Università degli Studi di Roma “La Sapienza”
Dipartimento di Informatica e Sistemistica “Antonio Ruberti”
Tesi di Dottorato in Ingegneria Informatica, XVII ciclo, 2005

Topological Mapping of Ambiguous Space: Combining Qualitative Biases and Metrical Information

Francesco Savelli

Doctoral Thesis in Computer Engineering.

Thesis Committee:
Prof. Fiora Pirri (Advisor), Univ. di Roma “La Sapienza”
Prof. Giuseppe De Giacomo, Univ. di Roma “La Sapienza”
Prof. Roberto Pirrone, Univ. di Palermo

External Reviewers:
Prof. Benjamin Kuipers, Univ. of Texas at Austin
Prof. Alessandro Saffiotti, Örebro Univ. Sweden

Abstract

The question “Have I already been here, or is it the first time I see this place?” offers a paradigmatic example of the *topological* spatial uncertainty that may arise when exploring a new environment. This type of uncertainty concerns both the number of places encountered, and the order in which they are visited. Its occurrence requires taking into account different hypotheses about the size and loops of the graph-model of places (nodes) and paths (edges) that abstracts the space. Because of their limited perceptual mechanisms, modern robotic systems are often faced with topological ambiguity, which makes the autonomous acquisition of a reliable map of the environment a particularly difficult task.

A different type of spatial uncertainty regards the exact geometrical layout of the environment in a single global frame of reference. Most approaches to map-building in robotics are primarily concerned with this second kind of uncertainty, and look at topological ambiguity as an additional adversity — the *correspondence* or *data association* problem — for which they either assume to be given a solution *a priori*, or devise *ad hoc* methods.

The main assumption that drives our work is that *metrical uncertainty* and *topological ambiguity* factorize the spatial uncertainty arising in the problem of map-building. They can thus be handled in isolation, but their partial solutions can also take advantage of each other and offer complementary benefits.

In this direction, we provide theory for combining modern metrical mapping methods, among the best of the state of the art, along with a customizable system of qualitative biases and ontological expectations that has proved relevant to topological mapping. The latter has been previously studied in the framework of the *Spatial Semantic Hierarchy* (SSH), which plays a foundational role in this thesis. To this purpose, we make use of probability theory, and in particular of the *Bayes Networks*, which we believe to account best for the different nature of metrical and topological uncertainty. We show experimentally the advantages of the approach. In the spirit of the SSH ontology, we also study the dramatic reduction of topological ambiguity that follows from enforcing the planarity constraint, which is adequate in many real-world cases.

Acknowledgments

I am deeply indebted to Fiora Pirri and Benjamin Kuipers, who have advised me during my studies and work leading to this thesis. Fiora is an inspiring example of how much one can devote to a research life, taking the necessary risk of following a scientific problem to its source, whatever the cost. Ben does an impressive job of advising and helping young students become mature scientists in every respect, encouraging them to select and address important research questions. This thesis owes a lot to his patient and intelligent guidance, his open-mindedness, and to the important research experience I had in his lab. I want to extend my thanks to his students Patrick Beeson, Philip Hendrix, Matt MacMahon, Joseph Modayil, and Jefferson Provost. They are all remarkable people, and working with them is fun as well as a source of inspiration.

Ben and Alessandro Saffiotti have served as reviewers of this thesis. I thank them very much for the many meaningful comments and careful proofreading.

Camil Demetrescu and my sister Alessandra Savelli have given me good suggestions concerning mathematical aspects of planarity of graphs, an important element in this thesis.

As a more experienced colleague and a person of deep intellect, Alberto Finzi has been a great fellow traveler along this path toward knowledge. Our countless and endless discussions have left me with clearer thinking and understanding of the research in our field.

Looking back, I see that it is impossible to disentangle my work from my life, especially in such crucial years as those of graduate studies. I want to thank two persons who I think have some responsibility for what I am now which I wasn't three years ago. They provided me with suggestions and support whenever the need arose. Marcelo Oglietti is a special and dear friend; much of his help unexpectedly came while the two of us were discussing any topic on Earth at any chance. Desirée Ward has brought brand new colors into my life.

Last, and most important, this thesis is dedicated to my parents. They have contributed to my intellectual life in the first place and taught me so much, and with so much love, about the struggles and joys of living.

To my parents

Contents

1 Introduction **13**
- 1.1 Motivations and Contributions . . . 13
- 1.2 Robot Map-Building . . . 17
  - 1.2.1 Metrical Maps . . . 18
  - 1.2.2 Topological Maps . . . 19
  - 1.2.3 Hybrid Maps . . . 19
  - 1.2.4 Cognitive Maps . . . 20
- 1.3 Thesis Outline . . . 22

2 The Hybrid SSH **23**
- 2.1 The Spatial Semantic Hierarchy . . . 23
  - 2.1.1 Sensory and Control Levels . . . 25
  - 2.1.2 Causal Level . . . 26
  - 2.1.3 Topological Level . . . 28
  - 2.1.4 Metrical Level — Global Metrical Mapping . . . 30
- 2.2 Accounting for Small-Scale Space . . . 30
  - 2.2.1 LPMs and Stars . . . 30
  - 2.2.2 Gateways and Stars from Occupancy-Grid Maps . . . 33
- 2.3 A Topological Mapper . . . 36
  - 2.3.1 Algorithm . . . 37
  - 2.3.2 Preference Policy . . . 43
  - 2.3.3 Advantages . . . 43
- 2.4 A Physical Exploration . . . 45
- 2.5 Related Work . . . 48
  - 2.5.1 Similar Approaches to Hybrid Mapping . . . 48
  - 2.5.2 Place Detection, Classification, and Recognition . . . 49

3 Loop-Closing and Planarity **51**
- 3.1 Introduction . . . 51

3.2 Embedded Graphs and Planarity . . . 52
3.3 Search-Space Reduction . . . 56
3.4 Experimental Results . . . 58
3.5 Related Work on Loop-Closing . . . 63

4 Topological Reasoning and Metrical Analysis 65
4.1 Introduction . . . 65
4.2 Topological Hypotheses as Bayes Networks . . . 67
4.3 Characterizing Maps Probability . . . 71
4.4 Computing the Metrical Likelihood . . . 74
4.4.1 More about the Approximation . . . 76
4.5 From the SSH to the Prior Probability . . . 77
4.6 Maximizing the Log-Posterior Probability . . . 78
4.7 Experimental Results . . . 79
4.7.1 Parameters . . . 79
4.7.2 Results . . . 82
4.8 Related Work . . . 86
4.8.1 Correspondence in Metrical Methods . . . 86
4.8.2 Correspondence in Hybrid Methods . . . 86

5 Conclusions 89

Bibliography 92

A Bayes Networks 101

B Global Poses Computation 103
B.1 Linearization . . . 103
B.2 Minimization . . . 104

Chapter 1

Introduction

1.1 Motivations and Contributions

While trying to build a map of an unknown environment under exploration, the place being currently visited may seem identical to one (or more) encountered earlier. When this is the case, a paradigmatic question arises: "Have I already been here (was this a loop), or is it the first time I see this place?". Obviously, the occurrence and extent of this ambiguity depend both on the perceptual processes employed, and the environment at hand. This problem is indeed a case of *perceptual aliasing* occurring over *large-scale space*. Perceptual aliasing regards two distinct entities when they cannot be distinguished by the perceptual mechanisms. "Large-scale" connotes a space that cannot be observed at once from a single viewpoint, and implies that its global description can be only built out of many observations collected over time [1].

The problem simply defined by the question above seriously affects the process of assimilating purposeful spatial representations of physical environments. Unfortunately, it has so far proved very hard for modern robotic systems, partially due to the their limited perceptual capabilities. Unless strong assumptions are made on the environment and the sensing model, for example the presence of properly located landmarks that are pairwise distinguishable, any kind of approach to map-building is prone to perceptual aliasing and must deal with large-scale spatial ambiguity.

This type of uncertainty is *topological*, in that it concerns the graph-model of places (nodes) and paths (edges) that abstracts the space. In particular, the number of places and the order in which they have been visited, which makes the loops of the graph, are uncertain. To deal with topological ambiguity, more hypotheses need to be explicitly considered. Unfortunately, the

[1] Quoting from Kuipers (1977, chapt. 1): "Large-scale space is defined by the perceptual mechanisms for exploring the space, rather than by its physical size. A large-scale space is defined as a space which cannot be perceived at once: its global structure must be derived from local observations over time. For example, a drawing is a large-scale space when viewed through a small movable hole, while a city can be small-scale when viewed from an airplane."

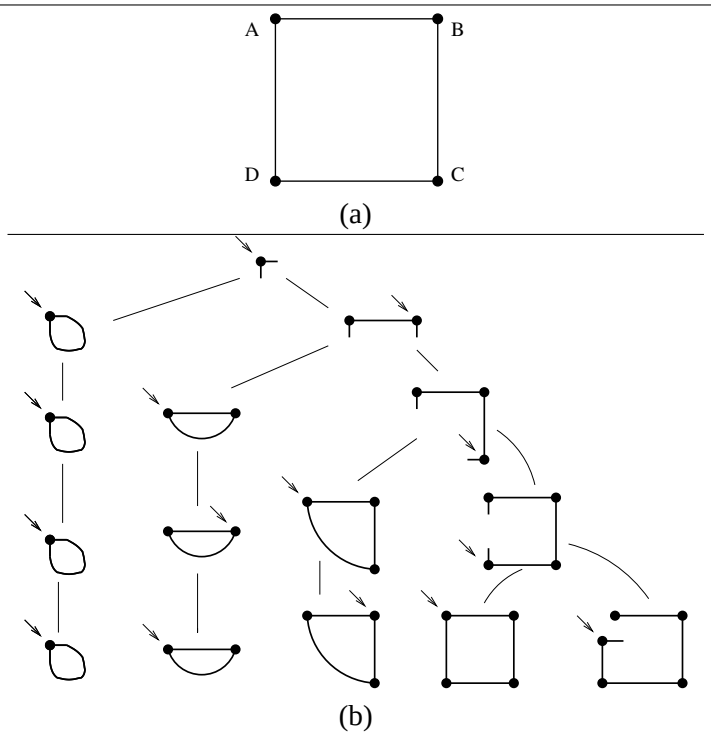

Figure 1.1: **Topological Search.** (a) The environment. The exploration considered is A-B-C-D. Assume the robot is able to follow corridors (lines) and recognize the corners as places. All the places are perceptually aliased as they appear identical (a corner). (b) Topological hypotheses from the exploration. Solid-filled points are the places assumed in each map. A little closed-headed arrow points to the assumed current place in each map.

space of hypotheses can grow exponentially, as shown in the simple case in Figure 1.1. This happens, for example, in the relevant class of environments whose topological abstraction is a grid-like pattern. In this case, discarding wrong models turns out to be particularly hard due to the symmetrical layout of perceptually aliased places, as illustrated in Figure 1.2.

Most current approaches to robotic map-building are primarily concerned with another kind of spatial uncertainty, about the precise geometrical layout of particular landmarks and features, or of any obstacle at a pixel resolution, in a single global frame of reference. The hard problem here is to enforce global metrical consistency despite the inevitable sensing and odometrical error, especially because the latter tends to scale up dramatically with the dimensions of the environment. Such approaches look at topological uncertainty as an additional adversity, usually referred to as the *data association* or *correspondence* problem. When they do not assume to be

given a solution for that in advance, these approaches devise *ad hoc* methods to deal with it. Sometimes a huge number of hypotheses are examined as samples of a probability distribution over the space of maps, and not directly as natural explanations of perceptual-aliasing events. A convergence behavior is usually pursued here. In other cases some greedy selection criteria are employed to overcome the ambiguity as soon as it arises. These techniques do not always commit to identifying the unique correct solution, and might end up with just an approximated description of the environment, for example hypothesizing a number of landmarks close to the correct one.

Our work builds on the idea that topological and metrical uncertainty can be addressed in isolation, but their partial solutions can take advantage of each other and offer complementary benefits. We adopt the *Spatial Semantic Hierarchy* (SSH) as foundational framework. The SSH represents different kinds of spatial knowledge — encompassing raw sensorimotor information and qualitative symbolic descriptions — over a hierarchy of ontologically distinct levels. It has proved relevant to a diverse body of research in the fields of Cognitive Science, Artificial Intelligence and Robotics, and has led to real robot implementations. A customizable system of qualitative biases and ontological expectations for typical classes of real-world topological configurations has been studied in previous work in the SSH. We address the problem of combining uncertain metrical information with such qualitative criteria, to the purpose of overcoming topological ambiguity more robustly. We make use of probability theory to provide a formal account of this integration, and to devise efficient algorithms after suitable approximations. Physical topological ambiguity is also ambiguity about the correct topology of the conditional (in)dependence relationships that characterize the probabilistic model for metrical uncertainty. This is best shown using *Bayes Networks* as graphical representations of the probabilistic models at hand. Thus, it can be pointed out that metrical uncertainty is represented within the model, and can be tackled by probabilistic inference, while topological uncertainty causes the probabilistic model itself to be uncertain, posing a harder challenge. On a more practical side, the framework proposed is compatible with some of the best modern metrical mapping methods. In the same spirit of the SSH, we also enhance its system of qualitative preferences studying the dramatic reduction of topological ambiguity when the planarity constraint — many interesting real-world environments meet the planarity assumption — is enforced.

In Section 1.2 we introduce the main paradigms of the approaches to the problem of robot map-building. In 1.3 we outline the organization of the thesis.

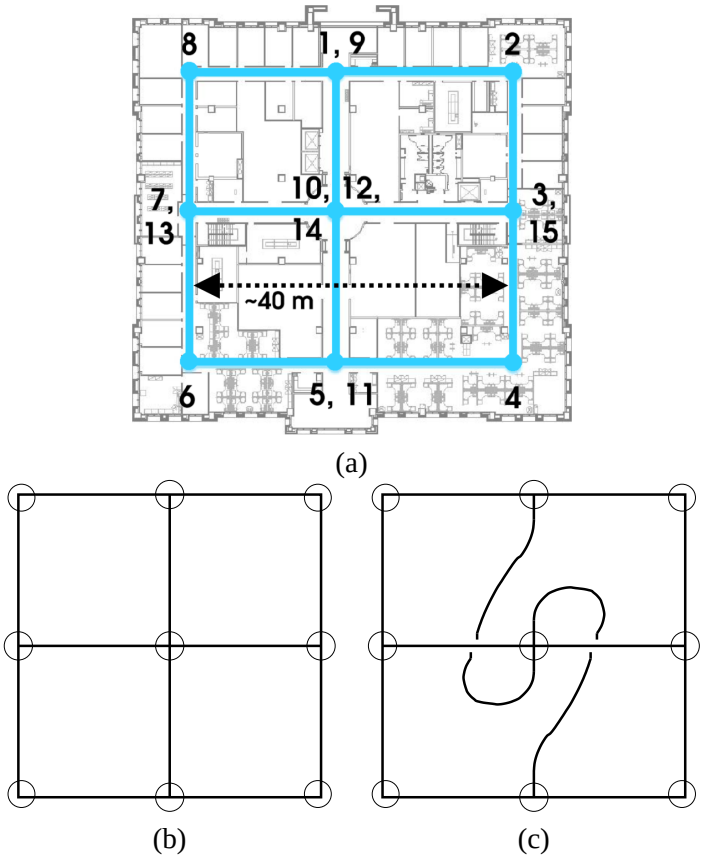

Figure 1.2: **Topological Ambiguity.** Both (b) and (c) are among the consistent topological hypotheses after the complete physical robot exploration enumerated in (a) (third floor of ACES building, University of Texas at Austin), and a topological search as in Figure 1.1. Places and their reliable local perceptual characteristics *at the topological abstraction level* are represented inside the small circles. Due to perceptual aliasing of places and the symmetrical structure of the environment, any further exploration will give the same sequence of observations, making it impossible to discriminate between the two maps. (This very case can be easily solved by enforcing the planarity constraint, see Chapter 3.)

1.2 Robot Map-Building

Two paradigms for map-building have been mainly pursued in Artificial Intelligence (AI) and Robotics, along with their integration.

A *metrical* map represents the environment by collecting the position of relevant landmarks, features, and objects w.r.t. a single metric frame of reference, *i.e.* a global coordinate system. In some cases, free and occupied space can be represented at high resolution, giving rise to a pixel map, called *occupancy grid*, which does not require a prior ontology and operational definition of objects or landmarks.

A *topological* map generally represents spatial knowledge as a graph, describing locations, places and objects of interest as nodes, and their spatial relations, such as proximity and links as edges. Often edges also reflect the procedural knowledge and control laws used to navigate between nodes.

Many hybrid approaches that integrate metrical and topological information have been proposed as well. They are usually aimed at extracting a map of one kind from one of the other. For example, a topological map can be used to drive the integration of local metrical maps in a single global frame of reference.

While seeking to build a correct map, whatever its nature, two types of uncertainty influence each other dramatically, making map-building a challenging task. One regards the current location in the map, which is necessary to correctly integrate new sensed spatial information, and update the map. The other concerns the map itself. Note that this distinction is different from the classification of spatial uncertainty as topological or metrical.

Current terminology, *Simultaneous Localization And Mapping* (SLAM) or *Concurrent Mapping and Localization* (CML), makes the interplay between localization and mapping explicit. Although this terminology is widely used in metrical mapping literature (for example, Thrun et al., 1998a; Dissanayake et al., 2001), it can apply to topological mapping as well (Choset and Nagatani, 2001).

In the metrical context, localization regards the problem of estimating the robot's pose (x, y, θ), *i.e.*, its position and orientation. Mapping itself consists of building the map relying on the exact pose, w.r.t. the global metrical frame of reference. The hard problem here is enforcing global metrical consistency, due to imprecision in sensing, motor control, and low-level estimation of the robot's own motion (odometry). The resulting accumulated error in the map tends to scale up dramatically with the dimensions of the explored area.

In a simple topological model, localization is deciding what node or edge the current place or path belongs to, while the mapping stage amounts to building the map as a graph.

However, "mapping" and "map-building" commonly refer to the entire SLAM problem, and we will use these two terms in this sense, unless otherwise specified.

Next sections discuss metrical, topological, and hybrid map-building in more detail.

1.2.1 Metrical Maps

Researchers have worked extensively on making metrical methods computationally feasible for building large maps within a single, global coordinate system. Among the first successful and influential approaches, Chatila and Laumond (1985) propose a metrical mapping algorithm that uses sets of polyehedra to describe the geometry of the environment, while Elfes (1989) and Moravec (1988) introduce the occupancy-grid framework. Methods to create metrical maps using a set of landmarks or geometric features have been widely explored (Smith et al., 1990; Dissanayake et al., 2001; Montemerlo et al., 2002).

Metrical maps excel in handling some of the initial, low-level problems encountered in robotics. Since these methods are often used with high-precision sensors, accurate localization and mapping can be efficiently accomplished in small-scale space (Gutmann and Schlegel, 1996; Gutmann et al., 1998; Gutmann and Fox, 2002). This reduces the effect of odometry error on pose estimation. Furthermore, recent probabilistic methods make metrical mapping robust and fast also in interestingly large environments, such as offices, under suitable operational conditions (Thrun, 2000, 2002).

Current metrical mapping methods, however, have several disadvantages. Metrical maps reduce pose error in local space, but errors dramatically propagate over large-scale space. As a consequence, metrical maps cannot easily handle cyclical environments once position estimates have drifted sufficiently. Mapping and planning in very large metrical maps can be time consuming, thus algorithms are often run offline. Metrical maps also suffer from the lack of a good interface for higher-level symbolic problem solvers. They are insufficient for a robot to reason about the layout of its environment, or to communicate route directions to another robot that lacks the same map.

Perceptual aliasing makes it difficult to handle loop-closure events. In principle, metrical mapping would be especially useful when the environment is perceptually poor w.r.t. the robot's perceptual skills. If successive robot's poses are known with sufficient accuracy, a loop can be detected when two such poses are close enough. However, most online successful approaches (see for example Paskin, 2003) assume to know which loops must be closed, for example by expecting that no two landmarks can be mistaken one for the other. This cycle in the structure of the problem confirms the importance of addressing perceptual aliasing in large-scale spatial knowledge acquisition.

Approaches that do not make use of probabilistic techniques have also been proposed, for example based on *fuzzy logic* (Driankov and Saffiotti, 2001). Fuzzy logic makes it possible to deal with heterogeneous uncertain spatial knowledge (metrical, topological etc.) in a very flexible fashion, at a low computational cost.

1.2.2 Topological Maps

A topological map represents an environment as a graph. Several distinct topological mapping frameworks have been proposed (Kuipers and Byun, 1991; Mataric, 1992; Shatkay and Kaelbling, 1997), which differ in semantics and ontology, in particular about what makes a place, and what should be considered a node or an edge in the graph. Some of them build topological maps autonomously, others only explore autonomously while the researcher provides place names to overcome perceptual aliasing.

Topological maps are more compact representations than global metrical maps. They allow high-level symbolic reasoning for map-building, navigation, planning, and communication. Since the environment is abstracted to a graph, movement errors that accumulate between graph nodes do not necessarily accumulate across a global frame of reference.

A recent topological mapping approach (Choset and Nagatani, 2001) employs Voronoi graphs to obtain the discrete representation of the environment from sensor data. Huang and Beevers (2004) address distributed multi-robot topological mapping; pattern matching on graphs is exploited to evaluate different hypotheses about how to merge different partial topological maps, built by different robots, that potentially overlap or match.

It is noteworthy that cognitive map research supports the creation of topological maps of large, complex environments (Yeap, 1988; Chown et al., 1995; Kuipers, 2000). Cognitive maps will be shortly introduced in Section 1.2.4.

1.2.3 Hybrid Maps

Arguably, the most promising approaches to robot map-building are based on hybrid topological/metrical maps. Kuipers and Byun (1991) introduce the concept of "patchwork metrical map", created using the topological map as a base for integrating metrical data gathered locally at places and along paths. Local frames of reference at place neighborhoods and along path segments are relaxed into a single global frame of reference, minimizing the "strain" at their joints. More sophisticated approaches, based on probabilistic techniques, have been later proposed (Thrun et al., 1998b; Duckett and Saffiotti, 2000; Modayil et al., 2004).

Thrun (1998) addresses the opposite direction, creating a topological map from a global metrical map, for indoor environments. This implies the aforementioned metrical map scaling problems. Related work connects the compact representations of rooms into both global metrical and topological maps (Yeap and Jefferies, 1999).

Fabrizi and Saffiotti (2000, 2002) employ fuzzy digital image processing to extract topological maps from occupancy grids, a mathematical technique rather more general and sophisticated than that proposed by Thrun (1998). It does not require a global metrical map, and can thus work online during exploration.

Most work on hybrid maps has dealt with generating topological and metrical maps as disjoint, sequential processes. Some recent research integrates metrical and topological mapping

by comparing local metrical models (Duckett and Nehmzow, 1999; Kuipers et al., 2004) or utilizing odometry (Tomatis et al., 2003; Bosse et al., 2003) to reduce topological place aliasing. Once a correct topological map is built, the local metrical models can be pieced together to make a global metrical map, by the methods mentioned above.

Following (Choset and Nagatani, 2001), Lisien et al. (2003) propose a hierarchical exploration and mapping framework which makes use of local metrical feature maps, and a global topological map. Although a promising approach, in the experiments loops were assumed to be notified externally during exploration.

Buschka and Saffiotti (2004) propose a principled ontology and theory of hybrid maps, discussing their use and properties from a general viewpoint.

1.2.4 Cognitive Maps

Cognitive Science has long been concerned with spatial reasoning and knowledge representation. Broadly speaking, this discipline aims at understanding and modeling the representational and computational mechanisms that underlie cognitive processes in humans and animals (*cognitive modeling*). In this respect, it is related to AI and Robotics, based on the long-term goal of reproducing human-level abilities and intelligence in agents that physically interact with the world.

In the light of their recent developments, AI, Robotics, and Cognitive Science appear to find an opportunity of taking further advantage from each other. In Cognitive Science, implementing a model and running it against data drawn from experiments with humans is a method commonly used to test the plausibility of such a model (*cognitive simulation*). On this side, advanced robotic platforms lend themselves to suitable testbeds for cognitive simulation of tasks involving physical interaction with the world, as in our case of spatial knowledge acquisition. On the other side, although AI and Robotics are not necessarily committed to having recourse to the very solutions Nature has devised, insights and ideas coming from the effort to model such solutions turn out to be of great value to artificial systems as well [2].

A *cognitive map* is a general cognitive model of large-scale spatial knowledge assimilation and representation, and of the mechanisms for reasoning with it. From related cognitive studies (Lynch, 1960; Piaget and Inhelder, 1967; Hart and Moore, 1973; Siegel and White, 1975), it emerges that cognitive maps have strong topological nature, and are grounded in sensorimotor experience and procedural knowledge of how to move between distinctive locations. Besides, humans make use of different kinds of spatial knowledge, depending on the particular task at hand.

Studies and models of cognitive mapping include (Kuipers, 1977; Yeap, 1988; Chown et al., 1995). We briefly survey the Spatial Semantic Hierarchy (SSH) below, because of its important

[2] Such a twofold statement of interest also offers a rather broad ontological and semantic frame of reference to the term *Cognitive Robotics*, currently used with different meanings in distinct research areas.

role in this thesis. Although primarily aimed at a general robotic framework, its scope significantly overlaps the theme of cognitive mapping. A more detailed description of the SSH will be given in Chapter 2.

The SSH

The Spatial Semantic Hierarchy (SSH) (Kuipers and Levitt, 1988; Kuipers and Byun, 1991; Kuipers, 2000) provides a framework in which spatial knowledge is represented over multiple levels. Such levels have different ontological nature, interact toward attaining a purposeful description of the environment and can work with partial knowledge. The different levels describe different kinds of information — applicable control laws, procedural knowledge, topological configurations — by different representational means, such as differential equations and symbolic descriptions, the latter including graphs and logical theories. This information has some degree of redundancy, because it is related to different aspects of the same phenomena, and is organized in a hierarchy of abstraction and dependencies over the different levels.

The SSH has proved relevant to a diverse body of research in AI, Robotics (Kuipers and Byun, 1991; Pierce and Kuipers, 1997) and Cognitive Science (Kuipers et al., 2003). The aspect that is most relevant to our purposes is the abduction of a topological map, possibly annotated with local metrical information, from the sensorimotor interaction of the robot with the environment.

Hybrid mapping methods based on the SSH allow metrical maps of local regions to be linked into topological maps of large-scale space. Increasingly efficient algorithms have been developed to exploit structure obtained from local metrical models (Remolina and Kuipers, 2004; Beeson et al., 2003; Kuipers et al., 2004; Savelli and Kuipers, 2004). These algorithms generate all possible topological models, filter out those inconsistent with the exploration and with the SSH theory, and provide a preference ordering on the remaining models.

