

# Graph Construction and Analysis as a Paradigm for Plan Recognition

**Jun Hong**

School of Information and Software Engineering  
University of Ulster at Jordanstown  
Newtownabbey, Co. Antrim BT37 0QB, United Kingdom  
j.hong@ulst.ac.uk

## Abstract

We present a novel approach to plan recognition in which graph construction and analysis is used as a paradigm. We use a graph structure called a Goal Graph for the plan recognition problem. The Goal Graph is first constructed to represent the observed actions, the state of the world, and the achieved goals at consecutive time steps. It also represents various connections between nodes in the Goal Graph. The Goal Graph can then be analysed at each time step to recognise those achieved goals that are consistent with the actions observed so far. The Goal Graph analysis can also reveal valid plans for the recognised goals or part of the recognised goals. We describe two algorithms, GoalGraphConstructor and GoalGraphAnalyser, based on this paradigm. These algorithms are sound, polynomial-time and polynomial-space. The algorithms have been tested in two domains with up to 245 goal schemata and 100000 possible goals. They perform well in these domains in terms of efficiency, accuracy and scalability.

## Introduction

Plan recognition involves inferring the goal of an agent from a set of observed actions and organising the observed actions into a plan structure for the goal. We introduce a novel approach to plan recognition, in which graph construction and analysis is used as a paradigm. Our attempt to do so is in spirit influenced by Blum and Furst's effort on planning with Planning Graphs (Blum & Furst 1995), (Blum & Furst 1997). They introduced a new graph-based approach to planning in STRIPS domains, in which a graph structure called a Planning Graph is first constructed explicitly rather than searching immediately for a plan as in standard planning methods. The Planning Graph is then analysed to generate possible plans.

Since being first introduced, further developments have been made with regard to handling more expressive representation languages (Gazen & Knoblock 1997), (Anderson, Smith, & Weld 1998), (Koehler *et al.* 1997), that allow the use of disjunctive preconditions, conditional effects, and universally quantified preconditions (goal descriptions) and effects in action and goal representation.

Copyright © 2000, American Association for Artificial Intelligence (www.aaai.org). All rights reserved.

We propose to use a different graph structure, called a Goal Graph, for the plan recognition problem. Instead of searching for a plan as in most plan recognition systems, a Goal Graph is first constructed to represent the observed actions, the state of the world as it is changed by these actions, and the fully or partially achieved goals at consecutive time steps. Connections are also made between different kinds of nodes in the Goal Graph. The constructed Goal Graph can then be analysed at each time step to recognise those fully or partially achieved goals that are consistent with the actions observed so far. The Goal Graph analysis also reveals causal links over actions and goals so that valid plans for the recognised goals or part of the recognised goals can be further recognised.

We describe two algorithms based on this paradigm. The GoalGraphConstructor takes a set of partially ordered actions as they are observed and constructs a Goal Graph. The GoalGraphAnalyser analyses the constructed Goal Graph to recognise consistent goals and valid plans. We prove that our algorithms are sound, polynomial-time and polynomial-space. The algorithms have been tested on a 500 MHz Pentium III in two domains. In the extended briefcase domain, we increase the number of locations and objects to create a series of sets of up to over 100000 possible goals for testing the scalability of our algorithms where the approximate linear time performance has been achieved. In the Unix domain, we use a set of data collected in the Unix domain at the University of Washington with over 245 goal schemata and over 10000 possible goals. In this domain, on average it only takes less than a CPU second to update the Goal Graph when an observed action is processed and usually only a very small number of consistent goals remain after a sequence of observed actions has been processed.

## The Domain Representation

We use an ADL-like representation (Pednault 1989), including actions with conditional and universally quantified effects, and existentially as well as universally quantified preconditions and goal descriptions. A plan recognition problem consists of

- A set of action schemata specifying primitive actions.
- A finite, dynamic universe of typed objects.
- A set of propositions called the Initial Conditions.

- A set of goal schemata specifying possible goals.
- A set of observed actions that are partially ordered.
- An explicit notion of discrete time.

