Belief Space Planning using Qualitative Spatial Relationships

Itai Zilberman

Belief Space Planning using Qualitative Spatial Relationships

Research Thesis

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

Itai Zilberman

Submitted to the Senate
of the Technion — Israel Institute of Technology
Adar 5782 Haifa February 2022

This research was carried out under the supervision of Associate Prof. Vadim Indelman and Prof. Ehud Rivelin in the Andrew and Erna Viterbi Faculty of Electrical Engineering.

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

1. I. Zilberman, E. Rivlin, and V. Indelman, "Incorporating compositions in qualitative approaches," IEEE Robotics and Automation Letters, vol. 7, no. 2, pp. 2660−2667, 2022.

2. I. Zilberman and V. Indelman, "Qualitative belief space planning via compositions," in IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS), 2020. submitted.

Acknowledgements

I would like to thank my advisor, Assoc. Prof. Vadim Indelman, for his enlightening and pleasant guidance. Also, I would like to thank my family for their support and for pushing me forward to success.

The generous financial help of the Technion is gratefully acknowledged.

Contents

List of Figures

List of Tables

Abstract

Planning under uncertainty, also known as belief space planning (BSP), is a crucial component in many autonomous systems and robotics applications. The aim is to find an optimal sequence of actions (i.e., plan) to accomplish a predefined task. Usually, the robot maintains a posterior distribution over the map and its self-locations, known as belief. A cost function is then defined on the belief as a penalty mechanism. The chosen plan is the one that most likely achieves the task's goal while minimizing the cost. Solving this problem becomes particularly challenging under conditions of uncertainty, that is, where there may be incomplete or faulty information, and the same actions may not always produce identical results. While most of the research done in the BSP domain formulates the problem via metric representation, other alternatives exist. A less common one is the Qualitative formulation, where the map is represented as a sparse set of Qualitative Spatial Relationships (QSRs) in independent and local coordinate frames. QSRs specify the robot's or landmarks' locations in terms of qualitative states ("Middle Right", "Top Left", etc.) rather than in an accurate manner. Since this approach does not aim to achieve an accurate description of the environment in the first place, it is more robust to noise compared to the metric. Thus, it becomes very useful in the absence of high-quality sensors. While several qualitative methods for localization and mapping exist in the literature, so far, only one qualitative work has addressed the planning problem.

Motivated by the above, in this research, we present a novel qualitative Belief Space Planning approach, highly suitable for platforms with low-cost sensors and particularly appealing in sparse environment scenarios. Innovatively, our planning algorithm smoothly incorporates compositions, which enable spatial information propagation between different QSRs to infer new ones.

However, incorporating compositions within a qualitative approach is not trivial. For instance, if the information required to perform a specific composition operation is unavailable, it must be inferred first, possibly via a preparatory composition operation. This recursive issue becomes more challenging as the amount of information grows. Two main questions are arising from the above, which remained open: 1. Given an initial set of qualitative spatial relationships, what new ones can be composed? 2. What is the optimal sequence of compositions operations to create a target QSR among all possible sequences? As a preparatory step for our qualitative BSP formulation, we address these

two questions and hence clarify how compositions should be incorporated in qualitative approaches in general.

Abbreviations and Notations

BSP	:	Belief Space Planning
SLAM	:	Simultaneous Localization and Mapping
QSR	:	Qualitative Spacial Relationship
QRM	:	Qualitative Relational Map
PQRM	:	Probabilistic Qualitative Relational Map
FSC	:	Freksa Single Cross partitioning
FDC	:	Freksa Double Cross partitioning
EDC	:	Extended Double Cross partitioning
WLOG	:	Without Loss Of Generality

$\mathbb{L}$	:	Landmark Space
L_i	:	The ith landmark under $\mathbb{L}$
$\mathbb{F}_t$	:	Frame Space at time step t
F_t	:	Frame at time step t
$F_t{:}L$	:	The triplet that locates L relative to frame F_t
$\mathcal{X}_{F_t}^{L}$	:	Metric state of $F_t{:}L$
$\mathcal{S}_{F_t}^{L}$	:	Qualitative state of $F_t{:}L$
$\mathcal{M}_t$	:	Set of all available landmark triplets from time step t
$\mathcal{S}^{\mathcal{M}_t}$	:	Set of qualitative state of elements from $\mathcal{M}_t$
$F_{t'}{:}X_t$	:	The triplet that locates the robot at time step t relative to $F_{t'}$
$\mathcal{X}_{F_{t'}}^{X_t}$	:	Metric state of $F_{t'}{:}X_t$
$\mathcal{S}_{F_{t'}}^{X_t}$	:	Qualitative state of $F_{t'}{:}X_t$
$\mathcal{S}^{X_{t':t}}$	:	Set of all consecutive robot states from time step $t^{'}$ to t
$\mathcal{X}^{F_t}$	:	Metric global frame scale of F_t
$\mathcal{S}^{F_t}$	:	Qualitative global frame scale of F_t
$\mathcal{S}^{\mathbb{F}_t}$	:	Set of qualitative scales of frames taken from $\mathbb{F}_t$
a_t^{q}	:	Qualitative action taken at time step t
a_t^{Link}	:	*Link* action taken at time step t
a_t	:	A tuple of qualitative and *Link* actions, both taken at time step t

β_t	:	Data Association from time step t
z_t	:	Measurement from time step t
$\square_{t':t}$	:	Sequence of consecutive elements from time index t' to t
$\mathcal{H}_t$	:	History of all actions measurements up to time t
$\mathcal{H}_t^-$	:	History of all actions up to time t and measurements up to time $t-1$
$\mathbb{P}(\cdot)$	:	Probability density function
b_t	:	Qualitative belief at time step t
$b_t^{\mathcal{M}_{t+1}}$	:	Qualitative belief at time step t augmented with composed triplets
R	:	Robot's sensing range
$c(\cdot)$	:	Cost function
J_t	:	Objective function at time step t
η	:	Normalization factor

Chapter 1

Introduction

1.1 Qualitative approaches

Many robotics applications rely on accurate metric estimations of the environment and robot's location to accomplish their aims.

However, in the absence of high-quality sensors, accurate robotics methods encounter significant difficulties. These methods, e.g., [1] and [2], are noise-sensitive and tend to accumulate errors as they rely on metrical estimates of map and robot's trajectory. Thus, noisy measurements can significantly impair their accuracy, cause undesirable drifts, and eventually lead to divergence if the loop-closer fails. Another concern that arises from the metric approaches is the need to maintain a dense, potentially large map representation, which often comes at the cost of substantial computational and memory resources.

While maintaining accurate information is often essential, it might be unnecessary in some cases, and therefore a burden. For instance, consider an autonomous cleaning robot operating in a living room, as illustrated in Fig. 1.1.

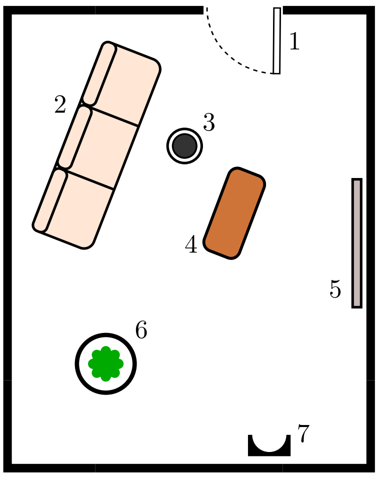

Figure 1.1: A simplified living room scene. 1-door; 2-sofa; 3-cleaning robot; 4-table; 5-TV; 6-flowerpot; 7-charger. Given known QSRs, the cleaning robot can compose new ones, which may aid it in accomplishing its task.

A typical living room contains a relatively small number of meaningful objects. Relying on rough relative relationships between the different objects, rather than on exact metric coordinates, may be sufficient for the robot to maneuver within the room successfully. E.g., if the robot seeks to clean under the table, it must pass safely between the table's legs. However, neither the exact metric coordinates of the legs nor the exact robot's location between them is required. Moreover, in long-term autonomous navigation missions, the robot is often required to travel long distances, so relying on a small number of critical landmarks along the way might be a good enough strategy. In other scenarios, the nature of the surrounding landscape is relatively monotonic and poor (for example, a desert or snowy terrain), and any methods that depend on finding dense features are prone to fail.

Qualitative approaches are motivated by the above. In contrast to the metrical methods, the environment and robot's poses are tracked using coarse, relative geometrical relations, known as qualitative spatial relationships (QSRs). Each QSR fixes a coordinate system based on a small set of landmarks and discretizes space into disjoint regions, called qualitative states. Then, the location of a target landmark or robot pose is described in terms of these states. See Fig. 1.2 below for illustration.

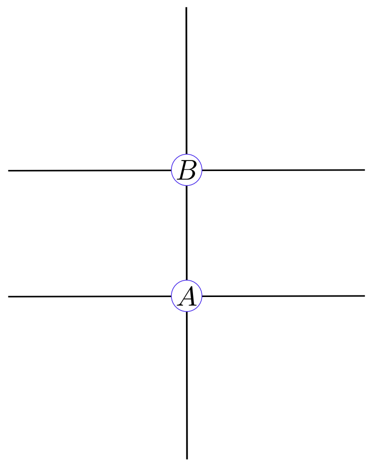

Figure 1.2: The Freksa Double Cross (FDC) partition divides the space into 6 disjoint qualitative states - "Top Left", "Top Right", "Middle Left", "Middle Right", "Bottom Left", "Bottom Right".

This coarse manner of reasoning about spatial information is potentially more noise-robust and suitable for low-cost platforms. Also, Qualitative Relational Mapping (QRM) algorithms produce QSR-based maps that sparsely represent the environment, in line with the motivation given above. Lastly, the robot can reason about the environment and even plan while accounting for partial information involved in a single or few QSR only, thus saving computational energy.

Since qualitative approaches represent spatial information through a set of QSRs, each in its own local independent coordinate frame, propagating information between different relationships may be essential. The composition operator was first introduced in 1992 out of this very need [3]. Given a pair of source QSRs, the latter aims to conclude the third one under some topological conditions. For example, consider a scenario where the above-mentioned cleaning robot has finished its work and is currently located between the sofa and the table. Suppose the robot seeks to return to its charger. Consider τ_1 to be a known QSR that describes the flowerpot's state relative to the sofa-table frame as "Middle Right". Furthermore, consider τ_2 to be another known QSR that describes the charger's state relative to the table-flowerpot frame as "Middle Left". Given τ_1 and τ_2, the robot can compose a new QSR, τ_3, which describes the charger's state relative to the sofa-table frame, as "Top Right". For illustration, see Fig. 1.3, where the sofa, table, flowerpot, and charger, correspond to landmarks A, B, C, and D. The robot then can use τ_3 to infer its target heading.

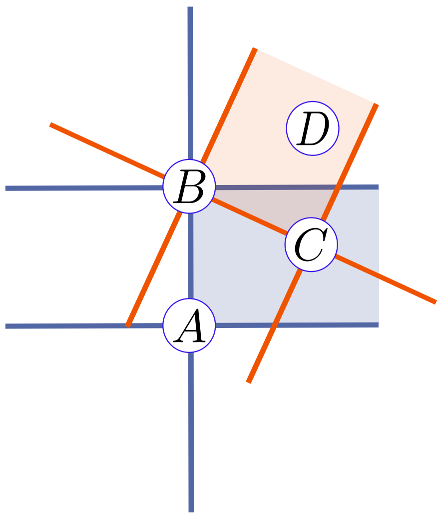

Figure 1.3: Illustration of the composition operation. Given evaluated qualitative states for $AB{:}C$ ("Middle Right") and $BC{:}D$ ("Middle Left"), the composition operator determines which qualitative states for $AB{:}D$ are feasible ("Top Right" and "Middle Right").

In the first part of this research, we focus on the composition operator alone and address two main theoretical issues that remain open until today. In the second part, we introduce a novel BSP approach adapted to the qualitative framework. Our approach utilizes compositions to further improve results.

Before stating our contributions, we briefly review the most relevant work done in the field to trace the existing gaps.

1.2 Related Work

QSR applications for various robotics tasks began to emerge about three decades ago.

Naturally, passive aspects were the first to be addressed. A pioneer work by [4] presented a novel approach for egocentric robot localization based on the relative ordering of observed landmarks. [5] and [6] further improved this idea by encoding ordering views in a more complex hence distinguishable fashion, enabling localization ambiguity reduction.

Freksa suggested in [3] to represent the qualitative location of a landmark relative to a boundary line settled by a pair of other landmarks used as a reference. The location is described as ”to the left” or ”to the right” of the boundary. Freksa further refined this binary partitioning of space in [7], into quadratic and hexagonal ones, by adding extra boundary lines perpendicularly crossing the original one. The latter mentioned partitioning forms known as the ”Freksa’s Single Cross” (FSC) and ”Freksa’s Double Cross” (FDC). Moreover, in [3], Freksa introduced the binary composition operator that allows inference about the qualitative relationships between landmarks not directly observed together.

Schlieder presented in [8] a set of qualitative spatial constraints between oriented straight line segments formed by pair of landmarks (dipoles) to describe the environment. [9] and [10] further extended this work and formulated a bipole-based composition operator. Another important QRM technique was suggested by [11], who developed a topological graph representation reflecting the adjacency between qualitative regions that divide the plane.

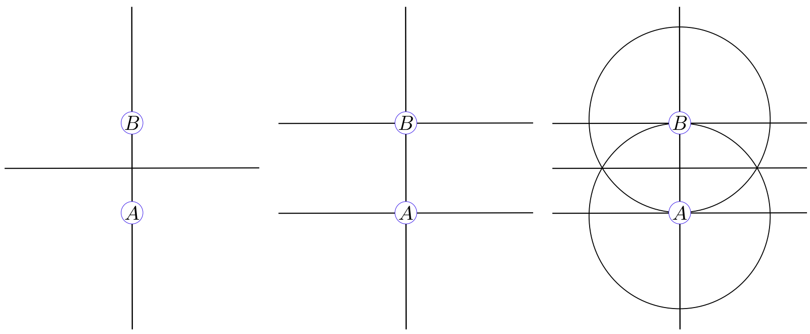

(a) (b) (c) (d)

Figure 1.4: Four different QSR partitions. (a) binary left-right [3]; (b) Freksa Single Cross (FSC, [7]); (c) Freksa Double Cross (FDC, [7]); (d) Extended Double Cross (EDC, [12]);

McClelland et al. took a step forward and introduced a more comprehensive QSR-based method for autonomous localization and mapping in [13]. The proposed algorithm constructs a graph-based map that encodes the environment using the relative geometrical layout of landmark triplets. For each triplet, one landmark is estimated in a local frame defined by the other two. The landmark is associated with one of several possible qualitative states, considering the FDC partitioning. McClelland extended his work in [12] by incorporating a method of determining the qualitative state for landmarks based on a novel set of geometric constraints. The above yielded a new qualitative spatial partitioning called the ”Extended Double Cross” (EDC). This partition and those mentioned above are illustrated in Fig. 1.4. In addition, this work contributed a new composition operator, suitable for triplets, formulated as a look-up table. Another followed-up paper generalized the latter by developing a probabilistic QRM method (PQRM, [14]). Mor and Indelman were the first to incorporate stochastic motion model constraints in his formulations, reducing uncertainty levels of both landmarks and robot trajectory estimations [15]. In addition, the authors contributed a novel derivation of a probabilistic composition.

While many qualitative methods for localization and mapping have been developed, active planning approaches are much rarer. Previously mentioned [11] and [12] suggested applying the Dijkstra algorithm on the graph-based maps they developed to find the shortest path to the desired destination. However, they both proposed only a general strategy rather than a detailed algorithm. In contrast, Padgett and Campbell developed a complete qualitative planning paradigm recently in [16]. The proposed *Q-Link* is a three-level planning architecture that generates high-level plans over QRM "links" (edges). It then uses local planners to execute trajectories to enable a robot to navigate from a start to a goal.

However, even though the *Q-Link* considered some stochastic aspects, its high-level mechanism is essentially deterministic and does not exploit the potential of using compositions.

1.3 Contributions

We shall now briefly review the main contributions of this research.

As we saw in Sec. 1.2, compositions have been formulated in numerous ways in past works. However, two fundamental questions remain open, which we aim to address in the first part of this research:

- Given a set of QSRs to start with, one seek to know what new QSRs can be formed using compositions. Consider the example from Sec. 1.1, given the initial set of τ_1, τ_2, and potentially more QSRs, τ_3 is only one possible target QSR among many others. We provide a theoretical derivation to address this question.

- Given a target triplet to compose, in general, they may be several alternatives to choose the source QSRs. Furthermore, at least one source might be unavailable and requires additional compositions to form it. This issue may occur recursively. Namely, there may be multiple composition sequences to form a target QSR, and it is unclear how to choose the optimal one. In the example, rather than composing τ_3 using τ_1 and τ_2, we could have chosen a different strategy if the initial QSRs set was richer. Thus, we develop an algorithm that finds the optimal sequence of composition operations to compose a target triplet, considering a cost criterion of interest.

In the second part, we aim to formulate a qualitative BSP approach that utilizes compositions. Our contributions in this part are:

- We develop a first-of-a-kind Qualitative Belief Space planning approach. Our planning mechanism accounts for possible future developments of the robot's and world's state by propagating a corresponding probabilistic distribution. It is not restricted by deterministic assumptions, in contrast to the methods discussed in Sec. 1.2, and thus more general.
- We innovatively incorporate compositions in our algorithm, allowing our planning process to consider qualitative relationships between landmarks that have never been observed together. We benefit from using these new relationships in two ways. Firstly, we can plan under challenging scenarios where other algorithms have failed. Secondly, we are able to find better plans in terms of expected cumulative cost, i.e., to find shortcuts.
- We introduce the concept of qualitatively estimating global scales of frames. Namely, we reason about the distances between pairs of landmarks, which form different reference frames. The distances are evaluated qualitatively in global terms. This capability is crucial when estimating the likelihood to observe a specific triplet of landmarks, which may not have been viewed together so far, given a candidate action to execute.
- We further take advantage of the latter capability and derive a new cost function, which globally measures metric path length.

Chapter 2

Background

In this chapter, we briefly explain several tools and concepts that serve as our work's basis. First, we introduce the Belief Space Planning problem considering the traditional metric formulation. Then, we elaborate on two concepts related to the qualitative field. The first is the *Link* action, which is unique to the qualitative framework and widely used in this work. The second is the Link-Graph, a useful topological representation for a qualitative map that reflects mobility in terms of *Link*s. Finally, we introduce the composition operator, allowing spatial information propagation between different QSRs.

2.1 Belief Space Planning

Belief Space Planning (BSP in short) is a probabilistic decision-making mechanism based on a *belief*, i.e., a posterior probability density function (pdf) maintained over variables of interest, such as the robot poses and the environment's landmarks.

We shall now concisely formulate the BSP problem, considering a metric space. Later on, in Sec. 4, we formulate a BSP approach considering a qualitative framework based on this formulation.

