A Survey on Goal Recognition as Planning

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

Abstract

Goal Recognition is the task of inferring an agent's goal, from a set of hypotheses, given a model of the environment dynamic, and a sequence of observations of such agent's behavior. While research on this problem gathered momentum as an offshoot of plan recognition, recent research has established it as a major subject of research on its own, leading to numerous new approaches that both expand the expressivity of domains in which to perform goal recognition and substantial advances to the state-of-the-art on established domain types. In this survey, we focus on the advances to goal recognition achieved in the last decade, categorizing the resulting techniques and identifying a number of opportunities for further breakthrough research.

1 Introduction

Goal Recognition is a task related to automated planning, where an agent employs abductive reasoning to infer the most likely desired goal from a series of observations of the observed agent's plan instead of deducing a plan from an initial state towards a goal using some kind of domain theory. While goal recognition is related to the problem of *Plan Recognition* [Mirsky *et al.*, 2021], which consists of trying to infer the actual plan being carried out by the observed agent, this survey focuses squarely on goal recognition. The task of goal recognition has a number of potential and actual applications, including assisting the handicapped [Geib, 2002], activities of daily living (e.g., cooking) [Granada *et al.*, 2017], workplace safety [Inam *et al.*, 2018], smart home [Hegde and Kenchannavar, 2019], among others [Singh *et al.*, 2020; Wayllace *et al.*, 2020a]. As research on goal recognition evolves into ever more complex domain models and better approaches, we expect the state-of-the-art to advance substantially and new application domains to be broken into.

Thus, we survey advances on goal recognition achieved primarily in the last decade, defining a common formal framework in Section 2, and exploring key aspects that distinguish current approaches. We formalize types of domain models that modern goal recognition approaches process in Section 3, followed by an analysis of the effect of different types of observations in Section 4. Section 5 looks into the assumptions each recognition approach makes about the agent being observed, especially regarding the awareness of the agent to the observer and how optimal the observed agent is. We evaluate the internal mechanism through which goal recognition approaches infer goals from observations in Section 6. Finally, we discuss various problems related to goal recognition in Section 7, and organize the approaches surveyed in this paper, pointing how future work can solve many of the limitations in the current state-of-the-art in Section 8.

2 Goal Recognition as Planning

We lay out the formal foundations of the problem we survey using a top-down approach and frame goal recognition in the context of automated planning, outlining the basic environment model, which we refine throughout the paper to match the specific approaches we discuss. We then formalize the task of goal recognition itself based on a planning model, which we also refine throughout the survey.

2.1 Automated Planning

Planning aims to select what an agent does next, given a model of the environment specifying how actions and sensors work, a current situation, and what is the goal to be achieved [Geffner and Bonet, 2013]. Such problem can be seen as a directed graph in which nodes represent states, edges represent the transition between states (caused by applying actions), and the solution is a path (i.e., plan) between two particular nodes (i.e., initial state and goal state) in this directed graph. Such graph is induced by a planning task following Definition 1.

Definition 1 (Planning Task) *A planning task* $\Pi = \langle \Xi, s_0, G \rangle$ *is a tuple composed of a domain definition* Ξ*, an initial state* s_0*, and a goal state specification* G*. A solution to a planning task is a plan or policy* π *that reaches a goal state* G *starting from the initial state* s_0 *by following the transitions defined in the domain definition* Ξ*.*

2.2 Goal Recognition

Goal Recognition is the task of recognizing which goal an agent aims to achieve by observing its interactions in an environment [Sukthankar *et al.*, 2014, Chapter 1, Page 3]. Such observed interactions (i.e., observations) constitute the evidence to recognize goals, and can be executed actions in the environment (e.g., a simple movement, cook, drive), and changing properties in an environment (e.g., at home, at work, resting). In this survey, we consider a goal recognition task by following the original problem defined by Ramírez and Geffner [2009]. As we formalize in Definition 2, a goal recognition task follows the same structure of a planning task (Definition 1), but instead of having a single goal condition, it has a set of *goal hypotheses*, and a sequence of observations induced by the actions executed by the observed agent. Definition 2 does not detail what a solution to a goal recognition problem is, as different approaches provide different solution concepts, which we provide in Section 6.2.

Definition 2 (Goal Recognition Task) *A goal recognition task $\Pi_{\mathcal{G}}^{\Omega} = \langle \Xi, s_0, \mathcal{G}, \Omega \rangle$ is a tuple composed of a domain definition Ξ, an initial state s_0, a set of goal hypotheses $\mathcal{G}$, and a sequence of observations Ω.*

Most goal recognition tasks assume there is at least one *correct intended goal* $G^* \in \mathcal{G}$ among the hypotheses. However, as we see in Section 6.1, how exactly to present such solutions varies with the specific approach. Our definitions so far are vague into exactly what constitutes a domain definition, and what form the observations take. This is a deliberate presentation choice, as goal recognition approaches assume a variety of different models and observation types. Thus, we use Ξ to refer to any type of domain model, and pin down these types in Section 3. Likewise, we use Ω to refer to any observation sequence so we can define recognition problems independently of the observations, which we detail in Section 4. Example 1 illustrates a goal recognition task.

Example 1 *Figure 1 illustrates a navigation environment where a robot can move and pick up objects to achieve its goals. The goal recognition task presented in the example is composed of an initial state s_0 (e.g., robot position), a set of goal hypotheses $\mathcal{G} = \langle G_0, G_1, G_2 \rangle$, and a sequence of observations Ω (e.g., a sequence of movements, represented by the dashed blue arrows). The domain definition Ξ can vary according to the domain type used to formalize the environment dynamics, e.g., a deterministic discrete domain model, a stochastic domain model, among others. For this example, we consider that Ξ is a deterministic discrete domain model. A goal state specification G for this example can be defined as the robot be at a certain location holding an object (e.g., G_1 represents the robot holding the ball in the upper right corner of the environment). Thus, based on the goal recognition task depicted in Figure 1, it is possible to infer that, most likely, the robot aims to achieve the goal condition G_2 (e.g., holding the box in the bottom right corner of the environment).*

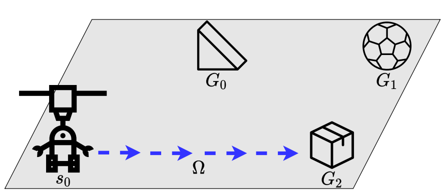

Figure 1: Goal Recognition example.

