STAIRS 2016
D. Pearce and H.S. Pinto (Eds.)

doi:10.3233/978-1-61499-682-8-87

Formal and Data Association Aware Robust Belief Space Planning

Shashank PATHAK [a,1], Sadegh SOUDJANI [b], Vadim INDELMAN [a] and Alessandro ABATE [b]

[a] *Technion - Israel Institute of Technology, Haifa 32000, Israel*
[b] *Department of Computer Science - University of Oxford, Oxford OX1 3QD, U.K.*

Abstract. State-of-the-art belief space planning (BSP) approaches assume data association to be solved or given. Some of the current authors have recently proposed a relaxation of this assumption, resulting in a more general framework of belief space planning where data association is incorporated within the belief (DA-BSP). Unfortunately, this can quickly become intractable under non-myopic planning. In this work, we seek to harness recent approaches in formal methods (specifically, linear temporal logic in the context of planning under uncertainty), to obtain formal-DA-BSP, an approach that incorporates high-level domain knowledge, to obtain more tractable planning. Thanks to generalised form of specification, the framework can also incorporate other complexities including explicit collision probability and determining planning horizon. The initial concepts are shown in an abstracted example of a robot janitor lost in one of the two floors.

Keywords. belief space planning, data association, probabilistic inference, formal methods

1. Introduction

Belief space planning (BSP) seeks to solve a planning problem on an underlying partially observable state space, through Markovian assumption i.e. a partially observable Markov decision process (POMDP). As is common in the literature(see [1]), we shall denote the observation by z and underlying unobservable states as x. Both of these are commonly assumed to be continuous variables. Evolution of x – the motion model – and of z conditioned on x – the observation model, are assumed to be given, along with an additive white Gaussian noise. Here, the belief is actually a probability distribution over states (strictly speaking a probability density and probability mass for continuous and discrete state cases respectively, with appropriate Lebesgue measure for the former), and the posterior of such a distribution is expressed as a function of control actions and of statistics (mean and covariance) of both the models mentioned before. Maximising the log-likelihood of this posterior, and averaging over all observations, planning could be framed as an optimisation problem in control actions. In general, this optimisation is multi-objective, catering to different needs such as reaching the goal state(s), reducing

[1] Corresponding Author: This research is partially funded by Israel Science Foundation (ISF) and by Technion Autonomous Systems Program (TASP); E-mail: shashank@technion.ac.il

the overall uncertainties, avoiding the collision with obstacles etc. State-of-the-art BSP approaches assume data association to be solved i.e., when an observation z is obtained while observing any scene A_i from unobservable pose x, we assume this A_i is known and given. This along with the Gaussian form of models, ensures that under *maximum likelihood assumption*, the posterior $\mathbb{P}(x|z)$ is still a Gaussian. Naturally, such a succinct representation is one of the prominent reason for scalable applications of these BSP approaches (as compared against the general POMDP approaches). However, assuming data association solved is a significant restriction which we seek to relax in an approach called data-association aware BSP or in short, DA-BSP (see [17]). But the focus of the current work is not DA-BSP, but rather how formal methods could be harnessed to provide a robust DA-BSP, where robustness implies both more scalable DA-BSP as well as being resilient to further complications such as wrong associations, presence of obstacles or infeasible planning under given parameters.

Effect of growing uncertainty (of state) could be implicitly incorporated in a cost model (see for e.g., [20]) or it could be (additionally) constrained to lie within a threshold determined *apriori*. Setting such a threshold is non-trivial as it often will be an artefact of the domain considered; for e.g., as discussed in [12] recently. Similarly, the effect of collision with the obstacles can be dealt either within the objective function (first proposed in [7]) or as a hard-constraint or even in probabilistic terms (called *chance constraints* [4]). Latter is more suitable when a completely collision-free path - one that deterministically avoids collision [8] - is hard (or impossible) to find *and* where a sufficiently safe (w.r.t. a threshold on collision probability) control strategy is acceptable. Recent instances of such an approach are [2], [3], [6], [22]. As a corollary, it is possible that for an arbitrarily small threshold, no safe control strategy exists.

In order to incorporate effect of uncertainty as well as of collision within the objective function, most current approaches take a convex combination of these two objectives where the weights are determined *apriori*. Though this results in an efficient method to the solution, it is undesirable (and even inaccurate) in many scenarios. Note that the constraints over uncertainty may render a particular desired plan infeasible, and it is only implicitly seen in quality of the plans. Consider, safety-first scenario where it would make more sense to have constraints, possibly probabilistic ones, based on collision. Another limitation of such approaches is that collision is often determined in simplistic low-level abstraction, for e.g. based on area of overlap of the two polyhedral shapes (see [19]). This is in contrast with high-level definition of safety that would be provided in many real-world robotics applications for e.g. a robot in the kitchen that needs to pass through several known objects, and avoid other objects with different degrees of caution (i.e. collision threshold is object dependent). Moreover, the presence of obstacles, may also assist in planning through more informed localisation (such as using the fact that current belief can not penetrate a dense obstacle or that complete occlusion implies that the object can not lie outside the shadows of the occluding obstacle).

