An LP-Based Approach for Goal Recognition as Planning

Felipe Rech Meneguzzi[1], **Giovanna Lazzari Miotto**[2], **Ramon Fraga Pereira**[1], **and André Grahl Pereira**[2]

[1]Pontifical Catholic University of Rio Grande do Sul (PUCRS), Brazil

[2]Federal University of Rio Grande do Sul (UFRGS), Brazil

felipe.meneguzzi@pucrs.br, ramon.pereira@edu.pucrs.br,

agpereira@inf.ufrgs.br

Abstract

Goal recognition is the task of inferring the intended goal of an agent given a sequence of observations. Advances in heuristics based on linear programming allows us to solve goal recognition tasks by encoding the declarative knowledge about such tasks resulting in two central contributions. First, we develop an approach that guarantees we select the actual hidden goal given the complete sequence of either optimal or suboptimal observations. Second, we automatically estimate the number of missing observations through a metric of uncertainty, which improves accuracy under very low observability. Experiments and evaluation show that the resulting approach is fast and dominates previous methods providing lower spread and higher accuracy on average.

1 Introduction

Goal recognition as planning (Ramírez and Geffner 2009; Ramírez and Geffner 2010) is the task of recognizing the actual goal from a set of hypotheses given a sequence of observations, an initial state, and a behavior model of the agent under observation. Approaches for goal recognition as planning have leveraged efficient planning technology and heuristic information to develop increasingly accurate and faster goal recognition approaches. Most approaches select goals based on metrics that compare the cost of optimal plans and the cost of plans constrained to comply or avoid observations. These approaches differ in how they compute or approximate these costs. While some approaches compute these costs using a planner (Ramírez and Geffner 2009; Ramírez and Geffner 2010), others approximate them using sophisticated structures from heuristic functions (E-Martín, R.-Moreno, and Smith 2015; Vered et al. 2018), or try to explicitly cope with missing and noisy observations by introducing weights in the action descriptions (Sohrabi, Riabov, and Udrea 2016). By contrast, recent work (Pereira, Oren, and Meneguzzi 2017) introduces recognition heuristics using information from the structure of the planning instances in order to recognize the actual goal from a set of goal hypothesis and observations. All of them try to balance speed

and accuracy: the more precise the information used to decide which goals should be selected, the higher the overall time.

The main challenge in goal recognition tasks is to achieve reasonable recognition time (i.e., a few seconds), high accuracy and low spread. In this paper, we address these challenges by leveraging recent advances on heuristics computed via linear programming (LP) (Pommerening et al. 2014). LP-based heuristics constitute a unifying framework for a variety of sources of information from planning tasks that provide both precise information about the plan cost to a goal, and fast computation time. This framework enables us to encode knowledge about planning and goal recognition tasks and use off-the-shelf LP solvers to combine this knowledge. This proves to be effective to build goal recognition heuristics, disambiguating between goal hypotheses.

We use the framework of LP-based heuristics from Pommerening et al. (2014) to encode knowledge about planning and goal recognition tasks and use off-the-shelf LP solvers to combine this knowledge in an optimal way. While previous approaches need to balance accuracy, speed, and spread, LP-based methods can satisfy constraints both about plans and about complying with optimal and suboptimal observations. The resulting approaches guarantee selection of the actual goal given the complete sequence of observations. In practice, such approaches lead to fast recognition time, low spread and high accuracy. Previous recognition approaches used the same method to select goals regardless of whether the number of available observation is low or high. We show that this is an undesirable property of these approaches, which leads to a higher spread and a lower accuracy. The LPs allow us to automatically estimate the number of missing observations, which, in turn, yields an approach that improves accuracy under very low observability scenarios. We empirically show that our LP-based heuristics are very effective at goal recognition, overcoming existing approaches in almost all domains in terms of accuracy while diminishing the spread of recognized goals. Such approaches are also substantially more effective for noisy settings, even without an explicit model of the noisy observations. Finally, we discuss how our approach can be further extended and how our uncertainty metric can be used to improve previous ap-

proaches.

2 Background

In this section, we review essential background on planning terminology and goal recognition as planning. Finally, we provide background on LP-based heuristics for planning, and discuss the framework we used to build our heuristics for goal recognition.

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 (τ_1, τ_2, ..., τ_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 $\pi = \langle a_1, a_2, ..., a_n \rangle$ 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 as Planning

Goal recognition is the task of recognizing agents' goals by observing their interactions in a particular 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 includes the actual hidden goal G^* (i.e., $G^* \in \mathcal{G}$); and $O = \langle o_1, o_2, \ldots, o_n \rangle$ is an observation sequence of executed actions, with each observation $o_i \in \mathcal{A}$.*

The ideal solution for a goal recognition problem is finding the single actual hidden goal $G^* \in \mathcal{G}$ that the observation sequence O of a plan execution achieves. Most approaches to goal and plan recognition return either a probability distribution over the goals (Ramírez and Geffner 2009; Ramírez and Geffner 2010; Sohrabi, Riabov, and Udrea 2016), or a score associated to the set of possible goals (Pereira, Oren, and Meneguzzi 2017). In Section 3, we explain how we infer the hidden goal from the observation sequence using LP. Note that an observation sequence can be either 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 Operator-Counting Framework

Recent work on linear programming (LP) based heuristics has generated a number of informative and efficient heuristics for optimal-cost planning (van den Briel et al. 2007; Pommerening et al. 2014; Bonet 2013a). These heuristics rely on constraints from different sources of information that every valid plan π must satisfy. All operator-counting constraints contain variables of the form Y_a for each operator a such that setting Y_a to the number of occurrences of a in π satisfies the constraints. In this paper, we adopt the formalism and definitions of Pommerening et al. (2014) for LP-based heuristics[1]. Definitions 2 and 3 formally define the concept of operator-counting constraints, and a planning heuristic based on operator-counting constraints (Pommerening et al. 2014).

Definition 2 (**Operator-Counting Constraints**). *Let Π be a planning instance with operator set $\mathcal{A}$ and let s be a reachable state in Π. Let $\mathcal{Y}$ be a set of non-negative real-valued and integer variables, including an integer variable Y_a for each operator $a \in \mathcal{A}$ along with any number of additional variables. The variables Y_a are called operator-counting variables. We say π is an s-plan in Π if it is a valid plan that leads from a state s to a goal G. If π is an s-plan, we denote the number of occurrences of operator $a \in \mathcal{A}$ in π with Y_a^π. A set of linear inequalities over $\mathcal{Y}$ is called an operator counting constraint for s if for every s-plan there exists a feasible solution with $\mathsf{Y}_a = \mathsf{Y}_a^\pi$ for all $a \in \mathcal{A}$. A constraint set for s is a set of operator-counting constraints for s where the only common variables between constraints are the operator-counting constraints.*

Definition 3 (**Operator-Counting IP/LP Heuristic**). *The operator-counting integer program IP_C for a set C of operator-counting constraints for state s is*

$$Minimize \sum_{a \in A} \mathsf{Y}_a \cdot cost(a) \text{ subject to}$$
$$C \text{ and } \mathsf{Y}_a \geq 0 \text{ for all } a \in A,$$

where A is the set of operators.

