

A unified framework for data association aware robust belief space planning and perception

The International Journal of
Robotics Research
1–29

© The Author(s) 2018
Reprints and permissions:
sagepub.co.uk/journalsPermissions.nav
DOI: 10.1177/0278364918759606
journals.sagepub.com/home/ijr



Shashank Pathak , **Antony Thomas**, and **Vadim Indelman**

Abstract

We develop a belief space planning approach that advances the state of the art by incorporating reasoning about data association within planning, while considering additional sources of uncertainty. Existing belief space planning approaches typically assume that data association is given and perfect, an assumption that can be harder to justify during operation in the presence of localization uncertainty, or in ambiguous and perceptually aliased environments. By contrast, our data association aware belief space planning (DA-BSP) approach explicitly reasons about data association within belief evolution owing to candidate actions, and as such can better accommodate these challenging real-world scenarios. In particular, we show that, owing to perceptual aliasing, a posterior belief can become a mixture of probability distribution functions and design cost functions, which measure the expected level of ambiguity and posterior uncertainty given candidate action. Furthermore, we also investigate more challenging situations, such as when prior belief is multimodal and when data association aware planning is performed over several look-ahead steps. Our framework models the belief as a Gaussian mixture model. Another unique aspect of this approach is that the number of components of this Gaussian mixture model can increase as well as decrease, thereby reflecting reality more accurately. Using these and standard costs (e.g. control penalty, distance to goal) within the objective function yields a general framework that reliably represents action impact and, in particular, is capable of active disambiguation. Our approach is thus applicable to both robust perception in a passive setting with data given a priori and in an active setting, such as in autonomous navigation in perceptually aliased environments. We demonstrate key aspects of DA-BSP in a theoretical example, in a Gazebo-based realistic simulation, and also on the real robotic platform using a Pioneer robot in an office environment.

Keywords

Localization, mobile and distributed SLAM, mapping, autonomous agents, cognitive robotics, AI reasoning methods

1. Introduction

Belief space planning and decision making under uncertainty are fundamental problems in robotics and artificial intelligence, with applications including autonomous navigation, object grasping and manipulation, active simultaneous localization and mapping (SLAM), and robotic surgery. In the presence of uncertainty, such as uncertainty 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 of possible states given available data.

Planning and decision making should therefore be performed over that distribution of possible states, the belief space, that 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) (Kaelbling et al., 1998), where, given an objective function,

one aims to determine an optimal control policy as a function of belief evolution over variables of interest, which are application-dependent.

However, state-of-the-art belief space planning approaches typically assume data association to be given and perfect (see Figure 1b), i.e. the robot is assumed to correctly perceive the environment to be observed by its sensors given a candidate action. Yet, the world is often full of ambiguity, which, together with other sources of uncertainty, make perception a challenging task. As an example, one might consider matching images from two different but similar-in-appearance places, or attempting

Department of Aerospace Engineering, Technion - Israel Institute of Technology, Haifa, Israel

Corresponding author:

Shashank Pathak, Department of Aerospace Engineering, Technion - Israel Institute of Technology, Haifa 32000, Israel.

Email: shashank@computer.org

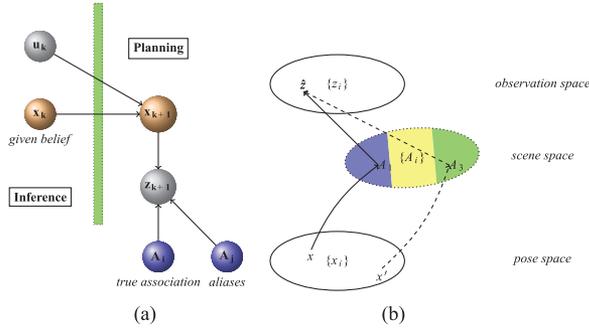


Fig. 1. (a) Generative graphical model. While standard belief space planning approaches typically assume that data association is given and perfect, we incorporate data association aspects within belief space planning and thus enable reasoning about ambiguity (e.g. perceptual aliasing) at a decision-making 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.

to recognize an object that is similar in appearance, from the current viewpoint, to another object. Both cases are examples of ambiguous situations, where naïve and straightforward approaches are likely to yield incorrect results, i.e. mistakenly considering the two places to be the same place, and incorrectly associating the observed object.

Considering data association to be solved and perfect within belief space planning can thus lead, in the presence of ambiguity, to incorrect posterior beliefs and, as a result, to sub-optimal actions that do not properly consider perceptual aliasing aspects. More advanced approaches are therefore required to enable reliable operation in ambiguous conditions, approaches often referred to as (active) robust perception. These approaches typically involve probabilistic data association and hypothesis tracking, given available data. Thus, for the object detection example, each hypothesis may represent a candidate object from a given database to which the current observation (e.g. image or point-cloud) is successfully registered. Similarly, one might reason probabilistically regarding perceptual aliasing, as in the first example, which would also involve probabilistic data association. Yet, existing robust perception approaches focus on the passive case, where robot actions are externally determined and given, while the closely related approaches for active object detection and classification consider the robot to be perfectly localized.

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 this end, we rigorously incorporate reasoning about data association within belief space planning, while also considering other sources of uncertainty (motion, sensing, and environment). In particular, we 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.

1.1. Related work

Calculating optimal solutions to POMDP is computationally intractable (PSPACE complete) (Papadimitriou and Tsitsiklis, 1987) for all but the smallest problems. The research community has thus been extensively investigating approximate approaches to reduce computational complexity and facilitate their application to real-world problems. These approaches can be roughly segmented into point-based value iteration methods (Kurniawati et al., 2008; Pineau et al., 2006), simulation-based (Stachniss et al., 2005) and sampling-based approaches (Agha-Mohammadi et al., 2014; Bry and Roy, 2011; Prentice and Roy, 2009), and direct trajectory optimization (Indelman et al., 2015; Patil et al., 2014; Van den Berg et al., 2012) methods. In all cases, finding the (locally) optimal actions involves the evaluation of a given objective function while considering future observations to be acquired as a result of each candidate action.

However, an underlying typical assumption in these approaches is that data association for these future observations is known and perfect. For example, it is typically assumed that the robot can be localized by making observations of known landmarks or beacons (see, e.g. Agha-Mohammadi et al. (2014) and Prentice and Roy (2009)), while assuming that each future measurement is correctly associated with an appropriate landmark. Although such an assumption is reasonable in certain scenarios, it becomes unrealistic when considering perceptually aliased environments (two scenes that look alike) and localization uncertainty, as we do in this study.

While belief space planning approaches typically assume the environment to be accurately known (e.g. a given map), recent works, including those of Chaves et al. (2014, 2015), Indelman et al. (2015), Kim and Eustice (2014), and Walls et al. (2015), 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 (Chaves et al., 2015; Indelman et al., 2015; Kim and Eustice, 2014; Walls et al., 2015) in this branch focused, in particular, on probabilistically modeling what future observations would be obtained given a candidate action. However, these approaches consider each such future observation to be correctly associated with 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. Carlone et al. (2014), Indelman

et al. (2014b, 2016), Olson and Agarwal (2013), and Sunderhauf and Protzel (2012). However these approaches, also known as robust graph optimization approaches, are developed only for the passive problem setting, i.e. robot actions are given and externally determined. By contrast, we consider a complimentary *active* framework and incorporate data association aspects within belief space planning.

Our approach is also tightly related to recent work on active hypothesis disambiguation in the context of object detection and classification (Atanasov et al., 2014; Lauri et al., 2015; Patten et al., 2016; Sankaran et al., 2015; Wong et al., 2015). 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 that the sensor is perfectly localized and thus the corresponding belief is only about the considered hypotheses. Also, during such data associations, a null hypothesis may be considered, where the scene or the object is assumed to be novel and undetected in the previous instances. See, e.g. Fourie et al. (2016), Indelman et al. (2016), and Pfungsthorh and Birk (2016). In this work, however, we assume that the set of scenes are known a priori, even though their locations might be uncertain. In a slightly different context, there are approaches to perform data-driven disambiguation, such as in multihypothesis tracking. Our approach, however, could be considered a more general one, where both passive and active robust perception occur within a unified framework.

Probably the closest work to our approach is that of Agarwal et al. (2015), who also consider hypotheses due to ambiguous data association and develop a belief space planning approach for active disambiguation. However, in their work, Agarwal et al. (2015) only consider ambiguous data association within the prior belief, modeling it as a mixture of Gaussians, and assume that there indeed exists an action that can yield complete disambiguation. By contrast, our framework is more general, since we additionally consider ambiguous data association within future belief (owing to future observations) about given candidate actions and do not assume that there is necessarily a fully disambiguating action. Consequently, in our approach, the number of modes may decrease or increase as is suitable for the given environment; ambiguity would persist wherever full disambiguation is not possible.

1.2. Contributions

This paper is an extension of the preliminary work presented in Pathak et al. (2017), where the relaxation of a known data association was first considered, albeit in a myopic setting. Further contributions of this manuscript are as follows:

1. We develop a unified framework for DA-BSP in both active and passive contexts. Here, the components of

our belief may both increase and decrease, thereby modeling the perceptually aliased environment more faithfully. Additionally, the framework does not require a fully disambiguating unique observation.

2. We extend DA-BSP by considering prior belief as a Gaussian mixture model, as well as by considering planning for several look-ahead steps.
3. We show how under helpful assumptions this general approach degenerates to known belief space planning approaches.
4. We present a complexity analysis of such an algorithm discrete state settings and comment on its correctness.
5. Finally, we analyze key aspects arising as a result of explicitly considering data association aspects within belief space planning in a simplified toy example, as well as in extensive experiments with a realistic synthetic simulation and in a real robotics scenario using a Pioneer robot.

1.3. Organization

The rest of this paper is organized as follows. We formulate the considered problem in Section 2 and then provide concept overview in Section 3. Thereafter, we consider a simpler case of myopic data association aware planning in Section 4. Later, in Section 5, we generalize the planning to consider a Gaussian mixture model prior and subsequently to the non-myopic setting. Then, Section 6 demonstrates key aspects in realistic simulations as well as a real-world scenario with a Pioneer robot platform. Section 7 concludes the discussion and suggests potential directions for future research.

2. Notation and problem formulation

Consider a robot, uncertain about its pose, operating in a partially known or premapped environment. The robot takes observations of different scenes or objects in the environment and uses these observations to infer random variables of interest that are application-dependent. Thus, in localization, these observations can be used to better estimate the robot pose, while in search and rescue missions, one is looking for survivors in a certain region.

A schematic equivalent to this is shown in Figure 1. As can be seen, it involves three spaces: *pose space*, *scene space*, and *observation space*. The pose space involves all the possible perspectives a robot can take with respect to a given world model and in the context of the task in hand.

We shall denote a particular pose at any time step k as x_k , and the sequence of these poses from 0 to k as $X_k \doteq \{x_0, \dots, x_k\}$. By uncertainty in the robot's pose, we mean that the current pose of the robot, at any step k , is known only through an a-posteriori probability distribution function $\mathbb{P}(X_k | u_{0:k-1}, Z_{0:k})$, given all controls $u_{0:k-1} \doteq \{u_0, \dots, u_{k-1}\}$ and observations $Z_{0:k} \doteq \{Z_0, \dots, Z_k\}$ up to time k . For notational convenience, we define histories \mathcal{H}_k

and \mathcal{H}_{k+1}^- as

$$\mathcal{H}_k \doteq \{u_{0:k-1}, Z_{0:k}\}, \quad \mathcal{H}_{k+1}^- \doteq \mathcal{H}_k \cup \{u_k\} \quad (1)$$

and we rewrite the posterior at time k as $b[X_k] \doteq \mathbb{P}(X_k|\mathcal{H}_k)$.

By contrast, the scene space involves a discrete set of objects (or a collection of objects)—called scenes and denoted by the set $\{A_{\mathbb{N}}\}$ —in the given world model. Scenes can be detected through the robot's sensors. 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. This is important when we consider the cases where the objects look similar from some perspectives. Finally, the 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

$$z_k = h(x_k, A_i) + v_k, \quad v_k \sim \mathcal{N}(0, \Sigma_v) \quad (2)$$

and represent it probabilistically as $\mathbb{P}(z_k|x_k, A_i)$. Here, we have assumed the same Gaussian noise for all observations, irrespective of the scenes being observed. This is a reasonable assumption, since such noise would be a typical property of the robotic sensors employed. Also, $h(x_k, A_i)$ is a noise-free observation, which we would refer to as a *nominal* observation \hat{z} .

For example, in the case of a camera, the function h could be defined as a pinhole projection operator, thereby projecting the object A_i onto the image plane, while in the case of a range sensor, this function calculates the range between (a particular point on) the object and the robot actual location.

Note that the exposition thus far is equivalently valid in the case where the environment model is given but uncertain, and when this model is unknown a priori and instead constructed online within a SLAM framework.

We also consider a standard motion model with Gaussian noise

$$x_{i+1} = f(x_i, u_i) + w_i, \quad w_i \sim \mathcal{N}(0, \Sigma_w) \quad (3)$$

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 distribution $\mathbb{P}(x_0)$ and motion and observation models, the joint posterior probability distribution function at the current time k can be written as

$$\mathbb{P}(X_k|\mathcal{H}) = \mathbb{P}(x_0) \prod_{i=1}^k \mathbb{P}(x_i|x_{i-1}, u_{i-1}) \mathbb{P}(Z_i|x_i, A_i) \quad (4)$$

This posterior probability distribution function is thus a multivariate 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 inference, see e.g. Kaess et al. (2012).

It is important to note that the underlying assumption in factorization (equation (4)) is that it is known which object

is being observed at each time i , i.e. data association is given and error-free. We will come back to this key point later.

Given the posterior probability distribution (equation (4)) at the current time k , one can reason about the robot's best future actions that would minimize (or maximize) a certain objective function. Such a function, for a single look-ahead step, is given by

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

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 realizations of a future observation z_{k+1} .

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

$$b[X_{k+1}] \doteq \mathbb{P}(X_{k+1}|\mathcal{H}_{k+1}^-, z_{k+1}) \quad (6)$$

Note that, according to equation (5), we need to calculate the posterior belief (equation (6)) for *each* possible value of z_{k+1} .

Similarly, we define the propagated joint belief as

$$b[X_{k+1}^-] \doteq \mathbb{P}(X_{k+1}|\mathcal{H}_{k+1}^-) = \mathbb{P}(X_k|\mathcal{H}_k) \mathbb{P}(x_{k+1}|x_k, u_k) \quad (7)$$

from which the marginal belief over the future pose x_{k+1} can be calculated as

$$b[x_{k+1}^-] \doteq \int_{-x_{k+1}} b[X_{k+1}^-] \quad (8)$$

As earlier, if data association is assumed given and perfect, as is commonly the case in belief space planning, then one can consider, for each specific value of z_{k+1} , the corresponding observed scene A_i , and express the posterior (equation (6)) as

$$b[X_{k+1}] = \eta \mathbb{P}(X_k|\mathcal{H}_k) \mathbb{P}(x_{k+1}|x_k, u_k) \mathbb{P}(z_{k+1}|x_{k+1}, A_i) \quad (9)$$

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} .

The objective function (equation (5)) can now be evaluated, given a candidate action u_k , by calculating the cost $c(\cdot)$ for each z_{k+1} . Finally, the optimal action u_k^* is defined as

$$u_k^* = \arg \min_{u_k} J(u_k) \quad (10)$$

Assuming data association to be given and perfect greatly simplifies the above formulation. Yet, in practice, determining data association reliably is often a non-trivial task in itself, especially when operating in perceptually aliased environments. An incorrect data association (wrong scene A_i in equation (9)) can lead to catastrophic results, see, e.g. Indelman et al. (2014a,b, 2016). In this work, we relax this restricting assumption and rigorously incorporate data association aspects within belief space planning.

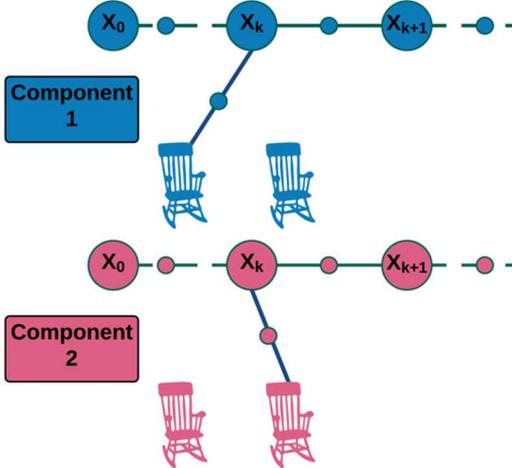


Fig. 2. Example of perceptual aliasing when observing one of two identical rocking chairs. Each rocking chair creates a Gaussian belief component; the corresponding factor graphs are shown here.

3. Concept and approach overview (DA-BSP)

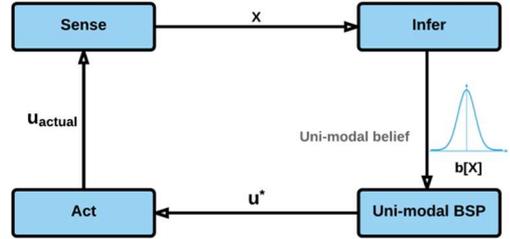
Consider Figure 2, where an observation regarding a rocking chair is made. Since the environment contains two identical rocking chairs, data association with either of these *scenes* cannot be deemed to be solved. Instead, we must consider *every* plausible data association and its corresponding component in the posterior belief. Note that what we will describe as a scene need not necessarily be a single object; e.g., later on we shall consider a cubicle with a desk and a chair as a single scene. We shall now discuss the overall approach by making the following observations, which will be transformed into a rigorous mathematical framework in the following sections.