3 Domain Model Types

To the best of our knowledge, the first practical goal recognition approach using a domain model recognizable by current approaches goes back to Lesh and Etzioni [1995]. This approach, instead of abstracting plans to strings from a grammar, uses a state transition system induced by a planning domain much like that of Definition 1. Much of the early plan and goal recognition approaches focused treating the problem as one of parsing, assuming the domain knowledge includes not only the state transition function, but also production rules restricting the space of possible plans [Geib and Steedman, 2007]. In effect, such planning domains induce graphs $\langle \mathcal{S}, \mathcal{A} \rangle$ via their transition function γ, with vertices $\mathcal{S}$ and edges $\mathcal{A}$ as in [Hong, 2001]. Definition 1 abstracts the specifics of the domain definition and what constitutes a state and a transition system. Indeed, Ξ induces a number of elements of a planning task, including a state space $\mathcal{S}$, an action space $\mathcal{A}$, and a transition function γ. Throughout this survey we refine domain model types defining what exactly states, actions, and transition systems can be.

3.1 Classical Planning Domain Models

The most common domain model in goal recognition research consists of the STRIPS fragment of PDDL [Fox and Long, 2003a], as used by the seminal work of Ramírez and Geffner [2009] and much subsequent work [Martín *et al.*, 2015; Sohrabi *et al.*, 2016; Pereira *et al.*, 2017; Pereira *et al.*, 2020][1]. In *Classical Planning*, the state space is often represented as a set of propositional facts $\mathcal{F}$ (i.e., instantiated predicates) denoting what is true in the environment, or, alternatively, as multivalued variables in a finite domain representation [Helmert, 2009]. Regardless of the specific representation, a state comprises a set of variables $\mathcal{V}$ with a finite domain and the state space $\mathcal{S}$ is finite and well-defined. Likewise, in this setting, the action space $\mathcal{A}$ comprises STRIPS-style actions of the form $a = \langle \mathrm{name}(a), \mathrm{pre}(a), \mathrm{post}(a), \mathrm{cost}(a) \rangle$ with conjunctive preconditions $\mathrm{pre}(a)$, and postconditions $\mathrm{post}(a) = \mathrm{post}^+(a) \cup \mathrm{post}^-(a)$ comprising both positive $\mathrm{post}^+(a)$ and negated facts $\mathrm{post}^-(a)$. Such actions induce a transition function $\gamma : \mathcal{S} \times \mathcal{A} \mapsto \mathcal{S}$ such that $\gamma(s, a) = s'$, where $s' = (s \cup \mathrm{post}^+(a)) - \mathrm{post}^-(a)$. For example, for the navigation environment depicted in Example 1, a *Classical Planning* task Π comprises a set of instantiated facts, representing the state properties of this navigation environment, where the robot and the objects are situated (e.g., `(at robot loc-0-0)`, `(at box loc-0-5)`), and what objects the robot is holding (e.g., `(holding ball)`), the action space $\mathcal{A} = \langle$`move, pick-up`$\rangle$ (we omit pre and post^+ due to lacking of space), $s_0 = \langle$`(at robot loc-0-0) (at box loc-0-5)`$\ldots\rangle$, and $G_0 = \langle$`(and (at robot loc-2-5) (holding ball))`$\rangle$.

3.2 Stochastic Domain Models

By contrast, *stochastic* domains are often modeled in terms of (Partially Observable) Markov Decision Processes (POMDPs and MDPs), as in the approaches of Ramírez

[1]While space limitations prevent us from providing details of the PDDL format, we refer to Haslum *et al.* [2019]

and Geffner [2011] and Oh *et al.* [2011]. In such models, MDPs transition functions γ are stochastic and such that $s' \sim \gamma(s, a)$, i.e., subsequent states are sampled following a stochastic function $\mathbb{P}[s' \mid s, a]$ for $s, s' \in \mathcal{S}$, $a \in \mathcal{A}$.

3.3 Continuous Domain Models

Vered and Kaminka [2017a; 2017b] formally define the task of goal recognition by using the formalism of *Continuous Motion Planning*. In this setting, the state space S of a domain model Ξ is represented in a multi-dimensional Euclidean space, where $S \subset \mathbb{R}^n$, such that $n \geq 2$. This type of domain model usually represents environments with two or three dimensions. The action space $\mathcal{A}$ is a discrete (possibly infinite) set of actions, encoding a transition function γ between states. This transition function γ allows transforming one state (s) into another (s') via *paths* through the state space (e.g., a non-negative straight-line from s to s'), rather than a discrete state, as defined in *Classical Planning*. We note that other recognition approaches in the literature also adopt this type of domain model, such as [Masters and Sardina, 2017; Vered *et al.*, 2018; Kaminka *et al.*, 2018; Masters and Sardina, 2019a; Masters and Sardina, 2019b]. For the navigation environment depicted in Example 1, a *Continuous Motion Planning* task Π could be defined with 3 dimensions (x, y, z), so the actions space $\mathcal{A} = \langle \texttt{move, pick-up} \rangle$ would be performed over continuous values in x, y, and z, whereas s_0 and G would be represented as a pose (e.g., position and orientation) of the robot and the objects in the environment, for instance, the robot pose would be $x(r) = 0.28, y(r) = 1.44, z(r) = 0.93, \theta(r) = 45°$.

3.4 Optimal Control Domain Models

Optimal Control aims to control a dynamical system such that its output follows a desired value, which may be a fixed or changing value [Bertsekas, 2017]. Pereira *et al.* [2019b] use concepts of *Optimal Control* for recognizing goals in *approximate* continuous domain models. To model *Optimal Control* problems and the range of possible agent behavior, the authors adopt the formalism of Finite-Horizon Optimal Control (FHOC) [Bertsekas, 2017] problems, incorporating and combining some terminologies of *Control* [Borrelli *et al.*, 2017] and *Automated Planning* [Geffner and Bonet, 2013] to account for constraints and goal conditions (also referred to as *target regions* in the *Control* literature).

