
Online Goal Recognition through Mirroring: Humans and Agents

Mor Vered
Gal A. Kaminka
Sivan Biham

VEREDM@CS.BIU.AC.IL
GALK@CS.BIU.AC.IL

Computer Science Dept and Brain Science Center, Bar Ilan University, Ramat Gan, Israel

Abstract

Goal recognition is the problem of inferring the (unobserved) goal of an agent, based on a sequence of its observed actions. Inspired by mirroring processes in human brains, we advocate *goal mirroring*, an online recognition approach that uses a black-box planner to generate recognition hypotheses. This approach avoids the prevalent assumption in current approaches, which rely on a dedicated *plan library*, representing all known plans to achieve known goals. Such methods are inherently limited to the knowledge represented in the library. In this paper, we (i) describe a novel online goal mirroring algorithm for continuous spaces; (ii) evaluate a novel heuristic for choosing between competing recognition hypotheses; (iii) contrast machine and human recognition in two challenging domains, revealing insights as to human capabilities; and (iv) compare mirroring to library-based methods.

1. Introduction

Plan, activity, and goal recognition is a fundamental research area in artificial intelligence. In goal recognition, the agent is given a sequence of observed actions, and the objective is to determine the (unobserved) goal of these actions, as early as possible, as the observations are incrementally revealed (Blaylock & Allen, 2004). The problem, closely related to plan, activity, and intent recognition (Sukthankar, Goldman, Geib, Pynadath, & Bui, 2014) has many applications in continuous environments, e.g., for recognizing intended gestures and sketches whether in air or on paper (Sezgin & Davis, 2005), anticipating user commands (Blaylock & Allen, 2004) or for recognizing intended navigational goals (Zhu, 1991).

The prevalent approach to goal recognition relies on a dedicated *plan recognition library*, which represents all known ways to achieve known goals (Sukthankar et al., 2014). Recognition methods—sometimes applicable to continuous domains—vary in the expressiveness of the representation and efficiency of the inference algorithms used. While powerful when the plans are known, these methods tend to 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.

Moreover, the use of a dedicated plan recognition library is space consuming and redundant as part of an integrated agent: An agent that plans and acts in an environment will need a separate

plan recognition library for recognizing the others’ actions, despite having implicit knowledge about what plans in the environment look like.

To address these challenges, we are inspired by mirroring processes in human brains. There has been evidence that humans ability to do online goal recognition comes from the *mirror neuron system* for matching the observation and execution of actions within the adult human brain (Rizzolatti, Fadiga, Gallese, & Fogassi, 1996; Rizzolatti, 2005). The mirror neuron system gives humans the ability to infer the intentions leading to an observed action using their own internal mechanism.

Analogously, we advocate *goal mirroring*, an *online* goal recognition approach which uses a planner as a black box, to generate recognition hypotheses on the fly, eliminating the need for storing plans in advance, in a plan recognition library. It is designed to efficiently handle incremental, continuous observations and tightly integrates planning and recognition: Whatever plan can be planned, can also be recognized.

We describe goal mirroring in detail, and report on extensive goal-recognition experiments in two challenging continuous domains (3D navigation goal recognition and hand-drawn shape recognition). In each, we compare the performance of goal mirroring to human subject recognition, and draw lessons as to human recognition capabilities. In particular, while *goal mirroring* often performs on-par with humans, or just below, it is more capable in recognizing non-optimal plans, and less capable in recognizing plans based on little evidence. This hints that humans employ additional sources of knowledge in their inference of goals. We additionally contribute a novel heuristic for ranking recognition hypotheses, showing it is superior to earlier work. We contrast the recognition results with those achievable with library-based methods and show that *goal mirroring*, utilizing no plan library, can recognize plans just as successfully.

2. Related Work

Prevalent approaches to plan- and goal- recognition rely on a dedicated *plan library* as the basis for the recognition process (see Sukthankar, Goldman, Geib, Pynadath, & Bui (2014) for a recent survey). The plan library efficiently represents all known plans to achieve known goals; methods vary in the representation and inference algorithms used. For instance, hidden Markov model variants (Blaylock & Allen, 2004). However, this comes at a cost: the requirement to store recognition knowledge separately from plan execution knowledge wastes space and novel plans, even for known goals, are difficult, if not impossible, to recognize.