Let x_i and W_i denote the robot state and the world state at time step i. Also, let z_i and u_i denote the available observation and the control action applied at time step i. The belief at present time k is defined as:

$$b_k \triangleq \mathbb{P}(x_{0:k},W_k|z_{1:k},u_{1:k-1}), \tag{2.1}$$

where $x_{0:k} \triangleq \{x_0,\ldots,x_k\}$ denotes all past robot states from time step 0 to k. Also, $z_{1:k}\triangleq\{z_0,z_1,\ldots,z_k\}$ and $u_{1:k-1}\triangleq\{u_0,u_1,\ldots,u_{k-1}\}$ represent all the available measurements and past controls until time steps k and $k-1$, respectively.

At planning, we assume a future horizon of L look ahead steps. Consider the future time step $t\in\{k+1,\ldots k+L\}$. we define the objective $J_k(u_{k:k+L})$ as the cumulative cost with respect to a candidate control sequence $u_{k:k+L-1}$:

$$J(u_{k:k+L})=\mathbb{E}_{z_{k+1:k+L}}\Big[\sum_{l=1}^{L} c_l(b_{k+l},u_{k+l-1})\Big], \tag{2.2}$$

where c_l is the cost considered at the lth future step.

At planning time t, the optimal control sequence is the one that minimizes the above objective, i.e.:

$$u^*_{k:k+L} = \underset{u_{k:k+L}}{\arg\min}\, J(u_{k:k+L}). \tag{2.3}$$

A more comprehensive formulation of the Belief Space Planning considering a metric framework can be found in [17, 18].

2.2 Link Action

Navigation with a Qualitative Relational Map (QRM) poses a unique challenge because all information about the environment is preserved as coarse relationships relative to local reference frames (QSRs). As the robot moves, it relies on the QSRs in its immediate surroundings to locate itself, update the map based on newly acquired measurements, or plan its next course of action. Thus, the ability to transition between these QSRs' coordinate frames is essential. First introduced in [16], the *Link* action, enables the robot to do exactly that, i.e., to switch between different reference frames. Note that it is not a natural action in the sense that the robot makes no actual movement but rather only translates its self-location from one frame to another.

A robot can execute a *Link* action between two frames only if it has sufficient information required for the transition. Topologically speaking, sufficient information might be the triplet relates between the frames, as. For instance, to link from AB to BC, the robot can rely on the triplet $AB{:}C$ (i.e., C's state relative to AB). In a deterministic setting, as was considered in [16], sufficient information means the triplet's state. In a probabilistic setting, it often means a belief over the triplet's state. The following Lemma formalizes the above topological rule:

Lemma 2.2.1. *A direct* Link *from F_1 to F_2 is feasible based on a triplet τ, if $F_i \subseteq \tau, \forall i \in \{1,2\}$.*

To conceptually understand how the robot uses *Link*s during navigation, consider the example shown in Figure 2.1.

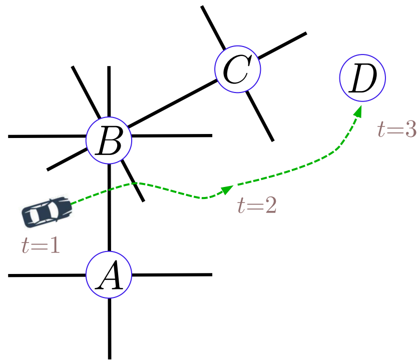

Figure 2.1: A simple example that illustrates how the robot navigates using *Link*s. At time step t=2, the robot links from AB to BC to further use BC:D's state and decide on its next action towards the goal landmark, D.

Suppose that at time step t=1, a robot is located at the "Middle Left" state relative to the frame AB and aims to reach landmark D, where the environment considered in this example contains the landmarks A,B,C, and D. Moreover, suppose that the only information the robot has on the environment is AB:C's and BC:D's qualitative states. As a first step, the robot computes an action towards landmark C. Given its initial state relative to AB, AB:X_1="Middle Left", and given that AB:C="Top Right", it infers that after "passing between A and B", C is likely to be at its left side. Suppose the robot has executed this action and reached near landmark C. Next, to reach its goal landmark D, the robot must rely on the only available triplet involving it, BC:D. Since the latter is given in BC coordinate system, the robot first needs to locate itself relative to this frame, i.e., to **link** from AB to BC. Given its current state, AB:X_2="Middle Right", and given that AB:C="Top Right", the robot deduces that after the *Link*, its new location is BC:X_2="Middle Right". Finally, using the state BC:D="Middle Right", the robot can continue in the same manner and compute its next action towards D. Such an action might be "continue straight while keeping C to your left".

In this research, we embrace the concept of *Link* action and widely use it, as we shall see in Section 4.

2.3 Link-Graph

A *Link-Graph* is a topological representation of a QRM. First introduced in [16], the *Link-Graph* was used for generating high-level plans over a QRM as part of a more comprehensive planning architecture called *Q-Link*.

Formally, the *Link-Graph* is defined as follows:

Definition 2.3.1. A *Link-Graph* is a graph $G = (V, E)$ where:

1. Each node $v \in V$ represents a triplet of landmarks, i.e., $v = \{L^1, L^2, L^3\}$.

2. There is an edge $e = (v_1, v_2) \in E$ if and only if $|v_1 \bigcap v_2| = 2$ (i.e., nodes v_1 and v_2 share exactly 2 landmarks in common).

*Link-Graph*s are useful representations, as they encode *Link*s' mobility (i.e., they visually encode Lemma 2.2.1). Recall from Section 2.2, a *Link* action between two frames can be taken (or considered during planning), based on the triplet relating between the frames. In *Link-Graph*'s terms, it means that a *Link* action between two frames is feasible only if the edges representing these frames are connected through a node in the graph (implying that the state of the triplet relates between them is known). For instance, consider a QRM represented by the *Link-Graph* shown in Figure 2.2.

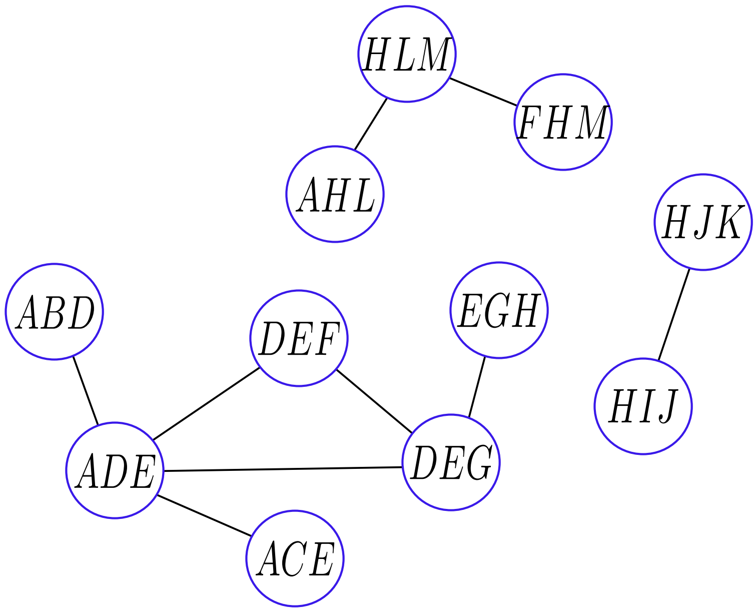

Figure 2.2: An example of a *Link-Graph*. An edge exists between every two nodes representing triplets that share exactly two landmarks in common. For example, ABD and ADE nodes are connected since A and D are common landmarks.

Based on this QRM, a *Link* from AD to AE is feasible via the triplet ADE. In contrast, a *Link* from AD to DF is not feasible, as the triplet ADF is not represented in the graph. We extend Lemma 2.2.1 to *Link-Graph*'s terms as well:

Lemma 2.3.2. *A direct* Link *from* F_1 *to* F_2 *is feasible based on a triplet* τ*, if* $F_i \subseteq \tau, \forall i \in \{1, 2\}$*, or, in* ***Link-Graph's terms****, if the edges representing* F_1 *and* F_2 *are connected to the node representing* τ*.*

One can further conclude from Lemma 2.3.2 that a *Link-Graph*'s path encodes a feasible sequence of link actions, where the edges along the path are the different frames, and the in-between nodes are the triplets the robot relies on to execute the *Link*s.

In this research, we use link graphs for explanatory purposes, as their topological structure reflects Lemma 2.2.1 (and its follow-up conclusion) visually.

2.4 Compositions - Spatial Data Propagation

The composition operator, first suggested in 1992 by [3], propagates data from two overlapping source triplets, τ_1=AB:C and τ_2=BC:D, to infer the third one τ_3=AB:D. [12] contributed a new composition operator, suitable for triplets, formulated as a look-up table, as illustrated in Fig. 2.3.

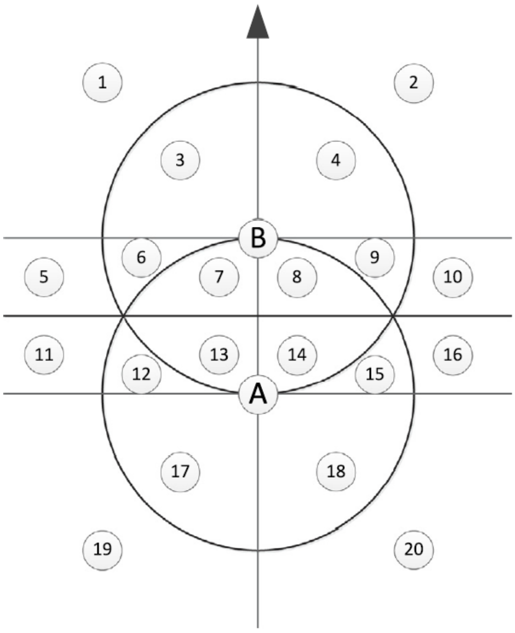

(a)

	1	2	3	4	5	$\cdots$	20
1					{1,5,11, 12,17,19}	$\cdots$	
2						$\cdots$	
3						$\cdots$	
4						$\cdots$	
5	{1,5,11, 12,17,19}				{12,17, 18,20}	$\cdots$	
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\ddots$	
20							

(b)

Figure 2.3: (a) The Extended Double Cross (EDC) partitioning, with enumerated qualitative states. (b) The composition operator as a look-up table, specifying the feasible qualitative states of the target triplet, given the states of the pair of source ones. For example, given the states AB:C=1 and BC:D=5, feasible state of AB:D are 1,5,11,12,17,19.

From a topological point of view, compositions must respect the following lemma:

Lemma 2.4.1. *A triplet τ can be composed using a single composition operation (or directly) based on the triplets τ_1 and τ_2, if the following hold:*

1. $|\tau_1 \cap \tau_2| = 2$
2. $\tau \subset \tau_1 \cup \tau_2$

For example, as was demonstrated above, we can directly compose AB:D using AB:C and BC:D. The mutual landmarks, B and C, allow us to fix both triplets relative to the same frame and, hence, infer (compose) relationships between new triplets' combinations.

[15] formulated the composition as a probabilistic operator, yielding a posterior distribution over the resulted state:

$$\mathbb{P}(\mathcal{S}^{\tau_3}|\mathcal{H}^1,\mathcal{H}^2) = \sum_{\mathcal{S}^{\tau_1}} \sum_{\mathcal{S}^{\tau_2}} \mathbb{P}(\mathcal{S}^{\tau_1}|\mathcal{H}^1)\mathbb{P}(\mathcal{S}^{\tau_2}|\mathcal{H}^2) \cdot \iint_{\mathcal{X}^{\tau_1} \in \mathcal{S}^{\tau_1} \mathcal{X}^{\tau_2} \in \mathcal{S}^{\tau_2}} \mathbb{P}(\mathcal{S}^{\tau_3}|\mathcal{X}^{\tau_1},\mathcal{X}^{\tau_2})\, d\mathcal{X}^{\tau_1}\, d\mathcal{X}^{\tau_2}, \tag{2.4}$$

where $\mathbb{P}(\mathcal{S}^{\tau_3}|\mathcal{X}^{\tau_1},\mathcal{X}^{\tau_2})$ is a simple deterministic geometric model.

In this research, we integrate compositions within our qualitative BSP approach, considering the formulation from Eq. (2.4). Additionally, we address two main theoretical issues concerning compositions calculi, stemming from Lemma. 2.4.1.

Chapter 3

Incorporating Compositions in Qualitative Approaches

3.1 Problem Statement

We consider a robot operating in a 2D environment, consisting of a known set of landmarks, denoted by $\mathcal{L}$. The robot is entrusted with a valuable task, such as planning its next course of action to reach a desirable destination. It maintains a map consisting of Qualitative Spatial Relationships (QSRs) between triplets of landmarks. For each triplet, a target landmark is localized relative to a local landmark-centric frame based on the other two. A predefined partition divides the space into a finite set of qualitative states (Fig. 1.4), and the target landmark is associated with one of them, using methods such as in [14] and [15]. We stress that these methods evaluate the triplets while assuming data association is solved. Accordingly, we assume the same in this work.

Let $\mathcal{T}$ denote the set of triplets evaluated by the robot. In the following, we shall refer to $\mathcal{T}$ as the set of source triplets. We consider only triplets in this work, as this is the most basic and standard case, and since richer QSRs can always be broken down into ternary ones. We exclude binary QSRs, as these are relevant when dealing with complex volumed landmarks (extended landmarks, see [11] and [19]), while we are assuming point landmarks in order to be aligned with the majority of works in this field.

As part of the robot's task, evaluating new triplets via composition may be helpful. Topologically speaking (recall Lemma. 2.4.1), to compose a target triplet, using a single composition operation, we need a pair of landmark triplets, such that their intersection size equals 2 (namely, they share two landmarks in common), and their union contains all three landmarks that form the target triplet (see [12]). The mutual landmarks allow us to fix both triplets relative to the same frame and hence to infer (compose) relationships between new combinations of triplets (consisting of landmarks taken from the union). For example (also given in Sec. 2.4), given the triplets $AB{:}C$ and $BC{:}D$, where the landmark after the colon is the target one, we can compose $AB{:}D$. Similarly, using AC:B and CB:D, we can compose AC:D. Note that all ordered permutations of a given triplet are available using additional unary manipulations, via the operators "INVERSE", "LEFT", and "RIGHT" defined in [12]. Consequently, given the triplets ABC and BCD (and by omitting the ':' notation, we mean, no matter the permutation), we can compose either ABD or ACD (all possible permutations). Hence, from this point on, we shall omit the colon notation.

One difficulty that may arise when performing compositions stems from recursive aspects. Consider a target triplet to compose τ. The triplets τ_1 and τ_2 required to perform the composition according to Lemma 2.4.1 may not be available (i.e., not in $\mathcal{T}$), and therefore, we would have to compose them first. This issue may occur repeatedly. Consequently, to compose τ, we need to perform a concrete sequence of composition operations, relying on source triplets from $\mathcal{T}$, that creates τ as a final outcome.

We aim to address two main issues arising from the above. The first refers to the initial set of source triplets. Intuitively, given only a sparse set of sources $\mathcal{T}$, we can compose only some triplets in $\mathcal{L}$ or worse, none. For that reason, we first identify and formulate a sufficient topological condition required for $\mathcal{T}$ so that any desired triplet in $\mathcal{L}$ would be feasible to compose (Sec. 3.2.1). The second issue is that there are often multiple alternatives for generating a sequence of composition operations to create a target triplet. We propose an algorithm that generates the optimal one in terms of a predefined cost (Sec. 3.2.2).

3.2 Approach

3.2.1 Composable Set of Triplets

In this section, we establish the term of a *Composable* set of triplets. We prove that we can compose any target triplet in the relevant landmark space given a source set $\mathcal{T}$ of this particular topological form.

As a preliminary step, we first define the following auxiliary terms:

Definition 3.2.1. Let $\mathcal{T}$ be a set of triplets. The *Landmark Space* of $\mathcal{T}$, denoted by $\mathcal{L}(\mathcal{T})$, is defined as:

$$\mathcal{L}(\mathcal{T})=\bigcup_{\tau\in\mathcal{T}}\tau \tag{3.1}$$

Note that the *Landmark Space* of a single triplet set is the triplet itself: $\mathcal{L}(\{\tau\})=\tau$.

Definition 3.2.2. Let $\mathcal{T}$ be a set of triplets. A *Cut* $C=(\mathcal{T}_L,\mathcal{T}_R)$ of $\mathcal{T}$, is a partition of $\mathcal{T}$ into two disjoint subsets, $\mathcal{T}_L$ and $\mathcal{T}_R$, s.t. $\forall\tau\in\mathcal{T}$, either $\tau\in\mathcal{T}_L$ or $\tau\in\mathcal{T}_R$, but not both.

Definition 3.2.3. Let $\mathcal{T}$ be a set of triplets and let $\alpha\in\mathbb{N}\cup\{0\}$. A *Cut* $C=(\mathcal{T}_L,\mathcal{T}_R)$ of $\mathcal{T}$ is called *α-common* if $|\mathcal{L}(\mathcal{T}_L)\cap\mathcal{L}(\mathcal{T}_R)|\geq\alpha$.

We are now ready to define the term of a *Composable* set of triplets.

Definition 3.2.4. Let $\mathcal{T}$ be a set of triplets and let $\mathcal{L}$ be a *Landmark Space*. We say that $\mathcal{T}$ is *Composable* under $\mathcal{L}$, if $\mathcal{L} \subseteq \mathcal{L}(\mathcal{T})$, and one of the following holds:

1. $|\mathcal{T}|$=1.

2. $|\mathcal{T}|$>1 and **there is** a 2-*common Cut* C=($\mathcal{T}_L, \mathcal{T}_R$) of $\mathcal{T}$, s.t. $\mathcal{T}_L$ is *Composable* under $\mathcal{L}(\mathcal{T}_L)$ and $\mathcal{T}_R$ is *Composable* under $\mathcal{L}(\mathcal{T}_R)$.

Fig. 3.1 demonstrates the above definitions.

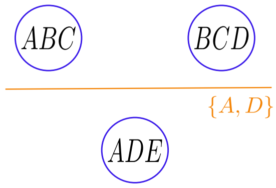

Figure 3.1: Demonstration of definitions 3.2.1-3.2.4. An example for a set of triplets, $\mathcal{T}$=$\{ABC, BCD, ADE\}$, is presented. The *Landmark Space* of $\mathcal{T}$ is $\mathcal{L}(\mathcal{T})$=$\{A, B, C, D, E\}$. The line represents the *Cut* C=($\mathcal{T}_L$=$\{ABC, BCD\}$, $\mathcal{T}_R$=$\{ADE\}$) of $\mathcal{T}$. Both landmarks A and D are common to the *Landmark Spaces* $\mathcal{L}(\mathcal{T}_L)$ and $\mathcal{L}(\mathcal{T}_R)$. Thus, C is 2-*common*. Note that C is also 1-*common* and 0-*common* by definition (but not 3-*common*). Finally, $\mathcal{T}$ is a *Composable* set under $\mathcal{L}(\mathcal{T})$ (or any $\mathcal{L} \subseteq \mathcal{L}(\mathcal{T})$), since C is 2-*common* such that: (1) $\mathcal{T}_L$ is *Composable* under $\mathcal{L}(\mathcal{T}_L)$, since the only non-trivial *Cut* C'=($\mathcal{T}_L'$=$\{ABC\}$, $\mathcal{T}_R'$=$\{BCD\}$) is 2-*common* and both $\mathcal{T}_L'$ and $\mathcal{T}_R'$ are *Composable* under $L(\mathcal{T}_L')$ and $\mathcal{L}(\mathcal{T}_R')$, respectively ($|\mathcal{T}_L'|$=$|\mathcal{T}_R'|$=1). (2) $\mathcal{T}_R$ is *Composable* under $\mathcal{L}(\mathcal{T}_R)$, since $|\mathcal{T}_R|$=1.