3.1. Data association while planning

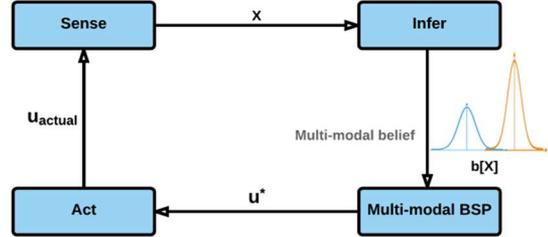
Given some candidate action (or sequence of actions) and the belief at planning time k , we can reason about a future observation z_{k+1} (e.g. an image) to be obtained once this action is executed. This future observation is yet to be acquired, and therefore its actual value is unknown. For this reason, all the possible values such an observation can assume should be taken into account while evaluating the objective function; hence, the expectation operator in equation (5). To see this, we write the expectation operator explicitly, which transforms equation (5) to

$$J(u_k) \doteq \int_{z_{k+1}} \overbrace{\mathbb{P}(z_{k+1} | \mathcal{H}_{k+1}^-)}^{(a)} c \left(\overbrace{\mathbb{P}(X_{k+1} | \mathcal{H}_{k+1}^-, z_{k+1})}^{(b)} \right) \quad (11)$$

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



(a) Usual belief space planning.



(b) Data association aware belief space planning.

Fig. 3. Relation of DA-BSP with the usual approaches within belief space planning (BSP). Note that in DA-BSP both planning and inference consider multimodal beliefs. Thus, DA-BSP is capable of handling a general Gaussian mixture model distributions in both *passive* and *active* data association setups.

to the posterior belief $b[X_{k+1}]$ from equation (6), *given* this specific observation z_{k+1} .

Considering that the data association is solved and perfect means we can correctly associate each possible measurement z_{k+1} with the corresponding scene A_i that it captures, as in equation (9).

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's actual (true) pose x_{k+1} can be anywhere within the propagated belief $b[x_{k+1}^-]$.

In typical cases, such as with navigation assisted through GPS, this *data association* is trivially known, since the scene coincides with the pose. However, in more complex applications, such as perceptual robotics, the observations could come from several different poses of viewing different scenes. In the belief space planning framework, such a data association is assumed to be solved. In other words, if \mathcal{A} represents the total space of scenes (or the real world) from where *all* observations $\{z\}$ are made and $\{A_N\}$ is the partitioning of this scene space, then belief space planning assumes that for each such observation $z \in \{z\}$ the corresponding observed scene $A_i \in \mathcal{A}$ is known.

In contrast, in this work, we do not assume that data association is solved, and instead reason about possible scenes or objects that the future observation z_{k+1} could be generated from, see Figure 1. Clearly, if the environment has only distinct scenes or objects, then for each specific value of

z_{k+1} , there will be only one scene A_i that can generate such an observation according to the model (equation (2)). However, in the case of perceptually aliased environments, there could also be several scenes (or objects) that are either completely identical, or have a similar visual appearance when observed from appropriate viewpoints that could equally well explain the considered observation z_{k+1} . In such a case, there are several possible associations $\{A_i\}$; owing to localization uncertainty, determining which association is the correct one is not trivial. As we show later, in these cases, the posterior belief $b[X_{k+1}]$ (term (b) in equation (11)) becomes a Gaussian mixture with appropriate weights that we rigorously compute (See Fig. 3 for the comparison).

In the following sections, we formalize these aspects probabilistically and develop an algorithm for data association aware belief space planning, capable of determining the best actions in perceptual aliased and distinct environments without considering that data association is solved. First, however, we formally define what we mean by perceptually aliasing.

3.1.1. Perceptual aliasing. Intuitively speaking, perceptual aliasing occurs when an object different from the actual one produces the same observation and thereby provides an alias, in the sense of perception, to the true object. We shall now define two notions of perceptual aliasing that we consider: *exact* and *probabilistic*. Exact perceptual aliasing of scenes A_i and A_j is defined as $\exists x, x', h(x, A_i) = h(x', A_j)$, and will be denoted in this paper by $\{A_i, A_j\}_{\text{alias}}$. In other words, the same nominal (noise-free) observation \hat{z} can be generated by observing different scenes, possibly from different viewpoints. Such a situation is depicted in Figure 1. A probabilistic perceptual aliasing is a more general form of aliasing, and can be defined as $\exists x, x', |\mathbb{P}(z|A_i, x) - \mathbb{P}(z|A_j, x')| < \epsilon$ for some small threshold ϵ .

If there is a unique feature in the environment that does not alias with any observation anywhere in the world and hence allows what we call *full disambiguation*, this would imply that eventually all multicomponent beliefs would be reduced to a single component belief. This can indeed be a very useful property. However, in some practical cases, such a guarantee cannot be made. In DA-BSP, we do not assume that a fully disambiguating feature exists in the environment. On the contrary, only *partial disambiguation* might be available. This hardens the planning problem but, as we shall see in this work, DA-BSP successfully tackles it by considering the data association (and its ambiguities) within the belief space planning framework.

3.2. Data association while perceiving

By inference, we have a similar situation, with the key difference that the observation z is given, i.e. it has been acquired. Let us now consider this setting for a moment.

Also here, the robot pose at measurement acquisition time is unknown—rather, we are trying to estimate it. To do so, we must first associate the captured measurement z with the scene or object A_i that it describes, i.e. write the appropriate measurement likelihood term in the posterior probability distribution (equation (4)).

A similar situation, however, also arises in our case: while the probability of acquiring a specific observation z_{k+1} is represented by the term (a) in equation (11), the posterior in the term (b) is *conditioned on this specific* observation z_{k+1} . As such, evaluating the posterior probability distribution given z_{k+1} involves inference, as if that observation were actually acquired. Thus, also here data association needs to be resolved or to be assumed given.

The DA-BSP framework can be used to calculate an optimal control u_k^* . On execution of the control u_k^* , the robot takes an *actual* observation z_{k+1} , which, given the Gaussian mixture model belief $b[X_k]$ (see equation (44)), can be used to calculate the posterior probability distribution at time t_{k+1} . Interestingly, while within planning at time t_k we considered all possible realizations $\{z_{k+1}\}$ of the future observation z_{k+1} , in inference we get some specific observation z_{k+1} . Consequently, the posterior probability distribution in inference is a specific instantiation of the different posterior probability distributions $\mathbb{P}(X_{k+1} | \mathcal{H}_{k+1}^-, z_{k+1})$ already calculated within planning, see Section 3.1. Thus, similar equations apply.

In particular, we first propagate the Gaussian mixture model belief based on control u_k^* from $b[X_k]$ to $b[X_{k+1}^-]$ via equation (36) and then calculate the posterior Gaussian mixture model $b[X_{k+1}]$ conditioned on the obtained observation z_{k+1} via equation (38). Observe that exactly the same insights regarding full or partial disambiguation discussed in Section 3.1 in the context of planning apply here as well. This forms a closed loop between planning and inference; for example, an optimal action calculated by the former, while considering the ambiguity cost c_w , should indeed lead to disambiguation on execution in inference.

4. Myopic data association aware belief space planning

For the sake of simplicity, in this section we would assume myopic planning i.e., with the planning horizon $L = 1$. Recall that we need to calculate each of the two terms, (a) and (b), in equation (11). For convenience, we specify the corresponding expressions again

$$(a): \mathbb{P}(z_{k+1} | \mathcal{H}_{k+1}^-) \quad (b): \mathbb{P}(X_{k+1} | \mathcal{H}_{k+1}^-, z_{k+1}) \quad (12)$$

Before proceeding further, recall the conceptual difference between the two terms: term (a) represents the likelihood of obtaining an observation z_{k+1} , while within term (b), the observation z_{k+1} is considered as given.

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

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

$$\mathbb{P}(z_{k+1} | \mathcal{H}_{k+1}^-) \equiv \sum_i \int_x \mathbb{P}(z_{k+1}, x, A_i | \mathcal{H}_{k+1}^-) \doteq \sum_i w_i \quad (13)$$

As seen from this 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, 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

$$\sum_i \int_x \mathbb{P}(z_{k+1} | x, A_i, \mathcal{H}_{k+1}^-) \mathbb{P}(A_i, x | \mathcal{H}_{k+1}^-) \quad (14)$$

However, since this integral could be over any arbitrary total distribution of x , we can use the propagated belief $b[x_{k+1}^-]$, see equation (8), to compute it as

$$\sum_i \int_x \mathbb{P}(z_{k+1} | x, A_i, \mathcal{H}_{k+1}^-) \mathbb{P}(A_i | \mathcal{H}_{k+1}^-, x) b[x_{k+1}^- = x] \quad (15)$$

Thus

$$w_i \doteq \int_x \mathbb{P}(z_{k+1} | x, A_i, \mathcal{H}_{k+1}^-) \mathbb{P}(A_i | \mathcal{H}_{k+1}^-, x) b[x_{k+1}^- = x] \quad (16)$$

Here, $\mathbb{P}(z_{k+1} | A_i, x, \mathcal{H}_{k+1}^-) \equiv \mathbb{P}(z_{k+1} | A_i, x)$ is the standard measurement likelihood term, while $\mathbb{P}(A_i | \mathcal{H}_{k+1}^-, x)$ represents the *event likelihood*, which denotes the probability that scene A_i will 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 (equation (16)) naturally capture the *perceptual aliasing* aspects discussed in Section 3.1.1: consider some observation z_{k+1} and the corresponding generative model $z_{k+1} = h(x^{\text{tr}}, A^{\text{tr}}) + v$ with appropriate unknown *true* robot pose x^{tr} and scene $A^{\text{tr}} \in \{A_{\mathbb{N}}\}$. Clearly, the measurement likelihood $\mathbb{P}(z_{k+1} | x, A_i, \mathcal{H}_{k+1}^-)$ will be high when evaluated for $A_i = A^{\text{tr}}$ and in the vicinity of x^{tr} . Note that we will necessarily consider such a case, since according to equation (13) we separately consider each scene A_i in $\{A_{\mathbb{N}}\}$, and, given A_i , we reason about all poses x in equation (16). In the case of perceptual aliasing, however, there will be other scenes, $A_j \neq A^{\text{tr}}$, which could generate the same observation z_{k+1} from the appropriate robot pose x' . Thus, the corresponding measurement likelihood term to A_j will also be high in the vicinity of x' .

However, the actual value of w_i (for each $A_i \in \{A_{\mathbb{N}}\}$) also depends, in addition to the measurement likelihood, on the aforementioned 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 $b[x_{k+1}^- = x^{\text{tr}}]$ will be unlikely to be actually acquired, leading to a low value of w_i with $A_i = A^{\text{tr}}$. However, the likelihood term (equation (13)) could still increase in the case of perceptual aliasing, where the aliased scene A_j generates a similar observation to z_{k+1} from viewpoint x' , with the latter being more probable, i.e. high probability $b[x_{k+1}^- = x']$.

In practice, the integral in equation (20) can be calculating efficiently if both the measurement likelihood $\mathbb{P}(z_{k+1} | A_i, x, \mathcal{H})$ and the predicted belief $b[x_{k+1}^-]$ are Gaussians, since a product of Gaussians remains Gaussian. The integral can then be only calculated for the window where the event likelihood is non-zero, i.e. $\mathbb{P}(A_i | x, \mathcal{H}) > 0$. In the absence of such assumptions, in general, the integral in equation (20) should be computed numerically. Since in practical applications $\mathbb{P}(A_i | x, \mathcal{H})$ is sparse with respect to x , this computational cost is not severe. For example, for a robot navigating in a two-floor environment, even under extreme uncertainty of pose, while reasoning for a scene such as a chair, we would only consider the viewpoints from which the latter is observable, instead of the entire belief space.

4.2. 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} . This term can be similarly calculated, with a key difference: since the observation z_{k+1} is given, it must have been generated by *one* specific (but unknown) scene A_i according to the measurement model (equation (2)). Hence, also here, we consider all possible such scenes and weight them accordingly, with weights \tilde{w}_i representing the probability that each scene A_i will have generated the observation z_{k+1} . As will be seen next, in both terms (a) and (b), the same weights are obtained; however, only in the latter case are the weights to be normalized such that $\sum_i \tilde{w}_i = 1$.

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

$$\begin{aligned} \sum_i \mathbb{P}(X_{k+1}, A_i | \mathcal{H}_{k+1}^-, z_{k+1}) = \\ \sum_i \mathbb{P}(X_{k+1} | \mathcal{H}_{k+1}^-, z_{k+1}, A_i) \cdot \mathbb{P}(A_i | \mathcal{H}_{k+1}^-, z_{k+1}) \quad (17) \end{aligned}$$

Here, the first term is the posterior belief, conditioned on observations and history, as well as a candidate scene A_i that supposedly generated the observation z_{k+1} . We discuss how this term can be calculated in Section 4.5.

The second term, $\mathbb{P}(A_i | \mathcal{H}_k, u_k, z_{k+1})$, is merely the likelihood of A_i being actually the one that generated the observation z_{k+1} . As will be seen now, this term is actually the

normalized weight w_i from Section 4.1. Marginalizing over all robot poses and applying Bayes rule yields

$$\mathbb{P}(A_i | \mathcal{H}_{k+1}^-, z_{k+1}) = \int_x \mathbb{P}(A_i, x | \mathcal{H}_{k+1}^-, z_{k+1}) \quad (18)$$

$$= \eta \int_x \mathbb{P}(z_{k+1} | A_i, x, \mathcal{H}_{k+1}^-) \mathbb{P}(A_i, x | \mathcal{H}_{k+1}^-) \quad (19)$$

with a normalization constant $\eta \doteq \mathbb{P}(z_{k+1} | \mathcal{H}_{k+1}^-)$.

Like the derivation in Section 4.1, since this integral could be over any arbitrary total distribution of x , we can use the propagated belief $b[x_{k+1}^-]$, to compute it as

$$\begin{aligned} \mathbb{P}(A_i | z_{k+1}, \mathcal{H}_{k+1}^-) &= \\ \eta \int_x \mathbb{P}(z_{k+1} | A_i, x, \mathcal{H}_{k+1}^-) \mathbb{P}(A_i | x, \mathcal{H}_{k+1}^-) b[x_{k+1}^- = x] & \quad (20) \end{aligned}$$

As seen, the same expression is obtained as in equation (16), except for the normalization constant η . Hence

$$\mathbb{P}(A_i | z_{k+1}, \mathcal{H}_{k+1}^-) = \eta w_i \doteq \tilde{w}_i \quad (21)$$

In practice, one can avoid calculation of η , and instead normalize the weights w_i such that $\sum_i \tilde{w}_i = 1$.

4.3. Summary thus far

To summarize the discussion thus far, we have shown that the objective function (equation (11)) can be rewritten as

$$J(u_k) = \int_{z_{k+1}} \left(\sum_i w_i \right) \cdot c \left(\sum_i \tilde{w}_i b[X_{k+1}^{i+}] \right) \quad (22)$$

with weights w_i and \tilde{w}_i defined in equations (16) and (21), and the posterior given scene A_i , defined as

$$b[X_{k+1}^{i+}] \doteq \mathbb{P}(X_{k+1} | \mathcal{H}_{k+1}^-, z_{k+1}, A_i) \quad (23)$$

Observe that, for each considered observation z_{k+1} , we get a *mixture probability distribution function* inside of the cost $c(\cdot)$, where each component represents the posterior probability distribution conditioned on the observation capturing scene A_i , and weighted by \tilde{w}_i . If there is no perceptual aliasing, there will be only one component with large weight \tilde{w}_i , that corresponds to the correct data association with scene A_i , with all other weights being negligible. Conversely, in the 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 normalized weights \tilde{w}_i for each $A_i \in \{A_{\mathbb{N}}\}$.

These insights also apply to the unnormalized weights w_i that appear outside of the cost, from which the likelihood of obtaining observation z_{k+1} is calculated. However, as already discussed in Section 4.1, this likelihood is calculated by summing over all such weights ($\sum_i w_i$), with each weight properly capturing the likelihood that a measurement z_{k+1} will be generated by scene A_i while taking into

account the probability of the corresponding robot pose x , given the propagated belief, i.e. $b[x_{k+1}^- = x]$.

For practical purposes, one can thus only consider viewpoints with non-negligible probability according to $b[x_{k+1}^-]$. Moreover, it is possible to threshold the weights in the mixture $\sum_i \tilde{w}_i b[X_{k+1}^{i+}]$, instead of always considering all scenes $\{A_{\mathbb{N}}\}$.

Having shown that incorporating data association within belief space planning leads to equation (22), we now proceed with the exposition of our approach.

4.4. Simulating future observations $\{z_{k+1}\}$ given $b[X_{k+1}^-]$

Calculating the objective function (equation (22)) for each candidate action u_k involves considering all possible realizations of z_{k+1} . One approach to perform this, in practice, is to simulate future observations $\{z_{k+1}\}$ given the propagated belief $b[X_{k+1}^-]$, scenes $\{A_{\mathbb{N}}\}$, and observation model (equation (2)). One can then evaluate equation (22) considering all observations in the set $\{z_{k+1}\}$.

We now briefly describe how this concept can be realized. First, viewpoints $\{x\}$ are sampled from $b[X_{k+1}^-]$. For each viewpoint $x \in \{x\}$, an observed scene A_i is determined according to event likelihood $\mathbb{P}(A_i | \mathcal{H}_k, x)$. Together, x and A_i are used to generate nominal $\hat{z} = h(x, A_i)$ and noise-corrupted observations $\{z\}$ according to the observation model (equation (2)): $z = h(x, A_i) + v$. The set $\{z_{k+1}\}$ is then the union of all such generated observations $\{z\}$. Note that while generating $\{z_{k+1}\}$, the true association is known (scene A_i) but it is unknown to our algorithm, i.e. while evaluating equation (22).

4.5. Computing a mixture of posterior beliefs $\sum_i \tilde{w}_i b[X_{k+1}^{i+}]$

As seen from equation (22), reasoning about data association aspects resulted in a mixture of posterior probability distributions within the cost $c(\cdot)$, i.e. $\sum_i \tilde{w}_i b[X_{k+1}^{i+}]$, for each possible observation $z_{k+1} \in \{z_{k+1}\}$. In this work, the set $\{z_{k+1}\}$ is simulated as discussed in Section 4.4; however, one could also consider treating future observation z_{k+1} as a random variable (Indelman et al., 2015; Platt et al., 2011; Van den Berg et al., 2012).

In this section, we briefly describe how one can actually calculate the corresponding posterior distributions, given some specific observation $z_{k+1} \in \{z_{k+1}\}$. For simplicity, we consider the belief at planning time k as a Gaussian $b[X_k] = \mathcal{N}(\hat{X}_k, \Sigma_k)$. However, our approach could also be applied to more general cases (e.g. a mixture of Gaussians) with a certain price in terms of computational complexity. Further investigation of these aspects is left to future research.

