

What’s the Context? Implicit and Explicit Assumptions in Model-Based Goal Recognition

Peta Masters¹ and Mor Vered²

¹School of Computing and Information Systems, The University of Melbourne

²Dept of Data Science and AI, Monash University

peta.masters@unimelb.edu.au, mor.vered@monash.edu

Abstract

Every model involves assumptions. While some are standard to all models that simulate intelligent decision-making (e.g., discrete/continuous, static/dynamic), goal recognition is well known also to involve choices about the observed agent: is it aware of being observed? cooperative or adversarial? In this paper, we examine not only these but the many other assumptions made in the context of model-based goal recognition. By exploring their meaning, the relationships between them and the confusions that can arise, we demonstrate their importance, shed light on the way trends emerge in AI, and suggest a novel means for researchers to uncover suitable avenues for future work.

1 Introduction

Goal recognition (GR) is the problem of determining an agent’s intent from observations of its behaviour. It falls under the scope of “plan, activity and intent recognition” and is that aspect of the problem which aims to identify top-level or end-goals rather than steps on the way to a goal’s achievement (though identification of the steps is often part of the solution).

GR is a fundamental problem in AI, tackled in numerous ways: traditionally using plan libraries against which observed actions are matched [Kautz and Allen, 1986; Charniak and Goldman, 1993]; more recently taking a model of the domain as problem input, then solving that problem using planning techniques [Baker *et al.*, 2007; Ramirez and Geffner, 2011; Pereira *et al.*, 2019a, etc.].

Recognition, in general, is a growing area of research, increasingly important as AI agents interact with humans and other AI systems for the purposes of real-world applications such as medical support, self-driving vehicles and social care. The field continues to expand, now including model recognition [Aineto *et al.*, 2019] and model-free GR, which uses machine learning and, rather than an explicit model, relies instead on combinations of features that may or may not include aspects of the domain description [Borrajo *et al.*, 2020].

In this paper, we restrict our attention to model-based GR whereby, instead of relying on pre-existing plans from plan libraries, either new plans or sets of landmarks (i.e., actions

or states that must be realised for a goal to be achieved) are generated over a model of the domain and compared with observed behaviour on-demand.

Model-based GR has flourished in recent years, partly thanks to its demonstrated success but also because it is conceptually simple, relatively easy to implement and because there is a wealth of alternative systems from which practitioners can choose and on which theorists can continue to build.

But every model is a simplification. The only reliably accurate model of the world is the world itself. Thus, every model used for model-based GR is constrained by numerous assumptions; some standard, familiar from countless AI research papers and text books [Russell and Norvig, 2010]; others are GR-specific and relate not only to the environment but to the point-of-view of the observed agent, the nature of the observations and the type of goals. Even the overarching purpose of the GR task may have an impact, implicitly evoking assumptions about the notional observer and their perspective on the observed agent and underlying domain.

For this paper, we examined research from consistent contributors to the field, dating from the inception of model-based GR (independently arrived at by Baker *et al.* [2009], Ramirez and Geffner [2009] and Pattison and Long [2010]). We identify 20 distinct assumptions that are frequently made, sometimes explicitly, often implicitly. Where previous surveys have focused on methodology [Sukthankar *et al.*, 2014; Chakraborti *et al.*, 2019], our focus is on the context within which the various models are expected to succeed. Our findings are illuminating, not only with respect to GR but for what they reveal about the dangers of implicit and imprecise assumptions; and for what examination of assumptions tells us about the way that trends emerge within an AI discipline.

There will never be one unified model of GR that represents everything it would be helpful to know because too much about the real world is unknowable so must be guessed or assumed [Kay and King, 2020]. In what follows, we examine assumptions commonly made in the context of model-based GR: their meaning, the relationships between them and the confusion that can arise when they are poorly stated or not stated at all. By this we hope: (a) to help theorists clarify the decisions they need to consider and ought to specify; (b) to help practitioners select the GR models most useful for their particular purpose; and (c) to help researchers identify assumptions they might usefully relax in future work.

2 Assumptions in Goal Recognition

GR involves *observations* of one or more *agent/s* aiming for one or more *goal/s* in an *environment*. Decisions made about each of these features directly impact the performance and suitability of the GR system.

In this section, we set out key assumptions we encountered in the literature and the ways that they relate to model-based GR. We have asterisked the *default* assumptions: those properties that (referring to the model definitions, experimental domains, etc.) seem to apply when not specified in the text. Unsurprisingly, the default assumption typically represents the greatest simplification; it is the non-default that tends to more closely match real world complexity and usually therefore represents the greater challenge.

2.1 Assumptions about the Environment

The terms we use for assumptions about the environment largely coincide with “Properties of task environments” as listed in [Russell and Norvig, 2010, p.41]. Although—and perhaps *because*—the terms are so familiar, they still leave room for misinterpretation: it is when we think everyone understands that we make the least effort to explain.

The primary source of confusion, here and throughout, is that there are at least two players in a GR environment: the GR system, for brevity the *observer*; and the observed agent whose goal is to be determined, the *actor*. We take the view that, unless otherwise stated, assumptions relate to *the observer’s model of the world that the actor inhabits* (itself a clarifying assumption perhaps worth stating in future work).

