An LP-Based Approach for Goal Recognition as Planning

Luísa Santos*, Felipe Meneguzzi†, Ramon Fraga Pereira‡, André Grahl Pereira*

✉ †felipe.meneguzzi@pucrs.br, *lrasantos,agpereira@inf.ufrgs.br, ‡rpereira@diag.uniroma1.it

What is Goal Recognition?

In a nutshell

- **Goal Recognition** is the task of recognizing agents' goal that explains a sequence of observations of its actions;
  - Related to plan recognition, i.e. recognizing a *top-level* action
  - A specific form of the problem of abduction

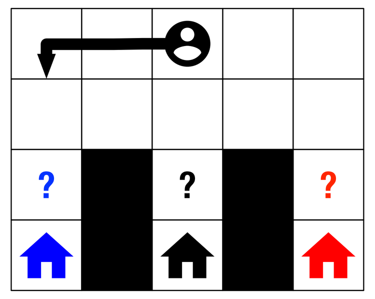

Automated Planning and Goal Recognition

Definition 1 (Planning). *A planning instance is represented by a triple* $\Pi = \langle \mathcal{V}, \mathcal{O}, s_0, s^*, \text{cost}\rangle$, *in which:*

- $\mathcal{V}$ *is a finite set of variables, each* $v \in \mathcal{V}$ *with domain* $D(v)$
- $\mathcal{O}$ *is a finite set of operators, where* $o \in \mathcal{O}$ *are tuples* $o = \langle \text{pre}(o), \text{post}(o)\rangle$ *each of which has cost* $\text{cost}(o)$
- s_0 *is the* ***initial state***
- s^* *is the* ***goal state***

Definition 2 (Goal Recognition Problem). *A goal recognition problem is a tuple* $P = \langle \Pi_\mathrm{P}, \Gamma, \Omega\rangle$, *where:*

- Π_P *is a planning task without a goal condition;*
- Γ *is a set of goal candidates; and*
- Ω *is a sequence* $\langle \vec{o}_1, \ldots \vec{o}_n\rangle$ *of observations, with each* $\vec{o}_i \in \mathcal{O}$
- Many solution concepts here (check the paper)
- Caveat: we may have other representations for the observations

A Running Example

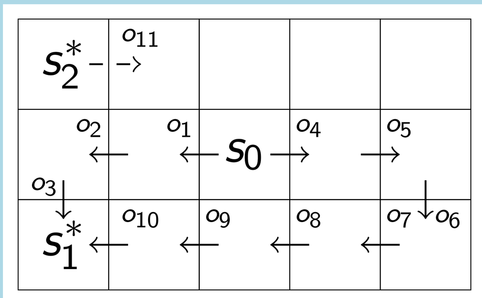

Goal Recognition task with:

- $\Gamma = \{s_1^*, s_2^*\}$
- Reference goal s_1^*

A number of observations

- $\Omega_1 = \langle \vec{o}_1\rangle$ is an optimal observation sequence from optimal plan $\pi_1 = \langle o_1, o_2, o_3\rangle$
- $\Omega_2 = \langle \vec{o}_5, \vec{o}_7, \vec{o}_9\rangle$ and $\Omega_3 = \langle \vec{o}_4, \ldots, \vec{o}_{10}\rangle$ are suboptimal observation sequences from suboptimal plan $\pi_2 = \langle o_4, \ldots, o_{10}\rangle$,
- $\Omega_4 = \langle \vec{o}_4, \ldots, \vec{o}_{10}, \vec{o}_{11}\rangle$ is a suboptimal and noisy observation sequence (with added $\vec{o}_{11}$)

Reference Solutions

- *Goal recognition task* $\langle \Pi_\mathrm{P}, \Gamma, \Omega\rangle$
- Π is a planning task with the reference goal $s^* \in \Gamma$
- π^* is the optimal plan for Π, and π plan for Π that generates Ω
- $h^*_\Omega(s_0, s_i^*)$ is the cost of an optimal plan for Π that complies with Ω, $h^*(s_0, s_i^*)$ is the cost of an optimal plan for Π, both with $s_i^* \in \Gamma$

The reference solution is

$$\Gamma^* = \{s_i^* \in \Gamma \mid \frac{h^*_\Omega(s_0, s_i^*)}{h^*(s_0, s_i^*)} \le \frac{\text{cost}(\pi)}{\text{cost}(\pi^*)} \wedge h^*_\Omega(s_0, s_i^*) \neq \infty\}$$

The *reference solution set* includes goal candidates that have plans as sub-optimal as or less than the plan that generated the observations for the reference goal.

