




Monte Carlo Planning in Hybrid Belief POMDPs

Moran Barenboim , Moshe Shienman , *Graduate Student Member, IEEE*, and Vadim Indelman 

Abstract—Real-world problems often require reasoning about hybrid beliefs, over both discrete and continuous random variables. Yet, such a setting has hardly been investigated in the context of planning. Moreover, existing online partially observable Markov decision processes (POMDPs) solvers do not support hybrid beliefs directly. In particular, these solvers do not address the added computational burden due to an increasing number of hypotheses with the planning horizon, which can grow exponentially. As part of this work, we present a novel algorithm, Hybrid Belief Monte Carlo Planning (HB-MCP) that utilizes the Monte Carlo Tree Search (MCTS) algorithm to solve a POMDP while maintaining a hybrid belief. We illustrate how the upper confidence bound (UCB) exploration bonus can be leveraged to guide the growth of hypotheses trees alongside the belief trees. We then evaluate our approach in highly aliased simulated environments where unresolved data association leads to multi-modal belief hypotheses.

Index Terms—Planning under uncertainty, autonomous agents.

I. INTRODUCTION

INTELLIGENT autonomous agents operating in real-world environments often need to reason about a hybrid belief containing discrete and continuous random variables. While the states of the agent and of the environment are commonly represented by continuous random variables, discrete random variables generally represent object classes, data association hypotheses or even transition models (e.g. due to slippage) and observation models. In ambiguous environments, where different objects or scenes can possibly be perceptually similar or identical, such discrete variables are particularly important, as wrong assignments can lead to a complete failure of the agent's task.

In general, all random variables in a hybrid belief are coupled, and the number of hypotheses, i.e. realizations of discrete variables may be combinatorially large with the number of ambiguous objects and classes or even develop exponentially with time given ambiguous data associations. Therefore, without any pruning or merging heuristic, the size of the considered belief quickly becomes prohibitively large and the computational

complexity of the corresponding problem becomes impossible to handle.

The research community has been extensively investigating passive inference approaches where the considered belief is hybrid. In [1] the authors proposed a message passing algorithm to correctly identify loop closures by optimizing a hybrid factor graph [2]. A convex relaxation approach over a discrete-continuous graphical model was presented in [3] to capture perceptual aliasing and find the maximal subset of internally coherent measurements, i.e. correct data association.

In spite of the significant progress made within the SLAM community, such a hybrid setting has received scant attention from the planning community. As such, most off-the-shelf, state-of-the-art POMDP online solvers do not directly support hybrid beliefs. Specifically, [4] introduced POMCP, an adaptation to Monte-Carlo Tree Search (MCTS) for POMDPs using the UCT algorithm [5] to guide the action selection process. POMCPOW and DESPOT [6], [7] employ transition and observation models to efficiently propagate particles from the prior belief, as an efficient approximation for belief update. However, in the context of hybrid beliefs, the belief update may not be as efficient, since it would require knowledge of the hypotheses' probabilities, which are not presumed to be given.

POMDPs can also be converted into belief Markov decision processes (BMDPs) to utilize MDP solvers. PFT-DPW [6] and AI-BSP [8] are two such solvers, where belief-states replace states in the original MDP algorithms. However, performing inference with hybrid belief is hardly efficient due to a large number of hypotheses. For instance, in ambiguous data association scenarios, the number of hypotheses grows exponentially with time, making full inference intractable.

Only recently have hybrid beliefs been explicitly considered in planning. In [9], the authors introduced DA-BSP, which allows reasoning about future data association hypotheses within a belief space planning framework for the first time. [10] suggested reducing the computational complexity of DA-BSP by selecting only a small subset of hypotheses and providing bounds over the loss in solution quality. [10] was later extended to a non-myopic setting, in [11]. The ARAS framework proposed in [12] leveraged the graphical model presented in [13] to reason about ambiguous data association in future beliefs using multi-modal factors to model discrete ambiguities. Due to its high computational burden, these approaches did not aim at closed-loop POMDP planning, neglecting its mathematical soundness.

In this paper we propose an approach to alleviate the computational complexity of planning with hybrid beliefs under the POMDP formulation. We show that previous algorithms result in biased estimators of the reward and value function, and suggest a different way for controlling the number of hypotheses to a manageable size. Utilizing sequential importance resampling (SIR) for hypothesis selection, we suggest an algorithm that results in

Manuscript received 14 November 2022; accepted 1 May 2023. Date of publication 5 June 2023; date of current version 14 June 2023. This letter was recommended for publication by Associate Editor Lawson L.S. Wong and Editor H. Kurniawati upon evaluation of the reviewers' comments. This work was supported in part by U.S. NSF/US-Israel BSF and in part by Israeli Smart Transportation Research Center (ISTRC). (*Corresponding author: Moran Barenboim.*)

Moran Barenboim and Moshe Shienman are with the Technion Autonomous Systems Program (TASP), Technion - Israel Institute of Technology, Haifa 32000, Israel (e-mail: moranbar@campus.technion.ac.il; smoshe@campus.technion.ac.il).

Vadim Indelman is with the Department of Aerospace Engineering, Technion - Israel Institute of Technology, Haifa 32000, Israel (e-mail: vadim.indelman@technion.ac.il).

Digital Object Identifier 10.1109/LRA.2023.3282773




Fig. 1. Nodes in the tree correspond to hybrid beliefs. Inner shapes illustrate different continuous distributions, each correspond to a different discrete variable. (a) A belief tree which computes a full hybrid belief (shown in blue) at each iteration, regardless of the hypotheses significance. (b) An adaptive approach (ours) that simultaneously generates both hypotheses and nodes in the belief tree sampled according to their probabilities.

an unbiased estimator and efficient belief tree construction, see Fig. 1 for an illustration. We show that the algorithm supports both state-dependent and belief-dependent rewards. We proceed with a contribution to inference in the setting of ambiguous data association, by introducing a natural way to incorporate negative information within Bayesian inference, and demonstrate how the hypotheses weights should be updated. Last, we demonstrate our approach on simulative environments to corroborate our findings.

Our contributions in this paper are as follows: (a) We introduce a novel algorithm that performs Monte-Carlo planning to solve a POMDP when the considered belief is hybrid. (b) We show that our algorithm, HB-MCP, leads to an unbiased utility estimate, in contrast to existing hybrid belief algorithms. (c) We introduce negative information to hybrid belief inference. (d) We demonstrate the effectiveness of our algorithm in extremely aliased simulated environments where unresolved data association leads to multi-modal belief hypotheses. This paper is accompanied by supplementary material [14] that provides proofs and further implementation and experimental details.

