

Landmark-Based Heuristics for Goal Recognition

Ramon Fraga Pereira[†] and Nir Oren[‡] and Felipe Meneguzzi[†]

[†]Pontifical Catholic University of Rio Grande do Sul, Brazil

ramon.pereira@acad.pucrs.br

felipe.meneguzzi@pucrs.br

[‡]University of Aberdeen, United Kingdom

n.oren@abdn.ac.uk

Abstract

Automated planning can be used to efficiently recognize goals and plans from partial or full observed action sequences. In this paper, we propose goal recognition heuristics that rely on information from planning landmarks — facts or actions that must occur if a plan is to achieve a goal when starting from some initial state. We develop two such heuristics: the first estimates goal completion by considering the ratio between achieved and extracted landmarks of a candidate goal, while the second takes into account how unique each landmark is among landmarks for all candidate goals. We empirically evaluate these heuristics over both standard goal/plan recognition problems, and a set of very large problems. We show that our heuristics can recognize goals more accurately, and run orders of magnitude faster, than the current state-of-the-art.

1 Introduction

Goal recognition is the task of recognizing agents’ goals from partial or full observed action sequences. Most goal and plan recognition approaches (Avrahami-Zilberbrand and Kaminka 2005; Geib and Goldman 2009; Mirsky et al. 2016) employ plan libraries to represent agent behavior (*i.e.*, a library that describes all plans for achieving goals), and plan recognition techniques which use such libraries are analogous to parsing. However, work that use classical planning domain definitions to represent potential agent behavior bring goal and plan recognition closer to automated planning (Ramírez and Geffner 2009; Ramírez and Geffner 2010; Pattison and Long 2010; Keren, Gal, and Karpas 2014; E.-Martín, R.-Moreno, and Smith 2015; Sohrabi, Ribov, and Udrea 2016; Pereira and Meneguzzi 2016). These approaches — which do not use plan libraries — show that automated planning techniques can be used to efficiently recognize goals and plans. In many domains where goal and plan recognition are important (*e.g.*, smart environments, user monitoring and crime detection), plan libraries may be unavailable, making this second class of approaches important. In this paper, we propose two goal recognition heuristics based on planning techniques that rely on planning landmarks (Hoffmann, Porteous, and Sebastia 2004). In automated planning, landmarks are properties (or actions) that

every plan must satisfy (or execute) at some point in every plan execution to achieve a goal. Although landmarks are often used to build planning algorithms (Richter and Westphal 2010), our approach allows goal recognition heuristics to rule out candidate goals whose landmarks are inconsistent with the observations.

Our contribution is twofold. First, we develop a new heuristic approach for goal recognition that obviates the need to execute a planner multiple times (Ramírez and Geffner 2009) yielding substantial runtime gains. Second, we develop two goal recognition heuristics which exploit landmarks. The first of our heuristics estimates goal completion by considering the ratio between achieved and extracted landmarks of a candidate goal. The second heuristic uses a *landmark uniqueness value*, representing the information value of the landmark for some specific candidate goal when compared to landmarks for all candidate goals. This latter heuristic then estimates which candidate goal is correct by summing the uniqueness value of the achieved landmarks from the observations. We empirically evaluate our goal recognition heuristics against the current state-of-the-art (2009) by using a dataset developed by Ramírez and Geffner (2009; 2010), and a dataset we generated for other planning domains with larger problems. We show that our heuristics are substantially faster and more accurate than the state-of-the-art for datasets that contain several domains and problems where recognizing the intended goal is non-trivial.

2 Background

2.1 Planning

Planning is the problem of finding a sequence of actions that achieves a goal from an initial state (Ghallab, Nau, and Traverso 2004). A *state* is a finite set of facts that represent logical values according to some interpretation. *Facts* can be either positive, or negated ground predicates. A predicate is denoted by an n -ary predicate symbol p applied to a sequence of zero or more terms $(\tau_1, \tau_2, \dots, \tau_n)$. An *operator* is represented by a triple $a = \langle name(a), pre(a), eff(a) \rangle$ where $name(a)$ represents the description or signature of a ; $pre(a)$ describes the preconditions of a — a set of facts or predicates that must exist in the current state for a to be executed; $eff(a) = eff(a)^+ \cup eff(a)^-$ represents the effects of a , with $eff(a)^+$ an *add-list* of positive facts or predicates, and

$eff(a)^-$ a delete-list of negative facts or predicates. When we instantiate an operator over its free variables, we call the resulting ground operator an *action*. A *planning instance* is represented by a triple $\Pi = \langle \Xi, \mathcal{I}, G \rangle$, in which $\Xi = \langle \Sigma, \mathcal{A} \rangle$ is a *planning domain definition*; Σ consists of a finite set of facts and \mathcal{A} a finite set of actions; $\mathcal{I} \subseteq \Sigma$ is the initial state; and $G \subseteq \Sigma$ is the goal state. A *plan* is a sequence of actions $\pi = \langle a_1, a_2, \dots, a_n \rangle$ that modifies the initial state \mathcal{I} into one in which the goal state G holds by the successive execution of actions in a plan π . While actions have an associated cost, as in classical planning, we assume that this cost is 1 for all actions. A plan π is considered optimal if its cost, and thus length, is minimal.

