Cost-Based Goal Recognition for the Path-Planning Domain

Peta Masters and **Sebastian Sardina**
RMIT University, Melbourne, Australia
{peta.masters, sebastian.sardina}@rmit.edu.au

Abstract

"Plan recognition as planning" uses an off-the-shelf planner to perform goal recognition. In this paper, we apply the technique to path-planning. We show that a simpler formula provides an identical result in all but one set of conditions and, further, that identical ranking of goals by probability can be achieved without using any observations other than the agent's start location and where she is "now".

1 Introduction

Goal recognition (GR) is the problem of identifying an agent's intent by observing her behaviour. The traditional approach involves matching observations to a pre-existing plan in a *plan library* [Kautz and Allen, 1986; Charniak and Goldman, 1991]. Recent developments, however, dispense with this overhead and treat the problem instead as one of "planning in reverse" [Ramirez and Geffner, 2009; Baker *et al.*, 2009]. Using a classical planner and linking a goal's probability to the cost of a plan that achieves it, this approach draws on the intuition that a rational agent is most likely to be following the optimal (i.e., minimum cost) or *least sub-optimal* [Ramirez and Geffner, 2010] plan for goal.

In this paper,[1] we examine the probabilistic GR model for general task-planning, developed in [Ramirez and Geffner, 2010], and reframe it in the strict context of path-planning. We prove that a simpler account yields an almost identical result in less than half the time and with less computational effort. More surprisingly, we also show that a probability distribution that ranks candidate goals in the same order can be obtained *without referencing any observations* (other than the agent's starting point and current location).

2 Goal Recognition as Planning

We begin by presenting the technical background.

Ramirez and Geffner's GR framework operates in a classical task-planning environment. In STRIPS, a ***planning domain*** is a tuple $D = \langle F, A \rangle$, where F is a set of fluents and A a set of actions a, each with a precondition, add and delete list $Pre(a)$, $Add(a)$ and $Del(a)$, all subsets of F. Action a can occur in state s if $Pre(a) \subseteq s$. The initial state is assumed fully observable and the domain deterministic: if a occurs in s, a new state s' results such that $s' = (s \cup Add(a)) \setminus Del(a)$. A ***planning problem*** $\langle F, I, A, G \rangle$ is a planning domain with specified initial and goal states $I, G \subseteq F$, and its ***solution*** is a plan $\pi = a_1, .., a_k$ that maps I to G. Typically, each action has a cost $c(a)$. The cost of π is defined $cost(\pi) = \sum c(a_i)$ and an ***optimal plan*** is a solution with the lowest cost.

Ramirez and Geffner (2009) defined a ***GR problem*** $T = \langle D, \mathcal{G}, I, O \rangle$, where: D is a planning domain; $\mathcal{G} \neq \emptyset$ a set of possible goal states; I the initial state; and $O = o_1, .., o_k, k \geq 0, o_i \in A$, a sequence of observations. The solution to T is a set of goals, the optimal plans for which satisfy observations, and a plan $\pi = a_1, .., a_n$ ***satisfies observations*** $o_1, .., o_m$ if it embeds them in a way that preserves the order of actions.

A major drawback of the above framework is that it only identifies a goal *if observations conform to an optimal plan*, whereas, realistically, agents behave *sub*optimally; and in [Ramirez and Geffner, 2010] the authors presented an alternative *probabilistic* framework which addresses the problem.

