

Incorporating Data Association Within Belief Space Planning For Robust Autonomous Navigation

Antony Thomas

Incorporating Data Association Within Belief Space Planning For Robust Autonomous Navigation

Research Thesis

Submitted in partial fulfillment of the requirements
for the degree of Master of Science in Aerospace Engineering

Antony Thomas

Submitted to the Senate
of the Technion — Israel Institute of Technology
Shevat 5777 Haifa February 2017

This research was carried out under the supervision of Assistant Prof. Vadim Indelman, in the Faculty of Aerospace.

Acknowledgements

I would like to express my sincere gratitude to my advisor Prof. Vadim Indelman for accepting me as his student and providing continuous support throughout my Master's study. He was a constant source of motivation and his guidance helped me throughout this research work. Passion for research is something that I have learned from Vadim and I have also drawn a lot of inspiration from his work ethic. I could not have imagined having a better advisor and mentor for my Master's study.

I would also like to thank my collaborator, Dr. Shashank Pathak, for all his help, patience, knowledge and valuable inputs. He was always ready to share ideas and discuss them at length and would patiently answer all my questions. I wish him all the best for his future career.

My sincere thanks also goes to Asaf Feniger, our lab engineer, for all his help, and my other lab-mates, especially Tal Regev, who have helped me during the course of this work.

I would also like to sincerely thank all my friends here at the Technion, especially my friends from the *shabbat group* without whom my stay here would not have been possible.

Finally, I would also like to thank my parents and my brother for the infinite support, care and love.

The generous financial help of the Technion is gratefully acknowledged.

Contents

List of Figures

List of Tables

Abstract	1
Abbreviations	3
1 Introduction	5
1.1 Related Work	6
1.2 Contributions	7
1.3 Organization	8
2 Notations and Problem Formulation	9
3 DA-BSP: Data Association aware Belief Space Planning	13
3.1 Concept and Approach Overview	13
3.2 Myopic DA-BSP	15
3.2.1 Computing term (a): $\mathbb{P}(z_{k+1} \mathcal{H}_{k+1}^-)$	15
3.2.2 Computing term (b): $\mathbb{P}(X_{k+1} \mathcal{H}_{k+1}^-, z_{k+1})$	16
3.2.3 Summary thus far	17
3.2.4 Simulating Future Observations $\{z_{k+1}\}$ given $b[X_{k+1}^-]$	18
3.2.5 Computing Mixture of Posterior Beliefs $\sum_i \tilde{w}_i b[X_{k+1}^{i+}]$	18
3.2.6 Designing a Specific Cost Function	19
3.2.7 Formal Algorithm for DA-BSP	20
3.3 Non-Myopic multi-modal DA-BSP	20
3.3.1 Prior belief as a mixture of Gaussian	21
3.3.2 Non-myopic DA-BSP	23
3.3.3 Overall algorithm	25
3.3.4 Effect of reducing a mixture of belief	26
3.3.5 Full vs. Partial Disambiguation	27
3.3.6 Degenerate cases of DA-BSP	27
3.3.7 On correctness of non-myopic DA-BSP	28

4	Results and Discussions	31
4.1	Implementation of data-association aware BSP	31
4.2	Metrics for evaluating DA-BSP	31
4.3	Real-world application with explicit scenes - octagonal corridor	32
4.4	Highly-aliased simulated office scenario	38
5	Conclusions and Future Work	43
5.1	Future Work	43
6	Appendix	45
	Hebrew Abstract	i

List of Figures

1.1	(a) Generative graphical model. While standard BSP approaches typically assume data association (DA) is given and perfect, we incorporate data association aspects within BSP and thus capable of reasoning about ambiguity (e.g. perceptual aliasing) at a decision-making level. (b) Schematic representation of pose, scene and observation spaces. Scenes A_1 and A_3 when viewed from perspective x and x' respectively, produce the same nominal observation \hat{z} , giving rise to <i>perceptual aliasing</i>	6
3.1	GMM posterior $b[X_{k+1}]$ given $z_{k+1} \in \{z_{k+1}\}$. The prior has two equi-probable components while the posterior has different weights for the two components.	25
4.1	The 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 on a real setup or a simulated one.	32
4.2	(left) Apriltag is detected, indicated with green patch at the 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 not detected since the non-centrality of the tag makes the observation highly untrustworthy. (right) No Apriltag lies within the field of view of the camera.	33
4.3	Real-world experimental setup (best viewed in colour). The images show the actual third-person view at different locations of the robot, while the two dimensional figure, shows semantic knowledge of the environment the robot possesses. Its current belief is a 4-modal GMM with mean position depicted by x_{init} . Ground truth robot position is indicated with \otimes ; arrows indicate orientation (and not motion). The actual scenario is depicted through the laser-scan of the environment shown in the green colour. This map is for representative purposes only and is not available to the robot. The zoomed-in picture of an AprilTag is also shown.	34

4.4	Evolution of belief at an epoch $E = 3$ under DA-BSP. Here tags are represented by shapes $\{*, +, \square, \triangle, \nabla\}$ while ground truth robot position is indicated with \star . (a) The prior belief is multimodal with four distinct modes as shown with the coloured ellipses. (b) After incorporating the motion model, the propagated belief is similarly a multi-modal 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 of them to vanish and components reduce from 4 to 3). Also new components in the posterior may emerge (not the case here). Here, $L = 3$	35
4.5	(a)-(c) Evolution of inferred belief as <i>decision epoch</i> progresses with $L = 3$; epochs depicted are $\{4, 18, 24\}$. They depict evolution of inferred belief, for different planning algorithms, i.e. DA-BSP, BSP-uni and BSP-mul, respectively. GMM components and associated weights are designated with different colors. Ground truth robot position is indicated with \star . For clarity, the detected scene(s) are shown in different colour. In case of BSP-mul and BSP-uni, this particular instance of planning leads to catastrophically bad data association. (d) Evolution of GMM components weights during these epochs. Note that the number of components <i>increases</i> as well as <i>decreases</i> and eventually goes to 1. Here, planning horizon is $L = 3$. . .	36
4.6	Evolution of belief as <i>decision epoch</i> progresses during DA-BSP planning. Average number of components in the belief mixtures and the KL_u metric are depicted in <i>left</i> and <i>right</i> respectively.	37
4.7	Fig. 4.7a Two-floor aliased office environment in a Gazebo simulator. p_1 and p_2 denote the printers while 1 and 6 are the mean positions in each floor for the initial four-component GMM belief. Figures 4.7b till 4.7h show the mean positions (modes) of the robot for each step of the DA-BSP path. Green denotes the ground truth while yellow the aliasing position.	40
4.8	(<i>left</i>) Evolution of weights of the components in the GMM after inference for $L = 2$. (<i>right</i>) Average number of components in the belief mixtures for different planning horizon. . . .	41
4.9	(<i>left</i>) Evolution of weights of the components of the belief when following the shortest path versus that following the DA-BSP path. (<i>right</i>) Laser scans at ground truth and aliased position (green and yellow respectively).	41
4.10	Prior belief is propagated according to the motion model. Within the subsequent propagated belief, there are perceptually aliased laser scans observed. Here, 2σ covariance is depicted with each ellipse.	41

List of Tables

4.1	Performing DA-BSP on a real corridor environment shown in the Fig. 4.3, with planning horizon $L = 4$. The times in seconds spent in planning and in inference is denoted by t , while \tilde{m} stands for average modes; refer Sec. 4.3.	34
4.2	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 \tilde{m} . DA signifies <i>correct</i> data association; refer Sec. 4.2. Values shown here are for average of 5 random runs.	37
4.3	Evaluating DA-BSP in several steps of planning, under different pruning thresholds of $\sigma = \{0.2, 0.05, 10^{-10}\}$. The times in seconds spent in planning is denoted by t , while η_{da} and \tilde{m} show weight of correct association and averaged number of modes, respectively. Values shown here are for average of 5 random runs while standard deviation is depicted within the parenthesis. Note that 5% threshold is sufficient in this case to perform equivalently with almost <i>unpruned</i> DA-BSP with $\sigma = 10^{-10}$	38
4.4	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 <i>plausible</i> ones. The times in seconds spent in planning is denoted by t , while average correct association is denoted by ξ_{ca} . Values shown here are for average of 5 random runs while standard deviation is depicted within the parenthesis.	38
4.5	Evaluating DA-BSP in several steps of planning and inference, for $L = 2$ and $L = 4$. The times in seconds spent in planning and in inference is denoted by t , while \tilde{m} denotes average modes. η_{da} measures the level of aliasing whereas DA is a binary variable denoting correct or wrong association.	42
4.6	BSP-uni in 5 different runs. BSP-uni can be seen as a very drastic pruning where data association may or may not be correct. This is seen from the ξ_{ca} values (for 5 random runs). Standard deviation are mentioned within the parenthesis.	42

Abstract

Belief space planning (BSP) and perception are fundamental problems in robotics and artificial intelligence, with applications including autonomous navigation and active SLAM. State-of-the-art BSP approaches assume that data association (DA), i.e. determining the correct correspondence between the observations and the landmarks, is given and perfect. However, real world environments are often ambiguous, which in the presence of different sources of uncertainty, make perception a challenging task. For example, an object might be similar in appearance from the current viewpoint to another object, while successfully matching images from two different but similar in appearance places (e.g. buildings that look alike) would incorrectly indicate the two places as one. An incorrect DA can lead to catastrophic results, e.g. a robot considering it is located in a wrong aliased corridor. Consequently, more advanced approaches, known as robust perception, are required. Yet, existing robust perception approaches focus on the passive case where robot actions are externally determined, while existing BSP methods assume data association to be given and perfect.

In this research we relax the above assumption and incorporate reasoning regarding DA aspects within BSP, while accounting for different sources of uncertainty (imperfect sensing, stochastic control, uncertain environment). We develop a data association aware belief space planning (DA-BSP) approach that explicitly reasons about DA within belief evolution while considering non-myopic planning and multi-modal beliefs represented by Gaussian Mixture Models (GMM). We envision such a framework to provide robust active perception and active disambiguation capabilities, in particular while operating in ambiguous and perceptually aliased environments. The approach is studied and proven effective using real-world experiments and synthetic simulations, carried out at the Autonomous Navigation and Perception Lab at the Technion.

Abbreviations

BSP	:	Belief Space Planning
DA	:	Data Association
GMM	:	Gaussian Mixture Model
POMDP	:	Partially Observable Markov Decision Process
SLAM	:	Simultaneous Localisation and Mapping
DA-BSP	:	Data Association aware Belief Space Planning
GT-SAM	:	Georgia Tech-Smoothing and Mapping
ROS	:	Robot Operating System

Chapter 1

Introduction

Belief space planning (BSP) and decision-making under uncertainty are fundamental problems in robotics and artificial intelligence, with applications including autonomous navigation, object grasping and manipulation, active SLAM, and robotic surgery. In presence of uncertainty, such as in robot motion and sensing, the true state of variables of interest (e.g. robot poses), is unknown and can only be represented by a probability distribution over possible states, given available data. Planning and decision-making should be therefore performed over this distribution, the belief space, which can be inferred using probabilistic approaches based on incoming sensor observations and prior knowledge. The corresponding problem is an instantiation of a partially observable Markov decision problem (POMDP) [15], where, given an objective function, one aims to determine an optimal control policy as a function of belief evolution over application-dependent variables of interest.

However, state-of-the-art BSP approaches typically assume data association to be given and perfect (see Figure 1.1b), 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, that together with other sources of uncertainty, make perception a challenging task. As an example, matching images from two different but similar in appearance places, or attempting to recognise 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 as same, and incorrectly associating the observed object.

Considering data association to be solved and perfect within BSP can thus lead, in presence of ambiguity, to incorrect posterior beliefs and as a result, to sub-optimal actions which 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

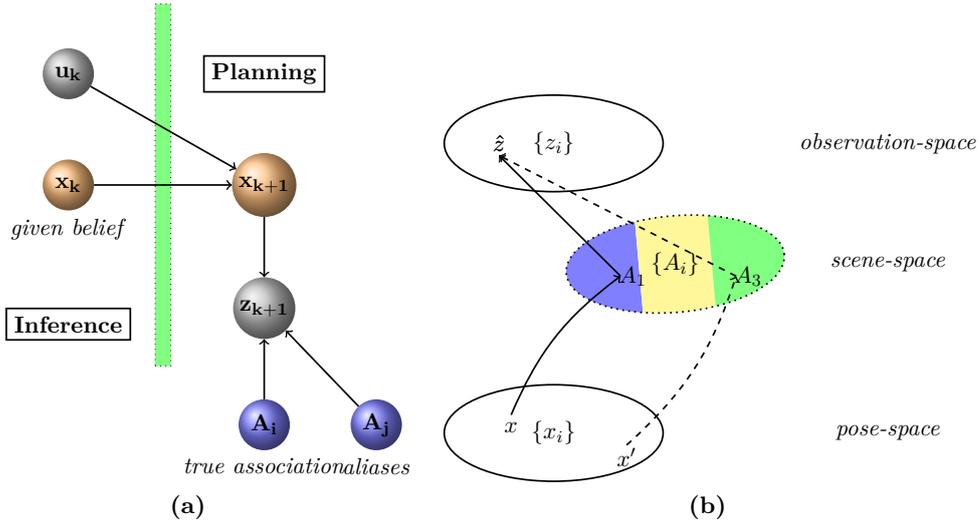


Figure 1.1: (a) Generative graphical model. While standard BSP approaches typically assume data association (DA) is given and perfect, we incorporate data association aspects within BSP and thus capable of reasoning about ambiguity (e.g. perceptual aliasing) at a decision-making level. (b) Schematic representation of pose, scene and observation spaces. Scenes A_1 and A_3 when viewed from perspective x and x' respectively, produce the same nominal observation \hat{z} , giving rise to *perceptual aliasing*.

a given database that the current observation (e.g. image or point-cloud) is successfully registered to. Similarly, one might reason probabilistically regarding perceptual aliasing, as in the first example above, 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 localised.