2.2 Goal Recognition

Goal recognition is the task of recognizing agents’ goals by observing their interactions in an environment (Sukthankar et al. 2014). In goal recognition, such observed interactions are defined as available evidence that can be used to recognize goals. As proposed by Ramírez and Geffner (2009; 2010), we formally define a goal recognition problem over a planning domain definition as follows.

Definition 1 (Goal Recognition Problem). *A goal recognition problem is a tuple $T_{GR} = \langle \Xi, \mathcal{I}, \mathcal{G}, O \rangle$, in which $\Xi = \langle \Sigma, \mathcal{A} \rangle$ is a planning domain definition; \mathcal{I} is the initial state; \mathcal{G} is the set of possible goals, which include a hidden goal G (i.e., $G \in \mathcal{G}$); and $O = \langle o_1, o_2, \dots, o_n \rangle$ is an observation sequence of executed actions, with each observation $o_i \in \mathcal{A}$.*

The solution for a goal recognition problem is the hidden goal $G \in \mathcal{G}$ that the observation sequence O of a plan execution achieves. Note that the observation sequence can be full or partial — in a full observation sequence we observe all actions of an agent’s plan; in a partial observation sequence, only a sub-sequence of actions are observed.

2.3 Landmarks

In the planning literature, landmarks are defined as necessary properties (actions) that must be true (executed) at some point in every valid plan to achieve a particular goal. Landmarks are often partially ordered based on the sequence in which they must be achieved. Hoffman et al. (2004) define landmarks as follows:

Definition 2 (Fact Landmark). *Given a planning instance $\Pi = \langle \Xi, \mathcal{I}, G \rangle$, a formula L is a fact landmark in Π iff L is true at some point along all valid plans that achieve G from \mathcal{I} . A landmark is a type of formula (e.g., a conjunctive or disjunctive formula) over a set of facts that must be satisfied at some point along all valid plan executions.*

Landmark extraction algorithms such as that of Hoffman et al. (2004) can extract conjunctive landmarks from a planning instance. To represent landmarks and their ordering, this algorithm uses a tree in which nodes represent landmarks and edges represent necessary prerequisites between landmarks. Each node in the tree represents a conjunction of facts that must be true simultaneously at some point during plan execution, and the root node is a landmark representing

the goal state. Hoffman et al. (2004) 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 (Bylander 1994). Thus, to operate efficiently, most landmark extraction algorithms extract only a subset of landmarks for a given planning instance.

For our approach, we use Hoffman et al. (2004)’s algorithm, which extracts a set of fact landmarks for every goal G in the set of candidate goals \mathcal{G} from the initial \mathcal{I} . We refer to this algorithm as the function `EXTRACTLANDMARKS`, which takes as input a planning domain definition $\Xi = \langle \Sigma, \mathcal{A} \rangle$; an initial state \mathcal{I} ; and a set of candidate goals. This function outputs a map \mathcal{L}_G that associates candidate goals to their respective ordered fact landmarks (i.e., a set of landmarks with an order relation).

3 Computing Achieved Landmarks

Our approach requires the ability to identify which fact landmarks were achieved based on a partial sequence of action observations. The evidence we use to infer fact landmarks consist of the preconditions and effects of observed actions during plan execution. This is achieved by function `COMPUTEACHIEVEDLANDMARKS` shown in Algorithm 1, which takes as input an initial state \mathcal{I} , a set of candidate goals \mathcal{G} , a sequence of observed actions O , and a map \mathcal{L}_G containing candidate goals and their extracted fact landmarks. Given these input parameters, the algorithm iterates over the set of candidate goals \mathcal{G} (Line 3) selecting the fact landmarks \mathcal{L}_G of each goal G in \mathcal{L}_G in Line 4 and computes the fact landmarks that are in the initial state in Line 5. With this information, the algorithm iterates over the observed actions O to compute the achieved fact landmarks of G in Lines 6 to 9. For each observed action o in O , the algorithm computes all fact landmarks of G that are in the preconditions and effects of o in Line 7. As we deal with partial observations in a plan execution some executed actions may be missing from the observation sequence, thus whenever we identify a fact landmark, we also infer that its predecessors must have been achieved in Line 8. For example, consider that the set of fact landmarks to achieve a goal from a state is represented by the following ordered facts: $(at\ A) \prec (at\ B) \prec (at\ C) \prec (at\ D)$, and we observe just one action during a plan execution, and this observed action contains the fact landmark $(at\ C)$ as an effect. From this observed action, we can infer that the predecessors of $(at\ C)$ must have been achieved before this observation (i.e., $(at\ A)$ and $(at\ B)$), and therefore, we also include them as achieved landmarks. At the end of each iteration over an observed action o , the algorithm stores the set of achieved landmarks of G in \mathcal{AL}_G in Line 9. Finally, after computing the evidence of achieved landmarks in the observations for a candidate goal G , the algorithm stores the set of achieved landmarks \mathcal{AL}_G of G in Λ_G (Line 11) and returns a map Λ_G containing all candidate goals and their respective achieved fact landmarks (Line 13).

