

**Structure Aware Probabilistic
Inference and Belief Space
Planning with Performance
Guarantees**

Moshe Shienman

Structure Aware Probabilistic Inference and Belief Space Planning with Performance Guarantees

Research Thesis

Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy

Moshe Shienman

Submitted to the Senate
of the Technion — Israel Institute of Technology
Tamuz 5784 Haifa July 2024

This research thesis was done under the supervision of Associate Professor Vadim Indelman, in the Technion Autonomous Systems Program.

The author of this thesis states that the research, including the collection, processing and presentation of data, addressing and comparing to previous research, etc., was done entirely in an honest way, as expected from scientific research that is conducted according to the ethical standards of the academic world. Also, reporting the research and its results in this thesis was done in an honest and complete manner, according to the same standards.

Some results in this thesis have been published as articles by the author and research collaborators in conferences and journals during the course of the author's doctoral research period, the most up-to-date versions of which being:

Journal Articles

- [1] M. Shienman, A. Kitanov and V. Indelman, "FT-BSP: Focused Topological Belief Space Planning", *IEEE Robotics and Automation Letters (RA-L)*, vol. 6, no. 3, pp. 4744–4751, July 2021. DOI: 10.1109/LRA.2021.3068947.
- [2] M. Barenboim, M. Shienman and V. Indelman, "Monte Carlo Planning in Hybrid Belief POMDPs", *IEEE Robotics and Automation Letters (RA-L)*, vol. 8, no. 8, pp. 4410–4417, Aug. 2023.
- [3] M. Shienman, O. Levy-Or, M. Kaess, and V. Indelman, "Towards Real-Time Nonparametric Inference in High Dimensional State Spaces via Slices", To be submitted to *IEEE Transactions on Robotics (T-RO)*, 2024.

In Conference Proceedings

- [1] M. Shienman, and V. Indelman, "Nonmyopic Distilled Data Association Belief Space Planning Under Budget Constraints", in *International Symposium on Robotics Research (ISRR)*, Switzerland, Sep. 2022, pp. 102–118.
- [2] M. Shienman, and V. Indelman, "D2A-BSP: Distilled Data Association Belief Space Planning with Performance Guarantees Under Budget Constraints", in *IEEE International Conference on Robotics and Automation (ICRA)*, USA, May 2022, pp. 11058–11065. **Outstanding Paper Award Finalist**
- [3] M. Shienman, O. Levy-Or, M. Kaess, and V. Indelman, "A Slices Perspective for Incremental Nonparametric Inference in High Dimensional State Spaces", Accepted to *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, UAE, Oct. 2024.

Acknowledgements

First, I would like to thank my supervisor, Professor Vadim Indelman, for providing me with the opportunity to work under his guidance. His unwavering support, vast expertise, and uncompromising standards have been instrumental in shaping my development as a researcher. I truly appreciate the trust he has placed in me and am thankful for the belief he has shown in my capabilities.

I would also like to thank the Technion Autonomous Systems Program (TASP) and the Autonomous Navigation and Perception Lab (ANPL), along with all their dedicated students and staff members, for fostering a supportive environment that has significantly contributed to my academic journey. Furthermore, I would like to express my appreciation to Dr. Andrej Kitanov, Moran Barenboim and Ohad Levy-Or, with whom I had the pleasure to collaborate. Their contributions have enriched my research experience, and I am thankful for the valuable cooperation.

I would like to express my gratitude and love to my parents. Their immense encouragement, unwavering support, love, and guidance have shaped the person I am today.

Finally, to the love of my life and my best friend, Lital. None of this would have been possible without you. Your support and encouragement have been indispensable in helping me complete my studies. This achievement is as much yours as it is mine, and I am profoundly thankful for having you, Mila and Toya in my life.

The Generous Financial Help of the Technion is Gratefully Acknowledged.

Contents

List of Figures

List of Tables

List of Algorithms

Abstract	1
Notation and Abbreviations	3
1 Introduction	5
1.1 Motivation	5
1.2 Inference and Belief Space Planning	6
1.3 Contributions	7
1.4 Thesis Structure	8
2 Literature Survey	11
2.1 Topological Structures and the Focused Case	11
2.2 Robust Perception and Planning	13
2.3 Simplification and Performance Guarantees in BSP	15
2.4 Non-Parametric Inference	16
3 Focused Topological BSP	19
3.1 Problem Formulation and Notations	20
3.1.1 Focused Belief Space Planning	20
3.1.2 Belief Topology	23
3.1.3 Weighted Tree Connectivity	24
3.2 Planning using Topological Signatures in the Focused Case	25
3.2.1 Manipulating the Focused Objective Function $J_{\mathcal{H}}^F$	26
3.2.2 The Unfocused Augmented Graph $\mathcal{G}^{A,U}$	27
3.2.3 Weighted Tree Connectivity Difference Signature	27
3.2.4 Von Neumann Difference Signature	28
3.3 Experimental Results	31

4	DA-BSP with Performance Guarantees Under Budget Constraints	35
4.1	Background and Notations	36
4.1.1	Inference with Data Association	36
4.1.2	Data Association BSP	37
4.2	One Look-Ahead Step: The Myopic Case	38
4.2.1	Bounding the cost function	40
4.2.2	Bounding η	42
4.2.3	Simulating future observations Z_{k+1}	42
4.2.4	Experimental Results	43
4.3	Longer Planning Horizons: The Nonmyopic Case	45
4.3.1	Constructing the belief tree <i>skeleton</i>	45
4.3.2	Methodology	46
4.3.3	Inference without Budget Constraints	49
4.3.4	Hard budget constraints in inference	51
4.3.5	Information Theoretic Cost Function	52
4.3.6	Experimental Results	54
5	A Slices Perspective for Incremental Nonparametric Inference	57
5.1	Background and Notations	58
5.2	Methodology	61
5.2.1	Inference Using Slices	62
5.2.2	Incremental Inference with Early Stopping Heuristic	65
5.3	Experimental Results	66
5.3.1	Synthetic Dataset - Multi Modal Four Doors	66
5.3.2	Real World Dataset - Plaza	68
6	Conclusion	71
6.1	Future Research Directions	72
A	Proofs for Chapter 3	73
A.1	Theorem 1	73
A.2	Theorem 2	73
A.3	Lemma 1	74
A.4	Theorem 3	74
B	Appendix for Chapter 4	77
B.1	Proofs	77
B.1.1	Corollary 1	77
B.1.2	Theorem 4	77
B.1.3	Theorem 5	78
B.1.4	Corollary 2	78
B.1.5	Theorem 6	78

B.1.6	Theorem 7	79
B.1.7	Corollary 3	80
B.1.8	Theorem 8	80
B.1.9	Theorem 9	80
B.1.10	Corollary 4	81
B.1.11	Theorem 10	82
B.1.12	Corollary 5	83
B.2	Incremental Bounds Updates: The Myopic Case	83
B.2.1	The Normalization Term Bounds $\mathcal{LB}[\eta], \mathcal{UB}[\eta]$	83
B.2.2	The Cost Function Bounds $\mathcal{LB}[\mathcal{H}], \mathcal{UB}[\mathcal{H}]$	84
B.3	Incremental Bounds Updates: The Nonmyopic Case	86
B.3.1	The Normalization Term Bounds $\mathcal{LB}[\eta_{k+n}], \mathcal{LB}[\eta_{k+n}]$	87
B.3.2	The Cost Function Bounds $\mathcal{LB}[\mathcal{H}_{k+n}], \mathcal{LB}[\mathcal{H}_{k+n}]$	87
C	Proofs for Chapter 5	91
C.1	Lemma 2	91
	Bibliography	93
	Hebrew Abstract	i

List of Figures

3.1	<i>focused</i> vs <i>unfocused</i> BSP problems	19
3.2	Factor graphs and topological signatures	23
3.3	Measurement selection: correlation between topological signatures and marginal entropy	31
3.4	Active 2D pose SLAM: candidate paths	32
3.5	Active 2D pose SLAM: correlation between topological signatures and a focused objective function including bounds	33
4.1	BSP considering data association: the belief tree and the exponential growth	38
4.2	Myopic 2D active disambiguation	43
4.3	Planning under hard budget constraints in the myopic case	44
4.4	Performance guarantees using bounds in BSP	47
4.5	No budget constraint in inference	50
4.6	Hard budget constraints in inference	51
4.7	Simulation environments for the nonmyopic BSP case	54
4.8	Results: BSP without budget constraints	55
4.9	Results: BSP with budget constraints	56
5.1	Joint probability: a <i>slices</i> perspective	58
5.2	Factor graph and a Bayes net	59
5.3	Generating samples during variable elimination	63
5.4	Marginal posterior distributions comparison in the four doors synthetic dataset	67
5.5	Results: RMSE and run time comparison for the four doors synthetic dataset	67
5.6	Marginal posterior distribution comparison in the Plaza2 dataset	69
5.7	Results: RMSE and run time comparison for the Plaza2 dataset	69

List of Tables

3.1	Runtime comparison: evaluating topological signatures vs evaluating the <i>focused</i> objective function	31
4.1	Budget constraint cases in inference and planning	49

List of Algorithms

1	<i>Focused</i> topological BSP	30
2	Constructing the belief tree <i>skeleton</i>	46
3	Generic nonmyopic distilled data association BSP	48

Abstract

Intelligent autonomous agents and robots are increasingly utilized across various fields, influencing numerous aspects of our daily lives. Employed in various applications such as autonomous navigation, robotic surgery and automated warehousing, these agents are expected to operate reliably and efficiently under different sources of uncertainty, often with limited knowledge about the environment. There are numerous possible sources for such uncertainty, including dynamic environments in which unpredictable events might occur; noisy or limited observations stemming from physical constraints; and ambiguous measurements, wherein a particular observation could have multiple plausible interpretations. To handle these real-world scenarios, autonomous agents are required to reason over high-dimensional probabilistic states that account for the associated uncertainty. Specifically, such agents should possess the capability to engage in both inference, i.e. maintain a posterior probability distribution over the high-dimensional state given available information, and decision making under uncertainty where they should autonomously determine their next best actions. However, solving both these problems is computationally expensive and practically infeasible in real-world autonomous systems, where the agent is required to operate in real time, often using inexpensive hardware with limited resources. In this work we address these two problems and develop different approaches to reduce the computational complexity towards real time operation by leveraging different structures found in the underlying representations of these problems. In decision making under uncertainty, we propose solving simplified problems that are computationally much easier to compute, while providing performance guarantees. These guarantees involve either ensuring the same optimal solution as in the original problem or providing information regarding the loss in solution quality. Furthermore, we demonstrate how our simplified solutions are tailored and adjusted to real-world scenarios, where the autonomous agent has limited resources and must operate under strict computational budget constraints. To that end, we utilize both topological structures, associated with underlying graph representations of the problem, and specific structures modeled in posterior probability distributions when considering ambiguous measurements. In inference, we extract slices from structures that represent high-dimensional posterior probability distributions. These slices are utilized to efficiently approximate posterior distributions of any shape, facilitating real time operation in real-world scenarios where posteriors distributions are nonparametric.

Notation and Abbreviations

A	Incidence matrix
H_k	History at time step k
L_w	The reduced weighted Laplacian matrix
X_k^*	The MAP estimate of X_k
X_k	The robot's state vector at time step k
X_k^F	<i>focused</i> subset of states
X_k^U	<i>unfocused</i> subset of states
$Z_{1:k}$	All observations taken up to and including time step k
Λ_k	The information matrix of the ML estimation at time step k
β_k	Data association realization vector at time step k
\hat{L}_w	The reduced weighted normalized Laplacian matrix
\mathcal{H}	Entropy
\mathcal{LB}	Lower bound
\mathcal{UB}	Upper bound
\mathcal{U}_k	Set of candidate actions at time step k
c_l	Cost function at the l th look-ahead step
u_k	The action taken at time step k
$u_{0:k-1}$	All actions taken up to time step k
x^l	A landmark pose
x_k	The robot's state at time step k
M	Measurement Jacobian matrix
Σ	Noise covariance matrix
\mathcal{C}	Budget constraint
\mathcal{E}	A set of graph edges
$\mathcal{G} \doteq (\mathcal{V}, \mathcal{E}, w)$	Weighted topological graph. Edges \mathcal{E} are weighted according to w
\mathcal{V}	A set of graph nodes
$\tau_w(\mathcal{G})$	The Weighted Tree Connectivity (WTC) for a graph \mathcal{G}
θ_k	The robot's orientation at time step k
$b[X_k]$	The robot's belief at time step k given all actions and observations. Shortly written as b_k
b_k^s	A simplified belief at time step k

p_k	The robot's position at time step k
$t_w(\mathcal{G})$	Weighted number of spanning tree for a graph \mathcal{G}
FG	A factor graph representing a belief
BA	Bundle Adjustment
BSP	Belief Space Planning
BT	Bayes Tree
DA	Data Association
EKF	Extended Kalman Filter
EM	Expectation Maximization
FIM	Fisher Information Matrix
KDE	Kernel Density Estimation
MCTS	Monte Carlo Tree Search
ML	Maximum Likelihood
MMD	Maximum Mean Discrepancy
pdf	probability density function
POMDP	Partially Observable Markov Decision Process
PRM	Probabilistic Road Map
RMSE	Root Mean Square Error
SfM	Structure from Motion
SLAM	Simultaneous Localization and Mapping
VND	Von Neumann Difference
WTCD	Weighted Tree Connectivity Difference

Chapter 1

Introduction

1.1 Motivation

As autonomous agents and robots take on ever more crucial roles, there is a growing demand to improve processes and integrate these agents further into our everyday routines. However, alongside this demand, there is a necessity for these agents to operate both reliably and efficiently. Unlike in closed, simulated environments, real-world scenarios inherently entail non-determinism, demanding the capability to navigate diverse sources of uncertainty and operate with limited knowledge about the environment. Such sources of uncertainty include dynamic environments prone to unpredictable events; limitations imposed by noisy or restricted observations due to physical constraints of various sensors; and ambiguous measurements, wherein a particular observation could have multiple plausible interpretations.

To handle these real-world scenarios, autonomous agents are required to reason over high-dimensional probabilistic states, also known as beliefs, that account for the associated uncertainty. Specifically, they should be able to perform two tasks, inference and decision making under uncertainty. In the process of inference, the agent maintains and propagates the uncertainty-aware, often high-dimensional state representation over time. In decision making under uncertainty the agent is expected to autonomously determine the future course of action based on the state estimate or uncertainty-aware representation, aiming to accomplish a specific objective. For a decision to be reliable and safely executable, the agent must anticipate the propagation of uncertainty into the future, considering various candidate actions or policies alongside possible future observations and environmental conditions. In high-dimensional state spaces, tackling these challenges poses significant computational demands.

Another layer of complexity is added when considering the challenge of ambiguous measurements. Such ambiguities arise when a particular observation has multiple possible interpretations. Some examples include the slip/grip behavior of odometry measurements; the loop closure problem in visual Simultaneous Localization and Mapping (SLAM); and unresolved Data Association (DA). The latter is defined as the process of

associating uncertain measurements to known tracks, e.g. determine if an observation corresponds to a specific landmark within a given map. Most existing inference and decision making algorithms assume DA to be given and perfect, i.e. assume a single hypothesis represented by a uni-modal state and map estimates. Yet, in perceptually aliased environments, this assumption is not reasonable and could lead to catastrophic results. Therefore, it is crucial to reason about DA, in both inference and decision making, while also considering other sources of uncertainty. However, explicitly reasoning about DA, the number of hypotheses grows exponentially with time, thereby further increasing computational complexity.

As if these computational challenges were not already hard enough, in real-world autonomous systems, the agent must also operate in real-time, often using inexpensive hardware with limited resources. In such settings, the agent is required to adhere to strict computational constraints. These may involve bounding the number of supported hypotheses and/or restricting computation time. State-of-the-art approaches therefore use different heuristics, e.g. pruning and merging, to relax the computational complexity. Yet, this loss of information incurs loss in solution quality and there are usually no performance guarantees. Moreover, inference and decision making are commonly treated separately and it is unclear how budget constraints in one process affect another.

1.2 Inference and Belief Space Planning

As previously discussed, inference and decision making are essential tasks that every autonomous agent must be capable of performing in real-world scenarios. We will briefly outline these tasks and explain their inter-dependencies.

In real-world scenarios, autonomous agents and robots gather information by acquiring measurements from various sensors. These measurements typically come with probabilistic constraints, forming a probability density function (pdf), the belief, over high-dimensional state spaces. The primary goal of inference is to find the maximum-a-posteriori (MAP) single point state estimate. For autonomous agents and robots, this state usually refers to the agent’s trajectory and/or the environment. In sequential probabilistic inference, new measurements are collected at each time step, and the belief is propagated over time. This process is also known as Bayesian inference, as it uses Bayes’ rule to update the belief. Sequential belief-based inference methods are commonly divided into filtering and smoothing. In filtering, the belief is maintained only over the present state by marginalizing (integrating out) all past states. In smoothing, the belief is maintained over the entire high-dimensional state, including past poses.

While filter-based approaches, such as the Extended Kalman Filter (EKF) (McElhoo, 1966; Smith et al., 1962), the Extended Information Filter (Thrun et al., 2004), and particle filters (Sim et al., 2005), typically have lower computational demands due to maintaining lower-dimensional beliefs, smoothing approaches (e.g., Dellaert and Kaess,

2006) are considered more accurate and robust. This is because smoothing maintains a belief over the entire trajectory, enabling loop closures and updates to past state estimates, re-linearization of past constraints with updated estimates for improved accuracy, and leveraging the topological structures of the constraint graph. However, with more constraints and possibly more variables added to the system at each time step, propagating the belief over high-dimensional state spaces becomes increasingly challenging due to the curse of dimensionality (Bellman et al., 1957). State-of-the-art smoothing approaches (e.g., Ila et al., 2017; Kaess et al., 2012; Mur-Artal and Tardós, 2017) address this issue by using efficient incremental belief updates, rather than re-computing everything from scratch.

Decision making under uncertainty is often modeled as a Partially Observable Markov Decision Process (POMDP) (Kaelbling et al., 1998) due to the consideration of stochastic actions and observations. The goal is to find the optimal action or policy from a set of candidates with respect to a specific task-related objective function. To achieve this, the belief is propagated through multiple time steps into the future according to different hypotheses for each candidate action. This process involves solving multiple state inference sessions and predicting the future belief development for each candidate, essentially performing a forward search in the belief space. This is why the problem is referred to as Belief Space Planning (BSP) (Bonet and Geffner, 2000).

In online autonomy for partially observable settings, the common approach includes a Plan-Act-Sense-Infer cycle at each time step, establishing a direct dependence between inference and BSP. Planning into the future relies on solving multiple inference problems, demonstrating BSP’s reliance on inference. Conversely, each inference problem within planning is based on a different hypothesis of the future. Additionally, each action decided by the planner leads to a different state, influencing future state inference in various ways. When an autonomous agent has limited resources and some information must be discarded, these inter-dependencies can have different effects, which we study here for the first time.

1.3 Contributions

This work primarily aims to enable computationally efficient inference and BSP for real-time operations in real-world scenarios by leveraging various structures found in the underlying representations of these problems. These structures can include graphical models that represent dependencies and relationships between different variables, and specific structures modeled in posterior probability distributions. By identifying and exploiting these structures, we can simplify complex problems and make them more tractable for real-time computation.

In BSP, we address computational budget constraints by solving simplified problems that are computationally much easier to handle. By focusing on these tractable problems, we aim to balance the need for real-time performance with the practical

limitations of computational resources. Furthermore, we mitigate the reduction in solution quality by offering robust performance guarantees. These guarantees involve either ensuring the same optimal solution as would be achieved in the original, more complex problem or providing detailed information regarding any potential loss in solution quality. To achieve this, we utilize both the topological structures associated with the underlying graph representations of the problem and the specific structures modeled in posterior probability distributions when dealing with ambiguous measurements. For the first time, we also rigorously analyze the impact of hard computational budget constraints in both planning and inference, shedding light on their implications for real-world applications.

In the context of inference, we focus on efficiently managing high-dimensional posterior probability distributions by extracting slices from these complex structures. These slices are then utilized to approximate posterior distributions of any shape accurately. This approach is particularly beneficial for facilitating real-time operation in real-world scenarios where posterior distributions are often nonparametric and can exhibit significant complexity. By approximating these distributions effectively, we ensure that the inference process remains computationally feasible even under the constraints of limited hardware resources.

Moreover, our approach emphasizes the practical applicability of these methods in real-world autonomous systems. By focusing on leveraging inherent problem structures and providing performance guarantees, we ensure that our methods are not only theoretically sound but also practically viable. This dual focus on theory and practice allows us to develop solutions that are both efficient and reliable, capable of operating within the strict constraints imposed by real-time applications and limited computational budgets. Ultimately, this work contributes to the advancement of autonomous systems by providing robust tools for efficient inference and planning, enabling these systems to perform effectively in complex, dynamic environments.

1.4 Thesis Structure

The structure of the thesis is as follows: Chapter 2 provides a thorough literature survey, setting the foundation by reviewing the state-of-the-art methods and highlighting the key challenges and gaps in current research.

In Chapter 3, we delve into the utilization of topological structures derived from factor graphs associated with posterior beliefs. We demonstrate how these topological signatures can be leveraged to approximate a *focused* objective function in BSP problems. This chapter introduces and details the concepts of our Weighted Tree Connectivity Difference (WTCD) and the Von Neumann Difference (VND) signatures, including proofs of asymptotic convergence for WTCD and computational advantages.

Chapter 4 presents our distilled DA-aware BSP approach, which capitalizes on specific structures modeled within posterior probability distributions. This chapter ad-

dresses the critical challenge of operating under hard computational budget constraints, where some information must be discarded. We showcase how our approach can maintain performance guarantees in such scenarios and provide an analysis of the impact of these constraints on both inference and planning processes.

In Chapter 5, we tackle nonparametric inference, accommodating posterior distributions of arbitrary shapes. We introduce our novel *slices* approach, which efficiently manages high-dimensional posterior probability distributions by extracting informative *slices* from these complex structures. This method offers a robust alternative to traditional iterative procedures, enhancing computational efficiency without compromising accuracy.

Finally, Chapter 6 concludes the thesis by summarizing our contributions, main findings, and the implications of our work. We further discuss interesting directions for future research, paving the way for further exploration and innovation in the domain of autonomous systems and BSP.

Chapter 2

Literature Survey

In an attempt to ensure reliable and efficient operation of autonomous systems, numerous studies have explored various structural properties for both inference and belief space planning. In this chapter, we aim to provide a comprehensive review of these works, focusing on the innovative approaches proposed to address the inherent computational complexities. By examining the evolution of these techniques, we highlight how different strategies have been employed to reduce dimensionality and offer performance guarantees, thereby facilitating more practical solutions for real-world applications.

2.1 Topological Structures and the Focused Case

Various topological aspects are utilized for inference and BSP. One aspect is the problem of topological mapping, which seeks to determine the graph structure of an environment from a sequence of measurements (Ranganathan and Dellaert, 2011). These topological maps of environments have been used for localization (e.g., in Angeli et al., 2009; Choset and Nagatani, 2001; Paul and Newman, 2010). Another aspect involves the topology of the configuration space, used to generate and classify candidate paths in robot motion planning (e.g., Bhattacharya et al., 2015; Kim et al., 2013). In our work, however, topology refers to the structure of the belief, which encodes dependency relationships among states and is derived from an associated probabilistic graphical model.

State-of-the-art approaches (e.g., Kaess et al., 2010; Kaess et al., 2012) utilize efficient probabilistic graph representations for inference. However, only recently, several works suggested utilizing topological structures associated with such graph representations to approximate the solution of the underlying estimation problems in order to reduce the computational complexity. Olson and Kaess (2009) described performance evaluation metrics for map optimization. They empirically observed that the impact of overfitting in pose graph optimization problems can be predicted using only the average node degree, a topological characteristic. Carlone (2013) calculated a conservative estimate of the basin of attraction for the Maximum Likelihood (ML) estimation in pose graph optimization via Gauss-Newton. This estimate was found to be related

to the smallest eigenvalue of the reduced Laplacian matrix, where the main diagonal elements correspond to the node degrees of the corresponding graph. Khosoussi et al. (2014) extended this line of work and empirically demonstrated that, specifically in linear-Gaussian models with isotropic noise (e.g., SLAM with known orientation), the Fisher information matrix is proportional to the reduced Laplacian matrix of the corresponding graph. This relationship enabled them to establish a connection between the determinant of the reduced Laplacian matrix, a topological signature given by the weighted number of spanning trees, and the determinant of the Fisher information matrix, which is closely related to the volume of confidence ellipsoids in the estimation problem. This connection was also extended to the case of non-isotropic noise covariance matrices (Khosoussi et al., 2015).

Several recent works have leveraged topological signatures for active pose graph SLAM and BSP. Chen et al. (2019) suggested using graph topologies, such as the weighted number of spanning trees and the weighted node tree (T-optimality metric), to reduce computational complexity in finding the best trajectory for loop-closures in active 3D pose-graph SLAM. It was later found (Chen et al., 2021) that the T-optimality metric is less computationally expensive than the D-optimality metric, though the D-optimality metric performs better in terms of uncertainty evaluation. Additionally, they proposed using a tight lower bound, representing the sum of the weighted number of spanning trees of two graphs, to replace the original objective function. Kitanov and Indelman (2018) introduced Topological Belief Space Planning (T-BSP) to address the computational complexity of BSP in high-dimensional state spaces, such as multi-robot active 2D SLAM. They demonstrated that topological properties, such as the normalized number of spanning trees and the Von Neumann graph entropy signatures (Mowshowitz and Dehmer, 2012; Passerini and Severini, 2009), of the underlying factor graph (Kschischang et al., 2001) representations of future posterior beliefs, are highly correlated with the information-theoretic cost. Additionally, they proposed using an approximation of the Von Neumann entropy (Han et al., 2012) to guide the search for an optimal BSP solution by sub-sampling the topological space as an anytime algorithm that eventually converges to the optimal solution. Their empirical results show that when a strong correlation exists between the objective and the topological metric, this convergence is significantly faster than existing state-of-the-art approaches. Kitanov and Indelman (2024) also derived error bounds for T-BSP that depend on topological parameters, measurement noise, and prior state estimates, which can be calculated with minimal additional computational cost. However, while all of these works leveraged topological aspects, none addressed the specific case where the objective is a function of only a predefined subset of variables.

Unlike the unfocused case, where the objective is a function of all variables, the focused case for both inference and BSP has been much less studied. Levine and How (2013) addressed a focused active inference problem of selecting a subset of observable random variables that are maximally informative with respect to a specified subset of

latent random variables. Specifically, they demonstrated that non-local Mutual Information (MI) can be decomposed into local information measures, i.e., between unfocused and focused variables, which can be computed efficiently using message-passing algorithms. Mu et al. (2017) proposed a two-stage approach to a focused inference problem. They calculated the posterior covariance matrix for each measurement to obtain the marginal over a focused set of variables in resource-constrained systems, where some data will eventually need to be discarded. The rAMD L approach (Kopitkov and Indelman, 2017; Kopitkov and Indelman, 2019) efficiently evaluated the information-theoretic cost in BSP for each candidate action in both focused and unfocused cases. This approach avoided the computationally expensive Schur complement operation for each candidate action by performing a one-time calculation of the marginal covariance associated with the variables involved in the candidate actions. While these works leverage various graphical model representations to reduce computational complexity in the focused case, none of them integrate topological aspects into inference nor planning. Moreover, a key advantage of a topological BSP approach is its ability to avoid belief propagation and marginal recovery during planning, potentially finding the optimal solution under certain conditions and tight bounds.

To the best of our knowledge, our work is the first to incorporate topological aspects specifically for the active focused case.

2.2 Robust Perception and Planning

Recent proposals aim to ensure reliable and efficient operation in ambiguous environments through various approaches. These strategies, commonly known as robust perception, typically involve maintaining probabilistic DA and tracking hypotheses based on available data.

An effective inference mechanism should be able to withstand false DA overlooked by front-end algorithms such as laser scans and image matching. However, explicitly considering DA leads to a growing number of hypotheses over time, making it challenging to maintain computational efficiency. To mitigate these issues, leveraging hypotheses’ weights from previous steps to enhance current-time hypotheses pruning was proposed (Shelly and Indelman, 2022) to reduce computational complexity. Convex relaxation approaches over graphs have been proposed (e.g, Carlone et al., 2014a; Lajoie et al., 2019) to handle perceptual aliasing, identifying the maximal subset of internally coherent measurements (i.e., correct DA) and discarding false ones. The max-mixture model (Olson and Agarwal, 2013 presented) was developed to allow fast ML inference on factor graphs containing complex probability distributions, such as the slip/grip multi-modal problem. An Expectation Maximization (EM) approach (Indelman et al., 2014; Indelman et al., 2016) with factor graphs was used to efficiently infer initial relative poses and solve multi-robot DA problems. Sunderhauf and Protzel (2012) modified parts of the graph’s topological structure during optimization to dis-

card false positive loop closures. The Bayes Tree (BT) algorithm (Kaess et al., 2012) was extended (e.g., Hsiao and Kaess, 2019; Jiang et al., 2021) to explicitly incorporate multi-modal measurements within the graph and generate multi-hypothesis outputs.

Considering ambiguous DA in planning for active disambiguation is even more challenging as the number of hypotheses can grow exponentially with the planning horizon. Consequently, finding the optimal action quickly becomes intractable. Some notable related works with reduced complexity (Atanasov et al., 2014; Lauri et al., 2015; Patten et al., 2016; Wong et al., 2015) are in the context of object detection and classification. Given hypotheses regarding object class and pose, these approaches aim to find a sequence of future viewpoints that will identify the correct hypothesis. However, these approaches assume perfect localization of the agent, neglecting the agent’s pose in the belief representation.

Considering the agent’s pose within the belief (Jensfelt and Kristensen, 2001) a greedy heuristic-based planning strategy that utilized multi-hypothesis Kalman filter-based pose tracking combined with a probabilistic formulation of hypothesis correctness to resolve multi-modal hypotheses in a kidnapped robot scenario. Gasparri et al. (2007) used a set of parallel EKFs for planning safe trajectories, selecting points near obstacles to disambiguate the hypotheses. Yet, these approaches did not explicitly consider the evolution of beliefs resulting from actions during the planning stage. In contrast, (Agarwal et al., 2016) addressed belief evolution resulting from actions. They modeled DA hypotheses within the prior belief as a mixture of Gaussians, assuming that certain actions could lead to complete disambiguation. However, their approach does not account for ambiguous DA within future beliefs resulting from future observations.

The DA-BSP framework (Pathak et al., 2018) was the first to incorporate reasoning about future data association hypotheses within BSP. The authors demonstrated that a posterior belief could become a mixture of pdfs and designed cost functions to measure the expected level of ambiguity and posterior uncertainty given candidate actions. Additionally, a unique aspect of this approach is that the number of mixture components can both increase and decrease, better reflecting reality and providing a general framework capable of active disambiguation. Subsequently, the ARAS framework (Hsiao et al., 2020) also addressed ambiguous DA within future beliefs, leveraging the MH-iSAM2 graphical model (Hsiao and Kaess, 2019), which is an extension of BT (Kaess et al., 2010), for multi-hypothesis inference. Yet, in this framework, the planning algorithm operated on each ambiguity-free hypothesis individually.

Perhaps the most significant distinction between all of these approaches and our work lies in how they manage the exponential growth in the number of hypotheses. They often resort to heuristics like pruning and merging to cope with this growth, which can degrade solution quality over time and consequently cannot guarantee performance. In contrast, our approach develops analytical bounds on the degradation in solution quality relative to the original problem, allowing us to provide performance guarantees.

2.3 Simplification and Performance Guarantees in BSP

Finding the optimal solution to a POMDP problem is computationally intractable in most real-world scenarios due to the curse of dimensionality and the curse of history (see, e.g., Pineau et al., 2006). Even with simplifying assumptions such as discrete states, actions, or observations, and finite planning horizons (Papadimitriou and Tsitsiklis, 1987), computational complexity remains a significant challenge. Consequently, the research community has extensively investigated solutions to enhance scalability for real-world problems, which are generally categorized into offline and online approaches (Ross et al., 2008).

Offline approaches involve solving the POMDP problem in its entirety before execution, usually without time constraints. This means the solution is computed in advance and used for decision-making throughout the execution phase. These methods typically perform global policy optimization using dynamic programming algorithms, such as value and policy iteration (e.g., Pineau et al., 2006; Porta et al., 2006). However, they are extremely computationally demanding, especially when considering high-dimensional state spaces. In contrast, online approaches solve the POMDP problem in real-time as new information becomes available. The solution is continuously updated based on current observations and decisions, aiming to infer a specific sequence of actions to be executed immediately.

