The Thirty-Second AAAI Conference
on Artificial Intelligence (AAAI-18)

Goal Recognition in Incomplete Domain Models

Ramon Fraga Pereira, Felipe Meneguzzi
Pontifical Catholic University of Rio Grande do Sul (PUCRS), Brazil
Postgraduate Programme in Computer Science, School of Informatics (FACIN)
Porto Alegre, RS, Av. Ipiranga 668, CEP: 90619-900
ramon.pereira@acad.pucrs.br
felipe.meneguzzi@pucrs.br
+55 51 3553 8622

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 observations: this is too strong for most real-world domains. In this work, we develop a goal recognition technique capable of recognizing goals using *incomplete* (and possibly incorrect) domain theories.

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 assume that the domain model is correct and complete (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 behavior under observation. Specifically, real world domains have two potential sources of uncertainty: (1) ambiguity in how actions performed by agents are realized; and (2) ambiguity from how imperfect sensor data reports features of the world. The former stems from an incomplete understanding of the action being modeled and requires a domain modeler to specify a number of alternate versions of the same action to cover the possibilities. For example, an action to turn on the gas burner in a cooker may or may not require the observed agent to press a spark button. The latter stems from imperfections in the way actions themselves may be interpreted from real-world noisy data, *e.g.*, if one uses machine learning algorithms to classify objects to be used as features (*e.g.*, logical facts) of the observations, certain features may not be recognizable reliably, so it is useful to model them as optional.

We develop a goal recognition approach that can cope with incomplete planning domain models (Nguyen, Sreedharan, and Kambhampati 2017). This paper has 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 *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 a heuristic to recognize goals that accounts for the various types of landmark as evidence.

2 Goal Recognition in Incomplete Domains

Our goal recognition approach assumes that the observer has an incomplete domain model while the observed agent is planning and acting with a complete domain model. To account for such uncertainty, the model available to the observer 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 5-tuple $\widetilde{\mathcal{T}} = \langle \widetilde{\mathcal{D}}, Z, \mathcal{I}, \mathcal{G}, O \rangle$, where: $\widetilde{\mathcal{D}} = \langle \mathcal{R}, \widetilde{\mathcal{O}} \rangle$ is an incomplete domain model, in which $\mathcal{R}$ is a set of predicates with typed variables, and $\widetilde{\mathcal{O}}$ is a set of incomplete operators with possible preconditions and effects; Z is the set of typed objects, in which $\mathcal{F}$ is the set instantiated predicates from $\mathcal{R}$ with objects from Z, and $\widetilde{\mathcal{A}}$ is the set of incomplete instantiated actions from $\widetilde{\mathcal{O}}$ with objects from Z; $\mathcal{I} \in \mathcal{F}$ is the initial state; $\mathcal{G}$ is the set of possible goals, including the correct hidden goal G ($G \in \mathcal{G}$); and $O = \langle o_1, o_2, ..., o_n \rangle$ is an observation sequence of executed actions, with each observation $o_i \in \widetilde{\mathcal{A}}$ (incomplete action with possible preconditions and effects). O is a plan that achieves the correct hidden goal G in a complete domain in $\langle\langle \widetilde{\mathcal{D}} \rangle\rangle$.*

A solution for such a goal recognition problem $\widetilde{\mathcal{T}}$ is the correct hidden goal $G \in \mathcal{G}$ that the observation sequence O of a plan execution achieves.

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). Landmarks are often used to build heuristics for planning algorithms using complete and correct domain models. 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. A *definite* landmark $L_{Definite}$ is a fact landmark extracted from a known add effect of an achiever action in the ORPG. A *possible* landmark $L_{Possible}$ is a fact landmark extracted from a possible add effect of an achiever action in the ORPG.

2.2 Heuristic Goal Recognition

Key to our goal recognition approach is observing the evidence of achieved landmarks during observations to recognize which goal is more consistent with the observations. To do so, we combine the concepts of *definite* and *possible* with that of *overlooked* landmarks. An *overlooked* landmark is an actual landmark, *i.e.*, a necessary fact for all valid plans towards a goal, that was not detected by approximate landmark extraction algorithms. Since we are dealing with incomplete domain models, and so 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. This *on the fly* landmark extraction checks if the facts in the known preconditions and known and possible add effects are not *definite* and *possible* landmarks, and if they are not, we check if these facts are *overlooked* landmarks. To do so, we develop a function called IS-LANDMARK that builds a new ORPG removing actions that achieve a fact (potential *overlooked* landmark) and checks the solvability over this modified problem. If the modified problem is unsolvable, then this fact is an *overlooked* landmark. We check every candidate goal using this function to extract additional landmarks.

Combining the concept of *definite*, *possible*, and *overlooked* landmarks, we develop an 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 achieved *definite* ($\mathcal{AL}_G$), *possible* ($\widetilde{\mathcal{AL}}_G$), *overlooked* ($\mathcal{ANL}_G$) landmarks and the amount of *definite* ($\mathcal{L}_G$), *possible* ($\widetilde{\mathcal{L}}_G$), and *overlooked* ($\mathcal{NL}_G$) landmarks. This estimate (heuristic) is formalized in Equation 1.

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

We evaluated our approach using a new dataset constructed by modifying an existing dataset with 15 domains and hundreds of problems, removing just information from the complete domain model and annotating it with possible preconditions and effects (incomplete domain model). Using these modified datasets (varying the percentage of incompleteness of each domain from 20% to 80%), our results show that this heuristic approach is fast and accurate for recognizing goals at most percentages of domain incompleteness for all domain models.

3 Conclusions

We have developed a novel goal recognition approach that deals with incomplete domains that have *possible*, rather than *known*, preconditions and effects. The main contributions of this work are: a landmark extraction algorithm that deals with incomplete domains; and a recognition heuristic for incomplete domains that relies on landmarks. Recent goal and plan recognition approaches differ from ours because they only deal with complete (even if modified) domain models, and most of them transform/compile the goal/plan recognition problem into a planning problem to take as input for a planner. Such transformation/compilation process may not necessarily work with incomplete domains, given the potentially very large number of potential models.

References

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

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*.