As an example of how the algorithm computes achieved fact landmarks from observations, let us consider the `BLOCKS-WORLD` example shown in Figure 1, which shows an initial state and a set of candidate goals. The initial state

Algorithm 1 Compute Achieved Landmarks From Observations.

Input: \mathcal{I} initial state, \mathcal{G} set of candidate goals, O observations, and \mathcal{L}_G goals and their extracted landmarks.

Output: A map of goals to their achieved landmarks.

```

1: function COMPUTEACHIEVEDLANDMARKS( $\mathcal{I}, \mathcal{G}, O, \mathcal{L}_G$ )
2:    $\Lambda_G \leftarrow \langle \rangle$   $\triangleright$  Map goals  $G$  to their respective achieved landmarks.
3:   for each goal  $G$  in  $\mathcal{G}$  do
4:      $\mathcal{L}_G \leftarrow$  fact landmarks of  $G$  s.t.  $\langle G, \mathcal{L}_G \rangle$  in  $\mathcal{L}_G$ 
5:      $\mathcal{L}_{\mathcal{I}} \leftarrow$  all fact landmarks  $L \in \mathcal{I}$ 
6:     for each observed action  $o$  in  $O$  do
7:        $\mathcal{L} \leftarrow$  all fact landmarks  $L$  in  $\mathcal{L}_G$  such that  $L$ 
          $\in \text{pre}(o) \cup \text{eff}(o)^+$  and  $L \notin \mathcal{L}$ 
8:        $\mathcal{L}_{\prec} \leftarrow$  predecessors  $L_{\prec}$  of all  $L$  in  $\mathcal{L}$ , such
         that  $L_{\prec} \notin \mathcal{L}$ 
9:        $\mathcal{AL}_G \leftarrow \mathcal{AL}_G \cup \{\mathcal{L}_{\mathcal{I}} \cup \mathcal{L} \cup \mathcal{L}_{\prec}\}$ 
10:    end for
11:     $\Lambda_G(G) \leftarrow \mathcal{AL}_G$   $\triangleright$  Achieved landmarks of  $G$ .
12:  end for
13:  return  $\Lambda_G$ 
14: end function

```




Figure 1: BLOCKS-WORLD example.

represents an initial configuration of stackable blocks, while the set of candidate goals is composed by the following stackable words: RED, BED, and SAD. As observed actions, assume we have: (unstack E A) and (stack E D). For this example, we take the candidate goal RED to show how the algorithm computes achieved landmarks. Thus, from these observed actions, the candidate goal RED, and the set of fact landmarks of this candidate goal (Figure 2), our algorithm computes that the following fact landmarks have been achieved:

- $\mathcal{AL}_{\text{RED}} = \{[(\text{clear R})], [(\text{on E D})], [(\text{clear R}) (\text{ontable R}) (\text{handempty})], [(\text{on E A}) (\text{clear E}) (\text{handempty})], [(\text{clear D}) (\text{holding E})], [(\text{on D B}) (\text{clear D}) (\text{handempty})]\}$

In the preconditions of (unstack E A) the algorithm computes [(on E A) (clear E) (handempty)]. Subsequently, in the preconditions and effects of (stack E D) the algorithm computes [(clear D) (holding E)] and [(on E D)], while it computes the other achieved landmarks for the word RED from the initial state. Figure 2 shows the set of achieved landmarks for the word RED in gray. Listing 1 shows in bold the set of achieved landmarks that our algorithm computes for the set of candidate goals in Figure 1.




Figure 2: Ordered fact landmarks extracted for the stacked blocks RED. Fact landmarks that must be true together are represented by connected boxes. Connected boxes in grey represent achieved fact landmarks. Edges represent prerequisites between landmarks.

4 Landmark-Based Goal Completion Heuristic

We now describe a goal recognition heuristic which selects a goal from a set of possible candidates based on the number of landmarks that have been detected, and are required to achieve that goal (computed using COMPUTEACHIEVEDLANDMARKS). We note that a candidate goal is composed of sub-goals comprised of the atomic facts that are part of a conjunction of facts.