In this work we develop a general data association aware belief space planning (DA-BSP) framework capable of better handling complexities arising in real world, possibly perceptually aliased, scenarios. To that end, we rigorously incorporate reasoning about data association within belief space planning, while also considering other sources of uncertainty (motion, sensing and environment). In particular, we show 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) [21] for all but the smallest problems. The vast research area of approximate approaches (with reduced computational complexity) can be roughly segmented into point-based value iteration methods [18, 24], simulation based [28] and sampling based approaches [2, 6, 26], and direct trajectory optimization [11, 22, 30] methods. In all cases, finding the (locally) optimal actions involves evaluating a given objective function while considering future observations to be acquired as a result of each candidate 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 localised by making observations of known landmarks or beacons (see, e.g. [2, 26]), while assuming to correctly associate each future measurement with an appropriate landmark. Though reasonable in certain scenarios, such an assumption becomes unrealistic in the presence of perceptually aliased environments (two scenes that look alike) and localisation uncertainty, as in this work.

While belief space planning approaches typically assume the environment to be accurately known (e.g. a given map), recent works, including [8, 9, 11, 17, 31], 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 [9, 11, 17, 31] in this branch focused in particular on probabilistically modelling what future observations will be obtained given a candidate action. However, these approaches consider each such future observation to be correctly associated to an appropriate scene, and hence, assume data association to be given and perfect.

In the last few years, the SLAM research community has investigated approaches to be resilient to false data association (outliers) overlooked by front-end algorithms (e.g. image matching), see e.g. [7, 13, 14, 20, 29]. 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. In contrast, we consider a complimentary *active* framework and incorporate data association aspects within BSP.

Our approach is also tightly related with recent work on active hypothesis disambiguation in the context object detection and classification [3, 19, 23, 27, 32]. Given hypotheses regarding object class and pose, these approaches aim to find a sequence future viewpoints that will lead to disambiguation, i.e. identifying the correct hypothesis. However, these approaches assume the sensor is perfectly localized and thus the corresponding belief is only about the considered hypotheses.

Probably the closest work to our approach is by Agarwal et al. [1], where the authors also consider hypotheses due to ambiguous data association and develop a BSP approach for active disambiguation. However, in that work the authors only consider ambiguous data association within the prior belief, modelling it as mixture of Gaussians, and assume there indeed exists an action that can yield complete disambiguation. In contrast, our framework is more general since we additionally consider ambiguous data association within future belief (due to future observations) given candidate action(s) and do not assume there is necessarily a fully-disambiguating action.

1.2 Contributions

As mentioned in chapter 1 we incorporate reasoning regarding data association (DA) in BSP while accounting for different sources of uncertainty (imperfect sensing, stochastic control, uncertain environment). As such this framework provides robust active perception and active disambiguation capabilities, in particular while operating in ambiguous

and perceptually aliased environments.

Main contributions of this thesis are as follows: We develop a unified framework for data association aware belief space planning (DA-BSP) in both active and passive context. Here, the components of our belief may both increase and decrease, thereby modeling the perceptually aliased environment more faithfully. Additionally, it does not require a fully disambiguating unique observation. We extend DA-BSP by considering prior belief as non-Gaussian as well as by considering planning for several lookahead steps. We show how under helpful assumptions this general approach degenerates to known BSP approaches. We present complexity analysis of such an algorithm as well as comment on its correctness. Finally, we analyze key aspects arising due to explicitly considering data association aspects within BSP in a realistic synthetic simulation and in a real robotics scenario using a Pioneer robot.

1.3 Organization

The rest of this thesis is organized as follows.

1. We formulate the problem and introduce the notations in **chapter 2**.
2. In **chapter 3** we provide the mathematical build-up for our method. This is done for both myopic and non-myopic planning.
3. We discuss the key aspects of our method using a synthetic simulation and real-world experiment in **chapter 4**.
4. **chapter 5** provides a conclusion and discusses potential future work.

Chapter 2

Notations and Problem Formulation

Planning and decision making under uncertainty are fundamental problems in the area of autonomous navigation. In the presence uncertainty arising, for example, due to stochastic robot motion and imperfect sensing, the true state over variables of interest, such as robot poses, is unknown and can be only represented by a probability distribution function (pdf), also known as the *belief*. The corresponding planning problem is known as belief space planning (BSP), which is an instantiation of a partially observable Markov decision problem (POMDP) [15]. Such a framework is applicable to numerous problem domains and applications such as active simultaneous localization and mapping (SLAM), active sensing, informative planning, and additional variants of autonomous navigation.

The Simultaneous Localisation and Mapping (SLAM) [15]) problem asks if it is possible for a mobile robot to be placed at an unknown location in an unknown environment and for the robot to incrementally build a consistent map of this environment while simultaneously determining its location within this map [10]. However in this work we assume a known environment and use this knowledge in our planning and inference. In the remaining part of this chapter we first motivate the problem and then go over the mathematical notations and models that will be used to formulate our problem. These models are general to any planning problem and we introduce our concept of incorporating reasoning regarding DA within BSP only in the next chapter.

Consider a robot, uncertain about its pose, operating in a partially known or pre-mapped environment. The robot takes observations of different scenes or objects in the environment and uses these observations to infer random variables of interest which are application-dependent. Thus, in localisation, 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 1.1. As can be seen, it involves three spaces: *pose-space*, *scene-space* and *observation-space*. *Pose-space* involves all possible perspectives a robot can take with respect to a given world model and in the

context of task at hand.

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

$$\mathcal{H}_k \doteq \{u_{0:k-1}, Z_{0:k}\} \quad , \quad \mathcal{H}_{k+1}^- \doteq \mathcal{H}_k \cup \{u_k\}. \quad (2.1)$$

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

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

We shall consider such an observation as the model:

$$z_k = h(x_k, A_i) + v_k \quad , \quad v_k \sim \mathcal{N}(0, \Sigma_v), \quad (2.2)$$

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 as *nominal* observation \hat{z} .

For example, in 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 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 also in case where the environment model is given but uncertain, and when this model is unknown a priori and instead is constructed on-line within SLAM framework.

We also consider a standard motion model with Gaussian noise,

$$x_{i+1} = f(x_i, u_i) + w_i \quad , \quad w_i \sim \mathcal{N}(0, \Sigma_w) \quad (2.3)$$

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 $\mathbb{P}(x_0)$ and motion and observation models, the joint posterior pdf at

the current time k can be written as

$$\mathbb{P}(X_k|\mathcal{H}) = \mathbb{P}(x_0) \prod_{i=1}^k \mathbb{P}(x_i|x_{i-1}, u_{i-1}) \mathbb{P}(Z_i|x_i, A_i). \quad (2.4)$$

This pdf is thus a Gaussian $\mathbb{P}(X_k|\mathcal{H}_k) = \mathcal{N}(\hat{X}_k, \Sigma_k)$ with mean \hat{X}_k and covariance Σ_k that can be efficiently calculated via maximum a posteriori (MAP) inference, see e.g. [16].

It is important to note that the underlying assumption in factorisation (2.4) 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 in the sequel.

Given the posterior (2.4) at the current time k , one can reason about the robot's best future actions that would minimise (or maximise) a certain objective function. Such a function, for a single look ahead step, is given by

$$J(u_k) = \mathbb{E}_{z_{k+1}} \{c(\mathbb{P}(X_{k+1}|\mathcal{H}_{k+1}^-, z_{k+1}))\}, \quad (2.5)$$

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

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

$$b[X_{k+1}] \doteq \mathbb{P}(X_{k+1}|\mathcal{H}_{k+1}^-, z_{k+1}). \quad (2.6)$$

Note that, according to Eq. (2.5), we need to calculate the posterior belief (2.6) for *each* possible value of z_{k+1} .

Similarly, we define the propagated joint belief as

$$b[X_{k+1}^-] \doteq \mathbb{P}(X_{k+1}|\mathcal{H}_{k+1}^-) = \mathbb{P}(X_k|\mathcal{H}_k) \mathbb{P}(x_{k+1}|x_k, u_k), \quad (2.7)$$

from which the marginal belief over the future pose x_{k+1} can be calculated as $b[x_{k+1}^-] \doteq \int_{-x_{k+1}} b[X_{k+1}^-]$.

As earlier, if data association is assumed given and perfect as commonly done in BSP, then one can consider for each specific value of z_{k+1} the corresponding observed scene A_i , and express the posterior (2.6) as

$$b[X_{k+1}] = \eta \mathbb{P}(X_k|\mathcal{H}_k) \mathbb{P}(x_{k+1}|x_k, u_k) \mathbb{P}(z_{k+1}|x_{k+1}, A_i), \quad (2.8)$$

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 (2.5) can be now 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

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

Assuming data association to be given and perfect simplifies greatly the above formulation. Yet, in practice, determining data association reliably is often a non trivial task by itself, especially when operating in perceptually aliased environments. An incorrect data association (wrong scene A_i in Eq. (2.8)) can lead to catastrophic results, see, e.g. [12–14]. In this work we relax this restricting assumption and rigorously incorporate data association aspects within belief space planning.

Chapter 3

DA-BSP: Data Association aware Belief Space Planning

3.1 Concept and Approach Overview

We begin by making the following observations that will be transformed into a rigorous mathematical framework in the following sections.

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 Eq. (2.5). To see that, we write the expectation operator explicitly which transforms Eq. (2.5) to

$$J(u_k) \doteq \int_{z_{k+1}} \overbrace{\mathbb{P}(z_{k+1} | \mathcal{H}_{k+1}^-)}^{(a)} c \left(\overbrace{\mathbb{P}(X_{k+1} | \mathcal{H}_{k+1}^-, z_{k+1})}^{(b)} \right) \quad (3.1)$$

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

Considering data association is solved and perfect then means we can correctly associate each possible measurement z_{k+1} with the corresponding scene A_i it captures, as in Eq. (2.8).

Yet, it is unknown from what future robot pose x_{k+1} the actual observation z_{k+1} will be acquired, since the *actual* robot pose x_k at time k is unknown and the control is stochastic. Indeed, as a result of action u_k , the robot actual (true) pose x_{k+1} can be anywhere within the propagated belief $b[x_{k+1}^-]$.

In 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, 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 it describes, i.e. write the appropriate measurement likelihood term in the posterior (2.4).

A similar situation, however, arises also in our case: while the probability of acquiring a specific observation z_{k+1} is represented by the term (a) in Eq. (3.1), the posterior in the term (b) is *conditioned on this specific* observation z_{k+1} . As such, evaluating the posterior given z_{k+1} involves inference, as if that observation was actually acquired. Thus, also here data association needs to be resolved or to be assumed given.

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 multiple different poses of viewing different scenes. In belief space planning (BSP) framework, such a data association is assumed to be solved. In other words, if \mathcal{A} represents the total space of scenes (or real world) from where *all* observations $\{z\}$ are made and $\{A_N\}$ be the partitioning of this scene space, then BSP 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 data association is solved, and instead reason about possible scenes or objects that the future observation z_{k+1} could be generated from, see Figures 1.1b and 1.1. 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 (2.2). However, in case of perceptually aliased environments, there could be also 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 $\{A_i\}$ and due to localisation uncertainty determining which association is the correct one is not trivial. As we show in the sequel, in these cases the posterior $b[X_{k+1}]$ (term (b) in Eq. (3.1)) becomes a Gaussian mixture with appropriate weights that we rigorously compute.

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

Perceptual aliasing Intuitively speaking, perceptual aliasing occurs when an object different from the actual one, produces the same observation and thereby seeks to provide 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 thesis 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 1.1. A probabilistic perceptual aliasing is a more general form of aliasing, which 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 .

3.2 Myopic DA-BSP

In this section we develop our approach for data association aware belief space planning, developing expressions for calculating each of the two terms (a) and (b) in Eq. (3.1) without assuming data association is solved, and discussing additional aspects. For simplicity, in this section we assume myopic planning i.e., with the planning horizon $L = 1$. For convenience, we specify the corresponding expressions again:

$$(a) : \mathbb{P}(z_{k+1} | \mathcal{H}_{k+1}^-) \quad , \quad (b) : \mathbb{P}(X_{k+1} | \mathcal{H}_{k+1}^-, z_{k+1}) \quad (3.2)$$

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.

3.2.1 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

$$\mathbb{P}(z_{k+1} | \mathcal{H}_{k+1}^-) \equiv \sum_i \int_x \mathbb{P}(z_{k+1}, x, A_i | \mathcal{H}_{k+1}^-) \doteq \sum_i w_i. \quad (3.3)$$

As seen from the above equation, to calculate the likelihood of obtaining some observation z_{k+1} , we consider separately, for each scene $A_i \in \{A_{\mathbb{N}}\}$, the likelihood that this observation was generated by scene A_i . This probability is captured for each scene A_i by a corresponding weight w_i ; these weights are then summed to get the actual likelihood of observation z_{k+1} . As will be seen below, these weights naturally account for perceptual aliasing aspects for each considered z_{k+1} .