The solution to a plan recognition problem consists of a set of recognised goals that are consistent with the observed actions together with the valid plans consisting of the observed actions for the recognised goals or part of the recognised goals.

The goal schema consists of a set of goal descriptions. The action schema consists of a set of preconditions and a set of effects. A goal is a ground instance of a goal schema. An action is a ground instance of an action schema. The set of goal descriptions for a goal must be satisfied in the state of the world when the goal is fully achieved. If some but not all goal descriptions are satisfied instead, the goal is partially achieved. The set of preconditions must be satisfied in the state of the world before an action can be executed. The set of effects are taken in the state of the world when an action is executed.

The representation of a simple example domain extended from Pednault's famous example (Pednault 1988) is shown in Figure 1. It involves transportation of two physical objects, a dictionary and a chequebook, between home and office using a briefcase. We assume that only one physical object can be carried in the briefcase at a time. The extended briefcase domain consists of three action schemata and three goal schemata.

In the actual implementation of our plan recognition algorithms, universally quantified preconditions and effects, and conditional effects in an action schema are eliminated and equivalent schemata are created. We use a particular approach we call *dynamic expansion*. Dynamic expansion involves two steps. In the first step, universally quantified preconditions and effects in an action schema are dynamically compiled into the corresponding Herbrand base taking into account the state of the universe at the current time step. In the second step, conditional effects are further eliminated. The universally quantified goal descriptions in a goal schema are treated in the same way as the universally quantified preconditions in an action schema.

## Goal Graphs, Valid Plans and Consistent Goals

### Goal Graphs

We first describe the structure of the Goal Graph. A Goal Graph is a directed, levelled graph. The levels alternate between proposition levels containing proposition nodes (each labelled with a proposition or negation of a proposition) representing the propositions true or explicitly known to be false in the state of the world at consecutive time steps, goal levels containing goal nodes (each labelled with a goal) representing goals fully or partially achieved at consecutive time steps and action levels containing action nodes (each labelled with an action) representing actions observed at consecutive time steps. The levels in a Goal Graph start with a proposition level at time step 1 that consists of one node for each proposition true in the Initial Conditions. They end

```
(:action mov-b
  :paras (?l ?m - loc)
  :pre (and (neg ?l ?m)(at B ?l))
  :eff (and (at B ?m) (neg (at B ?l))
    (forall (?z - physob)
      (when (in ?z)
        (and (at ?z ?m)
          (neg (at ?z ?l))))))) )

(:action put-in
  :paras (?x - physob ?l loc)
  :pre (and (neg ?x B)(at ?x ?l)(at B ?l))
  (forall (?z - physob)
    (not (in ?z))) )
  :eff (in ?x) )

(:action take-out
  :paras (?x - physob)
  :pre (in ?x)
  :eff (neg (in ?x)) )

(:goal move-object
  :paras (?x - physob ?l ?m - loc)
  :goal-des (and (neg ?l ?m)
    (neg ?x B)
    (imply (neg (at ?x ?l))
      (at ?x ?m))) )

(:goal keep-object-at
  :paras (?x - physob ?l - loc)
  :goal-des (and (neg ?x B)
    (imply (at ?x ?l)
      (not (in ?x)))) )

(:goal keep-object-in
  :paras (?x - physob)
  :goal-des (in ?x) )
```

Figure 1: The representation of the extended briefcase domain

with a goal level at the last time step that consists of a node for each of the goals fully or partially achieved so far.

The goal nodes in goal-level  $i$  are connected by description edges to their goal descriptions in proposition-level  $i$ . The action node in action-level  $i$  is connected by precondition edges to its preconditions in proposition-level  $i$ , and by effect edges to its effects in proposition-level  $i + 1$ . Those proposition nodes in proposition-level  $i$  are connected via persistence edges to the corresponding proposition nodes in proposition-level  $i + 1$  if their truth values have not been affected by the effects of the action in action-level  $i$ . In the Goal Graph shown in Figure 2, three actions have been observed at three consecutive time steps: (mov-b O H), (put-in D H), and (mov-b H O). The Initial Conditions consist of: (at B O), (at D H) and (at C H). Action and goal nodes are on the top and bottom parts of the graph respectively. The proposition nodes are in the middle part of the graph.

