

A Survey on Model-Free Goal Recognition

Leonardo Amado¹, Sveta Paster Shainkopf², Ramon Fraga Pereira³,
Reuth Mirsky² and Felipe Meneguzzi¹

¹University of Aberdeen, Scotland, UK

²Bar-Ilan University, Israel

³University of Manchester, England, UK

{leonardo.amado, felipe.meneguzzi}@abdn.ac.uk, mirskyr@cs.biu.ac.il,
pasters5@biu.ac.il, ramon.fragapereira@manchester.ac.uk

Abstract

Goal Recognition is the task of inferring an agent’s intentions from a set of observations. Existing recognition approaches have made considerable advances in domains such as human-robot interaction, intelligent tutoring systems, and surveillance. However, most approaches rely on explicit domain knowledge, often defined by a domain expert. Much recent research focus on mitigating the need for a domain expert while maintaining the ability to perform quality recognition, leading researchers to explore *Model-Free Goal Recognition* approaches. We comprehensively survey *Model-Free Goal Recognition*, and provide a perspective on the state-of-the-art approaches and their applications, showing recent advances. We categorize different approaches, introducing a taxonomy with a focus on their characteristics, strengths, weaknesses, and suitability for different scenarios. We compare the advances each approach made to the state-of-the-art and provide a direction for future research in *Model-Free Goal Recognition*.

1 Introduction

Goal Recognition is a problem in which one agent, a *recognizer*, aims to infer the most likely pursued goal of a set of candidate goals of an observed agent, a *subject*, based on a sequence of observations [Meneguzzi and Pereira, 2021; Mirsky et al., 2021]. Research in *Goal Recognition* has many similarities with *Plan Recognition*, as it entails computing the plan responsible for an agent’s observed behavior [Schmidt et al., 1978]. For the sake of this survey, we only focus on *Goal Recognition*, but we highlight that many of the surveyed approaches can also be adapted for *Plan Recognition*. Ramírez and Geffner [2009] defined the problem of *Plan Recognition as Planning* and developed approaches for the plan (and consequently goal) recognition in symbolic domains. These approaches serve as the cornerstone for subsequent approaches and as a baseline for evaluating many plan and goal recognition approaches. Their definition of a *Goal Recognition* problem assumes available domain theory. For example, in *Automated Planning* approaches, the domain knowledge is often

written in PDDL (Planning Domain Definition Language). This domain knowledge explicitly defines the environmental dynamics in which the observed agent acts. The availability of a domain theory allows recognizer agents to reason about the goals of the observed agent and how such an agent may want to achieve the possible candidate goals.

In practical real-world scenarios, the challenge of recognizing goals and plans is notably heightened when domain knowledge is either unknown or not readily accessible in a detailed form. Despite this complexity, there are two effective approaches that enable the recognition of goals even without explicit domain knowledge. The first is by *learning* or *approximating* domain knowledge, which can leverage recent research in *Model Acquisition in Automated Planning* [Kumar et al., 2023; Aineto et al., 2018; Aineto et al., 2019b; Aineto et al., 2019a; Aineto et al., 2022]. Second, one could design *Goal Recognition* approaches capable of recognizing goals without relying on explicit domain knowledge. These approaches are essential for adapting to the complexities inherent in real-world scenarios.

Recent advances in *Machine Learning* and *Neuro-Symbolic* approaches led to the development of many *model-free* recognition approaches. Learning-based approaches leverage data to facilitate the interface with the real physical world, while *Neuro-Symbolic* approaches aim to combine the strengths of modern *Machine Learning* and “classical” *Symbolic Artificial Intelligence* (AI) [Asai et al., 2021; Amado et al., 2023]. For instance, *Reinforcement Learning* enables agents to learn without explicitly knowing the domain model, empowering recent recognition approaches [Amado et al., 2022; Fang et al., 2023]. These advances pave the way for novel approaches to solving *Goal Recognition* problems by either learning the domain knowledge or obviating its need. While the vast majority of recognition approaches rely on engineered models, with most of their success in *Symbolic AI* approaches [Ramírez and Geffner, 2009; E.-Martín et al., 2015; Sohrabi et al., 2016; Pereira et al., 2020; Santos et al., 2021], these offer substantial opportunities to advance the field. In this survey, we take stock of recent advances in both *Machine Learning* and *Goal Recognition*, providing a formal underpinning for *Model-Free Goal Recognition*. We provide a rigorous taxonomy for the existing approaches, and discuss insights into their limitations and potential extensions, paving the way for further research in the area.

2 Preliminaries

2.1 Goal Recognition

Goal Recognition (GR) is the task of recognizing the intended goal that an observed agent aims to achieve based on the observations of the agent’s interactions in an environment [Schmidt *et al.*, 1978; Mirsky *et al.*, 2021]. *Observations* are key to the recognition process, and a recognition process may involve observations of one or multiple agent(s) that aim to achieve one or more goal(s) in a particular environment. There are at least two participants during the recognition process: the **observed** or **subject** agent, which acts in an environment to achieve its goals; and the **recognizer**, which performs the recognition process by observing the observed agent’s interactions in an environment. The subject agent can be of any nature, i.e., it can be a computer process or a human, and we make no distinction between these entities throughout this survey. Once a recognizer infers the goal of the subject agent, it can *anticipate* its behavior. Definition 1 formally lays down the *Goal Recognition* problem.

Definition 1 (Goal Recognition Problem). A *goal recognition problem* $\mathcal{P}_{GR} = \langle \mathcal{M}, s_{\mathcal{I}}, \mathcal{G}, \Omega \rangle$ is a tuple that comprises a model \mathcal{M} , representing the properties and actions of an environment, an initial state $s_{\mathcal{I}}$, a set of possible goals \mathcal{G} , including the intended goal $G^* \in \mathcal{G}$ unknown to the recognizer, and a sequence of observations Ω that projects a sequence of interactions in an environment to achieve G^* .