Our heuristic operates by aggregating the percentage of completion of each sub-goal into an overall percentage of completion for all facts of a candidate goal. This heuristic, denoted as h_{gc} , is formally defined 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 in \mathcal{AL}_G , and \mathcal{L}_g represents the number of necessary landmarks to achieve every sub-goal g of G in \mathcal{L}_G .

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

Thus, heuristic $h_{gc}(G)$ 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., $\sum_{g \in G} \frac{|\mathcal{AL}_g \cap \mathcal{AL}_G|}{|\mathcal{L}_g \cap \mathcal{L}_G|}$, and the size $|G|$ of the set of sub-goals, that is, the number of sub-goals in G . Algorithm 2 describes how to recognize goals using the h_{gc} heuristic and takes as input a goal recognition problem, as well as a threshold value θ . The θ threshold gives us flexibility to avoid eliminating candidate goals whose the percentage of goal completion are close to the highest completion value. In Line 2, the algorithm uses the EXTRACTLANDMARKS function to extract landmarks for all candidate goals. From the initial state \mathcal{I} , the observations O , and the extracted landmarks \mathcal{L}_G , in Line 3, our algorithm computes the set of achieved landmarks Λ_G for every candidate goal using Algorithm 1. Finally, the algorithm uses the heuristic h_{gc} to estimate goal completion for every candidate G in \mathcal{G} , and as output (Line 5), the algorithm returns those candidate goals with the highest estimated value within the threshold θ .

As an example of how heuristic h_{gc} estimates goal completion, recall the BLOCKS-WORLD example from Figure 1. Consider that among these candidate goals the hid-

Algorithm 2 Recognize goals/plans using the heuristic h_{gc} .

Input: Ξ planning domain definition, \mathcal{I} initial state, \mathcal{G} set of candidate goals, O observations, and θ threshold.

Output: Recognized goal(s).

```

1: function RECOGNIZE( $\Xi, \mathcal{I}, \mathcal{G}, O, \theta$ )
2:    $\mathcal{L}_{\mathcal{G}} \leftarrow$  EXTRACTLANDMARKS( $\Xi, \mathcal{I}, \mathcal{G}$ )
3:    $\Lambda_{\mathcal{G}} \leftarrow$  COMPUTEACHIEVEDLANDMARKS( $\mathcal{I}, \mathcal{G}, O, \mathcal{L}_{\mathcal{G}}$ )
4:    $maxh \leftarrow \max_{G' \in \mathcal{G}} h_{gc}(G', \Lambda_{\mathcal{G}}(G'), \mathcal{L}_{\mathcal{G}}(G'))$ 
5:   return all  $G$  s.t  $G \in \mathcal{G}$  and
       $h_{gc}(G, \Lambda_{\mathcal{G}}(G), \mathcal{L}_{\mathcal{G}}(G)) \geq (maxh - \theta)$ 
6: end function

```

den correct goal is RED, and we observe the following partial sequence of actions that achieve RED: (unstack E A) and (stack E D). Thus, based on the achieved landmarks $\mathcal{AL}_{\text{RED}}$ computed using Algorithm 1 (Figure 2), our heuristic h_{gc} estimates that the percentage of completion for the goal RED is 0.66: (clear R) = $1/1 +$ (on E D) = $3/3 +$ (on R E) = $1/3 +$ (ontable D) = $1/3$, and hence, $2.66/4 = 0.66$. For the words BED and SAD our heuristic h_{gc} estimates respectively, 0.54 and 0.58.

5 Landmark-Based Uniqueness Heuristic

Many goal recognition problems containing multiple ambiguous candidate goals have these goals sharing common fact landmarks. Clearly, landmarks that are common to multiple goals are less useful for recognizing a goal than landmarks that exist for only a single goal. As a consequence, computing how unique (and thus informative) each landmark is can help disambiguate similar goals for a set of candidate goals. We develop a second heuristic building on this intuition; to construct this heuristic, we introduce the concept of *landmark uniqueness*, which is the inverse frequency of a landmark among the landmarks found in a set of candidate goals. For example, consider a landmark L that occurs only for a single goal within a set of candidate goals; the uniqueness value for such a landmark is intuitively the maximum value of 1. Equation 2 formalizes this intuition, describing how the *landmark uniqueness value* is computed for a landmark L and a set of landmarks for goals $\mathcal{L}_{\mathcal{G}}$.

