Nonmyopic Data Association Aware Belief Space Planning for Robust Active Perception

Shashank Pathak*[1] and Antony Thomas*[1] and Vadim Indelman[1]

***Abstract*— One key assumption of Belief Space Planning (BSP) is that the data association is known perfectly. In this paper, we relax this assumption in the context of non-myopic planning as well as belief being a Gaussian Mixture Model (GMM). Interestingly, explicit reasoning about the data association within the belief enables our framework to have parsimonious data association, thereby resulting in a scalable solution compared with naïve permutational approaches. Unlike in some of the recent approaches where the number of components in a GMM belief can only be reduced, in our approach this can also go up such as due to perceptual aliasing present in the environment. Furthermore, our approach naturally integrates with inference, providing a unified framework for robust passive and active perception. We demonstrate key aspects of our approach and its comparison with the state of the art on a general abstract domain as well as in a real robot setup.**

I. INTRODUCTION

A key challenge in robotics is autonomous operation under different sources of uncertainty and in ambiguous situations. In such a setting, the true state of variables of interest (e.g. robot poses) is unknown and can be only inferred from a probability distribution conditioned on available data. Hence, planning and decision making should be performed considering how this distribution evolves due to candidate actions and the corresponding expected future observations. The corresponding problem, known as belief space planning (BSP), is an instantiation of a partially observable Markov decision problem (POMDP) [9], where, given an objective function over a suitable planning horizon, one aims to determine an optimal control policy as a function of belief evolution over application-dependent variables of interest.

In the last two decades, the research community has been actively developing BSP approaches that consider stochastic motion and observation models, and more recently also uncertainty in the environment. However, a typical assumption made in these approaches is that data association is given and perfect. In other words, the robot is assumed to correctly perceive the environment to be observed by its sensors, given a candidate action. Such an assumption simplifies the computation of posterior future beliefs, and is well motivated while operating in unique, unambiguous, environments.

However, real world scenarios often exhibit some level of ambiguity and perceptual aliasing (e.g. two objects that look alike from certain viewpoints), making the above assumption less appropriate. Assuming data association is given and perfect within BSP, i.e. neglecting ambiguity aspects, can lead to incorrect posterior beliefs, and consequently yield sub-optimal actions. Interestingly, ambiguity aspects have been recently accounted for in inference, e.g. in the context of SLAM, leading to robust perception and graph optimization approaches that aim to be resilient to incorrect data association due to perceptual aliasing. Yet, these approaches only tackle the passive instance of the problem, considering robot actions to be externally determined and given.

In recent work [14],[13] we developed a data association aware belief space planning (DA-BSP) framework, rigorously incorporating reasoning about data association within BSP, while also considering other sources of uncertainty (motion, sensing and environment). Such a framework is capable of better handling complexities arising in real world possibly ambiguous scenarios, and can be used e.g. for active disambiguation by determining appropriate actions for increasing confidence in a certain data association hypothesis. Yet, the approach in [14],[13] was formulated within a greedy decision making paradigm, i.e. only considering a single look ahead step.

Contributions: Our contributions in this paper are in: (1) Developing a nonmyopic formulation for DA-BSP, representing the belief as a Gaussian mixture model (GMM) (2) Demonstrating *parsimonious data association* in typical cases, where only the aliased scenes *within* the uncertainty region, as determined by a belief from appropriate look ahead step, should be accounted for (3) Generalizing the planning such that the number of components in a GMM belief can not only go down (due to full or partial disambiguation), but can also go up (4) Providing a unified framework for *robust passive and active perception* through seamless integration of our approach with the inference (5) Finally, evaluating the framework and its reasonable variation (that assumes the data association) on both a simulated as well as a real world robotic setup.

II. RELATED WORK

Calculating optimal solutions to POMDP is computationally intractable (PSPACE-complete) [12] for all but the smallest problems. The vast research area of approximate approaches (with reduced computational complexity) can be roughly segmented into point-based value iteration methods [16], simulation based [19] and sampling based approaches [17][2], and direct trajectory optimization [20][8] methods. In all cases, finding the (locally) optimal actions involves evaluating a given objective function while considering future observations to be acquired as a result of each candidate

*This work was supported by Technion Autonomous Systems Program and Israel Science Foundation. Corresponding authors: {antony.thomas,shashank}@technion.ac.il

[1]Department of Aerospace Engineering, Technion - Israel Institute of Technology http://vindelman.net.technion.ac.il

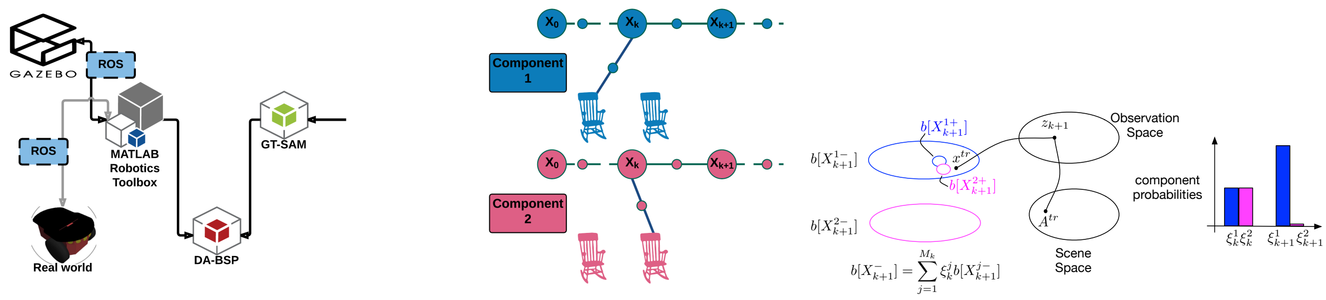

Fig. 1: (*left*) Overall implementation pipeline of DA-BSP. (*middle*) Factor-graph representation of an aliased world. Here, prior $b[X_k]$ is bi-modal due to the ambiguous data association with two similar-looking chairs. (*right*) Thus, the propagated GMM belief $b[X_{k+1}^-]$ has two components. It depicts an observation $z_{k+1} \in \{z_{k+1}\}$ of distinct scene (A^{tr}), the posterior belief $b[X_{k+1}]$ will also have two components, one of which may have subsequent negligible weight.

action. All these approaches assume data association is given and perfect. For example, it is typically assumed that the robot can be localized by making observations of known landmarks or beacons, while correctly associating future measurements with appropriate landmarks (see, e.g. [17][2]).

