

Partially Observable Monte-Carlo Graph Search

Yang You¹, Vincent Thomas³, Alex Schutz², Robert Skilton¹, Nick Hawes², Olivier Buffet³

¹ United Kingdom Atomic Energy Authority

² Oxford Robotics Institute, University of Oxford

³ Université de Lorraine, CNRS, Inria, LORIA, F-54000 Nancy, France

Abstract

Currently, large partially observable Markov decision processes (POMDPs) are often solved by sampling-based online methods which interleave planning and execution phases. However, a pre-computed offline policy is more desirable in POMDP applications with time or energy constraints. But previous offline algorithms are not able to scale up to large POMDPs. In this article, we propose a new sampling-based algorithm, the *partially observable Monte-Carlo graph search* (POMCGS) to solve large POMDPs offline. Different from many online POMDP methods, which progressively develop a tree while performing (Monte-Carlo) simulations, POMCGS folds this search tree on the fly to construct a policy graph, so that computations can be drastically reduced, and users can analyze and validate the policy prior to embedding and executing it. Moreover, POMCGS, together with action progressive widening and observation clustering methods provided in this article, is able to address certain continuous POMDPs. Through experiments, we demonstrate that POMCGS can generate policies on the most challenging POMDPs, which cannot be computed by previous offline algorithms, and these policies' values are competitive compared with the state-of-the-art online POMDP algorithms.

1 Introduction

The framework of partially observable Markov decision processes (POMDPs) (Åström 1965) provides a mathematical tool for modeling sequential decision-making problems in stochastic and partially observable environments. In a POMDP, the agent needs to take decisions based on its action-observation history or, equivalently, a distribution over possible states (known as a belief). The objective is to compute a policy that maximizes a given performance criterion. However, solving a POMDP is difficult due to the limited information available about the true state. It has been proven that solving a finite-horizon POMDP is PSPACE-complete (Papadimitriou and Tsitsiklis 1987), and finding an optimal policy is undecidable for infinite-horizon POMDPs (Madani, Hanks, and Condon 1999).

The majority of state-of-the-art POMDP planners are sampling-based online algorithms (Silver and Veness 2010; Somani et al. 2013) that perform Monte-Carlo simulations in

a search tree to plan the best action given the current belief. The agent then executes this action and updates its belief according to the new observation received before proceeding to the next planning phase. This interleaves the planning and execution processes, which saves computation resources significantly and enables sampling-based online approaches to address large, and even continuous POMDP problems (Sunberg and Kochenderfer 2017; Garg, Hsu, and Lee 2019; Hoerger and Kurniawati 2021; Wu et al. 2021). However, in applications with constraints such as limited computational power, time, or energy resources, it may not be feasible to perform extensive online computations (Grzes et al. 2015). In these situations, a pre-computed policy is often preferable, as it eliminates the need for further computations during execution and can provide immediate responses. However, offline planning methods encounter scalability challenges due to the computational and memory requirements needed to solve the entire POMDP problem upfront. This limits the practical use of offline planning for larger or more complex problems, where solving the problem entirely in advance may be infeasible.

In this article, taking advantage of successful ideas from existing methods, we propose a novel sampling-based offline algorithm called *Partially Observable Monte-Carlo Graph Search* (POMCGS) with following features: 1. POMCGS performs Monte-Carlo simulations to build and update a complete policy that can be executed offline without any further computations; and 2. POMCGS addresses continuous POMDPs using state discretization, action progressive widening, and observation clustering. Experiments show that POMCGS outperforms previous offline methods in scalability and is able to derive compact policies on the most challenging POMDPs, with competitive performance even compared to state-of-the-art online planners.

This paper is organized as follows: Section 2 describes related work on solving POMDPs, and background is presented in Section 3. We present our contribution in Section 4, and give experiments with comparisons to state-of-the-art POMDP solvers in Section 5. Section 6 discusses the strengths and limitations of our algorithm, as well as opportunities for related research fields. Finally, we conclude this work in Section 7.

2 Related Work

In this section, we review work related to our contribution. Lauri, Hsu, and Pajarinen (2022) and Kurniawati (2022) present more detailed surveys about POMDPs.

Offline POMDP solvers aim to compute a policy anticipating all possible situations. In this category, *point-based* methods (Pineau, Gordon, and Thrun 2003; Smith and Simmons 2004, 2005; Kurniawati, Hsu, and Lee 2008; Bai et al. 2010) approximate the optimal value function by exploiting its convexity, and focus on representative belief points that are reachable under an optimal policy; *policy iteration* algorithms (Hansen 1998a,b; Poupart and Boutilier 2003; Ji et al. 2007) iteratively refine a policy represented by a finite-state controller (FSC) until convergence in small scale POMDP problems; while *expectation-maximization* methods (Toussaint, Harmeling, and Storkey 2006; Pajarinen and Peltonen 2011) turn a POMDP into a probabilistic inference problem and estimate its parameters, which only gives local optima.

Conversely, state-of-the-art POMDP solvers are usually online methods based on Monte-Carlo Tree Search (Kocsis and Szepesvári 2006). In this category, *partially observable Monte-Carlo planning* (POMCP) (Silver and Veness 2010) extends MCTS to large POMDPs. POMCP performs forward simulations in a search tree, and each tree node corresponds to an action-observation history. Alternatively, *determined sparse partially observable tree* (DESPOT) (Somani et al. 2013) uses a fixed number of scenarios to sample observations and maintain a sparse belief-tree, which proves to be very efficient for solving POMDPs with many observations. These online algorithms outperform the previous offline methods in terms of scalability. Recent works (Sunberg and Kochenderfer 2017; Garg, Hsu, and Lee 2019; Hoerger and Kurniawati 2021; Wu et al. 2021) extend online POMDP planning further to continuous domains. For example, Sunberg and Kochenderfer (2017) extend POMCP by using double progress widening during the tree search, *i.e.*, slowly expanding the sets of applicable actions and considered observations.