We begin with the three least problematic assumptions. The environment may be *discrete** or *continuous*. That is, the observer may represent the actor’s world as a finite number of individual states (e.g., for a gridworld) or by “real” values on a continuous spectrum (e.g., to describe a robot’s moving parts). It may be *static**, that is, unchanging (e.g., a fixed maze) or *dynamic*, liable to change independently of the actor’s behaviour (e.g., a real-world road network navigated by an autonomous vehicle). It may be *deterministic** such that, if the observer knows the current state and the action taken, it knows precisely which state will result or *stochastic* with the observer only able to evaluate the *probable* result (e.g., given a physics engine, the actor drops a glass: usually the next state includes broken glass, but not always).

In a GR environment, the observer receives information as a set or sequence of *observations*. These may either be available to the observer before processing begins and handled *offline** or delivered incrementally and processed *online*, as they are received.¹ Online processing implies a temporal sequence of observations, which is sometimes specified explicitly [Shvo *et al.*, 2018]. In a GR context, some combinations of environmental assumptions tend to be related. For example, a dynamic environment tends to imply *online* processing: if the environment may change, it needs to be constantly monitored. Furthermore, if the environment is dynamic, it is likely also to be stochastic: any external force adds uncertainty to the potential outcomes of actions, which in turn suggests a need for online processing.

¹Russell and Norvig [2010] use the terms sequential/episodic.

The notions of *full* and *partial observability* can be confusing in a GR context, where—with the exception of prior knowledge—everything the observer knows is delivered via their observations. In most AI domains, these terms describe the environment in general (e.g., the state of a chessboard is fully observable whereas the hands in a Texas hold-em poker game are not). Consider this, however, from [Pereira *et al.*, 2019b, p.5548]: “We . . . define the sequence of observations *O* to be a *partially observed* trajectory of states. . .” (our emphasis). This is common usage in GR. It means that while some, perhaps many, of the events that occur in the environment may be observed, others may be missing, as is typically the case. For clarity, we use the term “observability” as an environmental assumption only, referring to the quality of information conveyed by *each individual observation*. We discuss partial sets and sequences of observations—the “partially observed trajectory” of Pereira *et al.*—under “missing” (see Section 2.3). We summarise this and other confusing assumptions in Table 1.

In the case of full/partial observability, there is no default as such. If not stated, we derive implicit assumptions from the domain (e.g., partial in a classical task-planning domain, where *actions* are observed rather than complete states, but full in a path-planning domain).

Further confusion arises with respect to whether the actor inhabits a *single-** or *multi-agent* environment. In contemporary work, by default, GR involves the intent of a *single* actor in a *single-agent* environment (that is, one actor, with no other agents present). In all other situations, conditions need to be stated more precisely. If a GR model is explicitly declared to be performing “multi-agent” recognition that *probably* means that the system will try to determine the purpose of common group behaviour or will take into account the behaviour of multiple agents and try to determine the intentions of them all. It is, however, entirely feasible for a system to perform GR with respect to a single agent, which itself happens to be operating in a multi-agent environment. Moreover, it is reasonable to regard the very act of GR as implying a second *observing* agent which, particularly in the context of human-aware computing, seems to place traditional, single-agent recognition in an apparently multi-agent environment. Though highly relevant to many real-world GR problems, we have found that this assumption is rarely made explicitly.

In the [Russell and Norvig, 2010] list, a domain model may be *known** or *unknown*. We had not intended to include this as an assumption since, although we are aware of model-free GR emerging from the machine learning community, it seemed self-evident that, in model-based GR, the model must be known. Recent work from Pereira *et al.* [2019a; 2019b] challenges our preconception, however, since it deals with models that are *partially* known. We felt it important, therefore, to include reference to this property as a clear recent example of the way that identification of assumptions not previously relaxed can open opportunities for valuable future research (as further discussed in Section 3).

2.2 Assumptions about the Goals

It seems obvious that *goal* recognition must be all about recognising goals. However, it falls within the broader scope

ASSUMPTION	THIS?	OR THIS?
partial observability equal priors multi-agent noise rational agent full observability goal recognition purpose of GR	partially-observed state equal priors throughout multiple actors false observations (didn't occur) cost-efficient fully observed, all remembered also known as plan recognition determine the goal	partial observation sequence equal priors at first timestep single-actor in multi-agent environment suboptimal behaviour (did occur) intentionally suboptimal unimpeded but may be forgotten may involve plan recognition determine observer's belief

Table 1: Confused/Confusing Assumptions.

of “plan, activity and intent recognition” [Sukthankar *et al.*, 2014]. Moreover, model-based GR uses *planning* techniques to solve the GR problem. This causes some confusion as to exactly what is being recognised. The seminal “Plan Recognition as Planning” [Ramirez and Geffner, 2009] concerns goal recognition whereas “Plan Recognition as Planning Revisited” [Sohrabi *et al.*, 2016] explicitly aims to identify plans *and* goals. Several papers explicitly state that they use the terms *plan recognition* and *goal recognition** interchangeably [Masters and Sardina, 2017a; Keren *et al.*, 2015]. Arguably, GR can be compiled into plan recognition (if you have identified the plan, you have identified the goal [Demolombe and Hamon, 2002]), but as Masters and Sardina [2017a] show in relation to path-planning, the reverse is not always the case. We argue, therefore, that for the sake of clarity and for the benefit of potential users, these distinctions should be maintained and the terms only declared interchangeable if both tasks are demonstrably supported.