We now aim to prove that we can compose any triplet $\tau \subseteq \mathcal{L}$, given a *Composable* set under $\mathcal{L}$.

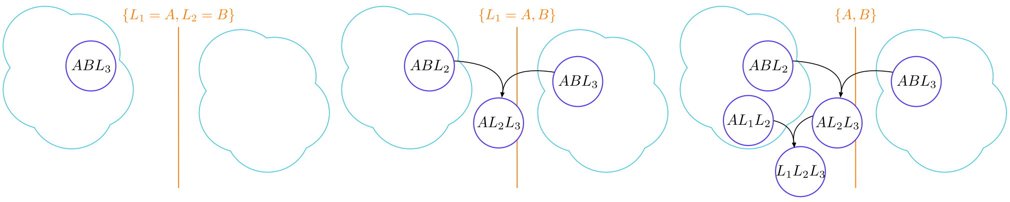

(a) (b) (c)

Figure 3.2: An illustration for the three cases described in Theorem 3.1's proof. (a) case 1; (b) case 2; (c) case 3; Triplets are shown inside circles. The line represents the *Cut* that splits the original set, $\mathcal{T}$, into two disjoint ones, $\mathcal{T}_L$ and $\mathcal{T}_R$ (represented by the clouds), with the {text} specifies landmarks that $\mathcal{L}(\mathcal{T}_L)$ and $\mathcal{L}(\mathcal{T}_R)$ share in common. Finally, a pair of triplets that points on a third one via black arrow represent that the latter is directly composed using the first two.

Theorem 3.1. *Let* $\mathcal{T}$ *be a* Composable *set of triplets under the* Landmark Space $\mathcal{L}$. *Then any triplet* $\tau \subseteq \mathcal{L}$ *can be composed based on triplets from* $\mathcal{T}$.

Proof.. We prove Theorem 3.1 using induction on number of set elements (triplets), $|\mathcal{T}|$. **Base step**: Suppose $|\mathcal{T}|$=2 ($|\mathcal{T}|$=1 is a trivial case). $\mathcal{T}$ is *Composable* under $\mathcal{L}$, thus, the only non-trivial *Cut* exists in this case is 2-*common*. Without loss of generality (WLOG), suppose $\mathcal{T}$=$\{ABC$,$BCD\}$. Indeed, according to Lemma 2.4.1, we can compose ABD and ACD, i.e., all other triplets exist in $\mathcal{L}(\mathcal{T})$ (and thus also in $\mathcal{L}$, since $\mathcal{L}\subseteq\mathcal{L}(\mathcal{T})$).

Induction step: Suppose any triplet $\tau\subseteq\mathcal{L}$ can be composed based on triples from $\mathcal{T}$, for all $1\leq|\mathcal{T}|\leq n$. We prove that the same is true for $|\mathcal{T}|=n+1$.

Suppose $|\mathcal{T}|=n+1$ and let $\tau=L_1L_2L_3$ be a triplet in $\mathcal{L}$. We show that τ can be composed using triplets from $\mathcal{T}$. Since $\mathcal{T}$ is *Composable* under $\mathcal{L}$, we are guaranteed that it has a 2-*common Cut*, ($\mathcal{T}_L$,$\mathcal{T}_R$), s.t. $\mathcal{T}_L$ is *Composable* under $\mathcal{L}(\mathcal{T}_L)$ and $\mathcal{T}_R$ is *Composable* under $\mathcal{L}(\mathcal{T}_R)$.

Suppose, WLOG, that $\{A,B\}\subseteq\mathcal{L}(\mathcal{T}_L)\cap\mathcal{L}(\mathcal{T}_R)$. We examine three possible cases (see illustration in Fig. 3.2).

<u>Case 1</u> : $|\{L_1,L_2,L_3\}\cap\{A,B\}|=2$. WLOG, we assume that $L_1=A$ and $L_2=B$ and continue examining L_3. The latter must be in $\mathcal{L}(\mathcal{T}_L)$, or $\mathcal{L}(\mathcal{T}_R)$, or both. WLOG, suppose $L_3\in\mathcal{L}(\mathcal{T}_L)$. Thus, we are guaranteed that $\{A,B,L_3\}\subseteq\mathcal{L}(\mathcal{T}_L)$, and consequently ABL_3 (namely, $L_1L_2L_3$) can be composed according to the assumption since $\mathcal{T}_L$ is *Composable* under $\mathcal{L}(\mathcal{T}_L)$ and $|\mathcal{T}_L|\leq n$.

<u>Case 2</u>:$|\{L_1,L_2,L_3\}\cap\{A,B\}|=1$. WLOG, we assume that $L_1=A$ and continue examining L_2,L_3. If they are both in $\mathcal{L}(\mathcal{T}_L)$ or both in $\mathcal{L}(\mathcal{T}_R)$, we finished (similarly to case 1). Otherwise, WLOG, we assume that L_2 is exclusively in $\mathcal{L}(\mathcal{T}_L)$ and L_3 is exclusively in $\mathcal{L}(\mathcal{T}_R)$. According to the assumption, we are guaranteed that ABL_2 and ABL_3 can be composed based on $\mathcal{T}_L$ and $\mathcal{T}_R$, respectively. Finally, using these two triplets, we can compose AL_2L_3 (Lemma 2.4.1), namely, $L_1L_2L_3$.

<u>Case 3</u> : $|\{L_1,L_2,L_3\}\cap\{A,B\}|=0$. If $\{L_1,L_2,L_3\}$ are all in $\mathcal{L}(\mathcal{T}_L)$ or all in $\mathcal{L}(\mathcal{T}_R)$, we finished (similarly to case 1). Otherwise, WLOG, we assume that L_1 and L_2 are exclusively in $\mathcal{L}(\mathcal{T}_L)$ and L_3 is exclusively in $\mathcal{L}(\mathcal{T}_R)$. According to the assumption, we are guaranteed that ABL_2 and AL_1L_2 can be composed based on $\mathcal{T}_L$, and that ABL_3 can be composed based on $\mathcal{T}_R$. Using ABL_2 and ABL_3, we can compose AL_2L_3 (Lemma 2.4.1).

Finally, using AL_1L_2 and AL_2L_3, we can compose $L_1L_2L_3$(Lemma 2.4.1). ∎

3.2.2 Optimal Composition Sequence

As we saw in the previous section, given a *Composable* set of source triplets under the *Landmark Space* $\mathcal{L}$, we can compose any triplet in $\mathcal{L}$. Once we chose a target triplet to compose, we aim to find an appropriate sequence of composition operations to create it, considering the recursive aspects discussed in Sec. 3.1. Topologically, such a sequence can be described as a binary tree representing all compositions required to create a target triplet. Each tree node represents a triplet to compose, with its child nodes representing the triplets pair required to composed it (directly). The root specifies the target triplet, while the leaves specify source triplets taken from $\mathcal{T}$ (the initial set of source triplets which considered given). We name this binary tree a *Composition-Tree*. Its formal definition is given below:

Definition 3.2.5. A *Composition-Tree* is a binary tree, $T{=}(V,E)$, where:

1. Each node $v_\tau{\in}V$ represents a triplet of landmarks τ.
2. Each node $v_\tau{\in}V$ has either two children, $v_{\tau_L},v_{\tau_R}{\in}V$, representing that τ is composed directly using τ_L and τ_R, or none, if τ is a source triplet (in which case, v_τ is a leaf).

See Figs. 3.7e-3.7f for illustration. In general, many topologically suitable trees may exist for a target triplet. So it is still a question, how to select the best one among all possibilities.

Formally, consider the set of source triplets $\mathcal{T}$, and a target triplet to compose, τ_o. We denote by $\mathbb{T}_{\tau_o}$ the set of all possible *Composition-Tree*s with τ_o being the root and leaves taken from $\mathcal{T}$. We aim to find the optimal *Composition-Tree*, considering all instances from $\mathbb{T}_{\tau_o}$, in terms of a predefined cost C (see cost specifications in Sec. 3.2.2). That is:

$$T^{*}{=}\underset{T\in\mathbb{T}_{\tau_o}}{\arg\min}\sum_{\tau\in T}C(\tau). \tag{3.2}$$

We propose a three-stage algorithm for solving the above problem, given an initial set of source triplets, $\mathcal{T}$. First, we initialize a unique topological structure called a *Composition-Graph* (Sec. 3.2.2). The graph sets the infrastructure for the cost maintenance regarding the different triplets that exist in $\mathcal{L}(\mathcal{T})$. Next, we show how to update the graph every time a new source triplet is being considered (Sec. 3.2.2). Finally, given a query triplet to compose, $\tau_o{\subseteq}\mathcal{L}(\mathcal{T})$, we extract from the graph the *Composition-Tree* representing the optimal sequence of compositions to form τ_o. In case no valid sequence exists, the *Composition-Tree* will be empty (Sec. 3.2.2).

Alg. 1 summarizes the above. Further aspects regarding the algorithm are analyzed in Sec. 3.2.3. A comprehensive running example can be found in Sec. 3.2.4.

Algorithm 1: Optimal Composition Sequence

Input: Set of source triplets $\mathcal{T}$, target triplet $\tau_o{\subseteq}\mathcal{L}(\mathcal{T})$

▷ Initialize empty graph
$G \leftarrow$ InitCompGraph$\big(\mathcal{L}(\mathcal{T})\big)$
▷ Update Graph
for $\tau{\in}\mathcal{T}$ **do**
    $G \leftarrow$ UpdateCompGraph(G,τ)
▷ Extract τ_o optimal *Composition-Tree*
$T \leftarrow$ ExtractCompTree(G,τ_o)
return T

Figure 3.3: Optimal Composition Sequence - Algorithm description.

Composition-Graph Initialization

The *Composition-Graph* is a useful topological representation that allows propagating costs between triplets easily. Its exact structure is defined as follows:

Definition 3.2.6. A *Composition-Graph* is a bipartite graph, $G{=}(V_\tau,V_q,E)$ where:

1. Each node $v_\tau{\in}V_\tau$ represents a triplet of landmarks τ.
2. Each node $v_q{\in}V_q$ represents a sub-space of four landmarks (a quartet) q.
3. There is an edge $e{=}(v_\tau,v_q){\in}E$, where $v_\tau{\in}V_\tau$ and $v_q{\in}V_q$, if and only if $\tau{\subset}q$.

The above representation is motivated by Lemma 2.4.1, where each quartet node connects exactly four triplets nodes with a direct composition relationship. Meaning, each pair of triplets taken from this foursome can be used to compose one of the remaining two, using a single composition. Alg. 2 shows how to construct and initialize a new *Composition-Graph*, considering a given *Landmark Space* $\mathcal{L}$. An illustration for the *Composition-Graph* initialization can be found in Fig. 3.7a. Note that we initialize a complete *Composition-Graph*, that is, it contains all possible $\binom{|\mathcal{L}|}{3}$ triplet and $\binom{|\mathcal{L}|}{4}$ quadrants that exist in $\mathcal{L}$, since potentially we aim to use it to compose any triplet in $\mathcal{L}$. Furthermore, each triplet node v_τ is initialized with a cumulative cost value (denoted by d) of ∞, which is later updated considering the arrival of new information regarding source triplets. The cumulative cost represents the total cost of the current optimal *Composition-Tree* of τ embedded in the graph. In addition, each v_τ is initialized with an empty parents list, later updated with the nodes representing the triplets designated to compose τ directly. Note that the parents (in terms of the *Composition-Graph*) are later referred to as children (in terms of the extracted *Composition-Tree*).

Algorithm 2: Initialize Composition Graph

Function INITCOMPGRAPH(*Landmark space* $\mathcal{L}$):
    ▷ Initialize empty graph
    $G \leftarrow \{V_\tau \leftarrow \emptyset, V_\mathcal{Q} \leftarrow \emptyset, E \leftarrow \emptyset\}$
    ▷ Initialize triplets nodes
    $\mathcal{T}_\mathcal{L} \leftarrow$ Set of all $\binom{|\mathcal{L}|}{3}$ triplets residing in $\mathcal{L}$
    for $\tau{\in}\mathcal{T}_\mathcal{L}$ **do**
        Add node v_τ to V_τ representing τ
        $d(v_\tau) \leftarrow \infty$ ▷ Cumulative Cost
        Parents$(v_\tau) \leftarrow \emptyset$
    ▷ Initialize quartets nodes and edges
    $\mathcal{Q}_\mathcal{L} \leftarrow$ Set of all $\binom{|\mathcal{L}|}{4}$ quartets residing in $\mathcal{L}$
    for $q{\in}\mathcal{Q}_\mathcal{L}$ **do**
        Add node v_q to $V_\mathcal{Q}$
        $\{v_{\tau_1},v_{\tau_2},v_{\tau_3},v_{\tau_4}\} \leftarrow$ nodes from V_τ representing triplets residing in q
        for $v{\in}\{v_{\tau_1},v_{\tau_2},v_{\tau_3},v_{\tau_4}\}$ **do**
            Add edge (v_q,v) to E
    return G

Figure 3.4: Initialize *Composition-Graph* - Algorithm description.

Composition-Graph Update

Given a source triplet $\tau \in \mathcal{T}$, we aim to properly update our *Composition-Graph* by updating the cumulative cost of v_τ and all affected triplets nodes. To that end, we define the following cost:

$$C(\tilde{\tau}) = \begin{cases} C^{\text{source}}(\tilde{\tau}), \text{ if } \tilde{\tau} \text{ is a source triplet} \\ C^{\text{comp}}(\tilde{\tau}_L, \tilde{\tau}_R), \text{ if } \tilde{\tau} \text{ is composed directly using } \tilde{\tau}_L, \tilde{\tau}_R \end{cases}$$

, where C^{source} can be any real function and where C^{comp} can be any symmetric non-negative real function. The cost is designed by the user to consider a desired criterion of interest. A specific selection example is given in Sec. 3.2.4.

We shall now explain in detail how we update the *Composition-Graph*, considering the arrival of new information regarding the source triplet τ.

First, we calculate $c = C^{\text{source}}(\tau)$, the new candidate cumulative cost of τ resulted from the newly arrived information. We then locate the node representing τ, v_τ, and check whether c is better (lower) than the cumulative cost currently assigned to it, $d(v_\tau)$. If so, we update v_τ with the new cost c and reset its parent list, stating that τ is qualified to be used as a source triplet.

In case of an update, we continue examining all quadrant nodes adjacent to v_τ. For each adjacent node v_q, we look at the other three triplets nodes connected to it. Let v_{τ_1},v_{τ_2}, and v_{τ_3} denote these nodes. We check for each one if we can improve its cumulative cost as a result of the recently updated one of v_τ. Consider v_{τ_1} for example, we compare its current cumulative cost, $d(v_{\tau_1})$, with the candidate ones $c_i = d(v_\tau) + d(v_{\tau_i}) + C^{\text{comp}}(\tau, \tau_i), \forall i \in \{2,3\}$. If for some $i \in \{2,3\}$ c_i improves $d(v_{\tau_1})$, we assign it to v_{τ_1} and update its parents to be v_τ and v_{τ_i}, stating that τ_1 is designated to be composed directly using v_τ and v_{τ_i}. For each newly updated triplet node, we go through the same process until there are no more updates.

The *Composition-Graph* is updated incrementally, each time based on a single source triplet from $\mathcal{T}$. Note that since the update rule obeys Lemma 2.4.1 topologically, Theorem 3.1 determines which nodes are expected to be updated based on the topological structure of $\mathcal{T}$ alone. That is, if $\mathcal{T}$ is *Composable* under the *Landmark Space* $\mathcal{L}$, then each node representing a triplet τ in $\mathcal{L}$ is guaranteed to be updated with a finite cost value. The latter means that we can extract a non-empty *Composition-Tree* describing the optimal sequence of composition operations required to compose τ. Of course, even when all the triplets nodes are assigned with finite costs, extra updates can only lead to further improvement.

Alg. 3 formalizes the above.

Algorithm 3: Composition Graph Update

Function UPDATECOMPGRAPH (*Composition Graph G, triplet to update τ*)**:**
- $\mathcal{Q} \leftarrow \emptyset$ ▷ Initialize empty set
- $v_\tau \leftarrow$ node from G representing τ
- ▷ Update v_τ
- c=$C^{\text{source}}(\tau)$
- **if** c<$d(v_\tau)$ **then**
  - $d(v_\tau) \leftarrow c$
  - $Parents(v_\tau) \leftarrow \emptyset$
  - $\mathcal{Q} \leftarrow \mathcal{Q} \cup v_\tau$
- ▷ Update graph (propagate)
- **while** $\mathcal{Q}$ *not empty* **do**
  - $v \leftarrow$ select and remove $v \in \arg\min_{v \in Q} d(v)$
  - **for** $v_q \in Adj(v)$ **do**
    - $V_\tau^{\text{Updated}} \leftarrow$ UPDATESTEP(G,v_q, v)
    - $\mathcal{Q} \leftarrow \mathcal{Q} \cup V_\tau^{\text{Updated}}$
- **return** G

Function UPDATESTEP (*Composition Graph G, Quartet node v_q, updated triplet node v_τ*)**:**
- $V_\tau^{\text{Updated}} \leftarrow \emptyset$
- $V_\tau^{\text{adj}} \leftarrow Adj(v_q) \backslash \{v_\tau\}$
- ▷ Propagate cost in case of improvement
- **for** $v \in V_\tau^{adj}$ **do**
  - **for** $\tilde{v} \in V_\tau^{adj} \backslash \{v\}$ **do**
    - $v_{\tau_L} \leftarrow v_\tau$,$v_{\tau_R} \leftarrow \tilde{v}$
    - $c \leftarrow d(v_{\tau_L})$+$d(v_{\tau_R})$+$C^{\text{Comp}}(\tau_L,\tau_R)$
    - **if** c<$d(v)$ **then**
      - $d(v) \leftarrow c$
      - $Parents(v) \leftarrow \{v_{\tau_L},v_{\tau_R}\}$
      - $V_\tau^{\text{Updated}} \leftarrow V_\tau^{\text{Updated}} \cup \{v\}$
- **return** V_τ^{Updated}

Figure 3.5: Composition Graph Update - Algorithm description.