### Valid Plans

We now define what we mean when we say a set of partially ordered actions forms a valid plan for a goal given the Initial Conditions.

**Definition 1 (Causal Link)** *Let  $a_1$  and  $a_2$  be two actions. There exists a causal link between  $a_1$  and  $a_2$ , written as*




Figure 2: A Goal Graph for an example of the extended briefcase domain

$a_1 \rightarrow a_2$ , if and only if one of the effects of  $a_1$  satisfies one of the preconditions of  $a_2$ .

A goal can be treated as an action with goal descriptions as its preconditions and an empty set of effects. Therefore causal links can also be established from actions to goals.

**Definition 2 (Valid Plan)** Let  $g$  be a goal, and  $P = \langle A, O, L \rangle$  where  $A$  is a set of actions,  $O$  is a set of temporal ordering constraints,  $\{a_i < a_j\}$ , over  $A$ , and  $L$  is a set of causal links,  $\{a_i \rightarrow a_j\}$ , over  $A$ . Let  $I$  be the Initial Conditions.  $P$  is a valid plan for  $g$ , given  $I$ , if and only if

1. the actions in  $A$  can be executed in  $I$  in any order consistent with  $O$ ;
2. the goal  $g$  is fully achieved after the actions in  $A$  are executed in  $I$  in that order.

**Consistent Goals**

We finally define what we mean when we say a goal is consistent with a set of partially ordered actions that have been observed so far.

**Definition 3 (Relevant Action)** Given a goal  $g$  and a set of partially ordered actions,  $\langle A, O \rangle$ , where  $A$  is a set of actions,  $O$  is a set of temporal ordering constraints,  $\{a_i < a_j\}$ , over  $A$ , an action  $a \in A$  is said to be relevant to  $g$  in the context of  $\langle A, O \rangle$ , if and only if

1. there exists a causal link,  $a \rightarrow g$ ; or
2. there exists a causal link,  $a \rightarrow b$ , where  $b \in A$  is a relevant action to  $g$  and  $a < b$  is consistent with  $O$ .

**Definition 4 (Consistent Goal)** A goal  $g$  is consistent with a set of partially ordered actions,  $\langle A, O \rangle$ , if and only if every  $a \in A$  is relevant to  $g$  in the context of  $\langle A, O \rangle$ .

**Proposition 1 (Valid Plan for Consistent Goal)** Let  $\langle A, O \rangle$  be a set of partially ordered actions that have been observed so far,  $I$  be the Initial Conditions before  $\langle A, O \rangle$ ,  $g$  be a goal consistent with  $\langle A, O \rangle$ . Given  $I$ ,  $P = \langle A, O, L \rangle$  where  $L$  is a set of causal links,  $\{a_i \rightarrow a_j\}$ , over  $A$ , is a valid plan for  $g$  when  $g$  is fully achieved after  $\langle A, O \rangle$ , or otherwise for the achieved part of  $g$  when  $g$  is partially achieved after  $\langle A, O \rangle$ .

Proposition 1 follows Definition 2, 3 and 4. Especially when  $g$  is partially achieved, let  $g'$  be the achieved part of  $g$ . So  $g'$  is fully achieved and  $P = \langle A, O, L \rangle$  is a valid plan for  $g'$ .

**Plan Recognition Algorithms**

Our plan recognition algorithms run in a two-stage cycle: Goal Graph construction and analysis. This two-stage cycle continues until no action is further observed.

**Constructing a Goal Graph**

The GoalGraphConstructor starts with a Goal Graph that consists of only proposition-level 1 with nodes representing the Initial Conditions.

Given a Goal Graph ending with proposition-level  $i$ , the GoalGraphConstructor first extends the Goal Graph to goal-level  $i$  with nodes representing goals fully or partially achieved at time step  $i$ . Meanwhile, if a node in proposition-level  $i$  satisfies a goal description, a description edge connecting the proposition node to the goal node is added onto the Goal Graph. We call this process *Goal Expansion*.

