Adaptive Online Packing-guided Search for POMDPs

Anonymous Author(s)
Affiliation
Address
email

Abstract

The partially observable Markov decision process (POMDP) provides a general framework for modeling an agent's decision process with state uncertainty, and online planning plays a pivotal role in solving it. A belief is a distribution of states representing state uncertainty. Methods for large-scale POMDP problems rely on the same idea of sampling both states and observations. That is, instead of exact belief updating, a collection of sampled states is used to approximate the belief; instead of considering all possible observations, only a set of sampled observations are considered. Inspired by this, we take one step further and propose an online planning algorithm, Adaptive Online Packing-guided Search (AdaOPS), to better approximate beliefs with adaptive particle filter technique and balance estimation bias and variance by fusing similar observation branches. Theoretically, our algorithm is guaranteed to find an ϵ-optimal policy with a high probability given enough planning time. We evaluate our algorithm on several tricky POMDP domains, and it outperforms the state-of-the-art in all of them.

1 Introduction

POMDPs generalize the MDPs by considering the state uncertainty [1]. In MDPs, the agent knows its state exactly. In POMDPs, though an MDP determines the underlying dynamics, the agent cannot access its actual state at each timestep and, instead, receives an observation serving as a clue. In order to make decisions in POMDPs, the state uncertainty must be handled. A popular method for tackling the uncertainty is to maintain a belief, a distribution on the state space, where the probability of a state indicates how possible the agent believes it is. However, an exact belief updating requires $O(|\mathcal{S}|^2)$ computations, which is unaffordable when the state space $\mathcal{S}$ is immense and impossible when $\mathcal{S}$ is continuous. A more practical way is to approximate the belief with a collection of weighted particles (or samples) and update it with particle filtering [2].

Finding an optimal solution for finite-horizon or discounted infinite-horizon POMDPs is proved to be PSPACE-complete [3]. Therefore, online planning, which computes a real-time solution, plays a key role in solving POMDPs, following from two reasons. The first is that online planning searches for a solution for not all possible beliefs but the current belief alone, relieving the computational load significantly. Second, online planning can take advantage of the approximate offline solution and other heuristics to expedite the tree search process.

Online planning algorithms have made substantial progress in solving large-scale POMDP problems [4–7]. These methods approximate the belief with a collection of sampled states instead of updating it exactly and consider a set of sampled observations instead of all possible observations. However, the estimation derived by simple belief approximation methods that have been applied, such as rejection sampling and sequential importance sampling [8], have a variance increasing exponentially with searching depth. Besides, direct sampling for observations, either the Monte Carlo sampling as in [6] or importance sampling as in [7], yields a large variance when the sample size is small. Inspired by this, we take one step further and propose an online planning algorithm, Adaptive Online

Submitted to 35th Conference on Neural Information Processing Systems (NeurIPS 2021). Do not distribute.

Packing-guided Search (AdaOPS), to better approximate beliefs and better balance estimation bias and variance. It features two innovations, adaptive particle filtering and belief packing.

The adaptive particle filter is designed for online planning. During the belief updating, it resamples adaptively and adjusts the sample size on the fly according to the dispersion of the belief distribution, which allows achieving high approximation accuracy with only a handful of particles.

AdaOPS forms a belief packing by merging similar beliefs. Only beliefs in the packing are explored, thus the name packing-guided search. Belief packing is meant to balance estimation bias and variance. With the use of beliefs in their exact form, it is guaranteed that similar beliefs yield a small difference in their optimal values [9]. We further extend this conclusion to cases where beliefs are approximated by a collection of weighted particles and prove that by merging similar beliefs, our algorithm is capable of converging to the optimal policy with high probability.

We present a discussion on the relation between AdaOPS and previous work in Appendix A.

2 Background

2.1 POMDPs

An MDP is mathematically defined as a tuple $(\mathcal{S}, \mathcal{A}, T, R, \gamma)$, where $\mathcal{S}$ and $\mathcal{A}$ are the state and action spaces, respectively, and T and R are the transition and reward functions, respectively. $\gamma \in [0, 1]$ is the discount factor. A POMDP problem can be seen as an extension to an MDP problem, which introduces two additional components, the observation space $\mathcal{O}$ and the observation function Z. At each timestep, instead of observing the actual state as an MDP agent does, a POMDP agent can only get an observation according to the observation function $Z(o \mid s', a) = \Pr(o \mid s', a)$, which determines the probability of observing $o \in \mathcal{O}$ after executing action $a \in \mathcal{A}$ and reaching state $s' \in \mathcal{S}$. The transition function T and observation function Z are sometimes hard to define explicitly. An alternative is the generative model $s', o, r \leftarrow G(s, a)$, which implicitly defines the transition and observation distributions by randomly generating a possible combination of next state s', observation o, and reward r given state s and action a. The POMDP agent maintains a belief b in states, a distribution on $\mathcal{S}$, and seeks to find a policy π mapping beliefs to actions, which maximizes the average discounted rewards, $\mathbb{E}[\sum_{t=0}^{\infty} \gamma^t R(s_t, \pi(b_t)]$. In this paper, assuming the reward function is Borel measurable and bounded, we denote $R_{\max} = \|R\|_\infty$.