Composition-Tree Extraction

Given an updated *Composition-Graph* and a finite-cost node v_{τ_o}, we aim to extract the optimal *Composition-Tree* of the corresponding triplet τ_o (Eq. 3.2).

Alg. 3 assigns to each node its optimal parents in terms of yielded cumulative cost. The latter converted to be the node's children in the embedded *Composition-Tree*. Accordingly, all we have left is to recursively extract the *Composition-Tree* embedded in the graph, starting from the root v_{τ_o}. We stop when we reach nodes representing source triplets, which are reflected as the tree leaves. Alg. 4 sums the recursive extraction process.

Algorithm 4: Extract *Composition-Tree*

Function EXTRACTCOMPTREE (*Composition Graph G, target triplet τ_o*)**:**
- Initialize empty tree T
- ▷ Stopping Criterion
- $v_{\tau_o} \leftarrow$ node from G representing τ_o
- **if** $d(v_{\tau_o})$=∞ **then**
  - **return** T
- ▷ Construct tree
- $Root(T) \leftarrow v_{\tau_o}$
- $\{v_{\tau_L},v_{\tau_R}\} \leftarrow Parents(v_{\tau_o})$
- $Left(T) \leftarrow$ EXTRACTCOMPTREE(G,τ_L)
- $Right(T) \leftarrow$ EXTRACTCOMPTREE(G,τ_R)
- **return** T

Figure 3.6: Extract *Composition-Tree* - Algorithm description.

3.2.3 Algorithm 3.3 Analysis

In this section, we analyze the correctness of Alg. 1 and explain why its convergence is guaranteed. In addition, we provide complexity analysis regarding time and place.

To analyze our algorithm, we shall first identify that the update step (Alg. 3) is an instance of a generalized version of *Dijkstra* [20] for directed *Hyper-Graph*s, called *SBT-Dijkstra* [21]. Recall that a directed *Hyper-Graph* is a pair G=(V,E), where $V = \{v_1,v_2,\ldots,v_n\}$ is the set of nodes and E=$\{e_1,e_2,\ldots,e_m\}$, with , $e_i = (T(e_i){\subseteq}V, H(e_i) \subseteq V), \forall i = 1, 2, \ldots, m$, is the set of directed *hyperedge*s. T and H known as the *tail* and *head* of the *hyperedge*. Note that the above definition generalizes the standard directed graph (if $|T(e_i)|$=$|H(e_i)|$=1, $\forall i$=1,2, $\ldots$,m, we get a standard directed graph). Alternatively, we could have defined the *Composition-Graph* as a directed *Hyper-Graph*, where each directed *hyperedge* connects a foursome of triplets having a direct composition relationship. Specifically, two out of the four triplets form the *tail*, whereas the remaining two form the *head*, stating that each *head* triplet can be composed directly using the pair of *tail* ones. Consequently, we would replace each quartet node in the original *Composition-Graph* with a set of 6 directed *hyperedge*s, since we have $\binom{4}{2}$=6 ways of choosing pair of tail triplets (the head is dictated given that choice). Moreover, accordingly, the cost C^{comp} would now be the weight of the *hyperedge*s.

Correctness. The correctness of each update step is guaranteed based on the correctness of the *SBT-Dijkstra* procedure. Similarly to the original *Dijkstra* algorithm, it can be proven using induction over the number of visited nodes (i.e., nodes pulled out from $\mathcal{Q}$). As explained in [21], a crucial key point to be ensured is that no negative cycle is detected during the algorithm's operations, as it might lead to cyclic costs improvements of nodes in the graph. Indeed, in our case, all cycles are nondecreasing, as node cumulative costs are being aggregated additively and since C^{comp} is nonnegative. While the original *Dijkstra* detects shortest paths, the *SBT-Dijkstra* detects shortest *hyperpath*s (i.e., sequences of nodes and *hyperedge*s). In our case, the *hyperpath* is equivalent to an embedded *Composition-Tree* branch. It is important to note that applying the update step multiple times, as we do in Alg. 1, does not affect its correctness (the same induction proof mentioned above still holds). As long as we keep the cumulative costs from previous update steps, Alg. 1 is guaranteed to terminate with the optimal embedded branches, considering all source triples (multiple sources shortest paths). For that reason, we only initialize the cumulative cost of each triplet node once, during Alg. 1. Consequently, at the end of all update steps, any finite cost node can be reversely opened into the shortest *Composition-Tree* rooted at that node.

Considering the above, and given a target triplet to compose, τ_o, we are guaranteed that Alg. 4 extracts the shortest *Composition-Tree* originating in v_{τ_o}, i.e., $T^*{=}\arg\min_{T\in\mathbb{T}_{\tau_o}} d(v_{\tau_o})$, where $\mathbb{T}_{\tau_o}$ is the set of all possible *Composition-Tree*s with v_{τ_o} as root. Assuming v_{τ_L} and v_{τ_R} are the parents (children in terms of the embedded

Composition-Tree) assigned to v_{τ_o}, we can rewrite the above as $T^*{=}\arg\min_{T\in\mathbb{T}_{\tau_o}} d(v_{\tau_L}){+}d(v_{\tau_R}){+}C(v_{\tau_o})$. We can repeat this recursively until we reach nodes representing source triplets, which have no parents (*Composition-Tree* leaves). We get $T^* = \arg\min_{T\in\mathbb{T}_{\tau_o}} \sum_{\tau\in T} C(\tau)$ (Eq. 3.2).

Complexity. Time complexity: *Composition-Graph* initialization is $\mathcal{O}(|V_\tau|+|V_q|+|E|)$. Then, each update step is $\mathcal{O}(\max\{|V_\tau| + |V_q| + |E|, |V_\tau|^2\})$, as the total cost of node selection and removals from Q is $\mathcal{O}(|V_\tau|^2)$ and the total cost of propagating costs in the graph is $\mathcal{O}(|V_\tau|+|V_q|+|E|)$ (based on [21]). Assuming a set of k source triplets to update the *Composition-Graph* with, we get $\mathcal{O}\big(\max\{|V_\tau| + |V_q| + |E|, |V_\tau|^2\}\cdot k\big)$. The extraction stage is $\mathcal{O}(|V_\tau|)$ at most, in case the target *Composition-Tree* consists of all the nodes in V_τ. In total, we have $\mathcal{O}\big(\max\{|V_\tau| + |V_q| + |E|, |V_\tau|^2\}\cdot k\big)$.

Space complexity: here, the initialization is the bottleneck. We have $\mathcal{O}(|V_\tau|+|V_q|+|E|)$.

3.2.4 Running Example

We shall now demonstrate Alg. 3.3 via a running example. Consider the initial set of source triplets, $\mathcal{T} = \{ABC, BCD, ADE\}$, and suppose we aim to compose the triplets ACE and BCE, using minimum composition operations, based on triplets from $\mathcal{T}$. Note that Theorem 3.1 guarantees that the above can be done since $\mathcal{T}$ is *Composable* under the *Landmark Space* $\mathcal{L}{=}\{A,B,C,D,E\}$, which contains both target triplets. We show how we infer the optimal (i.e., minimal) sequence of composition operations to create these triplets.

First, we formulate the following cost function:

$$C(\tau){=}\begin{cases} 0, & \text{if } \tau \text{ is a source triplet} \\ 1, & \text{if } \tau \text{ is composed directly using } \tau_L{,}\tau_R, \end{cases}$$

i.e., a triplet τ is assigned with the cost value of 0 if it is a source triplet (since we need 0 composition operations to form it), and with 1 if it is designated to be composed directly using the pair τ_L and τ_R (since we need a single composition to create τ, given τ_L and τ_R).

We start by initializing an appropriate *Composition-Graph* (Alg. 3.4), with $\mathcal{L}$ as the input *Landmark Space* (Fig. 3.7a). We then perform the update step (Alg. 3.5) for each source triplet in $\mathcal{T}$. The outcome graphs after update steps for ABC, BCD and ADE are illustrated in Fig. 3.7b-3.7d, respectively. At the end of these three update steps, every triplet node in the graph is updated with a finite cumulative cost, which indicates (due to the particular cost function we chose) the minimal amount of composition operations required to form it. As Fig. 3.7d shows, two compositions are required to create ACE and three to create BCE.

Finally, we extract the optimal (minimal) sequence of compositions, each represented via a *Composition-Tree*, to create these target triplets (Alg. 3.6). Fig. 3.7e and

Fig. 3.7f show the extracted trees for the target triplets. Recall that while Alg. 3.6 extracts the optimal *Composition-Tree* for a query triplet, many other non-optimal ones exist. For instance, BCE can also be composed directly using BDE and CDE, but such a choice would require at least five composition operations in total since both BDE and CDE require two composition operations each in the optimal case.

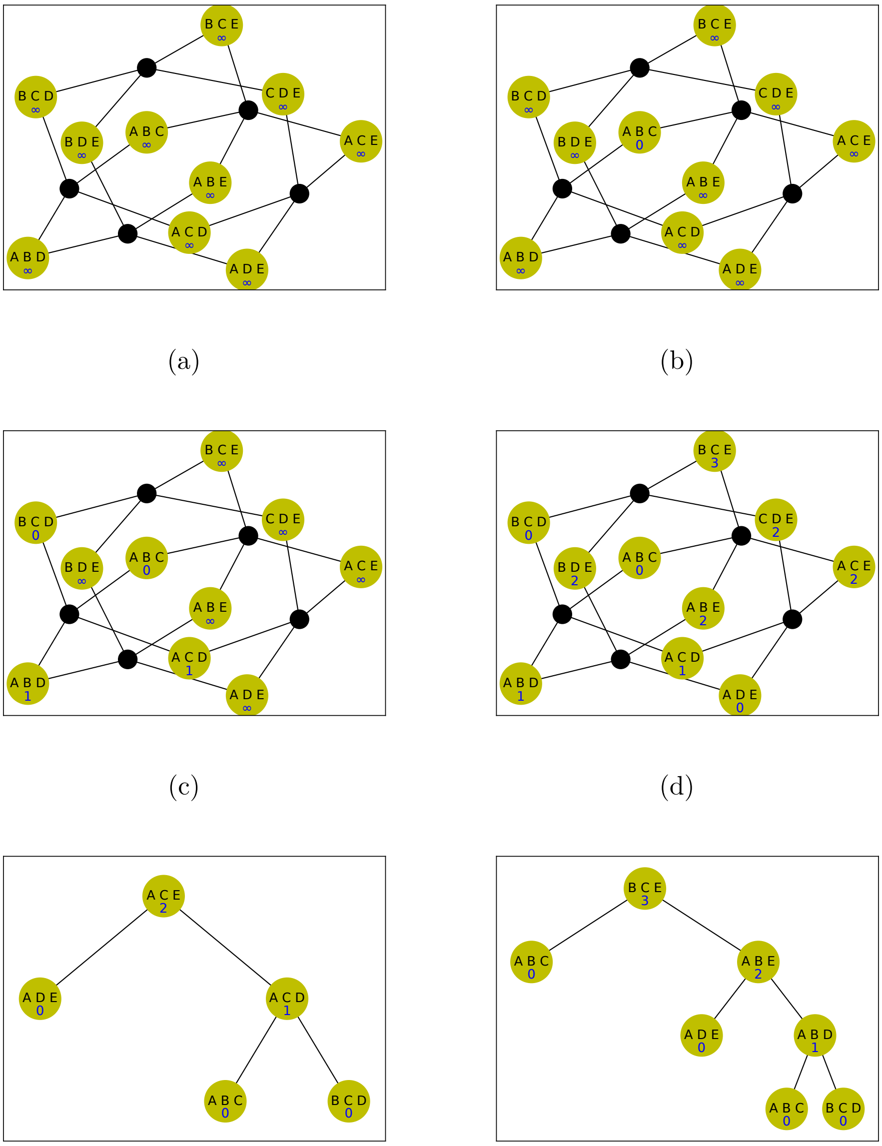

(a) (b)

(c) (d)

(e) (f)

Figure 3.7: A running example of Alg. 3.3, as described in Section 3.2.4. (a) illustrates *Composition-Graph* initialization, as an output of Alg. 3.4, with $\mathcal{L}$=$\{A$,B,C,D,$E\}$. Triplets nodes are shown in mustard color, whereas quartet nodes are in black; (b)-(d) demonstrate the graph update (Alg. 3.5), considering the cost described in Section 3.2.4. The outcome graphs of three consecutive update steps, in which the cost of ABC, BCD, and ADE is set to 0, are shown. Since the set of triplets mentioned above is *Composable* under $\mathcal{L}$, the update of the entire graph is guaranteed. Note that after the second step (c), all four triplets exist under the sub-space $\{A$,B,C,$D\}$ are updated, as $\{ABC$,$BCD\}$ is a *Composable* set under that space. Moreover, the same is true if we considering the first step (b) alone and the *Landmark Space* $\{A$,B,$C\}$ (however, the latter is a degenerate case, as no extra triplets exist in this sub-space); (e)-(f) Two *Composition-Trees* extracted from the *Composition-Graph* (Alg. 3.6);

Chapter 4

Qualitative Belief Space Planning via Compositions

4.1 Problem Statement

4.1.1 Qualitative Belief Definition

We consider a robot operating in a partially known environment, consisting of a known set of key landmarks, denoted by $\mathbb{L}$. As it travels in space, the robot moves between different landmark-centric frames (see e.g. [12, 14]). Each frame sets a local coordinate system determined by two reference landmarks. The first fixes the origin, i.e., $(0, 0)$, whereas the second fixes an additional coordinate, in most cases, $(0, 1)$. Then, a predefined partition divides space into a finite set of non-overlapping and complementary regions, known as qualitative states (see Fig. 1.4). We denote by $\mathbb{F}$ the set of all available frames, based on landmarks from $\mathbb{L}$.

The robot maintains its self-poses and map through Qualitative Spatial Relationships (QSRs). Each QSR localizes a target point, either a landmark or a robot pose, relative to a chosen frame $F{\in}\mathbb{F}$ by associating it with one of the qualitative states discussed above. This ternary type of QSRs is often referred to as triplets. In this work, we consider only triplets, as this is the most basic and standard case, and richer QSRs can always be split into triplets. We exclude binary QSRs, as these are relevant when dealing with complex volumed landmarks (extended landmarks, see [11] and [19]), while we assume point landmarks. In the following, we denote an ordered triplet by τ, or explicitly by $F{:}L$, where F and L are the triplet's reference frame and target landmark, respectively (for example, $\tau{=}AB{:}C$). We further denote by $\mathcal{S}^{\tau}$ and $\mathcal{X}^{\tau}$, or explicitly by $\mathcal{S}_F^L$ and $\mathcal{X}_F^L$ the qualitative and metric location of L relative to F (i.e., of $\tau{=}F{:}L$). Similarly, $\mathcal{S}_F^{X_t}$ and $\mathcal{X}_F^{X_t}$ denote the robot's qualitative and metric location at time step t, both relative to F as well. Note that $\mathcal{S}_F^L$ and $\mathcal{S}_F^{X_t}$ are discrete variables, while $\mathcal{X}_F^L$ and $\mathcal{X}_F^{X_t}$ are continuous.

Apart from self-poses and triplets, the robot accounts also for *Global Frame Scales*'.

A *Global Frame Scale* is a new concept suggested in this work, which refers to the global metric distance between the two landmarks creating the frame. For example, the global frame scale of AB is the global metric distance between the landmarks A and B. We further clarify the usage of this term in Sections 4.2.3 and 4.2.7. We denote the metric and qualitative scale of a frame F by $\mathcal{X}^F$ and $\mathcal{S}^F$, respectively. Furthermore, we denote by $n_F \in \mathbb{N}$ the resolution according to which we evaluate $\mathcal{S}^F$. Given a selected value of n_F (typical values are $4-6$), S^F is a discrete random variable (RV) equals to $s \in [1, 2, \ldots, n_F]$ if $\mathcal{X}^F \in [\frac{(s-1)\cdot 2R}{n_F}, \frac{s \cdot 2R}{n_F}]$ and equal to $n_F + 1$ if $\mathcal{X}^F \in [\frac{s \cdot 2R}{n_F}, \infty)$. Here, R is the robot's sensing range given in some global frame units (for example, meters). In the case where $s^F = n_F + 1$, the frame is considered unobservable since its global scale is bigger than $2R$.

Due to the nature of the robot's mobility, we consider two types of actions. The first, referred to as *qualitative action*, allows the robot to travel between different qualitative states w.r.t. a given frame. We shall denote by a_t^{q} a *qualitative action* taken at time step t. In contrast, the second type, referred to as *link action* [16], enables the robot to switch from one local frame to another. We shall denote by a_t^{Link} a *link action* taken at time step t. The robot executes the abovementioned actions alternately. Namely, consider the robot's frame at time step t, F_t, and the corresponding robot's state $\mathcal{S}_{F_t}^{X_t}$. After executing a_t^{q}, the robot moves to a new state, $\mathcal{S}_{F_t}^{X_{t+1}}$. Then, it links to a new frame by executing the action a_t^{Link}, which results in the state $\mathcal{S}_{F_{t+1}}^{X_{t+1}}$. Note that a_t^{Link} is not a natural action in the sense that the robot makes no actual movement. In fact, the latter can equivalently be written as a tuple of source and destination frames, i.e., $a_t^{Link} \triangleq (F_t, F_{t+1})$. Note that a_t^{Link} may be degenerated, in case $F_t = F_{t+1}$. Each time the robot completes a qualitative action, it acquires a new measurement. Let z_t denote a measurement acquired by the robot at time step t. We further denote by β_t the data association from time step t, that is, the identity of the landmarks captured in z_t. Of course, determining the data association is a very challenging problem in itself. In this work, we assume it to be solved.

Consider k as the current time step. We denote by $\mathcal{H}_k$ the history of applied actions, measurements and data associations up to that time step. That is, $\mathcal{H}_k \triangleq \{a_{1:k-1}, z_{1:k}, \beta_{1:k}\}$, where a_i represents a consecutive pair of qualitative and link actions $\{a_i^{\mathrm{q}}, a_i^{Link}\}, \forall i \in \{1, \ldots, k\}$. The index $t{:}t^{'}$ compactly refers to a series of elements between time steps t and $t^{'}$. Due to the stochastic nature of the problem, the robot maintains a qualitative belief, i.e., a posterior probability over the states of the robot, landmark triplets, and frames' scales, given by:

$$b_k \triangleq \mathbb{P}(\mathcal{S}^{X_{1:k}}, \mathcal{S}^{\mathcal{M}_k}, \mathcal{S}^{\mathbb{F}_k} | \mathcal{H}_k), \tag{4.1}$$

where $\mathcal{S}^{X_{1:k}}$ represents the set of robot states $\mathcal{S}^{X_{1:k}} \triangleq \{\mathcal{S}_{F_i}^{X_i}\}_{i=1}^{k}$, $\mathcal{S}^{\mathcal{M}_k}$ represents all available landmark triplets states at time step k, $\mathcal{S}^{\mathcal{M}_k} \triangleq \{\mathcal{S}^{\tau_j}\}_{j=1}^{m_k}$, with m_k being the set size, and finally, $\mathcal{S}^{\mathbb{F}_k}$ represents all available frames' scales at time step k, i.e.,

$\mathcal{S}^{\mathbb{F}_k} \triangleq \{\mathcal{S}^{F_q}\}_{q=1}^{p_k}$, with p_k being the set size.