\mathcal{M} represents the relationship between observed interactions in an environment and the possible goals an agent could achieve. The recognition process can use \mathcal{M} to reason over the possible behaviors of the agent for achieving its goals. The initial state $s_{\mathcal{I}}$ represents a “snapshot” of the environment at the beginning of the recognition process, before the subject agent executes any actions. The set of possible goals \mathcal{G} represents the potential intentions that an agent has within a given environment. Most *Goal Recognition* approaches in the literature assume there is at least one correct intended goal $G^* \in \mathcal{G}$ among the possible goals [Mirsky *et al.*, 2021]. The observations Ω can be seen as an indirect projection of an observed agent’s behavior, representing a sequence of interactions (e.g., a plan) for achieving its goals in an environment [Mirsky *et al.*, 2021; Meneguzzi and Pereira, 2021]. In the next sections, we delve into the pivotal role played by the presence, absence, or approximation of a model \mathcal{M} in the task of *Goal Recognition*.

Goal Recognition makes a few key assumptions [Masters and Vered, 2021] concerning the recognizer and the subject agent. Such key assumptions are essential and influence the capabilities and characteristics of both the recognizer and the subject agent during the recognition process. The specific types of assumptions lead to the most common types of *Goal Recognition* [Masters and Vered, 2021], as follows. In *Intended Recognition*, the subject agent is aware of the recognition process and actively cooperates by notifying the recognizer about its interactions via the observations. *Obstructed Recognition* involves the subject agent intentionally obstructing the recognition process. *Keyhole Recognition* assumes that the observed agent is unaware of the recognition process, and the interactions it performs serve as partially ob-

servable inputs. This type of recognition is common as it allows the recognizer to disregard any interpretation of the observed agent’s actions as adversarial or cooperative.

2.2 Model-Based Goal Recognition

In *Model-Based Goal Recognition*, the recognition process involves employing a **predefined** and **explicit model** that encapsulates the properties of an environment and the executable actions within that environment [Mirsky *et al.*, 2021; Meneguzzi and Pereira, 2021; Masters and Vered, 2021]. This *model* (commonly known as *domain model* or *domain theory*) describes the properties and dynamics of an environment in which agents act on, as well as the degrees of freedom of the observed agents in such an environment. The properties and dynamics formalized in a *model* are key to reason about the intentions and the possible behaviors of the observed agents in an environment. *Domain experts* typically engineer the *models* we refer to in this paper.

In this survey, we make no specific assumption about the nature of the model and the environment, be it about its observability or stochasticity. In the context of *Model-Based Goal Recognition*, a *model* \mathcal{M} follows a symbolic formalism [Mirsky *et al.*, 2021], e.g., *discrete* (predicate-based), *continuous*, or *mixed discrete-continuous*. *Model-Based Goal Recognition* approaches in the literature employ different types of domain models, such as *plan-libraries* (i.e., a collection of plans predefined to achieve a set of goals), *graph-based representations* (e.g., commonly used in *Path-Planning*), *Automated Planning* (*symbolic*) domain models, e.g., STRIPS (Stanford Research Institute Problem Solver) [Fikes and Nilsson, 1971], PDDL [McDermott *et al.*, 1998], and RDDDL (Relational Dynamic Influence Diagram Language) [Sanner and Boutilier, 2010], etc. For a more comprehensive understanding of *Model-Based Goal Recognition*, we refer the reader to the surveys of Meneguzzi and Pereira, and Masters and Vered. These surveys provide a comprehensive overview of seminal and recent developments in *Model-Based Goal Recognition*, as well as an insightful analysis of the most common assumptions for this problem.

3 Model-Free Goal Recognition

We define *Model-Free Goal Recognition* as the task of recognizing the intended goal of an observed agent (or human) **without relying on** a predefined and explicit *model* of the agent’s behavior or the environment. Definition 2 lays down *Model-Free Goal Recognition* problems.

Definition 2 (Model-Free Goal Recognition Problem). A *model-free goal recognition problem* $\mathcal{P}_{GR}^{\mathcal{M}} = \langle \mathcal{M}, s_{\mathcal{I}}, \mathcal{G}, \Omega \rangle$ is one in which the recognizer has no access to the underlying model \mathcal{M} that describes the properties and dynamics of the environment.

Note here that, while Definition 2 assumes the actual environment is consistent with some model \mathcal{M} , we do not yet specify the exact nature of this model, nor the relation between this model and the subject agent. The only assumption a $\mathcal{P}_{GR}^{\mathcal{M}}$ problem makes is that, whatever the actual environment model \mathcal{M} determines action outcomes and rewards, the

recognizer has no knowledge of its dynamics or reward function at the start of the recognition process. This assumption allows for multiple possibilities about the nature of the actions of the subject agent. It may be aware of the exact model \mathcal{M} , as in Definition 1, or it may have an approximation of the model (e.g., by acting through a learned policy). Indeed, from the recognizer perspective, virtually all approaches we survey disregard the specific nature of the subject agent’s understanding of the model, and the recognizer may perceive the subject agent’s actions as either noisy, sub-optimal, or adversarial during the recognition process.

4 Taxonomy of Model-Free Goal Recognition

In the previous sections, we defined the formalism and terminology to pin down *Model-Free Goal Recognition*. With this foundation, we now introduce a taxonomy of various types of *Model-Free Goal Recognition*, which comprises a series of attributes, and for each attribute, possibilities for its instantiation in a concrete approach.

Affinity to a Model. Existing recognizers often follow one of two paradigms: The first paradigm is *Model-Agnostic Goal Recognition*, meaning that the recognition process **does not rely on** any type of explicitly predefined model to recognize the intended goal G^* of an observed agent. The second paradigm is *Model-Approximate Goal Recognition*, the recognition process relies on an **approximate** model, which can be inferred and learned from data, or partially defined with incomplete information before the recognition process. This classification is elaborated in Section 4.1 and it stands as the central aspect of our taxonomy, defined as **Agnostic** and **Model-Approximate** in Table 1.

Environment. An *environment* is the external system in which agents operate [Mirsky *et al.*, 2021, Section 2.1], such a system can be physical, virtual, or a combination of both. Understanding the environment and its properties is a crucial aspect for a recognition process, as it defines the context in which the recognizer perceives information about the behavior of the subject agents. In our taxonomy, this attribute is referred to as **Environment**. We categorize environment types, characteristics, and assumptions into five different attributes: *discrete*, indicating that the environment is discrete; *continuous*, indicating that the environment is continuous; *image*, indicating that some visual information represents the environment, and it may involve *Computer Vision* techniques to interpret visual information; and *stochastic*, indicating that the environment interactions are partially random, introducing uncertainty¹. Environment attributes encompass more than the characteristics we enumerate in this survey [Meneguzzi and Pereira, 2021; Masters and Vered, 2021], such as partial observability, multiple-agents, episodic interactions, among others, our taxonomy ignores them, since existing research on *Model-Free Goal Recognition* does not use them.

