Towards Data Association Aware Belief Space Planning for Robust Active Perception

Shashank Pathak, Antony Thomas, Asaf Feniger and Vadim Indelman

***Abstract*—We develop a belief space planning (BSP) approach that advances the state of the art by incorporating reasoning about data association (DA) within planning (existing BSP approaches typically assume data association is given and perfect), while considering additional sources of uncertainty. Our data association aware belief space planning (DA-BSP) approach explicitly reasons about DA within belief evolution, and as such can better accommodate these challenging real world scenarios. Starting from a Gaussian prior, due to perceptual aliasing, we show that the posterior belief becomes a Gaussian mixture model. Overall, our approach is applicable to robust active perception and autonomous navigation in perceptually aliased environments.**

I. Introduction

Belief space planning (BSP) and decision-making under uncertainty are fundamental problems in robotics and artificial intelligence, with applications including autonomous navigation, object grasping and manipulation, active SLAM, and robotic surgery. In presence of uncertainty, such as in robot motion and sensing, the true state of variables of interest (e.g. robot poses), is unknown and can only be represented by a probability distribution over possible states, given available data. Planning and decision-making should be therefore performed over this distribution, the belief space, which can be inferred using probabilistic approaches based on incoming sensor observations and prior knowledge. The corresponding problem is an instantiation of a partially observable Markov decision problem (POMDP) [9], where, given an objective function, one aims to determine an optimal control policy as a function of belief evolution over application-dependent variables of interest.

However, state-of-the-art BSP approaches typically assume data association to be given and perfect (see Figure 1), i.e. the robot is assumed to correctly perceive the environment to be observed by its sensors, given a candidate action. Recent works, including [4]–[6], [11], [17], 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 work [5], [6], [11], [17] in this branch focused in particular on probabilistically modelling what future observations will be obtained given a candidate action. However, these approaches consider each such future observation to be correctly associated to an appropriate scene, and hence, assume data association to be given and perfect.

In the last few years, the SLAM research community has investigated approaches to be resilient to false data association (outliers) overlooked by front-end algorithms (e.g. image matching), see e.g. [3], [7], [8], [13], [15]. However these approaches, also known as robust graph optimization approaches, are developed only for the passive problem setting. In contrast, we consider a complimentary *active* framework and incorporate data association aspects within BSP.

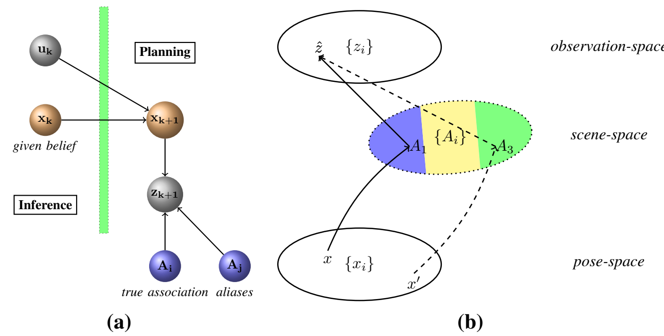

Figure 1: (a) Generative graphical model. While standard BSP approaches typically assume DA is given and perfect, we incorporate data association aspects within BSP and thus capable of reasoning about ambiguity (e.g. perceptual aliasing) at a decision-making level. (b) Schematic representation of pose, scene and observation spaces. Scenes A_1 and A_3 when viewed from perspective x and x' respectively, produce the same nominal observation $\hat{z}$, giving rise to *perceptual aliasing*.

Our approach is also tightly related with recent work on active hypothesis disambiguation in the context object detection and classification [2], [12], [14], [16], [18]. However, these approaches assume the sensor is perfectly localized and thus the corresponding belief is only about the considered hypotheses.