Inspired by [14, 16], we maintain marginals over qualitative states conditioned on only local information, prioritizing computational speedup over an accurate model. Accordingly, we maintain the belief as a product of individual posteriors:

$$b_k \approx \prod_{i=1}^{k} \mathbb{P}(\mathcal{S}_{F_i}^{X_i} | \mathcal{H}_k^{X_i}) \prod_{j=1}^{m_k} \mathbb{P}(\mathcal{S}^{\tau_j} | \mathcal{H}_k^{j}) \prod_{q=1}^{p_k} \mathbb{P}(\mathcal{S}^{F_q} | \mathcal{H}_k^{q}). \tag{4.2}$$

In the above, $\mathcal{H}_k^{X_i}$,$\mathcal{H}_k^{j}$,$\mathcal{H}_k^{q}$ denote relevant part of the history used to evaluate $\mathcal{S}_{F_i}^{X_i}$,$\mathcal{S}^{\tau_j}$, and $\mathcal{S}^{F_q}$, respectively. While Eq. (4.2) is beneficial for implementation needs, the theoretical formulations in this work are derived using Eq. (4.1) to stay as general as possible.

4.1.2 Qualitative Belief Space Planning

We now introduce a belief space planning (BSP) formulation considering the qualitative framework discussed above.

BSP, in essence, is the problem of finding an optimal sequence of actions, or policy, that minimizes a meaningful objective function. In this work, we consider actions sequences rather than policies. Assuming a future horizon of L look-ahead steps, we compactly represent by a_{k+} a candidate sequence of actions from time step k to the predefined horizon, that is, $a_{k+} \triangleq a_{k:k+L-1}$. The objective function maps the current belief, b_k, and a candidate actions sequence, a_{k+}, to an expected cumulative cost:

$$J(b_k, a_{k+}) \triangleq \mathop{\mathbb{E}}_{z_{k+1:k+L}} \left[\sum_{l=1}^{L} c_l(b_{k+l}, a_{k+l-1}) \right], \tag{4.3}$$

where c_l is the l-th cost function with the appropriate arguments, with $l \in \{1,2,\ldots,L\}$.

While the above formulation is expressed in terms of a general cost function, we choose a specific one that best serves our purposes. We elaborate on the different types of costs in the Approach section. The optimal action sequence is defined by:

$$a_{k+}^{*} = \arg\min_{a_{k+}} J(b_k, a_{k+}). \tag{4.4}$$

4.2 Approach

4.2.1 Approach Overview

We present an end-to-end algorithm to address the qualitative BSP problem defined in Sec. 4.1.2. Our algorithm operates in two steps.

First, we construct a belief tree, reflecting the future posterior beliefs considering various possible future developments. We describe the construction process and the different models assumed in this work in Secs. 4.2.2-4.2.4. We explain how to update the belief between two adjacent tree nodes in Sec. 4.2.5. In Sec. 4.2.6, we further explain how compositions can be incorporated within the updating step. We construct the tree considering a predefined depth $L \in \mathbb{N}$.

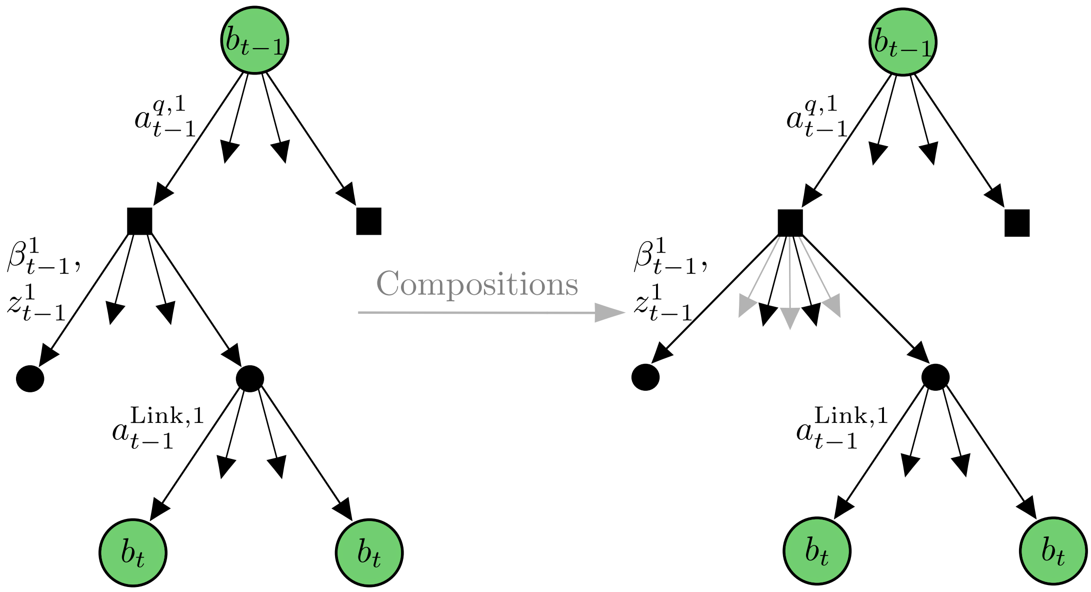

(a)

(b)

Figure 4.1: (a) A belief tree, suitable for the qualitative framework, is illustrated. Each construction step between two consecutive times $t-1$ and t required updating the belief according to Eq. (4.25). (b) The same tree is illustrated, this time with compositions incorporated within the planning process. As shown, we can consider a wider data association space via compositions, thereby improving planning performance.

In the second step, we utilize the constructed belief tree to evaluate the objective (4.7) for each candidate action sequence. Then, we search for the action sequence that minimizes the objective through a proper optimization process. In Sec. 4.2.7, we suggest two types of costs to consider while evaluating the objective.

4.2.2 Qualitative Action and Transition Model

The qualitative action enables the robot to move from one qualitative state to another, considering a specific reference frame. Mathematically, we assume a probabilistic transition model,

$$\mathbb{P}(\mathcal{S}_{F_{t-1}}^{X_t} | \mathcal{S}_{F_{t-1}}^{X_{t-1}}, a_t^{\mathrm{q}}), \tag{4.5}$$

which maps any realization of the pair $\mathcal{S}_{F_{t-1}}^{X_{t-1}}, a_t^{\mathrm{q}}$ to a Q dimensional vector which describes the outcome distribution of the new robot's state, $\mathcal{S}_{F_{t-1}}^{X_t}$. The values of each vector are chosen or learned offline according to the level of noise that characterizes the robot's control system. We consider this transition model to be available and assume that a proper low-level controller exists. Furthermore, we assume a finite set of $(Q-1)^2$ qualitative actions, where for any mutual realization $\mathcal{S}_{F_{t-1}}^{X_t}{=}i, \mathcal{S}_{F_{t-1}}^{X_{t-1}}{=}j$, where

$i{\neq}j$, there is a corresponding qualitative action that transforms the robot to state j with high probability, given that its current state is i.

4.2.3 Data Association and Measurement Likelihood

The objective requires evaluating the likelihood of capturing a measurement z_t and a corresponding data association β_t for any time step $t{\in}\{k{+}1,\ldots,k{+}L\}$. In this section, we rigorously derive the abovementioned likelihood term.

Formally, we can rewrite Eq. (4.3) recursively, as follows:

$$J(b_k,a_{k+}){=}\mathop{\mathbb{E}}_{z_{k+1}}\Big[c_1(b_{k+1},a_k){+}J(b_{k+1},a_{(k+1)+})\Big], \tag{4.6}$$

where the expectation is with respect to $\mathbb{P}(z_{k+1}|\mathcal{H}^-_{k+1})$. In the above, and throughout the rest of this work, $\mathcal{H}^-_t{\triangleq}\mathcal{H}_{t-1}{\cup}\{a_{t-1}\}, \forall t{\in}\{k{+}1,\ldots,k{+}L\}$. Uniquely in this work, data associations play a major role as they dictate link actions considered by the robot, as we shall see in Sec. 4.2.4. Therefore, explicitly incorporating them within the objective is useful. We can rewrite (4.6) as:

$$J(b_k,a_{k+}){=}\mathop{\mathbb{E}}_{\beta_{k+1}}\Bigg[\mathop{\mathbb{E}}_{z_{k+1}|\beta_{k+1}}\Big[c_1{+}J(b_{k+1},a_{(k+1)+})\Big]\Bigg], \tag{4.7}$$

where $c_1{\triangleq}c_1(b_{k+1},a_k)$. As the above implies, the law of total probability relates between z_t and β_t, for any $t{\in}\{k{+}1,\ldots,k{+}L\}$. That is:

$$\mathbb{P}(z_t|\mathcal{H}^-_t){=}\sum_{\beta_t}\mathbb{P}(z_t,\beta_t|\mathcal{H}^-_t). \tag{4.8}$$

We aim to derive the joint likelihood term $\mathbb{P}(z_t,\beta_t|\mathcal{H}^-_t)$, given specific realizations of z_t and β_t. In contrast to [22], we do so considering a qualitative framework.

First, we marginalize over the robot states $\mathcal{S}^{X_t}_{F_{t-1}}$ and $\mathcal{S}^{X_{t-1}}_{F_{t-1}}$, global frame scale $\mathcal{S}^{F_{t-1}}$ and, the qualitative state $\mathcal{S}^{\tau_{\beta_t}}$ that corresponds to the considered data association realization β_t. To simplify the exposition, we consider the latter to be a single landmark triplet, although this is not a limitation of our formulation:

$$\mathbb{P}(z_t,\beta_t \mid \mathcal{H}^-_t) = \sum_{\mathcal{S}^{X_t}_{F_{t-1}}}\sum_{\mathcal{S}^{X_{t-1}}_{F_{t-1}}}\sum_{\mathcal{S}^{F_{t-1}}}\sum_{\mathcal{S}^{\tau_{\beta_t}}}\mathbb{P}(z_t,\beta_t,\mathcal{S}^{X_t}_{F_{t-1}},\mathcal{S}^{X_{t-1}}_{F_{t-1}},\mathcal{S}^{F_{t-1}},\mathcal{S}^{\tau_{\beta_t}} \mid \mathcal{H}^-_t). \tag{4.9}$$

Continuing with chain rule over the inner expression, we get:

$$\begin{aligned}&\mathbb{P}(z_t,\beta_t,\mathcal{S}^{X_t}_{F_{t-1}},\mathcal{S}^{X_{t-1}}_{F_{t-1}},\mathcal{S}^{F_{t-1}},\mathcal{S}^{\tau_{\beta_t}}|\mathcal{H}^-_t){=}\\&\mathbb{P}(z_t|\mathcal{S}^{X_t}_{F_{t-1}},\mathcal{S}^{\tau_{\beta_t}},\beta_t,\mathcal{H}^-_t)\mathbb{P}(\beta_t|\mathcal{S}^{X_t}_{F_{t-1}},\mathcal{S}^{F_{t-1}},\mathcal{S}^{\tau_{\beta_t}},\mathcal{H}^-_t)\mathbb{P}(\mathcal{S}^{X_t}_{F_{t-1}}|\mathcal{S}^{X_{t-1}}_{F_{t-1}},a^{\mathrm{q}}_{t-1})\mathbb{P}(\mathcal{S}^{X_{t-1}}_{F_{t-1}},\mathcal{S}^{F_{t-1}},\mathcal{S}^{\tau_{\beta_t}}|\mathcal{H}_{t-1}).\end{aligned} \tag{4.10}$$

The term $\mathbb{P}(z_t|\mathcal{S}^{X_t}_{F_{t-1}},\mathcal{S}^{\tau_{\beta_t}},\beta_t,\mathcal{H}^-_t)$, known as the measurement model, describes the probability of capturing the measurement z_t, given the robot's state $\mathcal{S}^{X_t}_{F_{t-1}}$, the data association β_t with its corresponding state $\mathcal{S}^{\tau\beta_t}$, and history. We further develop it via marginalization over relevant metric realizations and considering dependencies:

$$\mathbb{P}(z_t|\mathcal{S}^{X_t}_{F_{t-1}},\mathcal{S}^{\tau\beta_t},\beta_t,\mathcal{H}^-_t)=\iint\limits_{x\in\mathcal{S}^{X_t}_{F_{t-1}},\mathcal{L}\in\mathcal{S}^{\tau\beta_t}} \mathbb{P}(z_t|x,\mathcal{L},F_{t-1})\mathbb{P}(x|\mathcal{S}^{X_t}_{F_{t-1}},\mathcal{H}^-_t)\mathbb{P}(\mathcal{L}|\mathcal{S}^{\tau\beta_t},\mathcal{H}^-_t)\,dx\,d\mathcal{L}, \quad (4.11)$$

The term $\mathbb{P}(z_t|x,\mathcal{L},F_{t-1})$ is the metric measurement model. The metric priors $\mathbb{P}(x \mid \mathcal{S}^{X_t}_{F_{t-1}},\mathcal{H}^-_t)$ and $\mathbb{P}(\mathcal{L} \mid \mathcal{S}^{\tau\beta_t},\mathcal{H}^-_t)$ can be further approximated as uniform distributions by neglecting the history term. Accordingly, (4.11) can be calculated offline.

The term $\mathbb{P}(\beta_t|\mathcal{S}^{X_t}_{F_{t-1}},\mathcal{S}^{F_{t-1}},\mathcal{S}^{\tau\beta_t},\mathcal{H}^-_t)$, known as the association model, describes the probability to observe the triplet β_t, given the robot's and β_t's states and F_{t-1}'s global scale. Intuitively, to evaluate this probability, the robot must estimate its sensing range, R, in local terms of F_{t-1}. Using F_{t-1}'s scale, this can be done via marginalizing over the relevant metric realizations and considering dependencies:

$$\begin{aligned}&\mathbb{P}(\beta_t|\mathcal{S}^{X_t}_{F_{t-1}},\mathcal{S}^{F_{t-1}},\mathcal{S}^{\tau\beta_t},\mathcal{H}^-_t)= \quad (4.12)\\ &\iiint\limits_{x\in\mathcal{S}^{X_t}_{F_{t-1}},d\in\mathcal{S}^{F_{t-1}},\mathcal{L}\in\mathcal{S}^{\tau\beta_t}}\mathbb{P}(\beta_t|x,d,\mathcal{L},F_{t-1})\mathbb{P}(x|\mathcal{S}^{X_t}_{F_{t-1}},\mathcal{H}^-_t)\mathbb{P}(d|\mathcal{S}^{F_{t-1}},\mathcal{H}^-_t)\mathbb{P}(\mathcal{L}|\mathcal{S}^{\tau\beta_t},\mathcal{H}^-_t)\,dx\,dd\,d\mathcal{L}.\end{aligned}$$

The term $\mathbb{P}(\beta_t|x,d,\mathcal{L},F_{t-1})$ is a deterministic geometric model that equals 1 if the metric hypotheses of β_t landmarks are inside the robot's sensing range, R (assumed to be a known hyperparameter), and 0 else, i.e.:

$$\mathbb{P}(\beta_t|x,d,\mathcal{L},F_{t-1})=\prod_{\mathcal{L}_i\in\{\mathcal{L}_1,\mathcal{L}_2,\mathcal{L}\}}\mathbb{I}\Big\{\,\|\mathcal{L}_i-x\|_2\leq\frac{R}{d}\Big\}, \quad (4.13)$$

where $\mathcal{L}_1$ and $\mathcal{L}_2$ are the local metric coordinate of the reference landmarks creating F_{t-1}. In most cases, $\mathcal{L}_1$=$(0,0)$ and $\mathcal{L}_2$=$(0,1)$. The metric priors can be further approximated as uniform distributions by neglecting the history term. Accordingly, (4.12) can be calculated offline.

Moreover, in (4.10), the term $\mathbb{P}(\mathcal{S}^{X_t}_{F_{t-1}} \mid \mathcal{S}^{X_{t-1}}_{F_{t-1}}, a^{\text{q}}_{t-1})$ is the qualitative motion model stated in (4.5), and finally, the term $\mathbb{P}(\mathcal{S}^{X_{t-1}}_{F_{t-1}},\mathcal{S}^{F_{t-1}},\mathcal{S}^{\tau\beta_t} \mid \mathcal{H}_{t-1})$ can be calculated via marginalization from the belief from time instant $t{-}1$, b_{t-1}.

Generally speaking, we need to consider all possible realizations for both z_t and β_t, for any $t\in\{k{+}1,\dots,k{+}L\}$, to calculate the objective precisely. However, in practice, this is intractable. Alternatively, we can sample a finite set of realizations and approximate the objective via proper averaging. Starting with β_t, we suggest considering only triplets' realizations containing the robot's current frame. This heuristic prioritizes triplets that are more likely to be observed, as those containing the current frame are more likely to be closer. For instance, in case F_k=AB, then ABC is being considered, also ABD, but not ACD. This heuristic yields N_β=$|\mathbb{L}|{-}2$ β_t's realizations, which is a significant reduction compared to all $\binom{|\mathbb{L}|}{3}$ triplets existing under the given landmark alphabet based on $\mathbb{L}$. In practice, among these, we can only consider those which are available in b_{t-1}, that is, n_β realizations, where $n_\beta \leq N_\beta$. Crucially, as we shall see in Sec. 4.2.6, utilizing composition we are able to extend n_β to come closer to, and under a certain assumption equal to, N_β. Considering also a finite set of n_z measurement realizations z_t for each realization of β (see Fig. 4.1 for illustration), the objective (4.7) can be approximated as follows:

$$
\begin{aligned}
J(b_k, a_{k+}) &= \underset{\beta_{k+1}}{\mathbb{E}}\Big[\underset{z_{k+1}|\beta_{k+1}}{\mathbb{E}} \big[c_1(b_{k+1}, a_k) + J(b_{k+1}, a_{(k+1)+}) \big] \Big] \\
&= \sum_{\beta_{k+1}} \mathbb{P}(\beta_{k+1}|b_k, a_k) \int_{z_{k+1}} \mathbb{P}(z_{k+1}|\beta_{k+1}, b_k, a_k) \cdot (c_1 + J(b_{k+1}, a_{(k+1)+}))\, dz_{k+1} \\
&\approx \sum_{m=1}^{N_\beta} \underbrace{\frac{\mathbb{P}(\beta^m_{k+1}|b_k, a_k)}{\sum_{q=1}^{N_\beta} \mathbb{P}(\beta^q_{k+1}|b_k, a_k)}}_{\tilde{w}^m} \int_{z_{k+1}} \mathbb{P}(z_{k+1}|\beta^m_{k+1}, b_k, a_k) \cdot (c_1 + J(b_{k+1}, a_{(k+1)+}))\, dz_{k+1} \\
&\approx \sum_{i=1}^{n_\beta} \underbrace{\frac{\tilde{w}^i}{\sum_{q=1}^{n_{\beta_k}} \tilde{w}^q}}_{w^i} \int_{z_{k+1}} \mathbb{P}(z_{k+1}|\beta^i_{k+1}, b_k, a_k) \cdot (c_1 + J(b_{k+1}, a_{(k+1)+}))\, dz_{k+1} \\
&\approx \sum_{i=1}^{n_\beta} \frac{w^i}{n_z} \sum_{j=1}^{n_z} \mathbb{P}(z^{i,j}_{k+1}|\beta^i_{k+1}, b_k, a_k) \cdot (c_1 + J(b_{k+1}, a_{(k+1)+})),
\end{aligned}
\tag{4.14}
$$

where in the 1st approximation, we consider only the subset of N_β β_{k+1}'s realizations containing triplets that involve the current frame ($N_\beta = |\mathbb{L}| - 2$), in the 2nd approximation, we further reduced this subset to the n_β triplets available in the belief ($n_\beta \leq N_\beta$), and finally, in the 3rd approximation, we show how the inner expectation term can be evaluated via averaging over a finite set of z_{k+1} samples.

