Simplifying Complex Observation Models in Continuous POMDP Planning with Probabilistic Guarantees and Practice

Idan Lev-Yehudi[1], **Moran Barenboim**[1], **Vadim Indelman**[2]

[1]Technion Autonomous Systems Program (TASP), Technion - Israel Institute of Technology, Haifa 32000, Israel
[2]Department of Aerospace Engineering, Technion - Israel Institute of Technology, Haifa 32000, Israel
moranbar, idanlev@campus.technion.ac.il, vadim.indelman@technion.ac.il

Abstract

Solving partially observable Markov decision processes (POMDPs) with high dimensional and continuous observations, such as camera images, is required for many real life robotics and planning problems. Recent researches suggested machine learned probabilistic models as observation models, but their use is currently too computationally expensive for online deployment. We deal with the question of what would be the implication of using simplified observation models for planning, while retaining formal guarantees on the quality of the solution. Our main contribution is a novel probabilistic bound based on a statistical total variation distance of the simplified model. We show that it bounds the theoretical POMDP value w.r.t. original model, from the empirical planned value with the simplified model, by generalizing recent results of particle-belief MDP concentration bounds. Our calculations can be separated into offline and online parts, and we arrive at formal guarantees without having to access the costly model at all during planning, which is also a novel result. Finally, we demonstrate in simulation how to integrate the bound into the routine of an existing continuous online POMDP solver.

Introduction

Partially observable Markov decision processes (POMDP) are a flexible mathematical framework for modeling real world decision-making and planning problems with inherent uncertainty. Yet, POMDPs are notoriously hard to solve (Papadimitriou and Tsitsiklis 1987). Online solvers like POMCP (Silver and Veness 2010) and DESPOT (Ye et al. 2017) focus on finding a solution for the current belief only, rather than solving for the entire belief space, hereby restricting the complexity of solution.

These planners are not trivially suitable for problems with continuous observation or action spaces. In recent years there have been several practical approaches for continuous solvers, such as POMCPOW, PFT-DPW (Sunberg and Kochenderfer 2018) and LABECOP (Hoerger and Kurniawati 2021). However the current state of the art is divided between practical algorithms and those with guarantees (Lim, Tomlin, and Sunberg 2021).

Recently were suggested deep neural networks for probabilistic visual observation models in AI. (Eslami et al. 2018) suggested a neural model for scene rendering. Both (Jonschkowski, Rastogi, and Brock 2018) and (Karkus, Hsu, and Lee 2018) suggested DPF and PF-NET respectively, which are models aimed at learning particle filtering. In the context of model free planning, examples include QMDP-net (Karkus, Hsu, and Lee 2017) and Deep Variational Reinforcement Learning (Igl et al. 2018). In the context of model based planning, recent approaches include DualSMC (Wang et al. 2020) and VTS (Deglurkar et al. 2023). All of these works demonstrate the computational challenges arising from incorporating deep neural visual models into practical POMDP planning.

Our research focuses on the implications of using a different observation model, less accurate but computationally favorable, during planning in continuous POMDPs. This could describe for example the case of using a shallower neural

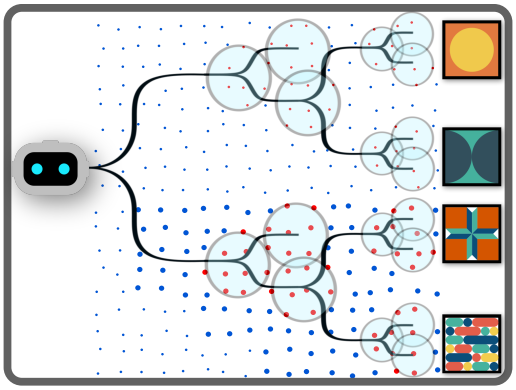

Figure 1: An illustration of a planning session with a simplified observation model. The scattered dots are the pre-sampled states, and the dot size is relative to Δ_Z, the estimated discrepancy between the simplified and original observation models. The simplified observation model is less accurate on the bottom where the surroundings are more visually complex. For the two policies, we compute the bound as a summation over Δ_Z weighted by the transition model. We bound the summation to a truncation distance indicated by the cyan circles, and Δ_Z within it is marked in red. The bottom policy chooses actions that give higher weights to states with greater Δ_Z, resulting in looser bounds.

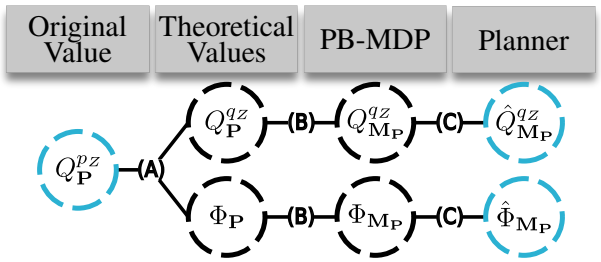

Figure 2: The various relationships between the action value functions in Corollary 3 for probably-approximately bounding $|Q^{pz}_{\mathbf{P}} - \hat{Q}^{qz}_{\mathbf{M_P}}| \leq \hat{\Phi}_{\mathbf{M_P}}$. (A) is given by Theorem 2, connecting theoretical value functions with the theoretical local state bound. (B) is given by Corollary 2, connecting theoretical action value functions with their PB-MDP approximation. (C) is given by any planner with performance guarantees, such as POWSS, approximating the PB-MDP values.

network for planning than for inference. To the best of our knowledge, there still isn't a clear-cut answer as to how accurate a learned observation model needs to be in order to attain a target performance in continuous POMDP planning. Our work provides insight into the complexity-performance trade-off in planning with different observation models.

Contribution

In this paper, we employ a simplified observation model and derive guarantees in the form of concentration inequalities. To the best of our knowledge, this is the first to work to do so. Specifically,

- We derive a novel non-parametric bound on the difference between value functions when planning with different observation models, and we observe that it can be formulated as a local state function.
- Our bound is practically computed by separating calculations to offline and online parts, such that the online estimator of the bounds refrains from accessing the computationally expensive original observation model.
- We derive concentration bounds on the online estimator of the bound. We do so by generalizing previous convergence results of particle-belief MDPs to state rewards under general policies.
- Finally, we demonstrate the practical computation of the bound's estimator in a simple simulated environment.

Related Work

Simplification of Probabilistic Models in POMDP Planning In (Hoerger, Kurniawati, and Elfes 2023) the authors introduce considered several simplification levels of the transition model, whereas we consider only two levels. Yet Hoerger, Kurniawati, and Elfes had to access the most complex model online, and their convergence guarantee is only asymptotic. (Ha and Schmidhuber 2018) consider learning a simplified generative model, for environments in which direct training would be infeasible. On this "world model" they train the policy, yet they too do not have any performance guarantees for the learned policy. A technique for measuring non-linearity based on total variation distance in POMDP planning was considered by (Hoerger et al. 2020). While there are some similarities in the approach to our work, they did not show how their practical estimator is related to the theoretical bound, nor did they give any performance guarantees.