In the SSH framework, metrical information can be very useful, but is not usually necessary for building a purposeful map, provided that the environment is not adversely pathological. This is not the case if structural ambiguity (*e.g.*, raised by symmetric perceptually-aliased places as in Figure 1.2) allows incorrect maps to be consistent with all past and all possible future experiences during exploration. In the worst cases, if the potential of odometrical information to rule out wrong place-association hypotheses is not exploited, coping with the combinatorial explosion of the search-space may not be practically possible. The recourse to such information, even when it is affected by strong uncertainty, can significantly help to tame the proliferation of candidates (we shall deal with this in Chapter 4).

1.3 Thesis Outline

This thesis is organized as follows.

In Chapter 2 we provide a thorough introduction of the Spatial Semantic Hierarchy, and of its recent extension, the Hybrid Spatial Semantic Hierarchy, to which we also contributed (Beeson et al., 2003; Kuipers et al., 2004). We improve and extend the mathematical formalization of the topological mapping function introduced in this previous work, and describe the topological mapper we have implemented, which will play an important role in the next chapters. This constitutes the foundations of the thesis.

In Chapter 3, we provide a theoretical and experimental analysis of the great impact that the planarity assumption (reasonable for many real-world environments) has on the reduction of topological ambiguity (Savelli and Kuipers, 2004). This lends a further argument to the use of discrete abstract representations of an environment: many (sometimes most) incorrect models can be discarded by a sublinear-time test on a graph *before* enforcing global metrical consistency, which is more computationally and representationally expensive.

In Chapter 4, we formalize the problem of solving topological ambiguity as a particular problem of Bayes-network learning, and derive a general framework in which metrical information and the qualitative biases provided by the SSH can be exploited together. Several approximations are proposed that make it possible to employ some techniques recently introduced in literature, among the best of the state of the art, for solving metrical pose uncertainty due to noisy odometry and sensing.

Finally, in Chapter 5, we provide some concluding remarks.

Chapter 2

The Hybrid SSH

The Spatial Semantic Hierarchy (SSH) plays a foundational role in this thesis. In this chapter we first introduce its framework (Section 2.1), then we describe a recent extension, called the hybrid SSH (from Section 2.2 on), to which we contributed (Beeson et al., 2003; Kuipers et al., 2004). Building accurate local representations is made computationally inexpensive by state-of-the-art methods such as metrical mapping techniques; the hybrid SSH takes advantage of such methods, and permits more effective large-scale topological mapping.

2.1 The Spatial Semantic Hierarchy

The Spatial Semantic Hierarchy (SSH) (Kuipers, 2000) provides a framework in which large-scale spatial knowledge is represented over different ontological levels. Inspired by cognitive findings about how humans assimilate and process large-scale spatial information, the SSH describes a computational model that includes topological maps abduction from the sensorimotor experience of a robot. It also allows for metric knowledge in the form of local annotations on the map and local frames of references integrated in a "patchwork". While the SSH makes some assumptions on the sensorimotor level of the robot, it does so in order to be independent from the details of such a level.

In the SSH, the explored space is partitioned into segments and regions that are qualitatively uniform w.r.t. the dynamical system composed of the robot, its control laws and sensor system, the conditions that trigger the application of the control laws, and the environment. In a sense, the SSH attempts to extract and exploit the fullest degree of determinism inherent in this system. The edges of a topological map thus provide spatial relations of proximity that are deterministic, in the sense that the robot can always move along them reproducing the same navigation results.

There are four levels of knowledge representation in the SSH, structured in a hierarchy of abstraction and dependencies, as illustrated in Figure 2.1. At the *control* level, the agent repeatedly selects a hill-climbing control law to converge to and localize at a *distinctive state* (dstate),

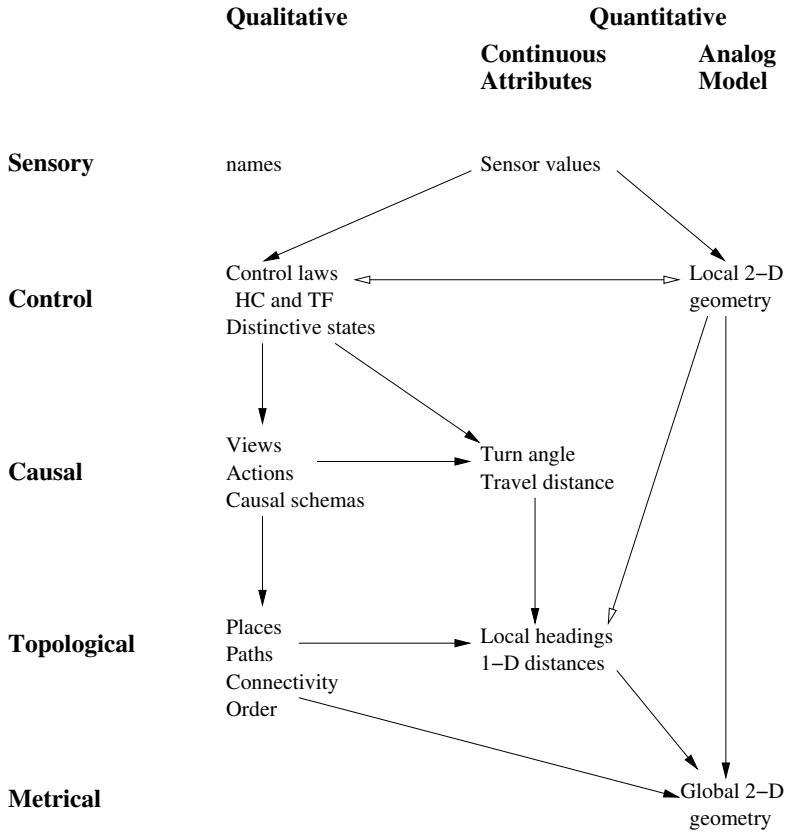

Figure 2.1: **The Spatial Semantic Hierarchy** Closed-headed arrows represent dependencies; open-headed arrows represent potential information flow without dependency. (Taken from Kuipers, 2000.)

and then a trajectory-following control law to move from the current dstate to the neighborhood of another, where hill-climbing converges to the next dstate, eliminating cumulative error. The *causal* level abstracts this pattern of behavior to a deterministic automaton, consisting of states (the distinctive ones), *actions* (sequences of control laws), *schemas* $\langle x, a, x' \rangle$ (asserting that state x' results from performing action a in state x), and *views* (the perceptual images of states, $view(x, v)$). The *topological* level distinguishes between turn and travel actions, and aggregates states into places, paths, and regions, related by connectivity, order, and containment. The *metrical* level consists of local metrical attributes for objects at the causal and topological levels, local metrical models of small-scale space in place neighborhoods, and (when resources permit) global metrical models of the large-scale environment (Modayil et al., 2004). A formalization of the topological map in non-monotonic logic, and an algorithm for identifying minimal models according to a prioritized circumscription policy is given in (Remolina and Kuipers, 2004). The circumscription policy formalizes preference and default criteria, which are well suited for the ontology of many real-world classes of environments such as offices or urban street networks.

We introduce the SSH's levels in more detail in the following sections. For a complete description see (Kuipers, 2000).

2.1.1 Sensory and Control Levels

The *sensory level* regards the interface with the agent's sensory system. Both continuous (from laser range finders, sonars etc.) and discrete (from names and directions provided, output from high-level vision) kinds of input fit this level.

The *control level* includes control laws, conditions for their applications and procedural knowledge for moving within small-scale space. It has the crucial role of abstracting the sensorimotor experience into a discrete description that permits symbolic reasoning at the causal and topological levels. It does so considering the agent and the environment a whole dynamic system w.r.t. the application of the control laws. A set of spatial segments that are qualitatively uniform for this system supports the construction of a net of distinctive states in the large-scale space. For a segment of the space to be qualitatively uniform means that the execution of a given law (e.g. "approach the visual target") always reproduces a predictable result, taking the agent to a distinctive location. In this sense, the resulting spatial description also reflects the procedural knowledge of the robot in the environment.

The identification of distinctive states and their connections is generally accomplished alternating two kinds of control laws: *hill-climbing* and *trajectory-following*.

Hill-climbing is meant to take the agent to a locally-distinctive state within the small-scale space surrounding the agent. Such a *distinctive state* is a point in the space where a properly chosen control law converges, usually by maximizing some kind of distinctiveness measure related to the perceived features in the surroundings. Examples include moving to the point where

the distances from the sensed features are maximal (e.g., “move to the center of the room”) or approaching a visual target. The high reproducibility of hill-climbing supports the extraction of the determinism intrinsic in the system <agent, environment, control laws/procedural knowledge>, when working in combination with trajectory-following.

Trajectory-following is used to pursue an option for leaving the current distinctive state and getting about another. Following a corridor is a typical example of this kind of control law.

Observe that both hill-climbing and trajectory-following can be performed through a behavior more complex than just the application of a single control law. They provide an ontological distinction between two kinds of procedural knowledge.

The partition of large-scale space into distinctive states, their surroundings and links induced by the control level allows the agent to localize itself at a local level. From a global standpoint, the uncertainty about its poses in the continuous space is reduced to uncertainty over a finite set of distinctive states. This in turn is handled by the causal and topological levels.

- Metrical Information

The agent may be able to acquire a precise metric/geometric representation of a *place-neighborhood* (the small-scale surroundings), for example by one of the recent successful probabilistic techniques (Thrun, 2002). If this is the case, the agent can use this representation as a local observer to avoid hill-climbing to dstates and to extract the perceptual characterization of its surroundings.

This is the basic way metrical information can be accommodated and exploited at the control level. The ‘patchwork’ map of local metric maps associated to the places of the final topological map (see topological level) might be used to build a global metrical map (see metrical level).

Technical details of possible implementations of the control level are provided in (Kuipers and Levitt, 1988; Kuipers and Byun, 1991; Lee, 1996; Beeson et al., 2003). A discussion of its assumptions and guarantees is given in (Kuipers, 2000).

2.1.2 Causal Level

A set of distinctive states, connected by the causal links that represent the control laws used to move between them, emerges from the control level. The causal level symbolically describes this net as an automaton. Besides distinctive states, the causal level introduces the notions of *view*, *action*, and *causal schema*.

An action summarizes the procedural knowledge for moving from a dstate to another, for example the sequence of hill-climbing/trajectory-following control laws needed.

Actions are often classified in *travel* and *turn* sorts. A travel action may change both the positions and the orientation of the agent, while a turn action only the latter. Thus, more dstates

linked by turn actions lie in the same location. The direct use of hill-climbing and the sharp distinction between turns and travels will be overcome to some extent in the extended framework of the hybrid SSH, in which our work takes place.

A view is a sensory image taken at a dstate. Since hill-climbing takes the agent into the same distinctive pose every time the place at hand is being visited, it is very likely that the agent always gets the same image. This eliminates or at least mitigates the problem of *image variability*, which occurs when the agent gets different sensory views of the same place upon different visits, and makes place-recognition very difficult. Examples of views are range-finders snapshots or high-level descriptions of the visual space.

Causal schemas are tuples of the form $< v,\ a,\ v' >$, which describe the event in which the agent performs action a in a dstate with view v, and ends up in a dstate with view v'. The association between dstates and their views is also registered. A chain of causal schemas of the kind:

$$< v_0,\ a_1,\ v_1 > \ < v_1,\ a_2,\ v_2 > \cdots < v_{i-1},\ a_i,\ v_i > \ldots$$

describes the whole experience of the agent at the causal level.

The causal level is compatible with the approaches based on graphs of views, as opposed to topological maps that are graphs of places, even though this distinction is not so sharp in literature. View-graphs have been proved consistent with human navigation abilities (Gillner and Mallot, 1998; Schölkopf and Mallot, 1995), and Franz et al. (1998) proposed them for robot navigation. Extraction of the spatial structure of the environment in the form of a stochastic automaton was addressed by Dean et al. (1995).

If two dstates never share the same view the view-graph is also a place-graph. If not, perceptual aliasing occurs and it must be figured out whether identical views, occurring twice or more times in the chain of causal schemas, correspond to different dstates or to the same. This problem is mostly tackled at the topological level, interpreting the structure of the environment according to the ontology. However, observe that the causal level provides enough information to deal with it to some extent.

Indeed, the chain of causal schemas built up to a certain point in the exploration may predict the sequence of views obtained during the execution of a particular action sequence, if it is assumed to be back again in a dstate with an identical view at an earlier point in the chain. If the actions really performed match this sequence, but views do not, such an assumption proves wrong.

This can be viewed as a form of subgraph matching in a graph of places. The principle here is very simple: two places are the same only if they, all their neighbors and neighbors of neighbors (and so no) appear respectively identical. This is, of course, only a necessary condition as the graph may exhibit genuinely isomorphic subgraphs, which is what happens in presence of two regions entirely perceptually aliased. In such a case, to detect the wrong assumption further exploration is needed — in the worst case, as long as the diameter of such regions — and, in presence of symmetry, no exploration will achieve this goal by this only

mechanism (see again Figure 1.2). Metrical information could help dramatically here, that is what we propose in Chapter 4.

- Metrical Information A causal schema of the kind $< v,\ travel,\ v' >$ might be annotated with metrical information about the net change in distance and orientation, one of the kind $< v,\ turn,\ v' >$ with just the net change in orientation.

A logical formalization of the causal level was given in (Remolina and Kuipers, 2001). More recently it has been expanded to take into account also the topological level. We shall briefly illustrate it in the topological level section.

2.1.3 Topological Level

The control and causal levels deal with the agent's experience. The topological level introduces an ontology of the external environment used to explain the agent experience in terms of *places*, *paths*, *regions*, and their connectivity, order and containment *relations*.

A place abstracts a part of the environment as a zero-dimensional point. It may lie on more paths. Every dstate is at a place, dstates linked by turn actions must be on the same place, while it is usually assumed that two dstates linked by a travel action are on different places.

A path abstract part of the environment as one-dimensional. It defines an ideal transitive closure and an order on a set of places linked by consecutive travel actions, which can be considered aligned (*e.g.*, if a turn action — not a turnaround — is performed between two travel actions, the two travel actions are not put on the same path). Each path has two directions denoting downstream ($+$) or upstream ($-$) in the order on its set of places.

It is important to avoid confusion with the notion of path in *graph theory*, which is based on the transitive closure of the connectivity relation. A SSH path is also a graph-theory path in the graph that underlies the topological map, but the inverse does not necessarily hold. Real-world examples of SSH paths include streets in a urban layout, corridors in an office, or tracks in a wood. These elements play a relevant role in human spatial cognition, in particular in the assimilation of internal representations for navigation (Kuipers, 1978; Kuipers et al., 2003).

A region represents a two-dimensional part of the environment. It may be defined by its boundaries, for example given by paths, by a set of places that is convenient to abstract as single spatial entity, or by a common metrical frame of reference.

A path divides the environment in two regions, left and right. A *boundary relation* involves a place and a path, and tells whether the place is on the right or on the left of the path. This is not always determined, in that paths in reality have not infinite length, and thus do not induce a perfect separation of the two regions.

A topological map is built from the minimal set of places and paths that explains the causal

schemas provided by the causal level, enforcing the ontological properties (for ex. see above the case of two travels split by a turn).

The properties enforced and the selection by minimality constrain the space of possible hypotheses, which is a consequence of perceptual aliasing, while interpreting the subjective experience of the agent.

Unfortunately, in some cases minimality may lead to over-reduced maps that do not correspond to reality (see for example Figure 3.4(c,d) of Chapter 3 page 61).

- Metrical Information The metrical annotation of causal schemas (see causal level) leads to the possibility of annotating the topological map as well. For example the headings at which paths lie in a place, or the distance between two places along a path can be annotated as attributes respectively of the place and of the path.

Logical formalization

Recently (Remolina and Kuipers, 2004) a formalization of the causal and topological level by means of circumscribed logical theories (McCarthy, 1980; Lifschitz, 1994) was proposed. Nested Abnormality Theories (Lifschitz, 1995) are used, because they allow for the application of more occurrences, possibly nested, of the circumscription operator to subsets of axioms in a modular fashion.

We can roughly sketch the overall structure of the theory as follows:

1. Some axioms assert the existence of each causal schema, involving actions, views and dstates as logical constants. These belong to different sorts of the language.

2. Some other axioms assert the existence of places, paths and (when considered) regions, upon the existence of the individuals in 1. (There must exist a place where a dstate lies. If there is a a schema like $< v,\ travel,\ v' >$, there must exist a path where it lies, and the two places to which the dstates associated with v and v' belong must be on this path as well. These are just two examples.)

3. Topological relations are axiomatically defined as well. Boundary relations, which play a fundamental role, are defined on the grounds of turn action experienced at the causal level.

4. Another group of axioms states the topological constraints assumed (for example, a place cannot be simultaneously on the left and on the right of the same path).

5. Some second order (induction-like) properties, like the fact that a path should contain only places actually linked by travel actions, are expressed using circumscription on a proper

block of axioms in isolation. In particular, these are also used to define the boundary relations. Every place encountered along a path p' after having turned right from path p will be on the right of p unless this generates an inconsistency with the axioms above.

6. The models of the theory so far represent all the SSH maps that are consistent with the experience and the causal and topological levels.

7. A prioritized circumscription policy is applied to the theory. This may take different forms. It usually establishes a lexicographical order based on the number of paths, places, boundary relations, *etc.* Unusual forms, such as self-intersecting paths, can be in principle modeled in the policy, so as to prefer models that do not exhibit such peculiarities.

2.1.4 Metrical Level — Global Metrical Mapping

A topological map of the environment can drive the integration of local metrical maps into a single global frame of reference. This helps to control the cumulative pose error, as shown in (Thrun et al., 1998b; Modayil et al., 2004). The SSH naturally allows for such approaches, provided that local metrical maps are built at places and metrical annotations are taken for causal schemas (as explained for metrical information in control and causal level sections).

2.2 Accounting for Small-Scale Space

The main purpose of hill-climbing in the control level is the reliable and unique localization of the agent in a place-neighborhood, over multiple visits of the same place. Thus, hill-climbing also allows the robot to get distinctive sensory snapshots, which are used to recognize and distinguish places during the exploration. If the agent can accomplish both these tasks through more general perceptual representations of small-scale space — such as visual panoramas, or local metrical maps — hill-climbing can be avoided.

This idea is supported by recent advancements in metrical mapping of environments of limited dimensions, and by those expected in the field of artificial vision, especially in scene interpretation and vision-based motion planning. These advancements are the technological motivation to the ontological and technical integration of small-scale space into the SSH, leading to an hybrid enhanced framework (Beeson et al., 2003; Kuipers et al., 2004).

2.2.1 LPMs and Stars

We term a general perceptual representation of a place-neighborhood *Local Perceptual Model* (LPM). A LPM is supposed to:

- Allow the robot to localize itself unambiguously in the place-neighborhood.

- Provide directly or support the extraction of a reliable and complete description of the features or views used to distinguish the place.

- Support the extraction of a reliable and complete description of the options available to leave and enter/approach the place.

Observe that the SSH is general enough to accommodate a LPM as a 360-degree sensory view of a dstate, provided that any turn action leaves the robot in this dstate, which must be the only one in the place at hand. However, the concept of LPM is general enough to cover the case of a place where more dstates are linked by turn actions.

In the SSH topological level a place is characterized by the paths on which it lies. Such paths locally correspond to the options for leaving or approaching and entering the place at the control and causal levels, traditionally by alternating trajectory-following and hill-climbing. We want to extract from a LPM the counterpart of this structure in the small-scale space. We call this new structure *star*, because of the shape the set of paths intersecting at a place exhibits. The paths in the star are called *local paths*, for they are intended to represent path segments of the large-scale topological map intersecting in the place at hand. Note the difference in terminology, a *path* is a (global) path of the SSH topological level, unless otherwise specified; a *local path* always refers to a path segment in the star.

We introduce some new terms and notations for this small-scale spatial ontology. LP_{id} is an identifier for a local path; $+$ or $-$ are used to disambiguate between the two possible directions of a local path. The two directions can be also viewed as distinguishing tags for the two endpoints of the local path that appear in the star. They do not necessarily have to be oriented in the same way as the $+$ and $-$ of the corresponding SSH global path. We call a local path annotated with a direction an *oriented local path*.

The notion of star can be now formally defined.

Definition 2.1 (Star). *A star S is a set of oriented local paths such that*

$$LP+ \in S \qquad \textit{iff} \qquad LP- \in S$$

and equipped with two functions:

(i) $next : S \rightarrow S$ *induces a circular order over the elements of* S*;* $next(x)$ *is the element clockwise next to* x*.*

(ii) $\alpha : S \rightarrow A$ *associates an attribute value* $\alpha(olp)$ *to* olp*;* A *is a given domain of possible values for the attribute.*

The attribute values in A carry additional information about each oriented local path. The basic version of A contains two values: $\{T,\ C\}$. T specifies that the local oriented path is actually *travelable*, in that the corresponding global path continues outside the place in the

specified direction. On the contrary, C specifies that the oriented local path is *closed*, because the global path terminates at the place at hand. A richer classification here is possible, as needed for the desired small-scale spatial ontology.

We conveniently represent a star S as a circular list of pairs of oriented local paths and their attributes

$$[< OLP_1,\ \alpha(OLP_1) >, < OLP_2,\ \alpha(OLP_2) >, \cdots < OLP_n,\ \alpha(OLP_n) >, \ \ldots]$$

Given an element olp of S, $next(olp)$ is its successor in the list, the successor of the last element being the first in the list.

Example 2.1.

The stars for a "+" intersection, for a "T" interesection, and for a "L" corner are:

$$\begin{aligned}
(+): &\quad [< LP_0+,\ T >, < LP_1+,\ T >, < LP_0-,\ T >, < LP_1-,\ T >] \\
(T): &\quad [< LP_0+,\ T >, < LP_1+,\ C >, < LP_0-,\ T >, < LP_1-,\ T >] \\
(L): &\quad [< LP_0+,\ C >, < LP_1+,\ T >, < LP_0-,\ T >, < LP_1-,\ C >]
\end{aligned}$$

Note that the star is independent of where the enumeration of oriented local paths begins. □

We introduce the notion of *isomorphism* between two stars.

Definition 2.2. *An isomorphism between two stars S and S' that share the same attribute domain A is a bijective function $\varphi : S \to S'$ such that*

(i) $\varphi(next(olp)) = next(\varphi(olp))$.

(ii) $\alpha(olp) = \alpha(\varphi(olp))$.

(iii) *$lp+$ and $lp-$ lie on the same local path in opposite directions iff $\varphi(lp+)$ and $\varphi(lp-)$ do so too.*

Intuitively, an isomorphism $\varphi : S \to S'$ such that $\varphi(olp) = olp'$ guarantees that S and S' appear identical when observed respectively from olp and olp'. Another intuitive way to look at this is that S and S' match when relatively rotated so as to have olp correspond to olp'.

Remark 2.1.

Note that there can be more than one isomorphism between two stars. For example there are four of them between two "+"-shaped stars. □

2.2.2 Gateways and Stars from Occupancy-Grid Maps

We now address the problem of how star-representations could be extracted from occupancy-grid maps that play the role of LPMs. To this purpose, we introduce the concept of *gateways*, intended as access points of a place neighborhood.

Possible operational definitions of gateways include:

1. Gateways can be defined as boundaries between regions of different control law applicability, when positions within a LPM are tagged with the control law applicability conditions they satisfy. In particular, the border between a hill-climbing region and a trajectory-following one should be considered a gateway.

2. Chown et al. (1995) define gateways as the locations of major changes in visibility:

   > *In buildings, these [gateways] are typically doorways.... Therefore, a gateway occurs where there is at least a partial visual separation between two neighboring areas and the gateway itself is a visual opening to a previously obscured area. At such a place, one has the option of entering the new area or staying in the previous area.*

3. In some cases, a geometric criterion for identifying gateways can be based on the medial axis of free space in the LPM, that is the one-dimensional set of points that are equally distant from nearby obstacles. All the such axes of a space make a graph, well known in as the *Voronoi diagram*. A gateway corresponds to a "constriction" (or "critical line" Thrun, 1998) along a medial axis edge, where the distance between the edge and obstacles is a local minimum near a larger maximum.

The principle in 2. does not apparently lend itself to a direct robot implementation, while 1. and 3. are the criteria illustrated in Figures 2.2 and 2.3 (1. is used for Figure 2.3(c)). In these examples, gateways are first detected and then used to identify the local paths making up the star of the place. Whether a local path terminates at or continues across a place must be decided. If a local path terminates at a place, it will have only one gateway; if it passes through, it will have two. The values of A are lists of control laws applicable on the oriented local path.

Figure 2.3 also provides some insights into how the small-scale and large-scale spatial ontologies can relate to each other, involving every level of the SSH:

- Gateways can be considered dstates. (Control and Causal Levels.)
- Such dstates would be linked by virtual turn actions, which do not correspond to physical turns at the control level, but to some short navigations planned in the occupancy-grid map. (Control and Causal Levels.)
- Attributes for oriented local paths can represent which control laws are applicable. (Control Level.)

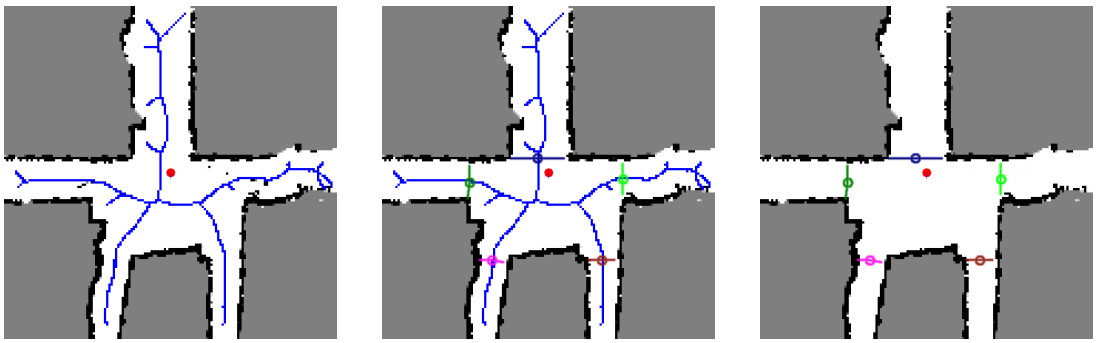

(a) (b) (c)

Figure 2.2: **Finding Gateways in a LPM.** A possible implementation of a local perceptual model (LPM) is a bounded occupancy-grid map. The robot is shown as a circle in the center of the LPM. **(a)** To find gateways in corridor environments, the algorithm computes the medial axis of the occupancy grid free space. **(b)** The maximum of the medial axis graph is found (where the distance of obstacles from the graph is maximal) and each edge is traversed, looking for "constrictions" (where the distance between the graph edge and obstacles is a local minimum). **(c)** The final gateways are drawn as lines connecting the graph edge minima (circle) with the closest obstacles. (Taken from Beeson et al., 2003.)