When an action is observed at time step  $i$ , the GoalGraphConstructor then extends the Goal Graph ending with goal-level  $i$ , to action-level  $i$  with a node representing the observed action. At the same time, the algorithm also extends the Goal Graph to proposition-level  $i + 1$  with nodes representing propositions true or explicitly known to be false after the action has been observed. Meanwhile, if a node in proposition-level  $i$  satisfies a precondition of the action, a precondition edge connecting the proposition node to the action node is added onto the Goal Graph. For every effect of the action, the GoalGraphConstructor simply adds a proposition node to proposition-level  $i + 1$ . The effect edge from the action node to the proposition node is also added onto the Goal Graph. Every proposition node at proposition-level  $i$  is brought forward to proposition-level  $i + 1$  by a maintenance action if its truth value has not been changed by the effect of the action observed at time step  $i$  (and it has not been added onto the Goal Graph by the effect of the action). Persistence edges connecting the proposition nodes at two proposition

levels are added onto the Goal Graph. We call this process *Action Expansion*.

**Theorem 1 (Polynomial Size and Time)** *Consider a plan recognition problem with  $t$  observed actions in  $t$  time steps, a finite number of objects at each time step,  $p$  propositions in the Initial Conditions, and  $m$  goal schemata each having a constant number of parameters. Let  $l_1$  be the largest number of the effects of any of the action schemata,  $l_2$  be the largest number of the goal descriptions of any of goal schemata. Let  $n$  be the largest number of objects at all time steps. Then, the size of the Goal Graph of  $t + 1$  levels created by the *GoalGraphConstructor*, and the time needed to create the graph, are polynomial in  $n, m, p, l_1, l_2$  and  $t$ .*

The maximum number of nodes in any proposition level is  $O(p + l_1 t)$ . Let  $k$  be the largest number of parameters in any goal schema. Since any goal schema can be instantiated in at most  $n^k$  distinct ways, the maximum numbers of nodes and edges in any goal level are  $O(mn^k)$  and  $O(l_2 mn^k)$  respectively. It is obvious that the time needed to create both nodes and edges in any level is polynomial in the number of nodes and edges in the level.

**Theorem 2** *The *GoalGraphConstructor* is sound: Any goal it adds to the Goal Graph at time step  $i$  is one either fully or partially achieved at time step  $i$  in the state of the world. The algorithm is complete: If a goal has been either fully or partially achieved by the observed actions up to time step  $i - 1$ , then the algorithm will add it to the Goal Graph at time step  $i$  under the assumption that all possible goals are restricted to the categories of goal schemata.*

Proposition-level  $i$  of the Goal Graph represents the state of the world at time step  $i$  that has been changed from the Initial Conditions after the actions have been observed at time step 1, ...,  $i - 1$ . A fully or partially achieved goal in goal-level  $i$  of the Goal Graph is one fully or partially achieved in the state of the world at time step  $i$ . On the other hand, goal-level  $i$  of the Goal Graph consists of all possible instances of the goal schemata that are fully or partially achieved in the state of the world at time step  $i$ .

### Recognising Consistent Goals and Valid Plans

We assume that every observed action is relevant to the goal intended by the agent in the context of the agent's actions. Therefore, the goal intended by the agent is consistent with the observed actions and a goal may be the intended goal if it is consistent with the set of the observed actions. Theorem 3 and Theorem 4 state how the recognition of the consistent goals can be achieved by the analysis of a constructed Goal Graph.

**Theorem 3** *Given a Goal Graph, there exists a causal link,  $a_i \rightarrow g_j$  between an action  $a_i$  at time step  $i$  and a goal  $g_j$  at time step  $j$ , where  $i < j$ , if  $a_i$  is connected to  $g_j$  via a path of an effect edge, zero or more persistence edges and a description edge. We call such a path a causal link path between  $a_i$  and  $g_j$ .*

**Theorem 4** *Given a Goal Graph, there exists a causal link,  $a_i \rightarrow a_j$ , and a temporal ordering constraint,  $a_i < a_j$ ,*

