

# Sparse Tree Search Optimality Guarantees in POMDPs with Continuous Observation Spaces

Michael H. Lim<sup>1</sup>, Claire J. Tomlin<sup>1</sup> and Zachary N. Sunberg<sup>2</sup>

<sup>1</sup>University of California, Berkeley

<sup>2</sup>University of Colorado, Boulder

michaelhlim@berkeley.edu, tomlin@eecs.berkeley.edu, zachary.sunberg@colorado.edu

## Abstract

Partially observable Markov decision processes (POMDPs) with continuous state and observation spaces have powerful flexibility for representing real-world decision and control problems but are notoriously difficult to solve. Recent online sampling-based algorithms that use observation likelihood weighting have shown unprecedented effectiveness in domains with continuous observation spaces. However there has been no formal theoretical justification for this technique. This work offers such a justification, proving that a simplified algorithm, partially observable weighted sparse sampling (POWSS), will estimate Q-values accurately with high probability and can be made to perform arbitrarily near the optimal solution by increasing computational power.

## 1 Introduction

The partially observable Markov decision process (POMDP) is a flexible mathematical framework for representing sequential decision problems where knowledge of the state is incomplete [Kaelbling *et al.*, 1998; Kochenderfer, 2015; Bertsekas, 2005]. The POMDP formalism can represent a wide range of real world problems including autonomous driving [Sunberg *et al.*, 2017; Bai *et al.*, 2015], cancer screening [Ayer *et al.*, 2012], spoken dialog systems [Young *et al.*, 2013], and others [Cassandra, 1998]. In one of the most successful applications, an approximate POMDP solution is being used in a new aircraft collision avoidance system that will be deployed worldwide [Holland *et al.*, 2013].

A POMDP is an optimization problem for which the goal is to find a policy that specifies actions that will control the state to maximize the expectation of a reward function. One of the most popular ways to deal with the challenging computational complexity of finding such a policy [Papadimitriou and Tsitsiklis, 1987] is to use online tree search algorithms [Silver and Veness, 2010; Ye *et al.*, 2017; Sunberg *et al.*, 2017; Kurniawati and Yadav, 2016]. Instead of attempting to find a global policy that specifies actions for every possible outcome of the problem, *online* algorithms look for local approximate solutions as the agent is interacting with the environment.




Figure 1: Trees generated from partially observable sparse sampling (POSS) algorithm (left), and partially observable weighted sparse sampling (POWSS) algorithm (right) with depth  $D = 2$  and width  $C = 2$ , for a continuous-observation POMDP. Nodes below the fading edges are omitted for clarity. Square nodes correspond to actions, filled circles to state particles with size representing weight, and unfilled circles to beliefs.

Previous online approaches such as ABT [Kurniawati and Yadav, 2016], POMCP [Silver and Veness, 2010], and DESPOT [Ye *et al.*, 2017] have exhibited good performance in large discrete domains. However, many real-world domains, notably when a robot interacts with the physical world, have continuous observation spaces, and the algorithms mentioned above will not always converge to an optimal policy in problems with continuous or naively discretized observation spaces [Sunberg and Kochenderfer, 2018].

Two recent approaches, POMCPOW [Sunberg and Kochenderfer, 2018] and DESPOT- $\alpha$  [Garg *et al.*, 2019], have employed a weighting scheme inspired by particle filtering to achieve good performance on realistic problems with large or continuous observation spaces. However, there are currently no theoretical guarantees that these algorithms will find optimal solutions in the limit of infinite computational resources.

A convergence proof for these algorithms must have the following two components: (1) A proof that the particle weighting scheme is sound, and (2) a proof that the heuristics used to focus search on important parts of the tree are sound. This paper tackles the first component by analyzing a new simplified algorithm that expands every node of a sparse search tree.

First, we naively extend the sparse sampling algorithm of Kearns *et al.* [2002] to the partially observable domain in the spirit of POMCP and explain why this algorithm, known as

partially observable sparse sampling (POSS), will converge to a suboptimal solution when the observation space is continuous. Then, we introduce appropriate weighting that results in the partially observable weighted sparse sampling (POWSS) algorithm. We prove that the value function estimated by POWSS converges to the optimal value function at a rate of  $\mathcal{O}(C^D \exp(-t \cdot C))$ , where  $C$  is the planning width and number of particles,  $D$  is the depth, and  $t$  is a constant specific to the problem and desired accuracy. This yields a policy that can be made arbitrarily close to optimal by increasing the computation.

To our knowledge, POWSS is the first algorithm proven to converge to a globally optimal policy for POMDPs with continuous observation spaces without relying on any discretization schemes. Since POWSS fully expands all nodes in the sparse tree, it is not computationally efficient and is only practically applicable to toy problems. However, the convergence guarantees justify the weighting schemes in state-of-the-art efficient algorithms like DESPOT- $\alpha$  and POMCPOW that solve realistic problems by only constructing the most important parts of the search tree.

The remainder of this paper proceeds as follows: First, Sections 2 and 3 review preliminary definitions and previous work. Then, Section 4 presents an overview of the POSS and POWSS algorithms. Next, Section 5 contains an importance sampling result used in subsequent sections. Section 6 contains the main contribution, a proof that POWSS converges to an optimal policy using induction from the leaves to the root to prove that the value function estimate will eventually be accurate with high probability at all nodes. Finally, Section 7 empirically shows convergence of POWSS on a modified tiger problem [Kaelbling *et al.*, 1998].

## 2 Preliminaries

### POMDP Formulation

A POMDP is defined by a 7-tuple  $(S, A, O, \mathcal{T}, \mathcal{Z}, R, \gamma)$ :  $S$  is the state space,  $A$  is the action space,  $O$  is the observation space,  $\mathcal{T}$  is the transition density  $\mathcal{T}(s'|s, a)$ ,  $\mathcal{Z}$  is the observation density  $\mathcal{Z}(o|a, s')$ ,  $R$  is the reward function, and  $\gamma \in [0, 1)$  is the discount factor [Kochenderfer, 2015; Bertsekas, 2005]. Since a POMDP agent receives only observations, the agent infers the state by maintaining a belief  $b_t$  at each step  $t$  and updating it with new action and observation pair  $(a_{t+1}, o_{t+1})$  via Bayesian updates [Kaelbling *et al.*, 1998]. A policy, denoted with  $\pi$ , maps beliefs  $b_t$  generated from histories  $h_t = (b_0, a_1, o_1, \dots, a_t, o_t)$  to actions  $a_t$ . Thus, to maximize the expected cumulative reward in POMDPs, the agent wants to find the optimal policy  $\pi^*(b_t)$ .

