Heuristic Goal Recognition in Incomplete Domain Models

Ramon Fraga Pereira

Pontifical Catholic University of Rio Grande do Sul (PUCRS)
School of Technology
Porto Alegre, Brazil

ramon.pereira@edu.pucrs.br

Abstract

Recent approaches to goal recognition have progressively relaxed the assumptions about the amount and correctness of domain knowledge and available observations, yielding accurate and efficient algorithms. These approaches, however, assume completeness and correctness of the domain theory against which their algorithms match the sequence of observations: this is too strong for most real-world domains. In this dissertation, we develop heuristic goal recognition techniques that are capable of recognizing goals using *incomplete* (and possibly incorrect) domain theories using new notions of planning landmarks.

1 Introduction

Goal recognition is the problem of recognizing the correct goal intended by an observed agent, given a sequence of observations as evidence of its behavior in an environment, and a domain model describing how the observed agent generates such behavior. Approaches to solve this problem vary on the amount and type of domain knowledge used in the agents' behavior (or plan generation). However, all recent planning-based approaches to goal and plan recognition assume that the domain model is complete and correct (Ramírez and Geffner 2009; 2010; Keren, Gal, and Karpas 2014; E.-Martín, R.-Moreno, and Smith 2015; Sohrabi, Riabov, and Udrea 2016; Pereira and Meneguzzi 2016; Pereira, Oren, and Meneguzzi 2017), preventing its application to realistic scenarios in which the domain modeler either has an incomplete or incorrect model of the agents' behavior under observation. Specifically, real-world domains often have two potential sources of uncertainty: (1) ambiguity in domain engineering either because of a noisy domain acquisition process or the nature of the actions being modeled; and (2) ambiguity from how imperfect sensor data reports features of the environment. The former stems from a possibly incomplete understanding of the actions being modeled, but more importantly, the inherently noisy and imperfect way in which automated domain acquisition through machine learning algorithms (Asai and Fukunaga 2018; Amado et al. 2018) we envision being the main source of real-world domain models. The latter stems from the potential unreliability in the interpretation of actions using real-world noisy data with learned sensor models being used to classify objects to be used as features (*e.g.*, logical facts) of the observations (Granada et al. 2017), so it is useful to model a domain with such feature as optional.

In this dissertation, we develop heuristic goal recognition approaches that can cope with incomplete planning domain models (Nguyen, Sreedharan, and Kambhampati 2017), and provide four main contributions. First, we formalize goal recognition in incomplete domains by combining the standard formalization of Ramírez and Geffner (2009; 2010) for plan recognition and that of Nguyen, Sreedharan, and Kambhampati (2017). Second, we develop an algorithm, adapted from (Hoffmann, Porteous, and Sebastia 2004), that extracts *definite* and *possible* landmarks in incomplete domain models. Third, we develop a notion of *overlooked* landmarks that we can extract online as we process (*on the fly*) observations. Fourth, we develop and enhance two heuristic approaches (Pereira, Oren, and Meneguzzi 2017) to recognize goals that account for the various types of landmark as evidence.

2 Goal Recognition in Incomplete Domains

Our heuristic approaches assume that the recognizer (observer) has an *incomplete* domain model while the observed agent is planning and acting in the environment with a *complete* domain model (Figure 1). To account for such domain uncertainty and incompleteness, the domain model available to the recognizer contains *possible preconditions* and *effects*, much like the incomplete domain models from previous planning approaches (Weber and Bryce 2011; Nguyen, Sreedharan, and Kambhampati 2017). We formalize the goal recognition problem over incomplete domain models in Definition 1.

Definition 1 (Goal Recognition Problem) *A goal recognition problem with an incomplete domain model is a quintuple $\widetilde{\mathcal{T}} = \langle \widetilde{\mathcal{D}}, Z, \mathcal{I}, \mathcal{G}, Obs \rangle$, where: $\widetilde{\mathcal{D}} = \langle \mathcal{R}, \widetilde{\mathcal{O}} \rangle$ is an incomplete domain model (with possible preconditions and effects); Z is the set of typed objects in the environment, in which $\mathcal{F}$ is the set of instantiated predicates from Z, and $\widetilde{\mathcal{A}}$ is the set of incomplete instantiated actions from $\widetilde{\mathcal{O}}$ with objects from Z; $\mathcal{I} \in \mathcal{F}$ an initial state; $\mathcal{G}$ is the set of possible goals, which include a correct hidden goal G^* (i.e., $G^* \in$*

$\mathcal{G}$); and $Obs = \langle o_1, o_2, ..., o_n \rangle$ *is an observation sequence of executed actions, with each observation* $o_i \in \widetilde{\mathcal{A}}$. *Obs corresponds to the sequence of actions (i.e., a plan) to solve a problem in a complete domain in* $\langle\langle\widetilde{\mathcal{D}}\rangle\rangle$.

