

Landmark-based Plan Recognition

Ramon Fraga Pereira and Felipe Meneguzzi¹

1 Introduction

As more computer systems require reasoning about what agents (both human and artificial) other than themselves are doing, the ability to accurately and efficiently recognize goals and plans from agent behavior becomes increasingly important. Plan recognition is the task of recognizing goals and plans based on often incomplete observations that include actions executed by agents and properties of agent behavior in an environment [10]. Accurate plan recognition is important to monitor and anticipate agent behavior, such as in crime detection and prevention, monitoring activities, and elderly-care. Most plan recognition approaches [3, 1] employ plan libraries (*i.e.*, a library with all plans for achieving a set of goals) to represent agent behavior, resulting in approaches to recognize plans that are analogous to language parsing. Recent work use planning domain definitions (domain theories) to represent potential agent behavior, bringing plan recognition closer to planning algorithms [8, 7, 6, 2].

In this paper, we develop a plan recognition approach that relies on planning landmarks [4] to filter candidate goals and plans from observations. Landmarks are properties (or actions) that every plan must satisfy (or execute) at some point in every plan execution to achieve a goal. In this way, we use this filtering algorithm in two settings. First, we build a landmark-based plan recognition heuristic that analyzes the amount of achieved landmarks to estimate the percentage of completion of each filtered candidate goal. Second, we show that the filter we develop can also be applied to other planning-based plan recognition approaches, such as the approach from Ramírez and Geffner [8]. We evaluate our approach empirically against the current state-of-the-art [8] using their own datasets [8, 7], and show that our approach has multiple advantages over existing approaches: it is more accurate than the state-of-the-art; it is substantially faster on its own; and it can also be used to speed up existing approaches.²

2 Filtering Candidate Goals from Landmarks in Observations

Key to our approach to plan recognition is the ability to filter candidate goals based on the evidence of fact landmarks and partitioned facts in preconditions and effects of observed actions in a plan execution. This filtering process analyzes fact landmarks inferred from observed actions, and selects goals from a set of candidate goals with the highest number of observed landmarks having been achieved. We take as input a plan recognition problem T_{PR} , which is composed of a planning domain definition Ξ , an initial state \mathcal{I} , a set of candidate goals \mathcal{G} , a set of observed actions O , and a filtering threshold θ . The

threshold gives us flexibility when dealing with incomplete observations and sub-optimal plans, which, when $\theta = 0$, may cause some potential goals to be filtered out before we get additional observations. Our algorithm iterates over the set of candidate goals \mathcal{G} , and, for each goal G in \mathcal{G} , it extracts and classifies fact landmarks and partitions for G from the initial state \mathcal{I} . We then check whether the observed actions O contain fact landmarks or partitioned facts of G in either their preconditions or effects. As we deal with partial observations in a plan execution some executed actions may be missing from the observation, thus whenever we identify a fact landmark, we also infer that its predecessors have been achieved. Given the number of achieved fact landmarks of G , we estimate the percentage of fact landmarks that the observed actions O have achieved according to the ratio between the amount of achieved fact landmarks and the total amount of landmarks. Finally, we return the goals from \mathcal{G} with the highest percentage of achieved landmarks within threshold θ .

3 Heuristic Plan Recognition using Landmarks

Our landmark-based heuristic method estimates the goal completion of every goal in the set of filtered goals. This estimate represents the percentage of sub-goals (atomic facts that are part of a conjunction of facts) in a goal that have been accomplished based on the evidence of achieved fact landmarks in observations. Our heuristic method estimates the percentage of completion towards a goal by using the set of achieved fact landmarks provided by the filtering process. We aggregate the percentage of completion of each sub-goal into an overall percentage of completion for all facts in a candidate goal. This heuristic, denoted as h_{prl} , is computed by Equation 1, where \mathcal{AL}_g is the number of achieved landmarks from observations of every sub-goal g of the candidate goal G , and \mathcal{L}_g represents the number of necessary landmarks to achieve every sub-goal g of G . Thus, heuristic $h_{prl}(G)$ estimates the completion of a goal G by calculating the ratio between the sum of the percentage of completion for every sub-goal $g \in G$, *i.e.*, $\sum_{g \in G} \frac{|\mathcal{AL}_g|}{|\mathcal{L}_g|}$, and the number of sub-goals in G .