A ***probabilistic GR problem*** is a GR problem plus a prior probability distribution. Its ***solution*** is a posterior probability distribution which prefers goals whose plans "best" satisfy observations. Basing their notion of "best" on the principle of rational action, Ramirez and Geffner (2010) capture this idea using the ***cost difference*** between the cheapest plan for goal, given the actions already observed, and the cheapest plan that *could* have achieved it, had the agent behaved differently. Comparing cost differences across goals, the authors arrive at a probability distribution with the following important property: *the lower the cost difference, the higher the probability*. Formally, cost difference is a function $costdif : 2^F \times A^* \mapsto \mathbb{R}$:

$$costdif(G, O) = optc(G, O) - optc^{\neg}(G, O), \qquad \text{(RG1)}$$

where $optc(G, O)$ and $optc^{\neg}(G, O)$ denote the optimal cost of plans that *do* and do *not* embed observations, respectively. The probability distribution itself is computed using:

$$P_X(G|O) = \alpha \frac{e^{-\beta X}}{1 + e^{-\beta X}}, \qquad \text{(RG2)}$$

where $X = costdif(G, O)$, α is a normalising constant across goals, and β a positive constant that modulates the shape of the distribution.[2]

[1] This paper is an abridgement of [Masters and Sardina, 2017] (with minor amendments), best student paper at AAMAS17.

[2] RG2, in code referenced from [Ramirez and Geffner, 2010], is

2.1 The Path-Planning Case

Reformulating the above approach in the context of path-planning, fluents become tied to locations and actions relate to movements. Whereas in task-planning, a solution plan is described as a sequence of actions, in path-planning, it becomes a *path*: a sequence of connected nodes in a graph.

Definition 1. A ***path-planning domain*** is a triple $\mathcal{D} = \langle N, E, c\rangle$ where: $N \neq \emptyset$ is a set of nodes (or locations); $E \subseteq N \times N$ is a set of edges between nodes; and $c : E \mapsto \mathbb{R}_0^+$ returns the non-negative cost of traversing each edge. ■

A ***path*** π in $\mathcal{D}$ is a sequence of locations $\pi = n_0, n_1, \ldots, n_k$ such that $(n_i, n_{i+1}) \in E$, for each $i \in \{0, 1, \ldots, k-1\}$. π^i denotes the i-th node n_i in π, and $|\pi|$ denotes the length of π, being the total number of edges k in π. $\pi(i,j) = \pi^i, \pi^{i+1}, \ldots \pi^j$ denotes the ***subpath*** of π from π^i to π^j (inclusive). The ***cost*** of a path is the cost of traversing all edges in π, that is, $\mathit{cost}(\pi) = \sum_{i=0}^{k-1} c(\pi^i, \pi^{i+1})$. The ***set of all paths*** in the domain is denoted by Π, and the set of all paths π starting at $\pi^0 = n_1$ and ending at $\pi^{|\pi|} = n_2$ is denoted by $\Pi(n_1, n_2)$.

Definition 2. A ***path-planning problem*** is a tuple $\mathcal{Q} = \langle \mathcal{D}, s, g\rangle$, where: $\mathcal{D}$ is the path-planning domain; $s \in N$ is the start location; and $g \in N$ is the goal location. ■

The ***solution*** to $\mathcal{Q}$ is a path π in $\mathcal{D}$ such that $\pi^0 = s, \pi^{|\pi|} = g$; the set of all of them being $\Pi(s, g)$. An ***optimal path*** is a solution with lowest cost. $\Pi^*(s, g)$ is the set of all such paths.

Waypoints are nodes that *must* be visited and a path *via waypoints* embeds them in a way that preserves their order. The ***optimal cost via waypoints*** W of a path from n_i to n_j is denoted by $\mathit{optc}(n_i, W, n_j)$. If $W = \emptyset$, we write $\mathit{optc}(n_i, n_j)$, and if $\pi^0 = w_0$ and $\pi^{|\pi|} = w_k$, we write $\mathit{optc}(W)$. We generalise the set of all solution paths $\Pi(s, g)$ to those embedding waypoints W as $\Pi(s, W, g)$. Similarly, $\Pi^*(s, W, g)$ denotes paths that are optimal w.r.t. cost among paths in $\Pi(s, W, g)$.

