Multi-Hypothesis Topological Mapping Using Qualitative Spatial Reasoning

Jan Oliver Wallgrün*
Transregional Collaborative Research Center SFB/TR 8 Spatial Cognition
Universität Bremen, Germany
wallgruen@sfbtr8.uni-bremen.de

Abstract

We describe a multi-hypothesis mapping system for mobile robots that learns graph-based topological representations. Our approach exploits direction information and the assumption of planarity to prune the space of possible map hypotheses. Qualitative spatial reasoning is used to check satisfiability of individual hypotheses. We evaluate the effects of absolute and relative direction information and incorporate the approach into a mapping system based on Voronoi graphs.

Introduction

Learning and maintaining a spatial model of an initially unknown environment is generally regarded as a fundamental problem of mobile robot research. During the last decades, work on this problem has focused on coordinate-based spatial representations like occupancy grids and feature-based representations. An alternative to these representation approaches are graph-based representations, often referred to as *topological maps*. In these approaches the environment is typically conceptualized as a *route graph* (Werner, Krieg-Brückner, and Herrmann 2000) consisting of nodes that stand for distinctive places or navigational decision points and edges that stand for the distinctive paths connecting these places. The problem of computing the correct graph model from a history of local observations has been investigated theoretically for graph environments without any geometric information. For instance, (Dudek et al. 1991) showed that without further information successful map learning cannot be guaranteed without the help of at least one movable marker.

In this text, we are concerned with the problem of making topological mapping robust in the presence of uncertainty and ambiguity in the perceived spatial information. In the majority of toplogical mapping approaches only a single map hypothesis is maintained with the consequence that the map construction process tends to fail as soon as a wrong decision is made. A more promising approach, for instance suggested in (Dudek, Freedman, and Hadjres 1996) and (Kuipers et al. 2004), is to keep track of all possible map hypotheses simultaneously. However, this approach can increase the computational costs dramatically as the number of possible topological map hypotheses can grow exponentially with the number of exploration steps. Hence, additional available information needs to be exploited in order to eliminate as many hypotheses as possible and make the multi-hypothesis approach feasible.

In this work, we extend multiple-hypothesis topological mapping and earlier work on qualitative spatial reasoning in route graphs (Moratz, Nebel, and Freksa 2003). We adopt Kuipers abductive learning approach (Kuipers et al. 2004) and prefer among all valid graph hypotheses one that has a minimal number of nodes. The resulting mapping approach incrementally incorporates observations performing a best-first search through the tree of possible graph hypotheses. In addition, we incorporate qualitative information about the directions of leaving hallways and the assumption that the environment is planar (a constraint which has already been individually investigated in Savelli and Kuipers 2004). Qualitative spatial reasoning and incremental planarity testing are used to discard invalid hypotheses and thereby prune the search space. We employ and compare information from two different qualitative spatial constraint calculi, the absolute cardinal direction calculus (Ligozat 1998) and the relative $\mathcal{OPRA}_2$ calculus (Moratz 2006). Furthermore, we combine the described approach with a topological mapping approach based on generalized Voronoi graphs (Wallgrün 2005) and extensively evaluate our approach using simulation experiments as well as real exploration data from a mobile robot. The experiments show that direction information and the planarity constraint lead to a huge increase in solution quality and decrease in search space. However, the application also reveals shortcomings of existing spatial calculi and shows that the overall problem of topological mapping could serve as a test bed for evaluating spatial reasoning methods.

We start by presenting our general multi-hypothesis topological mapping approach in the next section. We then explain the incorporation of spatial consistency and planarity checking. Finally, we describe the different experiments.

Multi-Hypothesis Topological Mapping

Let us consider the following scenario: A robot is roaming through a graph-like environment like the one shown

*Support by the Deutsche Forschungsgemeinschaft (DFG) is gratefully acknowledged.