Example Reference Solution

$\Gamma_i^* = \{s_1^*\}$, for $\Omega_1, \Omega_2, \Omega_3, \Omega_4$ since

- $h^*_{\Omega_4}(s_0, s_1^*) = 7$, $h^*_{\Omega_4}(s_0, s_2^*) = 9$
- $\text{cost}(\pi_2)/\text{cost}(\pi^*) = 7/3$

Using LP-Constraints for Goal Recognition

Satisfying IP/LP heuristic

The *satisfying* integer program IP^C_Ω for a set of operator-counting constraints C, a set of *observation-counting constraints*, and sequence of observations Ω for state s is

$$\begin{aligned}
&\text{minimize} \sum_{o \in \mathcal{O}} \text{cost}(o)\mathsf{Y}_o && \text{subject to } C, \\
&\mathsf{Y}_{\vec{o}} \le occur_\Omega(o) && \text{for all } o \in \mathcal{O} && (1)\\
&\mathsf{Y}_{\vec{o}} \le \mathsf{Y}_o && \text{for all } o \in \mathcal{O} && (2)\\
&\sum_{\mathsf{Y}_{\vec{o}} \in \mathcal{Y}^\Omega} \mathsf{Y}_{\vec{o}} \ge |\,\Omega\,| && && (3)\\
&\mathsf{Y}_o, \mathsf{Y}_{\vec{o}} \in \mathbb{Z}_0^+.
\end{aligned}$$

The *satisfying* IP *heuristic* h^{IP}_Ω is the objective value of IP^C_Ω, and the *satisfying* LP *heuristic* h_Ω is the objective value of its linear relaxation. If the IP or LP is infeasible, the heuristic estimate is ∞.

Constraints for Goal Recognition

We define a new heuristic h_Ω based on the existing operator counting framework using:

- Observations Ω for a state s, where $occur_\Omega(o)$ is the # of occurrences of $o \in \Omega$
- Variables $\mathsf{Y}_{\vec{o}}$ for each $\vec{o} \in \mathcal{O}$

with additional constraints:

- $\mathsf{Y}_{\vec{o}} \le occur_\Omega(o)$, for all $o \in \mathcal{O} \rightarrow$ limits occurrences of observations
- $\mathsf{Y}_{\vec{o}} \le \mathsf{Y}_o$ for all $o \in \mathcal{O} \rightarrow$ binds Ω to operators in the OC heuristic
- $\sum_{\mathsf{Y}_{\vec{o}} \in \mathcal{Y}^\Omega} \mathsf{Y}_{\vec{o}} \ge |\,\Omega\,| \rightarrow$ ensures observations are satisfied

The h_Ω heuristic:

- Y_o acts as an upper bound for $\mathsf{Y}_{\vec{o}}$
- The only difference of h_Ω to the OC heuristic are the *observation-counting constraints*
  - $h \rightarrow$ lower bound on optimal plans
  - $h_\Omega \rightarrow$ lower bound on optimal plans that satisfy observations

Computing solutions using h_Ω

- We compute the cost difference between observation-complying Operator Counts h_Ω and the OCs lower bound on optimal plan cost h

$$\delta_{\text{min}} = \min_{s_i^* \in \Gamma\,:\,h_\Omega(s_0, s_i^*) < \infty} \{h_\Omega(s_0, s_i^*) - h(s_0, s_i^*)\}$$

- And select goals for which the observation-complying plans have the least additional cost over the optimal plan

$$\Gamma^{\mathrm{LP}} = \{s_i^* \in \Gamma \mid h_\Omega(s_0, s_i^*) - h(s_0, s_i^*) = \delta_{\text{min}}\}$$

Example solution using h_Ω

For $\Omega = \langle \vec{o}_5, \vec{o}_7, \vec{o}_9\rangle$:

- $h_\Omega(s_0, s_1^*) = 7$ and $h(s_0, s_1^*) = 3$
- $h_\Omega(s_0, s_2^*) = 9$ and $h(s_0, s_2^*) = 3$
- $\delta_{\text{min}} = 4$, so $\Gamma^{\mathrm{LP}} = \{s_1^*\}$

Dealing with Noise and Uncertainty

Dealing with Noisy Observations

- Noisy Observations $\rightarrow$ Suboptimal or Spurious
  - Unlikely to be part of an optimal plan
  - Expensive to detect
- We estimate which observations are noisy in polynomial time in the linear relaxation in a new heuristic h^ϵ_Ω
- Relax h_Ω to ignore a fraction ϵ of the observations
  - ϵ corresponds to an **error rate**
  - Satisfy at least $|\,\Omega\,| - \lfloor |\,\Omega\,| * \epsilon \rfloor$ observations
  - This results in a new solution set Γ^ϵ

$$\Gamma^\epsilon = \{s_i^* \in \Gamma \mid h^\epsilon_\Omega(s_0, s_i^*) - h(s_0, s_i^*) = \delta_{\text{min}}\}$$