Achieving online performance is challenging, especially when autonomous systems use inexpensive hardware with limited resources and must adhere to strict computational constraints. Moreover, in continuous observation and state spaces often found in real-world scenarios, integration is impractical, necessitating approximate analytical solutions. Many prior works have focused on approximating solutions to the POMDP problem. Some methods use direct trajectory optimization (e.g., Indelman et al., 2015a; Van Den Berg et al., 2012). Other methods build on established motion planners (e.g., Karaman et al., 2011; Kavraki et al., 1996) by sub-sampling a finite graph in the belief space to search for a solution (e.g., Agha-Mohammadi et al., 2014; Prentice and Roy, 2009). While these methods focus on generating or sampling candidate trajectories, a complementary approach, and the focus of our work, is to efficiently compare objective function values by approximating the objective function or reducing the cost of belief propagation.

Belief sparsification is one approach to reducing the cost of belief propagation. Various sparsification methods have been used in inference (e.g., Dellaert and Kaess, 2006; Huang et al., 2012; Hsiung et al., 2018; Thrun et al., 2004) to lower this cost. However, Indelman (2016) was the first to introduce belief sparsification in planning to enable long-term operation. The author used a diagonal covariance approximation in a myopic setting with one-row unary Jacobians to significantly reduce the complexity of objective calculation. Inspired by this work, Elimelech and Indelman (2022) proposed identifying uninvolved variables and sparsifying the posterior information matrix for

each candidate action to reduce computation time in BSP problems. Alternatively, Davidson and Hutchinson (2009) suggested approximating the objective function using a bound on the maximum eigenvalue of the estimation error covariance matrix as the cost function for BSP. Nevertheless, whenever an approximation is presented, the question of optimality guarantees inevitably arises.

To provide performance guarantees when using any approximation, it is crucial first to analytically bound the loss in solution quality for each candidate action between the original problem and the simplified problem. When such bounds exist, the optimal action in the original problem can be recovered if a set of bounds, representing the error intervals in the simplified problem, does not overlap with any other interval associated with other actions. This concept, termed action consistency (see Elimelech and Indelman, 2017a; Elimelech and Indelman, 2017b; Elimelech and Indelman, 2017c), builds on observations from Indelman (2015) and Indelman (2016). Several works have applied this concept to maintain action consistency or bound the error loss. Examples include Elimelech and Indelman (2022), who approximated belief representations through sparsification; Szyglic and Indelman (2022), who used bounds as a function of simplified beliefs to reduce computational complexity in non-myopic BSP problems with general belief distributions; and Zhitnikov et al. (2024), who incorporated this concept within a Monte Carlo Tree Search (MCTS) planning framework. In our work, we also leverage similar concepts to provide performance guarantees and incorporate this concept for the first time in BSP problems while reasoning about DA.

2.4 Non-Parametric Inference

Research on inference algorithms, particularly for SLAM, has largely concentrated on parametric solutions, typically involving the approximation of posterior distributions with parametric Gaussian models. A prominent example is iSAM2 (Kaess et al., 2012), which introduced the BT graphical model. iSAM2 retrieves posterior distributions through incremental upward and downward passes over the BT, akin to the *forward-backward* algorithm. More recent parametric approaches aim to enhance robustness by addressing challenges such as unresolved DAs, multi-modal factors, and outliers (e.g., Hsiao and Kaess, 2019; Huang et al., 2013; Indelman et al., 2014; Indelman et al., 2016; Olson and Agarwal, 2013; Sünderhauf and Protzel, 2012). However, these methods still struggle with non-Gaussian posterior distributions.

Nonparametric inference methods, which use sampling techniques, are theoretically capable of approximating any posterior distribution. Early methods in this area include Markov Chain Monte Carlo (MCMC) algorithms (Hastings, 1970) and particle filters (Gordon et al., 1993). Briers et al. (2010) proposed an extension of the two-filter smoothing formula (Bresler, 1986; Kitagawa, 1994), which traditionally combines the outputs of two independent filters: the standard forward filter and the backward information filter. They introduced a generalized version of this formula that incorporates

Sequential Monte Carlo (SMC) (Doucet et al., 2001) techniques, enabling it to handle state-space models where a closed-form solution is not available. NBP (Sudderth et al., 2002) used concepts from regularized particle filters and belief propagation. This work introduced an algorithm designed for general graphs. In each iteration, NBP employs an efficient sampling procedure to update kernel-based approximations of the true continuous likelihoods. As a result, NBP extends the capabilities of particle filtering methods to handle the broader range of vision problems described by graphical models. Despite their theoretical flexibility, these approaches often become computationally impractical in high-dimensional scenarios like those found in SLAM.

FastSLAM (Montemerlo et al., 2002) was introduced to combine the advantages of particle filters and parametric techniques for effectively handling high-dimensional settings. However, it has its limitations. FastSLAM relies on EKF for landmark tracking and cannot handle general non-Gaussian distributions. Additionally, like many particle filters, FastSLAM can suffer from particle depletion, which may lead to degenerate estimates. GAPSLAM (Huang and Leonard, 2023) also integrates particle filters with parametric techniques to balance accuracy and efficiency. It uses particle filters to handle highly uncertain landmarks with sample-based density representations while maintaining a Gaussian approximation for computational efficiency. The approach computes an empirical covariance matrix and its largest eigenvalue for each landmark. When this eigenvalue drops below a certain threshold, indicating sufficient certainty, the landmark’s representation is switched to a Gaussian model. This combination allows the Gaussian model to streamline the particle filter’s computation and aids in optimizing the robot’s path by providing accurate linearization points for nonlinear solvers. However, GAPSLAM struggles when future observations cause previously Gaussian-represented landmarks to revert to non-parametric distributions.

Several recent state-of-the-art methods leverage the conditional independence structure within factor graphs (Kschischang et al., 2001) to handle nonparametric inference in SLAM. For instance, mm-iSAM (Fourie et al., 2016) approximates non-Gaussian posterior distributions using samples and Kernel Density Estimation (KDE). Specifically, mm-iSAM employs an iterative nested Gibbs sampling approach (Ihler et al., 2003) to generate samples from partial posteriors associated with each clique in the BT. These samples are then used to create new approximations of posterior distributions via KDEs. On the other hand, NF-iSAM (Huang et al., 2023) leverages the capabilities of neural networks by training normalizing flows (Rezende and Mohamed, 2015) to model conditionals that decompose non-Gaussian posterior distributions during the upward pass. After all cliques have trained their conditional samplers, NF-iSAM performs a downward pass to draw samples from posterior distributions in a root-to-leaf manner. However, this learning approach is computationally demanding and requires a large number of samples to achieve convergence.

NSFG (Huang et al., 2022), which uses nested sampling (Skilling, 2006), prioritizes higher accuracy over computational efficiency. It employs nested sampling techniques

to sample the posterior distribution directly across iterations until convergence. In this iterative process, the factor graph is initially divided into two components: an acyclic graph, referred to as the prior factor set, and a likelihood factor set. NSFG subsequently employs ancestral sampling to produce samples from the prior factor set, while likelihood evaluations are provided based on the likelihood factor set. Although NSFG delivers more precise estimates of non-Gaussian posterior distributions when compared to mm-iSAM and NF-iSAM, its iterative nature imposes significant computational demands. Consequently, it struggles to address high-dimensional, large-scale problems and cannot be employed in real-time applications.

Our work adopts a slices perspective for non-parametric inference. We use slices, derived from structures representing high-dimensional posterior probability distributions, to efficiently approximate posterior distributions of any shape. This method avoids the need for iterative processes like Gibbs sampling (Fourie et al., 2016), training neural networks (Huang et al., 2023), or NSFG (Huang et al., 2022), which inherently involve multiple iterations.

Chapter 3

Focused Topological BSP

BSP problems can be categorized based on the prioritization of variables in the planning process. In the *unfocused* case, the objective function includes a belief that represents all variables within the estimation problem. Conversely, in the *focused* case, the objective function specifically prioritizes or exclusively considers a subset of variables, typically those most relevant to the task at hand. Examples of *focused* BSP problems include: a navigation task in an uncertain environment where a robot must reach a goal position with maximum accuracy regardless of the chosen path; a reconstruction task where mapping a specific scene in the environment with high accuracy is prioritized; and collision avoidance, where localizing obstacles is of the greatest importance.




Figure 3.1: Chosen candidate paths in an active 2D pose SLAM scenario are shown with respect to different objectives. The red path aims to reduce uncertainty over the entire trajectory, representing the *unfocused* case, while the green path focuses on reducing uncertainty only over the final pose, representing the *focused* case. The yellow path indicates the locations previously visited by the robot, serving as the initial belief. The starting position is marked with a red star, and the goal is denoted by a circle. Black ellipses illustrate the marginal covariance at each step, and black dashed lines indicate loop closures. Notice how the marginal covariance decreases significantly in the green path right before reaching the goal due to a large loop closure.

In general, the optimal solution in the *focused* case can be significantly different from the one in the *unfocused* case (e.g., Levine and How, 2013). To solve these *focused* problems for Gaussian distributions, state-of-the-art algorithms calculate the marginal posterior covariance (or information) matrix for each candidate action. Despite the potentially small set of *focused* variables, these calculations involve a computationally expensive Schur complement operation. Moreover, many problems require a high n -dimensional state space model to represent poses, obstacles, and landmarks. In such cases, calculating the determinant of the marginal information matrix (the D-criterion), which is common in information-theoretic problems, has a general complexity of $O(n^3)$.

In this chapter, we examine the relationship between the graphical structure of SLAM and an information-theoretic objective function in the active *focused* case. We introduce a novel concept, Focused Topological Belief Space Planning (FT-BSP), which utilizes topological signatures associated with posterior probabilities to guide the search for an optimal action. Our contributions are as follows:

1. We introduce a novel approach, FT-BSP, that addresses the *focused* information-theoretic BSP problem through topological aspects.
2. Within FT-BSP, we derive two topological signatures to approximate the *focused* cost function: the WTCD signature and the VND signature.
3. We prove asymptotic convergence and develop bounds for the WTCD signature.
4. We demonstrate how the VND signature can be calculated incrementally to support sequential BSP problems through multiple time steps.
5. We provide empirical results showing that both topological signatures are highly correlated with the *focused* information-theoretic objective function and are significantly faster to calculate.

3.1 Problem Formulation and Notations

In this section we provide the theoretical background for *focused* BSP and briefly review the factor graph model (Kschischang et al., 2001) from which we induce the topological signatures associated with a posterior belief.

3.1.1 Focused Belief Space Planning

Consider a robot operating in a partially known environment, aiming to autonomously decide its future actions based on information accumulated thus far and a user defined objective function J . In this chapter, for simplicity, we assume a 2D pose SLAM framework with relative pose measurements.

Let $p \in \mathbb{R}^2$ and $\theta \in [-\pi, \pi)$ denote the robot position and orientation, respectively. Let x_k denote the robot's state at time instant k where the state vector is defined as $x_k = [p_k^T \ \theta_k]^T$. The joint state, up to and including time k , is defined as $X_k = \{x_0, x_1, \dots, x_k\}$. A *focused* subset of states is denoted by $X_k^F \subseteq X_k$ while the remaining *unfocused* states are denoted by $X_k^U = X_k / X_k^F$.

Let $z_{0:k}$ and $u_{0:k-1}$ denote, respectively, all observations and controls up to time k . The motion and observation models are given by

$$x_{k+1} = f(x_k, u_k) + w_k \quad , \quad z_{jk} = h(x_j, x_k) + v_{jk}, \quad (3.1)$$

where the terms $w_k \sim \mathcal{N}(0, \Sigma_w)$ and $v_{jk} \sim \mathcal{N}(0, \Sigma_{v_{jk}})$ represent the process and measurement noise, respectively, and are sampled from known Gaussian distributions, with zero means and specified covariance matrices $\Sigma_w, \Sigma_{v_{jk}}$.

The posterior pdf over the joint state, denoted as the belief, is given by

$$b[X_k] \doteq \mathbb{P}(X_k | Z_{1:k}, u_{0:k-1}) = \mathbb{P}(X_k | H_k), \quad (3.2)$$

where $H_k \doteq \{Z_{1:k}, u_{0:k-1}\}$ represent history at time k . Given a candidate action sequence $u_{k:k+L-1}$ and future observations $Z_{k+1:k+L}$, the future joint belief is given by

$$b[X_{k+L}] \doteq \mathbb{P}(X_{k+L} | Z_{1:k+L}, u_{0:k+L-1}). \quad (3.3)$$

The future joint belief can be expressed in terms of $b[X_k]$ and the corresponding motion and observation models

$$b[X_{k+L}] \propto \mathbb{P}(X_k | H_k) \prod_{l=k+1}^{k+L} \mathbb{P}(x_l | x_{l-1}, u_{l-1}) \mathbb{P}(Z_l | X_l), \quad (3.4)$$

where the measurement likelihood term $\mathbb{P}(Z_l | X_l) = \prod_{m=1}^{n_l} \mathbb{P}(z_{l,i_m} | x_l, x_{i_m})$ represents all n_l observations acquired at time l between involved variables x_l and x_{i_m} , $\{i_m\} \subseteq \{0, 1, \dots, l-1\}$.

Given a user defined objective function J , a belief $b[X_k]$ and a set of candidate actions \mathcal{U}_k , the goal of BSP is to find the optimal action given by

$$\mathcal{U}^* = \underset{\mathcal{U}}{\operatorname{argmin}} J(\mathcal{U}). \quad (3.5)$$

For the *unfocused* case, a general objective function in BSP can be written as

$$J(\mathcal{U}) = \mathbb{E}_{Z_{k+1:k+L}} \left[\sum_{l=0}^{L-1} c_l(b[X_{k+l}], u_{k+l}) + c_L(b[X_{k+L}]) \right], \quad (3.6)$$

where c_l represents a cost function for each look-ahead step, c_L represents the cost function of the terminal belief (at the end of the planning horizon), and the expectation

is taken with respect to future observations. While each cost function can include a number of different terms such as distance to goal, energy spent and information measures of future beliefs, in this work, we only consider the information theoretic term of a terminal belief (to be defined) at time step $k + L$. We note that different cost terms c_l can be treated in a similar manner. Specifically, given an appropriate posterior belief, we aim to minimize the differential entropy \mathcal{H} .

Evaluating the objective function (4.4) involves inference over the appropriate posterior beliefs. Within a pose SLAM framework with relative pose measurements, the Maximum Likelihood (ML) estimation is obtained by the optimal state X_k^* that maximizes the belief (4.2). By fixing an arbitrary pose as an anchor, e.g. x_0 , and considering the rest as unknown, X_k^* is obtained by minimizing the sum of weighted squared errors between the predicted and measured relative poses

$$X_k^* = \arg \min_{X_k} \|\Delta_k - h(X_k)\|_{\Sigma^{-1}}^2, \quad (3.7)$$

where the measurement model

$$\Delta_k = h(X_k) + \epsilon, \quad \epsilon \sim \mathcal{N}(0, \Sigma), \quad (3.8)$$

represents a vector of m stacked relative pose measurements $z_{ij}^r \in SE(2), r = 1, 2, \dots, m$ generated according to the motion and observation models (4.1). In this work, we assume the noise covariance matrices, for both the rotational and translational measurements, have a block-isotropic structure, i.e. $\Sigma = \text{diag}(\Sigma_p, \Sigma_\theta)$ where $\Sigma_p = \text{diag}(\sigma_{p_1}^2 \mathbf{I}_2, \dots, \sigma_{p_m}^2 \mathbf{I}_2)$ and $\Sigma_\theta = \text{diag}(\sigma_{\theta_1}^2, \dots, \sigma_{\theta_m}^2)$.

The information matrix $\Lambda(X_k)$, referred from here on simply as Λ_k , of the ML estimation is given by $\Lambda_k = M^T \Sigma^{-1} M$, where $M = \partial h / \partial X_k$ is a measurement Jacobian (Sorenson, 1980). Evaluating Λ_k at the true value of X_k is known as the Fisher Information Matrix (FIM) while evaluating Λ_k at X_k^* is used to approximate the FIM.

Given a multivariate Gaussian posterior belief and taking the ML observations assumption (Platt et al., 2010), the differential entropy, considering only the terminal joint belief, for the *unfocused* case is given by (log denotes the natural logarithm in this work)

$$J_{\mathcal{H}}(\mathcal{U}) = \mathcal{H}(b[X_{k+L}]) = \frac{n}{2} \log(2\pi e) - \frac{1}{2} \log |\Lambda_{k+L}|, \quad (3.9)$$

where n is the dimension of the joint state X_{k+L} and Λ_{k+L} is the estimated FIM associated with the joint posterior belief $b[X_{k+L}]$. As we are only interested in a *focused* set of variables X_{k+L}^F , the differential entropy, considering only the terminal marginal belief, over the *focused* set, is given by

$$J_{\mathcal{H}}^F(\mathcal{U}) = \mathcal{H}(b[X_{k+L}^F]) = \frac{n^F}{2} \log(2\pi e) - \frac{1}{2} \log |\Lambda_{k+L}^{M,F}|, \quad (3.10)$$

where n^F is the dimension of X_{k+L}^F and $\Lambda_{k+L}^{M,F}$ is the marginal FIM of the marginal posterior belief $b[X_{k+L}^F]$.




Figure 3.2: Graph representations of a posterior belief. (a) Factor graph with *unfocused* variables in blue, *focused* variables in purple and anchor variable in yellow; (b) The corresponding topological graph \mathcal{G} ; (c) The *unfocused* topological augmented graph $\mathcal{G}^{U,A}$. The blob encapsulating the *unfocused* nodes represents the reduced Laplacians in (3.26) and the edges in orange are new edges weighted according to (3.27). Specifically $w_{1a} = w_{10} + w_{15}$.

3.1.2 Belief Topology

The joint belief at time k can be represented by a factor graph (Kschischang et al., 2001). A factor graph (FG) is a probabilistic graphical model that represents a factorization of a pdf in terms of process and measurement models. It is a bipartite graph whose nodes consist of factors \mathcal{F} and variables \mathcal{X} . The variables \mathcal{X} represent the random variables in the estimation problem while the factors represent probabilistic information on those variables. The FG edges encode connectivity according to the variables involved in each factor. In this chapter we only consider pairwise factors without self-loops.

A topological representation of a joint belief is defined as a topological graph $\mathcal{G} \doteq (\mathcal{V}, \mathcal{E}, w)$ associated with a posterior FG where each node $v_i \in \mathcal{V}$ corresponds to a variable node $x_i \in \mathcal{X}$ (see Fig. 3.2). Specifically, if x_0 is fixed as an *anchor* pose then v_0 is the corresponding anchor node. We also note that \mathcal{V}^F and \mathcal{V}^U correspond to X^F and X^U , respectively. The edge set $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ is defined as

$$e_{ij} \doteq (v_i, v_j) \in \mathcal{E} \iff f_{ij} \doteq (x_i, x_j) \in \mathcal{F}, \quad (3.11)$$

and the weight function $w : \mathcal{E} \rightarrow \mathbb{R}_{>0}$ assigns each edge with a weight w_{ij} derived from the appropriate term in the associated noise model of the corresponding factor.

Let W be a diagonal matrix of size $|\mathcal{V}| \times |\mathcal{V}|$ defined as

$$W_{ii} = \sum_j w_{ij}, \quad (3.12)$$

i.e. each element on the main diagonal W_{ii} is equal to the sum of weights of all edges connected to node v_i . Chung (1997) defined the weighted Laplacian matrix \mathcal{L}_w

associated with a topological graph \mathcal{G} as

$$\mathcal{L}_w(i, j) = \begin{cases} W_{ii} & \text{if } i = j, \\ -w_{ij} & \text{if } v_i \text{ and } v_j \text{ are adjacent,} \\ 0 & \text{otherwise} \end{cases} \quad (3.13)$$

Chung and Richardson, 2006 defined the weighted normalized Laplacian matrix $\hat{\mathcal{L}}_w$ associated with a topological graph \mathcal{G} as

$$\hat{\mathcal{L}}_w(i, j) = \begin{cases} 1 & \text{if } i = j, \\ \frac{-w_{ij}}{\sqrt{W_{ii} \cdot W_{jj}}} & \text{if } v_i \text{ and } v_j \text{ are adjacent,} \\ 0 & \text{otherwise} \end{cases} \quad (3.14)$$

The reduced weighted Laplacian matrix L_w and the reduced weighted normalized Laplacian matrix \hat{L}_w are retrieved by removing the row and column that correspond to the anchor node from \mathcal{L}_w and $\hat{\mathcal{L}}_w$, respectively.

We also denote the reduced incidence matrix of \mathcal{G} by A . It is obtained by removing the row that corresponds to the anchor node from the incidence matrix of \mathcal{G} .

3.1.3 Weighted Tree Connectivity

Khosoussi et al. (2019) introduce a weighted version to Kirchoff's matrix tree theorem (Biggs, 1993). Given a weighted topological graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, w)$, the value of each spanning tree $T_{\mathcal{G}}$ is given by

$$\mathbb{V}_w(T_{\mathcal{G}}) = \prod_{e \in \mathcal{E}(T_{\mathcal{G}})} w(e), \quad (3.15)$$

and the weighted number of spanning trees is defined as

$$t_w(\mathcal{G}) = \sum_{t \in T_{\mathcal{G}}} \mathbb{V}_w(t). \quad (3.16)$$

The Weighted Tree Connectivity (WTC) of \mathcal{G} is defined as

$$\tau_w(\mathcal{G}) = \log t_w(\mathcal{G}). \quad (3.17)$$

Considering the 2D pose SLAM framework addressed in this chapter, the FIM of a posterior belief is given by (see eq. (18) in Khosoussi et al., 2019)

$$\Lambda = \begin{bmatrix} L_{w_p} \otimes I_2 & (A_{w_p} \otimes I_2) \Gamma \Delta_{w_p} \\ *^T & L_{w_\theta} + \Delta_{w_p}^T \Delta_{w_p} \end{bmatrix}, \quad (3.18)$$

where $*$ denotes the top-right block; I_2 is the identity matrix of size 2×2 and \otimes

denotes the Kronecker product; L_{w_p}, L_{w_θ} are the reduced weighted Laplacian matrices of \mathcal{G} when edges are weighted according to $w_p : e_i \rightarrow \sigma_{p_i}^{-2}, w_\theta : e_j \rightarrow \sigma_{\theta_j}^{-2}$ respectively; $A_{w_p} = A \cdot \text{diag}(\sigma_{p_1}, \dots, \sigma_{p_m})$ is the reduced weighted incidence matrix when edges are weighted by w_p ; Γ is defined as

$$\Gamma = I_{|\mathcal{E}|} \otimes \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}, \quad (3.19)$$

and $\Delta_{w_p}^T \Delta_{w_p} = D_{w_p}$ is a diagonal matrix where

$$D_{w_p}(i, i) = \sum_{j \in \mathcal{N}_{\text{out}}(i)} w_p(i, j) \|p_i - p_j\|^2, \quad (3.20)$$

i.e. $D_{w_p}(i, i)$ is equal to the weighted sum of squared distances between the i 'th robot pose and every node observed by it.

Lower and upper bounds on the actual D-criterion (see Theorem 3 in Khosoussi et al., 2019) are given by

$$\tau_w(\mathcal{G}) \leq \log |\Lambda| \leq \tau_w(\mathcal{G}) + n \cdot \log(1 + \delta/\lambda_1), \quad (3.21)$$

where $\tau_w(\mathcal{G}) = 2\tau_{w_p}(\mathcal{G}) + \tau_{w_\theta}(\mathcal{G})$; $\delta = \|\Delta_{w_p}^T \Delta_{w_p}\|_\infty$ and $\lambda_1 = \lambda_{\min}(L_{w_\theta})$. Moreover, they show that under some conditions, these bounds become asymptotically tight and the D-criterion is characterized solely by the weighted tree connectivity of \mathcal{G}

$$\lim_{\delta/\lambda_1 \rightarrow 0^+} \log |\Lambda| = 2\tau_{w_p}(\mathcal{G}) + \tau_{w_\theta}(\mathcal{G}). \quad (3.22)$$

From hereon, we denote $\tau_w(\mathcal{G}) \doteq \tau_w$.

3.2 Planning using Topological Signatures in the Focused Case

Our goal is to rank candidate actions by redefining the original *focused* problem (3.10) in a topological space, where it can be solved more efficiently. Given a belief $b[X_k]$ and a set of candidate actions \mathcal{U}_k , we look for a topological representation $\mathcal{G}(u_k)$ for each $u_k \in \mathcal{U}_k$ and a signature $S : \mathcal{G}(u_k) \rightarrow \mathbb{R}$, that would ideally yield a solution which is consistent with the optimal solution of the original *focused* problem (3.10) such that $\hat{\mathcal{U}} = \mathcal{U}^*$ where $\hat{\mathcal{U}} = \min_{\mathcal{U}} S(\mathcal{G}(\mathcal{U}))$.

Our derived topological signatures were inspired by the works of Khosoussi et al., 2019 and Kitanov and Indelman, 2018 on measurement selection and active SLAM in the *unfocused* case, respectively.

3.2.1 Manipulating the Focused Objective Function $J_{\mathcal{H}}^F$

In practice, at each time step, we maintain a factor graph that represents the joint posterior belief from which we can induce the topological graph. Specifically, at time step $k+L$ we have access to FG_{k+L} . As we are interested in a topological representation which corresponds to the marginal posterior belief over the *focused* variables X_{k+L}^F , we would first need to marginalize out the *unfocused* variables X_{k+L}^U . Although such algorithms exist, e.g. the sum-product algorithm (Kschischang et al., 2001), as the set of *focused* variables is often small with respect to the entire estimation problem, e.g. when we are interested in reducing the entropy only over the robot's last pose, calculating the marginal posterior information matrix involves an expensive Schur complement operation.

As the joint posterior information matrix Λ_{k+L} is positive-definite and symmetric, we can partition it such that

$$\Lambda_{k+L} = \begin{bmatrix} \Lambda_{k+L}^F & \Lambda_{k+L}^{F,U} \\ (\Lambda_{k+L}^{F,U})^T & \Lambda_{k+L}^U \end{bmatrix}, \quad (3.23)$$

where $\Lambda_{k+L}^F \in \mathbb{R}^{n^F \times n^F}$ and $\Lambda_{k+L}^U \in \mathbb{R}^{n^U \times n^U}$ are constructed from Λ_{k+L} by retrieving only the rows and columns related to X_{k+L}^F and X_{k+L}^U respectively. Notice that Λ_{k+L}^U is the conditional posterior information matrix of X_{k+L}^U , conditioned on the rest of the variables X_{k+L}^F . The remaining blocks, $\Lambda_{k+L}^{F,U}$, contain the mixed information between *focused* and *unfocused* variables.

While the marginal posterior information matrix $\Lambda_{k+L}^{M,F}$ can be calculated using the Schur complement, to calculate the *focused* objective function (3.10) we only need to evaluate $|\Lambda_{k+L}^{M,F}|$. Using theorem 2.1 from Ouellette, 1981 where

$$|\Lambda_{k+L}| = |\Lambda_{k+L}^{M,F}| \cdot |\Lambda_{k+L}^U| \Rightarrow |\Lambda_{k+L}^{M,F}| = |\Lambda_{k+L}| / |\Lambda_{k+L}^U|, \quad (3.24)$$

and substituting (3.24) into (3.10) we rewrite $J_{\mathcal{H}}^F$ as

$$J_{\mathcal{H}}^F(\mathcal{U}) = \frac{n^F}{2} \log(2\pi e) - \frac{1}{2} \log |\Lambda_{k+L}| + \frac{1}{2} \log |\Lambda_{k+L}^U|. \quad (3.25)$$

Using (3.21), we can approximate $\log |\Lambda_{k+L}|$. To derive an approximation for (3.25) that is characterized solely by structural properties, we now aim to approximate $\log |\Lambda_{k+L}^U|$ using topological aspects.

From here on, we drop the time indices and refer to a general posterior belief $b[X]$ with the associated FIM Λ .

3.2.2 The Unfocused Augmented Graph $\mathcal{G}^{A,U}$

Recall that our goal is to find a topological graph, and a corresponding signature that would approximate $\log |\Lambda^U|$.

According to (3.23), as Λ^U is a principal submatrix of Λ , it is not hard to see that by taking only the rows and columns that correspond to X^U , the *unfocused* block Λ^U is given by

$$\Lambda^U = \begin{bmatrix} L_{w_p}^U \otimes I_2 & (A_{w_p}^U \otimes I_2) \Gamma \Delta_{w_p}^U \\ *^T & L_{w_\theta}^U + (\Delta_{w_p}^U)^T \Delta_{w_p}^U \end{bmatrix}. \quad (3.26)$$

We observe that both $L_{w_p}^U$ and $L_{w_\theta}^U$ represent reduced weighted Laplacian matrices, that contain all *unfocused* variables and weighted edges connected to those variables (see e.g. Fig 3.2c). As such, to construct a topological graph, we can add a *virtual* node v_a , to create the augmented weighted topological graph $\mathcal{G}^{U,A}$. Note that the anchor node and all nodes that represent *focused* variables do not belong to $\mathcal{G}^{U,A}$.

This augmented weighted topological graph is formally defined as $\mathcal{G}^{U,A}(\mathcal{V}^U \cup v_a, \mathcal{E}^{U,A}, w^{U,A})$ where each node $v_i \in \mathcal{V}^U$ corresponds to a variable $x_i \in X^U$; the edge set $\mathcal{E}^{U,A}$ is defined as a union of three groups $\{e_{ij}^{U,A} | v_i, v_j \in \mathcal{V}^U \wedge e_{ij} \in \mathcal{E}\} \cup \{e_{ia}^{U,A} | v_i \in \mathcal{V}^U, v_j \in \mathcal{V}^F \wedge e_{ij} \in \mathcal{E}\} \cup \{e_{ia}^{U,A} | v_i \in \mathcal{V}^U \wedge e_{i0} \in \mathcal{E}\}$ where e_{i0} denotes an edge connected to the *anchor* node; each edge $e_{ia}^{U,A} \in \mathcal{E}^{U,A}$ connected to the *virtual* node v_a is weighted by

$$w_{ia}^{U,A} = \sum_j w_{ij} \text{ such that } v_i \in \mathcal{V}^U \text{ and } v_j \in \mathcal{V}^F, \quad (3.27)$$

while the remaining edges are weighted according to $w_{ij}^{U,A} = w_{ij}$. Note that in (3.27), edges to the *virtual* node are weighted such that we construct a Laplacian matrix, i.e. where the sum of each row and each column is exactly zero. We can now use the weighted version of Kirchhoff's theorem (3.17) as a topological signature for this graph. Specifically, we define

$$\tau_w^{U,A} = 2\tau_{w_p}(\mathcal{G}^{U,A}) + \tau_{w_\theta}(\mathcal{G}^{U,A}). \quad (3.28)$$

3.2.3 Weighted Tree Connectivity Difference Signature

We approximate the *focused* objective function (3.25) using the WTCD between \mathcal{G} and the augmented graph induced from it $\mathcal{G}^{U,A}$. This topological signature is defined as

$$S_{WTCD} = \frac{n^F}{2} \log(2\pi e) - \frac{1}{2} [\tau_w - \tau_w^{U,A}]. \quad (3.29)$$

Using this signature we can get an approximated solution \hat{U} whose properties we analyze from here on.

Theorem 1. *The WTC of $\mathcal{G}^{U,A}$ asymptotically bounds $\log|\Lambda^U|$.*

To allow fluid reading, proofs for all theorems, corollaries, and lemmas are provided in the Appendices. To maintain structural clarity, each Appendix corresponds to a specific Chapter.

Theorem 2. *Let $\epsilon(J_{\mathcal{H}}^F) \doteq J_{\mathcal{H}}^F - S_{WTC D}$ be the approximation error, then it is bounded by*

$$-\frac{n}{2} \log \left(1 + \frac{\delta}{\lambda_1} \right) \leq \epsilon(J_{\mathcal{H}}^F) \leq \frac{n^U}{2} \log \left(1 + \frac{\delta^U}{\lambda_1^U} \right).$$

Lemma 1.

$$\delta/\lambda_1 \rightarrow 0^+ \Rightarrow \delta^U/\lambda_1^U \rightarrow 0^+.$$

Lemma 1 implies that when $\delta/\lambda_1 \rightarrow 0^+$, the approximation error in (3.30) approaches zero. We also note that choosing the position of the *virtual* node within the Euclidean space, can be formulated as an optimization problem to lower the upper bound.