One can improve the efficiency of MCTS-based methods by exploiting the continuity of the optimal value function in state or belief space (Paludo Licks, Patrizi, and De Giacomo 2024). For instance, performance can be improved by merging similar belief nodes at the same depth along the search tree (Ballesteros et al. 2013; Wu et al. 2021; Kujanpää et al. 2023), which avoids expanding unnecessary branches. Leurent and Maillard (2020) discuss another variation called Monte-Carlo graph search in fully observable settings that merges nodes corresponding to similar states reached via different trajectories. An important insight for POMDPs is that merging similar nodes in a search tree leads to a structure similar to a finite state controller (FSC). In this article, we propose a Monte-Carlo graph search algorithm in POMDPs which, by merging similar nodes at possibly different depths, constructs an FSC on the fly.

3 Background

POMDPs

In this work, we consider one agent interacting with an environment under uncertain dynamics and partial observability. Such a problem is described as a *partially observable Markov decision process* (POMDP), which is formally defined by a tuple $\langle \mathcal{S}, \mathcal{A}, \Omega, T, O, r, b_0 \rangle$ where (i) \mathcal{S} is the set of states; (ii) \mathcal{A} is the set of actions; (iii) Ω is the set of observations; (iv) $T(s, a, s') = Pr(s'|s, a)$ is the transition function, which encodes the dynamics of the environment (the probability of reaching the next state s' given the current state s and an executed action a); (v) $O(o, a, s') = Pr(o|s', a)$ is the observation function; it indicates the probability of receiving an observation o given the next state s' reached while performing action a ; (vi) $r(s, a)$ is the reward function, which gives the instant reward when performing action a at state s ; (vii) b_0 is an initial probability distribution over states.

In POMDPs, a policy π maps action-observation histories $h_t = \langle a_1, o_1, \dots, a_t, o_t \rangle$ to actions, and its *value function* gives the expected discounted cumulative rewards under π from any reachable history: $V^\pi(h_t) \stackrel{\text{def}}{=} \sum_{k=t}^{\infty} \mathbb{E}_\pi [\gamma^{t-k} r(S_k, \pi(H_k)) | H_t = h_t]$, and the action-value is defined as $Q^\pi(h, a) \stackrel{\text{def}}{=} \mathbb{E}_\pi [r(S_t, A_t) + \gamma V^\pi(H_{t+1}) | H_t = h_t, A_t = a]$. There exists at least one *optimal policy* π^* , *i.e.*, a policy with *optimal value function* $V^*(h_t) = \max_\pi V^\pi(h_t)$ for any h_t . The probability distribution over states under history h_t (obtained through Bayesian inference) is its *belief state* b .

Explicit transition, observation and reward functions may not be available, or may be impractical due to time or memory requirements. In this case, one may rely on a generative model (black box simulator) \mathcal{G} , which samples a triplet $\langle s', o, r \rangle$ given a state-action pair $\langle s, a \rangle$.

POMCP

Partially observable Monte-Carlo planning (POMCP) (Silver and Veness 2010) is a sampling-based POMDP solver that extends MCTS to solve large discrete POMDPs. It maintains a partial search tree online during execution, each node corresponding to an action-observation history. Starting from the current root node under h_t , POMCP samples states from an estimated belief \hat{b}_t and simulates trajectories in the search tree. Along such a trajectory, it uses a bandit-based strategy to favor actions with high estimated value while still refining estimates of all action values. At the end of each trajectory, the search tree is locally expanded, and a heuristic is used to estimate the new leaf's value. When visiting a node, the current state is stored, thus contributing to that node's belief estimate \hat{b}_t , which will serve in particular if, after an actual transition, the node becomes the new root node.

Finite State Controllers

Sampling-based solvers such as POMCP typically build a search tree, and could thus be used in finite-horizon problems to compute policy trees offline. However, for an infinite-horizon setting, a policy tree will have infinitely many nodes, which cannot be all visited and stored. In infinite-horizon

POMDPs, solution policies are better represented as *finite state controllers* (FSC) (also called *policy graphs* (Meuleau et al. 1999)), i.e., automata whose transitions from one internal state to the next depend on the received observations and generate the actions to be performed.

Definition 1 For some POMDP sets \mathcal{A} and Ω , a (deterministic) FSC is specified by a tuple $fsc \equiv \langle \mathcal{N}, \eta, \psi \rangle$, where: (i) \mathcal{N} is a finite set of nodes, with n_0 the start node; (ii) $\eta : \mathcal{N} \times \mathcal{A} \times \Omega \rightarrow \mathcal{N}$ is the node transition function; $n' = \eta(n, \langle a, o \rangle)$ is the next node after executing action a and observing o from node n ; (iii) $\psi : \mathcal{N} \rightarrow \mathcal{A}$ is the action-selection function of the FSC; $a = \psi(n)$ is the action triggered when in node n .

In this article, we represent POMDP policies as FSCs (as described in the next section).

4 Monte-Carlo Graph Search for POMDPs

This section presents our algorithmic contribution. As an offline solver, POMCGS computes a compact policy represented by an FSC. Its main algorithm is given in Algorithm 1. First, POMCGS (i) initializes some parameters (line 1); (ii) solves the underlying MDP (i.e., assumes full observability of the POMDP problem) to provide an upper-bounding value function V_{MDP} (line 2); and (iii) creates an FSC policy π_{fsc} with a start node n_0 labeled by the initial belief b_0 (lines 3 to 5). POMCGS then iteratively improves (line 7) and evaluates (line 8) its FSC policy until convergence, before pruning unreachable nodes in the solution FSC using a graph traversal from n_0 to identify reachable nodes (line 9).

Section 4 will discuss how to efficiently identify similar beliefs, before describing how to construct an FSC on the fly in Sec. 4. Sec. 4 then explains POMCGS’s stopping criterion, and Sec. 4 presents the methods we use to address continuous domains.

Algorithm 1: POMCGS Main Algorithm

- 1 [Initialization:] Define ϵ and ξ for convergence detection and belief merging, and nb_{eval} and nb_{sim} for policy evaluation and improvement.
 - 2 Compute V_{MDP} by solving the underlying MDP
 - 3 $n_0 \leftarrow \text{node}(b_0)$
 - 4 $\mathcal{N}_{fsc} \leftarrow \{n_0\}$
 - 5 $\pi_{fsc} \leftarrow \langle \mathcal{N}_{fsc}, \psi, \eta \rangle$
 - 6 **while** $(\bar{V}(b_0)_{\pi_{fsc}} - \underline{V}(b_0)_{\pi_{fsc}}) > \epsilon$ **do**
 - 7 UpdateFSC($\pi_{fsc}, b_0, \xi, nb_{sim}$)
 - 8 EvaluateFSC($\pi_{fsc}, b_0, nb_{eval}$)
 - 9 Pruning(π_{fsc})
 - 10 **return** π_{fsc}
-

Merging Similar Belief Nodes Efficiently

A POMDP’s optimal value function V^* is continuous in belief space (Smallwood and Sondik 1973). Consequently, two nodes of the policy tree with similar beliefs can be merged with limited harm to the resulting policy (Hsu, Lee, and Rong 2007).