Definition 3. A ***path-planning GR problem*** is a tuple $\mathcal{P} = \langle \mathcal{D}, \mathcal{G}, s, O, \mathit{Prob}\rangle$, where: $\mathcal{D}$ is a path-planning domain; $\mathcal{G} \subseteq N$ is the set of possible goals; $s \in N$ is the start location; $O = o_1, \ldots, o_k$, $k \geq 0, o_i \in N$, is a sequence of locations (not actions, nor necessarily a path) where the agent has been observed; and *Prob* is the prior probabilities of the goals. ■

The solution to $\mathcal{P}$ is a probability distribution across $\mathcal{G}$ obtained using a reformulation of (RG1), grounded in path-planning as $\mathit{costdif}_{RG} : N \times N \times N^* \mapsto \mathbb{R}$:

$$\mathit{costdif}_{RG}(s, g, O) = \mathit{optc}(s, O, g) - \mathit{optc}^{\neg}(s, O, g).^3 \quad \text{(RG3)}$$

where $\mathit{optc}^{\neg}(s, O, g)$ denotes the optimal cost of navigating from s to g without embedding waypoints, that is:

$$\mathit{optc}^{\neg}(s, O, g) = \min_{\pi \in \Pi(s,g) \setminus \Pi(s,W,g)} \mathit{cost}(\pi).$$

Finally, the probability distribution is derived as for task-planning, by taking $X = \mathit{costdif}_{RG}(s, g, O)$ into Equation (RG2). For legibility, we call the resulting function P_{RG}:

$$P_{RG}(g|O, s) = \alpha \frac{e^{-\beta\, \mathit{costdif}_{RG}(s,g,O)}}{(1 + e^{-\beta\, \mathit{costdif}_{RG}(s,g,O)})}. \quad \text{(RG4)}$$

provably equivalent to the account there and in [Ramirez, 2012].

[3]We now make explicit the starting point s, implicit in (RG1).

3 A Simpler Cost Difference

We now present the first of our main technical contributions: a modified cost difference formula which is simpler to compute than (RG3) and faster to calculate. Based on the intuition that, in the great majority of cases, the optimal path that "does not pass through all observed locations" *is* the optimal path, instead of computing $\mathit{optc}^{\neg}(s, O, g)$ we deduct the more readily available "optimal path cost", arriving at:

$$\mathit{costdif}_1(s, g, O) = \mathit{optc}(s, O, g) - \mathit{optc}(s, g). \quad (1)$$

This alternative formulation is not only computationally less demanding (there is no requirement to reason negatively about the observations) but also, since the optimal path cost to each potential goal is not dependent on the observations, it can be pre-computed once at the outset. Note that if the potential start node and all candidate goal locations are known, as they are in the case of an airport terminal, for example, which has a fixed, finite number of entrances and boarding gates, then $\mathit{optc}(s, g)$ for all $g \in \mathcal{G}$ can be pre-computed and stored for retrieval as needed in constant time.

As the following theorems state, formula (1) generates an identical result to formula (RG3)in all cases bar one; and even then, the difference has minimal impact on the overall probability distribution. Moreover, in one corner-case, (1) actually enables calculation of a posterior probability distribution when the original, more involved, formula (RG3) may not.

Theorem 1. *Let O be an observation sequence such that $\mathit{optc}(s, O, g) > \mathit{optc}(s, g)$ (i.e., the observed behaviour is not optimal). Then, $\mathit{costdif}_{RG}(s, g, O) = \mathit{costdif}_1(s, g, O)$.*

In words, if the observed path is suboptimal—as, arguably, it would be most of the time—the simpler formula (1) yields *exactly the same* value as the original formula (RG3).

Theorem 2. *Let O be an observation sequence such that $\mathit{optc}(s, O, g) = \mathit{optc}(s, g)$ (i.e., the observed behaviour is optimal). If $\Pi^*(s, g) \setminus \Pi^*(s, O, g) \neq \emptyset$, then $\mathit{costdif}_{RG}(s, g, O) = \mathit{costdif}_1(s, g, O)$.*

In words, even if the observed behaviour is optimal, if there are other ways of behaving optimally, again formula (1) is *exactly equivalent* to the original formula (RG3).