We solve the finite-horizon problem of horizon length  $D$ . We formulate the state value function  $V$  and action value function  $Q$  for a given belief state  $b$  and policy  $\pi$  at step  $t$  by Bellman updates for  $t \in [0, D - 1]$ , where  $ba_o$  indicates the belief  $b$  updated with  $(a, o)$ :

$$V_t^\pi(b) = \mathbb{E} \left[ \sum_{i=t}^{D-1} \gamma^{i-t} R(s_i, \pi(s_i)) \middle| b \right], \quad V_D^\pi(b) = 0 \quad (1)$$

$$Q_t^\pi(b, a) = \mathbb{E} [R(s, a) + \gamma V_{t+1}^\pi(ba_o) | b] \quad (2)$$

Specifically, the optimal value functions satisfy the following:

$$V_t^*(b) = \max_{a \in A} Q_t^*(b, a) \quad (3)$$

$$\pi_t^*(b) = \arg \max_{a \in A} Q_t^*(b, a) \quad (4)$$

$$Q_t^*(b, a) = \mathbb{E} [R(s, a) + \gamma V_{t+1}^*(ba_o) | b] \quad (5)$$

**Generative models.** For many problems, the probability densities  $\mathcal{T}$  and  $\mathcal{Z}$  may be difficult to determine explicitly. Thus, some approaches only require that samples are generated according to the correct probability. In this case, a generative model  $G$  implicitly defines  $\mathcal{T}, \mathcal{Z}$  by generating a new state,  $s'$ , observation,  $o$ , and reward,  $r$ , given the current state  $s$  and action  $a$ .

**Probability notation.** We denote probability measures with calligraphic letters (e.g.  $\mathcal{P}, \mathcal{Q}$ ) to avoid confusion with the action value function  $Q(b, a)$ . Furthermore, for two probability measures  $\mathcal{P}, \mathcal{Q}$  defined on a  $\sigma$ -algebra  $\mathcal{F}$ , we denote  $\mathcal{P} \ll \mathcal{Q}$  to state that  $\mathcal{P}$  is absolutely continuous with respect to  $\mathcal{Q}$ ; for every measurable set  $A$ ,  $\mathcal{Q}(A) = 0$  implies that  $\mathcal{P}(A) = 0$ . Also, we use the abbreviations ‘‘a.s.’’ for almost surely, and ‘‘i.i.d.r.v.’’ for independent and identically distributed random variables.

## 3 Additional Related Work

In addition to the work mentioned in the introduction, there has been much work in similar areas. There are several on-line tree search techniques for fully observable Markov decision processes with continuous state spaces, most prominently Sparse-UCT [Bjarnason *et al.*, 2009], and double progressive widening [Couëtoux *et al.*, 2011].

There are also several approaches for solving POMDPs or belief-space MDPs with continuous observation spaces. For example, Monte Carlo Value Iteration (MCVI) can use a classifier to deal with continuous observation spaces [Bai *et al.*, 2014]. Others partition the observation space [Hoey and Poupart, 2005] or assume that the most likely observation is always received [Platt *et al.*, 2010]. Other approaches are based on motion planning [Bry and Roy, 2011; Agha-Mohammadi *et al.*, 2011], locally optimizing pre-computed trajectories [Van Den Berg *et al.*, 2012], or optimizing open-loop plans [Sunberg *et al.*, 2013]. McAllester and Singh [1999] also extend the sparse sampling algorithm of Kearns *et al.* [2002], but they use a belief simplification scheme instead of the particle sampling scheme.

## 4 Algorithms

We first define the algorithmic elements shared by POSS and POWSS, SELECTACTION and ESTIMATEV, in Algorithm 1. SELECTACTION is the entry point of the algorithm, which selects the best action for a belief  $b_0$  according to the  $Q$ -function by recursively calling ESTIMATEQ. ESTIMATEV is a subroutine that returns the value,  $V$ , for an estimated belief, by calling ESTIMATEQ for each action and returning the maximum. We use belief particle set  $\bar{b}$  at every step  $d$ , which contain pairs  $(s_i, w_i)$  that correspond to the generated sample and its corresponding weight. The weight at initial step is uniformly normalized to  $1/C$ , as the samples are drawn directly from  $b_0$ .

---

**Algorithm 1** Routines common to POWSS and POSS
 

---

**Algorithm:** SelectAction( $b_0, \gamma, G, C, D$ )

**Input:** Belief state  $b_0$ , discount  $\gamma$ , generative model  $G$ , width  $C$ , max depth  $D$ .

**Output:** An action  $a^*$ .

- 1: From the initial belief distribution  $b_0$ , sample  $C$  particles and store it in  $\bar{b}_0$ . For POWSS, weights are also initialized with  $w_i = 1/C$ .
- 2: For each of the actions  $a \in A$ , calculate:

$$\hat{Q}_0^*(\bar{b}_0, a) = \text{ESTIMATEQ}(\bar{b}_0, a, 0)$$

- 3: Return  $a^* = \arg \max_{a \in A} \hat{Q}_0^*(\bar{b}_0, a)$

**Algorithm:** EstimateV( $\bar{b}, d$ )

**Input:** Belief particles  $\bar{b}$ , current depth  $d$ .

**Output:** A scalar  $\hat{V}_d^*(\bar{b})$  that is an estimate of  $V_d^*(b)$ .

- 1: If  $d \geq D$  the max depth, then return 0.
- 2: For each of the actions  $a \in A$ , calculate:

$$\hat{Q}_d^*(\bar{b}, a) = \text{ESTIMATEQ}(\bar{b}, a, d)$$

- 3: Return  $\hat{V}_d^*(\bar{b}) = \max_{a \in A} \hat{Q}_d^*(\bar{b}, a)$
- 

In Algorithms 1 to 3, we omit  $\gamma, G, C, D$  in the subsequent recursive calls for convenience since they are fixed globally.

We define ESTIMATEQ functions in Algorithm 2 for POSS and Algorithm 3 for POWSS, where both methods perform sampling and recursive calls to ESTIMATEV to estimate the  $Q$ -function at a given step. The crucial difference between these algorithms is shown in Fig. 1.

POSS naively samples the next  $s'_i, o_i, r_i$  via the generating function for each state  $s_i$  in the belief particle set  $\bar{b}$ , at a given step  $d$ . Then, for each unique observation  $o_j$  generated from the sampling step, POSS inserts the states  $s'_i$  into the next-step belief particle set  $\overline{bao_j}$  only if the generated observation  $o_i$  matches  $o_j$ . This behavior is similar to POMCP, DESPOT, or a particle filter that uses rejection and can quickly lead to particle depletion when there are many unique observations. Finally, POSS returns the naive average of the  $Q$ -functions calculated via recursive calculation of ESTIMATEV for each of the next-step beliefs.