2.2 Particle Filters

Bayes filter recursively updates the posterior of the hidden states from a sequence of noisy and incomplete observations. In POMDPs, this technique is used to maintain a belief approximation. Particle filter [8] is an important class of Bayes filter, which approximates the belief with a collection of weighted particles, $\{(w_i, s_i)\}_{i \in [N]}$, where s_i is the i-th particle, and w_i is the corresponding normalized weight. It updates the belief by propagating particles with action a, $s'_i \sim T(\cdot \mid s_i, a)$, and reweighting particles with observation o, i.e., $w'_i = w_i Z(o \mid s'_i, a)/\rho$, where ρ is a normalizing factor. A resampling is then followed to sample N particles from $\{(w'_i, s'_i)\}_{i \in [N]}$ with probability proportional to the weights, and the weights of the resampled particles are set to be $1/N$.

Resampling is critical, without which we cannot derive time uniform convergence for particle filters [10]. However, frequent resampling induces sample impoverishment, which says that the diversity of particles is diminished, the whole collection of particles containing few distinct ones. In practice, resampling is only performed when necessary. A rule of thumb is to resample when the particle number N divided by the effective sample size (ESS) is greater than two, where ESS is the sample size required from the true distribution in order to achieve the same sampling error as sampling N samples from the proposal distribution and is approximately calculated as follows [11, 12],

$$\text{ESS} \approx \frac{1}{\sum_{n=1}^{N} w_n^2}. \tag{1}$$

For a comprehensive understanding of the notion, effective sample size, see Appendix B.

2.3 KLD-Sampling

KLD-Sampling is an approach proposed in [13] for adapting the sample size in particle filters according to the dispersion of the posterior distributions. Assuming the true posterior can be represented by a discrete piece-wise constant distribution consisting of k different multidimensional bins, it calculates the number of samples needed to ensure the Kullback-Leibler divergence (KLD) between the approximated distribution and the true distribution is less than an error ζ with a confidence level $1-\eta$. The number of samples N is given by

$$N \approx \frac{k-1}{2\zeta}\left(1-\frac{2}{9(k-1)}+\sqrt{\frac{2}{9(k-1)}}z_{1-\eta}\right), \tag{2}$$

where $z_{1-\eta}$ is the $1-\eta$ quantile of the standard normal distribution. During the sampling, one should keep track of the number of bins that the approximated distribution occupies and estimate the number of samples needed accordingly.

2.4 Belief Packing

Definition 1 (δ-packing)**.** Let $(X, \|\cdot\|)$ be a normed space and $Y \subseteq X$. $\{y_1, y_2, \ldots, y_n\} \subseteq Y$ is a δ-packing of Y if $\min_{i\neq j}\|y_i - y_j\| > \delta$. The δ-packing number of Y, denoted $P^{\delta}(Y)$, is the maximum cardinality of any δ-packing of Y. ■

Remark. If X is totally bounded and $Y \subseteq X$, then we have $P^{\delta}(Y) < \infty$ for all $\delta > 0$.

A belief packing is a special case of δ-packing, where the X is the belief space, a probability simplex of $|\mathcal{S}|-1$ dimensions, and the norm $\|\cdot\|$ is the $L1$ norm. Since a simplex is totally bounded, it follows that the packing number of a belief packing is finite.

3 Method

Online planning consists of two alternating stages, planning and execution. A look-ahead search is conducted during the planning stage to find the best action for the current belief. At the execution stage, the algorithm performs in the real environment the action found, gets a new observation, and updates the belief with a Bayes filter. Then a planning stage follows. Algorithm 1 shows the general procedure of online planning. AdaOPS follows the general procedure and makes changes in the planning part.

Algorithm 1 General Procedure

Input: initial belief b_0
$b \leftarrow b_0$
while True **do**
  $a^* \leftarrow$ PLANNING(b)
  Execute a^*
  Receive observation o
  Update the belief b with action a^* and observation o

Algorithm 2 PLANNING(bel)

bel $\leftarrow$ RESAMPLE(bel)
while time permitting **do**
  $b \leftarrow$ bel
  while $\mathrm{depth}(b) < D$ **do**
    if b is a leaf node **then**
      EXPAND(b)
      BACKUP(b)
      if best action branch changes for some ancestors of b or $\mathrm{EU}(b) \leq 0$ **then**
        break
    $a^* \leftarrow \arg\max_{a\in\mathcal{A}} u(b,a)$
    $o^* \leftarrow \arg\max_{o\in O_{b,a}} \hat{p}(o \mid b,a)\mathrm{EU}(\tau(b,a^*,o))$
    $b \leftarrow \tau(b,a^*,o^*)$
  if $\mathrm{depth}(b) \geq D$ **then**
    $u(b) \leftarrow l(b)$
return $\arg\max_{a\in\mathcal{A}} l(\mathrm{bel}, a)$

3.1 Planning

At each timestep, AdaOPS will gradually build up a belief tree consisting of alternating action and belief nodes, of which the root belief is resampled from the current belief bel using KLD-Sampling. A belief tree built is shown in Figure 1. Each belief node b is represented by a collection of weighted particles. Sibling beliefs share the particles and differ in the corresponding weights. Initially, the belief tree is comprised of the root node alone. A series of explorations will then be made, each traversing a path on the tree and expanding new nodes during the process. The complete pseudocode of AdaOPS is shown in Appendix C.