Theorem 3 sets out the only case where the two formulas return different results. Consider the following example.

Example 1. One goal location is a house with front and back doors. An agent is observed at the front gate which is on the only optimal path. Using formula (1) the cost difference for this goal is zero (the observed path is the optimal path) but using formula (RG3) cost difference—assuming cost equates to distance—is a negative value (the optimal path *not* embedding observations is longer than the path that *does* embed them because it involves a detour to reach the back door). □

Theorem 3. *Let O be an observation sequence and $g \in \mathcal{G}$. Then, $\mathit{costdif}_{RG}(s, g, O) \neq \mathit{costdif}_1(s, g, O)$ iff $\Pi^*(s, O, g) = \Pi^*(s, g)$ (i.e., all optimal paths embed the observations).*

This strengthens Theorem 2 by asserting that both formulas yield identical results in all cases bar one: when observations are not only sufficient for optimal behaviour, but also *necessary*, i.e., there is no other way of acting fully rationally.

In considering this corner case, recall that cost difference is just a stepping stone towards generating a probability distribution. Often, we do not need to know exactly *how* probable the goals are, only their *relative order* or, more particularly, which goal is *most* probable. With this in mind, we prove (in Theorem 4) that, even if an agent is observed taking an ***exclusively optimal*** path to a goal (i.e., all optimal paths embed the observations), unless observations also conform to an optimal path for some *other* goal, the relative *ranking* of goals by probability is unaffected by use of the simpler cost difference formula, which still results in successful identification of the most probable goal.

Before proceeding, we make the following auxiliary observation, which formally restates the intuition that the lower the cost difference, the more probable the goal.

Observation 1. *Let $f(g, O)$ be a cost difference function and P_X a template of the probability distribution in* (RG2). *If $f(g_1, O) < f(g_2, O)$, then $P_f(g_1|O) > P_f(g_2|O)$.*

Now, let $P_{RG}(\cdot)$ be the probability distribution obtained from (RG2) when $X = \mathit{costdif}_{RG}(s, g, O)$ (Equation (RG3)) and $P_1(\cdot)$ the distribution obtained when $X = \mathit{costdif}_1(s, g, O)$ (Equation (1)).

Theorem 4. *Let O be an observation sequence and suppose that, for some potential goal $g \in G$, it is the case that:*

1. *$\Pi^*(s, O, g) = \Pi^*(s, g)$, that is, observations are exclusively optimal; and*
2. *for every $g' \in \mathcal{G} \setminus \{g\}$, $\mathit{optc}(s, O, g') > \mathit{optc}(s, g')$, that is, observations would result in suboptimal paths to all the other possible goals.*

Then, for all $g_1, g_2 \in \mathcal{G}$ and $g_1 \neq g_2$, $P_1(g_1|O) > P_1(g_2|O)$ if and only if $P_{RG}(g_1|O) > P_{RG}(g_2|O)$.

There remains one variation of exclusive optimality so far excepted. It is the case where observations coincide with the only optimal path to *multiple* goals (rather than one). We have argued that exclusive optimality for one goal is unusual; clearly, for multiple goals, it is even more so. Should the situation arise, however, the complex formula (RG3) would return multiple (negative) cost differences, which could be ranked, whereas the simple formula (1) ranks all goals for which observations match the optimal path equally.

Arguably, this situation is not only extremely unlikely, it also concerns *the very set of goals in which this probabilistic account [Ramirez and Geffner, 2010] is least interested.*

Finally, in the extreme case, where observations conform not just to the only optimal path to a goal g but to the *only* path per se, the cost of a path that does not conform to observations is infinite (because no such path exists). In this case, as Ramirez and Geffner (2010) point out, (RG3) ought to return $-\infty$ giving g the highest possible probability within the distribution. However, since $-\infty$ is not a number, the result *may* be undefined with the flow-on effect that normalised scores for the rest of the distribution may also be undefined.