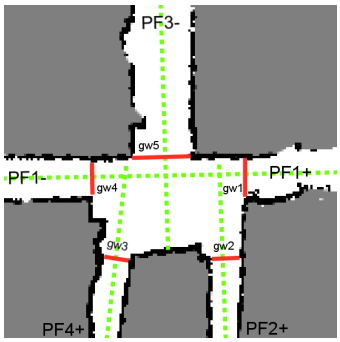

(a)

Star

[<PF1+, (gw1,out) & (gw4,in), Midline>,
<PF2+, (gw2,out), Midline>,
<PF3+, (gw5,in), Closed>,
<PF4+, (gw3,out), Midline>,
<PF1-, (gw4,out) & (gw1,in), Midline>,
<PF4-, (gw3,in), Closed>,
<PF3-, (gw5,out), Midline>,
<PF2-, (gw2,in), Closed>]

(b)

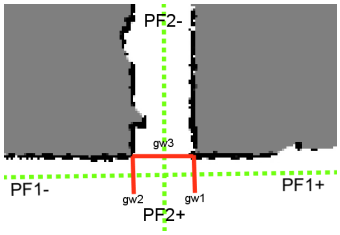

(c)

Star

[<PF1+, (gw1,out) & (gw2,in), LeftWall>,
<PF2+, (gw3,in), None>,
<PF1-, (gw2,out) & (gw1,in), RightWall>,
<PF2-, (gw3,out), Midline>]

(d)

Figure 2.3: **Local Extraction of the Star Description. (a-c)** Given the gateways and local paths of a place, the robot can extract the star. **(b-d)** The star enumerates all the local paths in clockwise order. Here the attribute for an oriented local path takes the form of the list of control laws (trajectory-following) applicable. The correspondence with the gateways is also annotated. (Adapted from Beeson et al., 2003.)

- Suitable portions of the occupancy-grid map can be considered views taken at the dstates. (Causal Level.)
- Local paths of the star can be annotated with identifiers for the corresponding global paths at the SSH topological level. (Topological Level.)

However, for our purposes, we will mostly focus on the star as the small-scale spatial counterpart of the topological level, in the simple form as in Example 2.2.1, page 32.

2.3 A Topological Mapper

Motivations for the ontological and technical integration of small-scale space in the SSH are provided in the previous section, and by the experimental results discussed later in Section 2.4. We have designed and implemented a topological mapper for the causal and topological level in this new context.

Instead of the traditional input as a sequence of causal schemas (see Section 2.1.2):

$$
\begin{aligned}
&< v_0,\ a_1,\ v_1 > \\
&< v_1,\ a_2,\ v_2 > \\
&\vdots \\
&< v_{i-1},\ a_i,\ v_i > \\
&< v_i,\ a_{i+1},\ v_{i+1} > \\
&\vdots
\end{aligned}
$$

the topological mapper takes a sequence of stars and oriented local paths:

$$
\begin{aligned}
&< S_0,\ S_0 :: OLP_{out},\ S_1 :: OLP_{in},\ S_1 > \\
&< S_1,\ S_1 :: OLP_{out},\ S_2 :: OLP_{in},\ S_2 > \\
&\vdots \\
&< S_{i-1},\ S_{i-1} :: OLP_{out},\ S_i :: OLP_{in},\ S_i > \\
&< S_i,\ S_i :: OLP_{out},\ S_{i+1} :: OLP_{in},\ S_{i+1} > \\
&\vdots
\end{aligned}
$$

S_i denotes a star, $S_i :: OLP_{out}$ and $S_i :: OLP_{in}$ denote the oriented local paths used respectively to leave and enter the star S_i, and refer to the internal representation of S_i.

We call a single tuple $< S_i,\ S_i :: OLP_{out},\ S_{i+1} :: OLP_{in},\ S_{i+1} >$ an *observation*. This observation describes the piece of exploration in which the robot (1) is at a place whose perceived star is S_i, (2) leaves the place through $S_i :: OLP_{out}$, (3) performs a travel action and (4) reaches a place whose perceived star is S_{i+1} accessing it through $S_{i+1} :: OLP_{in}$.

Except for the first observation, at every step the first star is identical to the second in the previous observation.

Figure 2.4 provides an example of exploration and its description as list of observations.

In this thesis, the only information characterizing a place that we consider in input is the star. This choice has four main reasons:

1. *Image variability* can cause the same place to appear differently due to differences in the LPMs built upon two visits. In a sense, image variability can be considered orthogonal to perceptual aliasing (Kuipers and Beeson, 2002). Considering only the star prevents or at least strongly reduces the chance of occurrences of image variability.

2. The star concisely represents the minimal information that a topological-mapping system needs: (a) a useful description of the options available to leave or enter/approach a place/an object of interest, and (b) a log of those such descriptions encountered and options taken.

3. Richer perceptual characterizations of places can be always accommodated on top of the minimal description given by the star, perhaps relaxing some of the SSH deterministic assumptions.

4. Retrieving and matching entire LPMs (for instance metrical maps if they are used as LPMs) would be far more expensive computationally.

2.3.1 Algorithm

The topological mapper takes into account one observation at a time and builds all the topological models consistent with the observations already considered. These models consist of a map and *a pointer to the current place*; the combination is called a *you-are-here* (YAH) map in the specialized literature. We shall denote such a map as $\dot{m} = \langle M, p \rangle$. M is the map, including places represented as stars, while p indicates the current place. Some of the oriented local paths in the stars of the map will be already pairwise linked along SSH paths. The remaining *travelable* oriented local paths correspond to options for further exploration that have not yet been pursued, according to the particular interpretation of the exploration so far given by the topological map $\dot{m}$. We will call such an unlinked oriented local path a *pending oriented local path*, abbreviated by *pending olp*. Likewise, we call a gateway that correspond to a pending olp a *pending gateway*.

The overall functioning of the topological mapper is based on the following recursive mechanism:

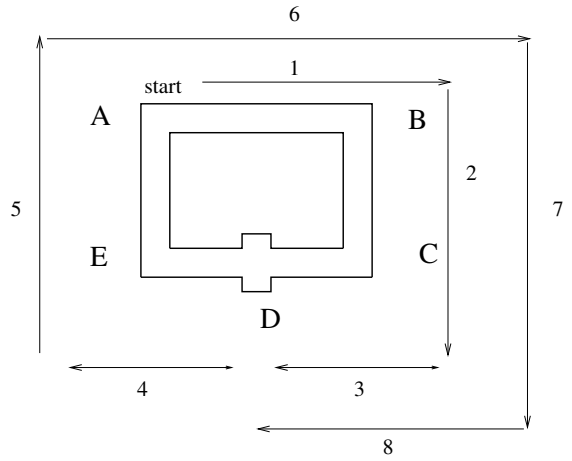

(a)

$$S_A = [< LP_0+,\ C >,\ < LP_1+,\ C >,\ < LP_0-,\ T >,\ < LP_1-,\ T >]$$
$$S_B = [< LP_0+,\ T >,\ < LP_1+,\ C >,\ < LP_0-,\ C >,\ < LP_1-,\ T >]$$
$$S_C = [< LP_0+,\ T >,\ < LP_1+,\ T >,\ < LP_0-,\ C >,\ < LP_1-,\ C >]$$
$$S_D = [< LP_0+,\ T >,\ < LP_1+,\ T >,\ < LP_0-,\ T >,\ < LP_1-,\ T >]$$
$$S_E = [< LP_0+,\ C >,\ < LP_1+,\ T >,\ < LP_0-,\ T >,\ < LP_1-,\ C >]$$

(b)

$$< S_A,\ S_A :: LP_0-,\ S_B :: LP_0+,\ S_B >$$
$$< S_B,\ S_B :: LP_1-,\ S_C :: LP_1+,\ S_C >$$
$$< S_C,\ S_C :: LP_0+,\ S_D :: LP_0-,\ S_D >$$
$$< S_D,\ S_D :: LP_0+,\ S_E :: LP_0-,\ S_E >$$
$$< S_E,\ S_E :: LP_1+,\ S_A :: LP_1-,\ S_A >$$
$$< S_A,\ S_A :: LP_0-,\ S_B :: LP_0+,\ S_B >$$
$$< S_B,\ S_B :: LP_1-,\ S_C :: LP_1+,\ S_C >$$
$$< S_C,\ S_C :: LP_0+,\ S_D :: LP_0-,\ S_D >$$

(c)

Figure 2.4: **Topological Input. (a)** The environment and the exploration. **(b)** Stars for the distinctive places. **(c)** List of observations. (The search-tree produced by the topological mapper for this input is given in Figure 2.5.)

- *Basic Step.* A trivial map with one single place is built from the first star in the first observation.
- *Recursive Step.* The maps that extend a given map consistently with a new observation are created. We call these new maps the *successors* of the given one; they represent all the different topological hypotheses that can explain the new observation in the given map. The number of successors can be 0 if the new observation is inconsistent with the assumed map, 1, or > 1 in case of perceptual aliasing.

Because of this recursive mechanism, the search space develops as a tree whose depth is the number of observations processed. We discuss the detail of these two steps.

Basic Step.
The first map ever created by the algorithm contains one single place whose star is given by S_0 in the first observation

$$< S_0,\ S_0 :: OLP_{out},\ S_1 :: OLP_{in},\ S_1 >$$

Of course, the current place is the only place of the map, while the next observation to take into account (in the first application of the recursive step explained below) is still the first observation above. □

Recursive Step.
Let $\dot{m} = \langle M, p \rangle$ be the map at hand, and

$$< S_i,\ S_i :: OLP_{out},\ S_{i+1} :: OLP_{in},\ S_{i+1} >$$

the next observation to take into account for $\dot{m}$. Let S_p be the star of p in M; assume (we will prove it by induction shortly) that there exists an isomorphism $\varphi : S_i \rightarrow S_p$, which intuitively guarantees that S_i and S_p are compatible.

There are three possibilities while creating the successor(s) of $\dot{m}$:

1. **Predicted-Consistent (pc).** M already includes a link from p through $\varphi(S_i :: OLP_{out})$ to another place, call it p' and let its star be $S_{p'}$, through one oriented local paths of $S_{p'}$, say $S_{p'} :: OLP_{in}$ (*i.e.*, the two local paths already belong to the same SSH path). Then $\dot{m}$ predicts the new position and the star expected in it (p' and $S_{p'}$). If there exists the isomorphism $\varphi' : S_{i+1} \rightarrow S_{p'}$ such that $\varphi'(S_{i+1} :: OLP_{in}) = S_{p'} :: OLP_{in}$ (intuitively, S_{i+1} and $S_{p'}$ match under the robot's viewpoint upon entering the place), then the observation actually confirms the prediction, and the only successor is the same map where p' is the robot's new current place. Finally, φ' will be the new isomorphism φ inherited by the next step to come.

2. **Predicted-Inconsistent (p-nc).** Same situation as in (pc), but here the second star given in the observation is not compatible with the predicted one. Formally, there is no isomorphism $\varphi' : S_{i+1} \rightarrow S_{p'}$ such that $\varphi'(S_{i+1} :: OLP_{in}) = S_{p'} :: OLP_{in}$. The map proves wrong and has no successor.

3. **Expand (e).** $\varphi(S_i :: OLP_{out})$ is a pending olp, then the observation is not predictable in the map. At least a successor is created, where a completely new place, say p', whose star will be S_{i+1}, is inserted. p' is the new current place of the robot, and is properly linked to p, by making $\varphi(S_i :: OLP_{out})$ and $S_{i+1} :: OLP_{in}$ lie on the same SSH path, and p and p' be one next the other. The new isomorphism φ for the next step to come is the identity function.

   If there are old places whose star is isomorphic to the observed one, this first successor assumes perceptual aliasing with all of them. In this case, as many other successors as the possible hypotheses of being back to an old place need to be created. This makes the original current map, intended as a node of the search-tree, branch.

   The set of such additional successors can be ideally obtained as follows. First, the set of all maps that can be created from $\dot{m}$ by linking $\varphi(S_p :: OLP_{out})$ to every other pending olp is considered. For a map of this set call such a pending olp $S_{p'} :: OLP_{in}$, p' being the arrival place and $S_{p'}$ its star. If there is no isomorphism $\varphi' : S_{i+1} \rightarrow S_{p'}$ such that $\varphi'(S_{i+1} :: OLP_{in}) = S_{p'} :: OLP_{in}$ the map is discarded. This consistency check is similar to that in the (p-c) and (p-nc) steps: the old star as perceived at the arrival must match the observed one. The new current place is p', while the new isomorphism φ for the next step is φ' that satisfies the consistency check.

   Note that, since there can be more isomorphisms between two stars (see Remark 2.2.1, page 32), an old place might have more pending olps to which the current place can be linked consistently with the observation.

The next observation to take into account after the completion of the recursive step is

$$< S_{i+1},\ S_{i+1} :: OLP_{out},\ S_{i+2} :: OLP_{in},\ S_{i+2} >$$

□

The practical role of φ is to keep track from step to step of the rotational offset between oriented local paths in the observed star and those in the star of the current place. The recursive assumption on the existence of this isomorphism is justified because the algorithm makes sure that the current place's star is always isomorphic to the second star of the last observation, which is the same as the first of the current observation.

Formally, the proof of this is by induction. After the basic step, φ is trivially the identity function $id : S_0 \rightarrow S_0$. At the recursive step, a proper isomorphism φ is assumed by inductive

hypotyhesis, and a new one for every successor is defined (see above) in the two cases (pc) and (e), while we do not need one for (p-nc) that corresponds to a branch termination in the tree of maps.

An example of search-tree is given in Figure 2.5 for the exploration in Figure 2.4.

The topological mapper also implements (and can optionally run) a test for each axiom of the SSH, like the one that rules out maps with self-intersecting paths. As a consequence, the models that do not meet common environmental defaults are filtered out, and the tree is pruned. Moreover, the topological mapper is also equipped with the algorithms for comparing two maps according to the preference criteria chosen at the SSH topological level, mostly based on the minimality of paths and places, and good-form quality of regions (see Section 2.3.2).

We take into consideration two policies for traversing the search-tree of topological maps.

A *breadth-first-search* can be performed along the whole sequence of observations. This yields all the leaves of the search-tree. Finding the optimal map(s) according to the given preference criteria is then the last step. This approach guarantees optimality, but suffers from combinatorial explosion of the number of maps, if perceptual aliasing occurs.

A *best-first-search* policy selects the best map, according to the preference criteria, among the maps already existing at the current step. Only the successors of this map are created at this step. This method is much more efficient in the average case, but may miss the optimal map. It always provides a consistent map, if such a map exists, but it does not guarantee optimality.

Observe the following problems, regardless of the method employed:

- In presence of symmetry and perceptual aliasing there may be more optimal maps, depending on how optimality was defined (for example the two maps in Figure 1.2 are both optimal according to the policy dicussed below in Section 2.3.2).

- When exploration is poor, minimality-based optimal selection is likely to output oversimplified models of the environment's topology. For example, consider the breadth-first search in Figure 2.5; it is easy to note that afer just two steps the minimal model would be the leftmost one with just two places, which has not been proved inconsistent yet. For the same reason, after 4–6 steps, the minimal model would be the 4-place one on the leftmost branch. In general, the preference policy is more complex than the plain minimization of the number of places. Nevertheless, these examples suggest what can happen during a best-first search, or when a premature commitment to a single model is made.

  There can be also situations where an overminimal map cannot be topologically ruled out, no matter how long and accurate the exploration. A simple example is given by exploring a square, where only corners are accounted as places, which all look the same. Remolina and Kuipers (2004) point out how a triangle could be an overminimal correct

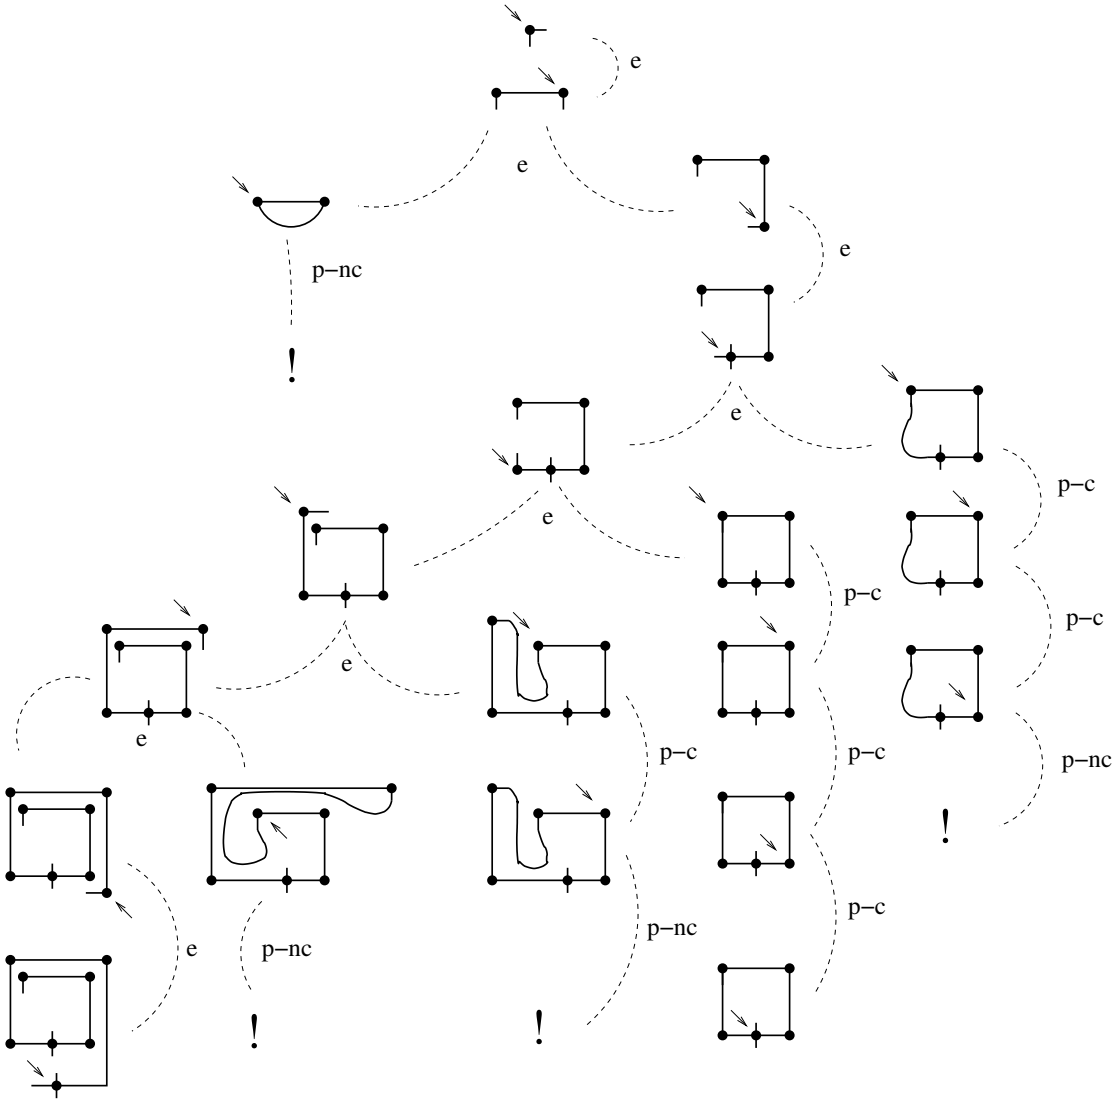

Figure 2.5: **Topological Search.** Search-tree of partial maps according to the exploration given in Figure 2.4. The solid-filled points represent the places assumed in each map. Here a topological axiom that does not allow paths to intersect themselves (self-intersections forbidden) is enforced while creating topological hypotheses. The sort of each derivation from map to map (e, p-c, p-nc as explained in the text) are annotated; ! indicates a branch termination after a (p-nc) step. A little closed-headed arrow points to the assumed current place in each map.

interpretation of this situation, from a pure topological viewpoint. One similar case for a 4 by 4 grid environment will be presented in Chapter 3, Figure 3.4(d).

We claim that all these problems will benefit from informing a best-first search also by a degree of metrical plausibility of each map created, evaluated on-line. We address this aspect in Chapter 4.

2.3.2 Preference Policy

In Section 2.1.3 we mentioned that the preference policy for ordering topological maps is specified by a prioritized circumscription operator in the logical formalization of the SSH. The result is a total lexicographical order over vectors of values, each such value being the number of occurrences of an element of the SSH topological ontology (paths, places, relations). This allows us to assign different importance to, say, the minimization of paths and places. (As an empirical property of common environments, preferring a map that has fewer paths over one that has fewer places usually yields better results.)

Unless otherwise specified, in this thesis we will also assume the prioritization to be "boundary relations win over paths, paths win over places". Formally

- $f_1(M)$ is the number of pairs $\langle path, place \rangle$ in M such that $place$ has no boundary relation with $path$; *i.e.*, $place$ cannot be told univocally either on the right or on the left of $path$.
- $f_2(M)$ is the number of paths in M.
- $f_3(M)$ is the number of places in M.

M_a is preferred to M_b *if and only if*

$$\begin{aligned} & f_1(M_a) < f_1(M_b) \ \vee \\ & f_1(M_a) = f_1(M_b) \ \wedge \ f_2(M_a) < f_2(M_b) \ \vee \\ & f_1(M_a) = f_1(M_b) \ \wedge \ f_2(M_a) = f_2(M_b) \ \wedge \ f_3(M_a) < f_3(M_b) \end{aligned} \tag{2.1}$$

Minimal maps according to this order will be considered the optimal ones, and preferred to the greater ones.

2.3.3 Advantages

In the hybrid SSH, the breadth-first and best-first policies for traversing the topological search-tree are implemented in the same way as in the non-hybrid formalization given in (Remolina, 2001; Remolina and Kuipers, 2004). Nevertheless, there is a difference that contributes to reducing the exponential growth of the tree.

Whether the last observation is consistent with the predicted one must be tested in a (p-c) or (p-nc) step, or after a loop-closure hypothesis, which may occur in a (e) step. Such a consistency check is carried out by matching entire LPMs or their topological abstractions, as opposed to simple view matching, which covers only part of a LPM. So, the unit of matching in the algorithm becomes the whole place rather than one of its distinctive states.

This has two important consequences. First, branching at step (e) is reduced as a result of matching larger portions of space at once. Second, topological ambiguity can only arise after travel actions. Indeed, by building a *complete* representation of a place-neighborhood (the LPM), it is possible to know about all the distinctive states at once. This allows the robot to predict the local result of any turn action; turn actions are implicitly fully represented in a LPM. This is not possible according to the logic formalization of the non-hybrid case, which indeed needs to allow for more hypotheses about the number of dstates at the same place, as a consequence of turn actions. For example, consider a robot that performs successive turn actions at a "+" intersection, each turn ending when the robot faces the next oriented local path. With no quantitative estimation of the rotation, if the four resulting views are summarized as we did by star-extraction and appear identical to the robot, there is no way to figure out whether the intersection has more than just two intersecting paths. In general, even if some views differ at a place, every model that is equivalent to the correct one up to a modulo operation on the circular sequence of views is legitimate. This dramatically affects the growth of the search-tree.

A related problem is the identification of boundary (left/right) relations between paths intersecting at a place. In the hybrid SSH the star provides a direct decision procedure to tell what is on the left or on the right of an oriented path. Starting from a local oriented path $LP+(-)$, the subset of oriented local paths encountered by moving down the clockwise order until observing $LP-(+)$ are on the right of $LP+(-)$. Likewise, moving down the counter-clockwise order accounts for the left side, and the two subsets are always disjoint by construction. Very differently, in order to keep the left and right subsets from overlapping in the the non-hybrid formalization, a more binding mechanism is adopted. A special *TurnAround* action of exactly 180 degrees is needed to associate two opposite local paths correctly. Performing such an action is the only way to define the continuation of a path across a place permitted by the theory. In practice, this makes the system utterly sensitive to the inevitable noise in odometry readings. To attain reasonable robustness, these data could be processed over time while turning, *e.g.* through repeated scan-matching. Such an approach would be very close to the idea of introducing, modeling and using LPMs that we have pursued.

In summary, both hybrid and non-hybrid frameworks are faced with a worst-case $O(b^d)$ number of leaves of the search-tree, where $b-1$ is an upper bound on the number of possible simultaneous perceptual aliasing occurrences, and d is the number of actions that can cause branching. In the hybrid case, b is reduced because matching places, using complete local topologies and LPMs, is more restrictive than matching views at dstates. Besides, d is reduced because only travel actions, rather than both travels and turns, can lead to branching.

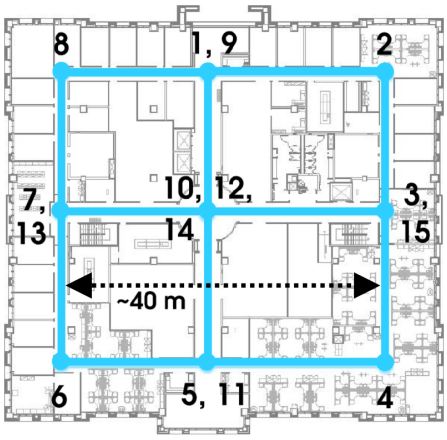

Figure 2.6: **Map Building I: Exploration.** The environment is the third floor of the ACES building of the University of Texas at Austin. The order of places in the exploration taken by the robot is enumerated on the plan. LPMs and the output of the mapping process are illustrated in Figure 2.7 and 2.8. (The three figures are adapted from Kuipers et al., 2004.)

2.4 A Physical Exploration

The hybrid SSH map-builder has been tested with an exploration of an office environment, shown in Figure 2.6. Figure 2.7 reports the sequence of LPMs observed at successive place neighborhoods. Figure 2.8 shows the correct topological map produced by the overall mapping system, with LPMs overlaid at corresponding places in the correct topological map. Here, no attempt is made to build a global metrical map of this environment within a single frame of reference. Modayil et al. (2004) address this task by using the final topological map to integrate metrical information locally gathered through a scrolling occupancy-grid map. The algorithms we propose in Chapter 4, whose aim is to evaluate the metrical likelihood of multiple topological hypotheses, also estimate the global metrical relations between the different places of a topological map as a side effect.

