Towards Online Goal Recognition Combining Goal Mirroring and Landmarks

Extended Abstract

Mor Vered
Bar-Ilan University
Ramat-Gan, Israel
veredm@cs.biu.ac.il

Ramon Fraga Pereira
PUCRS
Porto Alegre, Brazil
ramon.pereira@acad.pucrs.br

Maurício C. Magnaguagno
PUCRS
Porto Alegre, Brazil
mauricio@acad.pucrs.br

Gal A. Kaminka
Bar-Ilan University
Ramat-Gan, Israel
galk@cs.biu.ac.il

Felipe Meneguzzi
PUCRS
Porto Alegre, Brazil
felipe.meneguzzi@pucrs.br

KEYWORDS

Goal Recognition, Plan Recognition, Goal Mirroring, Landmarks

ACM Reference Format:

Mor Vered, Ramon Fraga Pereira, Maurício C. Magnaguagno, Gal A. Kaminka, and Felipe Meneguzzi. 2018. Towards Online Goal Recognition Combining Goal Mirroring and Landmarks. In *Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), Stockholm, Sweden, July 10–15, 2018,* IFAAMAS, 3 pages.

1 INTRODUCTION

Goal recognition is the problem of recognizing the goal of an agent based on a sequence of observations. The problem is at the core of many real-world applications such as human-robot interaction, intelligent user interfaces and recognizing navigation goals. The problem may be further categorized into two subsets; in *offline* goal recognition the set of observations, while in itself may be noisy and/or incomplete, is revealed ahead of time. Conversely, in *online* goal recognition the set of observations is revealed incrementally and an hypothesis must be made after each additional observation with no knowledge about which observation may be the final one.

Most approaches to goal recognition rely on a plan library describing the plans assumed known by the agent being observed to achieve its goals [13]. These approaches require substantial domain knowledge, and make strong assumptions about the preferences of observed agents. Conversely, *plan recognition as planning* (PRP) [8, 9], uses a planner in the recognition process to generate recognition hypotheses as needed, eliminating the need for a plan library. These approaches have shown that it is possible to carry out effective goal recognition using only a domain-theory describing actions in the environment as domain knowledge, with recent work showing the problem can be solved very efficiently [4, 6].

However, all of these approaches only apply to *offline* goal recognition. Indeed, Vered and Kaminka [3, 14, 15] have shown that the previously mentioned PRP approaches are not applicable to *online* recognition and must therefore be adapted to include multiple executions of a planning algorithm in order to compute alternative ways in which the observed agent can achieve a goal. Therefore, a straightforward implementation of these PRP methods adapted to *online* recognition may prove very computationally expensive.

We develop a procedure for extracting continuous-space landmarks and introduce an online goal recognition approach that combines online *Goal Mirroring* [14] and *recognition using landmarks* [6]. The results show superior efficiency and generally superior recognition performance over the state-of-the-art.

Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), M. Dastani, G. Sukthankar, E. André, S. Koenig (eds.), July 10–15, 2018, Stockholm, Sweden.

2 ONLINE GOAL RECOGNITION USING LANDMARKS

In the planning literature, *landmarks* are partially ordered facts that must be true at some point in every valid plan to achieve a particular goal from an initial state [2]. Given their usefulness for planners and planning heuristics [10], research has yielded multiple notions of landmarks [7], including that of disjunctive landmarks. Briefly, disjunctive landmarks represent an exclusive disjunction over possible instances of variables associated to predicates in the state representation.

Pereira et al. [4, 6] show that it is possible to carry out offline plan recognition by reasoning heuristically about landmarks. The key idea is to maintain a list of ordered landmarks associated with each goal, though partial overlaps are allowed. The *goal completion* heuristic from Pereira et al. [6] matches the observations against this list. This heuristic marks a landmark as *achieved* when facts in the observation match a landmark. The heuristic then uses the ratio of the number of landmarks achieved to the total number of landmarks associated with the goal, inducing a ranking of the goals.

