Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence (AAAI-17)

Plan Recognition Design

Reuth Mirsky, Roni Stern, Ya'akov (Kobi) Gal, Meir Kalech
Department of Information Systems Engineering
Ben-Gurion University of the Negev, Israel
+972-50-378-1474, dekelr@bgu.ac.il

Abstract

Goal Recognition Design (GRD) is the problem of designing a domain in a way that will allow easy identification of agents' goals. This work extends the original GRD problem to the Plan Recognition Design (PRD) problem which is the task of designing a domain using plan libraries in order to facilitate fast identification of an agent's plan. While GRD can help to explain faster *which* goal the agent is trying to achieve, PRD can help in faster understanding of *how* the agent is going to achieve its goal. we define a new measure that quantifies the worst-case distinctiveness of a given planning domain, propose a method to reduce it in a given domain and show the reduction of this new measure in three domains from the literature.

Introduction

Goal Recognition is the problem of inferring an agent's goal from observations given a domain description (Ramırez and Geffner 2009). Keren et al. (2014) define Goal Recognition Design (GRD) as the problem of building a domain in a way that will minimize the number of observations needed to recognize an agent's goal. They introduce a worst-case distinctiveness (wcd) measure that is an upper bound on the number of observations needed to solve the GRD problem for a given domain, and showed how to compute this measure in domains based on the STRIPS representations.

An alternative representation for a domain is plan libraries (Kabanza et al. 2013; Blaylock and Allen 2006) which describe agents' activities as hierarchies of basic and complex actions. Plan libraries provide a rich representation but require the GRD task to explicitly reason about the hierarchical structure of the domain.

This work defines the Plan Recognition Design (PRD) problem, which is the task of designing a domain using plan libraries in order to facilitate fast identification of an agent's plan. In this case, identifying an agent's plan requires to infer the complete hierarchy of activities the agent is doing, not just its goal. This is called the Plan Recognition problem, which is more challenging than the equivalent Goal Recognition problem (Blaylock and Allen 2006). The paper formally defines PRD as the problem of minimizing the number of actions required to observe before unambiguously identifying the agent's plan. The paper defines a new measure, called worst-case plan distinctiveness, $wcpd$, which is analogous to the wcd measure in the GRD setting, and presents the longest sequence of observations that are required to recognize the agent's plan given a domain. We provide means of computing and minimizing the $wcpd$ measure and show in an empirical evaluation the values of this measure in 3 known domains from the literature.

Plan Recognition Design

We define a plan library in the standard way it is defined in the plan recognition literature (Geib and Goldman 2009), where a plan library is a tuple $L = \langle B, C, G, R \rangle$, where B is a finite set of basic actions, C is a finite set of complex actions, $G \subseteq C$ the possible goals and R is a set of rules of the form $c \rightarrow \tau \mid O$, where $c \in C$, τ is a string from $(B \cap C)^*$ and $O = \{(i, j) \mid c_i \prec c_j\}$ where c_i, c_j refer to the i-th and j-th actions in τ respectively.

Intuitively, B represents all of the atomically observable actions an agent can execute, C represents more complex or abstract actions with G as the goals the agent can try to achieve and each $r \in R$ represents how a complex action from C can decompose to a sequence of other actions. For $c_i, c_j \in \tau$, we say that $c_i \prec c_j$ if there exists an ordering constraint $(i < j) \in O$.

A plan is a labeled tree $p = (V, E, \mathcal{L})$, where V and E are the nodes and edges of the tree, respectively, representing the actions and their decomposition from the goal to the observable actions and $\mathcal{L}$ is a labeling function $\mathcal{L} : V \rightarrow B \cup C$ mapping every node in the tree to either a basic or a complex action in the plan library. The root note is called the goal of the plan and is labeled with a complex action from G. Each inner node is labeled with a complex action such that its children nodes are a decomposition of its complex action into constituent actions according to one of the rules. Given a set of plans G, we define $Plans(g_i)$ to be the set of all plans in G with the root goal g_i and $Goal(p_i)$ to be the root goal of some plan p_i in G.

An *observation sequence* is an ordered set of basic actions that represents actions carried out by the observed agent. A plan p *describes* an observation sequence O iff every observation is mapped to a leaf in the tree in an order that does not collide with the ordering constraints of the plan.

For each plan p we define $OBS(p)$ as the set of observation sequences such that each sequence in this set (1) is described by p, and (2) is not a prefix of a different observation sequence in $OBS(p)$.

Plan Recognition Design (PRD) is the problem of designing a domain in a way that will allow faster recognition of the plan of an acting agent. While in GRD the design tries to minimize the number of observations required until there is no ambiguity regarding the goal of the acting agent, PRD tries to minimize the number of observations required until there is no ambiguity regarding the complete plan of the acting agent. Keren et al. (2014) defined the wcd measure for any two goals g, g' as the longest sequence of observations required to observe before there is no ambiguity regarding which of the goals the acting agents is pursuing. We define a new metric, *worst case plan distinctiveness (wcpd)*, which is defined as the number of observations we will need to see, in the worst case, until we know for certain what the agent's plan is. Formally, for each two plans p_1, p_2, we define $wcpd(p_1, p_2) = \max_{O \in (OBS(p_1) \cap OBS(p_2))} |O|$.