Using this uniqueness value, we estimate which candidate goal is the intended one by summing the uniqueness values of the landmarks achieved in the observations. Unlike our previous heuristic, which estimates progress towards goal completion by analyzing sub-goals and their achieved landmarks, the landmark-based uniqueness heuristic estimates the goal completion of a candidate goal G by calculating the ratio between the sum of the uniqueness value of the achieved landmarks of G and the sum of the uniqueness value of all landmarks of G . Effectively, this algorithm weighs the completion value by the informational value of a landmark so that unique landmarks have the highest weight. To estimate goal completion using the landmark uniqueness value, we must calculate the uniqueness value for every extracted landmark in the set of landmarks of the candidate goals. Using Equation 2, we compute the landmark uniqueness value of every landmark L of $\mathcal{L}_{\mathcal{G}}$ and store it into Υ_{uv} .

This heuristic is denoted as h_{uniq} and formally defined in Equation 3.

$$L_{Uniq}(L, \mathcal{L}_{\mathcal{G}}) = \left(\frac{1}{\sum_{L \in \mathcal{L}_{\mathcal{G}}} |\{L | L \in \mathcal{L}\}|} \right) \quad (2)$$

$$h_{uniq}(G, \mathcal{AL}_{\mathcal{G}}, \mathcal{L}_{\mathcal{G}}, \Upsilon_{uv}) = \left(\frac{\sum_{\mathcal{AL} \in \mathcal{AL}_{\mathcal{G}}} \Upsilon_{uv}(\mathcal{AL})}{\sum_{L \in \mathcal{L}_{\mathcal{G}}} \Upsilon_{uv}(L)} \right) \quad (3)$$

Algorithm 3 formalizes a goal recognition function that uses the h_{uniq} heuristic. This algorithm takes as input the same parameters as the previous approach: a goal recognition problem and a threshold θ . Like Algorithm 1, this algorithm extracts the set of landmarks for all candidate goals from the initial state \mathcal{I} , stores them in $\mathcal{L}_{\mathcal{G}}$ (Line 2), and computes the set of achieved landmarks based on the observations, storing these in $\Lambda_{\mathcal{G}}$. Unlike Algorithm 2, in Line 6 this algorithm computes the landmark uniqueness value for every landmark L in $\mathcal{L}_{\mathcal{G}}$ and stores it into Υ_{uv} . Finally, using these computed structures, the algorithm recognizes which candidate goal is being pursued from observations using the heuristic h_{uniq} , returning those candidate goals with the highest estimated value within the θ threshold.

Algorithm 3 Recognize goals/plans using the heuristic h_{uniq} .

Input: Ξ planning domain definition, \mathcal{I} initial state, \mathcal{G} set of candidate goals, O observations, and θ threshold.

Output: Recognized goal(s).

```

1: function RECOGNIZE( $\Xi, \mathcal{I}, \mathcal{G}, O, \theta$ )
2:    $\mathcal{L}_{\mathcal{G}} \leftarrow$  EXTRACTLANDMARKS( $\Xi, \mathcal{I}, \mathcal{G}$ )
3:    $\Lambda_{\mathcal{G}} \leftarrow$  COMPUTEACHIEVEDLANDMARKS( $\mathcal{I}, \mathcal{G}, O, \mathcal{L}_{\mathcal{G}}$ )
4:    $\Upsilon_{uv} \leftarrow \langle \rangle$   $\triangleright$  Map of landmarks to their uniqueness value.
5:   for each fact landmark  $L$  in  $\mathcal{L}_{\mathcal{G}}$  do
6:      $\Upsilon_{uv}(L) \leftarrow L_{Uniq}(L, \mathcal{L}_{\mathcal{G}})$ 
7:   end for
8:    $maxh \leftarrow \max_{G' \in \mathcal{G}} h_{uniq}(G', \Lambda_{\mathcal{G}}(G'), \mathcal{L}_{\mathcal{G}}(G'), \Upsilon_{uv})$ 
9:   return all  $G$  s.t  $G \in \mathcal{G}$  and
       $h_{uniq}(G, \Lambda_{\mathcal{G}}(G), \mathcal{L}_{\mathcal{G}}(G), \Upsilon_{uv}) \geq (maxh - \theta)$ 
10: end function

