Inferring Agents Preferences as Priors for Probabilistic Goal Recognition

Kin Max Gusmão[1], **Ramon Fraga Pereira**[2], **Felipe Meneguzzi**[3]
[1]Pontifical Catholic University of Rio Grande do Sul, Brazil
[2]Sapienza University of Rome, Italy
kin.gusmao@edu.pucrs.br, pereira@diag.uniroma1.it, felipe.meneguzzi@pucrs.br

Abstract

Recent approaches to goal recognition have leveraged the concept of planning landmarks to achieve high-accuracy with low run-time cost. These approaches, however, lack a probabilistic interpretation. Furthermore, while most probabilistic models to goal recognition assume that the recognizer has access to a prior probability representing, for example, an agent's preferences, virtually no goal recognition approach actually uses the prior in practice, simply assuming a uniform prior. In this paper, we provide a model to both extend landmark-based goal recognition with a probabilistic interpretation and allow the estimation of such prior probability and its usage to compute posterior probabilities after repeated interactions of observed agents. We empirically show that our model not only recognizes goals effectively but also successfully infers the correct prior probability distribution representing an agent's preferences.

1 Introduction

Goal Recognition is the task of inferring an agent's goals, given a potentially flawed observation of this agent's behavior (Sukthankar et al. 2014). The area of *Goal and Plan Recognition as Planning* (Ramírez and Geffner 2009) has advanced substantially over the past decade, yielding a number of approaches capable of coping with partial and noisy observations (E-Martín, R.-Moreno, and Smith 2015; Sohrabi, Riabov, and Udrea 2016), and doing this efficiently (Pereira, Oren, and Meneguzzi 2020).

Virtually, all such efforts use the model of Ramírez and Geffner (2010) as their underpinning, which defines via Bayes' Rule the probability of a goal, given observations in terms of the probability of the observations given the goal, and some prior probability of goals, representing an agent's preference. Comparatively, fewer efforts provide a probabilistic interpretation of the model defined by Ramírez and Geffner (Ramírez and Geffner 2010; Sohrabi, Riabov, and Udrea 2016; Kaminka, Vered, and Agmon 2018). Fewer efforts still actually use the prior probability on goals, assuming instead a uniform distribution for the goals, and ignoring the prior in their calculations. Ignoring the prior probability bakes into the goal recognition model the assumption that all goal recognition tasks are *one-shot*, such that agents pursue exactly one goal within a particular goal recognition domain exactly once. Such an assumption does not reflect many goal recognition tasks, such as intention recognition for elder care (Geib 2002), assistance for activities of daily living (Sim et al. 2010), proactive user interfaces (Amir and Gal 2013), among others.

In this paper, we expand goal recognition problems from the traditional *one-shot* setting used by all approaches so far, into problems that assume goal hypotheses have different probability distributions representing an agent's preferences and develop a solution for this problem by extending recent work on landmark-based goal recognition (Pereira, Oren, and Meneguzzi 2020). Our key contributions are twofold: (1) a novel definition of goal recognition problems with a distribution over goal preferences; and (2) a probabilistic interpretation that relies on the concept of landmarks.

2 Background

Planning

Planning is the problem of finding a sequence of actions (*i.e.*, a plan) that achieves a goal state 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_0, τ_1, ..., τ_n). An *operator* is represented by a triple $\mathrm{a} = \langle \mathrm{name(a)}, \mathrm{pre(a)}, \mathrm{eff(a)} \rangle$ where $\mathrm{name(a)}$ represents the description or signature of a; $\mathrm{pre(a)}$ describes the preconditions of a — a set of facts or predicates that must exist in the current state for a to be executed; $\mathrm{eff(a)} = \mathrm{eff(a)}^+ \cup \mathrm{eff(a)}^-$ represents the effects of a, with $\mathrm{eff(a)}^+$ an *add-list* of positive facts or predicates, and $\mathrm{eff(a)}^-$ 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 a triple $\Pi = \langle \Xi, \mathcal{I}, G \rangle$, where $\Xi = \langle \mathcal{F}, \mathcal{A} \rangle$ is a *planning domain definition*; $\mathcal{F}$ consists of a finite set of facts and $\mathcal{A}$ a finite set of actions; $\mathcal{I} \subseteq \mathcal{F}$ is the initial state; and $G \subseteq \mathcal{F}$ is the goal state. A *plan* is a sequence of actions $\pi = \langle \mathrm{a}_0, \mathrm{a}_1, ..., \mathrm{a}_n \rangle$ that modifies the initial state $\mathcal{I}$ into one in which the goal state G holds by the successive execution of actions in π. As in *Classical Planning*, actions have an associated cost, and here, we assume that this cost is 1 for all actions. A plan π is *optimal* if its cost is minimal.