While $S_{WTC D}$ asymptotically converges to the *focused* objective function (3.25), evaluating this signature still requires calculating the determinant of the associated Laplacian matrices. Moreover, if we would like to use these bounds to eliminate candidate actions, we need to perform eigenvalue decomposition to retrieve λ_1 , which cannot be done online.

Theorem 3. *The approximation error $\epsilon(J_{\mathcal{H}}^F)$ can also be bounded using topological aspects only, where*

$$\frac{1}{2} \left(\tau_{w_\theta} - \sum_{i=1}^n \log [W_\theta(i, i) + \delta] \right) \leq \epsilon(J_{\mathcal{H}}^F) \leq \frac{1}{2} \left(\sum_{i=1}^n \log [W_\theta^{U,A}(i, i) + \delta^{U,A}] - \tau_{w_\theta^{U,A}} \right). \quad (3.30)$$

While these bounds are somewhat more conservative, they are functions of topological aspects only. We avoid the eigenvalue decomposition and can use them online.

3.2.4 Von Neumann Difference Signature

We also propose a second topological signature (inspired by Kitanov and Indelman, 2018) to be evaluated online. The Von Neumann entropy $\hat{\mathcal{H}}_{VN}$ (Passerini and Severini, 2009) of a topological graph \mathcal{G} represents the Shannon entropy associated with the eigenvalues of its normalized Laplacian

$$\mathcal{H}_{VN}(\hat{\mathcal{L}}_w) = - \sum_{i=1}^n \frac{\hat{\lambda}_i}{2} \log \frac{\hat{\lambda}_i}{2}. \quad (3.31)$$

Using the quadratic approximation (Han et al., 2012), we rewrite (3.31) for the weighted case as

$$\mathcal{H}_{VN} \approx \tilde{\mathcal{H}}_{VN} = \frac{n \log 2}{2} - \frac{1}{2} \left(\text{Tr} \left[\hat{\mathcal{L}}_w^2 \right] - n \right). \quad (3.32)$$

Following a similar derivations (see eq. 8 in (Han et al., 2012), the trace of the square of the weighted normalized Laplacian is given by

$$\text{Tr} \left[\hat{\mathcal{L}}_w^2 \right] = n + 2 \sum_{e_{ij} \in \mathcal{E}} \frac{w_{ij}}{W_{ii} \cdot W_{jj}}. \quad (3.33)$$

Substituting (3.33) into (3.32) we get an expression for the approximated Von Neumann graph entropy for the weighted case

$$\tilde{\mathcal{H}}_{VN} \left(\hat{\mathcal{L}}_w \right) = \frac{n \log 2}{2} - \sum_{e_{ij} \in \mathcal{E}} \frac{w_{ij}}{W_{ii} \cdot W_{jj}}. \quad (3.34)$$

We denote

$$h_w = 2\tilde{\mathcal{H}}_{VN} \left(\hat{\mathcal{L}}_{w_p} \right) + \tilde{\mathcal{H}}_{VN} \left(\hat{\mathcal{L}}_{w_\theta} \right) \quad (3.35)$$

$$h_w^{U,A} = 2\tilde{\mathcal{H}}_{VN} \left(\hat{\mathcal{L}}_{w_p}^{U,A} \right) + \tilde{\mathcal{H}}_{VN} \left(\hat{\mathcal{L}}_{w_\theta}^{U,A} \right), \quad (3.36)$$

and formally define the VND topological signature as

$$S_{VND} = \frac{n^F}{2} \log(2\pi e) - \frac{1}{2} \left[h_w - h_w^{U,A} \right]. \quad (3.37)$$

Calculating (3.34) is dependent on the diagonal matrix W and generally has a quadratic complexity in the number of nodes $O(n^2)$. However, in the context of BSP, as the dimensionality of n grows with time, the information matrix and the topological representation become sparse. Evaluating (3.34), in this case, only depends on a small number of non-zero elements, i.e. the number of edges $|\mathcal{E}|$.

Incremental Aspects

Calculating (3.37) requires evaluating $\tilde{\mathcal{H}}_{VN}$ for the weighted Laplacian matrices of both \mathcal{G} and $\mathcal{G}^{U,A}$. However, we can re-use calculations rather than evaluating each from scratch. The approximated Von Neumann entropies (3.32) for a posterior belief, at time $k + L$, are given by

$$\tilde{\mathcal{H}}_{VN} \left(\hat{\mathcal{L}}_w \right) = \frac{(k + L + 1) \log 2}{2} - \sum_{e_{ij} \in \mathcal{E}} \frac{w_{ij}}{W_{ii} \cdot W_{jj}} \quad (3.38)$$

$$\tilde{\mathcal{H}}_{VN}(\hat{\mathcal{L}}_w^{U,A}) = \frac{(k + L + 2 - |n^F|) \log 2}{2} - \sum_{e_{ij} \in \mathcal{E}^{U,A}} \frac{w_{ij}^{U,A}}{W_{ii}^{U,A} \cdot W_{jj}^{U,A}} \quad (3.39)$$

Subtracting (3.39) from (3.38) and rearranging the result we get

$$\tilde{\mathcal{H}}_{VN}(\hat{\mathcal{L}}_w^{U,A}) = \tilde{\mathcal{H}}_{VN}(\hat{\mathcal{L}}_w) + \frac{(1 - |n^F|) \log 2}{2} + \Delta, \quad (3.40)$$

where

$$\Delta = \sum_{e_{ij} \in \mathcal{E}} \frac{w_{ij}}{W_{ii} \cdot W_{jj}} - \sum_{e_{ij} \in \mathcal{E}^{U,A}} \frac{w_{ij}^{U,A}}{W_{ii}^{U,A} \cdot W_{jj}^{U,A}}. \quad (3.41)$$

According to (3.27), all edges that are not connected to the *virtual* node in the augmented graph, share the same weights in both graphs. We denote an edge that does not connect to an *unfocused* node by e^{-U} and by e^V an edge which is connected to the *virtual* node. Reducing all shared terms between the two sums in (3.41), we rewrite Δ as

$$\Delta = \sum_{e_{ij}^{-U} \in \mathcal{E}} \frac{w_{ij}}{W_{ii} \cdot W_{jj}} - \sum_{e_{ij}^V \in \mathcal{E}^{U,A}} \frac{w_{ij}^{U,A}}{W_{ii}^{U,A} \cdot W_{jj}^{U,A}}. \quad (3.42)$$

Notice that evaluating Δ is dependent on the number of *focused* variables and their connectivity. The smaller they are, with respect to the entire problem, the more we gain in terms of computational costs.

Following similar steps, we can also derive a recursive update rule for calculating the VND signature for the posterior graph $S_{VND}(\mathcal{G}_{k+L})$ from the VND signature of the prior $S_{VND}(\mathcal{G}_k)$ for each candidate action. As such, calculating this signature is computationally very efficient.

Algorithm 1 Focused Topological BSP

input: Set of factor graphs, one for each candidate action \mathcal{U} and a graph signature S

output: approximate solution to FT-BSP: $\hat{\mathcal{U}}$

- 1: **for** each candidate action in \mathcal{U} **do**
 - 2: Infer topological graph \mathcal{G} (Section 3.1.2)
 - 3: Construct *unfocused* augmented topological graph $\mathcal{G}^{U,A}$ (Section 3.2.2)
 - 4: Evaluate S using \mathcal{G} and $\mathcal{G}^{U,A}$ (Eq. 3.29 or Eq. 3.37)
 - 5: **end for**
 - 6: **return** $\hat{\mathcal{U}} = \min_{\mathcal{U}} S(\mathcal{U})$
-

3.3 Experimental Results

We evaluate our approach considering a measurement selection problem and an active 2D pose SLAM simulation, to empirically study the two topological signatures.

signature	measurement selection	active SLAM
S_{WTCD}	18.88	1.21
S_{VND}	12.02	0.14
$J_{\mathcal{H}}^F$	146.24	6.34

Table 3.1: Average time in milliseconds for calculating both signatures and the *focused* objective function for each candidate action. We note that while the *WTCD* signature is calculated using a highly optimized Matlab code with a C back-end and that $J_{\mathcal{H}}^F$ is calculated in C++, we calculate the *VND* signature purely in Matlab without code optimization.

In our first experiment, we use the 2D pose SLAM Intel dataset (Carlone et al., 2014b), which contains $n = 1228$ robot poses and 278 loop-closure observations. We modified the original covariance matrix to have a block-isotropic structure. In this measurement selection problem, the goal is to find the most informative subset of observations, with respect to a random *focused* set of variables of size $n^F < n/2$. All sub-graphs share the same vertex set and contain all odometry edges. In total, we generate $n_L = 278$ such sub-graphs where, for each sub-graph we randomly choose a subset of loop closure edges. Given the *focused* set of variables, we evaluate the original *focused* objective function and the two topological signatures. This experiment was performed several times for different *focused* sets in different sizes. In Fig. 3.3 we see that both signatures are highly correlated with the *focused* objective function (3.25) given a specific *focused* set. In Table 3.1 we report the average run time statistics for all runs showing the improvement in computational cost.




Figure 3.3: Topological signatures vs marginal entropy for different loop-closure measurements given a random *focused* set of poses in the Intel dataset.

In our second experiment, we evaluate our approach in an active 2D pose SLAM

simulation. In this scenario, the robot’s objective is to reach a predefined goal with maximum accuracy, while navigating in an unknown environment. Using a Probabilistic Road Map (PRM) (Kavraki et al., 1996) we first discretize the environment. We assume that the robot previously visited some areas within the map and that all planning sessions start right after. We then randomly generate a set of candidate paths, over the road map, all ending at the predefined goal (see e.g. Fig 3.4). As such, the *focused* set of variables is defined as the last robot pose in each candidate path. i.e. the objective is to reduce the uncertainty over that position. Given all candidate paths, we evaluate the original *focused* objective function and the two topological signatures. As seen in Fig. 3.5 both signatures are highly correlated with the *focused* objective function (3.25). The figure also shows that in this setting, the bounds developed in Theorem 2 are sufficiently tight to allow action elimination. However, the Hadamard bounds developed in Theorem 3, are not informative in this case as they are not tight enough to allow action elimination. In general, Hadamard bounds are more informative in diagonally dominant matrices. In the specific case presented δ was relatively small.

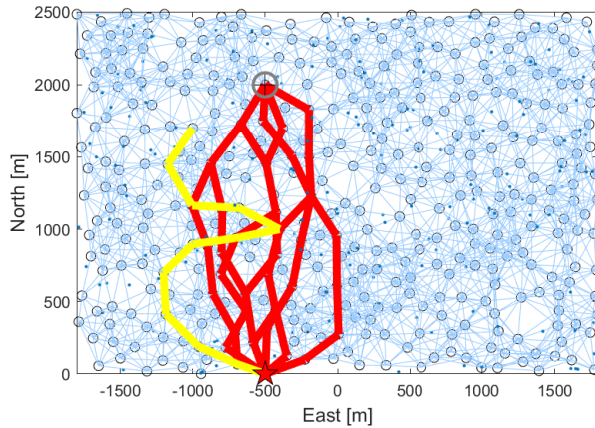


Figure 3.4: different candidate paths generated on top of a PRM in a single planning session. Starting position is denoted with a red star and the goal is denoted with a circle. The yellow path represents the locations that the robot had previously visited and acts as the initial belief.




Figure 3.5: Topological signatures in dashed lines vs $J_{\mathcal{H}}^F$ in a *focused* BSP problem. The solid lines represent the bounds developed in Theorem 2.

Chapter 4

DA-BSP with Performance Guarantees Under Budget Constraints

In the previous chapter we assumed DA to be given and perfect, i.e. assumed a single hypothesis represented by a uni-modal state and map estimates. However, in real-world scenarios, an autonomous agent should also be resilient to the problem of ambiguous measurements. These ambiguities occur when a certain observation has more than one possible interpretation. Some examples include the slip/grip behavior of odometry measurements; the loop closure problem in visual Simultaneous Localization and Mapping (SLAM); and unresolved DA. The latter is defined as the process of associating uncertain measurements to known tracks, e.g. determine if an observation corresponds to a specific landmark within a given map. While most existing inference and BSP algorithms assume DA to be given and perfect, in perceptually aliased environments, this assumption is not reasonable and could lead to catastrophic results. Therefore, it is crucial to reason about DA, in both inference and planning, while also considering other sources of uncertainty.

Explicitly reasoning about DA, the number of hypotheses grows exponentially with time. As such, when considering real time operation using inexpensive hardware, hard computational constraints are often required, e.g. bounding the number of supported hypotheses. State-of-the-art inference and planning approaches therefore use different heuristics, e.g. pruning and merging, to relax the computational complexity. However, this loss of information incurs loss in solution quality and there are usually no performance guarantees. Moreover, inference and planning are commonly treated separately and it is unclear how budget constraints in one process affect another.

Given a set of candidate actions, the main goal of BSP is to retrieve the optimal action with respect to a user defined objective function. Specifically, in the case of a multi-modal belief (corresponding to different hypotheses), a traditional solution

requires evaluating the objective function with respect to each hypothesis. Instead, we suggest to solve a *simplified* problem using only a distilled subset of hypotheses where the loss in solution quality can be bounded to provide performance guarantees.

Our contributions in this chapter are as follows:

1. We introduce a novel approach that utilizes a distilled subset of hypotheses in planning to solve a *simplified* problem and reduce computational complexity.
2. We develop the connection between our approach and the true analytical solution, owing to every possible DA.
3. We derive bounds over the true analytical solution, which can be incrementally adapted, and prove their convergence. Moreover, in a budget free scenario, these bounds are used to speed-up calculations while maintaining action consistency, i.e., preserving the same action selection as when considering all hypotheses.
4. Crucially, we address also the challenging setting of DA aware BSP with hard budget constraints, and show how these bounds provide performance guarantees.
5. We analyze the construction of a belief tree within planning given a mixture belief.
6. We show how to utilize the *skeleton* of such belief tree to reduce the computational complexity in BSP.
7. We study the impacts of hard budget constraints in both planning and inference.

4.1 Background and Notations

In this section, we review fundamental concepts from estimation theory and BSP, and formulate these problems with explicit consideration of DA.

4.1.1 Inference with Data Association

Some of the following notations are rewritten once again for convenience.

Consider an autonomous agent operating in a partially known or pre-mapped environment containing similar landmarks or scenes. The agent acquires observations and tries to infer random variables of interest that are application dependent while reasoning about DA.

We denote the agent’s state at time instant k by x_k . Let $Z_k \triangleq \{z_{k,1}, \dots, z_{k,n_k}\}$ denote the set of all n_k measurements and let u_k denote the agent’s action. $Z_{1:k}$ and $u_{0:k-1}$ denote all observations and actions up to time k , respectively. The motion and observation models are given by

$$x_{k+1} = f(x_k, u_k, w_k) \quad , \quad z_k = h(x_k, x^l, v_k), \quad (4.1)$$

where x^l is a landmark pose and w_k and v_k are noise terms, sampled from known motion and measurement distributions, respectively.

Given n_k observations, the DA realization vector is denoted by $\beta_k \in \mathbb{N}^{n_k}$. Elements in β_k are associated according to the given observation model and each element, e.g. landmark, is given a unique label. A specific DA hypothesis is thus given by a specific set j of associations up to and including time k and is denoted as $\beta_{1:k}^j$.

At each time step the agent maintains a posterior belief over both continuous and discrete variables given by

$$b[x_k, \beta_{1:k}] \triangleq \mathbb{P}(x_k, \beta_{1:k} | z_{0:k}, u_{0:k-1}) = \mathbb{P}(x_k, \beta_{1:k} | H_k), \quad (4.2)$$

where $H_k \triangleq \{Z_{1:k}, u_{0:k-1}\}$ represents history. Using the chain rule, the belief becomes a mixture and can be written as a linear combination of $|M_k|$ hypotheses

$$b_k = \sum_{j \in M_k} \underbrace{\mathbb{P}(x_k | \beta_{1:k}^j, H_k)}_{b_k^j} \underbrace{\mathbb{P}(\beta_{1:k}^j | H_k)}_{w_k^j}, \quad (4.3)$$

where b_k^j is a conditional belief, with some general distribution, and w_k^j is the associated weight. Therefore, M_k is a set of maintained weighted conditional beliefs, representing different DA hypotheses. In this chapter, we interchangeably refer to each b_k^j as both a hypothesis and a component.

Each conditional belief hypothesis b_k^j in (4.3) can be efficiently calculated by maximum a posteriori inference (e.g., as in Kaess et al., 2012) for the Gaussian case. Nevertheless, our formulation and approach also applies to a non-parametric setting. Each component weight w_k^j is calculated by marginalizing over the state space and applying the Bayes rule (Pathak et al., 2018).

4.1.2 Data Association BSP

Given a posterior belief (4.3) and a set of candidate action sequences \mathcal{U} the goal of BSP is to find the optimal action sequence that would minimize/maximize a certain objective function. We note that while in this work we consider, for simplicity, action sequences, our approach is applicable also to policies.

Reasoning about DA in planning, a user defined objective function J can be written as

$$J(b_k, u_{k:k+N-1}) = \mathbb{E}_{\beta_{(k+1)+}} \left[\mathbb{E}_{Z_{(k+1)+} | \beta_{(k+1)+}} \left[\sum_{n=1}^N c(b_{k+n}, u_{k+n-1}) \right] \right], \quad (4.4)$$

where $\beta_{(k+1)+} \triangleq \beta_{k+1:k+N}$, $Z_{(k+1)+} \triangleq Z_{k+1:k+N}$ and $c(\cdot)$ denotes a cost function. The expectation is taken with respect to both future DA realizations and observations. The

optimal action sequence $u_{k:k+N-1}^*$ is defined as

$$u_{k:k+N-1}^* = \underset{\mathcal{U}}{\operatorname{argmin}} J(b_k, u_{k:k+N-1}). \quad (4.5)$$

To solve (4.5) we need to consider all possible future realizations of Z_{k+n} for every $n \in [k+1, k+N]$ while marginalizing over all possible locations and DA realizations (see Section 5.2 in Pathak et al., 2018). However, solving these integrals analytically is typically not feasible. In practice, the solution should be approximated by sampling future observations from the relevant distributions. Using these samples, the agent constructs and traverses a belief tree (as shown in Fig. 4.1a) which branches according to future actions and observations.




Figure 4.1: (a) A belief tree constructed during planning. Each node represents a posterior belief (4.3); The number of belief components grows exponentially along the highlighted path as presented in (b).

Nevertheless, the number of hypotheses grows exponentially with the planning horizon (see Fig. 4.1b). Specifically, given $|M_k|$ hypotheses and $|D|$ DA realizations, i.e. different β_{k+i} at each look-ahead step, the number of belief components at the n th look-ahead step is $|M_{k+n}| = |M_k| |D|^n$. As such, considering every possible future hypothesis is not practical.

4.2 One Look-Ahead Step: The Myopic Case

As a first step towards applying our method for the general BSP problem (4.5), in this section we consider a myopic setting, i.e., one look-ahead step, which by itself can be computationally challenging in highly ambiguous scenarios.

We begin by formulating the belief update after performing control u_{k+1} and taking an observation Z_{k+1} . This process also requires reasoning about DA. Given M_k hypotheses from time k , marginalizing over all landmarks at time $k+1$ and using the

chain rule we explicitly write it as

$$b_{k+1} = \sum_{i=1}^{|L|} \sum_{j=1}^{M_k} \underbrace{\mathbb{P}\left(x_{k+1} | H_{k+1}, \beta_{k+1}^i, \beta_{1:k}^j\right)}_{b_{k+1}^{i,j}} \underbrace{\mathbb{P}\left(\beta_{k+1}^i, \beta_{1:k}^j | H_{k+1}\right)}_{w_{k+1}^{i,j}}, \quad (4.6)$$

where $|D|$ represents the number of different DA realizations considered at time $k+1$. The first term $b_{k+1}^{i,j}$ represents a conditional belief at time $k+1$ which originated from the j th hypothesis at time k and a specific DA realization β_{k+1}^i . The second term $w_{k+1}^{i,j}$ is the associated belief component weight.

Corollary 1. *Each posterior belief component weight $w_{k+1}^{i,j}$ can be written as*

$$w_{k+1}^{i,j} = \eta_{k+1}^{-1} \tilde{\zeta}_{k+1}^{i,j} w_k^j, \quad (4.7)$$

where w_k^j is the weight of the j th component from time k ; η_{k+1} is a normalization term; and $\tilde{\zeta}_{k+1}^{i,j}$ is the probability for the i th DA at time $k+1$ given the j th hypothesis from time k ,

$$\tilde{\zeta}_{k+1}^{i,j} \triangleq \mathbb{E}_{x_{k+1}} [\mathbb{P}\left(Z_{k+1} | \beta_{k+1}^i, x_{k+1}\right) \mathbb{P}\left(\beta_{k+1}^i | x_{k+1}\right)], \quad (4.8)$$

where the expectation is with respect to $\mathbb{P}\left(x_{k+1} | H_{k+1}^-, \beta_{1:k}^j\right)$.

The term $\mathbb{P}\left(Z_{k+1} | \beta_{k+1}^i, x_{k+1}\right)$ in (4.8) is the joint measurement likelihood for all observations obtained at time $k+1$ given the i th DA and state x_{k+1} . It can be explicitly written as

$$\mathbb{P}\left(Z_{k+1} | \beta_{k+1}^i, x_{k+1}\right) = \prod_{r=1}^{n_{k+1}} \mathbb{P}\left(z_{k+1,r} | x_{\beta_{k+1}^i(r)}^l, x_{k+1}\right), \quad (4.9)$$

where $x_{\beta_{k+1}^i(r)}^l$ denotes the landmark pose, corresponding to the r th measurement in the given DA realization vector β_{k+1}^i .

Writing the expectation operator in (4.5) explicitly, the objective function for the myopic setting is defined as

$$J(b_k, u_k) = \int_{Z_{k+1}} \eta_{k+1} c(b_{k+1}) dZ_{k+1}, \quad (4.10)$$

where $\eta_{k+1} \triangleq \mathbb{P}(Z_{k+1} | H_{k+1}^-)$ is the joint measurement likelihood, denoted from hereon simply as η . In this chapter we interchangeably refer to η as the normalization term and the measurement likelihood.

For this myopic setting, we consider using only a distilled subset of belief components $M_k^s \subseteq M_k$ from time k . We avoid calculating the posterior belief at time $k+1$ for components we do not consider in M_k^s . As such, the number of belief components

at time $k + 1$ reduces from $|M_k| |D|$ to $|M_k^s| |D|$ which also lowers the computational complexity of the considered cost function. When committing to a certain computational budget \mathcal{C} over the number of posterior belief components, the distilled subset M_k^s is subject to $|M_k^s| |D| \leq \mathcal{C}$. In this section, we only manipulate the size of M_k^s . Crucially, we analytically bound the loss in solution quality for every considered action with respect to (4.10).

We formally define a simplified belief at time k as

$$b_k^s \triangleq \sum_{j=1}^{M_k^s} w_k^{s,j} b_k^j, \quad w_k^{s,j} \triangleq \frac{w_k^j}{w_k^{m,s}}, \quad (4.11)$$

where weights are re-normalized with $w_k^{m,s} \triangleq \sum_{m \in M_k^s} w_k^m$.

To provide performance guarantees, we wish to bound (4.10), for each candidate action u_k , using b_k^s

$$\underline{J}(b_k, b_k^s, u_k) \leq J(b_k, u_k) \leq \bar{J}(b_k, b_k^s, u_k). \quad (4.12)$$

To efficiently evaluate these bounds, we simplify and analytically bound both η and the cost function terms in (4.10). Thus, we rewrite (4.12) as

$$\int_{Z_{k+1}} \mathcal{LB}[\eta] \mathcal{LB}[c(b_{k+1})] dZ_{k+1} \leq J(b_k, u_k) \leq \int_{Z_{k+1}} \mathcal{UB}[\eta] \mathcal{UB}[c(b_{k+1})] dZ_{k+1}, \quad (4.13)$$

where \mathcal{LB} , \mathcal{UB} denote lower and upper bounds, respectively. By definition, if $\mathcal{LB}[\eta]$ and $\mathcal{UB}[\eta]$ converge to η and $\mathcal{LB}[c(b_{k+1})]$, $\mathcal{UB}[c(b_{k+1})]$ converge to $c(b_{k+1})$, the bounds in (4.13) converge to $J(b_k, u_k)$.

4.2.1 Bounding the cost function

While the cost function in (4.10) can generally include a number of different terms, e.g. distance to goal, energy spent and information measures of future beliefs, in this section we only consider an information theoretic term over DA hypotheses weights that can be used for autonomous active disambiguation of hypotheses. We believe that conceptually similar derivations can also support other terms, e.g. distance to goal, and leave that for future research.

Specifically, to disambiguate between hypotheses, we utilize the Shannon entropy, defined as $\mathcal{H} \triangleq - \sum_{i=1}^n w^i \log(w^i)$, where each w^i corresponds to a belief component weight and $\sum_{i=1}^n w^i = 1$. Using Corollary 1, we rewrite \mathcal{H} as

$$c(b_{k+1}) \triangleq \mathcal{H} = - \sum_i^{|L|} \sum_j^{M_k} \frac{\tilde{\zeta}_{k+1}^{i,j} w_k^j}{\eta} \log \left(\frac{\tilde{\zeta}_{k+1}^{i,j} w_k^j}{\eta} \right). \quad (4.14)$$

To bound this cost function given a belief b_{k+1} using the same cost function given a simplified belief b_{k+1}^s , we first rigorously derive the analytic connection between the two.

Theorem 4. *Given a simplified belief b_k^s at time k , for every action u_k and considered future observation Z_{k+1} , the cost due to ambiguity (4.14) can be expressed by*

$$\mathcal{H} = \frac{w_k^{m,s}}{\eta} \left[\eta^s [\mathcal{H}^s - \log(\eta^s)] - \sum_i^{|L|} \sum_j^{M_k^s} \tilde{\zeta}_{k+1}^{i,j} w_k^{s,j} \log \left(\frac{w_k^{m,s}}{\eta} \right) \right] - \sum_i^{|L|} \sum_j^{\neg M_k^s} \frac{\tilde{\zeta}_{k+1}^{i,j} w_k^j}{\eta} \log \left(\frac{\tilde{\zeta}_{k+1}^{i,j} w_k^j}{\eta} \right), \quad (4.15)$$

where $\neg M_k^s \triangleq M_k \setminus M_k^s$; $\mathcal{H}^s \triangleq c(b_{k+1}^s)$; and $\eta^s \triangleq \mathbb{P}(Z_{k+1} | b_k^s, u_k)$.

We now use Theorem 4 to derive bounds for \mathcal{H} which are computationally more efficient to calculate as we only consider a subset of hypotheses. As can be seen in (B.5) (and also in Pathak et al., 2018 Section 4.1), evaluating η requires evaluating all posterior components weights $w_{k+1}^{i,j}$. As our considered cost is a function of these weights, simplifying and bounding \mathcal{H} has no computational merits without simplifying and bounding η (denoted below by η^s , $\mathcal{LB}[\eta]$ and $\mathcal{UB}[\eta]$).

Theorem 5. *Given a simplified belief b_k^s at time k , the cost due to ambiguity term in (4.10) is bounded by*

$$\mathcal{LB}[c(b_{k+1})] \triangleq \mathcal{LB}[\mathcal{H}] = \frac{\eta^s w_k^{m,s}}{\mathcal{UB}[\eta]} [\mathcal{H}^s - \log(\eta^s)] - \frac{w_k^{m,s}}{\mathcal{UB}[\eta]} \sum_i^{|L|} \sum_j^{M_k^s} \tilde{\zeta}_{k+1}^{i,j} w_k^{s,j} \log \left(\frac{w_k^{m,s}}{\mathcal{LB}[\eta]} \right), \quad (4.16)$$

$$\mathcal{UB}[c(b_{k+1})] \triangleq \mathcal{UB}[\mathcal{H}] = \frac{\eta^s w_k^{m,s}}{\mathcal{LB}[\eta]} [\mathcal{H}^s - \log(\eta^s)] - \frac{w_k^{m,s}}{\mathcal{LB}[\eta]} \sum_i^{|L|} \sum_j^{M_k^s} \tilde{\zeta}_{k+1}^{i,j} w_k^{s,j} \log \left(\frac{w_k^{m,s}}{\mathcal{UB}[\eta]} \right) - \gamma \log \left(\frac{\gamma}{|L| |\neg M_k^s|} \right), \quad (4.17)$$

where $\gamma \triangleq 1 - \frac{\eta^s w_k^{m,s}}{\mathcal{UB}[\eta]}$ and $|D| |\neg M_k^s| > 2$.

Furthermore, considering different levels of simplifications, i.e. adding belief components to M_k^s , these bounds become tighter.

Corollary 2. *Given a simplified belief b_k^s , the bounds developed in Theorem 5 converge to \mathcal{H} when $M_k^s = M_k$*

$$\lim_{M_k^s \rightarrow M_k} \mathcal{LB}[\mathcal{H}] = \mathcal{H} = \mathcal{UB}[\mathcal{H}]. \quad (4.18)$$

Using basic log properties, it is not hard to show that these bounds can be incrementally adapted if one chooses to add additional components to M_k^s . See full derivation in Appendix B.2.1.

4.2.2 Bounding η

In this section we derive the bounds $\mathcal{LB}[\eta]$ and $\mathcal{UB}[\eta]$ over η . We start by expressing η using η^s .

Theorem 6. *Given a simplified belief b_k^s at time k , for every action u_k and considered future observation z_{k+1} , the normalization term η in (4.10) can be expressed by*

$$\eta = w_k^{m,s} \eta^s + \sum_i^{|L|} \sum_j^{-M_k^s} \tilde{\zeta}_{k+1}^{i,j} w_k^j. \quad (4.19)$$

We can now use Theorem 6 to derive bounds for η .

Theorem 7. *Given a simplified belief b_k^s at time k , the measurement likelihood term η in (4.10) is bounded by*

$$\mathcal{LB}[\eta] = \eta^s w_k^{m,s}, \quad (4.20)$$

$$\mathcal{UB}[\eta] = \eta^s w_k^{m,s} + (1 - w_k^{m,s}) \sigma \sum_i^{|L|} \alpha^i, \quad (4.21)$$

where $\sigma \triangleq \max\left(\mathbb{P}\left(Z_{k+1}|\beta_{k+1}^i, x_{k+1}\right)\right)$ and $\alpha^i \triangleq \mathcal{UB}\left[\mathbb{P}\left(\beta_{k+1}^i|x_{k+1}\right)\right]$ is an indicator function.

As in Theorem 5, since we only consider a subset of hypotheses these bounds are also computationally more efficient to calculate and become tighter when adding belief components to M_k^s .

Corollary 3. *Given a simplified belief b_k^s , the bounds developed in Theorem 7 converge to η when $M_k^s = M_k$*

$$\lim_{M_k^s \rightarrow M_k} \mathcal{LB}[\eta] = \eta = \mathcal{UB}[\eta]. \quad (4.22)$$

Furthermore, these bounds can also be incrementally adapted if one chooses to add additional components to M_k^s . See full derivation in Appendix B.2.2.

4.2.3 Simulating future observations Z_{k+1}

Evaluating the objective function (4.10) is usually performed in two steps: we first simulate future observations Z_{k+1} by sampling from the measurement likelihood η using the generative model (4.1), and then calculate the measurement likelihood η for each such observation.

While previous works, that either use a ML assumption, (e.g., Elimelech and Indelman, 2022; Kitanov and Indelman, 2024), or do not use the ML assumption (e.g., Szttyglic and Indelman, 2022; Zhitnikov et al., 2024) all consider the likelihood terms η and η^s to be equal, to the best of our knowledge, we are the first to consider the impact of simplification on the normalization term in the myopic case.

Recall that in our proposed approach we only evaluate the bounds over η for each future observation. However, for the bounds in (4.13) to hold, we have to make sure we integrate over the same set of observations as in (4.10). To handle this issue in the myopic case, we propose propagating and sampling from the original belief rather than from the simplified belief. We note that the concept of simulating future observations is computationally not the same as calculating the measurement likelihood which requires marginalizing over all possible DAs realizations and states. As such, using the original belief to simulate future observations does not affect the computational complexity of our proposed approach.

4.2.4 Experimental Results

We evaluate the performance of our approach in a highly ambiguous environment comprising perceptually identical landmarks in different locations. Our prototype implementation uses the GTSAM library (Dellaert, 2012) with a python wrapper; all experiments were run on an Intel i7-7850 CPU running at 2200 GHz with 32GB RAM.




