*Appears at the AAMAS 2004 Workshop on Agent Tracking: Modeling Other Agents from Observations*

# A Utility-Based Approach to Intention Recognition

Wenji Mao
*University of Southern California*
*Institute for Creative Technologies*
13274 Fiji Way, Marina del Rey, CA 90292
mao@ict.usc.edu

Jonathan Gratch
*University of Southern California*
*Institute for Creative Technologies*
13274 Fiji Way, Marina del Rey, CA 90292
gratch@ict.usc.edu

## Abstract

*Based on the assumption that a rational agent will adopt a plan that maximizes the expected utility, we present a utility-based approach to plan recognition problem in this paper. The approach explicitly takes the observed agent's preferences into consideration, and computes the estimated expected utilities of plans to disambiguate competing hypotheses. Online plan recognition is realized by incrementally using plan knowledge and observations to change state probabilities. We also discuss the work and compare it with other probabilistic models in the paper.*

## 1. Overview

Since Schmidt *et al* [1978] first identified plan recognition as a problem in its own right, plan recognition has been applied widely to a variety of domains, including natural language understanding and generation [Allen and Perrault, 1980; Carberry, 1990], story understanding [Wilensky, 1978; Charniak and Goldman, 1989, 1993], multi-agent coordination [Huber *et al*, 1994], dynamic traffic monitoring [Pynadath and Wellman, 1995], collaborative systems [Ferguson *et al*, 1996, 1998], adventure game [Albrecht *et al*, 1998], network intrusion detection [Geib and Goldman, 2001], multiagent team monitoring [Kaminka *et al*, 2002], and so on.

Many plan recognition approaches have been proposed. Kautz and Allen [1986] presented the first formal theory of plan recognition, using McCarthy's circumscription. They define plan recognition problem as identifying a minimal set of top-level actions sufficient to explain the observed actions, and use minimal covering set as a principle for disambiguation. To deal with uncertainty inherently in plan inference, Charniak and Goldman [1989, 1993] built the first probabilistic model of plan recognition based on Bayesian reasoning. Their system supports automatically generation of a belief network (BN) from observed actions according to some network construction rules. The constructed belief network is then used for understanding a character's actions in a story. Huber, Durfee and Wellman [1994] used PRS as a general language for plan specification. They gave the dynamic mapping from PRS specification to belief networks, and applied the approach to coordinate multi-agent team. Pynadath and Wellman [2000] proposed a probabilistic method that was based on parsing. Their approach employs probabilistic state-dependent grammars (PSDGs) to represent an agent's plan generation process. The PSDG representation, together with inference algorithms supports efficient answering of restricted plan recognition queries. More recently, Bui *et al* [2002, 2003] proposed an online probabilistic policy recognition method based on the abstract hidden Markov model (AHMM) and the extension of AHMM allowing for policies with memories (AHMEM). In their frameworks, scalability in policy recognition in the models is achieved by using an approximate inference scheme (i.e., Rao-Blackwellised Particle Filter). Besides Bayesian models, some probabilistic approaches are based on Dempster-Shafer theory, e.g., Carberry [1990] and Bauer [1995, 1996].

Though the approaches differ, most plan recognition systems infer a hypothesized plan based on observed actions. World states and in particular, state desirability (typically represented as utilities of states) are rarely considered in the recognition. On the other hand, in many real-world applications, utilities of different outcomes are already known [Blythe, 1999]. A planning agent usually takes into account that actions may have different outcomes, and some outcomes are more desirable than the others. Therefore, when an agent makes decisions and acts on the world, the agent needs to balance between different possible outcomes in order to maximize the expected utility of overall goal attainment.