In principle, we can translate the same idea into recognition in continuous domains. In such domains, landmarks can be defined as areas surrounding goals and to achieve a goal would mean that the observed motion must intersect (go through) the corresponding landmark area. Naturally, we would prefer such areas to be maximal, but must maintain the restriction that landmarks cover only obstacle-free space, and do not intersect completely with other landmarks.

3 EXTRACTING LANDMARKS IN CONTINUOUS SPACE

We can use any one of a number of landmark extraction algorithms to extract landmarks in discrete environments. We have chosen to adapt the algorithm of Hoffman et al. in [2] since it efficiently approximates landmark sets that are good enough for the domains we use. This algorithm builds a graph in which nodes represent landmarks and edges represent necessary prerequisites between landmarks, thus representing the landmarks and their ordering. A node in this graph represents a conjunction of facts that must be true simultaneously at some point during the execution of a plan, and the root node is a landmark representing the goal state. Hoffman et al. [2] proves that the process of generating all landmarks and deciding their ordering is PSPACE-complete, which is exactly the same complexity as deciding plan existence [1].

Since the interpretation of landmarks we rely on for plan recognition is that of bottlenecks in the state space, we try to partition a continuous space so that such bottlenecks become identifiable areas in the continuous space. Specifically, to extract landmarks in continuous environments we partition the area using the wall corners as references, to eventually identify pathways between individual "rooms" in the space. Though we define a landmark generation algorithm for continuous path planning domains, our approach should work with any notion of numeric landmarks, e.g., recent work on landmarks for hybrid domains [11].

4 ONLINE GOAL MIRRORING

Online goal mirroring [14] uses the following procedure: For each goal the procedure compares the costs of an *ideal plan* and an *observation-matching plan.* Ramirez and Geffner [8, Theorem 7] show that necessarily, a goal for which the two plans have equal costs is a solution to the goal recognition problem. We use this to rank the goals. The closer two costs for both plans are are, the higher the likelihood of the goal.

The ideal plan is an optimal plan[1], computed once from the initial state to each of the goals. The observation-matching plan is constructed for each new observation such that it always visits the states included in all the observations thus far, and then optimally reaches the goal. For online operation, instead of calling the planner to re-generate the observation-matching-plan with each new observation, we construct it for each new observation by concatenating two parts: a plan prefix which is a concatenation of all observations received to date. This is very efficiently done by simply adding the latest observation to the current prefix; and a plan suffix which is a new plan, issued by a motion planner, from the last state of the prefix (after incorporating the observations), to the goal state. The bulk of the computation takes place here by calling the planner.

5 GOAL MIRRORING WITH LANDMARKS

In general, PRP recognizers repeatedly call a planner during recognition, and this is exacerbated in online recognition, as the goal recognizer previously described calls the planner to compute a new plan with every observation, and for every goal. By combining Goal Mirroring and the evidence provided by landmarks, we exploit both the flexibility of an online recognition approach that utilizes a planner within the recognition process and the efficiency of reasoning about landmarks.

We assume a single cached computation of domain specific landmarks for all monitored goals and use the information conveyed by the landmarks as a pruning mechanism with which we may rule out hypotheses, reducing the set of goals and therefore the number of calls to the planner and overall run-time.

For every newly available observation we ascertain whether this observation has caused any landmarks to be satisfied. If the observation has caused a landmark to be satisfied we may use the existing fact landmarks to prune unlikely goals, in which case we only call the planner to compute plans for those goals whose landmarks have not yet been satisfied.

6 EXPERIMENTS AND EVALUATION

We empirically evaluated our *online* goal recognition approach on both discrete and continuous environments, over hundreds of goal recognition problems while measuring both efficiency and performance. For our continuous environment we used the domain of 3D navigation, where the target is to recognize navigational goals as soon as possible while the observations, i.e., observed agents' positions, are incrementally revealed [12]. For our discrete environments, we used the openly available datasets [5] based on the ones developed by Ramírez and Geffner [8, 9].

We contrasted the performance and efficiency of our combined approach (Goal Mirroring with Landmarks) with the existing PRP approach (Goal Mirroring) and our newly presented online recognition approach using only the landmarks for ranking and pruning out goals (Online Recognition with Landmarks).

