Online Recognition of Navigation Goals Through Goal Mirroring (Extended Abstract)

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

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

1. INTRODUCTION

Goal recognition is the problem of inferring the (unobserved) goal of an agent, based on a sequence of its observed actions [4, 2, 1, 5]. It is a fundamental research problem in artificial intelligence, closely related to plan, activity, and intent recognition [10].

In *offline recognition* the entire sequence of observations is provided to the agent ahead of time. In contrast, in *online recognition* the sequence of observations is revealed incrementally instead of being known in advance, thus exacerbating an already hard problem.

The prevalent approach to goal recognition, both offline and online, relies on a dedicated *plan library*, a set of plans which represents all known ways to achieve known goals [10]. These recognition methods vary in the expressiveness of the representation and efficiency of the inference algorithms used. While powerful when the plans are known, these methods fail when the observations come from an unknown plan to achieve a known goal. An additional difficulty is raised when adding goals to the set of recognizable goals, as plans for them need to be inserted in the library, in order to be recognized.

One of the notable exceptions is plan recognition by planning (PRP), which focuses on library-free recognition, where a planner is used as a black box, to dynamically generate plans that are matched against the observations, eliminating the need for a plan library [6, 8]. This approach targets discrete domains only, and is inefficient for online recognition where it would produce $2|O||G|$ calls to the planner, where $|O|$ is the number of observations, and $|G|$ the number of goals.

We advocate *goal mirroring*. Like [6], goal mirroring uses a planner, to generate recognition hypotheses. However, it is designed for efficient, online recognition in continuous environments by using a motion planner, with a baseline number of calls to the planner of $(|O| + 1)|G|$.

We additionally identify two key decision points where, by inserting heuristics for navigational goal recognition, we can further influence the number of calls to the planner and the overall run-time.

2. GOAL MIRRORING

Goal Mirroring is an online goal recognition approach which utilizes a planner within the recognition process for every consecutive observation ([11]) inspired by mirroring processes hypothesized to take place within the human brain [7]. After each observation the recognizer utilizes a planner to generate possible plans to achieve each of the possible goals. Because the planners used are off-the-shelf planners and incorporating information from the observations as input to the planners is not a trivial task.

In general PRP recognizers avoid representing the plans explicitly, as a library of plans to be used for recognition. Instead, the set of plans is only implicitly represented, by using a planner to generate solution candidates on the fly. The planner is used at least twice, *for each goal*:

1. It is called to generate the ideal plan, which is a sequence of states (a path in continuous spaces) from the initial state to the goal, ignoring any knowledge of observations.

2. It is called a second time to generate a second plan that incorporates the observations seen so far.

In goal mirroring, this is done as follows. First, we observe that in both principle and practice, there are infinite plans that do not match the observations. Thus instead of generating *all* possible plans (a potentially infinite number) which will often be disqualified, mirroring algorithms use the planner to generate candidate plans that *always* match the observations, by folding the observations into the generated plan. In discrete domains, Ramirez and Geffner offer an elegant way for doing this by modifying the domain input of PDDL-type planners. In continuous domains, we propose folding the observations into the candidate solution by breaking the plan into two pieces:

- A plan prefix, is built by concatenating all seen observations into a single (possibly discontinuous) trajectory.

- A plan suffix which is generated by calling the planner, to generate a trajectory from the last state (point) in the prefix (the ending point of the last observation) to each goal.

A complete plan is then the sum of both these plans; i.e. a trajectory from the first observed point to each of the

goals. These complete plans necessarily *perfectly* match the observations, since they incorporate them.

The resulting plans then need to be ranked. In this we rely on we drew inspiration from studies of human estimates of intentionality and intended action [3]. Such studies have shown a strong bias on part of humans to prefer hypotheses that interpret motions as continuing in straight lines, i.e., without deviations from or corrections to, the heading of movements. Therefore our ranking is biased towards rational agents. We compare the resulting plans combined with the already seen observations to the ideal plan, calculated from the initial position to each of the goals. The closer the plans, the higher the corresponding goal is ranked. In this way our approach is able to work for continuous domains (navigational goals, shape recognition) as well as discrete.

2.1 Heuristic Recognition of Navigation Goals

Calling the planner for each new observation can be very expensive and inefficient in terms of run-time performance. We therefore introduce two heuristics applicable to the navigation goal recognition domain to be inserted in the key decision points in the process. These are again inspired by studies of human biases towards rationality.

The first heuristic is the *recomputation* heuristic. The purpose of this heuristic is not to call the planner when unnecessary, again reducing the overall run-time of the recognition process. In *Goal Mirroring*, for every new observation we are called to run the planner again, from the current state to each of the goal states. This could result in an additional $|G|$ calls to the planner for every added observation. However, for every new observation we have necessarily calculated the plans for the previous observation. By saving these previously calculated plans we may now consider whether the new observation is in agreement with previously calculated plans. If the observation matches it means we may continue to rely on former calculated plans and need not re-call the planner. Simply, it means that, if the agent is continuing to head in the same general direction, we may choose to keep the former goal rankings and not call the planner for recomputation at all assuming the rational agent is still advancing towards the same goal.