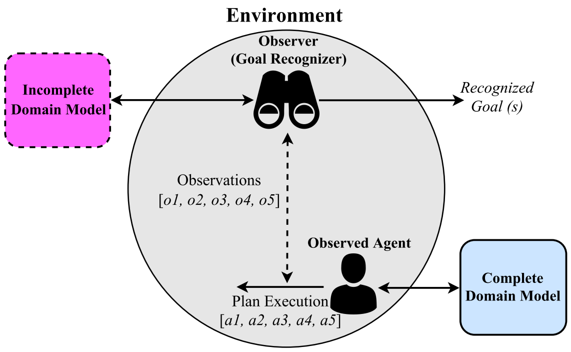

Figure 1: Goal Recognition in incomplete domain models.

A solution for a goal recognition problem in incomplete domain models $\widetilde{\mathcal{T}}$ is the correct hidden goal $G^* \in \mathcal{G}$ that the observation sequence Obs of a plan execution achieves, specifically, the correct hidden goal G^* is the intended goal that the observed agent wants to achieve. As most keyhole goal recognition approaches, observations consist of the actions of the underlying plan, *i.e.*, we observe incomplete actions with possible precondition and effects, in which some of the preconditions might be required and some effects might change the environment. A full (or complete) observation sequence contains all of the action signatures of the plan executed by the observed agent, while a partial observation sequence contains only a sub-sequence of actions of a plan and thus misses some of the actions actually executed in the environment. We note that our approaches are not limited to use just actions as observations and can also deal with logical facts as observations, *i.e.*, state observations, like (Sohrabi, Riabov, and Udrea 2016).

2.1 Landmark Extraction in Incomplete Domains

Planning landmarks are facts (or actions) that must be achieved (or executed) at some point along all valid plans to achieve a goal from an initial state (Hoffmann, Porteous, and Sebastia 2004). In Automated Planning, landmarks are often used to build heuristics for planning algorithms using complete and correct domain models. Here, we adapt the extraction algorithm from (Hoffmann, Porteous, and Sebastia 2004) to extract landmarks from incomplete domains by building an Optimistic Relaxed Planning Graph (ORPG) instead of the original Relaxed Planning Graph (RPG) (Hoffmann and Nebel 2001). An ORPG is a leveled graph that deals with incomplete domain models by assuming the most *optimistic* conditions. Thus, besides ignoring the delete-effects of all actions, this graph also ignores possible preconditions and possible delete-effects, considering that all possible add effects occur. Replacing an RPG for an ORPG allows us to extract *definite* and *possible* landmarks. We note that in this work we use the same landmark extraction process developed by Hoffman *et al.* (2004), the only adaptation is using of an ORPG instead the original RPG. We formally define the concept of *definite* and *possible* landmarks in Definitions 2 and 3, respectively.

Definition 2 (Definite Landmark) *A definite landmark* L_D *is a fact (landmark) that is extracted from a known add effect* $\mathit{eff}^+(a)$ *of an achiever*[1] a *(action) in the ORPG.*

Definition 3 (Possible Landmark) *A possible landmark* L_P *is a fact (landmark) that is extracted from a possible add effect* $\widetilde{\mathit{eff}}^+(a)$ *of an achiever* a *(action) in the ORPG and is such that* $L_P \cap L_D = \emptyset$.

2.2 Heuristic Goal Recognition Approaches

Key to our heuristic approaches is observing the evidence of achieved landmarks during observations to recognize which goal is more consistent with the observations. To do so, our approaches combine the concepts of *definite* and *possible* with that of *overlooked* landmarks (Definition 4).