Figure 4.2: Given a multi-modal initial belief, the agent’s goal is to fully disambiguate between all hypotheses. (a) Run-time [sec] as a function of number of prior hypotheses M_0 ; (b) A scenario with 8 prior hypotheses, each initialized in front of a blue square and denoted by a black ellipse. Headings are denoted with cyan triangles. The component that corresponds to the correct DA hypothesis, unknown to the agent, is highlighted in yellow.

In our experiment we specifically consider five different landmark types represented by Squares, Circles, Diamonds, Pentagons and Triangles, randomly placed within the environment. The agent is initially placed in front of a blue square. With no other prior information, the initial belief is multi-modal containing M_0 hypotheses, each associated

with a blue square. This scenario can be considered as a version of the kidnapped robot problem.

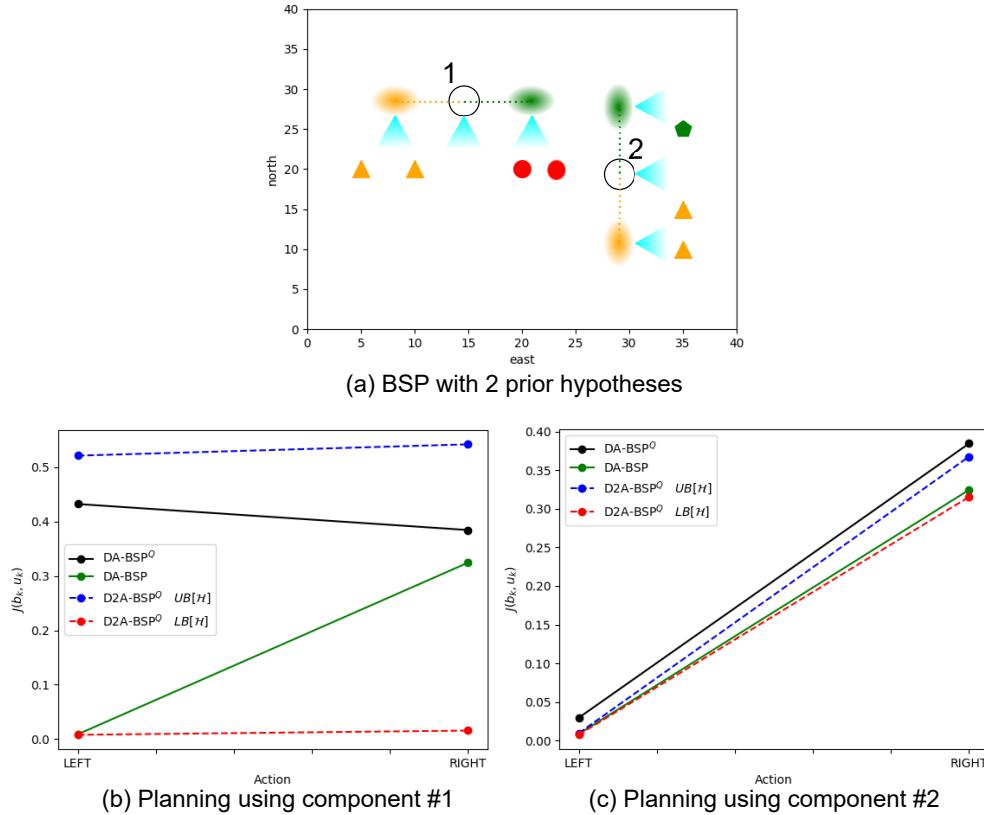


Figure 4.3: (a) A scenario under hard budget constraints of $\mathcal{C} = 6$. The agent can only move LEFT or RIGHT. Black ellipses and cyan triangles denote prior hypotheses and headings, respectively. Each prior weight equals $\frac{1}{2}$. Green and Orange ellipses denote the propagated belief after moving left or right, respectively. Moving LEFT is the best action as it is the only action that might lead to full disambiguation; (b) Objective function evaluations considering only component 1 (denoted by superscript Q). DA-BSP in green represents the solution with no budget constraints considering all hypotheses; (c) Objective function evaluations considering only component 2. Notations remain the same.

The agent’s goal is to fully disambiguate between hypotheses by solving the corresponding BSP problem (minimizing (4.10)) at each planning session, considering entropy over posterior belief components weights as a cost function. The considered actions set at each planning session contains predefined motion primitives in all four cardinal directions.

As we consider DA in inference as well, the number of belief components grows exponentially in time. For a fair comparison, we utilize the same pruning heuristics, based on a user defined weight threshold, for all approaches.

In Fig. 4.2 we see the computational merits of our approach, D2A-BSP, when there are no budget constraints. The higher the level of ambiguity within the environment, i.e. more hypotheses to reason about, the more prominent D2A-BSP becomes. In this scenario, the distilled subset M_k^s in each planning sessions is adapted greedily and

incrementally, based on prior components weights, until D2A-BSP can guarantee the same action selection as DA-BSP (Pathak et al., 2018).

Fig. 4.3 presents a scenario in which under hard budget constraints of $\mathcal{C} = 6$, DA-BSP is unable identify the best action while D2A-BSP can. Recall that we only control the size of the prior belief, i.e. as $|D| = 6$ only one component can be used each time. As can be seen, using component 1 DA-BSP^Q (under budget) selects the action RIGHT which is clearly not the best action. The bounds of D2A-BSP^Q in this case are uninformative. However, using component 2, D2A-BSP^Q guarantees that LEFT is the best action (as bounds do not overlap). While DA-BSP^Q yields LEFT as well in this case, it can only rely on heuristics to decide whether it should use component 1 or component 2.

4.3 Longer Planning Horizons: The Nonmyopic Case

We are now ready to generalize our approach and handle longer planning horizons, i.e., the nonmyopic case. We note that this case is harder since reasoning about DA, without any computational constraints, the number of considered hypotheses grows exponentially with time. In general, the belief becomes a function of $b_k = \psi_k(b_{k-1}, u_{k-1}, Z_k)$. However, under hard computational constraints, the number of hypotheses is bounded by $\mathcal{C} \in \mathbb{N}$. Therefore, the belief in each time step is a function of

$$b_k^\psi = \psi_k^{\mathcal{C}}(b_{k-1}, u_{k-1}, Z_k, \mathcal{C}), \quad (4.23)$$

where $\psi_k^{\mathcal{C}}$ contains some heuristic function h^{inf} such that $|M_k^\psi| \leq \mathcal{C}$. This has a direct impact on how the belief tree is constructed. This issue was never addressed while reasoning about DA.

In this section we first describe how to construct a belief tree *skeleton* during planning. We then present a general framework to reduce the computational complexity when solving a sampling based approximation of (4.4). Finally, we analyze the implications of using our proposed framework under different conditions.

4.3.1 Constructing the belief tree *skeleton*

Previous works addressed the exponential growth of the belief tree with the planning horizon without reasoning about DA. In this work we analyze and describe, for the first time, the structure of a belief tree given a mixture belief such as (4.3). In this setting there is an additional exponential growth in the number of belief components for every considered future observation realization (see Fig. 4.1). These realizations are functions of future beliefs (4.3), DA realizations and actions

$$\mathbb{P}(Z_{k+1:k+n} | b_k, u_{k:k+n-1}, \beta_{k+1:k+n}). \quad (4.24)$$

To construct the belief tree in practice, we sample states from beliefs, sample data association given states and finally sample observations from (4.24).

Our key observation is that in order to construct a belief tree *skeleton*, i.e. without explicitly calculating or holding posterior beliefs at each node, we can sample future observations in two different ways. We describe these two options for a planning horizon of $n = 2$. Specifically, we can either rewrite (4.24) as

$$\mathbb{P}(Z_{k+2}|b_{k+1}|\beta_{k+1}, u_{k+1}, \beta_{k+2})\mathbb{P}(Z_{k+1}|b_k, u_k, \beta_{k+1}), \quad (4.25)$$

where $b_{k+1}|\beta_{k+1}$ is a posterior belief and each term is evaluated by integrating over $x_{k+1:k+2}$, or, by first integrating and then applying the chain rule as

$$\int_{x_{k+2}} \mathbb{P}(Z_{k+2}|x_{k+2}, \beta_{k+2}) \int_{x_{k+1}} \mathbb{P}(x_{k+2}|x_{k+1}, u_{k+1})\mathbb{P}(Z_{k+1}|x_{k+1}, \beta_{k+1})\mathbb{P}(x_{k+1}|b_k, u_k). \quad (4.26)$$

While these two expressions are analytically identical, they represent two different processes of sampling. In the former observations are sampled from posterior beliefs, while in the latter observations are sampled using the motion and observation models, similar to the MCTS particle trajectories techniques (e.g., in Silver and Veness, 2010; Ye et al., 2017).

Algorithm 2 Construct belief tree *skeleton*

input: prior belief b_k , action sequence $u_{k:k+n-1}$

output: sampled future observations $Z_{k+1:k+n}$

```

1:  $Z = \emptyset$ 
2:  $x_k \sim b_k$ 
3: for  $i \in [1, n]$  do
4:    $x_{k+i} \sim \mathbb{P}(x_{k+i}|x_{k+i-1}, u_{k+i-1})$ 
5:   determine  $\beta_{k+i}$  based on  $x_{k+i}$ 
6:    $Z_{k+i} \sim \mathbb{P}(Z_{k+i}|x_{k+i}, \beta_{k+i})$ 
7:    $Z = Z \cup Z_{k+i}$ 
8: end for
9: return  $Z$ 

```

To avoid the explicit representation of the exponential number of belief components, in this work we sample future observations using (4.26) and bypass the inference stage. We formulate this sampling method in Algorithm 2.

Yet, this is of little help if the posterior belief is required for calculating the cost function itself. We next describe our approach to avoid these calculations.

4.3.2 Methodology

Recall that our goal is to reduce the computational complexity of nonmyopic BSP problems where ambiguous DA is explicitly considered, i.e. solving (4.5) efficiently. We

start by writing (4.4) in a recursive form

$$J(b_k, u_{k:k+N-1}) = c(b_k, u_k) + \mathbb{E}_{\beta_{k+1}} \left[\mathbb{E}_{Z_{k+1}|\beta_{k+1}} [J(b_{k+1}, u_{k+1:k+N-1})] \right]. \quad (4.27)$$

As in practice we approximate the solution via samples, we rewrite (4.27) as

$$\hat{J}(b_k, u_{k:k+N-1}) = c(b_k, u_k) + \hat{\mathbb{E}}_{\beta_{k+1}} \left[\hat{\mathbb{E}}_{Z_{k+1}|\beta_{k+1}} [\hat{J}(b_{k+1}, u_{k+1:k+N-1})] \right]. \quad (4.28)$$

Using Bellman's principle of optimality, the optimal solution for (4.28) is

$$\hat{J}(b_k, \hat{u}_{k:k+N-1}^*) = \min_{u_k} \left\{ c(b_k, u_k) + \hat{\mathbb{E}}_{\beta_{k+1}} \left[\hat{\mathbb{E}}_{Z_{k+1}|\beta_{k+1}} [\hat{J}(b_{k+1}, u_{k+1:k+N-1}^*)] \right] \right\}, \quad (4.29)$$

where $\hat{u}_{k:k+N-1}^* = \underset{u}{\operatorname{argmin}} \hat{J}(b_k, u_{k:k+N-1})$. To reduce the computational complexity in (4.29), we propose utilizing the belief tree *skeleton*, without having access to posterior beliefs, to solve an easier to compute version of the considered cost function. In general,

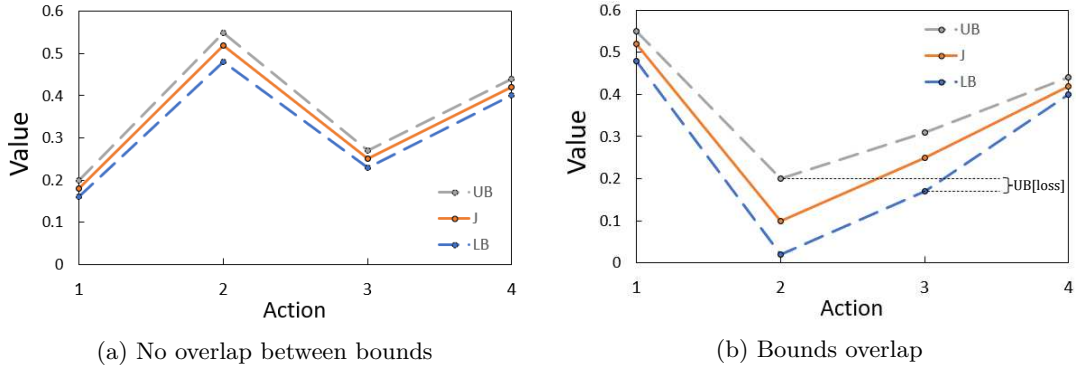


Figure 4.4: BSP using bounds over the objective function. In (a) choosing action #1 is guaranteed to be optimal as the corresponding upper bound is lower than all other lower bounds; In (b) choosing action #2 is not guaranteed to be optimal. The loss in solution quality, however, is upper bounded.

the cost function over the original beliefs can be bounded using a *simplified* belief b_k^s as

$$\underline{c}(b_k^s, u_k) \leq c(b_k, u_k) \leq \bar{c}(b_k^s, u_k). \quad (4.30)$$

We note that this formulation also supports replacing the cost function itself with a computationally simpler function (e.g., as in Kitanov and Indelman, 2018).

Using the belief tree *skeleton* and some method to calculate the *simplified* beliefs, to be defined, we now traverse the belief tree from the leaves upwards. At each belief tree node the bounds over the objective function (4.4) are calculated recursively using

the Bellman equation (4.29) and (4.30) for every $n \in [0, N - 1]$ such that

$$\begin{aligned} \underline{J}(b_{k+n}, u_{(k+n)+}) &= \underline{c}(b_{k+n}^s, u_{k+n}) + \hat{\mathbb{E}}_{\beta_{k+1}} \left[\hat{\mathbb{E}}_{Z_{k+n+1} | \beta_{k+1}} \left[\underline{J}(b_{k+n+1}, u_{(k+n)+}) \right] \right], \\ \bar{J}(b_{k+n}, u_{(k+n)+}) &= \bar{c}(b_{k+n}^s, u_{k+n}) + \hat{\mathbb{E}}_{\beta_{k+1}} \left[\hat{\mathbb{E}}_{Z_{k+n+1} | \beta_{k+1}} \left[\bar{J}(b_{k+n+1}, u_{(k+n)+}) \right] \right], \end{aligned} \quad (4.31)$$

where $u_{(k+n)+} \triangleq u_{k+n:k+N-1}$. If these bounds do not overlap (see Fig. 4.4a), one can guarantee to select the optimal action sequence as in (4.29).

Algorithm 3 Generic Nonmyopic Distilled Data Association BSP

input: belief tree *skeleton* T , simplification heuristic h , decision rule R

output: action sequence u^* , loss

```

1: Function ND2A-BSP( $T, h, R$ )
2:    $LB^*, UB^*, loss = \text{PLAN}(T.root, h, R)$ 
3:    $u^* \leftarrow$  corresponding to  $LB^*, UB^*$ 
4:   return  $u^*, loss$ 
5:
6: Function PLAN( $Node, h, R$ )
7:    $Node.b_{k+n}^s \leftarrow h(Node)$ 
8:   if  $Node$  is a leaf
9:     return  $\underline{c}(Node.b_{k+n}^s), \bar{c}(Node.b_{k+n}^s), 0$  // loss = 0 at leaf
10:   $Node.bounds = \emptyset$ 
11:  foreach child  $C$  of  $Node$  do
12:     $lb, ub, loss \leftarrow \text{ND2A-BSP}(C, h, R)$ 
13:     $LB \leftarrow \underline{c}(Node.b_{k+n}^s) + lb$  // objective lower bound (4.31)
14:     $UB \leftarrow \bar{c}(Node.b_{k+n}^s) + ub$  // objective upper bound (4.31)
15:     $Node.bounds = Node.bounds \cup (LB, UB)$ 
16:  while  $R(Node.bounds)$  is not satisfied do
17:    ND2A-BSP( $Node, h, R$ ) // further simplification is needed
18:   $LB^*, UB^*, loss \leftarrow Node.bounds$ 
19:  return  $LB^*, UB^*, loss$ 

```

Our general Nonmyopic Distilled Data Association BSP (ND2A-BSP) approach is presented in Algorithm 3. The algorithm receives a belief tree *skeleton*; a heuristic function h used to select the subsets of hypotheses in each belief tree node, i.e. defines b_{k+n}^s ; and a decision rule R which decides whether the considered subsets are enough, e.g. when no overlap between bounds is required or when calculations exceed a user defined time threshold, providing anytime performance guarantees. The algorithm returns the best action sequence, given the computational constraints, and an upper bound on the loss in solution quality.

It is worth mentioning that our approach can be adapted to a setting where the belief tree construction is coupled with Q function estimates, e.g. using MCTS and Upper Confidence Bound (UCB) techniques (Silver and Veness, 2010), following a similar approach to the one presented in Szytylic et al., 2021. However, we emphasize that

as the belief tree *skeleton* approximates (4.27) via samples, our method provides performance guarantees with respect to that specific *skeleton*, i.e. with respect to (4.29). Not to be confused with the asymptotic guarantees of MCTS approaches, with respect to the theoretical problem (4.27), which is an entirely different aspect not related to the approach presented in this work.

We now analyze different settings, within inference and planning, where the agent either has or does not have hard budget constraints. To the best of our knowledge, this is the first time that these aspects are addressed in works that attempt to reduce the computational complexity of the planning problem. The differences between the considered settings are summarized in Table 4.1.

	budget constraints in inference	budget constraints in planning
Case 1	✗	✗
Case 2	✗	✓
Case 3	✓	✗
Case 4	✓	✓

Table 4.1: A summary of the considered scenarios, with respect to budget constraints on the number of supported hypotheses in each algorithm, for each considered case. Cases 1&2 are presented in Section 4.3.3 while cases 3&4 are presented in section 4.3.4

4.3.3 Inference without Budget Constraints

In this section we assume that there are no constraints in inference, i.e. each belief tree node can theoretically hold every possible hypothesis within the planning horizon. The objective of inference however is different than the main goal of BSP. In inference the agent tries to represent the considered state as accurately as possible while in planning the goal is to retrieve the optimal action sequence or policy. As such, in this setting, the problems are decoupled (see Fig. 4.5a).

We now further separate between two cases, when the planning algorithm either has budget constraints or not. In both cases, each belief tree node still has an exponential number of components, which we avoid calculating explicitly.

Case 1

With no budget constraints in planning we propose bounding the cost function as

$$\underline{c}(b_k, u_{k+}, Z_{(k+1)+}, b_{k+n}^s) \leq c(b_{k+n}, u_{k+n}) \leq \bar{c}(b_k, u_{k+}, Z_{(k+1)+}, b_{k+n}^s). \quad (4.32)$$

where $u_{k+} \triangleq u_{k:k+n-1}$ and $Z_{(k+1)+} \triangleq Z_{k+1:k+n}$. A key difference from the approach presented in Sztyglic and Indelman, 2022 is that these bounds are not functions of b_{k+n} .

To manage the exponential growth in the number of belief components, we refrain from calculating $c(b_{k+n})$. Instead, we calculate a *simplified* belief b_{k+n}^s , using Bayesian

updates via $u_{k:k+n-1}$ and $Z_{k+1:k+n}$, only for specific components from the prior belief b_k . Each *simplified* belief is formally defined, using $M_{k+n}^s \subseteq M_{k+n}$ components, as

$$b_{k+n}^s \triangleq \sum_{r \in M_{k+n}^s} w_{k+n}^{s,r} b_{k+n}^r \quad , \quad w_{k+n}^{s,r} \triangleq \frac{w_{k+n}^r}{w_{k+n}^{m,s}}, \quad (4.33)$$

where $w_{k+n}^{m,s} \triangleq \sum_{m \in M_{k+n}^s} w_{k+n}^m$ is used to re-normalize each corresponding weight. Most importantly, a *simplified* belief b_{k+n}^s is calculated using only a subset of hypotheses, i.e. without calculating the posterior belief b_{k+n} .

Using Algorithm 3 given a decision rule R , with no overlap between bounds (4.31), and a heuristic h , e.g. which chooses hypotheses greedily based on prior weights, we guarantee the selection of the optimal actions sequence, with respect to the specific belief tree, while reducing the computational complexity.

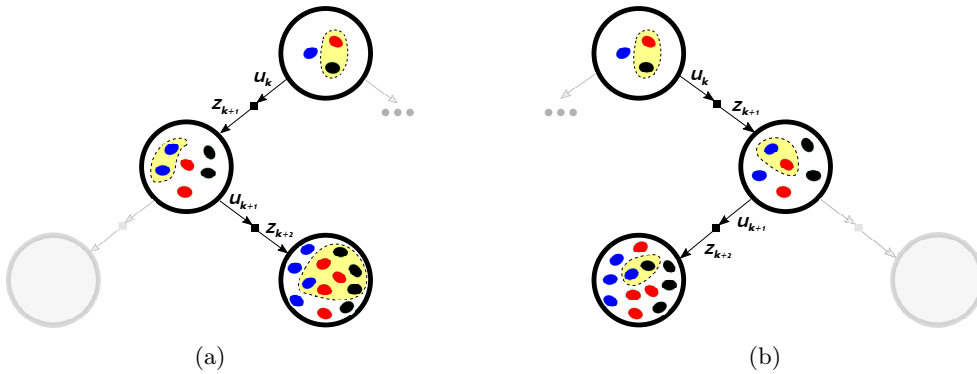


Figure 4.5: No budget constraints in inference. Each belief tree node holds different hypotheses denoted by colored particles. Different colors belong to hypotheses generated from previous time steps. (a) Planning without budget constraints, the algorithm can choose any subset of components, highlighted in yellow, in each node to evaluate the bounds; (b) With budget constraints in planning, each subset selection is bounded in size by $\mathcal{C} = 2$.

Case 2

Under budget constraints in planning, the algorithm can use up to \mathcal{C} components, in each *simplified* belief b_{k+n}^s , to calculate the bounds in (4.32). Yet, each subset of components is chosen independently w.r.t. b_{k+n} which develops exponentially, i.e. hypotheses chosen in time steps $k+n$ and $k+n+i$ are not necessarily related (see Fig. 4.5b).

In this setting, the number of possible distilled subsets for each b_{k+n}^s is $\binom{|M_{k+n}|}{\mathcal{C}}$ which can be very high. Moreover, there are no guarantees that the bounds between candidate actions would not overlap. However, using the bounds in (4.32), our proposed approach can yield the worst-case loss in solution quality, i.e. provide performance guarantees (see Fig. 4.4b).

4.3.4 Hard budget constraints in inference

In the previous section we only considered that the belief at the root of the tree is provided from inference. As the posterior beliefs within the constructed belief tree were with an exponentially increasing number of components, i.e. without budget constraints, the key idea was to avoid making explicit inferences. Instead, we calculated bounds that utilized, under budget constraints in planning, a fixed number of components. In practice, however, real world autonomous systems do not work that way. Instead, they are often required to operate in real time using inexpensive hardware with hard computational budget constraints in both inference and planning.

Under hard budget constraints on the number of considered hypotheses in inference, the posterior belief in each belief tree node is determined by (4.23), i.e. $|M_{k+n}^{\psi}| \leq \mathcal{C}$ under some heuristic h^{inf} . Moreover, once a hypothesis is discarded in time step k it is no longer considered in future time steps. Yet, the decision regarding which components to choose, while calculating the bounds in planning, depends on either if the heuristic in (4.23) is given or determined within planning. To the best of our knowledge, the latter is a novel concept never considered.




Figure 4.6: Hard budget constraints in inference. Each belief tree node holds different hypotheses denoted by colored particles. Different colors belong to hypotheses generated from previous time steps. (a) Planning given the heuristic in inference, the algorithm can only evaluate the bounds using components that represent how the belief would evolve in inference; (b) The planning algorithm is free to choose components under any valid heuristics in inference given the budget \mathcal{C} . Each selected component in time step $k + n + 1$ must originate from a selected component in time step $k + n$.

Case 3

In this setting we consider the heuristic in (4.23) to be given within planning, i.e. posterior belief tree nodes exactly represent how the belief would evolve in inference under (4.23). In contrast to Section 4.3.3, as the number of components does not grow exponentially, we sample future observations according to (4.25) and construct the belief tree explicitly, i.e. perform inference in each node. Therefore, the planning algorithm can no longer choose any subset of components for each b_{k+n}^s , i.e. hypotheses discarded in time step $k + n$ cannot be considered in time step $k + n + 1$ (see Fig. 4.6).

The bounds over the considered cost are now a function of the belief in the previous time step under (4.23). Specifically, we rewrite them as

$$\underline{c}\left(b_{k+n-1}^\psi, u_{k+n-1}, Z_{k+1}, b_{k+n}^s\right) \leq c\left(b_{k+n}^\psi, u_{k+n}\right) \leq \bar{c}\left(b_{k+n-1}^\psi, u_{k+n-1}, Z_{k+1}, b_{k+n}^s\right). \quad (4.34)$$

These bounds represent a recursive setting in contrast to the bounds in (4.32).

Using our approach iteratively in each time step, reduces the computational complexity of the considered cost function in planning while providing performance guarantees. As each posterior belief is determined by inference (Fig 4.6a), performance guarantees are with respect to the given heuristic in inference (4.23).

Case 4

We now relax the assumption that the planning algorithm is confined to the specific heuristic in (4.23). Unlike in Case 2, where each subset of components can be used in each node to calculate the bounds, this setting has an additional constraint. We formulate this by representing the bounds from (4.32) in two consecutive time steps

$$\begin{aligned} \underline{c}\left(b_k, u_{k+}, Z_{(k+1)+}, b_{k+n}^s\right) &\leq c\left(b_{k+n}\right) \leq \bar{c}\left(b_k, u_{k+}, Z_{(k+1)+}, b_{k+n}^s\right), \\ \underline{c}\left(b_k, u_{k+}, Z_{(k+1)+}, b_{k+n+1}^s\right) &\leq c\left(b_{k+n+1}\right) \leq \bar{c}\left(b_k, u_{k+}, Z_{(k+1)+}, b_{k+n+1}^s\right), \\ \text{s.t. } |M_{k+n}^s|, |M_{k+n+1}^s| &\leq \mathcal{C} \quad \text{and} \quad \forall b_{k+n+1}^{s,ij} \in b_{k+n+1}^s \Rightarrow b_{k+n}^{s,j} \in b_{k+n}^s, \end{aligned} \quad (4.35)$$

where $b_{k+n}^{s,j}$ denotes the j th hypothesis in the *simplified* subset b_{k+n}^s and $b_{k+n+1}^{s,ij}$ denotes the i th hypothesis in the *simplified* subset b_{k+n+1}^s , originated from $b_{k+n}^{s,j}$, i.e. as in Fig. 4.1b.

The components chosen in the sequence of bounds (4.35) which minimizes the loss, w.r.t. the original problem, define a heuristic h^{p^*} (see Fig. 4.9c), which is valid in inference. The heuristic h^{p^*} can be used with any BSP approach to solve (4.29) and to reduce computational complexity, using our approach, as described in Case 3. To the best of our knowledge, leveraging h^{p^*} is a novel concept. We note that while h^{p^*} minimizes the loss in planning, it is generally different than h^{inf} . As such, the implications of utilizing such heuristic in inference are not straightforward. The study of such mechanism is left for future research.

4.3.5 Information Theoretic Cost Function

While the formulation thus far was for a general cost function, in this section we focus on active disambiguation of hypotheses. Specifically, we utilize the Shannon entropy, defined over posterior belief components weights. The cost for a belief b_{k+n} with M_{k+n}

components is thus given by

$$\mathcal{H}_{k+n} \triangleq c(b_{k+n}) = - \sum_{r \in M_{k+n}} \frac{w_{k+n}^r}{\eta_{k+n}} \log \left(\frac{w_{k+n}^r}{\eta_{k+n}} \right), \quad (4.36)$$

where $\eta_{k+n} \triangleq \sum_{r \in M_{k+n}} w_{k+n}^r$. Similarly, for a *simplified* belief b_{k+n}^s with $M_{k+n}^s \subseteq M_{k+n}$ the cost is given by

$$\mathcal{H}_{k+n}^s \triangleq c(b_{k+n}^s) = - \sum_{r \in M_{k+n}^s} w_{k+n}^{s,r} \log \left(w_{k+n}^{s,r} \right). \quad (4.37)$$

Theorem 8. *For each belief tree node representing a belief b_{k+n} with M_{k+n} components and a subset $M_{k+n}^s \subseteq M_{k+n}$ the cost can be expressed by*

$$\mathcal{H}_{k+n} = \frac{w_{k+n}^{m,s}}{\eta_{k+n}} \left[\mathcal{H}_{k+n}^s + \log \left(\frac{\eta_{k+n}}{w_{k+n}^{m,s}} \right) \right] - \sum_{r \in \neg M_{k+n}^s} \frac{w_{k+n}^r}{\eta_{k+n}} \log \left(\frac{w_{k+n}^r}{\eta_{k+n}} \right), \quad (4.38)$$

where $\neg M_{k+n}^s \triangleq M_{k+n} \setminus M_{k+n}^s$.

Using Theorem 8, we derive bounds for \mathcal{H}_{k+n} which are computationally more efficient to calculate as we only consider a subset of hypotheses. However, as evaluating η_{k+n} requires by definition evaluating all posterior components weights, which we do not have access to, we need to bound this term as well (denoted below as $\mathcal{LB}[\eta_{k+n}]$ and $\mathcal{UB}[\eta_{k+n}]$).

Theorem 9. *Given a subset of components $M_{k+n}^s \subseteq M_{k+n}$, the cost term in each belief tree node is bounded by*

$$\mathcal{LB}[\mathcal{H}_{k+n}] = \frac{w_{k+n}^{m,s}}{\mathcal{UB}[\eta_{k+n}]} \left[\mathcal{H}_{k+n}^s + \log \left(\frac{\mathcal{LB}[\eta_{k+n}]}{w_{k+n}^{m,s}} \right) \right], \quad (4.39)$$

$$\mathcal{UB}[\mathcal{H}_{k+n}] = \frac{w_{k+n}^{m,s}}{\mathcal{LB}[\eta_{k+n}]} \left[\mathcal{H}_{k+n}^s + \log \left(\frac{\mathcal{UB}[\eta_{k+n}]}{w_{k+n}^{m,s}} \right) \right] - \bar{\gamma} \log \left(\frac{\bar{\gamma}}{|\neg M_{k+n}^s|} \right), \quad (4.40)$$

where $\bar{\gamma} = 1 - \sum_{r \in M_{k+n}^s} \frac{w_{k+n}^r}{\mathcal{UB}[\eta_{k+n}]}$ and $|\neg M_{k+n}^s| > 2$.

Furthermore, considering different levels of simplifications, i.e. adding belief components to M_{k+n}^s , these bounds converge.

Corollary 4. *The bounds in Theorem 9 converge to \mathcal{H}_{k+n} when $M_{k+n}^s \rightarrow M_{k+n}$*

$$\lim_{M_{k+n}^s \rightarrow M_{k+n}} \mathcal{LB}[\mathcal{H}_{k+n}] = \mathcal{H}_{k+n} = \mathcal{UB}[\mathcal{H}_{k+n}]. \quad (4.41)$$

A recursive update rule is given in Appendix B.3.1.

Theorem 10. Given a subset of components $M_{k+n}^s \subseteq M_{k+n}$, the term η_{k+n} , in each belief tree node, is bounded by

$$\mathcal{LB}[\eta_{k+n}] = w_{k+n}^{m,s} \leq \eta_{k+n} \leq w_{k+n}^{m,s} + \left(\frac{|M_{k+n}|}{|M_k|} - \sum_{r \in M_{k+n}^s} w_k^r \right) \prod_{i=1}^n \sigma^i = \mathcal{UB}[\eta_{k+n}], \quad (4.42)$$

where $\sigma^i \triangleq \max(\mathbb{P}(Z_{k+i}|x_{k+i}))$ and w_k^r is the prior weight at time k for every component in M_{k+n}^s at time $k+n$.

As in Theorem 9, since we only consider a subset of hypotheses, these bounds are also computationally more efficient to calculate and converge. We also note that specifically for Case 3, the bounds in Theorem 9 and Theorem 10 are calculated iteratively in each time step $k+n$ given the belief b_{k+n-1}^ψ as presented in (4.34).