On the other hand, POWSS uses particle weighting rather than using only unweighted particles with matching observation histories as in POSS. POWSS samples the next  $s'_i, o_i, r_i$  via the generating function for each state-weight pair  $(s_i, w_i)$  in the belief particle set  $\bar{b}$ . Now, for each observation  $o_j$  generated from the sampling step, POWSS inserts all the states  $s'_i$  and the new weights  $w'_i = w_i \cdot Z(o_j | a, s'_i)$  into the next-step belief particle set  $\overline{bao_j}$ . These weights are the adjusted probability of hypothetically obtaining  $o_j$  from state  $s'_i$ . POWSS then returns the weighted average of the  $Q$ -functions.

## 5 Importance Sampling

We begin the theoretical portion of this work by stating important properties about self-normalized importance sampling estimators (SN estimators). One goal of importance sampling is to estimate an expected value of a function  $f(x)$

---

**Algorithm 2** POSS
 

---

**Algorithm:** EstimateQ( $\bar{b}, a, d$ )

**Input:** Belief particles  $\bar{b}$ , action  $a$ , current depth  $d$ .

**Output:** A scalar  $\hat{Q}_d^*(\bar{b}, a)$  that is an estimate of  $Q_d^*(b, a)$ .

- 1: For each particle  $s_i$  in  $\bar{b}$ , generate  $s'_i, o_i, r_i = G(s_i, a)$ . If  $i > |\bar{b}|$ , use  $s_{i \bmod |\bar{b}|}$ .
- 2: For each unique observation  $o_j$  from previous step, insert all  $s'_i$  that satisfy  $o_i = o_j$  to a new belief particle set  $\overline{bao_j}$ .
- 3: Return

$$\hat{Q}_d^*(\bar{b}, a) = \frac{1}{C} \sum_{i=1}^C (r_i + \gamma \cdot \text{ESTIMATEV}(\overline{bao_i}, d + 1))$$


---

---

**Algorithm 3** POWSS
 

---

**Algorithm:** EstimateQ( $\bar{b}, a, d$ )

**Input:** Belief particles  $\bar{b}$ , action  $a$ , current depth  $d$ .

**Output:** A scalar  $\hat{Q}_d^*(\bar{b}, a)$  that is an estimate of  $Q_d^*(b, a)$ .

- 1: For each particle-weight pair  $(s_i, w_i)$  in  $\bar{b}$ , generate  $s'_i, o_i, r_i$  from  $G(s_i, a)$ .
- 2: For each observation  $o_j$  from previous step, iterate over  $i = \{1, \dots, C\}$  to insert  $(s'_i, w_i \cdot Z(o_j | a, s'_i))$  to a new belief particle set  $\overline{bao_j}$ .
- 3: Return

$$\hat{Q}_d^*(\bar{b}, a) = \frac{\sum_{i=1}^C w_i (r_i + \gamma \cdot \text{ESTIMATEV}(\overline{bao_i}, d + 1))}{\sum_{i=1}^C w_i}$$


---

where  $x$  is drawn from distribution  $\mathcal{P}$ , while the estimator only has access to another distribution  $\mathcal{Q}$  along with the importance weights  $w_{\mathcal{P}/\mathcal{Q}}(x) \propto \mathcal{P}(x)/\mathcal{Q}(x)$ . This is crucial for POWSS because we wish to estimate the reward for beliefs conditioned on observation sequences, while only being able to generate the marginal distribution of states with correct probability for an action sequence.

We define the following quantities:

$$\tilde{w}_{\mathcal{P}/\mathcal{Q}}(x) \equiv \frac{w_{\mathcal{P}/\mathcal{Q}}(x)}{\sum_{i=1}^N w_{\mathcal{P}/\mathcal{Q}}(x_i)} \quad (\text{SN Importance Weight})$$

$$d_\alpha(\mathcal{P} \parallel \mathcal{Q}) \equiv \mathbb{E}_{x \sim \mathcal{Q}} [w_{\mathcal{P}/\mathcal{Q}}(x)^\alpha] \quad (\text{Rényi Divergence})$$

$$\tilde{\mu}_{\mathcal{P}/\mathcal{Q}} \equiv \sum_{i=1}^N \tilde{w}_{\mathcal{P}/\mathcal{Q}}(x_i) f(x_i) \quad (\text{SN Estimator})$$

**Theorem 1** (SN  $d_\infty$ -Concentration Bound). *Let  $\mathcal{P}$  and  $\mathcal{Q}$  be two probability measures on the measurable space  $(\mathcal{X}, \mathcal{F})$  with  $\mathcal{P} \ll \mathcal{Q}$  and  $d_\infty(\mathcal{P} \parallel \mathcal{Q}) < +\infty$ . Let  $x_1, \dots, x_N$  be i.i.d.r.v. sampled from  $\mathcal{Q}$ , and  $f : \mathcal{X} \rightarrow \mathbb{R}$  be a bounded Borel function ( $\|f\|_\infty < +\infty$ ). Then, for any  $\lambda > 0$  and  $N$  large enough such that  $\lambda > \|f\|_\infty d_\infty(\mathcal{P} \parallel \mathcal{Q}) / \sqrt{N}$ , the following bound holds with probability at least  $1 - 3 \exp(-N \cdot t^2(\lambda, N))$ :*

$$|\mathbb{E}_{x \sim \mathcal{P}}[f(x)] - \tilde{\mu}_{\mathcal{P}/\mathcal{Q}}| \leq \lambda \quad (6)$$

where  $t(\lambda, N)$  is defined as:

$$t(\lambda, N) \equiv \frac{\lambda}{\|f\|_\infty d_\infty(\mathcal{P}||\mathcal{Q})} - \frac{1}{\sqrt{N}} \quad (7)$$

Theorem 1 builds upon the derivation in Proposition D.3 of Metelli *et al.* [2018], which provides a polynomially decaying bound by assuming  $d_2$  exists. Here, we compromise by further assuming that the infinite Rényi divergence  $d_\infty$  exists and is bounded to get an exponentially decaying bound:  $d_\infty(\mathcal{P}||\mathcal{Q}) = \text{ess sup}_{x \sim \mathcal{Q}} w_{\mathcal{P}/\mathcal{Q}}(x) < +\infty$ . The proof of Theorem 1 is given in Appendix A.<sup>1</sup>

This exponential decay is important for the proofs in Section 6.2. We need to ensure that all nodes of the POWSS tree at all depths  $d$  reach convergence. The branching of the tree induces a factor proportional to  $N^D$ . To offset this, we need a probabilistic bound at each depth that decays exponentially with  $N$ . Intuitive explanation of the  $d_\infty$  assumption is given at the beginning of Section 6.2.