In the last few years, the SLAM community has investigated approaches to be resilient to false data association (outliers) overlooked by front-end algorithms (e.g. image matching), see e.g. [4][15]. However these methods, also known as robust graph optimization approaches, are developed only for the passive problem setting, i.e. robot actions are given and externally determined. In contrast, we consider a complimentary *active* framework that incorporates data association aspects within BSP, while also coping with perceptual aliasing within inference. The latter aspect of our approach is closely related to the recent method by Pfingsthorn et al. [15] as in both cases the belief is represented by a GMM.

Coming back to scalable planning methods such as BSP, we note that while the traditional BSP approaches had typically assumed the environment to be accurately known (e.g. a given map), recent works, including [11][8][5], relax this assumption and model the uncertainty of the environment mapped thus far within the belief. The corresponding framework is thus tightly related to active SLAM, with the well known trade-off between exploration and exploitation. Recent works [11][8][5] in this branch focused in particular on probabilistically modeling what future observations will be obtained given a candidate action, but again, assuming data association is given and perfect.

The issue of perceptual aliasing has been considered in earlier works on POMDP planning (e.g. [6]), though with highly simplified scenarios, since data-association further complicates the problem, especially considering multi-modal prior belief as done herein. Approaches that study related aspects were in the context of multiple hypothesis tracking (MHT), see [18] for earliest work on MHT, or more recently, of robust passive perception [21][15]. However these works do not consider active perception. Probably the closest work to our approach is by Agarwal et al. [1], where the authors also consider hypotheses due to ambiguous data association and develop a BSP approach for active disambiguation. However, unlike them, DA-BSP considers ambiguous data association also in posterior and thus does not require a guarantee of fully disambiguating action in the future. Moreover, in [1], the number of GMM components representing the belief from either inference or planning can only be reduced, for example due to a disambiguating observation. In contrast, DA-BSP provides a more general formulation where the number of components can also grow as a result of making an observation of aliased scenes. We discuss the importance of this aspect in Section IV-B.

Our approach is also tightly related with recent work on active hypothesis disambiguation in the context object detection and classification (e.g. [3][21]). Given hypotheses regarding object class and pose, these approaches aim to find a sequence future viewpoints that will lead to disambiguation, i.e. identifying the correct hypothesis. However, these approaches assume the sensor is perfectly localized and can be shown to be a specific case of DA-BSP.

III. Belief Space Planning: Preliminaries and Notations

We consider a robot operating in a known or pre-mapped environment that can be ambiguous and perceptually aliased. The robot captures observations of different scenes (or objects) in the environment with its on-board sensors, and uses these observations to infer application-dependent random variables of interest (e.g. robot poses). Denote the robot pose at time step k by x_k and a sequence of poses by that time by $X_k \doteq \{x_0, \ldots, x_k\}$. Given all controls $u_{0:k-1} \doteq \{u_0, \ldots, u_{k-1}\}$ and observations $Z_{0:k} \doteq \{Z_0, \ldots, Z_k\}$ up to time step k, the posterior probability distribution function (pdf), the *belief*, is defined as $\mathbb{P}(X_k|u_{0:k-1}, Z_{0:k})$. For notational convenience, we define the histories $\mathcal{H}_k \doteq \{u_{0:k-1}, Z_{0:k}\}$ and $\mathcal{H}_{k+1}^- \doteq \mathcal{H}_k \cup \{u_k\}$, and rewrite $b[X_k]$ as $b[X_k] \doteq \mathbb{P}(X_k|\mathcal{H}_k)$.

Let $\{A_{\mathbb{N}}\}$ denote different scenes or objects A_i in the given environment map. Ambiguous scenes can be described as some specific scenes A_i and A_j that have similar appearance from certain viewpoints. In the case the two scenes are identical, they will have the same visual appearance regardless of the viewpoint.

We consider probabilistic motion and observation models

$\mathbb{P}(x_{k+1}|x_k, u_k)$ and $\mathbb{P}(z_k|x_k, A_i)$, respectively, that can be written explicitly as

$$x_{k+1} = f(x_k, u_k) + w_k \quad , \quad z_k = h(x_k, A_i) + v_k. \qquad (1)$$

As common in literature, we consider Gaussian zero-mean process and measurement noise $w_i \sim \mathcal{N}(0, \Sigma_w)$ and $v_k \sim \mathcal{N}(0, \Sigma_v)$, with known noise covariance matrices Σ_w and Σ_v.

Given a prior $\mathbb{P}(x_0)$ and motion and observation models (1), the joint posterior pdf at the current time k can be written as

$$\mathbb{P}(X_k|\mathcal{H}_k) = \mathbb{P}(x_0) \prod_{i=1}^{k} \mathbb{P}(x_i|x_{i-1}, u_{i-1})\mathbb{P}(Z_i|x_i, A_i). \qquad (2)$$

Note the above formulation assumes data association is given.

If the prior $\mathbb{P}(x_0)$ is Gaussian, it is not difficult to show that $b[X_k]$ is also a Gaussian with some mean $\hat{X}_k$ and covariance Σ_k that can be efficiently calculated, see e.g. [10]. This is also the case when the environment model is unknown a priori and instead is constructed on-line within SLAM framework.

However, in this paper we consider a more general case where the prior belief is modeled by a Gaussian mixture model (GMM). Such a situation can arise, for example, in the kidnapped robot problem in a perceptually aliased environment (e.g. look-alike rooms), where matching sensor observations against a given map would indicate several most probable robot locations. In such a case, the belief at time k can be represented by a GMM,

$$b[X_k] = \sum_{j=1}^{M_k} \xi_k^j \mathbb{P}(X_k|\mathcal{H}_k, \gamma = j), \qquad (3)$$

where M_k is the number of components, the jth component is represented by the weight $\xi_k^j \doteq \mathbb{P}(\gamma = j|\mathcal{H}_k)$, modeling the probability of the robot being in that component, and by the conditional Gaussian $b[X_k^j] \doteq \mathbb{P}(X_k|\mathcal{H}_k, \gamma = j) \equiv \mathcal{N}(\hat{X}_k^j, \Sigma_k^j)$, with appropriate mean $\hat{X}_k^j$ and covariance Σ_k^j. Here, γ is an indicator variable denoting the component number.

Given the belief at time k, one can reason about the robot's best future actions that would minimize an objective function J for L look-ahead steps,

$$J(u_{k:k+L-1}) = \mathop{\mathbb{E}}_{z_{k+1:k+L}} \{\sum_{l=1}^{L} c_l(\mathbb{P}(X_{k+l}|\mathcal{H}_k, u_{k:k+l-1}, z_{k+1:k+l}))\} \qquad (4)$$

where the expectation is over the (unknown) future observations $z_{k+1:k+L}$, and $c_l(.)$ is the immediate cost for the lth look ahead step.

