Landmark-Based Heuristics for Goal Recognition

Ramon Fraga Pereira†, 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

February, 2017

Introduction

- **Goal Recognition** is the task of recognizing agents' goals from either partial or full observations;
- In this work, we use a **planning domain definition** to represent agent behavior and environment properties;
- Our main contribution is twofold:
  - We **obviate the need to execute a planner multiple times** for recognizing goals; and
  - We develop novel goal recognition heuristics that **use planning landmarks**.
- We show that **our approaches are more accurate and orders of magnitude faster** than Ramírez and Geffner's approach.

Background: Goal Recognition Problem

Definition (Goal Recognition Problem)

- Domain Definition (Properties and Actions) $\Xi = \langle \Sigma, \mathcal{A} \rangle$;
- Initial State $\mathcal{I}$;
- A Set of Candidate Goals $\mathcal{G}$ (with a hidden goal G); and
- Sequence of Observations (*i.e.*, Observed Actions) O.
  - This observation sequence can be either **partial** or **full**.

- 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.

Background: Planning and Landmarks

Definition (Planning)

A planning instance is represented by a triple $\Pi = \langle \Xi, \mathcal{I}, G \rangle$*, in which:*

- $\Xi = \langle \Sigma, \mathcal{A} \rangle$ *is the* **domain definition***, and consists of a finite set of* **facts** Σ *and a finite set of* **actions** $\mathcal{A}$ *(action costs = 1);*
- $\mathcal{I}$ *and* G *represent the* **planning problem***, in which* $\mathcal{I} \subseteq \Sigma$ *is the* **initial state***, and* $G \subseteq \Sigma$ *is the* **goal state***.*

Definition (Landmarks)

Given a planning instance $\Pi = \langle \Xi, \mathcal{I}, G \rangle$*, a* **fact** *(or* **action***)* L *is a landmark in* Π *iff* L *must be* **satisfied** *(or* **executed***) at some point along all valid plans that achieve* G *from* $\mathcal{I}$*.*

- To extract landmarks and their ordering, we use an algorithm developed by Hoffman *et al.* (Ordered Landmarks in Planning. JAIR, 2004).

Computing Achieved Landmarks

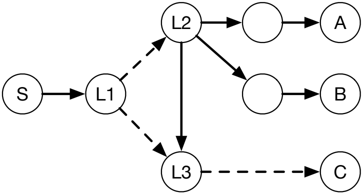

- Our heuristics require identifying which fact landmarks have been achieved during the observed plan execution for every candidate goal $G \in \mathcal{G}$;
- For every candidate goal $G \in \mathcal{G}$:
  - Extract *ordered* landmarks for G;
  - Analyze achieved landmarks of G in preconditions and effects of every observed action $o \in O$;
  - As we deal with partial observability, some observed actions may be missing, thus whenever we identify a fact landmark, we also infer that its predecessors must have been achieved;

Landmark-Based Goal Completion Heuristic

- Our first heuristic, called h_{gc}, operates by aggregating the percentage of completion of each sub-goal into an overall percentage of completion for all facts of a candidate goal;

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

where:

- $\mathcal{AL}_G$ achieved landmarks for goals in G
- $\mathcal{L}_G$ all landmarks for goals in G

Landmark-Based Goal Completion Heuristic: Algorithm

- Our approach allows the use of a threshold θ, giving us **flexibility to avoid eliminating candidate goals** whose the percentage of goal completion are close to the highest completion value;

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: $\quad \mathcal{L}_{\mathcal{G}} \leftarrow$ EXTRACTLANDMARKS($\Xi, \mathcal{I}, \mathcal{G}$)
3: $\quad \Lambda_{\mathcal{G}} \leftarrow$ COMPUTEACHIEVEDLANDMARKS($\mathcal{I}, \mathcal{G}, O, \mathcal{L}_{\mathcal{G}}$)
4: $\quad maxh \leftarrow \max_{G' \in \mathcal{G}} h_{gc}(G', \Lambda_{\mathcal{G}}(G'), \mathcal{L}_{\mathcal{G}}(G'))$
5: $\quad$ **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**

Landmark-Based Uniqueness Heuristic (1 of 2)

- To develop our second 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, *i.e.*, how unique (and thus informative) each landmark is among all landmarks;

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

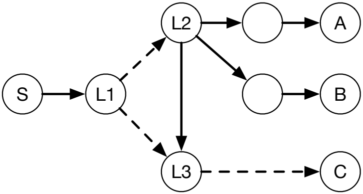

$L_{Uniq}(L2) = 1/2$
$L_{Uniq}(L1) = 1/3$
$L_{Uniq}(L3) = 1$

Landmark-Based Uniqueness Heuristic (2 of 2)

- Our second heuristic, called h_{uniq}, 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;

$$h_{uniq}(G, \mathcal{AL}_G, \mathcal{L}_G, \Upsilon_{uv}) = \left(\frac{\sum\limits_{\mathcal{A}_L \in \mathcal{AL}_G} \Upsilon_{uv}(\mathcal{A}_L)}{\sum\limits_{L \in \mathcal{L}_G} \Upsilon_{uv}(L)} \right) \tag{3}$$

where:

- Υ_{uv} is a table of uniqueness values
- $\mathcal{AL}_G$ achieved landmarks for goals in G
- $\mathcal{L}_G$ all landmarks for goals in G