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. In this work the authors only consider ambiguous data association within the prior belief, modelling it as mixture of Gaussians, and assume there indeed exists an action that yields complete disambiguation. In contrast, our framework is more general since we additionally consider ambiguous data association within future belief (due to future observations) given candidate action(s) and do not assume there is necessarily a fully-disambiguating action. In this work we develop a general data association aware belief space planning (DA-BSP) framework capable of better handling complexities arising in real world, possibly perceptually aliased, scenarios. To that end, we rigorously incorporate reasoning about data association within belief space planning and in particular show that our framework can be used for active disambiguation by determining appropriate actions, e.g. future viewpoints, for increasing confidence in a certain data association hypothesis.

*The authors are with Department of Aerospace Engineering, Technion - Israel Institute of Technology, Haifa 32000, Israel.

II. Problem Formulation

Consider a robot, uncertain about its pose, operating in a partially known or pre-mapped environment. The robot takes observations of different scenes or objects in the environment and uses these observations to infer random variables of interest which are application-dependent. A schematic equivalent to this is shown in Figure 1,which involves three spaces: *pose-space*, *scene-space* and *observation-space*. *Pose-space* involves all possible perspectives a robot can take with respect to a given world model and in the context of task at hand.

We shall denote a particular pose at any time step k as x_k, and the sequence of these poses from 0 up to k as $X_k \doteq \{x_0, \ldots, x_k\}$. By uncertainty in robot's pose, we mean that the current pose of robot at any step k, is known only through a posterior probability distribution function (pdf) $\mathbb{P}(X_k|u_{0:k-1}, Z_{0:k})$ 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 k. For notational convenience, we define 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 we rewrite the posterior at time k as $b[X_k] \doteq \mathbb{P}(X_k|\mathcal{H}_k)$.

In contrast, *scene-space* involves a discrete set of objects or scenes, denoted by the set $\{A_{\mathbb{N}}\}$, in the given world model, and which can be detected through the sensors of the robot. We shall use symbols A_i and A_j to denote such typical scenes. Note that even if the objects are identical, they are distinct in scene space. Finally, *observation-space* is the set of all possible observations that the robot is capable of obtaining when considering its mission and sensory capabilities. We shall consider such an observation as the model:

$$z_k = h(x_k, A_i) + v_k \quad , \quad v_k \sim \mathcal{N}(0, \Sigma_v), \tag{1}$$

and represent it probabilistically as $\mathbb{P}(z_k|x_k, A_i)$.

We also consider a standard motion model $x_{i+1} = f(x_i, u_i) + w_i$ with Gaussian noise $w_i \sim \mathcal{N}(0, \Sigma_w)$, where Σ_w is the process noise covariance, and denote this model probabilistically by $\mathbb{P}(x_{i+1}|x_i, u_i)$. Given a prior $\mathbb{P}(x_0)$ and motion and observation models, the joint posterior pdf at the current time k can be written as

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

This pdf is thus a Gaussian $\mathbb{P}(X_k|\mathcal{H}_k) = \mathcal{N}(\hat{X}_k, \Sigma_k)$ with mean $\hat{X}_k$ and covariance Σ_k that can be efficiently calculated via maximum a posteriori (MAP) inference, see e.g. [10]. It is important to note that the underlying assumption in factorisation (2) is that it is known which object is being observed at each time i, i.e. data association is given and error-free.

For notational convenience we will often represent the posterior $\mathbb{P}(X_{k+1}|\mathcal{H}^-_{k+1}, z_{k+1})$ as the *belief* $b[X_{k+1}]$, i.e.:

$$b[X_{k+1}] \doteq \mathbb{P}(X_{k+1}|\mathcal{H}^-_{k+1}, z_{k+1}). \tag{3}$$

Similarly, we define the propagated joint belief as $b[X^-_{k+1}] \doteq \mathbb{P}(X_{k+1}|\mathcal{H}^-_{k+1}) = \mathbb{P}(X_k|\mathcal{H}_k)\mathbb{P}(x_{k+1}|x_k, u_k)$, from which the marginal belief over the future pose x_{k+1} can be calculated as $b[x^-_{k+1}] \doteq \int_{\neg x_{k+1}} b[X^-_{k+1}]$.