Many approaches to GR give, as a solution to the problem, a probability distribution across the goals, indicating which is the most likely. All such papers that we examined used *Bayesian**-type reasoning, making it clearly the default. Some early practitioners, however, favoured Dempster-Shafer’s *possibilistic* approach [Carberry, 1996]. Despite the lack of recent representation, we mention the assumption, since it is one that might usefully be re-explored.

When probabilities are used, authors often make the “simplifying” assumption “without loss of generality” that the prior probability distribution across goals is uniform. The implication is that, although *equal priors** are explicitly assumed and have been used experimentally or for the purpose of examples, *unequal priors* are also supported. In this case, despite an *explicitly stated assumption*, the approach remains ambiguous since it is not always clear whether priors are (a) always equal and can therefore be ignored, (b) unequal and must therefore be incorporated, moreover whether they should be (c) updated following each new observation or (d) frozen throughout the process which, as Baker *et al.* [2011] point out, may produce a different result. This topic is thoroughly covered in [Masters *et al.*, 2021].

Other assumptions with respect to goals seem somewhat under-researched. By default the actor is expected to pursue a *single goal**. It may, however, be pursuing *multiple* (unordered) goals, which arguably could be compiled into one, or an ordered sequence of *sub-goals*, which could not. We have found that there exists an overlap between the notion of *sub-*

goals and the use of *landmarks*. Just as Sohrabi *et al.* [2016], for example, identify the most probable plan on the way to determining the most probable goal, so Pattison and Long [2010] (amongst others) identify landmarks—effectively *sub-goals*—on the way to identifying the final, top-level goal. Lastly, there is an implicit assumption that the actor’s goals remain *unchanging** throughout the process. Relaxation of this assumption was a key motivation for Baker *et al.*, who first explored the notion of *dynamic goals* in 2005, but that line of research does not appear to have gained traction.

2.3 Assumptions about the Observations

Observations are fundamental to GR. Directly or indirectly, they convey the actor’s behaviour, making it possible for the observer to formulate a hypothesis. We have already mentioned in online/offline and full/partial observability, environmental assumptions that apply to observations. We now consider three further conditions, specific to GR, which can be confusing if omitted or applied without clarification.

First, we need to know whether the observations are *complete* in the sense that all relevant behaviour was observed or if, as is usually the case, some observations may be *missing** in that some significant actions or events went unobserved. As discussed, this property is sometimes described as—and can be confused with—partial observability (see quotation from [Pereira *et al.*, 2019b] above).

Secondly, does the observer expect its observations to be *accurate** (i.e., to conform to real events) or does it accommodate *noisy* observations, which are essentially errors that, although registered as observations, did not actually occur? In the literature, we find “noise” used in two subtly different ways: to mean errors, as defined above, but also to describe *suboptimality*: observed behaviour that occurred and intended, but failed, to achieve a fully optimal plan [Ramirez and Geffner, 2010, p.1121]. Although the distinction may be apparent from the context, we take the view that misapplication of the term can lead researchers to believe that they have accommodated *both* properties when they have only accommodated one. Here we use—and recommend the use of—“noise” to describe behaviour that *did not occur*, caused perhaps by faulty recording equipment [Masters and Sardina, 2019b, 2021]; and reserve “suboptimality” to convey a property of the actor which reflects non-optimal behaviour that actually happened (see rationality, Section 2.4).

Thirdly, what is the nature of an observation? Model-based GR having arisen out of classical planning, the default ob-

ervation is an *action** but we also see *fluents* [Sohrabi *et al.*, 2016], full *states* [Masters and Sardina, 2017a] and *action/state* pairs [Shvo and McIlraith, 2020]. This assumption is critical to implementation. If a domain is stochastic, for example, and we observe fluents or states but not actions, we have no way of knowing whether the observed effect was caused deliberately. We may not even know what action the agent performed. This assumption is usually explicitly made by the modelers and largely depends on the empirical domain through which the approach is to be tested. Nevertheless, confusion does arise in connection with a popular technique introduced by Sohrabi *et al.* [2016] who showed that, while their model was extended to observe fluents (on the basis that the effects of an action are more likely to be observed than the action itself) it can be made apparently to observe actions by assigning a unique fluent per action that may be added only by that action. This can give the impression that the options are so closely linked that the distinction barely matters. That is not the case, however. The mechanism is introduced at the compilation stage and can only work post-hoc on an existing model that initially considers the observation of actions.

2.4 Assumptions about the Agents

GR can be regarded as involving at least two agents: the actor, whose goal is being recognized, and the observer, represented by the GR system itself, which formulates the hypothesis.