Corollary 5. The bounds in Theorem 10 converge to η_{k+n} when $M_{k+n}^s \rightarrow M_{k+n}$

$$\lim_{M_{k+n}^s \rightarrow M_{k+n}} \mathcal{LB}[\eta_{k+n}] = \eta_{k+n} = \mathcal{UB}[\eta_{k+n}]. \quad (4.43)$$

A recursive update rule is given in Appendix B.3.2.

4.3.6 Experimental Results

We evaluate the performance of our approach for the different cases presented. Our prototype implementation uses the GTSAM library (Dellaert, 2012). Our considered

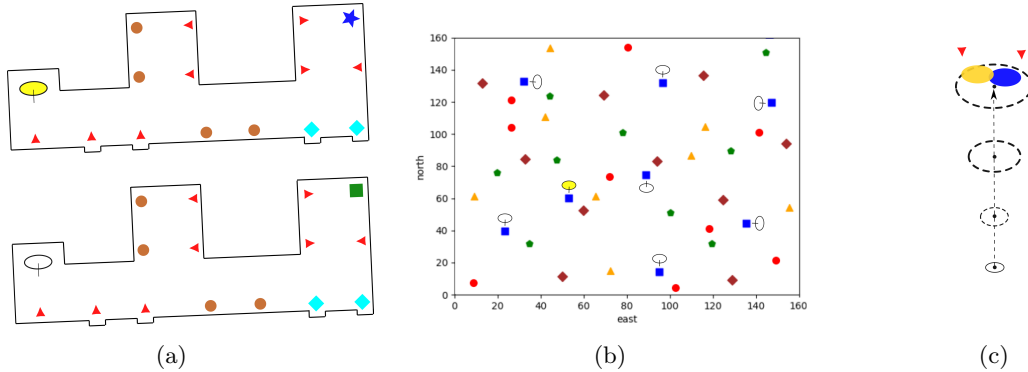


Figure 4.7: (a) The `floors` environment where F identical floors represent different prior hypotheses. Each floor contains a unique landmark. The true location of the agent is highlighted in yellow; (b) The `2d_random` environment with many identical landmarks. The agent is initially placed in front of a blue square with no other prior information; (c) A planning session where ambiguous DA results in two hypotheses denoted by the yellow and blue ellipses.

scenarios represent highly ambiguous environments containing perceptually identical landmarks in different locations. In our first scenario, `floors`, the agent is initially located in one of F floors such that each floor contains a unique landmark, specific to that floor (Fig. 4.7a). In our second scenario, `2d_random`, the agent is initially placed

in a random environment in front of a blue square (Fig. 4.7b). Both scenarios can be considered as versions of the kidnapped robot problem. With no other prior information, the initial belief, in both cases, is multi-modal containing $|M_0|$ hypotheses. The agent captures the environment using range measurements containing a class identifier, e.g. red triangle or green square. When the agent receives a measurement to some landmark which is ambiguous, i.e. it can theoretically be generated from more than one landmark, the number of hypotheses grows (see Fig. 4.7c). The number of identical landmarks can be adjusted to represent higher ambiguity, increasing the number of considered hypotheses. The agent’s goal is to disambiguate between hypotheses by solving the corresponding BSP problem (4.5) at each planning session using entropy over posterior belief components weights as a cost function.

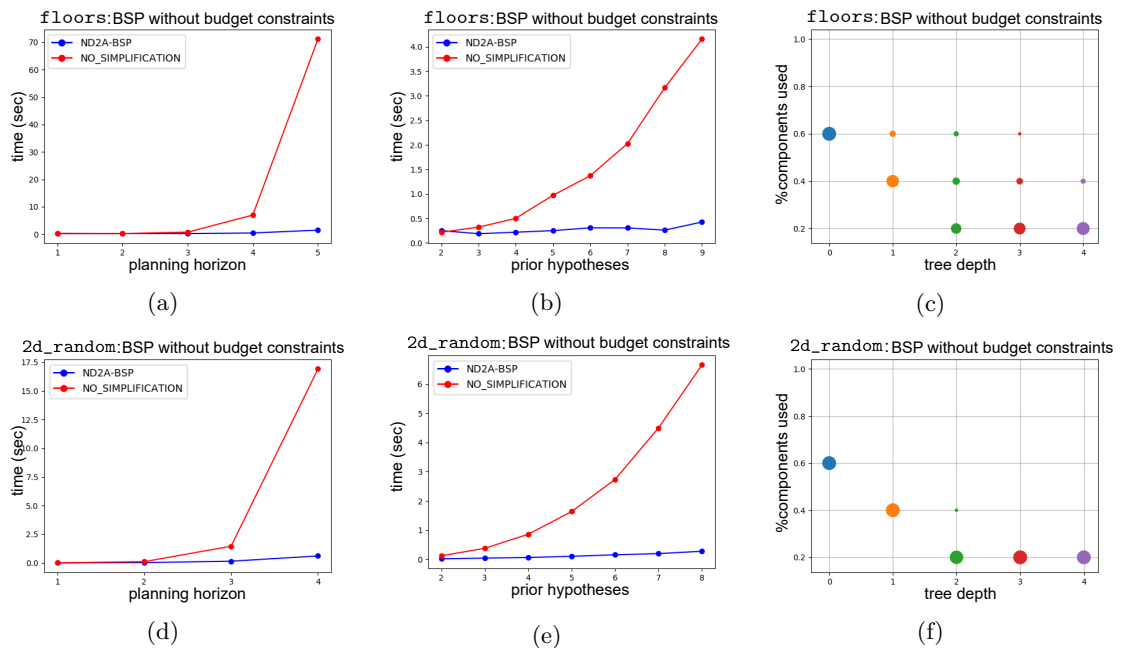


Figure 4.8: Case 1 study for `floors` and `2d_random` environments. All scenarios presented carry zero loss. (a),(d) Planning time as a function of the planning horizon. In both environments, all settings the considered 4 prior hypotheses; (b),(e) Planning time as a function of the number of prior hypotheses. In both environments, all settings considered a planning horizon of 3; (c),(f) % components used to calculate bounds in each level of the belief tree. Circles scales are normalized as the number of nodes grows exponentially going down the tree. Less components are used to calculate the bounds as tree depth increases. This is expected, as with longer horizons less components are needed for disambiguation.

In our first experiment we consider Case 1. We compare our approach with evaluating the cost function over the original belief, i.e. considering every possible future hypothesis. The heuristic in planning chooses the subset of components for each belief tree node greedily based on prior weights at time k . The decision rule R was set as no overlap, i.e. no loss with guaranteed optimal solution. The computational merits of our approach are presented in Fig. 4.8. Moreover, in Fig. 4.8c,4.8f we can see that with a longer planning horizon the subset of hypotheses used for disambiguation becomes

smaller. As more observations are utilized along the horizon, it is easier to discard wrong hypotheses in our considered cases.




Figure 4.9: (a) Normalized loss as a function of the size of the budget in Case 2. In all settings the number of floors, i.e. prior hypotheses was set to 12; (b) Normalized loss along the depth of a belief tree in Case 2 with $\mathcal{C} = 3$ components and a planning horizon of $n = 2$; (c) Normalized loss as a function of the size of the budget in Case 4, i.e. considering every valid heuristic in inference $h^{p_candidate}$. When $\mathcal{C} \leq 6$, the heuristic h^{p^*} induces a smaller loss than h^{inf} . When $\mathcal{C} > 6$, both h^{p^*} and h^{inf} induce zero loss, i.e. are optimal in this setting.

In our second experiment we consider Case 2. In Fig. 4.9a we present the loss as a function of the budget size. As expected, with higher budget constraints the loss in solution quality becomes smaller. Moreover, as can be seen in Fig. 4.9b the loss is higher closer to the root of the belief tree, as bounds are accumulated in the non-myopic setting, increasing the overlap.

Considering Case 3, our experiments did not show any computational improvements between calculating the original cost function and using our approach. We indicate that this is because there is no exponential growth in the number of hypotheses within the horizon and our considered cost function is linear w.r.t. the number of components. However, using a different cost, which is beyond the scope of this work, our approach can reduce the computational complexity while providing guarantees in Case 3 as well (as seen in Szttyglic and Indelman, 2022).

Finally, we consider Case 4. We first report that under this setting the computational complexity is high as every possible heuristic under the given budget is considered. In Fig. 4.9c preliminary results indicate that this process can improve the bounds over the loss in solution quality vs a given heuristic h^{inf} .

Chapter 5

A Slices Perspective for Incremental Nonparametric Inference

In the previous chapter, we assumed that beliefs could be represented as a weighted linear combination of parametric models. Our aim now is to extend this research to address nonparametric posterior beliefs, which can have arbitrary shapes. To achieve this, we begin by examining the nonparametric inference problem.

Substantial research efforts have been made in recent decades to develop probabilistic inference algorithms that are robust, accurate and capable of real time performance. These efforts often rely on various assumptions to achieve their objectives. Perhaps the most frequently employed assumption is that the actual posterior distribution can be approximated using a parametric Gaussian model. However, in real-world problems, the posterior distribution is often non-Gaussian, having multiple modes or a nonparametric structure. Due to the complex, non-Gaussian nature of such posterior distributions, obtaining closed-form analytical solutions is challenging and frequently impractical.

To approximate and model such non-Gaussian posterior distributions, modern non-parametric methods use different variants of the *forward-backward* algorithm, exploiting graphical models such as factor graphs (Kschischang et al., 2001) and BT (Kaess et al., 2012). In these methods, samples are generated at each step to reconstruct intermediate distributions through techniques such as KDE and various learning procedures.

Our key observation is that these distributions can be directly reconstructed without the need for any additional learning techniques or KDE. This is accomplished by accessing *slices* from high-dimensional surfaces that represent partial joint distributions (Fig. 5.1). By circumventing these processes, our approach proves to be computationally more efficient, requires significantly less samples and produces more accurate results, as supported by our experimental findings.

Our main contributions in this chapter are as follows:




Figure 5.1: (a) The factor graph of a small probabilistic inference problem containing two variable nodes X, Y and two factor nodes $f(X), f(X, Y)$ indicated as small solid black circles; (b) A high-dimensional surface representing the joint distribution $\mathbb{P}(X, Y)$. Each red slice, at a specific realization $Y = y$, represents a conditional distribution $\mathbb{P}(X|Y = y)$. The marginal distribution $\mathbb{P}(X)$, shown in blue, is calculated by integrating over all conditional slices.

1. We introduce a nonparametric inference approach which leverages slices from high-dimensional surfaces to approximate joint and marginal posterior distributions without any further intermediate reconstructions.
2. Unlike previous methods, ours does not require any iterative procedures nor breaking the underlying graphical model to generate samples in each step, even in cases where unary factors are not available.
3. We show how to utilize our slices perspective for nonparametric incremental inference and propose a novel early stopping heuristic criteria to further speed up calculations.
4. Our approach requires less samples and consistently outperforms state of the art nonparametric inference algorithms in terms of accuracy and computational complexity. This superiority is evident across evaluations conducted on both synthetic and real-world datasets, with improvements in time complexity reaching up to an order of magnitude.

Our proposed inference approach is general and can be effectively employed in various estimation problems such as tracking, sensor fusion, Bundle Adjustment (BA), Structure from Motion (SfM) and SLAM. In this particular work, we focus on demonstrating our approach within the framework of SLAM.

5.1 Background and Notations

Some of the following notations are rewritten once again for convenience.

Consider a SLAM framework in which an autonomous agent is operating in a partially known environment. At each time step k the agent takes an action u_k and acquires an observation $Z_k \triangleq \{z_k^1, \dots, z_k^n\}$ based on n measurements. The motion and observation models are given by

$$x_{k+1} = f(x_k, u_k, w_k), \quad z_k = h(x_k, x^l, v_k), \quad (5.1)$$

where x^l denotes a landmark pose and w_k, v_k are noise terms, sampled from known motion and measurement distributions, respectively. In this work specifically, while the motion and measurement distributions are known, they can have any arbitrary shape, i.e. not necessarily Gaussian.

We denote the set of all state variables, including all agent poses and observed landmarks, by Θ . Given all actions and measurements, the joint pdf, the belief, is given by

$$b_k[\Theta] \triangleq \mathbb{P}(\Theta | u_{0:k-1}, Z_{0:k}) = \mathbb{P}(\Theta | D), \quad (5.2)$$

where $D \triangleq \{b_0, u_{0:k-1}, Z_{0:k}\}$ represents all available data at time instant k . To reduce clutter, we omit time notations from hereon, and refer to the above joint density as $\mathbb{P}(\Theta | D)$. Our goal is to compute the joint posterior distribution (5.2) and marginal posterior distributions $\mathbb{P}(\theta | D)$, $\forall \theta \in \Theta$.




Figure 5.2: (a) A factor graph formulation of a SLAM problem with five variable nodes and seven factor nodes. Factor nodes represent probabilistic information over random variables. In this example, factor nodes include: a prior $f(x_1)$, odometry measurements $f(x_1, x_2)$, $f(x_2, x_3)$ and landmark measurements $f(x_1, l_1)$, $f(x_2, l_1)$, $f(x_2, l_2)$, $f(x_3, l_2)$. The factor graph represents a factorization of the joint distribution as a product of all factors; (b) The corresponding Bayes net after performing variable elimination on the factor graph using the elimination order $\mathcal{O} = \{\theta_1 = x_1, \theta_2 = l_1, \theta_3 = x_2, \theta_4 = l_2, \theta_5 = x_3\}$. The joint distribution is expressed as the product of conditionals produced in each step as $\mathbb{P}(x_1, x_2, x_3, l_1, l_2) = \mathbb{P}(x_1 | x_2, l_1) \cdot \mathbb{P}(l_1 | x_2) \cdot \mathbb{P}(x_2 | x_3, l_2) \cdot \mathbb{P}(l_2 | x_3) \cdot \mathbb{P}(x_3)$.

We use a factor graph model to represent the joint posterior distribution (Fig. 5.2a). A factor graph $G = (\mathcal{F}, \Theta, \mathcal{E})$ is a bipartite graph with two types of nodes: factor nodes $f \in \mathcal{F}$ and variable nodes $\theta \in \Theta$. The variable nodes represent the random variables in the estimation problem, whereas the factor nodes represent probabilistic information on those variables. The edges in \mathcal{E} encode connectivity based on the variables associated with each factor, with each edge $e(f, \theta) \in \mathcal{E}$ linking a factor node to a variable node. This probabilistic graphical model represents a factorization of the posterior

distribution in terms of process and measurement models

$$\mathbb{P}(\Theta|D) \propto \prod_i f_i(\Theta_i), \quad (5.3)$$

where $\Theta_i = \{\theta \mid \exists e(f_i, \theta) \in \mathcal{E}\}$ is the set of all variables adjacent to factor f_i . In this work we assume for simplicity that only unary and pairwise factors exist in G , i.e. each $f \in \mathcal{F}$ can be connected with up to two variable nodes.

A common approach for computing the posterior distribution (5.2) and the posterior marginals $\mathbb{P}(\theta|D)$, $\forall \theta \in \Theta$ is the *forward-backward* algorithm which operates in two passes given a specific variable ordering $\mathcal{O} \triangleq \text{ord}(\Theta)$.

During the *forward* pass, the factor graph is gradually transformed into a Bayes net (Pearl, 1988) through a bipartite elimination game (Heggernes and Matstoms, 1996). In each step, a single variable is eliminated from the factor graph following the elimination order \mathcal{O} , starting with a variable linked to a prior factor. Specifically, each step begins with a factor graph $G_{j-1} = (\mathcal{F}_{j-1}, \Theta_{j-1}, \mathcal{E}_{j-1})$, where $G_0 = G$ by convention. A variable node θ_j , representing the j th variable in \mathcal{O} , and all factors $\mathcal{F}_{j-1}(\theta_j) \triangleq \{f \mid \exists e(f, \theta_j) \in \mathcal{E}_{j-1}\}$ are first removed from G_{j-1} along with the corresponding edges, i.e. $\mathcal{E}_{j-1}(\theta_j) \triangleq \{e(f, \theta) \mid f \in \mathcal{F}_{j-1}(\theta_j), \theta \in \Theta_{j-1}\}$. All variables involved in $\mathcal{E}_{j-1}(\theta_j)$, except for θ_j , define a separator S_j representing the Markov blanket of variable node θ_j in G_{j-1} . Next, a joint density is defined by the product of all removed factors as

$$\mathbb{P}_{joint}(\theta_j, S_j | D_j) = \eta_j^{-1} \prod_{f_i \in \mathcal{F}_{j-1}(\theta_j)} f_i(\Theta_i), \quad (5.4)$$

where η_j^{-1} is a normalizing term and D_j represents all available data given by all factors removed up to and including the elimination of θ_j , i.e. all data in $\mathcal{F} \setminus \mathcal{F}_j$. Using the chain rule, the joint density (5.4) is then factorized as

$$\mathbb{P}_{joint}(\theta_j, S_j | D_j) = \mathbb{P}(\theta_j | S_j, D_j) f_{new}(S_j | D_j). \quad (5.5)$$

The conditional $\mathbb{P}(\theta_j | S_j, D_j)$ is added as a new node to the Bayes net and the factor $f_{new}(S_j | D_j)$ is added into the factor graph G_j . Formally, each $G_j = (\mathcal{F}_j, \Theta_j, \mathcal{E}_j)$ is recursively defined by

$$\begin{aligned} \mathcal{F}_j &= \mathcal{F}_{j-1} \setminus \mathcal{F}_{j-1}(\theta_j) \cup \{f_{new}(S_j | D_j)\}, \\ \Theta_j &= \Theta_{j-1} \setminus \{\theta_j\}, \\ \mathcal{E}_j &= \mathcal{E}_{j-1} \setminus \mathcal{E}_{j-1}(\theta_j) \cup \{e(f_{new}(S_j | D_j), \theta) \mid \theta \in S_j\}. \end{aligned}$$

Note that in each intermediate step we have both an incomplete Bayes net and a reduced factor graph which defines a density on the remaining variables.

Once all variables were eliminated the *forward* pass is completed. The joint distribution (5.3) can then be expressed as the product of conditionals produced in each step

as

$$\mathbb{P}(\Theta|D) = \prod_{j=1}^{|\Theta|} \mathbb{P}(\theta_j|S_j, D_j), \quad (5.6)$$

which defines the Bayes net (Fig. 5.2b). The marginal distribution $\mathbb{P}(\theta_j|D)$, for each variable θ_j , is calculated using backsubstitution in a second, *backward* pass. Starting from the last eliminated variable and following the elimination order \mathcal{O} in reverse, each marginal $\mathbb{P}(\theta_j|D)$ is calculated by integrating over S_j

$$\mathbb{P}(\theta_j|D) = \int_{S_j} \mathbb{P}(\theta_j, S_j|D) dS_j = \int_{S_j} \mathbb{P}(\theta_j|S_j, D) \cdot \mathbb{P}(S_j|D) dS_j. \quad (5.7)$$

Note that by definition, the separator for the last eliminated variable is an empty set, thereby yielding the direct marginal of the last eliminated variable. A key observation is that due to conditional independence $\mathbb{P}(\theta_j|S_j, D) = \mathbb{P}(\theta_j|S_j, D_j)$. As such, each marginal can be rewritten as

$$\mathbb{P}(\theta_j|D) = \int_{S_j} \mathbb{P}(\theta_j|S_j, D_j) \cdot \mathbb{P}(S_j|D) dS_j. \quad (5.8)$$

In this work, we put forth a solution that extends to cases where no closed form solutions for calculating the marginal posterior distributions exist and the noise terms are non-Gaussian. In such cases, posterior distributions must be approximated.

5.2 Methodology

Our key observation is that a joint probability function, for several random variables, can be seen as a *high-dimensional surface* from which conditional and marginal distributions can be calculated. For example, consider the joint distribution $\mathbb{P}(X, Y)$ over two random variables X and Y . Using the chain rule, we rewrite the joint distribution as

$$\mathbb{P}(X, Y) = \mathbb{P}(X|Y) \cdot \mathbb{P}(Y), \quad (5.9)$$

where the conditional $\mathbb{P}(X|Y)$ is given by a specific *slice* from the high-dimensional surface for each realization of Y . See illustration in Fig 5.1.

The marginal distribution $\mathbb{P}(X)$ is calculated using the Chapman-Kolmogorov transit integral

$$\mathbb{P}(X) = \int_Y \mathbb{P}(X|Y) \cdot \mathbb{P}(Y) dY. \quad (5.10)$$

As there is no closed form solution for (5.10) in the general case, it can be approximated using N samples of Y from $\mathbb{P}(Y)$. The corresponding estimated marginal is given by

$$\hat{\mathbb{P}}(X) = \hat{\mathbb{E}}_{y \sim \mathbb{P}(Y)} [\mathbb{P}(X|Y = y)] = \frac{1}{N} \sum_{i=1}^N \mathbb{P}(X|Y = y^i). \quad (5.11)$$

Put differently, evaluating (5.11) can be seen as accessing the high-dimensional surface at specific sample points of Y in order to extract a *mixture* of conditional *slices* over X . Moreover, when $N \rightarrow \infty$ this approximation is guaranteed to converge to the true analytical solution.

5.2.1 Inference Using Slices

For nonparametric inference, our proposed approach leverages the *forward-backward* algorithm presented in Sec. 5.1. For each eliminated variable θ_j in the *forward* pass, the joint density defined in (5.4) is factorized into a conditional and a marginal via (5.5) which is rewritten once again for convenience

$$\mathbb{P}(\theta_j, S_j | D_j) = \mathbb{P}(\theta_j | S_j, D_j) f_{new}(S_j | D_j). \quad (5.12)$$

This factorization serves as the pivotal step through which information propagates across the algorithm. In strike contrast to existing approaches that depend on intermediate density reconstructions, our primary contribution is the direct use of *slices* to approximate both $f_{new}(S_j | D_j)$ and the conditional $\mathbb{P}(\theta_j | S_j, D_j)$. We rigorously describe how our *slices* perspective is systematically employed throughout the algorithm.

Forward Pass

We start by explicitly writing the term $f_{new}(S_j | D_j) = \int_{\theta_j} \mathbb{P}_{joint}(\theta_j, S_j | D_j) d\theta_j$ in (5.12) as

$$f_{new}(S_j | D_j) = \eta_j^{-1} \int_{\theta_j} \prod_{f_i \in \mathcal{F}_{j-1}(\theta_j)} f_i(\Theta_i) d\theta_j, \quad (5.13)$$

where η_j^{-1} is a normalizing term. As there is no closed form solution for (5.13), it must be approximated. While other approaches use additional intermediate approximations, such as KDE (e.g, in Fourie et al., 2016) or learning techniques (e.g., in Huang et al., 2023), we employ *slices* as a more direct and effective means of approximation. We note that we assume the ability to access such high-dimensional *slices*, an assumption shared with mm-iSAM, NF-iSAM, and NSFG. By generating N samples of θ_j from some factor $f' \in \mathcal{F}_{j-1}(\theta_j)$, a standard Monte Carlo estimator of (5.13) is given by

$$\hat{f}_{new}(S_j | D_j) = \frac{\eta_j^{-1}}{N} \sum_{n=1}^N \prod_{f_i \in \mathcal{F}_{j-1}(\theta_j) \setminus \{f'\}} f_i(\theta_j^n, \Theta_i \setminus \{\theta_j\}), \quad (5.14)$$

where $\Theta_i \setminus \{\theta_j\} \subseteq S_j$ and each $f_i \in \mathcal{F}_{j-1}(\theta_j) \setminus \{f'\}$ is evaluated using the samples of θ_j . In a *slices* perspective, for each sample θ_j^n , a high-dimensional conditional *slice* over $\Theta_i \setminus \{\theta_j\}$ is retrieved from every non-unary factor $f_i \in \mathcal{F}_{j-1}(\theta_j)$. We acknowledge a slight misuse of notation, as D_j is no longer purely theoretical due to the fact that the conditioned data is approximated at each step by \hat{f}_{new} (see **Example** below).




Figure 5.3: Generating samples during variable elimination; (a) A subset of three variables x_1, x_2, l_1 from a larger factor graph. The given elimination order is $\mathcal{O} = \{\theta_1 = x_1, \theta_2 = l_1, \dots\}$. When eliminating x_1 , samples of x_1 are directly obtained from $f(x_1)$ to approximate the new factor $\hat{f}_{new}(l_1, x_2|D_1)$ with *slices* via (5.14); (b) When eliminating l_1 , samples of l_1 are obtained according to Lemma 2 from $\hat{f}_{new}(l_1, x_2|D_1)$. Even without a unary factor on l_1 , our *slices* approach provides a method to directly generate samples of l_1 .

Obtaining samples of θ_j is essential for approximating (5.13). However, generating these samples becomes non-trivial when no unary factor $f \in \mathcal{F}_{j-1}(\theta_j)$ is available for direct sampling. To handle this issue, mm-iSAM employs multiscale Gibbs sampling, which is computationally expensive. In contrast, both NF-iSAM and NSFG break the factor graph to chain like structures and generate samples using ancestral sampling, which does not leverage loop closure factors. Through direct approximation of (5.13) with *slices* in each step, our approach eliminates the need for iterative procedures or factor graph decomposition into chain-like structures for generating samples of θ_n . Unlike previous methods, this holds true even in the absence of unary factors.

Lemma 2. *Given a factor graph $G = (\mathcal{F}, \Theta, \mathcal{E})$ and an elimination order \mathcal{O} , if each eliminated variable $\theta_j \in \Theta$ either has a unary factor connected to it, i.e. $\exists f(\theta_j) \in \mathcal{F}_{j-1}(\theta_j)$ or, $\exists \theta_i \in \Theta$ such that θ_i was previously eliminated and $f(\theta_i, \theta_j) \in \mathcal{F}$, then samples of θ_j can be drawn from one of the factors $\mathcal{F}_{j-1}(\theta_j)$.*

Example: *We demonstrate how Lemma 2 is utilized to draw samples when there are no unary factors using the example in Fig. 5.3. After eliminating the first variable $\theta_1 = x_1$, a new factor $\hat{f}_{new}(l_1, x_2|D_1)$ is added to the factor graph. Consequently, the second variable to be eliminated, $\theta_2 = l_1$, lacks any unary factors at this stage. According to (5.4), the joint density $\mathbb{P}_{joint}(l_1, x_2|D_2)$ is given by*

$$\mathbb{P}_{joint}(l_1, x_2|D_2) = \hat{f}_{new}(l_1, x_2|D_1) \cdot f(l_1, x_2), \quad (5.15)$$

and is factorized, according to (5.5), as

$$\mathbb{P}_{\text{joint}}(l_1, x_2 | D_2) = \mathbb{P}(l_1 | x_2, D_2) \cdot f_{\text{new}}(x_2 | D_2), \quad (5.16)$$

given the separator $S_2 = \{x_2\}$. As previously stated, it is important to note that D_2 does not encompass all raw data up to this point. Instead of $D_2 = \{b_0, u_1, z_1, z_2\}$, where $b_0 \equiv f(x_1)$ is given by the prior, we have $D_2 = \{\hat{f}_{\text{new}}(l_1, x_2 | D_1), z_2\}$ due to the approximation from the previous elimination step. We write $f_{\text{new}}(x_2 | D_2)$ explicitly, following (5.13) and considering the estimator $\hat{f}_{\text{new}}(l_1, x_2 | D_1)$ obtained via (5.14). Denoting it intermediately by $\tilde{f}_{\text{new}}(x_2 | D_2)$ and changing the order of the integral and summation, yields

$$\tilde{f}_{\text{new}}(x_2 | D_2) = \frac{\eta^{-1}}{N} \sum_{n=1}^N f(x_1^n, x_2) \int_{l_1} f(x_1^n, l_1) \cdot f(l_1, x_2) dl_1. \quad (5.17)$$

We can now exploit the unary structure of $f(x_1^n, l_1)$, for each x_1^n , to generate samples of l_1 and approximate (5.17) as

$$\hat{f}_{\text{new}}(x_2 | D_2) = \frac{\eta^{-1}}{N^2} \sum_{n_1=1}^N f(x_1^{n_1}, x_2) \sum_{n_2=1}^N f(l_1^{n_1, n_2}, x_2).$$

Unlike other methods, as previously discussed, our slices approach exclusively utilizes structure for sample generation when there are no unary factors.

We next turn our attention to the conditional term $\mathbb{P}(\theta_j | S_j, D_j)$ in (5.12). By utilizing \hat{f}_{new} , given by (5.14), and replacing (5.4), an estimator of the conditional is obtained

$$\hat{\mathbb{P}}(\theta_j | S_j, D_j) = \frac{\prod_{f_i \in \mathcal{F}_{j-1}(\theta_j)} f_i(\Theta_i)}{\frac{1}{N} \sum_{n=1}^N \prod_{f_i \in \mathcal{F}_{j-1}(\theta_j) \setminus \{f'\}} f_i(\theta_j^n, \Theta_i \setminus \{\theta_j\})}. \quad (5.18)$$

We note that the normalization term η_j^{-1} cancels out. Moreover, there is no need to explicitly calculate η_j^{-1} in (5.14) as intermediate new factors created during the *forward* pass are only used to generate samples of S_j . During the *backward* pass, *slices* from the high-dimensional surface representing the numerator, are used to approximate the conditional (5.18) as described next.

Backward Pass

Once the *forward* pass concludes, we utilize a similar *slices* perspective, leveraging the constructed Bayes net, to retrieve the joint and marginal posterior distributions. Following the elimination order \mathcal{O} in reverse, we obtain an estimator of the marginal (5.7) for each variable θ_j . This is achieved by utilizing the estimator of the conditional $\mathbb{P}(\theta_j | S_j, D_j)$, which was already evaluated during the *forward* pass (5.18), and samples

of S_j directly obtained from the marginal $\hat{\mathbb{P}}(S_j|D)$ which is given by a *mixture* of *slices*. Specifically, given N samples of $S_j \sim \hat{\mathbb{P}}(S_j|D)$, the estimator is given by

$$\hat{\mathbb{P}}(\theta_j|D) = \frac{1}{N} \sum_{n=1}^N \hat{\mathbb{P}}(\theta_j|S_j^n, D_j), \quad (5.19)$$

which is by itself also a *mixture* of high-dimensional *slices*. For each sample S_j^n , a high-dimensional conditional *slice* over θ_j is retrieved from every non-unary factor $f_i \in \mathcal{F}_{j-1}(\theta_j)$. The weight of each *slice* is given by evaluating the denominator of (5.18) at sample point S_j^n . We note that (5.19) represents the marginal as a distribution. As such, we can both sample from the marginal and evaluate the likelihood for a given sample.

Similarly, the joint posterior distribution (5.2) can also be approximated. By replacing each conditional in (5.6) with (5.18), we can explicitly represent the joint distribution from which we can also both sample and evaluate the likelihood for a given sample. Sampling from the joint distribution is done by following the elimination order \mathcal{O} in reverse. We first generate a sample of the last eliminated variable directly from the corresponding marginal. Next, we recursively sample each variable from the corresponding conditional (5.18), given a sample of the separator, to obtain a single sample from the joint distribution.

Up to this point, our approach has focused on a batch processing method applied to a given factor graph. However, in real-world applications, a real-time online solution is needed. This requires continuous updates and providing estimates whenever new measurements are added.

5.2.2 Incremental Inference with Early Stopping Heuristic

Incremental inference was initially developed for the Gaussian case (Dellaert and Kaess, 2006; Kaess et al., 2008; Kaess et al., 2012). The underlining observation in these approaches is that the addition of new measurements affects only specific parts the original factor graphs. As a result, there’s no need to redo the entire *forward-backward* inference process. Instead, an incremental *forward* pass is conducted, focusing computations solely on the affected variables. This process, termed re-elimination, significantly reduces computational costs by reusing previously computed information and selectively updating relevant sections of the factor graph (See Algorithm 1 in Indelman et al., 2015b for a simplified version). Furthermore, performing full backsubstitution in every iteration can be computationally expensive, particularly in large-scale problems with complex factor graphs. Consequently, iSAM2 (Kaess et al., 2012) incorporated an early stopping heuristic in the downward pass to significantly reduce computational cost in the non-linear Gaussian case.

By utilizing the BT (Kaess et al., 2012), both mm-iSAM (Fourie et al., 2016) and NF-iSAM (Huang et al., 2023) achieve nonparametric incremental inference capabilities.

Nevertheless, these algorithms restrict the application of incremental aspects solely to the upward pass, while the entire downward pass is executed anew. Moreover, training new conditional samplers for affected variables in NF-iSAM, during re-elimination, must be performed from scratch, resulting in computational inefficiency.

To integrate incremental aspects into both the *forward* and *backward* passes, we suggest the adoption of re-elimination and an early stopping heuristic. This combination, which serves to notably diminish computational complexity, is a novel concept introduced for the first time in a nonparametric setting to the best of our knowledge.