*between an action  $a_i$  at time step  $i$  and another action  $a_j$  at time step  $j$ , where  $i < j$ , if  $a_i$  is connected to  $a_j$  via a path of an effect-edge, zero or more persistence-edges and a precondition-edge. We call such a path a causal link path between  $a_i$  and  $a_j$ .*

Based on the structure of the Goal Graph, we can prove the existence of the causal link,  $a_i \rightarrow g_j$ , in Theorem 3 and  $a_i \rightarrow a_j$ , in Theorem 4. It is obvious that a causal link path between  $a_i$  and  $a_j$  guarantees the temporal ordering constraint,  $a_i < a_j$ .

Given a constructed Goal Graph of  $t$  levels, the *GoalGraphAnalyser* recognises every consistent goal from the goals in goal-level  $t$  by deciding whether every observed action is relevant to it. This is done by first finding those relevant actions from the observed actions, that are connected to the goal by causal link paths. For each of the already-known relevant actions, the algorithm tries to find more relevant actions from the observed actions, that are connected to it by causal link paths. This continues until no more relevant action is found. The algorithm then organises the observed actions as well as temporal ordering constraints and causal links over these actions into a valid plan for the consistent goal.

**Proposition 2** *The *GoalGraphAnalyser* is sound: Any goal  $g$  it recognises at time step  $t$  is consistent with the observed actions so far, and the plan for  $g$  it organises is valid.*

In the example shown in Figure 2, the goal nodes in bold represent three consistent goals among which the goal node in italics represents a partially achieved goal while the other two represent two fully achieved goals. The edges in bold show causal link paths.

**Theorem 5 (Polynomial Space and Time)** *Consider a  $t$ -level Goal Graph. Let  $l_1$  be the number of fully or partially achieved goals at time step  $t$ ,  $m_1$  be the largest number of goal descriptions in any of these goals,  $l_2$  be the number of the observed actions, and  $m_2$  be the largest number of preconditions in any of these actions. The space size of possible causal link paths that connect the goals to the observed actions and that connect the observed actions to other observed actions, and the time needed to recognise all the consistent goals are polynomial in  $l_1, l_2, m_1$  and  $m_2$ .*

Persistence edges do not branch in a Goal Graph. The maximum number of paths that connect a goal to the observed actions is  $O(m_1)$ . The maximum number of paths that connect an observed action to other observed actions is  $O(m_2)$ . There are only at maximum  $l_1$  goals in goal-level  $t$  and  $l_2$  relevant actions to any of these goals. So the time needed to recognise all the consistent goals is polynomial in  $O(l_1(m_1 + l_2 m_2))$ .

### Experimental Results

Our algorithms have been implemented in Prolog and tested on a 500 MHz Pentium III in two domains in terms of efficiency, accuracy and scalability.

In the extended briefcase domain, we increase the number of locations to 50 and the number of objects up to 40




Figure 3: Experimental results of the extended briefcase domain

to create a series of spaces of 10,000, 20,000, up to 100,000 possible goals respectively. The same sequences of observed actions with the same Initial Conditions are used in the experiments in conjunction with these spaces of possible goals. Figure 3 shows that the average CPU time taken to process an observed action is approximately linear in the number of goals.

In the Unix domain, we tested our algorithms on a set of data collected at the University of Washington. To collect the data, the subjects are given goals described in English first and they then try to solve each goal by executing Unix commands. The executed Unix commands are recorded in the data set. We have 29 action schemata for the Unix commands including those executed by the subjects, 245 goal schemata and an estimate of 10000 possible goals. The results show that on average it only takes less than a CPU second to process an observed action and usually only a very small number of consistent goals remain after a sequence of observed actions have been processed.