4.2.4 ***Link*** Action and Transition Model

The *Link* action allows the robot to switch between different reference frames. Even though it is considered an action, in practice, the robot does not make any actual movement when executing it. Accordingly, a *Link* action taken at time step t is denoted by a_t^{Link}, but alternatively can be written as a tuple of source and target frames, (F_{t-1}, F_t). As Fig. 4.1 illustrates, the *Link* action adds another level of decision-making compared to the traditional BSP approaches, which use a single global frame.

We assume a probabilistic *Link* model,

$$
\mathbb{P}(\mathcal{S}^{X_t}_{F_t} | \mathcal{S}^{X_t}_{F_{t-1}}, \mathcal{S}^\tau, \mathcal{H}_t),
\tag{4.15}
$$

representing the probability of the robot being located at state $\mathcal{S}^{X_t}_{F_t}$, given its state relative to the former frame, $\mathcal{S}^{X_t}_{F_{t-1}}$, the state of the triplet relates between the frames, $\mathcal{S}^\tau$, with $\tau \triangleq F_{t-1}{:}F_t \backslash F_{t-1}$, and history, which includes the action a_t^{Link}. To understand why this model is a probabilistic one, consider the following example. Suppose that,

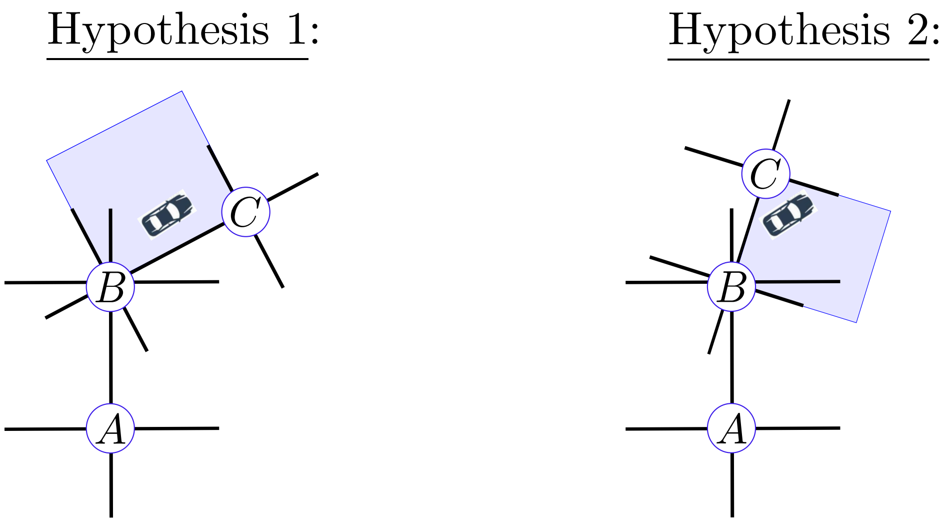

Figure 4.2: Illustration of two possible hypotheses for $\mathcal{X}_{AB}^{C}$. While both hypotheses yield the same qualitative state $\mathcal{S}_{AB}^{C}$, they yield different states for $\mathcal{S}_{BC}^{X_t}$.

at time step t, the robot is located relative to an old frame $F_{t-1}{=}AB$ and aims to link to a new one, $F_t{=}BC$. To that end, it must translate its location from AB's terms to BC's. Given its (unknown) current metric location $\mathcal{X}_{AB}^{X_t}$ and C's metric location in terms of the same frame, $\mathcal{X}_{AB}^{C}$ (i.e., $\tau{=}AB{:}C$), the robot can easily infer $\mathcal{X}_{BC}^{X_t}$, using a simple geometric transformation. Fig. 4.2 demonstrates two sets of metric realizations as described above. However, since the *Link* model considers only the corresponding qualitative locations rather than the actual metric ones, the robot must account for many possible metric hypotheses. For this reason, it can only infer a distribution over the resulting state $\mathcal{S}_{F_t}^{X_t}$. To calculate this model, we further marginalize over the relevant metric state:

$$\mathbb{P}(\mathcal{S}_{F_t}^{X_t}|\mathcal{S}_{F_{t-1}}^{X_t},\mathcal{S}^{\tau},\mathcal{H}_t)=\iint\limits_{\mathcal{X}_{F_{t-1}}^{X_t}\in\mathcal{S}_{F_{t-1}}^{X_t},\mathcal{X}^{\tau}\in\mathcal{S}^{\tau}} \mathbb{P}(\mathcal{S}_{F_t}^{X_t},\mathcal{X}_{F_{t-1}}^{X_t},\mathcal{X}^{\tau}|\mathcal{S}_{F_{t-1}}^{X_t},\mathcal{S}^{\tau},\mathcal{H}_t)\, d\mathcal{X}_{F_{t-1}}^{X_t}\, d\mathcal{X}^{\tau}. \tag{4.16}$$

We continue developing the inner term using chain rule:

$$\mathbb{P}(\mathcal{S}_{F_t}^{X_t},\mathcal{X}_{F_{t-1}}^{X_t},\mathcal{X}^{\tau}|\mathcal{S}_{F_{t-1}}^{X_t},\mathcal{S}^{\tau},\mathcal{H}_t)=\mathbb{P}(\mathcal{S}_{F_t}^{X_t}|\mathcal{X}_{F_{t-1}}^{X_t},\mathcal{X}^{\tau},a_t^{Link})\mathbb{P}(\mathcal{X}_{F_{t-1}}^{X_t}|\mathcal{S}_{F_{t-1}}^{X_t},\mathcal{H}_t)\mathbb{P}(\mathcal{X}^{\tau}|\mathcal{S}^{\tau},\mathcal{H}_t), \tag{4.17}$$

where $\mathbb{P}(\mathcal{S}_{F_t}^{X_t}|\mathcal{X}_{F_{t-1}}^{X_t},\mathcal{X}^{\tau},a_t^{Link})$ is a geometric model that deterministically determines the new state, given a metric realization of the former one and of the related triplet. The metric priors $\mathbb{P}(\mathcal{X}_{F_{t-1}}^{X_t}|\mathcal{S}_{F_{t-1}}^{X_t},\mathcal{H}_t)$ and $\mathbb{P}(\mathcal{X}^{\tau}|\mathcal{S}^{\tau},\mathcal{H}_t)$ can be further approximated as uniform distributions by neglecting the history term. Accordingly, (4.16) can be calculated offline.

4.2.5 Belief Update Step

This section focuses on updating the qualitative belief defined in Eq. (4.1) a single step into the future. Formally, consider the belief from time step $t{-}1{\in}\{k,\ldots,k{+}l{-}1\}$, b_{t-1}, candidate action, $a_{t-1}{\triangleq}\{a^{\mathrm{q}}_{t-1}, a^{Link}_{t-1}\}$, a measurement, z_t, and a corresponding data association, β_t. We aim to infer the belief at time step t, b_t, via a proper Bayesian update rule:

$$b_t = \psi(b_{t-1}, a_{t-1}, z_t, \beta_t). \tag{4.18}$$

We start by marginalizing over the next robot's position relative to the current frame, $\mathcal{S}^{X_t}_{F_{t-1}}$:

$$b_t{=}\sum_{\mathcal{S}^{X_t}_{F_{t-1}}} \mathbb{P}(\mathcal{S}^{X_{1:t}}, \mathcal{S}^{X_t}_{F_{t-1}}, \mathcal{S}^{\mathcal{M}_t}{,}\mathcal{S}^{\mathbb{F}_t}|\mathcal{H}_t). \tag{4.19}$$

We continue by breaking the inner term using chain rule:

$$\mathbb{P}(\mathcal{S}^{X_{1:t}}, \mathcal{S}^{X_t}_{F_{t-1}}, \mathcal{S}^{\mathcal{M}_t}, \mathcal{S}^{\mathbb{F}_t}|\mathcal{H}_t) = \mathbb{P}(\mathcal{S}^{X_t}_{F_t} \mid \mathcal{S}^{X_t}_{F_{t-1}}, \mathcal{S}^{\tau_{\beta_t}}, \mathcal{H}_t)\mathbb{P}(\mathcal{S}^{X_{1:t-1}}, \mathcal{S}^{X_t}_{F_{t-1}}, \mathcal{S}^{\mathcal{M}_t}, \mathcal{S}^{\mathbb{F}_t}|\mathcal{H}_t). \tag{4.20}$$

The left term, obtained after omitting all triplets' states from $\mathcal{S}^{\mathcal{M}_t}$ except for β_t's, is the *Link* Model stated in (4.15). We continue developing the right term via Bayes rule over z_t and β_t taken from $\mathcal{H}_t$ while omitting irrelevant information:

$$\begin{aligned}&\mathbb{P}(\mathcal{S}^{X_{1:t-1}}{,}\mathcal{S}^{X_t}_{F_{t-1}}{,}\mathcal{S}^{\mathcal{M}_t}{,}\mathcal{S}^{\mathbb{F}_t}|\mathcal{H}_t) = \\ &\eta_t\mathbb{P}(z_t|\mathcal{S}^{X_t}_{F_{t-1}}{,}\mathcal{S}^{\tau_{\beta_t}}{,}\beta_t{,}\mathcal{H}^-_t)\mathbb{P}(\beta_t|\mathcal{S}^{X_t}_{F_{t-1}}{,}\mathcal{S}^{\tau_{\beta_t}}{,}\mathcal{S}^{F_{t-1}}{,}\mathcal{H}^-_t)\mathbb{P}(\mathcal{S}^{X_t}_{F_{t-1}}|\mathcal{S}^{X_{t-1}}_{F_{t-1}}{,}a^{\mathrm{q}}_{t-1})\mathbb{P}(\mathcal{S}^{X_{1:t-1}}{,}\mathcal{S}^{\mathcal{M}_t}{,}\mathcal{S}^{\mathbb{F}_t}|\mathcal{H}^-_t),\end{aligned} \tag{4.21}$$

where $\eta_t{\triangleq}\mathbb{P}(z_t, \beta_t \mid \mathcal{H}^-_t)$ is a normalization term, $\mathbb{P}(z_t \mid \mathcal{S}^{X_t}_{F_{t-1}}, \mathcal{S}^{\tau_{\beta_t}}, \beta_t, \mathcal{H}^-_t)$ and $\mathbb{P}(\beta_t \mid \mathcal{S}^{X_t}_{F_{t-1}}, \mathcal{S}^{\tau_{\beta_t}}, \mathcal{S}^{F_t}, \mathcal{H}^-_t)$ are the measurement and association models discussed in Section 4.2.3, respectively, and finally, $\mathbb{P}(\mathcal{S}^{X_t}_{F_{t-1}} \mid \mathcal{S}^{X_{t-1}}_{F_{t-1}}, a^{\mathrm{q}}_{t-1})$ is the qualitative transition model (4.5). We continue by applying another chain rule over the remaining term in (4.21):

$$\mathbb{P}(\mathcal{S}^{X_{1:t-1}}{,}\mathcal{S}^{\mathcal{M}_t}{,}\mathcal{S}^{\mathbb{F}_t}|\mathcal{H}^-_t) = \mathbb{P}(\mathcal{S}^{F_{t-1}}|\mathcal{S}^{\mathcal{M}_t}{,}\mathcal{S}^{\mathbb{F}_{t-1}}{,}\mathcal{H}^-_t)\mathbb{P}(\mathcal{S}^{X_{1:t-1}}{,}\mathcal{S}^{\mathcal{M}_t}{,}\mathcal{S}^{\mathbb{F}_{t-1}}|\mathcal{H}^-_t). \tag{4.22}$$

In the above, we consider the new set of frames' scales at time step t, $\mathcal{S}^{\mathbb{F}_t}$, as the former set from time step $t{-}1, \mathcal{S}^{\mathbb{F}_{t-1}}$ unified with the currently considered scale, $\mathcal{S}^{F_{t-1}}$, i.e., $\mathcal{S}^{\mathbb{F}_t}{=}\mathcal{S}^{\mathbb{F}_{t-1}}{\cup}\mathcal{S}^{F_{t-1}}$. The term $\mathbb{P}(\mathcal{S}^{F_{t-1}}|\mathcal{S}^{\mathcal{M}_t}{,}\mathcal{S}^{\mathbb{F}_{t-1}}{,}\mathcal{H}^-_t)$ is the posterior over the relevant frame scale, $\mathcal{S}^{F_{t-1}}$, given the map, $\mathcal{S}^{\mathcal{M}_t}$, available frames' scales, $\mathcal{S}^{F_{t-1}}$, and history, $\mathcal{H}^-_t$. To further develop this term, we assume that the global scale of AB, $\mathcal{S}^{AB}$, is known. This assumption is fairly reasonable, as it requires prior knowledge of only one frame scale. Via this scale, we can evaluate the scale of $F_{t-1}{=}\{L_1{,}L_2\}$, using the states $\mathcal{S}^{L_1}_{AB}$ and $\mathcal{S}^{L_2}_{AB}$, which locate F_{t-1}'s landmarks relative to AB. Mathematically, we approximate the abovementioned term by neglecting unnecessary information:

$$\mathbb{P}(\mathcal{S}^{F_{t-1}}|\mathcal{S}^{\mathcal{M}_t},\mathcal{S}^{\mathbb{F}_{t-1}},\mathcal{H}_t^-) \approx \mathbb{P}(\mathcal{S}^{F_{t-1}}|F_{t-1},\mathcal{S}_{AB}^{L_1},\mathcal{S}_{AB}^{L_2},\mathcal{H}_t^{ABL_1},\mathcal{H}_t^{ABL_2},\mathcal{S}^{AB}), \tag{4.23}$$

where the triplet states $\mathcal{S}_{AB}^{L_1}$,$\mathcal{S}_{AB}^{L_2}$ were taken from $\mathcal{S}^{\mathcal{M}_t}$, the frame F_{t-1} and the relevant local histories $\mathcal{H}_t^{ABL_1}$,$\mathcal{H}_t^{ABL_2}$ were taken from $\mathcal{H}_t^-$, and finally, $\mathcal{S}^{AB}$ is the known global scale taken from $\mathcal{S}^{F_{t-1}}$. We can further develop this posterior term, assuming $F_{t-1}{=}L_1L_2$, via marginalization over the metric state of $AB{:}L_1$ and $AB{:}L_2$, followed by chain rule:

$$\begin{aligned}&\mathbb{P}(\mathcal{S}^{F_{t-1}}|F_{t-1}{=}L_1L_2,\mathcal{S}_{AB}^{L_1},\mathcal{S}_{AB}^{L_2},\mathcal{H}_t^{ABL_1},\mathcal{H}_t^{ABL_2}){=}\\ &\iint\limits_{\mathcal{X}_{AB}^{L_1},\mathcal{X}_{AB}^{L_2}}\mathbb{P}(\mathcal{S}^{F_{t-1}}|F_{t-1}{=}L_1L_2,\mathcal{X}_{AB}^{L_1},\mathcal{X}_{AB}^{L_2})\mathbb{P}(\mathcal{X}_{AB}^{L_1},\mathcal{X}_{AB}^{L_2}\mid\mathcal{S}_{AB}^{L_1},\mathcal{S}_{AB}^{L_2},\mathcal{H}_t^{ABL_1},\mathcal{H}_t^{ABL_2})\,d\mathcal{X}_{AB}^{L_1}\,d\mathcal{X}_{AB}^{L_2},\end{aligned} \tag{4.24}$$

where $\mathbb{P}(\mathcal{S}^{F_{t-1}}|F_{t-1}{=}L_1L_2,\mathcal{X}_{AB}^{L_1},\mathcal{X}_{AB}^{L_2})$ is a Dirac function equals to 1 if $\left\|\mathcal{X}_{AB}^{L_1}-\mathcal{X}_{AB}^{L_2}\right\|_2$ is in the interval represented by the value of $\mathcal{S}^{F_{t-1}}$ and to 0 otherwise. The metric prior term can be approximated via $\mathbb{P}(\mathcal{X}_{AB}^{L_1},\mathcal{X}_{AB}^{L_2}\mid\mathcal{S}_{AB}^{L_1},\mathcal{S}_{AB}^{L_2},\mathcal{H}_t^{ABL_1},\mathcal{H}_t^{ABL_2}) \approx \prod_{i=1}^{2}\mathbb{P}(\mathcal{X}_{AB}^{L_i}\mid\mathcal{S}_{AB}^{L_i},\mathcal{H}_t^{ABL_i})$, where $\forall i{\in}\{1{,}2\}$ the individual prior term can be further approximated via $\mathbb{P}(\mathcal{X}_{AB}^{L_i}\mid\mathcal{S}_{AB}^{L_i},\mathcal{H}_t^{ABL_i}){\approx}\mathbb{P}(\mathcal{X}_{AB}^{L_i}\mid\mathcal{S}_{AB}^{L_i})$, i.e., assuming a uniform distribution.

. In Sec. 4.2.6, we explain how to compose $\mathcal{S}_{AB}^{L_1}$ and $\mathcal{S}_{AB}^{L_2}$ in case they are not part of the current belief.