The exploration exhibits several interesting features that makes it well suited for experimentation. The metrical size of the environment would fit experiments commonly carried out to validate the success of modern metrical SLAM methods. The office contains multiple nested topological loops, which raise topological ambiguity in combination with perceptual aliasing. In particular, three non-isomorphic stars (one “+”, and four occurrences of “L” and “T” each) nine places and six paths (in this case respectively junctions and corridors) occur in the explored area. However, the office does not offer a crisp structure in terms of walls, corridors and junctions,

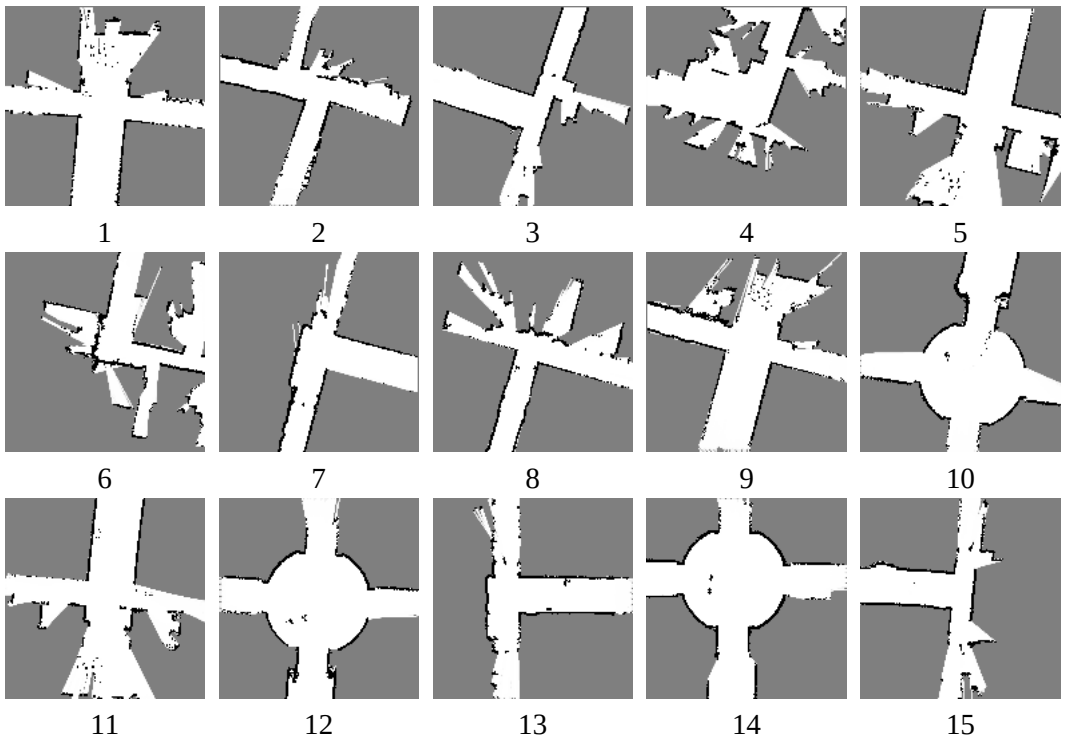

Figure 2.7: **Map Building II: LPMs.** The LPMs created at the places during the exploration shown in Figure 2.6. The stars generated from these LPMs are used to search through the space of consistent topological maps.

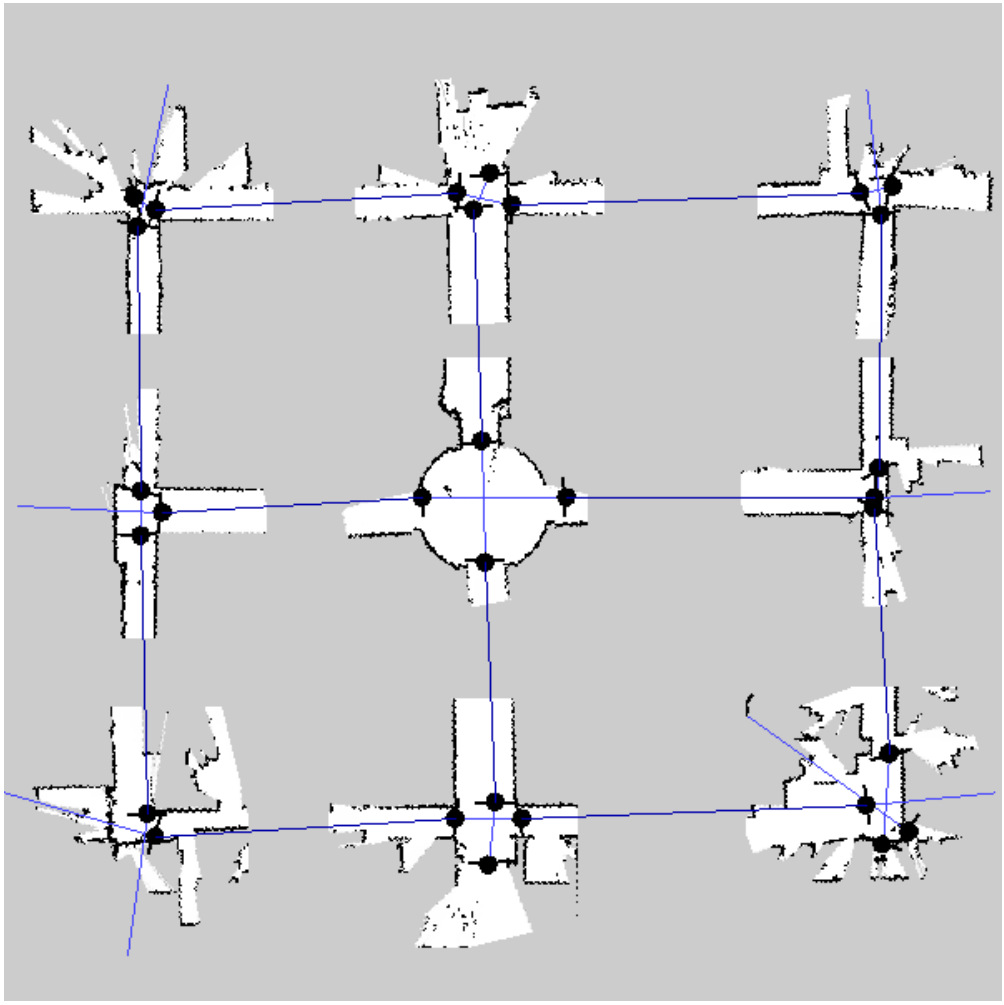

Figure 2.8: **Map Building III: Map.** The correct topological map is determined among the 4 final optimal ones either by the application of the planarity constraint (see Chapter 3), or through accurate full-LPM matching. The map shown here is overlaid with the LPMs generated at the places, with the gateways, and with the connections between gateways which lie on the same path.

because a quite complex and extended pattern of cubicles [1] or small lounge rooms occupy most space at the corners. Since the grid-map analysis system is not enough sophisticated to account for them, they appear as a fragmentation of the free space that affects the process of gateways detection. This can be observed in the grid map of the right-lowermost corner (Figure 2.8), where four gateways instead of just two are annotated. As a result, the topological mapper here works with four non-isomorphic stars instead of three.

After the exploration consisting of 14 travel actions, the topological mapper finds 83 possible configurations of the environment that are consistent with the observed stars and the topological constraints. The preference policy selects four final optimal models. The correct topological representation of the environment can then be obtained by matching the full LPMs in succession: the overall degree of correlation between the observed sequence of occupancy-grid maps and the one predicted by each of the four topological maps is maximized.

If we assume planarity of the environment, we can use a more sophisticated version of the topological map-building algorithm (which is the subject of Chapter 3) that rules out many more models as inconsistent. In this case, there are only 46 consistent configurations of the exploration experience, and the only optimal model selected by the preference policy is the correct map, so that full LPM-matching can be avoided. Currently, our implementation can build the tree of models and determine the unique minimal map of this environment in $\sim$200 ms on the robot's Pentium III 450MHz processor.

2.5 Related Work

2.5.1 Similar Approaches to Hybrid Mapping

We review related work on two important themes in hybrid metrical-topological mapping. First, the "patchwork metrical map" consists of local metrical maps with separate frames of reference, linked by topological relations. Second, ambiguity in the topological map means that multiple maps are consistent with the travel experience. Other approaches combine topological and metrical maps, for example creating a global metrical map first and then parsing it into topological regions (Thrun, 1998), but these are less related with the hybrid SSH.

In the "patchwork metrical map", places in the global topological map are annotated with local metrical maps with limited extent and separate frames of reference. In early work on the SSH, Kuipers and Byun (1991) associated topological places with local landmark maps and path segments with generalized cylinder models. The local frames of reference could be relaxed into a single global frame of reference by spreading metrical errors evenly across the topological map.

Yeap and Jefferies (1999) build maps consisting of adjacent metrical maps of rooms, which are directly connected by gateway-like entities. Bosse et al. (2003) link together perceptual

[1] A set of panels and desks arranged so as to extend the office environment in open space.

maps of a fixed number of landmark features. Similarly, Duckett and Saffiotti (2000) connect overlapping local occupancy-grid maps to form a dense topological network. Lankenau et al. (2002) create a topological graph of travel paths annotated with metrical maps at the "corners" where paths intersect.

Most mapping implementations only maintain a single map hypothesis, selected through greedy or maximum likelihood methods. Perceptual aliasing may lead to ambiguity about whether the robot has closed a loop or just reached a similar, nearby place. A single map hypothesis cannot model this ambiguity. However, a few hybrid mapping techniques do reason about structural ambiguity.

Kuipers and Byun (1991) detect perceptual aliasing and check for possible loop closures by performing physical motion to obtain more evidence. Assuming the environment does not contain large topologically identical substructures, this method handles the two cases of whether a loop is present or not. Other frameworks that make use of probablistic methods to solve topological uncertainty will be reviewed as related work of our approach described in Chapter 4.

As in the hybrid SSH, Dudek et al. (1993) construct an "exploration tree" of all possible consistent world models. To extend the exploration tree, they use the degree of each node (the number of graph edges) in the same way that we use local topology to match places. By using local topology (and even LPMs when necessary), our method has a smaller branching factor.

2.5.2 Place Detection, Classification, and Recognition

Place detection and classification are important steps for the correct functioning of an approach like the hybrid SSH.

The first approach to this problem that makes use of occupancy-grid maps was given by Yamauchi and Langley (1997). More recent techniques were investigated by Duckett and Nehmzow (2001), who make use of hystogram (from occupancy-grid maps) matching combined with multiple location hypotheses tracking. Fabrizi and Saffiotti (2000, 2002) adopt a different approach based on fuzzy digital image processing and mathematical morphology. Kuipers and Beeson (2002) approach the problem from a more general viewpoint, proposing unsupervised bootstrap learning methods for learning how to recognize places.

Chapter 3

Loop-Closing and Planarity

Loop-closing is at the core of topological mapping. While parsing the sequence of observations, when and where loops are closed determines a complete topological hypothesis. If the environment to map is known to be planar, the options to close loops are reduced, because the planarity constraint requires that the graph-model of the map can be drawn without crossing edges. In this chapter we formally analyze the impact of the planarity constraint on the search-tree of all topological maps consistent with exploration experience. Experiments demonstrate excellent results even in artificial environments where topological mapping is particularly difficult due to large amounts of perceptual aliasing and structural symmetry. This work is also reported in (Savelli and Kuipers, 2004).

3.1 Introduction

In Chapter 2 we developed a topological map-building system based on the hybrid SSH. Apart from the implemented preference policy, its functioning can be broadly summarized as (1) parsing the symbolic input sequence of observations (Figure 2.4, page 38), (2) pursuing all or some possibilities to close loops, and (3) enforcing consistency checks, both about observations and ontological constraints assumed *a priori*. Ambiguity about loop-closing is responsible for branching in the search-tree (step (e) in Section 2.3.1).

Planarity is a simple property of many real-worlds environments, which helps to reduce loop-closing hypotheses considerably. Planarity holds when the graph-abstraction of the environment can be drawn on a plane without crossing edges. Therefore, if we know that the environment is planar, there are fewer ways to allow spatially for a new loop during the topological search. If otherwise we believe that the environment is likely to be planar, but are not completely sure of that, the mathematical theory of graphs and topology provides a definition of the degree to which a graph is not planar.

These considerations show that the notion of planarity is well suited for the SSH theory

(Remolina and Kuipers, 2004). Planarity can play the role of strict constraint, or be just a default contributing to the preference policy.

Figure 3.1 illustrates a physical exploration scenario already introduced in the previous chapters, which motivates the application of the planarity constraint. In this case, the unintended consistent (but non-planar) model could be also discarded using odometrical information to enforce global metrical consistency, by means of state-of-the-art SLAM methods. However, the planarity test has the advantage of being independent of the actual geometric properties and metrical scale of the real environment, as it yields a boolean result on the available abstract topological representation. Besides, it can be performed in time and space *linear* in the number of topological links. If it is dynamically evaluated as the map grows during the topological search — exploiting past computations when a new element is inserted — a *sublinear* ($O(log^2(n))$ worst-case time) algorithm can be used (Italiano et al., 1993).

In this chapter we formally and experimentally analyze how the planarity constraint prune the search-tree. Experiments confirm the implications of the theoretical analysis, and show an impressive reduction of the space of topological hypotheses in particularly difficult types of simulated environments, where the agent is assumed to be faced with poor perceptual characterization of distinctive locations and landmarks, large amount of perceptual aliasing, multiple nested loops, and structural symmetry. In some relevant cases, the planarity test makes it possible to find unique optimal maps.

The rest of this chapter is organized as follows. In Section 3.2 we define the concept of planarity mathematically on embedded graphs, which are combinatorial structures well-suited to represent topological maps. Section 3.3 and 3.4 respectively provide formal and experimental investigation of the reduction of the topological search-space when the planarity constraint is enforced. Finally, Section 3.5 reviews some relevant related work.

3.2 Embedded Graphs and Planarity

We recall here basic mathematical concepts of graph theory and topology. These have been widely studied in mathematics and, more recently, computer science; for further details see for example (Gross and Tucker, 1987).

An *undirected graph* $G = (V, E)$ is composed of a set of vertices V and a set of undirected edges E, each edge linking two vertices. In a *directed* graph, each edge e is strictly directed; it starts from a vertex $source(e)$ and ends at a vertex $target(e)$. An *embedding* of an undirected graph G defines a clockwise circular order over all the edges incident on a vertex v, around each v of G. An *embedded graph* is an undirected graph equipped with an embedding. When embedded graphs are studied as special structures themselves, they are also called *maps*. Although this terminology witnesses how such structures are mathematically relevant to our

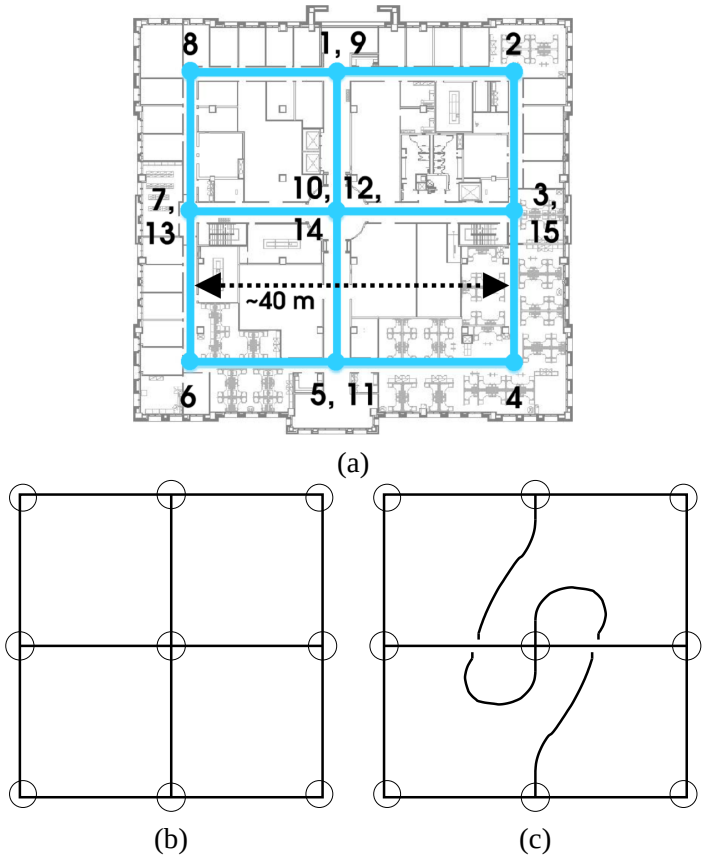

Figure 3.1: **Non-planar Hypothesis.** This figure is the same as Figure 1.2, and reported here just for convenience. Both (b) and (c) are among the consistent topological hypotheses after the complete physical robot exploration of the environment (a). Due to perceptual aliasing of places and the symmetrical structure of the environment, any further exploration will give the same sequence of observations, making it impossible to discriminate between the two maps. Planarity constraint is an inexpensive alternative to using metrical information or raw-data matching to discard (b), and applies directly to the topological abstraction. Indeed (a) and (b) represent the same graph, but considered as *embedded* graphs, (a) is planar, while (b) is not.

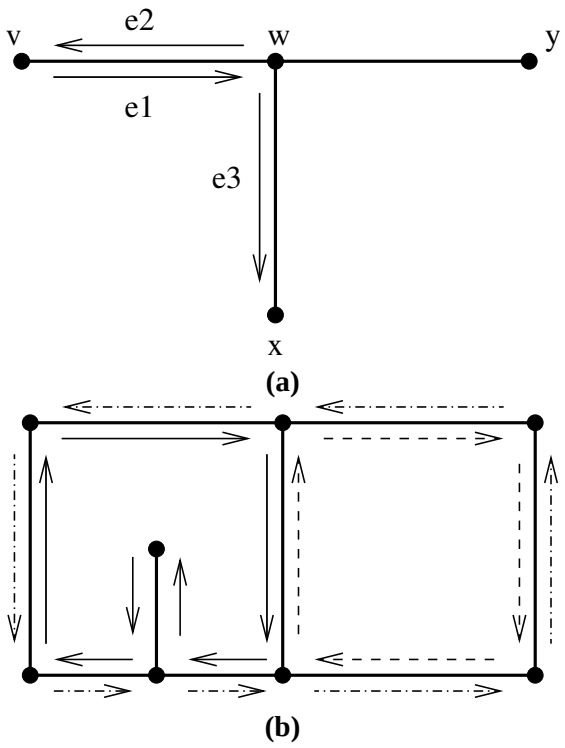

Figure 3.2: **Face Tracking.** If edges are imagined as corridors, face tracking can be intuitively viewed as walking along loops always following the right wall. **(a)** $e2 = rev(e1)$, $e3 = pred(e2) = pred(rev(1)) = next(e1)$. **(b)** the three faces of a simple embedded graph.

purposes, we shall keep using the term 'map' in the more general sense intended in "SSH *map*" or "robotic *map*-building".

An embedded graph is *planar* on an orientable surface S if it can be drawn on S enforcing its embedding without edge-crossings. The *genus* of an embedded graph G is the minimum genus among those of all the orientable surfaces on which G is planar. When G is planar on the plane, its genus is 0 and we simply say it is planar. Note that the plane and the sphere both have genus 0, so G is planar if and only if it is so on the sphere too. The torus is a surface with genus 1.

We describe here the classical method for deciding the planarity of an embedded graph.

For any embedded graph $G = (V, E)$, consider the transformation T such that:

1. $T(G) = (V, T(E))$ is a bidirected graph obtained replacing each undirected edge $e \in E$ that links vertices v and w by two directed edges e_v and e_w, where $source(e_v) = target(e_w) = v$ and $source(e_w) = target(e_v) = w$.

   $rev : T(E) \to T(E)$ is a *one-to-one* function such that $rev(e_v) = e_w$, and $rev(e_w) = e_v$. Observe that $rev(rev(x)) = x$ holds by construction.

2. The circular order of undirected edges $[e^1 \dots e^k]$ around each vertex v defined by the embedding in G is reflected by the circular order $[e_v^1 \dots e_v^k]$ of outward directed edges leaving v in $T(G)$. (By construction this will then hold for the inward directed edges as well.)

   The mapping $succ : T(E) \to T(E)$ is a *one-to-one* function that associates to any directed edge e_v the next outward directed edge in the circular clockwise order around v. Analogously for $pred : T(E) \to T(E)$, counterclockwise.

Computing the genus of an embedded graph G requires us to count its *faces*, which we define on the auxiliary structure $T(G)$. Consider the following function $next : T(E) \to T(E)$ such that $next(e_v) = pred(rev(e_v))$, i.e., the directed edge $next(e_v)$ is the predecessor of $rev(e_v) = e_w$ in the circular order of edges around $w = target(e_v)$ (Figure 3.2 (a)). Since both $pred$ and rev are total and bijective functions, $next$ will be so too, *i.e.* $next(e_v)$ always exists and is unique for any $e_v \in T(E)$. We can formally define a face as follows (Figure 3.2 (b)).

Definition 3.1. *Let $G = (V, E)$ be a directed graph, T the transformation specified above, and $[e_1, e_2 \dots e_n]$ a list of directed edges of $T(G)$ such that*

$$\begin{array}{ll} \textit{(i)} & e_{i+1} = next(e_i) \quad \textit{for all} \quad i = 1 \dots n-1 \\ \textit{(ii)} & e_i \neq e_j \quad \textit{for any} \quad 1 \leq i < j \leq n \\ \textit{(iii)} & next(e_n) = e_1 \end{array}$$

The directed loop formed by the sequence of edges in such a list is a face of $T(G)$.

This is a combinatorial definition of faces. If we consider a surface on which an embedded graph is planar, we can geometrically characterize faces as the partitions induced on the surface by the images of the edges. The infinite face of a planar embedded graph on the plane can be avoided if the drawing is considered on the sphere. Faces can then also be intuitively viewed as the inside of the "least" loops, that is, those loops that do not contain other loops.

The genus of an embedded graph G is given by the following (Euler-Poincaré) formula:

$$genus(G) = (|E| + 2c - |V| - |I| - f)/2,$$

where I is the set of isolated vertices, c the number of connected components, and f the number of faces. Since we shall only deal with totally connected ($c = 1,\ |I| = 0$) embedded graphs — they model maps from single-agent explorations — planarity (genus=0) can be tested by:

$$|V| - |E| + f = 2,$$

the well known Euler's formula on the number of faces in a polyhedron.

If unlikely non-planar maps must be taken into account, the planarity constraint might be relaxed, by simply preferring maps with smaller genus. Indeed, the genus of an orientable surface indicates the number of "holes" the surface exhibits; therefore the genus of an embedded graph is by definition the minimum number of such holes that are needed to avoid edge-crossings.

The number of faces can be counted in linear time and space (in the number of edges) following the mathematical formulation given above, based on the function $next$.

3.3 Search-Space Reduction

The number of maps that the planarity constraint discards among all the topological hypotheses depends on many factors, including the nature and structure of the environment, the number and location of perceptually aliased places and the particular exploration route. Nevertheless, it is possible to formalize the reduction of the branching factor at a node in the search-tree, in a fashion that sheds light on the role of the past choices that have shaped the map at that node.

To this purpose, we consider the embedded graph that underlies an hybrid SSH topological map, so that we can reason with the combinatorial structures and properties introduced in Section 3.2.

The relationship between an hybrid SSH topological map M and its underlying embedded graph G is straightforward:

(a) Each place p in M corresponds to a vertex v of G.

(b) Each travelable oriented local path olp (one with a gateway, see Section 2.2) in the star of a place p of M corresponds to an edge e around the vertex v corresponding to p in G. This can lead to two cases:

1. olp is linked to olp' of a place p', let v' be the corresponding vertex. Then e dubs both olp and olp', and links v and v'.
2. olp is a pending oriented local path (see Section 2.3.1), then we call e a *pending edge*, which is not connected to any vertex but v.[1]

(c) The circular clockwise orders in the stars of M are simultaneously preserved by the embedding of G — by (a) and (b).

Note that the notion of topological map strictly subsumes that of embedded graph, since a star adds to the clockwise order of its oriented local path also information about the continuation of a local path across the place, dividing the other oriented local paths into two disjoint sets: on the right and on the left (see Section 2.3.3). This is not the case with the edges around a vertex of an embedded graph.

Since pending edges are not regular edges of the embedded graph, they must be skipped by functions $succ$, $pred$, and $next$ while counting faces. When a pending oriented local path is linked, its corresponding pending edge will become a regular edge and will then have to be taken into account by these functions. In the meantime, the pending edge occupies and predicts the position of such a regular edge in the circular order around its only vertex.

A map (embedded graph) is *closed* if it has no pending oriented local paths (pending edges). A closed map implicitly represents the entire environment.

Every pending edge in a non-closed embedded graph G falls inside exactly one single face of $T(G)$. We formalize this notion as follows.

Definition 3.2. *Consider an adjacent pair $e_{i+1} = next(e_i)$ of directed edges of the loop defining a face F, and let v be the vertex $v = target(e_i) = source(e_{i+1})$. Any pending edge of v that comes after e_{i+1} and before $rev(e_i)$ in the circular clockwise order of outward directed edges around the vertex v belongs to F.*

As a particular case of the definition above, when only one edge is currently incident on v, *i.e.*, when $e_{i+1} = rev(e_i)$, then any pending edge of v trivially belongs to the face.

It follows that the set P of all the pending edges of a map is partitioned into the sets $P_1 \dots P_f$, each P_i collecting the pending edges belonging to face i, $i = 1 \dots f$, where f is the number of faces. Let $p_i = |P_i|$ and $p = |P|$.

Now, consider a node of the search-tree where the next exploration step leaves the current place through one of its pending oriented local path (step (e) of the algorithm). The branching factor of this node is given by the number of possible successors of this map that are consistent with the resulting observation. One successor is always the map where the arrival is at a totally new place. The others are maps where the move is to a place that is already represented in the

[1] Note that we are being slightly liberal with the notion of embedded graph here, since such structures do not traditionally include "pending edges". These are just meant to simplify the description of the next steps. The whole section can be rephrased avoiding the introduction of pending edges, though less concisely.

map. First, such a place must have been approached through one of its pending oriented local paths, and its star must be consistent with the observation made. Second, if planarity is assumed, in the underlying embedded graph *the two pending edges to be linked must belong to the same face*, otherwise they would be unified in an edge that would cross the loops defining their faces.

Assuming recursively that the map at the current node of the search tree is already planar, we are interested in the reduction of the branching factor when planarity is enforced also in its successors. In particular we consider the ratio b_p/b where b_p is the number of planar children of this node and b the number of all its children, including the non-planar ones.

Assume a worst-case scenario where every pending edge corresponds to a potential arrival compatible with the next observation. Then b is equal to p, minus the one pending edge where the move starts, plus the map with a totally new place, i.e., $b = p$. Analogously, if the starting pending edge belongs to the ith face, $b_p = p_i$. Then $b_p/b = p_i/p$. Denoting the average number of pending edges per face $\bar{p} = p/f$, we have $b_p/b = p_i/\bar{p} \cdot f$.