Under this setting, each of the components $b[X_{k+1}^{i+}]$ in the mixture probability distribution function can be written as

$$b[X_{k+1}^{i+}] \propto b[X_k] \mathbb{P}(x_{k+1} | x_k, u_k) \mathbb{P}(z_{k+1} | x_{k+1}, A_i) \quad (24)$$

It is then not difficult to show that this belief is a Gaussian and to find its first two moments via maximum a-posteriori inference

$$b[X_{k+1}^{i+}] = \mathcal{N}(\hat{X}_{k+1}^i, \Sigma_{k+1}^i) \quad (25)$$

with mean \hat{X}_{k+1}^i and covariance Σ_{k+1}^i .

Obviously, the mixture of posterior beliefs in the cost $c(\cdot)$ from equation (22) is now a mixture of Gaussians:

$$\sum_i \tilde{w}_i b[X_{k+1}^{i+}] = \sum_i \tilde{w}_i \mathcal{N}(\hat{X}_{k+1}^i, \Sigma_{k+1}^i) \quad (26)$$

At this point, it is useful to recall that the corresponding mean and covariance of a general Gaussian mixture probability distribution function

$$p(x) = \sum_{j=1}^n w_j \mathcal{N}(\hat{x}_j, \Sigma_j)$$

are given by (Bar-Shalom et al., 2004)

$$\hat{x} = \sum_{j=1}^n w_j \hat{x}_j, \quad \Sigma = \sum_{j=1}^n w_j \Sigma_j + \tilde{\Sigma} \quad (27)$$

with $\tilde{\Sigma} = \sum_{j=1}^n w_j (\hat{x}_j - \hat{x})(\hat{x}_j - \hat{x})^T$.

Thus, in cases where a combination of the conditional posterior probability distributions (equation (25)) makes sense, the overall combined belief can be collapsed to a Gaussian distribution according to equation (27). As will be seen later, this could indeed be required for certain cost functions $c(\cdot)$.

4.6. Under uncertain scenes

So far we have assumed that scenes were known fairly well i.e., A_i was known with little uncertainty. This might, however, not be true in practice, for example when such scenes arise out of a previous SLAM session. In such a scenario, it is more appropriate to consider the uncertainty of A_i , again within the DA-BSP framework. This implies that equation (20) should incorporate the actual model of variation of A_i . More precisely

$$\mathbb{P}(A_i | x, \mathcal{H}_{k+1}^-) = \mathcal{N}(\hat{x}_{A_i}, \Sigma_{A_i}) \quad (28)$$

where \hat{x}_{A_i} is the nominal position around which the scene A_i is distributed while Σ_{A_i} is the covariance. Hence we have

$$\begin{aligned} \mathbb{P}(A_i | z_{k+1}, \mathcal{H}_{k+1}^-) \\ = \eta \int_x \mathbb{P}(z_{k+1} | A_i, x, \mathcal{H}_{k+1}^-) \mathcal{N}(\hat{x}_{A_i}, \Sigma_{A_i}) b[x_{k+1}^- = x] \end{aligned} \quad (29)$$

Now consider that the prior (and hence also propagated) belief is a Gaussian. We have already seen that the observation likelihood (i.e., $\mathbb{P}(z_{k+1} | A_i, x, \mathcal{H}_{k+1}^-)$) is a Gaussian too. Note that by uncertain A_i , we imply that the position from

whence A_i can be viewed is uncertain, not the class of the scene itself. Therefore, we can still harness the observation model to compute the term $\mathbb{P}(z_{k+1} | A_i, x, \mathcal{H}_{k+1}^-)$. Furthermore, under the maximum likelihood assumption, we know that this observation is distributed normally around the most likely observation $\hat{z}_{k+1}^{\text{ml}}$. Therefore the right-hand side of the equation can be reduced to a single Gaussian. In other words, the weights w_i used in equations (16) and (20) can be computed efficiently through analytical approach.

Similarly, if the propagated belief $b[x_{k+1}^- = x]$ is a mixture of Gaussians, we can see that the expression in the left-hand side of equation (29) yields another mixture of Gaussians, with the same number of modes. Here too, computation could be performed analytically, by iterating over each Gaussian component. Obviously, the ease of computation of w_i is afforded not by the fact that A_i is uncertain, but by the fact that its uncertainty is of a specific structure (namely, a Gaussian distribution).

4.7. Designing a specific cost function

The treatment so far has been noncommittal to the structure of the cost function $c(\cdot)$. Recalling equation (22), we see that the belief over which the cost function is defined is multimodal in general. Standard cost functions in the literature typically include such terms as control usage c_u , distance to goal c_G and uncertainty c_Σ , see e.g. Indelman et al. (2015) or Van den Berg et al. (2012). In our case, however, the specific form of the latter should be re-examined and an additional term quantifying ambiguity level can be introduced. In this section, we thus briefly discuss these two terms, starting with the cost over posterior uncertainty.

Since, unlike usual belief space planning, the posterior belief in our case is multimodal and represented as a mixture of Gaussians $\sum_i \tilde{w}_i \mathcal{N}(\hat{X}_{k+1}^i, \Sigma_{k+1}^i)$ (see equation (26)), we could define several different cost structures, depending on how we treat the different modes. Two particular examples involve taking the worst-case covariance among all covariances Σ_{k+1}^i in the mixture, e.g. $\Sigma = \max_i \{\text{tr}(\Sigma_i)\}$, or collapsing the mixture into a single Gaussian $\mathcal{N}(\cdot, \Sigma)$, according to equation (27). In both cases, we can define the cost due to uncertainty as $c_\Sigma = \text{trace}(\hat{\Sigma})$.

The cost due to ambiguity, c_w , should penalize such ambiguities as those arising out of perceptual aliasing. Here, we note that the non-negligible weights w_i in equation (22) arise as a result of perceptual aliasing with respect to any scene A_i , whereas in the case of no aliasing, all but one of these weights are zero. In the most severe case of aliasing (all scenes or objects A_i are identical), all of these weights are comparable with each other. Thus, we take the Kullback–Leibler divergence $KL_u(\{\tilde{w}_i\})$ of these weights $\{\tilde{w}_i\}$ from a uniform distribution to penalize higher aliasing, and define

$$c_w(\{\tilde{w}_i\}) \doteq \frac{1}{KL_u(\{\tilde{w}_i\}) + \epsilon}$$

where ϵ is a small number, included to avoid division by zero in the case of extreme perceptual aliasing.

The overall cost can then be defined as a combination

$$c \doteq M_u c_u + M_G c_G + M_\Sigma c_\Sigma + M_w c_w \quad (30)$$

with user-defined weights M_u, M_G, M_Σ , and M_w .

4.8. Formal algorithm for DA-BSP

We now have all the ingredients to present the overall framework of data association aware belief space planning, DA-BSP for brevity. It is summarized in Algorithm 1 and briefly described next.

Given belief $b[X_k]$ and candidate action u_k , we first propagate the belief to get $b[X_{k+1}^-]$ and then simulate future observations $\{z_{k+1}\}$ (line 2), as described in Section 4.4. The algorithm then calculates the contribution of each observation $z_{k+1} \in \{z_{k+1}\}$ to the objective function (22). In particular, in lines 8 and 11 we calculate the weights w_i and the posterior beliefs $b[X_{k+1}^{i+}]$ for each $A_i \in \{A_{\mathbb{N}}\}$, respectively, according to Sections 4.1 and 4.5. Then, after weight normalization in line 13, we evaluate the cost $c(\cdot)$ (line 14) and use the accumulated unnormalized weights $w_s \equiv \sum_i w_i$ to update the value of the objective function J with the weighted cost for measurement z_{k+1} (line 15).

Later on, in Section 5.7, we comment more precisely on both the computational complexity of DA-BSP and its correctness when viewed from the perspective of data association.

4.9. An abstract example for DA-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

$$f(x, u) = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \cdot x + d \begin{cases} [0, 1]^T & \text{if } u = \textit{up} \\ [1, 0]^T & \text{if } u = \textit{right} \end{cases},$$

$$h(x, A_i) = h_i(x) = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \cdot (x - x_i) + s_i \quad (31)$$

where the observations, as well as the metric shift s_i , are in an object-centric frame, with x_i representing the location of A_i . Intuitively, s_i is a simple mechanism to model perceptual aliasing between objects; e.g., for identical objects, each A_i would have s_i that compensates for x_i in such a way that it results in the same observation h . The first row (panels a to d) of Figure 4 illustrates the process of simulating future observations $\{z_{k+1}\}$ for $u_k = \textit{up}$, considering unique and perceptually aliased scenes (Figure 4(c) and (d)). In particular, a sampled pose x^{tr} used to generate an observation $z_{k+1} \in \{z_{k+1}\}$ is shown in 4(b).

Figure 4 demonstrates key aspects in our approach, considering a single observation z_{k+1} each time. Our approach reasons about data association; hence, we consider that each z_{k+1} could have been generated by one of the three objects; each such association would fetch us a conditional posterior belief $b[X_{k+1}^{i+}]$, as denoted by small ellipses. Finally, we compute the total cost according to Algorithm 2.

Figure 4(e) to (h) denotes the situation when the true pose x^{tr} is close to the center and observes A_2 , while in Figure 4(i) to (l) the true pose is at the left side and observes 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 unnormalized weight w_i is greater when the object is at the center, because the overall likelihood of the observation is greater. Also, with no aliasing, for any scene A_j other than the true one, the normalized 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 the case of no aliasing, this can be negligibly small. This is crucial, since it implies that DA-BSP, in practical applications with infrequent aliasing, would not require any significant additional computational effort with respect to usual belief space planning.

Figure 4(f) to (h) depicts $\{A_1, A_2\}_{\text{alias}}$, $\{A_1, A_3\}_{\text{alias}}$, and $\{A_1, A_2, A_3\}_{\text{alias}}$. When $\{A_1, A_3\}_{\text{alias}}$, the weights w_i are similar, and indeed our cost c_w of weights (in equation (30)) is high. For similar uncertainty in pose, this cost would remain constant. Hence, in the presence of identical objects placed similarly within the current belief, optimization of general cost function would be guided toward *active localization*. Conversely, if one object j lies closer to the current nominal pose, it will have slightly higher w_j . If $\{A_1, A_2, A_3\}_{\text{alias}}$, i.e. *all* objects are identical, the weights w_i are simply an indication of the prior distribution. This is reasonable since, in such a case, considering different data association does not yield any new information.

5. Non-myopic data association aware multi-modal belief space planning

This section would generalize the DA-BSP that was developed in Section 4. We will start by considering a prior distribution that is non-Gaussian. In particular, we will assume our prior distribution to be a mixture of Gaussians and then follow a similar approach to compute the belief update and perform myopic planning, as done earlier. Once this is done and we have an approach that takes in a Gaussian mixture model belief and updates to another Gaussian mixture model belief, we will present the most general DA-BSP in a non-myopic setting of several look-ahead steps of planning.

5.1. Prior belief as a mixture of Gaussians

Let us assume that the prior belief is a Gaussian mixture model. In other words, our belief at time k is a linear

Algorithm 1 Myopic data association aware belief space planning.

Input: Current belief $b[X_k]$ at step k , history \mathcal{H}_k , action u_k , scenes $\{A_N\}$, event likelihood $\mathbb{P}(A_i | \mathcal{H}_k, x)$ for each $A_i \in \{A_N\}$

```

1:  $b[X_{k+1}^-] \leftarrow b[X_k] \mathbb{P}(x_{k+1} | x_k, u_k)$ 
2:  $\{z_{k+1}\} \leftarrow \text{SimulateObservations}(b[X_{k+1}^-], \{A_N\})$ 
3:  $J \leftarrow 0$ 
4: for  $\forall z_{k+1} \in \{z_{k+1}\}$  do
5:    $w_s \leftarrow 0$ 
6:   for  $i \in [1 \dots |A|]$  do
7:      $\triangleright$  compute weight, equation (16)
8:      $w_i \leftarrow \text{CalcWeights}(z_{k+1}, \mathbb{P}(A_i | \mathcal{H}_{k+1}^-, x), b[X_{k+1}^-])$ 
9:      $w_s \leftarrow w_s + w_i$ 
10:     $\triangleright$  Calculate posterior belief given  $A_i$ , Section 4.5
11:     $b[X_{k+1}^{i+}] \leftarrow \text{UpdateBelief}(b[X_{k+1}^-], z_{k+1}, A_i)$ 
12:  end for
13:   $\{\tilde{w}_i\} \leftarrow \text{NormalizeWeights}(\{w_i\})$ 
14:   $c \leftarrow \text{CalcCost}(\{\tilde{w}_i\}, \{b[X_{k+1}^{i+}]\})$ 
15:   $J \leftarrow J + w_s \cdot c$ 
16: end for
17: return  $J$ 

```

\triangleright Section 4.7

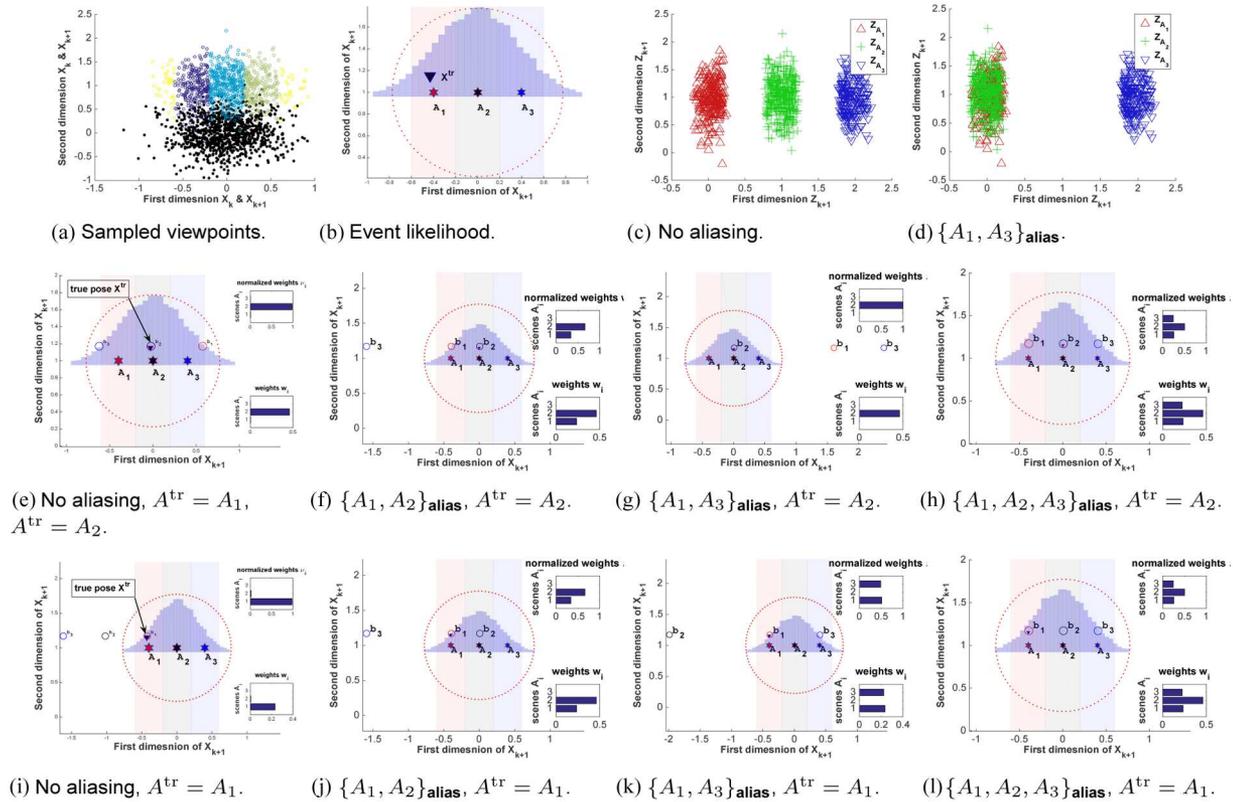


Fig. 4. (a) Black samples $\{x_k\}$ are drawn from $b[X_k] \doteq \mathcal{N}([0, 0]^T, \Sigma_k)$, from which, given control u_k , samples $\{x_{k+1}\}$ are computed, colored according to different scenes A_i being observed, and used to generate observations $\{z_{k+1}\}$. (b) Event likelihood, $\mathbb{P}(A_i | x, \mathcal{H}) \forall i$. Three vertical stripes represent locations from which each scene A_i is observable; histogram represents distribution of $\{x_{k+1}\}$, which corresponds to $b[X_{k+1}^-]$. (c,d) distributions of $\{z_{k+1}\}$ without aliasing and when $\{A_1, A_3\}$ alias. (e)–(h) 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 an inverted triangle. Smaller ellipses are the posterior beliefs $b[X_{k+1}^{i+}]$. 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.

Algorithm 2 Non-myopic data association aware belief space planning.