When data association is considered within the belief space planning, the resulting belief is no longer a Gaussian and is in fact a Gaussian Mixture Model (GMM). As a result, maintaining parametric representation of posteriors, especially under non-myopic planning becomes intractable quite quickly. Though usual techniques of pruning (say based on observation likelihood of z that generated the posterior) can ameliorate the problem, more efficient solutions can be obtained by harnessing semantic relations between the scenes against which data association is sought. Again, this can be represented in a principled approach using formal methods and some of the recent techniques, that we will be mentioning here.

Hence in this preliminary work, we propose to harness formal methods (such as formal verification and model-checking) to obtain a general framework to tackle challenges of BSP mentioned above, especially in the context of beliefs being GMM. Firstly, it is a more general approach for BSP, freeing the designer of the burden of low-level planning details. Secondly, it is a principled approach to reason about the belief in efficient as well as explicit fashion. Lastly, it can be easily extended to more complex domains (such as belief-space being a non-gaussian) as well as more expressive specifications for planning and its objective. [15] was one of the prominent work to utilise temporal logic for synthesising correct plan. In this approach, a high-level task specification based correctness is ensured for a low-level motion planning based control policy. [11] and [13] seeks to revise the plans as more information about a partially known environment is made available or sensor noise is accounted for, while [23] ensures robustness by considering interval MDP. Synthesis with linera temporal logic (LTL) for optimal control of MDP is proposed by [5] whereas [10] analyze this from probably approximately correct (PAC) setting. Very recently, [16] consider using co-safe LTL to synthesise infeasible plans by allowing for violations and obtaining least violating plans. [9] uses similar approach, though under active semantic SLAM, where the motion model is deterministic but the semantic map is uncertain. They convert the problem of partially observable landmarks (hence with uncertain pose after the inference step), to a deterministic problem over set of trajectories that lie within a δ-confidence region, that satisfies the specification.

Structure of the paper: First, we briefly draw link between POMDP and usual BSP as well as data association aware BSP (DA-BSP). Thereafter, we frame a more general form of belief space planning that harnesses formal methods approaches and techniques, named *formal-DA-BSP*. Here we also note its particular importance for DA-BSP, and summarise the over all algorithm. Finally, we take a simple but illustrative example of the kidnapped robot problem, analysing various facets of formal-DA-BSP.

2. Belief space planning and data association

A POMDP is a tuple $\langle \mathscr{S}, \mathscr{A}, \Omega, T, O, r, b_0 \rangle$ where symbols stand for state space, action space, observation space, transition function, observation function, reward function and initial belief (see any standard text or works like [14]). Since at any time-step k, the actual underlying state is unknown, the system should reason about probability distribution over all states. The belief of being in state x, is represented as $b(x)$. After each observation this belief (assuming $'$ denotes succession in the time-step) is updated as:

$$b^{u,z}(x') = \frac{O(x',u,z)}{\mathbb{P}(z|u,b)} \cdot \sum_{x \in \mathscr{S}} T(x,u,x')b(x) \quad (1)$$

which implies:

$$b^{u,z}(x') \propto O(x',u,z) \cdot \sum_{x \in \mathscr{S}} T(x,u,x')b(x) \quad (2)$$

Note that in the notations followed in this paper, this can be written as:

$$b[X_{k+1}|z_{k+1},u_k] \propto \mathbb{P}(z_{k+1}|x_{k+1}) \cdot \mathbb{P}(x_{k+1}|x_k,u_k)\mathbb{P}(x_k) \tag{3}$$

where X_k denotes all states up to the step k.

Thus, BSP can then be viewed as a special case of belief MDP where due to Gaussian nature of the prior $\mathbb{P}(x_{k+1}|x_k)$ and the likelihood $\mathbb{P}(z_{k+1}|x_{k+1})$, the posterior $\mathbb{P}(x_{k+1}|z_{k+1})$ remains Gaussian as well. This affords us an efficient and scalable method to perform (local) belief updates.

2.1. Updating belief locally

Due to the Gaussian assumption on motion model and observation model (denoted by f and h respectively), we can write the belief as:

$$b[X_{0:k}]\mathbb{P}(x_{k+1}|x_k,u_k)\mathbb{P}(z_{k+1}|x_{k+1}) \doteq \mathcal{N}(\hat{X}_{0:k+1|k},\Sigma_{0:k+1|k})$$

where $.|k$ denotes history of control actions, poses and observations up to the step k.

The maximum a posteriori (MAP) estimate of belief is then obtained as:

$$X_k^* = \underset{X_k}{\arg\min}\left(\| X_k - \hat{X}_k \|_{\Lambda_0}^2 + \| f(x_k,u_k) - x_{k+1} \|_{\Omega_w}^2 + \| h(x_k) - z_{k+1} \|_{\Omega_\omega}^2\right) \tag{4}$$

here $\{\Lambda_0, \Omega_w, \Omega_\omega\}$ are information matrices for goal attainment, motion model and observation model, respectively while $\|y\|_.$ denotes Mahalanobis norm of y with respect to these matrices.