Goal Recognition as Planning

Goal Recognition is the task of discerning the intended goal of autonomous agents or humans by observing their interactions in a particular environment (Sukthankar et al. 2014, Chapter 1). We formally define the problem of *Goal Recognition as Planning* by adopting the formalism of Ramírez and Geffner (2009; 2010), as follows in Definition 1.

Definition 1 (**Goal Recognition Problem**). *A goal recognition problem is a tuple $\Pi_{\mathcal{G}}^{\Omega} = \langle \Xi, \mathcal{I}, \mathcal{G}, \Omega \rangle$, where: $\Xi = \langle \mathcal{F}, \mathcal{A} \rangle$ is a planning domain definition; $\mathcal{I}$ is the initial state; $\mathcal{G} = \langle G_0, G_1, ..., G_n \rangle$ is the set of goal hypothesis, including the correct intended goal G^*, such that $G^* \in \mathcal{G}$; and $\Omega = \langle o_0, o_1, ..., o_n \rangle$ is an observation sequence of executed actions, with each observation $o_i \in \mathcal{A}$.*

The ideal solution for a goal recognition problem $\Pi_{\mathcal{G}}^{\Omega}$ is recognizing the correct intended goal $G^* \in \mathcal{G}$ that the observation sequence Ω of a plan execution achieves. An observation sequence Ω can be *full* or *partial*. A *full observation sequence* contains all actions of agents' plans, so all actions of a plan are observed, whereas in a *partial observation sequence*, only a sub-sequence of actions are observed.

Existing work on *Goal Recognition as Planning* considers the solution to a goal recognition problem to be either a *score system* associated to the goal hypothesis (Pereira, Oren, and Meneguzzi 2017, 2020), or a *probability distribution* for the goal hypothesis (Ramírez and Geffner 2009, 2010; E-Martín, R.-Moreno, and Smith 2015; Sohrabi, Riabov, and Udrea 2016). In this work, we extend a landmark-based approach for goal recognition and provide a probabilistic model that relies on the concept of *landmarks*.

3 Probabilistic Goal Recognition as Reasoning over Landmarks

Key to our probabilistic goal recognition approach is the concept of landmarks in planning, which has been extensively used in goal recognition approaches (Pereira, Oren, and Meneguzzi 2017; Vered et al. 2018; Pozanco et al. 2018; Shvo and McIlraith 2020). *Landmarks* are necessary facts (or actions) that must be true (or executed) at some point along all valid plans that achieve a particular goal from an initial state (Hoffmann, Porteous, and Sebastia 2004). Landmarks are often partially ordered based on the sequence in which they must be achieved. Hoffman *et al.* (2004) define fact 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 if and only if 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.*

The process of generating all landmarks and deciding their ordering is PSPACE-complete (Hoffmann, Porteous, and Sebastia 2004), which is exactly the same complexity as deciding plan existence (Bylander 1994). Thus, to operate efficiently, most landmark extraction algorithms (Hoffmann, Porteous, and Sebastia 2004; Silvia Richter 2008; Keyder, Richter, and Helmert 2010) extract only a subset of landmarks for a given planning instance.

