Intent Recognition Through Goal Mirroring (Doctoral Consortium)

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

1. INTRODUCTION

As the human population and life expectancy increases so does the need for integrating robots and virtual agents closely into everyday human life. These agents may provide care and attention where otherwise the manpower is lacking. In order to create better agents, that interact seamlessly with humans we need to draw lessons from what we know of human social cognition. Designing an agent inspired by these processes will provide an agent that is more predictable, less-threatening and overall welcome to its' human benefactor.

One important aspect of human social cognition is the innate ability to perform quick and efficient intention recognition. This ability enables humans to reason about the hidden goals of other agents around them through observations of their actions. In humans this ability is hypothesized to come from the existence of a *mirror neuron system*. Mirror neurons have first been discovered in the early 90's. These neurons were seen to fire both when a monkey manipulated an object and also when it saw another animal manipulate an object. Recent neuro-imaging data indicates that the adult human brain is also endowed with a *mirror neuron system* for matching the observation and execution of actions within the adult human brain [4, 6]. This system is hypothesized to give humans the ability to infer the intentions leading to an observed action using their own internal mechanism. It is also attributed to other high level cognitive functions such as imitation, action understanding, intention and language evolution. Consequently, the human mirror neuron system may be viewed as a part of the brains' very own plan/goal recognition module.

Inspired by mirroring processes we have developed *Goal Mirroring*. A fast, online method that works in continuous domains. *Goal Mirroring* uses a planner to dynamically generate plans for given goals, eliminating the need for the traditional plan library. In this we also build on previous approaches —plan recognition by planning (PRP)— [3]. However, while existing PRP based recognizers only operate in discrete domains and in an offline manner *Goal Mirroring* provides an efficient online PRP approach while operating in continuous domains. We have extensively evaluated this approach, over hundreds of experiments, while measuring recognition success over several different planners and three different continuous domains.

Appears in: *Proceedings of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2017), S. Das, E. Durfee, K. Larson, M. Winikoff (eds.), May 8–12, 2017, São Paulo, Brazil.*

2. GOAL MIRRORING

Goal recognition is the problem of inferring the unobserved goal of an agent from a set of observations of the agent's actions and their effects [5]. This set of observations is not necessarily complete or sequential; observations can be missing in the beginning, middle or end of the process. The goal recognition problem is then divided into two variants. In offline goal recognition the set of observations is revealed at once, while in online goal recognition the set of observations is revealed incrementally and there is merit in the identification of the goal as early in the recognition process as possible.

Goal Mirroring is an online goal recognition approach which utilizes a planner within the recognition process. 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, incorporating information from the observations as input to the planners is not a trivial task. Goal Mirroring achieves this by utilizing the planner to only calculate part of the plan, excluding the part already achieved as seen in the observations.

For example, for shape recognition this is accomplished by utilizing a shape planner to attempt to create only the remainder shape while for a goal navigation recognizer it can be accomplished by planning from the last seen observation point to each of the goals.

The resulting goal hypotheses then need to be ranked. In this we drew inspiration from studies of human estimates of intentionality and intended action [1]. 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.

We implemented online goal mirroring and introduced two suggested heuristics to improve recognition performance and efficiency. The first heuristic is the *pruning heuristic*. It aims to decrease overall run-time by reducing the number of times the planner it calls. It achieves this by pruning out goals that seem unlikely or impossible. In this, again, we are inspired by the rationality assumption. Once a rational agent is moving away or past a goal point, that goal is considered an unlikely target and may be pruned.

The second heuristic is the *recomputation* heuristic. The purpose of this heuristic is to refrain from calling the planner when unnecessary, again reducing the overall run-time of the recognition process. If we see that the agent is still heading in the same general direction, we may choose to keep the former goal rankings and not call the planner for recomputation at all.

In order to evaluate our approach we have contrasted our *mirroring* recognizer performance and efficiency, with and without the aforementioned heuristics, over continuous and discrete environments and several different challenging domains.

3. EVALUATION

We evaluated *Goal Mirroring* on three challenging continuous domains. The first is that of an online shape recognizer that identifies multi-stroke geometric shapes without a plan library, by utilizing a shape-drawing planner [6]. Humans increasingly use sketches as part of their communications with other agents, making the ability to recognize shapes, be it drawn on paper, on a computer, or via hand gestures in the air highly important and relevant. Most existing work focuses on offline (post-drawing) recognition methods, trained on large example sets. Given the infinite number of ways in which shapes can appear -rotated, scaled, translated- and inherent inaccuracies in the drawings, these methods do not allow on-line recognition, and require a very large library (or expensive pre-processing) in order to perform recognition.

We conducted a series of experiments utilizing a drawing planner for regular polygons. The basis for the experiment was a data-base of scanned hand-drawn regular polygons. Shapes were drawn in various scales, rotations, and translations with respect to the center of the page. Naturally, hand drawings, even under these ideal conditions, reflect quite a bit of inaccuracy. We ran these shapes through our recognizer and on a group of human participants. We instantiated the shape recognition approach in the recognition of regular polygons and evaluated the performance of different ranking and non-ranking variants of the recognizer against human subjects' recognition. The evaluation utilized several different evaluation criteria. Across the board, our *rational*-based ranking recognition proved superior to non-ranking recognition and other ranking variants. In some cases, the ranking recognizer surpassed human recognition results. However, in general the ranking recognizer performed on par, or just below, human levels of recognition.

The second domain is that of navigational goal recognition. The ability to recognize an agents' projected location target by a sequence of observed agent locations.

We experimented using the OMPL *cubicles* environment and default robot and tested over hundreds of goal recognition problems. The results demonstrated the power of our proposed heuristics and showed that, while powerful by themselves, a combination of them leads to a reduction of a substantial 63% of the calls the recognizer makes to the planner and in overall planner run-time in comparison with a proposed baseline PRP approach. In terms of convergence and overall first ranking, we saw an increase of over 20% in comparison with the baseline approach.

We further evaluated our approach on a simulated robot environment. We used ROS [2] to utilize our recognition algorithm to recognize the goals of navigation in 3D worlds using the ROS *MoveBase* default planner. We compared the recognition performance of our recognizer when using the continuous planner, to that when the recognition was carried out on a discretized grid. In particular, We divided the environment into robot-sized grid cells and converted all consecutive points along the path to the middle of each of the corresponding cells in the grid. In the same manner we also converted the goal locations.

For all problems the *continuous* recognizer ranked significantly higher than the *discrete* instance. This arises from the fact that the discrete recognizer may lose information in the discretization process. It is safe to say that with a reduction of the discretization factor these differences will decrease until the performance will be equivalent. However finding just the right amount of granularity could prove wasteful and domain specific.

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.

REFERENCES

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

[2] M. Quigley, K. Conley, B. Gerkey, J. Faust, T. Foote, J. Leibs, R. Wheeler, and A. Y. Ng. Ros: an open-source robot operating system. In *ICRA workshop on open source software*, volume 3, page 5. Kobe, 2009.

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

[4] G. Rizzolatti, L. Fadiga, V. Gallese, and L. Fogassi. Premotor cortex and the recognition of motor actions. *Cognitive brain research*, 3(2):131–141, 1996.

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

[6] 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.