Input: Current Gaussian mixture model belief $b[X_k]$ at step k , history \mathcal{H}_k , action u_k , scenes $\{A_{\mathbb{N}}\}$, event likelihood $\mathbb{P}(A_i | \mathcal{H}_k, x)$ for each $A_i \in \{A_{\mathbb{N}}\}$

- 1: $b[X_{k+1}^-] \leftarrow b[X_k] \mathbb{P}(x_{k+1} | x_k, u_k)$ ▷ equation (8)
- 2: $\{z_{k+1}\} \leftarrow \text{SimulateObservations}(b[X_{k+1}^-], \{A_{\mathbb{N}}\})$
- 3: $J \leftarrow 0$
- 4: **for** $\forall z_{k+1} \in \{z_{k+1}\}$ **do**
- 5: $w_s \leftarrow 0$
- 6: **for** $i \in [1 \dots |A|]$ **do**
- 7: ▷ compute weight, equation (16)
- 8: $w_{k+1}^i \leftarrow \text{CalcWeights}(z_{k+1}, \mathbb{P}(A_i | \mathcal{H}_{k+1}^-, x), b[X_{k+1}^-])$
- 9: $w_s \leftarrow w_s + w_i$
- 10: **for** $\forall j \in [1, \dots, M_k]$ **do**
- 11: ▷ compute weight \tilde{w}_{k+1}^{ij} for each Gaussian mixture model component, equation (37)
- 12: $\tilde{w}_{k+1}^{ij} \leftarrow \text{CalcWeights}(z_{k+1}, \mathbb{P}(A_i | \mathcal{H}_{k+1}^-, x), b[X_{k+1}^{j-}])$
- 13: $\xi_{k+1}^{ij} \leftarrow \xi_k^j \tilde{w}_{k+1}^{ij}$ ▷ equation (39)
- 14: ▷ Calculate posterior of $b[X_{k+1}^{j-}]$, given A_i , Section 5.2
- 15: $b[X_{k+1}^{ij+}] \leftarrow \text{UpdateBelief}(b[X_{k+1}^{j-}], z_{k+1}, A_i)$
- 16: **end for**
- 17: **end for**
- 18: Prune components with weights ξ_{k+1}^{ij} below a threshold
- 19: Construct $b[X_{k+1}^+]$ from the remaining M_{k+1} components via equation (38)
- 20: $c \leftarrow \text{CalcCost}(b[X_{k+1}^+])$ ▷ Section 4.7
- 21: $J \leftarrow J + w_s \cdot c$
- 22: **end for**
- 23: **return** J

combination of $M_k \in \mathbb{N}$ Gaussians, i.e.,

$$b[X_k] \doteq \mathbb{P}(X_k | \mathcal{H}_k^-, z_k) = \sum_{i=1}^{M_k} \xi_{k,i} \mathcal{N}(\hat{X}_{k,i}, \Sigma_{k,i}) \quad (32)$$

Since our motion model (see equation (3)) is still a Gaussian, the propagated belief is also a Gaussian mixture model with M_k components. More precisely

$$b[X_{k+1}^-] \doteq \mathbb{P}(X_{k+1} | \mathcal{H}_{k+1}^-) = \mathbb{P}(X_k | \mathcal{H}_k) \mathbb{P}(x_{k+1} | x_k, u_k) \\ = \sum_{i=1}^{M_k} \xi_{k,i} \mathcal{N}(\hat{X}_{k+1,i}^-, \Sigma_{k,i}^-) \quad (33)$$

Once the observation z_{k+1} is obtained, for each of the M_k components, we can consider all the aliased scenes $\{A_{\mathbb{N}}\}$. The derivation is very similar to the lines of the discussion in Section 4, with additional parameters introduced. For ease of disposition, let us reproduce the steps, such as equation (15), which we get after applying the chain rule and subsequent marginalization over all x and $A_i \in \{A_{\mathbb{N}}\}$

$$\sum_i^{\{A_{\mathbb{N}}\}} \int_x \mathbb{P}(z_{k+1} | x, A_i, \mathcal{H}_{k+1}^-) \mathbb{P}(A_i | \mathcal{H}_{k+1}^-, x) b[x_{k+1}^- = x]$$

Thus

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

Since the propagated belief (see equation (33)), from which $b[x_{k+1}^-]$ is calculated, is also a Gaussian mixture model, we can replace $b[x_{k+1}^- = x]$ with

$$b[x_{k+1}^- = x] = \sum_{j=1}^{M_k} \xi_{k+1,j} b[x_{k+1,j}^- = x] \quad (34)$$

However, the actual value of w_i (for each $A_i \in \{A_{\mathbb{N}}\}$) depends, in addition to the measurement likelihood and event likelihood, on the Gaussian mixture model belief $b[x_{k+1}^-]$, 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 $b[x_{k+1}^- = x^{\text{tr}}]$ will be unlikely to be actually acquired, leading to low values of w_i with $A_i = A^{\text{tr}}$. Since $b[x_{k+1}^-]$ is a Gaussian mixture model with M_k components, the low-probability pose x^{tr} corresponds to low probabilities $b[x_{k+1}^{j-} = x^{\text{tr}}]$ for each component $j \in \{1, \dots, M_k\}$. However, the likelihood term (equation (13)) could still increase in the case of perceptual aliasing, where the aliased scene A_j generates a similar observation to z_{k+1} from viewpoint x' with the latter being more probable, i.e. high probability $b[x_{k+1}^- = x']$.

In practice, the integral in equation (16) can be calculated efficiently by considering each component of the Gaussian mixture model $b[x_{k+1}^-]$ separately. Each such component is

a Gaussian that is multiplied by the measurement likelihood $\mathbb{P}(z_{k+1} | A_i, x, \mathcal{H})$, which is also a Gaussian; it is known that a product of Gaussians remains a Gaussian. The integral can then only be calculated for the window where event likelihood is non-zero, i.e. $\mathbb{P}(A_i | x, \mathcal{H}) > 0$. For general probability distributions, the integral in equation (16) should be computed numerically. Since, in practical applications, $\mathbb{P}(A_i | x, \mathcal{H})$ is sparse with respect to x , this computational cost is not severe.

Similarly for term (b), $\mathbb{P}(X_{k+1} | \mathcal{H}_{k+1}^-, z_{k+1})$, applying total probability over non-overlapping $\{A_{\mathbb{N}}\}$ as well as all the components of the propagated belief, we get

$$\begin{aligned} & \mathbb{P}(X_{k+1} | \mathcal{H}_{k+1}^-, z_{k+1}) \\ &= \sum_{j=1}^{M_k} \sum_{i=1}^{|A_{\mathbb{N}}|} \mathbb{P}(X_{k+1}, A_i, \gamma = j | \mathcal{H}_{k+1}^-, z_{k+1}) \end{aligned} \quad (35)$$

Proceeding as before, we split the term inside the summation using the chain rule as follows

$$\begin{aligned} & \mathbb{P}(X_{k+1}, A_i, \gamma = j | \mathcal{H}_{k+1}^-, z_{k+1}) \\ &= \mathbb{P}(X_{k+1} | \mathcal{H}_{k+1}^-, z_{k+1}, A_i, \gamma = j) \\ & \quad \cdot \mathbb{P}(A_i, \gamma = j | \mathcal{H}_{k+1}^-, z_{k+1}) \end{aligned}$$

The first term is the posterior obtained with the scene A_i while considering the j th propagated belief component; we denote this $b[X_{k+1}^{j+} | A_i]$.

For the second term, we again apply the chain rule, to obtain

$$\begin{aligned} & \mathbb{P}(A_i, \gamma = j | \mathcal{H}_{k+1}^-, z_{k+1}) \\ &= \mathbb{P}(A_i | \gamma = j, \mathcal{H}_{k+1}^-, z_{k+1}) \cdot \mathbb{P}(\gamma = j | \mathcal{H}_{k+1}^-, z_{k+1}) \end{aligned}$$

Here, $\mathbb{P}(\gamma = j | \mathcal{H}_{k+1}^-, z_{k+1})$ is equal to ξ_k^j , which is the weight of the j th component of the prior belief. For the first term, we marginalize over all x to obtain the weights \tilde{w}_{k+1}^{ij} . This is identical to marginalization done in the previous Section 4 (see equation (20)), with the only difference that here all x considered are from the j th component of the belief

$$\begin{aligned} & b[X_{k+1}] \\ &= \sum_{j=1}^{M_k} \sum_{i=1}^{|A_{\mathbb{N}}|} \xi_k^j \mathbb{P}(A_i | \mathcal{H}_{k+1}^-, z_{k+1}, \gamma = j) b[X_{k+1}^{j+} | A_i] \end{aligned} \quad (36)$$

$$\begin{aligned} & \tilde{w}_{k+1}^{ij} \\ & \doteq \eta' \int_x \mathbb{P}(z_{k+1} | A_i, x, \mathcal{H}_{k+1}^-) \mathbb{P}(A_i | \mathcal{H}_{k+1}^-, \gamma = j, x) b[x_{k+1}^{j-} = x] \end{aligned} \quad (37)$$

with $\eta' = 1/\mathbb{P}(z_{k+1} | \mathcal{H}_{k+1}^-)$. Note that for each component j , $\sum_i \tilde{w}_{k+1}^{ij} = 1$. Finally, we can rewrite equation (36) as

$$\mathbb{P}(X_{k+1} | \mathcal{H}_{k+1}^-, z_{k+1}) = \sum_{r=1}^{M_{k+1}} \xi_{k+1}^r \mathbb{P}(X_{k+1} | \mathcal{H}_{k+1}, \gamma = r) \quad (38)$$

or, in short

$$b[X_{k+1}] = \sum_{r=1}^{M_{k+1}} \xi_{k+1}^r b[X_{k+1}^{r+}]$$

where

$$\xi_{k+1}^r \doteq \xi_{k+1}^{ij} \equiv \xi_k^i \tilde{w}_{k+1}^{ij}, \quad b[X_{k+1}^{r+}] \doteq b[X_{k+1}^{j+} | A_i] \quad (39)$$

As seen, we get a new Gaussian mixture model with M_{k+1} components, where each component $r \in [1, M_{k+1}]$, with appropriate mapping to indices (i, j) from equation (36), is represented by weight ξ_{k+1}^r and posterior conditional belief $b[X_{k+1}^{r+}]$. The latter can be evaluated as the Gaussian

$$\begin{aligned} & b[X_{k+1}^{r+}] \propto b[X_{k+1}^{j-}] \mathbb{P}(z_{k+1} | x_{k+1}, A_i) \\ &= \mathcal{N}(\hat{X}_{k+1}^r, \Sigma_{k+1}^r) \end{aligned} \quad (40)$$

where the mean \hat{X}_{k+1}^r and covariance Σ_{k+1}^r can be efficiently recovered via maximum a-posteriori inference.

5.2. Non-myopic DA-BSP

It is easy to see that once the prior, as well as the posterior, belief is represented as a mixture of Gaussians, we can extend the DA-BSP to a non-myopic setting. Informally, for planning over a horizon of L steps, starting with a multimodal prior and a control sequence $u_{0:L-1}$, the planning would involve reasoning about the plausible data associations at each intermediate $l \in [1, L-1]$ step. To make it more concrete, consider a non-myopic cost function as

$$\begin{aligned} & J(u_{k:k+L-1}) \\ &= \int_{z_{k+1:k+L}} \sum_{l=1}^L \overbrace{\mathbb{P}(z_{k+l} | \mathcal{H}_{k+l}^-)}^{(a)} c_l \left(\overbrace{\mathbb{P}(X_{k+l} | \mathcal{H}_{k+l}^-, z_{k+l})}^{(b)} \right) \end{aligned} \quad (41)$$

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 l th look-ahead step.

Like the myopic case in Section 4, the two terms (a) and (b) in equation (41) have intuitive meaning: for each considered value of z_{k+l} , term (a) represents how likely is it to get such an observation, while term (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, this reasoning is valid for *all* possible realizations of $z_{k+1:k+l-1}$ and the corresponding posterior beliefs $\mathbb{P}(X_{k+l-1} | \mathcal{H}_{k+l-1})$.

It is not difficult to show that the posterior belief at each step k is actually the Gaussian mixture model

$$\mathbb{P}(X_{k+l} | \mathcal{H}_{k+l}^-, z_{k+l}, A_i) = \sum_{j=1}^{M_{k+l-1}} \xi_{k+l-1}^j b[X_{k+l}^{j+} | A_i] \quad (42)$$

where $b[X_{k+l}^{j+}|A_i] \doteq \mathbb{P}(X_{k+l}|\mathcal{H}_{k+l}^-, \gamma = j, A_i, z_{k+l})$ is the posterior belief of the j th Gaussian mixture model component of the propagated belief $b[X_{k+l}^-]$.

Plugging equation (42) into $\mathbb{P}(X_{k+l}|\mathcal{H}_{k+l}^-, z_{k+l}) \equiv b[X_{k+l}]$ from equation (9) yields

$$b[X_{k+l}] = \sum_{i=1}^{|\mathcal{A}_{\mathbb{N}}|} \sum_{j=1}^{M_{k+l-1}} \xi_{k+l-1}^j \mathbb{P}(A_i | \mathcal{H}_{k+l}^-, z_{k+l}) b[X_{k+l}^{j+}|A_i] \quad (43)$$

Accounting for $b[x_{k+l}^{j-}]$ for each considered j th component as $\mathbb{P}(A_i | \mathcal{H}_{k+l}^-, z_{k+l}) = \int_x \mathbb{P}(A_i, x | \mathcal{H}_{k+l}^-, z_{k+l})$, and applying Bayes' rule yields

$$\tilde{w}_{k+l}^{ij} \doteq \eta' \int_x \mathbb{P}(z_{k+l}|A_i, x, \mathcal{H}_{k+l}^-) \mathbb{P}(A_i | \mathcal{H}_{k+l}^-, x) b[x_{k+l}^{j-}=x] \quad (44)$$

with $\eta' = 1/\mathbb{P}(z_{k+l} | \mathcal{H}_{k+l}^-)$. Note that for each component j , $\sum_i \tilde{w}_{k+l}^{ij} = 1$. Finally, we can rewrite equation (36) as

$$b[X_{k+l}] = \sum_{r=1}^{M_{k+l}} \xi_{k+l}^r \mathbb{P}(X_{k+l}|\mathcal{H}_{k+l}^-, \gamma = r) = \sum_{r=1}^{M_{k+l}} \xi_{k+l}^r b[X_{k+l}^{r+}] \quad (45)$$

where $\xi_{k+l}^r \doteq \xi_{k+l}^{ij} \equiv \xi_{k+l-1}^j \tilde{w}_{k+l}^{ij}$ and $b[X_{k+l}^{r+}] \doteq \mathbb{P}(X_{k+l}|\mathcal{H}_{k+l}^-, \gamma = r)$. As seen, we get a new Gaussian mixture model with M_{k+l} components, where each component $r \in [1, M_{k+l}]$, with appropriate mapping to indices (i, j) from equation (36), is represented by weight ξ_{k+l}^r and posterior conditional belief $b[X_{k+l}^{r+}]$. The latter can be evaluated as the Gaussian $b[X_{k+l}^{r+}] = \mathcal{N}(\hat{X}_{k+l}^r, \Sigma_{k+l}^r)$, with mean \hat{X}_{k+l}^r and covariance Σ_{k+l}^r .

The associated cost of the overall posterior distribution of this L -step planning can then be compared with that of similar posterior distributions of other control sequences, enabling us to choose an optimal single step action. After the action is taken and a real observation is obtained, the inference over this observation allows us to update the posterior distribution, which then serves as a prior distribution for next- L -step planning. However, a naïve implementation of such a planning would likely suffer from the usual curses of *dimensionality* and *history*. Luckily, DA-BSP provides a principled way to strike a balance between the requirement for an efficient solution and not losing the correct data association in a challenging aliased environment.

5.3. Overall algorithm

We now have all the ingredients to present the overall framework of data association aware belief space planning, calling it DA-BSP for brevity. It is summarized in Algorithm 2 and briefly described next.

We shall, in general, refer to Section 5.2 now. Given a Gaussian mixture model belief $b[X_k]$ and candidate action

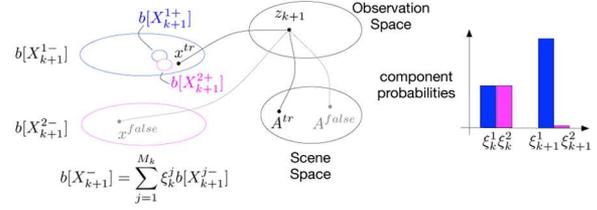


Fig. 5. Gaussian mixture model posterior $b[X_{k+1}]$ given $z_{k+1} \in \{z_{k+1}\}$. The prior has two equiprobable components while the posterior has different weights for the two components.

u_k , we first propagate the belief to get $b[X_{k+1}^-]$ and then simulate future observations $\{z_{k+1}\}$ (line 2), as described in Section 4.4. The algorithm then calculates the contribution of each observation $z_{k+1} \in \{z_{k+1}\}$ to the objective function (equation (22)). In particular, in lines 8 and 9, we calculate the weights w_{k+1}^i that are used in evaluating the likelihood w_s of obtaining observation z_{k+1} . In lines 10 to 16, we compute the posterior belief: this involves updating each j th component from the propagated belief $b[X_{k+1}^{j-}]$ with observation z_{k+1} , considering each of the possible scenes A_i . After pruning (line 18), this yields a posterior Gaussian mixture model with M_{k+1} components. We then evaluate the cost $c(\cdot)$ (line 20) and use w_s to update the value of the objective function J with the weighted cost for measurement z_{k+1} (line 21).

One can observe that, according to equation (43), each of the M_k components from the belief at a previous time is split into $|\mathcal{A}_{\mathbb{N}}|$ new components with appropriate weights. This would imply an explosion in the number of components, making the proposed framework hardly applicable. However, in practice, the majority of the weights will be negligible, and can therefore be pruned, while the remaining number of components is denoted by M_{k+1} in equation (38). Depending on the scenario and the degree of perceptual aliasing, this can correspond to *full* or *partial* disambiguation (see Figure 5).

5.4. Effect of reducing a mixture of belief

We have seen that DA-BSP, on account of considering all the possible data association, suffers from exponential blow up in a number of components. Using the discrete case as an example, it is easy to show that this—under a reasonable assumption that scene space is much smaller than state space—does not deteriorate the complexity of the underlying problem. Moreover, it is important to notice that each such association is accompanied with the weights, which reflect the significance of such a data association. In particular, if a scene is quite unique, it is unlikely to be aliased with any other; consequently, only the posterior conditioned on this correct association would have significant weight. A simple threshold-based pruning is then sufficient to discard insignificant modes, as shall also be evident from our extensive experiment in this regard later on.

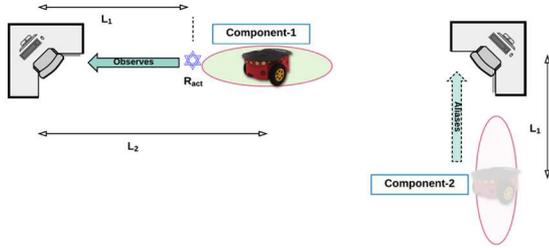


Fig. 6. When pose uncertainty is not taken into account, overly confident and wrong inference can be made. Here, in the absence of uncertainty and because $L_2 > L_1$, an observation from the ground truth (shown \star) is likely to be solely attributed to an incorrect component 2.