II. PRELIMINARIES

In this section we formally define a POMDP and a general hybrid belief, which will be used in the following sections.

A. Partially Observable Markov Decision Process

A discrete-time POMDP can be formally defined as a tuple $(\mathcal{X}, \mathcal{A}, \mathcal{Z}, T, O, \mathcal{R})$, where \mathcal{X} , \mathcal{A} and \mathcal{Z} denote the state, action and observation spaces respectively; $T(x, a, x') \triangleq \mathbb{P}(x'|x, a)$ is the transition density function which expresses the probability to move from state $x \in \mathcal{X}$ to state $x' \in \mathcal{X}$ by taking action $a \in \mathcal{A}$; $O(x, z) \triangleq \mathbb{P}(z|x)$ is the observation density function which expresses the probability to receive an observation $z \in \mathcal{Z}$ from state $x \in \mathcal{X}$; and \mathcal{R} is a user defined reward function.

As observations provide only partial information about the state, the true state of the agent is unknown. Therefore, the agent maintains a probability distribution function over the state space, also known as a belief. At each time step t the belief update is performed according to Bayes rule, using the transition and observation models, given the performed action a_{t-1} and the received observation z_t as $b_t(x') = \eta \int \mathbb{P}(z_t|x')\mathbb{P}(x'|x, a_{t-1})b_{t-1}(x)dx$, where η is a normalization constant.

Given a posterior belief b_t , a policy function $a_t = \pi(b_t)$ determines an action to be taken at time step t . For a finite horizon \mathcal{T} the value function for a policy π is defined as the expected cumulative reward received by executing π ,

$$V^\pi(b_t) = \mathcal{R}(b_t, \pi(b_t)) + \mathbb{E}_{z_{t+1}:\mathcal{T}} \left[\sum_{\tau=t+1}^{\mathcal{T}} \mathcal{R}(b_\tau, \pi(b_\tau)) \right]. \quad (1)$$

Similarly, an action-value function,

$$Q^\pi(b_t, a_t) = \mathcal{R}(b_t, a_t) + \mathbb{E}_{z_{t+1}} [V^\pi(b_{t+1})], \quad (2)$$

is defined by executing action a_t and then following the policy π for a finite horizon \mathcal{T} . At each planning session, the agent solves a POMDP by searching for the optimal policy π^* that maximizes (1).

B. Hybrid Belief

A hybrid belief is defined over both continuous and discrete random variables. The continuous random variables can represent the state of the agent and (possibly also) of the environment, as common in SLAM framework. The discrete random variables can represent, e.g., object classes and/or data association hypotheses. Nevertheless, the following definition is general and not restricted to these examples.

We formally define the hybrid belief at each time t as

$$b_t \triangleq \mathbb{P}(X_t, \beta_{0:t} | H_t) = \underbrace{\mathbb{P}(X_t | \beta_{0:t}, H_t)}_{b[X_t]_{\beta_{0:t}}} \underbrace{\mathbb{P}(\beta_{0:t} | H_t)}_{b[\beta_{0:t}]_{\omega_t}}, \quad (3)$$

where $X_t \triangleq \{x_0, \dots, x_t\}$, $\beta_{0:t}$ denote the discrete random variables and $H_t \triangleq \{z_{1:t}, a_{0:t-1}\}$ represents all past actions and observations. $b[X_t]_{\beta_{0:t}}$ is the conditional belief over continuous variables. ω_t is the marginal belief over discrete variables which can be considered as the hypothesis weight. We define $H_{t+1}^- \triangleq H_t \cup \{a_t\}$ and $b_{t+1}^- \triangleq \mathbb{P}(X_{t+1}, \beta_{0:t+1} | H_{t+1}^-)$ for notational convenience.

The marginal belief ω_t is updated for each realization of discrete random variables according to

$$\omega_t^{i,j} = \frac{\overbrace{\mathbb{P}(z_t | \beta_{0:t}^{i,j}, H_t^-)}^{\zeta_t^{i,j}} \overbrace{\mathbb{P}(\beta_{0:t-1}^j | H_t^-)}^{\omega_{t-1}^j} \mathbb{P}(\beta_{0:t-1}^j | H_t^-)}{\eta}, \quad (4)$$

which is obtained by Bayes rule followed by chain rule on ω_t . The un-normalized weight can be expressed recursively as $\tilde{\omega}_t^{i,j} = \zeta_t^{i,j} \omega_{t-1}^j$. We denote β_t^i and $\beta_{0:t-1}^j$ as the realization of β_t and $\beta_{0:t-1}$ respectively and $\beta_{0:t-1}^{i,j}$ denotes the realization of the joint variables $\beta_{0:t-1}$. $\zeta_t^{i,j}$ denotes the conditional dependence of ζ_t on β_t^i given $\beta_{0:t-1}^j$. This notation proceeds similarly for any random variable. The conditional belief $b[X_t]_{\beta_{0:t}}$ is updated for each realization of discrete random variables as

$$b[X_t]_{\beta_{0:t}}^{i,j} = \psi(b[X_t]_{\beta_{0:t-1}}^j, a_{t-1}, z_t), \quad (5)$$

where $\psi(\cdot)$ represents the Bayesian inference method.

Generally, when planning with hybrid beliefs the agent constructs both a belief tree and multiple hypotheses trees. Each

hypotheses tree represent the posterior hypotheses given a history. Since every node of the planning tree (i.e. belief tree) corresponds to a hypotheses tree, the computational complexity of the corresponding POMDP becomes a significant burden. In the following section we present a novel algorithm that circumvent this difficulty via Monte-Carlo sampling.

C. Ambiguous Data Association

One use case for hybrid beliefs is ambiguous data associations (DA), which introduces a significant computational burden in both inference and even more so in planning. In inference, the number of possible hypotheses grows exponentially with the number of steps the agent took, i.e. the length of the history array. In planning, the difficulty further increases due to the exponential number of histories considered by the agent.

To address ambiguous DA, (4) can be adapted to,

$$\omega_t^{i,j} = \tilde{\zeta}_t^{j|i} \omega_{t-1}^i, \quad (6)$$

where,

$$\tilde{\zeta}_t^{j|i} = \frac{\zeta_t^{j|i}}{\sum_i \zeta_t^{j|i} \omega_{t-1}^i},$$

$$\zeta_t^{j|i} = \int_{X_t} \mathbb{P}(z_t | X_t, \beta_t^i) \mathbb{P}(\beta_t^i | X_t) \mathbb{P}(X_t | \beta_{0:t-1}^j, H_t^-). \quad (7)$$

The latter is obtained by marginalizing $\zeta_t^{i|j}$ over the states, and adhering to the Markov assumption of the observation and association ($\mathbb{P}(\beta_t | X_t)$) models, as suggested in [9]. Our hybrid belief planner is tested in section VII using ambiguous DA as a challenging use case.