Table 1 gives a summary of the experimental results. We tested our system on four goals that were originally tested in (Lesh & Etzioni 1995). The CPU second per update is the average time it takes to process an observed action. The length of observation is the average number of observed actions executed by the subjects to achieve the given goal. The fully achieved goals are the goals fully achieved after the last action has been observed. The partially achieved goals are the goals partially achieved. The remaining goals are the goals recognised after the last observed action has been processed. On  $G_2$ ,  $G_3$  and  $G_4$ , our algorithms return single, consistent goals that are the same as the goals given to the subjects. On  $G_1$ , four goals are recognised including the goal given to the subjects. Our algorithms recognise that the subjects tried to find one of the four files in the directory but does not know which file it is. This is as good as a human observer can do because you simply can not tell from the observed actions which file the subject was trying to find. These four goals can be generalised into a single, consistent goal ‘finding a file in the directory’ where variables are allowed in the recognised goals. These results show that our

goal	cpu sec per update	length of observation	fully achieved goals	partially achieved goals	remaining goals
$G_1$	1.202	2.25	7	12	4
$G_2$	0.352	16	10	9	1
$G_3$	0.027	3.0	3	10	1
$G_4$	1.255	20.5	33	0	1

Table 1: Experimental results of the Unix domain

algorithms perform extremely well with regard to efficiency and accuracy. They also demonstrate a significant improvement on the performance of the goal recogniser (Lesh & Etzioni 1995) where 155, 37 and 15 goals remain on  $G_1$ ,  $G_2$  and  $G_4$  respectively.

## Related Work

Most plan recognition systems (e.g., (Allen & Perrault 1980), (Carberry 1986), (Pollack 1986), (Litman & Allen 1987), (Kautz 1987)) search a space of possible plans for candidate plans that account for the observations. To form the search space in a given domain, some kind of plan representation is required. For instance, in Kautz’s event hierarchy, plan decompositions are required that describe how low level actions make up complex actions. Even though the use of the plan representation has an obvious advantage of expressive richness, it has a serious limitation in its inability to deal with new plans whose types do not appear in the plan representation. Hand-coding the plan representation in a large and complex domain presents a tedious or impractical task. In some other domains, the knowledge about plans might not be readily available.

Some attempts (e.g., (Lesh & Etzioni 1996), (Bauer 1998)) have recently been made to address this issue in which machine learning techniques have been applied to automate acquisition of plan representation. Even when leaving aside the plan representation consideration, searching the plan space can be exponentially expensive because the number of possible plan hypotheses can be exponential in the number of actions (Kautz 1987). Most plan recognition systems have often been developed in domains in which there are fewer than 100 plans and goals (Lesh & Etzioni 1996).

Our graph construction and analysis approach to plan recognition differs significantly from these plan recognition systems. Instead of immediately searching for candidate plans, our approach explicitly constructs a graph structure in which temporal and causal constraints among the observed actions and the achieved goals are explicitly represented. Our definition of what constitutes a valid plan for a goal eliminates the plan representation in most plan recognition systems. Our plan recognition system only takes the goal schemata and action schemata as input. It recognises consistent goals only from fully or partially achieved goals and organises the observed actions into valid plans for the recognised goals or part of these goals. Under our formulation, the plan recogniser must consider how the observed actions can be composed into plans. Our formation is not

limited in its ability to recognise new plans. Our algorithms are sound, polynomial-time and polynomial-space. Our experimental results show that our algorithms can be scaled up and applied to domains in which there are tens of thousands of possible goals and plans. We have therefore accommodated both expressiveness and tractability without the use of plan representation in our system.

Lesh and Etzioni first tried to use a graph representation of actions and goals for the goal recognition problem (Lesh & Etzioni 1995). Their graph representation only consists of action and goal nodes that are fully connected to each other first and then inconsistent goals are repeatedly pruned from this graph representation. This will lead to a set of candidate goals that explain the observed actions. Their graph representation does not explicitly represent temporal constraints and causal links over actions and goals. So their system can only recognise goals rather than plans because it cannot organise the observed actions into plan structures for the recognised goals. Their system is sound and polynomial-time. They have however sacrificed expressiveness of the plan representation for tractability. This is not the case in our system that recognises both goals as well as plans and performs in polynomial time and space. As we indicated in the previous section the number of remaining goals after pruning in their system is usually large.

Our graph-based approach to plan recognition can be seen as a counterpart of planning with Planning Graph in planning (Blum & Furst 1995), (Blum & Furst 1997). Though graph structures are used in both approaches, they consist of different kinds of nodes and edges, take different inputs and aim at producing different outputs.