The IP heuristic h_C^{IP} is the objective value of IP_C, the LP heuristic h_C^{LP} is the objective value of its LP-relaxation. If the IP/LP is infeasible, the heuristic estimate is ∞.

While the framework from Pommerening, Röger, and Helmert (2013) unifies many types of constraints for operator-counting heuristics, we rely on three types of constraints for our goal recognition approaches:

---

[1] The only difference between their formalism and ours is that we refer to operators/actions with an $a/\mathcal{A}$ variable name to differentiate it from the observations o/O

landmarks (Bonet and van den Briel 2014), state-equations (van den Briel et al. 2007; Bonet 2013b), and post-hoc optimization (Pommerening, Röger, and Helmert 2013).

3 LP-Based Heuristics for Goal Recognition

We now develop an LP-based metric suitable for goal recognition. This metric basically adds a set of constraints to the LP of the operator-counting framework that enforce that solutions comply with all observations. We also use the information generated by the solution of the LP do estimate the number of missing observations. Thus, more missing observations lead to more relaxed criteria for our method to decide which goals should be selected as the correct one.

3.1 Hard Constraints

To develop our LP-based metric, we start with a basic operator-counting heuristic $h(s)$, which we define over the LP-heuristic of Definition 3 where C comprises the constraints generated by Landmarks, post-hoc optimization, and net change constraints as described by Pommerening et al. (2014). This heuristic yields important information about each goal hypothesis including the actual operator counts Y_a for all $a \in \mathcal{A}$ from Definition 3, whose minimization comprises the objective function $h(s)$. The heuristic value $h^G(s)$ for every candidate goal $G \in \mathcal{G}$ tells us about the optimal distance between the initial state $\mathcal{I}$ and G, while the operator counts indicate *possible* operators in a valid plan from $\mathcal{I}$ to G.

The h heuristic can be used to compute a heuristic conceptually similar to the *Goal Completion* heuristic of Pereira, Oren, and Meneguzzi (2017) if one tries to compare the operator counts and the observations. However, inferring goals using this overlap alone has a number of shortcomings in relation to their technique. First, while the landmarks themselves are enforced by the LP used to compute the operator counts (and thus observations that correspond to landmarks count as hits), the *overlap* computation loses the ordering of the landmarks that the Goal Completion heuristic uses to account for missing observations. Second, a solution to a set of operator-constraints, i.e., a set of operators with non-negative counts may not be a feasible plan for a planning instance. Third, if there are multiple valid plans to a goal hypothesis, there is no guarantee that the operator counts will actually correspond to the plan that generated the observations. Thus, these counts may not correspond to the plan that generated the observations.

While operator-counting heuristics on their own are fast and informative enough to help guide search when dealing with millions of nodes, goal recognition problems often require the disambiguation of a dozen or less goal hypotheses. Such goal hypotheses are often very similar so that the operator-counting heuristic value (i.e., the objective function over the operator counts) for each goal hypothesis is very similar, especially if the goals are more or less equidistant from the initial state.

Thus, we refine this heuristic by introducing additional constraints into the LP used to compute operator counts. Specifically, we force the operator counting heuristic to *only* *consider* operator counts that include every single observation $o \in O$. The resulting LP heuristic (which we call h_{HC}) then minimizes the cost of the operator counts for plans that necessarily agree with all observations.

We formally define these new constraints in Definition 4, and proceed to define the h_{HC} heuristic in Definition 5.

Definition 4 (**Hard Constraints for Observations**). *Let k_a be the number of occurrences of observations of the operator a in the sequence of observation O for the goal recognition task T. The* hard constraint $c_T^{hc,a}$ *for operator a is*

$$\mathsf{Y}_a \geq k_a$$

Definition 5 (h_{HC} **Heuristic**). *The integer program $\mathrm{IP}_C^{C^{hc,a}}$ for a set C of operator-counting constraints, and a set $C^{hc,a}$ of hard constraints for observations for state s is a operator-counting IP_C augmented with a set $C^{hc,a}$ of hard constraints for observations. The LP heuristic h_{HC} is the objective value of its LP-relaxation.*

3.2 Uncertainty

Up to this point, we have ignored a key challenge of goal recognition consisting of the unreliability of the observations. In most realistic settings, observations will either be noisy, incomplete, or both. Specifically, in settings where observability is low and goal hypotheses have similar distances from the initial state of the goal recognition problem, heuristically computed values will often agree. This limitation is particularly evident in the approach of Pereira, Oren, and Meneguzzi (2017), whose accuracy degrades substantially with lower observability for a number of domains. Unsurprisingly, the basic operator counting heuristic described so far also suffers from this limitation in such situations, as heuristic values for multiple goal hypotheses are either equal or very similar. This happens because when computing the overlap of unconstrained h heuristic, there might be a large number of possible solutions to the linear program with a similar objective value. This, in turn, might make these methods rule out the actual goal because of the specific solution returned by the LP-solver for h not having a substantial overlap with O, especially if the number of observations in O is small due to missing observations of the actual plan. However, as operator counting is an admissible heuristic, we know that the objective value for the constrained h_{HC} heuristic is necessarily a lower bound on the size of the optimal plan to achieve a specific goal hypothesis, regardless of the individual operator count values. This lower bound represents the minimal number of observations we expect to have received if the algorithm is under full observation, and the difference between the IP objective function and the length of the observation sequence constitutes the number of *missing observations* for a given goal recognition task. These missing observations allow us to account for a level of *uncertainty* in the expected value of the operator counts. In practice, we can compute an uncertainty ratio as shown in Equation 1 and infer the level of observability of a given goal recognition problem as we solve it, and return more goal hypotheses besides the least-cost one.

$$U \leftarrow 1 + \frac{\min_G h_{\text{HC}}^G - |O|}{\min_G h_{\text{HC}}^G} \qquad (1)$$

Algorithm 1 brings the h_{HC} heuristic and the notion of uncertainty together to perform goal recognition. It starts with the computation of the heuristic values for each goal in Lines 2–8, consisting of generating the operator counting constraints (Line 3), adding the constraints from Definition 4 (Lines 4–6) and computing h_{HC} (Line 8). Finally, we compute the uncertainty ratio using Equation 1 in Line 9 and return all goals that have either the least h_{HC} value, or are within the uncertainty ratio of this value in Line 10.

Algorithm 1 Goal Recognition using Linear Programming.