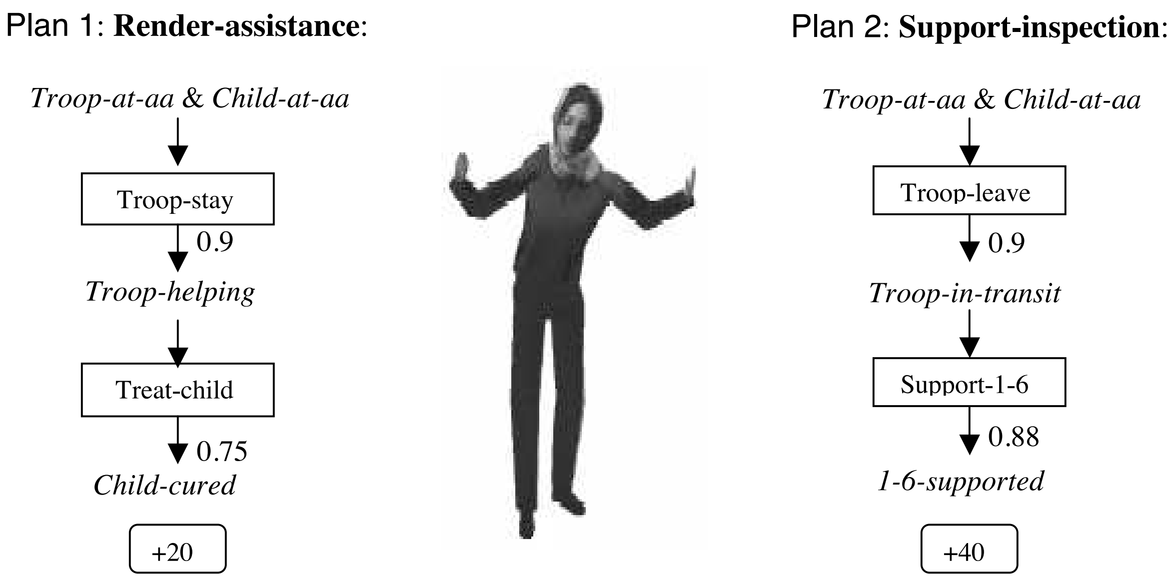


**Figure 1**. Competing Plans of the Troop from the Mother's Perspective

Utility and rationality issues have been explored in earlier work in AI (e.g., rational assumptions, [Doyle, 1992]). Plan recognition can be viewed as inferring the decision-making strategy of the observed agent. So it is natural to assume that a rational agent will adopt a plan that maximizes the expected utility. While current probabilistic approaches capture the fact of how well the observed actions support a hypothesized plan, the missing part is the utility computation.

In this paper, we take a decision-theoretic view and explicitly consider outcome utilities in plan recognition. There are different ways to address the utility issue in plan recognition, and we believe different approaches are appropriate for different problem domains. In [Mao and Gratch, 2004], we proposed three possible ways to refine the problem. Two of them are the extensions within probabilistic reasoning framework, by incorporating utility nodes into belief networks and using them to adjust the prior and conditional probabilities in the conditional probability tables (CPTs). The third proposal takes another viewpoint. It computes the expected utilities of plans and chooses the plan candidate with maximal expected utility (given the evidence so far). This paper focuses on the third solution.

The remainder of the paper is organized as follows. To motivate the work, in *Section 2*, we illustrate an example from our leadership training environment. *Section 3* introduces the plan representation we adopt in this work. In *Section 4*, we present our approach to intention recognition based on computing the expected utility of possible plans. We then illustrate the approach using the motivating example. *Section 5* discusses the related work and compares our approach with Bayesian probabilistic models. Finally, in *Section 6*, we summarize the paper and raise some future work.

## 2. Motivating Example

Consider an example in the context of the Mission Rehearsal Exercise (MRE) leadership trainer we are developing [Rickel *et al*, 2002]. A human trainee was in command of a troop in peacekeeping operation to support another unit, eagle 1-6. In route, one of the troop's vehicles severely injured a civilian child. The trainee must balance whether to continue the mission or render aid. Many decisions and outcomes are possible. In the scenario, the injured boy's mother (a virtual agent) observed the troop's actions, trying to infer the troop's plan and predict subsequent actions.

The mother had a simple task model of the troop. She kept two possible plans of the troop in her mind (see *Figure 1*). Plan *Render-assistance* is composed of *Troop-stay* and *Treat-child*. Plan *Support-inspection* consists of *Troop-leave* and *Support-eagle-1-6*. For simplification, *Troop-stay*, *Troop-leave*, *Treat-child* and *Support-eagle-1*-6 are all primitive actions. Each action has non-deterministic effects as shown in the figure. *Troop-helping*, *Troop-in-transit*, *Child-cured* and *Eagle-1-6-supported* are action effects. Outcomes *Child-cured* and *Eagle-1-6-supported* are goals of *Render-assistance* and *Support-inspection*, with utility values of 20 and 40, respectively (The example here is simplified, as a plan may have more than one outcome).