We extend the original PRD definition to plan libraries, such that the $wcpd$ of a plan library P_D is the maximal $wcpd$ of every pair of plans that can be generated for goals in P_D. Formally, $wcpd(P_D) = \max_{p_1, p_2 \in \bigcup_{g \in G} Plans(g)} wcpd(p_1, p_2)$.

Definition 1 *Plan Recognition Design (PRD) problem is defined as a tuple $D = \langle P_D, L_D \rangle$, where P_D is a planning domain represented by a plan library $\langle B, C, G, R \rangle$ and L_D is a set of possible plans such that $L_D = \bigcup_{g \in G} Plans(g)$. The output of a PRD problem is $P'_D = \langle B, C, G, R' \rangle$ such that $R' \subseteq R$ and $\forall P''_D \subseteq P_D$ $wcpd(P'_D) \leq wcpd(P''_D)$.*

The search after a minimal set of rules to remove in order to minimize the $wcpd$ can be done by traversing over all combinations of rules from the original plan library to find which combination provides the smallest $wcpd$ without restricting the acting agent from achieving each goal.

The stopping conditions are: (1) There are no more rules that can be removed, without hindering the ability of the acting agent to reach all goals, were tested; (2) We have reached a plan library with a $wcpd$ of 0.

Removing rules from the plan library will restrict the acting agent, and considering different combinations of rules can be costly. Therefore we propose an anytime search that gradually constrains the plan library, so that we first examine all removals of a single rule, then removals of two rules, etc.

Empirical Evaluation

We evaluate our work in three domains from the plan recognition literature: Monroe is a disaster management domain (Blaylock and Allen 2005). Soccer is a domain from Avrahami-Zilberbrand and Kaminka (2005). Simulated is a domain generated to evaluate plan recognition algorithms (Kabanza et al. 2013).

Table 1 shows the initial $wcpd$ values of the different domains, along with the reduced $wcpd$ value, achieved under the following conditions: removing a single rule, and removing all possible combinations of rules which represent the upper bound on the reduction potential of the domain. As seen in the table, the initial $wcpd$ value varies greatly across the different domains. We attribute the high $wcpd$ of the Monroe domain to the high flexibility of the possible plans.

	Original $wcpd$	1-Rule reduction	Time	All rule reduction	Time
Monroe	10	10	7.31	–	–
Soccer	2	1	1.26	1	6.53
Simulated	0.82	0.15	91.81	0.07	2327

Table 1: $wcpd$ reduction.

For the Soccer domain, the $wcpd$ is reduced by 50% by removing a single rule from the plan library. This is the maximal possible reduction for this domain. For the Simulated domain, we were able to achieve 90% of the possible reduction potential. Interestingly, in the Monroe domain, the $wcpd$ could not be reduced by removing one rule, as well as removing all combinations of up to four rules (not shown in the table). Since there are 49 rules in this domain, considering the removal of all possible 2^{49} combinations was not feasible.

These results provide means to compare plan recognition domains, and demonstrate the anytime capabilities of the reduction algorithm.

We are currently extending our work in several directions. First, removing rules from plan libraries might put heavy restrictions on agent's actions. We are considering other, less restrictive, manipulations on the domain, such as adding ordering constraints on actions. We are also designing heuristics for ordering which rule combinations to consider for removal from the plan libraries.

References

Avrahami-Zilberbrand, D., and Kaminka, G. 2005. Fast and complete symbolic plan recognition. In *International Joint Conference of Artificial Intelligence*.

Blaylock, N., and Allen, J. 2005. Generating artificial corpora for plan recognition. In *International Conference on User Modeling*, 179–188. Springer.

Blaylock, N., and Allen, J. 2006. Fast hierarchical goal schema recognition. In *National Conference on Artificial Intelligence*, volume 21, 796.

Geib, C. W., and Goldman, R. P. 2009. A probabilistic plan recognition algorithm based on plan tree grammars. *Artificial Intelligence* 173(11):1101–1132.

Kabanza, F.; Filion, J.; Benaskeur, A. R.; and Irandoust, H. 2013. Controlling the hypothesis space in probabilistic plan recognition. In *International Joint Conference of Artificial Intelligence (IJCAI)*, 2306–2312.

Keren, S.; Gal, A.; and Karpas, E. 2014. Goal recognition design. In *International Conference on Automated Planning and Scheduling (ICAPS)*.

Ramırez, M., and Geffner, H. 2009. Plan recognition as planning. In *AAAI*.