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

$$\sum_i \int_x \mathbb{P}(z_{k+1} | x, A_i, \mathcal{H}_{k+1}^-) \mathbb{P}(A_i, x | \mathcal{H}_{k+1}^-) \quad (3.4)$$

However, since this integral could be over any arbitrary total distribution of x , we can use the propagated belief $b[x_{k+1}^-]$, see Eq. (2.7), to compute it as:

$$\sum_i \int_x \mathbb{P}(z_{k+1} | x, A_i, \mathcal{H}_{k+1}^-) \mathbb{P}(A_i | \mathcal{H}_{k+1}^-, x) b[x_{k+1}^- = x]. \quad (3.5)$$

Thus,

$$w_i \doteq \int_x \mathbb{P}(z_{k+1}|x, A_i, \mathcal{H}_{k+1}^-) \mathbb{P}(A_i|\mathcal{H}_{k+1}^-, x) b[x_{k+1}^- = x]. \quad (3.6)$$

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 of scene A_i to be observed from viewpoint x . In other words, this scenario-dependent term encodes from what viewpoints each scene A_i is observable and could also model occlusion and additional aspects. As such, this term can be determined given a model of the environment and thus, in this work, we consider this term to be given.

The weights w_i (3.6) naturally capture *perceptual aliasing* aspects discussed in Section 3.1: consider some observation z_{k+1} and the corresponding generative model $z_{k+1} = h(x^{tr}, A^{tr}) + v$ with appropriate unknown *true* robot pose x^{tr} and scene $A^{tr} \in \{A_{\mathbb{N}}\}$. Clearly, the measurement likelihood $\mathbb{P}(z_{k+1} | x, A_i, \mathcal{H}_{k+1}^-)$ will be high when evaluated for $A_i = A^{tr}$ and in vicinity of x^{tr} . Note that we will necessarily consider such a case, since according to Eq. (3.3) we separately consider each scene A_i in $\{A_{\mathbb{N}}\}$, and, given A_i , we reason about all poses x in Eq. (3.6). In case of perceptual aliasing, however, there will be also another scene(s) A_j which could generate the same observation z_{k+1} from appropriate robot pose x' , i.e. $\{A_i, A_j\}_{\text{alias}}$. Thus, the corresponding measurement likelihood term to A_j will also be high for x' .

However, the actual value of w_i (for each $A_i \in \{A_{\mathbb{N}}\}$) depends, in addition to the measurement likelihood, also on the mentioned-above event likelihood and on the belief $b[x_{k+1}^-]$, with the latter weighting the probability of each considered robot pose. This correctly captures the intuition that those observations z with low-probability poses $b[x_{k+1}^- = x^{tr}]$ will be unlikely to be actually acquired, leading to low value of w_i with $A_i = A^{tr}$. However, the likelihood term (3.3) could still go up in case of perceptual aliasing, where the aliased scene A_j generates a similar observation to z_{k+1} from viewpoint x' with latter being more probable, i.e. high probability $b[x_{k+1}^- = x']$.

In practice, calculating the integral in Eq. (3.9) can be done 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 event likelihood is non-zero i.e. $\mathbb{P}(A_i | x, \mathcal{H}) > 0$. In absence of such assumptions, in general, the integral in Eq. (3.9) should be computed numerically. Since in practical applications $\mathbb{P}(A_i | x, \mathcal{H})$ is sparse w.r.t. 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.

3.2.2 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 measurement model (2.2). Hence, also here, we consider all possible such scenes and weight them accordingly, with weights \tilde{w}_i representing the probability of each scene A_i to have generated the observation z_{k+1} .

As will be seen next, in both terms (a) and (b) the same weights are obtained, however only in the latter case the weights are to be normalised such that $\sum_i \tilde{w}_i = 1$.

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

$$\sum_i \mathbb{P}(X_{k+1}, A_i | \mathcal{H}_{k+1}^-, z_{k+1}) = \sum_i \mathbb{P}(X_{k+1} | \mathcal{H}_{k+1}^-, z_{k+1}, A_i) \cdot \mathbb{P}(A_i | \mathcal{H}_{k+1}^-, z_{k+1}). \quad (3.7)$$

Here, the first term is the posterior belief conditioned on observations, history as well as a candidate scene A_i that supposedly generated the observation z_{k+1} . We discuss how this term can be calculated in Section 3.2.5.

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 which generated the observation z_{k+1} . As will be seen now, this term is actually the normalised weight w_i from Section (3.2.1). Marginalising over all robot poses and applying Bayes rule yields

$$\mathbb{P}(A_i | \mathcal{H}_{k+1}^-, z_{k+1}) = \int_x \mathbb{P}(A_i, x | \mathcal{H}_{k+1}^-, z_{k+1}) = \eta \int_x \mathbb{P}(z_{k+1} | A_i, x, \mathcal{H}_{k+1}^-) \mathbb{P}(A_i, x | \mathcal{H}_{k+1}^-), \quad (3.8)$$

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

Similarly to the derivation in Section (3.2.1), 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:

$$\eta \int_x \mathbb{P}(z_{k+1} | A_i, x, \mathcal{H}_{k+1}^-) \mathbb{P}(A_i | x, \mathcal{H}_{k+1}^-) b[x_{k+1}^- = x]. \quad (3.9)$$

As seen, the same expression is obtained as in Eq. (3.6), except for the normalisation constant η . Hence,

$$\mathbb{P}(A_i | z_{k+1}, \mathcal{H}_{k+1}^-) = \eta w_i \doteq \tilde{w}_i. \quad (3.10)$$

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

3.2.3 Summary thus far

To summarise the discussion thus far, we have shown that the objective function (3.1) can be re-written as

$$J(u_k) = \int_{z_{k+1}} \left(\sum_i w_i \right) \cdot c \left(\sum_i \tilde{w}_i b[X_{k+1}^{i+}] \right), \quad (3.11)$$

with the posterior given scene A_i defined as

$$b[X_{k+1}^{i+}] \doteq \mathbb{P}(X_{k+1} \mid \mathcal{H}_{k+1}^-, z_{k+1}, A_i). \quad (3.12)$$

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

The above insights also apply to the unnormalised 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 3.2.1, this likelihood is calculated by summing over all such weights ($\sum_i w_i$), with each weight properly capturing the likelihood of a measurement z_{k+1} to be generated by scene A_i while taking into account how probable is the corresponding robot pose x given $b[x_{k+1}^-]$. 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 incorporating data association within belief space planning leads to Eq. (3.11), we now proceed with the exposition of our approach.

3.2.4 Simulating Future Observations $\{z_{k+1}\}$ given $b[X_{k+1}^-]$

Calculating the objective function (3.11) for each candidate action u_k involves considering all possible realisations of z_{k+1} . One approach to perform this in practice, is to simulate future observations $\{z_{k+1}\}$ given propagated belief $b[X_{k+1}^-]$, scenes $\{A_{\mathbb{N}}\}$ and observation model (2.2). One can then evaluate Eq. (3.11) considering all observations in the set $\{z_{k+1}\}$.

We now briefly describe how this concept can be realised. 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 $\mathbb{P}(A_i \mid \mathcal{H}_k, x)$. Together, x and A_i are then used to generate nominal $\hat{z} = h(x, A_i)$ and noise-corrupted observations $\{z\}$ according to observation model (2.2): $z = h(x, A_i) + v$. 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), it is unknown to our algorithm, i.e. while evaluating Eq. (3.11).

3.2.5 Computing Mixture of Posterior Beliefs $\sum_i \tilde{w}_i b[X_{k+1}^{i+}]$

As seen from Eq. (3.11), reasoning about data association aspects resulted in a mixture of posteriors 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 Section 3.2.4; however, one could also consider treating future observation z_{k+1} as a random variable [11, 25, 30].

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 is a Gaussian $b[X_k] = \mathcal{N}(\hat{X}_k, \Sigma_k)$. However, our approach could be applied also to more general cases (e.g. 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 pdf can be written as

$$b[X_{k+1}^{i+}] \propto b[X_k] \mathbb{P}(x_{k+1} | x_k, u_k) \mathbb{P}(z_{k+1} | x_{k+1}, A_i). \quad (3.13)$$

It is then not difficult to show that the above belief is a Gaussian $b[X_{k+1}^{i+}] = \mathcal{N}(\hat{X}_{k+1}^i, \Sigma_{k+1}^i)$ and to find its first two moments via MAP inference.

Obviously, the mixture of posterior beliefs in the cost $c(\cdot)$ from Eq. (3.11) is now a mixture of Gaussians:

$$\sum_i \tilde{w}_i b[X_{k+1}^{i+}] = \sum_i \tilde{w}_i \mathcal{N}(\hat{X}_{k+1}^i, \Sigma_{k+1}^i). \quad (3.14)$$

3.2.6 Designing a Specific Cost Function

The treatment so far has been agnostic to the structure of the cost function $c(\cdot)$. Recalling Eq. (3.11) we see that the belief over which the cost function is defined, is multimodal in general. Standard cost functions in literature, typically include terms such as control usage c_u , distance to goal c_G and uncertainty c_Σ , see e.g. [11, 30]. 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 in usual BSP, the posterior belief in our case is multimodal and represented as mixture of Gaussians $\sum_i \tilde{w}_i \mathcal{N}(\hat{X}_{k+1}^i, \Sigma_{k+1}^i)$, see Eq. (3.14), we could define several different cost structures depending on how we treat the different modes. Two particular such costs are taking the worst-case covariance among all covariances Σ_{k+1}^i in the mixture, e.g. $\Sigma = \max_i \{tr(\Sigma_i)\}$, or to collapse the mixture into a single Gaussian $\mathcal{N}(\cdot, \Sigma)$, see e.g. [5]. In both cases, we can define the cost due to uncertainty as $c_\Sigma = trace(\hat{\Sigma})$.

The cost due to ambiguity, c_w , should penalise ambiguities such as those arising out of perceptual aliasing. Here, we note that non-negligible weights w_i in Eq. (3.11) arise due to perceptual aliasing with respect to any scene A_i , whereas in case of no aliasing, all but one of these weights are zero. In most severe case of aliasing (all scenes or objects A_i are identical), all of these weights are comparable among each other. Thus we take Kullback-Leibler divergence $KL_u(\{\tilde{w}_i\})$ of these weights $\{\tilde{w}_i\}$ from a uniform

distribution to penalise higher aliasing, and define $c_w(\{\tilde{w}_i\}) \doteq \frac{1}{KL_u(\{\tilde{w}_i\})+\epsilon}$, where ϵ is a small number to avoid division-by-zero in case of extreme perceptual aliasing. With user-defined weights M_u, M_G, M_Σ and M_w , the overall cost then can be defined as a combination

$$c \doteq M_u c_u + M_G c_G + M_\Sigma c_\Sigma + M_w c_w, \quad (3.15)$$

3.2.7 Formal Algorithm for DA-BSP

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 summarised in Algorithm 3.2 and briefly described below.

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 3.2.4. The algorithm then calculates the contribution of each observation $z_{k+1} \in \{z_{k+1}\}$ to the objective function (3.11). In particular, on lines 8 and 15 we calculate the weights w_i and the posterior beliefs $b[X_{k+1+i}]$ for each $A_i \in \{A_{\mathbb{N}}\}$, respectively. These calculations are according to Sections 3.2.1 and 3.2.5. Then, after weight normalisation on line 13, we evaluate the cost $c(\cdot)$ (line 20) and use the accumulated unnormalised 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 21).

Finally, a few words about computational complexity. To see the relation of DA-BSP with respect to general POMDP, we analyse the discrete space case and show in Appendix B in supplementary material [4] that under the reasonable assumption that the cardinality of the scene space is often much less than the cardinality of the state space, DA-BSP does not introduce significant additional computational complexity.

3.3 Non-Myopic multi-modal DA-BSP

This section would generalize the DA-BSP that was developed in the previous Section 3.2. We will start with considering a prior which is non-Gaussian. In particular, we will assume our prior to be a mixture of Gaussians and then follow a similar approach to compute belief update and perform myopic planning as done earlier. Once this is done and we have an approach that takes in a GMM belief and updates to another GMM belief, we will present the most general DA-BSP in a non-myopic setting of several look-ahead steps of planning.

Algorithm 3.1 Myopic data association aware belief-space planning

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

```
1:  $b[X_{k+1}^-] \leftarrow b[X_k] \mathbb{P}(x_{k+1} | x_k, u_k)$ 
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:      $\triangleright$  compute weight, Eq. 3.6
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  $A_i$ , Sec. 3.2.5
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:   $c \leftarrow \text{CalcCost}(\{\tilde{w}_i\}, \{b[X_{k+1}^{i+}]\})$   $\triangleright$  Sec. 3.2.6
15:   $J \leftarrow J + w_s \cdot c$ 
16: end for
17: return  $J$ 
```

3.3.1 Prior belief as a mixture of Gaussian

Let us assume that prior is Gaussian mixture model. In other words, our belief at time k is a linear combination of $M_k \in \mathbb{N}$ Gaussians i.e.,

$$b[X_k] \doteq \mathbb{P}(X_k | \mathcal{H}_k^-, z_k) = \sum_{i=1}^{M_k} \xi_{k,i} \mathcal{N}(\hat{X}_{k,i}, \Sigma_{k,i}) \quad (3.16)$$

Since our motion model (see 2.3) is still a Gaussian, the propagated belief is also a GMM with M_k components. More precisely,

$$b[X_{k+1}^-] \doteq \mathbb{P}(X_{k+1} | \mathcal{H}_{k+1}^-) = \mathbb{P}(X_k | \mathcal{H}_k) \mathbb{P}(x_{k+1} | x_k, u_k) = \sum_{i=1}^{M_k} \xi_{k,i} \mathcal{N}(\hat{X}_{k+1,i}^-, \Sigma_{k,i}^-) \quad (3.17)$$

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 along the lines of the discussion in the previous Section 3.2, with additional parameters introduced. For ease of disposition, let us reproduce the steps such as Eq. 3.5 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 Eq. 3.17), from which $b[x_{k+1}^-]$ is calculated, is also a GMM, we can replace $b[x_{k+1}^- = x]$ with

$$b[x_{k+1}^- = x] = \sum_{j=1}^{M_k} \xi_{k+1,j} b[x_{k+1,j}^- = x]. \quad (3.18)$$

However, the actual value of w_i (for each $A_i \in \{A_{\mathbb{N}j}\}$) depends, in addition to the measurement likelihood and event likelihood, also on the GMM 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 $b[x_{k+1}^- = x^{tr}]$ will be unlikely to be actually acquired, leading to low value of w_i with $A_i = A^{tr}$. Since $b[x_{k+1}^-]$ is a GMM with M_k components, low-probability pose x^{tr} corresponds to low probabilities $b[x_{k+1,j}^- = x^{tr}]$ for each component $j \in \{1, \dots, M_k\}$. However, the likelihood term (3.3) could still go up in case of perceptual aliasing, where the aliased scene A_j generates a similar observation to z_{k+1} from viewpoint x' with latter being more probable, i.e. high probability $b[x_{k+1}^- = x']$.

