Chapter 9
Conclusion

This book presented different approaches for adding semantic information to the representations of indoor environments. We concentrated on extending the information in the maps created by a mobile robot with labels that represent different places in the environment. These places have different functionalities, such as corridors, offices or kitchens. Moreover, throughout this book we have seen how the semantic information about places can improve the capabilities of mobile robots in different domains including human-robot interaction, localization, and exploration.

We first presented a technique based on supervised learning that enabled a mobile robot to recognize the different places in an indoor environment using a laser sensor. To carry out this classification the robot should first take observations and then extract some features from them. These features were later used to recognize the different places. As main observations we used the range measurements of laser range finders, from which several features were extracted that encoded their geometrical properties.

The learning method used for classifying the different places was based on the AdaBoost algorithm. The input to this algorithm was the features extracted from the observations, and as output we obtained a strong classifier which included the more informative features for each place.

The geometrical features are quite good candidates for generalization, since they encode space information. We saw in Chapts. 3 to 5 that the strong classifier created with geometrical features could successfully be transferred among different environments. The main reason is that indoor environments usually contain the same type of places, as for instance, corridors, doorways and rooms. These places share similar structures between the different indoor environments: corridors are typically elongated, and rooms are usually more compact and cluttered. These common characteristics permitted the robot to learn a classifier in one environment and recognized similar places in different ones.

In addition, we used vision sensors to increase the number of places to classify. The main problem with vision observations was to select the features that maintained a good generalization in the classifier over different places. We opted for counting the number of specific objects that appeared in a panoramic image

Ó.M. Mozos: Semantic Labeling of Places with Mobile Robots, STAR 61, pp. 109–111.
springerlink.com

taken by the robot. The selection of these features was motivated by the fact that typical objects appear at different places with different probabilities. For example, the probability of finding a computer monitor in an office is larger than finding one in a kitchen. Again these features are usually very common in several indoor environments.

The previous approach for semantic classification was used to classify a single pose of a mobile robot using the laser and image features. However, this method did not take into account the classification of neighboring poses when labeling the current one. To include this information, we extended the approach with some probabilistic techniques. We first smoothed the classification of all poses in an environment using probabilistic relaxation, and alternatively an instance-based associative Markov network. Both approaches improved the final classification using neighboring information, and allowed the robot to extract compact regions of the environment and create a topological map.

Mobile robots are dynamic agents that move along different trajectories. When operating in indoor environments, the robots usually have a moderate velocity and a relatively continuous movement. Furthermore, certain transitions between places in a trajectory are more likely. For example, to go from the kitchen to the office a robot should traverse a doorway first. This transitional information was encoded in a hidden Markov model and successfully applied to smooth the classification of the poses of a mobile robot along a trajectory. Different results using this approach were presented in Chapt. 5.

The semantic information about places can improve other typical robotics tasks. The main idea is that mobile robots can use the intrinsic information of human-made environments to improve their actions. In particular, we showed how the information about places could improve the performance of a team of mobile robots during exploration. The results of the experiments in Chapt. 6 demonstrated that places such as corridors are better exploration targets as they lead to other rooms.

Another typical problem is the localization of mobile robots. In this problem a robot must determine its pose relative to a given map. Recognizing the type of place in which the robot is located can be seen as a high level localization. If the robot is situated in an office, then other places can be discarded, and the robot can concentrate on hypotheses that belong only to office places. This idea was presented in Chapt. 6 together with experiments that corroborated its usefulness.

Since one of the main goals of the semantic labeling is to share spatial terms with humans, it seems necessary to develop robotic systems that can communicate these concepts to people. In Chapt. 7, we introduced an integrated system for conceptual representations of indoor environments. This system included a linguistic framework with capabilities for situated dialogue and map acquisition. The experiments of this chapter showed the interaction capabilities of the system, and demonstrated how a high level conceptual representation could be created based on language communication and semantic information about places.

The semantic information can also represent other kind of objects in the environment. In Chapt. 8, we presented an approach to include the semantic information directly in the data beams from a laser range finder. In this way, much richer information was available from the sensor.

In summary, this book presented many innovative approaches in respect to semantic information about places using mobile robots. As we described in the related literature of several chapters, various posterior works have applied and extended some of the ideas presented here. This indicates that a lot of work can still be done.

In future work it would be interesting to move from the supervised approach presented in this book to other methods with less supervision. One possibility could be to use semi-supervised techniques, in which the robot autonomously creates an initial classification of the environment. This classification can be corrected later on by the user. Another option could be to leave the robot to create a totally unsupervised classification of the places it visits.

In any case, it seems that information coming from the user is important, since someone has to decide how to name the different places. This last issue bring us to the problem of personalization: people can describe the same place with different terms. For example, what to one person might be a living room, could be a sitting room to another. It could be interesting to study approaches able to cope with this flexibility. Some solutions were already presented in Chapt. 7.

Finally, we think that the semantic labeling of places is a research area which will have a high impact in the future of mobile robotics.