One can notice that the objective of curtailing the complexity of data structure through some pruning heuristics is not a novel approach. Even in the context of data association, it occurs in slightly different form when the problem is posed as a multihypothesis tracking. Roughly speaking, in such a scenario, planning is through explicit instantiation of the trajectory of control and pruning decision is often based on the information-theoretic value a particular branch is expected to hold. Thus, multihypothesis tracking can handle the passive case of belief space planning where disambiguation is sought only in the inference step and not in the planning. By contrast, DA-BSP argues for data association within the belief space planning framework and can thus utilize weights that are shaped by the actual future associations. Nevertheless, we can harness similar approaches to curtailing the empirical complexity of DA-BSP, classifying them as *local* or *global* and *pruning* or *merging*. When the decision about how to reduce a given mode in DA-BSP is based on the overall likelihood of associations considered from the initial position, we call it global, while in the local approach, only local information of the conditional posterior is sufficient to decide on it being reduced via merging or pruning. As is evident from the name, pruning is the process of dropping a component in a conditional posterior distribution and re-normalizing the other weights, whereas merging is the process of combining two components to form a single component, which is an optimal (in some heuristic sense) representation of the both. Both pruning and merging can be recursive processes.

5.5. Full vs. partial disambiguation

In the context of selecting an appropriate horizon for planning, we can note that in most of the real-world examples, the greater the horizon, the greater the likelihood of having a unique observation that results in disambiguation between several or all components of the belief. However, in general, DA-BSP does not require a complete or full disambiguation for its correctness. Recall from Section 3.1 that by *full disambiguation* we mean that the posterior belief eventually has only a single component. For a usual forward L -step

planning, this cannot be guaranteed unless we assume an existence of a unique observation in the future. At best, there would be *partial disambiguation*, i.e. some components of the posterior belief vanish due to less aliased observations. On the contrary, in the cases where a full disambiguation does not occur within the planning horizon, DA-BSP would maintain all the components with appropriate weights. This not only allows for partial disambiguation in such a planning scenario, where only the aliased components remain in the posterior belief, but can also eventually result in a full disambiguation. Hence, DA-BSP captures the reality of a perceptually aliased environment quite well.

5.6. Degenerate cases of DA-BSP

Two prominent reasons for considering data association aware belief space planning are, firstly, that it accurately reflects the reality where, owing to pose uncertainty, the observation may no longer be associated with that from the nominal pose and, secondly, that it is a generalization of the usual belief space planning. To elucidate the latter, we shall consider three degenerate cases of data association aware belief space planning: without pose uncertainty, with data association solved and without perceptual aliasing.

5.6.1. Without pose uncertainty. Consider that for all practical purposes, the pose is known with certainty; hence, the belief is a Dirac pulse around the nominal \hat{x}_k . Since the scene space could still be uncertain, the belief space planning should consider all possible scenes being observed. However, in this case, the integral in equation (37) reduces to a single term. In a more realistic case of small variance in the pose, considering only the most likely data association may still lead to a reasonable performance. This is similar to many passive inference based approaches, where the most likely component is often sufficient to account for the overall posterior.

Note that when significant pose uncertainty exists, yet is not assumed to do so, wrong and catastrophic inference could be made. This is seen in Figure 6, where it is wrongly inferred with certainty that the robot is in a vertically aligned pathway corresponding to component 2, whereas the ground truth is that the robot is in a horizontally aligned pathway corresponding to component 1.

5.6.2. With data association solved. In this case, the scene that is captured from perspective x_{k+1} when observation z_{k+1} is obtained is known. More precisely

$$\exists t, \mathbb{P}(z_{k+1}|x_{k+1}, A_j) = \begin{cases} 1 & j = t \\ 0 & j \neq t \end{cases}$$

This implies that the summation over all $\{A_N\}$ is reduced to a single A_t , known a priori for each observation z_{k+1} . Therefore, with data association solved, the framework degenerates to the usual belief space planning.

5.6.3. *Without perceptual aliasing.* In the absence of perceptual aliasing, while considering the observation z_{k+1} , we are guaranteed to have only a single pose and scene pair (x_{k+1}, A_j) that generated it. This implies that if the observation z_{k+1} were given, the posterior beliefs would be all zero, except for the one corresponding to A_j . However, since while planning at step k , the observation z_{k+1} is an unknown random variable, we would still need to consider all possible events $\{A_N\}$ that generated it.

5.7. On complexity and correctness of non-myopic DA-BSP

In this section, we are interested in answering two questions: *first*, how much more computational effort would be required to consider uncertain data association in the context of belief space planning and, *second*, would such an approach always guarantee a correct association—in other words, will it guarantee that the component corresponding to the ground truth association will be recovered whenever possible. For the former, we need to analyze the computational complexity of DA-BSP, which, in general, would depend on size of the state space. Thus, we shall consider only the discrete state space where such a complexity can be derived. We first note that belief space planning can be framed as a POMDP. In a general infinite horizon case, the POMDP is undecidable (Madani et al., 1999), whereas for a finite horizon polynomial in $|S|$, it is known to be PSPACE complete (Papadimitriou and Tsitsiklis, 1987). Since DA-BSP seeks to solve a harder problem than the usual belief space planning, i.e. to perform belief space planning under uncertain data association, it is natural that DA-BSP incorporates all scalability issues involved in solving any POMDP. In this section, we seek to make it more explicit by considering a discrete state POMDP.

As is usually done, let the POMDP be represented by the tuple $\langle S, \mathcal{A}, \Omega, T, O, r, b_0 \rangle$ and, furthermore, let the index i represent the time-epoch at which decisions regarding control actions are executed. Let L be the horizon of the problem, i.e. $i \in \{j | 1 \leq j \leq L, j \in \mathbb{N}\}$. The tuple symbols are:

- S . The set of unobservable states into which the system can possible transit; $\forall i, x_i \in S$;
- \mathcal{A} . The set of all admissible control actions that the system can perform; $\forall i, u_i \in \mathcal{A}$;
- Ω . The set of all possible observations that can be obtained; $\forall i, z_i \in \Omega$;
- T . The transition function denoting the probability of transitioning from state x to state x' under action u ; $T : S \times \mathcal{A} \times S \mapsto \mathbb{R}_{[0,1]}$;
- O . The observation function denoting the probability of observation z under action u ; $O : S \times \mathcal{A} \times \Omega \mapsto \mathbb{R}_{[0,1]}$;
- r . The reward function for calculating the reward in each state x and action u ; $T : S \times \mathcal{A} \mapsto \mathbb{R} \setminus \{-\infty, \infty\}$;
- b_0 . The initial belief; the probability distribution over whole state space S .

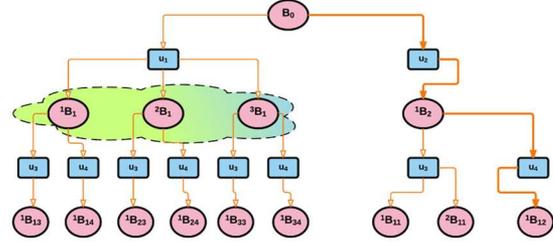


Fig. 7. Explicit representation of underlying POMDP as a *belief* Markov decision problem. Each circle represents a Gaussian component of the belief; rectangles denote actions u . Note that, owing to uncertain data association, more such Gaussian nodes may be created, as shown in the cloud here. In contrast, under unambiguous data association, such as that shown with bold-and-filled arrows, efficient planning can be performed by collapsing all the Gaussian beliefs into a single Gaussian. See Section 5.7 for the analysis in case of discrete state space.

Since, at any time step k , the actual underlying state is unknown, the system should reason about probability distribution over all states. The belief of being in state x is represented as $b(x)$. After each observation, this belief is updated as

$$b^{u,z}(x') = \frac{O(x', u, z)}{\mathbb{P}(z|u, b)} \cdot \sum_{x \in S} T(x, u, x') b(x) \quad (46)$$

which implies

$$b^{u,z}(x') \propto O(x', u, z) \cdot \sum_{x \in S} T(x, u, x') b(x) \quad (47)$$

Note that in the notation introduced in this paper, this can be written as:

$$b[X_{k+1} | z_{k+1}, u_k] \propto \mathbb{P}(z_{k+1} | x_{k+1}, u_k) \cdot \int_{x_k} \mathbb{P}(x_{k+1} | x_k, u_k) \mathbb{P}(x_k) \quad (48)$$

Visualizing equation (47) as a system of equations in $|S|$ variables (denoted by $b(x)$), we can see that each such update requires $O(|S|^2)$ computational effort, where O stands for big-O notation.

When data association is not solved, such a belief update must be made against each possible data association (see Figure 7), hence

$$b[X_{k+1} | A_i, z_{k+1}, u_k] \propto \mathbb{P}(z_{k+1} | A_i, x_{k+1}, u_k) \cdot \int_{x_k} \mathbb{P}(x_{k+1} | x_k, u_k) \mathbb{P}(x_k) \quad (49)$$

Consequently, the data-aware belief update would require $O((|S||\mathcal{E}|)^2)$ computational effort, where \mathcal{E} is the set of all events $\{A_N\}$. Assuming a finite horizon of size L , the overall complexity of the problem is $O((|\mathcal{O}||\mathcal{A}|)^L (|S||\mathcal{E}|)^2)$. Finally, we also note that, in practical applications, $|\mathcal{E}| \ll |S|$.

To reason about the *correctness* of DA-BSP; i.e. whenever there is a single disambiguating data association, the algorithm will recognize it and associate the observation correctly, we first define *pruned* and *unpruned* DA-BSP. Recall that DA-BSP adjusts the subsequent weights of the components based on the likelihood of the observation and of it being explained by the considered association. An unpruned DA-BSP considers all such associations, no matter how small the weights are (provided they are non-zero), while pruned DA-BSP has some reasonable threshold, below which all of the weights are pruned away. It is easy to see the correctness of unpruned DA-BSP. Consider that, at step $k \in [1, \infty)$, a full disambiguation occurs; then, by definition, the belief at $k - 1$, i.e. $b[X_{k-1}]$, will also contain the component corresponding to the ground truth. The subsequent computation of DA-BSP would yield weights that are all strictly zero, except the one corresponding to this ground truth. However, in the case of pruned DA-BSP, this might not necessarily be true, as the ground truth component might be pruned away in $b[X_{k-1}]$, possibly even leading to a catastrophic bad data association in the last step k . Note that this requires either the weight of the correct component to be too low or the pruning threshold to be too high. The former usually does not hold if we assume that, at the start, the multimodal belief contains the correct component as well. The latter can be avoided by judicious choice of pruning threshold. As shown in the experiments, DA-BSP is not sensitive to the choice of this threshold.

6. Experiments

In this section, we present an extensive analysis of the proposed approach. Throughout this section, we will assume that the set of action trajectories, from which the robot has to plan its future actions, is given a priori. DA-BSP is a general framework, which does not rely on any specific form of these trajectories. Typically, the nature of these actions is domain-specific, as we shall see in various setups analyzed in this section. How these actions are actually generated is outside of the scope of this work.

This section is organized into five different parts. First, we describe the relevant software and hardware infrastructure that was set up for experimenting with DA-BSP. Second, we mention some metrics that could be devised to evaluate the usefulness of considering data association within belief space planning. Since most state-of-the-art approaches assume the data association as given and perfect, such metrics are essential to evaluate DA-BSP. Third, we look into the abstract example introduced in Section 4.9, under the simplistic assumption of myopic planning, i.e., the planning horizon is one step. As we shall see, this allows us to understand conceptual differences that arise when data association is considered, even in this fairly simple example. Fourth, the case for explicit scenes is considered using AprilTags (Olson, 2011), in a perceptually aliased corridor environment with a Pioneer platform fitted with an

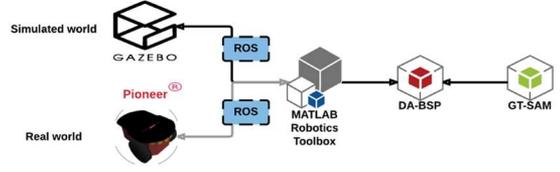


Fig. 8. Overall infrastructure of implementation of DA-BSP. Note that thanks to the middleware block of ROS, the algorithm is independent of whether it is applied to a real setup or a simulated one.

RGB camera. Here, we use non-myopic planning. Finally, to consider arbitrary levels of ambiguities, we create a simulated world model (in Gazebo) of two nearly identical office floors with various lookalike cubicles. In these simulations, the Pioneer robot is fitted with a laser scanner. Thus, scenes in this setup are implicit features of the world and may be a combination of several objects. In particular, each cubicle with walls and a chair would typically form a part of a scan that is treated as a single scene.

6.1. Implementation of data association aware belief space planning

Effective and realistic implementation of DA-BSP requires two separate threads of development. To be efficient, it is crucial that the algorithmic as well as the real-time cost of incorporating the data association within belief space planning remains as low as possible. We ensured this by representing each component of the Gaussian mixture model as a factor graph, so that state-of-the-art tool GT-SAM could be harnessed for a time-efficient inference. Conversely, to be realistic enough, we implemented it using a real robotic platform, Pioneer. Here, a propriety ROS Robotics Toolbox was used, which enabled our implementation to work seamlessly for both a simulated world and a real-world scenario. To simulate a complex world with arbitrary levels of ambiguity, we chose Gazebo, since it fits nicely with both the robotic platform and the ROS infrastructure. These two streams of development are shown in Figure 8. The DA-BSP algorithm itself was implemented in object-oriented MATLAB with the aim of striking a balance between rapid prototyping and obtaining a generalizable implementation that can be easily ported to such languages as C++.¹

6.2. Metrics for evaluating DA-BSP and compared approaches

Evaluation of DA-BSP is linked to the notion of data association, which is typically assumed to be solved in belief space planning. As mentioned before, accounting for data association within belief space planning does not come free. Conversely, assuming such an association randomly is bound to fail. For a very simple case of myopic planning, we compute the error in the posterior distribution as the metric

distance between the ground truth and the mean of the closest component. Suppose that the Gaussian mixture model belief has N components with weights w_i , means μ_i and covariances Σ_i for $i = [1, N]$. Now, if x is the ground truth position, then error δ is defined as

$$\delta = \min_i \|x - \mu_i\| \quad (50)$$

Such an error for a usual belief space planning approach shall be denoted δ_{BSP} . Note that since usual belief space planning has a single component, δ_{BSP} is the distance between the mean of this component and the ground truth. Similarly, δ_{DA} stands for the error in DA-BSP. To capture the probabilistic aspect of the problem, these metrics are averaged over several random trials.

Later, to evaluate the data association aware approach, we keep track of how many components the belief has and how many times it can make a correct association. We define \tilde{m} as an average number of components over all random runs, i.e. for R random runs where N_r is the number of components for run r

$$\tilde{m} = \frac{\sum_{r=1}^R N_r}{R} \quad (51)$$

Recall from Section 5.7 that, as long as the component corresponding to the ground truth is present in the belief, we define the belief to be *correct* (from the data association point of view). We denote this by Boolean symbol *data association*. Note that, in the case of no pruning, DA-BSP is guaranteed to be correct; hence, the data association is set as true. This is also true if the pruning is not detrimental to the correct component and association. ξ_{ca} measures the weight of the correct component in the belief, i.e. if, out of N components, the component c corresponds to the correct data association, then

$$\xi_{\text{ca}} = \omega_c \quad (52)$$

If the correct component is lost, the corresponding value of ξ_{ca} will be 0. Time taken by DA-BSP in any epoch is directly related to the number of the components in the belief. We keep track of this through the metric \tilde{m} .

Since usual belief space planning considers data association to be given and perfect, DA-BSP cannot be compared directly against the state-of-the-art approaches. To evaluate DA-BSP, we compare it against the approach where, with some probability $1 - \epsilon$, the true association is known and made by the planning while, in all other cases, a random choice from incorrect associations is made. This approach implies that the belief is always unimodal and is therefore named *BSP-un.i*. Similarly, in another approach, we assume that the correct association with the scene is made with probability $1 - \epsilon$. However, all components of the propagated belief are considered for such an association. This approach implies that a multimodal prior distribution also remains multimodal after inference. It is named *BSP-mu.l* here. In both of these variants, we are interested in correct

data associations being made out of many trials. This is measured by the metric ξ_{ca} , where the value 1 would indicate that, in all random trials, the component corresponding to the correct data association had weight 1 in the belief, i.e. the belief had a single component, corresponding to the correct association. This could happen, for example, when the belief is unimodal, owing to lack of ambiguity in the vicinity.

These metrics are summarized in Table 1.

6.3. An abstract example for DA-BSP

Referring to the abstract example mentioned in Section 4.9, we present in Table 2 a numerical analysis of the cost computation (see equation (30)) of these configurations, as well as a metric $\{\delta_{\text{BSP}}, \delta_{\text{DA}}\}$ quantifying the estimation error, defined over incorrect (with respect to ground truth) associations through $N = 100$ random samples of various modes. Note that DA-BSP is independent of the cost function considered while planning. Here, we consider three different costs. We refer the reader to Section 4.7 for details. The *worst* cost is the trace of the maximal covariance among all components while *maximum weight* considers the component with the greatest weight while computing the trace. In this case, both these result in the identical values of the cost. As mentioned before, KL_u is the Kullback–Leibler divergence of the weights of the components of the belief compared with the uniform distribution. Under no aliasing, i.e. $\{\Phi\}_{\text{alias}}$, the number of components denoted by *modes* does not increase.

Now, consider the action u_1 . Since it may lead to three scenes that are perceptually aliased, DA-BSP may have three components, whereas the usual belief space planning approaches are bound to commit to a single data association. Thus, as would be expected, this leads to a smaller error in the case of DA-BSP; e.g. when A_1 and A_2 alias ($\{A_1, A_2\}_{\text{alias}}$), δ_{BSP} is almost $6 \times \delta_{\text{DA}}$. Recall from equation (31) that, unlike action u_1 , action u_2 leads to fully unambiguous observations around the most likely values (see Figure 4) and, consequently, $\delta_{\text{BSP}} \simeq \delta_{\text{DA}}$. Thus, we can conclude that not only would DA-BSP lead to smaller estimation error when there is a perceptual aliasing but also that DA-BSP degenerates to usual belief space planning approaches when helpful assumptions could be made (such as in the case of no perceptual aliasing).