## 6 Convergence

### 6.1 POSS Convergence to QMDP

We present a short informal argument for the convergence of the  $Q$ -value estimates of POSS to the QMDP value (Definition 1) in continuous observation spaces. Sunberg and Kochenderfer [2018] provide a formal proof for a similar algorithm.

**Definition 1** (QMDP value). *Let  $Q_{\text{MDP}}(s, a)$  denote the optimal  $Q$ -function evaluated at state  $s$  and action  $a$  for the fully observable MDP relaxation of a POMDP. Then, the QMDP value at belief  $b$ ,  $Q_{\text{MDP}}(b, a)$ , is  $\mathbb{E}_{s \sim b} [Q_{\text{MDP}}(s, a)]$ .*

Since the observations  $o_i$  are drawn from a continuous distribution, the probability of obtaining duplicate  $o_i$  values in ESTIMATEQ, line 1 is 0. Consequently, when evaluating ESTIMATEQ, all the belief particle sets after the root node only contain a single state particle each (Fig. 1, left), which means that each belief node is merely an alias for a unique state particle. Therefore, ESTIMATEV performs a rollout exactly as if the current state became entirely known after taking a single action, identical to the QMDP approximation. Since QMDP is sometimes suboptimal [Kaelbling *et al.*, 1998], POSS is suboptimal for some continuous-observation POMDPs.

### 6.2 Near-Optimality of POWSS

On the other hand, we claim that the POWSS algorithm can be made to perform arbitrarily close to the optimal policy by increasing the width  $C$ .

In analyzing near-optimality of POWSS, we view POWSS  $Q$ -function estimates as SN estimators, and we apply the concentration inequality result from Theorem 1 to show that POWSS estimates at every node have small errors with high probability. Through the near-optimality of the  $Q$ -functions, we conclude that the value obtained by employing POWSS policy is also near-optimal with further assumptions on the closed-loop POMDP system.

<sup>1</sup>Appendices are available in the pre-print version of this paper: <https://arxiv.org/abs/1910.04332>.

### Assumptions for Analyzing POWSS

The following assumptions are needed for the proof:

- (i)  $S$  and  $O$  are continuous spaces, and the action space has a finite number of elements,  $|A| < +\infty$ .
- (ii) For any observation sequence  $\{o_n\}_j$ , the densities  $\mathcal{Z}, \mathcal{T}, b_0$  are chosen such that the Rényi divergence of the target distribution  $\mathcal{P}^d$  and sampling distribution  $\mathcal{Q}^d$  (Eqs. (20) and (21)) is bounded above by  $d_\infty^{\text{max}} < +\infty$  a.s. for all  $d = 0, \dots, D - 1$ :

$$d_\infty(\mathcal{P}^d||\mathcal{Q}^d) = \text{ess sup}_{x \sim \mathcal{Q}^d} w_{\mathcal{P}^d/\mathcal{Q}^d}(x) \leq d_\infty^{\text{max}}$$

- (iii) The reward function  $R$  is Borel and bounded by a finite constant  $\|R\|_\infty \leq R_{\text{max}} < +\infty$  a.s., and  $V_{\text{max}} \equiv \frac{R_{\text{max}}}{1-\gamma} < +\infty$ .
- (iv) We can sample from the generating function  $G$  and evaluate the observation probability density  $\mathcal{Z}$ .
- (v) The POMDP terminates after  $D < \infty$  steps.

Intuitively, condition (ii) means that the ratio of the conditional observation probability to the marginal observation probability cannot be too high. Additionally, our results still hold even when either of  $S, O$  are discrete, as long as it doesn't violate condition (ii), by appropriately switching the integrals to Riemann sums.

While we restrict our analysis to the case when  $\gamma < 1$  for a finite horizon problem, the authors believe that similar results can be derived for either when  $\gamma = 1$  or when dealing with infinite horizon problems.

**Theorem 2** (Accuracy of POWSS Q-Value Estimates). *Suppose conditions (i)-(v) are satisfied. Then, for a given  $\epsilon > 0$ , choosing constants  $C, \lambda, \delta$  that satisfy:*

$$\lambda = \epsilon(1 - \gamma)^2/5, \quad \delta = \lambda/(V_{\text{max}}D(1 - \gamma)^2) \quad (8)$$

$$\delta \geq 3|A|(3|A|C)^D \exp(-C \cdot t_{\text{max}}^2) \quad (9)$$

$$t_{\text{max}}(\lambda, C) = \frac{\lambda}{3V_{\text{max}}d_\infty^{\text{max}}} - \frac{1}{\sqrt{C}} > 0 \quad (10)$$

*The  $Q$ -function estimates obtained for all depths  $d = 0, \dots, D - 1$  and all actions  $a$  are near-optimal with probability at least  $1 - \delta$ :*

$$\left| Q_d^*(b_d, a) - \hat{Q}_d^*(\bar{b}_d, a) \right| \leq \frac{\lambda}{1 - \gamma} \quad (11)$$

**Theorem 3** (POWSS Policy Convergence). *In addition to conditions (i)-(v), assume that the closed-loop POMDP Bayesian belief update step is exact. Then, for any  $\epsilon > 0$ , we can choose a  $C$  such that the value obtained by POWSS is within  $\epsilon$  of the optimal value function at  $b_0$  a.s.:*

$$V^*(b_0) - V^{\text{POWSS}}(b_0) \leq \epsilon \quad (12)$$

Theorems 2 and 3 are proven sequentially in the following subsections. We generally follow the proof strategy of Kearns *et al.* [2002] but with significant additions to account for the belief-based POMDP calculations rather than state-based MDP calculations. We use induction to prove a concentration inequality for the value function at all nodes in the tree, starting at the leaves and proceeding up to the root.

### Value Convergence at Leaf Nodes

First, we reason about the convergence at nodes at depth  $D-1$  (leaf nodes). In the subsequent analysis, we abbreviate some terms of interest with the following notation:

$$\mathcal{T}_{1:d}^i \equiv \prod_{n=1}^d \mathcal{T}(s_{n,i}|s_{n-1,i}, a_n) \quad (13)$$

$$\mathcal{Z}_{1:d}^{i,j} \equiv \prod_{n=1}^d \mathcal{Z}(o_{n,j}|a_n, s_{n,i})$$

Here  $d$  denotes the depth,  $i$  denotes the index of the state sample, and  $j$  denotes the index of the observation sample. Absence of indices  $i, j$  means that  $\{s_n\}$  and/or  $\{o_n\}$  appear as regular variables. Intuitively,  $\mathcal{T}_{1:d}^i$  is the transition density of state sequence  $i$  from the root node to depth  $d$ , and  $\mathcal{Z}_{1:d}^{i,j}$  is the conditional density of observation sequence  $j$  given state sequence  $i$  from the root node to depth  $d$ . Additionally,  $b_d^i$  denotes  $b_d(s_{d,i})$ ,  $r_{d,i}$  the reward  $R(s_{d,i}, a_d)$ , and  $w_{d,i}$  the weight of  $s_{d,i}$ .