Typically, using belief networks in Bayesian reasoning, the first step is to create a random variable in the belief network representing the top-level plan. The rest of the variables, dependency arcs and probability values all provide evidence for or against the proposition that this top plan is being pursued by the observed agent. The next step is to create a random variable for each action in the plan. Because it is the adoption of the top plan that causes the execution of these actions, there is a dependency arc from the top plan to each of the action node in the belief network. Then the last step is to add evidence variables to the belief network, to represent the dependencies between actions and observed evidence or context.

Suppose in our example, initially plans *Render-assistance* and S*upport-inspection* have the same prior probability, and the CPTs of each node in the belief network of *Plan 1* are identical to those of corresponding nodes in *Plan 2*. But the two plans have different outcome utilities and different probabilities to achieve the outcomes as shown in *Figure 1*. As the scenario proceeded, assume the trainee decided to send two of his four squads moving forward to support eagle 1-6. The mother observed half of the troop staying and half leaving. Since the observed action equally support the two plans, Bayesian reasoning will infer that the two plans have the identical posterior probability (see results in [Mao and Gratch, 2004]).

But actually, in this example, the outcome of *Support-inspection* is more desirable to the troop, and more likely to be achieved. Intuitively, the troop should be more likely to actively pursue this plan. Current probabilistic models could not make this distinction, as neither outcome utility nor outcome probability is explicitly used in plan inference. To address the problem, we first introduce the plan representation used in our approach.

# 3. Plan Representation

We adopt probabilistic plan representation in our approach. Each action consists of a set of preconditions and effects. Actions can have non-deterministic effects (denoted as *Effect_prob*($A$, $e$), where $A$ is an action and $e$ is an effect of $A$), as well as conditional effects. To represent success or failure of action execution, actions have execution probability (denoted as *Execute_prob*($A$), where $A$ is an action). The likelihood of preconditions and effects is represented by probability values. The desirability of action effects is represented by *utility* values.

In a hierarchical plan representation, an action can be *primitive* (i.e., an action that is directly executed by an agent) or *abstract*. An abstract action may be decomposed in multiple ways and each decomposition consists of a sequence of primitive or abstract sub-actions. A *non-decision node* in plan structure is an action that can only be decomposed in one way. A *decision node*, on the other hand, can be decomposed in multiple ways and an agent must decide amongst the options. An *outcome* is a primitive action effect (or a group of primitive action effects when several action effects have a single utility value) with a non-zero utility value.

A *plan* is represented as an action sequence. Each plan is associated with an intended *goal* (with a positive utility value). When a plan contains abstract actions, this denotes a set of primitive plans that would result from decomposing these abstract actions into primitive ones. As there might be side effects in goal attainment, a plan may have more than one desirable/undesirable outcome including the goal itself. From a decision-theoretic point of view, the *expected utility* of a plan represents the overall benefit or disadvantage of the plan. We shall discuss how to compute plan utility in the next section.

# 4. Intention Recognition

Intention recognition is to infer another agent's goal/plan based on a perceiving agent's observations. The plan inference is from the perceiving agent's perspective, using the knowledge and information the perceiver has about the observed agent. We model the perceiver's inference process by computing the estimated expected utilities of the observed agent's possible plans, based on the perceiver's plan knowledge about and observations of the observed agent.

## 4.1 Computing Plan Utility

The computation of plan utility is similar to that in decision-theoretic planning (e.g., DRIPS, [Haddawy and Suwandi, 1994]), based on an abstraction hierarchy of operators. However, since plan recipes are already known, we compute an exact utility value rather than a range of utility values for searching the plan space as in decision-theoretic planning.

Since action theory (i.e., actions, their preconditions and effects) is known, this information can be utilized by the recognizer. In our approach, we take the observations of both actions and state information into consideration, and use them as evidence to incrementally update state probabilities. The updated state probabilities change the probabilities of other action preconditions, which in turn, will change the probabilities of other action execution and effect occurrences. Thus, the expected utilities of associated plans are updated incrementally with the changes of observations.

Let $E$ be the evidence. If an action $A$ is observed, the execution probability of $A$ is 1.0. The probability of each precondition of $A$ must be 1.0 (excluding those deleted by delete effects), and the probability of each effect of $A$ is equal to its effect probability. If $A$ has conditional effects, the probability of the consequent of each conditional effect of $A$ is equal to the probability of the antecedent of the conditional effect.