6.4. Real-world application with explicit scenes—octagonal corridor

To elucidate the crucial properties of non-myopic DA-BSP, we consider a real-world experiment, as shown in Figure 9 with a single Pioneer robot. The robot resides in an octagonal corridor with ample instances of ambiguous scenes. The abstracted schema of the world is shown in the center, while the surrounding figures are the third-person view of the environment. The state space $X \in \mathbb{R}^3$ consists of

Table 1. Metrics used in the experiments.

Metric	Description
δ_{BSP}	Distance between ground truth and closest component under usual belief space planning
δ_{DA}	Distance between ground truth and closest component under data association aware belief space planning
Data association	Boolean flag, which is set true if the ground truth component is within the belief
ϵ	Probability of incorrect data association
ξ_{ca}	Weighted probability of correct data association
\tilde{m}	Number of components in the belief

Table 2. Evaluating different cost functions for various configurations (see Figure 4).

Configuration	Cost Worst	Maximum weight	KL_u	Mode	Estimation error		\mathcal{U} u_1/u_2
					δ_{BSP}	δ_{DA}	
$\{\Phi\}_{\text{alias}}$	0.0977	0.0977	0.3496	1.1000	0.1851	0.1717	u_1
	0.1009	0.1009	na	1.0000	0.3461	0.3999	u_2
$\{A_1, A_2, A_3\}_{\text{alias}}$	0.0508	0.0508	0.5072	3.0000	1.1654	0.4772	u_1
	0.1009	0.1009	na	1.0000	0.3832	0.3990	u_2
$\{A_1, A_2\}_{\text{alias}}$	0.0833	0.0833	0.3757	1.5500	1.2197	0.2114	u_1
	0.1009	0.1009	na	1.0000	0.3912	0.3992	u_2
$\{A_1, A_3\}_{\text{alias}}$	0.0849	0.0849	0.3649	1.4000	1.0552	0.4197	u_1
	0.1009	0.1009	na	1.0000	0.4101	0.3940	u_2

na: not applicable.

2D coordinates, as shown, as well as the orientation of the robot. Here, the actual floor is shown via a laser scan. The map of the environment is given in terms of AprilTags and their corresponding ids. This enables us to simulate arbitrary levels of perceptual aliasing by AprilTags with identical tag ids. In particular, we have tags with ids 1 and 2 depicted in blue ‘+’ and red ‘*’ markers, respectively. Note that the laser scanned map is depicted here for representation purposes only. It is neither exact nor is it provided to the robot while planning. Because of this, some tag positions might appear non-aligned to the laser map. Initially, the belief of the robot is a multimodal distribution, represented by a Gaussian mixture model with four components having equal weights and shown by the ellipses in Figure 9. Typically, the tag detection is decided through the centrality of the tag in the image observed by the camera. Figure 10 shows the case where the off-center tag is not detected.

The objective of the robot is to both localize itself and to reach a specific elevator; all elevators are denoted ∇ . Initially, the planner is provided with a set \mathcal{T} of control trajectories; see Figure 11 for some of the trajectories considered. Consequently, depending on the planning algorithm used (DA-BSP, BSP-uni, or BSP-mul), as well as the planning horizon L , the cost of each trajectory $\tau \in \mathcal{T}$ is evaluated and the optimal (with respect to this cost) trajectory is chosen. Note that these trajectories are different in length and, owing to the specific nature of the problem, the robot will also witness different levels of perceptual aliasing while following them. Later on, in further experiments, the results would be shown for an arbitrarily chosen trajectory, viz. trajectory 1. 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 problems of longer time scales—can all be easily incorporated into it.

Since we model the visual observation via AprilTags, owing to sensory limitations (such as out-of-view or far-from-center tags), a reliable observation might not be available at each motion step. One such instance is depicted in Figure 10. In such conditions, no data association can be made (as there is no measurement) and consequently, DA-BSP behaves exactly like the usual belief space planning, albeit with a prior distribution that could have many components.

The purpose of Figure 12 is to illustrate the evolution of belief in the DA-BSP framework. The prior distribution is a multimodal belief with four components. When only the motion model is incorporated, the propagated belief still has four components but each has higher uncertainty. Finally, when a perceptually aliased observation is made, DA-BSP adjusts the weights of the components of the posterior distribution where, after a simple pruning, only three components remain. This reduction is because there is slight asymmetry at the ends of the corridors. However, when ambiguous data association occurs, DA-BSP considers all possible associations and weighs each new component of the posterior distribution according to equation (44). Figure 12 shows one such instance.

The result of running DA-BSP on this setup is shown in Table 3. As can be seen, the computational cost of

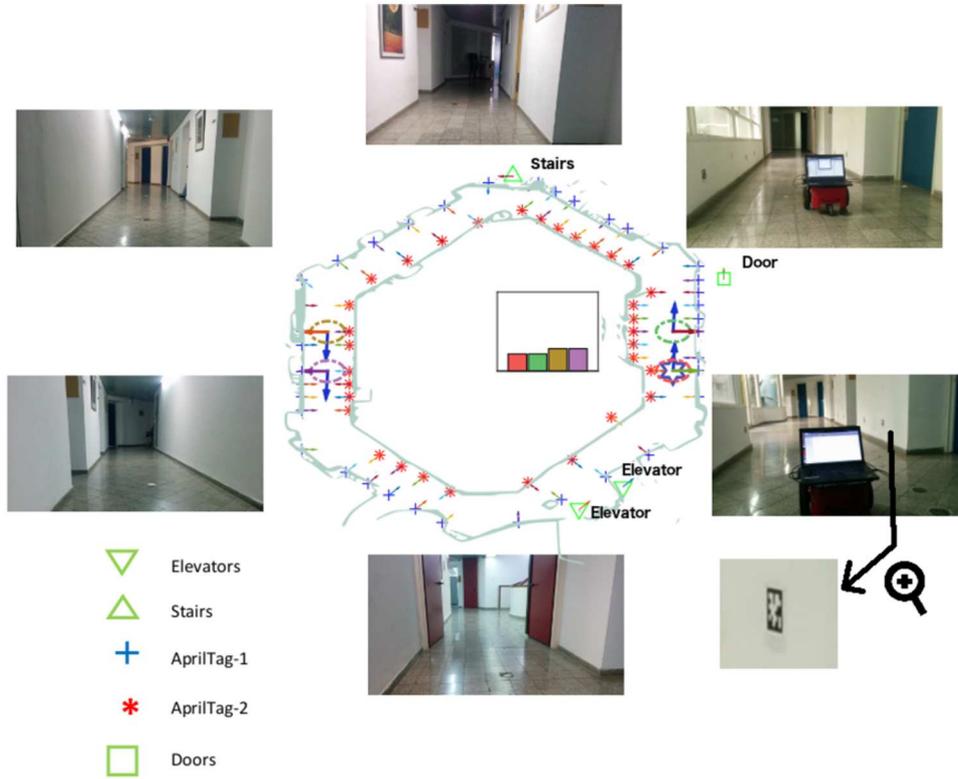
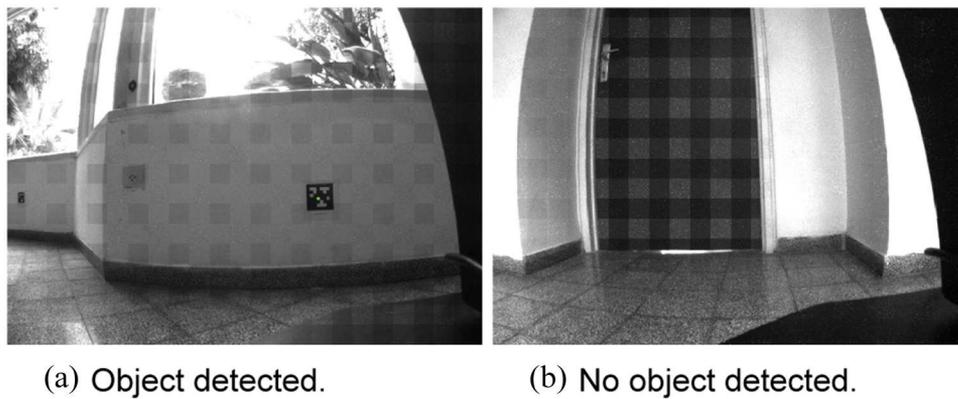


Fig. 9. Real-world experimental setup (best viewed in color). The images show the actual third-person view at different locations of the robot, while the two-dimensional figure shows the semantic knowledge of the environment that the robot possesses. Its current belief is a four-modal Gaussian mixture model with mean position depicted by x_{init} . The ground truth robot position is indicated by \odot , while two sets of aliasing scenes with tag ids 1 and 2 are shown with blue + and red * respectively; arrows indicate orientation (not motion). The actual scenario is depicted through the laser scan of the environment shown in green. This map is for representative purposes only and is not available to the robot. The zoomed-in picture of an AprilTag is also shown.



(a) Object detected. (b) No object detected.

Fig. 10. (a) AprilTag is detected, indicated by green patch at center. This provides the transformation matrix between the pose of the robot and the landmark pose. Note that a far-away AprilTag, though visible in this frame, is considered undetected, since the non-centrality of the tag makes the observation highly untrustworthy. (b) No AprilTag lies within the field of view of the camera.

DA-BSP planning depends on the number of components maintained in the posterior distribution, i.e. by \tilde{m} . \tilde{m} indicates the level of disambiguation with the belief while ξ_{ca} shows how well the correct data association is accounted for. For example, at epoch 20, both planning and inference have high values of ξ_{ca} (0.73 and 1.00, respectively) but full

disambiguation occurs for inference only ($\tilde{m} = 1$). Note that DA-BSP considers robust active and passive inference in a single framework. Hence, whether it is reasoning about future observations while planning (the left part of the table) or inferring based on a given observation (the right part of the table), the DA-BSP algorithm remains the same.

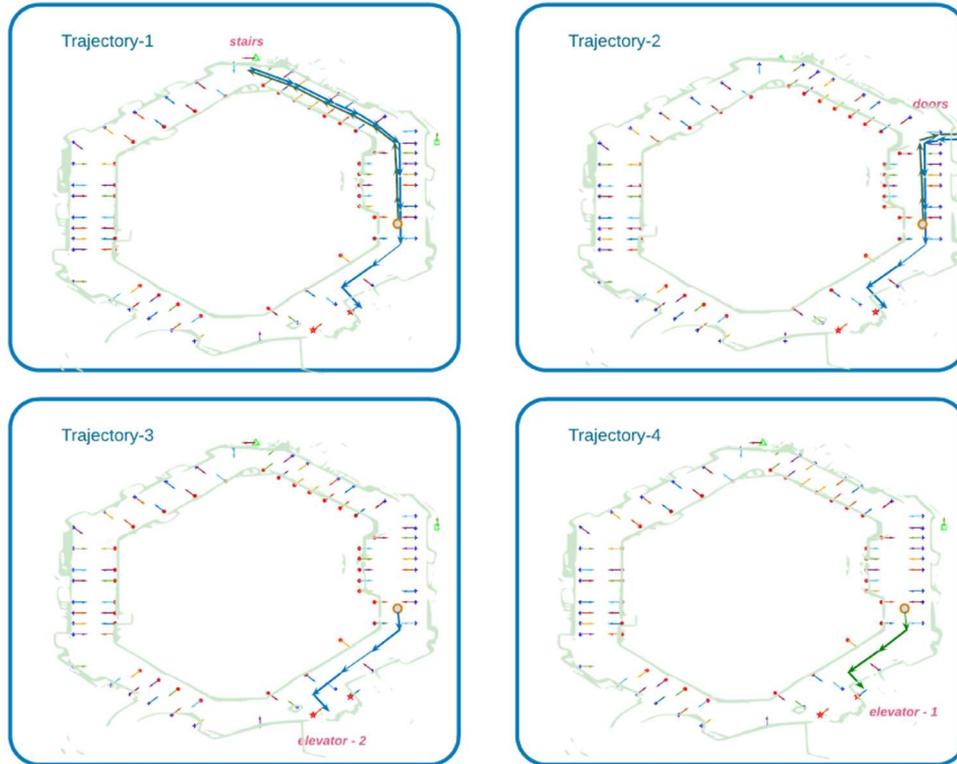


Fig. 11. Some of the trajectories considered by the robot (best viewed in color). The starting position is marked with a filled ellipse. The door on the right, the stairs, and the elevators, are shown by \square , Δ , and \times , respectively, and also marked by the text, where relevant.

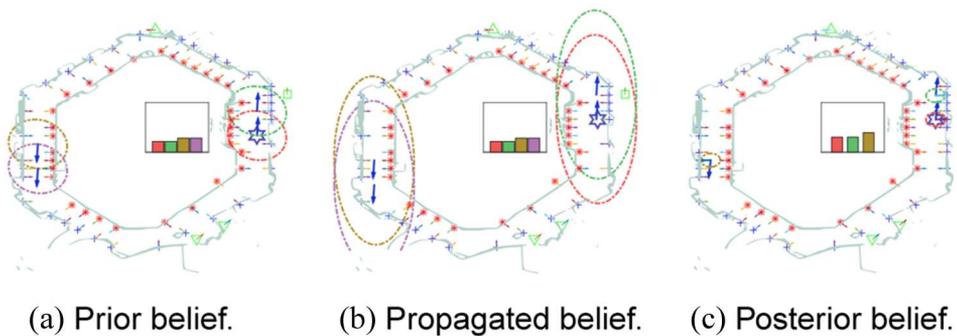


Fig. 12. Evolution of belief at epoch $E = 3$ under DA-BSP. Here, tags are represented by shapes $\{*, +, \square, \Delta, \nabla\}$, while the ground truth robot position is indicated by \star . (a) The prior belief is multimodal with four distinct modes, as shown by the colored ellipses. (b) After incorporating the motion model, the propagated belief is similarly a multimodal distribution. (c) When observation is accounted for and inference is performed, the posterior belief is as shown. Note that some of the earlier components of the prior might vanish (e.g. here, the slight asymmetry around the corner causes one component to vanish). Also new components in the posterior may emerge (not the case here). Here, $L = 3$.

DA-BSP incorporates planning and inference seamlessly under one framework, called an *epoch* earlier. The overall planning is performed as a model predictive control composed of several such steps. Figure 13 shows some of the epochs in DA-BSP, along with other approaches such as BSP-uni and BSP-mul. Here, the trajectory under evaluation has 24 actions. At the start, the belief has four modes with equal weights. DA-BSP results in these components being differently weighted in the next epoch, while at epoch

4 only two components remain. Later on (such as epoch 18), many more components arise, as the AprilTags are more densely distributed on the inner walls of the initial corridors. In contrast with this, BSP-uni and BSP-mul may result in catastrophically bad inference; one such instance of each is depicted in Figure 13.

Once the planning is performed under DA-BSP, the subsequent posterior distribution at the end of each epoch might have more or even fewer components than before.

Table 3. Performing DA-BSP on a real corridor environment shown in Figure 9, with planning horizon $L = 4$. The times in seconds spent in planning and in inference are denoted t , while \bar{m} stands for average modes; refer to Section 6.4.

Algorithm	Epoch	Planning			Inference		
		t (s)	ξ_{ca}	\bar{m}	t (s)	ξ_{ca}	\bar{m}
DA-BSP	1	21.81	0.09	6.00	0.84	0.22	4.00
	4	5.19	0.28	2.50	0.84	0.31	3.00
	8	8.66	–	1.00	0.80	1.00	1.00
	12	19.90	–	6.67	2.48	0.35	5.00
	16	3.50	0.16	2.00	0.14	–	10.00
	20	4.51	0.73	3.80	0.31	1.00	1.00

The former occurs when the presence of identical close tags causes perceptual aliasing, while the latter is the result of unlikely components being pruned away naturally, in the light of new observations. This is evident in Figure 13(d), where at epoch 8, the belief distribution is reduced to a single component and hence the correct corridor is determined but later on, by epoch 18, the belief distribution again has as many as eight components, arising from close identical AprilTags.

Consider Table 4, where we evaluate DA-BSP as well as other comparable alternatives, such as BSP-uni and BSP-mul, described previously. The metrics used here (such as ξ_{ca} and data association) were already described in detail in Section 6.2. The values are averaged over five random runs, where data association is set true only when the correct data association is accounted for in all of the random runs. When no plausible data association can be made, we represent these metrics as ‘–’. This can happen either because there is no AprilTag in the vicinity of the robot and hence no observation could be made (such as in epoch $E = 8$) or when the inference in the previous planning is too far way from the ground truth to make any plausible data association (such as several epochs in BSP-uni with $L = 3$). It could be argued that, in the presence of severe perceptual aliasing, both BSP-uni and BSP-mul are ill-equipped approaches for planning, as incorrect data association would imply almost certain catastrophic inference. To remedy this, we consider first the case where correct data association is provided to BSP-uni and BSP-mul with a high probability of 90% i.e., randomization is $\epsilon = 0.1$. Later on, in Table 5, we see how BSP-uni deteriorates (ξ_{ca} decreases) with increasing randomization, ϵ .