Employed Technique. The *Model-Free Goal Recognition* papers we survey employ different techniques to perform the

recognition process. In our taxonomy, we refer to this attribute as **Technique**, representing the main underlying technique used for the recognition process. We categorize techniques into four types: *Supervised Learning*, indicating the use of *supervised* learning models, such as a *Linear Regressor*, a *Support Vector Machine*, any one of a number of *Neural Network* architecture, etc; *Reinforcement Learning*, indicating the use of any RL technique; *Symbolic*, indicating the use of purely *symbolic* techniques, such as *Automated Planning* (heuristic search, optimization, etc); and *Neuro-Symbolic*, indicating the use of techniques that leverage a neural network to learn a symbolic model to perform the recognition process. As the deployed technique is inherently connected to the model’s affinity, we do not elaborate further on this attribute in this Section. Rather, Section 5 elaborates on specific *Model-Free Goal Recognition* approaches and how they employ the techniques we discussed above.

Input Data. Both *model-agnostic* and *model-approximate* approaches require data. The former uses it to directly **deduce** the likely goal from observations, while the latter uses it to first learn a model from which, given the observations, the goal can later be **abducted**. We categorize the *input data* for *Model-Free Goal Recognition* into three different categories: *traces*, which typically refers to sequences of actions and/or state properties that represent the executed behavior of an agent in an environment; *images*, representing sequences of images as transitions that are analogous to plans, i.e., sequence of images capturing transitional states that depict progression, transitioning from an initial state image to a goal state image; and *sampling*, representing input data generated from a *Reinforcement Learning* simulator [Amado *et al.*, 2022; Fang *et al.*, 2023]. Section 4.2 discusses the common existing data types used for *Model-Free Goal Recognition*.

Recognition Input. In our taxonomy, the information that a recognizer takes as input (and how it is revealed) is referred to as **Recognition Input**, and categorized as follows: observations as *actions*; observations as *states*; observations with *missing* information; observations with *noisy* information; and observations revealed incrementally *online*. While the initial state and the set of goals provide important input information to a recognizer, most approaches assume that they are usually well-defined and accurate. By contrast, we discuss in Section 4.3 how *observations* represent the key aspect of the recognition input, including its accuracy and how they are revealed to the recognizer.

Recognition Output. Existing *Model-Free Goal Recognition* approaches vary in how they output the most likely goal(s). The two most common ways of presenting the recognized goals as output are: *goal ranking*; or *probability distribution*. Note that the recognition output does not affect how accurate the recognition approaches are. However, it does influence how people or other systems interpret the most likely goal(s) of a subject agent.

In the remainder of this section, we dive deeper into some of the taxonomy’s attributes and values, and end with a summary of the surveyed papers over this taxonomy on Table 1.

¹Since most papers we survey rely on deterministic environments, if *stochastic* is not checked, the environment is deterministic.

4.1 Model Affinity

In what follows, we refine our assumptions about the level of knowledge on the part of the recognizer. Thus, we refine Definition 2 into two types of *Model-Free Goal Recognition* tasks: *Agnostic Model-Free Goal Recognition* and *Model-Approximate Goal Recognition*.

In *Agnostic Model-Free Goal Recognition*, the recognition process **does not rely on** any type of explicitly predefined model to recognize the intended goal G^* of an observed agent. Instead, it relies on the three pieces of information formalized in Definition 2: an initial state $s_{\mathcal{I}}$, a set of possible goals \mathcal{G} , and a sequence of observations Ω . This type of recognition task is particularly useful in settings in which constructing an accurate (or even approximate model) is challenging or impractical (e.g., limited data availability, low-quality data, etc). These approaches often rely on underlying data, such as plan traces, to train *Machine Learning* models.

Definition 3 (Agnostic Model-Free Goal Recognition). Let $\mathcal{P}_{\mathcal{GR}}^{\mathcal{M}}$ be a model-free goal recognition problem. A goal recognition process is **agnostic** if it is capable of recognizing the intended goal $G^* \in \mathcal{G}$ **without learning or approximating** the underlying model \mathcal{M} using the information provided by the initial state $s_{\mathcal{I}}$, the goals \mathcal{G} , and the observations Ω .

Conversely, in *Model-Approximate Goal Recognition*, the recognition process relies on an **approximate** model, which can be inferred and learned from data, or partially defined with incomplete information before the recognition process. Such an approximate model might rely not only on incomplete, but also incorrect information, leading to a potentially partial and *incorrect* model. In Definition 4, we formalize *Model-Approximate Goal Recognition*².

Definition 4 (Model-Approximate Goal Recognition). Let $\mathcal{P}_{\mathcal{GR}}^{\mathcal{M}}$ be a model-free goal recognition problem. A goal recognition process is **model-approximate** if it first **approximates** the underlying model \mathcal{M} as $\tilde{\mathcal{M}}$.

Once a recognizer has an approximate model $\tilde{\mathcal{M}}$, it can use it to solve a goal recognition problem analogous to Definition 1 such that $\mathcal{P}_{\mathcal{GR}}^{\tilde{\mathcal{M}}} = \langle \tilde{\mathcal{M}}, s_{\mathcal{I}}, \mathcal{G}, \Omega \rangle$ is a tuple that comprises an approximate model $\tilde{\mathcal{M}}$, an initial state $s_{\mathcal{I}}$, a set of possible goals \mathcal{G} , and a sequence of observations Ω . Thus, most approaches that can solve a model-based goal recognition $\mathcal{P}_{\mathcal{GR}}^{\mathcal{M}}$ problem could, in theory, solve model-approximate goal recognition $\mathcal{P}_{\mathcal{GR}}^{\tilde{\mathcal{M}}}$. It is important to highlight that there exist approaches in the literature that are explicitly developed for recognizing goals when dealing with approximate models [Pereira et al., 2019b], incomplete and possibly incorrect domain models [Pereira and Meneguzzi, 2018; Pereira et al., 2019a; Zhuo, 2019; Kerkez and Cox, 2002], imperfect domain models [Pereira, 2020], etc.

4.2 Input Data and Assumptions

Regardless of whether the technique used to solve a *Model-Free Goal Recognition* problem is *model-agnostic* or *model-*

approximate, the most common assumption is that there is some underlying data. In this section, we detail the different types of data often used in the literature to solve *Model-Free Goal Recognition* problems.