There are exceptions. Cox & Kerkez (2006) present a library-base technique that attempts to handle novel plans. However, the method uses a representation that is inappropriate for continuous environments where actions and ensuing states are not discrete.

Some prior investigations have begun to explore alternative methods. Hong (2001) uses a specialized representation and online algorithm to generate possible goals without a plan library. Some methods unify the plan-recognition and plan-execution libraries: *Agent tracking* (Tambe & Rosenbloom, 1995; Laird, 2001) uses an agent’s own BDI plan to recognize a BDI plan being executed by another. Similarly, Sadeghipour & Kopp (2011) represent (and store) shape drawing *plans*, that can be used both for recognition and execution by the agent. Most recently, ? advocated the use of combinatoric categorical grammars as a representation for both generating and recognizing plans.

In contrast, the key to *goal mirroring* is the use of an off-the-shelf *planner*—the same that would be used by the agent to drive execution—for plan recognition; it does not store plans in any form. In this, it is related to work on offline *plan recognition by planning* (Ramírez & Geffner, 2010), which works by assuming all observations are given at once to a planner-based recognizer. However, goal mirroring differs from Ramírez & Geffner (2010) in several important ways. First, it is intended for online recognition, where it is much more efficient, not having to re-calculate the initial paths to the goal for each iteration (see Section 3.2). Second, goal mirroring is defined for continuous and mixed continuous-discrete domains, while the previous method is defined for discrete domains only (e.g., using PDDL-capable planners), where there is no uncertainty in the observations, and observed actions are discretely defined. We therefore use a different ranking heuristic than earlier work, which we show is significantly superior in continuous domains (Section 4.5).

We depart from previous work also in that we contribute in this paper a comparison of goal mirroring with human recognition capabilities, in two challenging, continuous domains: recognition of shapes from incrementally drawn sketches, and recognition of navigation goals. These experiments allow some insight as to human recognition biases and capabilities.

3. Goal Recognition

We begin by giving a clear definition of the goal recognition problem (Section 3.1) followed by Section 3.2 with an in-depth portrayal of our mirroring algorithm.

3.1 Problem Definition

We define R , the *goal recognition* problem in continuous domains as a tuple $R = \langle W, G, T, T_o, O, M \rangle$. $W \subseteq \mathbb{R}^n$ is the world in which the recognition problem is contained. This includes the familiar work area in \mathbb{R}^n as defined in standard motion planning (LaValle, 2006). For robot poses on flat ground, for instance, W is the space of possible positions and angle in each position (i.e., defined over \mathbb{R}^3). It may also include additional dimensions, such as velocity, color (to capture drawings), etc.

G is a set of $k \geq 1$ goals g_1, \dots, g_k ; each goal $g_j \subset W$ represents a subset of the space, e.g., a point location, a polygon drawn in some color, a trajectory, etc. T limits the duration interval $[0, T]$ in which the observed agent was active, whether observed during this time or not.

The set of observations O is defined for a subset $T_o \subseteq [0, T]$. T_o may include *specific* times in which observations were made, or *continuous* intervals of time in which observations were made. We define the set of observations $O : T_o \rightarrow 2^W$ as a mapping such that for any observation time $t \in T_o$, there exists $O(t) \subset W$, i.e., each observation is of a specific subset of the work area, e.g., a point, an edge, a trajectory segment which includes velocity, etc.

Given the problem R , the task is to choose a specific goal $g \in G$ that *best matches* the observations O . We formulate this intuition by including M , a set of *plans*, in the definition of R , where at least one of the plans is assumed to account for the observations in O . Formally, a plan m_g^i , indexed by a goal $g \in G$, with $i \geq 1$, is a mapping $m_g^i : [0, T] \rightarrow 2^W$ from a time stamp $t \in [0, T]$ to a subset of the work area W , such that $m_g^i(1) = g$. In other words, a plan is a procedure that incrementally generates subsets of W , until the final subset at $t = 1$ is one of the goals $g \in G$. Intuitively, we are describing a plan by its effects on the world. This general definition of a plan avoids the question of the mechanisms by which effects are generated, and thus necessarily admits different approaches to representing or generating the plan set M (as we discuss below).