In practice, calculating the integral in Eq. 3.6 can be done efficiently considering separately each component of the GMM $b[x_{k+1}^-]$. 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 and it is known that a product of Gaussians remains a Gaussian. The integral can then be only 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 Eq. 3.6 should be computed numerically. Since in practical applications $\mathbb{P}(A_i | x, \mathcal{H})$ is sparse w.r.t. x , this computational cost is not severe.

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

$$\mathbb{P}(X_{k+1} | \mathcal{H}_{k+1}^-, z_{k+1}) = \sum_{j=1}^{M_k} \sum_{i=1}^{|A_{\mathbb{N}j}|} \mathbb{P}(X_{k+1}, A_i, \gamma = j | \mathcal{H}_{k+1}^-, z_{k+1}) \quad (3.19)$$

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

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

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

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

$$\mathbb{P}(A_i, \gamma = j \mid \mathcal{H}_{k+1}^-, z_{k+1}) = \mathbb{P}(A_i \mid \gamma = j, \mathcal{H}_{k+1}^-, z_{k+1}) \cdot \mathbb{P}(\gamma = j \mid \mathcal{H}_{k+1}^-, z_{k+1})$$

Here, $\mathbb{P}(\gamma = j \mid \mathcal{H}_{k+1}^-, z_{k+1})$ 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 3.2 (see Eq. 3.8) with the only difference that here all x considered are from the j^{th} component of the belief.

$$b[X_{k+1}] = \sum_{j=1}^{M_k} \sum_{i=1}^{|A_N|} \xi_k^j \mathbb{P}(A_i \mid \mathcal{H}_{k+1}^-, z_{k+1}, \gamma = j) b[X_{k+1}^{j+} \mid A_i]. \quad (3.20)$$

$$\tilde{w}_{k+1}^{ij} \doteq \eta' \int_x \mathbb{P}(z_{k+1} \mid A_i, x, \mathcal{H}_{k+1}^-) \mathbb{P}(A_i \mid \mathcal{H}_{k+1}^-, \gamma = j, x) b[x_{k+1}^{j-} = x], \quad (3.21)$$

with $\eta' = 1/\mathbb{P}(z_{k+1} \mid \mathcal{H}_{k+1}^-)$. Note that for each component j , $\sum_i \tilde{w}_{k+1}^{ij} = 1$. Finally, we can re-write Eq. 3.20 as

$$\mathbb{P}(X_{k+1} \mid \mathcal{H}_{k+1}^-, z_{k+1}) = \sum_{r=1}^{M_{k+1}} \xi_{k+1}^r \mathbb{P}(X_{k+1} \mid \mathcal{H}_{k+1}, \gamma = r), \quad (3.22)$$

or in short, $b[X_{k+1}] = \sum_{r=1}^{M_{k+1}} \xi_{k+1}^r b[X_{k+1}^{r+}]$, where

$$\xi_{k+1}^r \doteq \xi_{k+1}^{ij} \equiv \xi_k^j \tilde{w}_{k+1}^{ij} \quad , \quad b[X_{k+1}^{r+}] \doteq b[X_{k+1}^{j+} \mid A_i]. \quad (3.23)$$

As seen, we got a new GMM with M_{k+1} components, where each component $r \in [1, M_{k+1}]$, with appropriate mapping to indices (i, j) from Eq. 3.20, 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

$$b[X_{k+1}^{r+}] \propto b[X_{k+1}^{j-}] \mathbb{P}(z_{k+1} \mid x_{k+1}, A_i) = \mathcal{N}(\hat{X}_{k+1}^r, \Sigma_{k+1}^r), \quad (3.24)$$

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

3.3.2 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 step, 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:

$$J(u_{k:k+L-1}) = \int_{z_{k+1:k+L}} \sum_{l=1}^L \overbrace{\mathbb{P}(z_{k+l} | \mathcal{H}_{k+l}^-)}^{(a)} c_l \left(\overbrace{\mathbb{P}(X_{k+l} | \mathcal{H}_{k+l}^-, z_{k+l})}^{(b)} \right), \quad (3.25)$$

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.

Similar to the myopic case in Section 3.2, the two terms (a) and (b) in Eq. 3.25 have intuitive meaning: for each considered value of z_{k+l} , (a) represents how likely is it to get such an observation, while (b) corresponds to the posterior belief *given* this specific z_{k+l} . However, the difference in a non-myopic case is that both terms are conditioned on the history \mathcal{H}_{k+l}^- which is a function of $z_{k+1:k+l-1}$; hence, the above reasoning is valid for *all* possible realizations of $z_{k+1:k+l-1}$ and the corresponding posterior beliefs $\mathbb{P}(X_{k+l-1} | \mathcal{H}_{k+l-1})$.

It is not difficult to show that the posterior at each step k is actually the GMM

$$\mathbb{P}(X_{k+l} | \mathcal{H}_{k+l}^-, z_{k+l}, A_i) = \sum_{j=1}^{M_{k+l-1}} \xi_{k+l-1}^j b[X_{k+l}^{j+} | A_i], \quad (3.26)$$

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 of the j th GMM component of the propagated belief $b[X_{k+l}^-]$.

Plugging-in Eq. 3.26 into $\mathbb{P}(X_{k+l} | \mathcal{H}_{k+l}^-, z_{k+l}) \equiv b[X_{k+l}]$ from Eq. 2.8 yields:

$$b[X_{k+l}] = \sum_{i=1}^{|A_{\mathbb{N}}|} \sum_{j=1}^{M_{k+l-1}} \xi_{k+l-1}^j \mathbb{P}(A_i | \mathcal{H}_{k+l}^-, z_{k+l}) b[X_{k+l}^{j+} | A_i]. \quad (3.27)$$

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

$$\tilde{w}_{k+l}^{ij} \doteq \eta' \int_x \mathbb{P}(z_{k+l} | A_i, x, \mathcal{H}_{k+l}^-) \mathbb{P}(A_i | \mathcal{H}_{k+l}^-, x) b[x_{k+l}^{j-} = x], \quad (3.28)$$

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 re-write Eq. 3.20 as

$$b[X_{k+l}] = \sum_{r=1}^{M_{k+l}} \xi_{k+l}^r \mathbb{P}(X_{k+l} | \mathcal{H}_{k+l}, \gamma = r) = \sum_{r=1}^{M_{k+l}} \xi_{k+l}^r b[X_{k+l}^{r+}], \quad (3.29)$$

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 got a new GMM with M_{k+l} components, where each component $r \in [1, M_{k+l}]$, with appropriate mapping to indices (i, j) from Eq. 3.20, 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 .

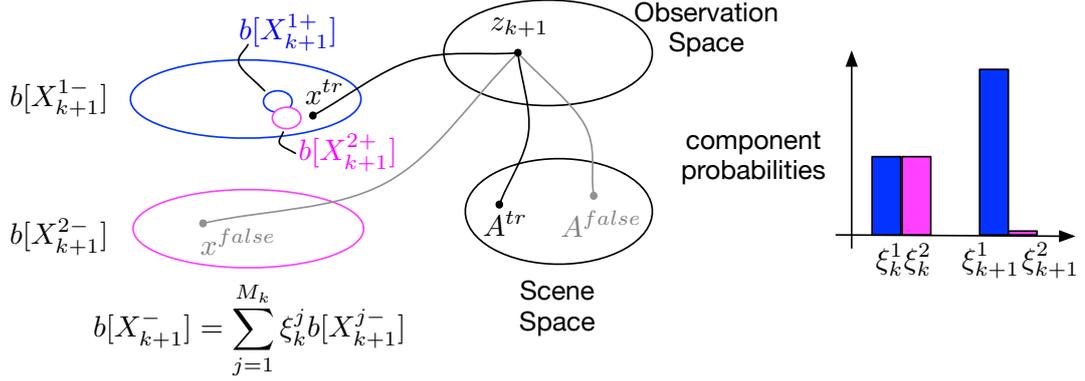


Figure 3.1: GMM posterior $b[X_{k+1}]$ given $z_{k+1} \in \{z_{k+1}\}$. The prior has two equi-probable components while the posterior has different weights for the two components.

The associated cost of the overall posterior of this L step planning can then be compared with that of similar posterior 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 which then serves as a prior for next L step planning. However, a naïve implementation of such a planning would likely suffer from the usual curses of *dimensionality* and of *history*. Luckily, DA-BSP provides a principled way to strike a balance between requirement of an efficient solution and not losing the correct data association in a challenging aliased environment.

3.3.3 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 summarised in Algorithm 3.2 and briefly described below.

Given a GMM 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 3.2.4. The algorithm then calculates the contribution of each observation $z_{k+1} \in \{z_{k+1}\}$ to the objective function (3.11). In particular, on 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} (see Section 3.2.1). On lines 10-16 we compute the posterior belief: according to Section 3.3.2, 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 GMM 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 Eq. 3.27, each of the M_k components from the belief at a previous time, is split into $|A_{\mathbb{N}}|$ new components with appropriate weights. This would imply an explosion in the number of components, making the proposed

Algorithm 3.2 Data association aware belief-space planning

Input: Current GMM 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: $b[X_{k+1}^-] \leftarrow b[X_k] \mathbb{P}(x_{k+1} | x_k, u_k)$ ▷ Eq. 2.7
- 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: ▷ compute weight, Eq. 3.6
- 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 GMM component, Eq. 3.21
- 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: $\xi_{k+1}^{ij} \leftarrow \xi_k^j \tilde{w}_{k+1}^{ij}$ ▷ Eq. 3.23
- 14: ▷ Calculate posterior of $b[X_{k+1}^{j-}]$, given A_i , Sec. 3.3.2
- 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 Eq. 3.22
- 20: $c \leftarrow \text{CalcCost}(b[X_{k+1}^+])$ ▷ Sec. 3.2.6
- 21: $J \leftarrow J + w_s \cdot c$
- 22: **end for**
- 23: **return** J

framework hardly applicable. However, in practice, the majority of the weights will be negligible, and therefore can be pruned, while the remaining number of components is denoted by M_{k+1} in Eq. 3.22. Depending on the scenario and the degree of perceptual aliasing, this can correspond to *full* or *partial* disambiguation (see Fig. 3.1).

3.3.4 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 number of components. Using 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, and 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 (see e.g., Tab. 4.3).

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 multi-hypothesis tracking (MHT). Roughly speaking, in such a scenario, planning is

through explicit instantiation of the trajectory of control and pruning decision is often based on information-theoretic value a particular branch is expected to hold. Thus MHT can handle passive case of BSP where disambiguation is sought after only in the inference step and not in planning. In contrast, DA-BSP argues for data-association within the BSP framework and thus can 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 overall likelihood of associations considered from the initial position, we call it global, while in 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 conditional posterior and re-normalizing the other weights whereas merging is the process of combining two components to form a single component, which is optimal (in some heuristic sense) representation of the both. Both pruning and merging can be recursive processes.

3.3.5 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 is the likelihood to have 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. Here, by *full disambiguation* we mean that eventually the posterior belief has only a single component. For a usual forward L step planning, this can not 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 result in a full disambiguation eventually. Hence, DA-BSP captures the reality of perceptually aliased environment quite well.

3.3.6 Degenerate cases of DA-BSP

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

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 3.21 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 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.

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 $\{A_{\mathbb{N}}\}$ 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 BSP.

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 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 $\{A_{\mathbb{N}}\}$, that generated it.

3.3.7 On correctness of non-myopic DA-BSP

In order to reason about the *correctness* of DA-BSP i.e., whenever there is a single disambiguating data association, the algorithm will recognise 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 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 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 0 except the one corresponding to this ground truth. However in the case of pruned DA-BSP this might not be true necessarily as the ground truth component might be pruned

away in $b[X_{k-1}]$ leading possibly to even 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 starting the multi-modal belief contains the correct component as well. The latter can be avoided by judicious choice of pruning threshold. As shown in the experiments later on, DA-BSP is not sensitive to the choice of this threshold.

Chapter 4

Results and Discussions

In this section we present an experimental analysis of the proposed approach. First, we talk about real experiment with a Pioneer platform fitted with an RGB camera in a perceptually aliased corridor environment. We use a non-myopic planning framework using explicit scenes simulated via AprilTags. To consider arbitrary levels of ambiguities, we then talk about a simulated world model (in Gazebo) of two nearly identical office floors with various look-alike cubicles. In these simulations, the Pioneer robot is fitted with a laser scanner.