Key elements in the relationship between the two are captured by the traditional classification of GR into three types: *keyhole**, where the actor is unaware of, or unaffected by, being observed; *intended*, where the actor is cooperative, aware of its observer and trying to make its intentions clear; or *deceptive*, attempting to obscure its intentions or implying one goal while striving for another [Carberry, 2001]. Recall that these types are all referenced from the point-of-view of the observer. Thus, a default observer that always believes itself to be performing keyhole recognition is regarded as *naïve*: a commonly cited simplifying assumption in deceptive planning [Masters and Sardina, 2017b, p.4374].

One frequently stated assumption is that the actor is *rational*. That is, the GR system/observer assumes that the actor is attempting to maximise some “performance measure” [Russell and Norvig, 2010, p.37]. Ostensibly straightforward, the term is confusing (a) because it is sometimes used to mean *fully optimal* [Ramirez and Geffner, 2009; Shvo and McIlraith, 2020], at others *suboptimal but cost-sensitive** [Vered *et al.*, 2016] (the definition we prefer) and (b) because the actor’s optimality is usually evaluated *with respect to the observer’s cost model* rather than its own. Yet, as we clearly see in the context of deception, an actor that is optimal/ rational with respect to its own cost model may appear *suboptimal* [Keren *et al.*, 2015] or downright *irrational* from the point-of-view of the observer.

As suggested by the above, for the most part, there is an assumption that actor and observer share the *same model**, not only in terms of costs but in its entirety [Vered and Kaminka, 2017; Vered *et al.*, 2018a]. The assumption has implications for implementation. If observer and actor share identical reasoning capabilities and knowledge of the domain then GR can be achieved self-referentially by asking, “If *I* were per-

forming those actions, what would *I* be trying to achieve?” If, on the other hand, the actor is assumed to have a *different model* from the observer—if, for example, a human is observed by a machine or vice versa—then the system must generate an independent model of the actor’s assumed internal beliefs [Baker *et al.*, 2009; Baker *et al.*, 2011].

The final assumption in this section relates to the nature of the observer rather than the actor. Recent work, particularly that which explores human-robot interaction, points up the distinction between a GR system that aims to determine the goal of the actor and one which aims to determine how the observed behaviour is most likely to be *interpreted* [Zhang *et al.*, 2017; Masters *et al.*, 2021]. Many of the papers grouped under “Extended GR” in Table 2 fall into this category. A key distinction is that, whereas GR is concerned with determining the ground truth (what is the agent’s goal?) extended GR is concerned with a hypothetical third party observer’s most likely *belief* (what would an observer believe to be the agent’s goal?)—which may be wrong. By default, the observer is assumed *competent**, that is, *not* wrong: fully aware of the observations it has registered and capable of processing them at least as efficiently as the actor it is observing. Whether human, human-like or robotic, however, the hypothetical observer may be known by the system to be *fallible*. That is, the observer may make mistakes—cognitive or processing errors—which an extended GR system can, if cooperative, accommodate or may, if adversarial, exploit [Kulkarni *et al.*, 2019].

3 Discussion

The list of assumptions examined in this survey is by no means exhaustive. The majority have been explicitly referenced by one or other of the consistent recent contributors to the field of model-based GR. Others—notably, possibilistic versus probabilistic reasoning and dynamic versus static goals—flag directions for research once keenly championed by prominent members of the plan recognition community.

In Section 2, we asterisked the default assumptions. In Table 2, we next highlight *contributions* by classifying papers according to the more challenging properties that are accommodated when the associated assumptions are relaxed. By this, we aim to provide a visual summary of the contributions and constraints of each approach, to assist practitioners in identifying which models best suit their needs and to help theorists identify gaps in the research.

The columns are organised according to assumptions made about the environment, observations, agents (both actor and observer) and goals; the rows represent distinct contributions. Wherever the paper accommodates (or claims to accommodate) a property, there is a tick; where it does not (or claims not), there is a cross. We write N/A to address cases where the assumption is irrelevant, such as non-equal priors when probabilistic reasoning is not involved. In classifying the papers, we considered assumptions which were addressed either explicitly in the text or implicitly via the empirical evaluation, framework definition, implication of the associated domain, and so on. Where an assumption has been explicitly referenced, we highlight the tick or cross with a circle.