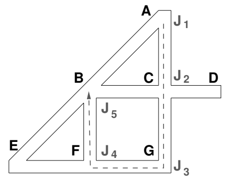

Figure 1: Walk of a robot through a graph-like environment

in Fig. 1. The environment consists of junctions and straight hallways connecting the junctions. For every passed junction, the robot stores a *junction observation* J_i consisting of a cyclically ordered set of leaving hallways $\langle l_1^{[J_i]}, l_2^{[J_i]}, ..., l_n^{[J_i]} \rangle$ and a spatial description consisting of spatial relations over the set of observed leaving hallways, e.g. $\{\text{southwest}(l_1^{[J_i]}), \text{south}(l_2^{[J_i]})\}$.

Junction observations are connected by *hallway traversal actions* consisting of leaving the current junction via one of the observed leaving hallways and arriving at the next junction via one of the leaving hallways belonging to the next junction observation, e.g., $l_2^{[J_1]} \rightarrow l_1^{[J_2]}$. A list $\langle J_1, T_1, J_2, T_2, ..., T_{n-1}, J_n \rangle$ of alternating junction observations J_i and hallway traversals T_j forms the *history* of one particular exploration run through the graph environment.

The goal of a topological mapping algorithm now is to incrementally process the history of observations and actions and for each step determine a route graph hypothesis that can be considered a valid explanation of the information processed so far. Each route graph hypothesis consists of an undirected graph with a combinatorial embedding into the plane (e.g., specified by cyclic orders for the leaving edges of each node in the graph) and the position and orientation of the robot at the beginning of the exploration run (e.g., given by a node and leaving edge).

During exploration, a currently valid hypothesis may turn out to be invalid when the next junction observation is processed. Hence, instead of committing to a single hypothesis, we track all valid hypotheses simultaneously. Fig. 2 shows in the top row three possible hypotheses assuming that the robot has just arrived at junction G in the example (and assuming that all hallways are straight and that the junction observations are given in terms of qualitative cardinal direction relations from Ligozat's cardinal direction calculus (Ligozat 1998)). Black nodes here stand for junctions that have been observed, while white nodes are introduced for the end points of hallways that have not been traversed so far. When moving on to F and processing the new observation J_4, the first hypothesis can be complemented in two different ways leading to two successors in the search tree. Similarly, the third hypothesis has five successors. For the second hypothesis, however, the new observation leads to a contradiction: no hallway leading northeast is observed and, hence, this hypothesis can be discarded completely.

Minimal Route Graph Model Finding

The approach sketched above performs an exhaustive search through the tree of possible hypotheses. A modification of

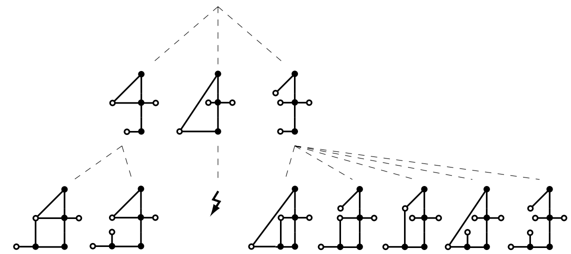

Figure 2: Part of the search space of valid route graph hypotheses for the example from Fig. 1

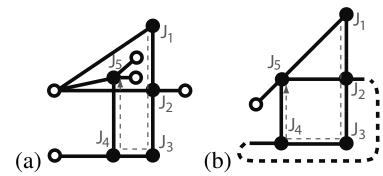

Figure 3: Two invalid hypotheses for the history from Fig. 1

this approach proposed by Kuipers (Kuipers et al. 2004) is to prefer among all valid hypotheses the one that offers the simplest explanation. In this text, we interpret *simplest* as meaning a hypothesis that contains a minimal number of nodes which we will call a *minimal route graph model*.

The number of nodes grows monotonically with increasing depth in the search tree because new nodes and edges will be added but never removed when a new observation is incorporated. As a result, we can search for the minimal route graph model in a best-first manner. This means that in the example the right hypothesis in the top row would not have been expanded because it already has the same number of nodes as the previously generated hypothesis at the bottom left.

Valid Route Graph Models

The search tree from Fig. 2 only contains *valid* route graph hypotheses, while all other hypotheses have already been discarded. Given the direction information contained in the junction observations and assuming that the mapped environment is planar, a hypothesis has to satisfy three conditions to be considered valid:

1. repeating the sequence of actions specified in the history but now within the hypothetical route graph yields a sequence of node degrees identical to the original sequence of leaving hallway numbers (*structural constraint*),
2. there must exist a way to draw the hypothetical route graph into the plane without crossing edges that is in accordance with the specified combinatorial embedding (*planarity constraint*), and
3. given this drawing, repeating the actions also reproduces the direction relations provided by the original junction observations (*direction constraints*).