To enable merging, POMCGS needs each new node to be equipped with a well-estimated belief (using many particles) from the beginning, in contrast to POMCP. When facing a new belief (estimate) b , the norm-1 distance between beliefs is used to search for an existing node n (if one exists) such that

$$\|b - n.b\|_1 \leq \xi, \quad (1)$$

with $\xi > 0$ a threshold. As a naive search would be computationally expensive for large FSCs, nodes are stored in a data structure similar to cover trees (Beygelzimer, Kakade, and Langford 2006; Izbicki and Shelton 2015), which significantly reduces the number of comparisons, and also allows for cheap insertions of new nodes. This process (called *SearchOrInsert* in Algorithm 2) takes the current set of FSC nodes \mathcal{N} and a new belief b' as input, then returns either: (i) an existing node n whose norm-1 distance satisfies Equation (1), or (ii) a new created node with b' since there is no similar belief found in the FSC. Another function, *SearchNode*(b', \mathcal{N}), searches for the node in \mathcal{N} closest to b' .

Constructing an FSC through Monte-Carlo Simulations

The policy improvement method *UpdateFSC* is given in Alg. 2. As in POMCP, we relax the need for an exact model by performing Monte-Carlo simulations with a black-box simulator. But, unlike POMCP, our method directly builds and searches in a graph (finite-state controller) instead of a tree. This process starts from the initial belief b_0 and repeatedly samples states to simulate nb_{sim} trajectories (lines 2–4), before returning the current FSC in line 5. Trajectories are generated by the recursive *Simulate* function, which stops when $\frac{\gamma^\delta}{1-\gamma}(r_{\max} - r_{\min}) \leq \epsilon$, where δ is the current depth, and the reward is bounded in $[r_{\min}, r_{\max}]$.

In *Simulate*, an action a is selected using the UCB rule in line 8. Visit counts $N(n)$ and $N(n, a)$ of the current node n are updated *before* moving on to the next node (lines 9–10). This is required to avoid repeating the same decision if node n is visited twice along the same trajectory, which may happen in such a graph with loops. Then, if a is performed for the first time in n , a call to *ProcessAction* (detailed in the next paragraph) creates outgoing connections for each observation and returns a first estimate of a ’s action value. If a has already been visited in n , POMCGS simulates a with the sampled state s , derives the next FSC node n' , updates n ’s Q -value, and assigns $\psi(n)$ with $\arg \max_a (Q(n, a))$ (lines 14–18) before returning the trajectory’s estimated return R .

ProcessAction simulates a in n ’s belief $n.b$ repeatedly until enough particles are collected to build the next nodes (lines 23–28). Each encountered observation (gathered in a set in line 30) is attached to its belief estimate (map created in line 31). Then, for each reachable belief b' , if the FSC has reached its size limit $N_{max-fsc}$, we choose the next node n' in \mathcal{N} minimizing $\|n'.b - b'\|_1$ (line 35). Otherwise, we search for a node whose belief is ξ -close to b' in norm-1 distance, and, if none exists, we create a next node n' with belief b' and add it to the FSC (line 37). If n' is created, then its value is initialized optimistically using a V_{MDP} heuristic (Hauskrecht

2000) (line 39). In line 41, node n 's Q -value is initialized with the estimate of the instant reward, plus the discounted future values.

Algorithm 2: FSC Improvement Procedure

```

1 Fct UpdateFSC( $\pi_{fsc}, b_0, \xi, nb_{sim}$ )
2   for  $i \in 1 : nb_{sim}$  do
3      $s \sim b_0$ 
4      $\text{Simulate}(s, n_0, 0)$ 
5   return  $\pi_{fsc}$ 
6 Fct Simulate( $s, n, \delta$ )
7   if  $\frac{\gamma^\delta}{1-\gamma}(r_{\max} - r_{\min}) < \epsilon$  then return 0
8    $a \leftarrow \arg \max_{a \in C(n)} \left[ Q(n, a) + c \sqrt{\frac{\log N(n)}{N(n, a)}} \right]$ 
9    $N(n) \leftarrow N(n) + 1$ 
10   $N(n, a) \leftarrow N(n, a) + 1$ 
11  if  $N(n, a) = 1$  then
12    return  $\text{ProcessAction}(n, a, \delta)$ 
13  else
14     $(s', o) \sim \mathcal{G}(s, a)$ 
15     $n' \leftarrow \eta(n, \langle a, o \rangle)$ 
16     $R \leftarrow r(n, a) + \gamma \text{Simulate}(s', n', \delta + 1)$ 
17     $Q(n, a) \leftarrow Q(n, a) + \frac{R - Q(n, a)}{N(n, a)}$ 
18     $\psi(n) \leftarrow \arg \max_a Q(n, a)$ 
19    return  $R$ 
20 Fct ProcessAction( $n, a, \delta$ )
21   $r(n, a) \leftarrow 0$ 
22   $\mathcal{P} \leftarrow \emptyset$ 
23  repeat
24     $s \sim n.b$ 
25     $\langle s', o, r \rangle \sim \mathcal{G}(s, a)$ 
26     $r(n, a) \leftarrow r(n, a) + r$ 
27     $\mathcal{P} \leftarrow \mathcal{P} \cup \{s', o\}$ 
28  until  $|\mathcal{P}| \geq nb_{particles}$ 
29   $r(n, a) \leftarrow \frac{r(n, a)}{|\mathcal{P}|}$ 
30   $\tilde{\Omega} \leftarrow \mathcal{P}.\{o\}$ 
31   $B' \leftarrow \text{BuildBeliefs}(\mathcal{P}.\{s'\}, \tilde{\Omega})$ 
32  for  $\tilde{o} \in \tilde{\Omega}$  do
33     $b' \leftarrow B'[\tilde{o}]$ 
34    if  $|\mathcal{N}| = \mathcal{N}_{max-fsc}$  then
35       $n' \leftarrow \text{SearchNode}(b', \mathcal{N})$ 
36    else
37       $n' \leftarrow \text{SearchOrInsert}(b', \mathcal{N}, \xi)$ 
38      if  $N(n') = 0$  then
39         $V(n') \leftarrow \sum_s b'(s) V_{MDP}(s)$ 
40       $\eta(n, \langle a, \tilde{o} \rangle) \leftarrow n'$ 
41   $Q(n, a) \leftarrow r(n, a) + \gamma \sum_{\tilde{o}} \frac{|B'[\tilde{o}]|}{|B'|} V(n')$ 
42  return  $Q(n, a)$ 

