Efficient Context Free Parsing of Multi-agent Activities for Team and Plan Recognition

(Extended Abstract)

Bikramjit Banerjee
School of Computing
The University of Southern Mississippi
Hattiesburg, MS 39406
Bikramjit.Banerjee@usm.edu

Jeremy Lyle
Dept. of Mathematics
The University of Southern Mississippi
Hattiesburg, MS 39406
Samuel.Lyle@usm.edu

Landon Kraemer
School of Computing
The University of Southern Mississippi
Hattiesburg, MS 39406
Landon.Kraemer@eagles.usm.edu

ABSTRACT

We extend a recent formalization of multi-agent plan recognition (MAPR), to accommodate compact multi-agent plan libraries in the form of context free grammars (CFG), incomplete plan executions, and uncertainty in the observation trace. Some existing approaches for single agent plan recognition cast it as a problem of parsing a single agent activity trace. With the help of our multi-agent CFG, we do the same for MAPR. However, known hardness results from multi-agent plan recognition constrain our options for efficient parsing, but we claim that *static* teams are a necessary (though not sufficient) condition for polynomial parsing. The necessity is supported by the fact that MAPR becomes NP-complete when teams can change dynamically. For sufficiency, we impose additional restrictions and claim that if the social structure among the agents is of certain types, then polynomial time parsing is possible.

Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence—*Multiagent Systems*

General Terms

Algorithms, Theory, Performance

Keywords

Plan recognition, Modeling other agents

1. INTRODUCTION

Multi-agent plan recognition (MAPR) refers to the problem of explaining the observed behavior of multiple agents by identifying the (dynamic) team-structures and the team plans (based on a given plan library) being executed, as well as predicting their future behavior. Applications of MAPR range from monitoring and surveillance, to automated sports commentary, to assistive technologies. Recently, Banerjee et. al. introduced a formal model for MAPR and used it to investigate the complexity of its simplest setting [1]. Zhuo and Li advanced this model to address missing observations in the activity traces as well as incompleteness of the plan library [5]. However, these models assume that the plan library is presented in an uncompact and non-hierarchical form, in particular as a set of team plans where each plan is a matrix of a fixed number of steps. Moreover, these models do not handle uncertainty in the observation trace in a general manner.

We have refined existing models from [1, 5] to make three important generalizations: allow compact, hierarchical, non-trivial plan libraries that correspond to infinite languages as opposed to the finite language in [1, 5], allow incomplete plan executions, and allow traces to be uncertain. Typically for plan recognition with single agents, a plan library is given in a compact hierarchical form, such as an HTN [2]. We have developed polynomial-time algorithms for a less expressive plan library, viz., context free grammars (CFGs) which incorporates some desirable features of HTN, e.g., recursiveness and hierarchies. This advances previous formalization in [1, 5] which accommodated none of these desirable features. Results from single agent plan recognition have shown that as long as partial ordering of plan steps is not allowed in the grammar, activity strings can be parsed in polynomial time [4]. However, even with the multi-agent CFG (i.e., with no partial ordering), MAPR would still be hard unless additional constraints are imposed. We identify specific types of social structures with static teams as such constraints.

In [5], missing steps in the trace as well as in the plan library were allowed. In constrast, assuming Σ to be the set of all possible observable activities, we model a missing observation as a uniform distribution over Σ to enunciate complete uncertainty. Unlike [5], this also allows for a varying degree of uncertainty on other observations that are not missed. A missing step in a plan is modeled as a don't care ($*$), similar to [5]. We include *noop* $\in \Sigma$ (i.e., "no operation") for cases when an agent is idling, which is often required for coordination among teammates. The CFG plan library is constructed in the same manner as [3, 4], except that each terminal activity represents a *vector* of activitites for the members of a *team* rather than a single agent.

2. ILLUSTRATIVE EXAMPLE

Figure 1 shows an illustrative example of the MAPR problem. The input is a (n =3)-agent trace, $\mathcal{T}$, that shows their

Appears in: *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012)*, Conitzer, Winikoff, Padgham, and van der Hoek (eds.), 4-8 June 2012, Valencia, Spain.