Using first order Taylor expansion around the nominal point $\hat{X}_k$, we have first term from $X_k = \hat{X} + \Delta X_k$, while the second approximates to $\| (f(\hat{x}_k,u_k) + \nabla_x f(\hat{x}_{k+1})\Delta x_k) - (\hat{x}_{k+1} + \Delta x_{k+1}) \|^2$ and similarly the third to $\| (h(\hat{x}_{k+1}) + \nabla_x h(\hat{x}_{k+1})\Delta x_{k+1}) - z_{k+1} \|^2$. Now rearranging this in matrix notation gives us an L_2-norm minimisation of form:

$$\begin{pmatrix} \Lambda_0{}^{\frac{1}{2}} & 0 \\ \Omega_w^{\frac{1}{2}}\nabla_x f_{k+1} & -1 \\ 0 & \Omega_v^{\frac{1}{2}}\nabla_x h_{k+1} \end{pmatrix} \begin{pmatrix} \Delta X_k \\ \Delta x_{k+1} \end{pmatrix} - \begin{pmatrix} 0 \\ \Omega_w^{\frac{1}{2}}(f(\hat{x}_k,u_k) - \hat{x}_{k+1}) \\ \Omega_v^{\frac{1}{2}}(h(\hat{x}_{k+1}) - z_{k+1}) \end{pmatrix}$$

Dropping the indices for clarity, we denote the minimisation as $\| \mathscr{A}\Delta X - b \|_2$, and use left inverse to get the solution as: $\Delta X = (\mathscr{A}^T\mathscr{A})^{-1}\mathscr{A}^T b$. Furthermore, ΔX is a function of both z as well as u where as $\mathscr{A}$ is independent of z. Since $\hat{X}$ is also a nominal point, the right most term becomes, $\begin{pmatrix} 0 \\ 0 \\ \Omega_v^{\frac{1}{2}}(h(\hat{x}_{k+1}) - z_{k+1}) \end{pmatrix}$.

Thus, any belief can be updated in the view of the control action u_k and the resulting future observation z_{k+1}. Once that is done, an objective function can be defined over such a belief. More precisely:

$$J(u_k) \doteq \int_{z_{k+1}} \overbrace{\mathbb{P}(z_{k+1} \mid \mathscr{H}_{k+1}^-)}^{(a)}\, c\left(\overbrace{\mathbb{P}(X_{k+1}|\mathscr{H}_{k+1}^-, z_{k+1})}^{(b)}\right) \tag{5}$$

Intuitively, for a particular value of z_{k+1}, (a) represents likelihood of such an observation when both the history $\mathcal{H}$ (denoting past observations and controls) and control u_k are known, while (b) is a conditioned posterior belief *given* this specific z_{k+1}.

2.2. Incorporating data association within BSP

We can now summarise the effect of considering data association within BSP. This relaxation under assumption of a unimodal prior and a GMM prior was proposed in [18] and [17] respectively. This section summarises parts taken from [17]. Note that the term (a) of the eq. 5 can be computed using the observation model and the motion model. In fact,

$$\mathbb{P}(z_{k+1}|\mathcal{H}_{k+1}^{-}) = \int_x \mathbb{P}(z_{k+1}|x, A^{tr}\mathcal{H}_{k+1}^{-}) b[x_{k+1}^{-}=x]. \tag{6}$$

where $\mathcal{H}_{k+1}^{-}$ denote history of control actions and of observations up to the step k, $b[x_{k+1}^{-}]$ is the propagated belief (utilising the motion model $\mathbb{P}(x_{k+1}|x_k)$) and A^{tr} denotes the true data association i.e., indeed A^{tr} is the correct scene which when viewed from x_{k+1} would likely result in the observation z_{k+1}. Similarly, the posterior i.e., the term (b) in the eq. 5 can be computed as:

$$b[X_{k+1}] \propto b[X_{k+1}]\mathbb{P}(z_{k+1} \mid x_{k+1}, A^{tr}) = \mathcal{N}(\hat{X}_{k+1}, \Sigma_{k+1}) \tag{7}$$

If we do not assume the data association to be solved, then in BSP the conditioning (of posterior in eq. 5) has important consequences; there is no longer a single true scene A^{tr}. Instead, we have to reason over all possible scenes, $A_i \in \{A_{\mathbb{N}}\}$. It can then be shown that these terms (such as eq. 7) are not a Gaussian but a mixture of Gaussians. Thus, the posterior is a Gaussian Mixture Model (GMM).

Therefore, for a single planning step, the objective function (5) can be re-written as

$$J(u_k) = \int_{z_{k+1}} \Big(\sum_{i\in\{A_{\mathbb{N}}\}} w_{k+1}^{i} \Big) \cdot c\left(\sum_{r}^{M_{k+1}} \xi_{k+1}^{r} b[X_{k+1}^{r+}] \right). \tag{8}$$

where $\{A_{\mathbb{N}}\}$ is the set of all possible data-associations while the prior $b[X_{k+1}]$ is a GMM with M_{k+1} components. The scalar weights w_{k+1}^{i} denote observation likelihood due to the data association with A_i. In order to make DA-BSP practical, it is crucial to keep the number of components, M_k low. However, this also depends on the degree of perceptual aliasing in the environment.