Estimating Uncertainty

- Recognizing goals is hard with low observability, most existing approaches have either
  - Low accuracy while maintaining low spread
  - High accuracy while having high spread
- Our approach modulates the accuracy/spread tradeoff in response to lower observability:
  - If $|\,\Omega\,| < h_\Omega$, then at least $h_\Omega - |\,\Omega\,|$ observations missing
- We estimate the degree of observability as follows:

$$\mu = 1 + \frac{\max_{s_i^* \in \Gamma^{\mathrm{LP}}} \{h_\Omega(s_0, s_i^*)\} - |\,\Omega\,|}{\max_{s_i^* \in \Gamma^{\mathrm{LP}}} \{h_\Omega(s_0, s_i^*)\}}$$

- And use the uncertainty when selecting goals

$$\Gamma^\mu = \{s_i^* \in \Gamma \mid h_\Omega(s_0, s_i^*) - h(s_0, s_i^*) \le \delta_{\text{min}}\mu\}$$

Experiments and Conclusions

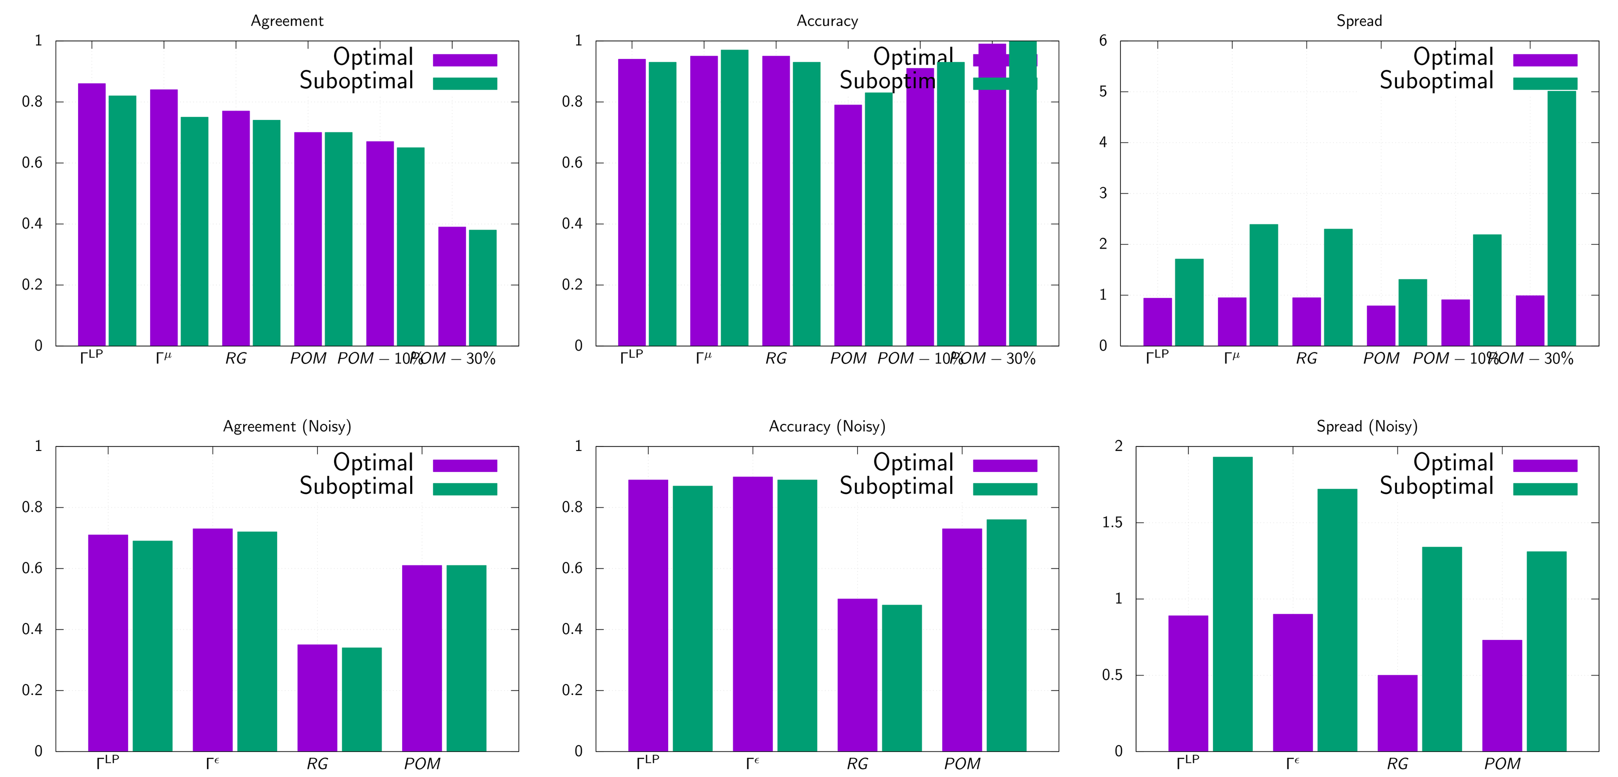

Conclusions

We developed a new class of goal recognition methods:

- Based on linear programming models with provably polynomial-time solutions
- Leverages operator counting framework

Code available at:

https://bit.ly/lp-goal-recognition