While calling the planner is wasteful when unnecessary, it is also wasteful to call the planner for goals that are highly improbable—or even impossible—given the observations. This leads us to the second heuristic, the *pruning heuristic*. The idea is to prune goals from being considered at all, reducing $|G|$ as observations come in. Here we again rely on the rationality of the observed agent, assuming that the observations and plans generated approximate the shortest possible paths between two coordinates. Once a rational agent is moving away or past a goal point, that goal is considered an unlikely target and may be pruned. In this manner we may decrease the overall number of calls to the planner and overall run-time of the recognition process.

3. EVALUATION

We empirically evaluated the performance of online goal mirroring along with the different heuristics over hundreds of goal recognition problems while measuring both the efficiency of the approach and the overall performance.

We implemented the different approaches of online goal mirroring to recognize the goals of navigation in 3D worlds. We used TRRT (Transition-based Rapidly-exploring Ran-

dom Trees), an off-the-shelf planner that guarantees asymptotic near-optimality by preferring shorter solutions, available as part of the Open Motion Planning Library (OMPL [9]) along with the OMPL *cubicles* environment and default robot (Figure 1). Each call to the planner was given a time limit of 1 sec. We used the OMPL *interpolate* method to generate between between 20 and 76 observed points for each path problem.

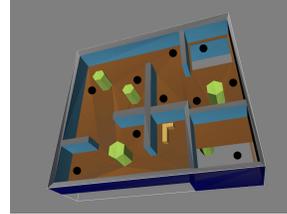**Figure 1: Visualization of 3D navigation goals environment**

We generated two observed paths from each point to all others, for *a total of* 110×2 *goal recognition problems*. The observations were obtained by running the RRT* planner on each pair of points, with a time limit of 5 minutes per run. RRT* was chosen because it is an optimized planner that guarantees asymptotic near-optimality. The longer the run-time the more optimal the path.

We saw that employing the heuristics makes a big impact on run-time and successfully reduced overall number of calls to the planner. While the *recomputation* heuristic outperformed the *pruning* heuristic, both in run-time and overall number of calls, utilizing both heuristics can reduce both run-time and number of calls made to the planner by over 80% from the naive approach.

4. FUTURE WORK

We next intend to examine the relation between the PRP recognizer and the planner used for recognition. We contend that recognition success relies heavily on a thorough knowledge of the observed agents' decision making process. In humans this is related to the rationality assumption. We intend to work with *Intelligent Tutoring Systems* where we will extend goal mirroring to recognize different strategies taken by students as they solve educational problems.

Acknowledgments

REFERENCES

[1] C. L. Baker, J. B. Tenenbaum, and R. R. Saxe. Goal inference as inverse planning. In *Proceedings of the 29th Annual Meeting of the Cognitive Science Society*, 2007.

[2] N. Blaylock and J. Allen. Fast hierarchical goal schema recognition. In *Proceedings of the Twenty-First National Conference on Artificial Intelligence (AAAI-06)*, pages 796–801, 2006.

[3] E. Bonchek-Dokow and G. A. Kaminka. Towards computational models of intention detection and intention prediction. *Cognitive Systems Research*, 28:44–79, 2014.

[4] J. Hong. Goal recognition through goal graph analysis. *Journal of Artificial Intelligence Research*, 15:1–30, 2001.

[5] N. Lesh and O. Etzioni. A sound and fast goal recognizer. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI-95)*, 1995.

[6] M. Ramırez and H. Geffner. Probabilistic plan recognition using off-the-shelf classical planners. In *International Joint Conference on Artificial Intelligence*, 2010.

[7] G. Rizzolatti. The mirror neuron system and its function in humans. *Anatomy and Embryology*, 210(5–6):419–421, 2005.

[8] S. Sohrabi, A. V. Riabov, and O. Udrea. Plan recognition as planning revisited. *The Twenty-Fifth International Joint Conference on Artificial Intelligence*, pages 3258–3264, 2016.

[9] I. A. Şucan, M. Moll, and L. E. Kavraki. The Open Motion Planning Library. *IEEE Robotics & Automation Magazine*, 19(4):72–82, December 2012.

[10] G. Sukthankar, R. P. Goldman, C. Geib, D. V. Pynadath, and H. Bui, editors. *Plan, Activity, and Intent Recognition*. Morgan Kaufmann, 2014.

[11] M. Vered, G. A. Kaminka, and S. Biham. Online goal recognition through mirroring: Humans and agents. *The Fourth Annual Conference on Advances in Cognitive Systems*, 2016.