```

Note that the solution FSC may be incomplete, *i.e.*, with leaves that have not been visited yet. If such a leaf is reached at execution time, then the blind policy is executed to guarantee a lower-bound value (see next section). If a pre-computed FSC is desired to have some understanding of the policy prior

to executing it, but computational resources are still available at execution time, then an online solver could be used once a leaf of the FSC is reached (taking over control from then on, using the leaf's attached belief as initial belief), rather than relying on a blind policy to achieve better performance.

Estimating Upper and Lower Bounds

We propose regularly estimating upper- and lower-bounds of the current FSC's value at b_0 (denoted \bar{V} and \underline{V}), using the function *EvaluateFSC* described in Algorithm 3, to stop POMCGS when the gap between bounds is small enough, despite the lack of guarantees w.r.t. $V^*(b_0)$. These estimates are obtained through nb_{eval} Monte-Carlo simulations of the FSC (line 3), stopping trajectories either when the remaining expected return is below a threshold $\epsilon > 0$ (line 7), or upon encountering a node n with an insufficient visit count (line 8), *i.e.*, $N(n) \geq N^*$ (with $N^* \geq 1$, so that the simulation stops before leaving the FSC if it is incomplete). In the latter case, the value at n is upper-bounded with a V_{MDP} heuristic (line 15), and lower-bounded using a blind policy approximation (line 16) (Hauskrecht 2000), which repeatedly selects the action that maximizes the worst-case reward.

Theoretically, both estimates asymptotically converge to $V_{fsc}(b_0)$ upon increasing the number of simulations nb_{eval} (and decreasing ϵ) only for FSCs where $N(n) \geq N^*$ for any reachable node n , implying that these FSCs are "closed", *i.e.*, without open nodes. We use large values of N^* expecting that decisions will be optimal in any node n with $N(n) \geq N^*$. However this also requires having good enough belief estimates (*i.e.*, large $nb_{particles}$) and few merges (*i.e.*, small ξ). Under these conditions, POMCGS shall stop with an FSC that is ϵ -optimal in discrete problems.

Algorithm 3: FSC Evaluation Procedure

```

1 Fct EvaluateFSC( $\pi_{fsc}, b_0, nb_{eval}$ )
2    $\bar{V}, \underline{V} \leftarrow 0$ 
3   for  $i \in 1 : nb_{eval}$  do
4      $s \sim b_0$ 
5      $\delta \leftarrow 0$ 
6      $n \leftarrow n_0$ 
7     while  $\frac{\gamma^\delta}{1-\gamma}(r_{\max} - r_{\min}) < \epsilon$  do
8       if  $N(n) > N^*$  then
9          $a \leftarrow \psi(n)$ 
10         $(s', o, r) \sim \mathcal{G}(s, a)$ 
11         $n \leftarrow \eta(n, \langle a, o \rangle)$ 
12         $\bar{V} \leftarrow \bar{V} + \gamma^\delta r$ 
13         $\underline{V} \leftarrow \underline{V} + \gamma^\delta r$ 
14      else
15         $\bar{V} \leftarrow \bar{V} + \gamma^\delta \sum_{s \in n.b} n.b(s) \cdot V_{MDP}(s)$ 
16         $\underline{V} \leftarrow \underline{V} + \gamma^\delta \frac{1}{1-\gamma} \max_a \min_s r(s, a)$ 
17        break
18       $\delta \leftarrow \delta + 1$ 
19  return  $\frac{\bar{V}}{nb_{eval}}, \frac{\underline{V}}{nb_{eval}}$ 

```




Figure 1: An illustration of observation clustering using K -Means in the Light Dark Problem (Sunberg and Kochenderfer 2017). Each point corresponds to a particle pair of observations (value on the y axis) and new states (value on the x -axis) collected by performing an action a from some belief b . Particle pairs are clustered according to the observation values into $K = 5$ labels represented by different colors, and beliefs are created with states attached to the same label.

Handling Continuous POMDPs

We now describe how POMCGS is modified to handle continuous POMDPs, where any combination of state, action and observation spaces could be continuous.

Continuous states are an issue only when it comes to measuring the distance between two belief estimates. In this case, the state space is discretized, so that particles are grouped in bins, which allows still using the same distance. A dedicated metric for continuous distributions, such as the Wasserstein distance, could allow for better merges, but may be much more time consuming.

Continuous actions are handled through action progressive widening (APW) as presented in Algorithm 4, which is called instead of line 8 of Algorithm 2. APW is stopped in nodes considered “finalized” ($N(n) \geq N^*$), so that the number of actions is bounded. One could also sample a set of actions initially, but it is preferable to use progressive widening to ensure that a few actions are well estimated rather than obtaining poor estimates for many actions.

Algorithm 4: Action Progressive Widening

```

1 Fact ActionProgressiveWidening( $n$ )
2   if  $|C(n)| \leq k_a N(n)^{\alpha_a}$  and  $N(n) < N^*$  then
3      $a \sim \mathcal{A}$ 
4      $C(n) \leftarrow C(n) \cup a$ 
5   else
6      $a \leftarrow \arg \max_{a \in C(n)} [Q(n, a) + c \sqrt{\frac{\log N(n)}{N(n, a)}}]$ 
7   return  $a$ 

```

Continuous observation spaces need to be discretized because the FSC needs to anticipate all possible observations.