The belief from the lth step, $b[X_{k+l}] \doteq \mathbb{P}(X_{k+l}|\mathcal{H}_{k+l})$, is a function of the history $\mathcal{H}_k$, actions $u_{k:k+l-1}$ and future observations $z_{k+1:k+l}$, i.e.

$$b[X_{k+l}] = \mathbb{P}(X_{k+l}|\mathcal{H}^-_{k+l}, z_{k+l}) = \mathbb{P}(X_{k+l}|\mathcal{H}_k, u_{k:k+l-1}, z_{k+1:k+l}). \qquad (5)$$

The propagated belief is defined as $b[X^-_{k+l}] \doteq b[X_{k+l-1}]\mathbb{P}(x_{k+l}|x_{k+l-1}, u_{k+l-1})$. As in the greedy case [14][13], one can calculate the marginal belief over the future pose x_{k+l} as $b[x^-_{k+l}] \doteq \int_{\neg x_{k+l}} b[X^-_{k+l}]$.

Given the GMM belief from the previous look ahead step, $b[X_{k+l-1}] = \sum_{j=1}^{M_{k+l-1}} \xi^j_{k+l-1} b[X^j_{k+l-1}]$, the propagated belief $b[X^-_{k+l}]$ becomes

$$b[X^-_{k+l}] = \sum_{j=1}^{M_{k+l-1}} \xi^{j-}_{k+l} b[X^{j-}_{k+l}], \qquad (6)$$

with $\xi^{j-}_{k+l} \doteq \mathbb{P}(\gamma_{k+l} = j|\mathcal{H}^-_{k+l}) \equiv \xi^j_{k+l-1}$, and $b[X^{j-}_{k+l}] \doteq \mathbb{P}(X_{k+l}|\mathcal{H}^-_{k+l}, \gamma_{k+l} = j) = b[X^j_{k+l-1}]\mathbb{P}(x_{k+l}|x_{k+l-1}, u_{k+l-1})$.

Assuming data association is known, one can consider for each specific value of z_{k+l} the corresponding observed scene A_i, and express the posterior (5) recursively in terms of $b[X^-_{k+l}]$ as

$$b[X_{k+l}] = \eta b[X^-_{k+l}]\mathbb{P}(z_{k+l}|x_{k+l}, A_i). \qquad (7)$$

Letting $u \doteq u_{k:k+L-1}$, the optimal control problem is defined as $u^\star \doteq \arg\min_u J(u)$.

IV. Nonmyopic Data Association Aware BSP

The non-myopic objective function (4) can be written as

$$J(u_{k:k+L-1}) = \int_{z_{k+1:k+L}} \sum_{l=1}^{L} \overbrace{\mathbb{P}(z_{k+l} \mid \mathcal{H}^-_{k+l})}^{(a)} c_l \left(\overbrace{\mathbb{P}(X_{k+l}|\mathcal{H}^-_{k+l}, z_{k+l})}^{(b)} \right), \qquad (8)$$

where the expectation over future observations is written explicitly, accounting for all possible realizations of these unknown observations. Although dropped to reduce clutter, the history $\mathcal{H}^-_{k+l}$ includes future observations $z_{k+1:k+l-1}$ up to the lth look ahead step (see definition in Section III).

Similarly to the myopic case [14][13], the two terms (a) and (b) in Eq. (8) have intuitive meaning: for each considered value of z_{k+l}, (a) represents how likely is it to get such an observation, while (b) corresponds to the posterior belief *given* this specific z_{k+l}. However, the difference in a non-myopic case is that both terms are conditioned on the history $\mathcal{H}^-_{k+l}$ which is a function of $z_{k+1:k+l-1}$; hence, the above reasoning is valid for *all* possible realizations of $z_{k+1:k+l-1}$ and the corresponding posterior beliefs $\mathbb{P}(X_{k+l-1}|\mathcal{H}_{k+l-1})$.

Assuming data association is given, implies that for each possible observation $z_{k+l} \in \{z_{k+l}\}$ the corresponding observed scene $A_i \in \mathcal{A}$ is known, making it possible to express the posterior recursively as in Eq. (7). Yet, it is unknown from what future robot pose x_{k+l} the actual observation z_{k+l} will be acquired, since the *actual* robot pose x_k at planning time k is unknown and the controls are stochastic. Indeed, as a result of actions $u_{k:k+l-1}$, the robot actual (true) pose x_{k+l} can be anywhere within the propagated belief $b[x^-_{k+l}]$, which according to Eq. (6) is a GMM.

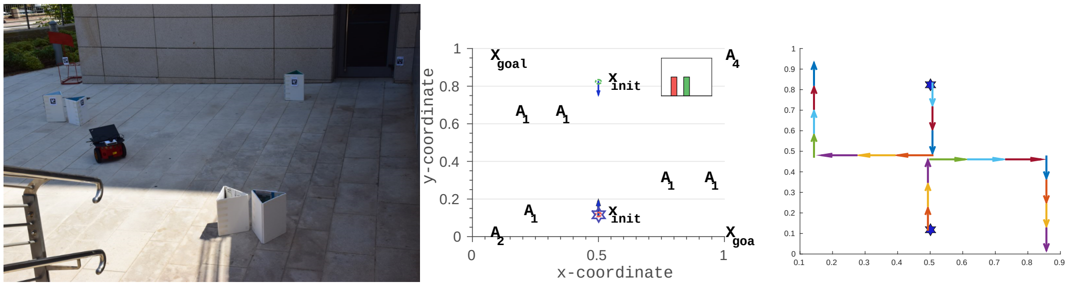

Fig. 2: (*left*) Real-world experimental setup. (*middle*) Schematics of the same world. Current belief is a 2-modal GMM with mean position depicted by x_{init}. Ground truth robot position is indicated with ✡; arrows indicate orientation (and not motion). Positions of goal x_{goal}, obstacles x_{ob} and scenes A_i are also shown. (*right*) An 11-step nominal trajectory, shown relative to two GMM components.