Continuous POMDP Planning With Guarantees (Lim, Tomlin, and Sunberg 2020) proved that POWSS, a Monte Carlo tree search algorithm for continuous POMDPs based on a modification of Sparse Sampling (Kearns, Mansour, and Ng 2002), converges to the optimal policy in high probability. However, is not efficient enough for practical use. Later in (Lim et al. 2023) the convergence results were extended to prove that generally the particle-belief MDP (PB-MDP) accurately approximates the original POMDP with high probability. Recently (Shienman and Indelman 2022; Barenboim, Lev-Yehudi, and Indelman 2023) developed pruning of data association hypotheses in continuous POMDPs with guarantees. The former focuses on a specific reward of entropy of hypotheses weights, while the latter provides looser bounds but for any general state reward. In both works, the hybrid belief setting and type of simplification are different from our work.

Preliminaries

A POMDP is the tuple $\langle \mathcal{X}, \mathcal{A}, \mathcal{Z}, p_T, p_Z, r, \gamma, L, b_0 \rangle$. $\mathcal{X}, \mathcal{A}, \mathcal{Z}$ are the state space, action space and observation space respectively. Throughout the paper, $x \in \mathcal{X}$, $a \in \mathcal{A}$, $z \in \mathcal{Z}$ denote individual states, actions and observations respectively. We consider $\mathcal{X}$ and $\mathcal{Z}$ that are continuous, while $\mathcal{A}$ can be either discrete or continuous

The conditional probability distributions p_T and p_Z are the transition and observation models respectively. $p_T(x' \mid x, a)$ models the uncertainty of taking action $a \in \mathcal{A}$ from state $x \in \mathcal{X}$, and $p_Z(z \mid x)$ models the uncertainty of receiving an observation $z \in \mathcal{Z}$ at state x. The reward function $r\colon \mathbb{N} \times \mathcal{X} \times \mathcal{A} \to \mathbb{R}$ gives the immediate reward of applying action a at state x and time t. We denote $r_t\colon \mathcal{X} \times \mathcal{A} \to \mathbb{R}$ as the reward function at time t. We assume that the POMDP starts at time 0 and terminates after $L \in \mathbb{N}$ steps. $\gamma \in (0, 1]$ is the discount factor.

Because of the partial observability, the agent has to maintain a probability distribution of the current state given past actions and observations, known as the belief. The initial belief of the agent, denoted as b_0, captures the uncertainty in the initial state. A history at time t is defined as a sequence of the starting belief, followed by actions taken and observations received until t: $H_t \triangleq (b_0, a_0, z_1, \ldots, a_{t-1}, z_t)$. The belief at time t is defined as the conditional distribution of the state given the history $b_t(x_t) \triangleq \mathbb{P}(x_t \mid H_t)$. We denote $H_t^- \triangleq (b_0, a_0, z_1, \ldots, a_{t-1})$ for the same history without the last measurement, and the propagated belief $b_t^-(x_t) \triangleq \mathbb{P}(x_t \mid H_t^-)$. The belief reward is the expected state reward: $r_i(b_i, a) \triangleq \mathbb{E}_{x_i \sim b_i}[r_i(x_i, a)]$.

A policy at time i, denoted as π_i, is a mapping from the space of all histories to the action space $\pi_i\colon \mathcal{H}_i \to \mathcal{A}$. A policy π is a collection of policies from the start time until the

POMDP terminates, i.e. $\pi \triangleq (\pi_t)_{t=0}^{t=L-1}$, and for brevity we denote π_i instead of $\pi_i(H_i)$. It can be shown that optimal decision-making can be made given the belief, instead of considering the entire history (Kaelbling, Littman, and Cassandra 1998). Since the belief and history fulfill the Markov property, a POMDP is a Markov decision process (MDP) on the belief space, commonly referred to as the belief MDP.

The value function of a policy from time t is the expected sum of discounted rewards until the horizon of the POMDP: $V_t^\pi(b_t) \triangleq \mathbb{E}_{z_{t+1:L}}[\sum_{i=t}^{L} \gamma^{i-t} r_i(b_i, \pi_i)]$. Most often the action value function, defined as $Q_t^\pi(b_t, a) \triangleq r(b_t, a) + \gamma \mathbb{E}_{z_{t+1}}[V_{t+1}^\pi(b_{t+1})]$, is used as an intermediate step in the estimation of the value function. The goal of planning is to compute a policy that would maximize the value function. The optimal policy is denoted $\pi^* = \arg\max_\pi V_t^\pi(b_t)$, and the corresponding value and action value functions for the optimal policy are V_t^* and Q_t^*.

Often in real world applications it is impractical to update the belief exactly. Particle filters are used instead to represent a belief non parametrically. The particle belief of C particles at time i is represented by a set state samples and weights $\bar{b}_t = \{(x_t^j, w_t^j)\}_{j=1}^C$, and is defined as the discrete distribution $\mathbb{P}(x \mid \bar{b}_t) = \frac{\sum_{j=1}^C w_t^j \cdot \delta(x - x_t^j)}{\sum_{j=1}^C w_t^j}$. The basic particle filtering algorithm, Sequential Importance Sampling (SIS) (Doucet, de Freitas, and Gordon 2001, 1.3.2), approximates the target posterior distribution by recursively updating the importance weights of the particles with the observation likelihood, and normalizes the weights accordingly. Sequential Importance Resampling (SIR) adds resampling of state particles after the weight updates, to avoid weight degeneracy (Kong, Liu, and Wong 1994).

The Monte Carlo (MC) estimator of an expectation, denoted with $\hat{\mathbb{E}}$, is its approximation with a finite number of samples. This applies to quantities that are defined as expectations, such as $\hat{V}_t^\pi$. For example, the empirical value function based on N samples $\{z_{t+1:L}^j\}_{j=1}^N$ would be $\hat{V}_t^\pi = \frac{1}{N}\sum_{j=1}^N \gamma^{i-t} r_i(b_i^j, \pi_i^j)$, $\pi_i^j = \pi(H_i^j)$, where b_i^j and H_i^j are updated with $z_{t+1:i}^j$, i.e. the sampled sequence of observations until time i.

(Lim et al. 2023) considered the PB-MDP, where the belief space is explicitly the space of particle beliefs $\bar{b}$ of C particles. They proved that with high probability the optimal value in the PB-MDP is bounded from the optimal POMDP value, and this bound converges as $C \to \infty$.

We similarly denote the original POMDP as $\mathbf{P}$ and the PB-MDP as $\mathbf{M_P}$. Quantities that are measured w.r.t. each are denoted accordingly, f.e. the value function starting at time t, given policy π and starting belief b_t in the original POMDP is denoted as $V_{\mathbf{P},t}^\pi(b_t)$. When a result applies to both $\mathbf{P}$ and $\mathbf{M_P}$, we omit this notation.