This is different from online algorithms such as POMCPOW, which can rely on observation progressive widening for exploration and then adapt to actual observations at execution time. Here, we propose to use K -means (MacQueen 1967) instead of line 30 of Algorithm 2 to cluster observations independently in each created node. When simulating an action a from some belief b in function *ProcessAction* (Alg. 2), observation and new state pairs (o, s') are collected and clustered by observation values according to a pre-defined parameter $K \in \mathbb{N}^*$, and new beliefs are created with their corresponding states (l. 31). An example of this clustering is illustrated in Fig. 1. Note that, unlike POMCPOW, we do not need an explicit observation model to query the relative likelihood of different observations, a black-box simulator suffices.

5 Experiments on POMDP Benchmarks

In this section, we evaluate POMCGS on various benchmarks provided by the POMDP.jl framework (Egorov et al. 2017), and compare our methods with other state-of-the-art offline solvers: SARSOP (Kurniawati, Hsu, and Lee 2008) and MCVI (Bai et al. 2010); and online planners: DESPOT (Somani et al. 2013), POMCPOW (Sunberg and Kochenderfer 2017), and AdaOPS (Wu et al. 2021). Note that SARSOP is usually considered as the best ϵ -optimal offline POMDP solver and often used to calibrate other online methods in small or moderate-size POMDPs. However SARSOP requires explicit POMDP models while other solvers do not.

All experiments for POMCGS were conducted on a computer with a 4.8 GHz i7 CPU and 64 GB RAM. Source code and parameter settings are provided in the supplementary material. In this article, V_{MDP} is computed by Q-learning on the underlying MDP, which is used to estimate POMCGS’ upper bound (for continuous domains, Q-learning requires domain discretization). One can also use other MDP techniques to derive V_{MDP} (Kochenderfer 2015), which is a mapping from states to values. In POMCGS, we set $nb_{\text{sim}} = 10^3$ in *UpdateFSC* for the policy improvement, and $nb_{\text{eval}} = 10^5$ in *EvaluateFSC* to evaluate the policy’s upper and lower bounds with $N^* = 50$. For each POMDP benchmark, we ran POMCGS 10 times to compute the average return. We set $\epsilon = 0.01$ to detect convergence.¹ All online solvers are limited to 3 seconds of computation time per step.

Table 1 presents the average discounted returns with associated standard errors for the 7 benchmark problems evaluated. SARSOP is labeled with a * exponent indicating that an explicit POMDP model is required. For POMCGS, we report its final average lower bound. Entries with a “—” symbol indicating that the algorithm is not able to run in this domain.

In terms of value achieved, on small and moderate-size domains, POMCGS is able to achieve near-optimal performance as SARSOP; on large and continuous problems (except the Laser Tag domain), POMCGS is the only offline algorithm to compete with state-of-the-art online algorithms. In particular, although both POMCGS and MCVI are able to derive FSC

¹The UCB constant c was set to 2.0 for RS(7,8), RS(11, 11), RS(15, 15) and Light-Dark; and to 1.0 for Bumper Roomba, Lidar Roomba and Laser Tag. We do not further discuss the influence of this parameter, which appeared not to be significant.

solutions among all offline solvers, there is no belief merging in MCVI. In each MCVI backup, it needs to add one node and compute a new FSC entirely, which is computationally expensive. Therefore, given a fixed timeout, we observed that MCVI’s performance degrades when the required planning horizon increases on large problems. In contrast, POMCGS merges similar nodes, which largely reduces unnecessary computations.

In later sections, we analyze POMCGS’s performance individually for each domain and discuss the influence of parameters.

Benchmark Problems

Rock Sample: Smith and Simmons (2004) proposed a large-state benchmark called *RockSample* (RS), in which a robot moves in a grid-world and tries to collect rocks. In $RS(n, k)$, the map is defined as an $n \times n$ grid and contains k rocks. Each rock may be good or bad while the robot is only equipped a noisy sensor to test it. The robot’s goal is to collect as many good rocks as possible and exit on the east boundary of the map.

We set offline solvers with timeouts of 1 hour, 3 hours, and 6 hours for $RS(7, 8)$, $RS(11, 11)$, and $RS(15, 15)$, respectively. POMCGS can converge to near-optimal solutions similar to SARSOP on $RS(7, 8)$ and $RS(11, 11)$. On the largest domain, $RS(15, 15)$, SARSOP cannot run successfully (the model being too large to generate and store), while POMCGS is able to obtain competitive performance compared with state-of-the-art online solvers.

Light Dark: Sunberg and Kochenderfer (2017) propose a continuous-observation problem where the agent moves in a one-dimensional space, trying to reach a target region. The observation of the location is only accurate in a specific “light” region and is much more noisy in dark regions.

We set a 1-hour timeout for offline solvers on the Light Dark problem. Under these conditions, POMCGS achieves a value of 3.74, closely matching the performance of AdaOPS and outperforming other offline and online algorithms. Due to the 1-dimensional nature of the observations in Light Dark, the K-means method effectively clusters continuous observations and constructs corresponding beliefs, contributing to POMCGS’s strong performance. Further improvements may be achievable by increasing the value of K , provided that sufficient computational resources are available.

Laser Tag: In Laser Tag (Somani et al. 2013), a robot moves in a 7×11 grid map (with randomly placed obstacles) and tries to tag an escaping target. In the beginning, the robot knows neither its own location nor the target location. Moreover, the robot is equipped with a laser sensor that reads noisy distances in eight directions. This creates a large number of possible eight-dimensional observations (around 1.5×10^6).

We use POMCGS together with observation clustering with a 3-hour timeout on Laser Tag due to its large observation space. Interestingly, although the state space in Laser Tag is not extremely large, it turns out to be the most difficult problem for our offline method. This may be due for instance 1. to the 8-dimensional observation space possibly preventing

from obtaining relevant clusters, and/or 2. to the merging of nodes, which may prevent necessary belief distinctions (while the agent needs to progressively refine its knowledge). As a result, given a bounded running time, POMCGS is not able to compute a complete policy to tag the target, which leads to a poor performance.

Bumper Roomba and Lidar Roomba: In the *Roomba* problem (Menda, Sunberg, and Kochenderfer 2023), a cleaning robot is placed in a known room with an unknown initial position. The robot needs to localize itself and reach a target region. There are two sensing options for the robot depending on the problem: a Bumper sensor that can detect collisions, or a Lidar sensor that reads a noisy distance to the front.