Input: Ξ *planning domain definition*, $\mathcal{I}$ *initial state*, $\mathcal{G}$ *set of candidate goals*, and O *observations*.
Output: *Recognized goal(s).*
1: **function** RECOGNIZEGOALS($\Xi, \mathcal{I}, \mathcal{G}, O$)
2:     **for all** $G \in \mathcal{G}$ **do**   ▷ Compute heuristics for G
3:         $C_{\text{HC}}^G \leftarrow$ GENERATECONSTRAINTS($\Xi, \mathcal{I}, G$)
4:         **for all** $o \in O$ **do**
5:             $k_o \leftarrow Count_{o \in O}$
6:             $C_{\text{HC}}^G \leftarrow C_{\text{HC}}^G \cup \mathsf{Y}_o > k_o$
7:         $IP^G \leftarrow$ GENERATEIP(C_{HC}^G, O)
8:         $h_{\text{HC}}^G \leftarrow$ LPSOLVER(IP^G)
9:     $U \leftarrow$ COMPUTEUNCERTAINTY(h_{HC}, O)
10:     **return** $\{G | G \in \mathcal{G} \wedge h_{\text{HC}}^G \leq \min_G h_{\text{HC}}^G * U\}$

3.3 Enforcement Delta

Although enforcing constraints to ensure that the LP heuristic computes only plans that do contain all observations helps us overcome the limitations of computing the overlap of the operator counts, this approach has a major shortcoming: it considers all observations as valid operators generated by the observed agent. Therefore, the heuristic resulting from the minimization of this LP might suffer from two problems of increasing severity. First, it might overestimate the actual length of the plan for the goal hypothesis due to noise. Second, depending on the problem and on the domain, enforcing an invalid action might even make the LP unsolvable. These problems may happen for one of two reasons: either the noise is simply a sub-optimal operator in a valid plan, or it is an operator that is completely unrelated to the plan that generated the observations. In both cases, the resulting heuristic value may prevent the algorithm from selecting the actual goal from the goal hypotheses. Proposition 1 states that this overestimation has the property that h_{HC} always dominates the operator counting heuristic h.

Proposition 1 (h_{HC} **dominates** h). *Let h be the basic operator-counting heuristic, h_{HC} be the over-constrained heuristic from Definition 5 that accounts for all observations $o \in O$, and s a state of Π. Then $h_{\text{HC}}(s) \geq h(s)$.*

Proof. Let C^h and $C^{h_{\text{HC}}}$ be sets of constraints used to compute, respectively, $h(s)$ and $h_{\text{HC}}(s)$. Every feasible solution to $C^{h_{\text{HC}}}$ is a solution to C^h. This is because to generate $C^{h_{\text{HC}}}$

we only *add* constraints to C^h. Thus, a solution to $C^{h_{\text{HC}}}$ has to satisfy all constraints in C^h. Therefore, since we are solving a minimization problem the value of the solution for C^h cannot be larger than the solution to $C^{h_{\text{HC}}}$. $\qquad \square$

Proposition 2. *The set of goals returned by h_{HC} with 100% of the observations always contains the actual goal.*

Proof. Suppose that the proposition is wrong. Hence there is a goal recognition task where h_{HC} with 100% of observations does not return the actual goal. Suppose that the G is the actual goal and that there are k observations thus $h_{\text{HC}}(G) = k$. Because the sequence of observations is a plan for k, and every plan satisfies all constraints of the IP of h_{HC}. If the proposition is wrong, this means that for another G' we have $h_{\text{HC}}(G') < h_{\text{HC}}(G)$. However, a solution for the IP of h_{HC} has to satisfy all hard constraints for observations and cannot cost less than k, which completes the proof. $\qquad \square$

The intuition here is that the operator-counting heuristic h estimates the total cost of any optimal plan, regardless of the observations, while h_{HC} estimates a plan following all observations, including noise, if any. If there is no noise, the sum of the counts must agree (even if the counts are different), whereas if there are noisy observations, there will be differences in all counts. Thus, our last approach consists of computing the difference between h_{HC} and h, and infer that the goal hypothesis for which these values are closer must be the actual goal. We call the resulting heuristic as δ_{HC} and formalize this approach in Algorithm 2. Here, we compute the LP twice, once with only the basic operator-counting constraints (Line 5), and once with the constraints enforcing the observations in the operator counts (Line 10), using these two values to compute δ_{HC} (Line 11). The algorithm then returns goal hypotheses that minimize δ_{HC} (Line 13) while being within the uncertainty induced by missing observations (Section 3.2). So, the combination of the heuristic δ_{HC} and the uncertainty ratio is denoted as δ_{HCU}.

Algorithm 2 Goal Recognition using Heuristic Difference of Operator Counts.

Input: Ξ *planning domain definition*, $\mathcal{I}$ *initial state*, $\mathcal{G}$ *set of candidate goals*, and O *observations*.
Output: *Recognized goal(s).*
1: **function** DELTARECOGNIZE($\Xi, \mathcal{I}, \mathcal{G}, O$)
2:     **for all** $G \in \mathcal{G}$ **do**   ▷ Compute $\delta_{\text{HC}}(\mathcal{I})$ for G
3:         $C_{\text{HC}}^G \leftarrow C^G \leftarrow$ GENERATECONSTRAINTS($\Xi, \mathcal{I}, G$)
4:         $IP_C^G \leftarrow$ GENERATEIP(C^G, O)
5:         $h \leftarrow$ LPSOLVER(IP_C^G)
6:         **for all** $o \in O$ **do**
7:             $k_o \leftarrow Count_{o \in O}$
8:             $C_{\text{HC}}^G \leftarrow C_{\text{HC}}^G \cup \mathsf{Y}_o > k_o$
9:         $IP_C^G \leftarrow$ GENERATEIP(C_{HC}^G, O)
10:         $h_{\text{HC}}^G \leftarrow$ LPSOLVER(IP_C^G)
11:         $\delta_{\text{HC}}^G \leftarrow h_{\text{HC}}^G - h$
12:     $U \leftarrow$ COMPUTEUNCERTAINTY(h^G, O)
13:     **return** $\{G | G \in \mathcal{G} \wedge \delta_{\text{HC}}^G \leq \min_G \delta_{\text{HC}}^G * U\}$

Figure 1: LP goal recognition example.

Example

Consider the example in Figure 1, where an agent can move up, right, down, and left. The agent has two possible goals G_1 and G_2, where the actual goal is G_1 achieved using the plan $\pi = \{o_1, o_2, o_3, o_4, o_5, o_6, o_7\}$ which is a non-optimal plan. Also, in this example, we have a noisy observation o_8 that is not part of the plan π. Suppose our method is given the complete sequence of observations $O = \{o_1, o_2, \ldots, o_7\}$, and that our heuristic can always compute the cost of the optimal path. Then, $h(G_1) = 3$, $h(G_2) = 3$, and the heuristics that have to comply with all observations have the value of $h_{\text{HC}}(G_1) = 7$ and $h_{\text{HC}}(G_2) = 9$. Thus, our δ values are $\delta_{\text{HC}}(G_1) = 4$ and $\delta_{\text{HC}}(G_2) = 6$. In this situation, our method would select the correct goal G_1.

