Goal Recognition as Reinforcement Learning

Leonardo Rosa Amado, Reuth Mirsky, Felipe Meneguzzi

Key takeaway: we're replacing the expert model in goal recognition with reinforcement learning

https://qrco.de/bcj1q0

Background

In this paper, we develop a framework that combines model-free reinforcement learning and goal recognition to alleviate the need for careful, manual domain design, and the need for costly online executions..

Goal Recognition as Reinforcement Learning

This approach consists of two main stages:

* **Stage 1** - For each $g \in G$, **learn** Q_g that represents the desired behavior to accomplish g
* **Stage 2** - **Infer** the goal of the actor by computing Given an observation sequence O=<s_0,a_0,s_1,a_1,...> find:

$$g^* = \operatorname*{argmin}_{g \in G} Distance(Q_g, O)$$

Goal Recognition as Q-Learning (GRAQL)

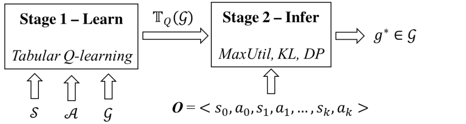

Learn using off-the-shelf Q-learning algorithms to get the Q-function, Q_g, for each $g \in G$.
Infer the goal of the actor by using a distance measure to compare each Q_g to the observation sequence. In this paper we discuss 3 measures: MaxUtil, KL-divergence, and Divergence Point (DP).

Empirical Results

We compared our recognizers to R&G [1] on 3 domains: Blocks, Hanoi, SkGrid.

More results with partial observability and noise are in the paper!

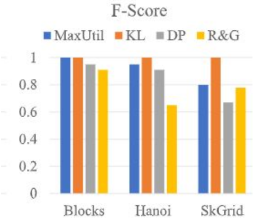

Figure: F-score with full observability and no noise, KL measure outperforms all other approaches.

Traditional Goal Recognition [1]

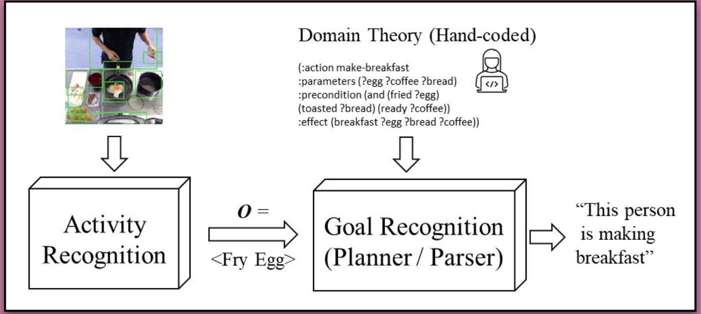

Our approach: Goal Recognition as Reinforcement Learning

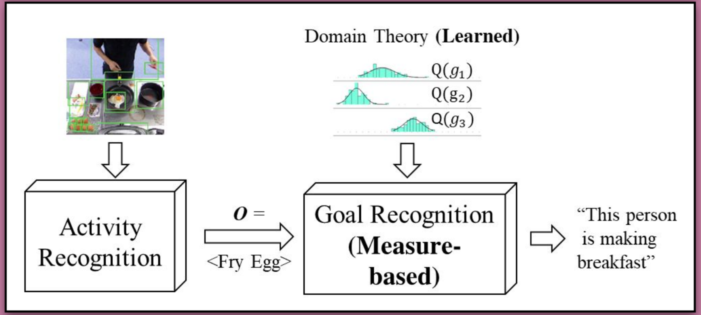

[1] Ramírez and Geffner. "Probabilistic plan recognition using off-the-shelf classical planners." AAAI. 2010.

Goal Recognition problem