Note that, if we consider the belief at planning time k to be a Gaussian $b[X_k] = \mathcal{N}(\hat{X}_k, \Sigma_k)$, then under this setting, each of the components $b[X_{k+1}^{i+}]$ in the mixture pdf can be written as $b[X_{k+1}^{i+}] \propto b[X_k]\mathbb{P}(x_{k+1} \mid x_k, u_k)\mathbb{P}(z_{k+1} \mid x_{k+1}, A_i)$. It can be easily shown that the above belief is a Gaussian $b[X_{k+1}^{i+}] = \mathcal{N}(\hat{X}_{k+1}^{i}, \Sigma_{k+1}^{i})$ and one can find its first two moments via MAP inference. Hence, the mixture of posterior beliefs in the cost $c(.)$ from Eq. (8) is now a mixture of Gaussians:

$$\sum_i \tilde{w}_i b[X_{k+1}^{i+}] = \sum_i \tilde{w}_i \mathcal{N}(\hat{X}_{k+1}^{i}, \Sigma_{k+1}^{i}). \tag{9}$$

2.3. Designing a Specific Cost Function

Standard cost functions typically have terms related to control actions c_u, distance to goal c_G and uncertainty c_Σ, see e.g. [21,12]. In contrast with usual BSP, the posterior belief in DA-BSP is a GMM $\sum_i \tilde{w}_i \mathcal{N}(\hat{X}^i_{k+1}, \Sigma^i_{k+1})$, see Eq. (9). Hence, based on different criteria to evaluate such mixtures, we can have different types of cost function, such as, taking the worst-case covariance among all covariances Σ^i_{k+1} in the mixture, e.g. $\Sigma = \max_i\{tr(\Sigma_i)\}$, or collapsing the mixture into a single Gaussian $\mathcal{N}(.,\Sigma)$ (see e.g. [1]) while maintaining the cost due to uncertainty as $c_\Sigma = trace(\hat{\Sigma})$.

In the presence of possibly multi-modal beliefs, we can prefer control actions that lead to unambiguous situation, by penalising those actions that lead to ambiguity. For a fixed number of modes, this cost to ambiguity should penalise a distribution for being closer to a uniform distribution. One such method is to take Kullback-Leibler divergence $KL_u(\{\tilde{w}_i\})$ of these weights, with respect to the uniform distribution such as by considering $c_w(\{\tilde{w}_i\}) \doteq \frac{1}{KL_u(\{\tilde{w}_i\})+\varepsilon}$ where $0 < \varepsilon \ll 1$ is to make the term well-defined. For a variable number of modes, higher cost should be associated with higher number of modes in a belief. With user-defined weights M_u, M_G, M_Σ and M_w, the overall cost then can be defined as a combination

$$c \doteq M_u c_u + M_G c_G + M_\Sigma c_\Sigma + M_w c_w, \tag{10}$$

3. Formal methods in belief space planning

Till now, we assumed that from the perspective of planning, all scenes $\{A_\mathbb{N}\}$ are equally significant, and naturally this leads us to computational hardships. However, explicit reasoning of these association would enable our optimisation routine to harness richer information available. For e.g., in search and rescue mission, it might be more important to associate humans or objects likely used by them, than to get natural features. We now first describe how a formal methods based approach can be augmented with the belief space planning in general and then in DA-BSP in particular. Let us first define very briefly the syntax and semantics of such a formal specification.

Syntactically, a set of LTL formula over a finite set of atomic propositions AP, is defined as:

- $p \in AP \doteq p$ is LTL formula.
- if Ψ and Φ are LTL formula $\neg\Psi, \Phi \vee \Psi, \mathscr{X}\Psi$ and $\Psi\mathscr{U}\Phi$ are LTL formulae.
- Boolean operators are $\neg, \vee, \wedge, \top, \bot$.

Semantics of an LTL formula is defined over infinite traces over 2^{AP} and satisfaction relation $\models$ is defined over an ω-word $w = a_0, a_1, \ldots$ (and $w_1 = a_1, \ldots$) as follows:

- $w \models p$ if $p \in a_0$
- $w \models \neg\Psi$ if $w \not\models \Psi$
- $w \models \Phi \vee \Psi$ if $w \models \Phi \vee w \models \Psi$
- $w \models \mathscr{X}\Psi$ if $w_1 \models \Psi$
- $w \models \Phi\mathscr{U}\Psi$, $\exists i,\ i \geq 0$ s.t. $w_i \models \Psi \wedge \forall k,\ 0 \leq k < i,\ w_k \models \Phi$

Additional operators are $\diamond\phi = \top\mathscr{U}\phi$ (eventually) and $\Box\phi = \neg\diamond\neg\phi$ (globally). Note that, if reaching the task is defined as being in the vicinity of a goal state x_G under a metric $d(.)$, then the corresponding LTL formula becomes $\phi = \diamond(\|x - x_G\|_d < \delta)$ where $\delta \in \mathbb{R}$ defines the neighbourhood. Analogously, a safety requirement could be avoiding a collision at all times of the trajectory. The corresponding LTL formula then can be $\phi = \Box\left(\|x - x_{obs,x}\|_d > \delta\right)$ where $\delta \in \mathbb{R}$ is the safe distance from the closest (in the sense of metric d) obstacle to x denoted by $x_{obs,x}$. Usually, a labeling function $L : X \mapsto 2^{AP}$ is defined over the states, to define set of states where the atomic proposition evaluates to true. This can be done analogously to the non-standard case such as in [9] where labels operate over set of all possible maps as well, or here, where labels operate over belief states.