III. POMDP PLANNING WITH HYBRID BELIEFS

In section III-A we start with a brief overview of how MCTS can be utilized to solve POMDPs with hybrid beliefs and its drawbacks. Then, in section III-B we present a novel approach to utilize the UCT exploration bonus to build an asymmetric hypotheses tree, which leads to better use of the computational resources by focusing on the most promising hypotheses according to the UCT bonus.

A. Vanilla Hybrid-Belief MCTS

For completeness, we first present a vanilla-HB-MCTS algorithm. Although the exact algorithm does not seem to exist in the literature, this is the ad-hoc way to interleave hybrid beliefs with state-of-the-art POMDP solvers. vanilla-HB-MCTS, can be seen as an adaptation of the state-dependent MCTS [4] algorithm to a (hybrid-)belief (3), by augmenting the belief to a belief-state. A similar approach was also taken by PFT-DPW [6], which utilized particle filters to approximate a posterior belief, over continuous variables. However, computing a full hybrid belief is a difficult and sometimes intractable task, even for particle-based solvers, and is thus prone to approximations.

Pruning: The number of hypotheses at each posterior node in the belief tree may be prohibitively large. To handle the infeasible number of the posterior hypotheses, vanilla-HB-MCTS utilizes a pruning mechanism similar to those suggested in [9], [13]. As a result, unlikely hypotheses are removed from the hypotheses tree.

In vanilla-HB-MCTS, each posterior node holds a fixed number of hypotheses once expanded, depending on a predefined hyperparameter. Such a method may sometimes be too harsh, pruning away hypotheses with high probability due to a limited hypotheses budget, or too loose, keeping highly unlikely hypotheses, thus wasting valuable computational time. Other approaches may also be applicable, such as fixing a probability threshold value, under which all hypotheses are pruned. However, the latter has its own deficiencies, such as hypothesis depletion. For completeness, we describe vanilla-HB-MCTS implementation details in the supplementary [14].

B. Hybrid Belief Monte-Carlo Planning

In contrast to vanilla-HB-MCTS, in HB-MCP, we do not use any pruning heuristic for two reasons: (1) this requires knowledge, or an insight, as to how many hypotheses would be sufficient for the specific POMDP; (2) Each posterior node in the belief tree maintains hypotheses based on a hyperparameter, regardless of how relevant this node may be for decision-making.

Conversely, we suggest an adaptive algorithm that focuses computational resources in proportion to their relevance in the belief tree, which circumvent the difficulty in full belief update. HB-MCP is recursively invoked with a single sampled hypothesis. Every such single hypothesis may evolve into multiple hypotheses. HB-MCP algorithm computes only the posterior weights (i.e. probability values) that are conditioned on that single hypothesis, followed by a random weight sample based on their categorical distribution. Then, only the hypothesis associated with the sampled weight is updated. This is in contrast to the full posterior update done in vanilla-HB-MCTS.

Additionally, to support belief-dependent rewards, the reward value is estimated based on state samples received across multiple visits to the belief node, i.e., state samples from multiple hypotheses. We describe the algorithm details in Section IV.

HB-MCP holds some desirable properties compared to the full belief update and pruning approaches. First, at each iteration of HB-MCP, a maximum of \mathcal{T} posterior hypotheses are computed, and a small subset of the weights. This is in contrast to the full posterior update, that would require the entire (or pruned-)set of the current posteriors, and compute all the posterior hypotheses of the next time-step, which is highly resource expensive for every iteration. Second, HB-MCP explores both the planning tree and the hypotheses trees by focusing its computational effort on the interesting parts, utilizing UCB to guide the search; this property is inspired by MCTS which builds the planning tree by focusing on the optimistic parts of the tree. In Section V, we show that this approach results in an unbiased estimator for the true value function.

IV. IMPLEMENTATION DETAILS

In this section we describe the implementation details of our approach, HB-MCP, as discussed in section III-B.

HB-MCP can be described as follows; first, it starts by receiving a single hypothesis and selecting an immediate action according to UCB exploration bonus. Then, samples are generated and appended to $B(h)$, which are later used for reward

Algorithm 1: HB-MCP.

Procedure: SIMULATE(b_t^j, h, d)

- 1: **if** $d = 0$ **then**
- 2: **return** 0
- 3: $a \leftarrow \arg \max_{\bar{a}} Q(h\bar{a}) + c\sqrt{\frac{\log(N(h))}{N(h\bar{a})}}$
- 4: $B(h) \leftarrow \text{GETSAMPLES}(b_t^j, B(h), N(h))$
- 5: $r \leftarrow \text{REWARD}(B(h), a)$
- 6: $r \leftarrow r + N(h)(r - r_{prev})$
- 7: **if** $|C(ha)| \leq k_o N(ha)^{\alpha_o}$ **then**
- 8: $z \leftarrow \text{SAMPLEOBSERVATION}(b_t^j, a)$
- 9: **else**
- 10: $z \leftarrow \text{Sample uniformly from } C(ha)$
- 11: $\{\omega_{t+1}^{i,j}\}_{i=1}^L \leftarrow \text{COMPUTEWEIGHTS}(b_t^j, a, z)$
- 12: $i \leftarrow \text{SAMPLECATEGORICAL}(\{\omega_{t+1}^{i,j}\}_{i=1}^L)$
- 13: $b_{t+1}^{i,j} \leftarrow \Psi(b_t^j, a, z, i) // (5)$
- 14: **if** $z \notin C(ha)$ **then**
- 15: $C(ha) \cup \{z\}$
- 16: $R \leftarrow r + \text{ROLLOUT}(b_{t+1}^{i,j}, d - 1)$
- 17: **else**
- 18: $R \leftarrow r + \text{SIMULATE}(b_{t+1}^{i,j}, haz, d - 1)$
- 19: $N(h) \leftarrow N(h) + 1$
- 20: $N(ha) \leftarrow N(ha) + 1$
- 21: $Q(ha) \leftarrow Q(ha) + \frac{R - Q(ha)}{N(ha)}$
- 22: **return** R

estimation (lines 4-6). Lines 7-10 perform observation progressive widening. Then, the approach for sampling hypotheses is shown in lines 11-13. Note that the algorithm directly computes *all* the weights conditioned on the hypothesis given as input (line 11). Then, we resample a *single* conditional belief, $b_{t+1}^{i,j}$, sampled according to the weights (line 12). We note that this is not a necessity, and different number of samples can be taken in those two steps to trade-off efficiency and accuracy. Depending on whether a new posterior node is sampled or not, lines 14-18 either call for rollout or continues recursively. Last, the action-value function and the counters are updated.