When generating the successor hypotheses in the search tree, we only generate hypotheses which satisfy the structural constraint. In addition, we take into account that two junction observations can only correspond to the same node in a hypothesis if the perceived directions match. As a result, we can simply store the direction information as constraints to the edges in the graphs.

In Fig. 3 we see two examples of invalid hypotheses, this time for the complete walk depicted in Fig. 1. Both are depicted by one particular drawing of the route graph into the plane and both satisfy the structural constraint. The drawing of the first hypothesis would also reproduce the observed direction relations. However, it has crossing edges and, more importantly, no drawing without crossing edges exists that is in accordance with the combinatorial embedding because the combinatorial embedding itself is not planar.

The drawing of the second hypothesis is planar but the positions assigned to the nodes do not reproduce the direction information correctly as the hallway that is supposed to directly connect the junctions labeled J_2 and J_4 is supposed to lead east from J_2 and arrive at J_4 from the west. Hence, J_4 would have to be to the east of J_2. However, from the knowledge that the hallway connecting J_2 with J_3 leads south and the hallway connecting J_3 with J_4 leads to the west it can be concluded that J_4 has to be somewhere to the west of J_2. As a result, no drawing satisfying the direction constraints can exist because the contained direction information is inconsistent.

The minimal route graph model finding problem we have described here is a combinatorial optimization problem. For deciding whether a given hypothesis is valid or not, we need to determine whether a drawing exists that satisfies the planarity as well as the direction constraints. This is a constraint satisfaction problem over infinite domains (points in the plane). However, as the two examples demonstrate that many structurally valid hypotheses generated during the search process can be ruled out by testing planarity of the combinatorial embedding and the global consistency of annotated direction constraints individually. This is the approach we will take in this work and it allows us to employ the efficient techniques for deciding consistency developed in the area of qualitative spatial reasoning. Nevertheless, the approach is incomplete in the sense that it may not filter out all invalid map hypotheses: There may exist a drawing for a given map hypothesis that is planar and one that is compliant with the direction constraints but none that is both. Results on how well this approach works in practice will be given in the section on experimental evaluation.

Two Mapping Variants

Up to now, we have described a version of the minimal model finding problem in which each model is a complete closed environment that might contain unvisited junctions which form the end points of perceived but never traversed hallways. A less complex version of the problem can be obtained by restricting the models to visited places and allowing hallways with open endings. We will investigate both variants and refer to them as CompEnv and VisOnly.

Rejection Based on Spatial Constraints

In the following, we briefly describe how checking of planarity and, in particular, of consistency of the direction constraints using the absolute cardinal direction calculus (Ligozat 1998) and the relative $\mathcal{OPRA}_2$ calculus (Moratz 2006) are realized in our mapping approach.

Planarity Constraint

Each graph hypothesis for which the cyclic order information derived from the cyclic orders of perceived hallways does not describe a planar embedding (which can be decided in linear time) can be immediately discarded. The criterion used is whether the genus of the graph given by Euler's formula is zero.

Our approach to planarity checking is similar to the one described in (Savelli and Kuipers 2004). We integrate planarity checking into our search algorithm by representing the route graph hypotheses as bidirected graphs and updating the information about faces of the embedding whenever we modify the graph structure. When the genus becomes non-zero, the hypothesis at hand can be discarded as the planarity constraint is violated.

Qualitative Direction Information

To incorporate direction constraints, we formulate observed directions by using the relations from a qualitative constraint calculus. To decide whether the direction information stored in a hypothesis is consistent, we first extract a constraint network from the route graph hypothesis and then employ the qualitative spatial reasoning toolbox SparQ (Wallgrün et al. 2007) for the consistency check using the standard algebraic closure algorithm which runs in cubic time.