Methodology

Problem Formulation

We consider cases in which the observation model p_Z is simplified for computational reasons. We denote the simplified model as q_Z, which is only used in planning. In our setting, inference is performed with the original model, such that the initial belief b_t for planning session starting at time t was updated with p_Z. We denote the future value function based on belief updates in planning with either the original or simplified observations models as V_t^{π,p_Z} and V_t^{π,q_Z} respectively, and similarly for Q_t^{π,p_Z} and Q_t^{π,q_Z}.

In the settings we consider, p_Z is computationally more expensive than q_Z, and both are considerably more expensive than p_T. In the example of a ground robot equipped with a camera, q_Z could be a neural network of similar architecture to p_Z but much shallower, whereas p_T might be a Gaussian or some other simple parametric distribution. Note that in part the difference in complexity arises from the difference in dimensions of the state and observation spaces.

Our goal is to derive calculable probabilistic bounds on $|V_t^{\pi,p_Z} - \hat{V}_t^{\pi,q_Z}|$, while restricting access to p_Z to offline computations only. Additionally, we wish to analyze the difference in using the bounds for post guarantees, i.e. after a policy has been extracted, to when they're used during the decision-making, i.e. policy computation.

We denote shortened notations for the following densities: $\tau_i \triangleq p_T(x_i \mid x_{i-1}, \pi_{i-1})$, $[\zeta/\xi]_i \triangleq [p_Z/q_Z](z_i \mid x_i)$, $[\zeta/\xi]_i^H \triangleq [p_Z/q_Z](z_i \mid H_i^-)$. The values in $[\cdot]$ can be replaced with either respective option. The product of densities is denoted $[\tau/\zeta/\xi]_{i:j} \triangleq \prod_{l=i}^{j} [\tau/\zeta/\xi]_l$. We denote the expectations:

$$\mathbb{E}_{b_i}[\star] \triangleq \int_{x_i} b_i(x_i) \star \,\mathrm{d}x_i \tag{1}$$

$$\mathbb{E}_{i:j}^{p_T}[\star] \triangleq \int_{x_{i:j}} \tau_{i:j} \star \,\mathrm{d}x_{i:j} \tag{2}$$

$$\mathbb{E}_{i:j}^{[p_Z/q_Z]}[\star] \triangleq \int_{z_{i:j}} [\zeta/\xi]_{i:j}^H \star \,\mathrm{d}z_{i:j} \tag{3}$$

$$\mathbb{E}_{i:j}^{p_T,[p_Z/q_Z]}[\star] \triangleq \int_{z_{i:j}} \mathbb{E}_{b_{i-1}}[\mathbb{E}_{i:j}^{p_T}[[\zeta/\xi]_{i:j}\star]] \,\mathrm{d}z_{i:j} \tag{4}$$

$$\mathbb{E}_{i:j+1^-}^{p_T,[p_Z/q_Z]}[\star] \triangleq \int_{z_{i:j}} \mathbb{E}_{b_{i-1}}[\mathbb{E}_{i:j+1}^{p_T}[[\zeta/\xi]_{i:j}\star]] \,\mathrm{d}z_{i:j}, \tag{5}$$

It holds that $[\zeta/\xi]_{i:j}^H = \mathbb{E}_{b_{i-1}}[\mathbb{E}_{i:j}^{p_T}[[\zeta/\xi]_{i:j}]]$ by the law of total probability, and it follows that $\mathbb{E}_{i:j}^{[p_Z/q_Z]}[\star] = \mathbb{E}_{i:j}^{p_T,[p_Z/q_Z]}[\star]$. Additionally, it holds that $\mathbb{E}_{i:j}^{p_T,[p_Z/q_Z]}[\mathbb{E}_{j+1}^{p_T}[\star]] = \mathbb{E}_{i:j+1^-}^{p_T,[p_Z/q_Z]}[\star]$ by Fubini's theorem (Durrett 2019, 1.7).

We denote for a bounded quantity its lower and upper bounds as $\mathcal{LB}$ and $\mathcal{UB}$ respectively. F.e, if $|A| \leq B$, then $\mathcal{LB}(A) = -B$ and $\mathcal{UB}(A) = B$.

For the formulation and practical computation of the bounds, we take the following assumptions:

i. The reward at each time step is bounded: $|r_i(x_i, a)| \leq R_i^{\max}$, for time index i. Hence, follows from the triangle inequality that the value function at a certain time step is bounded by $|V_t^\pi(b_t, a)| \leq \sum_{i=t}^{L} \gamma^{i-t} R_i^{\max} \triangleq V_t^{\max}$, and we define $V_{\max} \triangleq \max_t V_t^{\max}$.

ii. The reachable state space is totally bounded, i.e. for every $\varepsilon > 0$ it can be covered with a finite number of open balls of radius ε.

iii. The models p_T, p_Z, q_Z can be sampled from, and queried for PDF values given samples.

For the result of probabilistic convergence guarantees in Theorem 3, we additionally assume the following:

iv. Weights of particles in particle beliefs are updated via the Sequential Importance Sampling (SIS) algorithm, i.e. $w_t^j \propto [\zeta/\xi]_t \cdot w_{t-1}^j$ for the j'th observation sample z_t^j. Specifically, there is no resampling step.

v. The normalized observation likelihoods of the original and simplified models are bounded almost everywhere w.r.t. the particle filter's proposal distribution: $\sup_{z_{1:t}} \operatorname{ess\,sup}_{x_{0:t}\sim b_0\cdot\tau_{1:t}} \frac{[\zeta/\xi]_{1:t}}{[\zeta/\xi]_{1:t}^H} \leq d_\infty^{\max}$, where $d_\infty^{\max} \in \mathbb{R}$. See (Lim et al. 2023) for further details.

Bound for Simplified Observation Model

In this section we derive the theoretical bound for the value function with the original observation model and a given policy, w.r.t. the value with the simplified model. We start by stating the following known equivalence relationship between expectation over rewards.

Lemma 1. *The expected belief-dependent reward w.r.t. histories, is equivalent to the expected state-dependent reward w.r.t. the joint distribution of states and observations.*

$$\mathbb{E}_{t+1:i}^{[p_Z/q_Z]}[r_i(b_i,\pi_i)] = \mathbb{E}_{t+1:i}^{p_T,[p_Z/q_Z]}[r_i(x_i,\pi_i)]. \tag{6}$$

Proof. We refer to the supplementary material for all proofs (Lev-Yehudi, Barenboim, and Indelman 2023).