To estimate a belief-dependent reward, state samples should correspond to their likelihood in the full hybrid belief. In HB-MCP, hypotheses are generated iteratively, accumulating hypotheses (or, equivalently, state samples from those hypotheses), so that at each iteration the reward estimator is improved. Generally, a belief dependent reward is not a simple average over samples. However, as in MCTS, HB-MCP estimates the action value function, $Q(ha)$, as an average of all the cumulative returns passed through that node. To support belief dependent rewards, HB-MCP computes a new reward estimate based on all past samples, and replaces the previous reward estimate with the new one. To that end, a simple recursive subtraction and addition update is done for every node encountered along the path of the current iteration, described in line 6.

V. THEORETICAL ANALYSIS

In this section, we first claim that existing approximations, done in contemporary state-of-the-art multi-hypotheses planners, such as DA-BSP [9], ARAS [12] as well as vanilla-HB-MCTS (Section III-A), lead to a biased estimation of the

reward value, and therefore a biased value function. Further, we show that even if the reward value could be precisely recovered, the resultant value function is generally biased. Instead, HB-MCP performs sequential sampling which converges to the correct value. Then, we discuss how HB-MCP may also support belief-dependent reward functions and its applicability for value function estimation.

A. State-Dependent Rewards

State-dependent reward functions are defined as the expected reward value over the belief, i.e., $\mathcal{R}_X \triangleq \mathbb{E}_{X \sim b}[r(X, a)]$. Generally, state-dependent rewards cannot be computed analytically, thus, they are approximated using state samples. Since in a hybrid belief the number of hypotheses may be prohibitively expensive to compute, most existing algorithms approximate the belief, \hat{b} , by performing some heuristic pruning. As a consequence, the approximate distribution is shifted, and the reward value is biased even with an infinite number of state samples,

Lemma 1: The estimator $\mathbb{E}_{X \sim \hat{b}}[r(X, a)]$ is biased.

Proof: Assuming the weights of the pruned hypotheses are non-zero, the proof is immediate,

$$\begin{aligned}
 \mathbb{E}_{X \sim b}[r(X, a)] &= \int_X \sum_{\beta} b(X, \beta) r(X, a) dX \\
 &= \int_X \sum_{\beta \in A} b(X, \beta) r(X, a) dX + \sum_{\beta \in \bar{A}} b(X, \beta) r(X, a) dX \\
 &\neq \eta_A \int_X \sum_{\beta \in A} b(X, \beta) r(X, a) dX = \mathbb{E}_{X \sim \hat{b}}[r(X, a)]. \quad (8)
 \end{aligned}$$

where A denotes the set of un-pruned hypotheses, and η_A is their corresponding normalizer after pruning. ■

In contrast, HB-MCP samples hypotheses iteratively starting from the root node; it utilizes sequential importance resampling, which results in an unbiased estimator for the reward value. At every iteration, the new sampled states from the current hypothesis are added to the estimator from previous iterations, by averaging. The process for generating hypotheses can be described as follows; for any time t , a hypothesis is sampled i.i.d from a proposal-prior distribution, $\beta_0^i \sim \mathbb{Q}(\beta_0 | H_0)$. Then, hypotheses are recursively sampled from a proposal distribution, $\beta_\tau^i \sim \mathbb{Q}(\beta_\tau | \beta_{0:\tau-1})$ up to time $\tau = t$. We define $\mathbb{Q}(\beta_0 | H_0) \triangleq \mathbb{P}(\beta_0 | H_0)$, and $\mathbb{Q}(\beta_\tau | \beta_{0:\tau-1}) \triangleq \text{UNIFORM}[1, |\beta_\tau|]$. Then, for every time-step t , the corresponding importance weight is,

$$\begin{aligned}
 \lambda_t^{i,j} &= \frac{\mathbb{P}(\beta_{0:t}^{i,j} | H_t)}{\mathbb{Q}(\beta_{0:t}^{i,j} | H_0)} = \frac{\eta_t \zeta_t^{i|j} \mathbb{P}(\beta_{0:t-1}^j | H_{t-1})}{\mathbb{Q}(\beta_t^i | \beta_{0:t-1}^j) \mathbb{Q}(\beta_{0:t-1}^j | H_0)} \\
 &= \frac{\eta_t \zeta_t^{i|j} \mathbb{P}(\beta_{0:t-1}^j | H_{t-1})}{1/|\beta_t^{i|j}| \mathbb{Q}(\beta_{0:t-1}^j | H_0)} = \eta_t \zeta_t^{i|j} |\beta_t^{i|j}| \lambda_{t-1}^j, \quad (9)
 \end{aligned}$$

where $\lambda_0^j = 1$. As a consequence,

Lemma 2: HB-MCP state-dependent reward estimator, $\hat{\mathcal{R}}_X \triangleq \frac{1}{N} \sum_{i,j=1}^N \lambda_t^{i,j} \frac{1}{n_X} \sum_{k=1}^{n_X} r(X_t^{i,j,k}, a_t)$, is unbiased.

Proof: If states are sampled i.i.d. for each hypothesis, then the expected value of the reward estimator, $\hat{\mathcal{R}}_X$, is,

$$\begin{aligned} \mathbb{E}[\hat{\mathcal{R}}_X] &\triangleq \mathbb{E} \left[\frac{1}{N} \sum_{i,j=1}^N \lambda_t^{i,j} \frac{1}{n_X} \sum_{k=1}^{n_X} r(X_t^{i,j,k}, a_t) \right] \\ &= \mathbb{E}_{\mathbb{Q}} \left[\frac{1}{N} \sum_{i,j=1}^N \lambda_t^{i,j} \mathbb{E}_{b[X_t]_{\beta_{0:t}}^{i,j}} \left[\frac{1}{n_X} \sum_{k=1}^{n_X} r(X_t^{i,j,k}, a_t) \right] \right] \\ &= \frac{1}{N} \sum_{i,j=1}^N \mathbb{E}_{\mathbb{Q}} \left[\frac{\mathbb{P}}{\mathbb{Q}} \frac{1}{n_X} \sum_{k=1}^{n_X} \mathbb{E}_{b[X_t]_{\beta_{0:t}}} [r(X_t^{i,j,k}, a_t)] \right] \\ &= \mathbb{E}_{\mathbb{P}} [\mathbb{E}_{b[X_t]_{\beta_{0:t}}} r(X_t, a_t)] \triangleq \mathcal{R}_X \end{aligned} \quad (10)$$

where $\mathbb{P} = \mathbb{P}(\beta_{0:t} | H_t)$, $\mathbb{Q} = \mathbb{Q}(\beta_{0:t} | H_t)$, and N and n_X denote the number of samples from \mathbb{Q} and $b[X_t]_{\beta_{0:t}}^{i,j}$ respectively. ■