In this setting, transitions between states follow a *stationary*, *discrete–time* dynamical system $x_{k+1} = \gamma(x_k, u_k, w_k)$, where for each time point $k \in [0, N]$, x_k is the *state*, u_k is the *control* input and w_k is a *random variable* with a probability distribution that does not depend on past w_j, $j < k$. States $x_k \in S$, controls $u_k \in U$, and disturbances $w_k \in W$ are required to be part of $\mathbb{R}^n$. For extracting trajectories, agents seek to transform initial states x_0 into states x_N with specific properties. These properties are given as logical formulas over the components of states x_k, and the set of states $S_G \subseteq S$ are those where the desired property G, or goal, holds. The preferences of observed agents to pursue specific trajectories are accounted for with cost functions of the form $J(x_0) = \mathbb{E}\{g(x_N) + \sum_{k=0}^{N-1} g(x_k, u_k, w_k)\}$, where $g(x_N)$ is the *terminal* cost, and $g(x_k, u_k, w_k)$ is the *stage* cost. FHOC problems are defined as an optimization problem whose solutions are trajectories (i.e., policies) that describe the range of possible optimal behaviors of observed agents.

3.5 Approximate, Incomplete, and Learned Domain Models

Regardless of the nature of the underlying model, one of the key bottlenecks in planning, and by extension, goal recognition, is the acquisition of a model Π from either structured or unstructured data. While research on reinforcement learning has yielded reliable model acquisition techniques for stochastic domains driven by reward functions [Sutton and Barto, 2018], relatively less emphasis has been given to the automatic learning of *Classical Planning* models, with few techniques capable of generating the type of lifted PDDL models used by most of the approaches described in this survey. Recent work yielded mechanisms to learn classical planning models from structured [Aineto *et al.*, 2019b; Suárez-Hernández *et al.*, 2020] and unstructured data [Asai and Fukunaga, 2018].

These advances have led to a number of approaches for goal recognition using automatically derived models. Amado *et al.* [2018] develop a goal recognition approach that *learns* PDDL models from *unstructured data* allowing symbolic approaches [Pereira *et al.*, 2017] to quickly solve problems encoded in images. Similarly, Pereira *et al.* [2019a] adapt existing landmark-based goal recognition approaches to deal with *incomplete* STRIPS domain models [Weber and Bryce, 2011; Nguyen *et al.*, 2017], which assume that the model available to the recognizer might contain imprecisions (possible preconditions and effects) due to imperfect model acquisition. Pereira *et al.* [2019b] develop a goal recognition approach for continuous control domains that uses an *approximate learned* transition function and an optimal policy to infer goals by adapting existing symbolic techniques [Vered *et al.*, 2016]. Both approaches achieve a level of accuracy comparable to traditional approaches that rely on *authored domains*[2], indicating that traditional recognition approaches can be augmented to cope with learned models, even when the models themselves might have flaws from the learning process.

4 Observation Types

Observations constitute the key piece of evidence for agent behavior in goal recognition problems. In practice, observations are an indirect projection of an observed agent's behavior, and can be thought of as whatever the sensing capabilities of the agent performing goal recognition can provide. As we see in Section 6.2, the model we assume drives agent behavior influences the nature of how observations are generated. Such dependency on the model leads us to use Definition 3 as an intermediary structure between an agent's behavior model and the generation of observations.

Definition 3 (Trajectories) *Let π be a solution to a planning task Π, $\vec{\pi}$ is a trajectory induced by π. For Classical Planning tasks, $\vec{\pi}$ is a sequential plan π; for MDPs and stochastic*

[2]We refer to *authored domains* as domain models written by a human using a specification language like PDDL.

problems $\vec{\pi}$ is sequence of transitions from s_0 to G following policy π while complying with the transition function; and for Continuous Planning or Control tasks, $\vec{\pi}$ is a sequence of poses or configurations (e.g., coordinates $\langle x, y, z\rangle$ and a velocity) between which there are infinite intermediary poses.

Having defined trajectories, we can formally define how actual observations are generated in Definition 4.

Definition 4 (Observations) *Let $\vec{\pi} = \langle \vec{a}_1, \ldots, \vec{a}_n\rangle$ be a trajectory for planning task $\Pi = \langle \Xi, s_0, G\rangle$. An action projection function $tp(s, \vec{a}) : \mathcal{S} \times \mathcal{A} \mapsto \vec{\mathcal{A}}$ is a function that maps actions to sequences of zero or more observations. An observation sequence generation function $op(s_0, \vec{\pi})$ is a function mapping a trajectory $\vec{\pi}$ into an observation sequence:*

$$op(s_0, \vec{\pi}) = \begin{cases} \langle\rangle & \text{if } \vec{\pi} = \langle\rangle \\ \langle tp(s_0, \vec{a}_1)\rangle \cdot op(\langle \vec{a}_2, \ldots, \vec{a}_n\rangle) & \text{if } \vec{\pi} = \langle \vec{a}_1, \ldots, \vec{a}_n\rangle \end{cases}$$

Then, a sequence of observations Ω *is a sequence projected from $\vec{\pi}$ into observation state $\vec{\mathcal{A}}$ maintaining their order.*

The sensors available to the observer imposes limitations on types of observations. While most approaches construe observations as action descriptions of some sort, this is not necessarily the most realistic abstraction of how sensor data may be supplied to a goal recognizer. For instance, modern computer vision approaches to object detection [Wang *et al.*, 2020] may provide a more reliable representation of the current status of objects in an image than action identification. Indeed, recent work on goal recognition as planning [Sohrabi *et al.*, 2016] explicitly handles observations as states, the difference of which we formally define as follows.

Definition 5 (Action and State Observations) *Let $\vec{\pi} = \langle \vec{a}_1, \ldots, \vec{a}_n\rangle$ be a trajectory for the planning task Π and $tp(\cdot, \cdot)$ be an action projection. A sequence of observations $\Omega_{\vec{\pi}}$ is a sequence of actions from $\vec{\pi}$ if $tp(\cdot, \cdot) : \mathcal{S} \times \mathcal{A} \mapsto \mathcal{A}$, i.e., if $\vec{\mathcal{A}} = \mathcal{A}$. Alternatively, a sequence of state observations Ω_s is a sequence of states from $\mathrm{S}_{\vec{\pi}}$ if $tp(\cdot, \cdot) : \mathcal{S} \times \mathcal{A} \mapsto \mathcal{S}$, i.e., if $\vec{\mathcal{A}} = \mathcal{S}$. We denote the observation corresponding to action a_i as $\vec{a_i}$ and that corresponding to state s_i as $\vec{s_i}$*

For the navigation environment depicted in Example 1, action observations in the *Classical Planning* setting could be $\langle \vec{a_0} = \texttt{(move loc-0-0 loc-0-1)}, \vec{a_1} = \texttt{(move loc-0-1 loc-0-2)}, \ldots\rangle$, and state observations in *Continuous Motion Planning* setting could be $\langle \vec{s_0} = x(r) = 0.28, y(r) = 1.44, z(r) = 0.93, \theta(r) = 45°, \vec{s_1} = x(r) = 0.55, y(r) = 3.73, z(r) = 1.12, \theta(r) = 90°, \ldots\rangle$.