In what follows, we expand the *Landmark-Based Goal Recognition* framework of Pereira, Oren, and Meneguzzi (2020) with a probabilistic interpretation that allows us to perform recognition repeatedly refining estimated goal probabilities over time. The recognition framework of Pereira, Oren, and Meneguzzi (2020) provides a *scoring system* that ranks the goal hypotheses $\mathcal{G}$ according to the ratio between the *number of achieved* landmarks and the *total number* of landmarks. Our probabilistic interpretation model follows the well-known probabilistic model of Ramírez and Geffner (2010). The probabilistic model of (Ramírez and Geffner 2010) sets the probability distribution for every goal G in the set of goals $\mathcal{G}$, and the observation sequence Ω to be a Bayesian posterior conditional probability, as follows:

$$\mathbb{P}[G \mid \Omega] = \alpha * \mathbb{P}[\Omega \mid G] * \mathbb{P}[G] \tag{1}$$

where $\mathbb{P}[G]$ is a *prior probability* to goal G, α is a *normalizing factor*, and $\mathbb{P}[\Omega \mid G]$ is the probability of observing Ω when the goal is G. Ramírez and Geffner (2010) compute $\mathbb{P}[\Omega \mid G]$ by computing two plans for every goal G, and based on these two plans, they compute a *cost-difference* between these plans and plug it into a Boltzmann equation. Basically, they compute a plan that *complies* with the observations, and another a plan that *does not comply* with the observations. The intuition of Ramírez and Geffner probabilistic model is that the lower the *cost-difference* for a goal, the higher the probability for this goal.

In contrast, our probabilistic model reasons over the evidence of landmarks, and follows the intuition of Pereira, Oren, and Meneguzzi (2020), where goals $\mathcal{G}$ are ranked according to their score, namely, the most likely goals are the ones that have achieved most of their landmarks in the observations. Thus, replicating this ranking in a probabilistic setting entails assigning probabilities to the observation of landmarks. If we consider an arbitrary goal G and represent its landmarks as a set $\mathcal{L}_G$, where $L_G \in \mathcal{L}_G$ is an individual landmark for G, we can reason about the probabilistic properties of observing such landmarks. First, since landmarks are *necessary conditions* to achieve a goal, the probability of observing all landmarks in a set of observations for a given goal should be 1, as we formally define in Equation 2.

$$\mathbb{P}[\mathcal{L}_G \mid G] = \sum_{L_G \in \mathcal{L}_G} \mathbb{P}[L_G \mid G] = 1 \tag{2}$$

Without any additional evidence, we can also infer that the probability of observing any given individual landmark in an observation sequence Ω should be uniformly distributed as shown in Equation 3.

$$\mathbb{P}[L_G \mid G] = \frac{1}{|\mathcal{L}_G|} \tag{3}$$

If we completely ignore the ordering of the landmarks in observations, and consider only the probabilities of observing landmarks, we can compute the probability of a particular set of observations Ω towards a goal G using Equation 4.

$$\mathbb{P}[\Omega \mid G] = \sum_{L_G \in (\mathcal{L}_G \cap \Omega)} \mathbb{P}[L_G \mid G] \tag{4}$$

Thus, we use landmarks as a proxy for the probability of the entire set of observations Ω given a goal G. We can plug $\mathbb{P}[\Omega \mid G]$ defined in Equation 4 into the Bayesian formulation of Ramírez and Geffner from Equation 1. Since we assume the set of goal hypotheses to be exhaustive and mutually exclusive, we can compute instead a *normalizing factor* α, which we obtain from Equation 5.

$$\alpha = \frac{1}{\sum_{G \in \mathcal{G}} \mathbb{P}[\Omega \mid G] * \mathbb{P}[G]} \tag{5}$$

When no priors $\mathbb{P}[G]$ are informed, we can assume that their distribution is uniform, and compute them through $\mathbb{P}[G] = \frac{1}{|\mathcal{G}|}$. In Section 4, we show how we infer prior probabilities by observing repeated goal recognition episodes.

4 Prior Estimation by Repeated Episodes

We now expand the probabilistic model of Section 3 to compute posterior goal probabilities when the prior goal probabilities follow a *non-uniform* distribution over repeated goal-recognition episodes. The resulting model allows us to converge towards the actual probability distribution that can be used as a prior for further goal recognition episodes. We formalize the extended version of such problem in Definition 3.