As the planning horizon grows, sampling hypotheses uniformly quickly induce sample degeneracy. That is, the weights of most hypothesis samples become negligible, while only a few remain significant, which negatively affects the accuracy of the estimate. To avoid this issue, we perform resampling at every step, also known as sequential importance resampling (SIR). Before resampling, each hypothesis weight simply becomes, $\lambda_t^{i,j} = \eta_t \zeta_t^{i,j} |\beta_t^{i,j}|$, which is then updated to $1/N$ after resampling. Note that resampling does not introduce bias to the estimator [15]. To avoid repeated derivations, for the rest of this sequel we treat mathematical proofs as if hypotheses are directly sampled from distribution \mathbb{P} , even though they are in fact sampled from the proposal distribution, \mathbb{Q} . However, all derivations can be started by sampling from \mathbb{Q} , then follow similar steps of lemma 2 followed by resampling to arrive at the same result.

In some cases of interest, such as ambiguous DA, the normalizer η_t cannot be easily computed, and so the importance weight, λ_t , cannot be computed. A common practice is to use the self-normalized version of the estimator, i.e. $\tilde{\lambda}_t^{i,j} = \tilde{\lambda}_{t-1}^{i,j} \frac{\zeta_t^{i,j}}{\sum_{i,j} \zeta_t^{i,j}}$, which is no longer unbiased [15]. However, the self-normalizing variation is consistent, meaning it becomes less biased with more samples and converges in probability (denoted \rightarrow^p) to the theoretical value. This is a direct consequence of applying the weak law of large numbers on both the nominator and denominator of the self-normalized estimator,

$$\begin{aligned} \hat{\mathcal{R}}_X^{SN} &\triangleq \frac{\sum_{i,j=1}^N \zeta_t^{i,j} \omega_{t-1}^j \frac{1}{n_X} \sum_{k=1}^{n_X} r(X_t^{i,j,k}, a_t)}{\sum_{i,j=1}^N \zeta_t^{i,j} \omega_{t-1}^j} \\ &= \frac{\frac{1}{N} \sum_{i,j=1}^N \eta_t \zeta_t^{i,j} \omega_{t-1}^j \frac{1}{n_X} \sum_{k=1}^{n_X} r(X_t^{i,j,k}, a_t)}{\frac{1}{N} \sum_{i,j=1}^N \eta_t \zeta_t^{i,j} \omega_{t-1}^j} \rightarrow^p \frac{\mathcal{R}_X}{1}, \end{aligned} \quad (11)$$

where the denominator converges to the sum of weights, $\sum_{i,j} \omega_t^{i,j} = 1$ and the nominator to the reward value.

B. Belief-Dependent Rewards

Contrary to state-dependent rewards, belief dependent rewards are not necessarily linear in the belief, so averaging over state samples from different hypotheses does not guarantee convergence to the theoretical reward value. Moreover, different reward definitions may be functions of not only the states, but also the weights, the conditional beliefs, or the probability density values of the complete theoretical belief (such as Shannon's entropy [10] or differential entropy [8]). To support the various cases, we split our discussion into the parametric case, where the reward can be precisely calculated given a set of parametric conditional beliefs and the corresponding weights, and the nonparametric case, where the reward is estimated based on state and hypothesis samples.

HB-MCP supports belief-dependent rewards by accumulating conditional beliefs across multiple visitations of the same history (i.e. same node in the belief tree). The estimated weight of each conditional belief is the sample frequency of the corresponding

hypothesis. That is, $\hat{\mathbb{P}}(\beta_{0:t}^{i,j} | H_t) \triangleq \hat{\omega}_t^{i,j} = \frac{\sum_{i,j} \mathbf{1}_{\beta=\beta_{0:t}^{i,j}}}{N}$, where N is the number of hypothesis samples, $i, j \in [1, |\beta_{0:t}|]$, $|\beta_{0:t}|$ is the theoretical number of hypotheses at time t and $\mathbf{1}_{\square}$ denotes the indicator function.

Parametric: Assuming a parametric representation for the conditional beliefs, $b[X_t]_{\beta_{0:t}}^{i,j}$, the belief-dependent reward, $\mathcal{R}_b(b_t, a_t)$, is evaluated using the estimated hybrid belief, $\mathcal{R}_b(\hat{b}_t, a_t)$, where $\hat{b}_t = b[X_t]_{\beta_{0:t}} \hat{b}[\beta_{0:t}] \equiv b[X_t]_{\beta_{0:t}} \hat{\mathbb{P}}(\beta_{0:t} | H_t)$, and b_t defined in (3). Applying the hypothesis resampling approach as described in Section V-A, the sample frequency of each hypothesis in \hat{b}_t is unbiased, in other words, in expectation it equals the theoretical weights. Moreover,

Lemma 3: $\mathcal{R}_b(\hat{b}_t, a_t)$ converges in probability to $\mathcal{R}_b(b_t, a_t)$ for any continuous, real-valued function \mathcal{R}_b .

Proof: By the law of large numbers, $\hat{\omega}_t^{i,j}$ is consistent as $N \rightarrow \infty$ for all $i, j \in [1, |\beta_{0:t}|]$,

$$\hat{\omega}_t^{i,j} = \sum_{k=1}^N \frac{\mathbf{1}_{\beta^k=\beta_{0:t}^{i,j}}}{N} \rightarrow^p \mathbb{P}(\beta_{0:t}^{i,j} | H_t) = \omega_t^{i,j}, \quad (12)$$

then, due to the continuous mapping theorem,

$$\mathcal{R}_b(b[X_t]_{\beta_{0:t}} \hat{b}[\beta_{0:t}], a_t) \rightarrow^p \mathcal{R}_b(b[X_t]_{\beta_{0:t}} b[\beta_{0:t}], a_t),$$

that is, $\mathcal{R}_b(\hat{b}_t, a_t)$ is a consistent estimator for $\mathcal{R}_b(b_t, a_t)$. ■

Nonparametric: In the nonparametric case, the reward value is estimated based on state particles, which may correspond to conditional belief estimation via particle filters, or POMDPs with reward functions that have no close-form solution, and are thus approximated via Monte Carlo methods. Then, instead of $\mathcal{R}_b(b_t, a_t)$, an estimator over the reward is used, $\hat{\mathcal{R}}_b(\hat{b}[X_t]_{\beta_{0:t}} \hat{b}[\beta_{0:t}], a_t)$, where both the belief and the reward functions are estimators. We denote $\hat{b}[X_t]_{\beta_{0:t}} = \sum_{i=1}^{n_x} \alpha_t^{i,k} \delta(X - X_t^{i,k})$, where $\alpha_t^{i,k}$ is the weight of state particle i generated from conditional belief k and n_x is the number of particles used to approximate the conditional belief. To arrive at consistency results for an arbitrary nonparametric reward estimator, we assume that the reward estimator based

on samples from the full theoretical belief is consistent, i.e., $\hat{\mathcal{R}}_b(\hat{b}[X_t]_{\beta_{0:t}} b[\beta_{0:t}], a_t) \rightarrow^p \mathcal{R}_b(b_t, a_t)$.