```

As an example of how our heuristic h_{uniq} estimates goal completion using landmark uniqueness values, recall the BLOCKS-WORLD example from Figure 1. As previously, the correct hidden goal is RED and we observe the following actions: (unstack E A) and (stack E D). Listing 1 shows the set of extracted fact landmarks for the candidate goals in the BLOCKS-WORLD example and their respective uniqueness value. Based on the set of achieved landmarks (shown in bold in Listing 1), our heuristic h_{uniq} estimates the following percentage for each candidate goal: $h_{uniq}(\text{RED}) = 3.66/6.33 = 0.57$; $h_{uniq}(\text{BED}) = 1.66/6.33 = 0.27$; and $h_{uniq}(\text{SAD}) = 3.66/8.33 = 0.43$. In this case, Algorithm 3 correctly estimates RED to be the intended goal since it has the highest heuristic value.

```

- (and (clear B) (on B E) (on E D) (ontable D)) = 6.33
  [(on E D)] = 0.5, [(clear D) (holding E)] = 0.5,
  [(on E A) (clear E) (handempty)] = 0.33, [(ontable D)] = 0.33,
  [(on D B) (clear D) (handempty)] = 0.33, [(holding D)] = 0.33,
  [(clear B)] = 1.0, [(clear E) (holding B)] = 1.0,
  [(on B E)] = 1.0, [(clear B) (ontable B) (handempty)] = 1.0,

- (and (clear S) (on S A) (on A D) (ontable D)) = 8.33
  [(clear S)] = 1.0, [(on A D)] = 1.0, [(on S A)] = 1.0,
  [(clear A) (ontable A) (handempty)] = 1.0, [(ontable D)] = 0.33,
  [(clear S) (ontable S) (handempty)] = 1.0, [(holding D)] = 0.33,
  [(on E A) (clear E) (handempty)] = 0.33,
  [(on D B) (clear D) (handempty)] = 0.33,
  [(clear A) (holding S)] = 1.0, [(clear D) (holding A)] = 1.0

- (and (clear R) (on R E) (on E D) (ontable D)) = 6.33
  [(clear R)] = 1.0, [(clear R) (ontable R) (handempty)] = 1.0,
  [(clear D) (holding E)] = 0.5, [(on E D)] = 0.5,
  [(on E A) (clear E) (handempty)] = 0.33, [(ontable D)] = 0.33,
  [(on D B) (clear D) (handempty)] = 0.33, [(holding D)] = 0.33,
  [(on R E)] = 1.0, [(clear E) (holding R)] = 1.0

```

Listing 1: Extracted fact landmarks for the BLOCKS-WORLD example and their respective uniqueness value.

6 Experiments and Evaluation

We empirically evaluate our approach using datasets created using 15 domains from the planning literature¹. Six of these datasets are also used in the evaluation of Ramírez and Geffner (2009; 2010). In addition, we generate new datasets from the remaining nine planning domains (using open-source planners, such as FAST-DOWNWARD, FAST-FORWARD, and LAMA), each of which is based on planning problems containing both optimal and sub-optimal plans of various sizes, including large problems to test the scalability of the approach. These domains contain hundreds of non-trivial goal/plan recognition problems, *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 O . An observation sequence contains actions that represent an optimal plan or sub-optimal plan that achieves a hidden goal G , and this observation sequence can be partial or full. While the latter capture all actions used to achieve hidden goal G (*i.e.*, with 100% observed actions), partial observation sequences represent plans for G where 10%, 30%, 50% or 70% of actions are observed.

Our evaluation uses two metrics, namely accuracy of goal recognition and goal recognition speed. In this paper, we measure accuracy by counting the number of times an algorithm includes the hidden goal in the set of candidate goals generated by the algorithm. Note that in many domains, the number of candidate goals returned by all algorithms may be more than one. In the case of Ramírez and Geffner (2009), this may occur when goals have the same distance from their estimated state. Alternatively, for our heuristics this may occur when there are ties between the heuristic value of candidate goals within the threshold margin. We compare our heuristic approaches to the fastest and most accurate approach from Ramírez and Geffner (2009). To visualize the comparative performance of the multiple approaches we adapt the notation of the Receiver Operating Characteristic (ROC) curve, and, rather than plotting a single

¹<http://ipc.icaps-conference.org>




Figure 3: Comparative performance in ROC space.

curve per domain, we aggregate multiple domains and plot these results in ROC space. ROC space graphically shows the performance of a binary classifier system by evaluating the true positive rate against the false positive rate at various threshold settings. More specifically, we use ROC space to compare not only true positive predictions (*i.e.*, accuracy), but also to compare the false positive ratio of the evaluated goal/plan recognition approaches.

Table 1 compares the results of our heuristics h_{gc} and h_{uniq} against the Ramírez and Geffner’s approach (2009) (R&G in the table), by showing the total number of goal recognition problems used under each domain name. Each row of this table expresses averages for the number of candidate goals $|\mathcal{G}|$; the average number of extracted landmarks $|\mathcal{L}|$; the percentage of the plan that is actually observed; the average number of observations per problem $|O|$; and for each approach, the time in seconds to recognize the goal given the observations; and the accuracy with which the approaches correctly infer the goal. From this table, we can see that our landmark-based heuristics are both faster and more accurate than R&G (2009), and as we increase the threshold θ , our heuristics quickly surpass the state-of-the-art for all domains tested by yielding near perfect accuracy. We note that when measuring recognition time, we also include the time to extract the set of landmarks, so that landmark extraction is performed online, *i.e.*, during the process of goal recognition.