We set a timeout of 3 hours for POMCGS in these two benchmark domains². Both POMCGS and POMCPOW can successfully localize and navigate the robot to the target, achieving similar performance in the Lidar Roomba problem. In the Bumper Roomba problem, POMCGS converges to a positive value of 0.31, while we are unable to tune POMCPOW to have a similar performance. During simulations, we observe that our offline policy is able to take correct decisions earlier than POMCPOW, and, more importantly, has less failing executions that navigate the robot to the wrong target. POMCGS may benefit from never using belief estimates relying on very few particles, and from its FSC structure to update action values efficiently.

A Closer Look at POMCGS’s Behavior

Given a fixed computation time, we observe that parameters (ξ , K , and $nb_{\text{particles}}$), which influence belief formation, have a strong impact on the final results. Among these parameters, ξ represents the belief merging threshold and directly affects the accuracy of the belief comparison process; K controls observation clustering, influencing the selection of samples for constructing new beliefs in continuous-observation POMDPs; and $nb_{\text{particles}}$ defines the number of samples in a belief, impacting the accuracy of belief estimation. To better understand POMCGS’s behavior, this section provides an analysis of the influence of parameters ξ , K , and $nb_{\text{particles}}$.

We first investigate the influence of the number of particles ($nb_{\text{particles}}$) on the performance of POMCGS. We select two problems: $RS(7, 8)$ and Light Dark, fixing $\xi = 0.1$ for both problems and setting $K = 5$ for observation clustering in the Light Dark problem. All experiments are conducted with a timeout of 1 hour. We evaluate five different values of $nb_{\text{particles}}$. Our findings indicate that increasing the number of particles improves belief estimation accuracy. As shown in Figure 2, we observe that a higher number of particles consistently leads to better performance. In smaller domains, such as $RS(7, 8)$, even a relatively small number of particles per node can yield competitive solutions. However, in more complex problems, a larger $nb_{\text{particles}}$ is generally necessary to accurately represent beliefs. Based on these observations, we fix $nb_{\text{particles}}$ at 10^4 for subsequent tests.

²Wu et al. (2021) experimented with AdaOPS on Lidar Roomba and Bumper Roomba with finite-action sets and compared it with DESPOT.

	RS(7, 8)	RS(11, 11)	RS(15, 15)	Light Dark	Bumper Roomba	Lidar Roomba	Laser Tag
$ \mathcal{S} $	12, 544	247, 808	7, 372, 800	∞	∞	∞	5, 830
$ \mathcal{A} $	13	16	20	3	∞	∞	5
$ \mathcal{Z} $	3	3	3	∞	2	∞	$\sim 1.5 \times 10^6$
<i>Offline Solvers</i>							
POMCGS	21.13 ± 1.09	19.24 ± 1.31	16.28 ± 1.44	3.74 ± 0.06	0.31 ± 0.18	0.79 ± 0.12	-19.42 ± 0.15
MCVI	19.36 ± 1.21	18.09 ± 1.25	13.70 ± 0.96	2.94 ± 0.15	—	—	—
SARSOP*	21.45 ± 0.08	21.56 ± 0.12	—	—	—	—	—
<i>Online Planners</i>							
AdaOPS	20.93 ± 0.35	20.17 ± 0.41	17.81 ± 0.43	3.76 ± 0.10	—	—	-8.82 ± 0.21
DESPOT	21.08 ± 0.21	20.55 ± 0.26	18.32 ± 0.29	2.66 ± 0.14	—	—	-9.14 ± 0.27
POMCPOW	20.55 ± 0.17	19.51 ± 0.19	16.57 ± 0.23	3.29 ± 0.08	-0.41 ± 0.15	0.73 ± 0.10	-9.46 ± 0.24

Table 1: Performance Comparison of Different POMDP Planning Algorithms




Figure 2: POMCGS’s performances on RS(7, 8) and Light Dark problems with 5 different $nb_{\text{particles}}$ values.

We then study the POMCGS’s performance with three different belief merging thresholds on RS(7, 8). As illustrated in Figure 3, when ξ increases from 0.1 to 0.3, fewer iterations are required to converge, but this comes at the cost of decreased value because the belief merging becomes less accurate. Especially, POMCGS has the highest lower bound value when $\xi = 0.1$, but its lower bound and upper bound have not converged yet. We next study the influence of the parameter K for observation clustering. In Light Dark, we keep $\xi = 0.1$ and test three K values (5, 10, 20). The POMCGS’s behavior follows a similar trend where fewer observation clusters result in a lower convergence value and fewer iterations.

We observe that tuning ξ and K allows for a trade-off between producing high solution quality when ξ is small and K is large, which requires more planning time (iterations) to converge, and achieving faster convergence but lower solution quality when ξ is large and K is small. How to tune ξ and K thus usually depends on the problem at hand with respect to the solution quality and computational resources. In this article, we propose a standard parameter configuration with respect to the problem size as shown in Table 2, which is also used to report POMCGS’s performance in Table 1. One may achieve better performance by tuning these parameters when computational resources are available.

Problem Type	Examples	$nb_{\text{particles}}$	ξ	K (if needed)
Small-to-medium	RS(7,8)	5×10^3	0.1	10
	RS(11,11)			
	Light Dark Bumper Roomba			
Large	RS(15,15)	10^4	0.3	5
	Lidar Roomba			
	Laser Tag			

Table 2: Parameter Settings for Different Problems

6 Discussion

Contribution: POMCGS differs from previous state-of-the-art offline solvers, which perform backup operations on alpha-vectors or policy graphs, as it does not adhere to the value iteration scheme. Instead, POMCGS directly searches and updates within the same policy graph. The core innovation of POMCGS is its ability to merge similar beliefs within a Monte Carlo search tree, producing offline policies represented as finite-state controllers (FSCs). This approach reduces redundant computations and facilitates efficient updates of node values within the FSC structure. Our experiments demonstrate that POMCGS scales effectively across challenging benchmarks where prior offline solvers, such as SARSOP, cannot operate. Furthermore, unlike previous offline methods, POMCGS achieves competitive performance with state-of-the-art online methods on most complex POMDPs, particularly in scenarios with small or low-dimensional continuous observation spaces. The results also indicate that, despite the relatively coarse node merging, large state spaces do not pose significant challenges for POMCGS.