Now, consider the case where we have only one observation $O = \{o_3\}$. Again, in this case our δ values would be $\delta_{\text{HC}}(G_1) = 4$ and $\delta_{\text{HC}}(G_2) = 6$, and our method would select only G_1. However, we believe that in this, we should relax our decision because we have just one observation that is not part of an optimal path for any goal. Thus, we introduce uncertainty that for this case is $U = 1 + \frac{7-1}{7}$, and now we would return the two goals. Last, if we increase the number of observations provided to our method to four with $O = \{o_1, o_2, o_3, o_4\}$. All values of our heuristics would remain the same, but our metric our uncertainty would change to $U = 1 + \frac{7-4}{7}$ since we were provided more observations and we can be more confident of our decision. In this last case, we would select only the correct actual goal.

4 Experiments and Evaluation

We now report on the experiments we carried out to evaluate our LP-based heuristic approaches against the state-of-the-art in goal recognition as planning.

4.1 Domains

We implemented each of the heuristics described earlier and performed the goal recognition process over the large dataset introduced by Pereira, Oren, and Meneguzzi (2017)[2]. This dataset contains thousands of problems for goal and plan recognition under varying levels of observability (with optimal and suboptimal plans) for a number of traditional

IPC domains (Vallati, Chrpa, and McCluskey 2018), including BLOCKS-WORLD, CAMPUS, DEPOTS, DRIVERLOG, Dock-Worker Robots (DWR), IPC-GRID, FERRY, Intrusion Detection (INTRUSION), KITCHEN, LOGISTICS, MICONIC, ROVER, SATELLITE, SOKOBAN, and Zeno Travel (ZENO). This dataset also contains over a thousand goal recognition problems under partial observability *and* noisy observations in the CAMPUS, IPC-GRID, INTRUSION and KITCHEN domains.

To improve the experiments and evaluation for noisy observations, we extended the noisy dataset by creating new goal recognition problems for the other domains with 2 additional noisy actions for each level of observability, so with this, we now have a dataset with noisy observations for the same number of domains. For example, consider the average of number observations $|O|$ for DEPOTS for 25% of observability (Table 2), which is 4.4, so 2 of these 4.4 actions are noisy (i.e., spurious actions), and so on for all levels of observability in the noisy dataset.

4.2 Setup

We implemented our approach using PYTHON 2.7 for the main recognition algorithms with external calls to a customized version of the FAST-DOWNWARD (Helmert 2006) planning system[3] to compute the operator counts. Our customized planner not only returns the operator counts, but can also introduce additional constraints before running the CPLEX 12.9 optimization system. We ran experiments in a single core of a 24 core Intel® Xeon® CPU E5-2620 @2.00Ghz with 48GB of RAM, with a 2-minute time limit and a 2GB memory limit.

We note that, while the recognition time is competitive with R&G 2009 (and indeed other approaches), it can be substantially improved by implementing the operator counting heuristic within the main recognizer rather than relying on multiple calls to an external planner.

4.3 Benchmark

To evaluate the effectiveness of our heuristic approaches (δ_{SC} and δ_{HCU}), the baselines of our experimentation are the original approach from Ramírez and Geffner (2009) (R&G 2009) and the recent recognition heuristics from Pereira, Oren, and Meneguzzi (2017) (POM 2017 h_{gc} and POM 2017 h_{uniq}). Due to timeout and parsing issues that some approaches have faced during the experiments, we decided to show the comparison against such approaches (E-Martín, R.-Moreno, and Smith 2015; Sohrabi, Riabov, and Udrea 2016; Vered et al. 2018) in the supplement.

Table 1 shows the experimental results for the partially observable, non-noisy fragment of the dataset, whereas Table 2 shows the noisy fragment of the same dataset. For the noisy experiments, every sequence of observations contained at least 2 noisy (spurious) actions, which, while valid for the plan, were not actually executed by the agent being observed during the recognition process. Both tables show the number of recognition problems under the domain name, the average number of candidate goals $|\mathcal{G}|$, and the average

---

[2]https://zenodo.org/record/825878

[3]Revision 12828 from https://www.fast-downward.org

number of observations $|O|$. For evaluation, we used three metrics: recognition time in seconds (Time); accuracy – the fraction of times in which the actual hidden goals has been recognized correctly (Acc %); and spread in $\mathcal{G}$ – the average number of returned goals (S in $\mathcal{G}$). At the bottom of Tables 1 and 2, we provide the average for evaluated metrics we used over all domains and problems in the datasets.

The results in Tables 1 and 2 show that, while not nearly as fast as the heuristics approaches (POM 2017 h_{gc} and POM 2017 h_{uniq}) from Pereira, Oren, and Meneguzzi with a $\theta = 0$ recognition threshold, the accuracy (Acc %) of our δ_{HC} approach is either competitive or superior in virtually all domains (except for some levels of observability in IPC-GRID and DWR), and, even for the domains where the accuracy is similar, or lower, the spread (S in $\mathcal{G}$) of the resulting goals is very competitive against these landmark-based heuristics, i.e., the returned goals are unique for most problems. The accuracy of our approach with and without uncertainty (δ_{HC} and δ_{HCU}), thus, consistently matches or surpasses that of R&G 2009 for most domains and problems. Importantly, the cost of our approach is basically the same within each domain, regardless of the level of observability and noise, since our techniques rely on a single call to a planner that computes the operator counts for a single state and then stops the planner.

We note that the results for noisy observations (Table 2) show the greatest impact of our LP-based heuristics, especially δ_{HC}, showing an overall higher accuracy and lower spread across all domains. It is possible to see that δ_{HC} surpasses all other evaluated approaches in accuracy and spread when dealing with noisy observations. Thus, our LP-based heuristics show their effectiveness not only for dealing with partial and full observable plans (Table 1), but also for dealing with noisy, partial, and full observations (Table 2).

5 Related Work

Our LP-based heuristics follow the tradition of goal and plan recognition as planning algorithms as defined by Ramírez and Geffner (2009; 2010). The work developed by Ramírez and Geffner (2009) yields higher recognition accuracy in our settings (and hence we chose it as a baseline), whereas the work of Ramírez and Geffner (2010) models goal recognition as a problem of estimating the probability of a goal given the observations. The latter work uses a Bayesian framework to compute the probability of goals given observations by computing the probability of generating a plan given a goal, which they accomplish by running a planner multiple times to estimate the probability of the plans that either comply or not with the observations. By contrast, we do not try to provide a probabilistic interpretation for the recognized hypotheses.