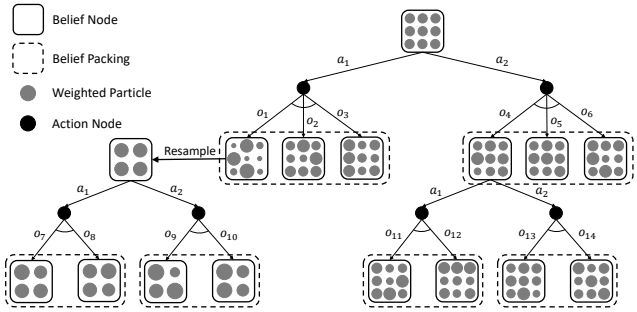

Figure 1: A belief tree built by AdaOPS. The sibling beliefs share the same particles and differ in the corresponding weights, forming a belief packing. Beliefs with high weight disparity will be resampled using KLD-Sampling.

AdaOPS maintains an upper bound $u(b, a)$ and a lower bound $l(b, a)$ of the optimal value for every action node a of belief b. During the exploration, the action a^* with the highest upper bound will be chosen (Line 10 in Alg. 2). This strategy ensures that an action branch gets explored only when its upper bound is higher than or equal to the optimal value. AdaOPS also maintains lower and upper bounds, $l(b)$ and $u(b)$, for each belief node b. After choosing the action, it will choose (Line 11 in Alg. 2) the child belief with the highest probability-weighted excess uncertainty (EU) [5, 14], $\hat{p}(o \mid b, a^*)\mathrm{EU}(\tau(b, a^*, o))$, where $\hat{p}(o \mid b, a^*)$ is an estimation of the probability $\Pr(o \mid b, a^*)$, and $\tau(b, a^*, o)$ is the belief reached after executing a^* and receiving o. The way to estimate $\Pr(o \mid b, a^*)$ is stated in Section 3.2.1. The excess uncertainty is defined as

$$\mathrm{EU}(b) = u(b) - l(b) - \frac{\xi \cdot (u(\mathrm{bel}) - l(\mathrm{bel}))}{\gamma^{\mathrm{depth}(b)}}, \tag{3}$$

where ξ is a parameter for adjusting the desired gap between upper and lower bounds at the root and defaults to 0.95, and $\mathrm{depth}(b)$ is the depth of belief b on the belief tree. The excess uncertainty measures how much uncertainty we need to eliminate at the current belief b if we want the gap at the root, $u(\mathrm{bel}) - l(\mathrm{bel})$, reduced to $\xi \cdot (u(\mathrm{bel}) - l(\mathrm{bel}))$. The excess uncertainty has a desirable property: $\mathrm{EU}(b) \leq \sum_{o \in O_{b,a^*}} \hat{p}(o \mid b, a^*)\mathrm{EU}(\tau(b, a^*, o))$, where O_{b,a^*} is the set of expanded observations under the action branch a^* of belief b. So, we can reduce the excess uncertainty on b by reducing the excess uncertainty on its child beliefs. Note that minimizing the excess uncertainty on the root belief is equivalent to minimizing $u(\mathrm{bel}) - l(\mathrm{bel})$.

Each time a leaf node is encountered (Line 6 in Alg. 2), all valid actions will be tried with new belief nodes expanded. For new belief nodes, the lower and upper bounds will be initialized with heuristics. This process of expanding leaf nodes is detailed in Section 3.2. After expanding a leaf node, a backup will be performed (Line 7 in Alg. 2) to update the bounds of ancestor nodes, which is realized by recursively applying the Bellman equation:

$$\begin{aligned} V(b) &= \max_a V(b, a) = \max_a \left\{ \int_{s \in \mathcal{S}} b(s) R(s, a) \mathrm{d}s + \int_{o \in \mathcal{O}} \Pr(o \mid b, a) V(\tau(b, a, o)) \mathrm{d}o \right\} \\ &= \max_a \left\{ \sum_{(w,s) \in b} w R(s, a) + \sum_{o \in O_{b,a}} \hat{p}(o \mid b, a) V(\tau(b, a, o)) \right\}, \end{aligned} \tag{4}$$

where $V(b)$ and $V(b, a)$ can be either upper or lower bound of the optimal value, $b(s) = \sum_i \mathbb{I}(s_i = s) w_i$ is the probability of state s estimated by the weighted particle collection of belief b, and $\mathbb{I}(\cdot)$ is the indicator function. After backup, the upper bounds of ancestor nodes tend to decrease, approaching the optimal value. According to Equation (4), the upper bound of a belief node is only determined by its best action branch, the action branch with the highest upper bound. Therefore, if the best action branch changes for some ancestor nodes during the backup due to the decreasing of the upper bound, expanding the current branch will no longer decrease the upper bound of the root.

In this case, AdaOPS will end the exploration and restart from the root node (Line 8 in Alg. 2). The exploration will also end if the excess uncertainty is less than or equal to 0 since we want to focus on branches with high uncertainty. Moreover, if a belief exceeds the maximum depth D (Line 13 in Alg. 2), its upper bound will be set to its lower bound so as to prevent further searching.

With the process of explorations, the bounds on the root will progressively approach the real one. Nevertheless, the planning time is often not adequate for it to converge. When the allocated time is up, the highest lower bound action at the root belief will be chosen for execution.