Given limitations in the sensing capability of the recognizer, observations may contain flaws, from missing to outright wrong/noisy observations, formalized in Definition 6.

Definition 6 (Missing and Noisy Observations) *Let Π be a planning task, π be a valid plan for Π and Ω be an observation sequence induced by an observation generation function $op(s_0, \cdot)$ with an action projection function $tp(\cdot, \cdot)$. An observation sequence Ω* ***misses*** *observations (is a partial observation sequence) with respect to the plan π if the $tp(\cdot, \cdot)$ function contains a mapping $a \mapsto \langle\rangle$ for some action a, i.e., it maps one or more actions into the empty sequence. An observation sequence Ω contains* ***noisy*** *observations with respect to the plan π if the $tp(s_{i-1}, a_i)$ function maps a_i into a non-empty sequence containing either one or more: (i) actions $a_j \neq a_i$ (for action observations); or (ii) states $s_i \notin \gamma(s_{i-1}, a_i)$.*

Indeed, most early approaches deal with missing observations implicitly by imposing a sequential constraint on the occurrence of elements of Ω in the plans considered plausible for a goal hypothesis G [Ramírez and Geffner, 2009], or ignoring most observations focusing on the overlap with necessary conditions [Pereira *et al.*, 2017]. Most approaches also deal with noisy observations either implicitly [Ramírez and Geffner, 2010; Pereira *et al.*, 2020], or not at all [Ramírez and Geffner, 2009]. Sohrabi *et al.* [2016], by contrast compiles Π into a single planning task with special predicates representing compliance with observations and achievement of goal hypotheses, as well as additional actions to achieve such predicates. This compilation ensures that goals whose plans have the least cost while complying with the most observations have higher probability. An *optimal* sequence of observations is extracted from an optimal plan and a *sub-optimal* sequence of observations from extracted of a sub-optimal plan.

5 Agent Assumptions

There are two key types of assumption regarding the observed agent that most approaches take into account during the recognition process: *awareness*, and *optimality*.

5.1 Awareness

The *awareness* assumption determines whether the agent is aware of it being monitored and what the attitude of the agent is towards the observer. Research on goal recognition characterize such awareness during the recognition process as follows [Armentano and Amandi, 2007; Sukthankar *et al.*, 2014]. *Intended Recognition* is the recognition process in which the observed agent is aware of the process of recognition. Therefore, in this type of recognition process the observed agent usually cooperates with the process by notifying the recognizer about its interactions in the environment. *Obstructed Recognition*, by contrast, is such that the observed agent is aware of the process of recognition and obstructs purposely the process, so the agent intentionally does not cooperate with the recognition process. *Keyhole Recognition* is a recognition process in which the observed agent is unaware of the process of recognition, namely, the interactions performed by the observed agent are partially observable inputs to the recognition process. This type of assumption on recognition process is the most common, since it allows the recognizer to ignore any process of interpretation of the actions by the observed agent as either adversarial or trying to be cooperative. Indeed, using such minimal level assumptions has led to many of the current state-of-the-art recognition algorithms to be vulnerable to either changes in target agent behavior or outright misdirection, especially when taking into consideration inferred agent preferences [Masters *et al.*, 2021].

5.2 Optimality

We now examine what is likely to be the most common assumption among goal recognition models, which is that the agent under observation either behaves *optimally* (i.e., perfectly rationally), or at least *approximately optimally*. The reason for this assumption is grounded on the fact that most early approaches to goal recognition, and indeed many more recent ones, carry out recognition by generating valid plans for all goal hypotheses $\mathcal{G}$. Since it is possible for planning problems to have infinitely many arbitrary plans that achieve G (e.g., by introducing irrelevant actions), methods that filter plausible goals by generating plans must assume optimality to bound the computation of such plans.

By contrast, methods that yield a probability distribution for $\mathbb{P}[G \mid \Omega]$ must have a criterion to assign probabilities to goal hypotheses, which often follows the assumption that agents under observation are approximately optimal and make $\mathbb{P}[G \mid \Omega]$ be proportional to how the plans that can project Ω are close to optimal plans towards G [Ramírez and Geffner, 2010; Sohrabi *et al.*, 2016], or, alternatively, optimal policies [Ramírez and Geffner, 2011]. Conversely, Pereira *et al.* [2017] reasons over necessary conditions to achieve goals G, instead of considering distributions over the possible plans an agent might choose to achieve G. In this case, no assumptions are made regarding agent optimality and probabilities are assigned by the degree to which the observations can account for such necessary conditions.

Indeed, Masters and Sardina [2019b] explore how the pervasiveness of the rationality assumption leads to errors in goal recognition, introducing a mechanism to automatically gauge rationality and modulate earlier formulas to compute $\mathbb{P}[\mathcal{G} \mid \Omega]$. Assuming rationality and honesty from the agents leads to a number of other weaknesses in recognition accuracy [Masters *et al.*, 2021], allowing agents to act irrationally to bias recognition away from the actual goal of the agent.

6 Goal Recognition Approaches

We now explore approaches in terms of the underlying mechanisms to infer the correct goal, the algorithms underpinning each approach, and how they process observations.

6.1 Filtering vs. Ranking and Probabilities

One of the key assumptions of goal recognition tasks is that there is a single goal G^* among the goal hypotheses $\mathcal{G}$. Nevertheless, most realistic goal recognition problems impose constraints on the quality of the observations and the level of observability afforded by the recognizer. This means that inferring goals can be fraught with uncertainty when observations suffer from flaws such as noisy and missing observations. Thus, while one may consider the problem of goal recognition as a *filtering process* aimed at providing all goal hypotheses G consistent with Ω [Ramírez and Geffner, 2009], such set of goals may be empty when Ω is too noisy, or encompass all hypothesis space $\mathcal{G}$ when Ω is too small. In such cases, it is useful to infer a set of potentially correct goals, either ranked by some preference relation, or, ideally, a *probability distribution* $\mathbb{P}\left[G^* = G\right], \forall G \in \mathcal{G}$.