Naturally, DA-BSP, when compared with BSP-uni and BSP-mul, trades computation efficiency with such correctness of data association. On the one hand, DA-BSP always has a component corresponding to correct data association (seen from the Boolean flag data association), while on the other hand, it is significantly slower than BSP-uni and BSP-mul, especially for non-myopic planning with larger horizons (seen from the time taken, denoted t). However, exponential blow up of computational complexity as the planning horizon increases is an issue not specific to DA-BSP. Unfortunately, DA-BSP cannot solve or

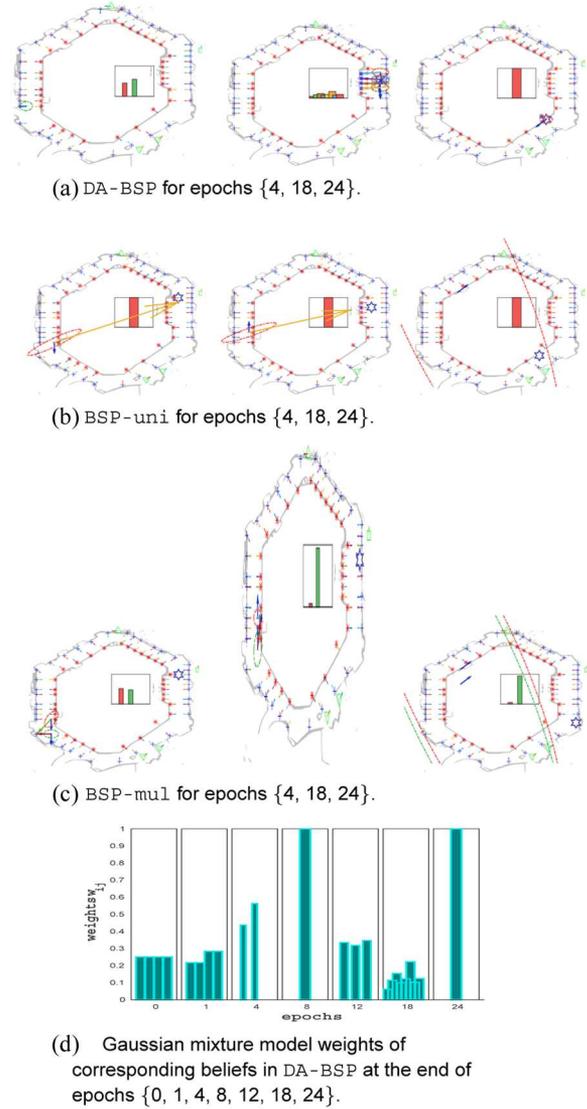


Fig. 13. (a–c) Evolution of inferred belief as *decision epoch* progresses with $L = 3$; epochs depicted are $\{4, 18, 24\}$. The trajectory of actions is *Trajectory 1* from Figure 11. These depict evolution of inferred belief, for different planning algorithms, DA-BSP, BSP-uni, and BSP-mul, respectively. Gaussian mixture model components and associated weights are shown in different colors. Ground truth robot position is indicated by \odot . For clarity, the detected scenes are shown in different colors. For BSP-mul and BSP-uni, this particular instance of planning leads to catastrophically bad data association. (d) Evolution of Gaussian mixture model component weights during these epochs. Note that the number of components *increases* and *decreases* and eventually goes to 1. Here, the planning horizon is $L = 3$.

even reduce this burden. On the contrary, if the world is replete with perceptually aliased scenes, computational effort would increase significantly. Fortunately, owing to *parsimonious data association* (grounded in the fact that the actual world is not replete with exact similarities) the additional cost of DA-BSP may not be significantly greater;

Table 4. Comparing DA-BSP against BSP-uni and BSP-mul in several steps of planning and inference, with $L = 1$ and $L = 3$. The times in seconds spent in planning and in inference is denoted by t , while average modes are shown by \bar{m} . Values shown here are for an average of five random runs. Here, data association signifies *correct* data association in *all* of the random runs, as defined in Section 6.2. Results for BSP-uni and BSP-mul were obtained for $\epsilon = 0.1$, i.e. correct data association was provided to BSP-uni and BSP-mul with a high probability of 90%. Performance of these approaches deteriorates for increasing randomization, ϵ ; see Table 5.

Algorithm	Epoch	$L = 1$				$L = 3$				Inference			
		t (s)	ξ_{ca}	\bar{m}	Data association	t (s)	ξ_{ca}	\bar{m}	Data association	t (s)	ξ_{ca}	\bar{m}	Data association
DA-BSP	1	2.60	0.11	4.00	✓	95.57	0.08	5.95	✓	0.80	0.22	4.00	✓
	2	1.21	0.29	2.00	✓	5.75	0.13	1.37	✓	0.05	–	4.00	–
	4	1.00	0.35	2.00	✓	4.29	–	1.00	–	0.61	0.50	2.00	✓
	8	0.11	–	1.00	–	0.35	–	1.00	–	0.02	–	1.00	–
	12	3.90	0.11	4.80	✓	191.48	0.08	6.79	✓	1.16	0.28	4.20	✓
	16	2.62	0.12	3.03	✓	3.58	–	3.02	–	0.60	0.11	4.60	✓
	19	3.14	0.09	2.60	✓	82.16	0.04	6.10	✓	0.94	0.14	6.60	✓
		t (s)	ξ_{ca}	Data association	t (s)	ξ_{ca}	Data association	t (s)	ξ_{ca}	Data association			
BSP-uni	1	0.43	0.90	×	2.19	–	–	0.20	1.00	✓			
	2	0.15	–	–	1.43	0.86	×	0.03	–	–			
	4	0.25	1.00	✓	4.51	0.98	×	0.17	1.00	✓			
	8	0.15	–	–	1.10	–	–	0.05	–	–			
	12	0.26	1.00	✓	3.90	–	–	0.17	1.00	✓			
	16	0.16	–	–	1.11	–	–	0.08	–	–			
	19	0.30	1.00	✓	1.24	–	–	0.17	–	–			
		t (s)	ξ_{ca}	Data association	t (s)	ξ_{ca}	Data association	t (s)	ξ_{ca}	Data association			
BSP-mul	1	2.74	0.15	×	34.33	0.18	×	0.86	0.80	×			
	2	2.01	0.27	×	20.84	0.40	×	0.03	–	–			
	4	1.66	0.23	×	4.14	–	–	0.77	0.20	×			
	8	0.77	–	–	1.54	–	–	0.18	–	–			
	12	0.80	0.80	×	1.52	–	–	0.81	0.20	×			
	16	2.33	0.27	×	14.39	–	–	0.33	–	–			
	19	1.70	0.63	×	38.33	0.82	×	0.48	–	–			

Table 5. Evaluating non-myopic BSP-uni in several steps of planning and inference, under different randomizations of $\epsilon = \{0.25, 0.5, 0.75, 1.0\}$. Recall that ϵ is the probability with which BSP-uni chooses a random association out of all *plausible* associations. The times in seconds spent in planning is denoted t , while the average correct association is denoted ξ_{ca} . Values shown here are for an average of five random runs while standard deviation is depicted within parentheses.

Algorithm	Epoch	$\epsilon = 0.25$		$\epsilon = 0.50$		$\epsilon = 0.75$		$\epsilon = 1.00$	
		t (s)	ξ_{ca}						
BSP-uni	1	1.14 (0.23)	0.78 (0.43)	1.08 (0.33)	0.90 (0.14)	1.14 (0.30)	0.96 (0.06)	1.17 (0.34)	0.92 (0.08)
	6	0.13 (0.08)	– (–)	0.38 (0.33)	– (–)	0.20 (0.10)	– (–)	0.23 (0.25)	– (–)
	12	1.17 (0.61)	0.76 (0.43)	1.30 (0.47)	0.59 (0.54)	0.97 (0.60)	0.54 (0.50)	1.02 (0.63)	0.36 (0.50)
	18	1.23 (0.60)	1.00 (0.00)	1.04 (0.74)	0.52 (0.49)	0.60 (0.58)	0.37 (0.51)	0.46 (0.23)	0.20 (0.45)

e.g. BSP-mul has similar magnitude of time taken to DA-BSP.

DA-BSP accounts for all plausible data association and hence might have many components in the belief (recall that the number of components is denoted \bar{m}). We depict in Table 4 epochs where ξ_{ca} might be low for DA-BSP, such as in epoch $E = 19$. Note that, even here, the correct association is being accounted for by DA-BSP. The low value of ξ_{ca} might arise from the exact nature of the observation,

which may be better explained by an aliasing scene than by the ground truth, in at least some of the random runs.

Another notable aspect is how often the belief incorporates the component corresponding to the correct data association. In the case of all of these algorithms, data association and ξ_{ca} quantify this aspect. Under highly uncertain data association, BSP-uni and BSP-mul would always come up with catastrophically bad inference. Recall that the results for BSP-uni and BSP-mul are shown in Table 4

for $\epsilon = 0.1$. Note that for BSP-uni under myopic planning ($L = 1$) correct association is often found and, correspondingly, $\xi_{ca} = 1$. However, with a non-myopic planning horizon ($L = 3$), this decreases and data association is never set as true. This also means that ξ_{ca} for BSP-uni is higher in inference than in the planning in general. On the contrary, BSP-mul never has data association set to true because in at least one of the random runs, it associates to the wrong scene. Consequently ξ_{ca} is always less than 1. These are highly suitable settings for BSP-uni and BSP-mul, since the randomization is low ($\epsilon = 0.1$). Table 5 shows an analysis for BSP-uni when the randomization is higher.

Table 5 depicts the performance of BSP-uni under different randomizations of ϵ with planning horizon $L = 2$. As can be expected, the data association performance of BSP-uni deteriorates gradually with increasing ϵ , while the computational effort remains the same. For example, for epoch 12, the metric ξ_{ca} decreases from 0.76 for $\epsilon = 0.25$ to 0.36 for $\epsilon = 1$. Note that, under mild randomization with $\epsilon = 0.25$, BSP-uni was successful in obtaining correct association in all five random runs, hence $\xi_{ca} = 1$. Also, to have sufficiently high ξ_{ca} at all epochs, the underlying algorithm should always guess the data association correctly. This illustrates the importance of DA-BSP in realistic scenarios, where this cannot be assumed.

Table 6 shows the effect of the pruning parameter on DA-BSP. As can be seen, DA-BSP is insensitive to such a pruning parameter σ . Even at a very high pruning threshold, DA-BSP may result in sufficiently decent planning, whereas a threshold of $\sigma_{prune} = 5\%$ was sufficient to obtain planning similar to the case of unpruned DA-BSP, which was simulated by using $\sigma_{prune} = 10^{-10}$. This is because realistic scenarios typically do not have persistent ambiguity at each step of navigation; hence, the weights of many components decrease naturally to afford an easy approach of pruning. Thus, at first, it might appear that DA-BSP is hopelessly expensive in terms of computational efforts and non-trivial pruning techniques might be required to make it applicable in any realistic scenario. However, quite the contrary is true.

Another unique aspect of DA-BSP is that the weights of the components are adjusted as is suitable after considering all future observations in both myopic and non-myopic settings. Based on the configuration of the environment, a longer planning horizon may enable quicker disambiguation and consequently reduced Kullback–Leibler cost. In Figure 14, we see how the number of components, as well as this cost, varies across different epochs of DA-BSP and also under various planning horizons.

6.5. Highly-aliased simulated office scenario

To demonstrate our concept in a more challenging scenario under a high level of perceptual aliasing, we considered a Gazebo-based simulation of a Pioneer robot in an aliased

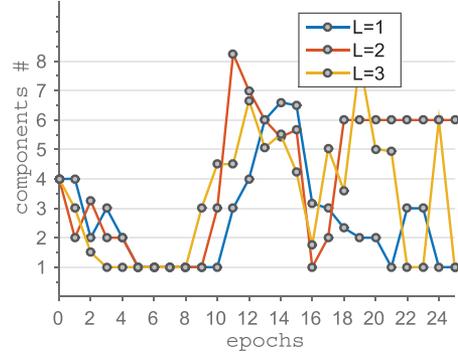


Fig. 14. Evolution of belief as *decision epoch* progresses during DA-BSP planning. The average number of components in the belief mixtures is depicted in the plot.

two-floor office room environment. This domain also illustrates the generality of DA-BSP, since it allows for a different representation for the scenes, as against the AprilTags considered before. The robot is fitted with realistic sensors enabling laser scans and odometry estimation. Apart from the implementation mentioned in Section 6.1, we use the iterative closest point algorithm for laser scan matching. This implies that the scenes against which DA-BSP considers the data association are reasoned implicitly through laser scans. In other words, a laser scan of an office cubicle with chair and desk is considered as a single scene. Based on sufficiently high iterative closest point matching, this scene aliases with the laser scan from another similar-looking cubicle.

The scenario is as shown in Figure 15(a). Unless stated otherwise, we will use natural numbers to denote specific places in this scenario as it is depicted and notation $x \rightarrow y$ to show a path from x toward y . The two floors are identical except that floor 2 has an additional printer p_1 (Figure 15(a)). Additionally, each floor has significant perceptual aliasing within itself, owing to identical cubicles and self-similar corridors. However, at the ends of the corridors, there could be a disambiguating feature present, such as a vending machine and sofa at one end and a printer at the other. The goal for the robot is to reach cabin c_1 (Figure 15(a)) and to disambiguate between the floors. Initially, the robot wakes up to find itself in either of the places 1 or 6 (facing places 2 or 7, respectively). Hence, its initial belief is modeled as a four-component Gaussian mixture model (two for each floor), whereas the ground truth is at position 1, i.e., the robot actually is at 1. Throughout this section, we use green and yellow to denote the ground truth and the aliasing, respectively.

Consider that the robot starts at 1 (Figure 15(a)); thus the initial belief has a mean at this position. A forward action to 2 can be used to propagate the initial belief. The subsequent prior and the propagated (means at 1 and 2, respectively) covariances are shown in Figure 16(a). The area of the ellipses equals the actual 2σ covariance. The laser scan obtained for the belief update is shown in

Table 6. Evaluating DA-BSP in several steps of planning with horizon $L = 2$, under different pruning thresholds of $\sigma = \{0.2, 0.05, 10^{-10}\}$. The time in seconds spent in planning is denoted t , while ξ_{ca} and \tilde{m} show weight of correct association and averaged number of modes, respectively. Values shown here are for an average of five random runs, while standard deviation is depicted within parentheses. Note that a 5% threshold is sufficient in this case to perform equivalently with almost *unpruned* DA-BSP with $\sigma = 10^{-10}$.

Algorithm	Epoch	$\sigma = 0.20$			$\sigma = 0.05$			$\sigma = 10^{-10}$		
		t (s)	ξ_{ca}	\tilde{m}	t (s)	ξ_{ca}	\tilde{m}	t (s)	ξ_{ca}	\tilde{m}
DA-BSP	1	14.88 (6.80)	0.11 (0.10)	2.60 (0.89)	16.02 (7.74)	0.12 (0.08)	4.70 (2.22)	14.27 (8.79)	0.15 (0.09)	4.40 (2.19)
	2	5.21 (4.78)	0.50 (0.34)	1.75 (0.75)	5.26 (4.87)	0.50 (0.34)	1.85 (0.93)	7.60 (5.74)	0.21 (0.10)	2.50 (0.68)
	6	0.37 (0.20)	0.20 (0.45)	1.00 (0.00)	0.38 (0.19)	0.20 (0.45)	1.00 (0.00)	0.50 (0.17)	0.60 (0.55)	1.00 (0.00)
	13	13.39 (4.02)	0.29 (0.23)	2.13 (0.51)	52.49 (41.63)	0.17 (0.13)	6.86 (1.58)	114.35 (71.18)	0.10 (0.09)	13.91 (6.80)
	16	2.43 (4.49)	0.18 (0.26)	1.70 (0.45)	8.58 (9.72)	0.34 (0.20)	3.18 (2.03)	9.49 (8.34)	0.12 (0.06)	3.37 (1.87)

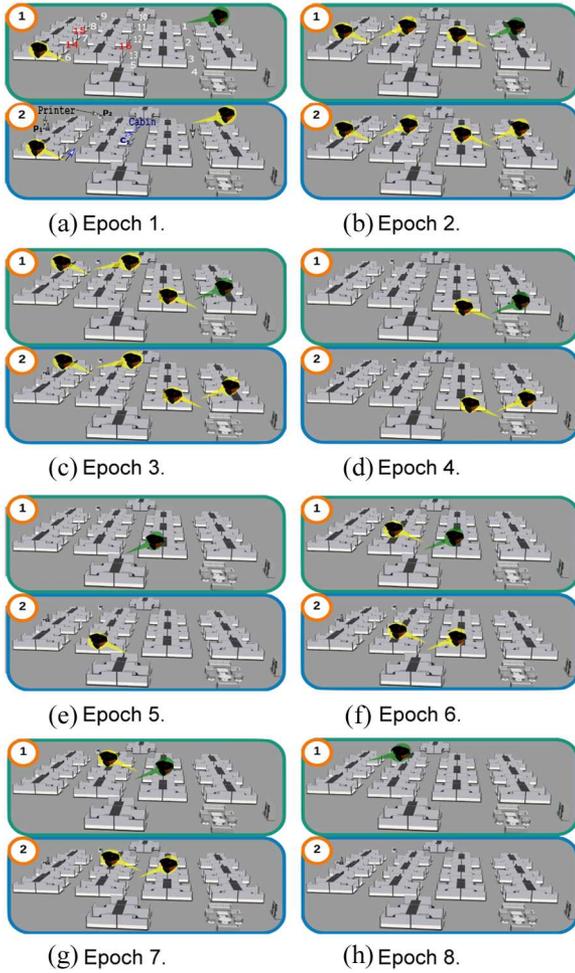


Fig. 15. (a) Two-floor aliased office environment in a Gazebo simulator: p_1 and p_2 denote printers, while 1 and 6 are the mean positions in each floor for the initial four-component Gaussian mixture model belief. (b–h) Mean positions (modes) of robot for each step of the DA-BSP path. Green denotes ground truth; yellow indicates the aliasing position.

Figure 17(b), where the green scan denotes the actual scan obtained. Note that from a different view point a similar scan is obtained (shown in yellow). This is due to the aliasing nature of the environment; considering this aliasing scan

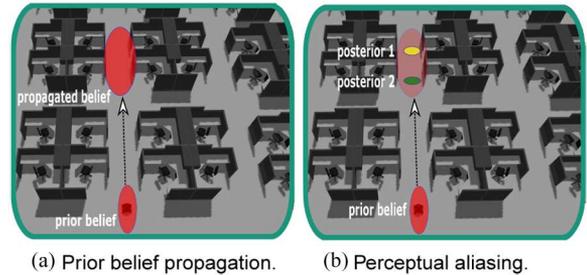


Fig. 16. Prior belief is propagated according to the motion model. Within the subsequent propagated belief, perceptually aliased laser scans are observed. Here, 2σ covariance is depicted with each ellipse.