- IF $x \in Precondition(A) - Del\_effect(A)$, $P(x|E) = 1.0$
- IF $x \in Add\_effect(A)$, $P(x|E) = Effect\_prob(A, x)$
- IF $x \in Del\_effect(A)$, $P(x|E) = 1.0 - Effect\_prob(A, \neg x)$
- IF $x \in Consequent(Cond\_effect(A))$
  $P(x|E) = P(Antecedent(Cond\_effect(A)))$

If an action effect $x$ is observed, then $P(x|E)$=1.0. If an action $A$ is observed, then $P(A|E)$=1.0 , otherwise the probability of the successful execution of $A$ given $E$ is computed as

$$P(A \mid E) = (\prod_{x \in precondition(A)} P(x \mid E)) \times Execute\_prob(A)$$

Let $O_i$ be the set of outcomes of a primitive plan $P_i$, an outcome $o_j \in O_i$. $A_1, \ldots, A_k$ is the action sequence in $P_i$ that leads to $o_j$. The probability of $o_j$ is computed as

$$P(o_j \mid E) = (\prod_{i=1,\ldots,k} P(A_i \mid E)) \times Effect\_prob(A_k, o_j)$$

The estimated expected utility of a primitive plan $P_i$ given $E$ is computed as

$$EU(P_i \mid E) = \sum_{o_j \in O_i} (P(o_j \mid E) \times Utility(o_j))$$

In a hierarchical plan representation, if an abstract action is a non-decision node, the expected utility of the abstract action is the sum of the utilities of its sub-actions. If an abstract action is a decision node, the expected utility of the abstract action is the maximum of the utilities of its sub-actions (because we assume that an agent tries to maximize the expected utility). The utility of the root node in the plan hierarchy is the estimated expected utility of the abstract plan.

After the computation, the plan with the highest estimated expected utility is chosen as the hypothesized plan given the evidence so far.

### 4.2 Illustration

Now we return to the example introduced in *Section 2*. Initially, the troop was at the accident area (*Troop-at-aa*) and the child was at the accident area (*Child-at-aa*). From the mother's perspective, assume the execution probability of each action is 0.95. The prior probabilities of states and the execution probabilities of actions are as follows.

$P(Troop\text{-}at\text{-}aa) = P(Child\text{-}at\text{-}aa) = 1.0$

$P(Troop\text{-}helping) = P(Troop\text{-}in\text{-}transit) = 0.5$
$P(Child\text{-}cured) = P(1\text{-}6\text{-}supported) = 0.5$

$Execute\_prob(Troop\text{-}stay) = 0.95$
$Execute\_prob(Troop\text{-}leave) = 0.95$
$Execute\_prob(Treat\text{-}child) = 0.95$
$Execute\_prob(Support\text{-}1\text{-}6) = 0.95$

From *Figure 1*, we also know

$Effect\_prob(Troop\text{-}stay, Troop\text{-}helping) = 0.9$
$Effect\_prob(Troop\text{-}leave, Troop\text{-}in\text{-}transit) = 0.9$
$Effect\_prob(Treat\text{-}child, Child\text{-}cured) = 0.75$
$Effect\_prob(Support\text{-}1\text{-}6, 1\text{-}6\text{-}supported) = 0.88$

The observation of the troop's action equally supports *Troop-stay* and *Troop-leave*, that is,

$P(Troop\text{-}stay|E) = P(Troop\text{-}leave|E) = 0.5$

So we have

$P(Troop\text{-}helping|E) = 0.45$
$P(Troop\text{-}in\text{-}transit|E) = 0.45$

Now compute the probabilities of executing *Treat-child* and *Support-1-6* given the evidence

$P(Treat\text{-}child|E) = 0.4275$
$P(Support\text{-}1\text{-}6|E) = 0.4275$

Now compute the outcome probabilities

$P(Child\text{-}cured|E) = 0.1603$
$P(1\text{-}6\text{-}supported|E) = 0.1881$

Now compute the estimated expected utilities of *Plan 1* and *Plan 2* given current observations

$EU(Plan_1|E) = 3.206$
$EU(Plan_2|E) = 7.524$

So the mother recognized that the troop was pursuing the plan *Support-eagle-1-6*.

Although this is an oversimplified example, for the more complex cases including multiple outcomes, conditional probabilities, abstract actions, etc, the algorithm can be applied the same way.

## 5. Discussions

Some probabilistic approaches have considered the influence of world states on plan recognition, e.g., Goldman *et al* [1999], Pynadath and Wellman [1995] and Bui *et al* [2002]. Goldman *et al* [1999] argued that the state of the world would influence an agent's decision to pursue plans. They proposed a plan execution model using probabilistic Horn abduction. In traffic monitoring in [Pynadath and

Wellman, 1995], actions themselves are as unobservable. The recognizer infers a driver's plan based on some observable action effects.