$$h_{prl}(G) = \left(\frac{\sum_{g \in G} \frac{|\mathcal{AL}_g|}{|\mathcal{L}_g|}}{|G|} \right) \quad (1)$$

4 Landmark-based Plan Recognition

Our plan recognition approach is detailed in Algorithm 1, which takes as input a plan recognition problem T_{PR} , and works in two stages. First, this algorithm filters candidate goals using the filtering process, which returns the candidate goals with the highest percentage of achieved landmarks within a given threshold θ . Second, from the filtered candidates, this algorithm then uses h_{prl} to return the recognized goals by estimating the percentage of completion using the set of achieved fact landmarks provided by the filter.

¹ Pontifical Catholic University of Rio Grande do Sul (PUCRS), Brazil. Contact: ramon.pereira@acad.pucrs.br and felipe.meneguzzi@pucrs.br

² A complete discussion of our plan recognition approach and experiments is provided in the full paper <https://arxiv.org/pdf/1604.01277v2.pdf>

Domain	G	L	%Obs	O	LANDMARK-BASED PLAN RECOGNITION				R&G		FILTER + R&G	
					Time		Accuracy		Time	Accuracy	Time	Accuracy
					θ (0 / 10 / 20 / 30)	θ (0 / 10 / 20 / 30)	θ (0 / 10 / 20 / 30)	θ (0 / 10 / 20 / 30)				
BLOCKS-WORLD (855)	20	15.6	10	1.1	0.99 / 0.100 / 0.105 / 0.111	36.1% / 38.8% / 70.0% / 89.4%	1.656	83.8%	0.452	52.7%		
			30	2.9	0.107 / 0.109 / 0.118 / 0.122	54.4% / 61.1% / 86.1% / 97.2%	1.735	90.0%	0.458	77.7%		
			50	4.2	0.113 / 0.113 / 0.120 / 0.127	63.8% / 83.8% / 98.3% / 100.0%	1.836	97.2%	0.462	94.4%		
			70	6.5	0.138 / 0.139 / 0.141 / 0.148	81.6% / 94.4% / 100.0% / 100.0%	2.056	98.8%	0.483	96.1%		
			100	8.5	0.163 / 0.166 / 0.172 / 0.185	100.0% / 100.0% / 100.0% / 100.0%	2.378	100.0%	0.494	100.0%		
CAMPUS (75)	2	8.5	10	1	0.038 / 0.039 / 0.042 / 0.044	93.3% / 100.0% / 100.0% / 100.0%	0.083	100.0%	0.090	100.0%		
			30	2	0.048 / 0.050 / 0.055 / 0.057	100.0% / 100.0% / 100.0% / 100.0%	0.091	100.0%	0.089	100.0%		
			50	3	0.063 / 0.062 / 0.066 / 0.068	93.3% / 100.0% / 100.0% / 100.0%	0.105	100.0%	0.092	100.0%		
			70	4.4	0.060 / 0.060 / 0.063 / 0.065	100.0% / 100.0% / 100.0% / 100.0%	0.112	100.0%	0.095	100.0%		