3.2 Expanding Leaf Nodes

During the expansion of leaf nodes, AdaOPS first checks if the resampling condition is met (Line 1 in Alg. 3). If it does, AdaOPS performs the KLD-Resampling [15], which means we apply the KLD-Sampling method to the resampling procedure. Then, for each action $a \in \mathcal{A}$, AdaOPS will propagate particles in belief b with a generative model (Line 4 in Alg. 3). This process is further discussed in Section 3.2.1. Particle propagation will produce a set of sampled observations and their corresponding beliefs, $B_{b,a}$, the average reward $R_{b,a}$ for performing action a at belief b, and an estimate $\hat{p}(\cdot \mid b, a)$ for the probability $\Pr(o \mid b, a)$. With the sample observations, AdaOPS generates a belief packing (Line 5 in Alg. 3), $P_{b,a}$, as described in Section 3.2.2. Finally, all beliefs in the belief packing will be expanded to the belief tree, and their upper and lower bounds will be initialized with some value approximation techniques [16], which are outlined in Appendix D.

Algorithm 3 EXPAND(b)

1: **if** $|b|/\mathrm{ESS}(b) > 2$ **then**
2: $\quad b \leftarrow$ RESAMPLE(b)
3: **for** $a \in \mathcal{A}$ **do**
4: $\quad B_{b,a}, R_{b,a}, \hat{p}(\cdot \mid b, a) \leftarrow$ PROPAGATE(b, a)
5: $\quad P_{b,a} \leftarrow$ GENERATEPACKING($B_{b,a}, \hat{p}(\cdot \mid b, a)$)
6: $\quad$ Expand action node a
7: $\quad$ Expand beliefs in $P_{b,a}$ as child beliefs of a

Algorithm 4 RESAMPLE(b)

1: Count the number of bins with support for belief b
2: Calculate the number N satisfying Equation (2)
3: Sample N particles in belief b

Algorithm 5 PROPAGATE(b, a)

1: $O \leftarrow \varnothing, B \leftarrow \varnothing, R \leftarrow 0, \hat{p}(\cdot) \leftarrow 0$
2: Initialize $\tilde{b}$ as an empty weighted particle collection
3: **for** $(w, s) \in b$ **do**
4: $\quad s', o, r \leftarrow G(s, a)$
5: $\quad \tilde{b} \leftarrow \tilde{b} \cup \{(w, s')\}, O \leftarrow O \cup \{o\}, R \leftarrow R + wr$
6: $\quad \hat{p}(o) \leftarrow \hat{p}(o) + w$
7: **for** $o \in O$ **do**
8: $\quad$ Initialize b' as an empty weighted particle collection
9: $\quad$ **for** $(w, s') \in \tilde{b}$ **do**
10: $\quad\quad b' \leftarrow b' \cup \{wZ(o \mid s', a)/\rho, s'\}$ $\quad \triangleright\ \rho = \sum_{(w,s') \in \tilde{b}} wZ(o \mid s', a)$
11: $\quad B \leftarrow B \cup \{(o, b')\}$
12: **return** $B, R, \hat{p}$

3.2.1 Adaptive Particle Filter

According to Equation (4), we need to approximate beliefs accurately so that the estimation for the first integral $\int_{s \in \mathcal{S}} b(s) R(s, a) \mathrm{d}s = \sum_{(w,s) \in b} wR(s, a)$ is accurate. However, in online planning, we cannot afford tons of particles for belief approximation and need a particle filter that works well with merely a few particles. That is why we introduce the adaptive resampling procedure and the KLD-Sampling method. When the sample size is small, the particle impoverishment issue will be prominent, which can be alleviated with adaptive resampling. Moreover, KLD-Sampling ensures the approximation quality while using fewer particles.

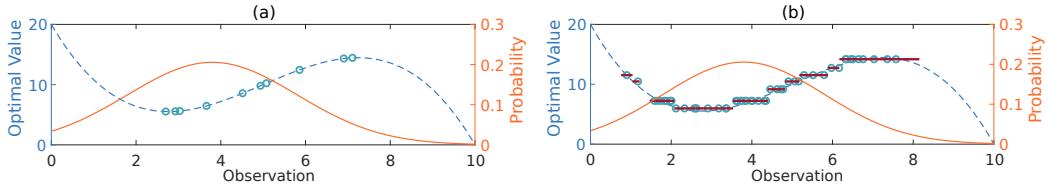

Figure 2: The solid orange line stands for an observation distribution $\Pr(\cdot \mid b, a)$ distributed between 0 to 10. The dashed blue line illustrates the optimal values of observations, $V^*(\tau(b, a, \cdot))$. (a) shows 10 samples from the observation distribution, forming an unbiased estimator for the first integral in Equation (4). Nonetheless, this estimator has a high variance since the sample size is too small. (b) shows 50 samples from the observation distribution. By building a belief packing, only 10 dissimilar beliefs having distinct optimal values are retained for evaluation. Each solid red line in the figure corresponds to a belief in the packing. The values of beliefs on the line are set to the value of the corresponding belief in the packing. By doing so, we obtain a low variance estimator with slight bias at the cost of 10 evaluations.