Realizing both mentioned calculi allows us to compare the effects of absolute and relative direction information. However, both calculi have their individual shortcomings. The cardinal direction calculus, on the one hand, while being rather efficient because a large tractable subset exists for which the algebraic closure algorithm decides consistency, does not allow for expressing the cyclic order information about the leaving edges in the route graph. As a result, it can happen that a constraint network deemed consistent by the consistency check, only has solutions for which the cyclic order information is not preserved.

$\mathcal{OPRA}_2$, on the other hand, can express the cyclic order information. However, for $\mathcal{OPRA}_2$ algebraic closure does not decide consistency even for atomic constraint networks which also means that inconsistent hypotheses may not be discovered. We still have chosen $\mathcal{OPRA}_2$ as to our knowledge no better suited relative direction calculus exists.

When using the absolute cardinal direction calculus, the extracted constraint network contains one variable for each node in the graph and the constraints holding between them are directly derived from the direction relations annotated to the edges. In contrast, the $\mathcal{OPRA}_2$ calculus is a relative calculus describing the relative orientation of two oriented points (points in the plane with an additional direction parameter) towards each other. In order to determine the right

$\mathcal{OPRA}_2$ relation, a robot would only need to be able to estimate the angles between the leaving hallways instead of needing a compass. In the constraint network, one oriented point variable is introduced for each pair of node and incident edge (meaning one for every leaving hallway). Hence, we end up with $2 \times n$ variables where n is the number of edges in the hypothesis.

Experimental Evaluation

To evaluate the application of qualitative spatial reasoning approaches to the topological mapping problem, we performed several simulated exploration experiments in randomly generated graph environments of varying size and using random walks of varying length through the graphs. In addition, we combined the overall approach with a mapping system based on Voronoi graphs and applied it to a data set from a real-world exploration run.

Simulation Experiments

In the simulation experiments, we investigated several aspects of our approach. The main results are summarized below.

Solution quality We first investigated how much the planarity constraint and qualitative direction information helps in order to improve the solution quality by ruling out incorrect hypotheses and, as a result, increase the frequency in which the correct solution is found by the minimal model approach. To measure the quality of a solution, we use a simple error measure: We count how often either two junction observations that correspond to different junctions have been mapped to the same node or two observations that correspond to the same junction have not been unified. To assure that even searching without pruning is possible in reasonable time, we used rather small problem instances varying the size of the environment between 4 and 16 junctions.

Table 1 shows the results of 15600 trials for each of the following settings and both the CompEnv and VisOnly variants of the minimal model algorithm: (1) only structural constraint, (2) structural constraint and planarity constraint, (3) structural constraint and cardinal direction constraints, (4) structural constraint, planarity constraint, and cardinal directions. In addition, Fig. 4 shows how the average error distances increase with the size of the correct model throughout the experiment.

As the average error distances show, the planarity constraint and in particular the direction constraints significantly improve the solution quality. For the CompEnv variant, the planarity constraint achieves a 26.27% reduction of error distance, while direction information decreases the error distance by 85.70%. Combining both planarity and direction constraints only gives slightly better results than without applying the planarity constraint. The application of the constraints is highly beneficial but in most cases is not sufficient to resolve all ambiguities.

For the VisOnly case in which unvisited junctions are not included in the model, the improvements are even more drastic. The application of cardinal direction information

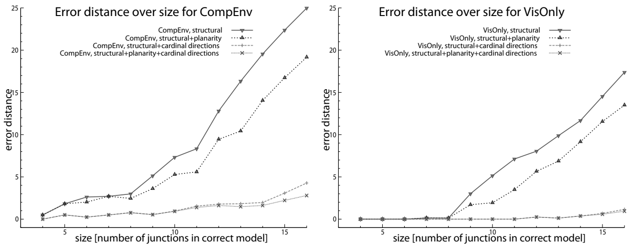

Figure 4: Average error distance depending on the size of the correct model for the CompEnv and VisOnly variants

leads to the correct model being found in 98.15% of all trials and has an extremely low average error distance of 0.20, or 0.17 when combined with the planarity constraint.

