Chapter 8

Conclusions and Outlook

We conclude this book by summarizing and discussing the results achieved and by providing an outlook on promising directions of future research arising from the work described in the previous chapters.

8.1 Summary and Conclusions

We started this work with the observation that most work on robot mapping is concentrated on developing powerful techniques for uncertainty handling for a small number of elementary low-level representation approaches: grid-based representations, geometric representations, and landmark-based representations. However, as we argued, in order to realize or explain high-level spatial cognitive abilities, more emphasis needs to be put on representational matters. Robot mapping needs to be studied in a general setting involving the perspectives of both spatial representation and uncertainty handling. In addition, it needs to explicitly take into account the operations that employ the environmental model to successfully solve spatial tasks.

Although more abstract forms of representation and more complex forms of organization have been proposed, we observed that only a small number of attempts have been made to achieve a high level of robustness in the presence of uncertainty for these approaches. This led to the main thesis of this work, namely that the combination of uncertainty handling methods developed for the traditional representations with abstract high-level representations is a very promising direction of research. It has the potential to lead to mobile robots with a much higher level of spatial capabilities and competences and, hence, to open up a new range of applications.

In the following, we directed our attention at one particular class of representations, namely route graph representations, derived as abstractions of the generalized Voronoi diagram of the environment. With the annotated generalized Voronoi graph (AGVG) containing a particular set of node and edge attributes, we introduced a concrete representation approach and proposed a hierarchical organization of this representation, the HAGVG. As we discussed, the advantages of this hierarchical AGVG representation

J.O. Wallgrün, *Hierarchical Voronoi Graphs: Spatial Representation and Reasoning for Mobile Robots*,
DOI 10.1007/978-3-642-10345-2_8, © Springer-Verlag Berlin Heidelberg 2010

are compactness and good scalability, well-suitedness for global route planning and route-based communication, the ability to consistently represent environments without requiring metrical consistency, and the ability to serve as a core structure for anchoring any kind of additional information in the form of further annotations. Furthermore, the approach directly facilitates systematic exploration of an environment in order to construct a complete environmental model. Parts of the proposed annotations are intended to help distinguish different places and, hence, improve localization without the need for a complete geometric model.

We also identified several limitations of the approach. The area of applicability is restricted as the approach presupposes that the environment shows a clearly pronounced route structure, e.g., indoor environments or path networks. Moreover, instability of the underlying generalized Voronoi diagram when the input data is noisy poses a challenge for the development of reliable construction and maintenance methods.

The focus of this book then has been on developing fundamental techniques required for the construction and maintenance of the HAGVG representation, mainly with regard to a mobile robot equipped with a laser scanner or similar range sensors which yield a 2D geometric description of the robot's surroundings. This work has been conducted on three levels.

First, we were concerned with the reliable autonomous extraction of AGVGs and hierarchical construction from a data structure point of view. The correct identification of corresponding features was taken for granted. Second, we studied the problem of robust data association for local matching and tracking of Voronoi nodes. Third, we turned to the global mapping problem, which we formulated as the problem of finding a minimal route graph model under spatial constraints. We studied this problem with a focus on effects of reliable qualitative directional information in combination with the planarity constraint.

In the following, we briefly summarize the results achieved for each of these three areas and with regard to the overall mapping approaches discussed in the evaluation chapter.

8.1.1 Extraction and HAGVG Construction

Concerning the extraction of AGVGs and the construction of a hierarchical representation, we proposed measures for the relevance of a Voronoi node and the regions accessible via its edges. The relevance values can be computed without requiring a complete geometric description of the environment, simply from the information stored in an AGVG. We developed efficient computation and incremental update algorithms. Based on the measures, we developed a simplification algorithm for AGVGs. This algorithm forms the core of our hierarchy generation scheme used to autonomously build up HAGVGs which represent the environment at different levels of granularity.

We also demonstrated that the concept of relevance coincides well with the stability of Voronoi nodes under noisy conditions. As a result, the approach can also be employed to deal with problems of the graph construction process caused by the instability of the underlying Voronoi diagram. The described approach worked very reliably and served as a basis for the other techniques developed in this book.