In contrast to the above, our approach relaxes the data association assumption and instead reasons about possible scenes that a future observation z_{k+l} could be generated from, see Figure 1. Such an explicit reasoning about data association within the belief would seem to incur significant additional computational complexity. However, while this would be the case with naïve permutational approaches, our framework is scalable in practice due to parsimonious data association (see Section IV-B). Moreover, our formulation can accommodate situations where due to perceptual aliasing, the number of components in a GMM belief (either in inference or planning) can increase. The weights of these components allow us to re-use some known merging and pruning approaches and hence further reduce the computational gap between data-association aware non-myopic BSP and the usual non-myopic BSP (Section IV-C). Finally, in Section IV-D we briefly describe the seamless integration of DA-BSP with passive inference, thus facilitating a *unified framework for robust passive and active perception*.

A. Calculating Terms (a) and (b) in Eq. (8)

The terms (a) and (b) in Eq. (8) can be calculated in a similar fashion to the myopic case [14][13]. We now discuss these calculations for the lth look ahead step ($l \in [1, L]$), given a propagated belief $b[X^-_{k+l}]$ and a future possible observation $z_{k+l} \in \{z_{k+l}\}$, where the set $\{z_{k+l}\}$ is generated as in [14].

1) Computing the term (a) - $\mathbb{P}(z_{k+l}|\mathcal{H}^-_{k+l})$:: Applying total probability over non-overlapping scene space $\{A_{\mathbb{N}}\}$ and marginalizing over all possible robot poses, yields

$$\mathbb{P}(z_{k+l}|\mathcal{H}^-_{k+l}) \equiv \sum_i^{|A_{\mathbb{N}}|} \int_x \mathbb{P}(z_{k+l}, x, A_i \,|\, \mathcal{H}^-_{k+l}) \doteq \sum_i^{|A_{\mathbb{N}}|} w^i_{k+l}. \quad (9)$$

As seen from the above equation, to calculate the likelihood of obtaining some observation z_{k+l}, we consider separately, for each scene $A_i \in \{A_{\mathbb{N}}\}$, the likelihood that this observation was generated by scene A_i. This probability is captured for each scene A_i by a corresponding weight w^i_{k+l}; these weights are then summed to get the actual likelihood of observation z_{k+l}. As will be seen below, these weights naturally account for perceptual aliasing aspects for each considered z_{k+l}.

In practice, instead of considering the entire scene space $\{A_{\mathbb{N}}\}$ that could be huge, the availability of the belief $b[X^-_{k+l}]$ makes it possible to consider only a small subset of $\{A_{\mathbb{N}}\}$. See further discussion in Section IV-B.

Proceeding with the derivation further, using the chain rule we compute

$$\sum_i \int_x \mathbb{P}(z_{k+l} \mid x, A_i, \mathcal{H}^-_{k+l}) \mathbb{P}(A_i, x \mid \mathcal{H}^-_{k+l}) \quad (10)$$

As $\mathbb{P}(A_i, x \mid \mathcal{H}^-_{k+l}) = \mathbb{P}(A_i | x \mid \mathcal{H}^-_{k+l}) b[x^-_{k+l} = x]$, we get $\sum_i^{|A_{\mathbb{N}}|} w^i_{k+l}$, with

$$w^i_{k+l} \doteq \int_x \mathbb{P}(z_{k+l}|x, A_i, \mathcal{H}^-_{k+l}) \mathbb{P}(A_i|\mathcal{H}^-_{k+l}, x) b[x^-_{k+l} = x]. \quad (11)$$

Since the propagated belief (6), from which $b[x^-_{k+l}]$ is calculated, is a GMM, we can replace $b[x^-_{k+l} = x]$ with $\sum_{j=1}^{M_{k+l-1}} \xi^{j-}_{k+l} b[x^-_{k+l,j} = x]$.

Here, $\mathbb{P}(z_{k+l} \mid A_i, x, \mathcal{H}^-_{k+l}) \equiv \mathbb{P}(z_{k+l} \mid A_i, x)$ is the standard measurement likelihood term, while $\mathbb{P}(A_i \mid \mathcal{H}^-_{k+l}, x)$ represents the *event likelihood*, which denotes the probability of scene A_i to be observed from viewpoint x. In other words, this scenario-dependent term encodes from what viewpoints each scene A_i is observable and could also model occlusion and additional aspects. As such, this term can be determined given a model of the environment; in this work, we consider this term to be given.

The weights w^i_{k+l} from Eq. (11) naturally capture *perceptual aliasing* aspects: consider some observation z_{k+l} and the corresponding generative model $z_{k+l} = h(x^{tr}, A^{tr}) + v$ with appropriate unknown *true* robot pose x^{tr} and scene $A^{tr} \in \{A_{\mathbb{N}}\}$. Clearly, the measurement likelihood $\mathbb{P}(z_{k+l} \mid x, A_i, \mathcal{H}^-_{k+l})$ will be high when evaluated for $A_i = A^{tr}$ and in vicinity of x^{tr}. In case of perceptual aliasing, however, there will be also another scene(s) A_j which could generate the same observation z_{k+l} from appropriate robot pose x'. Thus, the corresponding measurement likelihood term to A_j will also be high for x'. However, the actual value of w^i_{k+l} (for each $A_i \in \{A_{\mathbb{N}}\}$) depends, in addition to the measurement likelihood, also on the mentioned-above event likelihood and on the GMM belief $b[x^-_{k+l}]$, with the latter weighting the probability of each considered robot pose x. This correctly captures the intuition that those observations z with low-probability poses $b[x^-_{k+l} = x^{tr}]$ will be unlikely to be actually acquired, leading to low value of w^i_{k+l} with $A_i = A^{tr}$. See also the discussion in Section IV-B.

2) Computing the term (b) - $\mathbb{P}(X_{k+l}|\mathcal{H}^-_{k+l}, z_{k+l})$:: The term (b) represents the posterior probability conditioned on observation z_{k+l}. This term can be similarly calculated, with a key difference: since the observation z_{k+l} is given, it must have been generated by *one* specific (but unknown) scene A_i according to measurement model (1). Hence, also here, we consider all possible such scenes and weight them accordingly, with weights $\tilde{w}^i_{k+l}$ representing the probability of each scene A_i to have generated the observation z_{k+l}. As will be seen next, the posterior $\mathbb{P}(X_{k+l}|\mathcal{H}^-_{k+l}, z_{k+l})$ is a GMM with M_{k+l} components.