Setting		Correct model found	Average error distance
CompEnv	structural only	4.77%	9.86
	structural, planarity	5.97%	7.27
	structural, card. dir.	50.62%	1.41
	structural, planarity, card. dir.	50.92%	1.18
VisOnly	structural only	59.00%	5.63
	structural, planarity	64.77%	4.07
	structural, card. dir.	97.92%	0.20
	structural, planarity, card. dir.	98.15%	0.17

Table 1: Experimental results regarding the solution quality

The experiment shows that the planarity constraint and in particular the cardinal direction constraints are able to resolve most of the model ambiguities remaining on the structural level leading to a largely increased solution quality.

Pruning efficiency To investigate the effects of the individual settings on the size of the hypothesis space that has to be searched, we performed random experiments running the search until all solutions up to the size of the correct solution had been determined to mask out the effects of varying solution quality. We recorded (1) the number of expansion steps in which successors of a hypothesis are generated, (2) the average branching factor in the search tree, and (3) the maximal queue size occurring during the search.

The result of this experiment are summarized in Table 2. Fig. 5 shows how the number of expansions grows with increasing size of the environment (logarithmic scale is used for the y-axes).

Setting		Expansions	Branch. factor	Max. queue size
CompEnv	structural only	2407.09	4.49	833.96
	structural, planarity	284.97	2.38	86.17
	structural, card. dir.	39.84	2.48	13.58
	structural, planarity, card. dir.	21.85	1.64	6.10
VisOnly	structural only	790.61	3.19	160.88
	structural, planarity	254.25	2.00	47.87
	structural, card. dir.	20.72	1.18	2.95
	structural, planarity, card. dir.	20.11	1.15	2.76

Table 2: Results regarding the pruning efficiency

We clearly see that the CompEnv variant of the minimal model finding problem is much more complex than the VisOnly variant. The planarity constraint leads to an 88.16% decrease in expansion steps for CompEnv and 67.84% for VisOnly. The average branching factor has been decreased

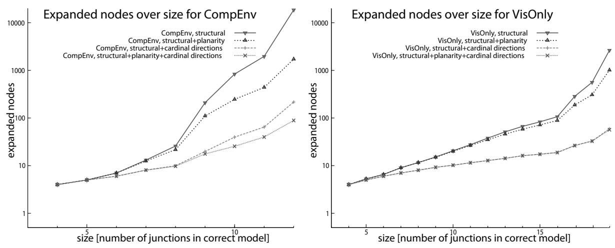

Figure 5: Comparison of expansion steps depending on the size of the correct model for CompEnv and VisOnly

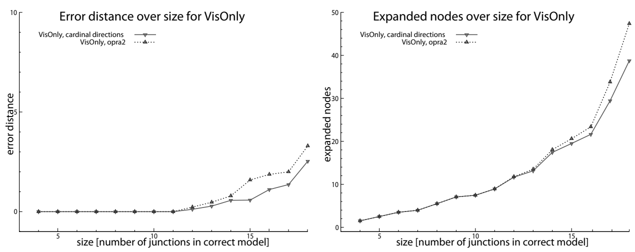

Figure 6: Error distance and expansion steps for the cardinal direction calculus and $\mathcal{OPRA}_2$ (VisOnly)

by 46.99% to 2.38 (CompEnv) and by 37.30% to 2.00 (VisOnly). For the cardinal direction constraints, we see a very high reduction of node expansions of 98.35% for CompEnv and 97.38% for VisOnly. By combining both, an extreme reduction in expansion steps of 99.99% was achieved for CompEnv which corresponds to an average branching factor of 1.64. For VisOnly the cardinal direction constraints yields a 99.97% reduction (branching factor 1.18).

We conclude that the planarity assumption and the coarse direction information given by the qualitative cardinal relations lead to a much increased efficiency of the minimal model finding approach which would otherwise only be feasible for very small problem instances.

Absolute vs. relative direction information One of the goals of our analysis was to compare the effects of employing absolute direction information (e.g., relations from the cardinal direction calculus) and relative direction information (e.g., $\mathcal{OPRA}_2$ relations). Therefore, we repeated the experiments for determining solution quality and pruning efficiency for both calculi using all constraints. Fig. 6 shows the diagrams for error distance and expansions steps for VisOnly. With regard to solution quality, the change from absolute to relative direction information increased the average error distance from 1.64 to 1.82 for CompEnv and from 0.44 to 0.68 for VisOnly. The average number of expansion steps increased from 49.12 to 82.60 (branching factor from 1.33 to 1.42) for CompEnv and from 12.79 to 13.93 (branching factor from 1.21 to 1.24) for VisOnly.