In total, we get the following update rule:

$$\begin{aligned}b_t{=}\sum_{\mathcal{S}_{F_{t-1}}^{X_t}}&\mathbb{P}(\mathcal{S}_{F_t}^{X_t}|\mathcal{S}_{F_{t-1}}^{X_t},\mathcal{S}^{\tau_{\beta_t}},\mathcal{H}_t)\eta_t\cdot\mathbb{P}(z_t|\mathcal{S}_{F_{t-1}}^{X_t},\mathcal{S}^{\tau_{\beta_t}},\mathcal{S}^{F_t},\beta_t,\mathcal{H}_t^-)\cdot\\ &\mathbb{P}(\beta_t|\mathcal{S}_{F_{t-1}}^{X_t},\mathcal{S}^{\tau_{\beta_t}},\mathcal{S}^{F_t},\mathcal{H}_t^-)\cdot\mathbb{P}(\mathcal{S}_{F_{t-1}}^{X_t}|\mathcal{S}_{F_{t-1}}^{X_{t-1}},a_{t-1}^{\mathrm{q}})\mathbb{P}(\mathcal{S}^{F_{t-1}}|\mathcal{S}^{\mathcal{M}_t},\mathcal{S}^{\mathbb{F}_{t-1}},\mathcal{H}_t^-)b_{t-1}.\end{aligned} \tag{4.25}$$

The above update step describes the operator $\psi(\cdot)$ from Eq. (4.18) explicitly. We apply it to update the belief between consecutive time steps when constructing the tree, as illustrated in Fig. 4.1a.

4.2.6 Incorporating Compositions

So far, we have described our qualitative BSP approach in its base form, without utilizing compositions at all. In this section, we present our *second key contribution* and show how compositions can be integrated within our algorithm to further improve planning in two ways. Firstly, it allows us to deal with a broader range of scenarios, i.e., in some cases, a plan can be found *only* via compositions. Secondly, using compositions, we can find better plans, i.e., ones with a lower objective. In this section, we provide theoretical justification for the above. Also, we formally show how to incorporate compositions within our algorithm.

Recall Sec. 2.4, and the topological rule declared in Lemma 2.4.1. According to the rule, two source triplets must share 2 landmarks in common to compose a third triplet. The mutual landmarks, allow us to fix both triplets relative to the same frame and, hence, infer (compose) relationships between new triplets' combinations. For example, we can directly compose $\tau_3{=}AB{:}D$ using $\tau_1{=}AB{:}C$ and $\tau_2{=}BC{:}D$.

In Eq. (4.1), we defined the belief as a posterior distribution over, among others, the different triplets in the map. Consider the map's state at the current time step, $\mathcal{S}^{\mathcal{M}_k}$. Via compositions, the robot can augment $\mathcal{M}_k$ with new triplets. However, compositions are allowed only under specific topological conditions, formulated in Lemma 2.4.1. For example, if $\mathcal{M}_k$ is too sparse, the ability to compose new triplets might be limited or even impossible. In Sec. 3.2.1 (also published in [23]), we have formulated a sufficient topological condition attributed to the set $\mathcal{M}_k$, whose existence ensures the ability to compose any desired triplet under the considered landmark space $\mathbb{L}$. We named such a sufficiently dense set a *Composable* set under $\mathbb{L}$ (the formal definition can be found in 3.2.4). As this topological aspect is not the focus of this part of the work, we assume that $\mathcal{M}_k$ is *Composable* under $\mathbb{L}$. However, this assumption is not a must. In Sec. 4.2.3, we suggested a heuristic that considers a maximum of N_β triplets as possible realizations for β_t at any planning time step $t{\in}\{k{+}1,\ldots,k{+}L\}$. We stressed that, in practice, we could consider only a subset of n_β triplets out of these N_β, which are available in b_{t-1}. Since we assumed $\mathcal{M}_k$ is *Composable* under $\mathbb{L}$, we are now guaranteed that all N_β realizations can be considered in planning via compositions, as illustrated in Fig. 4.1b.

We shall now prove that in some scenarios, a plan can be found **only** via compositions. We use *Link-Graph*s throughout the proof. As stated in Sec. 2.3, *Link-Graph*s are good representations for links mobility, since each triplet node enables a transition, or *Link*, between any two frames' edges connected to it. Consequently, a *Link-Graph*'s path encodes a feasible sequence of link actions, where the edges along the path are the different frames, and the in-between nodes are the triplets the robot relies on to execute the *Link*s.

Using this insight, we now aim to prove that in some cases, a plan can be found only via compositions.

Before approaching the formal proof, we provide some intuition. The key point of our explanation is simple. Via compositions, the robot can link to more frames than it could before. According to the conclusion from Lemma 2.2.1, the robot is allowed to link based on a path of a *Link-Graph*, whose nodes represent the set of available triplets, $\mathcal{M}_k$. Thus, without compositions, links are possible only based on existing paths. In contrast, using composition, we can create new triplets, i.e., augment the graph with new nodes, thus creating additional paths. Consequently, in cases where there is no path in the *Link-Graph* at planning time to a target triplet without compositions, we cannot find a valid plan towards the triplet.

Suppose that the robot's initial map, $\mathcal{M}_k$, is *Composable* under $\mathbb{L}$ (see Fig. 4.3a for illustration). Alternatively, we could assume that a *Link-Graph* whose nodes represent $\mathcal{M}_k$ is *Composable* under $\mathbb{L}$, considering the following definition:

Definition 4.2.1. Let $G{=}(V{,}E)$ be a *Link-Graph*. We say that G is *Composable* under $\mathbb{L}$ if V represents a *Composable* set of triplets under $\mathbb{L}$.

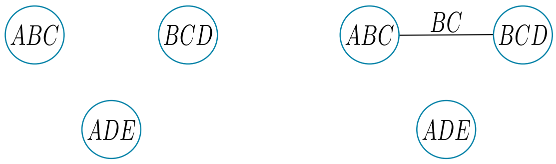

(a) (b)

Figure 4.3: (a) An illustration of a *Composable* set under $\mathbb{L}{=}\{A{,}B{,}C{,}D{,}E\}$, consists of three tripelts; (b) An illustration of a *Link-Graph*, based on the set from (a). The graph has a single edge connecting ABC with BCD, as B and C are mutual landmarks; There is an Invertible transformation between the two.

We aim to prove that any connected Link-Graph is, in particular, *Composable*, but not the other way around:

Theorem 4.1. *Let* $\mathbb{G}_{\text{cn}}$ *and* $\mathbb{G}_{\text{cm}}$ *be the sets of all connected and* Composable Link-Graph*s under the landmark space* $\mathbb{L}$*, respectively. Then* $\mathbb{G}_{\text{cn}}{\subsetneq}\mathbb{G}_{\text{cm}}$.

Proof.. First we show that $\mathbb{G}_{cn}{\subseteq}\mathbb{G}_{cm}$.
Let $G{=}(V{,}E)$ be a connected Link graph under $\mathbb{L}$. We prove that G is also *Composable* under $\mathbb{L}$ by induction on number of vertices in G, $|V|$.
Base step: When $|V|{=}1$, V is *Composable* under $\mathbb{L}$ by definition. Thus, G is also *Composable* under $\mathbb{L}$ by definition.
Induction step: Suppose G is *Composable* under $\mathbb{L}$ for all $1{\le}|V|{\le}n$. We show that G is *Composable* under $\mathbb{L}$ for $|V|{=}n{+}1$. We choose a cut in G, $C{=}(S{,}T)$, s.t. $G_S{\triangleq}(S,\{(u{,}v){\in}E|(u{,}v){\in}S^2\})$ and $G_T{\triangleq}(T,\{(u{,}v){\in}E|(u{,}v){\in}T^2\})$ are both connected graphs, where $S,T{\neq}\emptyset$. Note that such choise always exists for any $|V|{>}1$, since G is connected. Let us now observe the set of edges in G connecting S with T, that is, $E_{S{,}T}{\triangleq}\{(u{,}v){\in}E|u{\in}S \wedge v{\in}T\}$. Since G is connected, we are guaranteed that $E_{S{,}T}{\neq}\emptyset$. Thus, C is a 2-common cut in G. Finally, since $G_S{,}G_T$ are both connected subgraphs of G, they are both connected *Link-Graph*s, and since $1{\le}|S|{,}|T|{\le}n$, we further conclude that they are both *Composable* under $\mathbb{L}$, according to the assumption. That is to say, we showed by definition that for $|V|{=}n{+}1$, G is *Composable* under $\mathbb{L}$.
Conclusion: $\mathbb{G}_{cn}{\subseteq}\mathbb{G}_{cm}$.
We are left to show an instance of a *Composable Link-Graph* under $\mathbb{L}$ that is not connected. To that end, consider the landmark space $\{A{,}B{,}C{,}D{,}E\}$, and the *Link-Graph* from Fig. 4.3b.
Final conclusion: $\mathbb{G}_{cn}{\subsetneq}\mathbb{G}_{cm}$
■ ■

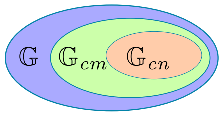

Figure 4.4: Relationships between general *Link-Graphs* ($\mathbb{G}$), *Composable Link-Graphs* ($\mathbb{G}_{cm}$), and connected *Link-Graphs*, all under the same landmark space, ($\mathbb{G}_{cn}$) are described through a Venn diagram.

Meaning, in some scenarios, where $\mathcal{M}_k$ creates a *Composable Link-Graph* that is disconnected, compositions are necessary to allow the robot to plan towards its goal.

For any planning time step $t \in \{k+1, \ldots, k+L\}$, we shall now explain how to incorporate a new set of triplets $\mathcal{M}_t^{\Delta}$, which can be chosen according to the heuristic described in Sec. 4.2.3, into the belief. We denote $\mathcal{M}_t \triangleq \mathcal{M}_{t-1} \cup \mathcal{M}_t^{\Delta}$ and express the belief after combining the set, given by $b_{t-1}^{\mathcal{M}} \triangleq \mathbb{P}(\mathcal{S}^{X_{1:t-1}}, \mathcal{S}^{\mathcal{M}_t}, \mathcal{S}^{\mathbb{F}_{t-1}} | \mathcal{H}_{t-1})$, in terms of the one before combining it, b_{t-1}, via chain rule and Markov assumption:

$$b_{t-1}^{\mathcal{M}} = \mathbb{P}(\mathcal{S}^{\mathcal{M}_t^{\Delta}} | \mathcal{S}^{\mathcal{M}_{t-1}}, \mathcal{H}_{t-1}) \cdot b_{t-1}. \tag{4.26}$$

Suppose $|\mathcal{M}_t^{\Delta}| = n_t$, the term $\mathbb{P}(\mathcal{S}^{\mathcal{M}_t^{\Delta}} | \mathcal{S}^{\mathcal{M}_{t-1}}, \mathcal{H}_{t-1})$ can be further broken down into a product of individuals posteriors, each aims to compose a single triplet from $\mathcal{M}_t^{\Delta}$:

$$\mathbb{P}(\mathcal{S}^{\mathcal{M}_t^{\Delta}} | \mathcal{S}^{\mathcal{M}_{t-1}}, \mathcal{H}_{t-1}) \approx \prod_{i=1}^{n_t} \mathbb{P}(\mathcal{S}^{\tau_i} | \mathcal{S}^{\mathcal{M}_{t-1}^i}, \mathcal{H}_{t-1}^i). \tag{4.27}$$

In the above, $\mathcal{S}^{\tau_i}$ denotes the state of the ith triplet taken from $\mathcal{S}^{\mathcal{M}_t^{\Delta}}$, $\mathcal{S}^{\mathcal{M}_{t-1}^i}$ denotes the minimal subset of $\mathcal{S}^{\mathcal{M}_{t-1}}$ required for the evaluation, and finally, $\mathcal{H}_{t-1}^i$ denotes the corresponding local history, where $\mathcal{S}^{\mathcal{M}_t^{\Delta}} = \bigcup_{i=1}^{n_t} \mathcal{S}^{\tau_i}, \mathcal{H}_{t-1} = \bigcup_{i=1}^{n_t} \mathcal{H}_{t-1}^i$. For instance, if $\mathcal{S}^{\tau_i} = \mathcal{S}_{AB}^{D}$ and $\mathcal{S}^{\mathcal{M}_{t-1}} = \{\mathcal{S}_{AB}^{C}, \mathcal{S}_{BC}^{D}, \mathcal{S}_{AD}^{E}\}$, then the subsets $\mathcal{S}^{\mathcal{M}_{t-1}^i} = \{\mathcal{S}_{AB}^{C}, \mathcal{S}_{BC}^{D}\}, \mathcal{H}_{t-1}^i = \{\mathcal{H}^{ABC}, \mathcal{H}^{BCD}\}$ are chosen as using the triplets $AB{:}C$ and $BC{:}D$, we can compose $AC{:}D$ via a single composition. We can describe the required composition using a binary tree, consisting of a root, representing the triplet to compose $AB{:}D$, and two leaves connected to it, representing the required source triplets, $AB{:}C$ and $BC{:}D$. The abovementioned tree, also known as *Composition Tree*, was first introduced in [23]. Depending on the identity of the target triplet and the available source set, we can generally get a large *Composition Tree*, representing the sequence of compositions required to form the target triplet. To find the tree with the minimal number of leaves (representing the set $\mathcal{S}^{\mathcal{M}_{t-1}^i}$), we use an Ad Hoc algorithm developed in [23]. The chosen tree dictates the required sequence of compositions, where each of them is then performed, inspired by 2.4, as follows:

$$\mathbb{P}(\mathcal{S}^{\tau_3}|\mathcal{S}^{\tau_1},\mathcal{S}^{\tau_2},\mathcal{H}^1,\mathcal{H}^2)=\iint\limits_{\mathcal{X}^{\tau_1}\in\mathcal{S}^{\tau_1}\ \mathcal{X}^{\tau_2}\in\mathcal{S}^{\tau_2}}\mathbb{P}(\mathcal{S}^{\tau_3}|\mathcal{X}^{\tau_1},\mathcal{X}^{\tau_2})\mathbb{P}(\mathcal{X}^{\tau_1},\mathcal{X}^{\tau_2}|\mathcal{S}^{\tau_1},\mathcal{S}^{\tau_1},\mathcal{H}_t^{\tau_1},\mathcal{H}_t^{\tau_2})\, d\mathcal{X}^{\tau_1}\, d\mathcal{X}^{\tau_2}, \tag{4.28}$$

where $\mathbb{P}(\mathcal{S}^{\tau_3}|\mathcal{X}^{\tau_1},\mathcal{X}^{\tau_2})$ is a simple deterministic geometric model. The metric prior term can be approximated via $\mathbb{P}(\mathcal{X}^{\tau_1},\mathcal{X}^{\tau_2}|\mathcal{S}^{\tau_1},\mathcal{S}^{\tau_1},\mathcal{H}_t^{\tau_1},\mathcal{H}_t^{\tau_2}) \approx \prod_{i=1}^{2}\mathbb{P}(\mathcal{X}^{\tau_i}|\mathcal{S}^{\tau_i},\mathcal{H}_t^{\tau_i})$, where $\forall i \in \{1,2\}$ the individual prior term can be further approximated via $\mathbb{P}(\mathcal{X}^{\tau_i}|\mathcal{S}^{\tau_i},\mathcal{H}_t^{\tau_i}) \approx \mathbb{P}(\mathcal{X}^{\tau_i}|\mathcal{S}^{\tau_i})$, i.e., assuming a uniform distribution.

We can now incorporate compositions within the belief update step from (4.25) by replacing b_{t-1} with $b_{t-1}^{\mathcal{M}}$ from (4.26).

4.2.7 Cost Function

The objective function in (4.3) describes for each candidate action sequence $a_{k:k+L-1}$ its expected accumulated cost. Having access to a qualitative belief allows defining both state-dependent and belief-dependent cost functions. The cost function is chosen according to the task's nature. For instance, if we aim to find the shortest path between some initial and goal key points, then a distance-based cost can be used, whereas if we wish to reduce uncertainty, then an entropy-based cost may be a good choice.

In this work we focus on the first type, and suggest now two alternatives for distance-based costs. The first is to evaluate the expected number of qualitative states traversals:

$$c_t(b_t,a_{t-1}) = \mathop{\mathbb{E}}_{s_1,s_2}\left[d(s_1,s_2)\right], \tag{4.29}$$

where $s_1 \triangleq \mathcal{S}_{F_{t-1}}^{X_{t-1}}$, $s_2 \triangleq \mathcal{S}_{F_{t-1}}^{X_t}$, and where $d(s_1, s_2)$ is a metric that returns the minimum number of qualitative states traversals required to travel from state s_1 to s_2. For a given QSR representation, this metric is a simple lookup table.

While the above cost measures distance qualitatively, there is no proven correlation to the true metric distance traveled. To bridge this gap, we propose another option of evaluating the traveled distance metrically, given the qualitative belief b_t. The idea is simple, at each time step, we evaluate the expected traveled distance in terms of the current local coordinates system, multiplied by the appropriate global scale taken from the belief. Mathematically:

$$c_t(b_t, a_{t-1})=\mathop{\mathbb{E}}_{s}\left[\mathop{\mathbb{E}}_{x|s}\left[\left\|\mathcal{X}_{F_{t-1}}^{X_t}-\mathcal{X}_{F_{t-1}}^{X_{t-1}}\right\|_2\cdot\mathcal{X}^{F_{t-1}}\right]\right], \tag{4.30}$$

where we denoted $x \triangleq \{\mathcal{X}_{F_{t-1}}^{X_t},\mathcal{X}_{F_{t-1}}^{X_{t-1}},\mathcal{X}^{F_{t-1}}\}$, and $s \triangleq \{\mathcal{S}_{F_{t-1}}^{X_t},\mathcal{S}_{F_{t-1}}^{X_{t-1}},\mathcal{S}^{F_{t-1}}\}$. Moreover, based on the Law of Total Expectation, we can further simplify the above expression

		Cost 1 (# q-states)		Cost 2 (metric path length)	
		W/O Comp	**W Comp**	**W/O Comp**	**W Comp**
All Tests (3000)	**Plan exists**	68.4%	81.6%	68.4%	81.6%
Comparable & Different Tests (13%)	**Average executed cost**	6.45	4.97	2.76	2.32

Table 4.1: The first row shows the percentage of tests (considering all 3000 tests) in which the robot was able to plan towards its goal. The second row shows the average cost among all different and comparable tests. The two main columns are divided by the different costs (Sec. 4.2.7), where each considers two planning modes: without and with compositions.

and replace it with a pure metric version:

$$c_t(b_t, a_{t-1}) = \mathop{\mathbb{E}}_{x}\left[\left\|\mathcal{X}_{F_{t-1}}^{X_t} - \mathcal{X}_{F_{t-1}}^{X_{t-1}}\right\|_2 \cdot \mathcal{X}^{F_{t-1}}\right]. \tag{4.31}$$

To calculate this cost in practice, we must use (4.30) as we only have access to the metric priors conditioned on the corresponding qualitative states.

4.3 Results

We evaluate our approach using a simulation developed in *Python* 3.7. We don't compare our performances to other algorithms, as there are no comparable ones in the literature. The only other existing work that considers planning under the qualitative framework, [16], assumed a deterministic framework making any comparison irrelevant. For simplicity, we consider the following: (1) Motion and Measurement models with additive gaussian white noise. (2) Small environments with 8−12 landmarks. (3) Data-Association is solved. (4) The initial belief considers a *Composable* set of triplets under $\mathbb{L}$.