Definition 4 (Overlooked Landmark) *An overlooked landmark* L_O *is an actual landmark, a necessary fact for all valid plans towards a goal from an initial state, that was not detected by approximate landmark extraction algorithms.*

Most landmark extraction algorithms extract only a subset of landmarks for a given planning problem, and to overcome this problem, we aim to extract *overlooked* landmarks by analyzing preconditions and effects in the observed actions of an observation sequence. Since we are dealing with incomplete domain models, and it is possible that they have few (or no) *definite* and/or *possible* landmarks, we extract *overlooked* landmarks from the evidence in the observations as we process them in order to enhance the set of landmarks useable by our heuristic. To extract such landmarks, we build a new ORPG removing all action achievers that achieve a potentially *overlooked* fact landmark and checks the solvability of this modified problem. If the modified problem is indeed unsolvable, then this fact is an *overlooked* landmark.

2.3 Enhanced Goal Completion Heuristic

We now combine our new notions of landmarks to develop a goal recognition heuristic for recognizing goals in incomplete domain models. Our heuristic estimates the correct goal in the set of candidate goals by calculating the ratio between the amount of achieved *definite* ($\mathcal{AL}_G$), *possible* ($\widetilde{\mathcal{AL}}_G$), and *overlooked* ($\mathcal{ANL}_G$) landmarks and the amount of *definite* ($\mathcal{L}_G$), *possible* ($\widetilde{\mathcal{L}}_G$), and *overlooked* ($\mathcal{NL}_G$) landmarks. The estimate computed using Equation 1 represents the percentage of achieved landmarks for a candidate goal from observations.

$$h_{\widetilde{GC}}(G) = \left(\frac{\mathcal{AL}_G + \widetilde{\mathcal{AL}}_G + \mathcal{ANL}_G}{\mathcal{L}_G + \widetilde{\mathcal{L}}_G + \mathcal{NL}_G} \right) \quad (1)$$

[1] An achiever is an action at the level before a candidate landmark in the ORPG (or RPG) that can be used to achieve this candidate landmark.

2.4 Enhanced Uniqueness Heuristic

Most goal recognition problems contain multiple candidate goals that share common fact landmarks, generating ambiguity that jeopardizes the goal completion heuristic. Clearly, landmarks that are common to multiple candidate goals are less useful for recognizing a goal than landmarks that exist for only a single goal. As a consequence, computing how unique (and thus informative) each landmark is can help disambiguate similar goals for a set of candidate goals. Our uniqueness heuristic is based on this intuition, using the concept of *landmark uniqueness*, which is the inverse frequency of a landmark among the landmarks found in a set of candidate goals. Intuitively, a landmark L that occurs only for a single goal within a set of candidate goals has the maximum uniqueness value of 1. We calculate the *landmark uniqueness value* for a landmark L and a set of landmarks for all candidate goals $K_{\mathcal{G}}$ using the following equation:

$$L_{Uniq}(L, K_{\mathcal{G}}) = \left(\frac{1}{\sum\limits_{\mathcal{L} \in K_{\mathcal{G}}} |\{L | L \in \mathcal{L}\}|} \right) \tag{2}$$

Using the concept of *landmark uniqueness value*, we estimate which candidate goal is the intended one by summing the uniqueness values of the landmarks achieved in the observations. Unlike our previous heuristic, which estimates progress towards goal completion by analyzing just the set of achieved landmarks, the landmark-based uniqueness heuristic estimates the goal completion of a candidate goal G by calculating the ratio between the sum of the uniqueness value of the achieved landmarks of G and the sum of the uniqueness value of all landmarks of a goal G. Our new uniqueness heuristic also uses the concepts of *definite*, *possible*, and *overlooked* landmarks. We store the set of *definite* and *possible* landmarks of a goal G separately into $\mathcal{L}_G$ and $\widetilde{\mathcal{L}}_G$, and the set of *overlooked* landmarks into $\mathcal{NL}_G$. Thus, the uniqueness heuristic effectively weighs the completion value of a goal by the informational value of a landmark so that unique landmarks have the highest weight. To estimate goal completion using the *landmark uniqueness value*, we calculate the uniqueness value for every extracted (*definite*, *possible*, and *overlooked*) landmark in the set of landmarks of the candidate goals using the equation we mentioned before. Since we use three types of landmarks and they are stored in three different sets, we compute the *landmark uniqueness value* separately for them, storing the landmark uniqueness value of *definite* landmarks $\mathcal{L}_G$ into $\Upsilon_{\mathcal{L}}$, the landmark uniqueness value of *possible* landmarks $\widetilde{\mathcal{L}}_G$ into $\Upsilon_{\widetilde{\mathcal{L}}}$, and the landmark uniqueness value of *overlooked* landmarks $\mathcal{NL}_G$ into $\Upsilon_{\mathcal{NL}_G}$. Our uniqueness heuristic $h_{\widetilde{UNIQ}}$ is computed and formally defined in Equation 3.