4.1 Implementation of data-association aware BSP

Effective and realistic implementation of DA-BSP requires two separate threads of development. In order 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 GMM as a factor graph so that state-of-the-art tool GT-SAM could be harnessed for a time-efficient inference. On the other hand, in order to be realistic enough, we implemented it on 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 as well as a real world scenario. In order to simulate a complex world with arbitrary levels of ambiguity, we chose Gazebo since it fits nicely to both the robotic platform as well as the ROS infrastructure. These two streams of development are shown in left and right parts of Fig. 4.1. The DA-BSP algorithm itself was implemented in object-oriented MATLAB with the aim of striking a balance between rapid prototyping and obtaining a generalisable implementation that can be easily ported to languages such as C++.

4.2 Metrics for evaluating DA-BSP

Evaluating DA-BSP is linked to the notion of data association which is typically assumed to be solved in BSP. As mentioned before, accounting for data association within belief

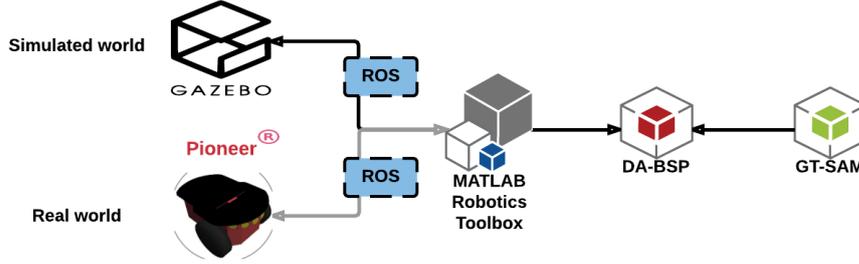


Figure 4.1: The overall infrastructure of implementation of DA-BSP. Note that thanks to the middleware block of ROS, the algorithm is independent of whether it is applied on a real setup or a simulated one.

space planning does not come free. On the other hand, assuming such an association randomly is bound to fail. In order to evaluate an approach, we keep track of how many components the belief has and how many times it can make a correct association. Recall from Section 3.3.7 that as long as the component corresponding to the ground truth is present in the belief, the correctness can be guaranteed. We denote this by boolean symbol DA. Note that in case of no pruning, DA-BSP is guaranteed to be correct and hence DA is set as true. This is true also if the pruning is not detrimental to the correct component and association. η_{da} measures the weight of the correct component in the belief; in case, the correct component is lost, the corresponding η_{da} 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} . In order 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. In another approach, we assume the similar correct association with the scene however instead of the correct ground truth component, all components of the belief are considered. This approach implies that a multi-modal prior remains multi-modal after inference too. It is named BSP-mul here. In both of these variants, we are interested in correct associations being made out of many trials. This is measured by the metric ξ_{ca} where value 1 would indicate that correct associations were made in all of the random trials. This could happen when the belief is unimodal due to lack of ambiguity in the vicinity.

4.3 Real-world application with explicit scenes - octagonal corridor

In order to elucidate the crucial properties of non-myopic DA-BSP, we consider a real world experiment as shown in the Fig. 4.3 with a single robot R. The abstracted schema of the world is shown in the center. The state space $X \in \mathbb{R}^3$ consists of 2D coordinates as shown, as well as the orientation of the robot. Here, A_i denotes an Apriltag with

the index i . This enables us to simulate arbitrary levels of perceptual aliasing.¹ To ensure robustness, the tag A_i is considered detected only if it is also within a closed sub-space $X_{A_i} \subset X$. Typically, this is decided through the centrality of the detected tag in the image observed by the camera. Figure 4.2 shows two sampled cases where the tags are detected and not detected respectively. Initially, the belief of the robot is a multi-modal distribution, represented by a GMM with 4 components having equal weights and centered around each x_{init} . The objective of the robot is to both localize itself and to reach the elevator denoted by x_{goal} . Initially, the planner is provided with a set \mathcal{T} of control trajectories. Consequently, depending on the planning algorithm used (i.e., DA-BSP, or BSP-uni) as well the planning horizon L , the cost of each trajectory $\tau \in \mathcal{T}$ is evaluated and the optimal (w.r.t this cost) trajectory is chosen. The L -step planning, followed by enacting one optimal control action and the consequent inference, shall together be called an *epoch*. Note that this simple representation of the world is very general. Indeed, real world complications – such as the state space being of higher dimension, different levels of ambiguities between the scenes and planning problem of longer time-scales – can all be easily incorporated into it.

Since we model the visual observation via AprilTags, due to sensory limitations (such as out-of-view or far-from-center tags), a reliable observation might not be available at each step of motion. One such instance is depicted in the Fig. 4.2. In such conditions, no data association can be made and consequently, DA-BSP behaves exactly like the usual belief space planning.

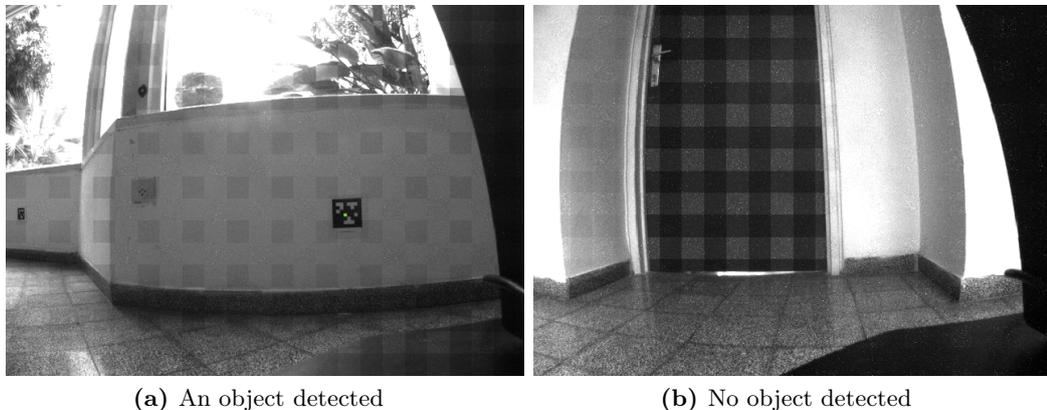


Figure 4.2: (*left*) Apriltag is detected, indicated with green patch at the center. This provides the transformation matrix between the pose of the robot and the landmark pose. Note that a far-away Apriltag, though visible in this frame, is considered not detected since the non-centrality of the tag makes the observation highly untrustworthy. (*right*) No Apriltag lies within the field of view of the camera.

As shown in the Fig. 4.3, the robot resides in an octagonal corridor with ample instances of ambiguous scenes. Here, the actual floor is shown via laser scan while the 3D view of some typical locations are depicted in the inset figures. Note that the robot does not have these information, and actually has a semantic representation of the map

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

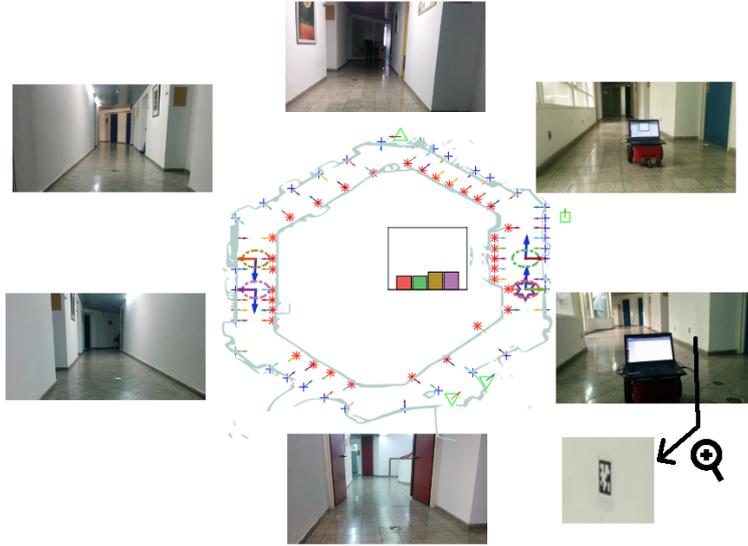


Figure 4.3: Real-world experimental setup (best viewed in colour). The images show the actual third-person view at different locations of the robot, while the two dimensional figure, shows semantic knowledge of the environment the robot possesses. Its current belief is a 4-modal GMM with mean position depicted by x_{init} . Ground truth robot position is indicated with \odot ; arrows indicate orientation (and not motion). The actual scenario is depicted through the laser-scan of the environment shown in the green colour. This map is for representative purposes only and is not available to the robot. The zoomed-in picture of an AprilTag is also shown.

Algorithm	Epoch	Planning			Inference		
		t(s)	η_{da}	\bar{m}	t(s)	η_{da}	\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

Table 4.1: Performing DA-BSP on a real corridor environment shown in the Fig. 4.3, with planning horizon $L = 4$. The times in seconds spent in planning and in inference is denoted by t, while \bar{m} stands for average modes; refer Sec. 4.3.

where perceptual aliasing is accounted for by identical Apriltags. The result of running the DA-BSP on this setup is shown in the Fig. 4.1.

However, when ambiguous data association occurs, DA-BSP considers all possible associations and weighs each new component of the posterior according to the equation 3.28. Fig. 4.4 shows one such instance.

DA-BSP incorporates planning and inference seamlessly under one framework, called *epoch* earlier. The overall planning is performed as a model predictive control composed of several such steps. Fig. 4.5 shows some of the epochs in DA-BSP along with other approaches such as BSP-uni and BSP-mul.

Once such planning is performed under DA-BSP, the subsequent posterior at the end of each epoch might have more or even lesser number of components than before. The former occurs when presence of identical close by tags causes perceptual aliasing

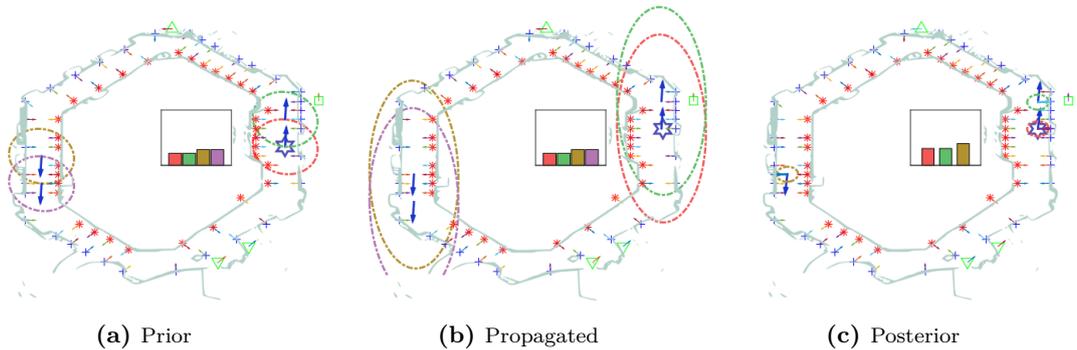


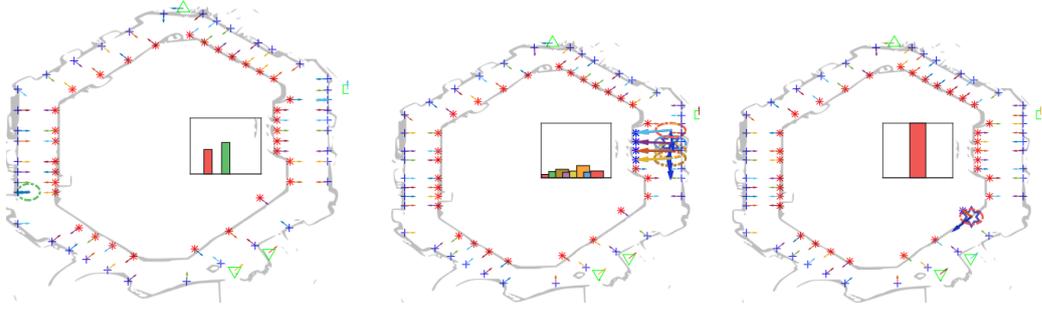
Figure 4.4: Evolution of belief at an epoch $E = 3$ under DA-BSP. Here tags are represented by shapes $\{*, +, \square, \triangle, \nabla\}$ while ground truth robot position is indicated with \star . (a) The prior belief is multimodal with four distinct modes as shown with the coloured ellipses. (b) After incorporating the motion model, the propagated belief is similarly a multi-modal 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 of them to vanish and components reduce from 4 to 3). Also new components in the posterior may emerge (not the case here). Here, $L = 3$.

while the latter is the result of unlikely components being pruned away naturally, in the light of new observations. This evident in Fig. 4.5d.

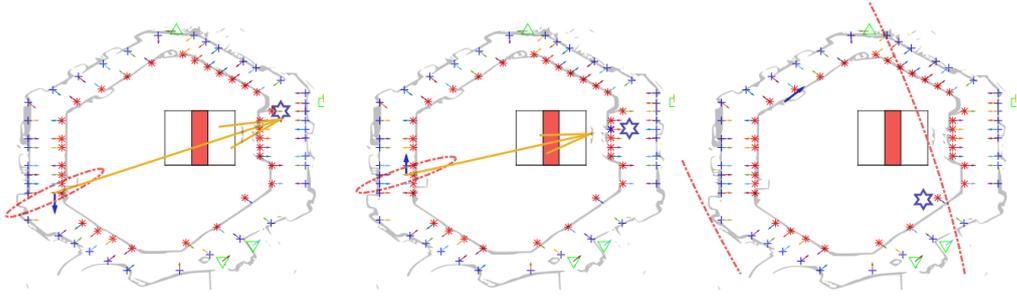
DA-BSP considers all possible associations and then adjusts the weights of the components accordingly whereas other approaches like BSP-uni and BSP-mul perform some kind of simplifying assumption on data association. BSP-uni assumes that with probability $1 - \epsilon$ an oracle tells it the correct component of the prior as well as correct data association. Under extremely mild or no perceptual aliasing, ϵ is close to 0. BSP-mul on the other hand considers correct data association but with respect to all components of the prior. Here, the number of components can not increase. Naturally, DA-BSP when compared with BSP-uni and BSP-mul, trades computation efficiency with correctness of data association. The quantitative aspect of such a comparison is shown in the Tab. 4.2. Another criticism against DA-BSP could be its prohibitive cost when non-myopic planning is considered. Exponential blowup of computational complexity as the planning horizon increases is an issue not specific to DA-BSP. Unfortunately, DA-BSP can not solve or even reduce this burden. Nevertheless, due to *parsimonious data association* the additional cost of DA-BSP may not be significantly more. This is also depicted in the Tab. 4.2.