Definition 3 (**Repeated Goal Recognition Problem**). *A repeated goal recognition problem is a tuple* $\Pi^{\circlearrowleft}_{\mathcal{G}} = \langle \Xi, \mathcal{I}, \mathcal{G}, \Omega^{\mathcal{G}} \rangle$*, where:* $\Xi = \langle \mathcal{F}, \mathcal{A} \rangle$ *is a planning domain definition;* $\mathcal{I}$ *is the initial state;* $\mathcal{G} = \langle G_0, G_1, ..., G_n \rangle$ *is the set of goal hypothesis; and* $\Omega^{\mathcal{G}} = \{\Omega_0, \ldots, \Omega_n\}$ *is a set of observation sequences, where each* $\Omega_i \in \Omega^{\mathcal{G}}$ *is an observation sequence* $\langle o_0, o_1, \ldots, o_m \rangle$ *of executed actions, with each observation* $o_i \in \mathcal{A}$*. Observation sequences* Ω_i *are projections of plans* π_i *for planning tasks* $\langle \Xi, \mathcal{I}, G_i \rangle$ *such that the intended goal* $G_i \in \mathcal{G}$ *is drawn from a probability distribution* $\mathbb{P}[\mathcal{G}]$ *with probability* $\mathbb{P}[\mathcal{G} = G_i]$.

The solution for a repeated goal recognition problem is the correct probability distribution $\mathbb{P}[\mathcal{G}]$ that generated the set of observation sequences Ω in the problem of Definition 3. Here, $\mathbb{P}[\mathcal{G}]$ does not represent the result of a single episode of goal recognition, but rather the goal preferences of the agent under observation under repeated episodes.

Our prior estimation consists of processing each observation sequence $\Omega_i \in \Omega^{\mathcal{G}}$ and count the number of times we recognize each candidate goal as the actual goal of an observation sequence Ω_i. We recognize the goals of each observation sequence independently, ignoring any priors in order to avoid biasing the count, by simply running our probabilistic goal recognition model against each individual observation sequence. Thus, during the estimation of the priors, we do not use any partially computed prior probabilities in the recognizer (Line 4). After each run, we check whether we correctly recognize the goal for sample Ω_i (Line 5), which we do in a supervised way. Each correctly recognized goal G for a sample results in an increment of the corresponding counter $\mathcal{C}_G$. After repeating the process for all samples, we compute the prior for every candidate using the counter values and a form of *Laplace smoothing* (marquis de Laplace 1825) shown in Line 6, where k is the number of ghost samples we include to prevent any goal from having a probability of exactly 0. Algorithm 1 formally describes how our prior estimation process works.

Algorithm 1 Prior Estimation.

1: **function** ESTIMATEPRIOR($\Pi^{\circlearrowleft}_{\mathcal{G}}$)
2: $\quad\mathcal{C}_G \leftarrow 0$ for all $G \in \mathcal{G}$
3: $\quad$**for** $\Omega \in \Omega^{\mathcal{G}}$ **do**
4: $\quad\quad\mathbf{G} \leftarrow$ RECOGNIZE($\Pi^{\Omega}_{\mathcal{G}}$)
5: $\quad\quad$**if** $G^* \in \mathbf{G}$ **then** $\mathcal{C}_G \leftarrow \mathcal{C}_G + 1$ for all $G \in \mathbf{G}$
6: $\quad\mathbb{P}[G] \leftarrow \frac{k + \mathcal{C}_G}{(k * |\mathcal{G}|) + \sum_{G \in \mathcal{G}} \mathcal{C}_G}$ for all $G \in \mathcal{G}$
7: $\quad$**return** $\mathbb{P}[\mathcal{G}]$ $\quad\triangleright$ Return probability distribution.

5 Experiments and Evaluation

We empirically evaluate our probabilistic model over the recognition datasets from (Ramírez and Geffner 2009). These datasets comprise hundreds of recognition problems from four planning domains (BLOCKS-WORLD, EASY-IPC-GRID, INTRUSION-DETECTION, and LOGISTICS), having recognition problems with both partial and full observability. Recognition problems with partial observability have four observation levels: 10%, 30%, 50% and 70%.

Repeated Goal Recognition Setup

To evaluate our repeated goal recognition algorithm, we develop a recognition problem generator that generates a set of *samples* that comprises $\Omega^{\mathcal{G}}$ from a set of possible goal hypothesis $\mathcal{G}$. We use the datasets from (Ramírez and Geffner 2009) to generate the samples. The generator receives problems from that dataset, which serves as the basis for the samples, the number of samples to be generated, the observability level, and the probability distribution $\mathbb{P}[\mathcal{G}]$ over the set of goal hypotheses. This probability distribution guides the selection of the goal state for each sample. Each sample generated from a given problem has the same domain definition Ξ, initial state $\mathcal{I}$ and goal hypotheses $\mathcal{G}$ as the original problem from the dataset. We sample the goal state for each plan sample from the goal hypotheses using the probability distribution $\mathbb{P}[\mathcal{G}]$, which is only known to the generator.