Since the problem ends after  $D$  steps, the  $Q$ -function for nodes at depth  $D-1$  is simply the expectation of final reward and the POWSS estimate has the following form:

$$Q_{D-1}^*(b_{D-1}, a) = \int_S R(s_{D-1}, a) b_{D-1} ds_{D-1} \quad (14)$$

$$\hat{Q}_{D-1}^*(\bar{b}_{D-1}, a) = \frac{\sum_{i=1}^C w_{D-1,i} r_{D-1,i}}{\sum_{i=1}^C w_{D-1,i}} \quad (15)$$

**Lemma 1** (SN Estimator Leaf Node Convergence).

$\hat{Q}_{D-1}^*(\bar{b}_{D-1}, a)$  is an SN estimator of  $Q_{D-1}^*(b_{D-1}, a)$ , and the following leaf-node concentration bound holds with probability at least  $1 - 3 \exp(-C \cdot t_{\max}^2(\lambda, C))$ ,

$$|Q_{D-1}^*(b_{D-1}, a) - \hat{Q}_{D-1}^*(\bar{b}_{D-1}, a)| \leq \lambda \quad (16)$$

*Proof.* First, we show that  $\hat{Q}_{D-1}^*(\bar{b}_{D-1}, a)$  is an SN estimator of  $Q_{D-1}^*(b_{D-1}, a)$ . By following the recursive belief update, the belief term can be fully expanded:

$$b_{D-1}(s_{D-1}) = \frac{\int_{S^{D-1}} (\mathcal{Z}_{1:D-1})(\mathcal{T}_{1:D-1}) b_0 ds_{0:D-2}}{\int_{S^D} (\mathcal{Z}_{1:D-1})(\mathcal{T}_{1:D-1}) b_0 ds_{0:D-1}} \quad (17)$$

Then,  $Q_{D-1}^*(b_{D-1}, a)$  is equal to the following:

$$\begin{aligned} Q_{D-1}^*(b_{D-1}, a) &= \int_S R(s_{D-1}, a) b_{D-1} ds_{D-1} \quad (18) \\ &= \frac{\int_{S^D} R(s_{D-1}, a) (\mathcal{Z}_{1:D-1})(\mathcal{T}_{1:D-1}) b_0 ds_{0:D-1}}{\int_{S^D} (\mathcal{Z}_{1:D-1})(\mathcal{T}_{1:D-1}) b_0 ds_{0:D-1}} \quad (19) \end{aligned}$$

We approximate the  $Q^*$  function with importance sampling by utilizing problem requirement (iv), where the target density is  $b_{D-1}$ . First, we sample the sequences  $\{s_n\}_i$  according to the joint probability  $(\mathcal{T}_{1:D-1}) b_0$ . Afterwards, we weight the sequences by the corresponding observation density  $\mathcal{Z}_{1:D-1}$ , obtained from the generated observation sequences  $\{o_n\}_j$ . For now, we assume the observation sequences  $\{o_n\}_j$  are fixed.

Applying the importance sampling formalism to our system for all depths  $d = 0, \dots, D-1$ ,  $\mathcal{P}^d$  is the normalized measure incorporating the probability of observation sequence  $j$  on top of the state sequence  $i$  until the node at depth  $d$ , and  $\mathcal{Q}^d$  is the measure of the state sequence. We can think of  $\mathcal{P}^d$  being indexed by the observation sequence  $\{o_n\}_j$ .

$$\mathcal{P}^d = \mathcal{P}_{\{o_n\}_j}^d(\{s_n\}_i) = \frac{(\mathcal{Z}_{1:d}^{i,j})(\mathcal{T}_{1:d}^i) b_0^i}{\int_{S^{d+1}} (\mathcal{Z}_{1:d}^{j,i})(\mathcal{T}_{1:d}) b_0 ds_{0:d}} \quad (20)$$

$$\mathcal{Q}^d = \mathcal{Q}^d(\{s_n\}_i) = (\mathcal{T}_{1:d}^i) b_0^i \quad (21)$$

$$w_{\mathcal{P}^d/\mathcal{Q}^d}(\{s_n\}_i) = \frac{(\mathcal{Z}_{1:d}^{i,j})}{\int_{S^{d+1}} (\mathcal{Z}_{1:d}^{j,i})(\mathcal{T}_{1:d}) b_0 ds_{0:d}} \quad (22)$$

The weighing step is done by updating the self-normalized weights given in POWSS algorithm. We define  $w_{d,i}$  and  $r_{d,i}$  as the weights and rewards obtained at step  $d$  for state sequence  $i$  from POWSS simulation. With our recursive definition of the empirical weights, we obtain the full weight of each state sequence  $i$  for a fixed observation sequence  $j$ :

$$w_{D-1,i} = w_{D-2,i} \cdot Z(o_{D-1,j}|a_{D-1}, s_{D-1,i}) \quad (23)$$

$$\propto \mathcal{Z}_{1:D-1}^{i,j} \quad (24)$$

Realizing that the marginal observation probability is independent of indexing by  $i$ , we show that  $\hat{Q}_{D-1}^*(\bar{b}_{D-1}, a)$  is an SN estimator of  $Q_{D-1}^*(b_{D-1}, a)$ :

$$\hat{Q}_{D-1}^*(\bar{b}_{D-1}, a) = \frac{\sum_{i=1}^C (\mathcal{Z}_{1:D-1}^{i,j}) R(s_{D-1,i}, a)}{\sum_{i=1}^C (\mathcal{Z}_{1:D-1}^{i,j})} \quad (25)$$

$$= \frac{\sum_{i=1}^C \frac{(\mathcal{Z}_{1:D-1}^{i,j})}{\int_{S^D} (\mathcal{Z}_{1:D-1}^{j,i})(\mathcal{T}_{1:D-1}) b_0 ds_{0:D-1}} R(s_{D-1,i}, a)}{\sum_{i=1}^C \frac{(\mathcal{Z}_{1:D-1}^{i,j})}{\int_{S^D} (\mathcal{Z}_{1:D-1}^{j,i})(\mathcal{T}_{1:D-1}) b_0 ds_{0:D-1}}} \quad (26)$$

$$= \frac{\sum_{i=1}^C w_{\mathcal{P}^{D-1}/\mathcal{Q}^{D-1}}(\{s_n\}_i) R(s_{D-1,i}, a)}{\sum_{i=1}^C w_{\mathcal{P}^{D-1}/\mathcal{Q}^{D-1}}(\{s_n\}_i)} \quad (27)$$

$$= \sum_{i=1}^C \tilde{w}_{\mathcal{P}^{D-1}/\mathcal{Q}^{D-1}}(\{s_n\}_i) R(s_{D-1,i}, a) \quad (28)$$

Since  $\{s_n\}_1, \dots, \{s_n\}_C$  are i.i.d.r.v. sequences of depth  $D$  sampled from  $\mathcal{Q}^{D-1}$ , and  $R$  is a bounded function from problem requirement (iii), we can apply the SN concentration bound in Theorem 1 to prove Lemma 1. Detailed finishing steps of the proof are given in Appendix B.  $\square$

### Induction from Leaf to Root Nodes

Now, we want to show that nodes at all depths have convergence guarantees via induction.