Lemma 4: If $\hat{\mathcal{R}}_b(\hat{b}[X_t]_{\beta_{0:t}} b[\beta_{0:t}], a_t) \rightarrow^p \mathcal{R}_b(b_t, a_t)$, then $\hat{\mathcal{R}}_b(b[X_t]_{\beta_{0:t}} \hat{b}[\beta_{0:t}], a_t) \rightarrow^p \mathcal{R}_b(b_t, a_t)$.

Proof: The proof follows similar steps to lemma 3. ■

C. Value Function

When using the existing hypotheses pruning approximations, the estimated value function converges to the wrong value even when some external source provides the exact reward value. This is due to the way observations are generated. The value function is defined as

$$V^\pi(b_t) = \int_z \mathbb{P}(z_{t+1:\tau} | H_t^-) \sum_{\tau=t}^{\tau} \mathcal{R}(b_\tau, \pi_\tau) dz, \quad (13)$$

and since there is usually no direct access to observations given history, first state-samples are generated, then observations are sampled using the observation model, that is, $\mathbb{P}(z_t | H_t^-) = \sum_{\beta} \int_X \mathbb{P}(z_t | X_t, \beta_{0:t}) b^-(X_t, \beta_{0:t})$. Replacing b^- with its pruned counterpart, \hat{b}^- , results in a shifted distribution for both the belief and the measurements, which impacts the value function estimation. Proof of this claim is similar to that of lemma 1 and skipped here for conciseness.

Instead, HB-MCP generates observations by first receiving a hypothesis from the belief at the current node, $\beta_{0:t}^j$. Conditioned on $\beta_{0:t}^j$ and the history, HB-MCP samples a new plausible hypothesis, β_{t+1}^i . Then, an observation is sampled based on the posterior hypothesis. More formally,

$$\begin{aligned} \mathbb{E}_{z_{t+1:\tau}} \left[\sum_{\tau=t+1}^{\tau} \mathcal{R}_\tau \right] &= \mathbb{E}_{z_{t+1}} [\mathcal{R}_{t+1} + \mathbb{E}_{z_{t+2:\tau}} [V_{t+2}^\pi]] \\ &= \underbrace{\mathbb{E}_{\beta_{0:t}} \mathbb{E}_{\beta_{t+1} | \beta_{0:t}} \mathbb{E}_{z_{t+1} | \beta_{0:t+1}} [\mathcal{R}_{t+1}]}_{\triangleq \alpha_{t+1}} + \mathbb{E} [V_{t+2}^\pi]. \end{aligned} \quad (14)$$

We then define the estimator for the expected reward, $\hat{\alpha}_{t+1}$,

$$\hat{\mathbb{E}}_{\mathbb{Q}} \left[\frac{\mathbb{P}(\beta_{t+1}^i | \beta_{0:t}^j, H_{t+1}^-)}{\mathbb{Q}(\beta_{t+1}^i | \beta_{0:t}^j, H_0)} \lambda_t^j \hat{\mathbb{E}}_{z_{t+1} | \beta_{0:t+1}, H_{t+1}^-} [\hat{\mathcal{R}}_{t+1}] \right] \quad (15)$$

Lemma 5: Given an unbiased reward estimator, $\hat{\mathcal{R}}$, the value-function estimator used in HB-MCP is unbiased.

Proof: Applying similar steps from the proof of lemma 2 on $\hat{\alpha}_{t+1}$, leads to an unbiased value, α_{t+1} . Continuing recursively on the value function yields the desired result. See [14] for further details. ■

VI. NEGATIVE INFORMATION IN AMBIGUOUS DATA ASSOCIATION

Just like observations affect the hypotheses' weights, not receiving an expected observation also affects the weights, commonly known as negative information. We build on previous work [9] which addresses hybrid Bayesian inference for ambiguous DA and shows how the mathematical formulation naturally extends to include negative information. We limit our discussion of negative information to the context of landmark-based observations. We conjecture that this formulation can also

be adapted to arbitrary observations, but is out of the scope of this paper.

Negative information is based on not receiving an observation from a mapped landmark. We denote $|L_t| \in \mathbb{N}$ as the number of mapped landmarks at time instant t . This usually refers to the number of landmarks that already exist in the agent state (but can be defined otherwise). We also define observation as, $z_t = [z_t^1, \dots, z_t^{|L_t|}]$. Note that there are $|L_t|$ observation elements in the observation, even though usually not all landmarks can be observed at a single time step, as some might be out of the sensing range due to limited field of view, occlusions, and so on. If at time t only $n_{z_t} < |L_t|$ landmarks are observed, we fill the rest of the observation array with $z_t^k = \infty$, i.e., out of sensing range. Then, the observation array becomes $z_t = [z_t^1, \dots, z_t^{n_{z_t}}, \infty, \dots, \infty]_{1 \times |L_t|}$. The reason for such uncommon inflation of the observation array will become clear shortly.

We define $\beta_t = [\beta_{t,1}, \dots, \beta_{t,|L_t|}]$ as an array that subscribes each landmark with some observation. For example, $\beta_{t,k} = 1$ associates landmark l^k with observation-element z_t^1 from z_t . Note that by the definition of the observation array, $z_t^{\beta_{t,k}} = \infty$ for all $\beta_{t,k} > n_{z_t}$, which does not correspond to any real observation.