We define the following state dependent total variation distance (TV-distance) function,

$$\Delta_Z(x) \triangleq \int_{\mathcal{Z}} |p_Z(z \mid x) - q_Z(z \mid x)| \, \mathrm{d}z. \tag{7}$$

There are several motivations for considering $\Delta_Z(x)$. First, we can estimate it via samples for any density from which we can sample and evaluate for its PDF. It does not require any parametric form of the density, which is useful when considering general learned models. Second, it is state-dependent, hence can be computed locally as we show later for a given belief or trajectory in the belief tree. Different actions might result in different bounds and the locality helps differentiate actions that result in tighter or looser bounds. Third, as given by Pinsker's lemma, the TV-distance is bounded from above by the KL-divergence (Tsybakov 2009). Thus, current approaches that learn probabilistic models with ELBO or directly minimize empirical KL-divergence (Kingma and Welling 2014; Sohn, Lee, and Yan 2015; Rezende and Mohamed 2015; Winkler et al. 2019) also indirectly minimize the TV-distance, therefore it is appropriate to assume that models trained with these objectives will also indirectly minimize Δ_Z.

Theorem 1. *Assume current belief is b_t and a given policy is π. Denote with $b_i^{p_Z}$, $b_i^{q_Z}$ the future belief at time step i updated with either p_Z or q_Z, respectively. Then it holds that*

$$\begin{aligned} &|\mathbb{E}_{t+1:i}^{p_Z}[r_i(b_i^{p_Z},\pi_i)] - \mathbb{E}_{t+1:i}^{q_Z}[r_i(b_i^{q_Z},\pi_i)]| \\ &\quad \leq R_i^{\max} \textstyle\sum_{l=t+1}^{i} \mathbb{E}_{t+1:l^-}^{p_T,q_Z}[\Delta_Z(x_l)], \end{aligned} \tag{8}$$

i.e. the difference between the expected reward at future time i for the original and simplified POMDP is bounded by the maximum reward, times the sum of the expected state-dependent TV-distances between the observation models.

By applying to the value function with the triangle inequality, and changing the order of summation, we arrive at the following result.

Corollary 1. *The difference between the original and simplified value functions can be bounded by the following sum of scaled expected TV-distance terms:*

$$\begin{aligned} &|V_t^{\pi,p_Z}(b_t) - V_t^{\pi,q_Z}(b_t)| \\ &\leq \textstyle\sum_{i=t+1}^{L} \mathbb{E}_{t+1:i^-}^{p_T,q_Z}[\Delta_Z(x_i)] \sum_{l=i}^{L} \gamma^{l-t} R_l^{\max} \qquad (9) \\ &= \textstyle\sum_{i=t+1}^{L} V_i^{\max} \cdot \mathbb{E}_{t+1:i^-}^{p_T,q_Z}[\Delta_Z(x_i)]. \qquad (10) \end{aligned}$$

Equivalence to State-Action Local Bound

Corollary 1 requires computing the theoretical expectations $\mathbb{E}_{t+1:i^-}^{p_T,q_Z}[\Delta_Z(x_i)]$, which is not trivial with continuous states and observations. Our key insight is that the bound can be viewed as the expected sum of a state-action function. This serves two purposes. The first is that the bounds become more easily estimated, as we can integrate our calculations in existing POMDP planners by adding a secondary reward-like function to compute. The second is that we can extend convergence results of POMDP algorithms to the empirical estimate of the bound.

First we define the following time dependent state-action function, and its extension to a belief function,

$$m_i(x_i,\pi_i) \triangleq V_{i+1}^{\max} \cdot \mathbb{E}_{i+1}^{p_T}[\Delta_Z(x_{i+1})], \tag{11}$$

$$m_i(b_i,\pi_i) \triangleq \mathbb{E}_{b_i}[m(x_i,\pi_i)]. \tag{12}$$

Intuitively speaking, m_i bounds the loss of the value function at time i when considering the action of π_i from state x_i or belief b_i, based on Δ_Z. It is then natural to extend this definition with the following:

1. Cumulative bound for a given policy and initial belief, $M_t^\pi(b_t) \triangleq \mathbb{E}_{t+1:L-1}^{q_Z}[\sum_{i=t}^{L-1} m_i(b_i,\pi_i)]$, analogous to $V_t^\pi(b_t)$. Note that the time horizon has decreased by 1, and there is no discount factor.
2. Action cumulative bound for a given policy and initial belief, $\Phi_t^\pi(b_t,a) \triangleq m_t(b_t,a) + \mathbb{E}_{t+1}^{q_Z}[M_{t+1}^\pi(b_{t+1})]$, analogous to $Q_t^\pi(b_t,a)$.

Theorem 2. *Under the conditions of Theorem 1, the difference between the original and simplified value function can be bounded by*

$$|V_t^{\pi,p_Z}(b_t) - V_t^{\pi,q_Z}(b_t)| \leq M_t^\pi(b_t). \tag{13}$$

In addition, the respective difference in action value function can be bounded by the action cumulative bound,

$$|Q_t^{\pi,p_Z}(b_t,a) - Q_t^{\pi,q_Z}(b_t,a)| \leq \Phi_t^\pi(b_t,a). \tag{14}$$

Computing $M_t^\pi(b_t)$ would be useful when trying to arrive at a bound for a specific policy, usually as post-guarantees after a planning session. On the other hand, the computation of $\Phi_t^\pi(b_t,a)$ can be defined recursively over an entire belief tree, meaning it extends the original definition of the bound to all extractable policies from a given tree. Hence, the action cumulative bound can be used during planning, when computing a policy or pruning low value branches like in (Sztyglic and Indelman 2022).

Online Estimator of Local Bound

We now show how we estimate m_i during online planning without requiring access to p_Z. We do this by offline pre-sampling state samples $\{x_n^\Delta\}_{n=1}^{N_\Delta} \sim Q_0(x)$, named delta states, and evaluating $\Delta_Z(x_n^\Delta)$ for each. During online planning, we simply reweight Δ_Z using the importance sampling formalism for state sample x_i^j. We use this to redefine m_i and also to define the empirical estimate $\tilde{m}_i$.

$$m_i(x_i^j, a) = V_{i+1}^{\max} \cdot \mathbb{E}_{x' \sim Q_0}[\tfrac{p_T(x'|x_i^j,a)}{Q_0(x')} \Delta_Z(x')], \quad (15)$$

$$\tilde{m}_i(x_i^j, a) \triangleq V_{i+1}^{\max} \cdot \tfrac{1}{N_\Delta} \textstyle\sum_{i=1}^{N_\Delta} \tfrac{p_T(x_n^\Delta|x_i^j,a)}{Q_0(x_n^\Delta)} \Delta_Z(x_n^\Delta). \quad (16)$$

Based on this, we define $\tilde{\Phi}_t^\pi(b_t, a)$ when computed with $\tilde{m}_i$ instead of m_i. We explicitly define this notation to differentiate $\tilde{\Phi}$ from $\hat{\tilde{\Phi}}$; The former is defined with a theoretical expectation over the observation space, while the latter is an empirical estimate with sampled belief trajectories. For notational brevity, in the rest of the paper we denote $\hat{\Phi} \triangleq \hat{\tilde{\Phi}}$.