	Papers	Environment				Observations			Agents			Goals						
		Continuous	Online	Dynamic	Stochastic	Partial Observability	Multi Agent	Missing	Noisy	Fluents/States	Non-Keyhole	Different Model	Suboptimal	Fallible Observer	Non-Equal Priors	Multi-Goals	Sub-Goals	Dynamic Goals
Approach	Bayesian Theory of Mind	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×
		×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×
Landmarks	[Pattison and Long, 2010]	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×
	[Pattison and Long, 2011]	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×
	[Pereira <i>et al.</i> , 2017]	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×
	[Vered <i>et al.</i> , 2018a; 2018b]	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×
	[Ramirez and Geffner, 2009]	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×
	[Ramirez and Geffner, 2010]	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×
	[Ramirez and Geffner, 2011]	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×
	[Zhuo <i>et al.</i> , 2012]	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×
	[E-Martín <i>et al.</i> , 2015]	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×
	[Sohrabi <i>et al.</i> , 2016]	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×
Plan Recognition As Planning	[Vered <i>et al.</i> , 2016]	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×
	[Vered and Kaminka, 2017]	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×
	[Masters and Sardina, 2017a]	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×
	[Kaminka <i>et al.</i> , 2018]	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×
	[Shvo <i>et al.</i> , 2018]	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×
	[Masters and Sardina, 2019a]	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×
	[Masters and Sardina, 2019b]	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×
	[Shvo and McIlraith, 2020]	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×
	[Keren <i>et al.</i> , 2014]	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×
	[Keren <i>et al.</i> , 2015]	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×
Extended GR	[Keren <i>et al.</i> , 2016; 2019]	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×
	[Wayllace <i>et al.</i> , 2016; 2017]	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×
	[Zhang <i>et al.</i> , 2017]	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×
	[Masters and Sardina, 2017b]	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×
	[Pozanco <i>et al.</i> , 2018]	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×
	[Kulkarni <i>et al.</i> , 2019]	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×
	[Masters <i>et al.</i> , 2021]	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×	×

Table 2: Categorisation of papers by contribution. The columns represent properties accommodated when simplifying assumptions are relaxed. Circles indicate explicit assumptions. Ambiguity (e.g., where claims are made that are not empirically demonstrated) is not currently shown; in such cases, we have favoured the authors' declared assumption. A practitioner seeking an appropriate model should look for ticks in columns of interest; a researcher seeking avenues for future work should look for crosses. We note that the list of assumptions examined in this survey is not exhaustive; they were selected on the basis of having been explicitly mentioned by one or more consistent contributors to the field. As discussed inline, fresh assumptions could be added by reference to related disciplines and/or real-life scenarios. Not all assumptions discussed in this table.

The field of model-based GR is largely built on independent contributions from Baker *et al.* [2007], Ramirez and Geffner [2009] and Pattison and Long [2010]. All three recognised the potential relationship between plan recognition (of which goal recognition is a sub-problem) and planning, which had previously seemed to be a separate branch of AI. Their approaches are quite different, however, and we have categorised them separately in the table. Baker *et al.* took inspiration from cognitive science and, using POMDPs, developed the idea of “Bayesian Theory-of-Mind”, Ramirez and Geffner focused on classical planning transformations, while Pattison and Long responded to the need for domain-independence by developing, in a PDDL environment, perhaps the first GR approach to involve the use of landmarks.

Of the three approaches, the table shows Ramirez and Geffner’s Plan Recognition as Planning (PRP) approach, in particular, has been extended by many subsequent authors. This can undoubtedly be attributed partly to the accessibility of the underlying framework—and to that of their probabilistic model which soon followed [Ramirez and Geffner, 2010]—but partly also to the visibility of their work to the broader AI community via IJCAI and AAAI.

Many conclusions can be drawn from this method of tabulation. First, it appears that, when a paper’s contribution is the invention of a novel approach, researchers tend to experiment with a model close to the default. As the approach gains traction, the same or subsequent researchers look for non-default directions into which the new ideas can be extended.

This clear progression is readily seen in the PRP stream of work. In the first of those papers [Ramirez and Geffner, 2009], almost every column shows a cross. That is, almost every assumption, explicit or non-explicit, was a default assumption.² In the second paper, we again see primarily defaults, but now suboptimality is addressed with the introduction of probabilistic reasoning (implied by consideration of priors). In subsequent papers, we see the relaxing of first one assumption, then another, as authors explore distinct aspects of the problem. Sohrabi *et al.* [2016], for example, are the first to look at noise in this context and incorporate fluents into their definition of the observations. Vered *et al.* [2016] continue what becomes a trend of using fluents over actions (later *in addition* to actions), and is the first to focus on continuous, online domains; in the most-recently listed paper under PRP, Shvo and McIlraith [2020] begin to look at dynamic environments.

The progressive effect is magnified when limitations that have *not* been addressed—which show up in our table as crosses—are *explicitly* mentioned by a researcher. When this occurs, the issue becomes more likely to be handled in future work. Consider, for example, how Baker *et al.* [2009] explicitly reference partial observability, missing and noisy observations, then address precisely these issues in [Baker *et al.*, 2011]. Similarly, Keren *et al.* repeatedly assert reliance on a deterministic (non-stochastic) domain for their work on GR design until the issue is apparently registered

²Ticks only appear against *partial observability* and *missing*: the more challenging properties (and therefore those shown in the table) but the defaults for GR in a task-planning environment.

and addressed [Wayllace *et al.*, 2016; Wayllace *et al.*, 2017]. Table 2 thus illustrates the power and importance of explicit assumptions—including the seemingly negative assumption of default properties—to drive research.