Equipped with the definitions of β_t and z_t , we now discuss the adaptation of the observation and association models. We drop the $\square^{i,j}$ notation to avoid notation overloading, the derivations below are true for each hypothesis separately. In the landmark-based context, it is common to further simplify the expression in (7) by assuming conditional independency of an observation given the state variables, to a product of observation models, $\mathbb{P}(z_t | X_t, \beta_t) = \prod_{k=1}^{|L_t|} \mathbb{P}(z_t^{\beta_{t,k}} | x_t^r, l^k)$, where x_t^r and l^k are the current pose of the agent and landmark k . For simplicity, we assume in this paper an ideal detection sensor, in the sense that if a landmark is within range, the sensor will detect it. Under this assumption, likelihood of obtaining an out-of-range observation ($z_t^{\beta_{t,k}} = \infty$), given that the landmark is within the sensing range (denoted $S.R.$), is $\mathbb{P}(z_t^{\beta_{t,k}} = \infty | x_t^r, l^k \in S.R.) = 0$. However, obtaining an out-of-range observation given that the landmark is indeed out of the sensing range, is $\mathbb{P}(z_t^{\beta_{t,k}} = \infty | x_t^r, l^k \notin S.R.) = 1$.

The association model, $\mathbb{P}(\beta_{t,k} | x_t^k, l^k)$, assigns a probability to associate a landmark, l^k , with a specific observation index, $\beta_{t,k}$. We define the likelihood of associating an out-of-sensing-range landmark to an actual observation element (i.e. $\beta_{t,k} \leq n_{z_t}$), as $\mathbb{P}(\beta_{t,k} \leq n_{z_t} | x_t^k, l^k \notin S.R.) = 0$. Conversely, associating a landmark that is within the sensing range, equals a nonzero value, for simplicity defined here as a uniform distribution across all feasible associations, $\frac{1}{n_{z_t}}$. We explicitly state all possible combinations of state, association, and observation in table I.

VII. EXPERIMENTS

In this section we evaluate our approach, HB-MCP, considering multiple hypotheses due to ambiguous DA. We compare our approach with the state of the art algorithms, DA-BSP [9] and PFT-DPW [6]. PFT-DPW is utilized here as a single hypothesis solver, as it does not explicitly support multiple hypotheses beliefs. Its hypothesis is chosen based on the hypotheses weights through sampling. While it is possible to modify PFT-DPW to accommodate multiple hypotheses, we

TABLE I
POSSIBLE COMBINATIONS WHEN CONSIDERING NEGATIVE INFORMATION

$z^{\beta_{t,k}} = \infty$	$\beta_{t,k} > n_{z_t}$	$(x^r, l^k) \in S.R.$	$\mathbb{P}(z x, l)$	$\mathbb{P}(\beta x, l)$
no	no	yes	$f(\cdot)$	1
no	no	no	0	0
yes	yes	no	1	1
yes	yes	yes	0	0
no	yes	yes	$f(\cdot)$	0
no	yes	no	0	1
yes	no	no	1	0
yes	no	yes	0	1

$Z^{\beta_{t,k}} = \infty$ indicates no observation. Hypothesis element $\beta_{t,k} > n_{z_t}$ assumes that x^r, l^k are out of the sensing range. $(x^r, l^k) \in S.R.$ indicates that a specific realization is within the sensing range. $\mathbb{P}(z^{\beta_{t,k}} | x^r, l^k)$ and $\mathbb{P}(\beta_{t,k} | x^r, l^k)$ indicate the likelihood of the Models. Last, $f(\cdot)$ denotes the likelihood value of the observation Sensor (E.G. Gaussian).

TABLE II
ALGORITHMS EXAMINED IN OUR EXPERIMENTS

Algorithm	Hypotheses control	Estimator
vanilla-HB-MCTS (III-A)	pruning	biased
PFT-DPW [6]	single hypo.	biased
DA-BSP [9]	pruning	biased
HB-MCP (ours)	sampling	unbiased

TABLE III
COMPARISON OF ALGORITHM PERFORMANCES ON DIFFERENT SCENARIOS

	Aliased matrix	Goal reaching	Kidnapped robot
HB-MCP (ours)	-585.2	-716.8	-323.7
vanilla-HB-MCTS	-909.6	-939.4	-349.5
PFT-DPW	-961.8	-1009.8	-327.8
DA-BSP	-979.5	-931.5	-330.4

Results are based on a simulation study with 100 trials per Scenario and algorithm.

leave this for future research. To make DA-BSP comparable to other algorithms, we adapted the algorithm to support anytime planning by utilizing Monte-Carlo trajectory samples instead of a full tree traversal. We also evaluated vanilla-HB-MCTS as the ad-hoc baseline for MCTS implementation with hybrid beliefs, see table II. A summary of the performance of each algorithm is given in table III.

In all cases, the experiments were done using GTSAM library [16] with a python wrapper as an inference engine for each of the hypotheses. Most current state-of-the-art online tree search planners rely on particle filters as an inference mechanism. However, particle filters are limited in their ability to support high-dimensional and correlated state spaces efficiently. Instead, through GTSAM we modeled each conditional belief as nonlinear state space model corrupted with multivariate Gaussian noise. We give more information of the hyperparameter choice in the supplementary file [14]. In the experiments, we assumed a SLAM setting, in which the map is not perfectly known, and the agent is only given a noisy prior on the map and its own pose. Due to ambiguous data associations, each measurement may be obtained from any of the surrounding landmarks within the sensing range of the agent. As a result of the ambiguous data associations, the full posterior belief becomes multi-modal, with discrete variables representing different possible associations.

Aliased matrix: The first environment is a highly aliased map, depicted in Fig. 2(b). The task of the agent is to reduce the uncertainty of its pose and all landmarks of the map, measured by the (negative-) \mathcal{A} -optimality criteria. The \mathcal{A} -optimality is the




Fig. 2. *Aliased matrix.* The goal of the agent is to minimize the uncertainty of its pose and the location of all landmarks. (a) Mean and standard deviation of the cumulative reward, over 100 trials (higher is better). (b) Illustration of the initial belief of the agent. x^* denotes the ground truth pose of the agent. l^* denotes a unique landmark. The agent receives as a prior three hypotheses at different locations, drawn as blue ellipses.




Fig. 3. *Kidnapped robot.* The goal of the agent is to minimize the uncertainty of its pose. (a) Mean and standard deviation of the cumulative reward, over 100 trials. (b) Illustration of the initial belief of the agent, blue circles illustrate conditional beliefs, crosses denote landmarks.

trace for the belief covariance matrix, commonly used as uncertainty measure. The state of the agent is its trajectory and prior landmarks. The agent is initially given three possible hypotheses for its pose, and 24 aliased landmarks evenly scattered across the map and a unique landmark, given as noisy prior to the agent. The unique landmark breaks the symmetry and may be used by the agent to disambiguate hypotheses. The action space is defined as a straight 4-directional open-loop actions, consisting of 12 intermediate steps, each of 4[m]. Each planning session was limited to 40 seconds.