We use two different probability distributions to generate such samples: a *normal* distribution with $\mu = 1$ and $\sigma = 0$, which we denote as NORMAL-SINGLE distribution, where all samples have the same goal state as the original problem; and a *normal* distribution, such that a single (preferred) goal G_i has $\mathbb{P}[G_i] = 0.5$, and the probabilities for other candidates follow a normal distribution with goals more similar to G_i have higher probability, resulting in a distribution with $\mu \approx 1.7$ and $\sigma \approx 2.4$. We denote this second probability distribution as NORMAL-DIVERSE.

Finally, we generate the observations for each sample. We use Fast Downward (Helmert 2011) to generate the plans from which we project the observations. For each sample i, we feed the Fast Downward planner with a planning task comprised of the domain definition Ξ and initial state $\mathcal{I}$ from the original problem that is being used to generate the samples and the goal state G_i drawn from the probability distribution $\mathbb{P}[\mathcal{G}]$ for that sample. We then take the plan returned by Fast Downward and randomly select actions comprising the observations for the sample. The number of selected actions is defined by the observability level informed to the generator.

For the experiments described in this paper, we generate 250 samples for each problem in the dataset. This number of samples shall provide a good diversity of combinations between observation sets (when dealing with partial observations) and goal states over the samples, which allows us to better evaluate our solution.

After all the samples have been generated, we run the recognizer for each sample, using the result to infer our priors, as explained in Section 4. We smooth the priors using the process described in Algorithm 1 with a k value of 1. We then insert the inferred priors into the original problem that was used to generate the samples, whose correct intended goal is the agent's preferred, and run the recognizer for that problem. We use the result from this run to compute our *recognition time*, *Accuracy*, *Spread in* $\mathcal{G}$ and Δ metrics, which we explain in the following section.

Evaluation Metrics

We use three metrics in our evaluation: *Accuracy* (**Acc** %), representing the fraction of problems in which the correct intended goal is among the goals with the highest posterior probability; *Spread in* $\mathcal{G}$ (**S in** $\mathcal{G}$), representing the average number of goals recognized as the most likely; and *recognition time* (**Time**) in seconds, representing the recognition time including the landmark extraction process.

We use two additional metrics when evaluating our probabilistic model with prior probabilities. **Max-Norm** is the largest difference between corresponding probabilities in the distribution that generated the samples and the estimated distribution of priors, used to evaluate the distance between these two distributions. If we can infer the priors exactly right, **Max-Norm** $= 0$. The second metric is a Δ metric, which is the difference between the $\mathbb{P}[G \mid \Omega]$ of the real goal when using priors and when not using priors and gives us an insight on how helpful the priors are in one-shot recognition.

Finally, we evaluate how the repeated process affects the accuracy over time, by running the recognizer for the original problem with priors after each sample is executed. We compute the average accuracy per number of samples aggregating results for all domains, and plot a graph that indicates the change in accuracy as the number of samples grows.

Goal Recognition Results

Table 1 shows the results for executions with **no priors** (traditional *one-shot* recognition, denoted as NO PRIORS), with priors generated through a **single-goal** samples distribution (NORMAL-SINGLE), and with priors generated through **normal** samples distribution (NORMAL-DIVERSE). We show the results for all four domains using the recognition datasets from (Ramírez and Geffner 2009). For each domain, we show the number of problems (under the domain name), the average number of candidate goals $|\mathcal{G}|$, the average number of extracted landmarks $|\mathcal{L}|$, and the average number of observations $|\Omega|$. For each of the three prior setups, we show recognition time, accuracy, and Spread in $\mathcal{G}$. As for the prior setups that use priors, we show results for two additional metrics: **Max-Norm** and Δ. We can see that when using **no priors** we achieve similar results (in terms of accuracy and Spread in $\mathcal{G}$) to the landmark-based approaches in (Pereira, Oren, and Meneguzzi 2020). However, we achieve much better results when using prior probabilities (NORMAL-SINGLE and NORMAL-DIVERSE columns in Table 1), as it simulates agents' preference using our prior estimation process. Naturally, the NORMAL-SINGLE distribution yields better results, as the agent always chooses the same intended goal in the samples.