Existing *model-agnostic* approaches often require training data or have access to a simulator of the environment to sample data. *Model-agnostic* recognition uses two types of data: *structured data* and *unstructured data*. In *Model-Free Goal Recognition*, we consider structured data to be plan traces, either as sequences of actions or states, as these can provide a direct path from an initial state to a goal [Chiari et al., 2023]. Existing approaches that rely on structured data often require substantial amounts of data, and they assume that the traces used for training are optimal (or at least rational). Unstructured data would be data that does not explicitly contain domain information, such as images [Amado et al., 2018a], video streams, natural language descriptions, or even multi-modal (e.g., images and descriptions) information. Based on the literature, only model-agnostic approaches rely on images as unstructured data, and they often need a substantial amount of data [Min et al., 2016]. Although most approaches assume available data, *Reinforcement Learning* ones assume a simulator to sample “unlimited” data by interacting with a simulator [Amado et al., 2022; Fang et al., 2023].

Approaches that compute an approximate domain model \mathcal{M} often need a substantial amount of data, much like *model-agnostic* approaches. However, approximating a model requires not only data quantity, but also quality, in the form of varied state-transitions. If such data lacks samples of certain transitions, these approaches can offer very few guarantees. Once again, we consider *model-approximate* approaches to **structured data** and **unstructured data**. The ideal structured data to learn a domain is action traces and state traces that enable reasoning about the actions of an environment [McDermott et al., 1998; Aineto et al., 2018]. As unstructured data, we consider approaches that approximate models from images [Asai and Fukunaga, 2018; Amado et al., 2018b; Asai et al., 2021] and very recently text [Guan et al., 2023] (although not used explicitly for goal recognition). These approaches require even more data to approximate domain models, and most of the time they provide no guarantee about the completeness of the model. For example, to ensure a complete model, some approaches require images of all possible transitions of the environment [Amado et al., 2018b].

4.3 Recognition Input and Output

Observations play a fundamental role in *Goal Recognition*, providing the necessary input with key information for recognition approaches to infer the underlying intended goal of an observed agent [Meneguzzi and Pereira, 2021; Masters and Vered, 2021]. *Observations* can be a sequence of interactions in an environment, which are usually generated as a consequence of a plan execution (i.e., a sequence of executed actions) to achieve a particular goal, representing the agent’s behavior. The notion of *observations* Ω can vary based on the characteristics of the environment, the type of domain model employed, the rationality of the observed agents, and the level observability to which the recognizer is capable of perceiving

²Our definition of *Model-Approximate Goal Recognition* is analogous to model-based approaches in *Reinforcement Learning*, where the agent first learns an approximate model, and then plans on it.

	Papers	Environment				Technique				Input Data			Recognition Input (Observations)						Recognition Output	
		Discrete	Continuous	Image	Stochastic	Supervised Learning	Reinforcement Learning	Symbolic	Neuro-Symbolic	Traces	Images	Sampling	Actions	States	Images	Missing Observations	Noisy Observations	Online	Probability Distribution	Goal Ranking
Agnostic	[Min <i>et al.</i> , 2014]	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	[Min <i>et al.</i> , 2016]	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	[Amado <i>et al.</i> , 2018a]	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	[Maynard <i>et al.</i> , 2019]	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	[Borraj <i>et al.</i> , 2020]	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	[Amado <i>et al.</i> , 2022]	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	[Fang <i>et al.</i> , 2023]	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	[Chiari <i>et al.</i> , 2023]	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Model-Approximate	[Bauer, 1998]	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	[Bisson <i>et al.</i> , 2015]	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	[Geib and Kantharaju, 2018]	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	[Amado <i>et al.</i> , 2018b]	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	[Kantharaju <i>et al.</i> , 2019a]	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	[Kantharaju <i>et al.</i> , 2019b]	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	[Pereira <i>et al.</i> , 2019b]	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	[Polyvyanyy <i>et al.</i> , 2020]	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	[Shvo <i>et al.</i> , 2021]	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	[Hu <i>et al.</i> , 2021]	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	[Su <i>et al.</i> , 2023a]	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	[Ko <i>et al.</i> , 2023]	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	[Su <i>et al.</i> , 2023b]	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	[Amado <i>et al.</i> , 2023]	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

Table 1: *Model-Free Goal Recognition* approaches organized and characterized according to our taxonomy.

the behavior of the agents being observed.

A *sequence of observations* represents the behavior of the observed agent and can be *complete* from the perspective of the recognizer if it represents a **complete sequence of observations** that contains *all* interactions of the behavior of the agent for achieving its goals [Mirsky *et al.*, 2021, Section 2.3.2]. In contrast, a **partial sequence of observations** is an observation sequence *misses* interactions and information on the behavior of the agent for achieving its goals. Depending on the sensing capabilities of the recognizer, a *sequence of observations* can be precise and **accurate**, projecting the exact interactions observed from the behavior of the agent, or they can be **noisy** and represent spurious information regarding the behavior of the agent [Sohrabi *et al.*, 2016].

Representations for observations cover a wide spectrum, and the formalization of the observations depends directly on the environment and the domain model employed. In *Goal Recognition*, the observations fall into various types of representations, such as **symbolic observations** [Ramírez and Geffner, 2009; Sohrabi *et al.*, 2016; Mirsky and Gal, 2016; Pereira *et al.*, 2020] (e.g., *discrete*, *continuous* spatial information [Kaminka *et al.*, 2018; Pereira *et al.*, 2019b], mixed *discrete-continuous*, etc), **image-based observations** [Amado *et al.*, 2018b] (*pixel-level information within images*), **video streaming observations** [Granada *et al.*, 2017] (*visual input from a sequence of events via video*). *Symbolic observations* can be captured by logs, sensors, motion detectors, etc, whereas *image-based observations* can be captured by video cameras, pictures, etc.

The way that the observations are perceived by the recognizer plays a very important role in the recognition process. The observations are perceived in two ways: **offline** and **online**. *Offline* recognition involves a “retrospective” analysis of a sequence of observations, which is usually given at once for a recognizer. On the other hand, *online recognition* [Vered *et al.*, 2016] involves real-time analysis of a sequence of observations as the agent’s interactions occur within an environment, and the observations are usually given and revealed incrementally (i.e., one by one), and the recognition process

is performed multiple times, as the observations are revealed.