Bui *et al* [2002] use the abstract hidden Markov model for online policy recognition. While their framework is a stochastic model, they exploit the special properties of the AHMM structure that lead to efficient plan recognition algorithms. We did not adopt a Markov model in our work for some considerations. A Markov-based approach requires relatively large state space, and assumes fixed goals. The core technologies of our application system center on a common representation of plan knowledge, which is shared and reused among different system components. Besides, in modeling realistic virtual agents, we would like to give our agents the flexibility of strategically varying their interpretations of outcome desirability, as a result of coping with specific situations [Marsella and Gratch, 2003].

In Pynadath and Wellman [2000] and Bui *et al* [2002], an agent's utility functions are implicitly taken into account, as their approaches capture the likelihood that the agent will expand a plan in a particular way (e.g., the expansion probabilities of Pynadath and Wellman's PSDGs). The main difference between our approach and most other probabilistic approaches is that we explicitly take an agent's preferences into consideration. However, in doing so, we are not claiming that a recognizer must always know the exact utility functions of the observed agent, but rather, we think that if the recognizer does know (or partially know), this information can be utilized as evidence to impact the recognition process. Indeed, in many real world applications as well as in our own, utilities of states are already there [Blythe, 1999].

Bayesian probabilistic models view plan recognition as abduction, and use Bayesian rules to compute the best candidate plan. Though probabilistic reasoning is advantageous in accounting for how well the observed actions support a hypothesized plan, the inference itself requires large numbers of prior and conditional probabilities. In many situations, these probabilities are hard to obtain. There is no good answer for where the numbers come from.

We view plan recognition as recognizing the decision-making strategy of the observed agent, and use maximizing expected utilities of plans as criterion for disambiguation. Our approach also needs prior probabilities, such as prior probabilities of states, action success/failure and probabilities of action effects. Some probabilities, like non-deterministic and/or conditional action effects are already available in many systems with a planning component. State probabilities and probabilities of action execution are relatively easier to obtain comparing with the CPTs required in belief networks. So it partly eases the burden of defining large numbers of prior and conditional probabilities as in the probabilistic models, but the trade-off is that our approach is an approximate one, and it does not consider state dependencies.

The knowledge about actions, their preconditions and effects are typically available in a plan-based system. Our approach makes use of this knowledge, using observations of actions and effects to change the probabilities of states. However, there is no strong assumption of the observability of actions or effects in our approach, and a sequence of observations can be processed incrementally in the same way. Finally, our approach is compatible with the idea of decision-theoretic planning. It helps computer systems share representation, intermediate results and underlying techniques for both planning and recognition, and allows systems to interleave between planning and inferring plans depending on the tasks at hand.

# 6. Summary and Future Work

Based on the assumption that a rational agent will adopt a plan that maximizes the expected utility, we view plan recognition as inferring the decision-making strategy of the observed agent. In this paper, we present a utility-based approach to plan recognition problem. The approach explicitly takes the observed agent's preferences into consideration, and computes the estimated expected utilities of plans to disambiguate competing hypotheses. We consider both actions and state information in the recognition process. Online plan recognition is realized by incrementally using plan knowledge and observations to change state probabilities.

We discuss the work and compare it with other probabilistic models in the paper. We point out some limitations in Bayesian inference, and show where our approach might help. Though the approach seems sufficient for our practical application and compatible with the existing system representation, the heuristic of narrowing the hypotheses space based on the rational assumption is not evaluated. In the future, we need to collect experimental data and run experiments in realistic scenarios to test the effectiveness of the approach. As our virtual environment supports face-to-face interactions of humans and virtual agents, it provides an ideal testbed for evaluating the work.

# Acknowledgement

This paper was developed with funds of the Department of the Army under contract number DAAD 19-99-D-0046. We thank our colleagues on the MRE project for their collaboration. Any opinions, findings and conclusions or recommendations in this paper are those of the authors and do not necessarily reflect the views of the Department of the Army.

## References


D. W. Albrecht, I. Zukerman and A. E. Nicholson. Bayesian Models for Keyhole Plan Recognition in an Adventure Game. *User Modeling and Use-Adapted Interaction*, 8(1-2):5-47, 1998.