The support of Q_0 must cover all of $\mathcal{X}$ for the estimator in (16) to be consistent (Doucet, de Freitas, and Gordon 2001, 1.3.2). Although it is possible to construct such densities over unbounded state spaces, we find that in practice it is often not required as $\mathcal{X}$ is either naturally bounded, such as the configuration space of a manipulator, or can be defined to be large enough but bounded to cover the relevant domain of the planning problem. Hence, it could be sufficient to choose Q_0 with a finite support. To the end of simplifying the discussion, we took assumption ii.

Under this condition, we can establish with Hoeffding's inequality the following simple probabilistic bound:

$$\mathbb{P}(|m_i(x, a) - \tilde{m}_i(x, a)| \geq \nu) \leq 2 \exp(-2N_\Delta \tfrac{\nu^2}{B_i^2}), \quad (17)$$

where $B_i = 2 \cdot V_i^{\max} \cdot \max_{x,x',a} \frac{p_T(x'|x,a)}{Q_0(x')}$ since $0 \leq \Delta_Z(x) \leq 2$. We assume B_i to be computable offline. This is sufficient for obtaining the concentration bounds of Theorem 3, we approximate this even further N_Δ may be very large, as we describe in section .

Convergence Guarantees

Our major convergence result is based on an adaptation of Lemma 1 of (Lim et al. 2023). Lemma 1 shows that the theoretical optimal POMDP action value $Q_{\mathbf{P}}^*$ is bounded with high probability from the theoretical optimal PB-MDP action value $Q_{\mathbf{M_P}}^*$. We extend their result in two ways: First, proving the bound for general policies, rather than only the optimal one, allows for proving convergence of Φ_t^π w.r.t. a policy that is chosen to maximize V_t^{π,q_Z}, or any other objective. Second, by introducing an additional assumption of a probabilistically bounded reward, denoted with $\tilde{r}$, we are able to apply Theorem 3 to $\tilde{m}_i$.

Theorem 3 (Generalized PB-MDP Convergence)**.** *Assume that the immediate state reward estimate is probabilistically bounded such that $\mathbb{P}(|r_i^j - \tilde{r}_i^j| \geq \nu) \leq \delta_r(\nu, N_r)$, for a number of reward samples N_r and state sample x_i^j. Assume that $\delta_r(\nu, N_r) \to 0$ as $N_r \to \infty$. For all policies π, $t = 0, \ldots, L$ and $a \in \mathcal{A}$, the following bounds hold with probability of at least $1 - 5(4C)^{L+1}(\exp(-C \cdot \acute{k}^2) + \delta_r(\nu, N_r))$:*

$$|Q_{\mathbf{P},t}^{\pi,[p_Z/q_Z]}(b_t, a) - Q_{\mathbf{M_P},t}^{\pi,[p_Z/q_Z]}(\bar{b}_t, a)| \leq \alpha_t + \beta_t, \quad (18)$$

where,

$$\alpha_t = (1 + \gamma)\lambda + \gamma\alpha_{t+1},\ \alpha_L = \lambda \geq 0, \quad (19)$$

$$\beta_t = 2\nu + \gamma\beta_{t+1},\ \beta_L = 2\nu \geq 0, \quad (20)$$

$$k_{\max}(\lambda, C) = \frac{\lambda}{4V_{\max} d_\infty^{\max}} - \frac{1}{\sqrt{C}} > 0, \quad (21)$$

$$\acute{k} = \min\{k_{\max}, \lambda/4\sqrt{2}V_{\max}\}. \quad (22)$$

If we require the bound to hold for all possible policies that can be extracted from a given belief tree simultaneously, then under the assumption of a finite action space, the probability is at least $1 - 5(4|\mathcal{A}|C)^{L+1}(\exp(-C \cdot \acute{k}^2) + \delta_r(\nu, N_r))$.

Corollary 2. *For arbitrary precision ε and accuracy δ we can choose constants λ, ν, C, N_r such that the following holds with probability of at least $1 - \delta$:*

$$|Q_{\mathbf{P},t}^{\pi,[p_Z/q_Z]}(\bar{b}_t, a) - Q_{\mathbf{M_P},t}^{\pi,[p_Z/q_Z]}(\bar{b}_t, a)| \leq \varepsilon. \quad (23)$$

By applying Corollary 2 twice, once to the original reward function then once to the action cumulative bound, we can conclude that the estimated simplified PB-MDP value is probabilistically bounded from the original theoretical POMDP value.

We now arrive at our key theoretical result.

Corollary 3. *Assuming that $\mathcal{P}$ is an MDP planner that can approximate Q-values with arbitrary precision $\varepsilon^{\mathcal{P}}$ at an accuracy $1 - \delta^{\mathcal{P}}$, we denote the precision and accuracy of the action value and action cumulative bound functions:*

$$\mathbb{P}(|Q_{\mathbf{M_P},t}^{\pi,q_Z}(\bar{b}_t, a) - \hat{Q}_{\mathbf{M_P},t}^{\pi,q_Z}(\bar{b}_t, a)| \leq \varepsilon_Q^{\mathcal{P}}) \geq 1 - \delta_Q^{\mathcal{P}} \quad (24)$$

$$\mathbb{P}(|\Phi_{\mathbf{M_P},t}^{\pi}(\bar{b}_t, a) - \hat{\Phi}_{\mathbf{M_P},t}^{\pi}(\bar{b}_t, a)| \leq \varepsilon_\Phi^{\mathcal{P}}) \geq 1 - \delta_\Phi^{\mathcal{P}} \quad (25)$$

From Corollary 2 it holds that we can choose constants λ, ν, C, N_r such that the following holds,

$$\mathbb{P}(|Q_{\mathbf{P},t}^{\pi,q_Z}(b_t, a) - Q_{\mathbf{M_P},t}^{\pi,q_Z}(\bar{b}_t, a)| \leq \varepsilon_Q) \geq 1 - \delta_Q, \quad (26)$$

$$\mathbb{P}(|\Phi_{\mathbf{P},t}^{\pi}(b_t, a) - \tilde{\Phi}_{\mathbf{M_P},d}^{\pi}(\bar{b}_t, a)| \leq \varepsilon_\Phi) \geq 1 - \delta_\Phi. \quad (27)$$

Then with probability of at least $1 - (\delta_Q + \delta_Q^{\mathcal{P}} + \delta_\Phi + \delta_\Phi^{\mathcal{P}})$

$$|Q_{\mathbf{P},t}^{\pi,p_Z}(b_t, a) - \hat{Q}_{\mathbf{M_P},t}^{\pi,q_Z}(\bar{b}_t, a)| \leq \quad (28)$$

$$\hat{\Phi}_{\mathbf{M_P},t}^{\pi}(\bar{b}_t, a) + \varepsilon_Q + \varepsilon_Q^{\mathcal{P}} + \varepsilon_\Phi + \varepsilon_\Phi^{\mathcal{P}}. \quad (29)$$

In summary, any planner that can approximate the PB-MDP values $Q_{\mathbf{M_P},t}^{\pi,q_Z}(\bar{b}_t, a)$ with high probability can also approximate the cumulative bound $\Phi_{\mathbf{M_P},t}^{\pi}(\bar{b}_t, a)$ with high probability, since it is mathematically formulated like a state reward. Therefore, we can construct a bound for the theoretical value $Q_{\mathbf{P},t}^{\pi,p_Z}(b_t, a)$ with high probability from online calculations that do not involve access to p_Z.