From the recursive assumption of planarity, we have $f = |E| - |V| + 2$ (before including the new edge in E). The interesting point here is that the difference $|E| - |V|$ depends on the choices made at past branching points leading to this map. Every time it was decided to explain the observation at hand by linking the place with one already encountered, unifying two pending edges, $|E| - |V|$ increased by one unit. If instead it was decided to build a totally new place, then a vertex and an edge were inserted at once, and $|E| - |V|$ did not change. Considering that at the root of the search tree the first embedded graph has just one vertex and no edge, $|E| - |V| = m - 1$ where m is the number of matches ever made between existing places, as in the first case above. Thus, finally, the dynamic ratio that quantifies the branching factor reduction is:

$$\frac{b_p}{b} = \frac{p_i}{\bar{p}} \cdot \frac{1}{m+1}, \tag{3.1}$$

The first factor normalizes according to the relative number of consistent pending edges in the current face w.r.t. the average face. The second factor confirms the intuition that the more loops have been closed, the more topologically compact the map must be, and therefore the fewer ways there are to close new loops while preserving planarity.

Note that we get the same result if we replace the worst-case assumption that all pending edges are consistent with the observation with a more moderate assumption that pending edges consistent with the observation are uniformly distributed throughout the map. In this case, we restrict the numbers p_i, p, and $\bar{p}$ to only those pending edges consistent with the observation.

3.4 Experimental Results

We have carried out several experiments to investigate in detail how planarity testing improves the topological search.

Environment Structure:	3×3			3×4		
Planarity Pruning:	no	yes	$\sim$ red.	no	yes	$\sim$ red.
Search-tree nodes built:	317	240	24%	4,239	1,722	59%
Maps built:	188	147	22%	3,400	1,425	58%
Final maps (all):	88	62	29%	1,465	507	65%
Final maps (closed):	6	3	50%	16	3	81%
Optimal Final Maps:	2	1	/	8	1	/

(a)

Environment Structure:	4×4 (up to 16 places)		
Planarity Pruning:	no	yes	$\sim$ red.
Search-tree nodes built:	192,569	8,452	95%
Maps built:	93,034	4,951	94%
Final maps (all):	23,685	545	97%
Final maps (closed):	6992	58	99%
Optimal Final Maps:	2	2	/

(b)

Figure 3.3: **Search-Space Reduction.** Breadth-first search expands the entire space of topological hypotheses permitting to observe the exact reduction of the search-space when planarity is assumed. The 4×4 results regard only the portion of space of maps with less than 17 places, and show that the planarity constraint is more effective as more loop-closures are forced. Since minimality is related to loop-closing, the planarity constraint will prove especially useful in practical cases when a best-first search is employed.

The topological map-builder does not assume a particular global structure for the environment. However, for any topological mapping algorithm, the worst-case environments will be those with large amounts of perceptual aliasing and structural symmetry — because it may be difficult or impossible to refute incorrect hypotheses.

We evaluate this algorithm using simulated square and rectangular grids, to maximize the difficulty facing the algorithm. The only places are corridor intersections, with "L", "T", or "+" structure. As can be easily noted, here every place is perceptually aliased multiple times throughout the environment, and the global structure of the environment has multiple embedded loops and strong symmetries along multiple axes. This kind of abstract environment allows a fair and straight comparison of experimental results as the environment scales topologically. Besides, its pattern is relevant to several classes of real-world environments, such as outdoor urban layouts, or indoor large libraries with long corridors and shelves that strongly limit the agent's sensory horizon in most of the locations.

The results we report are for a "snake" pattern of exploration. The agent starts from a corner and walks along all the horizontal corridors, in alternate directions, moving from a corridor to the next parallel one when it reaches a "T" (just a corner the first time) intersection. Then it starts an analogous exploration of the vertical corridors. We believe that this kind of exploration would also prove hard when a pure metrical mapping method is used that closes loops based on a maximum-likelihood choice. Indeed, if corridors are long enough, the inevitable angular odometrical error might often support the hypothesis that the agent is back to a "+" intersection on a parallel corridor. If the features available at intersections in combination with the robot's perceptual capabilities do not make possible to distinguish different "+" junctions (because of poor sensory system, or high image variability over multiple visits of the same place), the robot could not know of the error, as in our simulation. In this case, if maximum-likelihood is used for a greedy on-line search, the map would be irreversibly affected by such an error.

We assume the agent can acquire the correct star ("L", "T", or "+") of the place being visited.

We have run the experiments on an implementation of the SSH topological breadth-first search, with and without planarity testing. The reason for experimenting with breadth-first search is that it builds the complete set of the current topological hypotheses on-line: this allows us to observe the exact reduction of the whole search-space when planarity pruning is applied. We have gathered data about the number of (1) search-tree nodes, (2) maps ever built (these are usually fewer than search-tree nodes because when a map already correctly predicts the result of an action in step (p-c) of the algorithm in Section 2.3.1, only the assumed current position changes from node to node), (3) final maps (those on the consistent leaves of the search-tree at the end of the exploration), (4) final maps that are also closed, and (5) final maps that are also optimal (according to the preference policy, see Section 2.3.2). The tables in Figure 3.3 collect the results for the environments discussed below.

In Table (a) the first column illustrates the results for a 3×3 grid (see Figure 3.1(a)). Planarity testing discards half of the final closed maps, and makes it possible to determine the

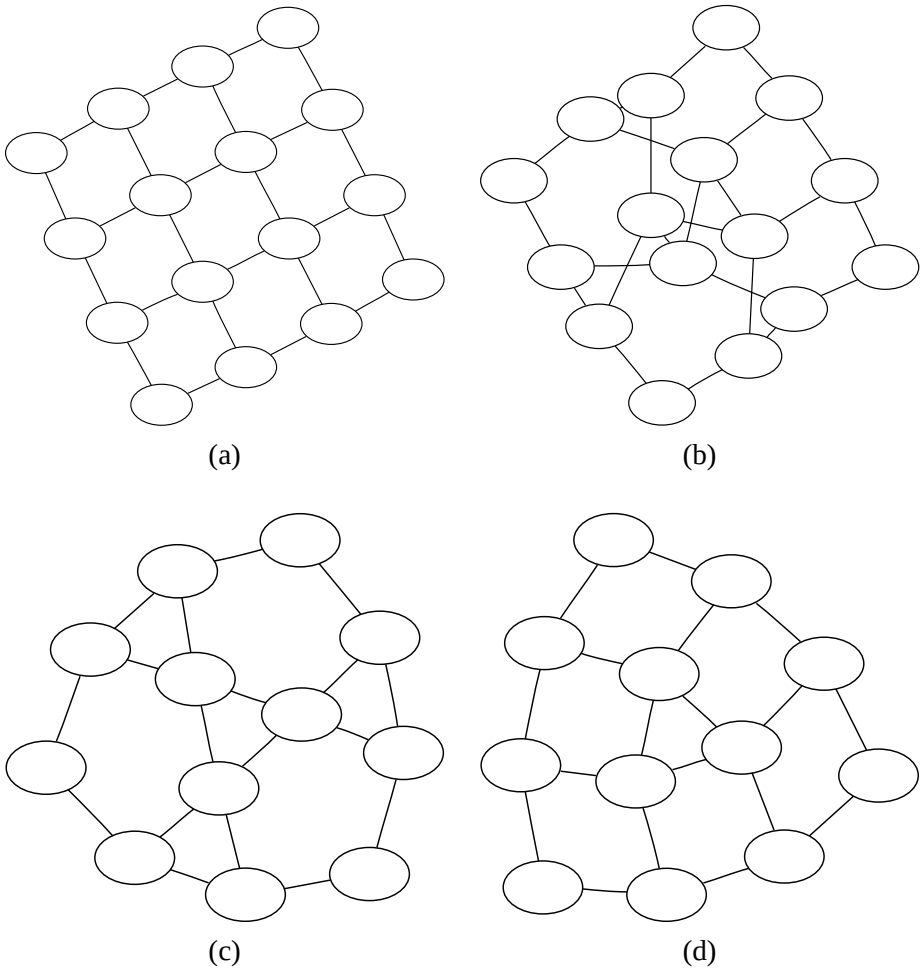

Figure 3.4: 4×4 **Grid Models.** After the "snake" exploration of the 4×4 grid, and considering only closed and qualitatively perpendicular maps, (a) is the unique (and correct) planar map. (b) is one example of the 255 non-planar "relatives" of (a) that are filtered out by the planarity constraint. Without the qualitative perpendicular requirement, there are two consistent "over-minimal" maps with only 12 states, (c) and (d). Sufficient further exploration can rule out (c), but not (d).

correct model as the unique optimal final map. (The two optimal maps obtained when planarity is not assumed are those in Figure 3.1(b,c).)

The second column presents results from a single complete exploration of a more topologically complex environment: a 3×4 grid of places. The gain in reducing the search is larger than in the single exploration of the 3×3 grid. This is due to the branching-factor reduction, which makes planarity testing *exponentially more efficient* as the exploration proceeds.

In Table (b) we address a 4×4 grid. Note that, since now the two axes have same length, the environment has one more degree of symmetry, and the chances of ruling out wrong models by inconsistencies along observation-sequences is much reduced. A breadth-first search of the maps consistent with a complete exploration does not terminate in reasonable time. However, we have driven the search so as to gather statistics significant for our purposes. We have asserted an upper bound of 16 on the number of places a map can have. In our case this was intentionally chosen equal to the actual number of places in the assumed physical environment. During the search any partial map that grows over this limit is discarded. Since during the search the number of places in a partial map never decreases, *no final map with less than 17 places that would be produced with the complete search is lost.* (Note that having an upper bound — not necessarily picking the exact correct number of places — could be a reasonable strategy in some practical situations. The completeness of the search under this limit would then prove useful.)

Besides the previous considerations we can observe here that as the search proceeds, only those maps that account for a certain minimum number of loop-closures can remain under the upper-bound and not be ruled out. This means that the number of matches m in the denominator of the dynamic ratio formalized in Equation (3.1) must grow at a certain average rate. Therefore, the dramatic performance of planarity pruning in this last experiment is coherent with our formalization, and with the intuition that in more compact maps there are far fewer ways to allow spatially for a new link while avoiding edge-crosses.

The two optimal final maps found with planarity pruning are "overminimal". That is, due to symmetry and aliasing in the actual environment, the maps are complete and consistent with exploration experience, but they have fewer places than the actual environment (Figure 3.4(c,d)). Note that this example illustrates the kind of extreme structural symmetry that requires a portable marker to find the correct map, as in (Dudek et al., 1991).

We went further in making sense of the 4×4 grid exploration data. Among *all the final maps* we have considered those closed maps that are "qualitatively perpendicular", i.e., no two global paths intersect each other in more than one place and no three global paths form a triangle. There is only one such map (the correct model of the environment) in the case of planarity pruning, as opposed to 256 such maps in the other case (all with 16 places). The unique planar solution is shown in Figure 3.4(a), while a map among the 255 non-planar ones is shown in Figure 3.4 (b). These numbers provide a concrete insight into how many non-planar "close relatives" of the correct topological model can arise in such a symmetrical environment.

We have observed that a *best-first* search, where the search-tree expansion is prioritized by the optimality policy, in some cases produces the same results while reducing the time and space

used dramatically. Furthermore, since minimality of the number of places is a component of the optimality policy, by similar reasoning as with the upper bound above, we expect the maps thus prioritized to be those where m is larger and so the planarity constraint tends to do more work (by Equation (3.1)). That is, best-first search better leverages the potential of the planarity assumption. Since best-first policy is more likely to be employed, perhaps additionally informed by metrical consideration as we do in Chapter 4, to deal with practical cases, planarity testing can have an even more effective role in such cases. Applied to the snake-exploration of the 4×4 grid *without upper bound on the number of places*, with pruning of non-planar and non-qualitatively-perpendicular maps, and with backtracking if the final solution is not closed, our C++ implementation of best-first search determines the correct model in 0.84 sec. on an Intel Pentium 1.5 Ghz. However, it does not always ensure similar benefits over breadth-first search as the exploration/environment grows in size and topological ambiguity.

The relationship between the particular preference policy chosen and optimal-completeness of the search (which is guaranteed by exhaustive searches such as breadth-first) requires more investigation (although a consistent solution is always found, if one exists).

The cases we have considered here are extreme in the weak perceptual characterization and adversity of the environments being explored. They are meant to investigate the potential of topological mapping. In real cases, *any* kind of additional feature or irregularity in the structure can be exploited to test and refute those incorrect models that we have assumed impossible to discriminate (see for example Kuipers et al., 2004; Kuipers and Beeson, 2002).

3.5 Related Work on Loop-Closing

Loop-closing has long been recognized as a central problem in the topological mapping literature (Kuipers and Byun, 1991; Dudek et al., 1991; Choset and Nagatani, 2001).

Following Kuipers and Levitt (1988), where the topological nature of cognitive maps is pointed out, embedded graphs as representations of topological maps were proposed by Dudek et al. (1991). They show that correct map-building is impossible in general if only the cyclic order of the incident edges is used to recognize a place, unless the agent is provided with a portable marker it can drop and pick up. In this case the agent can cope with perceptual aliasing and symmetries, and learn the correct topological model. The upper bound on the length of the exploration with a portable marker reduces from polynomial to almost linear in the size of the graph, when planarity is assumed (Rekleitis et al., 1999); a similar improvement is obtained also for the map-validation problem (Deng et al., 2001).

The approach presented by Dudek et al. (1993) is closer to our work in that there is no recourse to portable markers. It provides an algorithm that expands a tree of hypotheses about loop-closures, but the role of planarity in branching-factor reduction is not investigated, and a

rich spatial ontology that allows for a preference policy is not used.

Because of the negative result above, several works have addressed topological map-building as the problem of learning the minimal "discernable" structure of the environment, i.e., its smallest underlying automaton: for example, Rivest and Schapire (1989) address the deterministic case, while the case of stochastic/noisy observations is introduced by Dean et al. (1995). These and smilar approaches based on the minimal automaton representation, however, would yield wrong results with environments such as grids, which can have overminimal isomorphic models (Figure 3.4 (d)). Dealing with such ambiguous topologies is impossible without considering metrical information or matching richer perceptual representations (or using a portable marker, see above). The work presented in this chapter, however, demonstrates that ontological considerations can be exploited to drive and prune the topological search inexpensively, with impressive results.

Chapter 4

Topological Reasoning and Metrical Analysis

In this chapter we propose a probabilistic framework for modeling and solving large-scale topological ambiguity. The method introduced combines topological mapping and reasoning, as presented in the previous chapters, with modern approaches to metrical mapping. We present formal analysis, underlying intuitions and rationale, and preliminary experimental studies.

4.1 Introduction

Most approaches to SLAM (Simultaneous Localization And Mapping) employ probabilistic methods to build an accurate metrical model of the environment from noisy metrical data, provided by motion and range-finder sensors. This task is usually accomplished by determining the correct pose of some distinctive landmarks in a cartesian frame of reference, or by integrating the sensorial snapshots collected along the exploration into a large-scale occupancy-grid map, which represents the occupied/free state of space at a pixel resolution level.

In this context, ambiguity and possible aliasing in sensing and perception is referred to as the *data association* or *correspondence* problem. The first term borrows from sonar and radar literature (Bar-Shalom and Fortmann, 1988), where it refers to the problem of associating intermittent measures with the correct objects being tracked, while also allowing for objects falling in and out the scene. A very similar problem faces a robot that needs to correctly map landmarks that appear identical, and is also related to the uncertainty about whether the territory being explored is new or overlaps some already visited region, when building an occupancy-grid map.

Pure topological mapping, including our work presented in the previous chapters, is complementary of such metrical approaches in more than one way. Metrical methods very effectively cope with metrical uncertainty when data association is known. Data-association ambiguity,

when it is at all addressed, is usually solved by *ad hoc* techniques "plugged in" the framework at hand. Conversely, pure topological mapping is primarily concerned with the number of the distinctive places of the environment, and their correct qualitative relations — connection, order, containment — but does not utilize metrical information.

Obviously, if the metrical layout of a topological map is estimated by applying SLAM methods, this could be used to evaluate the geometrical plausibility of the closed loops, according to the triangle property. At the same time, purely qualitative constraints and preferences in topological mapping might help to further discriminate between different topological maps, when metrical relations are too uncertain. They can also improve the search process; for instance, non-planar maps can be inexpensively discarded beforehand, where this makes sense.

We point out that topological ambiguity and data-association uncertainty affect the model within which metrical uncertainty is traditionally modeled and solved. In particular, this model can be represented as a *Bayes network* (Pearl, 1988), which is determined by the assumed topological map. The problem of finding the right topological map can then be viewed as a problem of model selection, or Bayes-network learning. This is widely recognized in the AI community as a very hard problem (Friedman, 1998; Jordan, 1999; Russell and Norvig, 2002), and we observe that this is consistent with the well-known difficulty of coping with data-association uncertainty and topological ambiguity.

The formalization proposed represents the potential solution of the problem as the model maximizing the *a posteriori* probability conditioned to the observed metrical relations. The qualitative ontological and cognitive biases, such as the planarity constraint and the SSH preference criteria, form the *a priori* probability of the model. This allows us to combine the two worlds in a rigorous way, while taking advantage of the complementary benefits. After suitable approximations, the framework is compatible with some of the most successful relaxation-based probabilistic techniques recently introduced for SLAM (Frese et al., 2004; Konolige, 2004), which can thus be utilized towards effective implementations.

The rest of this chapter is organized as follows. We first model topological ambiguity and metrical uncertainty in the framework of the Bayes networks (Section 4.2). In Section 4.3 we propose a way to compute the *a posteriori* probability of a topological map (and its related Bayes network). This in turn requires us to compute the metrical likelihood of the map (Section 4.4) and its *a priori* probability (Section 4.5). Section 4.6 brings everything together and provides some more operational details, while Section 4.7 describes experimental results. Section 4.8 concludes the chapter with a discussion of related work that addresses the problem of spatial correspondence in robot map-building.

4.2 Topological Hypotheses as Bayes Networks

Consider a robot that can reliably detect distinctive places (some approaches are discussed in Section 2.5.2, page 49) and navigate between them when exploring a large-scale environment, as assumed by the hybrid SSH and more in general by any topological mapping system. The robot is supposed to estimate the distance and change of orientation between any two distinctive poses, in addition to their uncertainty. This estimate is expected in the form of mean values and covariance matrix, or a fully specified probability density distribution. To this purpose, well-known probabilistic filtering techniques can be applied to sensory readings, based on motion and sensor probabilistic models. We do not address here how this will be accomplished in the specific implementation, and we assume such information to be available as input. Virtually every metrical SLAM method that correct odometry at a local level might be used. Best candidates seem to be scan-matching (Lu and Milios, 1997), or the use of a "scrolling" local occupancy-grid map that allows continuous relocalization, as proposed by (Modayil et al., 2004) in the framework of the hybrid SSH. We focus on a more abstract level, more similarly to Smith et al. (1990).

We start with a very simple model.

For the n places hypothesized in the environment by a topological map, let $\overline{\pi} = \pi_1 \dots \pi_n$ be the random vectors of pose variables $\pi = (x_\pi,\ y_\pi,\ \theta_\pi)^T$ w.r.t. a global frame of reference. Each such pose can be also viewed as the origin and orientation of a local frame of reference. We assume π_1 to be the starting pose without loss of generality. The $\overline{\pi}$ contain hidden variables of our model, because they cannot be directly observed, except π_1, which we assume to be the origin of the global frame of reference (*i.e.*, $\pi_1 = (0, 0, 0)$). At each step $i = 1 \dots p$, according to the topological map at hand, the robot moves from a place with global pose $\pi_{s(i)}$ ('s' for start) to another with global pose $\pi_{e(i)}$ ('e' for end), with $\pi_{e(i)} = \pi_{s(i+1)}$. We call $\overline{\tau} = \tau_1 \dots \tau_p$ the random vectors of coordinate transformation variables for every such a move along the exploration; $\tau_i = (x_{\tau_i},\ y_{\tau_i},\ \theta_{\tau_i})^T$ represents the move between $\pi_{s(i)}$ and $\pi_{e(i)}$, and can be also viewed as the pose of $\pi_{e(i)}$ w.r.t. the local frame of reference at $\pi_{s(i)}$. The variables of $\overline{\tau}$ account for the observations, since their continuous probability distributions are estimated from sensorimotor data and models. This simple model is illustrated in Figure 4.1.

From a geometrical viewpoint, each τ is completely determined by the initial and final poses according to the following transformation:

$$\tau = (\ominus \pi_s) \oplus \pi_e$$

that is:

$$\begin{pmatrix} x_\tau \\ y_\tau \\ \theta_\tau \end{pmatrix} = \begin{pmatrix} (x_{\pi_e} - x_{\pi_s}) \cos\theta_{\pi_s} \ + \ (y_{\pi_e} - y_{\pi_s}) \sin\theta_{\pi_s} \\ -(x_{\pi_e} - x_{\pi_s}) \sin\theta_{\pi_s} \ + \ (y_{\pi_e} - y_{\pi_s}) \cos\theta_{\pi_s} \\ \theta_{\pi_e} \ - \ \theta_{\pi_s} \end{pmatrix} \tag{4.1}$$

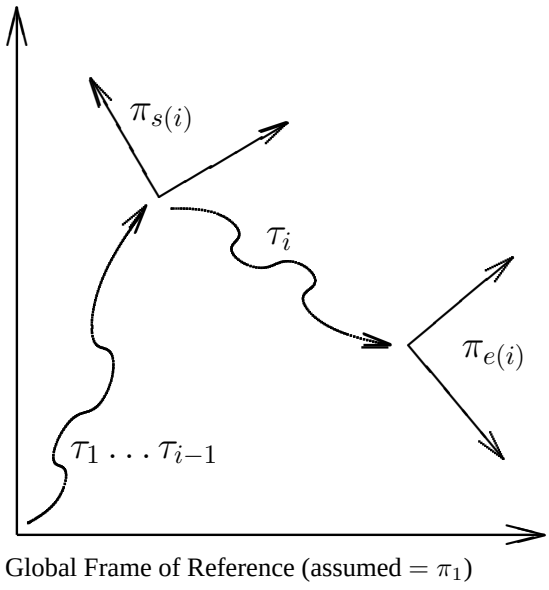

Figure 4.1: **Geometrical Model.** Every place in the topological map is associated with an unknown pose π w.r.t. a global frame of reference. At the ith step of exploration, depending on the topological hypothesis, the agent is assumed to move from one such pose $\pi_{s(i)}$ to another $\pi_{e(i)}$. The agent observes (with uncertainty) the related coordinate transformation τ_i, which is the pose $\pi_{e(i)}$ w.r.t. the local frame at $\pi_{s(i)}$.

Therefore, *each τ_i conditionally depends on $\pi_{s(i)}$ and $\pi_{e(i)}$, while it is conditionally independent of any other variable given these two*. The vectorial function $(\ominus\pi_s) \oplus \pi_e$ is non-linear; given a possible linearization

$$\tau = (\ominus\pi_s) \oplus \pi_e \approx \lambda(\pi_s, \pi_e) \tag{4.2}$$

we can specify the conditional distribution as a linear Gaussian

$$p(\tau \mid \pi_s, \pi_e) = \frac{1}{\sqrt{(2\pi)^3 \det C}} \exp\left\{-\frac{1}{2}\,(\lambda(\pi_s, \pi_e) - \tau)^T\, C^{-1}\,(\lambda(\pi_s, \pi_e) - \tau)\right\} \tag{4.3}$$

Here C is a 3×3 positive semidefinite covariance matrix that quantifies the uncertainty on τ. Although most of our analysis holds independently of the particular choice of $p(\tau \mid \pi_s,\ \pi_e)$, we will commit to (4.3).

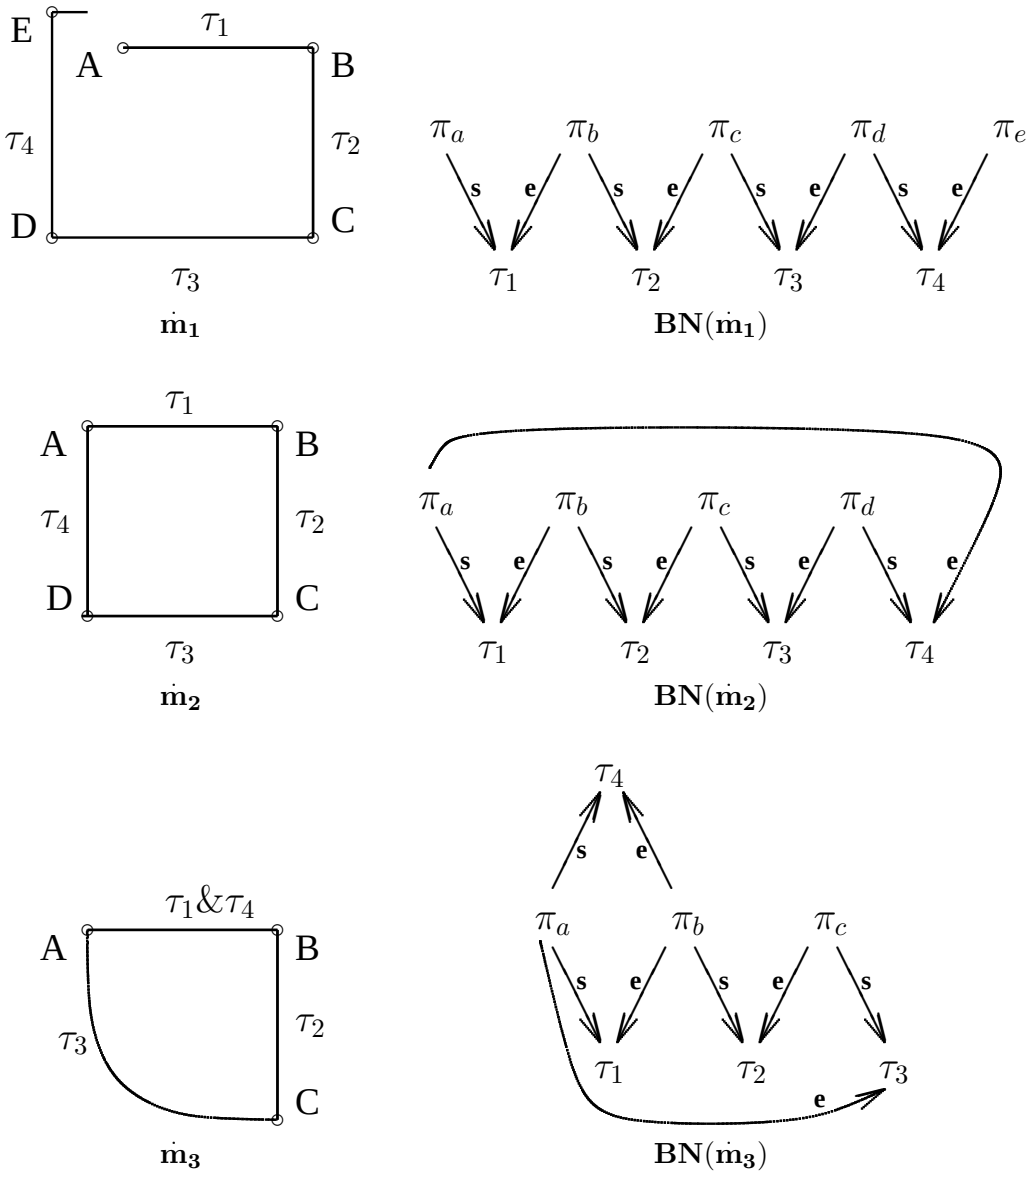

Figure 4.2: **Maps and Bayes Networks.** Maps $\dot{m}_1$, $\dot{m}_2$, and $\dot{m}_3$ are among the topological hypotheses after a square exploration; $\dot{m}_2$ is the correct interpretation. $BN(\dot{m}_1)$, $BN(\dot{m}_2)$, and $BN(\dot{m}_3)$ represent the full joint distributions over the domain variables for $\dot{m}_1$, $\dot{m}_2$, and $\dot{m}_3$. Spatial uncertainty is twofold. Metrical uncertainty is modeled and handled within one $BN(\dot{m})$. Topological ambiguity affects the topology of $BN(\dot{m})$, *i.e.*, the conditional (in)dependency assumptions for poses and metrical relations. Determining the correct topological map then amounts to identifying the right probabilistic model of metrical uncertainty.