5 Model-Free Goal Recognition Approaches

In this section, we survey and review the existing approaches in the literature on *Model-Free Goal Recognition*. We organize this section in two parts: *Model-Agnostic* and *Model-Approximate* approaches. We explain in detail each of the approaches and then highlight their employed techniques and applications for *Model-Free Goal Recognition*, and conclude this section with a discussion regarding the limitations and constraints of the existing approaches.

5.1 Model-Agnostic Approaches

Agnostic Model-Free Goal Recognition can be viewed as a *multi-class classification* task. Indeed, most approaches we surveyed use supervised learning models at their core. To the best of our knowledge, the first approach that uses a recognition pipeline based on *Neural Networks* is the work of Min *et al.* [2014]. This approach uses *feed-forward n-gram* models to learn the player’s objective/intentions from a sequence of its actions in a game. In follow-on research, Min *et al.* [2016] develops an approach that uses Long Short-Term Memory (LSTM) to learn patterns in sequences of actions in a game. In both approaches, they use specific learned features as input to the Neural Networks instead of raw game-play events. Similarly, Amado *et al.* [2018a] employs LSTM networks for recognizing goals, using a pre-trained encoder and an LSTM network for representing and analyzing a sequence of observed states as images in image-based domains.

Maynard *et al.* [2019] compare symbolic inverse planning to *Deep Learning* (DL) using CNN and LSTM architectures using five synthetic benchmarks often used in the literature. It shows that it is possible to use deep learning approaches to predict goals quickly while remaining competitive with the baseline of Ramírez and Geffner in a narrow number of domains, if a very large training set is available. Borrajo *et al.* [2020], in a similar class of approach compares model-based (planning-based) and model-free (learning-based) approaches in various domains. They specifically discuss goal

recognition approaches in the context of finance-related tasks, finding that model-based approaches perform better when there is a partial order of actions in plans, while learning-based approaches excel when there is a correlation between actions and goals. The results highlight the trade-offs between the two approaches and suggest potential applications in the financial industry. Chiari *et al.* [2023] adopts LSTM in their GRNet approach, which addresses the challenge of limited observability in the actions of an agent's plan. It develops a more general approach to solve different goal recognition instances within the same domain using a singularly trained network. Finally, Amado *et al.* [2023] combine machine learning (ML) and symbolic reasoning by using ML to pre-process input data for planning algorithms. Their Predictive Plan Recognition (PPR) approach combines machine learning statistical prediction with domain knowledge within planning techniques, mitigating low and faulty observability, and solving both goal and plan recognition problems simultaneously. On the learning side they train a predictive statistical model of the most likely next states given a set of state observations. They combine such predictive models with symbolic heuristics for goal recognition to predict relevant states towards a goal hypothesis given a sequence of observations.

The final two approaches leverage *Reinforcement Learning*. *Goal Recognition as Q-Learning* (GRAQL) [Amado *et al.*, 2022], focuses on a two-stage process of offline learning and online inference. This approach uses learned Q-values implicitly representing the agents under observation instead of explicit goals from traditional GR. Goal inference works by minimizing the distance between an observation sequence and Q-values representing goal hypotheses or policies extracted from them. Fang *et al.* [2023] extend GRAQL to continuous domains, employing Twin Delayed Deep Deterministic Policy Gradient (TD3) *Deep Reinforcement Learning*. This approach models opponent behavior and learns policies in continuous environments, effectively addressing the challenges of representing infinite action-state pairs.

5.2 Model-Approximate Approaches

The seminal research on *model-approximate* recognition, such as Bauer's paper [1998], lays the foundation by detailing how to create plan-libraries from logged actions. Complementing this, Granada *et al.* [2017] and Bisson *et al.* [2015] innovate in integrating Neural Networks with plan-libraries, but with several distinct focuses. The former employs CNN for analyzing video streams, blending *Activity Recognition* with *Plan Recognition* to interpret real-time visual data. The latter, on the other hand, utilizes RNN to delve into an agent's decision-making process, aiming to predict actions based on learned behavioral patterns. Amado *et al.* [2018b] have extended Asai and Fukunaga's architecture by allowing both *Planning* and *Plan Recognition* tasks over the latent vectors. In this work, PDDL domain models are generated from latent vectors and used for *Goal Recognition*.

Geib and Kantharaju [2018], Kantharaju *et al.* [2019a] and Kantharaju *et al.* [2019b] explore several different approaches using *Combinatory Categorical Grammar* (CCG), a linguistic framework, for recognizing intentions and plans from action sequences. The first paper introduces LEXlearn, an algorithm

for learning probabilistic CCGs, providing a structural basis for understanding plans. The second paper enhances scalability in complex domains like Real-Time Strategy (RTS) games using Monte-Carlo Tree Search (MCTS) with CCGs, demonstrating an effective approach to handle larger, more intricate plan structures. Finally, the third paper focuses on extracting CCGs specifically for RTS games, showing how CCGs can be tailored for specific application areas.

Pereira *et al.* [2019b] develop recognition approaches over *Nominal Models*, i.e., continuous control domains with approximate transition functions. Their recognition approaches adapt existing recognition approaches for recognizing goals when dealing with approximate transition functions, and focus on understanding the influence of predictive errors, which are typical in identifying system dynamics, on the rates of recognition errors. This paper addresses the impact of inaccuracies (imperfections) in learned control domain models on the effectiveness of recognizing goals.

Shvo *et al.* [2021] present an approach, denoted as DISC (Discrete Optimization for Interpretable Sequence Classification), that learns interpretable sequence classifiers using finite state automata. The classifiers are learned through *Mixed Integer Linear Programming* (MILP) and offer interpretability, explanation, and counterfactual reasoning. The approach achieves comparable performance to LSTM-based classifiers while being more interpretable. Hu *et al.* [2021] develop a novel self-organizing *Neural Network* based inference model, which is able to learn compact rule sets through generalizing the streaming observations of an evader to perform intention recognition in navigation domains.