AdaOPS performs the KLD-Resampling adaptively (Line 1 in Alg. 3) when the particle number of the leaf belief b is greater than twice the effective sample size, $\mathrm{ESS}(b)$. According to Equation (1), a low effective sample size indicates a high level of weight disparity. It is also illustrated in Figure 1 that AdaOPS resamples when the weight disparity of a belief is high. In order to determine the particle number for resampling, the state space is partitioned into a grid. Then, we can count the number of bins occupied by belief b and calculate the particle number for resampling with Equation (2). The hyperparameter η is set to be 0.05, and the ζ is tuned such that $m_{\min}$ is the minimal number satisfying Equation (2) when $k = 2$, i.e., $\zeta = \frac{1}{2m_{\min}}\left(1 - \frac{2}{9} + \sqrt{\frac{2}{9}}z_{0.95}\right)$, where $m_{\min}$ is another hyperparameter representing the minimum number of particles needed for approximating the belief.

For each action $a \in \mathcal{A}$, AdaOPS propagates all particles $(w_i, s_i) \in b$ through a generative model, i.e., $s'_i, o_i, r_i \leftarrow G(s_i, a)$ (Line 4 in Alg. 5) and calculates the average reward, $R = \sum_i w_i r_i$ (Line 5 in Alg. 5). We estimate the probability of generating an observation o, $\Pr(o \mid b, a)$, by the sum of weights of particles generating it (Line 6 in Alg. 5), i.e., $\Pr(o \mid b, a) \approx \hat{p}(o \mid b, a) = \sum_i \mathbb{I}(o = o_i) w_i$.

The belief $\tau(b, a, o)$ for sampled observation o is approximated by the same collection of particles but assigns a different weight for each (Line 10 in Alg. 5). The weight for the i-th particle under observation o is $w'_i = w_i Z(o \mid s'_i, a)/\rho$, where $\rho = \sum_i w_i Z(o \mid s'_i, a)$ is a normalizing factor. The set of sampled observations and their corresponding beliefs is denoted as $B_{b,a}$.

3.2.2 Belief Packing

Equation (4) reveals that we can approximate the integral $\int_{o \in \mathcal{O}} \Pr(o \mid b, a) V(\tau(b, a, o)) \mathrm{d}o$ with the value of some sampled observations. The more observations we sample, the lower the estimation error is. However, each sampled observation corresponds to an observation branch on the tree. If we expand all observations in $B_{b,a}$ as new branches, the induced branching factor is often beyond our computational capability. Nevertheless, restricting the number of observations renders high estimation variance. Inspired by [17], we introduce belief packing to fuse similar beliefs and only evaluate the distinct ones, balancing estimation bias and variance.

Algorithm 6 GENERATEPACKING($B, \hat{p}$)

```
P ← ∅
for (o, b) ∈ B do
    covered ← false
    for (o', b') ∈ P do
        if ||b' − b||_1 ≤ δ then
            p̂(o') ← p̂(o') + p̂(o)
            p̂(o) ← 0
            covered ← true
            break
    if not covered then
        P ← P ∪ {(o, b)}
return P
```

Since the sibling beliefs b and b' share the same particles, we can calculate their $L1$ distance with $\|b - b'\|_1 = \sum_{s \in \mathcal{S}} |b(s) - b(s')| = \sum_i |w_i - w'_i|$, where w_i and w'_i are the weights of the i-th particle for beliefs b and b', respectively. As will be shown in Theorem 1, when the $L1$ distance between two approximated beliefs is close, the true optimal value of one is close to the estimated optimal value of the other with high probability. In other words, it is safe to fuse similar beliefs. The formal convergence analysis is given in Section 4.

In AdaOPS, a belief packing is maintained under each action branch a of belief b as illustrated in Figure 1. We merge two observation branches when their beliefs are within a distance of δ (Line 5 in Alg. 6). Specifically, if b_1 is within a distance δ of a belief b_2 in the packing, we set $\hat{p}(o_1 \mid b, a) \leftarrow 0$ and $\hat{p}(o_2 \mid b, a) \leftarrow \hat{p}(o_1 \mid b, a) + \hat{p}(o_2 \mid b, a)$ for their corresponding observations o_1 and o_2, which is equivalent to set the optimal value of b_2, $V^*(b_2)$, to be the optimal value of b_1, $V^*(b_1)$. If the belief b_2 is not covered by any δ-ball centered at beliefs in the packing, then we add it to the packing (Line 10 in Alg. 6).

As shown in Figure 2, by setting the value of a belief to that of the closest belief in the δ-packing, we obtain a piece-wise constant approximation to $V^*(\tau(b, a, \cdot))$, which avoids evaluating $V^*(\cdot)$ at similar beliefs at the cost of limited bias.

4 Theoretical Analysis

In this section, we demonstrate the convergence of AdaOPS to the optimal policy. The proof, along with a running time analysis, is detailed in Appendix E.

As the resampling will cause the particles to be interrelated and significantly increase the difficulty of analysis, we follow the proof in [18] and do not factor in the adaptive particle filter. The effectiveness of the adaptive particle filter is empirically demonstrated in Section 5.2.