As earlier, if data association is assumed given and perfect as commonly done in BSP, then one can consider for each specific value of z_{k+1} the corresponding observed scene A_i, and express the posterior (3) as

$$b[X_{k+1}] = \eta \mathbb{P}(X_k|\mathcal{H}_k)\mathbb{P}(x_{k+1}|x_k, u_k)\mathbb{P}(z_{k+1}|x_{k+1}, A_i), \tag{4}$$

which can be represented as $b[X_{k+1}] = \mathcal{N}(\hat{X}_{k+1}, \Sigma_{k+1})$ with appropriate mean $\hat{X}_{k+1}$ and covariance Σ_{k+1}. Yet, it is unknown from what future robot pose x_{k+1} the actual observation z_{k+1} will be acquired, since the *actual* robot pose x_k at time k is unknown and the control is stochastic. Indeed, as a result of action u_k, the robot actual (true) pose x_{k+1} can be anywhere within the propagated belief $b[x^-_{k+1}]$.

III. Concept and Approach

Given the posterior (2) at the current time k, one can reason about the robot's best future actions that would minimise (or maximise) a certain objective function. Such a function, for a single look ahead step, is given by

$$J(u_k) = \mathop{\mathbb{E}}_{z_{k+1}} \{c(\mathbb{P}(X_{k+1}|\mathcal{H}^-_{k+1}, z_{k+1}))\}, \tag{5}$$

where the expectation is taken about the random variable z_{k+1} with respect to the propagated belief $\mathbb{P}(X_{k+1}|\mathcal{H}^-_{k+1})$ to consider all possible realisations of a future observation z_{k+1}.

To see that, we write the expectation operator explicitly which transforms Eq. (5) to

$$J(u_k) \doteq \int_{z_{k+1}} \overbrace{\mathbb{P}(z_{k+1} \mid \mathcal{H}^-_{k+1})}^{(a)} \; c \left(\overbrace{\mathbb{P}(X_{k+1}|\mathcal{H}^-_{k+1}, z_{k+1})}^{(b)} \right) \tag{6}$$

The two terms (a) and (b) in the above equation have intuitive meaning: for each considered value of z_{k+1}, (a) represents how likely is it to get such an observation when both the history $\mathcal{H}$ and control u_k are known, while (b) corresponds to the posterior belief *given* this specific z_{k+1}.

A. Computing term (a): $\mathbb{P}(z_{k+1} \mid \mathcal{H}^-_{k+1})$

Applying total probability over non-overlapping $\{A_{\mathbb{N}}\}$ and marginalizing over all possible robot poses, yields

$$\mathbb{P}(z_{k+1} \mid \mathcal{H}^-_{k+1}) \equiv \sum_i \int_x \mathbb{P}(z_{k+1}, x, A_i \mid \mathcal{H}^-_{k+1}) \doteq \sum_i w_i. \tag{7}$$

As seen from the above equation, to calculate the likelihood of obtaining some observation z_{k+1}, 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; these weights are then summed to get the actual likelihood of observation z_{k+1}. As will be seen below, these weights naturally account for perceptual aliasing aspects for each considered z_{k+1}.

Proceeding with the derivation further, using the chain rule we get $\sum_i \int_x \mathbb{P}(z_{k+1} \mid x, A_i, \mathcal{H}^-_{k+1})\mathbb{P}(A_i, x \mid \mathcal{H}^-_{k+1})$. Since this integral could be over any arbitrary total distribution of x, we can use the propagated belief $b[x^-_{k+1}]$, giving:

$$w_i \doteq \int_x \mathbb{P}(z_{k+1}|x, A_i, \mathcal{H}^-_{k+1})\mathbb{P}(A_i|\mathcal{H}^-_{k+1}, x) b[x^-_{k+1} = x]. \tag{8}$$

Here, $\mathbb{P}(z_{k+1} \mid A_i, x, \mathcal{H}_{k+1}^-) \equiv \mathbb{P}(z_{k+1} \mid A_i, x)$ is the standard measurement likelihood term, while $\mathbb{P}(A_i \mid \mathcal{H}_{k+1}^-, x)$ 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 and thus, in this work, we consider this term to be given.