Recent papers employ *Process Mining* techniques for model-approximate recognition, involving the extraction of crucial information from event logs and traces to discover models for recognizing goals. *Process Mining* can help to understand how actual processes are performed, identifying bottlenecks, deviations, and opportunities for efficiency improvements. Polyvyanyy *et al.* [2020] develop a probabilistic recognition approach that relies on *Process Mining* techniques (for discovering models from plan *traces* or event logs), and they evaluate their recognition approach over well-known recognition benchmarks, as well as *Process Mining* benchmarks. Su *et al.* [2023a] extend and improve the approach in [Polyvyanyy *et al.*, 2020], and evaluate the new approach through a wide-ranging and detailed empirical evaluation with several different benchmarks against the state-of-the-art approaches in *Goal Recognition*. Su *et al.* [2023b] use recognition approaches based *Process Mining* in the context of transhumeral prostheses, demonstrating the applicability of *Goal Recognition* in biomedical applications. Ko *et al.* [2023] also employ *Process Mining* techniques but for recognizing both goals and plans, leveraging existing probabilistic *Trace Alignment* algorithms.

5.3 Limitations and Constraints

Deep Learning (DL) models, notably, require a substantial amount of training data. This demand can be particularly challenging in environments where data is limited or difficult to obtain emphasized also by [Amado *et al.*, 2018a; Borrajo *et al.*, 2020]. Moreover, the opacity of DL models

complicates their interpretability, making it difficult to explain the rationale behind their decisions. This challenge, as highlighted by Maynard *et al.* [2019], underscores a significant limitation in the application of Deep Learning to areas requiring transparent decision-making processes. Furthermore, the generalization capabilities of DL are often limited to the scope of the training data, posing difficulties in adapting to novel or significantly different scenarios.

Reinforcement Learning (RL) introduces a distinct set of challenges. The role of reward function design and the necessity for extensive interaction with the environment poses significant challenges. Inadequate reward functions can lead to undesirable behaviors, and RL’s sample inefficiency escalates computational demands, particularly in dynamic settings [Fang *et al.*, 2023] emphasized in *Deep RL*. Balancing the exploration of new actions with the exploitation of known rewards is complex, affecting the effectiveness and efficiency of RL. Amado *et al.* [2022] further emphasize these challenges, noting the difficulties RL faces in dynamic environments and its dependence on extensive offline learning, which complicates adaptation to frequently changing goal scenarios.

Symbolic approaches pose several different challenges. They rely on clear and well-defined logical constructs, requiring additional engineering to deal with incomplete or uncertain information. While machine learning can deal with such situations with additional data, symbolic systems trade data for additional domain expertise and knowledge engineering. While there is no obvious trade-off between these two approaches, there is potential for leveraging the advantages of both types of approaches. To this end, *Neuro-Symbolic* approaches, which blend learning via *Neural Networks* with *Symbolic Reasoning*, seek to harness the strengths of both paradigms. One of the challenges of such approaches is achieving scalability, particularly for large-scale problems. Amado *et al.* [2023] the data requirement is a limitation, which contrasts with the data independence typically seen in standard goal recognition methods, introducing complexity in scenarios where data is sparse or not readily available.

6 Related Problems

Model-Free Goal Recognition is closely related to several other problems and research topics, such as *Activity Recognition* [Van-Horenbeke and Peer, 2021], *Plan Recognition*, *Behavior Recognition*, and tangentially related to *Model Recognition*. There is an overlap between some of these related topics (including *Model-Free Goal Recognition*) in their similarities to classification problems, i.e., *Activity Recognition* and *Behavior Recognition*. Existing approaches to *Activity Recognition* [Dhattarwal and Ratnoo, 2023] and *Behavior Recognition* [Sur, 2021] leverage *Deep Learning* techniques, demonstrating significant success in achieving high accuracy across diverse tasks and datasets. Establishing a closer connection between *Model-Free Goal Recognition* and recent strides in *Activity Recognition* and *Behavior Recognition* could be beneficial to improve the development of more efficient *Model-Free Goal Recognition*.

Plan Recognition and *Goal Recognition* are often used interchangeably, but they differ in the sense that *Plan Recog-*

nition is a problem of *abduction* rather than deduction (like *Goal Recognition*), such that it aims to recognize a plan that best *explains* the observations. Regardless of the contrast and similarities of these tasks, we have seen few *Model-Free Plan Recognition* approaches [Amado *et al.*, 2023], establishing a pathway for research endeavors in this area.

Model Recognition [Aineto *et al.*, 2019b; Aineto *et al.*, 2020] is the task of identifying the model that “best” explains and captures a sequence of observations. We believe that, although it is not directly related to *Model-Free Plan Recognition*, this task, along with its existing symbolic solution approaches, could be used to underpin and improve *Model-Approximate Recognition* approaches.

As for *Goal Recognition* in *imperfect domain models* [Pereira *et al.*, 2019a; Zhuo, 2019; Pereira, 2020; Zhuo *et al.*, 2020], while it is not directly connected to *Model-Free Goal Recognition*, this task, along with existing approaches, could be the cornerstone and foundation for future developments in *Model-Free Goal Recognition*, paving the way to develop approaches that could ignore possible imperfections in a model, and perhaps even disregard the absence of a model.

7 Conclusions and Future Directions

Model-Free Goal Recognition is a flexible approach that mitigates the need for detailed models, which often requires domain experts. It is valuable in domains like robotics, human-computer interaction, and personalized learning systems where building such domain models is challenging. In this survey, we explored and defined new concepts, approaches, and applications of *Model-Free Goal Recognition*. We enumerate them and highlight their strengths and limitations. From data-driven and *Neuro-Symbolic* approaches to *Reinforcement Learning* and symbolic approaches, various approaches tackle the challenge of model-free goal recognition under different assumptions and constraints. We systematically organize such approaches, emphasizing the significant impact of observations, environmental properties, and recognition input on performance, showcasing the multifaceted nature of *Model-Free Goal Recognition*.

Recent advancements in *Machine Learning* are poised to produce more efficient approaches for *Model-Free Goal Recognition*. More sophisticated learning models (i.e., *Transformers*), could yield more robust *Model-Agnostic* recognition approaches capable of dealing with different types of data. Moreover, fusing data from various modalities, such as images, text, and sensor data, will yield new approaches capable of accommodating different types of observations. While still limited by *hallucination* issues, LLMs [Min *et al.*, 2023] could be used to yield a new class of *Model-Approximate* approaches that rely only on natural language to approximate a model, mitigating even more the need for a domain expert.

In closure, we expect that *Goal Recognition* will be a key driver in improving human-computer interaction, especially nowadays, when most of our interaction with others is non-verbal. Thus, understanding the goals aimed by our collaborators (i.e., humans, *Dialog Systems*, LLMs, etc.) and how they intend to achieve them allows anticipation of their behavior, leading to more enriching interactions.