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 setting. Based on the configuration of the environment, a longer planning horizon may enable quicker disambiguation and consequently reduced KL-cost. In Fig. 4.6, 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.

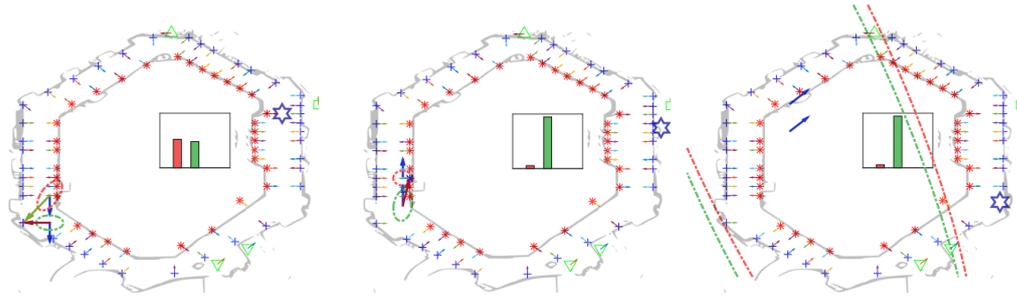
It might appear that DA-BSP is hopelessly expensive in terms of computational efforts and non-trivial pruning techniques might be required to make it applicable in any realistic scenario. However, quite the contrary is true. Realistic scenarios typically



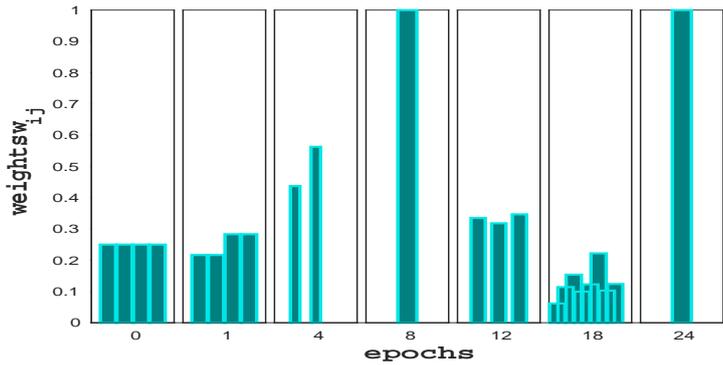
(a) DA-BSP for epochs {4,18,24}.



(b) BSP-uni for epochs {4,18,24}



(c) BSP-mul for epochs {4,18,24}



(d) GMM weights of corresponding beliefs in DA-BSP at the end of epochs {0,1,4,8,12,18,24}.

Figure 4.5: (a)-(c) Evolution of inferred belief as *decision epoch* progresses with $L = 3$; epochs depicted are {4,18,24}. They depict evolution of inferred belief, for different planning algorithms, i.e. DA-BSP, BSP-uni and BSP-mul, respectively. GMM components and associated weights are designated with different colors. Ground truth robot position is indicated with \star . For clarity, the detected scene(s) are shown in different colour. In case of BSP-mul and BSP-uni, this particular instance of planning leads to catastrophically bad data association. (d) Evolution of GMM components weights during these epochs. Note that the number of components *increases* as well as *decreases* and eventually goes to 1. Here, planning horizon is $L = 3$.

Algorithm	Epoch	$L = 1$				$L = 3$				Inference			
		t(s)	η_{da}	\bar{m}	DA	t(s)	η_{da}	\bar{m}	DA	t(s)	η_{da}	\bar{m}	DA
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	✓
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	-	-	-	
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 4.2: 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} . DA signifies *correct* data association; refer Sec. 4.2. Values shown here are for average of 5 random runs.

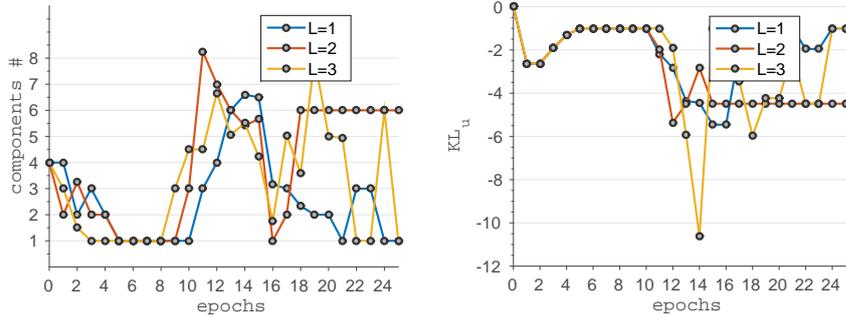


Figure 4.6: Evolution of belief as *decision epoch* progresses during DA-BSP planning. Average number of components in the belief mixtures and the KL_u metric are depicted in *left* and *right* respectively.

Algorithm	Epoch	$\sigma = 0.20$			$\sigma = 0.05$			$\sigma = 10^{-10}$		
		t(s)	η_{da}	\bar{m}	t(s)	η_{da}	\bar{m}	t(s)	η_{da}	\bar{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)

Table 4.3: Evaluating DA-BSP in several steps of planning, under different pruning thresholds of $\sigma = \{0.2, 0.05, 10^{-10}\}$. The times in seconds spent in planning is denoted by t, while η_{da} and \bar{m} show weight of correct association and averaged number of modes, respectively. Values shown here are for average of 5 random runs while standard deviation is depicted within the parenthesis. Note that 5% threshold is sufficient in this case to perform equivalently with almost *unpruned* DA-BSP with $\sigma = 10^{-10}$.

Algorithm	Epoch	$\epsilon = 0.25$		$\epsilon = 0.50$		$\epsilon = 0.75$		$\epsilon = 1.00$	
		t(s)	ξ_{ca}	t(s)	ξ_{ca}	t(s)	ξ_{ca}	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)

Table 4.4: 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* ones. The times in seconds spent in planning is denoted by t, while average correct association is denoted by ξ_{ca} . Values shown here are for average of 5 random runs while standard deviation is depicted within the parenthesis.

do not have persistent ambiguity at each step of navigation, hence the weights of many components drop down naturally to afford an easy approach of pruning. Also DA-BSP is not sensitive to such a pruning parameter σ . Effect of σ on DA-BSP is shown in the Tab. 4.3.

In the presence of data association challenges, the quality of planning can be roughly assessed by considering if at least one of the posterior contains correct data association. This is represented by DA in the Tab. 4.2. Here, η_{da} which also considers the weight of such associations, is also shown. Naturally, reasoning over all possible associations results in greater computational effort. We measure the run-time of the algorithm as a proxy for *effectiveness*. Both these measures along with the number of hypotheses in the beliefs are shown in the Tab. 4.2 where we can see the effect of non-myopic DA-BSP with two different planning horizons.

4.4 Highly-aliased simulated office scenario

To demonstrate our concept in a more challenging scenario under high level of perceptual aliasing, we considered a Gazebo-based simulation of a Pioneer robot in an aliased two-floor office room environment. The robot is fitted with realistic sensors enabling laser scans and odometry estimation. Apart from the implementation mentioned in the Section 4.1, we use ICP for laser scan matching.

The scenario is as shown in Fig. 4.7a. 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 towards y . The two floors are identical except that the floor-2 has an additional printer p_1 (Fig. 4.7a). Additionally, each floor has significant perceptual aliasing within itself due to identical cubicles and self-similar corridors. However, at the

end of the corridors, there could be 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 the cabin c_1 (Fig. 4.7a) and to disambiguate between the floors. Initially the robot wakes up to find itself either of the places 1 or 6 (facing 2 and 7 respectively). Hence its initial belief is modelled as a four-component GMM (two for each floor) whereas the ground truth is at 1 i.e., the robot actually is at 1. Throughout this section, we will use green colour and yellow colour to denote the ground truth and the aliasing respectively.

Consider that the robot starts at 1 (Fig. 4.7a), thus initial belief has mean at this position. A forward action to 2 (Fig. 4.7a). This action can be used to propagate the initial belief. Subsequently, the prior and the propagated (means at 1 and 2 respectively) covariances are shown in Fig. 4.10a. The area of the ellipses equals the actual 2σ covariance. The laser scan obtained for the belief update is shown in Fig. 4.9b, where the green coloured 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 and considering this aliasing scan 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 Fig. 4.10b.

Starting from the positions 1 or 6 there are many possible paths to reach the goal (cabin c_1 (Fig. 4.7a)). We would like to show two such paths. The shortest path is $6 \rightarrow 14 \rightarrow 16 \rightarrow 12$ (Fig. 4.7a). However it leads to an increase in the number of modes and on reaching the the goal, robot is uncertain of the floor it is in. As seen previously (e.g., Fig. 4.10), the modes increase due to the highly aliasing environment. Now consider a longer path $1 \rightarrow 2 \rightarrow 3 \rightarrow 4 \rightarrow 5 \rightarrow 13 \rightarrow 12 \rightarrow 15$ (Fig. 4.7a). Let us call this the DA-BSP path. While following $1 \rightarrow 2$ and $2 \rightarrow 3$, due to the aliased cubicles the number of components increase from 4 to 8. See Figures 4.7a to 4.7c) for the corresponding mean positions of the robot. Intra-floor 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 2 (in Fig. 4.7e) and then increases again to 4 along the paths $5 \rightarrow 13$ and $13 \rightarrow 12$ (in Figures 4.7f and 4.7g respectively). Full disambiguation resulting in a uni-modal belief occurs at 15 due to the presence of the unique printer, p_1 (Figure 4.7a). Fig. 4.9 depicts the evolution of weights along both the paths.

In Fig. 4.8, we see how the different components of the belief and the respective weights evolve when following 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 ($E=8$). Fig. 4.8b shows the cardinality of components in the GMM 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

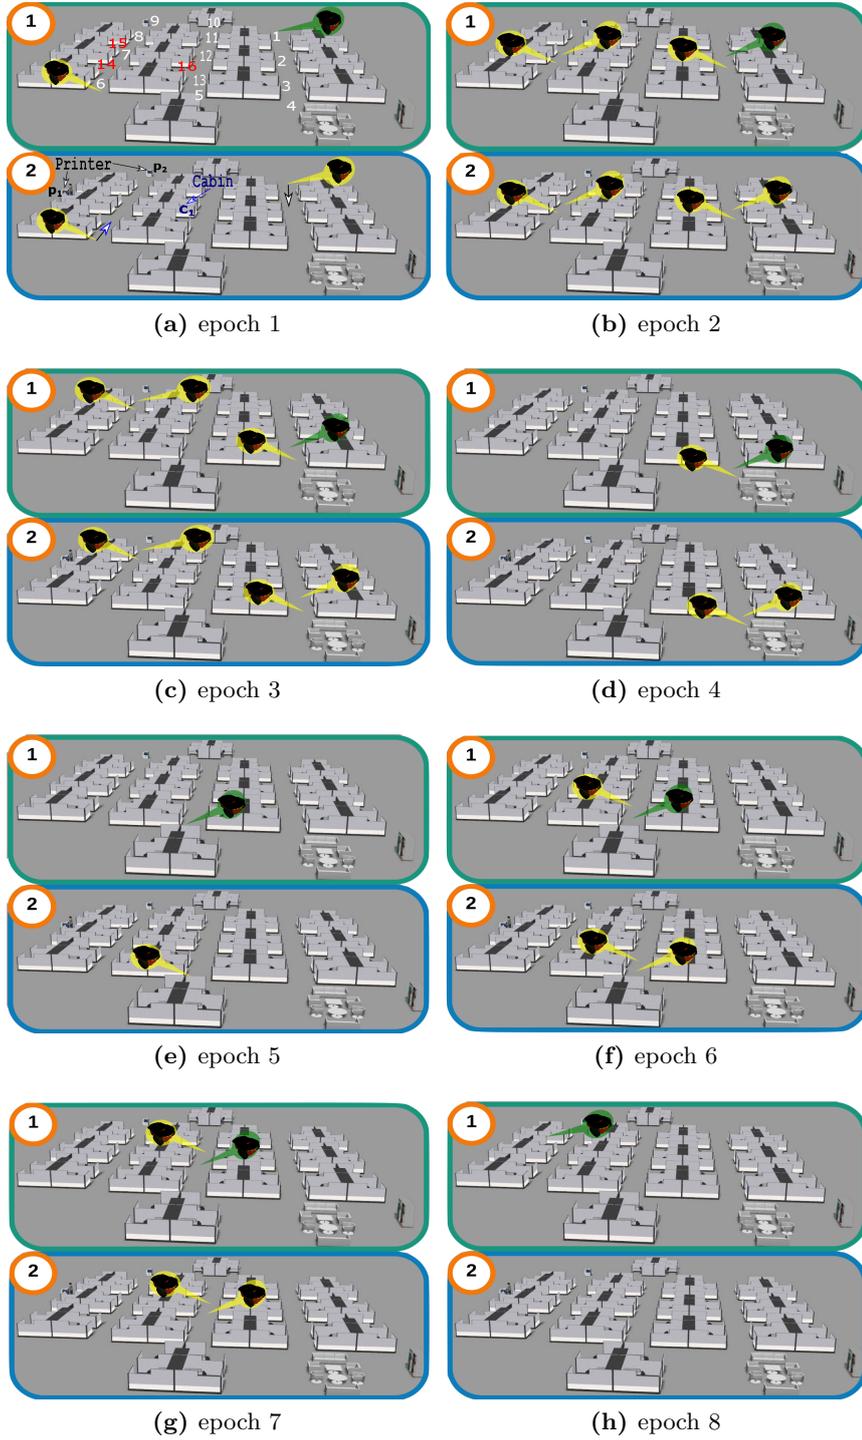


Figure 4.7: Fig. 4.7a Two-floor aliased office environment in a Gazebo simulator. p_1 and p_2 denote the printers while 1 and 6 are the mean positions in each floor for the initial four-component GMM belief. Figures 4.7b till 4.7h show the mean positions (modes) of the robot for each step of the DA-BSP path. Green denotes the ground truth while yellow the aliasing position.