The weights w_i (8) naturally capture *perceptual aliasing* aspects: consider some observation z_{k+1} and the corresponding generative model $z_{k+1} = h(x^{tr}, A^{tr}) + v$ with appropriate unknown *true* robot pose x^{tr} and scene $A^{tr} \in \{A_{\mathbb{N}}\}$. Clearly, the measurement likelihood $\mathbb{P}(z_{k+1} \mid x, A_i, \mathcal{H}_{k+1}^-)$ will be high when evaluated for $A_i = A^{tr}$ and in vicinity of x^{tr}. Note that we will necessarily consider such a case, since according to Eq. (7) we separately consider each scene A_i in $\{A_{\mathbb{N}}\}$, and, given A_i, we reason about all poses x in Eq. (8). In case of perceptual aliasing, however, there will be also another scene(s) A_j which could generate the same observation z_{k+1} from appropriate robot pose x', i.e. $\{A_i, A_j\}_{\textbf{aliased}}$. Thus, the corresponding measurement likelihood term to A_j will also be high for x'.

However, the actual value of w_i (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 belief $b[x_{k+1}^-]$, with the latter weighting the probability of each considered robot pose. This correctly captures the intuition that those observations z with low-probability poses $b[x_{k+1}^- = x^{tr}]$ will be unlikely to be actually acquired, leading to low value of w_i with $A_i = A^{tr}$. However, the likelihood term (7) could still go up in case of perceptual aliasing, where the aliased scene A_j generates a similar observation to z_{k+1} from viewpoint x' with latter being more probable, i.e. high probability $b[x_{k+1}^- = x']$.

B. Computing term (b): $\mathbb{P}(X_{k+1}|\mathcal{H}_{k+1}^-, z_{k+1})$

The term (b), $\mathbb{P}(X_{k+1}|\mathcal{H}_{k+1}^-, z_{k+1})$, represents the posterior probability conditioned on observation z_{k+1}. Since the observation z_{k+1} 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 $\tilde{w}_i$ representing the probability of each scene A_i to have generated the observation z_{k+1}.

Applying total probability over non-overlapping $\{A_{\mathbb{N}}\}$ and chain-rule, we get:

$$\sum_i \mathbb{P}(X_{k+1} \mid \mathcal{H}_{k+1}^-, z_{k+1}, A_i) \cdot \mathbb{P}(A_i \mid \mathcal{H}_{k+1}^-, z_{k+1}). \quad (9)$$

Here, the first term is the posterior belief conditioned on observations, history as well as a candidate scene A_i that supposedly generated the observation z_{k+1}. The second term, $\mathbb{P}(A_i \mid \mathcal{H}_k, u_k, z_{k+1})$, is merely the likelihood of A_i being actually the one which generated the observation z_{k+1}. Marginalising over all robot poses and applying Bayes rule yields

$$\mathbb{P}(A_i \mid \mathcal{H}_{k+1}^-, z_{k+1}) = \eta \int_x \mathbb{P}(z_{k+1} \mid A_i, x, \mathcal{H}_{k+1}^-)\mathbb{P}(A_i, x \mid \mathcal{H}_{k+1}^-) \quad (10)$$

Proceeding in a similar way as in section section (III-A),it can be shown that this term is actually the normalised weight w_i, such that $\sum_i \tilde{w}_i = 1$. Hence, $\mathbb{P}(A_i \mid z_{k+1}, \mathcal{H}_{k+1}^-) = \eta w_i \doteq \tilde{w}_i$.