**Lemma 2** (SN Estimator Step-by-Step Convergence).

$\hat{Q}_d^*(\bar{b}_d, a)$  is an SN estimator of  $Q_d^*(b_d, a)$ , and for all  $d = 0, \dots, D-1$  and  $a$ , the following holds with probability at least  $1 - 3|A|(3|A|C)^D \exp(-C \cdot t_{\max}^2)$ :

$$|Q_d^*(b_d, a) - \hat{Q}_d^*(\bar{b}_d, a)| \leq \alpha_d \quad (29)$$

$$\alpha_d \equiv \lambda + \gamma \alpha_{d+1}; \alpha_{D-1} = \lambda \quad (30)$$

*Proof.* First, we set  $C$  such that  $C > (3V_{\max}d_{\infty}^{\max}/\lambda)^2$  to satisfy  $t_{\max}(\lambda, C) > 0$ , which ensures that the SN concentration inequality holds with probability  $1 - 3\exp(-C \cdot t_{\max}^2(\lambda, C))$  at any given step  $d$  and action  $a$ . Furthermore, we multiply the worst-case union bound factor  $(3|A|C)^D$ , since we want the function estimates to be within their respective concentration bounds for all the actions  $|A|$  and child nodes  $C$  at each step  $d = 0, \dots, D - 1$ , for the 3 times we use SN concentration bound in the induction step. We once again multiply the final  $\delta$  by  $|A|$  to account for the root node  $Q$ -value estimates also satisfying their respective concentration bounds for all actions.

Following our definition of ESTIMATEQ, the value function estimates at step  $d$  are given as the following:

$$\hat{V}_d^*(\bar{b}_d) = \max_{a \in A} \hat{Q}_d^*(\bar{b}_d, a) \quad (31)$$

$$\hat{Q}_d^*(\bar{b}_d, a) = \frac{\sum_{i=1}^C w_{d,i} (r_{d,i} + \gamma \hat{V}_{d+1}^*(\overline{b_d a o_i}))}{\sum_{i=1}^C w_{d,i}} \quad (32)$$

The base case  $d = D - 1$  holds by Lemma 1. Then for the inductive step, we assume Eq. (29) holds for all actions at step  $d + 1$ . Using the triangle inequality for step  $d$ , we split the difference into two terms, the reward estimation error (A) and the next-step value estimation error (B):

$$\begin{aligned} & |Q_d^*(b_d, a) - \hat{Q}_d^*(\bar{b}_d, a)| \quad (33) \\ & \leq \underbrace{\left| \mathbb{E}[R(s_d, a)|b_d] - \frac{\sum_{i=1}^C w_{d,i} r_{d,i}}{\sum_{i=1}^C w_{d,i}} \right|}_{(A)} \\ & \quad + \gamma \underbrace{\left| \mathbb{E}[V_{d+1}^*(b a o)|b_d] - \frac{\sum_{i=1}^C w_{d,i} \hat{V}_{d+1}^*(\overline{b_d a o_i})}{\sum_{i=1}^C w_{d,i}} \right|}_{(B)} \end{aligned}$$

Each of the error terms are bound by (A)  $\leq \frac{R_{\max}}{3V_{\max}} \lambda$  and (B)  $\leq \frac{1}{3} \lambda + \frac{2}{3\gamma} \lambda + \alpha_{d+1}$ . We provide a detailed justification of these bounds in Appendix C, which uses the SN concentration bound 3 times. Combining (A) and (B), we prove the inductive hypothesis:

$$\begin{aligned} |Q_d^*(b_d, a) - \hat{Q}_d^*(\bar{b}_d, a)| & \leq \frac{R_{\max}}{3V_{\max}} \lambda + \gamma \left[ \frac{1}{3} \lambda + \frac{2}{3\gamma} \lambda + \alpha_{d+1} \right] \\ & \leq \lambda + \gamma \alpha_{d+1} = \alpha_d \quad (34) \end{aligned}$$

Therefore, Eq. (29) holds for all  $d = 0, \dots, D - 1$  with probability at least  $1 - 3|A|(3|A|C)^D \exp(-C \cdot t_{\max}^2)$ .  $\square$

*Proof.* (Theorem 2) First, we choose constants  $C, \lambda, \delta$  and densities  $\mathcal{Z}, \mathcal{T}, b_0$  that satisfy the conditions in Theorem 2. Since  $\alpha_d \leq \alpha_0$ , the following holds for all  $d = 0, \dots, D - 1$  with probability at least  $1 - \delta$  through Lemmas 1 and 2:

$$|Q_d^*(b_d, a) - \hat{Q}_d^*(\bar{b}_d, a)| \leq \alpha_0 = \sum_{d=0}^{D-1} \gamma^d \lambda \leq \frac{\lambda}{1 - \gamma} \quad (35)$$

Note that the convergence rate  $\delta$  is  $\mathcal{O}(C^D \exp(-tC))$ , where  $t = (\lambda/(3V_{\max}d_{\infty}^{\max}))^2$ .  $\square$

## Near-Optimal Policy Performance

We have proven in the previous subsection that the planning step results in a near-optimal  $Q$ -value for a given belief. Assuming further that we have a perfect Bayesian belief update in the outer observe-plan-act loop, we prove Theorem 3, which states that the closed-loop POMDP policy generated by POWSS at each planning step results in a near-optimal policy. The proof given in Appendix D combines Theorem 2 with results from Kearns *et al.* [2002]; Singh and Yee [1994]:

$$V^*(b_0) - V^{\text{POWSS}}(b_0) \leq \epsilon \quad (36)$$

## 7 Experiments

The simple numerical experiments in this section confirm the theoretical results of Section 6. Specifically, they show that the value function estimates of POSS converge to the QMDP approximation and the value function estimates of POWSS converge to the optimal value function for a toy problem.