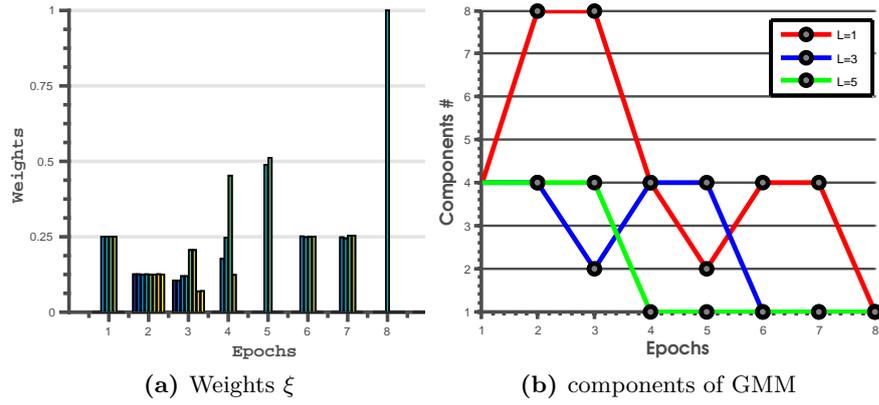


Figure 4.8: (left) Evolution of weights of the components in the GMM after inference for $L = 2$. (right) Average number of components in the belief mixtures for different planning horizon.

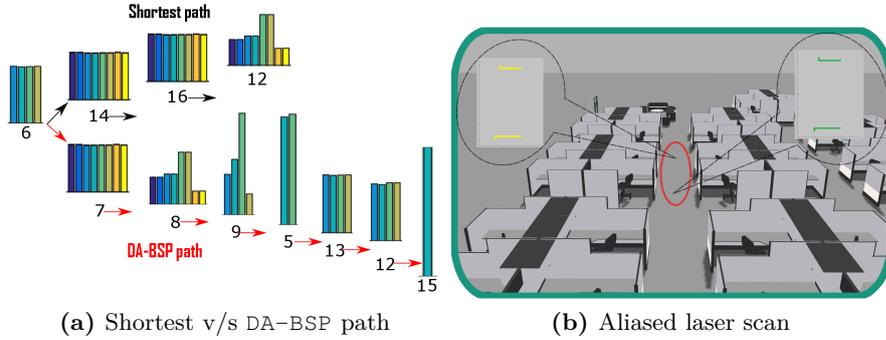


Figure 4.9: (left) Evolution of weights of the components of the belief when following the shortest path versus that following the DA-BSP path. (right) Laser scans at ground truth and aliased position (green and yellow respectively).

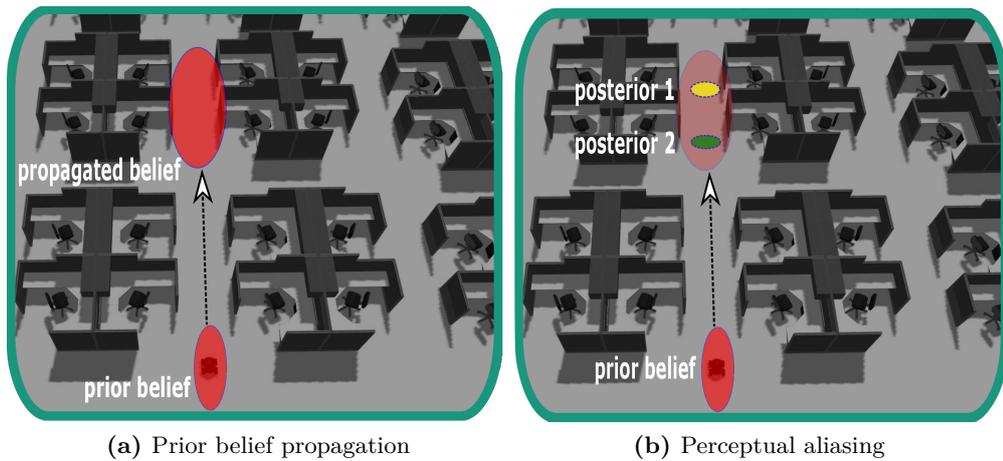


Figure 4.10: Prior belief is propagated according to the motion model. Within the subsequent propagated belief, there are perceptually aliased laser scans observed. Here, 2σ covariance is depicted with each ellipse.

Algorithm	Epoch	$L = 2$		$L = 4$		Inference	
		t(s)	(η_{da}, \bar{m})	t(s)	(η_{da}, \bar{m})	t(s)	(η_{da}, \bar{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

Table 4.5: Evaluating DA-BSP in several steps of planning and inference, for $L = 2$ and $L = 4$. The times in seconds spent in planning and in inference is denoted by t, while \bar{m} denotes average modes. η_{da} measures the level of aliasing whereas DA is a binary variable denoting correct or wrong association.

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)

Table 4.6: BSP-uni in 5 different runs. BSP-uni can be seen as a very drastic pruning where data association may or may not be correct. This is seen from the ξ_{ca} values (for 5 random runs). Standard deviation are mentioned within the parenthesis.

faster as can be seen from Fig. 4.8b. Note that full disambiguation occurs at $E=8$ for myopic planning ($L = 1$). Thus, for $L = 5$ which can project 5 steps in the future, such a disambiguation occurs from $E=4$ onwards. $L = 3$ lies somewhere in between where the full disambiguation occurs from $E=6$ onwards.

Tab. 4.5 compares DA-BSP with BSP-uni at different epochs of planning and inference for planning horizons of $L = \{2, 4\}$. Here, the DA-BSP path is considered. Recall that η_{da} stands for the weight of the component corresponding to the ground truth. For example for $E=2$, DA-BSP inference results in 8 modes arising from 7 other observations that alias the ground truth. Subsequently, $\eta_{da} = 0.12$. 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 due 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.

Chapter 5

Conclusions and Future Work

In this thesis, we presented a unified framework for robust perception in planning as well as inference. State-of-the-art belief space planning (BSP) approaches typically consider data association to be given and perfect. However, such an assumption is less appropriate in 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 data association aware belief space planning (DA-BSP) approach that relaxes the data-association assumption by incorporating reasoning regarding DA within BSP. In the context of passive approach with the observations provided, it results in more robust inference. On the other hand, in the context of 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. As such, this work is a unified framework for robust active and passive perception. DA-BSP considers data-association in a principled rigorous way with the belief space planning. It is a more faithful representation of an aliased environment since the number of components can increase as well as decrease. Though this increases the computational burden of planning, it is both necessary for ambiguous environments and is still practically applicable, as shown through numerous experiments in both a realistic Gazebo simulation as well as in real experiment with the Pioneer robot platform. Additionally, DA-BSP degenerates to usual approaches in the presence of helpful assumptions such as under very small localization uncertainty and under lack of 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.

5.1 Future Work

In this thesis we incorporated reasoning regarding DA within BSP. While this is a novel work by itself there is scope for further improvement/extension of the framework developed. Given below are few such directions.

1. In this work, we assume that the environment that the robot operates upon is known/pre-mapped. The DA-BSP framework can be extended to a general SLAM framework where the environment itself is partially mapped or unknown.
2. Multi-robot DA-BSP is a natural extension of the single robot case, with numerous advantages compared to the latter. Yet, multi-robot collaboration is a challenging problem: To perform cooperative inference, each robot needs to determine what information to share with other robots, and to correctly associate information received from nearby robots with appropriate random variables. Building upon our DA-BSP approach, we propose an *active* case, reasoning about robot actions that will facilitate better collaboration between robots in previously unknown perceptually aliasing environments.
3. In aliasing environments, it becomes important to reason about the probability of each scene or object to be observed from a view point. This is because only objects that are likely to be observed should be considered while we reason about DA. This probability, also called the *event likelihood* depends on the environment model we assumed this term to be given. However, a more general approach is to determine this term from the environment that is being operated upon. A possible avenue for future work is to investigate such a general representation.

Chapter 6

Appendix

List of Publications

Publications based on this Master's research:

1. S. Pathak, A. Thomas, A. Feniger, and V. Indelman. "Robust Active Perception for Belief Space Planning in Perceptually Aliased and Uncertain Environments," in *The 5th Israeli Conference on Robotics (ICR)*, Herzliya, Israel, April 2016
2. S. Pathak, A. Thomas, A. Feniger, and V. Indelman. "Towards data association aware belief space planning for robust active perception," in *AI for Long-term Autonomy, workshop in conjunction with IEEE International Conference on Robotics and Automation (ICRA)*, Stockholm, Sweden, May 2016
3. S. Pathak, A. Thomas, A. Feniger, and V. Indelman. "DA-BSP: Towards Data Association Aware Belief Space Planning for Robust Active Perception," in *European Conference on Artificial Intelligence (ECAI)*, The Hague, Netherlands, September 2016
4. S. Pathak, A. Thomas, A. Feniger, and V. Indelman. "Robust Active Perception via Data-association aware Belief Space Planning," *arXiv:1606.05124*, 2016
5. S. Pathak, A. Thomas and V. Indelman. "Nonmyopic Data Association Aware Belief Space Planning for Robust Active Perception," in *International Conference on Robotics and Automation (ICRA 2017)*
6. S. Pathak, A. Thomas and V. Indelman. "A Unified Framework for Data Association Aware Robust Belief Space Planning and Perception," in *International Journal of Robotics Research*, **submitted**

Bibliography

- [1] S. Agarwal, A. Tamjidi, and S. Chakravorty. Motion planning in non-gaussian belief spaces for mobile robots. *arXiv preprint arXiv:1511.04634*, 2015.
- [2] A.-A. Agha-Mohammadi, S. Chakravorty, and N. M. Amato. Firm: Sampling-based feedback motion planning under motion uncertainty and imperfect measurements. *Intl. J. of Robotics Research*, 2014.
- [3] N. Atanasov, B. Sankaran, J.L. Ny, G. J. Pappas, and K. Daniilidis. Nonmyopic view planning for active object classification and pose estimation. *IEEE Trans. Robotics*, 30:1078–1090, 2014.
- [4] Anonymous authors. Robust autonomous navigation in perceptually aliased environments via belief space planning - supplementary material.
- [5] Yaakov Bar-Shalom, X Rong Li, and Thiagalingam Kirubarajan. *Estimation with applications to tracking and navigation: theory algorithms and software*. John Wiley & Sons, 2004.
- [6] A. Bry and N. Roy. Rapidly-exploring random belief trees for motion planning under uncertainty. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pages 723–730, 2011.
- [7] L. Carlone, A. Censi, and F. Dellaert. Selecting good measurements via l1 relaxation: A convex approach for robust estimation over graphs. In *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, pages 2667–2674, 2014.
- [8] S. M. Chaves, A. Kim, and R. M. Eustice. Opportunistic sampling-based planning for active visual slam. In *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, pages 3073–3080. IEEE, 2014.
- [9] S. M. Chaves, J. M. Walls, E. Galceran, and R. M. Eustice. Risk aversion in belief-space planning under measurement acquisition uncertainty. In *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2015.

- [10] H.F. Durrant-Whyte and T. Bailey. Simultaneous localisation and mapping (SLAM): Part I the essential algorithms. *Robotics & Automation Magazine*, Jun 2006.
- [11] V. Indelman, L. Carlone, and F. Dellaert. Planning in the continuous domain: a generalized belief space approach for autonomous navigation in unknown environments. *Intl. J. of Robotics Research*, 34(7):849–882, 2015.
- [12] V. Indelman, N. Michael, and F. Dellaert. 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*, July 2014.
- [13] V. Indelman, E. Nelson, J. Dong, N. Michael, and F. Dellaert. Incremental distributed inference from arbitrary poses and unknown data association: Using collaborating robots to establish a common reference. *IEEE Control Systems Magazine (CSM), Special Issue on Distributed Control and Estimation for Robotic Vehicle Networks*, 36(2):41–74, 2016.
- [14] V. Indelman, E. Nelson, N. Michael, and F. Dellaert. Multi-robot pose graph localization and data association from unknown initial relative poses via expectation maximization. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2014.
- [15] L. P. Kaelbling, M. L. Littman, and A. R. Cassandra. Planning and acting in partially observable stochastic domains. *Artificial intelligence*, 101(1):99–134, 1998.
- [16] M. Kaess, H. Johannsson, R. Roberts, V. Ila, J. Leonard, and F. Dellaert. iSAM2: Incremental smoothing and mapping using the Bayes tree. *Intl. J. of Robotics Research*, 31:217–236, Feb 2012.
- [17] A. Kim and R. M. Eustice. Active visual slam for robotic area coverage: Theory and experiment. *Intl. J. of Robotics Research*, 2014.
- [18] H. Kurniawati, D. Hsu, and W. S. Lee. Sarsop: Efficient point-based pomdp planning by approximating optimally reachable belief spaces. In *Robotics: Science and Systems (RSS)*, volume 2008, 2008.
- [19] M. Lauri, N. A. Atanasov, G. Pappas, and R. Ritala. Active object recognition via monte carlo tree search. In *Workshop on Beyond Geometric Constraints at the International Conference on Robotics and Automation (ICRA)*, 2015.
- [20] E. Olson and P. Agarwal. Inference on networks of mixtures for robust robot mapping. *Intl. J. of Robotics Research*, 32(7):826–840, 2013.

