

Probabilistic Plan Recognition for Hostile Agents

Christopher W. Geib and Robert P. Goldman
May 21st, 2001



Honeywell Laboratories

Chris Geib, 2001 FLAIRS

Background: definitions

- ï What is plan recognition: The process of recognizing the intentions or goals of an agent on the basis of observations of their actions.
 - ñ AKA intent recognition, task tracking

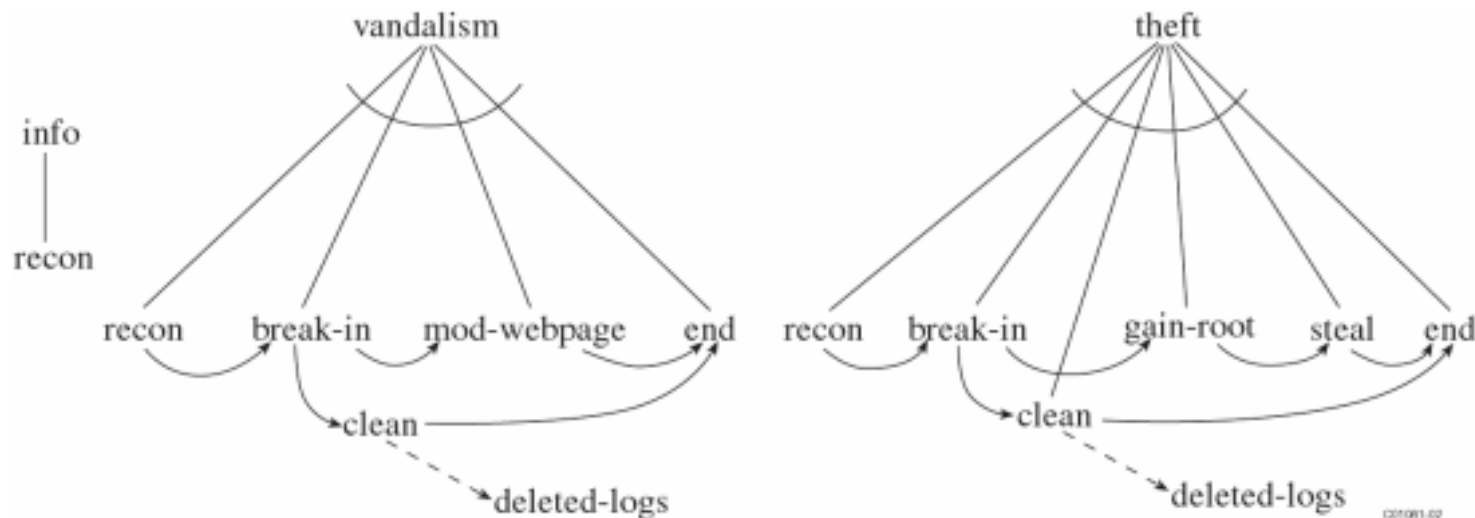
- ï Substantive work has been done in the area making different assumptions about the agent's attitude:
 - ñ Agnostic: The agent doesn't care or is completely unaware of the fact that you are observing their actions.(monitoring systems)
 - ñ Cooperative: The agent is aware and wants you to observe and understand the actions they are performing (collaborative or advisory systems)

 - ñ Hostile: The agent is hostile to the inference of their intent.
 - (not studied in the AI literature)



Background: plan library

- ï And/Or tree representation of the set of possible plans that could be used by the agent to be watched.
- ï Distinguished nodes of the trees (most often the roots) represent the agent's possible goals,
- ï Partial ordering constraints
- ï Effects of select actions



Background: previous approaches

- Set covering (Kautz and Allen 86)
- Probabilistic abduction (Charniak and Goldman 93)
- Grammar formalisms (Sidner 85, Vilain 90, Pynadath and Wellman 00)
- Reactive systems (Huber 94, Rao 94)
- Game theory, Mini-max search (Carmel and Markovitch 96)