Recent research on goal recognition has yielded a number of approaches to deal with partial observability and noisy observations, of which we single out three key contributions. First, E-Martín, R.-Moreno, and Smith (2015) developed a goal recognition approach based on constructing a planning graph and propagating operator costs and the interaction among operators to provide an estimate of the probabilities of each goal hypothesis. While their approach pro-

vides probabilistic estimates for each goal, its precision in inferring the most likely goals is consistently lower than ours, often ranking multiple goals with equal probabilities (i.e., having a large spread). Second, Sohrabi, Riabov, and Udrea (2016) develop an approach that also provides a probabilistic interpretation and explicitly deals with missing and noisy observations by adding weights in the domain description. Their approach works through a compilation of the recognition problem into a planning problem that is processed by a planner that computes a number of approximately optimal plans to compute goal probabilities under R&G's Bayesian framework (Ramírez and Geffner 2010). Third, Masters and Sardiña (2017) developed a fast and accurate approach to goal recognition that works strictly in the context of path-planning, providing a new probabilistic framework for goal recognition in path planning. Finally, Pereira, Oren, and Meneguzzi (2017) develop heuristic goal recognition approaches using landmark information. While the approach of Pereira, Oren, and Meneguzzi is conceptually closer to ours in that we also compute heuristics, our approach is distinct and novel in at least two ways: 1) we overcome the sparsity of landmarks in most domain by using operator-count information; and 2) we explicitly handle partial observability by estimating the missing observations. Whereas their approach tries to overcome low accuracy under low observability by introducing a θ parameter to relax the number of goals ranked as most likely, we automatically infer this relaxation by computing an uncertainty ratio. The result is that our approach consistently outperforms both their heuristics at low observability, especially at low (10%) observability. Finally, while we do not explicitly try to overcome noise with the constraints, we prove to be substantially more accurate than the state-of-the-art when dealing with noisy observations.

6 Conclusions

In this paper, we develop a novel class of goal recognition techniques based on operator-counting heuristics from classical planning (Pommerening et al. 2014) which, themselves rely on Integer LP constraints to estimate which operators occur in valid optimal plans towards a goal. The resulting heuristic approach outperforms the state-of-the-art in terms of high accuracy and low false positive rate (i.e., the spread of returned goals), at a moderate computational cost (i.e., recognition time). We have shown empirically that the overall accuracy of our approach is substantially higher to the state-of-the-art over a large dataset with noisy, partial, and full observable plans.

The technique described in this paper uses a set of simple additional constraints in the Integer LP formulation to improve performance, so we expect substantial future work towards further goal recognition approaches and heuristics that explore more refined constraints to improve accuracy and reduce spread, as well as deriving a probabilistic approach using operator-counting information. Examples of such work include using the constraints to force the LP to generate the counterfactual operator-counts (i.e., non-compliant with the observations) used by the R&G approach, or, given an estimate of the noise, relax the obser-

Table 1:

	$|\mathcal{G}|$			δ_{HC}			δ_{HCU}			R&G 2009			POM 2017 h_{gc}			POM 2017 h_{uniq}		
#		% Obs	$|O|$	Time	Acc %	S in $\mathcal{G}$	Time	Acc %	S in $\mathcal{G}$	Time	Acc %	S in $\mathcal{G}$	Time	Acc %	S in $\mathcal{G}$	Time	Acc %	S in $\mathcal{G}$
BLOCKS (1076)	20.3	10	1.8	11.575	95.5%	7.81	11.574	95.9%	8.7	1.222	86.8%	7.84	0.144	39.9%	1.05	0.131	31.7%	1.04
		30	4.9	11.570	89.4%	2.79	11.567	94.3%	5.12	1.639	87.2%	3.57	0.156	50.6%	1.09	0.144	51.4%	1.06
		50	7.6	11.581	92.7%	1.78	11.587	93.9%	3.24	2.402	97.9%	2.63	0.179	65.0%	1.09	0.168	60.1%	1.08
		70	11.1	11.594	98.8%	1.41	11.595	97.5%	1.8	3.785	97.7%	1.83	0.192	84.8%	1.12	0.192	79.0%	1.14
		100	14.5	11.904	100.0%	1.21	11.937	100.0%	1.21	6.791	100.0%	1.46	0.246	100.0%	1.36	0.239	100.0%	1.09
DEPOTS (364)	8.5	10	3.1	8.299	61.9%	2.15	8.305	71.4%	3.51	1.496	77.4%	3.99	0.369	35.7%	1.18	0.393	32.1%	1.1
		30	8.6	8.292	71.4%	1.4	8.298	88.1%	4.17	2.309	77.4%	2.39	0.357	58.3%	1.06	0.393	47.6%	1.07
		50	14.1	8.282	92.9%	1.38	8.280	98.8%	3.65	3.411	84.5%	1.92	0.369	76.2%	1.06	0.405	71.4%	1.02
		70	19.7	8.282	97.6%	1.06	8.279	98.8%	1.75	5.271	91.7%	1.68	0.393	89.3%	1.01	0.444	84.5%	1.01
		100	24.4	8.296	100.0%	1.0	8.296	100.0%	1.0	7.117	92.9%	1.46	0.464	100.0%	1.04	0.502	100.0%	1.04
DRIVERLOG (364)	10.5	10	2.6	5.159	77.4%	2.61	5.168	78.6%	3.17	1.169	96.4%	4.71	0.333	41.7%	1.04	0.321	35.7%	1.11
		30	6.9	5.149	83.3%	1.69	5.163	91.7%	2.76	1.411	92.9%	3.35	0.311	54.8%	1.13	0.310	47.6%	1.1
		50	11.1	5.167	92.9%	1.23	5.146	97.6%	2.0	1.694	94.0%	2.88	0.321	72.6%	1.17	0.310	64.3%	1.14
		70	15.6	5.148	95.2%	1.14	5.146	95.2%	1.5	1.973	89.3%	2.46	0.333	90.5%	1.14	0.321	90.5%	1.17
		100	21.7	5.166	100.0%	1.04	5.157	100.0%	1.04	2.821	89.3%	2.14	0.321	100.0%	1.21	0.321	100.0%	1.18
DWR (364)	7.3	10	5.7	5.759	54.8%	2.21	5.745	95.2%	5.45	1.767	83.3%	4.21	0.452	36.9%	1.1	0.512	33.3%	1.06
		30	16.0	5.745	83.3%	1.58	5.753	100.0%	4.92	2.723	81.0%	3.35	0.452	60.7%	1.04	0.504	51.2%	1.06
		50	26.2	5.742	90.5%	1.21	5.734	100.0%	3.98	4.822	72.6%	2.27	0.488	66.7%	1.0	0.548	61.9%	1.01
		70	36.8	5.752	97.6%	1.07	5.735	100.0%	2.26	10.914	70.2%	2.05	0.536	89.3%	1.0	0.607	78.6%	1.05
		100	51.9	5.763	100.0%	1.0	5.75	100.0%	1.0	25.092	67.9%	1.68	0.643	100.0%	1.0	0.751	96.4%	1.04
DPC-GRID (673)	9.0	10	2.9	6.216	92.8%	1.92	6.209	94.8%	2.32	1.091	96.1%	2.46	0.248	66.7%	2.58	0.242	62.7%	2.58
		30	7.8	6.051	95.4%	1.29	6.044	98.0%	1.48	1.476	97.4%	1.42	0.242	81.7%	1.65	0.242	83.7%	1.66
		50	12.7	6.123	98.7%	1.11	6.122	100.0%	1.25	1.905	100.0%	1.16	0.261	90.8%	1.18	0.248	90.8%	1.18
		70	17.9	6.251	99.4%	1.1	6.259	100.0%	1.19	2.552	100.0%	1.05	0.268	97.4%	1.07	0.268	97.4%	1.07
		100	24.8	5.825	100.0%	1.03	5.826	100.0%	1.0	4.057	100.0%	1.0	0.262	100.0%	1.0	0.262	100.0%	1.0
FERRY (364)	7.5	10	2.9	4.201	100.0%	3.17	4.201	100.0%	3.2	0.491	98.8%	3.37	0.071	58.3%	1.26	0.071	58.3%	1.18
		30	7.6	4.072	100.0%	1.56	4.075	100.0%	1.76	0.677	100.0%	1.76	0.061	85.7%	1.12	0.060	83.3%	1.06
		50	12.3	4.141	100.0%	1.29	4.141	100.0%	1.44	0.795	100.0%	1.42	0.062	95.2%	1.07	0.060	91.7%	1.01
		70	17.3	4.199	100.0%	1.1	4.197	100.0%	1.12	1.253	98.8%	1.14	0.071	100.0%	1.01	0.071	100.0%	1.0
		100	24.2	4.196	100.0%	1.07	4.204	100.0%	1.07	1.631	100.0%	1.07	0.071	100.0%	1.0	0.071	100.0%	1.0
LOGISTICS (673)	10.5	10	2.9	6.792	100.0%	2.5	6.789	100.0%	2.8	1.201	99.3%	2.98	0.641	55.6%	1.73	0.641	49.0%	1.24
		30	8.2	6.953	98.0%	1.3	6.944	98.0%	1.76	1.799	98.7%	1.39	0.621	80.4%	1.21	0.634	76.5%	1.12
		50	13.4	6.944	98.7%	1.13	6.959	98.7%	1.37	2.509	98.7%	1.29	0.641	90.2%	1.1	0.647	86.3%	1.05
		70	18.9	6.938	100.0%	1.08	6.951	100.0%	1.15	3.461	100.0%	1.13	0.667	96.7%	1.06	0.661	96.7%	1.02
		100	26.5	6.632	100.0%	1.0	6.633	100.0%	1.0	4.832	100.0%	1.0	0.607	100.0%	1.0	0.607	100.0%	1.0
MICONIC (364)	6.0	10	3.9	4.905	100.0%	2.12	4.886	100.0%	2.29	0.813	100.0%	3.26	0.464	67.9%	1.33	0.352	54.8%	1.26
		30	11.1	4.897	100.0%	1.19	4.895	100.0%	1.46	1.191	100.0%	1.58	0.452	96.4%	1.11	0.364	90.5%	1.08
		50	18.1	4.901	100.0%	1.1	4.891	100.0%	1.32	1.722	100.0%	1.29	0.452	96.4%	1.01	0.352	96.4%	1.0
		70	25.3	4.892	100.0%	1.01	4.891	100.0%	1.02	2.591	100.0%	1.04	0.452	100.0%	1.01	0.376	100.0%	1.01
		100	35.6	4.894	100.0%	1.0	4.894	100.0%	1.0	5.107	100.0%	1.0	0.464	100.0%	1.0	0.364	100.0%	1.0
ROVERS (364)	6.0	10	3.0	4.863	98.8%	2.71	4.858	100.0%	2.94	0.745	98.8%	2.86	0.348	64.3%	1.73	0.371	51.2%	1.11
		30	7.9	4.881	85.7%	1.17	4.871	91.7%	1.83	1.031	100.0%	1.67	0.348	83.3%	1.24	0.348	69.0%	1.07
		50	12.7	4.870	98.8%	1.14	4.856	98.8%	1.44	1.345	100.0%	1.3	0.336	92.9%	1.08	0.348	85.7%	1.01
		70	17.9	4.858	98.8%	1.01	4.838	98.8%	1.06	1.177	100.0%	1.07	0.348	98.8%	1.01	0.362	91.7%	1.0
		100	24.9	4.823	100.0%	1.0	4.884	100.0%	1.0	2.298	100.0%	1.07	0.371	100.0%	1.0	0.571	100.0%	1.0
SATELLITE (364)	6.5	10	2.1	5.196	91.7%	2.73	5.204	92.9%	2.92	1.076	97.6%	3.42	0.451	57.1%	1.56	0.414	47.6%	1.21
		30	5.4	5.178	92.9%	1.76	5.201	96.4%	2.31	1.183	97.6%	2.4	0.451	76.2%	1.31	0.414	69.0%	1.14
		50	8.7	5.197	96.4%	1.32	5.194	98.8%	1.77	1.328	97.6%	1.69	0.426	85.7%	1.1	0.414	81.0%	1.11
		70	12.2	5.191	97.6%	1.11	5.193	97.6%	1.21	1.841	96.4%	1.52	0.402	97.6%	1.02	0.414	94.0%	1.04
		100	16.8	5.212	100.0%	1.07	5.205	100.0%	1.07	2.045	96.4%	1.32	0.414	100.0%	1.07	0.414	100.0%	1.07
SOKOBAN (364)	7.3	10	3.1	7.827	72.6%	1.61	7.849	78.6%	2.39	3.153	69.0%	4.02	0.607	53.6%	2.06	0.607	51.2%	1.86
		30	8.7	7.744	89.3%	1.11	7.742	95.2%	1.75	4.622	89.3%	4.17	0.595	57.1%	1.37	0.607	56.0%	1.21
		50	14.1	7.701	95.2%	1.08	7.709	100.0%	1.46	7.441	89.3%	4.11	0.595	71.4%	1.32	0.607	69.0%	1.2
		70	19.8	7.690	97.6%	1.04	7.677	98.8%	1.15	9.877	89.3%	4.18	0.608	83.3%	1.05	0.607	86.9%	1.08
		100	35.5	7.671	100.0%	1.0	7.658	100.0%	1.0	12.996	89.3%	4.54	0.607	100.0%	1.0	0.643	100.0%	1.0
ZENO (364)	7.5	10	2.6	6.843	86.9%	2.71	6.871	88.1%	3.12	1.814	96.4%	3.4	0.567	39.3%	1.11	0.555	36.9%	1.05
		30	6.7	6.838	90.5%	1.61	6.854	96.4%	2.56	2.539	88.1%	2.12	0.555	70.2%	1.15	0.531	60.7%	1.02
		50	10.8	6.845	95.2%	1.15	6.851	96.4%	1.83	3.079	92.9%	1.42	0.543	78.6%	1.07	0.555	76.2%	1.0
		70	15.2	6.855	100.0%	1.0	6.851	100.0%	1.04	3.907	96.4%	1.13	0.567	97.6%	1.05	0.555	90.5%	1.0
		100	21.1	6.852	100.0%	1.0	6.866	100.0%	1.0	4.866	100.0%	1.07	0.543	100.0%	1.0	0.543	100.0%	1.0
Average				6.456	94.11%	1.55	6.457	96.94%	2.14	3.322	93.38%	2.30	0.397	79.66%	1.18	0.400	75.87%	1.12

Table 1: Experimental results comparing our lp-based heuristics against the state-of-the-art under partial, and full observable plans.

