 2008. Dynamic Experience Management in Virtual Worlds for Entertainment, Education, and Training. *International Transactions on Systems Science and Applications*, Special Issue on Agent Based Systems for Human Learning, 3(1): 23-42.
- 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. Helsinki, Finland: AUAI Press.
- Robison, J., McQuiggan, S. and Lester, J. (2009). Evaluating the Consequences of Affective Feedback in Intelligent Tutoring Systems. In *Proceedings of the International Conference on Affective Computing and Intelligent Interaction*, 37-42. IEEE.
- Rowe, J. and Lester, J. 2010. Modeling User Knowledge with Dynamic Bayesian Networks in Interactive Narrative Environments. In *Proceedings of the Sixth International Conference on Artificial Intelligence in Interactive Digital Entertainment*, 57-62. Menlo Park, CA: AAAI Press.
- Rowe, J., Shores, L., Mott, B., and Lester, J. (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. 2010. Recognizing Multi-Agent Activities from GPS Data. In *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence*, 1134-1139. Menlo Park, CA: AAAI Press.
- Singla, P. and Mooney, R. 2011. Abductive Markov Logic for Plan Recognition. In *Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence*, 1069-1075. Menlo Park, CA: AAAI Press.
- Sturtevant, N. 2011. A Sparse Grid Representation for Dynamic Three-Dimensional Worlds. In *Proceedings of the Seventh AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, 73-78. Menlo Park, CA: AAAI Press.
- Yannakakis, G. 2012. Game AI Revisited. In *Proceedings of ACM Computing Frontier Conference*.