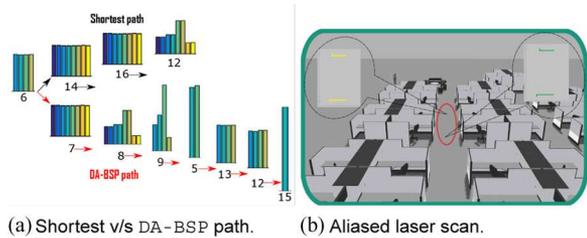


Fig. 17. (a) Evolution of weights of the components of the belief when following the shortest path versus that following the DA-BSP path. (b) Laser scans at ground truth and aliased position (green and yellow respectively).

within our planning-inference framework (DA-BSP) gives rise to two components in the posterior belief, each of which are weighted according to the corresponding likelihood for the respective scans to be obtained. See Figure 16(b).

Starting from positions 1 or 6, there are many possible paths to reach the goal (cabin c_1 (Figure 15(a)). We would like to show two such paths. The shortest path is $6 \rightarrow 14 \rightarrow 16 \rightarrow 12$ (Figure 15(a)). However, it leads to an increase in the number of modes; on reaching the goal, the robot is uncertain of the floor it is in. As seen previously (e.g., Figure 16), the modes increase, owing to the highly aliasing environment. Now consider a longer path, $1 \rightarrow 2 \rightarrow 3 \rightarrow 4 \rightarrow 5 \rightarrow 13 \rightarrow 12 \rightarrow 15$ (Figure 15(a)). Let us call this the DA-BSP path. While following $1 \rightarrow 2$ and $2 \rightarrow 3$, owing to the aliased cubicles, the number of

components increases from four to eight. See Figure 15(a) to (c) for the corresponding mean positions of the robot. Intrafloor disambiguation occurs along the paths $3 \rightarrow 4$ and $8 \rightarrow 9$. This is because of unique features present, viz. the sofa and the printer, for these respective paths. Similarly, along $4 \rightarrow 5$, the components are reduced to two (in Figure 15(e)) and then increase again to four along the paths $5 \rightarrow 13$ and $13 \rightarrow 12$ (in Figures Figure 15(f) and (g), respectively). Full disambiguation, resulting in a unimodal belief, occurs at 15, owing to the presence of the unique printer, p_1 (Figure 15(a)). Figure 17 depicts the evolution of weights along both paths.

In Figure 18, we see how the different components of the belief and the respective weights evolve when following the DA-BSP path. When planning with horizon $L = 2$, the components increase in number and retain similar weights ($E = 2$), while subsequent discrimination ($E = 3$) and reduction within the components (e.g. $E = 5$) leads eventually to full disambiguation ($E = 8$). Figure 18(b) shows the cardinality of components in the Gaussian mixture model during planning with different horizons viz., $L = \{1, 3, 5\}$. It can be seen that the graph gets steeper with increasing L . For a specific path and depending on the configuration of the environment, a longer planning horizon might help us disambiguate faster, as can be seen from Figure 18(b). Note that full disambiguation occurs at $E = 8$ for myopic planning ($L = 1$). Thus, for $L = 5$, which can project five steps in the future, such as disambiguation occurs from $E = 4$ onwards. $L = 3$ lies somewhere in between where the full disambiguation occurs, from $E = 6$ onwards.

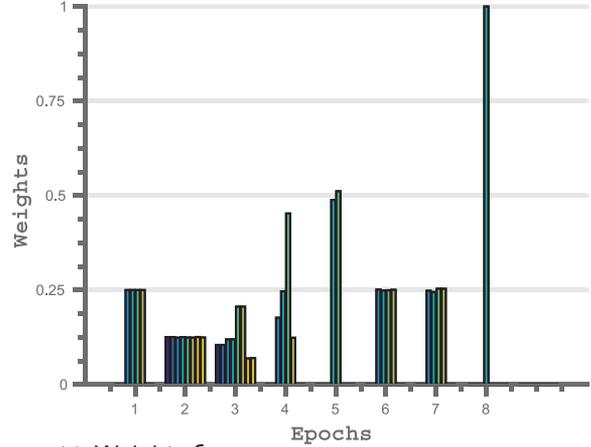
Table 7 compares DA-BSP with BSP-uni at different epochs of planning and inferences for planning horizons of $L = \{2, 4\}$. Here, the DA-BSP path is considered. Recall that ξ_{ca} stands for the weight of the component corresponding to the ground truth. For example, for $E = 2$, DA-BSP inference results in eight modes arising from seven other observations that alias the ground truth. Subsequently, $\xi_{ca} = 0.12$. In the case of BSP-uni, the metric ξ_{ca} measures how many times the correct association was made. Thus, the table shows that for all random runs there are instances where BSP-uni fails, owing to catastrophically bad data association. For example, at $E = 7$, where $\xi_{ca} = 0$, the robot *always* infers itself to be at a wrong place.

Table 8 shows how BSP-uni copes with uncertain data association. As expected, for a randomization value of $\epsilon = 0.3$, the overall correctness of the association is very low; e.g., at the epoch $E = 7$, $\xi_{ca} = 0.2$, meaning that wrong and possibly catastrophic inference will be made with 80% probability.

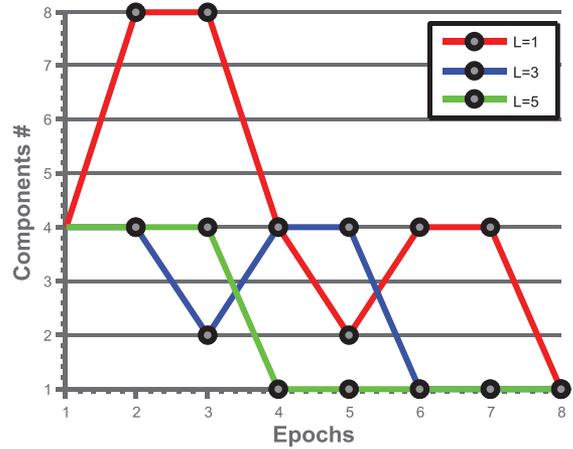
7. Conclusions and future work

7.1. Conclusions

We presented a unified framework for robust perception in planning as well as inference. State-of-the-art belief space planning approaches typically consider data association to



(a) Weights ξ .



(b) Components of Gaussian mixture model.

Fig. 18. (a) Evolution of weights of the components in the Gaussian mixture model after inference for $L = 2$. (b) Average number of components in belief mixtures for different planning horizons.

be given and perfect. However, such an assumption is less appropriate in the presence of localization uncertainty while operating in ambiguous environments, where two scenes could be similar in appearance when observed from appropriate viewpoints. In contrast with such state of the art, here we developed a DA-BSP approach that relaxes the data association assumption. In the context of a passive approach with the observations provided, the DA-BSP results in more robust inference. Conversely, in the context of an active approach, where planning needs to consider all possible future observations, this provides *better* action selection, such that catastrophically bad inferences and (if possible) actions leading to ambiguities are avoided. DA-BSP considers data association in a principled rigorous way with belief space planning. It is a more faithful representation of an aliased environment, since the number of components can increase as well as decrease. Although this increases the computational burden of planning, it is both necessary for ambiguous environments and still practically applicable, as shown through numerous experiments in both a realistic

Table 7. 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 are denoted t , while \tilde{m} denotes average modes. ξ_{ca} measures the level of aliasing, whereas data association is a binary variable denoting correct or wrong association.

Algorithm	Epoch	$L = 2$		$L = 4$		Inference	
		t (s)	(ξ_{ca}, \tilde{m})	t (s)	(ξ_{ca}, \tilde{m})	t (s)	(ξ_{ca}, \tilde{m})
DA-BSP	2	293.45	(0.13,8)	733.67	(0.49,2)	29.40	(0.12,8)
	3	262.37	(0.25,4)	557.57	(0.25,4)	26.80	(0.12,8)
	5	10.05	(0.25,4)	115.95	(1,1)	2.40	(0.26,4)
	7	2.47	(1,1)	2.57	(1,1)	1.46	(1,1)
BSP-uni		$t(s)$	η	t (s)	η	t (s)	η
	2	7.04	1	18.96	1	4.17	1
	3	1.23	1	2.20	0	0.77	0
	5	1.04	0	1.90	0	0.56	1
	7	0.47	0	0.50	0	0.46	0

Gazebo simulation and real experiments with the Pioneer robot platform. Moreover, DA-BSP degenerates to usual approaches in the presence of helpful assumptions, such as under very small localization uncertainty and under lack of perceptual ambiguities in the environment. In other words, DA-BSP is a rigorous holistic approach to consider data association in the context of belief space planning.

7.2. Future work

When DA-BSP is applied in the context of non-myopic planning, parametric solutions, such as incorporating all the associations explicitly, suffer from a scalability problem, owing to exponential blow up of the components of the beliefs. It was argued in Section 5.4 that components in DA-BSP arise out of data association ambiguities, where the observations that are not aliased drastically get increasingly smaller weights. As shown in the experiments, a short planning horizon is therefore often sufficient to collapse a highly multimodal belief into one with a very small number of components. It is interesting to note that this very same reasoning could also be harnessed to represent the resulting multimodal belief in a non-parametric fashion, such as through a Gibbs sampling approach, or a more efficient approximate method, harnessing Dirichlet processes. It was shown recently by Fourie et al. (2016), albeit in a passive setting, that such a non-parametric representation can tackle realistic data association issues in the context of SLAM. Thus, comparison of the proposed DA-BSP with those passive non-parametric approaches is a lucrative research direction, while another worthy pursuit is to obtain a non-parametric version of the proposed DA-BSP.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: This work was supported by the Technion Autonomous Systems Program and the Israel Science Foundation.

Table 8. BSP-uni in five different runs. BSP-uni can be seen as a very drastic pruning, where data association may or may not be correct. Randomization is $\epsilon = 0.3$. This is seen from the ξ_{ca} values (for five random runs). Standard deviations are given within the parentheses.

Algorithm	Epoch	Inference t (s)	ξ_{ca}
BSP-uni	2	4.10 (0.42)	0.60 (0.54)
	3	0.80 (0.12)	0.60 (0.54)
	5	0.53 (0.14)	0.80 (0.44)
	7	0.42 (0.09)	0.20 (0.44)

ORCID iD

Shashank Patha  <http://orcid.org/0000-0002-2888-6750>

Notes

1. The opensource code will be provided.

References

- Agarwal S, Tamjidi A, and Chakravorty S (2015) Motion planning in non-Gaussian belief spaces for mobile robots. *arXiv:1511.04634*.
- Agha-Mohammadi AA, Chakravorty S, and Amato NM (2014) FIRM: Sampling-based feedback motion planning under motion uncertainty and imperfect measurements. *International Journal of Robotics Research* 33(2): 268–304.
- Atanov N, Sankaran B, Ny J, et al. (2014) Nonmyopic view planning for active object classification and pose estimation. *IEEE Transactions on Robotics* 30: 1078–1090.
- Bar-Shalom Y, Li XR, and Kirubarajan T (2004) *Estimation with applications to tracking and navigation: Theory algorithms and software*. New York, NY: John Wiley & Sons.
- Bry A and Roy N (2011) Rapidly-exploring random belief trees for motion planning under uncertainty. In: *IEEE international conference on robotics and automation (ICRA)*, Shanghai, China, 9–13 May 2011, pp. 723–730. Piscataway, NJ: IEEE.
- Carlone L, Censi A, and Dellaert F (2014) Selecting good measurements via ℓ_1 relaxation: A convex approach for robust estimation over graphs. In: *IEEE/RSJ international conference*

- on intelligent robots and systems (IROS), Chicago, IL, 14–18 September 2014, pp. 2667–2674. Piscataway, NJ: IEEE.
- Chaves SM, Kim A, and Eustice RM (2014) Opportunistic sampling-based planning for active visual SLAM. In: *IEEE/RSJ international conference on intelligent robots and systems (IROS)*, Chicago, IL, 14–18 September 2014, pp. 3073–3080. Piscataway, NJ: IEEE.
- Chaves SM, Walls JM, Galceran E, et al. (2015) Risk aversion in belief-space planning under measurement acquisition uncertainty. In: *IEEE/RSJ international conference on intelligent robots and systems (IROS)*, Hamburg, Germany, 28 September–2 October 2015, pp. 2079–2086. Piscataway, NJ: IEEE.
- Fourie D, Leonard J, and Kaess M (2016) A nonparametric belief solution to the Bayes tree. In: *IEEE/RSJ international conference on intelligent robots and systems (IROS)*, Daejeon, South Korea, 9–14 October 2016, pp. 2189–2196. Piscataway, NJ: IEEE.
- Indelman V, Carlone L, and Dellaert F (2015) Planning in the continuous domain: A generalized belief space approach for autonomous navigation in unknown environments. *International Journal of Robotics Research* 34(7): 849–882.
- Indelman V, Michael N, and Dellaert F (2014a) Incremental distributed robust inference from arbitrary robot poses via EM and model selection. In: *RSS workshop on distributed control and estimation for robotic vehicle networks*, Berkeley, CA, 12 July 2014. Non-archived workshop proceedings.
- Indelman V, Nelson E, Dong J, et al. (2016) Incremental distributed inference from arbitrary poses and unknown data association: Using collaborating robots to establish a common reference. *IEEE Control Systems Magazine (CSM)* 36(2): 41–74.
- Indelman V, Nelson E, Michael N, et al. (2014b) Multi-robot pose graph localization and data association from unknown initial relative poses via expectation maximization. In: *IEEE international conference on robotics and automation (ICRA)*, Hong Kong, China, 31 May–7 June 2014, pp. 593–600. Piscataway, NJ: IEEE.
- Kaelbling LP, Littman ML, and Cassandra AR (1998) Planning and acting in partially observable stochastic domains. *Artificial Intelligence* 101(1): 99–134.
- Kaess M, Johannsson H, Roberts R, et al. (2012) iSAM2: Incremental smoothing and mapping using the Bayes tree. *International Journal of Robotics Research* 31: 217–236.
- Kim A and Eustice RM (2014) Active visual SLAM for robotic area coverage: Theory and experiment. *International Journal of Robotics Research* 34(4–5): 457–475.
- Kurniawati H, Hsu D, and Lee WS (2008) SARSOP: Efficient point-based POMDP planning by approximating optimally reachable belief spaces. In: *Robotics: Science and systems IV* (eds. O Brock, J Trinkle and F Ramos), 25–28 June 2008, Zurich, Switzerland. Cambridge, MA: MIT Press.
- Lauri M, Atanasov NA, Pappas G, et al. (2015) Active object recognition via Monte Carlo tree search. In: *Workshop on beyond geometric constraints at the international conference on robotics and automation (ICRA)*, Seattle, WA, 26–30 May 2015. Piscataway, NJ: IEEE.
- Madani O, Hanks S, and Condon A (1999) On the undecidability of probabilistic planning and infinite-horizon partially observable Markov decision problems. In: *AAAI '99/IAAI '99 proceedings of the sixteenth national conference on Artificial intelligence and the eleventh innovative applications of artificial intelligence conference innovative applications of artificial intelligence*, Orlando, FL, 18–22 July 1999, pp. 541–548. Menlo Park, CA: American Association for Artificial Intelligence.
- Olson E (2011) AprilTag: A robust and flexible visual fiducial system. In: *Proceedings of the IEEE international conference on robotics and automation (ICRA)*, Shanghai, China, 9–13 May 2011, pp. 3400–3407. Piscataway, NJ: IEEE.
- Olson E and Agarwal P (2013) Inference on networks of mixtures for robust robot mapping. *International Journal of Robotics Research* 32(7): 826–840.
- Papadimitriou C and Tsitsiklis J (1987) The complexity of Markov decision processes. *Mathematics of Operations Research* 12(3): 441–450.
- Pathak S, Thomas A, and Indelman V (2017) Nonmyopic data association aware belief space planning for robust active perception. In: *IEEE international conference on robotics and automation (ICRA)*, Singapore, 29 May–3 June 2017, pp. 4487–4494. Piscataway, NJ: IEEE.
- Patil S, Kahn G, Laskey M, et al. (2014) Scaling up Gaussian belief space planning through covariance-free trajectory optimization and automatic differentiation. In: Akin H, Amato N, Isler V, et al. (eds.) *Algorithmic Foundations of Robotics XI*. (Springer Tracts in Advanced Robotics, vol 107). Cham: Springer, pp. 515–533.
- Patten T, Zillich M, Fitch R, et al. (2016) Viewpoint evaluation for online 3-D active object classification. *IEEE Robotics and Automation Letters* 1(1): 73–81.
- Pfingsthorn M and Birk A (2016) Generalized graph SLAM: Solving local and global ambiguities through multimodal and hyperedge constraints. *International Journal of Robotics Research* 35(6): 601–630.
- Pineau J, Gordon GJ, and Thrun S (2006) Anytime point-based approximations for large POMDPs. *Journal of Artificial Intelligence Research* 27: 335–380.
- Platt R, Kaelbling L, Lozano-Perez T, et al. (2011) Efficient planning in non-Gaussian belief spaces and its application to robot grasping. In: Christensen H and Khatib O (eds.) *Robotics Research*. (Springer Tracts in Advanced Robotics, vol. 100). Cham: Springer, pp. 253–269.
- Prentice S and Roy N (2009) The belief roadmap: Efficient planning in belief space by factoring the covariance. *International Journal of Robotics Research* 28(11–12): 1448–1465.
- Sankaran B, Bohg J, Ratliff N, et al. (2015) Policy learning with hypothesis based local action selection. *arXiv:1503.06375*.
- Stachniss C, Grisetti G, and Burgard W (2005) Information gain-based exploration using Rao-Blackwellized particle filters. In: *Robotics: Science and systems (RSS)*, Cambridge, MA, 8–11 Jun 2005, pp. 65–72. Cambridge, MA: MIT Press.
- Sunderhauf N and Protzel P (2012) Towards a robust back-end for pose graph SLAM. In: *IEEE International conference on robotics and automation (ICRA)*, Saint Paul, MN, 14–18 May 2012, pp. 1254–1261. Piscataway, NJ: IEEE.
- Van den Berg J, Patil S, and Alterovitz R (2012) Motion planning under uncertainty using iterative local optimization in belief space. *International Journal of Robotics Research* 31(11): 1263–1278.

Walls JM, Chaves SM, Galceran E, et al. (2015) Belief space planning for underwater cooperative localization. In: *IEEE/RSJ international conference on intelligent robots and systems (IROS)*, Hamburg, Germany, 28 September–2 October 2015, pp. 2264–2271. Piscataway, NJ: IEEE.

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