Landmark-Based Uniqueness Heuristic: Algorithm

- Our second heuristic is called h_{uniq};

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).*

function RECOGNIZE($\Xi, \mathcal{I}, \mathcal{G}, O, \theta$)
    $\mathcal{L}_{\mathcal{G}} \leftarrow$ EXTRACTLANDMARKS($\Xi, \mathcal{I}, \mathcal{G}$)
    $\Lambda_{\mathcal{G}} \leftarrow$ COMPUTEACHIEVEDLANDMARKS($\mathcal{I}, \mathcal{G}, O, \mathcal{L}_{\mathcal{G}}$)
    $\Upsilon_{uv} \leftarrow \langle\rangle$ ▷ *Map of landmarks to their uniqueness value*.
    for each fact landmark L in $\mathcal{L}_{\mathcal{G}}$ **do**
        $\Upsilon_{uv}(L) \leftarrow L_{Uniq}(L, \mathcal{L}_{\mathcal{G}})$
    end for
    $maxh \leftarrow \max_{G' \in \mathcal{G}} h_{uniq}(G', \Lambda_{\mathcal{G}}(G'), \mathcal{L}_{\mathcal{G}}(G'), \Upsilon_{uv})$
    return all G s.t $G \in \mathcal{G}$ and
        $h_{uniq}(G, \Lambda_{\mathcal{G}}(G), \mathcal{L}_{\mathcal{G}}(G), \Upsilon_{uv}) \geq (maxh - \theta)$
end function

Experiments and Evaluation

- We evaluate our heuristics over datasets with 15 planning domains (6 of these domains are provided by Ramírez and Geffner), such as:
  - Blocks-World, Campus, Depots, Driver-Log, Dock-Worker-Robots, Easy-IPC-Grid, Ferry, Intrusion-Detection, Kitchen, Logistics, Miconic, Rovers, Satellite, Sokoban, and Zeno-Travel;
- These datasets contain hundreds of goal recognition problems, varying the observability (10%, 30%, 50%, 70%, and 100%);
- We compared our heuristics against the heuristic-based approach of Ramírez and Geffner (Plan Recognition as Planning. IJCAI, 2009), which is their fastest and most accurate approach;

- We select the results of our heuristics using threshold $\theta = 30\%$;
- To evaluate our heuristics agains Ramírez and Geffner's approach, we use the ROC space, which shows the trade-off between True Positive results and False Positive results;
- We aggregate multiple domains and plot these goal recognition results in ROC space, we aim to show how good each approach is in general (all domains);

Experiments and Evaluation - ROC Space (2 of 2)

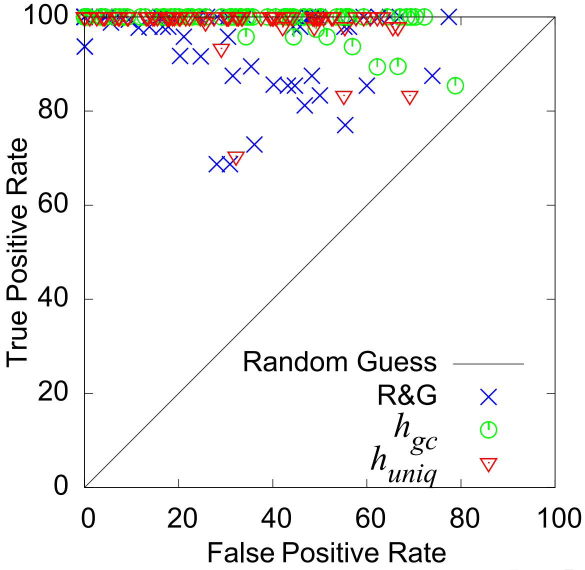

Experiments and Evaluation - Recognition Time

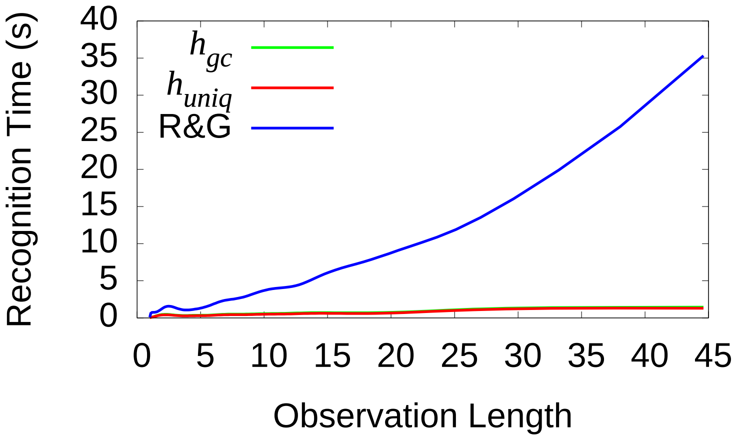

Conclusions

- **Contribution:**
  - Use planning landmarks for goal recognition; and
  - Obviate the need to run a planner during goal recognition, resulting in much faster and highly accurate recognition.
- **Limitations:**
  - Sensitive to the presence of landmarks
  - Low accuracy with very few observations, *i.e.*, 10% of observability;
- **Future Work:**
  - Use different landmark extraction algorithms;
  - Use goal ordering techniques;
  - Derive a probabilistic interpretation for the landmarks; and
  - Apply our landmark-based heuristics to continuous and temporal domains.

Thank you!
Questions?