Kidnapped robot: The goal of the agent is to minimize the uncertainty about the agent's pose. The environment has 16 randomly scattered landmarks on a 160 m \times 160 m grid, with added Gaussian noise given as prior. The prior pose of the agent is three hypotheses randomly scattered within the grid boundaries. The action space is defined similarly to aliased matrix environment. The reward function is defined by the \mathcal{A} -optimality criteria on the robot's pose. Each planning session was limited to 20 seconds.

Goal reaching. The goal of the agent is to reach a predefined target region. The agent prior belief is given as three hypotheses, located at different directions with respect to the target. To ensure that the right hypothesis gets to the target, the agent must first disambiguate some of the hypotheses (using the unique landmark shown in Fig. 4), and only then attempt to reach the goal. The reward function is defined as the negative sum of the Euclidean distance to goal and the \mathcal{A} -optimality criteria. Each planning session was limited to 20 seconds.

HB-MCP received the highest expected cumulative reward in both the ambiguous matrix and goal reaching scenarios. Note how in the ambiguous matrix scenario, HB-MCP achieves significant improvement in cumulative reward from step number 2. The reason for that is the agent's ability to spot the unique landmark, which is two open-loop steps away when




Fig. 4. *Goal reaching*. The goal of the agent is to reach the target location while minimizing uncertainty. (a) Mean and standard deviation of the cumulative reward, over 100 trials. (b) Illustration of the initial belief of the agent. x^* denotes the ground truth pose of the agent. l^* denotes a unique landmark. The agent receives as a prior three hypotheses at different locations.

$t = 0$, see Fig. 2(a). Due to restricted planning time, vanilla-HB-MCTS and DA-BSP fail to identify and utilize the reduction in uncertainty via disambiguation using the unique landmark. In all cases a single-hypothesis PFT-DPW is unaware of the multi-modality of the problem, and has no incentive to prioritize the unique landmark over any other (ambiguous) landmark. In case of PFT-DPW, this statement is true for all the experiments.

In the kidnapped robot scenario the algorithms performed almost equally well, with slight superiority to HB-MCP. Although PFT-DPW is mathematically inaccurate due to the choice of merely a single hypothesis, it enjoys higher inference and planning efficiency which might translate in some cases to good performance. Although the kidnapped robot reward punishes for high uncertainty, the random scatter of landmarks and poses did not lead to any strong preference of a single policy for disambiguation, which can be clearly seen from the cumulative reward of all algorithms in Fig. 3(a). Clearly, depending on the scenario, even a heuristic, single-hypothesis solver might lead to good performance. For more details, please refer to [14].

VIII. CONCLUSION

In this work, we introduced HB-MCP, a novel algorithm to handle the significant increase in computational effort of planning with hybrid beliefs. We showed that current state-of-the-art algorithms rely on an approximation, namely hypotheses pruning, that leads to a biased and inconsistent reward and value function estimate. We proposed and analyzed a different approach, namely HB-MCP, which utilizes sequential importance resampling to converge to the correct value. Additionally, instead of building symmetric hypotheses trees, HB-MCP focuses computations on the promising branches corresponding to the UCB bonus. We demonstrated how HB-MCP could be

used for planning in ambiguous scenarios and derived a simple extension to Bayesian inference to handle negative information naturally. Last, we demonstrated our approach in a simulated environment. In our experiments, HB-MCP outperformed the current state-of-the-art hybrid belief space planning algorithms.

ACKNOWLEDGMENT

The author's would like to thank Andrey Zhitnikov for helpful discussions regarding negative information.

REFERENCES

- [1] A. V. Segal and I. D. Reid, "Hybrid inference optimization for robust pose graph estimation," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2014, pp. 2675–2682.
- [2] F. Kschischang, B. Frey, and H.-A. Loeliger, "Factor graphs and the sum-product algorithm," *IEEE Trans. Inf. Theory*, vol. 47, no. 2, pp. 498–519, Feb. 2001.
- [3] P.-Y. Lajoie, S. Hu, G. Beltrame, and L. Carlone, "Modeling perceptual aliasing in slam via discrete–continuous graphical models," *IEEE Robot. Automat. Lett.*, vol. 4, no. 2, pp. 1232–1239, Apr. 2019.
- [4] D. Silver and J. Veness, "Monte-Carlo planning in large POMDPs," in *Proc. Adv. Neural Inf. Process. Syst.*, 2010, pp. 2164–2172.
- [5] L. Kocsis and C. Szepesvári, "Bandit based Monte-Carlo planning," in *Proc. Eur. Conf. Mach. Learn.*, 2006, pp. 282–293.
- [6] Z. Sunberg and M. Kochenderfer, "Online algorithms for POMDPs with continuous state, action, and observation spaces," in *Proc. Int. Conf. Automated Plan. Scheduling*, 2018, pp. 259–263.
- [7] A. Somani, N. Ye, D. Hsu, and W. S. Lee, "DESPOT: Online POMDP planning with regularization," in *Proc. NIPS*, 2013, pp. 1772–1780.
- [8] M. Barenboim and V. Indelman, "Adaptive information belief space planning," in *Proc. 31st Int. Joint Conf. Artif. Intell.*, Jul. 2022, pp. 4588–4596.
- [9] S. Pathak, A. Thomas, and V. Indelman, "A unified framework for data association aware robust belief space planning and perception," *Int. J. Robot. Res.*, vol. 32, no. 2/3, pp. 287–315, 2018.
- [10] M. Shienman and V. Indelman, "D2A-BSP: Distilled data association belief space planning with performance guarantees under budget constraints," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2022, pp. 11058–11065.
- [11] M. Shienman and V. Indelman, "Nonmyopic distilled data association belief space planning under budget constraints," in *Proc. Int. Symp. Robot. Res.*, 2022, pp. 102–118.
- [12] M. Hsiao, J. G. Mangelson, S. Suresh, C. Debrunner, and M. Kaess, "ARAS: Ambiguity-aware robust active slam based on multi-hypothesis state and map estimations," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2020, pp. 5037–5044.
- [13] M. Hsiao and M. Kaess, "MH-iSAM2: Multi-hypothesis iSAM using Bayes Tree and Hypo-tree," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2019, pp. 1274–1280.
- [14] M. Barenboim, M. Shienman, and V. Indelman, "Monte Carlo planning in hybrid belief POMDPs - supplementary material," Technion - Israel Inst. Technol., Haifa, Israel, Tech. Rep. [Online]. Available: <https://tinyurl.com/59euan3r>
- [15] T. Kennedy, *Monte Carlo Methods-A Special Topics Course*. 2016.
- [16] F. Dellaert, "Factor graphs and GTSAM: A hands-on introduction," Georgia Inst. Technol., Atlanta, GA, USA, Tech. Rep. GT-RIM-CP&R-2012-002, 2012.