All this information is systematically represented in the framework of the Bayes networks, which permits compact representations of joint probability distributions by exploiting conditional independence between variables (Pearl 1988, see also Appendix A). In our case, a particular topological map $\dot{m}$, defining a hypothesis about the number of distinctive places/poses and the order in which they have been visited, can be always associated with a Bayes network, call it $BN(\dot{m})$, as shown in Figure 4.2 and explained below. The correspondence between such maps and Bayes networks is *one-to-one*. By $BN(\dot{m}, \overline{C})$, besides denoting $BN(\dot{m})$, we additionally specify that the covariance matrices $\overline{C} = C_1 \dots C_p$ for $\overline{\tau} = \tau_1 \dots \tau_p$ are used for defining the conditional distributions, for example as in Equation (4.3).

Nodes of $BN(\dot{m})$ are both observed variables $\tau_1, \dots \tau_p, \pi_1$ and hidden ones $\pi_2 \dots \pi_n$, one for each different place hypothesized. Arrows represent conditional dependency relationships, which reflect the particular route and loops hypothesized: for each τ_i, we have two arrows from $\pi_{s(i)}$ and $\pi_{e(i)}$ to τ_i.

This class of Bayes networks naturally represents the type of probabilistic model in which an instance of SLAM is solved according to the approaches in the tradition of (Lu and Milios, 1997), when data association, *i.e.* the topological structure, is known. Likewise, with little additional sophistication, it can be extended to model probabilistic-filtering based approaches as well.

Most important, Bayes networks lead to a formal and explicit account of the fundamental difference between the two factors of uncertainty involved in map-building: the number and topological relationships of the observed places, and their global metrical layout. The latter is modeled as a problem of Bayesian inference given the observed data, and can be handled within the probabilistic model defined by the Bayes network. The former affects the probabilistic model itself, posing a much harder challenge. In fact, uncertainty about the topological configuration of physical space is reflected in the topological layout of the network, raising ambiguity concerning the actual conditional (in)dependence relationships (Figure 4.2). This means that the Bayes network, besides being the tool for representing and solving *metrical* uncertainty, becomes *topologically* uncertain itself, and in turn must be probabilistically estimated.

In this respect, identifying the correct topological map belongs to a class of problems well known to be very hard to deal with, usually referred to as *model selection*, *statistical learning*, or, when Bayes networks are employed, *Bayes-network learning* [1].

Three cases are usually distinguished for a Bayes network to learn, listed below in order of increasing difficulty:

1. The structure, (nodes and arrows) is known, all variables represent observable (with uncertainty) quantities (case of complete data), and only the conditional probability distributions associated with the arrows need to be learned.

[1] While primarily studied in the community concerned with uncertainty in AI, their applications have recently grown fast, now covering areas such as bioinformatics, data mining, and many others.

2. Same as 1., but there are hidden variables, representing quantities that cannot be directly measured.

3. Same as 2., but the structure in not known in advance either.

The last case is the most adverse, and no general solution has so far been devised (but see Friedman, 1998). Although our problem falls into this category, it is mitigated by the fact that conditional probability distributions are actually known from the observations of the robot, according to Equation (4.1), for example in the form (4.3). Moreover, the Bayes network must belong to the particular class described above, which considerably constrain the hypotheses that can be made about its structure.

These arguments support the idea of working toward a computational framework that can leverage the peculiarities of our problem, which make possible to take advantage of a particular family of SLAM algorithms (Gutmann and Konolige, 1999; Konolige, 2004; Duckett et al., 2002; Frese et al., 2004), rather than resorting to general-purpose methods (Jordan, 1999).

Next section will address the characterization of the probability of a map and its Bayes network, to the purpose of comparing and selecting maps.

4.3 Characterizing Maps Probability

Given enough perceptual aliasing and metrical error, identifying the right topological map may be impossible. What we can do is to estimate the probability of each hypothesis conditioned to the observations, in the hope that the correct one will eventually dominate.

Selecting the most probable model given the observations is a criterion commonly adopted in many statistically modeled problems, sometimes called MAP (maximum *a posteriori*). It is not fully Bayesian, because later answer to spatial queries will not be computed in the form of expected values, over all the models weighted by their respective probability. However, in our case the topological model that correctly represents the physical world at the desired abstraction level is actually unique, and MAP selection also well fits a single-model committment, when this is required.

We need first understand how to assess and compute useful *a posteriori* probabilities for the maps at hand. To this purpose, we want to combine ontological biases and expectations, like those involved in the SSH and presented in the previous chapters, with the map's metrical plausibility evaluated from the (noisy) metrical information available.

Given the estimated mean coordinate tranformations τ_i^o ('o' for observed), their covariance matrices C_i^o, $i = 1 \dots p$ and a topological map $\dot{m}$ among those built by the topological mapping algorithm, we are interested in the following conditional probability

$$p(\dot{m} \mid \text{metrical information}) = p(\dot{m} \mid \overline{\tau^o},\ \overline{C^o}) \tag{4.4}$$

According to the Bayes rule, this can be factorized as

$$p(\dot{m} \mid \overline{\tau^o},\ \overline{C^o}) = \frac{p(\overline{\tau^o} \mid \dot{m},\ \overline{C^o})\ p(\overline{C^o} \mid \dot{m})\ p(\dot{m})}{p(\overline{\tau^o},\ \overline{C^o})}$$

Since the covariance C^o is independent of the global topological structure assumed, we have

$$p(\overline{C^o} \mid \dot{m}) = p(\overline{C^o})$$

and thus

$$p(\dot{m} \mid \overline{\tau^o},\ \overline{C^o}) = \alpha\ p(\overline{\tau^o} \mid \dot{m},\ \overline{C^o})\ p(\dot{m}) \tag{4.5}$$

where α denotes a constant quantity independent of $\dot{m}$

$$\frac{p(\overline{C^o})}{p(\overline{\tau^o},\ \overline{C^o})}$$

This factor can be obtained by renormalizing over the set of maps on the search-tree fringe that is developed by the topological mapping algorithm (see Section 2.3), *i.e.*, the entire space of hypotheses.

However, as long as we use probability (4.4) to *compare* different maps, or to find the map that maximizes such a probability, we can forget about α, that does not affect the outcome of such operations.

The factor $p(\overline{\tau^o} \mid \dot{m},\ \overline{C^o})$ is the likelihood, that is the probability that the observed $\overline{\tau^o}$ are generated along the hypothesized exploration and map $\dot{m}$. We will mostly refer to this term as *metrical likelihood*, because of the metrical nature of the observations here involved, and propose a way to compute it in the next section.

The rightmost factor $p(\dot{m})$ in (4.5) is the prior probability of the topological hypothesis being the true one, *a priori* of any metrical observations, *i.e.*, before the information provided by observing τ_i is ever taken into consideration. We claim that such a probability can play the role of quantitative account of the SSH preference policy and constraints introduced in the previous chapters. Those maps that violate the SSH constraints are implicitly interpreted as having 0-probability. Those left on the search-tree fringe have probabilities that are consistent with the ordering determined by the preference policy.

This approach might be faulted to translate a qualitative notion of preference into a real number, since the possibilities to do so are infinite and may lead to different outcomes when multiplied to the metrical likelihood in (4.5). Indeed, when comparing two maps $\dot{m}_1$ and $\dot{m}_2$, the metrical likelihood and the qualitative preference ordering might disagree about the most plausible model. When this is the case, say

$$\dot{m}_1 \text{ preferred to } \dot{m}_2 \quad \text{and} \quad p(\overline{\tau^o} \mid \dot{m}_1,\ \overline{C^o}) < p(\overline{\tau^o} \mid \dot{m}_2,\ \overline{C^o})$$

the actual mapping from the qualitative preference to the real values for $p(\dot{m}_1)$ and $p(\dot{m}_2)$ determines which map will dominate

$$p(\dot{m}_1 \mid \overline{\tau^o},\ \overline{C^o}) > p(\dot{m}_2 \mid \overline{\tau^o},\ \overline{C^o}) \quad \text{iff} \quad \frac{p(\dot{m}_1)}{p(\dot{m}_2)} > \frac{p(\overline{\tau^o} \mid \dot{m}_2,\ \overline{C^o})}{p(\overline{\tau^o} \mid \dot{m}_1,\ \overline{C^o})}$$

This means that two different numerical implementations $p(\dot{m}_1)$ and $p(\dot{m}_2)$ of the same qualitative ordering $p(\dot{m}_1) > p(\dot{m}_2)$ might yield two different final results.

We believe that there is no way round this "non-deterministic" situation, which is not to suggest that the approach is ill-posed. Indeed, the case above is inevitably part of the game, due to using probability theory to represent and evaluate beliefs by real numbers. Besides, there are two points worth noting.

First, the preference policy order is usually based on "the number of" occurrences of some spatial entity, such as places, paths, boundary relations (see Section 2.3.2). That is, it already owns some countably quantitive nature, making it easier to move into the realm of the real numbers.

Second, the preference policy is strongly based on the notion of minimality, that drives the focus on the current simplest, most compact, and consistent model, according to the well-known Occam's Razor principle. This type of bias is common in machine learning systems, because larger and more complex explanations can accommodate errors more easily, in the extreme case by listing all the past inputs. This situation is called *overfitting*, and requires a careful tradeoff between complexity/size and degree of fit. In our case, the maximally accommodating hypothesis is the topological map where no loop has ever been closed, always hypothesizing yet another new place at every step.

Whether Occam's Razor principle, more related to epistemological issues, can be legitimately expressed in probabilistic terms, by having the prior distribution penalize the complexity and size of a model, is currently debated. Some insights into this can be provided by an information-theoretic account of the relationship between the *minimum description length* principle and prior probabilities (a brief and gentle introduction is given in Russell and Norvig, 2002). We refrain from addressing this problem[2], far beyond the scope of our present work, although we adopt probabilistic penalization of complexity, when deriving the prior distribution from the SSH preference policy.

In the next section we make use of the formalization of $BN(\dot{m})$ presented in Section 4.2 for computing the metrical likelihood.

[2]Though, you may ask yourself if you would ever bet your money on a city map with a sole utterly long street and no loops at all, the gamble being on whether it represents a real city that was picked fairly randomly from a nation-wide geographic map. If you wouldn't, you are arguably closer to accepting that you have a prior in your mind about what a full city-layout should look like (*before* ever stepping in that city), that penalizes complexity based on your ontological expectations, at least consistently with the SSH way (for example, complex = many places with few loops).

4.4 Computing the Metrical Likelihood

We now focus on the metrical likelihood $p(\overline{\tau^o}|\dot{m},\ \overline{C^o})$. By our Bayes-network formalization of the joint probability distribution over $\overline{\pi}$ and $\overline{\tau}$, we define

$$p(\overline{\tau^o} \mid \dot{m},\ \overline{C^o}) \triangleq p^{BN(\dot{m},\ \overline{C^o})}(\overline{\tau} = \overline{\tau^o})$$

Here $p^{BN}(x)$ denotes the marginal probability distribution of the random variable x in the Bayes network BN, while $BN(\dot{m},\ \overline{C^o})$ denotes the Bayes network whose structure and conditional distributions are obtained from $\dot{m}$ and $\overline{C^o}$ as illustrated in Section 4.2.

Intuitively, a probability distribution conditioned to $\dot{m}$ and $\overline{C^o}$ is the same as the unconditioned evaluated in the probabilistic model provided by $BN(\dot{m},\ \overline{C^o})$. Indeed, $\dot{m}$ and $\overline{C^o}$ determine the probabilistic model $BN(\dot{m},\ \overline{C^o})$, but do not belong to its set of random variables [3].

By Bayes rule and marginalization we can write

$$p^{BN(\dot{m},\ \overline{C^o})}(\overline{\tau} = \overline{\tau^o}) = \int_{\mathbb{R}^{3n}} p^{BN(\dot{m},\ \overline{C^o})}(\overline{\tau} = \overline{\tau^o} \mid \overline{\pi})\ p^{BN(\dot{m},\ \overline{C^o})}(\overline{\pi}) d\overline{\pi} \tag{4.6}$$

$p^{BN(\dot{m},\ \overline{C^o})}(\overline{\pi})$ should be considered a uniform distribution, because there is no reason why one particular metrical layout for the distinctive poses ($\overline{\pi}$) should be more probable than another, if there is no conditioning to any τ. However, the marginal distribution in the analytic form above would be hard to turn into a computationally effective implementation, either in a closed form or by approximating techniques such as Monte Carlo integration. We then introduce a preliminary approximation, that will lead to an effective class of algorithms.

$$p^{BN(\dot{m},\ \overline{C^o})}(\overline{\tau} = \overline{\tau^o}) \approx \max_{\overline{\pi}} \left\{ p^{BN(\dot{m},\ \overline{C^o})}(\overline{\tau} = \overline{\tau^o} \mid \overline{\pi}) \right\} \tag{4.7}$$

Instead of computing the exact marginal distribution in the point $\overline{\tau} = \overline{\tau^o}$, in the approximation above the Bayes network is partially sampled assigning to $\pi_1 \ldots \pi_n$ those values that maximize the conditioned distribution in the same point. We provide more mathematical insights into this approximation in Section 4.4.1.

Exploiting conditional independence (see also Appendix A), (4.7) can be factorized and we have:

$$\max_{\overline{\pi}} \left\{ p^{BN(\dot{m},\ \overline{C^o})}(\overline{\tau} = \overline{\tau^o} \mid \overline{\pi}) \right\} = \max_{\overline{\pi}} \left\{ \prod_{i=1}^{p} p^{BN(\dot{m},\ \overline{C^o})}(\tau_i = \tau_i^o \mid \pi_{s(i)},\ \pi_{e(i)}) \right\} = \tag{4.8}$$

[3] As a matter of fact, there is no reason why we should not consider them random variables, since they appear as arguments of probability expressions, and range over specific domains. Yet, we want to set the uncertainty on the model $BN(\dot{m},\ \overline{C})$ apart from the probabilistic relationships embodied in $BN(\dot{m},\ \overline{C})$. One is the ambiguity about the topological map, while the probabilistic relationships concern this map's metrical layout and the observed metrical relations. Thus, if $\dot{m}$ and $\overline{C^o}$ were explicitly modeled as random variables, a *meta*-level should be distinguished in the overall ontology.

if we evaluate the metrical likelihood as in (4.3),

$$\max_{\overline{\pi}} \left\{ \prod_{i=1}^{p} \frac{1}{\sqrt{(2\pi)^3 \det C_i^o}} \exp\left\{ -\frac{1}{2} \left(\lambda(\pi_{s(i)}, \pi_{e(i)}) - \tau_i^o\right)^T C_i^{o-1} \left(\lambda(\pi_{s(i)}, \pi_{e(i)}) - \tau_i^o\right) \right\} \right\} \tag{4.9}$$

We first turn to find the point $\overline{\pi}$ in which the expression above is maximal. If we apply a monotonic logarithmic transformation, this amounts to maximize the *log-likelihood*:

$$\begin{aligned}
&\arg\max_{\overline{\pi}} \left\{ \log \left[\prod_{i=1}^{p} \frac{1}{\sqrt{(2\pi)^3 \det C_i^o}} \exp\left\{ -\frac{1}{2} \left(\lambda(\pi_{s(i)}, \pi_{e(i)}) - \tau_i^o\right)^T C_i^{o-1} \left(\lambda(\pi_{s(i)}, \pi_{e(i)}) - \tau_i^o\right) \right\} \right] \right\} = \\
&\arg\max_{\overline{\pi}} \left\{ \sum_{i=1}^{p} \log \left[\frac{1}{\sqrt{(2\pi)^3 \det C_i^o}} \right] - \frac{1}{2} \left(\lambda(\pi_{s(i)}, \pi_{e(i)}) - \tau_i^o\right)^T C_i^{o-1} \left(\lambda(\pi_{s(i)}, \pi_{e(i)}) - \tau_i^o\right) \right\} = \\
&\arg\max_{\overline{\pi}} \left\{ -\sum_{i=1}^{p} \left(\lambda(\pi_{s(i)}, \pi_{e(i)}) - \tau_i^o\right)^T C_i^{o-1} \left(\lambda(\pi_{s(i)}, \pi_{e(i)}) - \tau_i^o\right) \right\} = \\
&\arg\min_{\overline{\pi}} \left\{ \sum_{i=1}^{p} \left(\lambda(\pi_{s(i)}, \pi_{e(i)}) - \tau_i^o\right)^T C_i^{o-1} \left(\lambda(\pi_{s(i)}, \pi_{e(i)}) - \tau_i^o\right) \right\}
\end{aligned} \tag{4.10}$$

The expression to minimize in (4.10) has a physical analogy. It represents the energy of a network of $i = 1 \ldots p$ springs, each one with stiffness given by C_i^o. The solution for $\overline{\pi}$ that maximizes our metrical likelihood represents the geometrical configuration with minimal energy of such a system.

The final expression (4.10) can also be efficiently computed. In fact, it takes the form of a *linear* (following from $\lambda(\pi_s, \pi_e)$) *regression* problem, and in particular of a *least square errors/fitting* problem. Its solution is classically posed as that of a particular linear system (Press et al., 1992). Mathematical details, also covering the linearization $\lambda(\pi_s, \pi_e)$, are reported in Appendix B. Least square fitting is a well-known statistical technique, extensively studied and applied. Its application to the SLAM problem (Lu and Milios, 1997) has been recently addressed with excellent results by exploiting sparseness in the linearized system (Gutmann and Konolige, 1999; Konolige, 2004), or by relaxation methods (Duckett et al., 2002; Frese et al., 2004), which solve the resulting linear system in almost-linear time.

Finally, once the solution is obtained, it can be substituted in the argument of the $\max$ operator in (4.9) to obtain the desired metrical likelihood of m.

Next, Section 4.4.1 provides more mathematical insights into the metrical likelihood approximation used in (4.7).

4.4.1 More about the Approximation

We show that the approximation in (4.7) can be derived from the following one:

$$p^{BN(\dot{m},\,\overline{C^o})}(\overline{\pi}) \approx \delta^{\,\arg\max_{\overline{\pi}}\left\{p^{BN(\dot{m},\,\overline{C^o})}(\overline{\pi}\,\mid\,\overline{\tau}=\overline{\tau^o})\right\}}(\overline{\pi}) \tag{4.11}$$

By $\delta^{\,c}(x)$ we denote the Dirac function centered in c; the following properties hold for such a function:

1. $$\delta^{\,c}(x) = 0 \quad \forall x \neq c$$

2. In c there is concentrated a unit mass of probability, *i.e.*,

   a. $$\int_{c-\epsilon}^{c+\epsilon} \delta^{\,c}(x)dx = 1 \quad \forall \epsilon > 0$$

   b. $$\int_{c-\epsilon}^{c+\epsilon} \delta^{\,c}(x)f(x)dx = f(c) \quad \forall \epsilon > 0$$

Because of the approximation (4.11), the approach is not being purely Bayesian. Indeed, the marginal $p^{BN(\dot{m})}(\overline{\pi})$ — which in principle is supposed to cover uniformly the entire domain — is replaced by a distribution that concentrates all the probability in the point where $p^{BN(\dot{m},\,\overline{C^o})}(\overline{\pi}\,\mid\,\overline{\tau}=\overline{\tau^o})$ is maximal. The approximation is twofold: (1) a spurious conditioning is added to an unconditioned marginal distribution, and (2) the whole distribution is "shrinked" in a single point of the domain. Intuitively, this can be explained as getting over-confident in the observations $\overline{\tau^o}$.

Now, consider that

$$\begin{aligned}
&\arg\max_{\overline{\pi}} \left\{p^{BN(\dot{m},\,\overline{C^o})}(\overline{\pi}\,\mid\,\overline{\tau}=\overline{\tau^o})\right\} = \\
&\quad \text{(Bayes Rule)} \\
&\arg\max_{\overline{\pi}} \left\{\frac{p^{BN(\dot{m},\,\overline{C^o})}(\overline{\tau}=\overline{\tau^o}\,\mid\,\overline{\pi})\; p^{BN(\dot{m},\,\overline{C^o})}(\overline{\pi})}{p^{BN(\dot{m},\,\overline{C^o})}(\overline{\tau}=\overline{\tau^o})}\right\} = \\
&\quad (p^{BN(\dot{m},\,\overline{C^o})}(\overline{\tau}=\overline{\tau^o}) \text{ independent of } \overline{\pi} \text{ and } p^{BN(\dot{m},\,\overline{C^o})}(\overline{\pi}) \text{ uniform}) \\
&\arg\max_{\overline{\pi}} \left\{p^{BN(\dot{m},\,\overline{C^o})}(\overline{\tau}=\overline{\tau^o}\,\mid\,\overline{\pi})\right\}
\end{aligned} \tag{4.12}$$

Combining (4.6), (4.11) and (4.12) we get

$$\begin{aligned}
&p^{BN(\dot{m},\,\overline{C^o})}(\overline{\tau}=\overline{\tau^o}) \approx \\
&\quad \int_{\mathbb{R}^{3n}} p^{BN(\dot{m},\,\overline{C^o})}(\overline{\tau}=\overline{\tau^o}\,\mid\,\overline{\pi})\; \delta^{\,\arg\max_{\overline{\pi}}\left\{p^{BN(\dot{m},\,\overline{C^o})}(\overline{\tau}=\overline{\tau^o}\,\mid\,\overline{\pi})\right\}}(\overline{\pi})d\overline{\pi}
\end{aligned} \tag{4.13}$$

and by property 2b. above of Dirac functions this is equal to

$$p^{BN(\dot{m},\,\overline{C^o})}\left(\overline{\tau} = \overline{\tau^o} \mid \overline{\pi} = \arg\max_{\overline{\pi}}\left\{p^{BN(\dot{m},\,\overline{C^o})}(\overline{\tau} = \overline{\tau^o} \mid \overline{\pi})\right\}\right) \tag{4.14}$$

that is equal to the right-hand side of (4.7)

$$\max_{\overline{\pi}}\left\{p^{BN(\dot{m},\,\overline{C^o})}(\overline{\tau} = \overline{\tau^o} \mid \overline{\pi})\right\}$$

4.5 From the SSH to the Prior Probability

We now turn to the prior $p(\dot{m})$. In Section 4.3 we proposed that this should be derived from the semi-qualitative criteria involved in the SSH preference policy. We pointed out that there are infinite possibilities to do so, which may lead to different results, but we did not go into further details.

In our preliminary implementation we adopt a straightforward approach.

First of all we relax the mechanism for preference ordering. Remember (Section 2.3.2, page 43) that this is given by minimality according to a lexicographical order based on the numbers of (1) missing boundary relations, (2) paths, and (3) places. For a map $\dot{m}$, these numbers are $f_1(\dot{m}),\ f_2(\dot{m}),\ f_3(\dot{m})$. The different importance of these values, accounted by prioritized comparison, will be now handed over to the weights $c_1 > c_2 > c_3$ of the following linear combination to minimize

$$F(\dot{m}) = c_1 \cdot f_1(\dot{m}) \ + \ c_2 \cdot f_2(\dot{m}) \ + \ c_3 \cdot f_3(\dot{m}) \tag{4.15}$$

Observe that it is always possible to find c_1, c_2, c_3 such that the combination above induce the same ordering as the prioritized comparison, over a *given* set of maps. Of course, this may not hold for any set of maps once the weights are given. For example consider the situation

$$\begin{aligned} f_1(\dot{m}_a) &= f_1(\dot{m}_b) \\ f_2(\dot{m}_a) &> f_2(\dot{m}_b) \\ f_3(\dot{m}_a) &< f_3(\dot{m}_b) \end{aligned} \tag{4.16}$$

By the lexicographical order, $\dot{m}_b$ will be certainly preferred to $\dot{m}_a$, while the same happens with the linear combination only if c_2 and c_3 satisfy

$$\frac{c_2}{c_3} > \frac{f_3(\dot{m}_b) - f_3(\dot{m}_a)}{f_2(\dot{m}_a) - f_2(\dot{m}_b)}$$

which may not be the case.

We accept this "risk" for three reasons. First, it is a remote risk if one chooses c_1 enough greater than c_2, and c_2 enough greater than c_3. Indeed, at least for paths and places, it is reasonable to assume that they grow linearly together in non-pathological environments. Second,

although the linear combination does not provide such an appealing logical mechanism to compare maps as the circumscription does, we believe that the important empirical nature of such a mechanism (roughly speaking, prefer minimal maps in "good form") is almost entirely preserved. Third, the linear combination greatly simplifies the estimation of real numbers for probabilities. In addition, *preliminary* experiments have shown that setting $c_2 \approx 2.0 \cdot c_1$ yields better results than choosing c_2 an order of magnitude greater than c_1.

Intuitively, the probability of every map $\dot{m}_i$ in a set $\dot{m}_1, \ldots \dot{m}_n$ should be inversely proportional to $F(\dot{m}_i)$. Taking into account the necessary normalization we get:

$$p(\dot{m}_i) = \frac{[F(\dot{m}_i)]^{-1}}{\sum_{i=1}^{n}[F(\dot{m}_i)]^{-1}} \tag{4.17}$$

which will be our probability assessment for the prior $p(\dot{m})$.