Limitation: While POMCGS enhances offline planning performance on large problems, we acknowledge that no single algorithm can address all possible scenarios in this field. One notable limitation of POMCGS, as well as other potential offline methods, is handling high-dimensional observation spaces (e.g., in the Laser Tag problem). Unlike online planners, which can employ observation progressive widening and reason about actual observations during execution, offline algorithms like POMCGS must discretize observations for both planning and execution. This constraint presents significant challenges for high-dimensional data, as effective

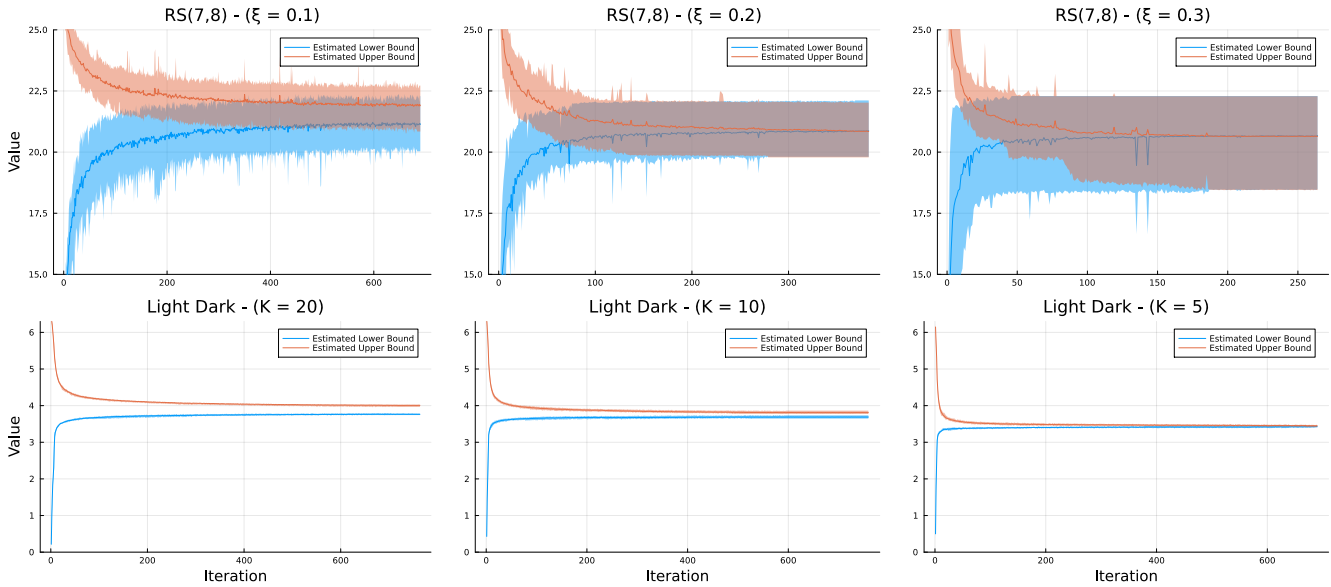


Figure 3: Figures present the evolution of the POMCGS’s upper and lower bounds at each iteration for the RS(7,8) and Light Dark with different parameters.

observation discretization (e.g., using K-Means clustering) becomes increasingly difficult.

Convergence: Lim et al. (2023) demonstrate that, when any online sampling-based solving algorithm for MDPs can be adapted to solving a POMDP through deriving an associated *Particle Belief MDP* (PB-MDP), and convergence guarantees of the algorithm (*i.e.*, concentration bounds exploiting the Lipschitz continuity of the optimal value function) extend to POMDPs. In this paper, Algorithm 2 can also be interpreted as relying on a PB-MDP, but as an offline algorithm, it must anticipate and discretize all possible continuous observations, employing a node merging process that results in an FSC rather than a tree or directed acyclic graph. Proving convergence guarantees thus remains an open question.

Future Work: We believe that offline planning in POMDPs with high-dimensional observations is a critical area deserving of dedicated research. In this direction, we aim to improve the observation clustering process by incorporating dimensionality reduction techniques. However, care must be taken regarding potential information loss, which could impair the algorithm’s ability to distinguish between beliefs. Another area of future work could explore avoiding state discretization by using more effective belief comparison metrics.

7 Conclusion

In the last decade, most POMDP planning methods have focused on solving large problems online. However, in extreme environments or under severe constraints, embedded systems often lack the resources needed for online computation. At the same time, traditional offline solvers struggle to scale up for large problems or to provide effective solutions in continuous domains. In this article, we present a novel offline POMDP algorithm, POMCGS, to address these challenging

POMDPs and obtain compact policies that can be directly executed without further computations. Relying on Monte-Carlo simulations, we relax the requirements of explicit models and successfully turn the search tree into a policy-graph (*i.e.*, a finite-state controller). In particular, we pave the way for a new approach, distinct from previous (point-based) value iteration methods, to compute complete policies in large, continuous domains, which are long-standing challenges in offline POMDP planning.

Future work will focus on improving the observation clustering process in high-dimensional observation domains and providing better belief comparison techniques to avoid discretizing states, so that we push the offline POMDP planning to a new level.

Acknowledgments

This work has been supported by the EPSRC Energy Programme under UKAEA/EPSC Fusion Grant 2022/2027 No. EP/W006839/1 and by French National Research Agency (ANR) through the “Flying Coworker” Project under Grant 18-CE33-0001.

References

- Bai, H.; Hsu, D.; Lee, W. S.; and Ngo, V. A. 2010. Monte Carlo Value Iteration for Continuous-State POMDPs. In *Algorithmic Foundations of Robotics IX: Selected Contributions of the Ninth International Workshop on the Algorithmic Foundations of Robotics*, 175–191. Springer.
- Ballesteros, J.; Merino, L.; Trujillo, M. A.; Viguria, A.; and Ollero, A. 2013. Improving the efficiency of online POMDPs by using belief similarity measures. In *ICRA-13*.
- Beygelzimer, A.; Kakade, S.; and Langford, J. 2006. Cover Trees for Nearest Neighbor. In *ICML-06*.