Theorem 1. $\bar{b}_d$ and $\bar{b}'_d$ are two estimated belief points at depth d represented by N particles, and $\|\bar{b}_d - \bar{b}'_d\|_1 \leq \delta$. For a given $\lambda > 0$, let $\eta = \frac{\lambda}{R_{\max} D(1-\gamma)}$. The number of particles N is chosen such that $3|\mathcal{A}|(3|\mathcal{A}|N)^D \exp(-N \cdot t_{\max}) \leq \eta$, where $t_{\max} = \frac{\lambda(1-\gamma)}{3R_{\max} d_\infty^{\max}} - \frac{1}{\sqrt{N}} > 0$, and $d_\infty^{\max}$ is the essential supremum of the Rényi divergences between the sampling and target distributions of all depths $d = 0, 1, \ldots, D-1$. The corresponding true beliefs of $\bar{b}_d$ and $\bar{b}'_d$ are b_d and b'_d. When the assumptions in [18] hold, it follows that

$$|V_d^*(b'_d) - \hat{V}_d^*(\bar{b}_d)| \leq \frac{R_{\max}}{1-\gamma}\delta + \frac{\lambda}{1-\gamma},$$

where $V_d^*(b'_d)$ denotes the optimal value at the belief b'_d, and $\hat{V}_d^*(\bar{b}_d)$ stands for the estimated optimal value at the estimated belief $\bar{b}_d$. ■

Intuitively, Theorem 1 says that if $\bar{b}_d$ and $\bar{b}'_d$ are close enough, we can replace $\bar{b}'_d$ with $\bar{b}_d$, and the error caused in estimating the optimal value will not be too large with a probability of at least $1 - \eta$.

Theorem 2. Suppose that all assumptions in Theorem 1 hold. Let N in Theorem 1 be $m_{\min}$, $\epsilon > 0$, $\delta = \frac{(1-\gamma)^2\epsilon}{3\gamma R_{\max}}$, $\lambda = \frac{(1-\gamma)^2\epsilon}{3\gamma}$, and $D = \log_\gamma \frac{(1-\gamma)\epsilon}{3R_{\max}}$. Then $|V^*(\text{bel}) - V^*_{\text{AdaOPS}}(\text{bel})| \leq \epsilon$ holds with a probability of at least $1 - \eta$, where $V^*_{\text{AdaOPS}}(\text{bel})$ is the estimated optimal value given by AdaOPS. ■

Theorem 2 guarantees that given unlimited time, AdaOPS will terminate obtaining a near-optimal policy with a probability of at least $1 - \eta$.

5 Experiments

This section evaluates our method on four domains based on the POMDPs.jl framework [19] and conducts an ablation study showing the contribution of each component. Besides, we compare the adaptive particle filter with other particle filter variants demonstrating that the adaptive particle filter is able to approximate the belief better.

5.1 Performance Comparison

In what follows, we compare AdaOPS with ARDESPOT [5] and POMCPOW [6] on four domains, Laser Tag [5], Rock Sample (RS) [14], Roomba [20], and Light Dark [6]. These domains are carefully chosen so that they can cover various types of POMDPs. The details of these domains are shown in Appendix F.1. In order to understand the contribution from adaptive resampling (AR), KLD-Sampling (KLD-S), and belief packing (BP), respectively, we remove one component at a time and tune the hyperparameters to achieve the best performance for others. Specifically, the removal of adaptive resampling means we resample at each timestep as a standard particle filter does.

5.1.1 Experiment settings

In all these tasks, solvers are allocated $1s$ for planning per step. Since the obstacle positions in Laser Tag and the rock positions in Rock Sample greatly influence the optimal value, tests on these domains are conducted with one hundred randomly generated maps so that the result is trustworthy. Hyperparameters for each algorithm are tuned via grid search, and another independent experiment is conducted to test their performance. ARDESPOT and AdaOPS initialize upper and lower bounds with heuristics, and POMCPOW evaluates nodes it encountered for the first time with a value estimator. All experiments were conducted on a computer with Intel(R) Core(TM) i7-10750H, 6vCPUs running at 2.6GHz, and 32G main memory. Please refer to Appendix F.2 for the ranges of grid search and the hyperparameters and heuristics selected for each algorithm.

5.1.2 Results

In Table 1, the average discounted return is shown along with the corresponding standard error of mean (SEM) in the form of $\text{Return} \pm \text{SEM}$. In all these domains, AdaOPS outperforms other solvers with statistical significance ($p < 0.0001$). A hyperparameter sensitivity analysis in Light Dark is presented in Appendix F.3.

Table 1: Performance Comparison

	Laser Tag	RS(15,15)	Bumper Roomba	Lidar Roomba	Light Dark
$\|\mathcal{S}\|$	$4,830$	$7,372,800$	∞	∞	∞
$\|\mathcal{A}\|$	5	20	9	6	3
$\|\mathcal{O}\|$	$\sim 1.5 \times 10^6$	3	2	∞	∞
DESPOT	-10.71 ± 0.22	15.67 ± 0.20	-2.18 ± 0.11	-0.84 ± 0.06	2.50 ± 0.10
POMCPOW	-10.15 ± 0.17	10.12 ± 0.19	-2.28 ± 0.09	-0.50 ± 0.07	3.23 ± 0.11
AdaOPS	$\mathbf{-8.31 \pm 0.18}$	$\mathbf{17.16 \pm 0.21}$	$\mathbf{-1.48 \pm 0.11}$	$\mathbf{-0.09 \pm 0.07}$	$\mathbf{3.79 \pm 0.07}$
AdaOPS(AR)	-8.79 ± 0.17	16.73 ± 0.21	-1.63 ± 0.11	-0.21 ± 0.06	3.62 ± 0.07
AdaOPS(KLD-S)	-8.79 ± 0.18	\	-1.71 ± 0.11	-0.20 ± 0.07	3.72 ± 0.07
AdaOPS(BP)	-9.84 ± 0.18	17.10 ± 0.21	\	-0.51 ± 0.07	1.50 ± 0.20