4.6 Maximizing the Log-Posterior Probability

In Section 4.3 we proposed to compute the probability of a topological map through the Bayesian posterior (4.5):

$$p(\dot{m} \mid \overline{\tau^o},\ \overline{C^o}) = \alpha\ p(\overline{\tau^o} \mid \dot{m},\ \overline{C^o})\ p(\dot{m})$$

In Section 4.4, and Section 4.5 we showed how the metrical likelihood

$$p(\overline{\tau^o} \mid \dot{m},\ \overline{C^o}) = p^{BN(\dot{m},\ \overline{C^o})}(\overline{\tau} = \overline{\tau^o})$$

and the prior probability $p(\dot{m})$ can be respectively computed by linear regression after suitable approximations, and from the SSH preference policy.

We now introduce two simplifications for the computation of the posterior, which preserve the results of maximization/minimization and comparison.

1. We apply a monotonic logarithmic transformation:

$$\log[p(\dot{m} \mid \overline{\tau^o},\ \overline{C^o})] = \log[\alpha] + \log[p(\overline{\tau^o} \mid \dot{m},\ \overline{C^o})] + \log[p(\dot{m})]$$

   This is common practice, because of two computational benefits:

   - Multiplications are replaced by additions.
   - Log-values computations usually stand better the consequences of machine approximations.

2. We can ignore all the constant factors. In particular we do not need normalization. However, the trade-off between the metrical likelihood and the prior probability *will matter*.
   - The denominator $[\sum_i F(\dot{m}_i)]^{-1}$ in (4.17) for the prior probability can be neglected.
   - We can neglect the values $\frac{1}{\sqrt{(2\pi)^3 \det C_i^o}}$ in (4.9).

Therefore, the map that is most probable *a posteriori* is given by

$$\max_{\dot{m}} \left\{ \log[F(\dot{m})^{-1}] + \log[p^{BN(\dot{m},\, \overline{C^o})}(\overline{\tau} = \overline{\tau^o})] \right\}$$

equivalent to the nicer form

$$\max_{\dot{m}} \left\{ \log[p^{BN(\dot{m},\, \overline{C^o})}(\overline{\tau} = \overline{\tau^o})] - \log[F(\dot{m})] \right\} \tag{4.18}$$

4.7 Experimental Results

We have implemented the framework proposed in this chapter, and carried out a preliminary yet accurate experimentation. The experiments we present here regard the 4 by 4 grid we addressed in Section 3.4, where we investigated the effect of the planarity constraint. We observed that this environment is particularly hard for pure topological mapping, both because an exhaustive (breadth-first) search is computationally unfeasible (on a Pentium 1.5 Ghz, 512 Mb Ram), and because there are overminimal maps consistent with the exploration that deceptively attract a best-first search. We could solve it by applying a strong structural bias, which we called qualitative perpendicularity.

By contrast, this environment can be easily mapped by a best-first search that traverses the search-tree maximizing the probability as computed in Section 4.6. However, as we will show, things get worse as topological reasoning and constraints are reduced. This means that the advantages of the approach cannot be solely ascribed to the utilization of metrical information, and rather resides in the integration of topological reasoning and metrical analysis, which is the central theme of this thesis.

In Section 4.7.1, we collect a number of details concerning the particular conditions and assumptions in our experiments. In Section 4.7.2 we report the results.

4.7.1 Parameters

One issue that arose in our first experiments is that some of the topological biases used in the SSH can determine sudden variations in $F(\dot{m})$, when some loops are closed. This is not an ideal behavior for a best-first search, which works better when such variations are smooth.

When closing a loop, say a wrong one, the number of overall missing boundary (left/right) relations can considerably reduce, as a result of unifying two paths. The new path will inherit

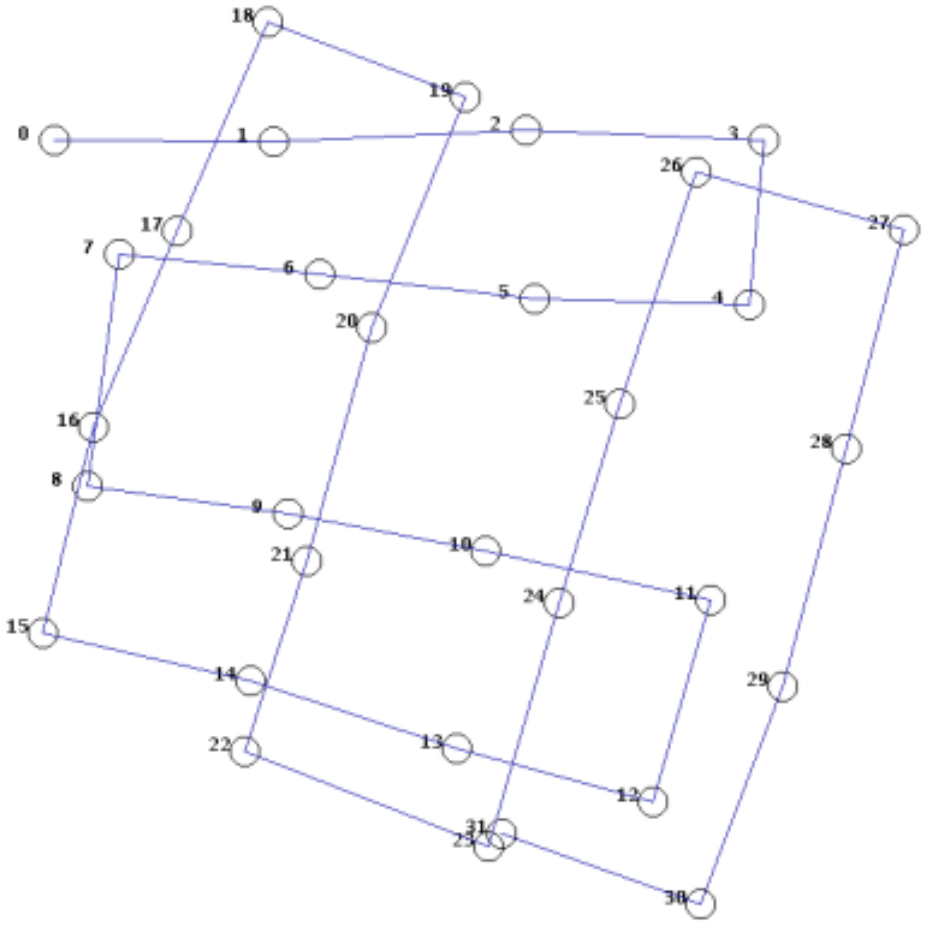

Figure 4.3: **Metrical Input.** The sequence of metrical relations that the algorithm takes in input. They are not to be considered pure odometry, but probabilistic estimates of the coordinate transformations relative to the moves from place to place. As explained in the text, this is also the result output by the algorithm, when only the metrical likelihood is used to inform the best-first search.

the boundary relations of both original paths, whereas the original paths were very likely to miss each other's boundary relations. If the metrical likelihood fails to compensate the advantage that this map gains (from the sudden reduction of $f_1(\dot{m})$), the correct hypothesis might not have a later chance of being resumed.

An analogous problem arises when a new place is inserted. Its star is instantiated in the map, along with all its new local-to-global paths, except the one along which the place is assumed to be accessed (which already exists in the map). Although such new paths are likely to be later unified with other paths already existing in the map, at the time of place insertion they add to $f_2(\dot{m})$. Even worse, such new paths lack boundary relations with all but one place; this increases $f_1(\dot{m})$ dramatically. It is easy to imagine a common situation in which the robot moves across many intersections without yet closing a physical loop; the correct map would score very low compared to wrong models that prematurely close loops.

These problems make it difficult to draw general insights and conclusions, since they cause some instability in the results obtained. For this reason, we have decided (1) to count as paths only those containing at least two places, and (2) not to involve $f_1(\dot{m})$ in the preliminary experiments below.

Note that we do not intend this as an argument against the use of boundary relations as qualitative biases. They can play an important role in driving the search for correct models. However, more investigation is needed to establish how they can be employed more effectively in a best-first search.

The particular linear combination that we used for F is

$$F(\dot{m}) = 2.0 \cdot f_2(\dot{m}) + 1.0 \cdot f_3(\dot{m})$$

The exploration assumed is a "snake" over a 4 by 4 topological grid. We discussed the motivations for this testbed in Section 3.4.

Regarding the metrical observations τ_i^o and C_i^o, we have obtained them adding random drift and rotational error to the transformed polar coordinates for every move in the simulated environment. Transforming back the (linearized) system it is possible to compute both the error on τ_i^o and the covariance matrix of the hypothesized error.

The whole sequence of all the τ_i^o in input for the experiments described below is given in Figure 4.3.

Apart from the simulations run here, we have already pointed out (Section 4.2) that metrical observations are supposed to come from a real robot system, as a result of a lower level process that partially corrects the odometrical error (scan-matching, scrolling occupancy-grid maps etc.). This makes the odometry assumed in 4.3 plausible.

As for other assumptions not specified here, they are as in Chapters 2 and 3.

4.7.2 Results

We discuss some different relevant cases. Figures of maps are automatically obtained as a side-effect of the algorithm, which needs to compute the poses π_i that minimize (4.10).

Only Topological Mapping

When the metrical likelihood is totally neglected, the best-first search is informed only by the SSH preference policy. We have seen in Section 3.4 that in this case an overminimal map is eventually selected. The only way round this problem was to enforce an almost *ad hoc* structural bias, which we called qualitative perpendicularity.

This suggests the use of metrical information to overcome the problem.

Only Metrical Likelihood

If $p(\dot{m})$ is suppressed from the posterior probability informing the search, in presence of even little metrical errors the metrical likelihood is maximal in the map where no loop is closed. This is actually a problem of overfitting, if the entire process is viewed as model learning: the map where no loop has been closed avoids any geometrical constraint, each place visited is considered new, and its pose is computed by adding last τ_i^o to the pose of the precedent place in a recursive way.

The layout produced equals the metrical input (Figure 4.3).

For the same reason, if backtracking is forced when the selected final map is not closed (see Section 3.3, this is a *topological* bias in its own respect, provided by the hybrid SSH framework) the search does not end in reasonable time, since it will sweep the branches where maps have fewer loops, which are not closed, before focusing on the region of search-space where more compact maps are built.

This suggests that at least a bias toward minimality is needed. This can be provided by the prior probability, derived by the SSH preference policy.

Integrated Framework

The correct topological model is selected as final map when the planarity constraint is enforced, only closed final maps are accepted, and the best-first search is driven by the maximization of the posterior probability, taking into account both the prior derived from the SSH preference policy and the metrical likelihood. Its layout is shown in Figure 4.4.

If planarity is not enforced, a wrong model is selected, analogously to what happens with the non-planar model in Figure 1.2(c), page 16. Indeed non-planarity arises due to two places connected in a “twisted way” (Figure 4.5 and 4.6).

This shows how uncertain metrical information, when utilized by a best-first search may end up with wrong models that can be easily ruled out by topological reasoning.

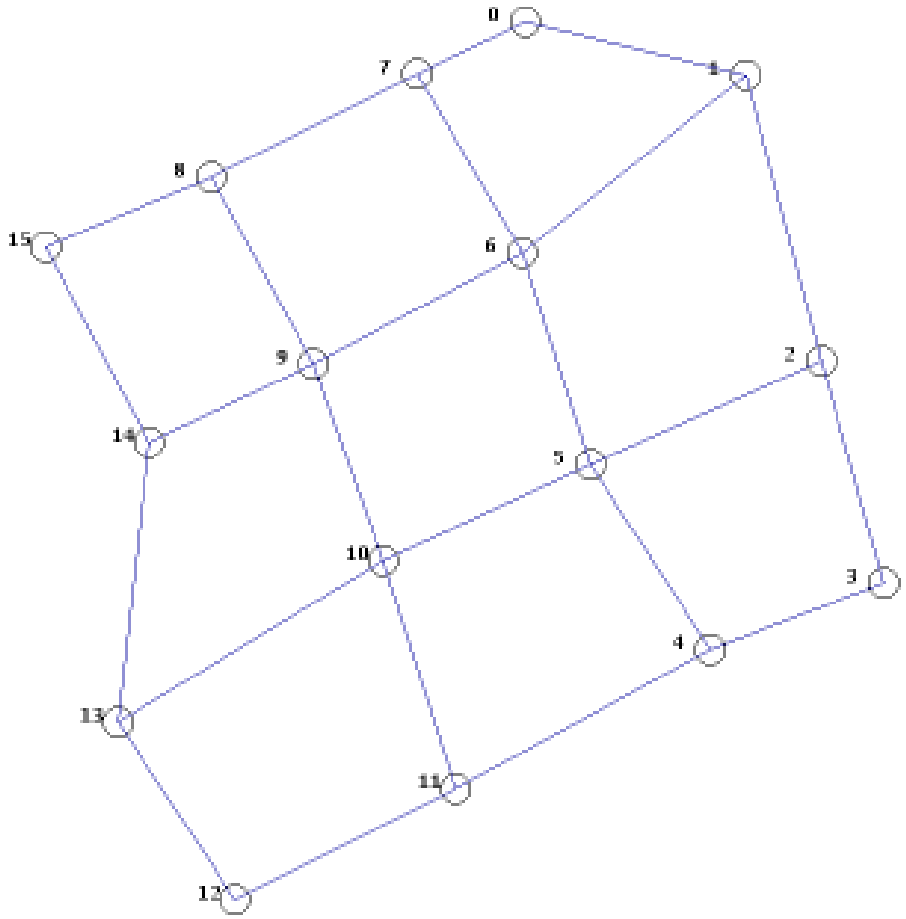

Figure 4.4: **Correct Result.** This is the topologically correct result obtained when the full posterior probability is used to inform the search, and in addition planarity is enforced.

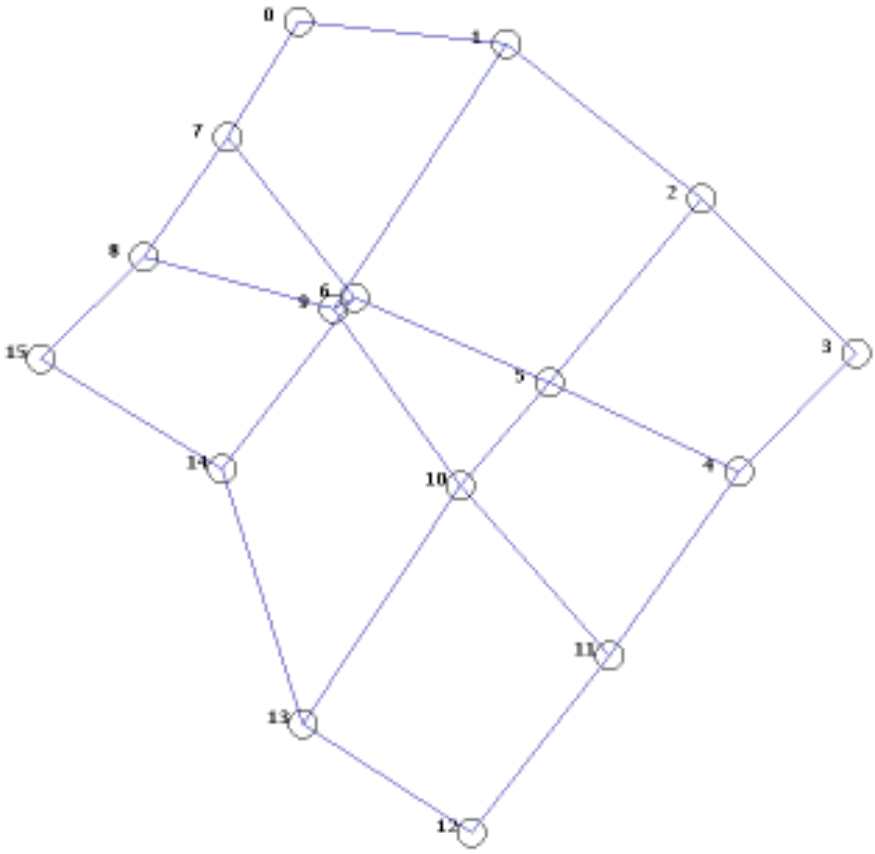

Figure 4.5: **Non-Planar Result I.** This is the (wrong) result that is obtained when planarity is not enforced, despite the full posterior probability is utilized. Indeed, place 6 and place 9, are connected by "twisted" edges, in the same way as it happens in Figure 1.2(c), page 16. The detail is reported in Figure 4.6. Compared with Figure 4.4, this case shows how pure topological considerations can be very valuable when metrical analysis fails to recover the correct correspondence between places over large-scale space.

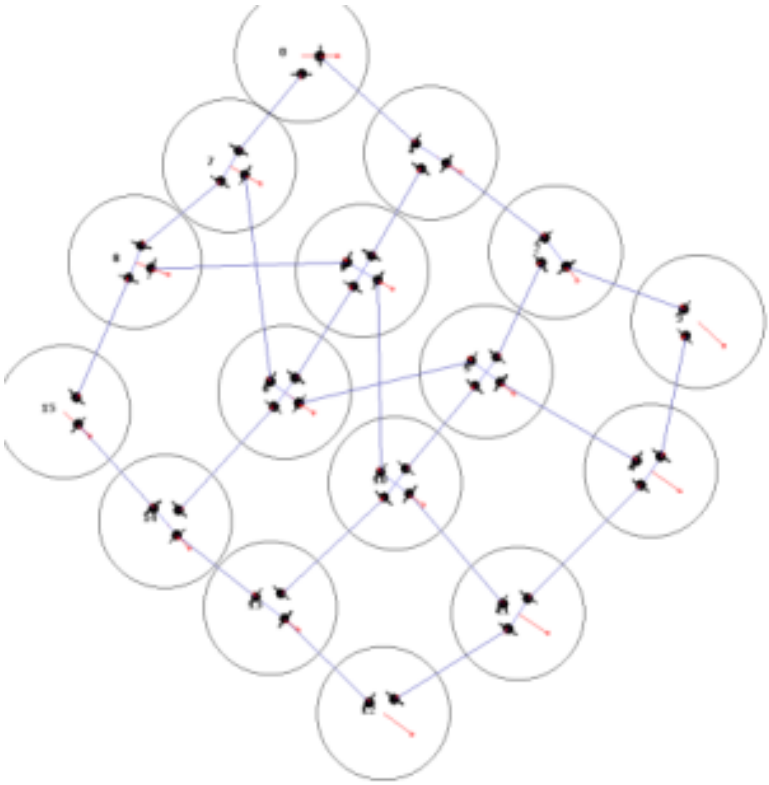

Figure 4.6: **Non-Planar Result II.** Stars and gateways of the map in Figure 4.5. Here it is possible to observe that the connections are arranged in a non-planar layout. In this picture places have been manually shifted to help visual inspection. The actual metrical layout output by the algorithm is given in Figure 4.5.

The search expands 2934 nodes in the first case (2098 when planarity is not applied), in a few seconds (C++ implementation running on a Pentium 2.8 Ghz, 512 Mb RAM).

4.8 Related Work

We review here related work that addresses the correspondence problem in the context of robot map-building.

4.8.1 Correspondence in Metrical Methods

In metrical mapping, the term *data association* has been borrowed from the literature concerned with the problem of tracking targets in raw data. It stands for the problem, also called the *correspondence problem*, of deciding which object/feature/target, among multiple candidates, is associated with the data obtained from the sensors. The method usually employed is to select one of the candidates, according to some maximum likelihood evaluation (among most recent work Montemerlo et al. (2002, 2003)), or the Nearest Neighbor Filter (NNF, Bar-Shalom and Fortmann (1988)). Limits of this kind of approaches have been pointed out by Neira and Tardós (2001); the most likely hypothesis at selection time might prove wrong later, affecting the whole mapping process. More sophisticated approaches (Leonard et al., 2001; Nieto et al., 2003) combine the NNF with the Multiple Hypothesis Tracking (MHT, Reid (1980)) method, with great advantages. Here the decision is delayed; before selecting a hypothesis some look-ahead is performed. In the approach proposed by Cox and Leonard (1994), very close to our framework in some respects, a probability-guided search is performed over a tree of hypotheses. The differences with our work are: (1) tree-pruning is also guided by probability considerations, whereas we discard branches only if they prove inconsistent, (2) certainty on robot poses is assumed, and then the SLAM (CML) problem is not entirely addressed.

In (Montemerlo et al., 2002, 2003; Eliazar and Parr, 2003) and similar work the probability distribution over the map and the pose of the robot is updated through particle-filtering (Doucet et al., 2001); roughly speaking, a population of samples (particles) approximates such a distribution, and evolves as more data are made available. If particles are looked at as hypotheses, this approach is clearly related to ours, even though creation of hypotheses is not explicitly driven by data association ambiguity or view comparison. The above mentioned work (Nieto et al., 2003) goes in this direction creating also particles for the hypotheses originated by the NNF and MHT combined together.

4.8.2 Correspondence in Hybrid Methods

In the framework of the SSH, Remolina and Kuipers (2004) consider local metrical annotations, whose uncertainty is represented by intervals. Based on these annotations, metrical constraints

are axiomatized in the logical theory we sketched in Section 2.1.3. Therefore maps are discarded irreversibly by metrical considerations, whereas we rank them by a degree of confidence.

Atlas (Bosse et al., 2003) is an hybrid approach very close to ours. It is strongly based on the Smith and Cheeseman's framework (Smith and Cheeseman, 1986). It builds a graph of local metrical maps in local frames of reference, and exploits the transformations from frame to frame, to evaluate global positional relationships according to the Smith and Cheseeman's method. They call this a Dijkstra projection, because they use a path that is minimal w.r.t. the covariance of the final compound transformation. They use the Dikstra projection between two local maps and some criteria of map-matching to evaluate hypotheses about whether closing a loop. More hypotheses are kept simultaneously and their performance in explaining few next steps of sensor measurements are evaluated, before the selection. Beyond some important technical details, the main difference here is that the policy employed in Atlas to handle loop-closure events is the exact opposite of ours. In Atlas a very conservative strategy is adopted, which attempts to avoid wrong loop-closure at the expense of accepting the risk of missing some genuine loop-closure event. We adopt a very conservative strategy that potentially creates the whole set of loop-closure topological candidates.

Tomatis et al. (2002, 2003) use a probabilistic framework to localize and detect loops. When the estimate of the robot's location has two probable hypotheses, the framework assumes it is recreating a previously known portion of the topological map. The robot will then physically backtrack until the location estimate converges to a single hypothesis, producing a simpler topological map.

Hähnel et al. (2003) use metrical likelihood to prioritize a best-first search on the tree of correspondence hypotheses, what they call "Lazy Data-Association". This approach is very close to ours, but for the fact that topological reasoning and biases are not utilized. The same argument applies to (Shatkay and Kaelbling, 1997), in addition to the difference in the probabilistic framework they adopt. They use Partially Observable Markov Decision Processes (POMDPs) as a framework to integrate a topological map and metrical information, to the purpose of determining the correct topological model.

Chapter 5

Conclusions

Topological ambiguity is a type of spatial uncertainty that faces an agent, human or robot, whose tasks or goals have to be accomplished over unknown large-scale space. While trying to build a useful map of the environment, whatever the nature of such a map, uncertainty might arise about whether the region currently being explored has already been visited, a problem sometime known as the data-association (or correspondence) problem. This is clearly captured in a topological representation, where the graph of spatial relationships is laid down. Topological ambiguity makes the loops of such a graph uncertain.

The amount of topological ambiguity depends on the environment at hand, in particular on how many regions appear identical, and on where these regions are w.r.t. each other, especially in presence of symmetry. Note, however, that these factors also depend on what the agent "sees" of such an environment.

Therefore, one important research direction towards making robots more robust against topological uncertainty is to improve their perceptual mechanisms. This also includes devising new ways a robot can extract the distinctive features of the places of a given environment, which allow it to reliably and conveniently recognize and distinguish such places. These features need to be extracted from the massive amount of raw data made available by modern sensors. It seems that relatively little effort has so far been made in this direction, and researchers are now increasingly focusing on it (see Section 2.5.2).

Another way to address topological uncertainty is to build a very precise and reliable geometrical description of the environment, in a single global frame of reference. When this is possible, topological uncertainty is implicitly solved, since the identity of a region can be uniquely determined by its position. This direction has received much attention in the SLAM community, with great success especially over the last decade. Impressive advancements have been made possible by new families of algorithms, primarily based on probabilistic techniques, and by the computing power now available.

However, there are several reasons why it is important to make explicit use of topological

representations (graphs). First, they are necessary regardless of the need of dealing with topological ambiguity. For example, topological representations better serve cheap and hierarchical navigation planning, and human-robot interaction, because they allow for symbolic computation. Second, the amount of data that have to be retrieved, and whose global consistency has to be enforced at every step may reduce the appeal of pure metrical mapping for the purpose of dealing with topological ambiguity. Third, there are situations where even the most successful techniques for global metrical mapping would fail, for example in presence of unexpectedly large errors, well beyond the typical range assumed by the probabilistic model. Last, but not least, explicit topological reasoning can be considerably cost-effective for reducing topological uncertainty, even in presence of a selective use of metrical information. The thesis contributes to this last point by studying the planarity constraint (Chapter 3), and the integration of the SSH qualitative reasoning with metrical estimation methods in a probabilistic framework (Chapter 4).

We have addressed the problem of topological ambiguity in its natural context, that is topological mapping. Our work builds on and contributes to the topological mapping framework of the hybrid SSH (Chapter 2), which has been demonstrated with robot explorations of office-like environments.

Abstracting a real environment to a graph, however, is not usually so straightforward. The SSH provides general (and cognitively plausible) principles for doing this. When this comes to practice, and to the peculiar sensing capabilities of a robot, what should be considered a place or a path, for building a consistent and reliable topological map, depends on the task at hand. Hierarchical representations seem to be crucial here, but research with real robots in this direction has so far been quite poor.

For the same reasons, it might be necessary to move beyond the concept of topological maps as (built on) graphs, perhaps introducing more sophisticated combinatorial structures. Such structures should preserve the same advantages, for example they should allow for effective symbolic spatial reasoning as exploited by the planarity test. It can be envisaged that these research efforts would be strictly related to the aforementioned investigation of methods for detecting, classifying, and recognizing distinctive places.

In Chapter 4 we combined topological reasoning, and the SSH ontology, with a family of metrical mapping methods. The contribution here was twofold.