Figure 3 shows the trade-off between true positive results and false positive results for the tested goal/plan recognition approaches. In the ROC curve, the diagonal line represents a random guess to recognize a goal from observations, *i.e.*, points above the diagonal represent good classification results (better than random), whereas points below the line represent poor results (worse than random). The best possible (perfect) prediction for recognizing goals must be a point in the upper left corner (*i.e.*, coordinate $x = 0$ and $y = 100$)

Domain	G	L	%Obs	O	h_{gc}		h_{uniq}		R&G			
					Time		Accuracy		Time		Accuracy	
					$\theta (0/10/20/30)$	$\theta (0/10/20/30)$	$\theta (0/10/20/30)$	$\theta (0/10/20/30)$	$\theta (0/10/20/30)$	$\theta (0/10/20/30)$	Time	Accuracy
BLOCKS-WORLD (855)	20	15.6	10	1.1	0.097/0.100/0.105/0.111	36.9%/97.4%/70.2%/89.2%	0.096/0.098/0.101/0.105	31.9%/37.4%/59.4%/70.4%	1.656	83.8%		
				30	2.9	0.107/0.109/0.113/0.122	54.4%/61.5%/86.1%/97.4%	0.103/0.106/0.107/0.111	56.6%/62.5%/82.5%/93.3%	1.735	90.0%	
				50	4.2	0.113/0.113/0.120/0.127	62.5%/82.5%/98.3%/100.0%	0.108/0.109/0.112/0.115	70.4%/86.1%/93.3%/98.9%	1.836	97.2%	
				70	6.5	0.138/0.139/0.141/0.148	83.5%/94.8%/100.0%/100.0%	0.118/0.121/0.125/0.129	76.6%/86.1%/96.9%/100.0%	2.056	98.8%	
				100	8.5	0.163/0.166/0.172/0.185	100.0%/100.0%/100.0%/100.0%	0.136/0.142/0.146/0.151	100.0%/100.0%/100.0%/100.0%	2.378	100.0%	

Table 1: Comparison of accuracy and recognition time against Ramirez and Geffner’s approach (2009).

in the ROC space. The closer a goal recognition approach (point) is to the upper left corner, the better it is for recognizing goals and plans. To compare our recognition results against R&G in the ROC curve, we select the results of our heuristics using the $\theta = 30\%$ threshold. For each approach, we plot its recognition results for all domains into a cloud of points, which represents (in general) how well each approach recognizes the correct goal from observations. The points in ROC space show that our heuristics are not only competitive with R&G for all variations of observability, but also surpass R&G in a substantial number of domains.

Finally, we compare the time that each approach takes to recognize the hidden goal for different sizes of the observation sequence. We illustrate runtime in Figure 4, which summarizes, for the three evaluated approaches, the runtime (Time columns in Table 1) as a function of the average size of the observations (|O| column in Table 1). Curves in the graph were generated by averaging the runtime when observation sizes were the same and smoothing over the resulting points. The graph shows the scalability of the three evaluated approaches. Our heuristics never take more than two seconds to compute the hidden goal in the set of can-

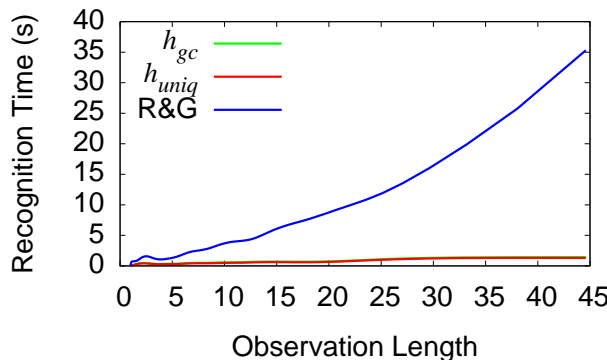


Figure 4: Recognition time comparison.

didate goals, while R&G’s approach appears to grow super-linearly. As shown for the DOCK-WORKER-ROBOTS and SOKOBAN domains, larger plan lengths also lead R&G’s approach to lose accuracy.

7 Related Work

Pattison and Long (2010) propose AUTOGRAPH (Automatic Goal Recognition with A Planning Heuristic), a probabilistic heuristic-based goal recognition over planning domains. AUTOGRAPH uses heuristic estimation and domain analysis to determine which goals an agent is pursuing. Ramírez and Geffner (2009) propose planning approaches for plan recognition, and instead of using plan-libraries, they model the problem as a planning domain theory with respect to a known set of goals. Their work uses a heuristic, an optimal and modified sub-optimal planner to determine the distance to every goal in a set of goals after an observation. In this paper, we compare their most accurate approach directly with ours. Follow-up work (2010) proposes a probabilistic plan recognition approach using off-the-shelf planners. E.-Martín *et al.* (2015) propose a planning-based goal recognition approach that propagates cost and interaction information in a plan graph, and uses this information to estimate goal probabilities over the set of candidate goals. Sohrabi *et al.* (2016) propose a probabilistic plan recognition approach that deals with unreliable observations (*i.e.*, noisy or missing observations), and recognizes both goals and plans. Unlike these last three approaches, which provide a probabilistic interpretation of the recognition problem, we do not deal with probabilities yet. Keren *et al.* (2014) develop an alternate view of the goal recognition problem, and rather than developing new goal recognition algorithms, they develop an approach that changes the domain definition to facilitate the goal recognition process. Their work could potentially be used alongside our techniques. Most recently, Pereira and Meneguzzi (2016) describe a landmark-based plan recognition approach, which was not as accurate and scalable as the one shown in this paper and did not use the notion of landmark uniqueness to improve recognition.