The usefulness of GR in general, combined with the accessibility of the model-based approach, has led to its incorporation into related areas of AI as a means of accessing what would otherwise be the internal workings (i.e., goals and intentions) of agents operating in known domains. The papers listed under “Extended GR” in the lower part of Table 2 fall into this category. They all incorporate model-based GR, not to determine an agent’s goal per se but in order to perform some other task. In the context of human-aware planning, Zhang *et al.* [2017] describe the relationship between GR and this “meta” approach like this: “In human-aware planning, the challenge is to obtain the human model which captures human capabilities, intents and etc. The modeling in this work is one level deeper: it is about the interpretation of the agent model from the human’s perspective...” [p.1314]

The extended GR contributions touch on the keyhole/intentional/adversarial assumption, since the relationship between observer and actor is paramount. The approach here, however, is *from the point-of-view of the actor*. Extended GR models use GR to assess the impact of their plans on a hypothetical observer. Thus, those that deal with *intentional* recognition [Zhang *et al.*, 2017; Kulkarni *et al.*, 2019], *flip* the “non-keyhole” assumption, since the actor and observer *both* expect intended recognition to occur. But those attempting *deception* [Masters and Sardina, 2017b; Masters *et al.*, 2021] assume a *naïve* observer—that is, an unsuspecting observer, performing keyhole recognition in the usual way—so the default for that property stands. We are unaware of any GR-based deceptive framework involving a *non-naïve* observer. Kulkarni *et al.* [2018] offer such an approach, but it explicitly avoids the use of model-based GR.

We have included Goal Recognition Design (GRD) under the umbrella of Extended GR (the first four papers in that category), while recognising that it approaches the problem from a different, unique angle. GRD asks, how can we modify the domain so that an agent acting within it reveals its objective as early as possible? Papers on this theme again demonstrate the way research progresses by relaxing assumption after assumption. Where the topic was first introduced [Keren *et al.*, 2014], the table shows a predominance of crosses, indicating that challenges were minimised in order to focus on the novelty of the approach itself. Subsequent papers relax different assumptions: [Keren *et al.*, 2015] relaxes the optimality assumption; [Keren *et al.*, 2016] relaxes the need for complete observations; [Wayllace *et al.*, 2016] restores the optimality assumption but addresses GRD in a stochastic environment.

Close examination of the table reveals the strengths and limitations of numerous contemporary approaches to model-based GR. We see that Sohrabi *et al.* [2016] and Shvo and McIlraith [2020] are highly focused on aspects of GR associated with observations. Zhang *et al.* [2017] dominates with respect to their treatment of observer-agent interaction. Baker *et al.* [2009; 2011] remains the only practitioner that we encountered whose model explicitly handles dynamic goals. Shvo and McIlraith [2020] and Pozanco *et al.* [2018] are

alone in considering dynamic environments; and it is tantalising to consider how their work might be combined: the one on active goal recognition, whereby the observer can intervene directly in the actor's world; the other on counter-planning, which uses GR to anticipate which landmarks the actor will need to achieve, in order to block or (potentially) facilitate their progress.

Relaxation of the *known environment* assumption, discussed in Section 2.1, is omitted from the table since it is largely the province of model-free machine learning techniques. As noted, however, recent work by Pereira *et al.* [2019a; 2019b] addresses the possibility of a *partially* known domain, where the couplings between state variables, actions and perturbations were not represented explicitly. They were able to address it by incorporating deep learning methods. We note that Zhang *et al.* [2017] also uses learning to label sub-goals. Indeed, this is a strength of machine learning and we are likely to see a lot more crossover between model-free and model-based approaches to GR in the future.

For the benefit of those inclined to adopt our approach to identify avenues for future work in GR or other AI disciplines, we suggest three ways to proceed: (1) generate a table based on recent contributions to the field, such as Table 2, and check for a predominance of crosses in any given column; (2) add columns for assumptions *not* currently tackled by reference to assumptions that arise in related disciplines (in our case, for example, machine learning or network security); (3) consider the key elements being modelled (in our case goals, observations, actor/observer and the environment), identify what it is about that element that makes *real-world* scenarios more challenging than those seen in contemporary models and add further columns to represent those properties. As an example of (2), consider assumptions in the machine-learning domain with respect to independent versus non-independent data [Mountrakis *et al.*, 2011]. In GR, data is conveyed to the system via observations. To date (to our knowledge) no GR model has explicitly considered independent observations. It is the sort of assumption that might drive a new line of research sparking insights into methods for dealing with missing observations or deciding when to repeat an online GR process. As an example of (3), consider, in the context of GR, an actor that changes its mind; or a goal that is not one of a known set but instead has to be inferred.

Finally—whilst we acknowledge that many of the papers we have discussed are conference papers with a limited page count—we suggest that assumptions are just as important to a paper's technical quality as its formulae and theorems. Every model involves assumptions; but they are of no use or interest if they are not stated, not clear or not correct.

Acknowledgements

We thank Miquel Ramirez and our anonymous reviewers for their comments and suggestions. The first author was supported by grant DP180101215 from the Australian Research Council.

References

[Aineto *et al.*, 2019] Diego Aineto, Sergio Jiménez, Eva Onaindia, and Miquel Ramírez. Model recognition as planning. In *Pro-*

ceedings of the International Conference on Automated Planning and Scheduling, volume 29, pages 13–21, 2019.