To summarise the discussion thus far, we have shown that the objective function (6) can be re-written as

$$J(u_k) = \int_{z_{k+1}} (\sum_i w_i) \cdot c\left(\sum_i \tilde{w}_i b[X_{k+1}^{i+}]\right), \quad (11)$$

with the posterior given scene A_i defined as

$$b[X_{k+1}^{i+}] \doteq \mathbb{P}(X_{k+1} \mid \mathcal{H}_{k+1}^-, z_{k+1}, A_i). \quad (12)$$

Observe, that for each considered observation z_{k+1}, we get a *mixture pdf* inside of the cost $c(.)$, where each component represents the posterior conditioned on the observation capturing scene A_i, and weighted by $\tilde{w}_i$. In case there is no perceptual aliasing, there will be only one component with high weight $\tilde{w}_i$, that corresponds to the correct data association to scene A_i, with all other weights being negligible. On the other hand, in presence of perceptual aliasing, we expect to see numerous non-negligible weights. In the extreme case, where all scenes (objects) are identical, we will get equal normalised weights $\tilde{w}_i$ for each $A_i \in \{A_{\mathbb{N}}\}$.

C. An abstract example of data-association aware BSP

Consider the problem of robotic manipulation of objects in the kitchen. For simplicity, let us abstract it to a simpler domain of three objects, $|\{A_{\mathbb{N}}\}| = 3$. We consider a single step control at time step k, from a given belief $b[X_k]$, as well as that of one step ahead $b[X_{k+1}^-]$, and assume the following motion and observation models f and h

$$f(x,u) = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \cdot x + d \begin{cases} [0,1]^T & \text{if } u = \textit{up} \\ [1,0]^T & \text{if } u = \textit{right} \end{cases}, \quad h(x, A_i) = h_i(x) = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \cdot (x - x_i) + s_i. \quad (13)$$

where observations as well as the shift s_i is in an object-centric frame, with x_i representing location of A_i. Intuitively, s_i is a simple mechanism to model perceptual aliasing between objects; e.g., identical objects A_i would have the same s_i.

Figures 2a-2d denote the situation when the true pose x^{tr} is close to center and observe A_2, while in Figures 2e-2h it is at the left side and observe A_1. Different degrees of aliasing are considered. Both weights w_i and $\tilde{w}_i$ are shown in the inset histograms. Note that the unnormalised weight w_i is higher when the object is at the centre, because the overall likelihood of the observation is higher. Also, with no aliasing, for any other scene A_j than the true one, the normalised weight w_j is small irrespective of where x^{tr} is. In other words, weights are also related to how likely the objects are to be the causes behind an observation; in case of no aliasing, this can be negligibly small. Thus, DA-BSP in practical applications with infrequent aliasing, would not require any significant additional computational effort w.r.t. usual BSP.