8.1.2 Data Association and Matching

To achieve reliable data association, we extended constraint-based branch and bound approaches by incorporating structural constraints arising from the graph topology and the combinatorial embedding of AGVGs. We developed a dynamic programming approach based on edit distance to match AGVGs. The relevance values of nodes were used to determine the costs of the different edit operations. We incorporated several types of constraints in order to reduce the search space: (1) unary constraints derived from node similarity and, if available, position estimates, (2) binary distance constraints, and (3) ternary angle constraints. In order to achieve a high level of efficiency, we restricted the approach to tree-formed AGVGs (which we enforced if needed) and exploited the fact that the resulting trees are ordered as specified by the planar embedding of the AGVGs.

When we evaluated and compared our approach with a standard individual compatibility nearest neighbor approach, it proved to yield very low error rates with only a mild increase in computational costs. The approach worked very well to locally track Voronoi features, even when no odometry information is available. However, the approach in its current form does not support loop closing.

8.1.3 Minimal Route Graph Model Finding

Bringing together work from topological mapping and qualitative spatial reasoning, we finally formulated global mapping of a graph-like environment as the problem of finding a minimal route graph model that explains a sequence of junction observations and hallway traversal actions. A solution to the problem has been developed consisting of a best-first branch and bound search through the interpretation tree of possible associations of junctions.

We integrated planarity and spatial consistency checking into the search algorithm in order to prune the search space. For spatial consistency checking, we employed constraint-based reasoning techniques based on qualitative spatial calculi. In our evaluation we focused on direction information and compared the effects of absolute and relative direction information. We also compared two versions of the algorithm, one that computes complete environment models (CompEnv) and one that only models the junctions that have actually been observed (VisOnly).

The results of the evaluation using randomly generated graph worlds as well as data extracted from a few selected real environments showed that the combination of the planarity constraint and in particular the direction constraints significantly reduced the search space and ambiguity of the history. This resulted in a very high rate of correctly mapped environments once the exploration paths covered a large enough fraction of the environments. One general advantage of the approach is that as soon as the correct model has been determined, the time needed for incorporating new observations becomes constant.

When combining the VisOnly variant with absolute direction information, the approach was efficient enough for all the environments we used in the experiments. However, employing relative direction information in the form of relations from the $\mathcal{OPRA}_2$ calculus quickly led to runtime problems caused by the global consistency test, which does not scale sufficiently well to larger environments.

The CompEnv variant, on the other hand, allows making predictions about connections and junctions that have not been observed directly. This can, for instance, help to find shortcuts. However, the experiments showed that this approach is limited to smaller environments because of the increased complexity of the problem. Hence, it can only be employed beneficially in suitable environments (e.g., a network of hallways) and when observations are provided on a very high level of abstraction (e.g., the robot can detect junctions, rooms, etc.).

In order to apply the theoretical minimal model finding framework to our Voronoi-based representations, a set of modifications and relaxations was required. However, it turned out that these adaptations only lead to a small decrease in overall performance.

8.1.4 Complete Mapping Approaches

The evaluation of the techniques developed in this work also involved three different examples of combining the (H)AGVG representation with uncertainty handling methods which demonstrated that this approach is indeed fruitful. The overlay approach consisting of a grid map representation generated via a FastSLAM approach and an HAGVG demonstrated the robustness of the route graph extraction and abstraction methods. The approach, however, does not scale well to very large environments because of the high space consumption and the increased computation time needed to extract the underlying Voronoi graph.

By employing a feature-based FastSLAM approach that uses the developed extraction and data association techniques and builds up the route graph representation directly, we were able to avoid these problems, but our approach does not facilitate loop closing. In addition, the approach is subject to the particle depletion problem.

In the final mapping system, we used the extraction methods and the data association techniques in order to derive a history of observations and actions at a high level of abstraction. Applying the minimal route graph model finding approach then allowed us to determine the correct route graph hypothesis. This approach has the advantage that no absolute coordinates need to be assigned to the Voronoi nodes and the relative metric information contained in the route graph may be globally inconsistent without diminishing its applicability. However, like many other approaches to robot mapping, this approach still depends on the correctness of the extracted history and, hence, the extraction and data association processes. Unfortunately, errors at these levels cannot be ruled out entirely and extensions of the approach are required, for instance involving multiple alternative histories.