Problems

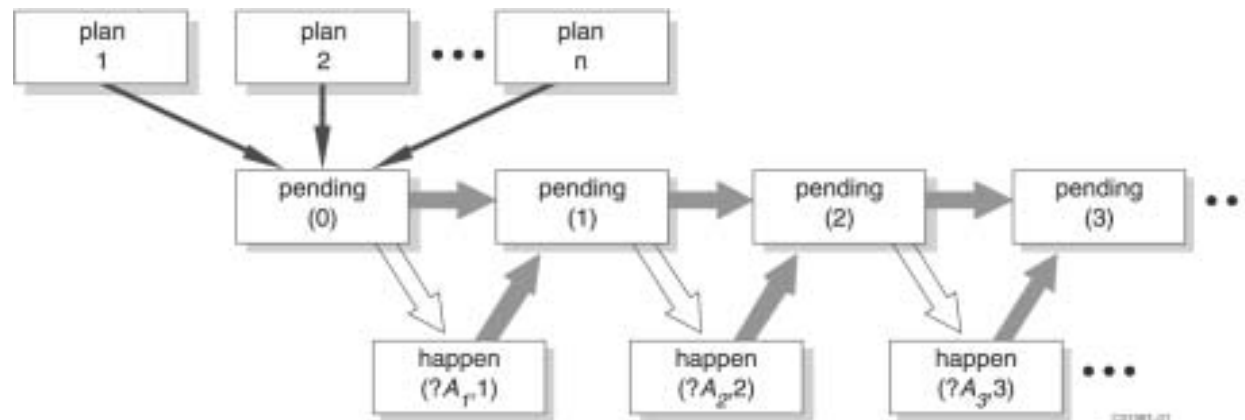
- ï Previous work suffers from a number of limitations that make them problematic for the domains we want to look at:
 - ñ partial order plans (peanut butter or jelly first?)
 - ñ effect of world state on goals (plant monitoring Vs. shutting down)
 - ñ overloading of actions (you only need to scan the network once)
 - ñ multiple goal Vs. single goal explanations
 - ñ cumulative effect of not seeing something (The dog didn't bark)

- ï Hostile agents are generally not covered since most work assumes agnostic or cooperative agents.
 - ñ Hidden actions/missing observations
 - ñ Observation of state changes rather than actions




Solution: Pending sets

- ï Central Insight: the agent will only execute actions that are consistent with their goals and are enabled by the previous actions they have performed.
- ï We define a **pending set** as the set of actions that are currently enabled by the agent's hypothesized goals and the actions the agent has taken.
- ï We use the sequence of observed actions to generate the pending set and as a result the hypothesized goals of the agent.



Solution: explanations

- ï Define: an ***Explanation*** for a set of observations as the substructure of the plan library with pending sets and hypothesized goals sufficient to explain the current set of observed actions.
- ï Example: observed **recon, break-in, mod-webpage**



Solution: algorithm intuition

- ï Build an exclusive and exhaustive set of the explanations for the observed actions.
- ï Establish the probability of each of the explanations ($P(\text{Exp})$.)
- ï The conditional probability of the goal given the observations is just the sum of those explanations that have the goal divided by the probability of the observations. ($P(\text{Goal} \mid \text{Observations})$)



Solution: what probabilities do we need?

- ï Not many that you have to know and they are easy to get
 - ñ Prior probability of a given intention or
 - ñ The conditional probability of the intention given a state of the world

- ï The other probabilities that we need we can compute:
 - ñ The probability of choosing a given action from the pending set.
 - ñ The probability of choosing a specific method
(when you have a choice of plans.)


 - ñ Our implementation assumes a uniform distribution for these
(but this is not required)




Solution: example

ï Observations: **recon, break-in, clean**

ï Explanation1: $cp(\text{theft} \mid \text{recon, break-in, clean}) = 0.5$



ï Explanation2: $cp(\text{vandalism} \mid \text{recon, break-in, clean}) = 0.5$



Solution: midpoint summary

- ï This approach will handle:
 - ñ partial order plans
 - ñ effect of world state on goals
 - ñ overloading of actions
 - ñ multiple goal Vs. single goal explanations
 - ñ cumulative effect of not seeing something

- ï But we can also extend it to handle hostile agents



Solution: hostile agent problems

- ï This algorithm won't work for hostile agents.
- ï Hostile agents like to hide their actions so that you don't know what they are doing. (Think of your dog or cat or small child (or even big child))
- ï The given algorithm relies on the fact that we have a complete record of the actions performed by the agent to compute the pending sets.
- ï This assumption doesn't hold anymore.
- ï In the parlance of Markov Decision Processes (MDP) we have moved from a fully observable case to the partially observable case (POMDP).



Solution: general approach

- ï If we just had a complete set of observations then we wouldn't have a problem.
- ï Intuition:
 - ñ Infer the unobserved actions from other information.
 - ñ Use the set of inferred unobserved actions to insert actions into the observations sequence to produce the complete set.
 - ñ Use the algorithm as before.
- ï How can we infer the execution of unobserved actions?
 - ñ Observations of unenabled actions
 - ñ Observations of unexplained state changes



Solution: intuition unenabled actions

- ï Definition: an ***unenabled action*** is an executed action that is not enabled by the agents previous actions. The observation of an unenabled action implies the execution of enabling actions that were performed unobserved.
- ï Example: all you observe is a **clean** action you can infer the execution of a **recon** and a **break-in**.
- ï The intuitive implementation of this would be to reason backwards once an unenabled action is found. This is NOT what we do. Instead we work forward inserting possible unobserved actions. We will return to discuss this later.