A matching between a specific plan $m \in M$ and the observations O is defined by the *matching error* $e(m, O) = \int_{t \in T_o} D(O(t), m(t))$ for a distance metric D , such that $D \geq 0$ for all $t \in T_o$ (i.e., for all observations), and specifically $D = 0$ if $O(t) = m(t)$. m is said to *perfectly match* the observations if for any $t \in T_o$, $m(t) = O(t)$, in which case $e(m, O)$ is 0.

A solution to the goal recognition problem R minimizes $e(m_g, O)$; it is a member of the subset $S_R \subseteq G$ minimizing $e(m_s, O)$: $S_R = \{s \mid \operatorname{argmin}_{s \in G} e(m_s, O)\}$.

In general, this condition (minimizing $e(m_g, O)$) is *necessary*, but *not sufficient*. Any number of potential plans, especially in continuous domains, may perfectly match the observations, but differ in the unobserved parts. In general, a plan can perfectly match the observations, and then still achieve any goal g . For instance, in navigation goal recognition, suppose we observe points leading to a goal in the north. A path planner may generate a path that goes through all the observed points, and then doubles back to the south. Such a path will perfectly match all observations, but add an arbitrary suffix for any goal g .

This necessary—but insufficient—condition can be understood as a result of the abductive nature of plan-recognition; reasoning to the best explanation out of a potentially infinite set of explanations cannot be done without defining the necessary condition that allows to filter out non-explanations. This is what the condition above does. A second—separate—condition must define sufficiency; we discuss this in Section 3.2.

Online goal recognition. In this paper, we specifically address the online version of this problem, where the set O is revealed incrementally. Specifically, we set $t = 0$, and increment it until $t = T$. For every value of t , we may induce a goal recognition problem $R(t) = \langle W, G, t, T_o^t, O^t, M^t \rangle$, where T_o^t, O^t, M^t are defined as the respective subsets of T_o, O and M induced over the duration $[0, t]$. We denote the solution of $R(t)$ by $S_R(t)$. The objective of the online goal recognition problem is to minimize $t \in [0, T]$ such that $S_R(t) = S_R$.

In principle, it is possible to naively solve the online goal recognition problem by repeatedly calling an offline goal recognizer with the problem $R(t)$, as t increases and the latest new observation $O^t(t)$ is made available. However, this is quite inefficient, as the growing set O^t is processed anew with every call. Thus the challenge is to determine an efficient solution.

3.2 Online Goal Mirroring

Goal mirroring uses a planner to generate M dynamically in the recognition process, instead of representing the plans explicitly, as a library of plans. The planner is used to dynamically generate plans $m \in M$ as needed. This raises two key challenges.

First, the planner needs to maintain the *necessity condition*; it must generate a plan m_g that agrees with the observations, in order to minimize the matching error $e(m_g, O)$ defined above. It therefore needs to *fold the observation history* into m_g . For STRIPS-like discrete planners, Ramírez & Geffner (2010) have shown an elegant way to do this, by changing the domain theory used by the planner. But in continuous spaces, e.g., by most motion planners, this cannot be done. We therefore build a plan m_g that includes the observations, by appending a prefix plan (composed of the observations in time t) and a suffix plan (composed of a new generated plan, from the latest observed state to the goal g). We denote the suffix plan m'_g .

The second challenge is that the plan m_g must also meet a *sufficiency condition*. To address this, goal mirroring is biased towards rationality. It uses the planner to also generate \bar{m}_g , an *ideal plan* that ignores all observations, and simply reaches g from the *initial* observed state, denoted $O(\emptyset)$. If the cost $cost(m_g)$ —where m_g is made from the observations thus far and the suffix plan m'_g —greatly exceeds the $cost(\bar{m}_g)$, we rank g lower (or eliminate it from S_R). The underlying assumption here is that the ideal plan is optimal; if the observed plan is far from the ideal plan, then the agent must not be rational, and is likely pursuing an alternative goal altogether.

Algorithm 1 integrates these insights for online goal mirroring. It accepts as input a recognition problem R and a *planner* to be used as a black box. It then works as follows. First, in lines 1–2 we call the planner to compute the plan \bar{m}_g (initial state $O(\emptyset)$) to each goal g , ignoring observations). This avoids re-computation of \bar{m}_g with every loop iteration.