References

- [Aineto *et al.*, 2018] Diego Aineto, Sergio Jiménez, and Eva Onaíndia. Learning STRIPS action models with classical planning. In *ICAPS*, 2018.
- [Aineto *et al.*, 2019a] Diego Aineto, Sergio Jiménez Celorio, and Eva Onaíndia. Learning Action Models with Minimal Observability. *Artificial Intelligence*, 275, 2019.
- [Aineto *et al.*, 2019b] Diego Aineto, Sergio Jiménez, Eva Onaíndia, and Miquel Ramírez. Model Recognition as Planning. In *ICAPS*, 2019.
- [Aineto *et al.*, 2020] Diego Aineto, Sergio Jiménez, and Eva Onaíndia. Observation decoding with sensor models: Recognition tasks via classical planning. In *ICAPS*, 2020.
- [Aineto *et al.*, 2022] Diego Aineto, Sergio Jiménez, and Eva Onaíndia. A Comprehensive Framework for Learning Declarative Action Models. *Journal of Artificial Intelligence Research*, 74:1091–1123, 2022.
- [Amado *et al.*, 2018a] Leonardo Amado, João Paulo Aires, Ramon Fraga Pereira, Maurício C. Magnaguagno, Roger Granada, and Felipe Meneguzzi. Lstm-based goal recognition in latent space. In *AAAI Workshop on Plan, Activity, and Intent Recognition (PAIR)*, 2018.
- [Amado *et al.*, 2018b] Leonardo Amado, Ramon Fraga Pereira, João Paulo Aires, Mauricio Magnaguagno, Roger Granada, and Felipe Meneguzzi. Goal recognition in latent space. In *IJCNN*, 2018.
- [Amado *et al.*, 2022] Leonardo R. Amado, Reuth Mirsky, and Felipe Meneguzzi. Goal Recognition as Reinforcement Learning. In *AAAI*, 2022.
- [Amado *et al.*, 2023] Leonardo Amado, Ramon Fraga Pereira, and Felipe Meneguzzi. Robust Neuro-Symbolic Goal and Plan Recognition. In *AAAI*, 2023.
- [Asai and Fukunaga, 2018] Masataro Asai and Alex Fukunaga. Classical planning in deep latent space: Bridging the subsymbolic-symbolic boundary. In *AAAI*, 2018.
- [Asai *et al.*, 2021] Masataro Asai, Hiroshi Kajino, Alex S. Fukunaga, and Christian Muise. Classical planning in deep latent space. *Journal of Artificial Intelligence Research*, 74:1599–1686, 2021.
- [Bauer, 1998] Mathias Bauer. Acquisition of abstract plan descriptions for plan recognition. In *AAAI*, 1998.
- [Bisson *et al.*, 2015] Francis Bisson, Hugo Larochelle, and Froduald Kabanza. Using a recursive neural network to learn an agent’s decision model for plan recognition. In *IJCAI*, 2015.
- [Borrajó *et al.*, 2020] Daniel Borrajó, Sriram Gopalakrishnan, and Vamsi K. Potluru. Goal recognition via model-based and model-free techniques. *arXiv/CoRR*, abs/2011.01832, 2020.
- [Chiari *et al.*, 2023] Mattia Chiari, Alfonso Emilio Gerevini, Francesco Percassi, Luca Putelli, Ivan Serina, and Matteo Olivato. Goal Recognition as a Deep Learning Task: The GRNet Approach. In *ICAPS*, 2023.
- [Dhatarwal and Ratnoo, 2023] Aayush Dhatarwal and Saroj Ratnoo. A review of deep learning techniques for human activity recognition. In *Hybrid Intelligent Systems*, 2023.
- [E.-Martín *et al.*, 2015] Yolanda E.-Martín, María D. R.-Moreno, and David E. Smith. A fast goal recognition technique based on interaction estimates. In *IJCAI*, 2015.
- [Fang *et al.*, 2023] Zihao Fang, Dejun Chen, Yunxiu Zeng, Tao Wang, and Kai Xu. Real-Time Online Goal Recognition in Continuous Domains via Deep Reinforcement Learning. *Entropy*, 25, 2023.
- [Fikes and Nilsson, 1971] Richard E. Fikes and Nils J. Nilsson. A formal basis for the heuristic determination of minimum cost paths. *IEEE Transactions on Systems Science and Cybernetics*, 4(2):152–162, 1971.
- [Geib and Kantharaju, 2018] Christopher W. Geib and Pavan Kantharaju. Learning combinatory categorial grammars for plan recognition. In *AAAI*, 2018.
- [Granada *et al.*, 2017] Roger Leitzke Granada, Ramon Fraga Pereira, Juarez Monteiro, Rodrigo Coelho Barros, Duncan D. Ruiz, and Felipe Meneguzzi. Hybrid activity and plan recognition for video streams. In *AAAI Workshop on Plan, Activity, and Intent Recognition (PAIR)*, 2017.
- [Guan *et al.*, 2023] Lin Guan, Karthik Valmeekam, Sarath Sreedharan, and Subbarao Kambhampati. Leveraging Pre-trained Large Language Models to Construct and Utilize World Models for Model-based Task Planning. *arXiv/CoRR*, abs/2305.14909, 2023.
- [Hu *et al.*, 2021] Yue Hu, Kai Xu, Budhitama Subagdja, Ah-Hwee Tan, and Quanjun Yin. Interpretable goal recognition for path planning with art networks. In *Proceedings of the International Joint Conference on Neural Networks (IJCNN)*, 2021.
- [Kaminka *et al.*, 2018] Gal A. Kaminka, Mor Vered, and Noa Agmon. Plan recognition in continuous domains. In *AAAI*, 2018.
- [Kantharaju *et al.*, 2019a] Pavan Kantharaju, Santiago Ontañón, and Christopher W. Geib. Extracting CCGs for Plan Recognition in RTS Games. In *AAAI Workshop on Knowledge Extraction from Games*, 2019.
- [Kantharaju *et al.*, 2019b] Pavan Kantharaju, Santiago Ontañón, and Christopher W. Geib. Scaling up CCG-Based Plan Recognition via Monte-Carlo Tree Search. In *IEEE Conference on Games (CoG)*, 2019.
- [Kerkez and Cox, 2002] Boris Kerkez and Michael T Cox. Case-Based Plan Recognition with Incomplete Plan Libraries. In *AAAI Fall Symposium on Intent Inference*, 2002.
- [Ko *et al.*, 2023] Jonghyeon Ko, Fabrizio Maggi, Marco Montali, Rafael Peñaloza, and Ramon Pereira. Plan Recognition as Probabilistic Trace Alignment. In *ICPM*, 2023.
- [Kumar *et al.*, 2023] Nishanth Kumar, Willie McClinton, Rohan Chitnis, Tom Silver, Tomás Lozano-Pérez, and