We also see that the average **Max-Norm** value is higher for lower observability, especially for NORMAL-SINGLE distribution (on average). This difference might be due to the higher variability in the observations at lower observability, leading to multiple goals being recognized as possible in each recognition episode. This may result in the prior of an incorrect goal becoming higher, which affects the **Max-Norm**. Since in the NORMAL-SINGLE every goal that is not the preferred one has a probability of 0 in the generator's distribution, this causes the **Max-Norm** metric to be higher for this distribution. The Δ metric increases with the observability level. As the accuracy increases with the increase in observations, the probabilistic model is correct more often during the prior estimation process, which helps to increase the probability of the correct intended goal in the prior.

Note that, for BLOCKS-WORLD domain, our probabilistic model has lower accuracy in problems under the **normal** than the **no priors** in problems with full observations. We believe that it happens due to the problem being actually much harder under the normal distribution. For example, consider two goals: A and B. Consider that goal A is the most likely intended goal for a repeated recognition problem. Assume a set of samples in which A is the actual intended goal, and goals A and B are considered the most likely ones by our model during training. However, for the samples where B is the intended goal, only goal B is the most likely one. Therefore, in the resulting probability distribution, B's prior probability will tend to be greater than A's prior probability. The result will be a prior probability skewed towards B, misleading the model.

Finally, we analyze how the accuracy changes as the number of samples grows. Figure 1 shows the accuracy over time for the NORMAL-SINGLE distribution for each observability level. We can see an increase in accuracy right from the start as well as quick stabilization of the average accuracy over time, meaning that even with a moderate number of samples we can increase the accuracy significantly.

Figure 2, on the other hand, shows the accuracy over time for NORMAL-DIVERSE distribution. Although the increase in accuracy isn't as significant as in NORMAL-SINGLE dis-

					No Priors			Normal-Single					Normal-Diverse				
Domain	$\|\mathcal{G}\|$	$\|\mathcal{L}\|$	**% Obs**	Ω	**Time**	**Acc %**	**S in $\mathcal{G}$**	**Time**	**Acc %**	**S in $\mathcal{G}$**	**Max Norm**	Δ	**Time**	**Acc %**	**S in $\mathcal{G}$**	**Max Norm**	Δ
Blocks-World (793)	20.3	12.0	10	1.1	0.23	21.9%	1.3	0.265	68.3%	1.2	0.588	0.371	0.255	45.9%	1.1	0.312	0.193
			30	2.9	0.352	39.3%	1.2	0.204	96.7%	1.0	0.334	0.633	0.192	82.0%	1.0	0.197	0.323
			50	4.3	0.346	59.0%	1.2	0.159	96.7%	1.0	0.259	0.699	0.16	88.5%	1.0	0.155	0.365
			70	6.4	0.174	80.9%	1.2	0.286	97.8%	1.0	0.215	0.73	0.28	88.0%	1.0	0.124	0.398
			100	8.6	0.358	100.0%	1.5	0.292	100.0%	1.5	0.246	0.694	0.269	65.6%	1.0	0.162	0.361
Easy-Ipc-Grid (390)	8.3	6.8	10	1.8	0.413	71.1%	2.7	0.593	98.9%	1.1	0.336	0.494	0.634	73.3%	1.0	0.239	0.149
			30	4.4	0.474	86.7%	1.6	0.586	97.8%	1.0	0.213	0.584	0.567	96.7%	1.0	0.143	0.285
			50	7.0	0.637	96.7%	1.2	0.496	100.0%	1.0	0.156	0.613	0.465	100.0%	1.0	0.097	0.344
			70	9.8	0.379	98.9%	1.0	0.616	100.0%	1.0	0.09	0.639	0.639	100.0%	1.0	0.057	0.388
			100	13.4	0.438	100.0%	1.0	0.641	100.0%	1.0	0.028	0.662	0.61	100.0%	1.0	0.039	0.408
Intrusion-Detection (390)	16.7	13.8	10	1.9	0.478	75.6%	1.4	0.343	100.0%	1.0	0.274	0.591	0.368	100.0%	1.0	0.139	0.367
			30	4.5	0.491	94.4%	1.0	0.367	100.0%	1.0	0.082	0.733	0.377	100.0%	1.0	0.051	0.46
			50	6.7	0.467	100.0%	1.0	0.354	100.0%	1.0	0.06	0.745	0.385	100.0%	1.0	0.042	0.466
			70	9.5	0.46	100.0%	1.0	0.377	100.0%	1.0	0.058	0.745	0.4	100.0%	1.0	0.042	0.463
			100	13.1	0.524	100.0%	1.0	0.375	100.0%	1.0	0.058	0.735	0.441	100.0%	1.0	0.043	0.464
Logistics (390)	10.0	14.3	10	2.0	0.544	62.2%	2.0	0.497	100.0%	1.0	0.473	0.426	0.542	80.0%	1.0	0.262	0.151
			30	5.9	0.666	86.7%	1.3	0.508	100.0%	1.0	0.244	0.63	0.527	98.9%	1.0	0.132	0.321
			50	9.6	0.701	94.4%	1.1	0.463	100.0%	1.0	0.131	0.714	0.482	100.0%	1.0	0.076	0.396
			70	13.5	0.459	97.8%	1.0	0.548	100.0%	1.0	0.07	0.754	0.567	100.0%	1.0	0.048	0.429
			100	18.7	0.675	100.0%	1.0	0.563	100.0%	1.0	0.035	0.776	0.571	100.0%	1.0	0.031	0.461