An important observation is that the probabilistic bound obtained by Theorem 3 is independent of the chosen policy, i.e. constant w.r.t. the actions. Therefore, when using

the bound in Corollary 3 for decision-making, if our goal is to distinguish between actions that result in maximal lower or upper bound, i.e. $\max_a[\mathcal{LB}/\mathcal{UB}](Q^{\pi,p_Z}_{\mathbf{P},t}(b_t,a))$, we will obtain the same action choice by computation of $\max_a \hat{Q}^{\pi,q_Z}_{\mathbf{M_P},t}(\bar{b}_t,a)[-/+]\hat{\Phi}^{\pi}_{\mathbf{M_P},t}(\bar{b}_t,a)$.

Implementation

We verify the computability of our approach in a 2D beacons POMDP. We integrate the computation of $\tilde{m}_i$ and $\hat{\Phi}_{\mathbf{M_P},t}$ into PFT-DPW, and showcase an example of where the bounds could affect a policy's decision-making.

Simulative Setting

Our experimental setting is a 2D light-dark inspired simulation, shown in figure 3.

An agent is in a wall-surrounded arena, with a gate at the bottom to a goal region. The agent's starting location is either to the left or to the right of the goal region. The agent's task is to enter the goal region without colliding with a wall. The observations are noisy measurements of the agent's location, being more certain when in the "light" region, and less when in the "dark" region.

The light region $\mathcal{X}_{\text{light}}$ is defined by circles centered at each of the 6 beacons located at the top of the arena, and $\mathcal{X}_{\text{dark}} = \mathcal{X} \setminus \mathcal{X}_{\text{light}}$. The observation model in the light region $p_Z(z \mid x \in \mathcal{X}_{\text{light}})$ is defined as a Gaussian mixture model (GMM) with 1126 components arranged such that they approximately form a truncated Gaussian distribution centered at x. The simplified observation model differs only in the light region, and it approximates p_Z with a single Gaussian: $q_Z(z \mid x \in \mathcal{X}_{\text{light}}) \sim \mathcal{N}(x, \Sigma^{q_Z}_{\text{light}})$. We refer readers to (Bar-Shalom, Li, and Kirubarajan 2004, 1.4.16) on approximating GMMs with a single Gaussian. In the dark region, $p_Z(z \mid x \in \mathcal{X}_{\text{dark}}) = q_Z(z \mid x \in \mathcal{X}_{\text{dark}}) = \mathcal{N}(x, \Sigma_{dark})$

This setting demonstrates a case where the original observation model is an overly parameterized model, like a complex neural network, and one would like to replace it with a less parameterized albeit similar model.

The action space is $\mathcal{A} = \{(\pm 1, 0), (0, \pm 1)\}$. The transition model $p_T(x' \mid x, a)$ is a Gaussian centered at the agent's location plus the action. The horizon is $L = 15$, and the POMDP will terminate early if the agent enters the goal region or collides with a wall.

The reward is only state and time dependent, and is a sum of three indicators: $r_t(x) = R_{hit} \cdot \mathbf{1}_{x \in \mathcal{X}_{goal}} + R_{miss} \cdot \mathbf{1}_{x \notin \mathcal{X}_{goal}} + R_{collide} \cdot \mathbf{1}_{x \in \mathcal{X}_{collision}}$. In all time steps, $R_{hit} = 100$, $R_{collide} = -50$. The miss reward is $R_{miss} = -50$ if $t = L$ and is -1 otherwise. The discount factor is $\gamma = 1$.

Further details of the experimental setup can be found in the supplementary material.

Implementation of Bounds

When sampling delta states $\{x^{\Delta}_n\}_{n=1}^{N_\Delta}$ offline, without prior knowledge of which states are more likely, we choose a uniform Q_0. In order to assure even coverage of the state space, we chose to sample them from the quasi-random sequence R_2, which has been shown empirically to minimize the discrepancy (Roberts 2018). Hence, we perform in practice quasi Monte Carlo method (QMCM) for the computation of $\tilde{m}_i$ (Caflisch 1998). It has been shown in many empirical examples that QMCM obtains faster convergence in practice than regular MC with an equivalent number of samples. In theory, it is possible to provide a deterministic upper bound to the QMCM integration error via the Koksma-Hlwaka inequality (Lemieux 2009, 5.6). However, it is generally hard to compute, and infinite for several very simple functions.

For each sampled x^{Δ}_n we estimate Δ_Z via observation samples, based on assumption iii. We perform the following importance sampling estimation w.r.t. $(p_Z + q_Z)/2$:

$$\hat{\Delta}_Z(x^{\Delta}_n) = \sum_{j=1}^{N_Z} 2 \cdot \frac{|p_Z(z^n_j \mid x^{\Delta}_n) - q_Z(z^n_j \mid x^{\Delta}_n)|}{p_Z(z^n_j \mid x^{\Delta}_n) + q_Z(z^n_j \mid x^{\Delta}_n)} \quad (30)$$

where $\{z^n_j\}_{n=1}^{N_Z} \overset{i.i.d.}{\sim} (p_Z + q_Z)/2$. It is possible to quantify the MC estimation error from this step, however we assume that with enough offline compute power it could be made negligible, such that $\Delta_Z \approx \hat{\Delta}_Z$.

We did several optimizations in order to compute $\tilde{m}_i$ in a time efficient manner. The first is by pre-filtering to only keep x^{Δ}_n for which $\Delta_Z(x^{\Delta}_n) > \Delta_{Thresh}$. The second optimization is to only consider sampled states within a truncation distance of d_T from state particle x^j_i in (16). We implemented this by keeping all x^{Δ}_n in a KD-tree for efficient radius queries (Maneewongvatana and Mount 1999). We chose Δ_{Thresh} and d_T such that the error in computing $\tilde{m}_i$ is at most $V^{\max} \cdot 10^{-4}$. Lastly, since the runtime complexity of $\tilde{m}_i$ grows linearly with C, the number of particles in the belief $\bar{b}_i$, we limit the number of particles used to N_x by performing MC estimation w.r.t. the particle belief: $\hat{\tilde{m}}_i(\bar{b}_i, a) = \frac{1}{N_x}\sum_{j=1}^{N_X} \tilde{m}_i(x^j_i, a)$ where $\{x^j_i\}_{j=1}^{N_X} \overset{i.i.d.}{\sim} \bar{b}_i$.

Empirical Bound Evaluation

In all scenarios, a planning session is fixed at 500 simulations of PFT-DPW with the same parameters, as described in the supplementary material.