Table 2:

	$|\mathcal{G}|$			δ_{HC}			δ_{HCU}			R&G 2009			POM 2017 h_{gc}			POM 2017 h_{uniq}		
#		% Obs	$|O|$	Time	Acc %	S in $\mathcal{G}$	Time	Acc %	S in $\mathcal{G}$	Time	Acc %	S in $\mathcal{G}$	Time	Acc %	S in $\mathcal{G}$	Time	Acc %	S in $\mathcal{G}$
BLOCKS (144)	20.3	25	2.4	16.033	58.3%	6.22	16.019	75.0%	10.39	1.045	38.9%	5.39	0.083	2.8%	1.22	0.083	8.3%	1.0
		50	4.4	12.626	52.8%	3.31	12.675	88.9%	12.44	1.122	52.8%	4.61	0.083	13.9%	1.09	0.083	13.9%	1.08
		75	6.8	11.282	80.6%	2.11	11.235	91.7%	7.36	1.405	75.0%	2.72	0.056	47.2%	1.19	0.056	38.9%	1.25
		100	8.8	10.576	88.9%	1.92	10.629	97.2%	2.86	1.652	86.1%	2.03	0.083	77.8%	1.36	0.056	75.0%	1.33
DEPOTS (144)	9.3	25	4.4	11.495	52.8%	2.19	11.445	63.9%	4.47	0.284	5.6%	9.17	0.528	38.9%	1.64	0.528	27.8%	1.22
		50	8.4	10.299	61.1%	1.67	10.263	83.3%	4.14	0.189	0.0%	9.33	0.472	52.8%	1.2	0.472	41.7%	1.19
		75	12.7	8.671	88.9%	1.19	8.674	94.4%	2.14	0.361	11.1%	8.28	0.472	80.6%	1.11	0.500	75.0%	1.06
		100	16.2	8.098	94.4%	1.19	8.075	91.7%	1.22	0.292	5.6%	8.83	0.472	88.9%	1.11	0.472	86.1%	1.11
DRIVERLOG (144)	6.6	25	3.5	5.383	55.6%	2.61	5.362	83.3%	4.69	0.234	44.4%	5.89	0.111	36.1%	1.33	0.111	25.0%	1.08
		50	6.7	4.982	77.8%	1.72	4.971	91.7%	3.47	0.282	38.9%	4.72	0.083	58.3%	1.28	0.083	52.8%	1.11
		75	10.0	4.825	86.1%	1.25	4.855	94.4%	2.31	0.239	30.6%	5.47	0.083	61.1%	1.33	0.111	52.8%	1.14
		100	12.8	4.511	97.2%	1.31	4.511	97.2%	1.47	0.358	44.4%	4.42	0.084	94.4%	1.47	0.083	97.2%	1.42
DWR (144)	7.0	25	9.2	7.546	72.2%	2.0	7.614	97.2%	5.11	0.808	41.7%	5.67	0.444	44.4%	1.14	0.501	33.3%	1.0
		50	17.8	7.108	80.6%	1.67	7.153	94.4%	4.47	1.569	22.2%	5.39	0.417	63.9%	1.08	0.444	50.0%	1.06
		75	26.6	6.509	91.7%	1.22	6.482	94.4%	1.78	2.793	19.4%	5.5	0.417	94.4%	1.06	0.472	69.4%	1.08
		100	34.9	5.803	100.0%	1.0	5.768	97.2%	1.0	7.392	30.6%	4.42	0.444	94.4%	1.0	0.472	94.4%	1.03
DPC-GRID (300)	8.3	25	4.0	7.871	81.1%	1.67	7.889	85.6%	2.61	0.265	12.2%	7.56	0.244	58.9%	1.78	0.233	53.3%	1.72
		50	7.7	6.031	94.4%	1.14	6.011	94.4%	1.71	0.240	4.4%	8.07	0.222	85.6%	1.33	0.211	83.3%	1.32
		75	11.5	5.481	98.9%	1.1	5.497	97.8%	1.13	0.223	6.7%	7.89	0.213	94.4%	1.09	0.211	94.4%	1.09
		100	16.9	5.011	100.0%	1.0	4.988	90.0%	0.9	0.266	10.0%	7.77	0.202	100.0%	1.0	0.200	100.0%	1.0
FERRY (144)	7.0	25	5.8	4.745	80.6%	2.78	4.736	86.1%	5.11	0.256	75.0%	2.92	0.028	27.8%	1.08	0.028	27.8%	1.08
		50	11.2	4.402	97.2%	1.72	4.365	97.2%	3.83	0.375	94.4%	1.94	0.028	88.9%	1.31	0.028	77.8%	1.06
		75	16.6	4.236	94.4%	1.56	4.223	97.2%	2.42	0.585	88.9%	1.53	0.028	97.2%	1.17	0.028	83.3%	1.08
		100	21.9	4.187	97.2%	1.17	4.183	97.2%	1.17	0.906	97.2%	1.25	0.029	100.0%	1.08	0.028	97.2%	1.06
LOGISTICS (144)	10.0	25	4.8	7.461	91.7%	2.11	7.446	97.2%	3.42	0.203	5.6%	9.42	0.222	61.1%	1.56	0.222	38.9%	1.08
		50	9.4	6.875	97.2%	1.11	6.848	97.2%	1.36	0.214	5.6%	9.33	0.194	83.3%	1.17	0.167	75.0%	1.06
		75	14.1	6.253	100.0%	1.03	6.234	100.0%	1.06	0.259	13.9%	8.78	0.194	97.2%	1.03	0.167	100.0%	1.03
		100	18.1	5.436	100.0%	1.08	5.438	100.0%	1.08	0.321	13.9%	8.78	0.195	100.0%	1.06	0.194	100.0%	1.03
MICONIC (144)	6.0	25	4.4	4.892	52.8%	1.97	4.908	83.3%	4.17	0.318	88.9%	2.89	0.111	50.0%	1.39	0.111	33.3%	1.2
		50	8.4	4.533	80.6%	1.17	4.566	97.2%	2.17	0.387	100.0%	1.75	0.083	83.3%	1.11	0.083	80.6%	1.11
		75	12.6	4.387	91.7%	1.06	4.379	100.0%	1.31	0.483	100.0%	1.19	0.083	97.2%	1.03	0.083	88.9%	1.03
		100	16.3	4.294	100.0%	1.03	4.311	100.0%	1.08	0.628	100.0%	1.0	0.056	100.0%	1.0	0.056	100.0%	1.0
ROVERS (144)	6.0	25	3.1	5.346	72.2%	2.22	5.359	75.0%	2.78	0.264	33.3%	4.78	0.083	52.8%	1.14	0.083	50.0%	1.14
		50	5.7	4.973	80.6%	1.61	4.941	88.9%	2.92	0.296	50.0%	3.81	0.056	69.4%	1.31	0.056	58.3%	1.08
		75	8.4	4.889	91.7%	1.14	4.905	100.0%	1.47	0.301	44.4%	3.92	0.083	86.1%	1.14	0.056	75.0%	1.08
		100	10.8	4.513	100.0%	1.03	4.531	100.0%	1.03	0.305	38.9%	4.06	0.083	97.2%	1.11	0.056	86.1%	1.03
SATELLITE (144)	6.0	25	3.3	4.511	75.0%	3.31	4.444	88.9%	4.28	0.224	58.3%	4.53	0.056	30.6%	2.42	0.056	30.6%	1.33
		50	5.7	4.056	72.2%	2.44	4.091	83.3%	3.92	0.255	72.2%	3.58	0.028	72.2%	2.08	0.028	44.4%	1.31
		75	8.4	3.915	83.3%	1.44	3.966	88.9%	2.83	0.292	77.8%	2.75	0.028	69.4%	1.28	0.028	69.4%	1.19
		100	10.7	3.929	94.4%	1.47	3.935	94.4%	1.86	0.298	72.2%	3.0	0.056	94.4%	1.31	0.028	91.7%	1.19
SOKOBAN (144)	8.6	25	5.3	13.104	36.1%	1.64	13.101	72.2%	4.69	1.953	25.0%	7.28	0.751	41.7%	1.75	0.750	38.9%	1.56
		50	10.3	11.319	50.0%	1.17	11.323	58.3%	1.94	2.086	19.4%	6.67	0.667	66.7%	1.44	0.694	58.3%	1.08
		75	15.6	9.527	36.1%	1.0	9.461	36.1%	0.56	2.121	19.4%	7.69	0.694	80.6%	1.28	0.667	72.2%	1.03
		100	20.1	8.946	33.3%	1.11	9.029	36.1%	0.5	5.878	33.3%	6.47	0.694	94.4%	1.22	0.694	86.1%	1.06
ZENO (144)	6.6	25	3.0	8.209	44.4%	2.72	8.271	80.6%	5.25	0.946	72.2%	3.92	0.417	55.6%	1.92	0.417	33.3%	1.03
		50	5.8	7.712	91.7%	1.61	7.653	97.2%	4.03	1.025	88.9%	1.78	0.361	77.8%	1.67	0.361	61.1%	1.08
		75	8.8	6.785	91.7%	1.08	6.766	100.0%	2.44	1.167	100.0%	1.22	0.361	88.9%	1.28	0.361	77.8%	1.03
		100	11.3	6.177	97.2%	1.0	6.141	100.0%	1.42	1.289	97.2%	1.08	0.361	97.2%	1.08	0.333	97.2%	1.03
Average				6.870	80.14%	1.69	6.868	88.68%	3.04	0.929	47.22%	5.01	0.236	73.21%	1.30	0.234	64.71%	1.13