In any practical implementation, of course, the problem is easily rectified by allocating some minimum value instead of $-\infty$ or treating this case separately. In an identical situation, however, Equation (1) (based on optimal cost from start to goal rather than "optimal cost given not the observations") returns zero and the issue does not arise.

To summarise the advantages of the simpler formula (1):

- $\mathit{costdif}_1(s, g, O)$ returns the same result as $\mathit{costdif}_{RG}(s, g, O)$ in all cases bar one;
- the one case where the formulas do not return identical results relates to fully rational behaviour, which is precisely *not* the motivation of the probabilistic GR model;
- even when $\mathit{costdif}_1(s, g, O)$ returns a different result, it is unlikely to impact the overall probability distribution;
- in the most extreme case—when no path to goal avoids the observations— it returns a meaningful result when $\mathit{costdif}_{RG}(s, g, O)$ may return *no* result at all; and
- $\mathit{costdif}_1(s, g, O)$ is computationally advantageous: it requires no 'negative' reasoning (so any standard path-planner can be used off-the-shelf); furthermore, in many domains, its second term may be pre-computed.

4 Single-Observation[4] Recognition

We now come to our second core contribution. As shown, in all but one extreme corner case, formula (1) can be used interchangeably with (RG3). Here, we go further and demonstrate that—if the starting location is known (e.g., because there is one common entrance to a building or terminal)—the ranking of goals, as judged by the probability distribution $P_1(\cdot)$ and generated using formula (1), can be achieved *without reference to the observation sequence*.

At first sight, the finding is counter-intuitive. Indeed, it implies that we can perform goal recognition without observing how the agent behaves over time! Nevertheless, if we know an agent's start location and the location of each candidate goal, we require only her *current* location in order to calculate a probability distribution within which goals are ranked in *exactly the same order* as if we had used formula (1).

Our single-observation formula, $\mathit{costdif}_2 : N \times N \times N \mapsto \mathbb{R}$ is defined as:

$$\mathit{costdif}_2(s, g, n) = \mathit{optc}(n, g) - \mathit{optc}(s, g), \qquad (2)$$

where n stands for the most recently observed location of the agent whose destination we are trying to determine (i.e., $n = O^{|O|}$). Let $P_2(\cdot)$ be the probability function obtained by taking $X = \mathit{costdif}_2(s, g, n)$ in (RG3).

Theorem 5. *Let O be an observation sequence. For all $g_1, g_2 \in \mathcal{G}$, $P_1(g_1|O) > P_1(g_2|O)$ iff $P_2(g_1|O) > P_2(g_2)$.*

Proof. From Observation 1, $P_1(g_1|O) > P_1(g_2|O)$ if and only if $\mathit{costdif}_1(s, g_1, O) < \mathit{costdif}_1(s, g_2, O)$. Recall, from Equation (1), that for each $i \in \{1, 2\}$:

$$\mathit{costdif}_1(s, g_i, O) = \mathit{optc}(s, O, g_i) - \mathit{optc}(s, g_i),$$

where the first term can be written as:

$$\mathit{optc}(s, O, g_i) = \mathit{optc}(s, O^0) + \mathit{optc}(O) + \mathit{optc}(O^{|O|}, g_i).$$

[4] We previously used the term "observation-free". However, the method does depend on observing one location (usually the agent's current location) with respect to which probabilities are required.