Table 1: Experimental results comparing our landmark-based probabilistic model with *no priors*, priors estimating through repeated episodes with Normal-Single distribution and priors estimated through repeated episodes with Normal-Diverse distribution

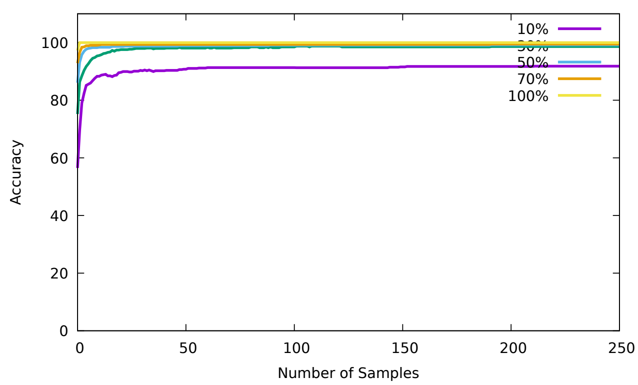

Figure 1: Accuracy over time (number of samples) with Normal-Single distribution.

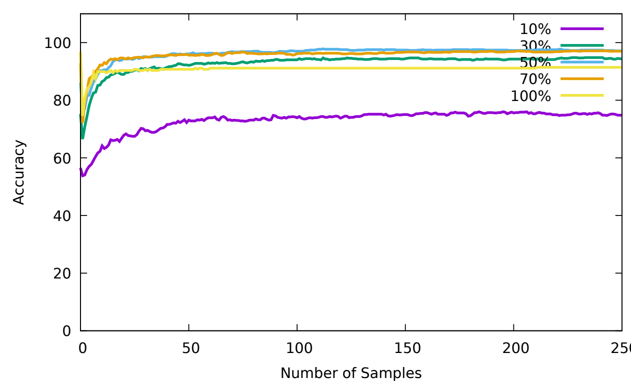

Figure 2: Accuracy over time (number of samples) with Normal-Diverse distribution.

tribution, the accuracy increases in the very early stages of the repeated goal recognition process. We note that for Normal-Diverse, the accuracy curve takes more samples to stabilize, showing that dealing with a more diverse distribution of goals takes more samples with a constant preference to obtain the preference insight.

6 Conclusions

In this paper, we have developed a novel probabilistic model for *Goal Recognition as Planning* that relies on the concept of landmarks, and a prior estimation process that infers prior probabilities from past recognition episodes. We have shown that our probabilistic model clearly benefits when using prior probabilities that have been inferred from past recognition episodes.