Applying total probability over non-overlapping $\{A_{\mathbb{N}}\}$ and chain-rule, we get $\mathbb{P}(X_{k+l}|\mathcal{H}^-_{k+l}, z_{k+l}) = \sum_{i=1}^{|A_{\mathbb{N}}|} \mathbb{P}(X_{k+l} \mid \mathcal{H}^-_{k+l}, z_{k+l}, A_i) \cdot \mathbb{P}(A_i \mid \mathcal{H}^-_{k+l}, z_{k+l})$. The first term, $\mathbb{P}(X_{k+l} \mid \mathcal{H}^-_{k+l}, z_{k+l}, A_i)$, is the posterior belief conditioned on observation z_{k+l}, history $\mathcal{H}^-_{k+l}$, as well as a candidate scene A_i that supposedly generated z_{k+l}. It is not difficult to show that this posterior is actually the GMM

$$\mathbb{P}(X_{k+l} \mid \mathcal{H}^-_{k+l}, z_{k+l}, A_i) = \sum_{j=1}^{M_{k+l-1}} \xi^j_{k+l-1} b[X^{j+}_{k+l}|A_i], \tag{12}$$

where $b[X^{j+}_{k+l}|A_i] \doteq \mathbb{P}(X_{k+l}|\mathcal{H}^-_{k+l}, \gamma = j, A_i, z_{k+l})$ is the posterior of the jth GMM component of the propagated belief $b[X^-_{k+l}]$, see Eq. (6).

Plugging-in Eq. (12) into $\mathbb{P}(X_{k+l}|\mathcal{H}^-_{k+l}, z_{k+l}) \equiv b[X_{k+l}]$ from Eq. (7) yields:

$$b[X_{k+l}] = \sum_{i=1}^{|A_{\mathbb{N}}|} \sum_{j=1}^{M_{k+l-1}} \xi^j_{k+l-1} \mathbb{P}(A_i \mid \mathcal{H}^-_{k+l}, z_{k+l}) b[X^{j+}_{k+l}|A_i]. \tag{13}$$

The term, $\mathbb{P}(A_i \mid \mathcal{H}_k, u_k, z_{k+l})$, is merely the likelihood of A_i being actually the one which generated the observation z_{k+l}. This term can be evaluated, similarly to Section IV-A.1. Accounting for $b[x^{j-}_{k+l}]$ for each considered jth component as $\mathbb{P}(A_i \mid \mathcal{H}^-_{k+l}, z_{k+l}) = \int_x \mathbb{P}(A_i, x \mid \mathcal{H}^-_{k+l}, z_{k+l})$, and applying Bayes' rule yields

$$\tilde{w}^{ij}_{k+l} \doteq \eta' \int_x \mathbb{P}(z_{k+l}|A_i, x, \mathcal{H}^-_{k+l}) \mathbb{P}(A_i|\mathcal{H}^-_{k+l}, x) b[x^{j-}_{k+l} = x], \tag{14}$$

with $\eta' = 1/\mathbb{P}(z_{k+l} \mid \mathcal{H}^-_{k+l})$. Note that for each component j, $\sum_i \tilde{w}^{ij}_{k+l} = 1$. Finally, we can re-write Eq. (13) as

$$b[X_{k+l}] = \sum_{r=1}^{M_{k+l}} \xi^r_{k+l} \mathbb{P}(X_{k+l}|\mathcal{H}_{k+l}, \gamma = r) = \sum_{r=1}^{M_{k+l}} \xi^r_{k+l} b[X^{r+}_{k+l}], \tag{15}$$

where $\xi^r_{k+l} \doteq \xi^{ij}_{k+l} \equiv \xi^j_{k+l-1} \tilde{w}^{ij}_{k+l}$ and $b[X^{r+}_{k+l}] \doteq \mathbb{P}(X_{k+l}|\mathcal{H}_{k+l}, \gamma = r)$. As seen, we got a new GMM with M_{k+l} components, where each component $r \in [1, M_{k+l}]$, with appropriate mapping to indices (i, j) from Eq. (13), is represented by weight ξ^r_{k+l} and posterior conditional belief $b[X^{r+}_{k+l}]$. The latter can be evaluated as the Gaussian $b[X^{r+}_{k+l}] = \mathcal{N}(\hat{X}^r_{k+l}, \Sigma^r_{k+l})$, with mean $\hat{X}^r_{k+l}$ and covariance Σ^r_{k+l}.

B. Parsimonious Data Association

According to Eq. (15) one may think the number of components grows unboundedly with the planning horizon. However, in practice this is not the case due to parsimonious data association (see Figure 1): Only the aliased scenes within the uncertainty region according to the propagated belief $b[X^-_{k+l}]$ should be accounted for. This can be also seen from Eq. (14), where distant viewpoints x from the mean of the Gaussian $b[x^{j-}_{k+l}]$ naturally get negligible probabilities, limiting the scope of aliased scenes to be considered from $\{A_{\mathbb{N}}\}$ in practice.

Moreover, only some of the weights of the M_{k+l} components of posterior $b[X_{k+l}]$ will be typically non-negligible, which corresponds to full or partial disambiguation. In particular, if the environment has only distinct scenes, then for each specific value of z_{k+l}, there will be only one scene A_i that can generate such an observation according to the model (1). From Eq. (15) it can be seen that in this fully-disambiguating case $M_{k+l} = M_{k+l-1}$; yet, only *one* of the weights ξ^r_{k+l} will be non-negligible. In Section IV-C we exploit this observation to prune negligible belief components.

To conclude this section, we note that a unique aspect of our approach is that the number of components adaptively changes (in planning and, similarly, in inference), i.e. the number of components can go down in the case of full or partial disambiguation, and can also go up when incorporating a new observation of aliased scenes. This is in contrast to existing approaches (e.g. [1]) where the number of components can only be reduced.

C. Evaluating data associations

In order to evaluate the effectiveness of data association, we note the averaged number of component in the mixture, denoted by $\tilde{m}$. Also, we define a suitable metric η_{da} as follows: $\eta_{da}(z) = \sum_{M_{k+l}} w_{ij} \cdot c(z, A_i)$ where $c(z, A_i)$ is an indicator function of the true data association, given the observation z and considering a scene A_i i.e., $c(z, A_i)$ is 1 if the data association is correct and 0 otherwise. Since during planning we simulate the observations, $\eta_{da}(z)$ can be computed and averaged over all observations i.e., $\eta_{da} = \frac{\sum_{z \in Z} \eta_{da}(z)}{|Z|}$. For example, under absence of perceptual aliasing and with only correct associations being made, $\eta_{da} = 1$, whereas with increasing aliasing this quantity tends to 0. Similarly, with growing number of incorrect associations (provided their respective posteriors are with non-negligible weights) η_{da} diminishes.