$$h_{\widetilde{UNIQ}}(G) = \left(\frac{\sum\limits_{\mathcal{A}_L \in \mathcal{AL}_G} \Upsilon_{\mathcal{L}}(\mathcal{A}_L) + \sum\limits_{\widetilde{\mathcal{A}_L} \in \widetilde{\mathcal{AL}}_G} \Upsilon_{\widetilde{\mathcal{L}}}(\widetilde{\mathcal{A}_L}) + \sum\limits_{\mathcal{ANL} \in \mathcal{ANL}_G} \Upsilon_{\mathcal{NL}_G}(\mathcal{ANL})}{\sum\limits_{L \in \mathcal{L}_G} \Upsilon_{\mathcal{L}}(L) + \sum\limits_{\tilde{L} \in \widetilde{\mathcal{L}}_G} \Upsilon_{\widetilde{\mathcal{L}}}(\tilde{L}) + \sum\limits_{\mathcal{NL} \in \mathcal{NL}_G} \Upsilon_{\mathcal{NL}_G}(\mathcal{NL})} \right) \tag{3}$$

3 Conclusions

We have developed novel goal recognition approaches that deal with incomplete domain models that represent *possible* preconditions and effects besides traditional models where such information is assumed to be *known*. The main contributions of this work include the formalization of goal recognition in incomplete domains, two (enhanced) heuristic approaches for such goal recognition, and novel notions of landmarks for incomplete domain models. Recent approaches differ from ours in that they only deal with complete and correct (even if modified) domain models, and most of them transform or compile the recognition problems into planning problems for a classical planner.

Acknowledgements: Ramon thanks his advisor Felipe Meneguzzi for his mentoring over past three years. Ramon also thanks Miquel Ramírez and André Grahl Pereira for the invaluable discussions about this work over the past year.

References

Amado, L.; Pereira, R. F.; Aires, J. P.; Magnaguagno, M.; Granada, R.; and Meneguzzi, F. 2018. Goal recognition in latent space. In *IJCNN*.

Asai, M., and Fukunaga, A. 2018. Classical planning in deep latent space: Bridging the subsymbolic-symbolic boundary. In *AAAI*.

E.-Martín, Y.; R.-Moreno, M. D.; and Smith, D. E. 2015. A Fast Goal Recognition Technique Based on Interaction Estimates. In *IJCAI*.

Granada, R.; Pereira, R. F.; Monteiro, J.; Barros, R.; Ruiz, D.; and Meneguzzi, F. 2017. Hybrid activity and plan recognition for video streams. In *The AAAI 2017 Workshop on Plan, Activity, and Intent Recognition*.

Hoffmann, J., and Nebel, B. 2001. The FF Planning System: Fast Plan Generation Through Heuristic Search. *Journal of Artificial Intelligence Research (JAIR)*.

Hoffmann, J.; Porteous, J.; and Sebastia, L. 2004. Ordered Landmarks in Planning. *Journal of Artificial Intelligence Research (JAIR)*.

Keren, S.; Gal, A.; and Karpas, E. 2014. Goal Recognition Design. In *ICAPS*.

Nguyen, T.; Sreedharan, S.; and Kambhampati, S. 2017. Robust Planning with Incomplete Domain Models. *Artificial Intelligence*.

Pereira, R. F., and Meneguzzi, F. 2016. Landmark-Based Plan Recognition. In *ECAI*.

Pereira, R. F.; Oren, N.; and Meneguzzi, F. 2017. Landmark-Based Heuristics for Goal Recognition. In *AAAI*.

Ramírez, M., and Geffner, H. 2009. Plan Recognition as Planning. In *IJCAI*.

Ramírez, M., and Geffner, H. 2010. Probabilistic Plan Recognition using off-the-shelf Classical Planners. In *AAAI*.

Sohrabi, S.; Riabov, A. V.; and Udrea, O. 2016. Plan Recognition as Planning Revisited. In *IJCAI*.

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