Ramírez and Geffner [2010] formulate a probabilistic interpretation of goal recognition whereby the idea is computing the probability $\mathbb{P}[G \mid \Omega]$ of a given goal G, given observations Ω. Instead of computing that probability directly, most approaches use Bayes rule, which stipulates that $\mathbb{P}[G \mid \Omega] = \frac{\mathbb{P}[\Omega \mid G]\mathbb{P}[G]}{\mathbb{P}[\Omega]}$, and compute $\mathbb{P}[\Omega \mid G]$. The conditional probability of observations given a goal is directly related to the probability of the observed agent choosing a plan π for a particular goal G, which leads to the most common assumption for plan preference, which is that observed agents are approximately optimal and prefer close to optimal, and thus $\mathbb{P}[G \mid \Omega] \propto (\mathrm{cost}(\Omega) - \mathrm{cost}(\pi^*))$. Sohrabi *et al.* [2016] compute the probability $\mathbb{P}[\pi \mid G]$ of the top-k plans π consistent with G while complying with observations Ω as a proxy for $\mathbb{P}[\Omega \mid G]$. Masters and Sardina [2017; 2019a] reformulate the probabilistic interpretation of Ramírez and Geffner [2010] in the context of path-planning, and show that a single-observation recognition yields similar results in less than half of the recognition time. Martín *et al.* [2015] develop a probabilistic recognition approach that relies on planning graphs. In contrast, Pereira *et al.* [2020] *rank* hypotheses $\mathcal{G}$ following a heuristic based on the number of landmarks inferred from observations, returning a set of goals within a θ threshold of the highest *ranking* hypothesis. While probabilistic approaches notionally use a prior probability $\mathbb{P}[G]$ indicating agent preferences over goals, virtually no approach actually uses such prior, and instead assumes a uniform prior over the goals, leading to potential vulnerability to deception [Masters *et al.*, 2021].

6.2 Search vs. Heuristic Computation

While seminal goal recognition approaches rely on various graph analysis algorithms [Lesh and Etzioni, 1995; Hong, 2001], the approaches that have dominated the state-of-the-art in recent years can be broadly divided into two categories.

The first category relies on *compilations* of the goal recognition task into *Classical Planning* tasks, running the *search procedure* of a planner once or more and comparing the resulting plans to extract G^* [Ramírez and Geffner, 2010; Sohrabi *et al.*, 2016]. The second category uses the structure of planning tasks derived from $\Pi_{\mathcal{G}}^{\Omega}$ by adapting planning techniques to compute G^* either through the *computation of the heuristic* or by comparing heuristic values computed for each goal hypothesis. Typical examples of the first category include Ramírez and Geffner [2009] and Sohrabi *et al.* [2016]. Ramírez and Geffner [2009] generate a planning task Π from $\Pi_{\mathcal{G}}^{\Omega}$ by compiling away the observations such that when a plan tries to solve Π, it must generate plans compatible with the observations Ω. By contrast, typical examples of the second category include the approaches developed by Pattison and Long [2010], Martín *et al.* [2015], Pereira *et al.* [2017], Vered and Kaminka [2017b], and Santos [2021]. Pattison and Long [2010] propose a probabilistic recognition approach that uses heuristic estimation and domain analysis to determine which goals an agent is pursuing. Martín *et al.* [2015] develop a heuristic estimation approach that relies on planning graphs. Pereira *et al.* [2017] compute the set of landmarks [Hoffmann *et al.*, 2004] for each goal hypothesis $G \in \mathcal{G}$ and computing which landmarks have definitely been reached by plans that comply with the observations. Vered and Kaminka [2017b] propose a very effective recognition heuristic for continuous domain models. Recently, Santos *et al.* [2021] computes a *goal ranking* using linear programming over operator counting constraints on the problems induced by the goal hypothesis.

6.3 Online and Offline

Approaches to goal recognition differ in the way that the observations are perceived and revealed from the perspective of the recognizer. Most seminal approaches perform the recognition task *offline* [Ramírez and Geffner, 2009; Ramírez and Geffner, 2010; Martín *et al.*, 2015; Sohrabi *et al.*, 2016; Pereira *et al.*, 2017; Masters and Sardina, 2019a], in which the observations Ω are *all* revealed *at once* and *up front*, before starting the recognition process. The converse way to perform the recognition task is *online*, in which the observations Ω are revealed *incrementally* and recognition takes place multiple times, as each observation is revealed.

Online and *offline* goal recognition are not only different in the way that the observations are perceived and revealed, but also in the way the approaches are designed to perform the recognition task. *Online* goal recognition approaches are designed to be efficient and deal with incremental observations. Essentially, it is possible to perform *online* goal recognition by repeatedly calling an *offline* recognition approach for every new revealed observation. However, as proven by Vered and Kaminka in [2017a; 2017b], this is quite inefficient. In particular, Vered and Kaminka [2017a; 2017b] developed an *online* goal recognition that requires *at most* $|\mathcal{G}| * (|\Omega| + 1)$ calls to a planner, and *at best* $|\mathcal{G}|$ calls, whereas the seminal *offline* approach of Ramírez and Geffner [2010] requires $2 * |\mathcal{G}| * |\Omega|$ calls to a planner.

7 Related Problems

Goal Recognition can be seen as a subproblem of *Plan Recognition*, and it is closely related to a variety of problems in the literature. Baker *et al.* [2009] propose a probabilistic framework based on Bayesian inverse planning for modeling human action understanding. This probabilistic framework attempts to approximate the *principle of rationality*, which expects that the observed agents will plan approximately rationally to achieve their goals. Ramírez and Geffner [2009; 2010] borrowed the idea of Baker *et al.* [2009], and formally define *Plan Recognition as Planning*, claiming that *Plan Recognition* can be defined as *Planning* in reverse. Recent research extends the problem formulation of Ramírez and Geffner for recognizing plans in a variety of domain models, such as continuous domain models [Kaminka *et al.*, 2018] and epistemic planning problems [Shvo *et al.*, 2020]. Alternatively, Aineto *et al.* [2019a] introduce *Model Recognition as Planning*, a novel recognition task that aims to identify the *model* that explains a sequence of observations.