8 Conclusions

We developed a novel goal recognition approach that employs planning landmarks to compute progress towards a set of candidate goals from observations. This paper provides two contributions. First, a goal recognition algorithm that can leverage any landmark extraction technique to efficiently infer hidden goals; and second, we show that our heuristics surpass the state-of-the-art across a variety of planning domains. Both heuristics yield high accuracy consistent with the state-of-the-art, while taking significantly less time than earlier approaches to recognize the hidden goal. Importantly, our approach scales much better for larger plan lengths while maintaining accuracy, unlike the current state-of-the-art in the field (Ramírez and Geffner 2009).

We intend to explore multiple avenues for future work. First, we intend to explore other landmark extraction algorithms to obtain additional information from planning domains, such as temporal landmarks (Karpas *et al.* 2015). Second, we aim to associate a probabilistic interpretation to the observed landmarks and compare these to recent work, such as (E.-Martín, R.-Moreno, and Smith 2015) and (Sohrabi, Riabov, and Udrea 2016).

References

- Avrahami-Zilberbrand, D., and Kaminka, G. A. 2005. Fast and Complete Symbolic Plan Recognition. In *Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence (IJCAI 2005)*.
- Bylander, T. 1994. The Computational Complexity of Propositional STRIPS Planning. *Journal of Artificial Intelligence Research (JAIR)* 69:165–204.
- E.-Martín, Y.; R.-Moreno, M. D.; and Smith, D. E. 2015. A Fast Goal Recognition Technique Based on Interaction Estimates. In *Proceedings of the 24th International Conference on Artificial Intelligence (IJCAI 2015)*.
- Geib, C. W., and Goldman, R. P. 2009. A Probabilistic Plan Recognition Algorithm Based on Plan Tree Grammars. *Artificial Intelligence* 173(11):1101–1132.
- Ghallab, M.; Nau, D. S.; and Traverso, P. 2004. *Automated Planning - Theory and Practice*. Elsevier.
- Hoffmann, J.; Porteous, J.; and Sebastia, L. 2004. Ordered Landmarks in Planning. *Journal of Artificial Intelligence Research (JAIR)* 22(1):215–278.
- Karpas, E.; Wang, D.; Williams, B. C.; and Haslum, P. 2015. Temporal Landmarks: What Must Happen, and When. In *Proceedings of the Twenty-Fifth International Conference on Automated Planning and Scheduling (ICAPS 2015)*.
- Keren, S.; Gal, A.; and Karpas, E. 2014. Goal Recognition Design. In *Proceedings of the 24th International Conference on Automated Planning and Scheduling (ICAPS 2014)*.
- Mirsky, R.; Stern, R.; Gal, Y. K.; and Kalech, M. 2016. Sequential Plan Recognition. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI 2016)*.
- Pattison, D., and Long, D. 2010. Domain Independent Goal Recognition. In *Proceedings of the Fifth European Starting AI Researcher Symposium (STAIRS)*.
- Pereira, R. F., and Meneguzzi, F. 2016. Landmark-Based Plan Recognition. In *Proceedings of the 22nd European Conference on Artificial Intelligence (ECAI 2016)*.
- Ramírez, M., and Geffner, H. 2010. Probabilistic Plan Recognition Using Off-the-Shelf Classical Planners. In *Proceedings of the Conference of the American Association of Artificial Intelligence (AAAI 2010)*.
- Ramírez, M., and Geffner, H. 2009. Plan Recognition as Planning. In *Proceedings of the Twenty-First International Joint Conference on Artificial Intelligence (IJCAI 2009)*.
- Richter, S., and Westphal, M. 2010. The LAMA Planner: Guiding Cost-based Anytime Planning with Landmarks. *Journal of Artificial Intelligence Research (JAIR)* 39(1):127–177.
- Sohrabi, S.; Riabov, A. V.; and Udrea, O. 2016. Plan Recognition as Planning Revisited. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI 2016)*.
- Sukthankar, G.; Goldman, R. P.; Geib, C.; Pynadath, D. V.; and Bui, H. H. 2014. *Plan, Activity, and Intent Recognition: Theory and Practice*. Elsevier.