The loop in lines 3–7 is comprised of two steps. The first step (lines 4–6), centers on approximating the cost of plan m_g (which folds the observations), from the cost of the prefix (maintained by Δ), and the cost of a suffix plan m'_g .

The second step then ranks the goal hypotheses. For each goal, line 7 assigns a score, which is the ratio of the costs of \bar{m}_g and the approximated m_g . As differences between them grow, the ratio of the costs decreases, resulting in a lower score. In line 8, we transform these scores into probabilities $P(G|O)$ via the normalizing factor $\eta = 1 / \sum_{g \in G} score(g)$.

Algorithm 1 ONLINE GOAL MIRRORING ($R, planner$)

```

1: for all  $g \in G$  do
2:    $\bar{m}_g \leftarrow planner(W, g, O(\emptyset))$ 
3: for  $t = 0$  to  $T$  do
4:    $\Delta \leftarrow cost(O^t)$ 
5:   for all  $g \in G$  do
6:      $m'_g \leftarrow planner(W, g, O^t(t))$ 
7:      $score(g) \leftarrow cost(\bar{m}_g) / (\Delta + cost(m'_g))$ 
8:    $P(G|O(t)) \leftarrow \eta \cdot score(g)$ 

```

The choice to use this ratio in line 7 is not arbitrary. There exists evidence that human estimates of intentionality in action are heavily biased towards motion that is efficient (or rational), i.e. preferring hypotheses that do not deviate from the optimal, rational plan. A cost ratio between this plan (\bar{m}_g if the planner is optimal) and the observed plan m_g is known to agree with human judgement of intentions (Bonchek-Dokow & Kaminka, 2014).

A different heuristic, though motivated by the same principle, is suggested by Ramírez & Geffner (2010): they propose looking at the *difference* in costs, i.e., a score inversely proportional to $|cost(m_g) - [\Delta + cost(m'_g)]|$. We believe that a difference may be biased when dealing with larger cost values, where small differences may still be very large and skew results. In Section 4 we empirically contrast these two heuristics.

4. Experiments

We empirically evaluate goal mirroring in two challenging continuous domains, over multiple goal recognition problems. Section 4.1 presents the domains and measures used in evaluation. Section 4.2 reports on the main set of recognition results of goal mirroring, contrasting them with those

of human subjects. Section 4.3 contrasts human and goal mirroring performance when the observed plans are not optimal, directly touching on key assumptions in goal mirroring and other work. Finally, we contrast goal mirroring with a method based on a plan library (Section 4.4), and with previous work on plan recognition by planning (Section 4.5).

4.1 Experiment setup

4.1.1 Two recognition domains

We evaluate the performance of goal mirroring in two vastly different, continuous domains; sketch recognition and navigational goal recognition.

Recognizing sketches of regular polygons. Here the task is to recognize 2D hand-drawn regular polygons. The use of sketches, drawn on paper, on a computer, or via hand gestures in the air, as part of every day communications is continually increasing. By evaluating the recognition in this domain we might be able to gather some information as to how humans perform this recognition and what biases they encounter.

We had three people draw (by hand) equilateral triangles, squares, pentagons, hexagons, septagons, and octagons, for a total of 18 drawings. Shapes were drawn in various scales and rotations. Naturally, hand drawings, even under ideal conditions, reflect quite a bit of inaccuracy.

Each of the 18 drawings was revealed one edge at a time, with the goal of correctly identifying the goal shape, i.e., there are 18 online recognition problems. Observations, of the edges were generated by using machine vision to analyze the drawings. Specifically, we used OpenCV to implement a Hough-transform edge detector (Duda & Hart, 1972) which we used to identify coordinates of the initial and last points in the drawing, marking the limits of the initial and current observed edge. To overcome scanning noise and drawing inaccuracy (which causes the Hough transform to generate multiple candidate edges) hierarchical clustering (de Hoon, Imoto, Nolan, & Miyano, 2004) was used to estimate the actual number of edges. See (Vered & Kaminka, 2015) for more details.

We developed a shape-drawing planner, which takes a partial drawing (as an initial state), and a goal shape type, and attempts to complete the drawing to the goal shape (or report failure if cannot be done, e.g., attempting to complete a 4-edge open polygon into a triangle). This is the planner we utilize in the recognition process. To rank hypotheses (Algorithm 1, line 7), we looked at the ratio between the ideal internal angle size for the goal shape, and the mean observed internal angle.