Another consequence of Gaussian beliefs in BSP is that certain properties like 95% confidence region for a simple reachability, can be determined efficiently and analytically. Therefore, in contrast with the cost function mentioned in Sec. 2.3, we can frame a more general optimisation problem as:

$$\begin{aligned} \underset{u}{\text{minimize}} \quad & \mathrm{J}(b_u^L, \Psi_1) \\ \text{subject to} \quad & b_u^k \models \Psi_2,\ k \in \{1, 2, \ldots, L\} \end{aligned} \tag{11}$$

where b_u^i represents a belief at step i, under control action u. In other words, we incorporate a hard-constraint via an LTL formula Ψ_2 whereas infeasible formula Ψ_1 may have an associated cost included explicitly in the objective function.

3.1. Incorporating collisions

Under the presence of obstacles, the local optimisation approach mentioned earlier in Sec. 2.1, needs to consider the homotopy class of non-colliding trajectories. This, like many other approaches, considers collision implicitly. The basic idea is to ensure no collision along a given control trajectory; one computes analytically the probability of remaining safe (i.e., away from collision) and maximises it for the path. Like in most multi-objective optimisation, a convex combination of this along with other costs can be taken. In order to improve the efficiency, [21] considers second-order approximation of this analytical function. They also relax the maximum likelihood assumption when incorporating the observation. Since in DA-BSP, each posterior conditioned on a specific association, is still a Gaussian (see eq. 9), this can be done similarly.

However, we suggest that this can be generalised using formal methods, where the objective or the constraints or both are specified using LTL. For e.g., [5] uses this approach for optimal planning in MDP considering LTL constraints. Taking cue from that, in BSP, when the underlying POMDP is viewed as the belief-MDP, a rich set of reachability properties can be defined in LTL. Thus, at any planning iteration, given the composed belief $b[X_k]$, we can define the LTL formula Ψ which encapsulates such notion of safety. For e.g., to ensure that the belief is safe within 2-standard deviation from the nominal value, one creates a corresponding labelling function for each node of belief-MDP. In case of DA-BSP, this corresponds to conjunction $\bigwedge_{i \in \{A_{\mathbb{N}}\}} \phi(b[X_{k+1}|\mathscr{H}, A_i])$ where labelling can be computed efficiently due to Gaussian nature of the conditional probabil-

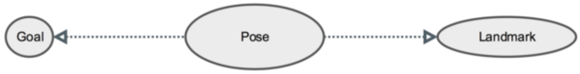

Figure 1. Schematic of an infeasible planning problem. The landmark uncertainty is high enough to deny the robot sufficient localization after loop-closure, while moving directly towards the goal fails to contain the uncertainty within a maximum allowed limit.

ity distribution. This is similar to the numerous approaches that seek to integrate high-level specification-based task planner and low-level motion planning. More specifically, the automaton encoding the LTL specification is composed with the underlying system dynamic. Planning in this product space, can then yield conformant plans.

3.2. Incorporating uncertainty budget and infeasibility

In many scenarios, the objective function has competing goals of uncertainty reduction and reaching to the target as soon as possible. For e.g., a robot may be able to reduce its pose uncertainty by visiting a landmark, whereas it might need to move away from this landmark in order to reach the target. As is seen in the Fig. 1, under a maximum *uncertainty budget* allowed for any control trajectory, simultaneously achieving these two objectives might be infeasible. Recognising this, [12] suggests ad-hoc approach of setting the budget apriori. Very similar to the approach mentioned previously in Sec. 3.1, this can be obtained through integration of high-level task planner (where such requirements are specified formally) and low-level motion planner.

In more interesting case, when such a specification is infeasible, a recently proposed approach of *co-safe* LTL can be harnessed to obtain least violating policy (see e.g., [16] in this context). [9] propose it for active semantic SLAM i.e., where the classes as well as probability distribution of uncertain landmarks are known, but the objective is localization of the pose as well as these landmarks. They show that the general stochastic optimization problem, similar to eq. 5 with an additional scalar term proportional to probability of satisfying the specification, can be transformed into a deterministic optimization when the candidate paths belong to a δ-confidence region of nominal positions of these landmarks (such that the nominal path satisfies the specification). This approach is applicable where the only source of uncertainty is in the location of landmarks. In case of DA-BSP, the uncertainty arises out of motion model, and in absence of this uncertainty, data association can be resolved trivially. However, like in [9], we can convert a multi-objective (convex in objectives) optimization problem to one over general cost minimization and probabilistic satisfaction of the safety specification.