- [Baker and Tenenbaum, 2014] Chris L Baker and Joshua B Tenenbaum. Modeling human plan recognition using bayesian theory of mind. *Plan, activity, and intent recognition: Theory and practice*, pages 177–204, 2014.
- [Baker *et al.*, 2005] Chris Baker, Rebecca Saxe, and Joshua Tenenbaum. Bayesian models of human action understanding. *Advances in neural information processing systems*, 18:99–106, 2005.
- [Baker *et al.*, 2007] Chris L Baker, Joshua B Tenenbaum, and Rebecca R Saxe. Goal inference as inverse planning. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, volume 29, 2007.
- [Baker *et al.*, 2009] Chris L Baker, Rebecca Saxe, and Joshua B Tenenbaum. Action understanding as inverse planning. *Cognition*, 113(3):329–349, 2009.
- [Baker *et al.*, 2011] Chris Baker, Rebecca Saxe, and Joshua Tenenbaum. Bayesian theory of mind: Modeling joint belief-desire attribution. In *Proceedings of the annual meeting of the cognitive science society*, volume 33, 2011.
- [Borrajo *et al.*, 2020] Daniel Borrajo, Sriram Gopalakrishnan, and Vamsi K Potluru. Goal recognition via model-based and model-free techniques. *arXiv preprint arXiv:2011.01832*, 2020.
- [Carberry, 1996] Sandra Carberry. Plan recognition: achievements, problems, and prospects. *A A*, 5:6, 1996.
- [Carberry, 2001] Sandra Carberry. Techniques for plan recognition. *User modeling and user-adapted interaction*, 11(1-2):31–48, 2001.
- [Chakraborti *et al.*, 2019] Tathagata Chakraborti, Anagha Kulkarri, Sarath Sreedharan, David E Smith, and Subbarao Kambhampati. Explicability? legibility? predictability? transparency? privacy? security? the emerging landscape of interpretable agent behavior. In *Proceedings of the international conference on automated planning and scheduling*, volume 29, pages 86–96, 2019.
- [Charniak and Goldman, 1993] Eugene Charniak and Robert P Goldman. A bayesian model of plan recognition. *Artificial Intelligence*, 64(1):53–79, 1993.
- [Demolombe and Hamon, 2002] Robert Demolombe and Erwan Hamon. What does it mean that an agent is performing a typical procedure? A formal definition in the situation calculus. pages 905–911, 2002.
- [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 *Proceedings of the 24th International Conference on Artificial Intelligence*, pages 761–768, 2015.
- [Kaminka *et al.*, 2018] Gal A Kaminka, Mor Vered, and Noa Agmon. Plan recognition in continuous domains. In *AAAI*, pages 6202–6210, 2018.
- [Kautz and Allen, 1986] Henry A. Kautz and James F. Allen. Generalized plan recognition. In *AAAI*, pages 32–37, 1986.
- [Kay and King, 2020] John Anderson Kay and Mervyn A King. *Radical uncertainty*. Bridge Street Press, 2020.
- [Keren *et al.*, 2014] Sarah Keren, Avigdor Gal, and Erez Karpas. Goal recognition design. In *Twenty-Fourth International Conference on Automated Planning and Scheduling*, 2014.
- [Keren *et al.*, 2015] Sarah Keren, Avigdor Gal, and Erez Karpas. Goal recognition design for non-optimal agents. In *AAAI*, pages 3298–3304, 2015.