In order to compare `DA-BSP` with the approaches that consider a given data-association, we consider a configuration for *known* data association `BSP-u`. Here, we determine (heuristically and simulataneously) the chosen scene as well as the chosen component of the propagated belief. We take the closest (under a suitable distance metric) scene $\tilde{A}_i$ from all the components of the propagated belief and their respective detectable scenes. Even for a GMM prior and under presence of aliasing, this approach yields a unimodal posterior. We define $\eta(z) = c(z, \tilde{A}_i)$. Like before, in case of many observations (such as during planning), we average $\eta(z)$ over all observations. Furthermore, we define a parameter ϵ such that with a probability of ϵ, the associated scene (and the component) is guaranteed to be the ground truth.

D. Robust Passive Perception

Our approach seamlessly integrates within inference. To see that, consider the GMM belief at the current time step k, which is used by DA-BSP to determine the best actions $u^\star_{k:k+L-1}$. Given the latter, the robot executes the first action $u^\star_k$ (or a sequence of actions), and gets a new observation z_{k+1}. While in planning we had to consider, conceptually, all possible realizations of *future* observations, in inference the observation is acquired in practice. However, calculating the posterior belief $b[X_{k+1}]$ in inference involves *exactly* the same equations as in planning (term (b) in Section IV-A) considering the acquired observation z_{k+1} and the propagated belief due to action $u^\star_k$. Once $b[X_{k+1}]$ is available, one can either continue execution of the next optimal action(s) from $u^\star_{k:k+L-1}$, or resort to model predictive control (MPC) and use DA-BSP to calculate an updated sequence of optimal actions. In our implementation we follow the latter case.

V. Experiments

In this section, we seek to evaluate various facets of data aware belief space planning or `DA-BSP`. We use heuristic `BSP-u` which uses minimum distance to chose the correct association. Also, we consider $\epsilon = 0.5$ i.e., at each detection, with equal probability the planner is either given the correct association or has to choose using `BSP-u` heuristic.

`DA-BSP` is a general framework and can handle any cost function of the form Eq. (8). However, for the sake of simplicity, in the computation of cost due to belief, we consider the usual reaching of the goal as well as data association. The latter is measured by noting how far (the more is this divergence, the lower is the cost) a multimodal distribution is from a perfectly uniform one, through measuring negative KL divergence of the weights from the uniform weights, and is denoted by KL_u.

A. Real world outdoor

In order to elucidate the crucial properties of non-myopic `DA-BSP`, we consider a real world experiment as shown in the Figure 2 (left) with a single robot R. The set of control trajectories is finite and known a priori; one of which is shown in the Figure 2 (right). The state space $X \in \mathbb{R}^3$ consists of 2D coordinates as shown, as well as the orientation of the robot. Here, A_i denotes an Apriltag with the index i. This enables us to simulate perceptual aliasing.[1] To ensure robustness, the tag A_i is considered detected only if it is also within a closed sub-space $X_{A_i} \subset X$. The action space of the robot comprises of $a \in \{$north, east, west, south$\}$, along with the known motion uncertainty. Initially, the belief of the robot is a multi-modal distribution, represented by a GMM with 2 components having equal weights and centered around each x_{init}. The objective of the robot is to both localize itself and to reach the goal x_{goal}. The L-step planning, followed by enacting one optimal control action and the consequent inference, shall together be called an *epoch*. Note that this simple representation of the world is very general. Indeed, real world complications – such as the state space being of higher dimension, different levels of ambiguities between the scenes and planning problem of longer time-scales – can all be easily incorporated into it.

[1]Though not the focus here, any object detector can be easily incorporated in our general framework of `DA-BSP`.

Algorithm	Epoch	$L=2$			$L=4$			Inference		
		t(s)	$(\eta_{da},\tilde{m})$	DA	t(s)	$(\eta_{da},\tilde{m})$	DA	t(s)	$(\eta_{da},\tilde{m})$	DA
`DA-BSP`	1	0.19	(-,2)	-	0.36	(-,2)	-	0.03	(-,2)	-
	6	6.24	(0.18,4)	✓	13.81	(-,1)	-	0.03	(-,2)	-
	7	23.70	(0.26,4)	✓	8.06	(-,1)	-	3.09	(0.56,4)	✓
	10	1.00	(1,1)	✓	1.71	(1,1)	✓	0.99	(1,1)	✓
		t(s)	η	DA	t(s)	η	DA	t(s)	η	DA
`BSP-u`	1	0.19	-	-	0.36	-	-	0.03	-	-
	6	0.34	0.5	×	1.94	1	✓	0.03	-	-
	7	0.34	0	×	0.87	0	×	0.27	0	×
	10	0.25	0	×	0.03	0	×	0.24	0	×

TABLE I: Evaluating `DA-BSP` in several steps of planning and inference, with $L = 2$ and $L = 4$. The times in seconds spent in planning and in inference is denoted by t, while average modes are denoted by $\tilde{m}$. DA denotes *correct* data association; refer Sec. V-A.

The left column of Figures 3a–3c shows the evolution of belief at the end of various decision epochs of `DA-BSP`. Initially, the belief is a multi-modal GMM with progressive steps reducing the number of components of this GMM. The right column of Figures 3a–3c show the evolution of the posterior under a different assumption of data association – `BSP-u`. Full disambiguation might occur sooner here, however unless the ϵ is very high i.e., the correct ground truth association is made frequently, the posterior may get really far away from the actual position of the robot, such as seen in Figure 3c for `BSP-u`, where $\epsilon = 0.5$.

Since our approach provides a uniform framework for *plan-act-infer*, we can study the hypotheses generation during planning in similar ways to inference. As seen in Figure 4 (left), the cardinality of components in the GMM varies during planning, across the different samples considered. Given a particular GMM during planning, we calculate the KL-divergence of it from a uniform distribution and denote it by KL_u, as seen in Figure 4 (right).

In the presence of data association challenges, the quality of planning can be roughly assessed by considering if at least one of the posterior contains correct data association. This is represented by DA in the Table I. Here, η_{da} which also considers the weight of such associations, is also shown. Naturally, reasoning over all possible associations results in greater computational effort. We measure the run-time of the algorithm as a proxy for *effectiveness*. Both these measures along with the number of hypotheses in the beliefs are shown in the Table I where we see the effect of non-myopic `DA-BSP` with two different planning horizons.