Table 2: Experimental results comparing our LP-based heuristics against the state-of-the-art under noisy, partial, and full observable plans.

vation constraints to discard observations from the resulting operator-counts. Finally, we argue that casting goal recognition as an ILP problem is both natural, and a promising avenue for further research not only in symbolic (STRIPS-like) domains, but also for path planning and continuous domains (Masters and Sardiña 2017).

References

[Bonet and van den Briel 2014] Bonet, B., and van den Briel, M. 2014. Flow-based heuristics for optimal planning: Landmarks and merges. In *Proceedings of the 24th International Conference on Automated Planning and Scheduling (ICAPS)*.

[Bonet 2013a] Bonet, B. 2013a. An Admissible Heuristic for SAS+ Planning Obtained from the State Equation. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*.

[Bonet 2013b] Bonet, B. 2013b. An Admissible Heuristic for SAS+ Planning obtained from the State Equation. In *Proceedings of the 23nd International Joint Conference on Artificial Intelligence (IJCAI)*.

[E-Martín, R.-Moreno, and Smith 2015] E-Martín, Y.; R.-Moreno, M. D.; and Smith, D. E. 2015. A Fast Goal Recognition Technique Based on Interaction Estimates. In *Proceedings of 24th the International Joint Conference on Artificial Intelligence (IJCAI)*.

[Ghallab, Nau, and Traverso 2004] Ghallab, M.; Nau, D.; and Traverso, P. 2004. *Automated Planning: Theory and Practice*. Burlington, MA: Elsevier.

[Helmert 2006] Helmert, M. 2006. The Fast Downward Planning System. *Journal of Artificial Intelligence Research* 26:191–246.

[Masters and Sardiña 2017] Masters, P., and Sardiña, S. 2017. Cost-Based Goal Recognition for Path-Planning. In *Proceedings of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.

[Pereira and Meneguzzi 2017] Pereira, R. F., and Meneguzzi, F. 2017. Goal and plan recognition datasets using classical planning domains.

[Pereira, Oren, and Meneguzzi 2017] Pereira, R. F.; Oren, N.; and Meneguzzi, F. 2017. Landmark-Based Heuristics for Goal Recognition. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence*.

[Pommerening et al. 2014] Pommerening, F.; Röger, G.; Helmert, M.; and Bonet, B. 2014. LP-Based Heuristics for Cost-Optimal Planning. *Proceedings of the 24th International Conference on Automated Planning and Scheduling (ICAPS)*.

[Pommerening, Röger, and Helmert 2013] Pommerening, F.; Röger, G.; and Helmert, M. 2013. Getting the Most Out of Pattern Databases for Classical Planning. *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*.

[Ramírez and Geffner 2009] Ramírez, M., and Geffner, H. 2009. Plan recognition as planning. In *Proceeding of the International Joint Conference on Artificial Intelligence (IJCAI)*.

[Ramírez and Geffner 2010] Ramírez, M., and Geffner, H. 2010. Probabilistic Plan Recognition Using Off-the-Shelf Classical Planners. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence*, 1121–1126.

[Sohrabi, Riabov, and Udrea 2016] 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)*.

[Sukthankar et al. 2014] Sukthankar, G.; Goldman, R. P.; Geib, C.; Pynadath, D. V.; and Bui, H. H. 2014. *Plan, Activity, and Intent Recognition: Theory and Practice*. Elsevier.

[Vallati, Chrpa, and McCluskey 2018] Vallati, M.; Chrpa, L.; and McCluskey, T. L. 2018. What you always wanted to know about the deterministic part of the International Planning Competition (IPC) 2014 (but were too afraid to ask). *Knowledge Engineering Review* 33:383.

[van den Briel et al. 2007] van den Briel, M.; Benton, J.; Kambhampati, S.; and Vossen, T. 2007. An LP-Based Heuristic for Optimal Planning. *International Conference on Principles and Practice of Constraint Programming (CP)*.

[Vered et al. 2018] Vered, M.; Pereira, R. F.; Magnaguagno, M.; Meneguzzi, F.; and Kaminka, G. A. 2018. Towards Online Goal Recognition Combining Goal Mirroring and Landmarks. In *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.