Nevertheless, this overall mapping system demonstrates that with the developed techniques we have achieved a significant progress towards the robust learning of Voronoi-based route graph representations for mobile robots. Part of this is due to the fact that fine metric relations are only considered at a local level, while at the top level we were more concerned with abstract relations of connectivity and coarse direction relations. This made the discrete multiple-hypothesis approach to the global mapping problem feasible. Techniques such as the data association based on ordered tree matching also could easily be adapted for other kinds of route graph representations used for robot mapping. The minimal route graph model finding framework is very general with regard to the actual representation approach and application area, and can be employed whenever local observations of decision points should be integrated into a global graph model. To give an example, an outdoor robot able to reliably detect ways and paths in a park could very well use the minimal route graph model framework for global mapping without employing a Voronoi-based approach.

8.2 Outlook

We close this text by discussing open questions and promising continuations arising from the work presented here. We start at the technical level and then gradually broaden our view coming back to the main theses of the work, resulting research directions, and general challenges for the field of mobile robot research.

8.2.1 Extensions of the Work Described in Chaps. 3–6

In Chap. 3, we argued that it is one advantage of the route graph representation that it makes it easy to anchor additional information in the form of further attributes to its nodes and edges. From a representational point of view, one such kind of information that is very relevant for navigation and route-based communication in humans is landmark information, e.g., information about salient objects encountered along a route segment or at a decision point (Richter, 2007). In addition, distinct landmarks have the potential to greatly improve localization and further reduce ambiguities with regard to different map hypotheses. Hence, it would be desirable to include and exploit landmark information in the Voronoi-based mapping approach. However, to make the step from simple geometric features to distinct landmarks that can support human-robot communication, largely increased sensor and object recognition capabilities are required on the side of the robot.

With regard to the relevance measure for Voronoi nodes and the automatic hierarchy generation approach, one thing that could be improved is the treatment of cycles. As we explained, we treat leaving edges that are part of cycles as maximally relevant, which is reasonable when we consider real loops within the environmental structure.

For small cycles caused by clutter or small obstacles, this approach can yield non-intuitive results. While it should be possible to adapt the relevance measures to distinguish these two situations, we believe that aiming at a way to compute the Voronoi graph from a local description which excludes all kinds of clutter and objects (e.g., only with respect to the walls in an indoor environment) will be more advantageous. We will come back to this point again below.

In the work on minimal route graph model finding under hard spatial constraints in Chap. 6, we focused on directional information, restricting the relative positions of the junctions in a route graph model. The reason for doing so is that directions of leaving hallways can often be perceived very reliably, which makes them a good candidate for pruning. Still, other kinds of spatial information can be available as well, and while direction information here is treated as information about junctions, information regarding the traversed hallways can also be used.

The first thing that comes to mind here is information about the length of a traversed hallway, which would restrict the distance between the connected junctions. While extending the theoretical framework for this kind of information is straightforward, the lack of qualitative distance calculi or real positional calculi comprising distance and direction information already shows that distance information is not such a good candidate for use as a hard constraint for pruning.

As we also discussed in Chap. 6, the individual checking of planarity and consistency of the direction information means that the entire procedure has to be classified as an incomplete method that may not discover and discard all inconsistent hypotheses. The facts that the cardinal direction calculus cannot express the cyclic edge orderings of the combinatorial embedded AGVGs and that the consistency check based on algebraic closure is incomplete in the case of the $\mathcal{OPRA}_2$ calculus can cause additional undiscovered inconsistencies. Although we found no indication that this is a problem in practice during the experimental evaluation, a complete approach would naturally be desirable. However, it seems doubtful that a complete approach can be achieved without increasing the computational costs in a way that makes the application within the search algorithm infeasible.

One thing we did not investigate closer in this work is the effects of applying finer-grained qualitative information. For relative direction information, this would directly be possible by switching to a higher granularity parameter within the $\mathcal{OPRA}_m$ calculi family. For absolute direction information, one option would be to employ a variant of the Star calculus (Renz & Mitra, 2004). In general, if the sensor system can reliably provide information at this finer granularity, a noticeable increase in pruning efficiency and solution quality can be expected.