The observed decrease in performance is not surprising as relative direction information in general allows for more perceptual aliasing. Taking this into account, the decrease in performance seems to be rather mild, especially for the VisOnly variant, and still much lower than for the less constrained settings investigated in the previous experiments. The main advantage of relative information is that it often can be obtained more easily and more reliably.

Overall computational costs When investigating the pruning efficiency, we restricted ourselves to small problem instances. In addition, we focused on the effects of the different settings on the search space. For the complete minimal model finding approaches featuring all kinds of constraints, we further investigated how the approaches perform for larger problem instances. This investigation yielded two main results: First, even applying all constraints is not sufficient to conquer the combinatorial explosion for the CompEnv variant. Second, the computational costs of global consistency checking when employing the relative $\mathcal{OPRA}_2$ calculus quickly becomes excessive, making this approach infeasible for large environments for both variants.

As a result of the first observation, the CompEnv variant seems limited to scenarios with a rather small number of junctions in which the ability to predict the structure of unvisited parts is worth the increased computational costs. Taking a closer look at the second issue revealed that the computation times spent on global consistency checking for $\mathcal{OPRA}_2$ rise sharply, in some cases making up 90% of the overall computation time. We believe that there are two issues that contribute to this explosion in computational costs: the large number of base relations in the $\mathcal{OPRA}_2$ calculus which makes it impossible to store the complete composition table of general relations, and the size of constraint networks contains $2\times$ the number of edges as variables.

As a result, it seems that currently the VisOnly variant in combination with absolute direction information is the only one that scales sufficiently well to larger environments.

Real-World Experiment

In a last experiment, we integrated the minimal model framework into a topological mapping system based on generalized Voronoi graphs (see Wallgrün 2005 for details) to demonstrate its ability to map real environments. The environment and the trajectory of the robot during the experiment is shown in Fig. 7(a)[1]. The mapping system extracts local Voronoi graphs from local grid maps of the robot's immediate surrounding and incrementally generates the history information about the observed Voronoi nodes and traversed Voronoi edges. The minimal model finding module updates the search tree based on new history information and computes a new minimal graph model using the VisOnly variant.

To apply the minimal model approach to the Voronoi graph representation several adaptations were made:

1. Multiple connections between two nodes are allowed.
2. Observed local Voronoi graphs can contain multiple nodes and edges which are translated into history information without actually traversing the edges.

[1]The data set has been recorded at the Intel Research Lab, Seattle, and is available at (`http://radish.sourceforge.net/`), courtesy of D. Hähnel.

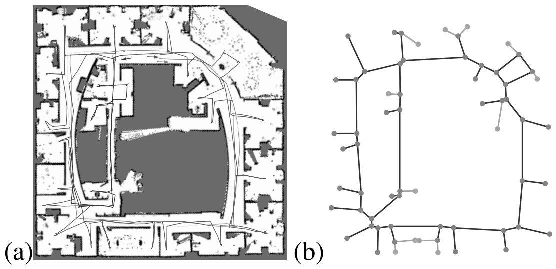

Figure 7: (a) Environment used in the real-world experiment. (b) Minimal route graph hypothesis computed

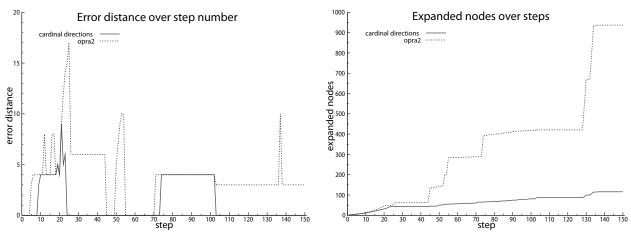

Figure 8: Error distance and expansion steps over the steps of the real-world exploration experiment

3. In practice, it may not be possible to reliably determine the exact direction relations. Therefore, we utilize disjunctions of base relations when the perceived direction is a linear relation or lies close to the boundary of a relation sector (e.g., $\{ne, n, nw\}$ for observed relation $\{n\}$).
4. Voronoi curves are typically not straight line connections. Hence, we only employ direction constraints in the global consistency check if both connected Voronoi nodes have been perceived simultaneously. Otherwise, the direction information for this edge is only used for matching junction observations.