J. F. Allen and R. Perrault. Analyzing Intention in Utterances. *Artificial Intelligence*, 15(3):143-178, 1980.

M. Bauer. A Dempster-Shafer Approach to Modeling Agent Preferences for Plan Recognition. *User Modeling and User-Adapted Interaction*, 5(3-4):317-348, 1995.

M. Bauer. Acquisition of User Preferences for Plan Recognition. *Proceedings of the Fifth International Conference on User Modeling*, 1996.

J. Blythe. Decision-Theoretic Planning. *AI Magazine*, 20(2):37-54, 1999.

H. H. Bui, S. Venkatesh and G. West. Policy Recognition in the Abstract Hidden Markov Model. *Journal of Artificial Intelligence Research*, 17:451-499, 2002.

H. H. Bui. A General Model for Online Probabilistic Plan Recognition. *Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence*, 2003.

S. Carberry, *Plan Recognition on Natural Language Dialogue*. The MIT Press, 1990.

E. Charniak and R. Goldman. A Semantics for Probabilistic Quantifier-Free First-Order Languages, with Particular Application to Story Understanding. *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*. 1989.

E. Charniak and R. Goldman. A Bayesian Model of Plan Recognition. *Artificial Intelligence*, 64(1):53-79, 1993.

J. Doyle. Rationality and Its Roles in Reasoning. *Computational Intelligence*, 8(2):376-409, 1992.

G. Ferguson, J. F. Allen and B. Milller. TRAINS-95: Towards a Mixed-Initiative Planning Assistant. *Proceedings of the Third Conference on Artificial Intelligence Planning Systems*, 1996.

G. Ferguson and J. F. Allen. TRIPS: An Intelligent Integrated Problem-solving Assistant. *Proceedings of the Fifteenth National Conference on Artificial Intelligence*, 1998.

C. W. Geib and R. P. Goldman. Plan Recognition in Intrusion Detection Systems. *Proceedings of the Second DARPA Information Survivability Conference and Exposition*, 2001.

R. P. Goldman, C. W. Geib and C. A. Miller. A New Model of Plan Recognition. *Proceedings of the Fifteenth Conference on Uncertainty in Artificial Intelligence*, 1999.

P. Haddawy and M. Suwandi. Decision-Theoretic Refinement Planning Using Inheritance Abstraction. *Proceedings of the Second International Conference on Artificial Intelligence Planning*, 1994.

M. J. Huber, E. H. Durfee and M. P. Wellman. The Automated Mapping of Plans for Plan Recognition. *Proceedings of the Tenth International Conference on Uncertainty in Artificial Intelligence*, 1994.

G. Kaminka, D. V. Pynadath and M. Tambe. Monitoring Teams by Overhearing: A Multiagent Plan Recognition Approach. *Journal of Artificial Intelligence Research*, 17:83-135, 2002.

H. A. Kautz and J. F. Allen. Generalized Plan Recognition. *Proceedings of the Fifth National Conference on Artificial Intelligence*, 1986.

W. Mao and J. Gratch. Decision-Theoretic Approaches to Plan Recognition. *ICT Technical Report* (http://www.ict.usc.edu/publications/ICT-TR-01-2004.pdf), 2004.

S. Marsella and J. Gratch. Modeling Coping Behavior in Virtual Humans: Don't Worry, Be Happy. *Proceedings of the Second International Joint Conference on Autonomous Agents and Multiagent Systems*, 2003.

D. V. Pynadath and M. P. Wellman. Accounting for Context in Plan Recognition, with Application to Traffic Monitoring. *Proceedings of the Eleventh International Conference on Uncertainty in Artificial Intelligence*, 1995.

D. V. Pynadath and M. P. Wellman. Probabilistic State-Dependent Grammars for Plan Recognitions. *Proceedings of the Sixteenth Conference on Uncertainty in Artificial Intelligence*, 2000.

J. Rickel, S. Marsella, J. Gratch, R. Hill, D. Traum and W. Swartout. Toward a New Generation of Virtual Humans for Interactive Experiences. *IEEE Intelligent Systems*, 17(4):32-38, 2002.

C. F. Schmidt, N. S. Sridharan and J. L. Goodson. The Plan Recognition Problem: An Intersection of Psychology and Artificial Intelligence. *Artificial Intelligence*, 11(1-2):45-83, 1978.

R. Wilensky. Understanding Stories Involving Recurring Goals. *Cognitive Science*, 2:235-266, 1978.