3D Navigation Recognition. In this domain the task is to identify the goal location of an object observed moving in a 3D continuous world. The observations are made incrementally, as a sequence of waypoints reached by the object as it moves. This domain, represents a common goal recognition problem, applicable to a broad range of scenarios; be it in teamwork, where one participant needs to anticipate the direction of the other or in adversarial scenarios where one agent tries to intercept another. It is especially interesting to witness the comparison between human recognition and the goal-mirroring recognizer in this domain due to the formidable spatial reasoning abilities of humans.

To carry out the experiments, we utilized the Open Motion Planning Library (OMPL; Şucan, Moll, & Kavraki (2012)), in particular its *cubicles* 3D environment, the default robot, and the TRRT planner that comes with OMPL. The cost measure (Algorithm 1, lines 4 and 7) is simply the length of the path.

We selected six points spread out over the environment as the goal set, G . The points are shown in Figure 1(a). Four of the points were chosen arbitrarily over the larger visible surface of the cubicles environment, the same surface where the observed object starts; these points were picked such that they might be intuitive or easy goal positions for humans (points A - D). Two additional points were specifically chosen to be harder, based on some pilot studies. One point requires humans to think about the other side of the environment (a point on the “other floor” of the environment), point F . The other point hangs in mid-air in the opening between the two “floors” of the environments, point E .

To generate recognition problems, we generated paths from a fixed starting position to all six goal points. Two paths were generated for each goal, for a total of 12 problems: one path generated by the asymptotically-optimal RRT* algorithm implemented in OMPL (Şucan et al., 2012), and another hand-modified to deviate from this optimal path by taking a longer route. The motivation for generating such paths is to examine recognition performance with observations of non-optimal plans, which test the rationality assumption of goal mirroring and previous works (see Section 4.3).

4.1.2 Evaluating recognition across domains

We evaluated goal mirroring in both domains, demonstrating the general applicability of the technique. However, to gain insight as to its strengths and weaknesses and to the general performance of humans in the problem of goal recognition, we require recognition performance measures that are neutral with respect to the domain and any specific problem. The performance measures must allow comparison across a wide variety of observation sequence lengths, sizes of the goal set G , indeed the specifics of any domain and recognition problems within it.

Measuring recognition results across domains and problems. We define three performance measures below, using an example run to assist in the presentation. Let us examine the recognizer output on a specific problem, here in the 3D navigation domain. Figure 1(b) shows the recognition results for *goal mirroring* and one human participant in the navigation domain. The X-axis marks the observations coming in incrementally. The Y axis measures the rank of the correct goal hypothesis among all the goals ranked by the recognizer, thus lower is better (rank 1 indicates that the correct goal was ranked as the top hypothesis). Naturally, this is post-hoc analysis; the recognizer does not have access to the ground truth during the run.


After making two observations we can see that the correct goal was ranked 2 (out of 6) by the goal mirroring recognizer and 4 by human participant. As more and more observations come in, it is only natural that the recognition problem becomes easier, and indeed towards the middle of the observation sequence, both instances converge to ranking the correct goal at the top of their ranking (i.e., rank 1).

Such graphs can be drawn for any specific online recognition problem instance, to compare the performance of different recognizers. Recognizers may vary in three measures: (1) the time (measured by number of observations from the end) in which the recognizer converged to the correct hypothesis (including 0 if it failed); (2) the area under the curve drawn in this graph, the false-positive response (greater area means recognizer tended to rank the correct hypothesis lower, farther from top); and (3) the number of times they ranked the correct hypothesis at the top (i.e., rank 1), which indicates their general accuracy.

For example in Figure 1(b), the goal mirroring recognizer converges at observation 6, whereas the human participant converges earlier, at observation 5, i.e., was quicker to converge. Normalizing for the observation sequence length along that specific path, to allow comparison across different recognition problems, we measure the normalized convergence of the goal mirroring recognizer at 40%, and of the human participant at 50%. Higher results indicate earlier convergence, thus better.