Note: RS(n,m) stands for the Rock Sample with $n \times n$ map and m rocks. ∞ means continuous state (or observation) space.
AdaOPS(XX) denotes AdaOPS without component XX. AR = Adaptive Resampling. KLD-S = KLD-Sampling. BP = Belief Packing.

Laser Tag In Laser Tag, an agent tries to tag an escaping target. Initially, the agent knows nothing about its own position and the target's position, and the belief is spread over the entire grid. It can then access the map and infer the positions from the sensor information. After several steps, the belief will shrink into a small region. Since the dispersion of beliefs varies considerably, the KLD-Sampling is advantageous. The adaptive resampling avoids sample impoverishment and contributes equally with KLD-Sampling. The belief packing contributes most by allowing AdaOPS to handle massive observation space effectively.

Rock Sample In Rock Sample, a Mars rover intends to sample as many good rocks as possible while only inferring the goodness of rocks from its noisy sensor measurements. It maintains a belief of the goodness for each rock, which always is comprised of two bins, good and bad. Thus, the particle number given by Equation (2) is always $m_{\min}$, i.e., KLD-Sampling makes no difference in the Rock Sample problem. Besides, belief packing is also unproductive since there are only two observations. The improvements of AdaOPS mainly come from the fact that AdaOPS maintains an accurate belief approximation in the depth of the tree using particle filter techniques. Contrarily, DESPOT and POMCPOW tend to be overoptimistic on future uncertainty.

Roomba Roomba is a robotic vacuum cleaner that attempts to locate itself in a familiar room and reach the target region. It can equip either a Bumper sensor that senses the collision or a Lidar sensor that measures the distance to the obstacle in the front. As a state is either collided or not, the beliefs corresponding to different observations in Bumper Roomba always have an $L1$ distance 2, which means there are no similar beliefs, and belief packing is of little avail in this case. However, in Lidar Roomba, belief packing handles continuous observation gracefully, making a significant contribution.

The KLD-Sampling fails to contribute much as expected in the Roomba domain. A possible reason is that according to Equation (2), high state uncertainty comes with a large number of particles and, consequently, observations, making online planning extremely hard given limited planning time. The above analysis is corroborated by an additional result presented in Appendix F.4.

Light Dark Light Dark is an information-gathering task. With the aim of reaching the target area, the agent needs to find its way to the light region, gathering intelligence. This requires the agent to be fully aware of the state uncertainty while factoring in all possible observations. The belief packing empowers AdaOPS to consider only distinct observations. On the KLD-Sampling, the problem in Roomba also occurs in Light Dark, resulting in little improvement.

5.2 Benefits of Adaptive Particle Filters

We present a comparison between different particle filters in this section. The sequential importance resampling (SIR) particle filter is the standard particle filter that resamples at each timestep. The belief updating in POMCPOW can be considered as a SIR particle filtering biased towards early sampled particles. The standard particle filter without the resampling procedure becomes a sequential importance sampling (SIS) particle filter, which is adopted in DESPOT-α [7]. DESPOT utilizes rejection sampling for belief updating, which approximates a belief with particles accompanying conforming observations. This method is doomed to fail in large or continuous observation space because the same observation is unlikely to be produced twice.

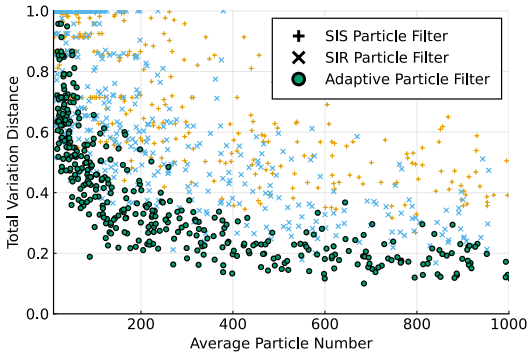

Figure 3: The belief approximated via the adaptive particle filter is the closest to the real posterior, using the same amount of particles.

This experiment is conducted on Laser Tag. Each particle filter is given the same sequence of action-observation history for belief updating. The approximation error (measured by total variation distance) is checked after 10 steps. As illustrated in Figure 3, the adaptive particle filter gets the lowest approximation error and the lowest variance using the same amount of particles. This result supplements the convergence analysis in Section 4 since a better approximation of beliefs will not undermine the convergence to the optimal policy.

6 Discussion

This paper presents an online planning algorithm, AdaOPS, featuring adaptive particle filter and belief packing, which is able to solve most POMDP problems optimally, given enough planning time. The success of the adaptive particle filter points out that in previous methods [5, 6], the inaccurate belief updating is a detriment to the performance, outweighing the potential benefit of efficiency. Furthermore, it is demonstrated that forming a belief packing with approximated beliefs is an effective tool for handling large observation spaces. We believe it could find itself extensive use in POMDP planning.