### 7.1 Continuous Observation Tiger Problem

We consider a simple modification of the classic tiger problem [Kaelbling *et al.*, 1998] that we refer to as the continuous observation tiger (CO-tiger) problem. In the CO-tiger problem, the agent is presented with two doors, left (L) and right (R). One door has a tiger behind it ( $S = \{\text{Tiger L}, \text{Tiger R}\}$ ). In the classic problem, the agent can either open one of the doors or listen, and the CO-tiger problem has an additional wait action to illustrate the suboptimality of QMDP estimates ( $A = \{\text{Open L}, \text{Open R}, \text{Wait}, \text{Listen}\}$ ). If the agent opens a door, the problem terminates immediately; If the tiger is behind that door, a penalty of -10 is received, but if not, a reward of 10 is given. Waiting has a penalty of -1 and listening has a penalty of -2. If the agent waits or listens, a noisy continuous observation between 0 and 1 is received ( $O = [0, 1]$ ). In the wait case, this observation is uniformly distributed, independent of the tiger's position, yielding no information. In the listen case, the observation distribution is piecewise uniform. An observation in  $[0, 0.5]$  corresponds to a tiger behind the left door and  $(0.5, 1]$  the right door. Listening yields an observation in the correct range 85% of the time. The discount is 0.95, and the terminal depth is 3.

The optimal solution to this problem may be found by simply discretizing the observation space so that any continuous observation in  $[0, 0.5]$  is treated as a Tiger L observation, and any continuous observation in  $(0.5, 1]$  is treated as a Tiger R observation. This fully discrete version of the problem may be easily solved by a classical solution method such as the incremental pruning method of Cassandra *et al.* [1997]. Given an evenly-distributed initial belief, the optimal action is Listen with a value of 4.65, and the Wait action has a value of 3.42. The QMDP estimate for Wait is 8.5 and for Listen is 7.5.

While the CO-tiger problem is too small to be of practical significance, it serves as an empirical demonstration that POWSS converges toward the optimal value estimates and that POSS converges toward the QMDP estimates. In fact, the QMDP estimates generated by POSS are suboptimal in this example and lead to picking the suboptimal Wait

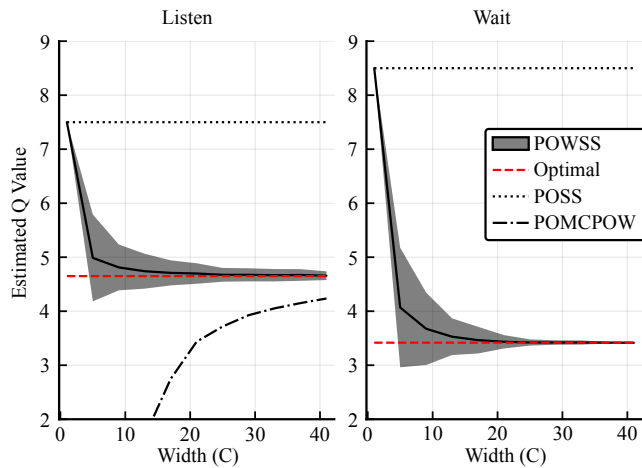


Figure 2: Numerical convergence of  $Q$ -value estimates for POSS, POWSS, and POMCPOW in the CO-tiger problem. Ribbons indicate standard deviation.

action. Both POWSS and POSS were implemented using the POMDPs.jl framework, [Egorov *et al.*, 2017] and open-source code can be found at <https://github.com/JuliaPOMDP/SparseSampling.jl>.

## 7.2 Results

The results plotted in Fig. 2 show the  $Q$ -value estimates of POWSS converging toward the optimal  $Q$ -values as the width  $C$  is increased. Each data point represents the mean  $Q$ -value from 200 runs of the algorithm from a uniformly-distributed belief, with the standard deviation plotted as a ribbon. The estimates for POSS have no uncertainty bounds since the estimates in this problem are the same for all  $C$ .

With  $C = 1$ , POWSS suffers from particle depletion and, because of the particular structure of this problem, finds the QMDP  $Q$ -values. As  $C$  increases, one can observe that both bias and variance in the  $Q$ -value estimates significantly decrease in agreement with our theoretical results, while POSS continues to yield incorrect estimates.

Some estimates by POMCPOW are also included. These are not directly comparable since POMCPOW is parameterized differently. For these tests, the double progressive widening parameters  $k_o = C$ ,  $\alpha_o = 0$  were used to limit the tree width, with  $n = C^3$  iterations to keep the particle density high in wider trees (see Sunberg and Kochenderfer [2018] for parameter definitions). POMCPOW’s value estimates are strongly biased downwards by exploration actions, but the estimated value for Listen action is much higher than the estimated value for the Wait action, which is too low to appear on the plot. Thus the correct action will usually still be chosen. At  $C = 41$ , POMCPOW is about an order of magnitude faster than POWSS.

## 8 Conclusion

This work has proposed two new POMDP algorithms and analyzed their convergence in POMDPs with continuous observation spaces. Though these algorithms are not computationally

efficient and thus not suitable for realistic problems, this work lays the foundation for analysis of more complex algorithms, rigorously justifying the observation likelihood weighting used in POWSS, POMCPOW, and DESPOT- $\alpha$ .

There is a great deal of future work to be done along this path. Most importantly, the theory presented in this work should be extended to more computationally efficient and hence practical algorithms. Before extending to POMCPOW and DESPOT- $\alpha$ , it may be beneficial to apply these techniques to an algorithm that is less conceptually complex, such as a modification of Sparse-UCT [Bjarnason *et al.*, 2009] extended to partially observable domains. Such an algorithm could enjoy strong theoretic guarantees, ease of implementation, and good performance on large problems.

Moreover, the proof techniques in this work may yield insight into which problems are difficult for sparse tree search techniques. For example, the Rényi divergence between the marginal and conditional state distributions (assumption (ii)) may be a difficulty indicator for likelihood-weighted sparse tree solvers, similar to the covering number of the optimal reachable belief space for point-based solvers [Lee *et al.*, 2008].

## Acknowledgements

This material is based upon work supported by a DARPA Assured Autonomy Grant, the SRC CONIX program, NSF CPS Frontiers, the ONR Embedded Humans MURI, and the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE 1752814. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of any aforementioned organizations. The authors also thank Lasse Peters for valuable manuscript feedback.