## Conclusions

In this paper, we introduced a new approach to plan recognition in which a graph structure called a Goal Graph is constructed and analysed for plan recognition. We described two algorithms for constructing and analysing a Goal Graph. Our algorithms recognise both goals and plans. They allow redundant and partially ordered actions. They are sound, polynomial-time and polynomial-space. Our empirical experiments show that our algorithms are computationally efficient and they can be scaled up and applied to domains where there are tens of thousands of goals and plans. They recognise goals and plans with great accuracy. Since our new graph-based approach to plan recognition is fundamentally different from the existing methods for plan recognition, it provides an alternative to these methods and shows a new perspective of research into plan recognition.

Our plan recognition system is limited in its ability to recognise every type of erroneous plans. For instance, if an erroneous plan involves an observed action that is completely irrelevant to the intended goal, our system fails to recognise the goal as a consistent one. The GoalGraphAnalyser is not complete: it may not immediately recognise the intended goal as a consistent one when the action currently observed has a causal link with a relevant action that has not yet been observed. So it may temporarily miss the intended goal if it is not yet in the set of consistent goals. This is of course just a delay on updating the set of consistent goals because

as soon as the relevant action is observed, the intended goal will be recognised as a consistent one and the set of consistent goals will be updated accordingly. This is natural and inevitable to the human observer: when an observed action is not yet found relevant to a goal in a consistent way, we can either make an unsound guess that it could be relevant to some of the consistent goals we have at moment, or delay the decision for a little while until more actions are observed and this currently observed action is found relevant to a goal in a consistent way. Despite these limitations, our system performs extremely well in our two test domains.

## References

- Allen, J., and Perrault, C. 1980. Analysing intention in utterances. *Artificial Intelligence* 15:143–178.
- Anderson, C.; Smith, D.; and Weld, D. 1998. Conditional effects in graphplan. In *Proc. of 4th Intl. Conf. AI Planning Systems*.
- Bauer, M. 1998. Acquisition of abstract plan descriptions for plan recognition. In *Proceedings of AAAI-98*, 936–941.
- Blum, A., and Furst, M. 1995. Fast planning through planning graph analysis. In *Proc. of IJCAI-95*, 1636–1642.
- Blum, A., and Furst, M. 1997. Fast planning through planning graph analysis. *Artificial Intelligence* 90:281–300.
- Carberry, S. 1986. User models: the problem of disparity. In *Proceedings of the 11th International Conference on Computational Linguistics*, 29–34.
- Gazen, B., and Knoblock, C. 1997. Combining the expressivity of ucpop with the efficiency of graphplan. In *Proc. 4th Euro. Conf. on Planning*.
- Kautz, H. 1987. *A Formal Theory of Plan Recognition*. PhD Thesis, University of Rochester.
- Koehler, J.; Nebel, B.; Hoffmann, J.; and Dimopoulos, Y. 1997. Extending planing graphs to an adl subset. In *Proc. of 4th Euro. Conf. on Planning*, 273–285.
- Lesh, N., and Etzioni, O. 1995. A sound and fast goal recognizer. In *Proc. 15th Int. Joint Conf. on AI*, 1704–1710.
- Lesh, N., and Etzioni, O. 1996. Scaling up goal recognition. In *Proc. of Int. Conf. on Principles of Knowledge Representation and Reasoning*.
- Litman, D., and Allen, J. 1987. A plan recognition model for sub-dialogues in conversation. *Cognitive Science* 11:163–200.
- Pednault, E. 1988. Synthesizing plans that contain actions with context-dependent effects. *Computational Intelligence* 4(4):356–372.
- Pednault, E. 1989. Adl: Exploring the middle ground between strips and the situation calculus. In *Proc. of KR-89*, 324–332. Morgan Kaufman.
- Pollack, M. 1986. A model of plan inference that distinguishes between the beliefs of actors and observers. In *Proceedings of the 24th Annual Meeting of the Association for Computational linguistics*, 207–214.