From a representational viewpoint, we formalized metrical and topological uncertainty in a probabilistic framework. We emphasized that topological uncertainty affects the conditional dependence assumptions in the probabilistic model of metrical uncertainty. Thus, while metrical uncertainty could be handled by probabilistic inference in a Bayes network, topological ambiguity makes such a Bayes network uncertain itself. Topological mapping can then be viewed as a Bayes-network learning problem, which in AI and statistics is well known to be very hard.

The structure of the Bayes network to learn, however, is not free to vary arbitrarily. It must belong to a particular set of candidates, which reflect the loops assumed in the map (Figure 4.2, page 69). From the computational viewpoint, this allowed us to demonstrate that after suitable approximations the problem can be efficiently addressed. The SSH qualitative preference policy

is used to derive the probability a priori of the topological model, while a family of state-of-the-art solutions to the SLAM problem, recently introduced in literature, can be used to compute its metrical likelihood. Experimental results show that this integrated approach can work better than either approach in isolation.

The planarity constraint filters out most incorrect topological loop-closure hypotheses inexpensively, and independently of the actual metrical scale and geometrical appearance of the loops at issue. In Chapter 3 we provided a formal analysis of the reduction of topological uncertainty that follows from enforcing the planarity constraint. We also experimentally showed that this constraint can make a dramatic difference in particularly adverse (because of perceptual aliasing, multiple nested loops, symmetry) environments, both when topological mapping is used in isolation, and in the extended framework that takes advantage of metrical information (Chapter 4).

The contributions of this thesis add to the arguments in favor of topological and qualitative reasoning for robot applications faced with spatial uncertainty, and to the techniques available for overcoming this uncertainty.

Bibliography

Y. Bar-Shalom and T. E. Fortmann. *Tracking and Data Association*. Academic Press, 1988.

P. Beeson, M. MacMahon, J. Modayil, J. Provost, F. Savelli, and B. Kuipers. Exploiting local perceptual models for topological map-building. In *Proc. of IJCAI-2003 Workshop on Reasoning with Uncertainty in Robotics (RUR-03)*, pages 15–22, Acapulco, Mexico, August 2003.

M. Bosse, P. Newman, J. Leonard, M. Soika, W. Feiten, and S. Teller. An Atlas framework for scalable mapping. In *Proc. IEEE Int. Conf. on Robotics and Automation*, 2003.

P. Buschka and A. Saffiotti. Some notes on the use of hybrid maps for mobile robots. In *Proc. of the 8th Int. Conf. on Intelligent Autonomous Systems (IAS)*, pages 547–556, Amsterdam, NL, 2004.

R. Chatila and J.-P. Laumond. Position referencing and consistent world modeling for mobile robots. In *Proc. IEEE Int. Conf. on Robotics and Automation*, pages 138–170, 1985.

H. Choset and K. Nagatani. Topological simultaneous localization and mapping (SLAM): toward exact localization without explicit localization. *IEEE Trans. on Robotics and Automation*, 17(2):125–137, April 2001.

E. Chown, S. Kaplan, and D. Kortenkamp. Prototypes, location, and associative networks (PLAN): Towards a unified theory of cognitive mapping. *Cognitive Science*, 19(1):1–51, 1995.

I. J. Cox and J. J. Leonard. Modeling a dynamic environment using a bayesian multiple hypothesis approach. *Artificial Intelligence*, 66(2):311–344, 1994.

T. Dean, D. Angluin, K. Basye, S. Engelson, L. Kaelbling, E. Kokkevis, and O. Maron. Inferring finite automata with stochastic output functions and an application to map learning. *Machine Learning*, 18(1):81–108, 1995.

X. Deng, E. E. Milios, and A. Mirzaian. Robot map verification of a graph world. *Journal of Combinatorial Optimization*, 5(4):383–395, 2001.

G. Dissanayake, P. Newman, H. F. Durrant-Whyte, S. Clark, and M. Csorba. A solution to the simultaneous localization and map building (SLAM) problem. *IEEE Trans. on Robotics and Automation*, 17(3):229–241, 2001.

A. Doucet, J. F. G. de Freitas, and N. J. Gordon. *Sequential Monte Carlo Methods In Practice*. Springer Verlag, New York, 2001.

D. Driankov and A. Saffiotti, editors. *Fuzzy Logic Techniques for Autonomous Vehicle Navigation*. Studies in Fuzziness and Soft Computing. Springer-Phisica Verlag, 2001.

T. Duckett, S. Marsland, and J. Shapiro. Fast, on-line learning of globally consistent maps. *Autonomous Robots*, 12(3):287–300, 2002.

T. Duckett and U. Nehmzow. Exploration of unknown environments using a compass, topological map and neural network. In *Proc. of the 1999 IEEE Int. Symp. on Computational Intelligence in Robotics and Automation (CIRA'99)*, Monterey, CA, 1999.

T. Duckett and U. Nehmzow. Mobile robot self-localisation using occupancy histograms and a mixture of gaussian location hypotheses. *Robotics and Autonomous Systems*, 34(2–3):119–130, 2001.

T. Duckett and A. Saffiotti. Building globally consistent gridmaps from topologies. In *Proc. of the 6th Int. IFAC Symp. on Robot Control (SYRORO-00)*, Vienna, Austria, September 2000. Elsevier.

G. Dudek, P. Freedman, and S. Hadjres. Using local information in a non-local way for mapping graph-like worlds. In *Proc. of the 13th Int. Joint Conf. on Artificial Intelligence (IJCAI-93)*, pages 1639–1645, 1993.

G. Dudek, M. Jenkin, E. Milios, and D. Wilkes. Robotic exploration as graph construction. *IEEE Trans. on Robotics and Automation*, 7(6):859–865, 1991.

A. Elfes. *Occupancy Grids: A Probabilistic Framework for Robot Perception and Navigation*. PhD thesis, Carnegie Mellon Univ. 1989.

A. Eliazar and R. Parr. DP-SLAM: Fast, robust simultaneous localization and mapping without predetermined landmarks. In *Proc. of the 18th Int. Joint Conf. on Artificial Intelligence (IJCAI-03)*, pages 1135–1142. Morgan Kaufmann, 2003.

E. Fabrizi and A. Saffiotti. Extracting topology-based maps from gridmaps. In *Proc. IEEE Int. Conf. on Robotics and Automation*, pages 2972–2978, San Francisco, CA, 2000.

E. Fabrizi and A. Saffiotti. Augmenting topology-based maps with geometric information. *Robotics and Autonomous Systems*, 40(2):91–97, 2002.

M. O. Franz, B. Schölkopf, H. A. Mallot, and H. H. Bülthoff. Learning view graphs for robot navigation. *Autonomous Robots*, 5(1):111–125, 1998.

U. Frese, P. Larsson, and T. Duckett. A multilevel relaxation algorithm for simultaneous localisation and mapping. *IEEE Transactions on Robotics*, 2004. to appear.

N. Friedman. The Bayesian structural EM algorithm. In *Proc. of the Fourteenth Conf. on Uncertainty in Artificial Intelligence (UAI-98)*, 1998.

S. Gillner and H. A. Mallot. Navigation and acquisition of spatial knowledge in a virtual maze. *Journ. of Cognitive Neuroscience*, 10(4):445–463, 1998.

J. L. Gross and T. W. Tucker. *Topological Graph Theory*. John Wiley and Sons, New York, 1987.

J. S. Gutmann, W. Burgard, D. Fox, and K. Konolige. An experimental comparison of localization methods. In *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, Victoria, Canada, 1998.

J. S. Gutmann and D. Fox. An experimental comparison of localization methods continued. In *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, Lusanne, Switzerland, 2002.

J. S. Gutmann and K. Konolige. Incremental mapping of large cyclic environments. In *Proc. of the Conf. on Intelligent Robots and Applications (CIRA-99)*, Monterey, CA, 1999.

J. S. Gutmann and C. Schlegel. AMOS: Comparison of scan matching approaches for self-localization in indoor environments. In *Proc. of the 1st Euromicro Workshop on Advanced Mobile Robots*, 1996.

D. Hähnel, W. Burgard, B. Wegbreit, and S. Thrun. Towards lazy data association in SLAM. In *Proceedings of the 11th International Symposium of Robotics Research (ISRR'03)*, Siena, Italy, 2003. Springer.

R. A. Hart and G. T. Moore. The development of spatial cognition: A review. In R. M. Downs and D. Stea, editors, *Image and Environment*, pages 246–288. Aldine Publishing Company, Chicago, 1973.

W. H. Huang and K. R. Beevers. Topological map merging. In *Distributed Autonomous Robotic Systems (DARS)*, pages 91–100, 2004.

G. F. Italiano, J. A. L. Poutre, and M. Rauch. Fully dynamic planarity testing in planar embedded graphs (extended abstract). In *European Symposium on Algorithms*, pages 212–223, 1993.

M. I. Jordan, editor. *Learning in Graphical Models*. MIT Press, Cambridge, Massachusetts, 1999.

K. Konolige. Large-scale map-making. In *Proc. of the National Conference on AI (AAAI-04)*, San Jose, CA, 2004.

B. Kuipers and P. Beeson. Bootstrap learning for place recognition. In *Proc. of the 18th National Conf. on Artificial Intelligence (AAAI-2002)*, Edmonton, Canada, 2002.

B. Kuipers, J. Modayil, P. Beeson, M. MacMahon, and F. Savelli. Local metrical and global topological maps in the hybrid spatial semantic hierarchy. In *Proc. IEEE Int. Conf. on Robotics and Automation*, New Orleans, Louisiana, 2004.

B. Kuipers, D. G. Tecuci, and B. Stankiewicz. The skeleton in the cognitive map: a computational and empirical exploration. *Environment and Behavior*, 35(1):80–106, 2003.

B. J. Kuipers. *Representing Knowledge of Large-Scale Space*. PhD thesis, Mathematics Department, Massachusetts Institute of Technology, Cambridge, Massachusetts, June 1977. Also MIT-AI TR-418, 1977.

B. J. Kuipers. Modeling spatial knowledge. *Cognitive Science*, 2:129–153, 1978. Reprinted in *Advances in Spatial Reasoning, Volume 2*, Su-Shing Chen (Ed.), Norwood NJ: Ablex Publishing, 1990, pages 171–198.

B. J. Kuipers. The spatial semantic hierarchy. *Artificial Intelligence*, 119(1-2):191–233, 2000.

B. J. Kuipers and Y.-T. Byun. A robot exploration and mapping strategy based on a semantic hierarchy of spatial representations. *Journ. of Robotics and Autonomous Systems*, 8:47–63, 1991.

B. J. Kuipers and T. Levitt. Navigation and mapping in large scale space. *AI Magazine*, 9(2): 25–43, 1988. Reprinted in *Advances in Spatial Reasoning, Volume 2*, Su-shing Chen (Ed.), Norwood NJ: Ablex Publishing, 1990, pages 207–251.

A. Lankenau, T. Röfer, and B. Krieg-Brückner. Self-localization in large-scale environments for the bremen autonomous wheelchair. In *Spatial Cognition III*, volume 2685 of *Lecture Notes in Artificial Intelligence*, pages 34–61, Berlin, Germany, 2002. Springer-Verlag.

W. Y. Lee. *Spatial Semantic Hierarchy for a Physical Mobile Robot*. PhD thesis, Univ. of Texas at Austin, Computer Science Department, December 1996.

J. J. Leonard, P. M. Newman, R. J. Rikoski, J. Neira, and J. D. Tardós. Towards robust data association and feature modeling for concurrent mapping and localization. In *Proceedings of the Tenth International Symposium on Robotics Research*, Lorne, Victoria, Australia, 2001.

V. Lifschitz. Circumscription. In D. M. Gabbay, C. J. Hogger, and J. A. Robinson, editors, *Handbook of Logic in AI and Logic Programming*, volume 3, pages 298–352. Oxford Univ. Press, 1994.

V. Lifschitz. Nested abnormality theories. *Artificial Intelligence*, 74(2):351–365, 1995.

B. Lisien, D. Morales, D. Silver, G. Kantor, I. Rekleitis, and H. Choset. Hierarchical simultaneous localization and mapping. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 448–453, Las Vegas, NV, 2003. IEEE/RSJ.

F. Lu and E. Milios. Globally consistent range scan alignment for environment mapping. *Autonomous Robots*, 4(4):333–349, 1997.

K. Lynch. *The Image of the City*. MIT Press, Cambridge, MA, 1960.

M. J. Mataric. Integration of representation into goal-driven behavior-based robots. *IEEE Trans. on Robotics and Automation*, 8(3):304–312, 1992.

J. McCarthy. Circumscription—a form of non-monotonic reasoning. *Artificial Intelligence*, 13(1,2):27–39, 171–172, 1980.

J. Modayil, P. Beeson, and B. Kuipers. Using the topological skeleton for scalable global metrical map-building. In *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2004.

M. Montemerlo, S. Thrun, D. Koller, and B. Wegbreit. FastSLAM: A factored solution to the simultaneous localization and mapping problem. In *Proc. of the AAAI National Conf. on Artificial Intelligence*, Edmonton, Canada, 2002. AAAI.

M. Montemerlo, S. Thrun, D. Koller, and B. Wegbreit. FastSLAM 2.0 An improved particle filtering algorithm for simultaneous localization and mapping that provably converges. In *Proc. of the 18th Int. Joint Conf. on Artificial Intelligence (IJCAI-03)*, Acapulco, Mexico, 2003.

H. P. Moravec. Sensor fusion in certainty grids for mobile robots. *AI Magazine*, Summer 1988.

J. Neira and J. D. Tardós. Data association in stochastic mapping using the joint compatibility test. *IEEE Trans. on Robotics and Automation*, 17(6):890–897, 2001.

J. Nieto, J. Guivant, E. Nebot, and S. Thrun. Real time data association for FastSLAM. In *Proc. IEEE Int. Conf. on Robotics and Automation*, Taiwan, 2003.

M. A. Paskin. The junction filters for simultaneous localization and mapping. In *Proc. of the 18th Int. Joint Conf. on Artificial Intelligence (IJCAI-03)*, Acapulco, Mexico, 2003.

J. Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann, 1988.

J. Piaget and B. Inhelder. *The Child's Conception of Space*. Norton, New York, 1967. First published in French, 1948.

D. M. Pierce and B. J. Kuipers. Map learning with uninterpreted sensors and effectors. *Artificial Intelligence*, 92(1-2):169–227, 1997.

W. Press, S. Tukolsky, W. Vetterling, and B. Flannery. *Numerical Recipes, Second Edition*. Cambridge University Press, Cambridge, 1992.

D. B. Reid. An algorithm for tracking multiple targets. *IEEE Trans. on Aerospace and Electronic Systems*, 1980.

I. M. Rekleitis, V. Dujmović, and G. Dudek. Efficient topological exploration. In *Proc. IEEE Int. Conf. on Robotics and Automation*, pages 676–681, Detroit, USA, May 1999.

E. Remolina. *Formalizing the Spatial Semantic Hierarchy*. PhD thesis, Univ. of Texas at Austin, Department of Computer Sciences, 2001.

E. Remolina and B. Kuipers. A logical account of causal and topological maps. In *Proc. of the 17th Int. Joint Conf. on Artificial Intelligence (IJCAI-01)*, pages 5–11, Menlo Park, CA, 2001.

E. Remolina and B. Kuipers. Towards a general theory of topological maps. *Artificial Intelligence*, 152(1):47–104, 2004.

R. L. Rivest and R. E. Schapire. Inference of finite automata using homing sequences. In *Proc. 21st ACM Symposium on Theory of Computing*, pages 411–420. ACM, 1989.

S. Russell and P. Norvig. *Artificial Intelligence: A Modern Approach, Second Edition*. Prentice Hall, 2002.

F. Savelli and B. Kuipers. Loop-closing and planarity in topological map-building. In *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2004.

B. Schölkopf and H. A. Mallot. View-based cognitive mapping and path planning. *Adaptive Behavior*, 3:311–348, 1995.

H. Shatkay and L. P. Kaelbling. Learning topological maps with weak local odometric information. In *Proc. of the 15th Int. Joint Conf. on Artificial Intelligence (IJCAI-97)*, pages 920–927, 1997.

A. W. Siegel and S. H. White. The development of spatial representations of large-scale environments. In H. W. Reese, editor, *Advances in Child Development and Behavior*, volume 10. Academic Press, New York, 1975.

R. Smith and P. Cheeseman. On the representation and estimation of spatial uncertainty. *Int. Journ. of Robotics Research*, 5:56–68, 1986.

R. Smith, M. Self, and P. Cheeseman. Estimating uncertain spatial relationships in robotics. In I. J. Cox and G. T. Wilfong, editors, *Autonomous Robot Vehicles*, pages 167–193. Springer, 1990.

S. Thrun. Learning metric-topological maps for indoor mobile robot navigation. *Artificial Intelligence*, 99(1):21–71, 1998.

S. Thrun. Probabilistic algorithms in robotics. *AI Magazine*, 21(4):93–109, 2000.

S. Thrun. Robotic mapping: A survey. In G. Lakemeyer and B. Nebel, editors, *Exploring Artificial Intelligence in the New Millenium*. Morgan Kaufmann, 2002.

S. Thrun, D. Fox, and W. Burgard. A probabilistic approach to concurrent mapping and localization for mobile robots. *Machine Learning*, 31(1–3):29–53, 1998a.

S. Thrun, S. Gutmann, D. Fox, W. Burgard, and B. J. Kuipers. Integrating topological and metric maps for mobile robot navigation: A statistical approach. In *Proc. of the 15th National Conf. on Artificial Intelligence (AAAI-98)*, pages 989–995. AAAI/MIT Press, 1998b.

N. Tomatis, I. Nourbakhsh, and R. Siegwart. Hybrid simultaneous localization and map building: Closing the loop with multi-hypotheses tracking. In *Proc. IEEE Int. Conf. on Robotics and Automation*, pages 2749–2754, Washington, DC, May 2002.

N. Tomatis, I. Nourbakhsh, and R. Siegwart. Hybrid simultaneous localization and map building: a natural integration of topological and metric. *Robotics and Autonomous Systems*, 44 (1):3–14, 2003.

B. Yamauchi and P. Langley. Place recognition in dynamic environments. *Journal of Robotic Systems*, 14(2):107–120, February 1997.

W. K. Yeap. Towards a computational theory of cognitive maps. *Artificial Intelligence*, 34: 297–360, 1988.

W. K. Yeap and M. E. Jefferies. Computing a representation of the local environment. *Artificial Intelligence*, 107(2):265–301, 1999.

Appendix A

Bayes Networks

A Bayes network is a partially graphical representation of a joint probability distribution over a set of random variables. The set of conditional dependencies between the random variables is represented by a graph, while quantitative information about the respective conditional (probability, or probability density) distributions are specified apart.

The advantages of Bayes networks are twofold. On one hand, they make possible to examine the conditional dependency relationships qualitatively, by visual inspection. On the other hand, algorithm design for inference and learning benefits from the explicit combinatorial description of such relationships.

In a Bayes network, random variables are represented as the nodes of an acyclic directed graph. Given two variables x, y, an arrow from x to y represents the fact that y conditionally depends on x [1].

A variable x is a *parent* of a variable y if there is an arrow from x to y. A variable w is a *descendant* of a variable v if there is a directed path of arrows from v to w. An important representational property of Bayes networks is the following:

A variable is conditionally independent of its non-descendants given its parents.

This property is at the heart of the many computational techniques available for efficient probabilistic inference in Bayes networks. These are beyond the scope of this Appendix, since they are not used in this thesis. We point out, however, a simple consequence that lends a Bayes-network perspective to the factorization (4.8) used in Section 4.4 (page 74).

[1] This implies that a single direction must be chosen for denoting correlation, which may not exhibit such an asymmetric nature in the domain at hand. The reason for this is that Bayes networks are primarily meant to capture conditional dependencies that account for *causal* relationships between the elements of the domain at hand. Other types of *graphical models* exist that make use of undirected graphs, and also allow cycles in them.

The marginal distribution $p^{BN(\dot{m})}(\overline{\tau} \mid \overline{\pi})$ can be factorized by the Bayes rule as

$$p^{BN(\dot{m})}(\overline{\tau} \mid \overline{\pi}) = p^{BN(\dot{m})}(\tau_1 \mid \tau_2, \ldots \tau_p, \overline{\pi}) \cdot p^{BN(\dot{m})}(\tau_2, \ldots \tau_p \mid \overline{\pi}) \tag{A.1}$$

and then, by applying the same step recursively,

$$p^{BN(\dot{m})}(\overline{\tau} \mid \overline{\pi}) = \prod_{i=1}^{p} p^{BN(\dot{m})}(\tau_i \mid \tau_{i+1}, \ldots \tau_p, \overline{\pi}) \tag{A.2}$$

Because of the type of conditional independence assumptions that hold in our domain, illustrated in Section 4.2, every variable of the vector τ_i has no descendant in $BN(\dot{m})$, and its parents are given by the union of $\pi_{s(i)}$ and $\pi_{e(i)}$. Therefore, by the property above,

$$p^{BN(\dot{m})}(\tau_i \mid \tau_{i+1}, \ldots \tau_p, \overline{\pi}) = p^{BN(\dot{m})}(\tau_i \mid \pi_{s(i)}, \pi_{e(i)}) \tag{A.3}$$

which leads to the factorization used:

$$p^{BN(\dot{m})}(\overline{\tau} \mid \overline{\pi}) = \prod_{i=1}^{p} p^{BN(\dot{m})}(\tau_i \mid \pi_{s(i)},\ \pi_{e(i)}) \tag{A.4}$$

Note that, although this could serve as a formal justification of the factorization used in Chapter 4, a traditional account of the conditional independence assumptions is sufficient to this purpose.

Appendix B

Global Poses Computation

We report here the mathematical details of the linearization (4.2)

$$\tau = (\ominus \pi_s) \oplus \pi_e \approx \lambda(\pi_s, \pi_e) \tag{B.1}$$

with (4.1)

$$\begin{pmatrix} x_\tau \\ y_\tau \\ \theta_\tau \end{pmatrix} = \begin{pmatrix} (x_{\pi_e} - x_{\pi_s}) \cos\theta_{\pi_s} \ + \ (y_{\pi_e} - y_{\pi_s}) \sin\theta_{\pi_s} \\ -(x_{\pi_e} - x_{\pi_s}) \sin\theta_{\pi_s} \ + \ (y_{\pi_e} - y_{\pi_s}) \cos\theta_{\pi_s} \\ \theta_{\pi_e} \ - \ \theta_{\pi_s} \end{pmatrix} \tag{B.2}$$

and of the minimization (4.10) $\arg\min_{\overline{\pi}} \{f(\overline{\pi})\}$, with

$$f(\overline{\pi}) = \sum_{i=1}^{p} (\lambda(\pi_{s(i)}, \pi_{e(i)}) - \tau_i^o)^T \ C_i^{o-1} \ (\lambda(\pi_{s(i)}, \pi_{e(i)}) - \tau_i^o) \tag{B.3}$$

We follow (Press et al., 1992; Lu and Milios, 1997; Frese et al., 2004).

B.1 Linearization

We linearize the function $\tau = (\ominus \pi_s) \oplus \pi_e \approx \lambda(\pi_s, \pi_e)$ by a first-order Taylor expansion. Let $\pi_s = l_s, \ \pi_e = l_e$ be the linearization point. We assume that l_s is given by the last estimation, or by initialization from the sequence of $\overline{\tau}$. We choose l_e such that

$$l_e = l_s \oplus \tau \tag{B.4}$$

that is

$$\lambda(l_s, l_e) = (\ominus l_s) \oplus l_e = \tau \tag{B.5}$$

Thus, the first-order Taylor expansion is

$$\lambda(\pi_s, \pi_e) = \tau \ + \ S \ (\pi_s \ - \ l_s) \ + \ E \ (\pi_e \ - \ l_e) \tag{B.6}$$

with

$$S = \begin{pmatrix} -\cos\theta_{l_s} & -\sin\theta_{l_s} & \tau_y \\ \sin\theta_{l_s} & -\cos\theta_{l_s} & -\tau_x \\ 0 & 0 & -1 \end{pmatrix} \tag{B.7}$$

$$E = \begin{pmatrix} \cos\theta_{l_s} & \sin\theta_{l_s} & 0 \\ -\sin\theta_{l_s} & \cos\theta_{l_s} & 0 \\ 0 & 0 & 1 \end{pmatrix} \tag{B.8}$$

B.2 Minimization

By the linearization in equation (B.6), we can write

$$f(\overline{\pi}) = \sum_{i=1}^{p} [S_i\ (\pi_{s(i)} - l_{s(i)}) + E_i\ (\pi_{e(i)} - l_{e(i)})]^T\ C_i^{o-1}\ [S_i\ (\pi_{s(i)} - l_{s(i)}) + E_i\ (\pi_{e(i)} - l_{e(i)})] \tag{B.9}$$

which after a suitable regrouping of terms is equal to

$$\overline{\pi}^T\ A\ \overline{\pi} - 2\ \overline{\pi}^T\ b + constant \tag{B.10}$$

where A and b are the following sparse matrix and vector

$$A = \sum_{i=1}^{p} \left(\begin{array}{c|c|c|c|c} & & \mathbf{0} & & \\ \hline \mathbf{0} & S_i^T\ C_i^{o-1}\ S_i & \mathbf{0} & S_i^T\ C_i^{o-1}\ E_i & \mathbf{0} \\ \hline & & \mathbf{0} & & \\ \hline \mathbf{0} & E_i^T\ C_i^{o-1}\ S_i & \mathbf{0} & E_i^T\ C_i^{o-1}\ E_i & \mathbf{0} \\ \hline & & \mathbf{0} & & \end{array}\right) \tag{B.11}$$

$$b = \sum_{i=1}^{p} \left(\begin{array}{c} \mathbf{0} \\ \hline S_i^T\ C_i^{o-1}\ (S_i\ l_{s(i)} + E_i\ l_{e(i)}) \\ \hline \mathbf{0} \\ \hline E_i^T\ C_i^{o-1}\ (S_i\ l_{s(i)} + E_i\ l_{e(i)}) \\ \hline \mathbf{0} \end{array}\right) \tag{B.12}$$

(Here we have assumed $s < e$ without loss of generality.)

The minimization can be accomplished by solving for $\overline{\pi}$

$$\frac{\partial f(\overline{\pi})}{\partial \overline{\pi}} = \frac{\partial(\overline{\pi}^T\ A\ \overline{\pi} - 2\ \overline{\pi}^T\ b + constant)}{\partial \overline{\pi}} = 0 \tag{B.13}$$

that is is equivalent to (A is positive, symmetric, and definite)

$$2(A\,\overline{\pi} - b) = 0 \tag{B.14}$$

Therefore, the values for $\overline{\pi}$ that minimize $f(\overline{\pi})$ can be found by solving the linear system

$$A\,\overline{\pi} = b \tag{B.15}$$