8.2.2 Combining Voronoi Graphs and Uncertainty Handling

The three different ways of combining Voronoi-based route graph representations and uncertainty handling methods considered in Chap. 7 of this book have to be seen more as case studies in order to evaluate the developed techniques rather than as perfect solutions. For instance, the final overall mapping system discussed in Sect. 7.4, as mentioned, still has the shortcoming that it relies on a correct sequence of Voronoi node observations and Voronoi curve traversal actions, which in turn is reliant on correct tracking of these features while the robot moves through the environment. Hence, robustness could be further increased by considering alternative histories whenever the data association becomes ambiguous. Alternatively, uncertainty arising from extraction and data association might be expressed in the history and incorporated into the search algorithm.

Another promising approach would be to combine the constraint-based framework with probabilistic methods to estimate the likelihood of the individual hypotheses by incorporating information that is currently not used on the global level, e.g., global position estimates. This likelihood assignments could be used in multiple ways: to distinguish between multiple existing minimal hypotheses, to reject hypotheses for which the likelihood falls below a certain threshold, or to limit the number of hypotheses tracked by discarding the least likely ones when the number of hypotheses starts to exceed the limit.

In principle, it would also be possible to drop the minimal model idea and turn the algorithm into one that searches for the most likely route graph hypothesis as in the lazy data association approach described by Hähnel et al. (2003), while still enforcing the hard constraints. The crucial problem of these approaches is to formulate an adequate probabilistic model of the processes involved in the Voronoi graph computation. In how far rough approximations, e.g., based on similarity, can be employed here remains to be investigated.

8.2.3 Challenges for Voronoi-Based Representation Approaches

One downside of Voronoi-based representations we have already mentioned multiple times is their restricted applicability as the environment needs to suit a route-based approach. For robots equipped with laser range finders, problems can arise even in indoor environments, either when the environment contains very large open areas or when it is extremely cluttered.

Large open areas can cause problems because when the limited range of the laser scanner prevents the detection of even the nearest obstacle boundaries, no local GVD can be computed. With the range of today's laser range finders and the computation of the GVD from a local metric map this problem becomes less severe. Alternatively, Beeson et al. (2005) recently proposed an extension of the GVG called the extended Voronoi graph in which open areas are represented by edges enclosing them at a certain distance from the boundaries. Incorporating this approach into our system seems in principle possible, but would require us to rework some of the techniques.

Highly cluttered environments, on the other hand, can cause a lot of difficulties because they can lead to a high density of nodes and mini-cycles (e.g., when a robot only sees the legs of chairs and tables in a cafeteria). Again, we suggest that the most adequate approach is to deal with these kinds of problems on the extraction level by generating an intermediate representation only containing the most stable obstacles and computing the underlying GVD from there. However, this requires scene understanding capabilities which are beyond current technology.

In recent years, a lot of research has been undertaken to extend 2D mapping approaches to full 3D mapping. From the perspective of a route graph representation, this step only would make sense if the agent is also able to freely move in 3D space. As this is not the case for the wheeled mobile robots we have been considering here, it seems more reasonable to exploit the 3D sensor information in order to improve the computation of a normal 2D embedded route graph. For instance, for the extraction of the underlying GVD, a local metric 3D representation would be advantageous. It would for instance remove problems caused by situations in which the robot sometimes does and sometimes does not see certain obstacles because of varying height or inclination of the laser scanner. In addition, local information to improve the localization could be extracted from the 3D map as well and then annotated to the 2D route graph. The resulting system would consolidate local 3D geometric information into local 2D information and then further into local route graphs which are integrated in order to produce the global route graph representation. If indeed a 3D embedded equivalent of the (H)AGVG is needed, a good starting point would be an extension of the GVG to higher dimensions proposed by Choset which ensures that the resulting structure is still a connected network of one-dimensional curves (Choset & Burdick, 2000; Choset et al., 2000).