Measuring the area-under-curve (AUC) gives us a measure of the uncertainty of the recognizer during the recognition process. Here, a lower value is considered better indicating that the recognizer was closer to the correct ranking along most of the process. For instance, in Figure 1(b) it is clear that the AUC for the *Goal-Mirroring* recognizer is smaller than for the *Human* recognizer. We can again normalize to allow comparison between different recognition problems, even normalizing for the number of potential goals considered. We compute the ratio of the AUC to the worst-case scenario, where a recognizer consistently ranked the correct hypothesis as the lowest rank (i.e., at $\text{rank}=|G|$). A smaller percentage indicates fewer false positives considered by the recognizer. To be consistent with the other measures (where a larger result is better), we consider the complimentary normalized value (1-normalized AUC).

Finally, counting the amount of times the planner ranked the correct goal as the top hypothesis ($\text{rank}=1$) gives us an overall measure of the reliability of the recognizer. The more frequently the recognizer ranked the correct hypothesis at the top, the more reliable it is, hence a larger value is better. We again normalize using the length of the observation sequence. In Figure 1(b), the goal mirroring recognizer ranked the correct goal at the top 5 times, whereas the human participant ranked it a total of 6 times. In this instance, the human participant is the better recognizer.



(a) 3D navigation domain: *Cubicles* environment, with goals and initial position of the object. Both surfaces (“floors”) shown.



(b) Recognition results of goal mirroring vs. a human subject, for one problem with 10 observations, in the 3D navigation domain.

Figure 1.

4.2 Goal mirroring and human recognition results

We conducted recognition experiments in the two domains, measuring performance of goal mirroring as well as human subjects using the three measures described. The exact same online goal recognition problems were given to the human subjects as to the mirroring recognizer. In both domains, humans had immediate access to the goal library; they were shown the possible goals at all time, visually, so they did not have to rely on memory. After each observation was revealed, human subjects were asked to provide a ranking for the goals, and to rule out any goals which they felt were no longer possible.

In the 3D navigation domain, we tested 19 human subjects (8 women; ages 17–51, mean 27.5). Results for this domain are shown in Figure 2; the X-axis denotes the 12 paths, organized in pairs: 1–2 for goal point A, 3–4 for goal point B, etc. In the shape recognition domain, we tested 20 human subjects (14 men; ages 19–52, mean 29). Results are shown in Figure 3; the X-axis denotes the 6 goal shapes. Both figures show the results for the three measures separately; light bars indicate goal mirroring, dark bars represent the human mean results. Higher bars indicate better results.




Figure 2. Recognition results in the 3D navigational domain. X-axis:12 paths, two to each goal. Y-axis: Success ratio, higher is better.




Figure 3. Recognition results in the shape recognition domain.

The figures show that in the shape recognition domain (Figure 3), goal mirroring generally performs on par with human recognition performance—or close to it—in all three measures. However, it is very different in recognizing triangles. In the 3D navigation domains (Figure 2), the results of mirroring and humans are not as close. For instance, humans are better at the convergence measure for goals A,C (paths 1–2, 5–6, resp.). They are worse for goals E,F (paths 9–10, 11–12, resp.).

We draw several lessons from these results. First, humans likely use additional knowledge (not part of the recognition problem as defined), to make recognition decisions. This is evident when we examine the greatly inferior recognition results of humans for goals E, F in the 3D navigation domain. Humans repeatedly ignored these goal points in their ranking (despite being presented to them throughout the experiment). We believe this hints at additional knowledge being used about what constitutes a goal, coinciding with work on action parsing (Baird & Baldwin, 2001).

A related second lesson w.r.t using knowledge of goals is that humans do not just rule-out goals too quickly, they also commit to them too early. In the shape recognition domain, the superior recognition of humans in identifying triangles was achieved often after having seen only a single edge. Somehow, humans successfully inferred that this single edge was part of a triangle, despite the fact that at that point all shapes were equally possible. Analysis of errors—which we do not present here for lack of space—shows that this early commitment came at a cost of making more errors, ruling out the correct goal shapes early on.

4.3 Optimality assumption

As discussed in Section 3, generating a candidate hypothesized plan that matches the observations is only a necessary step in plan recognition. Multiple such hypotheses exist, and the key is to determine a sufficient condition for selecting between them. In library-based methods, a plan can be a solution if it can be found in the library of plans. But in planner-based methods, where candidates are generated from scratch, dynamically, a different approach is needed.

