Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence (IJCAI-24)
Journal Track

Optimality Guarantees for Particle Belief Approximation of POMDPs (Abstract Reprint)

Michael Lim[1] , **Tyler Becker**[2] , **Mykel Kochenderfer**[3] , **Claire Tomlin**[1] and **Zachary Sunberg**[2]
[1]University of California, Berkeley,Electrical Engineering and Computer Sciences Department
[2]University of Colorado Boulder,Aerospace Engineering Sciences Department
[3]Stanford University,Aeronautics and Astronautics Department
michaelhlim@berkeley.edu, tyler.becker-1@colorado.edu, mikel@stanford.edu,
tomlin@EECS.berkeley.edu, zachary.sunberg@colorado.edu

Abstract Reprint. This is an abstract reprint of a journal by [Lim *et al.*, 2023].

Abstract

Partially observable Markov decision processes (POMDPs) provide a flexible representation for real-world decision and control problems. However, POMDPs are notoriously difficult to solve, especially when the state and observation spaces are continuous or hybrid, which is often the case for physical systems. While recent online sampling-based POMDP algorithms that plan with observation likelihood weighting have shown practical effectiveness, a general theory characterizing the approximation error of the particle filtering techniques that these algorithms use has not previously been proposed. Our main contribution is bounding the error between any POMDP and its corresponding finite sample particle belief MDP (PB-MDP) approximation. This fundamental bridge between PB-MDPs and POMDPs allows us to adapt any sampling-based MDP algorithm to a POMDP by solving the corresponding particle belief MDP, thereby extending the convergence guarantees of the MDP algorithm to the POMDP. Practically, this is implemented by using the particle filter belief transition model as the generative model for the MDP solver. While this requires access to the observation density model from the POMDP, it only increases the transition sampling complexity of the MDP solver by a factor of O(C), where C is the number of particles. Thus, when combined with sparse sampling MDP algorithms, this approach can yield algorithms for POMDPs that have no direct theoretical dependence on the size of the state and observation spaces. In addition to our theoretical contribution, we perform five numerical experiments on benchmark POMDPs to demonstrate that a simple MDP algorithm adapted using PB-MDP approximation, Sparse-PFT, achieves performance competitive with other leading continuous observation POMDP solvers.

References

[Lim *et al.*, 2023] Michael H. Lim, Tyler J. Becker, Mykel J. Kochenderfer, Claire J. Tomlin, and Zachary N. Sunberg. Optimality guarantees for particle belief approximation of pomdps. *J. Artif. Intell. Res.*, 77:1591–1636, 2023.