For both continuous and discrete domains, the combined Goal Mirroring with Landmarks approach achieved the best performance and proved just as reliable as Online Recognition with Landmarks. However, it was not as reliable as Goal Mirroring, which does not prune out goals at all, incurring no risk of overlooking the correct goal. There were several instances where the dataset was so complex that both the Goal Mirroring and Goal Mirroring with Landmarks approaches failed. Due to the repeated calls to the planner these approaches timed-out without results. These problems were considerably more complex with a larger number of objects and instantiated actions.

7 CONCLUSIONS AND FUTURE WORK

We have developed an efficient online goal recognition approach which works in both continuous and discrete domains using a combination of Goal Mirroring and reasoning over a generalized notion of landmarks. We have shown that not only is our approach more efficient than the existing online recognizer but also outperforms both other approaches.

However, as our technique continually calls a planner within the recognition process, it might have limitations in recognizing very complex problems, specifically, those for which current planning algorithms are not efficient. Among some of the other limitations is its use of relatively simple landmarks for spatial domains, as well as the assumption that landmarks do not change over the course

[1] We assume knowledge of plan cost, such as the length of the plan trajectory in continuous spaces, or the number of plan steps in discrete environments, etc.

of the recognition, which would not be realistic for dynamically changing environments.

REFERENCES

[1] Tom Bylander. 1994. The Computational Complexity of Propositional STRIPS Planning. *JAIR* 69 (1994), 165–204.

[2] Jörg Hoffmann, Julie Porteous, and Laura Sebastia. 2004. Ordered Landmarks in Planning. *JAIR* 22, 1 (2004), 215–278.

[3] Gal A Kaminka, Mor Vered, and Noa Agmon. 2018. Plan Recognition in Continuous Domains. *In Proceedings of the 32nd Conference on Artificial Intelligence (AAAI)* (2018).

[4] Ramon Fraga Pereira and Felipe Meneguzzi. 2016. Landmark-Based Plan Recognition. In *ECAI*.

[5] Ramon Fraga Pereira and Felipe Meneguzzi. 2017. Goal and Plan Recognition Datasets using Classical Planning Domains. Zenodo. (July 2017). https://doi.org/10.5281/zenodo.825878

[6] Ramon Fraga Pereira, Nir Oren, and Felipe Meneguzzi. 2017. Landmark-Based Heuristics for Goal Recognition. In *AAAI*. AAAI Press.

[7] J. Porteous and S. Cresswell. 2002. Extending Landmarks Analysis to Reason about Resources and Repetition. In *Proceedings of the 21st Workshop of the UK Planning and Scheduling Special Interest Group (PLANSIG '02)*.

[8] Miquel Ramírez and Hector Geffner. 2009. Plan recognition as planning. In *IJCAI*. 1778–1783.

[9] Miquel Ramírez and Hector Geffner. 2010. Probabilistic plan recognition using off-the-shelf classical planners. In *AAAI*.

[10] Silvia Richter and Matthias Westphal. 2010. The LAMA Planner: Guiding Cost-based Anytime Planning with Landmarks. *JAIR* 39, 1 (2010), 127–177.

[11] Enrico Scala, Patrik Haslum, Daniele Magazzeni, and Sylvie Thiébaux. 2017. Landmarks for Numeric Planning Problems.. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence*. 4384–4390.

[12] Ioan A Sucan, Mark Moll, and Lydia E Kavraki. 2012. The open motion planning library. *IEEE Robotics & Automation Magazine* 19, 4 (2012), 72–82.

[13] Gita Sukthankar, Robert P. Goldman, Christopher Geib, David V. Pynadath, and Hung Bui (Eds.). 2014. *Plan, Activity, and Intent Recognition.* Morgan Kaufmann.

[14] Mor Vered and Gal A Kaminka. 2017. Heuristic Online Goal Recognition in Continuous Domains. *In Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)* (2017).

[15] Mor Vered, Gal A Kaminka, and Sivan Biham. 2016. Online Goal Recognition through Mirroring: Humans and Agents. *The Fourth Annual Conference on Advances in Cognitive Systems (ACS)* (2016).