Previous works (Ramírez & Geffner, 2010; Bonchek-Dokow & Kaminka, 2014) have proposed to assume that the observed plan is carried out by a (bounded) rational agent, which means that the observed plan should approximate the optimal plan for getting the initial state to the goal state. Recognizers based on this assumption prefer hypothesized plans that better match the optimal plan.

To evaluate this, we purposely generated two different observed paths to each possible goal in the 3D navigation domain, as described above. One path was the optimal. The other was modified from it by introducing a detour which made it longer, though still smooth and executable by the moving object. This detour purposely made the path to the goal pass near other goals on the way. In this way we were able to evaluate how much both recognizers would deteriorate due to rational spatial reasoning. By comparing the performance of goal mirroring and humans on each of these two paths we hope to gain insight as to the importance of the rationality assumption in recognition.

Figure 4 presents the deterioration in both human and goal mirroring recognition when observing an optimal and a non-optimal plan (path) to the same goal, across the three criteria. The X-axis in each figure shows the goal. The bars mark improvement or degradation: values that are larger than zero indicate an improvement in performance when observing the non-optimal plan. Values smaller than zero indicate degradation in recognition. Higher values are better: Larger positive values indicate more improvement; smaller (closer to 0) negative values show less degradation.




Figure 4. Comparison of degradation from optimal to non optimal paths.

The figure shows that overall, mirroring deteriorates less, or improves more, than humans. This happens across all goals except D, for the convergence measure (Figure 4(a)). It occurs for many of the goals in AUC (Figure 4(b)), and for all goals but one (A) for the ranked-first measure (Figure 4(c)). The results are especially striking in the ranked first and AUC measure, where sometimes mirroring results positively improve when observing non-optimal plans, while humans results degrade.

From the perspective of building automated recognizers, these are promising results, demonstrating the robustness of goal mirroring. From a perspective of human cognitive modeling, these results hint at a very strong reliance of humans on the rationality of the observed plan. As a result,


goal mirroring as presented in this paper is still not a good enough model of human recognition capabilities.

4.4 Mirroring vs. library-based methods


Goal mirroring has the principled advantage over library-based methods in terms of storage, and in being able to handle any arbitrary initial observed state; no need to add possible plans to a plan library and the only information saved between processing one observation and the next is the *cost* of each path (see Section 3.2). However, this is also a disadvantage, in principle: goal mirroring does not utilize prior knowledge even when it can be made available. To evaluate this aspect, we contrasted mirroring with a hidden Markov model (HMM), a popular library-based technique, often used as a standard technique, e.g., (Blaylock & Allen, 2004).

Testing the HMM on plans unknown to it is a valid, but futile exercise, where the superiority of goal mirroring would be obvious. We therefore evaluate HMMs vs goal mirroring when the plans are known to the HMM. To do this, we first needed to discretize the navigation problem. we created a robot-size cell grid in the 3D environment, each cell represented by a state in a hidden Markov model. To generate goal recognition problems, we arbitrarily selected 11 points in the environment. We generated observed paths from each point to all others, for *a total of 110 goal recognition problems*. For each such problem, we trained (MATLAB HMM package) one HMM, using 20 paths generated by the RRT* planner as training data. In other words, we created a specialized HMM, trained on 20 examples of asymptotically-optimal data, for each recognition problem.

Figure 5(a) contrasts the recognition results of HMMs and goal mirroring. Even without any prior knowledge, goal mirroring is on-par with the HMM results, even better in the *AUC* measure. Obviously, as more prior knowledge is available, this can change. Our conclusion is that goal mirroring should be preferred when relatively less data is available, or when the number of possible plans is very large (or infinite, as in these two domains).



(a) Goal mirroring vs HMM



(b) Goal mirroring vs Plan Recognition By Planning

Figure 5. Comparison to two different approaches.

4.5 The effects of the ranking heuristic

Finally, we empirically contrast the ranking heuristic we propose (ratio of costs), to that of (Ramírez & Geffner, 2010) (difference of costs). Figure 5(b) shows the mean results when applying the different heuristics on the 12 3D navigation problems. In all three measures, the ratio heuristic is

clearly superior. We believe that this is because paths (to different goals) can substantially vary in length. The difference heuristic compares absolute differences between paths of different lengths, while the ratio heuristic compares relative differences.