We employ the Maximum Mean Discrepancy (MMD) (Gretton et al., 2012) metric as a heuristic for early stopping of the *backward* pass. The MMD is a metric used to assess the dissimilarity between two probability distributions, relying on samples. During the *backward* pass of each incremental step, we calculate and cache the most recent marginal distributions. Given a new marginal and a previously found marginal for a given variable, we use the MMD metric to evaluate the distance between the two distributions. If this distance is below a user defined threshold, we stop the *backward* pass. Notably, we have two hyper-parameters available for fine-tuning, providing a balance between accuracy and efficiency: the number of samples used for evaluating the MMD distance and the threshold criteria.

5.3 Experimental Results

We evaluate our proposed *slices* approach using both synthetic and real-world datasets. To assess its performance, we compare our results with those obtained using mm-iSAM (Fourie et al., 2016), NF-iSAM (Huang et al., 2023), and NSFSG (Huang et al., 2022), employing the open-source code provided by the respective authors. We use the Root Mean Square Error (RMSE) metric to measure the discrepancies between ground truth and samples from posterior distributions (As in Following Huang et al., 2022; Huang et al., 2023). For robustness, all reported results for each method were averaged over ten independent runs.

5.3.1 Synthetic Dataset - Multi Modal Four Doors

The synthetic four doors localization example (Fourie et al., 2016) is a one-dimensional SLAM problem. The robot is aware of a map with four identical doors but initially observes only one of them, without knowing which specific door it is. Consequently, the prior distribution in this scenario is multi-modal, with four modes of equal weight. As the robot progresses through several steps, it gathers loop closure measurements related to a landmark within its environment. Additionally, it acquires two more measurements, each corresponding to different doors. These latter measurements play a critical role in resolving the ambiguity, leading to the collapse of the posterior distribution into a single mode.

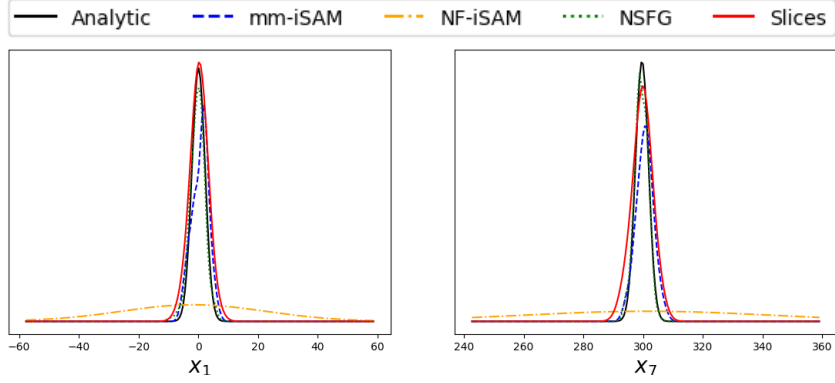


Figure 5.4: Four doors synthetic SLAM dataset (Fourie et al., 2016) marginal posterior distributions for the first and last robot poses. A black solid line indicates the analytic solution when considering only the mode which corresponds to the ground truth. Our *slices* approach directly approximates these marginal via (5.18). In mm-ISAM, KDEs are used to approximate these marginal distributions, whereas in NF-iSAM and NSFG, they are derived solely from samples and approximated using KDEs based on these samples (as demonstrated in Huang et al., 2022). In each method, 200 samples were utilized to approximate the distributions.

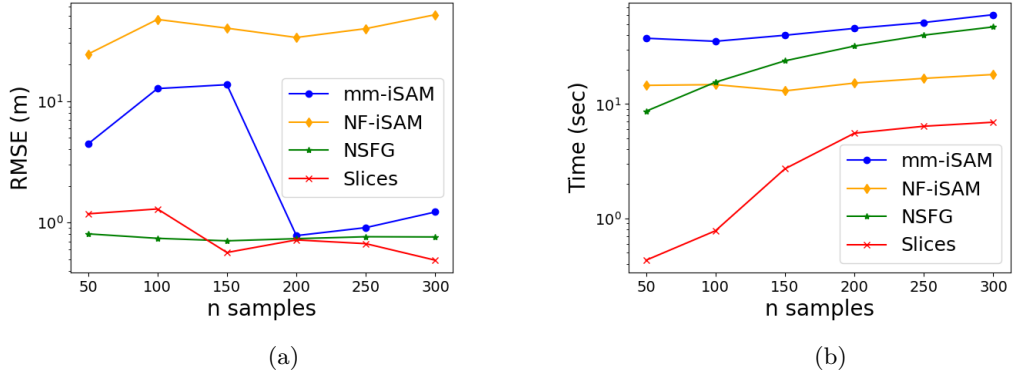


Figure 5.5: Four doors synthetic SLAM dataset (Fourie et al., 2016) results. We report RMSE (m) and run time (sec) as functions of the number of samples used by all methods.

This one dimensional problem is interesting as the marginal posterior distributions for all robot poses and the landmark can be analytically calculated (as discussed in Fourie et al., 2016). Our *slices* approach directly approximates these marginal distributions via (5.19). In mm-ISAM these marginal distributions are approximated with KDEs while in NS-iSAM and NSFG they are only available through samples (Fig. 5.4). Note how NF-ISAM fails to recover the correct mode, even in this small scenario, due to its difficulty in converging when the number of samples is relatively small.

In this experiment, we assess run time and RMSE as functions of the number of samples generated by the algorithms in each step, employing a batch computation approach across the entire factor graph. Our *slices* approach demonstrates superior accuracy compared to both mm-iSAM and NF-iSAM, achieving results on par with NSFG (Fig. 5.5), despite the latter being recognized as an offline algorithm. Furthermore, our approach shows significantly improved computational efficiency when using

the same number of samples.

5.3.2 Real World Dataset - Plaza

The Plaza2 dataset (Djugash et al., 2009) comprises a sequence of odometry and range measurements to four unknown landmarks gathered by a vehicle navigating a planar environment. This range measurement SLAM problem is challenging because it involves both nonlinear measurements and non-Gaussian likelihood models.

In this experiment, we perform incremental inference and evaluate both the run time and RMSE at each step. We adopt the same noise models across all methods (following Huang et al., 2023), for a fair comparison. We also set our hyperparameters, for the early stopping heuristic, to $N_M = 100$ and $\delta = 1e^{-4}$, representing the number of samples used for evaluating the MMD distance and the threshold criteria, respectively. The chosen re-elimination order for affected variables, in each incremental step, is such that landmark variables are eliminated last.

We configured our *slices* approach to generate 150 samples in each step. For NF-iSAM, we followed the recommended setting of 2000 samples and used the same hyperparameters as published by the authors in Huang et al., 2023. As for mm-iSAM, we employed the default value of 100 samples, as advised in the open source implementation of Fourie et al., 2016. Notably, when we attempted to use 200 samples, the mm-iSAM run time became prohibitively long for this specific problem. The scale and dimensionality of this problem were too demanding for NSFG (Huang et al., 2022) where the authors only presented results for a limited number of time steps. Notably, our *slices* approach achieves superior accuracy compared to NF-iSAM and mm-iSAM (Fig. 5.7). Furthermore, our method significantly enhances computational efficiency, reducing complexity by an order of magnitude compared to NF-iSAM and mm-iSAM. An illustration depicting the marginal distribution across different time steps for the second landmark is shown in Fig. 5.6.

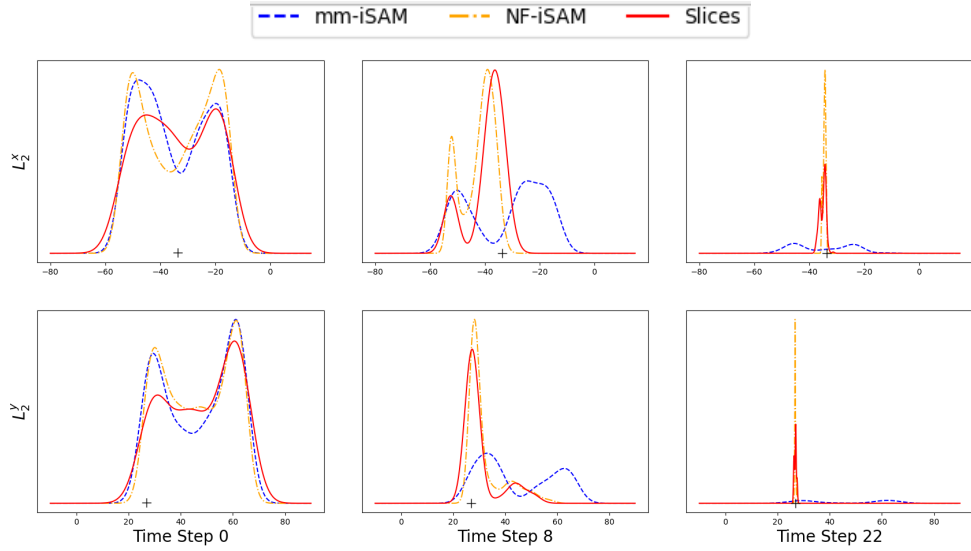


Figure 5.6: Plaza2 dataset (Djugash et al., 2009) marginal posterior distributions for landmark number 2 across different time steps, presented separately for the x and y axes. The robot captures the first, fifth and ninth range measurement to L_2 at time steps 0, 8 and 22 respectively. The ground truth position of L_2 is denoted by a black + sign. The hyperparameters used for each method are described in Section V.B .

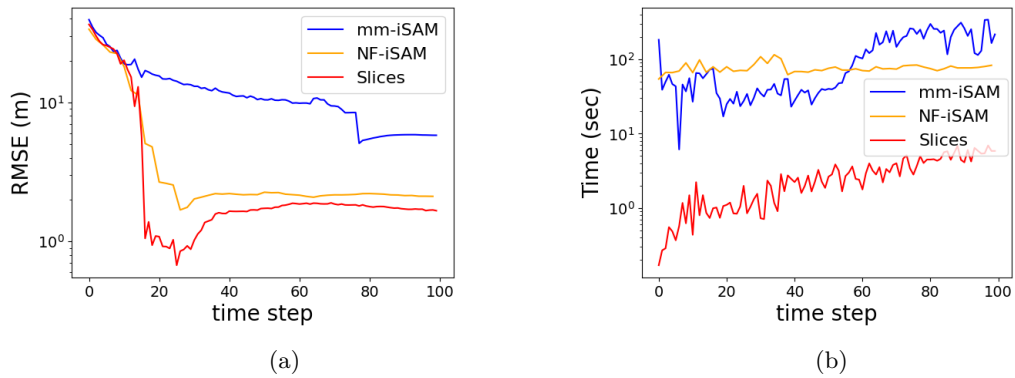


Figure 5.7: Plaza2 dataset (Djugash et al., 2009) with range measurements. We report RMSE (m) and run time (sec) for each incremental step.

Chapter 6

Conclusion

To enable efficient solutions for inference and BSP in high-dimensional state spaces, this thesis has introduced methods that utilize various structures to alleviate computational complexity in real-world problems characterized by ambiguous environments and strict computational budget constraints. In conclusion, we will summarize our contributions and explore potential directions for extending this work.

In Chapter 3, we introduced a novel concept focusing on topological aspects for decision-making under uncertainty concerning a *focused* set of variables. We developed two topological signatures, WTCD and VND, tailored for information-theoretic problems, and demonstrated a strong empirical correlation between these signatures and the *focused* information-theoretic objective function in two distinct scenarios. We established bounds for the approximation error of the WTCD signature, showing that under specific conditions, it converges to the optimal solution, enabling effective discrimination between candidate actions. Furthermore, we introduced the Hadamard bounds for WTCD, which can be calculated online and rely solely on topological aspects. Although these bounds are generally less informative and primarily useful for action elimination in cases with diagonally dominant matrices, they offer valuable insights. For the VND signature, we demonstrated its computational efficiency and its capability for incremental updates, supporting real-time sequential decision-making problems. Our results indicated that both signatures provide a substantial speed advantage over computing the objective function, with notable improvements observed in active SLAM experiments. Yet, in this chapter we assumed DA to be given and perfect, i.e. assumed a single hypothesis represented by a uni-modal state and map estimates.

Chapter 4 introduced an innovative approach that leverages a distilled subset of hypotheses to mitigate computational complexity in DA-aware BSP problems. We developed analytical bounds to quantify the loss in quality of our proposed solution, starting with the myopic case. We demonstrated how our approach can achieve optimal solutions even under strict budget constraints, unlike other methods that depend solely on heuristics and cannot guarantee such results. We then addressed the more complex non-myopic case, where the exponential growth in the number of hypotheses with the

planning horizon poses additional challenges. We examined the distinct impacts of hard budget constraints on both inference and planning in this context, redeveloped analytical bounds for each case, and introduced a general algorithm capable of handling such constraints with performance guarantees on the loss in solution quality. Our approach was validated in extremely aliased simulated scenarios, where it achieved a substantial reduction in computational complexity compared to existing methods. In this chapter, we assumed that beliefs could be expressed as a weighted linear combination of parametric models. Our goal was to expand this research to handle nonparametric posterior beliefs, which can take on any arbitrary shape.

In Chapter 5, we introduced a novel approach using *slices* from high-dimensional surfaces to efficiently approximate nonparametric posterior distributions. Unlike existing methods that rely on generated samples and intermediate approximations, our *slices* perspective avoids such procedures and iterative methods. We also introduced an early stopping heuristic during the backward pass, reducing computational complexity and enabling real-time operation. Our *slices* perspective demonstrated superior accuracy compared to other online nonparametric inference methods and matched the performance of state-of-the-art offline methods while achieving significant computational complexity reductions in both synthetic and real-world datasets.

6.1 Future Research Directions

Future research could explore integrating the various methods proposed in this thesis to enhance their effectiveness. For instance, combining topological signatures with DA reasoning in planning may further reduce computational complexity. This integration could leverage the strengths of both approaches, offering a more streamlined and efficient solution. Additionally, applying our *slices* perspective from Chapter 4 to DA-aware BSP could remove the dependency on parametric models, allowing for more flexible and real-time nonparametric DA BSP.

Another promising direction is to simplify the number of landmarks considered in planning to manage hypothesis growth more effectively. This approach could mitigate the exponential increase in hypotheses and improve computational efficiency. Additionally, exploring how information from planning can be utilized in inference, especially when faced with strict computational budget constraints, presents a valuable research opportunity. By leveraging bounds derived from different simplified beliefs, it may be possible to optimize both planning and inference processes, leading to more robust and adaptable autonomous systems.

Appendix A

Proofs for Chapter 3

A.1 Theorem 1

Proof .

Following a similar proof to theorem 3 in Khosoussi et al., 2019 we get

$$\tau_w^{U,A} \leq \log|\Lambda^U| \leq \tau_w^{U,A} + n^U \cdot \log(1 + \delta^U/\lambda_1^U), \quad (\text{A.1})$$

where $\delta^U = \|(\Delta_{w_p}^U)^T \Delta_{w_p}^U\|_\infty$ and $\lambda_1^U = \lambda_{\min}(L_{w_\theta}^U)$. It is easy to see that

$$\lim_{\delta^U/\lambda_1^U \rightarrow 0^+} \log|\Lambda^U| = \tau_w^{U,A}. \quad (\text{A.2})$$

■

A.2 Theorem 2

Proof .

Using inequalities (3.21), (A.1) and eq. (3.25) we get

$$\mathcal{UB}[J_{\mathcal{H}}^F] = \frac{n^F}{2} \log(2\pi e) - \frac{1}{2} \tau_w + \frac{1}{2} \left[\tau_w^{U,A} + n^U \cdot \log \left(1 + \frac{\delta^U}{\lambda_1^U} \right) \right] \quad (\text{A.3})$$

$$= S_{W_{TCD}} + \frac{n^U}{2} \log \left(1 + \frac{\delta^U}{\lambda_1^U} \right). \quad (\text{A.4})$$

Similarly, the lower bound is given by

$$\begin{aligned} \mathcal{LB}[J_{\mathcal{H}}^F] &= \frac{n^F}{2} \log(2\pi e) - \frac{1}{2} \left[\tau_w + n \cdot \log \left(1 + \frac{\delta}{\lambda_1} \right) \right] + \frac{1}{2} \tau_w^{U,A} \\ &= S_{W_{TCD}} - \frac{n}{2} \log \left(1 + \frac{\delta}{\lambda_1} \right). \end{aligned} \quad (\text{A.5})$$

■

A.3 Lemma 1

Proof .

Since the FIM is Hermitian and since Λ^U is a principal submatrix of Λ , according to Cauchy's interlacing theorem (Horn and Johnson, 2012), the eigenvalues must satisfy

$$\lambda_1 \leq \lambda_1^U \leq \lambda_2 \leq \dots \leq \lambda_n \leq \lambda_n^U \leq \lambda_{n+1}. \quad (\text{A.6})$$

In addition, as n^U is a finite number, i.e. represents the size of a finite graph, and as $w^{U,A}$ is defined over finite non negative weights, $\exists \alpha < \infty$ such that $\delta^U < \alpha\delta$. Combined with (A.6), we get

$$\delta/\lambda_1 \rightarrow 0^+ \Rightarrow \alpha\delta/\lambda_1 \rightarrow 0^+ \Rightarrow \delta^U/\lambda_1^U \rightarrow 0^+. \quad (\text{A.7})$$

■

A.4 Theorem 3

Proof .

We follow Kitanov and Indelman (2024) and derive the Hadamard bounds for the *focused* case. Using the proof to Theorem 3 in Khosoussi et al. (2019), we know that

$$\log|\Lambda| \leq 2\tau_{w_p} + \log|L_{w_\theta} + \delta I|. \quad (\text{A.8})$$

We denote by W_θ the diagonal matrix defined in (3.12) based on weights w_θ defined in (3.18). Since $L_{w_\theta} + \delta I$ is a positive-definite matrix, applying Hadamard inequality we get

$$\log|L_{w_\theta} + \delta I| \leq \sum_{i=1}^n \log [W_\theta(i, i) + \delta]. \quad (\text{A.9})$$

Replacing (A.9) into (A.8) we get

$$\log|\Lambda| \leq 2\tau_{w_p} + \sum_{i=1}^n \log [W_\theta(i, i) + \delta]. \quad (\text{A.10})$$

Similarly, for $\log|\Lambda^U|$ we get

$$\log|\Lambda^U| \leq 2\tau_{w_p^{U,A}} + \sum_{i=1}^n \log [W_\theta^{U,A}(i, i) + \delta^{U,A}]. \quad (\text{A.11})$$

Replacing (A.10) and (A.11) into the definition of $S_{WTC D}$ we get new bounds for $\epsilon \left(J_{\mathcal{H}}^F \right)$.
■

Appendix B

Appendix for Chapter 4

B.1 Proofs

B.1.1 Corollary 1

Proof .

We follow a similar derivation to the one presented in Pathak et al., 2018 and factorize $w_{k+1}^{i,j}$ by first marginalizing over x_{k+1} and then by applying the Bayes rule

$$w_{k+1}^{i,j} = \int_{x_{k+1}} \frac{\mathbb{P}(Z_{k+1} | \beta_{k+1}^i, \beta_{1:k}^j, x_{k+1}, H_{k+1}^-) \mathbb{P}(\beta_{k+1}^i, \beta_{1:k}^j, x_{k+1} | H_{k+1}^-)}{\mathbb{P}(Z_{k+1} | H_{k+1}^-)}.$$

Using the chain rule multiple times over the second term in the numerator completes the proof. ■

B.1.2 Theorem 4

Proof .

We split (4.14) based on belief components from M_k^s and use (4.11) to rewrite \mathcal{H} as

$$\mathcal{H} = - \sum_i^{|L|} \sum_j^{M_k^s} \frac{\tilde{\zeta}_{k+1}^{i,j} w_k^{s,j} w_k^{m,s}}{\eta} \log \left(\frac{\tilde{\zeta}_{k+1}^{i,j} w_k^{s,j} w_k^{m,s}}{\eta} \right) - \sum_i^{|L|} \sum_j^{\neg M_k^s} \frac{\tilde{\zeta}_{k+1}^{i,j} w_k^j}{\eta} \log \left(\frac{\tilde{\zeta}_{k+1}^{i,j} w_k^j}{\eta} \right). \quad (\text{B.1})$$

Using the key observation that $\tilde{\zeta}_{k+1}^{s,ij} = \tilde{\zeta}_{k+1}^{i,j}$, basic log properties and that by definition all posterior weights sum to 1, we write the cost for a simplified belief as

$$\mathcal{H}^s = - \frac{1}{\eta^s} \sum_i^{|L|} \sum_j^{M_k^s} \left[\tilde{\zeta}_{k+1}^{i,j} w_k^{s,j} \log \left(\tilde{\zeta}_{k+1}^{i,j} w_k^{s,j} \right) \right] + \log(\eta^s). \quad (\text{B.2})$$

Replacing (B.2) back into (B.1) and using basic log properties completes the proof. ■

B.1.3 Theorem 5

Proof .

The last term in (4.15) is non negative as all posterior weights are at most 1 by definition. Thus, removing this term and using Theorem 7 we immediately get the lower bound. For the upper bound, we revisit the last term in (4.15). We first define γ using (4.11) and (B.6)

$$\gamma \triangleq \sum_i^{|L|} \sum_j^{\neg M_k^s} \frac{\tilde{\zeta}_{k+1}^{i,j} w_k^j}{\eta} = 1 - \sum_i^{|L|} \sum_j^{M_k^s} \frac{\tilde{\zeta}_{k+1}^{i,j} w_k^j}{\eta} = 1 - \frac{\eta^s w_k^{m,s}}{\eta}.$$

Using the log sum inequality (Cover and Thomas, 1991)

$$\sum_i^n a_i \cdot \log \left(\frac{a_i}{b_i} \right) \geq a \cdot \log \left(\frac{a}{b} \right) \text{ where } \sum_i^n a_i = a, \sum_i^n b_i = b,$$

with $a_i = \frac{\tilde{\zeta}_{k+1}^{i,j} w_k^j}{\eta}$, $\sum_i^{|L|} \sum_j^{\neg M_k^s} \frac{\tilde{\zeta}_{k+1}^{i,j} w_k^j}{\eta} = \gamma$ and $b_i = 1$, we bound the last term in (4.15)

$$\sum_i^{|L|} \sum_j^{\neg M_k^s} \frac{\tilde{\zeta}_{k+1}^{i,j} w_k^j}{\eta} \log \left(\frac{\tilde{\zeta}_{k+1}^{i,j} w_k^j}{\eta} \right) \geq \gamma \log \left(\frac{\gamma}{|L| |\neg M_k^s|} \right). \quad (\text{B.3})$$

Substituting (B.3) into (4.15) and using Theorem 7 completes the proof. \blacksquare

B.1.4 Corollary 2

Proof .

Given $M_k^s = M_k$ it holds by definition that $w_k^{m,s} = 1$ and $\mathcal{H} = \mathcal{H}^s$ as $b_{k+1} = b_{k+1}^s$. Substituting these back into (B.10) and using (4.11) and Corollary 3, the lower bound becomes

$$\mathcal{LB}[\mathcal{H}] = \mathcal{H} - \log(\eta) - \sum_i^{|L|} \sum_j^{M_k} \frac{\tilde{\zeta}_{k+1}^{i,j} w_k^{s,j}}{\eta} \log \left(\frac{1}{\eta} \right) = \mathcal{H}. \quad (\text{B.4})$$

Given $M_k^s = M_k$ it is also straightforward by Corollary 3 that $\gamma = 0$. As such, using exactly the same derivations as for the lower bound, it immediately holds that $\mathcal{H} = \mathcal{UB}[\mathcal{H}]$. This completes the proof. \blacksquare

B.1.5 Theorem 6

Proof .

We write η explicitly and first marginalize over all DA realizations and states at time $k + 1$. We then marginalize over all hypotheses from time k and apply the chain rule

multiple times

$$\begin{aligned}
\eta &\triangleq \mathbb{P}\left(Z_{k+1}|H_{k+1}^-\right) \\
&= \sum_i^{|L|} \sum_j^{M_k} \int_{x_{k+1}} \mathbb{P}\left(Z_{k+1}|x_{k+1}\right) \mathbb{P}\left(\beta_{k+1}^i|x_{k+1}\right) \cdot \mathbb{P}\left(x_{k+1}|\beta_{1:k}^j, H_{k+1}^-\right) \mathbb{P}\left(\beta_{1:k}^j|H_{k+1}^-\right) \\
&= \sum_i^{|L|} \sum_j^{M_k} \tilde{\zeta}_{k+1}^{i,j} w_k^j.
\end{aligned} \tag{B.5}$$

Using similar derivations and the key observation that $\tilde{\zeta}_{k+1}^{s,ij} = \tilde{\zeta}_{k+1}^{i,j}$, we also write η^s as

$$\eta^s = \sum_i^{|L|} \sum_j^{M_k^s} \tilde{\zeta}_{k+1}^{i,j} w_k^{s,j}. \tag{B.6}$$

Splitting (B.5) based on belief components from M_k^s and using (4.11) and (B.6) completes the proof. \blacksquare

B.1.6 Theorem 7

Proof .

As all weights are positive by definition, removing the last term in (4.19) we immediately get the lower bound. For the upper bound, we rewrite the second term in (4.19) using $\tilde{\zeta}_{k+1}^{i,j}$

$$\begin{aligned}
&\sum_i^{|L|} \sum_j^{\neg M_k^s} \tilde{\zeta}_{k+1}^{i,j} w_k^j = \\
&\sum_j^{\neg M_k^s} w_k^j \sum_i^{|L|} \int_{x_{k+1}} \mathbb{P}\left(Z_{k+1}|\beta_{k+1}^i, x_{k+1}\right) \cdot \mathbb{P}\left(\beta_{k+1}^i|x_{k+1}\right) \mathbb{P}\left(x_{k+1}|H_{k+1}^-, \beta_{1:k}^j\right).
\end{aligned} \tag{B.7}$$

The joint measurement likelihood term, given in (4.9), is a product of probability distribution functions, all given by (4.1) and can thus be bounded using an a priori known maximum value σ . The term $\mathbb{P}\left(\beta_{k+1}^i|x_{k+1}\right)$ represents the probability for the i th DA realization given x_{k+1} , i.e. the probability of observing a specific set of landmarks. As we assume the map to be given, it can be bounded using some constant α^i , e.g. in the case of a camera, it can be an indicator function for landmarks that are within the field of view. Finally, for every hypothesis j it holds that $\int_{x_{k+1}} \mathbb{P}\left(x_{k+1}|H_{k+1}^-, \beta_{1:k}^j\right) = 1$. Substituting these and $w_k^{m,s}$ back into (4.19) completes the proof. \blacksquare

B.1.7 Corollary 3

Proof .

Given $M_k^s = M_k$ it holds by definition that $w_k^{m,s} = 1$ and $\eta = \eta^s$ as $b_{k+1} = b_{k+1}^s$. Replacing these back into (4.42), (4.21) immediately completes the proof. \blacksquare

B.1.8 Theorem 8

Proof .

Given M_{k+n} belief components, we split the cost function (4.36) based on components in and outside M_{k+n}^s

$$\mathcal{H}_{k+n} = - \sum_{r \in M_{k+n}^s} \frac{w_{k+n}^r}{\eta_{k+n}} \log \left(\frac{w_{k+n}^r}{\eta_{k+n}} \right) - \sum_{r \in \neg M_{k+n}^s} \frac{w_{k+n}^r}{\eta_{k+n}} \log \left(\frac{w_{k+n}^r}{\eta_{k+n}} \right). \quad (\text{B.8})$$

Using basic log properties and $w_{k+n}^{s,r} \triangleq \frac{w_{k+n}^r}{w_{k+n}^{m,s}}$ where $w_{k+n}^{m,s} \triangleq \sum_{m \in M_{k+n}^s} w_{k+n}^m$ completes the proof

$$\begin{aligned} & \mathcal{H}_{k+n} \\ &= - \sum_{r \in M_{k+n}^s} \frac{w_{k+n}^{m,s} w_{k+n}^{s,r}}{\eta_{k+n}} \log \left(w_{k+n}^{m,s} w_{k+n}^{s,r} \right) + \sum_{r \in M_{k+n}^s} \frac{w_{k+n}^{m,s} w_{k+n}^{s,r}}{\eta_{k+n}} \log (\eta_{k+n}) - \\ & \sum_{r \in \neg M_{k+n}^s} \frac{w_{k+n}^r}{\eta_{k+n}} \log \left(\frac{w_{k+n}^r}{\eta_{k+n}} \right) \\ &= \frac{w_{k+n}^{m,s}}{\eta_{k+n}} \left[- \sum_{r \in M_{k+n}^s} w_{k+n}^{s,r} \log \left(w_{k+n}^{m,s} w_{k+n}^{s,r} \right) + \sum_{r \in M_{k+n}^s} w_{k+n}^{s,r} \log (\eta_{k+n}) \right] - \\ & \sum_{r \in \neg M_{k+n}^s} \frac{w_{k+n}^r}{\eta_{k+n}} \log \left(\frac{w_{k+n}^r}{\eta_{k+n}} \right) \quad (\text{B.9}) \\ &= \frac{w_{k+n}^{m,s}}{\eta_{k+n}} \left[- \sum_{r \in M_{k+n}^s} w_{k+n}^{s,r} \log \left(w_{k+n}^{s,r} \right) - \log \left(w_{k+n}^{m,s} \right) + \log (\eta_{k+n}) \right] - \\ & \sum_{r \in \neg M_{k+n}^s} \frac{w_{k+n}^r}{\eta_{k+n}} \log \left(\frac{w_{k+n}^r}{\eta_{k+n}} \right) \\ &= \frac{w_{k+n}^{m,s}}{\eta_{k+n}} \left[\mathcal{H}_{k+n}^s + \log \left(\frac{\eta_{k+n}}{w_{k+n}^{m,s}} \right) \right] - \sum_{r \in \neg M_{k+n}^s} \frac{w_{k+n}^r}{\eta_{k+n}} \log \left(\frac{w_{k+n}^r}{\eta_{k+n}} \right). \end{aligned}$$

\blacksquare

B.1.9 Theorem 9

Proof .

The last term in (B.9) is non negative as all posterior weights are at most 1 by definition. Thus, removing this term and using the bounds over η_{k+n} in Theorem 3 we immediately

get the lower bound

$$\mathcal{LB}[\mathcal{H}_{k+n}] = \frac{w_{k+n}^{m,s}}{\mathcal{UB}[\eta_{k+n}]} \left[\mathcal{H}_{k+n}^s + \log \left(\frac{\mathcal{LB}[\eta_{k+n}]}{w_{k+n}^{m,s}} \right) \right]. \quad (\text{B.10})$$

For the upper bound, we first define

$$\gamma \triangleq \sum_{r \in \neg M_{k+n}^s} \frac{w_{k+n}^r}{\eta_{k+n}} = 1 - \sum_{r \in M_{k+n}^s} \frac{w_{k+n}^r}{\eta_{k+n}}. \quad (\text{B.11})$$

Using the log sum inequality (Cover and Thomas, 1991)

$$\sum_i^n a_i \cdot \log \left(\frac{a_i}{b_i} \right) \geq a \cdot \log \left(\frac{a}{b} \right) \text{ where } \sum_i^n a_i = a, \sum_i^n b_i = b \quad (\text{B.12})$$

with $a_i = \frac{w_{k+n}^r}{\eta_{k+n}}$ and $b_i = 1$, we bound the last term in (B.9)

$$\sum_{r \in \neg M_{k+n}^s} \frac{w_{k+n}^r}{\eta_{k+n}} \log \left(\frac{w_{k+n}^r}{\eta_{k+n}} \right) \geq \gamma \log \left(\frac{\gamma}{|\neg M_{k+n}^s|} \right). \quad (\text{B.13})$$

Substituting (B.13) into (B.9); using the bounds over η_{k+n} from Theorem 3; and since by definition $0 \leq \gamma \leq 1$, we get the upper bound

$$\mathcal{UB}[\mathcal{H}_{k+n}] = \frac{w_{k+n}^{m,s}}{\mathcal{LB}[\eta_{k+n}]} \left[\mathcal{H}_{k+n}^s + \log \left(\frac{\mathcal{UB}[\eta_{k+n}]}{w_{k+n}^{m,s}} \right) \right] - \bar{\gamma} \log \left(\frac{\bar{\gamma}}{|\neg M_{k+n}^s|} \right), \quad (\text{B.14})$$

where $\bar{\gamma} = 1 - \sum_{r \in M_{k+n}^s} \frac{w_{k+n}^r}{\mathcal{UB}[\eta_{k+n}]}$ and $|\neg M_{k+n}^s| > 2$. ■

B.1.10 Corollary 4

Proof .

