The Thirty-Ninth AAAI Conference on Artificial Intelligence (AAAI-25)

Learning Logic Specifications for Policy Guidance in POMDPs: an Inductive Logic Programming Approach

Daniele Meli, Alberto Castellini, Alessandro Farinelli

University of Verona

Abstract Reprint. This is an abstract reprint of a journal article by Meli, Castellini, and Farinelli (2024).

Abstract

Partially Observable Markov Decision Processes (POMDPs) are a powerful framework for planning under uncertainty. They allow to model state uncertainty as a belief probability distribution. Approximate solvers based on Monte Carlo sampling show great success to relax the computational demand and perform online planning. However, scaling to complex realistic domains with many actions and long planning horizons is still a major challenge, and a key point to achieve good performance is guiding the action-selection process with domain-dependent policy heuristics which are tailored for the specific application domain. We propose to learn high-quality heuristics from POMDP traces of executions generated by any solver. We convert the belief-action pairs to a logical semantics, and exploit data- and time-efficient Inductive Logic Programming (ILP) to generate interpretable belief-based policy specifications, which are then used as online heuristics. We evaluate thoroughly our methodology on two notoriously challenging POMDP problems, involving large action spaces and long planning horizons, namely, rocksample and pocman. Considering different state-of-the-art online POMDP solvers, including POMCP, DESPOT and AdaOPS, we show that learned heuristics expressed in Answer Set Programming (ASP) yield performance superior to neural networks and similar to optimal handcrafted task-specific heuristics within lower computational time. Moreover, they well generalize to more challenging scenarios not experienced in the training phase (e.g., increasing rocks and grid size in rocksample, incrementing the size of the map and the aggressivity of ghosts in pocman).

References

Meli, D.; Castellini, A.; and Farinelli, A. 2024. Learning logic specifications for policy guidance in pomdps: an inductive logic programming approach. *Journal of Artificial Intelligence Research*, 79: 725–776.