3.3. Reasoning over the associations

Note that since labelling of the belief does not require either a specific formula (as long as it is within the language) or a specific structure of the belief (though assuming such a structure could make labelling easier), the formal methods approach can be harnessed to also reason over data association. As before, if certain specification in terms of data association is infeasible, we can resort to co-safe LTL specification to obtain minimally violating control policy.

3.4. Summary of formal BSP

In summary of the discussion on formal BSP and also DA-BSP, we note that LTL specification can be incorporated into belief space planning, to account for collisions, uncertainty budget, infeasibility as well as various data association. Table 1 shows how each of them can be specified. Note that the state $\hat{x}$ denotes the mean value around which the current belief is distributed. Here, x^i_{Ob} denotes the position of i^{th} obstacle while $\sigma_.$ denotes a threshold fixed a-priori. Even though the belief is over entire state space of x, its succinct Gaussian form enables us to compute these labels over the belief, quite efficiently. In the case of DA-BSP, since the belief is a GMM, the properties would need to be changed slightly. For e.g., reachability to the goal, may become $p_{goal} = \underset{j}{min}\ \|\hat{x}_j, x_G\|_d < \sigma_g$, implying one of the modes is close enough to the goal state. In addition to this, DA-BSP allows us to reason about the associations directly within the belief state. For e.g., to allow active disambiguation, the LTL specification should guide towards unimodal beliefs. Similarly, to avoid having numerous data-association beliefs, a safety specification of the form $p_{pruned}|\{A_{\mathbb{N}}\}| < \sigma_N,\ \sigma_N \in \mathbb{N}$ can be used.

Table 1. Examples of formal BSP and DA-BSP

Planning	Property	LTL-formula	Comment	Example
BSP	Reaching target	$\diamond p_{goal}$	eventually the goal is reached	$p_{goal} = \|\hat{x} - x_G\|_d < \sigma_g$
	Avoiding obstacle	$\square p_{safe}$	obstacles are avoided at each step	$p_{safe} = \underset{i}{min}\ \|\hat{x} - x^i_{Ob}\|_d > \sigma_{safe}$
	Bounded uncertainty	$\square p_{unc}$	pose uncertainty within a bound	$p_{unc} = tr(\Sigma_x) < \sigma_\Sigma$
DA-BSP	Reaching target	$\diamond p_{goal}$	goal is reached	$p_{goal} = \underset{j}{min}\ \|\hat{x}_j, x_G\|_d < \sigma_g$
	Avoiding obstacle	$\square p_{safe}$	obstacles are avoided at each step	$p_{safe} = \underset{i,j}{min}\ \|\hat{x}_j - x^i_{Ob}\|_d > \sigma_{safe}$
	Active disambiguation	$\diamond p_{disamg}$	eventually, disambiguation	$p_{disamg} = |\{A_{\mathbb{N}}\}| = 1$
	Efficient propagation	$\square p_{pruned}$	parsimonious data association	$p_{pruned}|\{A_{\mathbb{N}}\}| < \sigma_N$

3.5. Overall algorithm

The algorithm that incorporates DA-BSP as well as formal methods approaches is shown in Alg. 1. For brevity, it is shown for a single step of planning, though it is straightforward to apply it in non-myopic setting by calling it recursively. First the given prior is composed with the motion model to obtain the propagated belief. The observations are samples drawn from this propagated belief. The cost function is set to 0 and specification is assumed to be initially satisfied. Consequently, for each sampled observation, all possible data association in the current step k and future $k+1$ is considered, resulting in the nested loops shown. The weights are computed as described before. Each conditioned posterior is a Gaussian belief against which specification could checked and appropriate labels be attached. Finally, when all associations are considered, we obtain a multi-modal mixture for each observation z.

4. Simulated examples to illustrate formal DA-BSP

We consider a janitor robot that is lost within one of the two office floors, as shown in the Fig. 2. Each floor has cubicles at the top corners (close to B and D) and the elevator (close to F) at the bottom middle. The two floors are significantly aliased, except the left-

Algorithm 1 Formal and data association aware belief-space planning (formal-DA-BSP)

Input: Current GMM belief $b[X_k]$ at step-k, history $\mathscr{H}_k$, action u_k, scenes $\{A_\mathbb{N}\}$, event likelihood $\mathbb{P}(A_i \mid \mathscr{H}_k, x)$ for each $A_i \in \{A_\mathbb{N}\}$, labelling F_{ϕ_1}, F_{ϕ_2}