One major problem of the Voronoi-based approach and similar route graph approaches that we have not touched on in this work is dealing with dynamics. As the Voronoi-based approach attempts to determine the topology of free space, it is particularly sensitive to changes that affect this topology. For instance, an object temporarily placed in the middle of a hallway will induce a change from a single Voronoi curve to two new ones passing this object to the left and right. Such changes can significantly complicate the matching and may require a modification of the graph structure of already mapped areas when they are detected.

Overall, different kinds of dynamics will need to be tackled at different levels, and most of them are more adequately dealt with by working around the sensor and recognition limitations of current robots instead of within the context of mapping. Changes caused by objects moving around the robot, for instance, may be filtered from the sensor data or their effects could be completely masked out by following our long-term suggestion of basing the Voronoi computation only on stable objects. The latter approach could also deal with changes caused by the replacement of movable objects. Adequately dealing with topological changes as caused, for instance, by doors will only be possible by incorporating improved object recognition capabilities and background knowledge into the mapping process. From the mapping point of view, the hardest problem is probably caused by long-term changes to the main structure of an environment, e.g., by moving the inner walls within a configurable building. However, these kinds of changes also tend to confuse humans and may well just trigger the learning of a new model from scratch.

Finally, while in this work we have only considered mapping as a passive procedure that processes the history of observations and actions, another interesting direction of research would be to use the state of the multi-hypothesis search tree used in the minimal model approach in order to actively guide the exploration behavior of the robot. For instance, the robot could be deliberately moved to a place that allows deciding between two possible hypotheses. As a result, the number of hypotheses that need to be tracked simultaneously could be reduced significantly.

8.2.4 Challenges for Qualitative Spatial Reasoning

Aside from the robot mapping problem, this book has also resulted in some interesting observations with regard to abstract spatial reasoning. The results with regard to spatial consistency checking based on qualitative direction information can well be seen as a challenge for future research on qualitative spatial reasoning. As we have argued, none of the currently existing directional spatial calculi ideally fit the demands arising from our theoretical problem.

The cardinal direction calculus, for instance, while having good computational properties cannot express cyclic ordering information. For relative directional calculi like the $\mathcal{OPRA}_2$ calculus, the standard constraint reasoning techniques like algebraic closure are typically incomplete. In addition, even the application of the standard algebraic closure algorithm quickly became infeasible for $\mathcal{OPRA}_2$.

Therefore, either new qualitative calculi or new reasoning techniques for consistency checking are needed. Alternatively, if there is a fundamental conflict between the different demands, the involved trade-offs have to be investigated further. Finally, calculi integrating different aspects of space would ultimately be required if the theoretical problem is to be extended by other qualitative spatial information in addition to the directional information.

8.2.5 The Future: Towards Spatially Competent Mobile Robots

At the beginning of this text, we argued that in order to advance towards more capable robots able to operate in large-scale heterogeneous environments, to communicate about space with other robots or with humans, and to interpret and utilize externalized spatial information in the form of sketches or real maps, a turning away from the simple homogeneous and sensor-near representation approaches is required. The hierarchical route graph representations considered in this work naturally can only be seen as a tiny step in this direction, providing a suitable conceptualization for indoor environments, street networks, etc. And while realization of higher cognitive abilities based on this kind of representation was outside the scope of this text, work in this direction has already started, e.g., with regard to carrying out verbal route instructions (Shi et al., 2007). However, to achieve the overall goal, it seems likely that such an approach will only be one building block in the pool of representation schemes or conceptualizations, from which a robot can choose the most suitable one based on the properties of the considered part of the environment or the task at hand.

From the previous discussions concerning the general improvement of robustness and dealing with dynamics, it should have become clear that we think that progress with regard to more abstract spatial representations and complex forms of organization depends to a large degree on progress achieved in other areas of robotics and AI research. The key components in our opinion are improved object recognition and scene analysis capabilities and the inclusion of background knowledge, spatial and also non-spatial (e.g., regarding stability, temporal aspects, or typical behavior), into the mapping process. As we pointed out, the techniques developed in this work can directly benefit from advancement in these fields.

In the long term, such a development could mean that robot mapping will turn more and more into an all-embracing AI-complete problem. At any rate, there can be little doubt that robot mapping will remain a challenging and exciting field of research for years to come.