Fig. 7(b) shows the minimal model computed using the cardinal direction calculus which is indeed the correct graph model for this exploration run. For $\mathcal{OPRA}_2$, the resulting model was correct as well except for two wrongly merged nodes in a room that was entered via two different doors. However, while the computation took 16 seconds using cardinal directions, it took over 10 hours for $\mathcal{OPRA}_2$ because of the issues described in the previous section.

Fig. 8 shows how error distance of the current minimal model and number of expansion steps develop over the 150 exploration steps for both spatial calculi. The diagram for the error distance shows that the variant using cardinal directions immediately settles for the correct hypothesis when the first loop traversal is completed in step 23, while this takes almost the entire second loop for $\mathcal{OPRA}_2$. We later see another increase in error distance caused by entering the new rooms. We also see that the number of expansion steps required and the number of tracked alternative hypotheses is significantly higher for $\mathcal{OPRA}_2$. However, the main reason for the hugely increased computation time again is the time spent on the global consistency checking.

Conclusions

We formulated topological mapping as the problem of finding a minimal route graph model that explains a sequence of observations and actions. Our solution consists of a search through the tree of possible graph hypotheses exploiting qualitative direction information (absolute or relative) and the planarity assumption. The experimental evaluation showed that this approach leads to a significantly reduced search space and improved solution quality. The approach has also been incorporated into a Voronoi-based mapping system and been applied to real exploration data.

The results with regard to spatial consistency checking based on qualitative direction information can be seen as a challenge for future research on qualitative spatial reasoning to improve current reasoning methods. We also think that the described problem can serve as a test bed to evaluate and compare different approaches.

References

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

Dudek, G.; Freedman, P.; and Hadjres, S. 1996. Using multiple models for environmental mapping. *Journal of Robotic Systems* 13(8):539–559.

Kuipers, B.; Modayil, J.; Beeson, P.; MacMahon, M.; and Savelli, F. 2004. Local metrical and global topological maps in the hybrid Spatial Semantic Hierarchy. In *Proceedings IEEE International Conference on Robotics and Automation 2004 (ICRA-04)*, 4845–4851.

Ligozat, G. 1998. Reasoning about cardinal directions. *Journal of Visual Languages and Computing* 9:23–44.

Moratz, R.; Nebel, B.; and Freksa, C. 2003. Qualitative spatial reasoning about relative position: The tradeoff between strong formal properties and successful reasoning about route graphs. In Freksa, C.; Brauer, W.; Habel, C.; and Wender, K. F., eds., *Spatial Cognition III*. Springer.

Moratz, R. 2006. Representing relative direction as binary relation of oriented points. In *Proceedings of the 17th European Conference on Artificial Intelligence (ECAI 2006)*.

Savelli, F., and Kuipers, B. 2004. Loop-closing and planarity in topological map-building. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS-04)*, 1511–1517.

Wallgrün, J. O.; Frommberger, L.; Wolter, D.; Dylla, F.; and Freksa, C. 2007. Qualitative spatial representation and reasoning in the SparQ-toolbox. In Barkowsky, T.; Knauff, M.; Ligozat, G.; and Montello, D., eds., *Spatial Cognition V*. Springer. 39–58.

Wallgrün, J. O. 2005. Autonomous construction of hierarchical Voronoi-based route graph representations. In Freksa, C.; Knauff, M.; Krieg-Brückner, B.; Nebel, B.; and Barkowsky, T., eds., *Spatial Cognition IV*. Springer.

Werner, S.; Krieg-Brückner, B.; and Herrmann, T. 2000. Modelling navigational knowledge by route graphs. In Freksa, C.; Habel, C.; Brauer, W.; and Wender, K. F., eds., *Spatial Cognition II*. Springer. 295–316.