Given that $M_{k+n}^s = M_{k+n}$ it holds by definition that $\eta_{k+n} = w_{k+n}^{m,s}$ and $\mathcal{H}_{k+n} = \mathcal{H}_{k+n}^s$. Substituting back into (B.10) and using Corollary 2 we get

$$\lim_{M_{k+n}^s \rightarrow M_{k+n}} \mathcal{LB}[\mathcal{H}_{k+n}] = \frac{w_{k+n}^{m,s}}{\eta_{k+n}} \left[\mathcal{H}_{k+n} + \log \left(\frac{\eta_{k+n}}{w_{k+n}^{m,s}} \right) \right] = \mathcal{H}_{k+n}. \quad (\text{B.15})$$

It is also straightforward that $M_{k+n}^s = M_{k+n} \Rightarrow \bar{\gamma} = 0$. As such, similarly to the lower bound, it immediately holds that $\mathcal{H}_{k+n} = \lim_{M_{k+n}^s \rightarrow M_{k+n}} \mathcal{UB}[\mathcal{H}_{k+n}]$. ■

B.1.11 Theorem 10

Proof .

By definition $\eta_{k+n} = \sum_{r \in M_{k+n}} w_{k+n}^r$. Splitting this sum, we rewrite η_{k+n} as

$$\eta_{k+n} = \sum_{r \in M_{k+n}^s} w_{k+n}^r + \sum_{r \in \neg M_{k+n}^s} w_{k+n}^r. \quad (\text{B.16})$$

The second term in (B.16) is positive by definition. As such, removing it we immediately get the lower bound

$$\mathcal{LB}[\eta_{k+n}] = \sum_{r \in M_{k+n}^s} w_{k+n}^r. \quad (\text{B.17})$$

For the upper bound, we first rewrite the second term in (B.16) as

$$\sum_{r \in \neg M_{k+n}^s} w_{k+n}^r = \sum_{r \in \neg M_{k+n}^s} \prod_{i=0}^n w_{k+i}^r = \sum_{r \in \neg M_{k+n}^s} w_k^r \prod_{i=1}^n w_{k+i}^r. \quad (\text{B.18})$$

Each w_{k+i}^r is defined as

$$\int_{x_{k+i}} \mathbb{P}(Z_{k+i} | \beta_{k+i}^r, x_{k+i}) \mathbb{P}(\beta_{k+i}^r | x_{k+i}) \mathbb{P}(x_{k+i} | H_{k+i}^-, \beta_{1:k+i}^r), \quad (\text{B.19})$$

which can also be bounded as presented in Theorem 7. The joint measurement likelihood term is a product of pdfs, all given a priori, and can be bounded using a known maximum value σ^i . The term $\mathbb{P}(\beta_{k+i}^r | x_{k+i})$ represents the probability for the r th DA realization given x_{k+i} and can be bounded by 1 representing, for example, an indicator function for landmarks that are within the field of view. Finally, for every hypothesis r it holds that $\int_{x_{k+i}} \mathbb{P}(x_{k+i} | H_{k+i}^-, \beta_{1:k+i}^r) = 1$. By definition, $\sum_{r \in M_k} w_k^r = 1$ and each component at time k generates $\frac{|M_{k+n}|}{|M_k|}$ at time $k+n$, thus

$$\sum_{r \in M_{k+n}} w_k^r = \sum_{r \in M_{k+n}^s} w_k^r + \sum_{r \in \neg M_{k+n}^s} w_k^r = \frac{|M_{k+n}|}{|M_k|}. \quad (\text{B.20})$$

As such, we can bound (B.18) as

$$\sum_{r \in \neg M_{k+n}^s} w_{k+n}^r \leq \sum_{r \in \neg M_{k+n}^s} w_k^r \prod_{i=1}^n \sigma^i = \left(\frac{|M_{k+n}|}{|M_k|} - \sum_{r \in M_{k+n}^s} w_k^r \right) \prod_{i=1}^n \sigma^i. \quad (\text{B.21})$$

Substituting back into (B.16) we get the upper bound

$$\mathcal{UB}[\eta_{k+n}] = \sum_{r \in M_{k+n}^s} w_{k+n}^r + \left(\frac{|M_{k+n}|}{|M_k|} - \sum_{r \in M_{k+n}^s} w_k^r \right) \prod_{i=1}^n \sigma^i. \quad (\text{B.22})$$

■

B.1.12 Corollary 5

Proof .

Given that $M_{k+n}^s = M_{k+n}$ it holds by definition that $\eta_{k+n} = w_{k+n}^{m,s}$. Substituting back into (B.17) we get

$$\lim_{M_{k+n}^s \rightarrow M_{k+n}} \mathcal{LB} [\eta_{k+n}] = \sum_{r \in M_{k+n}^s} w_{k+n}^r = w_{k+n}^{m,s} = \eta_{k+n}. \quad (\text{B.23})$$

It is also straightforward that $M_{k+n}^s = M_{k+n} \Rightarrow \frac{|M_{k+n}|}{|M_k|} - \sum_{r \in M_{k+n}^s} w_k^r = 0$ As such, similarly to the lower bound, it immediately holds that $\eta_{k+n} = \lim_{M_{k+n}^s \rightarrow M_{k+n}} \mathcal{UB} [\eta_{k+n}]$.

■

B.2 Incremental Bounds Updates: The Myopic Case

In this section, we rigorously derive the method for incrementally updating the bounds presented in Section 4.2 for the myopic case.

B.2.1 The Normalization Term Bounds $\mathcal{LB} [\eta], \mathcal{UB} [\eta]$

We denote the bounds presented in Theorem 4.42 as $\mathcal{LB} [\eta | b_k^s], \mathcal{UB} [\eta | b_k^s]$, i.e. with respect to a simplified belief b_k^s with M_k^s components. Given a belief component $r_k \notin M_k^s$ with associated weight w_k^r , we denote $M_k^{s+1} \triangleq M_k^s \cup r_k$. Using (4.11) the simplified belief at time k for M_k^{s+1} components is given by

$$b_k^{s+1} \triangleq \sum_{j=1}^{M_k^{s+1}} w_k^{s+1,j} b_k^j, \quad w_k^{s+1,j} \triangleq \frac{w_k^j}{w_k^{m,s+1}}, \quad (\text{B.24})$$

where w_k^j corresponds to the original belief component weight (4.3) and $w_k^{m,s+1} = w_k^{m,s} + w_k^r$. As such, $\mathcal{LB} [\eta | b_k^{s+1}], \mathcal{UB} [\eta | b_k^{s+1}]$ represent the bounds for the measurement likelihood η given a simplified belief b_k^{s+1} with M_k^{s+1} components. Using (B.6) and (B.24) we define

$$\eta^{s+1} \triangleq \sum_i^{|L|} \sum_j^{M_k^{s+1}} \tilde{\zeta}_{k+1}^{i,j} w_k^{s+1,j}. \quad (\text{B.25})$$

We now present how to incrementally adapt the lower and upper bounds. We begin by writing the lower bound with respect to the simplified belief b_k^{s+1} using (B.25) and get

the recursive update rule

$$\begin{aligned}
\mathcal{LB} \left[\eta | b_k^{s+1} \right] &= \eta^{s+1} w_k^{m,s+1} \\
&= \sum_i^{|L|} \sum_j^{M_k^{s+1}} \tilde{\zeta}_{k+1}^{i,j} w_k^j \\
&= \sum_i^{|L|} \sum_j^{M_k^s} \tilde{\zeta}_{k+1}^{i,j} w_k^j + \sum_i^{|L|} \tilde{\zeta}_{k+1}^{i,r} w_k^r \\
&= \eta^s w_k^{m,s} + \sum_i^{|L|} \tilde{\zeta}_{k+1}^{i,r} w_k^r \\
&= \mathcal{LB} [\eta | b_k^s] + \sum_i^{|L|} \tilde{\zeta}_{k+1}^{i,r} w_k^r.
\end{aligned} \tag{B.26}$$

Using similar derivations the recursive update rule for the upper bound is given by

$$\begin{aligned}
\mathcal{UB} \left[\eta | b_k^{s+1} \right] &= \eta^{s+1} w_k^{m,s+1} + (1 - w_k^{m,s+1}) \sigma \sum_i^{|L|} \alpha^i \\
&= \eta^s w_k^{m,s} + \sum_i^{|L|} \tilde{\zeta}_{k+1}^{i,r} w_k^r + (1 - w_k^{m,s} - w_k^r) \sigma \sum_i^{|L|} \alpha^i \\
&= \eta^s w_k^{m,s} + (1 - w_k^{m,s}) \sigma \sum_i^{|L|} \alpha^i + \sum_i^{|L|} \tilde{\zeta}_{k+1}^{i,r} w_k^r - w_k^r \sigma \sum_i^{|L|} \alpha^i \\
&= \mathcal{UB} [\eta | b_k^s] + w_k^r \sum_i^{|L|} \left[\tilde{\zeta}_{k+1}^{i,r} - \sigma \alpha^i \right].
\end{aligned} \tag{B.27}$$

B.2.2 The Cost Function Bounds $\mathcal{LB} [\mathcal{H}]$, $\mathcal{UB} [\mathcal{H}]$

We follow similar derivations as in Section B.2.1 and denote the bounds presented in Theorem 5 as $\mathcal{LB} [\mathcal{H} | b_k^s]$, $\mathcal{UB} [\mathcal{H} | b_k^s]$, i.e. with respect to a simplified belief b_k^s with M_k^s components. Given a belief component $r_k \notin M_k^s$ with associated weight w_k^r , we denote $M_k^{s+1} \triangleq M_k^s \cup r_k$. Using (B.24) we also denote the bounds over the cost term, given a simplified belief b_k^{s+1} with M_k^{s+1} components, as $\mathcal{LB} [\mathcal{H} | b_k^{s+1}]$, $\mathcal{UB} [\mathcal{H} | b_k^{s+1}]$. Deriving a direct recursive update rule for these bounds is not trivial. Instead, we show how each term in $\mathcal{LB} [\mathcal{H} | b_k^{s+1}]$, $\mathcal{UB} [\mathcal{H} | b_k^{s+1}]$ can be incrementally updated individually. Using (B.25) we begin with a recursive update rule for η^{s+1} given by

$$\eta^{s+1} = \frac{1}{w_k^{m,s+1}} \sum_i^{|L|} \sum_j^{M_k^{s+1}} \tilde{\zeta}_{k+1}^{i,j} w_k^{s+1,j} = \frac{1}{w_k^{m,s+1}} \left[\eta^s w_k^{m,s} + \sum_i^{|L|} \tilde{\zeta}_{k+1}^{i,r} w_k^r \right]. \tag{B.28}$$

Using (B.2) and (B.6) we write the recursive update rule for \mathcal{H}^{s+1} , i.e. the cost given a simplified belief b_k^{s+1}

$$\begin{aligned}
& \mathcal{H}^{s+1} \\
&= -\frac{1}{\eta^{s+1}} \sum_i^{|L|} \sum_j^{M_k^{s+1}} \left[\tilde{\zeta}_{k+1}^{i,j} w_k^{s+1,j} \log \left(\tilde{\zeta}_{k+1}^{i,j} w_k^{s+1,j} \right) \right] + \log \left(\eta^{s+1} \right) \\
&= -\frac{1}{\eta^{s+1}} \left[\sum_i^{|L|} \sum_j^{M_k^s} \left[\frac{\tilde{\zeta}_{k+1}^{i,j} w_k^j}{w_k^{m,s+1}} \log \left(\frac{\tilde{\zeta}_{k+1}^{i,j} w_k^j}{w_k^{m,s+1}} \right) \right] + \sum_i^{|L|} \left[\frac{\tilde{\zeta}_{k+1}^{i,r} w_k^r}{w_k^{m,s+1}} \log \left(\frac{\tilde{\zeta}_{k+1}^{i,r} w_k^r}{w_k^{m,s+1}} \right) \right] \right] + \\
&\log \left(\eta^{s+1} \right) \\
&= -\frac{1}{\eta^{s+1}} \left[\frac{w_k^{m,s}}{w_k^{m,s+1}} \sum_i^{|L|} \sum_j^{M_k^s} \left[\tilde{\zeta}_{k+1}^{i,j} w_k^{s,j} \log \left(\frac{\tilde{\zeta}_{k+1}^{i,j} w_k^{s,j} w_k^{m,s}}{w_k^{m,s+1}} \right) \right] + \right. \\
&\left. \sum_i^{|L|} \left[\frac{\tilde{\zeta}_{k+1}^{i,r} w_k^r}{w_k^{m,s+1}} \log \left(\frac{\tilde{\zeta}_{k+1}^{i,r} w_k^r}{w_k^{m,s+1}} \right) \right] \right] + \log \left(\eta^{s+1} \right) \\
&= -\frac{1}{\eta^{s+1}} \left[\frac{w_k^{m,s}}{w_k^{m,s+1}} \left[\sum_i^{|L|} \sum_j^{M_k^s} \left[\tilde{\zeta}_{k+1}^{i,j} w_k^{s,j} \log \left(\tilde{\zeta}_{k+1}^{i,j} w_k^{s,j} \right) \right] + \right. \right. \\
&\left. \left. \sum_i^{|L|} \sum_j^{M_k^s} \left[\tilde{\zeta}_{k+1}^{i,j} w_k^{s,j} \log \left(\frac{w_k^{m,s}}{w_k^{m,s+1}} \right) \right] \right] + \sum_i^{|L|} \left[\frac{\tilde{\zeta}_{k+1}^{i,r} w_k^r}{w_k^{m,s+1}} \log \left(\frac{\tilde{\zeta}_{k+1}^{i,r} w_k^r}{w_k^{m,s+1}} \right) \right] \right] + \log \left(\eta^{s+1} \right) \\
&= -\frac{1}{\eta^{s+1}} \left[\frac{w_k^{m,s}}{w_k^{m,s+1}} \left[-\eta^s [\mathcal{H}^s - \log(\eta^s)] + \eta^s \log \left(\frac{w_k^{m,s}}{w_k^{m,s+1}} \right) \right] + \right. \\
&\left. \sum_i^{|L|} \left[\frac{\tilde{\zeta}_{k+1}^{i,r} w_k^r}{w_k^{m,s+1}} \log \left(\frac{\tilde{\zeta}_{k+1}^{i,r} w_k^r}{w_k^{m,s+1}} \right) \right] \right] + \log \left(\eta^{s+1} \right) \\
&= \frac{\eta^s}{\eta^{s+1}} \frac{w_k^{m,s}}{w_k^{m,s+1}} \left[\mathcal{H}^s - \log(\eta^s) - \log \left(\frac{w_k^{m,s}}{w_k^{m,s+1}} \right) \right] - \frac{1}{\eta^{s+1}} \sum_i^{|L|} \left[\frac{\tilde{\zeta}_{k+1}^{i,r} w_k^r}{w_k^{m,s+1}} \log \left(\frac{\tilde{\zeta}_{k+1}^{i,r} w_k^r}{w_k^{m,s+1}} \right) \right] + \\
&\log \left(\eta^{s+1} \right).
\end{aligned}$$

(B.29)

Using Theorem 5 we explicitly write the lower bound with respect to the simplified belief b_k^{s+1}

$$\begin{aligned}
& \mathcal{LB} \left[\mathcal{H} | b_k^{s+1} \right] \\
&= \frac{\eta^{s+1} w_k^{m,s+1}}{\mathcal{UB} \left[\eta | b_k^{s+1} \right]} \left[\mathcal{H}^{s+1} - \log(\eta^{s+1}) \right] - \\
& \frac{w_k^{m,s+1}}{\mathcal{UB} \left[\eta | b_k^{s+1} \right]} \sum_i^{|L|} \sum_j^{M_k^{s+1}} \tilde{\zeta}_{k+1}^{i,j} w_k^{s+1,j} \log \left(\frac{w_k^{m,s+1}}{\mathcal{LB} \left[\eta | b_k^{s+1} \right]} \right) \\
&= \frac{\eta^{s+1} w_k^{m,s+1}}{\mathcal{UB} \left[\eta | b_k^{s+1} \right]} \left[\mathcal{H}^{s+1} - \log(\eta^{s+1}) \right] - \frac{w_k^{m,s+1} \eta^{s+1}}{\mathcal{UB} \left[\eta | b_k^{s+1} \right]} \log \left(\frac{w_k^{m,s+1}}{\mathcal{LB} \left[\eta | b_k^{s+1} \right]} \right),
\end{aligned} \tag{B.30}$$

and observe that each term can be incrementally updated individually using (B.28), (B.29) and Section B.2.1. Similarly, using Theorem 5, we explicitly write the upper bound with respect to the simplified belief b_k^{s+1}

$$\begin{aligned}
& \mathcal{UB} \left[\mathcal{H} | b_k^{s+1} \right] \\
&= \frac{\eta^{s+1} w_k^{m,s+1}}{\mathcal{LB} \left[\eta | b_k^{s+1} \right]} \left[\mathcal{H}^{s+1} - \log(\eta^{s+1}) \right] - \\
& \frac{w_k^{m,s+1}}{\mathcal{LB} \left[\eta | b_k^{s+1} \right]} \sum_i^{|L|} \sum_j^{M_k^{s+1}} \tilde{\zeta}_{k+1}^{i,j} w_k^{s+1,j} \log \left(\frac{w_k^{m,s+1}}{\mathcal{UB} \left[\eta | b_k^{s+1} \right]} \right) - \gamma \log \left(\frac{\gamma}{|L| \left| \neg M_k^{s+1} \right|} \right) \\
&= \frac{\eta^{s+1} w_k^{m,s+1}}{\mathcal{LB} \left[\eta | b_k^{s+1} \right]} \left[\mathcal{H}^{s+1} - \log(\eta^{s+1}) \right] - \\
& \frac{w_k^{m,s+1} \eta^{s+1}}{\mathcal{LB} \left[\eta | b_k^{s+1} \right]} \log \left(\frac{w_k^{m,s+1}}{\mathcal{UB} \left[\eta | b_k^{s+1} \right]} \right) - \gamma \log \left(\frac{\gamma}{|L| \left| \neg M_k^{s+1} \right|} \right),
\end{aligned} \tag{B.31}$$

where $\gamma \triangleq 1 - \frac{\eta^{s+1} w_k^{m,s}}{\mathcal{UB} \left[\eta | b_k^{s+1} \right]}$. Since $0 \leq \gamma \leq 1$ by definition, the upper bound (B.31) holds when $|L| \left| \neg M_k^{s+1} \right| > 2$. We observe that each term can be incrementally updated individually using (B.28), (B.29) and Section B.2.1.

B.3 Incremental Bounds Updates: The Nonmyopic Case

In this section, we rigorously derive the method for incrementally updating the bounds presented in Section 4.3.5 for the nonmyopic case.

Given a simplified belief b_{k+n}^s with M_{k+n}^s components and a belief component $p_{k+n} \notin M_{k+n}^s$ with associated weight w_{k+n}^p , we denote $M_{k+n}^{s+1} \triangleq M_{k+n}^s \cup p_{k+n}$. We also further denote the bounds in Theorem 9 as $\mathcal{LB} \left[\mathcal{H}_{k+n} | M_{k+n}^s \right], \mathcal{UB} \left[\mathcal{H}_{k+n} | M_{k+n}^s \right]$, and the bounds in Theorem 10 as $\mathcal{LB} \left[\eta_{k+n} | M_{k+n}^s \right], \mathcal{UB} \left[\eta_{k+n} | M_{k+n}^s \right]$, i.e. with respect

to M_{k+n}^s components of the simplified belief b_{k+n}^s .

B.3.1 The Normalization Term Bounds $\mathcal{LB}[\eta_{k+n}]$, $\mathcal{LB}[\eta_{k+n}]$

Given a belief component $p_{k+n} \notin M_{k+n}^s$ with associated weight w_{k+n}^p , we first derive a recursive update rule for the lower bound

$$\mathcal{LB}[\eta_{k+n}|M_{k+n}^{s+1}] = \sum_{r \in M_{k+n}^{s+1}} w_{k+n}^r = w_{k+n}^p + \sum_{r \in M_{k+n}^s} w_{k+n}^r = w_{k+n}^p + \mathcal{LB}[\eta_{k+n}|M_{k+n}^s]. \quad (\text{B.32})$$

The recursive update rule for the upper bound is given by

$$\begin{aligned} \mathcal{UB}[\eta_{k+n}|M_{k+n}^{s+1}] &= \sum_{r \in M_{k+n}^{s+1}} w_{k+n}^r + \left(\frac{|M_{k+n}|}{|M_k|} - \sum_{r \in M_{k+n}^{s+1}} w_k^r \right) \prod_{i=1}^n \sigma^i \\ &= w_{k+n}^p + \sum_{r \in M_{k+n}^s} w_{k+n}^r + \left(\frac{|M_{k+n}|}{|M_k|} - w_k^p - \sum_{r \in M_{k+n}^s} w_k^r \right) \prod_{i=1}^n \sigma^i \\ &= w_{k+n}^p - w_k^p \prod_{i=1}^n \sigma^i + \mathcal{UB}[\eta_{k+n}|M_{k+n}^s], \end{aligned} \quad (\text{B.33})$$

where $\sigma^i \triangleq \max(\mathbb{P}(Z_{k+i}|x_{k+i}))$ and w_k^p is the prior weight at time k of the belief component p_{k+n} . ■

B.3.2 The Cost Function Bounds $\mathcal{LB}[\mathcal{H}_{k+n}]$, $\mathcal{LB}[\mathcal{H}_{k+n}]$

Deriving a direct recursive update rule for these bounds is not trivial. Instead, we show how each term in $\mathcal{LB}[\mathcal{H}_{k+n}|M_{k+n}^{s+1}]$, $\mathcal{UB}[\mathcal{H}_{k+n}|M_{k+n}^{s+1}]$ can be incrementally updated individually.

Given a belief component $p_{k+n} \notin M_{k+n}^s$ with associated weight w_{k+n}^p , we first derive a recursive update rule for the cost over the simplified belief b_{k+n}^{s+1} , i.e. containing

$M_{k+n}^{s+1} \triangleq M_{k+n}^s \cup p_{k+n}$ components

$$\begin{aligned}
\mathcal{H}_{k+n}^{s+1} &\triangleq c \left(b_{k+n}^{s+1} \right) \\
&= - \sum_{r \in M_{k+n}^{s+1}} \frac{w_{k+n}^r}{\sum_{r \in M_{k+n}^{s+1}} w_{k+n}^r} \log \left(\frac{w_{k+n}^r}{\sum_{r \in M_{k+n}^{s+1}} w_{k+n}^r} \right) \\
&= - \frac{w_{k+n}^p}{\sum_{r \in M_{k+n}^{s+1}} w_{k+n}^r} \log \left(\frac{w_{k+n}^p}{\sum_{r \in M_{k+n}^{s+1}} w_{k+n}^r} \right) - \\
&\quad \sum_{r \in M_{k+n}^s} \frac{w_{k+n}^r}{\sum_{r \in M_{k+n}^{s+1}} w_{k+n}^r} \log \left(\frac{w_{k+n}^r}{\sum_{r \in M_{k+n}^{s+1}} w_{k+n}^r} \right) \\
&= - \frac{w_{k+n}^p}{w_{k+n}^p + \sum_{r \in M_{k+n}^s} w_{k+n}^r} \log \left(\frac{w_{k+n}^p}{w_{k+n}^p + \sum_{r \in M_{k+n}^s} w_{k+n}^r} \right) - \\
&\quad \sum_{r \in M_{k+n}^s} \frac{w_{k+n}^r}{w_{k+n}^p + \sum_{r \in M_{k+n}^s} w_{k+n}^r} \log \left(\frac{w_{k+n}^r}{w_{k+n}^p + \sum_{r \in M_{k+n}^s} w_{k+n}^r} \right) \\
&= - \frac{w_{k+n}^p}{w_{k+n}^p + w_{k+n}^{m,s}} \log \left(\frac{w_{k+n}^p}{w_{k+n}^p + w_{k+n}^{m,s}} \right) - \sum_{r \in M_{k+n}^s} \frac{w_{k+n}^r}{w_{k+n}^p + w_{k+n}^{m,s}} \log \left(\frac{w_{k+n}^r}{w_{k+n}^p + w_{k+n}^{m,s}} \right) \\
&= - \frac{w_{k+n}^p}{w_{k+n}^p + w_{k+n}^{m,s}} \log \left(\frac{w_{k+n}^p}{w_{k+n}^p + w_{k+n}^{m,s}} \right) - \\
&\quad \frac{w_{k+n}^{m,s}}{w_{k+n}^p + w_{k+n}^{m,s}} \sum_{r \in M_{k+n}^s} \frac{w_{k+n}^r}{w_{k+n}^{m,s}} \left[\log \left(\frac{w_{k+n}^r}{w_{k+n}^{m,s}} \right) + \log \left(\frac{w_{k+n}^{m,s}}{w_{k+n}^p + w_{k+n}^{m,s}} \right) \right] \\
&= - \frac{w_{k+n}^p}{w_{k+n}^p + w_{k+n}^{m,s}} \log \left(\frac{w_{k+n}^p}{w_{k+n}^p + w_{k+n}^{m,s}} \right) - \\
&\quad \frac{w_{k+n}^{m,s}}{w_{k+n}^p + w_{k+n}^{m,s}} \left[-\mathcal{H}_{k+n}^s + \sum_{r \in M_{k+n}^s} \frac{w_{k+n}^r}{w_{k+n}^{m,s}} \log \left(\frac{w_{k+n}^r}{w_{k+n}^p + w_{k+n}^{m,s}} \right) \right] \\
&= - \frac{w_{k+n}^p}{w_{k+n}^p + w_{k+n}^{m,s}} \log \left(\frac{w_{k+n}^p}{w_{k+n}^p + w_{k+n}^{m,s}} \right) - \\
&\quad \frac{w_{k+n}^{m,s}}{w_{k+n}^p + w_{k+n}^{m,s}} \left[-\mathcal{H}_{k+n}^s + \log \left(\frac{w_{k+n}^{m,s}}{w_{k+n}^p + w_{k+n}^{m,s}} \right) \sum_{r \in M_{k+n}^s} \frac{w_{k+n}^r}{w_{k+n}^{m,s}} \right] \\
&= - \frac{w_{k+n}^p}{w_{k+n}^p + w_{k+n}^{m,s}} \log \left(\frac{w_{k+n}^p}{w_{k+n}^p + w_{k+n}^{m,s}} \right) - \frac{w_{k+n}^{m,s}}{w_{k+n}^p + w_{k+n}^{m,s}} \left[\log \left(\frac{w_{k+n}^{m,s}}{w_{k+n}^p + w_{k+n}^{m,s}} \right) - \mathcal{H}_{k+n}^s \right].
\end{aligned} \tag{B.34}$$

Using (B.34) and the recursive update rules derived in Section B.3.1, we get a recursive

update rule for the lower bound

$$\begin{aligned}
& \mathcal{LB} \left[\mathcal{H}_{k+n} | M_{k+n}^{s+1} \right] \\
&= \frac{\sum_{r \in M_{k+n}^{s+1}} w_{k+n}^r}{\mathcal{UB} \left[\eta_{k+n} | M_{k+n}^{s+1} \right]} \left[\mathcal{H}_{k+n}^{s+1} + \log \left(\frac{\mathcal{LB} \left[\eta_{k+n} | M_{k+n}^{s+1} \right]}{\sum_{r \in M_{k+n}^{s+1}} w_{k+n}^r} \right) \right] \\
&= \frac{w_{k+n}^p + w_{k+n}^{m,s}}{\mathcal{UB} \left[\eta_{k+n} | M_{k+n}^{s+1} \right]} \left[\mathcal{H}_{k+n}^{s+1} + \log \left(\frac{\mathcal{LB} \left[\eta_{k+n} | M_{k+n}^{s+1} \right]}{w_{k+n}^p + w_{k+n}^{m,s}} \right) \right].
\end{aligned} \tag{B.35}$$

Similarly, we also get a recursive update rule for the upper bound

$$\begin{aligned}
& \mathcal{UB} \left[\mathcal{H}_{k+n} | M_{k+n}^{s+1} \right] \\
&= \frac{\sum_{r \in M_{k+n}^{s+1}} w_{k+n}^r}{\mathcal{LB} \left[\eta_{k+n} | M_{k+n}^{s+1} \right]} \left[\mathcal{H}_{k+n}^{s+1} + \log \left(\frac{\mathcal{UB} \left[\eta_{k+n} | M_{k+n}^{s+1} \right]}{\sum_{r \in M_{k+n}^{s+1}} w_{k+n}^r} \right) \right] - \bar{\gamma} \log \left(\frac{\bar{\gamma}}{|\neg M_{k+n}^{s+1}|} \right) \\
&= \frac{w_{k+n}^p + w_{k+n}^{m,s}}{\mathcal{LB} \left[\eta_{k+n} | M_{k+n}^{s+1} \right]} \left[\mathcal{H}_{k+n}^{s+1} + \log \left(\frac{\mathcal{UB} \left[\eta_{k+n} | M_{k+n}^{s+1} \right]}{w_{k+n}^p + w_{k+n}^{m,s}} \right) \right] - \bar{\gamma} \log \left(\frac{\bar{\gamma}}{|\neg M_{k+n}^{s+1}|} \right),
\end{aligned} \tag{B.36}$$

where $\bar{\gamma} = 1 - \sum_{r \in M_{k+n}^{s+1}} \frac{w_{k+n}^r}{\mathcal{UB} \left[\eta_{k+n} | M_{k+n}^{s+1} \right]} = 1 - \frac{w_{k+n}^p + w_{k+n}^{m,s}}{\mathcal{UB} \left[\eta_{k+n} | M_{k+n}^{s+1} \right]}$ and $|\neg M_{k+n}^{s+1}| > 2$. ■

Appendix C

Proofs for Chapter 5

C.1 Lemma 2

Proof .

We prove this lemma by induction.

base case: Let θ_1 be the first eliminated variable. Due to the properties of the elimination order given in Sec. 5.1, $\exists f(\theta_1) \in \mathcal{F}$ from which samples of θ_1 can be generated.

induction step: Eliminating θ_j , if $\exists f(\theta_j) \in \mathcal{F}$ then according to the elimination algorithm it holds that $f(\theta_j) \in \mathcal{F}_{j-1}$ from which samples of θ_j can be generated. Else, according to the given elimination order $\exists \theta_i \in \Theta$ such that θ_i was previously eliminated and $f(\theta_i, \theta_j) \in \mathcal{F}$. Since $f(\theta_i, \theta_j) \in \mathcal{F}$, according to the elimination algorithm it holds that $f(\theta_i, \theta_j) \in \mathcal{F}_{i-1}$ and thus $\theta_j \in S_i$. As such, according to (5.13), when θ_i was eliminated, a new factor

$$f_{new}(S_i|D_i) = \eta^{-1} \int_{\theta_i} \prod_{f_k \in \mathcal{F}_{i-1}(\theta_i)} f_k(\Theta_k) d\theta_i, \quad (\text{C.1})$$

was added. Using the induction assumption, samples of θ_i were generated from one of the factors in $\mathcal{F}_{i-1}(\theta_i)$ to approximate this integral. Without loss of generality, we denote the factor from which samples of θ_i were generated as $f_{\bar{k}}$ and write the approximation to (C.1) as

$$\hat{f}_{new}(S_i|D_i) = \frac{\eta^{-1}}{N} \sum_{n=1}^N f(\theta_i^n, \theta_j) \prod_{f_k \in \mathcal{F}_{i-1}(\theta_i) \setminus \{f(\theta_i^n, \theta_j), f_{\bar{k}}\}} f_k(\theta_i^n, \Theta_{\bar{k}}^{-i}). \quad (\text{C.2})$$

According to the elimination algorithm, for any variable $\theta_m \in S_i$ eliminated after θ_i and before θ_j , it must hold that the new factor $f_{new}(S_m|D_m)$ must contain $f(\theta_i^n, \theta_j)$. Thus, \mathcal{F}_{j-1} also contains $f(\theta_i^n, \theta_j)$ from which samples of θ_j can be generated. ■

Bibliography

- [1] Saurav Agarwal, Amirhossein Tamjidi, and Suman Chakravorty. “Motion planning for active data association and localization in non-Gaussian belief spaces.” In: *Intl. Workshop on the Algorithmic Foundations of Robotics (WAFR)*. Vol. 13. Springer Proceedings in Advanced Robotics. Springer, 2016, pp. 288–303.
- [2] A.-A. Agha-Mohammadi, S. Chakravorty, and N. M. Amato. “FIRM: Sampling-based feedback motion planning under motion uncertainty and imperfect measurements.” In: *Intl. J. of Robotics Research* 33.2 (2014), pp. 268–304.
- [3] Adrien Angeli, Stéphane Doncieux, Jean-Arcady Meyer, and David Filliat. “Visual topological SLAM and global localization.” In: *IEEE Intl. Conf. on Robotics and Automation (ICRA)*. IEEE. 2009, pp. 4300–4305.
- [4] N. Atanasov, B. Sankaran, J.L. Ny, G. J. Pappas, and K. Daniilidis. “Nonmyopic View Planning for Active Object Classification and Pose Estimation.” In: *IEEE Trans. Robotics* 30 (2014), pp. 1078–1090.
- [5] R. Bellman, Rand Corporation, and Karreman Mathematics Research Collection. *Dynamic Programming*. Rand Corporation research study. Princeton University Press, 1957.
- [6] Subhrajit Bhattacharya, Robert Ghrist, and Vijay Kumar. “Persistent homology for path planning in uncertain environments.” In: *IEEE Trans. Robotics* 31.3 (2015), pp. 578–590.
- [7] Norman Biggs. *Algebraic Graph Theory*. 2nd edition. Cambridge University Press, 1993, p. 205.
- [8] Blai Bonet and Héctor Geffner. “Planning with Incomplete Information As Heuristic Search in Belief Space.” In: *Intl. Conf. on Artificial Intelligence Planning Systems*. Breckenridge, CO, USA: AAAI Press, 2000, pp. 52–61.
- [9] Yoram Bresler. “Two-filter formulae for discrete-time non-linear Bayesian smoothing.” In: *International Journal of Control* 43.2 (1986), pp. 629–641.
- [10] Mark Briers, Arnaud Doucet, and Simon Maskell. “Smoothing algorithms for state-space models.” In: *Annals of the Institute of Statistical Mathematics* 62 (2010), pp. 61–89.