			100	5.5	0.068 / 0.069 / 0.073 / 0.072	100.0% / 100.0% / 100.0% / 100.0%	0.126	100.0%	0.097	100.0%		
EASY-IPC-GRID (465)	7.5	11.3	10	1.8	0.585 / 0.588 / 0.609 / 0.623	82.2% / 85.5% / 97.7% / 100.0%	1.206	97.7%	0.770	97.7%		
			30	4.3	0.597 / 0.600 / 0.614 / 0.644	86.6% / 93.3% / 97.7% / 100.0%	1.291	98.8%	0.790	98.8%		
			50	6.9	0.608 / 0.609 / 0.627 / 0.656	94.4% / 97.7% / 97.7% / 100.0%	1.306	98.8%	0.860	100.0%		
			70	9.8	0.629 / 0.628 / 0.661 / 0.715	95.5% / 98.8% / 98.8% / 100.0%	1.715	100.0%	0.932	100.0%		
			100	13.3	0.630 / 0.632 / 0.685 / 0.759	100.0% / 100.0% / 100.0% / 100.0%	2.263	100.0%	1.091	100.0%		
INTRUSION-DETECTION (465)	15	16	10	1.9	0.197 / 0.200 / 0.211 / 0.233	76.4% / 96.6% / 100.0% / 100.0%	1.130	98.8%	0.506	98.8%		
			30	4.5	0.214 / 0.219 / 0.227 / 0.241	94.4% / 100.0% / 100.0% / 100.0%	1.142	100.0%	0.521	100.0%		
			50	6.7	0.218 / 0.221 / 0.246 / 0.269	100.0% / 100.0% / 100.0% / 100.0%	1.203	100.0%	0.531	100.0%		
			70	9.5	0.219 / 0.223 / 0.258 / 0.274	100.0% / 100.0% / 100.0% / 100.0%	1.482	100.0%	0.568	100.0%		
			100	13.1	0.277 / 0.281 / 0.303 / 0.325	100.0% / 100.0% / 100.0% / 100.0%	1.567	100.0%	0.566	100.0%		
KITCHEN (75)	3	5	10	1.3	0.003 / 0.003 / 0.002 / 0.004	93.3% / 100.0% / 100.0% / 100.0%	0.099	100.0%	0.093	100.0%		
			30	3.5	0.003 / 0.004 / 0.005 / 0.005	93.3% / 100.0% / 100.0% / 100.0%	0.111	100.0%	0.107	100.0%		
			50	4	0.004 / 0.004 / 0.006 / 0.006	93.3% / 100.0% / 100.0% / 100.0%	0.112	100.0%	0.111	100.0%		
			70	5	0.006 / 0.007 / 0.007 / 0.008	93.3% / 93.3% / 100.0% / 100.0%	0.111	100.0%	0.110	100.0%		
			100	7.4	0.007 / 0.008 / 0.008 / 0.009	100.0% / 100.0% / 100.0% / 100.0%	0.118	100.0%	0.112	100.0%		
LOGISTICS (465)	10	18.7	10	2	0.441 / 0.449 / 0.455 / 0.458	73.3% / 96.6% / 100.0% / 100.0%	1.125	100.0%	0.615	98.8%		
			30	5.9	0.447 / 0.452 / 0.461 / 0.466	88.7% / 100.0% / 100.0% / 100.0%	1.195	100.0%	0.663	100.0%		
			50	9.5	0.457 / 0.469 / 0.474 / 0.488	96.6% / 100.0% / 100.0% / 100.0%	1.248	98.8%	0.712	98.8%		
			70	13.4	0.474 / 0.481 / 0.490 / 0.497	100.0% / 100.0% / 100.0% / 100.0%	1.507	100.0%	0.786	100.0%		
			100	18.7	0.498 / 0.505 / 0.513 / 0.522	100.0% / 100.0% / 100.0% / 100.0%	1.984	100.0%	0.918	100.0%		