- Egorov, M.; Sunberg, Z. N.; Balaban, E.; Wheeler, T. A.; Gupta, J. K.; and Kochenderfer, M. J. 2017. POMDPs.jl: A Framework for Sequential Decision Making under Uncertainty. *JMLR*, 18(26): 1–5.
- Garg, N. P.; Hsu, D.; and Lee, W. S. 2019. DESPOT-Alpha: Online POMDP Planning with Large State and Observation Spaces. *RSS*.
- Grzes, M.; Poupart, P.; Yang, X.; and Hoey, J. 2015. Energy Efficient Execution of POMDP Policies. *IEEE Trans. on Cybernetics*, 45: 2484–2497.
- Hansen, E. A. 1998a. An Improved Policy Iteration Algorithm for Partially Observable MDPs. In *NIPS-97*.
- Hansen, E. A. 1998b. Solving POMDPs by Searching in Policy Space. In *UAI-98*.
- Hauskrecht, M. 2000. Value-Function Approximations for Partially Observable Markov Decision Processes. *JAIR*, 13: 33–94.
- Hoerger, M.; and Kurniawati, H. 2021. An On-Line POMDP Solver for Continuous Observation Spaces. In *ICRA*, 7643–7649.
- Hsu, D.; Lee, W. S.; and Rong, N. 2007. What makes some POMDP problems easy to approximate? In *Proceedings of the 20th International Conference on Neural Information Processing Systems*, 689–696.
- Izbicki, M.; and Shelton, C. R. 2015. Faster Cover Trees. In Bach, F. R.; and Blei, D. M., eds., *ICML-15*.
- Ji, S.; Parr, R. E.; Li, H.; Liao, X.; and Carin, L. 2007. Point-Based Policy Iteration. In *AAAI-07*.
- Kochenderfer, M. J. 2015. *Decision making under uncertainty: theory and application*. MIT press.
- Kocsis, L.; and Szepesvári, C. 2006. Bandit Based Monte-Carlo Planning. In *ECML-06*.
- Kujanpää, K.; Babadi, A.; Zhao, Y.; Kannala, J.; Ilin, A.; and Pajarinen, J. 2023. Continuous Monte Carlo Graph Search. arXiv:2210.01426.
- Kurniawati, H. 2022. Partially observable Markov decision processes and robotics. *Annual Review of Control, Robotics, and Autonomous Systems*, 5: 253–277.
- Kurniawati, H.; Hsu, D.; and Lee, W. S. 2008. SARSOP: Efficient point-based POMDP planning by approximating optimally reachable belief spaces. In *RSS-08*.
- Lauri, M.; Hsu, D.; and Pajarinen, J. 2022. Partially observable Markov decision processes in robotics: A survey. *IEEE Trans. on Robotics*, 39(1): 21–40.
- Leurent, E.; and Maillard, O.-A. 2020. Monte-Carlo Graph Search: the Value of Merging Similar States. In *Proc. of the 12th Asian Conf. on ML*.
- Lim, M. H.; Becker, T. J.; Kochenderfer, M. J.; Tomlin, C. J.; and Sunberg, Z. N. 2023. Optimality Guarantees for Particle Belief Approximation of POMDPs. *J. Artif. Intell. Res.*, 77: 1591–1636.
- MacQueen, J. B. 1967. Some Methods for Classification and Analysis of MultiVariate Observations. In *Proc. of the 5th Berkeley Symposium on Math. Statistics and Probability*.
- Madani, O.; Hanks, S.; and Condon, A. 1999. On the Undecidability of Probabilistic Planning and Infinite-Horizon Partially Observable Markov Decision Problems. In *AAAI-99*.
- Menda, K.; Sunberg, Z.; and Kochenderfer, M. 2023. RoombaPOMDPs.jl. <https://github.com/sisl/RoombaPOMDPs.jl/>. [Online; accessed 03-Oct-2023].
- Meuleau, N.; Kim, K.-E.; Kaelbling, L. P.; and Cassandra, A. R. 1999. Solving POMDPs by searching the space of finite policies. In *UAI-99*.
- Pajarinen, J.; and Peltonen, J. 2011. Periodic Finite State Controllers for Efficient POMDP and DEC-POMDP Planning. In *NIPS-11*.
- Paludo Licks, G.; Patrizi, F.; and De Giacomo, G. 2024. Monte Carlo Tree Search with State Merging for Reinforcement Learning in Regular Decision Processes. In *27th European Conference on Artificial Intelligence*.
- Papadimitriou, C. H.; and Tsitsiklis, J. N. 1987. The Complexity of Markov Decision Processes. *Mathematics of operations research*, 12(3): 441–450.
- Pineau, J.; Gordon, G.; and Thrun, S. 2003. Point-Based Value Iteration: An Anytime Algorithm for POMDPs. In *IJCAI-03*.
- Poupart, P.; and Boutilier, C. 2003. Bounded Finite State Controllers. In *NIPS-03*.
- Silver, D.; and Veness, J. 2010. Monte-Carlo Planning in Large POMDPs. In *NIPS-10*, volume 23.
- Smallwood, R. D.; and Sondik, E. J. 1973. The Optimal Control of Partially Observable Markov Decision Processes over a Finite Horizon. *Operation Research*, 21: 1071–1088.
- Smith, T.; and Simmons, R. 2004. Heuristic Search Value Iteration for POMDPs. In *UAI-04*.
- Smith, T.; and Simmons, R. 2005. Point-Based POMDP Algorithms: Improved Analysis and Implementation. In *UAI-05*.
- Somani, A.; Ye, N.; Hsu, D.; and Lee, W. S. 2013. DESPOT: Online POMDP Planning with Regularization. In *NIPS-13*.
- Sunberg, Z.; and Kochenderfer, M. 2017. POMCPOW: An online algorithm for POMDPs with continuous state, action, and observation spaces. In *ICAPS-17*.
- Toussaint, M.; Harmeling, S.; and Storkey, A. 2006. Probabilistic inference for solving (PO)MDPs. Technical Report EDI-INF-RR-0934, School of Informatics, Univ. of Edinburgh.
- Wu, C.; Yang, G.; Zhang, Z.; Yu, Y.; Li, D.; Liu, W.; and Hao, J. 2021. Adaptive Online Packing-guided Search for POMDPs. In *NIPS-21*.
- Åström, K. 1965. Optimal control of Markov processes with incomplete state information. *Journal of Mathematical Analysis and Applications*, 10(1): 174 – 205.