We performed 3000 tests, each randomizing a different environment and choosing a target triplet randomly. We repeated each test four times, once for each cost suggested in 4.2.7, with and without the robot using compositions. Table 4.1 summarizes the overall statistics, where Fig. 4.5 emphasizes the advantage of using compositions.

Two specific scenarios are demonstrated in Fig. 4.6, where the robot aims to reach and observe the goal triplet ABC.

In the first scenario (Fig. 4.6a), which considered the cost formulated in Eq. (4.29), the robot has succeeded in reaching and observing ABC only via compositions. Specifically, the robot's initial frame was $F_1{=}EG$, and its initial belief was over the set of source landmark triplets $\mathcal{M}_1{=}\{ABC,ABD,ABE,ABH,ACI,DEF,DEG,EGJ\}$. At time step $t{=}2$, the robot linked to a new frame, $F_2{=}AE$. This link, which was necessary to reach ABC, was feasible exclusively based on the triplet AEG, which was composed using source triplets from $\mathcal{M}_1$. Meaning, in this case, without using compositions, no feasible path would have been found.

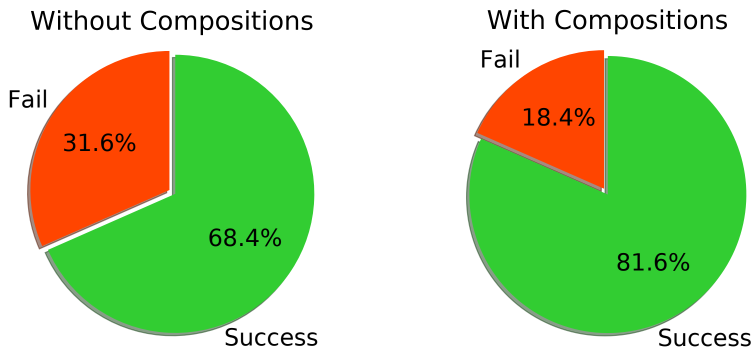

(a) (b)

Figure 4.5: Improvement in results due to the use of compositions. (a) Without compositions, the robot succeeded in planning towards its goal in 68.4% of the tests.; (b) Via compositions, the robot improved and planned towards its goal in 81.6% of the tests.

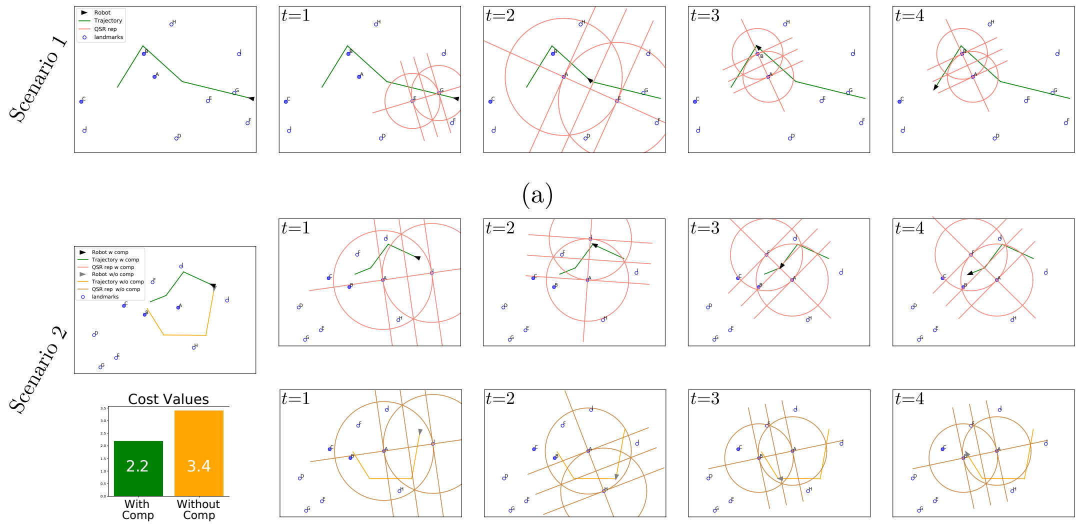

(b)

Figure 4.6: Two scenarios from our simulations are presented. We show the execution of the calculated paths with replanning between consecutive time steps. For each time step, the local frame of the robot, considering the EDC partitioning, is displayed. In both scenarios, the robot aims to reach and observe the goal triplet ABC (filled blue circles). (a) In scenario 1, the robot succeeded in reaching its goal only via compositions. (b) in scenario 2, the robot was able to find a path towards its goal with (upper row) and without (bottom row) compositions. However, via compositions, the path is shorter.

In the second scenario (Fig. 4.6b), which considered the cost formulated in Eq. 4.30, the robot was able to find a plan without using composition. However, by composing the triplet AIJ, it was able to link to AJ at time step t=2 and via that link to find a shorter path towards ABC's vicinity.

Chapter 5

Conclusions

In the first part of this research, we addressed two main issues that arise regarding compositions calculi. First, given a set of prior qualitative spatial relationships between triplets of landmarks (source triplets), we have formulated a sufficient topological condition (*Composability*) attributed to the set, whose existence ensures the ability to compose an entire space of spatial relationships between new triplets. Secondly, we addressed the question of how these triplets should be composed. We presented a novel algorithm that finds the optimal way to compose a target triplet under an optimality criterion defined by the user. To the best of our knowledge, this algorithm is the first of its kind.

This encourages incorporating our algorithm as a component in future qualitative approaches for a variety of tasks, such as localization, mapping, and active planning. The latter is addressed in the second part of this research. one can also formulate a cost function designed to turn a prior map (i.e., the set of qualitative relationships that form it) into a *Composable* one, which would allow it to be further expanded through composition. Future research works may also generalize our work to include qualitative relationships between pairs of landmarks as well (relaxing the assumption of point landmarks).

In the second part of this research, we presented a novel algorithm to address the problem of Belief Space Planning, considering a qualitative framework. Our algorithm operates in two steps. Given an initial qualitative belief and a target triplet, it first constructs a belief tree that accounts for multiple possibilities for future developments, where each corresponds to a candidate plan. This step is the main focus of this work, where the belief tree is specifically designed to support qualitative belief propagation. Then, it chooses the best plan, i.e., that minimizes a meaningful objective function. This step is a standard one. Moreover, our mechanism enables incorporating compositions to improve planning results. Finally, we suggested a novel cost function, which considers metric path length, thus being more realistic. We believe this first work on qualitative BSP opens new research opportunities to follow.

Chapter 6

Future Work

We believe that our research opens various future possibilities and ideas to investigate. We detail below several possible directions for expanding our work.

1. **Proving (or rejecting) the opposite direction of Theorem 3.1.**
   In Sec. 4.2.6, we proved that given a *Composable* set of source triplets under some *Landmark Space*, one could compose any triplet existing within this space (see Proof 1). We suspect that the opposite direction of this theorem is also true, i.e., that *Composability* is not only a sufficient condition to compose all triplets in a given *Landmark Space* but also a necessary one. However, we did not manage to prove it within the time limit of this research scope. We encourage others to try.

2. **Incorporating Alg. 3.3 as a component in future qualitative approaches.**
   Incorporating our mechanism to compose triplets given an initial set of source ones can be integrated as a component in future qualitative approaches for a variety of tasks, such as localization, mapping, and active planning. The latter is addressed in the second part of this research but can be incorporated within other qualitative planning paradigms as well.

3. **A cost function designed to turn a qualitative map *Composable*.**
   One can also formulate a cost function designed to turn a prior map (i.e., the set of qualitative relationships that form it) into a *Composable* one, which would allow it to be further expanded through composition.

4. **Study the effect of compositions on entropy.**
   In our research, we empirically saw that the entropy of a posterior distribution over a composed triplet is increasing as the number of composition operations required to create it grows. We encourage future studies to try and find a mathematical derivation that explains the above.

5. **Relaxing the assumption of point landmarks**
   Future research works may also generalize our work (in both parts) to include

qualitative relationships between pairs of landmarks (possible with volumed landmarks) as well (relaxing the assumption of point landmarks).

6. **Real life experiment using our qualitative BSP algorithm.**
   We highly encourage taking our qualitative belief space planning algorithm a step forward by operating it on a real-life platform.

Bibliography

[1] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics.* The MIT press, Cambridge, MA, 2005.

[2] M. Kaess, H. Johannsson, R. Roberts, V. Ila, J. Leonard, and F. Dellaert, "iSAM2: Incremental smoothing and mapping using the Bayes tree," *Intl. J. of Robotics Research*, vol. 31, pp. 217–236, Feb 2012.

[3] C. Freksa, "Using orientation information for qualitative spatial reasoning," in *Theories and methods of spatio-temporal reasoning in geographic space*, pp. 162–178, Springer, 1992.

[4] T. S. Levitt and D. T. Lawton, "Qualitative navigation for mobile robots," *Artificial intelligence*, vol. 44, no. 3, pp. 305–360, 1990.

[5] C. Schlieder, "Representing visible locations for qualitative navigation," in *Qualitative Reasoning and Decision Technologies*, pp. 523–532, CIMNE, 1993.

[6] T. Wagner, U. Visser, and O. Herzog, "Egocentric qualitative spatial knowledge representation for physical robots," *Robotics and Autonomous Systems*, vol. 49, no. 1, pp. 25–42, 2004.

[7] C. Freksa and K. Zimmermann, "On the utilization of spatial structures for cognitively plausible and efficient reasoning," in *IEEE International Conference on Systems, Man and Cybernetics*, pp. 261–266, 1992.

[8] C. Schlieder, "Reasoning about ordering," in *Spatial Information Theory A Theoretical Basis for GIS* (A. U. Frank and W. Kuhn, eds.), (Berlin, Heidelberg), pp. 341–349, Springer Berlin Heidelberg, 1995.

[9] R. Moratz, L. Dominik, and M. Till, "A condensed semantics for qualitative spatial reasoning about oriented straight line segments," 2011.

[10] T. Mossakowski and R. Moratz, "Qualitative reasoning about relative direction of oriented points," *Artificial Intelligence*, vol. 180-181, pp. 34–45, 2012.

[11] P. Fogliaroni, J. O. Wallgrün, E. Clementini, F. Tarquini, and D. Wolter, "A qualitative approach to localization and navigation based on visibility information," in

International Conference on Spatial Information Theory, pp. 312–329, Springer, 2009.

[12] M. McClelland, M. Campbell, and T. Estlin, "Qualitative relational mapping and navigation for planetary rovers," *Robotics and Autonomous Systems*, vol. 83, pp. 73–86, 2016.

[13] M. McClelland, M. Campbell, and T. Estlin, "Qualitative relational mapping for mobile robots with minimal sensing," *Journal of Aerospace Information Systems*, vol. 11, no. 8, pp. 497–511, 2014.

[14] J. Padgett and M. Campbell, "Probabilistic qualitative mapping for robots," *Robotics and Autonomous Systems*, vol. 98, pp. 292–306, 2017.

[15] R. Mor and V. Indelman, "Probabilistic qualitative localization and mapping," in *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2020.

[16] J. Padgett and M. Campbell, "Q-link: A general planning architecture for navigation with qualitative relational information," *Robotics and Autonomous Systems*, vol. 108, pp. 51–65, 2018.

[17] V. Indelman, L. Carlone, and F. Dellaert, "Planning in the continuous domain: A generalized belief space approach for autonomous navigation in unknown environments," *The International Journal of Robotics Research*, vol. 34, no. 7, pp. 849–882, 2015.

[18] E. I. Farhi and V. Indelman, "Towards efficient inference update through planning via jip—joint inference and belief space planning," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 4479–4486, IEEE, 2017.

[19] R. K. Goyal, *Similarity assessment for cardinal directions between extended spatial objects*. PhD thesis, in Spatial Information Science and Engineering, University of Maine, 2000.

[20] E. Dijkstra, "A note on two problems in connexion with graphs," *Numerische Mathematik*, vol. 1, pp. 269–271, 1959.

[21] G. Gallo, G. Longo, S. Pallottino, and S. Nguyen, "Directed hypergraphs and applications," *Discrete applied mathematics*, vol. 42, no. 2-3, pp. 177–201, 1993.

[22] S. Pathak, A. Thomas, and V. Indelman, "A unified framework for data association aware robust belief space planning and perception," *Intl. J. of Robotics Research*, vol. 32, no. 2-3, pp. 287–315, 2018.

[23] I. Zilberman, E. Rivlin, and V. Indelman, "Incorporating compositions in qualitative approaches," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 2660–2667, 2022.

ביצוע פעולת קומפוזיציה ספציפית אינו זמין, עליו להיות מוסק תחילה, ייתכן ע"י פעולת קומפוזיציה מקדימה. בעיה זו עשויה לחזור על עצמה שוב ושוב באופן רקורסיבי, והיא נעשית מאתגרת יותר ככל שכמות המידע גדלה. שתי שאלות מרכזיות עולות מן האמור לעיל, אשר נותרו פתוחות: 1. בהינתן סט התחלתי של קשרים מרחביים איכותיים, אילו קשרים חדשים ניתן ליצור בעזרת קומפוזיציה? 2. מהו רצף פעולות הקומפוזיציה האופטימלי ליצירת קשר נתון מתוך כל הרצפים האפשריים? בטרם ננסח את אלגוריתם התכנון שלנו, אנו עונים על שתי שאלות אלה ובכך מבהירים כיצד יש לשלב קומפוזיציות בגישות איכותיות באופן כללי.

תקציר

תכנון תחת אי-וודאות, הידוע גם כתכנון במרחב האמונה, הנו רכיב חיוני במערכות אוטונומיות ואפליקציות רובוטיות רבות, כדוגמת רכבים אוטונומיים, שואבי אבק ביתיים ומיכשור רפואי מתקדם. המטרה היא למצוא רצף אופטימלי של פעולות (דהיינו, תכנית) על מנת להשלים משימה מוגדרת מראש. לרוב, הרובוט מתחזק פילוג אוחר על המפה ומיקומיו העצמיים, הידוע גם כאמונה. פונקציית מחיר אזי מוגדרת על האמונה כמנגנון ענישה. התכנית הנבחרת הנה זו שככל הנראה תשיג את מטרת המשימה תוך מזעור פונקציית המחיר. לדוגמה, אם המטרה היא למצוא את המסלול הקצר ביותר אל היעד, נבחר פונקציית מחיר אשר מודדת את המרחק אל היעד. פתרון בעיה זו נעשה מאתגר במיוחד תחת תנאי אי-וודאות, כלומר, כאשר ייתכן שהמידע אינו שלם או שגוי בחלקו, ואותן פעולות עשויות לעיתים להוביל לתוצאות שאינן זהות. בעוד שרוב המחקר שנעשה על תכנון במרחב האמונה ניסח את הבעיה דרך ייצוג מטרי ישנן אלטרנטיבות אחרות. אחת פחות נפוצה הנה הניסוח האיכותי, בו המפה מיוצגת ע"י סט של קשרים מרחביים איכותיים במערכות צירים מקומיות ובלתי-תלויות. קשרים אלה מתארים את מיקום הרובוט או נקודות הציון במונחים של מצבים איכותיים ("מימין ובאמצע", "משמאל ולמעלה", וכו'.) ולא באופן מטרי מדויק. מאחר וגישה זו אינה מכוונת להשיג תיאור מדויק של הסביבה מלכתחילה, היא חסינה יותר בפני רעשים לעומת זאת המטרית. על כן, היא נעשית שימושית מאוד בהיעדר חיישנים באיכות גבוהה. בעוד מספר שיטות איכותיות ללוקליזציה ומיפוי קיימות בספרות, עד כה, ישנה רק גישה איכותית אחת עבור בעיית התכנון.

מונעים מהאמור לעיל, במחקר זה, אנו מציגים גישה חדשנית לתכנון במרחב האמונה תוך שימוש בקשרים איכותיים מרחביים, אשר מתאימה לפלטפורמות בעלות סנסורים זולים ורלוונטית במיוחד במקרים בהם הסביבה דלילה. באלגוריתם המוצע, אנו מגדירים תחילה אמונה המותאמת למסגרת העבודה האיכותית, כלומר, כזו המוגדרת על סט קשרים מרחביים איכותיים המתארים את המפה ומיקומי הרובוט. על בסיס אמונה זו, הרובוט בונה עץ אמונה, בו כל ענף מתאר התפתחות אחת אפשרית עבור סט פעולות נבחן. נציין כי סטי הפעולות הנבחנים כוללים סוג ייחודי של פעולות הנקראות "לינק", אשר מאפשרות לרובוט לעבור ממערכת ייחוס אחרת לאחרת. הענפים בעץ מוענשים ע"י פונקציית המחיר ובכך מודדים עד כמה סטי הפעולות אותם הם מייצגים הנם כדאיים. לבסוף, תוך ביצוע תהליך אופטימיזציה המנצל את מבנה העץ, סט הפעולות הכדאי ביותר נבחר. נציין גם כי כחלק ממחקר זה פותחה פונקציית מחיר ייחודית, אשר משערכת את המרחק המטרי אל היעד באופן גס על בסיס האמונה איכותית בלבד. בנוסף, באופן חסר תקדים, אלגוריתם התכנון שלנו משלב באופן אינטגרטיבי קומפוזיציות, אשר מאפשרות חלחול מידע מרחבי בין קשרים איכותיים שונים לצורך הסקת חדשים.

אולם, שילוב קומפוזיציות בגישה איכותית אינו דבר של מה בכך. לדוגמה, אם מידע אשר דרוש לצורך

המחקר בוצע בהנחייתו של פרופ' חבר ואדים אינדלמן ופרופ' אהוד ריבלין בפקולטה להנדסת חשמל.

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

1. I. Zilberman, E. Rivlin, and V. Indelman, “Incorporating compositions in qualitative approaches,” IEEE Robotics and Automation Letters, vol. 7, no. 2, pp. 2660−2667, 2022.
2. I. Zilberman and V. Indelman, “Qualitative belief space planning via compositions,” in IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS), 2020. submitted.

תודות

ברצוני להודות לפרופ' חבר ואדים אינדלמן על הנחייתו במהלך המאסטר, שהיתה מלמדת ונעימה גם יחדיו. בנוסף, ברצוני להודות למשפחתי על תמיכתם ועל שדחפו אותי קדימה להצליח.

אני מודה לטכניון על התמיכה הכספית הנדיבה בהשתלמותי.

תכנון במרחב אי-וודאות בעזרת קשרים מרחביים איכותיים

חיבור על מחקר

לשם מילוי חלקי של הדרישות לקבלת התואר
מגיסטר למדעים בהנדסת חשמל

איתי זילברמן

הוגש לסנט הטכניון — מכון טכנולוגי לישראל
אדר התשפ"ב חיפה פברואר 2022

תכנון במרחב אי-וודאות בעזרת קשרים מרחביים איכותיים

איתי זילברמן