- [21] C. Papadimitriou and J. Tsitsiklis. The complexity of markov decision processes. *Mathematics of operations research*, 12(3):441–450, 1987.
- [22] S. Patil, G. Kahn, M. Laskey, J. Schulman, K. Goldberg, and P. Abbeel. Scaling up gaussian belief space planning through covariance-free trajectory optimization and automatic differentiation. In *Intl. Workshop on the Algorithmic Foundations of Robotics*, 2014.
- [23] T. Patten, M. Zillich, R. Fitch, M. Vincze, and S. Sukkariéh. Viewpoint evaluation for online 3-d active object classification. *IEEE Robotics and Automation Letters (RA-L)*, 1(1):73–81, January.
- [24] J. Pineau, G. J. Gordon, and S. Thrun. Anytime point-based approximations for large pomdps. *J. of Artificial Intelligence Research*, 27:335–380, 2006.
- [25] R. Platt, L. Kaelbling, T. Lozano-Perez, and R. Tedrake. Efficient planning in non-gaussian belief spaces and its application to robot grasping. In *Proc. of the Intl. Symp. of Robotics Research (ISRR)*, 2011.
- [26] S. Prentice and N. Roy. The belief roadmap: Efficient planning in belief space by factoring the covariance. *Intl. J. of Robotics Research*, 2009.
- [27] Bharath Sankaran, Jeannette Bohg, Nathan Ratliff, and Stefan Schaal. Policy learning with hypothesis based local action selection. *arXiv preprint arXiv:1503.06375*, 2015.
- [28] C. Stachniss, G. Grisetti, and W. Burgard. Information gain-based exploration using rao-blackwellized particle filters. In *Robotics: Science and Systems (RSS)*, pages 65–72, 2005.
- [29] N. Sunderhauf and P. Protzel. Towards a robust back-end for pose graph slam. In *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, pages 1254–1261. IEEE, 2012.
- [30] J. Van Den Berg, S. Patil, and R. Alterovitz. Motion planning under uncertainty using iterative local optimization in belief space. *Intl. J. of Robotics Research*, 31(11):1263–1278, 2012.
- [31] J. M. Walls, S. M. Chaves, E. Galceran, and R. M. Eustice. Belief space planning for underwater cooperative localization. In *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2015.
- [32] L. Wong, L. Kaelbling, and T. Lozano-Pérez. Data association for semantic world modeling from partial views. *Intl. J. of Robotics Research*, 34(7):1064–1082, 2015.

DA-BSP מצליח לתכנן את מסלול הרובוט שמוביל ל-disambiguation לפני ההגעה אל היעד. הניסויים והסימולציה הדגימו היבטים ותכונות עיקריות של הגישה שפותחה; בפרט, הראינו כי מספר הרכיבים ב-GMM יכול לקטון וגם לגדול, אי הודאות המיוצגת על ידי כל רכיב ב-GMM מעודכנת בצורה מתאימה, וניתחנו את זמני הריצה של האלגוריתם. למעשה, גישת DA-BSP יכולה להיחשב כשיטה מאוחדת לחישה חסינה (לדו משמעות) פאסיבית ואקטיבית.

אפשריות ובהינתן פונקציית מחיר. עבור פעולה אפשרית כלשהי מאוסף זה, הדבר מצריך להתחשב בכל הערכים האפשריים של מדידה עתידית (למשל, תמונה) שיכולים להתקבל לאחר ביצוע הפעולה, היות והמדידה העתידית הזאת טרם התקבלה.

אינטואיטיבית, לכל ערך אפשרי של תצפית, אנחנו מחשבים שני גורמים אשר נדרשים לכימות פונקציית המחיר: (א) אנחנו מחשבים את הסבירות לקבלת התצפית הזו בהינתן היסטוריה של מדידות ופעולות עד כה, ו(ב) אנחנו מחשבים את האמונה האפוסטריורית (posterior belief), בהינתן תצפית זו. בשני המקרים, ההנחה של שיוך מידע (DA) נתון ומושלם, כפי שנעשה בגישות BSP קיימות, משמעותה שהאובייקט שנצפה מתוך התצפית הזאת הינו ידוע (ונכון). הנחה זו מפשטת רבות את חישוב שני הגורמים שצוינו. עם זאת, תחת בקרה סטוכסטית ואי ודאות במצב הרובוט, המצב האמיתי של הרובוט אינו ידוע; כל שניתן הינו לאפיין את הפילוג ההסתברותי (belief) המתאים עבור מצב הרובוט. אם ישנה דו משמעות בסביבה בה אמונה זו מקבלת ערכים לא זניחים, משמעות הדבר היא שעלול להתבצע שיוך שגוי.

בגישה שלנו, אנחנו לוקחים בחשבון מצבים מסוג זה ומתארים את כל השייכים האפשריים (עבור ערך מסוים של התצפית העתידית) כחלק מהאמונה. תהליך זה מוביל ישירות להתפלגות אפוסטריורית שהיא GMM, המתקבלת מהאמונה האפריורית (prior belief), שכאמור גם היא יכולה להיות GMM. בגישה שלנו מספר הרכיבים ב-GMM של האמונה האפוסטריורית יכול להשתנות בהתאם לרמת הדו משמעות בסביבה: לקטון במקרה של disambiguation אבל גם לגדול במקרה והפעולה העתידית מובילה את הרובוט לאזור עם דרגה גבוהה של דו משמעות. אנחנו מחשבים את משקלי רכיבי ה-GMM בצורה מפורשת, על סמך משקלי ה-GMM של ה-prior belief והתצפית שהתקבלה. זה מהווה גישה פורמאלית לטיפול בבעיית שיוך המידע כחלק מאלגוריתם לתכנון במרחב הסתברותי, כלומר גישת (DA-BSP) (data association aware space planning). בפרט, גישה זו מאפשרת למצוא פעולות מיטביות בסביבה דו משמעית וגם בסביבה ייחודית, ללא ההנחה המקלה של שיוך מידע פתור ונתון.

DA-BSP הינה גישה כללית שיכולה להתמודד עם מגוון רחב של פונקציות מחיר. בעבודה זו, פונקציית המחיר כוללת מספר איברים, ובהם הגעה ליעד, אי ודאות ואיבר המכמת דו משמעות. שני האיברים הראשונים הינם סטנדרטיים בגישות BSP. האביר האחרון מתייחס לאופי האמונה בהקשר של דו משמעות. בפרט, אנחנו מסתכלים על מספר הרכיבים ב-GMM ועל משקלי הרכיבים. אינטואיטיבית, פעולה שמובילה ל-disambiguation משמעותה שה-GMM של ה-posterior belief יהיה בעל מרכיב אחד בעל משקל גבוה ואילו משקלם של כל שאר הרכיבים יהיו זניחים. ישנן מספר אפשרויות לפונקציית מחיר המייצגות היבטים אלו. בעבודה זו, למשל, נעשה שימוש ב-KL divergence) עם סימן מינוס) ביחס להתפלגות אחידה, כך שככל שמשקלי האמונה מתרחקים מהתפלגות זו, המחיר (penalty) קטן. בפרט, עבור מצב של disambiguation שתואר קודם לכן, פונקציה זו תקבל ערך נמוך מאוד. כתוצאה מכך, ניתן לזהות פעולות אשר צפויות להוביל ל-disambiguation.

במסגרת עבודה זו בוצע גם ניתוח ביצועים של גישת DA-BSP; הניתוח בוצע בסימולציה סינטטית מציאותית (Gazebo) וגם בניסויים אמיתיים במעבדת ניווט אוטונומי וחישת עולם (Perception Lab and ANPL-Autonomous Navigation) בטכניון. בשני המקרים נעשה שימוש ברובוט קרקעי מסוג Pioneer, המצויד במד לייזר או במצלמה, אשר פועל בסביבה דו משמעית ידועה (כלומר, מפה נתונה). הרובוט מתעורר לחיים בלי לדעת את מיקומו על גבי המפה, ועקב דו המשמעות, ה-prior belief מתקבל להיות מסוג GMM. על הרובוט להגיע לנקודה מסוימת במפה, אשר נקבעת מראש. אלגוריתם ה-

תקציר

תכנון וקבלת החלטות תחת אי ודאות הינן בעיות יסוד בתחום הניווט האוטונומי. בנוכחות אי ודאות, שיכולה לנבוע למשל מתנועה סטוכסטית של הרובוט וחישה לא מושלמת, וקטור המצב של הנעלמים בבעיה, כגון מצב הרובוט במונחים של מיקום ואוריינטציה, אינו ידוע וניתן לייצוג על ידי פונקציית פילוג הסתברותית (pdf), הידועה גם כאמונה (belief). בעיית התכנון המתאימה הינה תכנון במרחב הסתברותי (BSP, belief space planning). מסגרת זו ישימה למגוון רב של תחומים ואפליקציות, כגון מיפוי ואיכון בו זמני אקטיבי (Active SLAM), חישה אקטיבית, תכנון אינפורמטיבי, וואריאציות נוספות של ניווט אוטונומי.

מרכיב חיוני הן בהסקה/חישה והן בתכנון במרחב הסתברותי הינו שיוך מידע (DA, data association), כלומר, קביעת שיוך (או התאמה) נכון בין תצפיות למשתנים אקראיים מתאימים. עם זאת, סביבות בעולם האמיתי נוטות להיות דו משמעויות, דבר, אשר בנוכחות גורמים שונים של אי וודאות, הופך את מטלת החישה (perception) למאתגרת. לדוגמה, אובייקט יכול להראות דומה לאובייקט אחר מנקודות מבט מסוימות, ואילו התאמה מוצלחת של שתי תמונות משני מקומות שונים בעלי חזות דומה (למשל שני בניינים או מסדרונות שנראים דומה) תצביע באופן שגוי על שני המקומות כעל מקום אחד. שיוך מידע לא נכון, אם כן, יכול להוביל לתוצאות הרסניות, ולכן ישנה חשיבות רבה בפיתוח גישות להסקה ולתכנון במרחב הסתברותי אשר אינן מניחות ששיוך המידע פתור ומושלם.

למרות זאת, עד כה, קהיליית המחקר ברובוטיקה התרכזה בפיתוח גישות מסוג זה רק עבור בעיית ההסקה, דבר שהניב גישות אופטימיזציה מבוססת מודלים גרפיים המתאימות להיות חסינות לשיוך מידע שגוי (outliers). כאמור, גישות אלו מטפלות בבעיית ההסקה, כלומר, בהנחה שפעולות הרובוט נתונות (למשל מחושבות על ידי תהליך נפרד או נקבעות על ידי מפעיל אנושי). מצד שני, גישות BSP קיימות מניחות שיוך המידע נתון ומושלם; הנחה זו מאפשרת לחזות איך האמונה (belief) תתפתח בעקבות פעולה עתידית אפשרית כזו או אחרת של הרובוט, ובעקבות כך לבחור את הפעולה המיטבית עבור פונקציית מחיר נתונה. אף על פי כן, בתרחישים מציאותיים יש לעתים איזה שהיא דרגה של דו משמעות ו־perceptual aliasing, דבר המקשה להצדיק את ההנחה לגבי AD מושלם.

במחקר זה מפותחת גישת BSP בה אין הנחה ש־AD נתון ומושלם, וזאת בניגוד לגישות BSP הקיימות כיום. אנחנו מניחים כי הרובוט פועל בסביבה ידועה חלקית או ממופת מראש, אשר יכולה להיות דו משמעית (למשל שני מסדרונות שנראים דומה). הרובוט מקבל תצפיות של אזורים או אובייקטים בסביבה, ומשתמש בתצפיות אלו על מנת לעדכן את האמונה (belief) שלו לגבי מיקום ומצב זוויתי (ואולי נעלמים נוספים בבעיה). עקב דו משמעות ו־perceptual aliasing, ישנן מספר היפותזות לגבי המצב האמיתי של הרובוט, ועל כן אנחנו ממדלים את האמונה תוך שימוש ב־Gaussian mixture model (GMM). כמו בכל בעיית BSP, אנחנו מעוניינים למצוא את הפעולה המיטבית מתוך אוסף פעולות

המחקר נעשה בהנחיית פרופסור ואדים אינדלמן בפקולטה להנדסת אירונותיקה וחלל

תודות

אני מודה לטכניון על התמיכה הכספית הנדיבה בהשתלמותי.

שילוב אסוציאציית מדידות בתכנון במרחב הסתברותי עבור ניווט אוטונומי רובסטי

חיבור על מחקר

לשם מילוי חלקי של הדרישות לקבלת התואר
מגיסטר למדעים בהנדסת חלל

אנטוני טומאס

הוגש לסנט הטכניון – מכון טכנולוגי לישראל
שבט תשע"ז חיפה פברואר 2017

שילוב אסוציאציית מדידות בתכנון במרחב הסתברותי עבור ניוט אוטונומי רובסטי

אנטוני טומאס