Let $n_l = O^{|O|}$ be the last observation in O. From Observation 1, recall that the relative ranking between g_1 and g_2 with respect to their posterior probabilities can be deduced directly from the relative value of their cost difference formulas. So, let us expand that value:

$$\begin{aligned}
&\mathit{costdif}_1(s, g_1, O) - \mathit{costdif}_1(s, g_2, O) \\
&\quad = [\mathit{optc}(s, O^0) + \mathit{optc}(O) + \mathit{optc}(n_l, g_1) - \mathit{optc}(s, g_1)] - \\
&\qquad\quad [\mathit{optc}(s, O^0) + \mathit{optc}(O) + \mathit{optc}(n_l, g_2) - \mathit{optc}(s, g_2)] \\
&\quad = \mathit{optc}(s, O^0) + \mathit{optc}(O) + \mathit{optc}(n_l, g_1) - \mathit{optc}(s, g_1) - \\
&\qquad\quad \mathit{optc}(s, O^0) - \mathit{optc}(O) - \mathit{optc}(n_l, g_2) + \mathit{optc}(s, g_2) \\
&\quad = \mathit{optc}(n_l, g_1) - \mathit{optc}(s, g_1) - \mathit{optc}(n_l, g_2) + \mathit{optc}(s, g_2) \\
&\quad = [\mathit{optc}(n_l, g_1) - \mathit{optc}(s, g_1)] - [\mathit{optc}(n_l, g_2) - \mathit{optc}(s, g_2)] \\
&\quad = \mathit{costdif}_2(s, g_1, n_l) - \mathit{costdif}_2(s, g_2, n_l).
\end{aligned}$$

It follows then that $\mathit{costdif}_1(s, g_1, O) > \mathit{costdif}_1(s, g_2, O)$ iff $\mathit{costdif}_2(s, g_1, O) > \mathit{costdif}_2(s, g_2, O)$. Thus, from Observation 1, $P_1(g_1|O) > P_1(g_2|O)$ iff $P_2(g_1|O) > P_2(g_2)$. □

The finding is useful and unexpected. All parameters are independent of the observation sequence (modulo where the agent is "*now*") and can be obtained using a *standard* path-planner: no specialised path-finding system is needed to reason about observations. Furthermore, if all start and candidate goal locations are known—as would typically be the case in most domains—formula $\mathit{costdif}_2(s, g, n)$ can be fully *pre-computed offline* for any node $n \in N$ in the domain.

The implications are significant. We can create a sort of "heat map" of the domain, showing the probability of each goal according to where the agent entered. If we have a goal of interest (e.g., a valuable location to monitor and protect), we can focus attention fully on *locations where that goal becomes the most probable*. That is, rather than tracking an agent's movements all over the terrain, we can just monitor the "hot" spots and only start tracking her in earnest if she arrives at one of them.

We close by noting that the above result is not dependent on formula (RG4) but is applicable whatever manipulation is used to derive the probability distribution, provided it satisfies the property that the lower the cost difference, the higher the probability, and relative cost differences are preserved.

5 Experimental Results

We tested formulas (RG3), (1) and (2) on 990 problems adapted from the well-known Moving-AI[5] path-planning benchmarks [Sturtevant, 2012] to confirm *(a)* that the case of exclusive optimality (as in Theorem 3) is rare and that otherwise the simpler formula yields identical posterior probability distributions to the more complex formula; *(b)* that all three accounts return posterior probability distributions that rank goals the same; and *(c)* using either of the modified formulas presented here cuts processing time by more than half.

We adapted the Moving-AI problems for GR by adding candidate goals at random locations, used Weighted-A* [Pohl, 1970] to generate a continuous path to the real goal then extracted observation sequences varied by: path quality (optimal, suboptimal, greedy), observation density (sparse 20%, medium 50%, dense 80%) and distribution (*random* locations along the path or a consecutive *prefix* beginning at the start).

[5]http://movingai.com/

	P_{RG}	P_1		P_2^*		
Obs	Time	Time	Match	Time	Match	Δ
20%P	24.596	3.835	100%	1.780	34.6%	0.040
20%R	10.346	2.112	100%	2.062	53.8%	0.025
50%P	31.934	1.706	100%	1.681	38.5%	0.034
50%R	18.962	2.110	100%	2.073	57.7%	0.024
80%P	33.274	1.876	100%	1.848	42.3%	0.035
80%R	26.073	2.111	100%	2.062	57.7%	0.023

Table 1: Landscapes - suboptimal

156 problems. Average goals: 4.35. Average path cost: 280.74. Δ displays average difference for non-matches. We obtain P_2^* as P_2 but adding a constant to each cost difference.