- Leslie Pack Kaelbling. Learning Efficient Abstract Planning Models that Choose What to Predict. *arXiv/CoRR*, 2208.07737, 2023.
- [Masters and Vered, 2021] Peta Masters and Mor Vered. What’s the Context? Implicit and Explicit Assumptions in Model-Based Goal Recognition. In *IJCAI*, 2021.
- [Maynard *et al.*, 2019] Mariane Maynard, Thibault Duhamel, and Froduald Kabanza. Cost-Based Goal Recognition Meets Deep Learning. In *AAAI Workshop on Plan, Activity, and Intent Recognition (PAIR)*, 2019.
- [McDermott *et al.*, 1998] Drew McDermott, Malik Ghallab, Adele Howe, Craig Knoblock, Ashwin Ram, Manuela Veloso, Daniel Weld, and David Wilkins. Pddl-the planning domain definition language. In *Conference on Artificial Intelligence Planning Systems (AIPS)*, 1998.
- [Meneguzzi and Pereira, 2021] Felipe Meneguzzi and Ramon Fraga Pereira. A Survey on Goal Recognition as Planning. In *IJCAI*, 2021.
- [Min *et al.*, 2014] Wookhee Min, E.Y. Ha, Jonathan Rowe, Bradford Mott, and James Lester. Deep Learning-Based Goal Recognition in Open-Ended Digital Games. In *AI-IDE*, 2014.
- [Min *et al.*, 2016] Wookhee Min, Bradford W. Mott, Jonathan P. Rowe, Barry Liu, and James C. Lester. Player Goal Recognition in Open-World Digital Games with Long Short-Term Memory Networks. In *IJCAI*, 2016.
- [Min *et al.*, 2023] Bonan Min, Hayley Ross, Elinor Sulem, Amir Pouran Ben Veyseh, Thien Huu Nguyen, Oscar Sainz, Eneko Agirre, Ilana Heintz, and Dan Roth. Recent Advances in Natural Language Processing via Large Pre-trained Language Models: A Survey. *ACM Computing Surveys*, 56(2), 2023.
- [Mirsky and Gal, 2016] Reuth Mirsky and Ya’akov Gal. Slim: semi-lazy inference mechanism for plan recognition. In *IJCAI*, 2016.
- [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.
- [Pereira and Meneguzzi, 2018] Ramon Fraga Pereira and Felipe Meneguzzi. Goal recognition in incomplete domain models. In *AAAI*, 2018.
- [Pereira *et al.*, 2019a] Ramon Fraga Pereira, André 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.
- [Pereira, 2020] Ramon Fraga Pereira. Goal recognition over imperfect domain models. *Ph.D. Thesis*, 2020.
- [Polyvyanyy *et al.*, 2020] Artem Polyvyanyy, Zihang Su, Nir Lipovetzky, and Sebastian Sardina. Goal Recognition Using Off-The-Shelf Process Mining Techniques. In *AA-MAS*, 2020.
- [Ramírez and Geffner, 2009] Miquel Ramírez and Hector Geffner. Plan Recognition as Planning. In *IJCAI*, 2009.
- [Sanner and Boutilier, 2010] Scott Sanner and Craig Boutilier. Relational dynamic influence diagrams and the decision-theoretic approach to rddl. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence (UAI)*, 2010.
- [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.
- [Schmidt *et al.*, 1978] Charles F. Schmidt, N. S. Sridharan, and John L. Goodson. The Plan Recognition Problem: An Intersection of Psychology and Artificial Intelligence. *Artificial Intelligence*, 11, 1978.
- [Shvo *et al.*, 2021] Maayan Shvo, Andrew C. Li, Rodrigo Toro Icarte, and Sheila A. McIlraith. Interpretable Sequence Classification via Discrete Optimization. In *AAAI*, 2021.
- [Sohrabi *et al.*, 2016] Shirin Sohrabi, Anton V. Riabov, and Octavian Udrea. Plan Recognition as Planning Revisited. In *IJCAI*, 2016.
- [Su *et al.*, 2023a] Zihang Su, Artem Polyvyanyy, Nir Lipovetzky, Sebastian Sardina, and Nick van Beest. Fast and accurate data-driven goal recognition using process mining techniques. *Artificial Intelligence*, 323, 2023.
- [Su *et al.*, 2023b] Zihang Su, Tianshi Yu, Nir Lipovetzky, Alireza Mohammadi, Denny Oetomo, Artem Polyvyanyy, Sebastian Sardina, Ying Tan, and Nick Beest. Data-Driven Goal Recognition in Transhumeral Prostheses Using Process Mining Techniques. In *ICPM*, 2023.
- [Sur, 2021] Behaviour recognition of pigs and cattle: Journey from computer vision to deep learning. *Computers and Electronics in Agriculture*, 187, 2021.
- [Van-Horenbeke and Peer, 2021] Franz A. Van-Horenbeke and Angelika Peer. Activity, plan, and goal recognition: A review. *Frontiers Robotics AI*, 8:643010, 2021.
- [Vered *et al.*, 2016] Mor Vered, Gal A Kaminka, and Sivan Biham. Online goal recognition through mirroring: Humans and agents. In *The Fourth Annual Conference on Advances in Cognitive Systems*, volume 4, 2016.
- [Zhuo *et al.*, 2020] Hankz Hankui Zhuo, Yantian Zha, Subbarao Kambhampati, and Xin Tian. Discovering underlying plans based on shallow models. *ACM Transaction Intelligent Systems and Technology*, 11(2), 2020.
- [Zhuo, 2019] Hankz Hankui Zhuo. Recognizing Multi-Agent Plans When Action Models and Team Plans Are Both Incomplete. *ACM Transactions on Intelligent Systems and Technology*, (3), 2019.