We record for each time step the estimated expected value $\hat{Q}^{[p_Z/q_Z]}_t$ and the plan session duration. In simplified planning, we also record the expected action cumulative bound $\hat{\Phi}_{\mathbf{M_P},t}$. The policy is maximization of the estimated action value function based on the original or simplified model, i.e. $\pi^{[p_Z/q_Z]}_t \triangleq \max_a \hat{Q}^{[p_Z/q_Z]}_t(\bar{b}_t, a)$, and we respectively name them as the original or simplified value policy. After each planning session, we apply the selected action, update the particle belief according to the observation received, and continue the scenario until the POMDP terminates.

In our first test, we run 100 scenarios of each planning scheme - once with the original observation model, and once with the simplified with the additional computation of $\hat{\Phi}_{\mathbf{M_P},t}$. As shown in figure 4, the simplified planning time is greatly reduced for the same number of simulations compared to planning with the original model. These results were expected because of increased overhead for sampling and evaluating the PDF of the original observation model, leading to longer planning times even when additionally computing $\hat{\Phi}$ in the simplified planning.

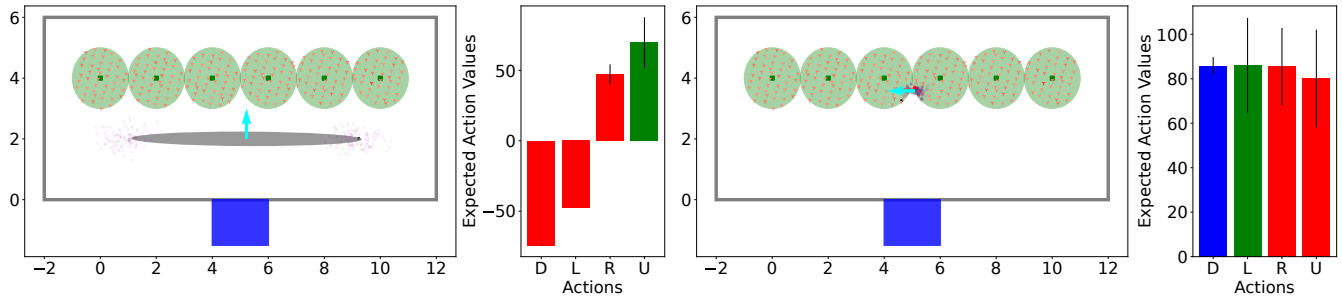

(a) Start of scenario, $t = 0$.

(b) Middle of scenario, $t = 4$.

Figure 3: The results of two planning sessions in 2D beacons. The goal is indicated by the blue rectangle, the beacons and their radii by the green squares and circles, and the outer walls by the grey outer rectangle. The filtered delta states $\{x_n^{\Delta}\}_{n=1}^{N_{\Delta}^{kept}}$ are indicated by the tri-downs, with color relative to estimated TV-distance of the simplified observation model $7.06 \leq \hat{\Delta}_Z \cdot 10^2 \leq 12.08$. The colored dots indicate: the true state in black, the observation in red and the particle belief in purple. The belief empirical mean and covariance are the grey ellipse. The bars on the right depict $\hat{Q}_t^{q_Z}$ for all actions, and $\hat{\Phi}_t$ as symmetric error bars. The action chosen by the simplified value policy $\pi_t^{q_Z}$ is colored in green, and the lower bound policy $\pi_t^{\mathcal{LB}}$ in blue if different. At $t = 4$ we can see an inconsistency of action order between $\pi_t^{q_Z}$ that chooses left, whereas $\pi_t^{\mathcal{LB}}$ chooses down.

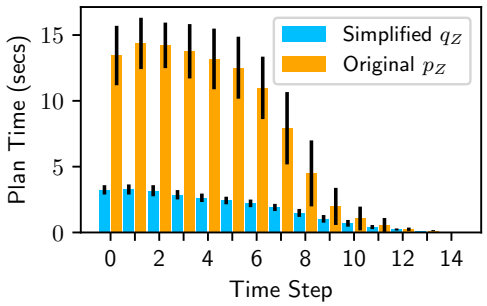

Figure 4: Mean and standard deviation of planning duration over 100 scenarios vs. scenario time step, with the original observation model p_Z or simplified model q_Z.

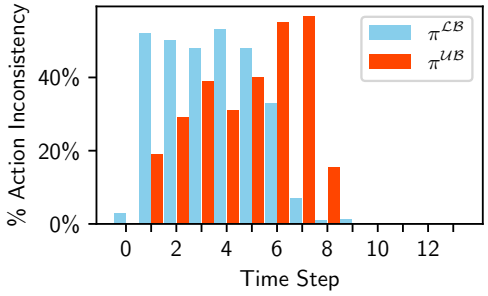

Figure 5: Percentage of scenarios in each time step in which the lower or upper bound policies, $\pi^{\mathcal{LB}}$ or $\pi^{\mathcal{UB}}$, chose an action different from the simplified value policy π^{q_Z}.

In our second test, we quantified how often the simplified policy is different from the lower or upper bound policies. We denote the lower bound policy as $\pi_t^{\mathcal{LB}} \triangleq \arg\max_a \{\hat{Q}_t^{q_Z}(\bar{b}_t, a) - \hat{\Phi}_t(\bar{b}_t, a)\}$ and the upper bound policy as $\pi_t^{\mathcal{UB}} \triangleq \arg\max_a \{\hat{Q}_t^{q_Z}(\bar{b}_t, a) + \hat{\Phi}_t(\bar{b}_t, a)\}$. In our results in figure 5, we can see that there is a great difference between the policies, in particular in time steps 1-7, which is when the agent is mostly around the light region. $\pi^{\mathcal{LB}}$ tends to steer away from the light region, hence it differs mostly in the earlier time steps, whereas $\pi^{\mathcal{UB}}$ prefers the light region, hence it chooses to stay there during the descent of the agent towards the goal. These results indicate that the bounds are non-trivial, and corresponding policies do point towards different objectives as the scenario progresses.

Conclusion

This paper builds upon the paradigm of solving POMDP planning problems with simplification for adhering to computational limitations. We suggest a planning framework with a simplified observation model, to the end of practically solving POMDPs with complex high dimensional observations like visual observations, with finite time performance guarantees. We formulate a novel bound based on local reweighting of pre-calculated TV-distances at pre-sampled states, and show that its estimator bounds with high probability the theoretical value function of the original problem. Finally, an example showcasing how the bounds can influence the decision-making process is presented for both the lower and upper bound policies. In future research, we envision the use of our bounds during the planning process itself, for pruning of action branches to the end of runtime improvement, or for certifying performance when they're computed explicitly.

Acknowledgements

This work was supported by the Israel Science Foundation (ISF) and by US NSF/US-Israel BSF.

References

Bar-Shalom, Y.; Li, X. R.; and Kirubarajan, T. 2004. *Estimation with applications to tracking and navigation: theory algorithms and software*. John Wiley & Sons.