- [11] 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)*. IEEE. 2014, pp. 2667–2674.
- [12] Luca Carlone. “A convergence analysis for pose graph optimization via Gauss-Newton methods.” In: *IEEE Intl. Conf. on Robotics and Automation (ICRA)*. IEEE. 2013, pp. 965–972.
- [13] Luca Carlone, Rosario Aragues, José A Castellanos, and Basilio Bona. “A fast and accurate approximation for planar pose graph optimization.” In: *Intl. J. of Robotics Research* 33.7 (2014), pp. 965–987.
- [14] Yongbo Chen, Shoudong Huang, Robert Fitch, Liang Zhao, Huan Yu, and Di Yang. “On-line 3D active pose-graph SLAM based on key poses using graph topology and sub-maps.” In: *IEEE Intl. Conf. on Robotics and Automation (ICRA)*. 2019, pp. 169–175.
- [15] Yongbo Chen, Shoudong Huang, Liang Zhao, and Gamini Dissanayake. “Cramér–Rao bounds and optimal design metrics for pose-graph SLAM.” In: *IEEE Trans. Robotics* 37.2 (2021), pp. 627–641.
- [16] H. Choset and K. Nagatani. “Topological simultaneous localization and mapping (SLAM): toward exact localization without explicit localization.” In: *IEEE Trans. Robot. Automat.* 17.2 (2001), pp. 125–137.
- [17] Fan Chung and Ross M Richardson. “Weighted Laplacians and the sigma function of a graph.” In: *Contemporary Mathematics* 415 (2006), pp. 93–108.
- [18] Fan RK Chung. *Spectral graph theory*. Vol. 92. American Mathematical Soc., 1997.
- [19] T.M. Cover and J.A. Thomas. *Elements of Information Theory*. John Wiley & Sons, New York, NY, 1991.
- [20] James C Davidson and Seth A Hutchinson. “A sampling hyperbelief optimization technique for stochastic systems.” In: *Algorithmic Foundation of Robotics VIII: Selected Contributions of the Eight International Workshop on the Algorithmic Foundations of Robotics*. Springer. 2009, pp. 217–231.
- [21] F. Dellaert. *Factor Graphs and GTSAM: A Hands-on Introduction*. Tech. rep. GT-RIM-CP&R-2012-002. Georgia Institute of Technology, 2012.
- [22] F. Dellaert and M. Kaess. “Square Root SAM: Simultaneous Localization and Mapping via Square Root Information Smoothing.” In: *Intl. J. of Robotics Research* 25.12 (2006), pp. 1181–1203.
- [23] Joseph Djugash, Bradley Hamner, and Stephan Roth. “Navigating with ranging radios: Five data sets with ground truth.” In: *Journal of Field Robotics* 26.9 (2009), pp. 689–695.

- [24] Arnaud Doucet, Nando De Freitas, Neil James Gordon, et al. *Sequential Monte Carlo methods in practice*. Vol. 1. 2. Springer, 2001.
- [25] Khen Elimelech and Vadim Indelman. “Consistent Sparsification for Efficient Decision Making Under Uncertainty in High Dimensional State Spaces.” In: *IEEE Intl. Conf. on Robotics and Automation (ICRA)*. Singapore, May 2017, pp. 3786–3791.
- [26] Khen Elimelech and Vadim Indelman. “Fast Action Elimination for Efficient Decision Making and Belief Space Planning Using Bounded Approximations.” In: *Proc. of the Intl. Symp. of Robotics Research (ISRR)*. Chile, Dec. 2017.
- [27] Khen Elimelech and Vadim Indelman. “Scalable Sparsification for Efficient Decision Making Under Uncertainty in High Dimensional State Spaces.” In: *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*. Canada, Sept. 2017, pp. 5668–5673.
- [28] Khen Elimelech and Vadim Indelman. “Simplified decision making in the belief space using belief sparsification.” In: *The International Journal of Robotics Research* 41.5 (2022), pp. 470–496.
- [29] D. Fourie, J. Leonard, and M. Kaess. “A Nonparametric Belief Solution to the Bayes Tree.” In: *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*. 2016.
- [30] Andrea Gasparri, Stefano Panzieri, Federica Pascucci, and Giovanni Ulivi. “A hybrid active global localisation algorithm for mobile robots.” In: *IEEE Intl. Conf. on Robotics and Automation (ICRA)*. IEEE. 2007, pp. 3148–3153.
- [31] N.J. Gordon, D.J. Salmond, and A.F.M. Smith. “Novel approach to nonlinear/non-Gaussian Bayesian state estimation.” In: *IEE Proceedings F* 140.2 (1993), pp. 107–113.
- [32] Arthur Gretton, Karsten M Borgwardt, Malte J Rasch, Bernhard Schölkopf, and Alexander Smola. “A kernel two-sample test.” In: *J. of Machine Learning Research* 13.1 (2012), pp. 723–773.
- [33] Lin Han, Francisco Escolano, Edwin R. Hancock, and Richard C. Wilson. “Graph characterizations from von Neumann entropy.” In: *Pattern Recognition Letters* 33.15 (2012). Graph-Based Representations in Pattern Recognition, pp. 1958–1967. URL: <http://www.sciencedirect.com/science/article/pii/S0167865512000980>.
- [34] W.K. Hastings. “Monte Carlo sampling methods using Markov chains and their applications.” In: *Biometrika* 57 (1970), pp. 97–109.
- [35] P. Heggernes and P. Matstoms. “Finding Good Column Orderings for Sparse QR Factorization.” In: *Second SIAM Conference on Sparse Matrices*. 1996.

- [36] Roger A Horn and Charles R Johnson. *Matrix analysis*. Cambridge university press, 2012.
- [37] M. Hsiao and M. Kaess. “MH-iSAM2: Multi-hypothesis iSAM using Bayes Tree and Hypo-tree.” In: *IEEE Intl. Conf. on Robotics and Automation (ICRA)*. 2019.
- [38] Ming Hsiao, Joshua G Mangelson, Sudharshan Suresh, Christian Debrunner, and Michael Kaess. “ARAS: Ambiguity-aware Robust Active SLAM based on Multi-hypothesis State and Map Estimations.” In: *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*. IEEE. 2020, pp. 5037–5044.
- [39] J. Hsiung, M. Hsiao, E. Westman, R. Valencia, and M. Kaess. “Information Sparsification in Visual-Inertial Odometry.” In: *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*. 2018, pp. 1146–1153.
- [40] G. Huang, M. Kaess, and J.J. Leonard. “Consistent Sparsification for Graph Optimization.” In: *Proc. of the European Conference on Mobile Robots (ECMR)*. 2012, pp. 150 –157.
- [41] Guoquan Huang, Michael Kaess, John J Leonard, and Stergios I Roumeliotis. “Analytically-selected multi-hypothesis incremental MAP estimation.” In: *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*. 2013, pp. 6481–6485.
- [42] Qiangqiang Huang and John J Leonard. “GAPSLAM: Blending Gaussian Approximation and Particle Filters for Real-Time Non-Gaussian SLAM.” In: *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE. 2023, pp. 9183–9190.
- [43] Qiangqiang Huang, Alan Papalia, and John J Leonard. “Nested Sampling for Non-Gaussian Inference in SLAM Factor Graphs.” In: *IEEE Robotics and Automation Letters (RA-L)* 7.4 (2022), pp. 9232–9239.
- [44] Qiangqiang Huang, Can Pu, Kasra Khosoussi, David M Rosen, Dehann Fourie, Jonathan P How, and John J Leonard. “Incremental non-Gaussian inference for SLAM using normalizing flows.” In: *IEEE Trans. Robotics* 39.2 (2023), pp. 1458–1475.
- [45] Alexander Ihler, Erik Sudderth, William Freeman, and Alan Willsky. “Efficient multiscale sampling from products of Gaussian mixtures.” In: *Advances in Neural Information Processing Systems (NIPS)* 16 (2003).
- [46] Viorela Ila, Lukas Polok, Marek Solony, and Pavel Svoboda. “SLAM++ - A highly efficient and temporally scalable incremental SLAM framework.” In: *Intl. J. of Robotics Research* 36.2 (2017), pp. 210–230.
- [47] V. Indelman. “No Correlations Involved: Decision Making Under Uncertainty in a Conservative Sparse Information Space.” In: *IEEE Robotics and Automation Letters (RA-L)* 1.1 (2016), pp. 407–414.

- [48] V. Indelman. “Towards Information-Theoretic Decision Making in a Conservative Information Space.” In: *American Control Conference*. Chicago, USA, 2015, pp. 2420–2426.
- [49] V. Indelman, L. Carlone, and F. Dellaert. “Planning in the Continuous Domain: a Generalized Belief Space Approach for Autonomous Navigation in Unknown Environments.” In: *Intl. J. of Robotics Research* 34.7 (2015), pp. 849–882.
- [50] 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.” In: *IEEE Control Systems Magazine (CSM), Special Issue on Distributed Control and Estimation for Robotic Vehicle Networks* 36.2 (2016), pp. 41–74.
- [51] 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.
- [52] V. Indelman, R. Roberts, and F. Dellaert. “Incremental Light Bundle Adjustment for Structure From Motion and Robotics.” In: *Robotics and Autonomous Systems* 70 (2015), pp. 63–82.
- [53] Patric Jensfelt and Steen Kristensen. “Active global localization for a mobile robot using multiple hypothesis tracking.” In: *IEEE Transactions on Robotics and Automation* 17.5 (2001), pp. 748–760.
- [54] Fan Jiang, Varun Agrawal, Russell Buchanan, Maurice Fallon, and Frank Dellaert. “iMHS: An Incremental Multi-Hypothesis Smoother.” In: *arXiv preprint arXiv:2103.13178* (2021).
- [55] L. P. Kaelbling, M. L. Littman, and A. R. Cassandra. “Planning and acting in partially observable stochastic domains.” In: *Artificial intelligence* 101.1 (1998), pp. 99–134.
- [56] M. Kaess, V. Ila, R. Roberts, and F. Dellaert. “The Bayes Tree: An Algorithmic Foundation for Probabilistic Robot Mapping.” In: *Intl. Workshop on the Algorithmic Foundations of Robotics*. 2010.
- [57] M. Kaess, H. Johannsson, R. Roberts, V. Ila, J. Leonard, and F. Dellaert. “iSAM2: Incremental Smoothing and Mapping Using the Bayes Tree.” In: *Intl. J. of Robotics Research* 31.2 (2 2012), pp. 217–236.
- [58] M. Kaess, A. Ranganathan, and F. Dellaert. “iSAM: Incremental Smoothing and Mapping.” In: *IEEE Trans. Robotics* 24.6 (2008), pp. 1365–1378.
- [59] Sertac Karaman, Matthew R Walter, Alejandro Perez, Emilio Frazzoli, and Seth Teller. “Anytime motion planning using the RRT*.” In: *IEEE Intl. Conf. on Robotics and Automation (ICRA)*. IEEE. 2011.

- [60] L.E. Kavraki, P. Svestka, J.-C. Latombe, and M.H. Overmars. “Probabilistic Roadmaps for Path Planning in High-Dimensional Configuration Spaces.” In: *IEEE Trans. Robot. Automat.* 12.4 (1996), pp. 566–580.
- [61] Kasra Khosoussi, Matthew Giamou, Gaurav S Sukhatme, Shoudong Huang, Gamini Dissanayake, and Jonathan P How. “Reliable graph topologies for SLAM.” In: *Intl. J. of Robotics Research* 38.2-3 (2019), pp. 260–298.
- [62] Kasra Khosoussi, Shoudong Huang, and Gamini Dissanayake. “Good, bad and ugly graphs for SLAM.” In: *Robotics: Science and Systems (RSS)*. Workshop on the problem of mobile sensors. 2015.
- [63] Kasra Khosoussi, Shoudong Huang, and Gamini Dissanayake. “Novel insights into the impact of graph structure on SLAM.” In: *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*. 2014, pp. 2707–2714.
- [64] Soonkyum Kim, Subhrajit Bhattacharya, Robert Ghrist, and Vijay Kumar. “Topological exploration of unknown and partially known environments.” In: *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*. IEEE. 2013, pp. 3851–3858.
- [65] Genshiro Kitagawa. “The two-filter formula for smoothing and an implementation of the Gaussian-sum smoother.” In: *Annals of the Institute of Statistical Mathematics* 46 (1994), pp. 605–623.
- [66] Andrej Kitanov and Vadim Indelman. “Topological Belief Space Planning for Active SLAM with Pairwise Gaussian Potentials and Performance Guarantees.” In: *Intl. J. of Robotics Research* 43.1 (2024), pp. 69–97.
- [67] Andrej Kitanov and Vadim Indelman. “Topological Multi-Robot Belief Space Planning in Unknown Environments.” In: *IEEE Intl. Conf. on Robotics and Automation (ICRA)*. 2018, pp. 5726–5732.
- [68] D. Kopitkov and V. Indelman. “No Belief Propagation Required: Belief Space Planning in High-Dimensional State Spaces via Factor Graphs, Matrix Determinant Lemma and Re-use of Calculation.” In: *Intl. J. of Robotics Research* 36.10 (2017), pp. 1088–1130.
- [69] Dmitry Kopitkov and Vadim Indelman. “General Purpose Incremental Covariance Update and Efficient Belief Space Planning via Factor-Graph Propagation Action Tree.” In: *Intl. J. of Robotics Research* 38.14 (2019), pp. 1644–1673.
- [70] F.R. Kschischang, B.J. Frey, and H-A. Loeliger. “Factor Graphs and the Sum-Product Algorithm.” In: *IEEE Trans. Inform. Theory* 47.2 (2001), pp. 498–519.
- [71] Pierre-Yves Lajoie, Siyi Hu, Giovanni Beltrame, and Luca Carlone. “Modeling Perceptual Aliasing in SLAM via Discrete-Continuous Graphical Models.” In: *IEEE Robotics and Automation Letters (RA-L)* (2019).

- [72] 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.
- [73] D. Levine and J. P. How. “Sensor selection in high-dimensional Gaussian trees with nuisances.” In: *Advances in Neural Information Processing Systems (NIPS)*. 2013, pp. 2211–2219.
- [74] Bruce A McElhoe. “An assessment of the navigation and course corrections for a manned flyby of mars or venus.” In: *IEEE Transactions on Aerospace and Electronic Systems* 4 (1966), pp. 613–623.
- [75] M. Montemerlo, S. Thrun, D. Koller, and B. Wegbreit. “FastSLAM: A Factored Solution to the Simultaneous Localization and Mapping Problem.” In: *Proc. 19th AAAI National Conference on AI*. Edmonton, Alberta, Canada, 2002.
- [76] Abbe Mowshowitz and Matthias Dehmer. “Entropy and the complexity of graphs revisited.” In: *Entropy* 14.3 (2012), pp. 559–570.
- [77] Beipeng Mu, Liam Paull, Ali-Akbar Agha-Mohammadi, John J Leonard, and Jonathan P How. “Two-stage focused inference for resource-constrained minimal collision navigation.” In: *IEEE Transactions on Robotics* 33.1 (2017), pp. 124–140.
- [78] Raul Mur-Artal and Juan D Tardós. “Orb-slam2: An open-source slam system for monocular, stereo, and rgb-d cameras.” In: *IEEE Trans. Robotics* 33.5 (2017), pp. 1255–1262.
- [79] E. Olson and P. Agarwal. “Inference on networks of mixtures for robust robot mapping.” In: *Intl. J. of Robotics Research* 32.7 (2013), pp. 826–840.
- [80] E. Olson and M. Kaess. “Evaluating the Performance of Robot Mapping Systems.” In: *Workshop on Good Experimental Methodology in Robotics*. 2009.
- [81] Diane Valerie Ouellette. “Schur complements and statistics.” In: *Linear Algebra and its Applications* 36 (1981), pp. 187–295.
- [82] C. Papadimitriou and J. Tsitsiklis. “The complexity of Markov decision processes.” In: *Mathematics of operations research* 12.3 (1987), pp. 441–450.
- [83] Filippo Passerini and Simone Severini. “Quantifying Complexity in Networks: The Von Neumann Entropy.” In: *Int. J. Agent Technol. Syst.* 1.4 (2009), pp. 58–67. URL: <http://dx.doi.org/10.4018/jats.2009071005>.
- [84] S. Pathak, A. Thomas, and V. Indelman. “A Unified Framework for Data Association Aware Robust Belief Space Planning and Perception.” In: *Intl. J. of Robotics Research* 32.2-3 (2018), pp. 287–315.
- [85] T. Patten, M. Zillich, R. Fitch, M. Vincze, and S. Sukkarieh. “Viewpoint Evaluation for Online 3-D Active Object Classification.” In: *IEEE Robotics and Automation Letters (RA-L)* 1.1 (2016), pp. 73–81.

- [86] Rohan Paul and Paul Newman. “FAB-MAP 3D: Topological mapping with spatial and visual appearance.” In: *IEEE Intl. Conf. on Robotics and Automation (ICRA)*. IEEE. 2010, pp. 2649–2656.
- [87] J. Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann, 1988.
- [88] J. Pineau, G. J. Gordon, and S. Thrun. “Anytime Point-Based Approximations for Large POMDPs.” In: *J. of Artificial Intelligence Research* 27 (2006), pp. 335–380.
- [89] R. Platt, R. Tedrake, L.P. Kaelbling, and T. Lozano-Pérez. “Belief space planning assuming maximum likelihood observations.” In: *Robotics: Science and Systems (RSS)*. Zaragoza, Spain, 2010, pp. 587–593.
- [90] J. M. Porta, N. Vlassis, M. T. Spaan, and P. Poupart. “Point-based value iteration for continuous POMDPs.” In: *J. of Machine Learning Research* 7 (2006), pp. 2329–2367.
- [91] S. Prentice and N. Roy. “The belief roadmap: Efficient planning in belief space by factoring the covariance.” In: *Intl. J. of Robotics Research* 28.11-12 (2009), pp. 1448–1465.
- [92] A. Ranganathan and F. Dellaert. “Online probabilistic topological mapping.” In: *Intl. J. of Robotics Research* 30.6 (2011), pp. 755–771.
- [93] Danilo Rezende and Shakir Mohamed. “Variational inference with normalizing flows.” In: *Intl. Conf. on Machine Learning (ICML)*. PMLR. 2015, pp. 1530–1538.
- [94] Stéphane Ross, Joelle Pineau, Sébastien Paquet, and Brahim Chaib-Draa. “Online planning algorithms for POMDPs.” In: *J. of Artificial Intelligence Research* 32 (2008), pp. 663–704.
- [95] O. Shelly and V. Indelman. “Hypotheses Disambiguation in Retrospective.” In: *IEEE Robotics and Automation Letters (RA-L)* (2022). Accepted.
- [96] David Silver and Joel Veness. “Monte-Carlo planning in large POMDPs.” In: *Advances in Neural Information Processing Systems (NIPS)*. 2010, pp. 2164–2172.
- [97] R. Sim, P. Elinas, M. Griffin, and J.J. Little. “Vision-based SLAM using the Rao-Blackwellised Particle Filter.” In: *Proc. of the IJCAI Workshop on Reasoning with Uncertainty in Robotics (RUR)*. Edinburgh, Scotland, 2005.
- [98] John Skilling. “Nested sampling for general Bayesian computation.” In: *Bayesian analysis* 1(4) (2006), pp. 833–859.

- [99] Gerald L Smith, Stanley F Schmidt, and Leonard A McGee. *Application of statistical filter theory to the optimal estimation of position and velocity on board a circumlunar vehicle*. Vol. 135. National Aeronautics and Space Administration, 1962.
- [100] Harold W Sorenson. *Parameter estimation, volume 9 of Control and system theory*. Marcel Dekker, Inc., New-York and Basel, 1980.
- [101] Erik B Sudderth, Alexander T Ihler, William T Freeman, and Alan S Willsky. “Nonparametric belief propagation and facial appearance estimation.” In: (2002).
- [102] N. Sünderhauf and P. Protzel. “Switchable Constraints for Robust Pose Graph SLAM.” In: *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*. 2012.
- [103] N. Sunderhauf and P. Protzel. “Towards a robust back-end for pose graph SLAM.” In: *IEEE Intl. Conf. on Robotics and Automation (ICRA)*. IEEE. 2012, pp. 1254–1261.
- [104] Ori Sztyglic and Vadim Indelman. “Speeding up Online POMDP Planning via Simplification.” In: *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*. 2022.
- [105] Ori Sztyglic, Andrey Zhitnikov, and Vadim Indelman. “Simplified Belief-Dependent Reward MCTS Planning with Guaranteed Tree Consistency.” In: *arXiv preprint arXiv:2105.14239* (2021).
- [106] S. Thrun, Y. Liu, D. Koller, A.Y. Ng, Z. Ghahramani, and H. Durrant-Whyte. “Simultaneous Localization and Mapping With Sparse Extended Information Filters.” In: *Intl. J. of Robotics Research* 23.7-8 (2004), pp. 693–716.
- [107] J. Van Den Berg, S. Patil, and R. Alterovitz. “Motion planning under uncertainty using iterative local optimization in belief space.” In: *Intl. J. of Robotics Research* 31.11 (2012), pp. 1263–1278.
- [108] L. Wong, L. Kaelbling, and T. Lozano-Pérez. “Data association for semantic world modeling from partial views.” In: *Intl. J. of Robotics Research* 34.7 (2015), pp. 1064–1082.
- [109] Nan Ye, Adhiraj Somani, David Hsu, and Wee Sun Lee. “DESPOT: Online POMDP planning with regularization.” In: *JAIR* 58 (2017), pp. 231–266.
- [110] Andrey Zhitnikov, Ori Sztyglic, and Vadim Indelman. “No Compromise in Solution Quality: Speeding Up Belief-dependent Continuous POMDPs via Adaptive Multilevel Simplification.” In: *Intl. J. of Robotics Research* (2024). Accepted.

בין החסמים, נוכל לתת מידע לגבי הירידה באיכות הפתרון. אנו מדגימים כיצד הפתרונות המפושטים שלנו מותאמים לתרחישים בעולם האמיתי, בהם קיימת בעיית משאבים והסוכן נדרש לפעול תחת מגבלות של תקציב חישובי אשר מחייבות אותו להשליך חלק מן מידע בכדי לעמוד בדרישות עיבוד בזמן אמת. ייתרה מכך, אנו אף נראה כי במקרים מסוימים בהם ישנם מדידות זו משמעותיות ומגבלות תקציב גם יחד, השיטה שאנו מציעים יכולה לתת הבטחות ביצועים שלא היו יכולות להתקבל אלמלא פישוט הבעיה המקורית. לצורך פתרונות אלו, אנו מנצלים מבנים טופולוגיים, הקשורים לייצוגים גרפים בסיסיים של הבעיה, וכן מבנים ספציפיים אשר ניתן לחלץ מתוך ההתפלגויות הסתברותיות. בבעיית ההיסק, אנו מציעים גישה חדשנית לטיפול בהתפלגויות הסתברותיות מרובות ממדים המייצגת התפלגויות אלו בעזרת "פרוסות" מהמרחב הרב ממדי. על ידי שימוש במבנה של הפרוסות הללו, אנו יכולים להעריך ביעילות התפלגויות מרובות ממדים מכל צורה שהיא על מנת להוריד את הסיבוכיות החישובית אל עבר פעולה בזמן אמת.

תקציר

השימוש בסוכנים ורובוטים אוטונומיים חכמים הולך וגובר ואלו במהרה משפיעים על היבטים רבים של חיי היומיום שלנו. בשימושים שונים בעולם האמיתי, כגון ניווט אוטונומי, ניתוחים רובוטיים ומחסנים אוטומטיים, מצופה מסוכנים אלו לפעול בצורה אמינה ויעילה תחת מקורות שונים של אי ודאות, לעיתים עם ידע מוגבל על הסביבה. מקורות אי ודאות אלה רבים ומגוונים. הם כוללים סביבות דינמיות שבהן עלולים להתרחש אירועים בלתי צפויים, חיישנים רועשים עקב מגבלות פיזיקליות ומדידות בעלות זו משמעות בהן עצם אחד יכול להתפרש כעצם אחר בגלל דמיון מסוים ביניהם. כדי להתמודד בצורה אפקטיבית עם תרחישים שכאלו בעולם האמיתי, הסוכנים האוטונומיים צריכים לתחזק התפלגויות הסתברותיות מרובות ממדים הלוקחות בחשבון את אי הוודאות הנלווית. באופן מפורש, הסוכנים הללו צריכים להיות מסוגלים לפתור שתי בעיות, היסק וקבלת החלטות תחת אי ודאות. כאשר הסוכנים מסיקים על מצבם ומצב הסביבה, עליהם לעדכן את ההתפלגויות ההסתברותיות מרובות הממדים בהתבסס על המידע הזמין כך שייצגו את המצב האמיתי בצורה הטובה ביותר. לעומת זאת, בתהליך קבלת ההחלטות תחת אי ודאות, על הסוכן לקבוע באופן אוטונומי מהן הפעולות הבאות הטובות ביותר לביצוע בהתאם לפונקציית מטרה כלשהי. פתרון שתי הבעיות הללו - היסק וקבלת החלטות תחת אי ודאות - הוא יקר חישובית ולרוב בלתי מעשי במערכות אוטונומיות בעולם האמיתי אשר צריכות לפעול בזמן אמת תוך שימוש בחומרה זולה ומשאבים מוגבלים. הבעיות הללו אף הופכות למורכבות עוד יותר כאשר הסוכן צריך לקחת בחשבון מדידות בעלות זו משמעות. כאשר לא ניתן לשייך מדידה מסוימת באופן וודאי למקור כלשהו, על הסוכן לתחזק מספר השערות בהתאם למספר האפשרויות. ייתרה מכך, כאשר יתקבלו מדידות זו משמעותיות נוספות וזאת מבלי שיפתרו מדידות זו משמעותיות קודמות, מספר ההשערות אותן יידרש הסוכן לתחזק יגדל בצורה מעריכית עם הזמן. גידול שכזה, בעבור שתי הבעיות, יגביר את העומס החישובי עוד יותר. בעבודה זו אנו מתמודדים עם שתי הבעיות הקריטיות הללו ומפתחים גישות שונות להפחתת המורכבות החישובית במטרה לפעול בזמן אמת תחת מגבלות תקציב. הפתרונות שאנו מציעים מתבססים על שימוש במבנים שונים הנמצאים בייצוגים הבסיסיים של בעיות אלו בכדי ליעל את החישוב. בעבור תהליך קבלת החלטות תחת אי ודאות, אנו מציעים להשתמש במבנים אלו על מנת לפתור בעיות ממושטות בעלות סיבוכיות חישובית פחותה. הרעיון המרכזי הוא להפוך את הבעיה המורכבת, מרובת הממדים, לבעיה פשוטה יותר שמחזיקה במאפיינים החיוניים הדרושים לקבלת החלטות ובכך, להפחית את העומס החישובי ולהציע פתרון מעשי ליישומים בזמן אמת. בנוסף להיותה בעלת חישוביות סיבוכיות פחותה, אנו נדרוש כי בעיה פשוטה שכזו תוכל לאפיין בצורה כמותית כל אחת מן הפעולות אותן שוקל הסוכן לבצע. את הפתרון של בעיה זו, בעבור כל אחת מן הפעולות, נקשור בצורה אנליטית באמצעות חסמים, אל מול הפתרון של הבעיה המקורית, אותו אנו נמנעים מלחשב בצורה זו, נראה כיצד ניתן לתת הבטחות ביצועים של פתרון הבעיה המפושטת אל מול הבעיה המקורית. במקרה בו החסמים של הפתרון האופטימלי בבעיה המפושטת אינם חופפים כלל עם אף אחד מהחסמים של הפעולות האחרות, נוכל להבטיח כי הפתרון האופטימלי זהה לזה שבבעיה המקורית. במקרה וישנה חפיפה

המחקר בוצע בהנחייתו של פרופסור חבר ואדים אינדלמן, במסגרת התוכנית למערכות אוטונומיות ורובוטיקה.

מחבר/ת חיבור זה מצהיר/ה כי המחקר, כולל איסוף הנתונים, עיבודם והצגתם, התייחסות והשוואה למחקרים קודמים וכו', נעשה כולו בצורה ישרה, כמצופה ממחקר מדעי המבוצע לפי אמות המידה האתיות של העולם האקדמי. כמו כן, הדיווח על המחקר ותוצאותיו בחיבור זה נעשה בצורה ישרה ומלאה, לפי אותן אמות מידה.

חלק מן התוצאות בחיבור זה פורסמו כמאמרים מאת המחבר ושותפיו למחקר בכנסים ובכתבי-עת במהלך תקופת מחקר הדוקטורט של המחבר, אשר גרסאותיהם העדכניות ביותר הינן:

בכתבי עת

- [1] M. Shienman, A. Kitanov and V. Indelman, "FT-BSP: Focused Topological Belief Space Planning", *IEEE Robotics and Automation Letters (RA-L)*, vol. 6, no. 3, pp. 4744–4751, July 2021. DOI: 10.1109/LRA.2021.3068947.
- [2] M. Barenboim, M. Shienman and V. Indelman, "Monte Carlo Planning in Hybrid Belief POMDPs", *IEEE Robotics and Automation Letters (RA-L)*, vol. 8, no. 8, pp. 4410–4417, Aug. 2023.
- [3] M. Shienman, O. Levy-Or, M. Kaess, and V. Indelman, "Towards Real-Time Nonparametric Inference in High Dimensional State Spaces via Slices", To be submitted to *IEEE Transactions on Robotics (T-RO)*, 2024.

בכנסים

- [1] M. Shienman, and V. Indelman, "Nonmyopic Distilled Data Association Belief Space Planning Under Budget Constraints", in *International Symposium on Robotics Research (ISRR)*, Switzerland, Sep. 2022, pp. 102–118.
- [2] M. Shienman, and V. Indelman, "D2A-BSP: Distilled Data Association Belief Space Planning with Performance Guarantees Under Budget Constraints", in *IEEE International Conference on Robotics and Automation (ICRA)*, USA, May 2022, pp. 11058–11065. *Outstanding Paper Award Finalist*
- [3] M. Shienman, O. Levy-Or, M. Kaess, and V. Indelman, "A Slices Perspective for Incremental Nonparametric Inference in High Dimensional State Spaces", Accepted to *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, UAE, Oct. 2024.

אני מודה לטכניון על התמיכה הכספית הנדיבה בהשתלמותי.

הסקה ותכנון הסתברותיים מיודעי מבנה עם הבטחות ביצועים

חיבור על מחקר

לשם מילוי חלקי של הדרישות לקבלת התואר
דוקטור לפילוסופיה

משה שיינמן

הסקה ותכנון הסתברותיים מיודעי מבנה עם הבטחות ביצועים

משה שיינמן