Input trace

Steps	Agent activities		
	Agent 1	Agent 2	Agent 3
1	guard	collect	threaten
2	guard	collect	threaten
3	guard	collect	threaten
4	ride	drive	ride
5	ride	drive	ride
6	ride	drive	ride

Plan library

Plan: P_1^2 (2-agent Money Pickup)
$P_1^2 \rightarrow \langle collect, guard\rangle Q_1^2$ $Q_1^2 \rightarrow \langle collect, guard\rangle Q_1^2 \vert \langle drive, ride\rangle Q_2^2$ $Q_2^2 \rightarrow \langle drive, ride\rangle Q_2^2 \vert \epsilon$

Plan: P_1^3 (3-agent Bank Robbery)
$P_1^3 \rightarrow \langle collect, guard, threaten\rangle Q_1^3$ $Q_1^3 \rightarrow \langle collect, guard, threaten\rangle Q_1^3 \vert \langle drive, ride, ride\rangle Q_2^3$ $Q_2^3 \rightarrow \langle drive, ride, ride\rangle Q_2^3 \vert \epsilon$

Figure 1: Illustrative example of MAPR with a CFG plan library.

recognized activities for ($T = 6$) steps, where n is the number of agents and T is the observation horizon. Suppose

$$\Sigma = \{collect, guard, threaten, drive, ride, noop\},$$

and each activity in $\mathcal{T}$ is associated with likelihood $1 - \delta$, with the probability that each could be some other activity in Σ being $\delta/5$, for some small $\delta > 0$. The input also contains the CFG shown in Figure 1(right) as the plan library.

Given the CFG plan library and the trace, the activities of agents 2 and 1 (in that order) could be parsed as fitting plan P_1^2 (2-agent money pickup in an armored car) with a high probability. However, activities of agents 2, 1, and 3 (in that order) might also be parsed as fitting plan P_1^3 (3-agent bank robbery) with a high probability. As in [1, 5], this ambiguity is resolved by noting that if the first hypothesis is accepted then it would be difficult to explain the activity of the remaining agent 3, and any explanation (provided there are other plans in the library that can explain agent 3's actions) could have a very low probability. In other words, the trace does not *partition* well. On the other hand, accepting the second hypothesis explains all agents' activities with a high probability. This also illustrates the power of partitioning the trace. Since money pickup is a more commonly observed plan, it has a high prior likelihood compared to bank robbery. Thus, if partitioning was not required and we were allowed to leave some activities unexplained, then bank robbery would be consistently missed.

3. MAIN CONTRIBUTIONS

Due to space limitation, we highlight the main contributions of our work.

- We have formally extended the definition of an *occurrence* [1, 5] to account for plans whose executions have not been completed by the observation horizon T. Such occurrences are called *partial* occurrences. As a consequence of the partiality of occurrences, we have also adapted the formal definitions of notions of conflict of (complete/partial) occurrences (i.e., hypotheses competing to explain some observations), and the definition of MAPR as a partition of the trace ($\mathcal{T}$) into complete or partial occurrences such that partial occurrences can only end at step T. In order for such a solution concept to apply, we have extended the closed world assumptions of [3] for multi-agent systems enunciating that every team-member of a team plan are observed.

- We have specified an algorithm (Parser) that extends Villain's Earley parser, for the determination of the maximum likelihood parse of a certain set of columns of the trace, $\mathcal{T}$. That is, given a team hypothesis, the parser yields the most likely sequence of high-level plans that the team might be executing. This algorithm has complexity $O(s^{2.5}G^2t^3)$, where s is the size of the hypothesized team, G is the size of the grammar, and t is the number of steps explained.

- We have revisited the notion of social structures that was used in the past to facilitate team hypotheses formation for MAPR. Instead of the hierarchical decomposition of teams into subteams, as done in the past, our notion of social structures captures the actual hierarchical organization of the agents. We allow teams to form only along paths in the social structure graph.

- Using Parser as a subroutine, we have formally proved that MAPR can be solved in polynomial time when the social structure graph is a star, a path, or a tree of bounded depth, and the teams are static.

- We have formally proved that when the team structures can vary with time, then MAPR is NP-complete even when the social stucture graph is as simple as a path. This proof is based on a reduction from the rectangular tiling problem (RTILE). Together with the previous result, this means that the staticity of teams is a necessary, but insufficient condition for the polynomial solvability of MAPR. Furthermore, social structure graphs such as star, path and trees, impose additional structure that turn out to be sufficient for polynomial solvability, in conjunction with static teams.

Acknowledgment: This work was supported in part by a start-up grant from the University of Southern Mississippi.

4. REFERENCES

[1] B. Banerjee, L. Kraemer, and J. Lyle. Multi-agent plan recognition: Formalization and algorithms. In *Proceedings of AAAI-10*, pages 1059–1064, Atlanta, GA, 2010.

[2] K. Erol, J. Hendler, and D. S. Nau. HTN planning: Complexity and expressivity. In *Proc. AAAI-94*, pages 1123–1128. AAAI Press, 1994.

[3] H. A. Kautz and J. F. Allen. Generalized plan recognition. In *Proc. AAAI*, 1986.

[4] M. Vilain. Getting serious about parsing plans: a grammatical analysis of plan recognition. In *Proc. of AAAI-90*, 1990.

[5] H. H. Zhuo and L. Li. Multi-agent plan recognition with partial team traces and plan libraries. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI-11)*, pages 484–489, 2011.