Table 1: Comparison and experimental results of our landmark-based approach against Ramirez and Geffner [8] approach. R&G denotes their plan recognition approach and Filter + R&G denotes the same approach but using our filtering algorithm.

Algorithm 1 Recognize goals and plans using the filtering method and the landmark-based heuristic.

Input: $\Xi = \langle \Sigma, \mathcal{A} \rangle$ planning domain, \mathcal{I} initial state, \mathcal{G} set of candidate goals, \mathcal{O} observations, and θ threshold.

Output: Recognized goal(s).

- 1: **function** RECOGNIZE($\Xi, \mathcal{I}, \mathcal{G}, \mathcal{O}, \theta$)
- 2: $\Lambda_{\mathcal{G}} := \langle \rangle$ \triangleright Map goals to % of landmarks achieved.
- 3: $\Lambda_{\mathcal{G}} := \text{FILTERCANDIDATEGOALS}(\Xi, \mathcal{I}, \mathcal{G}, \mathcal{O}, \theta)$
- 4: **return** $\arg \max_{G \in \Lambda_{\mathcal{G}}} h_{prl}(G)$

Table 1 shows the result of our experiments, which uses six domains from datasets provided by Ramirez and Geffner [8, 7], comprising hundreds of plan recognition problems T_{PR} , i.e., a domain description as well as an initial state, a set of candidate goals \mathcal{G} , a hidden goal G in \mathcal{G} , and an observation sequence \mathcal{O} (10%, 30%, 50%, 70%, or 100% of observability). We use two metrics, the accuracy of the recognition and the speed to recognize a goal, and compare our approach to two other approaches: the approach of Ramirez and Geffner [8] on its own, and this approach combined with our filter. More specifically, we use their faster and most accurate approach; as well as a combination of their approach and our filter. For our approach, we show the accuracy under different filtering thresholds (0%, 10%, 20% and 30%). If threshold $\theta = 0$, our approach gives no flexibility for filtering candidate goals, returning only the goals with the highest percentage of achieved landmarks. Each row of this table shows the observability (% Obs) and averages of the number of candidate goals $|\mathcal{G}|$, the number of observed actions $|\mathcal{O}|$, recognition time (seconds), and accuracy. We can see from the table that our approach is both faster and more accurate than Ramirez and Geffner [8], and, when we combine their algorithm with our filter, the resulting approach gets a substantial speedup.

5 Conclusion

We have developed an approach for plan recognition that relies on planning landmarks and a new heuristic based on these landmarks. Landmarks provide key information about what cannot be avoided to achieve a goal, and we show that landmarks can be used efficiently

for very accurate plan recognition. We have shown empirically that our approach yields not only superior accuracy results but also substantially faster recognition times for all domains used in evaluating against the state of the art [8] at varying observation levels.

There are multiple avenues for future work, such as: evaluating heuristics and symmetries in classical planning [9]; other landmark extraction techniques [5]; adding a probability interpretation to the observed landmarks and comparing to a recent work [2]; and account for information gain over multiple competing plan hypotheses.

Acknowledgments: This research was carried out in cooperation with HP Brazil using incentives of the Brazilian Informatics Law (# 8.2.48 of 1991).

REFERENCES

- [1] Dorit Avrahami-Zilberbrand and Gal A. Kaminka, ‘Fast and Complete Symbolic Plan Recognition’, in *IJCAI 2005, Edinburgh, Scotland*, pp. 653–658, (2005).
- [2] Yolanda E.-Martín, María D. R.-Moreno, and David E. Smith, ‘A Fast Goal Recognition Technique Based on Interaction Estimates’, in *IJCAI 2015, Buenos Aires, Argentina, July 25-31, 2015*, pp. 761–768, (2015).
- [3] Christopher W. Geib and Robert P. Goldman, ‘Partial Observability and Probabilistic Plan/Goal Recognition’, in *Workshop on Modeling Others from Observations (MOO-2005)*, (2005).
- [4] Jörg Hoffmann, Julie Porteous, and Laura Sebastia, ‘Ordered Landmarks in Planning’, *Journal of Artificial Intelligence Research (JAIR)*, **22**(1), 215–278, (November 2004).
- [5] Erez Karpas, David Wang, Brian C. Williams, and Patrik Haslum, ‘Temporal landmarks: What must happen, and when’, in *ICAPS 2015, Jerusalem, Israel, June 7-11, 2015*, pp. 138–146, (2015).
- [6] David Pattison and Derek Long, ‘Domain Independent Goal Recognition’, in *STAIRS*, pp. 238–250, (2010).
- [7] Miquel Ramirez and Hector Geffner, ‘Probabilistic Plan Recognition Using Off-the-Shelf Classical Planners’, in *AAAI 2010, Atlanta, Georgia, USA, July 11-15, 2010*, (2010).
- [8] Miquel Ramirez and Hector Geffner, ‘Plan Recognition as Planning’, in *IJCAI 2009*, pp. 1778–1783, (2009).
- [9] Alexander Shleyfman, Michael Katz, Malte Helmert, Silvan Sievers, and Martin Wehrle, ‘Heuristics and Symmetries in Classical Planning’, in *AAAI 2015, Austin, Texas, USA*, pp. 3371–3377, (2015).
- [10] Gita Sukthankar, Robert P. Goldman, Christopher Geib, David V. Pynadath, and Hung Hai Bui, *Plan, Activity, and Intent Recognition: Theory and Practice*, Elsevier, 2014.