Solution: intuition state change

- ï Definition: an ***unexplained state change*** is an observed state change that does not have a causal explanation within the observed actions.
- ï Example: All you observe is **deleted-logs** you can infer the execution of a **clean** action.
- ï The addition of reports of state change to intent recognition is not common in the AI literature. This kind of information isn't important when you have a complete set of observations.
- ï Interestingly reports of state changes are often the most common in real domains. Hiding the effects of actions can be a great deal more difficult than hiding the actions. (think of nuclear weapons testing)




Solution: algorithm intuition

- Insert unobserved actions that are consistent with the
 - n unobserved actions
 - n unexplained state changes
 - n number of unobserved actions is less than a chosen threshold
- Produces a complete execution trace
- Use the algorithm as before to compute the pending sets and explanations on the basis of the new 'complete' set of observations.
- Example: observations: **recon, break-in, steal, s(deleted-logs)**

$$cp(\text{theft} \mid \text{recon, break-in, steal, s(deleted-logs)}) = 1.0$$

Intend(theft)




Implementation


- ï Based on Poole's Probabilistic Horn Clause Abduction
 - ñ Push all probability out to axioms of a proof
 - ñ Use first order logic to do your reasoning
 - ñ Requires an exclusive and exhaustive set of explanations
 - ñ Compute the probability of a single explanation by considering the 'axioms' of the explanation and their likelihood.
- ï Our current implementation
 - ñ in prolog
 - ñ not as efficient as it could be.



Implementation with no unobserveds



Implementation with unobserved actions



Implementation: analysis

- ï HEY WAIT A MINUTE! That looks exponential!

- ï Preliminary analysis: the good news -
 - ñ The length of the observations is NOT the problem

- ï Preliminary analysis: the bad news -
 - ñ The number of possible non-differentiated plans IS the problem

- ï Factors related to the number of non-differentiated plans
 - ñ Unobserved action threshold (too low and too high both bad)
 - ñ Branching factor of the plan library
 - ñ Longest non differentiating prefix



Conclusions

- A new, abductive, probabilistic, theory of plan recognition that is based on taking the sequential nature of observations of actions and state changes seriously. It is able to handle a wide variety of phenomena that other work finds problematic and we have applied it to the problem of inferring the intent of hostile agents.
- Implementation that is efficient.
- Preliminary analysis of the algorithm



End



Background: applications

- NIST Abnormal Situation Management 09: study project for assistant systems for oil refineries. Watch plan operator's actions and infer their goals. When possible suggest actions that might assist them in their task.
- DARPA CyberPanel project 00-01: Use reports from existing computer network intrusion detection systems as observations of
- Associate systems
- Educational or mentoring systems
- Intelligence analysis applications (business or governmental)



Solution: algorithm

ï Inputs: a sequence of action observations, and a plan library

ï Output: the conditional probability for each root intention

ï Code:

For each observation

Progress pending set

. remove executed action

. add newly enabled actions

Add to possible hypotheses set any new explanations indicated by the actions

Remove from possible hypothesis set any explanations inconsistent with the executed action.

Loop

For each explanation compute the explanations probability

For each root compute the conditional probability



Solution: more probabilities needed.

- ï To do this we will need more probabilities.
- ï Again its just priors of the kind you would expect to get from simple observations of the functioning system.
- ï Prior probability of the action being performed unobserved.
- ï Threshold of how many unobserved actions you are willing to accept.



Future work

- ï More complete analysis of the features affecting the runtime

- ï Industrial grade implementation.
 - ñ More efficient.
 - ñ Incremental algorithm
 - ñ precompilation of sequences and use of string matching algorithms.
 - ñ application of Tree-Adjoining Grammars and $O(n^6)$ parsing methods.

- ï Misdirection
 - ñ Execution of actions simply to mislead the observer.

- ï Multiple agents

- ï Observations of failed actions



Solution: bounding the explanations

- ï Problem: You can't allow the insertion of unobserved actions to go on forever. You would generate ever more unlikely explanations for the same set of observations.
- ï Realization: Some actions are more likely to be done unobserved than others. (a stealthy port scan Vs. a denial of service)
- ï Solution:
 - ñ Represent how likely it is that a given action is performed and not observed.
 - ñ For each explanation compute the probability that the unobserved actions could have been performed unobserved.
 - ñ Bound how unlikely an explanation we are willing to accept.
 - ñ Reject all explanations that are below the threshold.



Talk Outline

- ï Background
- ï Problem
- ï Solution
 - ñ General approach
 - ñ Extension to hostile agents
- ï Implementation
 - ñ Implementation details
 - ñ Happy graphs
 - ñ Analysis
- ï Conclusions

- ï Joint work with Robert Goldman