Barenboim, M.; Lev-Yehudi, I.; and Indelman, V. 2023. Data Association Aware POMDP Planning with Hypothesis Pruning Performance Guarantees. *IEEE Robotics and Automation Letters (RA-L)*.

Caflisch, R. E. 1998. Monte carlo and quasi-monte carlo methods. *Acta numerica*, 7: 1–49.

Deglurkar, S.; Lim, M. H.; Tucker, J.; Sunberg, Z. N.; Faust, A.; and Tomlin, C. 2023. Compositional Learning-based Planning for Vision POMDPs. In *Learning for Dynamics and Control Conference*, 469–482. PMLR.

Doucet, A.; de Freitas, N.; and Gordon, N., eds. 2001. *Sequential Monte Carlo Methods In Practice*. New York: Springer-Verlag.

Durrett, R. 2019. *Probability: theory and examples*, volume 49. Cambridge university press.

Eslami, S. A.; Rezende, D. J.; Besse, F.; Viola, F.; Morcos, A. S.; Garnelo, M.; Ruderman, A.; Rusu, A. A.; Danihelka, I.; Gregor, K.; et al. 2018. Neural scene representation and rendering. *Science*, 360(6394): 1204–1210.

Ha, D.; and Schmidhuber, J. 2018. Recurrent world models facilitate policy evolution. *Advances in Neural Information Processing Systems (NIPS)*, 31.

Hoerger, M.; and Kurniawati, H. 2021. An On-Line POMDP Solver for Continuous Observation Spaces. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 7643–7649. IEEE.

Hoerger, M.; Kurniawati, H.; Bandyopadhyay, T.; and Elfes, A. 2020. Linearization in motion planning under uncertainty. In *Intl. Workshop on the Algorithmic Foundations of Robotics (WAFR)*, 272–287. Springer.

Hoerger, M.; Kurniawati, H.; and Elfes, A. 2023. Multilevel Monte Carlo for solving POMDPs on-line. In *Intl. J. of Robotics Research*, volume 42, 196–213. Sage Publications Sage UK: London, England.

Igl, M.; Zintgraf, L.; Le, T. A.; Wood, F.; and Whiteson, S. 2018. Deep variational reinforcement learning for POMDPs. In *International Conference on Machine Learning*, 2117–2126. PMLR.

Jonschkowski, R.; Rastogi, D.; and Brock, O. 2018. Differentiable Particle Filters: End-to-End Learning with Algorithmic Priors. In *Robotics: Science and Systems (RSS)*.

Kaelbling, L. P.; Littman, M. L.; and Cassandra, A. R. 1998. Planning and acting in partially observable stochastic domains. *Artificial intelligence*, 101(1): 99–134.

Karkus, P.; Hsu, D.; and Lee, W. S. 2017. Qmdp-net: Deep learning for planning under partial observability. In *Advances in Neural Information Processing Systems (NIPS)*, 4694–4704.

Karkus, P.; Hsu, D.; and Lee, W. S. 2018. Particle Filter Networks with Application to Visual Localization. In *Conference on Robot Learning*.

Kearns, M.; Mansour, Y.; and Ng, A. Y. 2002. A sparse sampling algorithm for near-optimal planning in large Markov decision processes. *Machine learning*, 49(2): 193–208.

Kingma, D. P.; and Welling, M. 2014. Auto-encoding Variational Bayes. In *International Conference on Learning Representations (ICLR)*.

Kong, A.; Liu, J. S.; and Wong, W. H. 1994. Sequential imputations and Bayesian missing data problems. *Journal of the American Statistical Association*, 89(425): 278–288.

Lemieux, C. 2009. *Monte Carlo and Quasi-Monte Carlo Sampling*. Springer.

Lev-Yehudi, I.; Barenboim, M.; and Indelman, V. 2023. Simplifying Complex Observation Models in Continuous POMDP Planning with Probabilistic Guarantees and Practice. arXiv:2311.07745.

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. *Journal of Artificial Intelligence Research*, 77: 1591–1636.

Lim, M. H.; Tomlin, C.; and Sunberg, Z. N. 2020. Sparse Tree Search Optimality Guarantees in POMDPs with Continuous Observation Spaces. In *Intl. Joint Conf. on AI (IJCAI)*, 4135–4142.

Lim, M. H.; Tomlin, C. J.; and Sunberg, Z. N. 2021. Voronoi progressive widening: efficient online solvers for continuous state, action, and observation POMDPs. In *2021 60th IEEE conference on decision and control (CDC)*, 4493–4500. IEEE.

Maneewongvatana, S.; and Mount, D. M. 1999. It's okay to be skinny, if your friends are fat. In *Center for geometric computing 4th annual workshop on computational geometry*, volume 2, 1–8. Citeseer.

Papadimitriou, C.; and Tsitsiklis, J. 1987. The complexity of Markov decision processes. *Mathematics of operations research*, 12(3): 441–450.

Rezende, D.; and Mohamed, S. 2015. Variational inference with normalizing flows. In *Intl. Conf. on Machine Learning (ICML)*, 1530–1538. PMLR.

Roberts, M. 2018. The Unreasonable Effectiveness of Quasirandom Sequences. https://extremelearning.com.au/unreasonable-effectiveness-of-quasirandom-sequences/. Accessed: 2023-08-12.

Shienman, M.; and Indelman, V. 2022. D2A-BSP: Distilled Data Association Belief Space Planning with Performance Guarantees Under Budget Constraints. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*.

Silver, D.; and Veness, J. 2010. Monte-Carlo planning in large POMDPs. In *Advances in Neural Information Processing Systems (NIPS)*, 2164–2172.

Sohn, K.; Lee, H.; and Yan, X. 2015. Learning structured output representation using deep conditional generative models. In *Advances in neural information processing systems*, 3483–3491.

Sunberg, Z.; and Kochenderfer, M. 2018. Online algorithms for POMDPs with continuous state, action, and observation

spaces. In *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 28.

Sztyglic, O.; and Indelman, V. 2022. Speeding up Online POMDP Planning via Simplification. In *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*.

Tsybakov, A. B. 2009. *Introduction to nonparametric estimation*. New York ; London : Springer.

Wang, Y.; Liu, B.; Wu, J.; Zhu, Y.; Du, S. S.; Fei-Fei, L.; and Tenenbaum, J. B. 2020. DualSMC: Tunneling Differentiable Filtering and Planning under Continuous POMDPs. In Bessiere, C., ed., *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20*, 4190–4198. International Joint Conferences on Artificial Intelligence Organization. Main track.

Winkler, C.; Worrall, D.; Hoogeboom, E.; and Welling, M. 2019. Learning likelihoods with conditional normalizing flows. *arXiv preprint arXiv:1912.00042*.

Ye, N.; Somani, A.; Hsu, D.; and Lee, W. S. 2017. DESPOT: Online POMDP planning with regularization. *JAIR*, 58: 231–266.