We calculated optimal path costs using standard A* [Hart *et al.*, 1968][6] and obtained optimal costs "given *not* the observations" by modifying A* so that it pruned from consideration any path that encountered *all* observed locations.

Our results confirmed the hypotheses. Importantly, the corner case—where observations conform to the only optimal path—did not arise in any randomly generated scenario. Table 1 summarises a representative subset of our results. We found that formula (RG3) performed even more slowly than expected (perhaps reflecting the calculation's inherent complexity): in some cases the simpler formula (1) performed twelve times faster. Probabilities based on formula (1) always exactly matched those based on (RG3); probabilities generated using formula (2) were usually different but *in all cases* use of the single-observation formula successfully identified the same goal as having the highest, or equal highest, posterior probability as either other formula.

6 Conclusion

We have examined GR techniques introduced by Ramirez and Geffner (2010) and applied them in the context of path-planning. We have shown that a simpler cost difference formula returns an *identical result* to the original in *all but one case*, which we characterise. We argue, in line with intuitions expressed in [Ramirez and Geffner, 2010], that this is a case of little interest and, in fact, it did not even arise in automated tests. Further, we have presented an alternative formula (2) that *does not depend on the observation sequence* but nevertheless generates a posterior probability distribution that *exactly preserves* the ranking of goals from the simplified account and, by extension, results in an *identical ordering* to formula (RG3) in all cases bar one. This formula has the benefit that it can be pre-computed in many realistic domains. So one can create a sort of "heat map" of posterior goal probabilities from which to identify a perimeter that should be monitored around any goal of interest, knowing that, at all locations within the perimeter, that goal is the most probable.

Acknowledgements

Australian Research Council (DP120100332) and Australian Government Research Training Program Scholarship.

[6]We used a Python-based infrastructure, designed as a testbed for path-planning algorithms, which already included implementations of A* and Weighted A* in its library (https://tinyurl.com/p4sim).

References

[Baker *et al.*, 2009] C. Baker, R. Saxe, and J.B. Tenenbaum. Action understanding as inverse planning. *Cognition*, 113(3):329–349, 2009.

[Charniak and Goldman, 1991] E. Charniak and R. Goldman. Probabilistic abduction for plan recognition. Technical report, Brown University and Tulane University, 1991.

[Hart *et al.*, 1968] P. E. Hart, N. J. Nilsson, and B. Raphael. A formal basis for the heuristic determination of minimum cost paths. *IEEE Transactions of Systems Science and Cybernetics*, 4(2):100–107, 1968.

[Kautz and Allen, 1986] H. Kautz and J. Allen. Generalized plan recognition. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 32–37, 1986.

[Masters and Sardina, 2017] Peta Masters and Sebastian Sardina. Cost-based goal recognition for path-planning. In *Proceedings of the International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 750–758, 2017.

[Pohl, 1970] Ira Pohl. Heuristic search viewed as path finding in a graph. *Artificial Intelligence*, 1(3-4):193–204, 1970.

[Ramirez and Geffner, 2009] Miquel Ramirez and Hector Geffner. Plan recognition as planning. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1778–1783, 2009.

[Ramirez and Geffner, 2010] Miquel Ramirez and Hector Geffner. Probabilistic plan recognition using off-the-shelf classical planners. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 1121–1126, 2010.

[Ramirez, 2012] Miquel Ramirez. *Plan Recognition as Planning*. PhD thesis, Univeristat Pompeu Fabra, 2012.

[Sturtevant, 2012] Nathan R Sturtevant. Benchmarks for grid-based pathfinding. *IEEE Transactions on Computational Intelligence and AI in Games*, 4(2):144–148, 2012.