$b[X_{k+1}^-] \leftarrow b[X_k]\mathbb{P}(x_{k+1} \mid x_k, u_k)$
$\{z_{k+1}\} \leftarrow$ `SimulateObservations`$(b[X_{k+1}^-], \{A_\mathbb{N}\})$
$J \leftarrow 0$
$F_{\phi_1} \leftarrow \top$
for $\forall z_{k+1} \in \{z_{k+1}\}$ **do**
  $w_s \leftarrow 0$
  for $i \in [1 \ldots |A|]$ **do**
    ▷ compute weight
    $w_{k+1}^i \leftarrow$ `CalcWeights`$(z_{k+1}, \mathbb{P}(A_i \mid \mathscr{H}_{k+1}^-, x), b[X_{k+1}^-])$
    $w_s \leftarrow w_s + w_i$
    for $\forall j \in [1, \ldots, M_k]$ **do**
      ▷ compute weight $\tilde{w}_{k+1}^{ij}$ for each GMM component
      $\tilde{w}_{k+1}^{ij} \leftarrow$ `CalcWeights`$(z_{k+1}, \mathbb{P}(A_i \mid \mathscr{H}_{k+1}^-, x), b[X_{k+1}^{j-}])$
      $\xi_{k+1}^{ij} \leftarrow \xi_k^j \tilde{w}_{k+1}^{ij}$
      ▷ Calculate posterior of $b[X_{k+1}^{j-}]$, given A_i
      $b[X_{k+1}^{ij+}] \leftarrow$ `UpdateBelief`$(b[X_{k+1}^{j-}], z_{k+1}, A_i)$
      $F_{\phi_1,k+1} \leftarrow$ `Check`$(b[X_{k+1}^{j-}], z_{k+1}, A_i)$
      $F_{\phi_1} \leftarrow F_{\phi_1} \wedge F_{\phi_1,k+1}$
    end for
  end for
  Prune components with weights ξ_{k+1}^{ij} below a threshold
  Construct $b[X_{k+1}^+]$ from the remaining M_{k+1} components
  $F_{\phi_2} \leftarrow$ `Check`$(b[X_{k+1}^+])$
  $c \leftarrow$ `CalcCost`$(b[X_{k+1}^+], F_{\phi_1}, F_{\phi_2})$
  $J \leftarrow J + w_s \cdot c$
end for
return J

middle of corridor which has a unique feature (a sofa) in the second floor. The objective of the robot is to localise itself, both within the floor (intra-floor) and between the floors (inter-floor). Initially, there are two intra-floor nominal positions at each floor, shown with letters A and C. In other words, the prior for robot pose is a multi-modal Gaussian with 4 modes. The set of robot control actions consist of $\mathscr{U} = \{u_{fwd}, u_{bwd}, u_{lft}, u_{rgt}\}$. For simplicity, these actions are assumed to be suitably abstracted so that non-holonomic constraints are not looked into. Also the nominal positions out of such control actions are in close vicinity of $\{A, B, \ldots, F\}$.

This simplistic example is sufficient to illustrate the effect of DA-BSP and formal reasoning, both in qualitative sense (through this schematics) as well as quantitative sense (through actual planning in a similar Gazebo based robotic world). For realistic simulation of the problem, we use Gazebo simulator along with the Pioneer robotic platform. This is also to facilitate the experiment with real Pioneer robot, which is planned in the future works. With reference to the examples in Table 1, note how each aspect of this example, can be framed using formal BSP and DA-BSP. For e.g., to ensure that the uncertainty is within the narrow corridor, a safety property $\square\left(\|\hat{x} - x_0\|_d < \frac{w_{corr}}{2}\right)$ is sufficient.

Adapting planning horizon: Initially, given an approximate map of the environment (say from pervious SLAM session), the robot could decide to plan for multiple steps

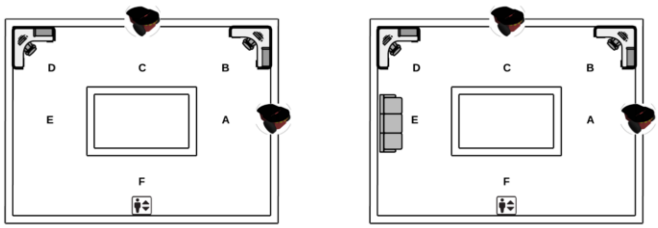

Figure 2. Schematics of a lost janitor robot (figure not to scale). The prior belief is a multimodal Gaussian, with 4 modes, two each floor. Note that there is significant aliasing between the floors.

before actually taking an action. If there were no inference involved, each motion would increase the uncertainty of the position and this may even have significant overlap with the narrow corridors. However, note that, in such planning scenario, inference is often called as a sub-routine and thereby reducing the uncertainty. Since in each planning step, we consider all possible data-association, such a collision is envisioned in some of the modes of the multi-modal posterior belief. This proves to be quite easy to be incorporated using the labelling function described in the section 3; for e.g. using L_2-norm as the metric. Once an appropriate safety criteria is formulated, the optimisation is performed within this feasible space of policies. In case of Fig. 2, this means that rather than moving from A to B (or C to D) in one planning action, the robot may perform smaller actions, to attain better localisation. However, due to the particular nature of the problem, this will not make planning infeasible in the sense of conflicting goals and uncertainty budget, discussed in Sec. 3.2.