Our landmark-based probabilistic model can be used not only in *Classical Planning* settings, but also in other planning settings that define the concept of landmarks, i.e., *Temporal Planning* landmarks (Karpas et al. 2015), *Numeric Planning* landmarks (Scala et al. 2017). Our prior estimation mechanism is completely independent of the underlying goal recognition algorithm, and any such algorithm (even a non-probabilistic one) could be used in estimating the priors.

As future work, we intend to expand our prior estimation algorithm to non-classical planning settings, as well as to settings where the agent under observation is adversarial. An example of adversarial setting involves the agent deliberately choosing undesired goals to skew the prior probability away from the preference relation.

Acknowledgments

Felipe Meneguzzi acknowledges support from CNPq with projects 407058/2018-4 (Universal) and 302773/2019-3 (PQ Fellowship). Ramon Fraga Pereira acknowledges support from the ERC Advanced Grant WhiteMech (No. 834228) and the EU ICT-48 2020 project TAILOR (No. 952215).

References

Amir, O.; and Gal, Y. K. 2013. Plan Recognition and Visualization in Exploratory Learning Environments. *ACM Transactions on Interactive Intelligent Systems (TiiS)* 3(3): 16:1–16:23.

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 International Joint Conference on Artificial Intelligence (IJCAI)*.

Geib, C. W. 2002. Problems with Intent Recognition for Elder Care. In *Proceedings of the Conference of the Association for the Advancement of Artificial Intelligence (AAAI)*.

Ghallab, M.; Nau, D. S.; and Traverso, P. 2004. *Automated Planning - Theory and Practice.* Elsevier.

Helmert, M. 2011. The Fast Downward Planning System. *Computing Research Repository (CoRR)* abs/1109.6051.

Hoffmann, J.; Porteous, J.; and Sebastia, L. 2004. Ordered Landmarks in Planning. *Journal of Artificial Intelligence Research (JAIR)* 22(1): 215–278.

Kaminka, G. A.; Vered, M.; and Agmon, N. 2018. Plan Recognition in Continuous Domains. In *Proceedings of the Conference of the Association for the Advancement of Artificial Intelligence (AAAI)*.

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)*.

Keyder, E.; Richter, S.; and Helmert, M. 2010. Sound and Complete Landmarks for And/Or Graphs. In *ECAI*.

marquis de Laplace, P. S. 1825. *Essai philosophique sur les probabilités*. Bachelier.

Pereira, R. F.; Oren, N.; and Meneguzzi, F. 2017. Landmark-Based Heuristics for Goal Recognition. In *Proceedings of the Conference of the Association for the Advancement of Artificial Intelligence (AAAI)*.

Pereira, R. F.; Oren, N.; and Meneguzzi, F. 2020. Landmark-Based Approaches for Goal Recognition as Planning. *Artificial Intelligence* 279.

Pozanco, A.; E-Martín, Y.; Fernández, S.; and Borrajo, D. 2018. Counterplanning using Goal Recognition and Landmarks. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence*.

Ramírez, M.; and Geffner, H. 2009. Plan Recognition as Planning. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*.

Ramírez, M.; and Geffner, H. 2010. Probabilistic Plan Recognition Using Off-the-Shelf Classical Planners. In *Proceedings of the Conference of the Association for the Advancement of Artificial Intelligence (AAAI)*.

Scala, E.; Haslum, P.; Magazzeni, D.; and Thiébaux, S. 2017. Landmarks for Numeric Planning Problems. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence (IJCAI)*.

Shvo, M.; and McIlraith, S. A. 2020. Active Goal Recognition. In *Proceedings of the Conference of the Association for the Advancement of Artificial Intelligence (AAAI)*.

Silvia Richter, M. H. e. M. W. 2008. Landmarks Revisited. In *Proceedings of the Twenty-Third AAAI Conference on Artificial Intelligence (2008)*.

Sim, K.; Yap, G.; Phua, C.; Biswas, J.; Phyo Wai, A. A.; Tolstikov, A.; Huang, W.; and Yap, P. 2010. Improving the accuracy of erroneous-plan recognition system for Activities of Daily Living. In *The IEEE International Conference on e-Health Networking, Applications and Services*.

Sohrabi, S.; Riabov, A. V.; and Udrea, O. 2016. Plan Recognition as Planning Revisited. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*.

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

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 International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.