B. Gazebo Simulation

To demonstrate our concept in a more realistic simulation, we considered a Pioneer robot in an aliased 2 floor

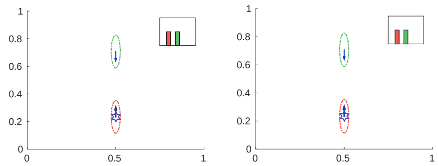

(a) Epoch 1

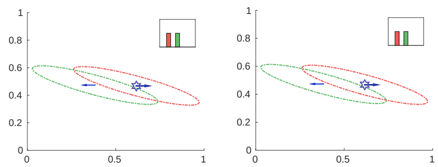

(b) Epoch 4

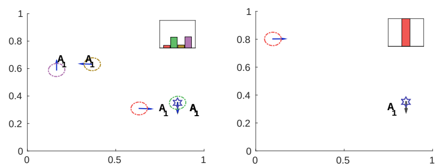

(c) Epoch 7

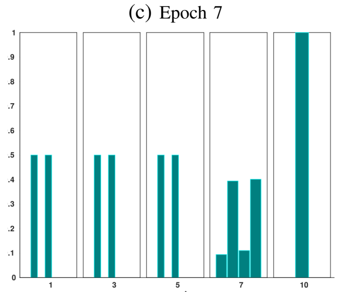

(d) GMM weights of corresponding beliefs in `DA-BSP`. Epochs {1,3,5,7,10}.

Fig. 3: (a)-(c) Evolution of inferred belief as *decision epoch* progresses with $L = 3$; epochs depicted are $\{1,4,7\}$. Left and right columns depict evolution of propagated beliefs as well as inferred one (the larger and the smaller ellipses respectively), for different planning algorithms, i.e. `DA-BSP` and `BSP-u`, respectively. GMM components and associated weights are designated with different colors. For clarity, only the detected scene(s) are shown. In (c), while `BSP-u` fully disambiguates (only one component), the chosen component is wrong due to incorrect data association. (d) Evolution of GMM components weights during some epochs of these approaches. Note that the number of components *increases* from 2 to 4 and then decreases to 1. Here, $L = 2$.

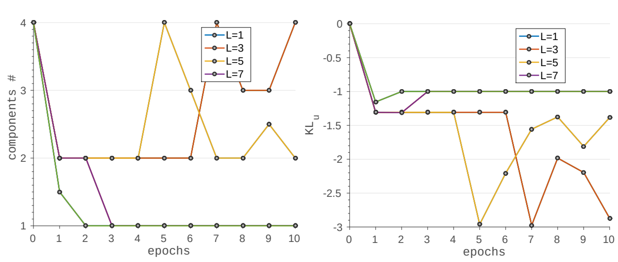

Fig. 4: Evolution of belief as *decision epoch* progresses during `DA-BSP` planning. Average number of components in the belief mixtures and the KL_u metric are depicted in *left* and *right* respectively.

office room environment within the Gazebo simulator. The robot is fitted with realistic sensors enabling laser scans and odometric estimation. Our implementation uses ICP for laser scan matching and GTSAM [7] for inference (both in passive and active perception) within ROS framework. This implementation is sufficiently realistic to be used also in real world experiments, which are planned in the near future.

The considered scenario is shown in Figure 5 (left), with numbers indicating different places. The two floors are very similar in appearance except for the printer p_1 (Figure 5, left). Also there is aliasing within each floor due to the way the cubicles are arranged. The goal for the robot is to reach the seating area near the vending machine and to disambiguate between the floors. Initially the robot wakes up to find itself either at place 1 (in the direction 1–2) or 6 (in the direction 6–7). Hence its initial belief is modeled as a 4-component GMM (due to inter-floor aliasing).

In total the robot takes 7 steps (actions) to fully disambiguate. From the places 1 and 6 the set of trajectories taken are respectively given by the segments, 1-2-3-4-5-12-11-8 and 6-7-8-9-10-11-12-3. Intra-floor disambiguation occurs along the path 4–5 (or 9–10) and inter-floor (or full disambiguation) occurs at 8 (or 3). For a planning horizon of 2 look ahead steps ($L = 2$), the evolution of weights of the components in the GMM after inference are shown in Figure 5 (center). As seen in the figure, there are 4 components initially, due to intra and inter-floor aliasing, which disambiguates to 2 after the 3^{rd} step (intra-floor disambiguation) and finally to 1 component (full disambiguation) at the 7-th step. Figure 5 (right) shows the cardinality of components in the GMM during planning, for $L = 1, 3$ and 5. If we have a longer horizon, then planning leads to reduction in the number of components in a fewer steps and this is seen in Figure 5; graph gets steeper as L increases.

Table II gives a comparison of `DA-BSP` and `BSP-u` (see section IV-C for definition) at 4 different steps of planning and inference for $L = 2$ and 3. In `DA-BSP` after the first step we have $\eta_{da} = 0.25$ as we consider all the four associations that are equally likely. η_{da} is close to 0.5 when intra-floor disambiguation occurs (see step 5 for $L = 2$ and step 2 for $L = 3$) and finally $\eta_{da} = 1$ once the robot fully disambiguates and identifies the floor it is in. However `BSP-u` is sometimes correct and sometimes wrongs depending on which association it chooses as the correct one. As seen in the table II, $\eta = 0$ at the end of the 7-th step in all the columns, indicating that with the `BSP-u` approach, the robot would have inferred itself to be at the wrong floor and/or wrong place. `DA-BSP` is computationally more expensive, however it is much more likely to perform the correct data association as seen from the η_{da} metric.

VI. Conclusions

State-of-the-art belief space planning (BSP) approaches typically assume data association to be given and perfect. In this work, we developed a nonmyopic data association aware belief space planning (DA-BSP) approach that relaxes

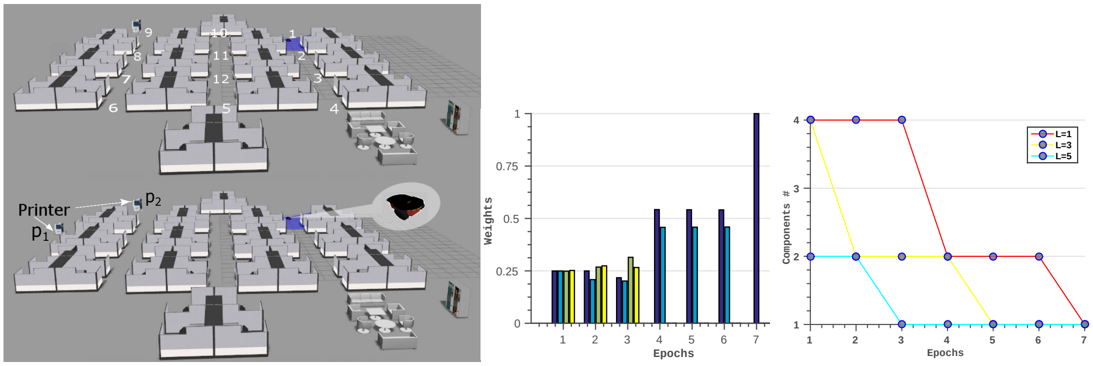

Fig. 5: (*top*) The aliased 2 floor office room environment in a Gazebo simulator with Pioneer. Here, p_1 and p_2 denote the printers; 1, 6 are the mean positions in each floor for the 4-modal initial GMM belief. (*bottm left*) Evolution of weights of the components in the GMM after inference for $L = 2$. (*bottom right*) Average number of components in the belief mixtures for different planning horizon.