5. Summary

We have presented *online goal mirroring*, a goal recognition approach for continuous domains that does not rely on a plan library, but instead uses a planner to generate recognition hypotheses that are continually matched against incremental observations. We demonstrated the generality of goal mirroring by performing extensive experiments in two separate challenging domains. We have shown the improved recognition performance of goal mirroring over earlier attempts. We further contrasted the recognition performance of goal mirroring with humans showing that while goal mirroring often performs on-par with humans, or just below, it is more robust to observing non-optimal plans. Finally, we demonstrated that, in essence, mirroring can recognize plans as successfully as library-based methods (up to a limit).

Acknowledgements

We thank Kobi Gal for valuable advice. This research was supported by ISF Grant 1511/12. Thanks to K. Ushi.

References

- Baird, J. A., & Baldwin, D. A. (2001). Making sense of human behavior: Action parsing and intentional inference. *Intentions and intentionality: Foundations of social cognition*, (pp. 193–206).
- Blaylock, N., & Allen, J. (2004). Statistical goal parameter recognition. *Proceedings of the Fourteenth International Conference on Automated Planning and Scheduling (ICAPS-04)* (pp. 297–304).
- Bonchek-Dokow, E., & Kaminka, G. A. (2014). Towards computational models of intention detection and intention prediction. *Cognitive Systems Research*, 28, 44–79.
- Cox, M. T., & Kerkez, B. (2006). Case-based plan recognition with novel input. *Control and intelligent systems*, 34, 96.
- Duda, R. O., & Hart, P. E. (1972). Use of the hough transformation to detect lines and curves in pictures. *Communications of the ACM*, 15, 11–15.
- Hong, J. (2001). Goal recognition through goal graph analysis. *Journal of Artificial Intelligence Research*, 15, 1–30.
- de Hoon, M. J., Imoto, S., Nolan, J., & Miyano, S. (2004). Open source clustering software. *Bioinformatics*, 20, 1453–1454.
- Laird, J. E. (2001). It knows what you’re going to do: adding anticipation to a quakebot. *Proceedings of the Fifth International Conference on Autonomous Agents (Agents-01)* (pp. 385–392). Montreal, Canada: ACM Press.

- LaValle, S. M. (2006). *Planning Algorithms*. Cambridge University Press.
- Ramírez, M., & Geffner, H. (2010). Probabilistic plan recognition using off-the-shelf classical planners. *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence (AAAI-10)*.
- Rizzolatti, G. (2005). The mirror neuron system and its function in humans. *Anatomy and Embryology*, 210, 419–421.
- Rizzolatti, G., Fadiga, L., Gallese, V., & Fogassi, L. (1996). Premotor cortex and the recognition of motor actions. *Cognitive brain research*, 3, 131–141.
- Sadeghipour, A., & Kopp, S. (2011). Embodied gesture processing: Motor-based integration of perception and action in social artificial agents. *Cognitive Computation*, 3, 419–435.
- Sezgin, T. M., & Davis, R. (2005). Hmm-based efficient sketch recognition. *Proceedings of the 10th International Conference on Intelligent User Interfaces* (pp. 281–283). ACM.
- Şucan, I. A., Moll, M., & Kavraki, L. E. (2012). The Open Motion Planning Library. *IEEE Robotics & Automation Magazine*, 19, 72–82.
- Sukthankar, G., Goldman, R. P., Geib, C., Pynadath, D. V., & Bui, H. (Eds.). (2014). *Plan, Activity, and Intent Recognition*. Morgan Kaufmann.
- Tambe, M., & Rosenbloom, P. (1995). RESC: An approach for real-time, dynamic agent tracking. *International Joint Conference on Artificial Intelligence* (pp. 103–111).
- Vered, M., & Kaminka, G. A. (2015). If you can draw it, you can recognize it: Mirroring for sketch recognition. *Proceedings of the AAMAS Workshop on Human-Agent Interaction Design and Models*.
- Zhu, Q. (1991). Hidden markov model for dynamic obstacle avoidance of mobile robot navigation. *Robotics and Automation, IEEE Transactions on*, 7, 390–397.