Keren *et al.* [2014; 2020] provide an alternate view of the goal recognition task, introducing a novel approach that modifies the domain model in order to facilitate the goal recognition process, so-called *Goal Recognition Design*. The aim of this task is optimizing the domain design so that recognition approaches can provide inferences with as few observations as possible. Over the past years, a variety of goal recognition design (and derived) approaches have been developed

Approach	Technique	Model	Obs. Flaws	Agent Assumption	Solution
Offline					
Ramírez and Geffner [2009]	Search Compilation	Authored Classical	Missing	Optimal Keyhole	Filtering
Ramírez and Geffner [2010]	Search Compilation	Authored Classical	Missing	≈Optimal Keyhole	Probability
Pattison and Long [2010]	Heuristic	Authored Classical	Missing	Optimal Keyhole	Probability
Ramírez and Geffner [2011]	Search	Authored Stochastic	Missing & Noisy	Optimal Keyhole	Probability
Martín *et al.* [2015]	Heuristic	Authored Classical	Missing	≈Optimal Keyhole	Ranking Probability
Sohrabi *et al.* [2016]	Search Compilation	Authored Classical	Missing & Noisy	≈Optimal Keyhole	Ranking Probability
Masters and Sardina [2017; 2019a]	Search	Authored Discrete /Continuous	Missing	≈Optimal Keyhole	Ranking Probability
Pereira *et al.* [2017; 2020]	Heuristic	Authored Classical	Missing & Noisy	None Keyhole	Ranking
Amado *et al.* [2018]	Heuristic /Search	Learned Classical	Missing	None Keyhole	Ranking Probability
Masters and Sardina [2019b]	Search	Authored Discrete /Continuous	Missing	≈Optimal Adversarial	Ranking Probability
Pereira *et al.* [2019a]	Heuristic	Incomplete Classical	Missing	None Keyhole	Ranking
Santos *et al.* [2021]	Heuristic	Authored Classical	Missing & Noisy	None Keyhole	Ranking
Online					
Vered and Kaminka [2017a]	Heuristic	Authored Continuous	Missing	Optimal Keyhole	Ranking Probability
Vered and Kaminka [2017b]	Search	Authored Continuous	Missing	Optimal Keyhole	Ranking Raking Probability
Vered *et al.* [2018]	Heuristic /Search	Authored Discrete /Continuous	Missing	≈Optimal Keyhole	Ranking
Pereira *et al.* [2019b]	Search	Learned Control	Missing	Optimal Keyhole	Ranking Probability

Table 1: Goal Recognition landscape.

over distinct types of domain models and settings, such as [Keren *et al.*, 2016; Wayllace *et al.*, 2016; Keren *et al.*, 2018; Wayllace *et al.*, 2020b; Shvo and McIlraith, 2020].

Pozanco *et al.* [2018] develop a *Counterplanning* approach that relies on techniques from goal recognition and planning to *prevent* observed agents from achieving their goals. Most recently, Bernardini *et al.* [2020] develop a set of strategies for *Goal Obfuscation*, in which an actor agent aims to maintain its goal private, without revealing it to an observer.

8 Conclusions and Perspectives

We have surveyed advances in goal recognition as planning from the last years, classifying the various aspects that distinguish the approaches, including the types of domains being handled, how observations are handled, what assumptions are being made about the agent under observation, culminating in the algorithms that underpin such approaches. We summarize this classification in Table 1, where *Technique* refers to the underlying approach, *Model* refers to the type of model that describes the tasks (whether authored or learned, and the domain description), *Observation Flaws* refers to the potential flaws in the observations, *Agent Assumption* refers to the assumptions made about the agent, and finally, *Solution* refers to the type of solution provided by the approach.

Despite the substantial progress in the approaches we survey, many research challenges remain. Existing approaches assume *pure* models in some sense, so models are either continuous or symbolic. Thus, future research should focus on what Ramírez calls *interesting* domains, i.e., models described in PDDL+ [Fox and Long, 2006], or temporal planning [Fox and Long, 2003b]. Similarly, few approaches deal with multi-agent domains, which we do not survey [Shvo *et al.*, 2017; Argenta *et al.*, 2017; Zhuo, 2019].

The key challenge to current incremental research is comparing existing approaches in a principled way. Current resources could be improved in two ways. First, most approaches use variations of the original problem description and dataset developed by Ramírez and Geffner [2009], all of them derived from problems in the International Planning Competition (IPC). While IPC domains constitute challenging settings for planning implementations, these domains are not necessarily realistic or representative of goal recognition problems. Second, and unlike planning algorithms, many of the approaches surveyed lack an openly accessible reference implementation amenable to objective comparison. Thus, developing a unified benchmark in a common format using representative goal recognition problems, and making all efforts openly available constitute key targets for future efforts.

Acknowledgements

Felipe acknowledges support from CNPq with projects 407058/2018-4 (Universal) and 302773/2019-3 (PQ). Ramon acknowledges support from ERC Advanced Grant White-Mech (No. 834228) and EU ICT-48 2020 TAILOR (No. 952215).

References

[Aineto *et al.*, 2019a] Diego Aineto, Sergio Jiménez, Eva Onaindia, and Miquel Ramírez. Model Recognition as Planning. In *ICAPS*, 2019.

[Aineto *et al.*, 2019b] Diego Aineto, Sergio Jiménez Celorrio, and Eva Onaindia. Learning action models with minimal observability. *Artificial Intelligence*, 275:104–137, October 2019.

[Amado *et al.*, 2018] Leonardo Amado, Ramon Fraga Pereira, Joao Paulo Aires, Mauricio Magnaguagno, Roger Granada, and Felipe Meneguzzi. Goal recognition in latent space. In *IJCNN*, 2018.

[Argenta *et al.*, 2017] C Argenta, J Doyle, and 2017. Probabilistic Multi-Agent Plan Recognition as Planning (P-Maprap): Recognizing Teams, Goals, and Plans from Action Sequences. *scitepress.org*, 2017.