Algorithm	Epoch	$L=2$			$L=4$			Inference		
		t(s)	$(\eta_{da},\tilde{m})$	DA	t(s)	$(\eta_{da},\tilde{m})$	DA	t(s)	$(\eta_{da},\tilde{m})$	DA
`DA-BSP`	1	81.04	(0.25,4)	✓	307.95	(0.25,4)	✓	4.95	(0.25,4)	✓
	2	64.41	(0.27,4)	✓	97.60	(0.35,2)	✓	5.05	(0.25,4)	✓
	5	4.91	(0.51,2)	✓	8.56	(1,1)	✓	1.03	(0.49,2)	✓
	7	0.57	(1,1)	✓	0.57	(1,1)	✓	0.54	(1,1)	✓
		t(s)	η	DA	t(s)	η	DA	t(s)	η	DA
`BSP-u`	1	2.86	1	✓	4.84	1	✓	0.88	1	✓
	2	0.56	1	✓	0.91	0	×	0.32	1	✓
	5	0.54	1	✓	0.89	0	×	0.33	1	✓
	7	0.20	0	×	0.24	0	×	0.16	0	×

TABLE II: Evaluating `DA-BSP` in several steps of planning and inference, for $L = 2$ and $L = 4$. The times in seconds spent in planning and in inference is denoted by t, while average modes are denoted by $\tilde{m}$. DA denotes *correct* data association; refer Sec. V-A

this assumption while considering different sources of uncertainty (uncertainty in robot motion, sensing and possibly in the observed environment). As such, it is capable of better coping with ambiguous, perceptually aliased, situations by appropriately calculating belief evolution and expected cost due to candidate actions, and in particular, could be used for active disambiguation. Importantly, DA-BSP seamlessly integrates with inference, thereby providing a *unified framework for robust active and passive perception*. We show key aspects of DA-BSP and compare it with the state-of-the-art through experiment using Pioneer robot as well as through the Gazebo simulation. We are currently investigating the theoretical aspects of greedy data association by reasoning over relevant properties – such as *submodularity* and *optimal substructure* – of an appropriate objective function in this context.

References

[1] S. Agarwal, A. Tamjidi, and S. Chakravorty. Motion planning in non-gaussian belief spaces for mobile robots. *arXiv preprint arXiv:1511.04634*, 2015.

[2] A.-A. Agha-Mohammadi, S. Chakravorty, and N. M. Amato. Firm: Sampling-based feedback motion planning under motion uncertainty and imperfect measurements. *Intl. J. of Robotics Research*, 2014.

[3] N. Atanasov, B. Sankaran, J.L. Ny, G. J. Pappas, and K. Daniilidis. Nonmyopic view planning for active object classification and pose estimation. *IEEE Trans. Robotics*, 30:1078–1090, 2014.

[4] L. Carlone, A. Censi, and F. Dellaert. Selecting good measurements via l1 relaxation: A convex approach for robust estimation over graphs. In *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, pages 2667–2674, 2014.

[5] Stephen M Chaves, Jeffrey M Walls, Enric Galceran, and Ryan M Eustice. Risk aversion in belief-space planning under measurement acquisition uncertainty. In *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2015.

[6] L. Chrisman. Reinforcement learning with perceptual aliasing: The perceptual distinctions approach. In *Nat. Conf. on Artificial Intelligence (AAAI)*, pages 183–188, 1992.

[7] F. Dellaert. Factor graphs and GTSAM: A hands-on introduction. Technical Report GT-RIM-CP&R-2012-002, Georgia Institute of Technology, September 2012.

[8] V. Indelman, L. Carlone, and F. Dellaert. Planning in the continuous domain: a generalized belief space approach for autonomous navigation in unknown environments. *Intl. J. of Robotics Research*, 34(7):849–882, 2015.

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

[10] M. Kaess, H. Johannsson, R. Roberts, V. Ila, J. Leonard, and F. Dellaert. iSAM2: Incremental smoothing and mapping using the Bayes tree. *Intl. J. of Robotics Research*, 31:217–236, Feb 2012.

[11] A. Kim and R. M. Eustice. Active visual slam for robotic area coverage: Theory and experiment. *Intl. J. of Robotics Research*, 2014.

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

[13] S. Pathak, A. Thomas, A. Feniger, and V. Indelman. Da-bsp: Towards data association aware belief space planning for robust active perception. In *Eur. Conf. on AI (ECAI)*, September 2016. Accepted.

[14] S. Pathak, A. Thomas, A. Feniger, and V. Indelman. Robust active perception via data-association aware belief space planning. *arXiv preprint: 1606.05124*, 2016.

[15] M. Pfingsthorn and A. Birk. Generalized graph slam: Solving local and global ambiguities through multimodal and hyperedge constraints. *Intl. J. of Robotics Research*, 35(6):601–630, 2016.

[16] J. Pineau, G. J. Gordon, and S. Thrun. Anytime point-based approximations for large pomdps. *J. of Artificial Intelligence Research*, 27:335–380, 2006.

[17] S. Prentice and N. Roy. The belief roadmap: Efficient planning in belief space by factoring the covariance. *Intl. J. of Robotics Research*, 2009.

[18] D.B. Reid. An algorithm for tracking multiple targets. *IEEE Trans. Automat. Contr.*, AC-24(6):84–90, December 1979.

[19] C. Stachniss, G. Grisetti, and W. Burgard. Information gain-based exploration using rao-blackwellized particle filters. In *Robotics: Science and Systems (RSS)*, pages 65–72, 2005.

[20] J. Van Den Berg, S. Patil, and R. Alterovitz. Motion planning under uncertainty using iterative local optimization in belief space. *Intl. J. of Robotics Research*, 31(11):1263–1278, 2012.

[21] L. Wong, L. Kaelbling, and T. Lozano-Pérez. Data association for semantic world modeling from partial views. *Intl. J. of Robotics Research*, 34(7):1064–1082, 2015.