5. Conclusion

Usual belief space planning assumes data association to be solved. However, in realistic applications such as navigation under significant uncertainty and perceptual aliasing, this is too restrictive an assumption and can also lead to catastrophic failure. Unfortunately, considering data association within belief space planning results in posterior becoming a non-Gaussian multi-modal distribution. Though under realistic assumption of mild perceptual aliasing, this approach can still be tractable, we propose using formal methods based approach to harness a richer class of dependency between scenes and differentiating important association from unimportant ones. Interestingly, this also provides a framework for a more general class of belief space planning where the cost function is aware of obstacles as well as the uncertainty budget, thereby avoiding infeasible planning. Though interesting, this is a very preliminary work towards an integrated approach of formal-DA-BSP, hence the future work is targeted towards application of these ideas on a real robotic platform while fully harnessing recent advances in formal methods in the context of path planning and advances in generalisation of belief space planning.

References

[1] Yaakov Bar-Shalom. *Tracking and data association*. Academic Press Professional, Inc., 1987.

[2] Lars Blackmore, Hui Li, and Brian Williams. A probabilistic approach to optimal robust path planning with obstacles. In *American Control Conference, 2006*, pages 7–pp. IEEE, 2006.

[3] Lars Blackmore, Masahiro Ono, Askar Bektassov, and Brian C Williams. A probabilistic particle-control approximation of chance-constrained stochastic predictive control. *Robotics, IEEE Transactions on*, 26(3):502–517, 2010.

[4] Abraham Charnes and William W Cooper. Chance-constrained programming. *Management science*, 6(1):73–79, 1959.

[5] Xuchu Ding, Stephen L Smith, Calin Belta, and Daniela Rus. Optimal control of markov decision processes with linear temporal logic constraints. *Automatic Control, IEEE Transactions on*, 59(5):1244–1257, 2014.

[6] Noel E Du Toit and Joel W Burdick. Probabilistic collision checking with chance constraints. *Robotics, IEEE Transactions on*, 27(4):809–815, 2011.

[7] Dieter Fox, Wolfram Burgard, and Sebastian Thrun. Active markov localization for mobile robots. *Robotics and Autonomous Systems*, 25(3):195–207, 1998.

[8] Thierry Fraichard and Hajime Asama. Inevitable collision states?a step towards safer robots? *Advanced Robotics*, 18(10):1001–1024, 2004.

[9] Jie Fu, Nikolay Atanasov, Ufuk Topcu, and George J Pappas. Optimal temporal logic planning in probabilistic semantic maps. *arXiv preprint arXiv:1510.06469*, 2015.

[10] Jie Fu and Ufuk Topcu. Probably approximately correct mdp learning and control with temporal logic constraints. *arXiv preprint arXiv:1404.7073*, 2014.

[11] Meng Guo, Karl H Johansson, and Dimos V Dimarogonas. Revising motion planning under linear temporal logic specifications in partially known workspaces. In *Robotics and Automation (ICRA), 2013 IEEE International Conference on*, pages 5025–5032. IEEE, 2013.

[12] Vadim Indelman, Luca Carlone, and Frank Dellaert. Planning in the continuous domain: A generalized belief space approach for autonomous navigation in unknown environments. *The International Journal of Robotics Research*, 34(7):849–882, 2015.

[13] Benjamin Johnson and Hadas Kress-Gazit. Analyzing and revising synthesized controllers for robots with sensing and actuation errors. *The International Journal of Robotics Research*, 34(6):816–832, 2015.

[14] Leslie Pack Kaelbling, Michael L Littman, and Anthony R Cassandra. Planning and acting in partially observable stochastic domains. *Artificial intelligence*, 101(1):99–134, 1998.

[15] Hadas Kress-Gazit, Georgios E Fainekos, and George J Pappas. Temporal-logic-based reactive mission and motion planning. *Robotics, IEEE Transactions on*, 25(6):1370–1381, 2009.

[16] Morteza Lahijanian, Shaull Almagor, Dror Fried, Lydia E Kavraki, and Moshe Y Vardi. This time the robot settles for a cost: A quantitative approach to temporal logic planning with partial satisfaction. In *AAAI*, pages 3664–3671, 2015.

[17] 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)*, September 2016.

[18] 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)*, May 2016.

[19] Sachin Patil, Jur Van Den Berg, and Ron Alterovitz. Estimating probability of collision for safe motion planning under gaussian motion and sensing uncertainty. In *Robotics and Automation (ICRA), 2012 IEEE International Conference on*, pages 3238–3244. IEEE, 2012.

[20] Robert Platt Jr, Russ Tedrake, Leslie Kaelbling, and Tomas Lozano-Perez. Belief space planning assuming maximum likelihood observations. 2010.

[21] Jur Van Den Berg, Sachin Patil, and Ron Alterovitz. Motion planning under uncertainty using iterative local optimization in belief space. *The International Journal of Robotics Research*, 31(11):1263–1278, 2012.

[22] Michael P Vitus, Wei Zhang, and Claire J Tomlin. A hierarchical method for stochastic motion planning in uncertain environments. In *Intelligent Robots and Systems (IROS), 2012 IEEE/RSJ International Conference on*, pages 2263–2268. IEEE, 2012.

[23] Eric M Wolff, Ufuk Topcu, and Richard M Murray. Robust control of uncertain markov decision processes with temporal logic specifications. In *Decision and Control (CDC), 2012 IEEE 51st Annual Conference on*, pages 3372–3379. IEEE, 2012.