[Armentano and Amandi, 2007] Gabriel Armentano and Analía Amandi. Plan recognition for interface agents. *Artificial Intelligence Review*, 28(2):131–162, 2007.

[Asai and Fukunaga, 2018] Masataro Asai and Alex Fukunaga. Classical Planning in Deep Latent Space - Bridging the Subsymbolic-Symbolic Boundary. *AAAI*, 2018.

[Baker *et al.*, 2009] Chris L Baker, Rebecca Saxe, and Joshua B Tenenbaum. Action understanding as inverse planning. *Cognition*, 113(3):329–349, December 2009.

[Bernardini *et al.*, 2020] Sara Bernardini, Fabio Fagnani, and Santiago Franco. An Optimization Approach to Robust Goal Obfuscation. In *KR*, 2020.

[Bertsekas, 2017] Dmitri P. Bertsekas. *Dynamic Programming and Optimal Control*. Athena Scientific, 4th edition, 2017. 520p.

[Borrelli *et al.*, 2017] Francesco Borrelli, Alberto Bemporad, and Manfred Morari. *Predictive control for linear and hybrid systems*. Cambridge University Press, 1st edition, 2017. 440p.

[Fox and Long, 2003a] Maria Fox and D Long. PDDL2. 1: An extension to PDDL for expressing temporal planning domains. *Journal of Artificial Intelligence Research*, 2003.

[Fox and Long, 2003b] Maria Fox and Derek Long. PDDL2.1: an extension to PDDL for expressing temporal planning domains. *Journal of Artificial Intelligence Research*, 20:61–124, 2003.

[Fox and Long, 2006] Maria Fox and Derek Long. Modelling mixed discrete-continuous domains for planning. *Journal of Artificial Intelligence Research*, 27:235–297, 2006.

[Geffner and Bonet, 2013] Hector Geffner and Blai Bonet. *A Concise Introduction to Models and Methods for Automated Planning*. Morgan & Claypool, 1st edition, 2013.

[Geib and Steedman, 2007] Christopher W Geib and Mark Steedman. On natural language processing and plan recognition. In *Proceedings of the 20th IJCAI*, pages 1612–1617, 2007.

[Geib, 2002] Christopher Geib. Problems with intent recognition for elder care. In *AAAI*, 2002.

[Granada *et al.*, 2017] Roger Granada, Ramon Fraga Pereira, Juarez Monteiro, Rodrigo Barros, Duncan Ruiz, and Felipe Meneguzzi. Hybrid Activity and Plan Recognition for Video Streams. In *AAAI Workshop on Plan, Activity, and Intent Recognition*, 2017.

[Haslum *et al.*, 2019] Patrik Haslum, Nir Lipovetzky, Daniele Magazzeni, and Christian Muise. *An Introduction to the Planning Domain Definition Language*, volume 13. Morgan & Claypool Publishers, 3 edition, April 2019.

[Hegde and Kenchannavar, 2019] Rakshith M. D. Hegde and Harish H. Kenchannavar. A survey on predicting resident intentions using contextual modalities in smart home. *International Journal of Advanced Pervasive and Ubiquitous Computing*, 11(4):44–59, 2019.

[Helmert, 2009] Malte Helmert. Concise finite-domain representations for PDDL planning tasks. *Artificial Intelligence*, 173(5-6):503–535, April 2009.

[Hoffmann *et al.*, 2004] J Hoffmann, J Porteous, and L Sebastia. Ordered Landmarks in Planning. *Journal of Artificial Intelligence Research*, 22(1):215–278, April 2004.

[Hong, 2001] Jun Hong. Goal Recognition through Goal Graph Analysis. *Journal of Artificial Intelligence Research*, 15:1–30, 2001.

[Inam *et al.*, 2018] Rafia Inam, Klaus Raizer, Alberto Y Hata, Ricardo S Souza, Elena Fersman, Enyu Cao, and

Shaolei Wang. Risk Assessment for Human-Robot Collaboration in an Automated Warehouse Scenario. *ETFA*, pages 743–751, 2018.

[Kaminka *et al.*, 2018] Gal A Kaminka, Mor Vered, and Noa Agmon. Plan Recognition in Continuous Domains. In *AAAI*, 2018.

[Keren *et al.*, 2014] Sarah Keren, Avigdor Gal, and Erez Karpas. Goal Recognition Design. In *ICAPS*, 2014.

[Keren *et al.*, 2016] Sarah Keren, Avigdor Gal, and Erez Karpas. Goal Recognition Design with Non-Observable Actions. In *AAAI*, 2016.

[Keren *et al.*, 2018] Sarah Keren, Avigdor Gal, and Erez Karpas. Strong Stubborn Sets for Efficient Goal Recognition Design. *ICAPS*, 2018.

[Keren *et al.*, 2020] Sarah Keren, Avigdor Gal, and Erez Karpas. Goal Recognition Design - Survey. In *IJCAI*, 2020.

[Lesh and Etzioni, 1995] Neal Lesh and Oren Etzioni. A Sound and Fast Goal Recognizer. In *IJCAI*, 1995.

[Martín *et al.*, 2015] Yolanda E Martín, María D R Moreno, and David E Smith. A Fast Goal Recognition Technique Based on Interaction Estimates. *IJCAI*, 2015.

[Masters and Sardina, 2017] Peta Masters and Sebastian Sardina. Cost-Based Goal Recognition for Path-Planning. In *AAMAS*, 2017.

[Masters and Sardina, 2019a] Peta Masters and Sebastian Sardina. Cost-based goal recognition in navigational domains. *Journal of Artificial Intelligence Research*, 64:197–242, 2019.

[Masters and Sardina, 2019b] Peta Masters and Sebastian Sardina. Goal Recognition for Rational and Irrational Agents. In *AAMAS*, 2019.

[Masters *et al.*, 2021] Peta Masters, Michael Kirley, and Wally Smith. Extended goal recognition: a planning-based model for strategic deception. In *AAMAS*, 2021.

[Mirsky *et al.*, 2021] Reuth Mirsky, Sarah Keren, and Christopher Geib. Introduction to symbolic plan and goal recognition. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 16(1):1–190, 2021.

[Nguyen *et al.*, 2017] Tuan Nguyen, Sarath Sreedharan, and Subbarao Kambhampati. Robust planning with incomplete domain models. *Artificial Intelligence*, 245:134–161, April 2017.