During planning, AdaOPS builds a tree of size $O((|\mathcal{A}|P^{\delta}_{\max})^D)$, where $P^{\delta}_{\max} = \sup_{b,a} P^{\delta}(Y_{b,a})$, and $Y_{b,a}$ is the set of reachable beliefs after executing action a at belief b. However, the tree size goes unmanageable when $|\mathcal{A}|$ and $P^{\delta}_{\max}$ are large. Though it is caused by the complexity of the problem itself [9], we can abate this problem by implementing a parallel version following [21]. Furthermore, it is possible to improve the search efficiency with sawtooth approximation [16] and alpha vectors as described in [7]. We will also explore the possibility of handling large-scale or continuous action spaces.

References

[1] L. P. Kaelbling, M. L. Littman, and A. R. Cassandra, "Planning and Acting in Partially Observable Stochastic Domains," *Artificial Intelligence*, vol. 101, no. 1-2, pp. 99–134, 1998.

[2] M. J. Kochenderfer, *Decision Making Under Uncertainty: Theory and Application*. MIT press, 2015.

[3] C. H. Papadimitriou and J. N. Tsitsiklis, "The Complexity of Markov Decision Processes," *Mathematics of Operations Research*, vol. 12, no. 3, pp. 441–450, 1987.

[4] D. Silver and J. Veness, "Monte-Carlo Planning in Large POMDPs," in *NIPS*, pp. 2164–2172, 2010.

[5] N. Ye, A. Somani, D. Hsu, and W. S. Lee, "DESPOT: Online POMDP Planning with Regularization," *Journal of Artificial Intelligence Research*, vol. 58, pp. 231–266, 2017.

[6] Z. N. Sunberg and M. J. Kochenderfer, "Online Algorithms for POMDPs with Continuous State, Action, and Observation Spaces," in *ICAPS*, pp. 259–263, 2018.

[7] N. Priyadarshini Garg, D. Hsu, and W. Sun Lee, "DESPOT-Alpha: Online POMDP Planning with Large State and Observation Spaces," in *RSS*, 2019.

[8] A. Doucet and A. M. Johansen, "A Tutorial on Particle Filtering and Smoothing: Fifteen Years Later," *Handbook of Nonlinear Filtering*, vol. 12, no. 3, pp. 656–704, 2009.

[9] D. Hsu, W. S. Lee, and N. Rong, "What Makes Some POMDP Problems Easy to Approximate?," in *NIPS*, pp. 689–696, 2007.

[10] P. Del Moral, A. Doucet, and A. Jasra, "On Adaptive Resampling Procedures for Sequential Monte Carlo Methods," *Bernoulli*, vol. 18, no. 1, pp. 252–278, 2012.

[11] L. Kish, *Survey Sampling*. John Wiley & Sons Inc, 1965.

[12] J. S. Liu, *Monte Carlo Strategies in Scientific Computing*. Springer Science & Business Media, 2008.

[13] D. Fox, "KLD-Sampling: Adaptive Particle Filters," in *NIPS*, pp. 713–720, 2001.

[14] T. Smith and R. Simmons, "Heuristic Search Value Iteration for POMDPs," in *UAI*, pp. 520–527, 2004.

[15] T. Li, S. Sun, and T. P. Sattar, "Adapting Sample Size in Particle Filters through KLD-Resampling," *Electronics Letters*, vol. 49, pp. 740–742, June 2013.

[16] M. Hauskrecht, "Value-Function Approximations for Partially Observable Markov Decision Processes," *Journal of Artificial Intelligence Research*, vol. 13, pp. 33–94, 2000.

[17] Z. Zhang, D. Hsu, and W. S. Lee, "Covering Number for Efficient Heuristic-based POMDP Planning," in *ICML*, pp. 48–60, 2014.

[18] M. H. Lim, C. Tomlin, and Z. N. Sunberg, "Sparse Tree Search Optimality Guarantees in POMDPs with Continuous Observation Spaces," in *IJCAI*, pp. 4135–4142, 2020.

[19] M. Egorov, Z. N. Sunberg, E. Balaban, T. A. Wheeler, J. K. Gupta, and M. J. Kochenderfer, "POMDPs.jl: A Framework for Sequential Decision Making under Uncertainty," *Journal of Machine Learning Research*, vol. 18, no. 26, pp. 1–5, 2017.

[20] K. Menda, Z. Sunberg, and M. Kochenderfer, "RoombaPOMDPs.jl." https://github.com/sisl/RoombaPOMDPs.jl, 2021.

[21] P. Cai, Y. Luo, D. Hsu, and W. S. Lee, "HyP-DESPOT: A Hybrid Parallel Algorithm for Online Planning under Uncertainty," in *RSS*, 2018.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes] Section 6
   (c) Did you discuss any potential negative societal impacts of your work? [No] We see no potential negative societal impacts.
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes]
   (b) Did you include complete proofs of all theoretical results? [Yes]
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes]
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes]
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes]
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes]
   (b) Did you mention the license of the assets? [Yes]
   (c) Did you include any new assets either in the supplemental material or as a URL? [Yes] RoombaPOMDPs.jl is included as a URL.
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [Yes]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A] We do not use any data in this paper.
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]