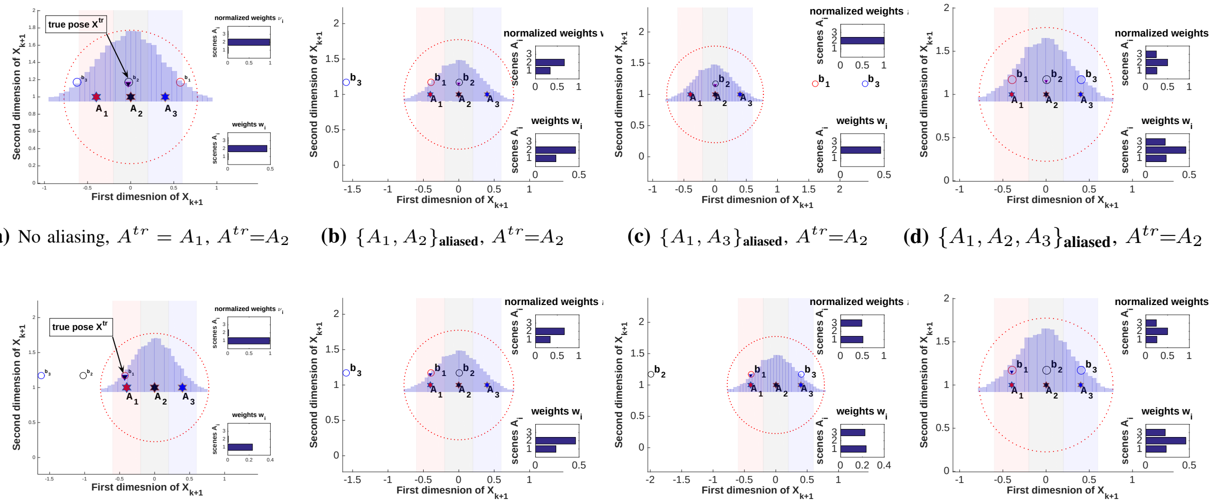

(a) No aliasing, $A^{tr} = A_1$, A^{tr}=A_2 **(b)** $\{A_1, A_2\}_{\textbf{aliased}}$, A^{tr}=A_2 **(c)** $\{A_1, A_3\}_{\textbf{aliased}}$, A^{tr}=A_2 **(d)** $\{A_1, A_2, A_3\}_{\textbf{aliased}}$, A^{tr}=A_2

(e) No aliasing, $A^{tr} = A_1$ **(f)** $\{A_1, A_2\}_{\textbf{aliased}}$, A^{tr}=A_1 **(g)** $\{A_1, A_3\}_{\textbf{aliased}}$, A^{tr}=A_1 **(h)** $\{A_1, A_2, A_3\}_{\textbf{aliased}}$, A^{tr}=A_1

Figure 2: DA-BSP for a single observation z_{k+1}. Red-dotted ellipse denotes $b[X^-_{k+1}]$, while the true pose that generated z_{k+1} is shown by inverted triangle. Smaller ellipses are the posterior beliefs $b[X^{i+}_{k+1}]$. *Top row* x^{tr} is near center, observing A_2; *bottom row* x^{tr} is on the left, observing A_1. Columns represent different perceptual aliasing cases. Weights w_i and $\tilde{w}_i$, corresponding to each scene A_i are shown in the inset bar-graphs.

IV. Acknowledgments

This work was partially supported by the Technion Autonomous Systems Program (TASP).

V. Conclusions

State-of-the-art BSP approaches typically consider data association to be given and perfect. In this work, we developed a DA-BSP approach that relaxes this assumption. Our framework rigorously incorporates data association aspects within BSP, while considering different sources of uncertainty (uncertainty in robot motion, sensing and possibly in the observed environment). As such, it is better suited to cope 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. Potential directions for future research include extension to non-myopic planning and quantitative evaluation through real-world experiments.

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] 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.

[3] 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.

[4] S. M. Chaves, A. Kim, and R. M. Eustice. Opportunistic sampling-based planning for active visual slam. In *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, pages 3073–3080. IEEE, 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] 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.

[7] V. Indelman, E. Nelson, J. Dong, N. Michael, and F. Dellaert. Incremental distributed inference from arbitrary poses and unknown data association: Using collaborating robots to establish a common reference. *IEEE Control Systems Magazine (CSM), Special Issue on Distributed Control and Estimation for Robotic Vehicle Networks*, 36(2):41–74, 2016.

[8] V. Indelman, E. Nelson, N. Michael, and F. Dellaert. Multi-robot pose graph localization and data association from unknown initial relative poses via expectation maximization. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2014.

[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] M. Lauri, N. A. Atanasov, G. Pappas, and R. Ritala. Active object recognition via monte carlo tree search. In *Workshop on Beyond Geometric Constraints at the International Conference on Robotics and Automation (ICRA)*, 2015.

[13] Edwin Olson and Pratik Agarwal. Inference on networks of mixtures for robust robot mapping. *Intl. J. of Robotics Research*, 32(7):826–840, 2013.

[14] Bharath Sankaran, Jeannette Bohg, Nathan Ratliff, and Stefan Schaal. Policy learning with hypothesis based local action selection. *arXiv preprint arXiv:1503.06375*, 2015.

[15] N. Sunderhauf and P. Protzel. Towards a robust back-end for pose graph slam. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pages 1254–1261. IEEE, 2012.

[16] R. Fitch M. Vincze T. Patten, M. Zillich and S. Sukkarieh. Viewpoint evaluation for online 3-d active object classification. *IEEE Robotics and Automation Letters (RA-L)*, 1(1):73–81, January.

[17] Jeffrey M Walls, Stephen M Chaves, Enric Galceran, and Ryan M Eustice. Belief space planning for underwater cooperative localization. In *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2015.

[18] 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.