## References

- [Agha-Mohammadi *et al.*, 2011] Ali-Akbar Agha-Mohammadi, Suman Chakravorty, and Nancy M. Amato. FIRM: Feedback controller-based information-state roadmap - a framework for motion planning under uncertainty. In *IEEE/RSJ IROS*, 2011.
- [Ayer *et al.*, 2012] Turgay Ayer, Oguzhan Alagoz, and Natasha K. Stout. A POMDP approach to personalize mammography screening decisions. *Operations Research*, 60(5):1019–1034, 2012.
- [Bai *et al.*, 2014] Haoyu Bai, David Hsu, and Wee Sun Lee. Integrated perception and planning in the continuous space: A POMDP approach. *International Journal of Robotics Research*, 33(9):1288–1302, 2014.
- [Bai *et al.*, 2015] Haoyu Bai, Shaojun Cai, Nan Ye, David Hsu, and Wee Sun Lee. Intention-aware online POMDP planning for autonomous driving in a crowd. In *IEEE ICRA*, pages 454–460, 2015.
- [Bertsekas, 2005] D. Bertsekas. *Dynamic Programming and Optimal Control*. Athena, 2005.
- [Bjarnason *et al.*, 2009] Ronald Bjarnason, Alan Fern, and Prasad Tadepalli. Lower bounding Klondike solitaire with Monte-Carlo planning. In *ICAPS*, 2009.

- [Bry and Roy, 2011] Adam Bry and Nicholas Roy. Rapidly-exploring random belief trees for motion planning under uncertainty. In *IEEE ICRA*, pages 723–730, 2011.
- [Cassandra *et al.*, 1997] Anthony Cassandra, Michael L. Littman, and Nevin L. Zhang. Incremental pruning: A simple, fast, exact method for partially observable Markov decision processes. In *UAI*, pages 54–61, 1997.
- [Cassandra, 1998] Anthony R. Cassandra. A survey of POMDP applications. In *AAAI Fall Symposium: Planning with POMDPs*, 1998.
- [Couëtoux *et al.*, 2011] Adrien Couëtoux, Jean-Baptiste Hoock, Nataliya Sokolovska, Olivier Teytaud, and Nicolas Bonnard. Continuous upper confidence trees. In *Learning and Intelligent Optimization*, Rome, Italy, 2011.
- [Egorov *et al.*, 2017] Maxim Egorov, Zachary N. Sunberg, Edward Balaban, Tim A. Wheeler, Jayesh K. Gupta, and Mykel J. Kochenderfer. POMDPs.jl: A framework for sequential decision making under uncertainty. *Journal of Machine Learning Research*, 18(26):1–5, 2017.
- [Garg *et al.*, 2019] Neha P. Garg, David Hsu, and Wee Sun Lee. DESPOT- $\alpha$ : Online POMDP planning with large state and observation spaces. In *RSS*, 2019.
- [Hoey and Poupart, 2005] Jesse Hoey and Pascal Poupart. Solving POMDPs with continuous or large discrete observation spaces. In *IJCAI*, pages 1332–1338, 2005.
- [Holland *et al.*, 2013] Jessica E. Holland, Mykel J. Kochenderfer, and Wesley A. Olson. Optimizing the next generation collision avoidance system for safe, suitable, and acceptable operational performance. *Air Traffic Control Quarterly*, 21(3):275–297, 2013.
- [Kaelbling *et al.*, 1998] Leslie Pack Kaelbling, Michael L. Littman, and Anthony R. Cassandra. Planning and acting in partially observable stochastic domains. *Artificial Intelligence*, 101(1):99 – 134, 1998.
- [Kearns *et al.*, 2002] Michael Kearns, Yishay Mansour, and Andrew Y. Ng. A sparse sampling algorithm for near-optimal planning in large Markov decision processes. *Machine Learning*, 49(2):193–208, 2002.
- [Kochenderfer, 2015] Mykel J. Kochenderfer. *Decision Making Under Uncertainty: Theory and Application*. MIT Press, 2015.
- [Kurniawati and Yadav, 2016] Hanna Kurniawati and Vinay Yadav. An online POMDP solver for uncertainty planning in dynamic environment. In *Robotics Research*, pages 611–629. Springer, 2016.
- [Lee *et al.*, 2008] Wee Sun Lee, Nan Rong, and David Hsu. What makes some POMDP problems easy to approximate? In *NeurIPS*, pages 689–696, 2008.
- [McAllester and Singh, 1999] David A. McAllester and Satinder Singh. Approximate planning for factored POMDPs using belief state simplification. In *UAI*, pages 409–416, 1999.
- [Metelli *et al.*, 2018] Alberto Maria Metelli, Matteo Papini, Francesco Faccio, and Marcello Restelli. Policy optimization via importance sampling. In *NeurIPS*, pages 5442–5454, 2018.
- [Papadimitriou and Tsitsiklis, 1987] Christos H. Papadimitriou and John N. Tsitsiklis. The complexity of Markov decision processes. *Mathematics of Operations Research*, 12(3):441–450, 1987.
- [Platt *et al.*, 2010] Robert Platt, Jr., Russ Tedrake, Leslie Kaelbling, and Tomas Lozano-Perez. Belief space planning assuming maximum likelihood observations. In *RSS*, 2010.
- [Silver and Veness, 2010] David Silver and Joel Veness. Monte-Carlo planning in large POMDPs. In *NeurIPS*, pages 2164–2172, 2010.
- [Singh and Yee, 1994] Satinder P. Singh and Richard C. Yee. An upper bound on the loss from approximate optimal-value functions. *Machine Learning*, 16(3):227–233, 1994.
- [Sunberg and Kochenderfer, 2018] Zachary Sunberg and Mykel J. Kochenderfer. Online algorithms for POMDPs with continuous state, action, and observation spaces. In *ICAPS*, 2018.
- [Sunberg *et al.*, 2013] Zachary Sunberg, Suman Chakravorty, and R Scott Erwin. Information space receding horizon control. *IEEE Transactions on Cybernetics*, 43(6):2255–2260, 2013.
- [Sunberg *et al.*, 2017] Zachary N. Sunberg, Christopher J. Ho, and Mykel J. Kochenderfer. The value of inferring the internal state of traffic participants for autonomous free-way driving. In *ACC*, 2017.
- [Van Den Berg *et al.*, 2012] Jur Van Den Berg, Sachin Patil, and Ron Alterovitz. Motion planning under uncertainty using iterative local optimization in belief space. *International Journal of Robotics Research*, 31(11):1263–1278, 2012.
- [Ye *et al.*, 2017] Nan Ye, Adhiraj Somani, David Hsu, and Wee Sun Lee. DESPOT: Online POMDP planning with regularization. *Journal of Artificial Intelligence Research*, 58:231–266, 2017.
- [Young *et al.*, 2013] Steve Young, Milica Gašić, Blaise Thomson, and Jason D Williams. POMDP-based statistical spoken dialog systems: A review. *IEEE*, 101(5):1160–1179, 2013.