[Oh *et al.*, 2011] Jean Oh, Felipe Meneguzzi, Katia P Sycara, and Timothy J Norman. An Agent Architecture for Prognostic Assistance. In *IJCAI*, 2011.

[Pattison and Long, 2010] David Pattison and Derek Long. Domain Independent Goal Recognition. *STAIRS*, 2010.

[Pereira *et al.*, 2017] Ramon Fraga Pereira, Nir Oren, and Felipe Meneguzzi. Landmark-Based Heuristics for Goal Recognition. In *AAAI*, 2017.

[Pereira *et al.*, 2019a] Ramon Fraga Pereira, Andr e Grahl Pereira, and Felipe Meneguzzi. Landmark-Enhanced Heuristics for Goal Recognition in Incomplete Domain Models. In *ICAPS*, 2019.

[Pereira *et al.*, 2019b] Ramon Fraga Pereira, Mor Vered, Felipe Meneguzzi, and Miquel Ramírez. Online Probabilistic Goal Recognition over Nominal Models. In *IJCAI*, 2019.

[Pereira *et al.*, 2020] Ramon Fraga Pereira, Nir Oren, and Felipe Meneguzzi. Landmark-based approaches for goal recognition as planning. *Artificial Intelligence*, 279:103217, 2020.

[Pozanco *et al.*, 2018] Alberto Pozanco, Yolanda E-Martín, Susana Fernández, and Daniel Borrajo. Counterplanning using goal recognition and landmarks. In Jérôme Lang, editor, *IJCAI*, 2018.

[Ramírez and Geffner, 2009] Miquel Ramírez and Hector Geffner. Plan recognition as planning. In *IJCAI*, 2009.

[Ramírez and Geffner, 2010] Miquel Ramírez and Hector Geffner. Probabilistic Plan Recognition Using Off-the-Shelf Classical Planners. In *AAAI*, 2010.

[Ramírez and Geffner, 2011] Miquel Ramírez and Hector Geffner. Goal Recognition over POMDPs - Inferring the Intention of a POMDP Agent. In *IJCAI*, 2011.

[Santos *et al.*, 2021] Luisa R de A Santos, Felipe Meneguzzi, Ramon Fraga Pereira, and Andre Pereira. An LP-Based Approach for Goal Recognition as Planning. In *AAAI*, 2021.

[Shvo and McIlraith, 2020] Maayan Shvo and Sheila A McIlraith. Active Goal Recognition. In *AAAI*, 2020.

[Shvo *et al.*, 2017] Maayan Shvo, Shirin Sohrabi, and Sheila A McIlraith. An AI Planning-Based Approach to the Multi-Agent Plan Recognition Problem. In *AAAI Workshop on Plan, Activity, and Intent Recognition*, 2017.

[Shvo *et al.*, 2020] Maayan Shvo, Toryn Q Klassen, Shirin Sohrabi, and Sheila A McIlraith. Epistemic Plan Recognition. In *AAMAS*, 2020.

[Singh *et al.*, 2020] Ronal Rajneshwar Singh, Tim Miller, Joshua Newn, Eduardo Velloso, Frank Vetere, and Liz Sonenberg. Combining gaze and AI planning for online human intention recognition. *Artificial Intelligence*, 284:103275, 2020.

[Sohrabi *et al.*, 2016] S Sohrabi, A V Riabov, and O Udrea. Plan Recognition as Planning Revisited. In *IJCAI*, 2016.

[Suárez-Hernández *et al.*, 2020] Alejandro Suárez-Hernández, Javier Segovia-Aguas, Carme Torras, and Guillem Alenyà. STRIPS Action Discovery. *arXiv.org*, January 2020.

[Sukthankar *et al.*, 2014] Gita Sukthankar, Robert P Goldman, Christopher Geib, David V Pynadath, and Hung Hai Bui. *Plan, Activity, and Intent Recognition: Theory and Practice*. Elsevier, 2014.

[Sutton and Barto, 2018] Richard S Sutton and Andrew G Barto. *Reinforcement Learning*. An Introduction. The MIT Press, 2018.

[Vered and Kaminka, 2017a] Mor Vered and Gal A Kaminka. Heuristic Online Goal Recognition in Continuous Domains. In *IJCAI*, 2017.

[Vered and Kaminka, 2017b] Mor Vered and Gal A Kaminka. Online Recognition of Navigation Goals Through Goal Mirroring. *AAMAS*, 2017.

[Vered *et al.*, 2016] M Vered, G A Kaminka, and Sivan Biham. Online goal recognition through mirroring: Humans and agents. In *ACS*, 2016.

[Vered *et al.*, 2018] Mor Vered, Ramon Fraga Pereira, Mauricio Cecilio Magnaguagno, Gal A Kaminka, and Felipe Meneguzzi. Towards Online Goal Recognition Combining Goal Mirroring and Landmarks. In *AAMAS*, 2018.

[Wang *et al.*, 2020] Chien-Yao Wang, Alexey Bochkovskiy, and Hong-Yuan Mark Liao. Scaled-YOLOv4: Scaling Cross Stage Partial Network. *arXiv*, November 2020.

[Wayllace *et al.*, 2016] Christabel Wayllace, Ping Hou, William Yeoh, and Tran Cao Son. Goal recognition design with stochastic agent action outcomes. In Subbarao Kambhampati, editor, *IJCAI*, 2016.

[Wayllace *et al.*, 2020a] Christabel Wayllace, Sunwoo Ha, Yuchen Han, Jiaming Hu, Shayan Monadjemi, William Yeoh, and Alvitta Ottley. DRAGON-V: detection and recognition of airplane goals with navigational visualization. In *AAAI*, 2020.

[Wayllace *et al.*, 2020b] Christabel Wayllace, Sarah Keren, Avigdor Gal, Erez Karpas, William Yeoh, and Shlomo Zilberstein. Accounting for observer's partial observability in stochastic goal recognition design. In *ECAI*, 2020.

[Weber and Bryce, 2011] C Weber and D Bryce. Planning and Acting in Incomplete Domains. *ICAPS*, 2011.

[Zhuo, 2019] Hankz Hankui Zhuo. Recognizing Multi-Agent Plans When Action Models and Team Plans Are Both Incomplete. *ACM TIST*, 10(3):1–24, 2019.