- [Keren *et al.*, 2016] Sarah Keren, Avigdor Gal, and Erez Karpas. Goal recognition design with non-observable actions. In *Thirtieth AAAI Conference on Artificial Intelligence*, pages 3152–3158, 2016.
- [Keren *et al.*, 2019] Sarah Keren, Avigdor Gal, and Erez Karpas. Goal recognition design in deterministic environments. *Journal of Artificial Intelligence Research*, 65:209–269, 2019.
- [Kulkarni *et al.*, 2018] Anagha Kulkarni, Matthew Klenk, Shantanu Rane, and Hamed Soroush. Resource bounded secure goal obfuscation. In *AAAI Fall Symposium on Integrating Planning, Diagnosis and Causal Reasoning*, 2018.
- [Kulkarni *et al.*, 2019] Anagha Kulkarni, Siddharth Srivastava, and Subbarao Kambhampati. A unified framework for planning in adversarial and cooperative environments. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 2479–2487, 2019.
- [Masters and Sardina, 2017a] Peta Masters and Sebastian Sardina. Cost-based goal recognition for path-planning. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*, pages 750–758, 2017.
- [Masters and Sardina, 2017b] Peta Masters and Sebastian Sardina. Deceptive path-planning. In *IJCAI*, pages 4368–4375, 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 *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*, pages 440–448, 2019.
- [Masters and Sardina, 2021] Peta Masters and Sebastian Sardina. Expecting the unexpected: Goal recognition for rational and irrational agents. *Artificial Intelligence*, 297:103490, 2021.
- [Masters *et al.*, 2021] Peta Masters, Michael Kirley, and Wally Smith. Extended goal recognition: A planning-based model for strategic deception. In *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems*, pages 871–879, 2021.
- [Mountrakis *et al.*, 2011] Giorgos Mountrakis, Jung-ho Im, and Caesar Ogole. Support vector machines in remote sensing: A review. *ISPRS Journal of Photogrammetry and Remote Sensing*, 66(3):247–259, 2011.
- [Pattison and Long, 2010] David Pattison and Derek Long. Domain independent goal recognition. In *Frontiers in Artificial Intelligence and Applications*, volume 222, pages 238–250, 2010.
- [Pattison and Long, 2011] David Pattison and Derek Long. Accurately determining intermediate and terminal plan states using bayesian goal recognition. ICAPS, 2011.
- [Pereira *et al.*, 2017] Ramon Pereira, Nir Oren, and Felipe Meneguzzi. Landmark-based heuristics for goal recognition. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 31, 2017.
- [Pereira *et al.*, 2019a] Ramon Fraga Pereira, André Grahl Pereira, and Felipe Meneguzzi. Landmark-enhanced heuristics for goal recognition in incomplete domain models. In *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 29, pages 329–337, 2019.
- [Pereira *et al.*, 2019b] Ramon Fraga Pereira, Mor Vered, Felipe Rech Meneguzzi, and Miquel Ramírez. Online probabilistic goal recognition over nominal models. In *IJCAI*, 2019.
- [Pozanco *et al.*, 2018] Alberto Pozanco, E Yolanda, Susana Fernández, and Daniel Borrajo. Counterplanning using goal recognition and landmarks. In *IJCAI*, pages 4808–4814, 2018.
- [Ramirez and Geffner, 2009] Miquel Ramirez and Hector Geffner. Plan recognition as planning. In *IJCAI*, pages 1778–1783, 2009.
- [Ramirez and Geffner, 2010] Miquel Ramirez and Hector Geffner. Probabilistic plan recognition using off-the-shelf classical planners. In *AAAI*, pages 1121–1126. Citeseer, 2010.
- [Ramirez and Geffner, 2011] Miquel Ramirez and Hector Geffner. Goal recognition over pomdps: Inferring the intention of a POMDP agent. In *IJCAI*, pages 2009–2014. Citeseer, 2011.
- [Russell and Norvig, 2010] Stuart Russell and Peter Norvig. *Artificial intelligence: a modern approach*. Pearson Education, Inc., 2010.
- [Shvo and McIlraith, 2020] Maayan Shvo and Sheila A McIlraith. Active goal recognition. In *AAAI*, pages 9957–9966, 2020.
- [Shvo *et al.*, 2018] Maayan Shvo, Shirin Sohrabi, and Sheila A McIlraith. An ai planning-based approach to the multi-agent plan recognition problem. In *Canadian Conference on Artificial Intelligence*, pages 253–258. Springer, 2018.
- [Sohrabi *et al.*, 2016] Shirin Sohrabi, Anton V Riabov, and Octavian Udrea. Plan recognition as planning revisited. In *IJCAI*, pages 3258–3264, 2016.
- [Sukthankar *et al.*, 2014] Gita Sukthankar, Christopher Geib, Hung Hai Bui, David Pynadath, and Robert P Goldman. *Plan, activity, and intent recognition: Theory and practice*. Newnes, 2014.
- [Vered and Kaminka, 2017] Mor Vered and Gal A Kaminka. Heuristic online goal recognition in continuous domains. *arXiv preprint arXiv:1709.09839*, 2017.
- [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*, 2016.
- [Vered *et al.*, 2018a] Mor Vered, Ramon Fraga Pereira, Gal Kaminka, and Felipe Rech Meneguzzi. Online goal recognition as reasoning over landmarks. In *Proceedings of the 2018 AAAI Workshop on Plan, Activity, and Intent Recognition (PAIR), 2018, Estados Unidos.*, 2018.
- [Vered *et al.*, 2018b] Mor Vered, Ramon Fraga Pereira, Gal Kaminka, and Felipe Rech Meneguzzi. Towards online goal recognition combining goal mirroring and landmarks. In *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems, 2018, Suécia.*, 2018.
- [Wayllace *et al.*, 2016] Christabel Wayllace, Ping Hou, William Yeoh, and Tran Cao Son. Goal recognition design with stochastic agent action outcomes. In *IJCAI*, 2016.
- [Wayllace *et al.*, 2017] Christabel Wayllace, Ping Hou, and William Yeoh. New metrics and algorithms for stochastic goal recognition design problems. In *IJCAI*, pages 4455–4462, 2017.
- [Zhang *et al.*, 2017] Yu Zhang, Sarath Sreedharan, Anagha Kulkarni, Tathagata Chakraborti, Hankz Hankui Zhuo, and Subbarao Kambhampati. Plan explicability and predictability for robot task planning. In *2017 IEEE international conference on robotics and automation (ICRA)*, pages 1313–1320. IEEE, 2017.
- [Zhuo *et al.*, 2012] Hankzhankui Zhuo, Qiang Yang, and Subbarao Kambhampati. Action-model based multi-agent plan recognition. *Advances in Neural Information Processing Systems 25*, 1:368, 2012.